Types of DSLs

As the processing power of digital signal processors has grown, so have DSL bit rates. DSL technology began with 144 kb/s basic rate ISDN (BRI), and has evolved to 1.5 and 2.0 Mb/s versions of HDSL, 7 Mb/s ADSL, and now 52 Mb/s VDSL.

2.1 DSL Design Margin

DSLs are designed with a 6 dB SNR margin. This means that the DSL will provide 10^{-7} bit error rate (BER) when the crosstalk signal power is 6 dB greater than the defined “worst-case” crosstalk model. In many cases, the worst-case crosstalk model is a 50-pair binder group filled with 49 self crosstalkers. With pure Gaussian noise, a 6 dB SNR margin would result in a 10^{-24} BER. However, in the real world, noise is often non-Gaussian. Thus, for typical conditions, the 6 dB margin provides assurance that DSLs usually operate at a BER of better than 10^{-9} and that DSLs will provide reliable service even when the transmission environment is worse than normal.

The 6 dB value originated during work on the ANSI basic rate ISDN standard in T1D1.3 (the predecessor of T1E1.4) with a 1985 contribution from Richard McDonald of Bellcore. As described in T1E1.4/95-133, the 6 dB design margin still serves as an appropriate value. The design margin provides for cable variations (aging, splices, wet cable), additional noise in CO and customer premises wiring, other noise sources, imperfect transceiver designs, and manufacturing variations. The amount of design margin is a tradeoff between assuring reliable operation in all conditions and permitting the use of the technology on the longest possible loops.

More sophisticated transmission methods can achieve higher performance, but the need for design margin remains. However, systems that measure margin at start-up can provide the installer an instant indication if the loop has inadequate margin. The installer can then take cor-

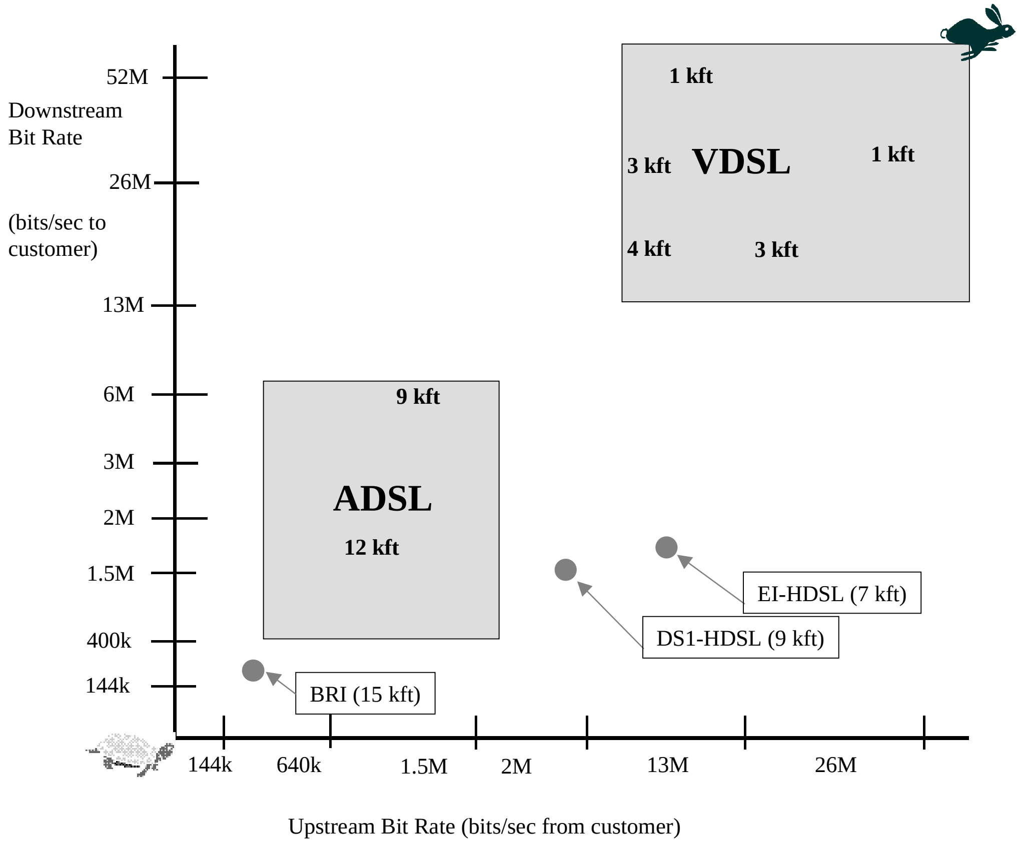

Figure 2.1 DSL bit rates with loop reach shown for 26 AWG wire (kilofeet — no repeater).

rective actions such as finding a better wire pair or removing bridged taps. An argument can be made that systems that provide a real-time indication of transmission margin could reasonably be used with a margin threshold of 5 dB. However, relaxing the design margin by one or two decibels translates into expanding the size of accessible loop population by only about 1% of loops.

2.2 DSL Precursors

One could argue that T1 trunks, E1 trunks, and DDS (digital data service) lines were the first DSLs. Although T1 (1.544 Mb/s Alternate Mark Inversion (AMI) used primarily in North America), and E1 (2.048 Mb/s HDB3) transmission systems were originally intended for use as trunks between Central Offices (COs), they later proved useful as high-speed links from COs to customer sites. T1 carrier was first used by AT&T in 1962. CO-to-CO trunks today are entirely fiber and microwave based. T1/E1 lines are not used today for their original purpose. T1/E1 lines are still used on subscriber lines, but they have their drawbacks. They are expensive and time consuming to install and are usually segregated into binder groups (wire pair bundles) separate from other types of transmission systems. A T1 line consists of four wires. Two wires convey information to the customer, and another two wires convey information from the customer. To reduce near-end crosstalk between the two directions of transmission, one cable binder group (a bundle of wires) carries only outbound T1 pairs, and a different binder group carries only inbound T1 pairs. T1 lines are designed with a maximum of 15 dB (e.g., 2 to 3 kft) of line loss at 772 kHz for the CO end section (CO-to-first repeater), a maximum of 36 dB (e.g., 3 to 6 kft) loss for repeater-to-repeater sections, and up to 22.5 dB of line loss from the last repeater to the customer premises. T1 lines must be unloaded and have no bridged taps. Distances of many miles may be covered by the use of many repeaters. T1 repeaters are powered via +/– 130 volt DC line power. For the purposes of this book, we shall consider T1/E1 and DDS not to be DSLs.

The AMI line code for T1 transmission is simple to implement but is inefficient by today's standards. AMI sends one bit per baud; a baud is one signal element. T1 transmission uses a high transmitted signal power, which generates high levels of crosstalk from 100 kHz to 2 MHz. Other DSLs, which use these same frequencies, can be affected if placed in a binder group with T1 lines. In extreme cases, T1 crosstalk can affect loops in other binder groups.

2.3 Basic Rate ISDN

2.3.1 ISDN Basic Rate Origins

In this book, we shall consider basic rate ISDN (BRI) to be the first in the family of DSLs. Integrated services digital network (ISDN) was first conceived in 1976 and was largely defined by Recommendations developed within the CCITT (now called the ITU, International Telecommunications Union). The ISDN vision was ambitious: a uniform global network for data communications and telephony. Development of the ISDN transmission, switching, signaling, and operations systems required a herculean effort reminiscent of the construction of a continental railway network, only to be followed by the invention of the airplane. The effort to develop ISDN spanned a decade, with efforts of thousands of people from hundreds of companies in more than 20 countries. We estimate that the development of ISDN cost over $50 billion; it is not known if this investment will be fully recovered. ISDN was focused on telephony services and lower-speed packet-switched data. This focus ultimately became a major weakness. ISDN networks were poorly suited for the high-speed packet switching and long holding-time sessions that characterize Internet access. Nonetheless, those who would claim the failure of ISDN must not forget the millions of happy ISDN customers.

ISDN service trials began in 1985. The first North American ISDN service was provided in 1986 by AT&T–Illinois Bell (now called Ameritech) in Oakbrook, Illinois. Early trial BRI systems employed TCM (ping-pong), or alternate mark inversion (AMI) transmission techniques. These early systems were simpler to implement, but 2B1Q (2 binary, 1 quaternary) transmission was selected for the standard transmission techniques for nearly all parts of the world except the Federal Republic of Germany and Austria, which use 4B3T (4 Binary, 3 Ternary), and Japan, which uses a ping-pong AMI transmission method. The loop reach of the 2B1Q and 4B3T systems is greater than the prestandard systems, which quickly faded from use.

The total number of BRI lines in service worldwide grew from 1.7 million in 1994 to nearly 6 million by the end of 1996. The approximate number of ISDN lines for the countries with the largest ISDN deployments are provided below. The 1994 information is based on ITU statistics. The 1996 values are based on information provided by experts from the respective countries. The U.S. 1996 number is from FCC statistics. ISDN deployment is growing at 30% to 50% per year in many countries.

Table 2.1 Basic Rate ISDN Lines in Service

Country	1994 BRI lines	1996 BRI lines
Germany	428,000	2,000,000
United States	352,000	843,115
Japan	320,000*	1,000,000
France	240,000*	1,400,000
United Kingdom	75,000*	200,000

*Extrapolated values.

The deployment of ISDN in Germany was accelerated by government mandate, whereas other countries have followed a market demand deployment model. ISDN service was available in 1996 to about 90% of telephone customers in the countries listed in Table 2.1.

2.3.2 Basic Rate ISDN Capabilities and Applications

BRI transports a total of 160 kb/s of symmetric digital information over loops up to approximately 18 kft (5.5 km, or up to 42 dB of loss at 40 kHz). This is channelized as two 64 k/s B channels, one 16 kb/s D channel, and 16 kb/s for framing and line control. The B channels may be circuit switched or packet switched. The D channel carries signaling and user data packets. An embedded operations channel (eoc) and indicator bits are contained within the 8 kb/s of overhead. The eoc conveys messages used to diagnose the line and the transceivers. The indicator bits identify block errors so that the transmission performance of the line may be measured.

2.3.3 Basic ISDN Rate Transmission

BRI modulates data using one four-level pulse (a quat) to represent two binary bits, hence 2 Binary one Quarternary (2B1Q). Data is sent in both directions simultaneously using echo-canceled hybrid (ECH) transmission. The simple 2B1Q baseband transmission technique sends 160 kb/s using 80 kHz of bandwidth, yielding a modest bandwidth efficiency of 2 bits/s per Hz. Adaptive equalization automatically compensates for attenuation across the transmission band.

BRI can work on a loop with bridged taps, providing the total loss is less than 42 dB at 40 kHz. BRI loops must be unloaded.

2.3.4 Extended-Range Basic Rate ISDN

Loops beyond direct BRI reach of 5.5 km (18 kft) from the CO may be served via alternative methods: BRITE, mid-span repeater, and extended-range BRI.

2.3.4.1 BRITE

Basic rate ISDN transmission extension (BRITE) (see Figure 2.2) uses digital channel banks (e.g., D4 and D5 type multiplexers, which time division multiplex 24 DS0 channels into one 1.544 Mb/s line) and digital loop carriers (DLC) as a means to extend ISDN service to areas served by these channel banks. Special ISDN channel units use three DS0s in the channel bank to transport BRI. Due to the additional channel units, the BRITE configuration has a relatively high cost per line. However, when using preexisting SLC or channel bank equipment, the low start-up cost of BRITE is ideal for serving very small numbers of lines in remote areas.

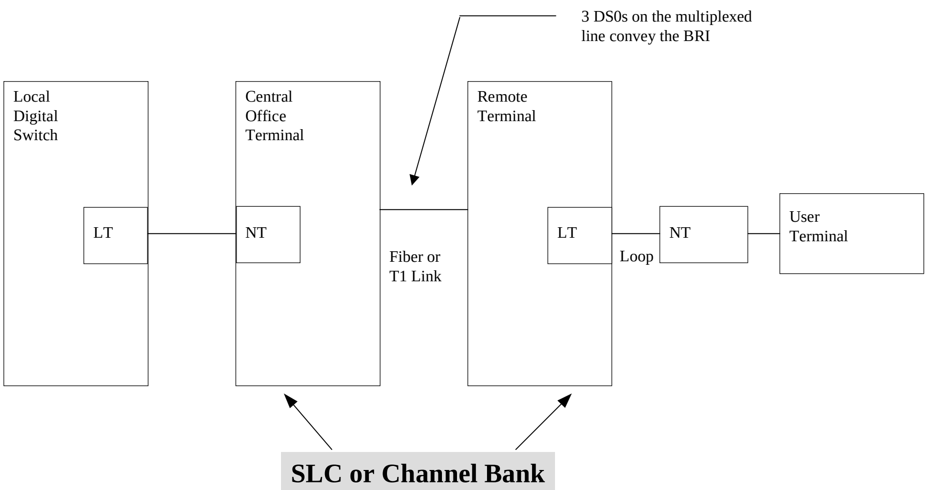

Figure 2.2 BRITE.

2.3.4.2 Mid-Span Repeater

The loop reach may be nearly doubled by placing a repeater in the middle of the loop. See Figure 2.3. Since the repeater is essentially a back-to-back NT and LT, the loop is divided in a tandem pair of DSLs. Each of the two loops may have up to 42 dB of loss at 40 kHz, which corresponds to a total reach of approximately 30 kft (2 × 15). Repeaters are typically located in a multislot repeater apparatus case located in a manhole or mounted on a pole. Since a manhole with available space may not be located at the exact midpoint of the loop, the repeater often is located somewhat off-center. As a result, the attainable loop reach of a repeatered line may be slightly less than twice the unrepeatered reach. Loading coils must be removed from the loop for BRI operations with or without repeaters.

Mid-span repeaters are typically powered via a DC voltage (usually –130 volts DC) in the United States supplied from a CO power feed circuit. For yet longer reach, a second repeater may be employed. The two-repeater configuration is rarely used due to power feeding and administrative complexities. The cost of a repeatered line is dominated by the labor for loop design, the apparatus case, and installation of the apparatus case (including cable splicing). The cost of the repeater electronics is relatively small in comparison.

The repeatered configuration and the BRITE configuration have twice the signal transfer delay (2.5 ms one way) of the direct DSL configuration (1.25 ms).

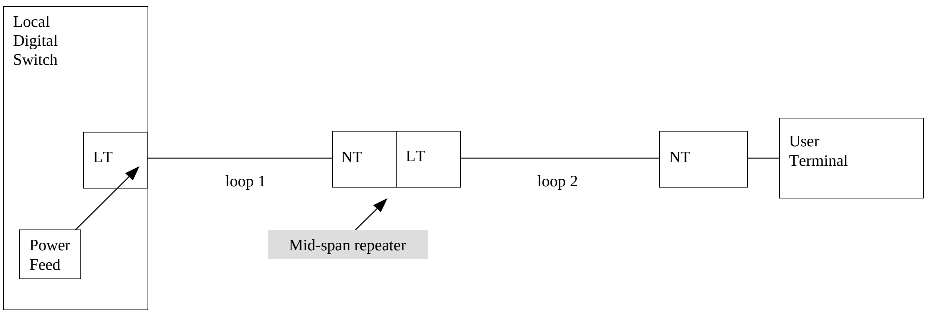

Figure 2.3 Mid-span repeater configuration.

2.3.4.3 Extended-Range BRI

Transmission techniques have advanced since the creation of the BRI standard (ANSI T1.601). Techniques, such as trellis coding, permit 160 kb/s to be transmitted over loops up to 8.5 km (28 kft) without the need for mid-span repeaters. For backward compatibility, the extended-range BRI systems present the standard ANSI T1.601 interface to the LT in the CO switch and also to the customer’s NT1. See Figure 2.4. Normally, a conversion unit is located in a miscellaneous equipment bay in the CO, and the other conversion unit is located in an enclosure located at the outside of the customer premises. However, locating the remote conversion unit in a mid-span location may extend the loop reach further. As a result, a total reach of approximately 43 kft (15 + 28) may be attained. Furthermore, the network-side converter may also be placed remotely, provided that local power is available at the site.

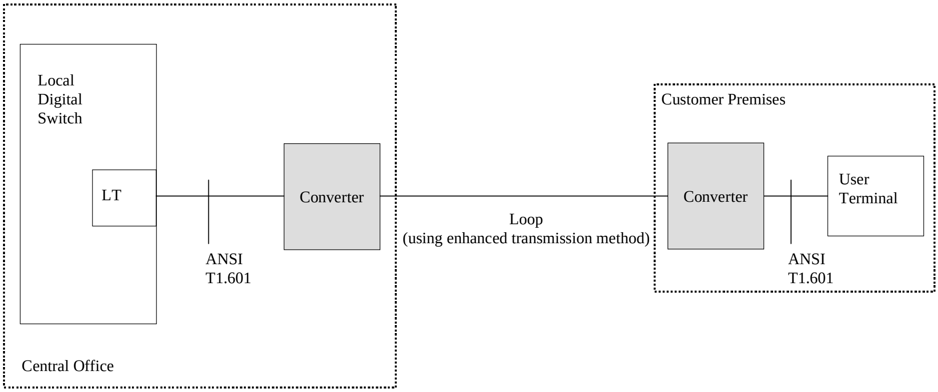

Figure 2.4 Extended-range ISDN configuration.

2.3.5 Digital Added Main Line

BRI transceivers are also used for non-ISDN applications — most notably digital added main line (DAML). DAML systems permits one loop to convey two voice telephone circuits. See Figure 2.5. Voice coder/decoders (CODECs) at each end of the DAML system convert the 64 kb/s BRI B channel to an analog telephone interface. Thus, the traditional voice telephony interface is provided to the CO switch and the customer's phones. DAML systems are used to provide additional telephone service to sites in an area having a shortage of spare wire pairs between the CO and the customers. The DAML unit at the customer's premises is usually powered from CO power fed via the loop. DAML systems using BRI technology have a maximum loop reach of 5.5 km (18 kft). HDSL-based DAML systems can convey more than two voice circuits via one pair of wires.

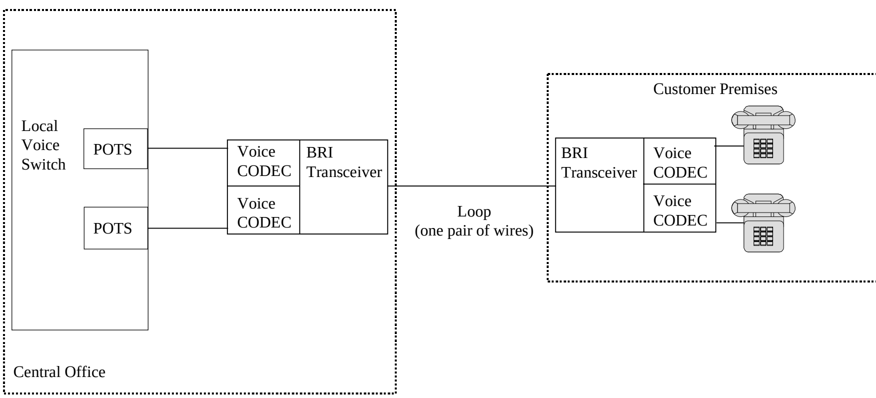

Figure 2.5 Digital added mail line.

2.3.6 IDSL

Another non-ISDN application of BRI transceivers is IDSL (ISDN DSL). The BRI symmetric channels (128 kb/s or 144 kb/s) are concatenated to form one channel for transmission of packet data between a router and the customer's computer. Most forms of IDSL will work with a conventional ISDN NT at the customer end of the line. Thus, with IDSL the ISDN local switch is replaced by a packet router. This configuration is used for Internet access.

2.4 HDSL

2.4.1 HDSL Origins

The early concept definition of HDSL (high-bit-rate digital subscriber line) took place in late 1986 at AT&T Bell Laboratories and Bellcore. HDSL transceiver designs were essentially up-scale basic rate ISDN designs. Laboratory prototype HDSL systems appeared in 1989. The first HDSL was placed into service in March 1992 by Bell Canada using equipment manufactured by Tellabs Operations Inc. in Lisle, Illinois. Nearly every major telephone company in the world now uses HDSL. In 1997, approximately 450,000 HDSL lines were in service worldwide, with approximately 350,000 lines of these lines in North America. HDSL deployment is growing at more than 150,000 lines per year. In October 1998, the ITU approved Recommendation G.991.1 for first generation HDSL; this is based closely on the ETSI Technical Specification TM-03036. The ITU has started work on a second generation HDSL (HDSL2) Recommendation that will be called G.991.2.

The need for HDSL became evident as T1 and E1 transmission systems ceased to be used for their original purpose as interoffice trunks and saw rapid growth as private lines from Central Office to customer premises. T1/E1 transmission systems operated over the existing telephone wires, but at a large cost for special engineering, loop conditioning (removal of bridged taps and loading coils), and splicing for apparatus cases to hold the repeaters that were required every 3,000 to 5,000 feet. The transmission methods used for T1/E1 lines placed high levels of transmit signal power at frequencies from 100 kHz to above 2 MHz; this required the segregation of T1/E1 lines into binder groups separate from many other services. In addition to being expensive to install and maintain, T1/E1 lines often took many weeks from service order to service turn-up. What was needed was a *plug-and-play* transmission system that could quickly and easily provide 1.5 or 2 Mb/s transport over most subscriber lines, thus HDSL.

HDSL's benefits are largely due to the elimination of mid-span repeaters. Each repeater site must be custom-engineered to assure that each section of the line remains within the limits for signal loss. The repeated signals can cause severe crosstalk; thus special care must be taken in the design of repeatered facilities to avoid excessive crosstalk to other transmission systems. The repeater is placed in an environmentally hardened apparatus case in a manhole or on a pole. The apparatus case must be spliced into the cable. The apparatus case costs far more than the repeaters it holds. A repeater failure results in a field service visit. Repeaters are usually line powered; this requires a special line feed power supply at the CO. Most of the power fed by the CO power supply is wasted due to loop resistance and power supply inefficiencies.

HDSL is also preferred over traditional T1 carriers because HDSL provides more extensive diagnostic features (including SNR measurement) and HDSL causes less crosstalk to other transmission systems because its transmit signal is confined to narrower bandwidth than the traditional T1 carrier.

2.4.2 HDSL Capabilities and Application

HDSL provides two-way 1.544 or 2.048 Mb/s transport over telephone lines up to 3.7 km (12 kft) of 0.5 mm (24 AWG) twisted pair without a mid-span repeater and up to nearly twice this distance with one mid-span repeater. More than 95% of HDSL lines have no repeater. As a rule, no line conditioning or binder group segregation is required for HDSL. HDSL provides reliable transmission over all carrier serving area (CSA) lines with a typical bit error rate of 10^{-9} to 10^{-10}. DS1 (1.544 Mb/s) HDSL systems use two pairs of wires, with each pair conveying 768 kb/s of payload (784 kb/s net) in both directions. Thus, the term *dual duplex* is used to describe HDSL transmission. See Figure 2.6. E1 (2.048 Mb/s) HDSL systems have the option of using two or three wire pairs, with each wire pair using full-duplex transmission. The three-pair 2.048 Mb/s HDSL uses the very same 784 kb/s transceivers as the 1.544 Mb/s systems. HDSL loops may have bridged taps, but no loading coils.

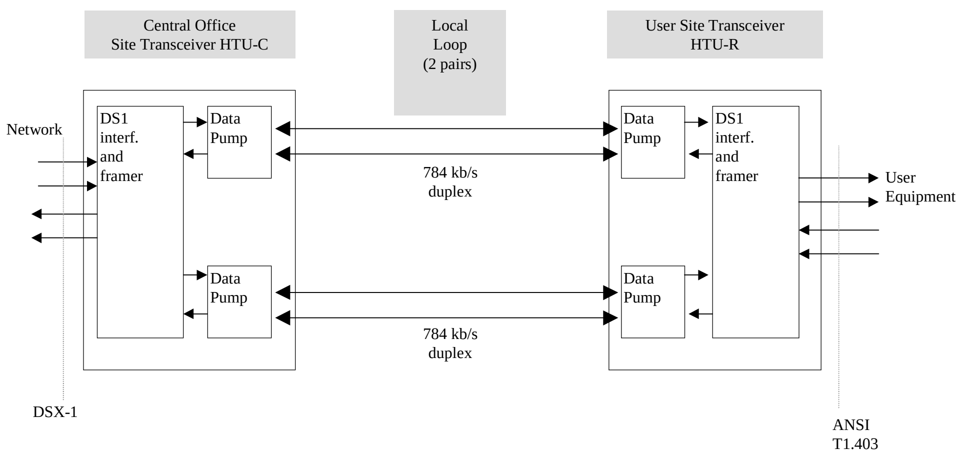

Figure 2.6 HDSL Dual duplex transmission system.

Despite early descriptions of HDSL as a “repeaterless technology,” HDSL repeaters are commonly used for lines beyond HDSL’s nonrepeatered reach of 2.75 to 3.7 km (9 to 12 kft). For 24 AWG wire, up to 7.3 km (24 kft) can be reached with one repeater and up to 11 km (36 kft) with two repeaters. The actual reach can be less where it is not possible to place the repeater at the precise midpoint. Early two-repeater HDSL systems powered the first repeater via line power from the CO, and the second repeater was powered from the customer site. Power feeding from the customer site poses maintenance and administrative drawbacks. With recent reductions in transceiver power consumption, it has become possible to line power two tandem HDSL repeaters from the CO power source.

Primary rate (1.544 or 2.048 Mb/s) private line circuits from a user to the network is the leading HDSL application. HDSL is a popular means to connect private branch exchange (PBX) and packet/ATM data equipment to the public network. HDSL links are used to link wireless radio sites into the landline network. HDSL is used to connect small digital loop carrier (DLC) sites to the CO. During its first few years of use, the high cost of HDSL equipment limited the use of HDSL to situations where there was no economical site to place a repeater apparatus case. By the end of 1994, the price of HDSL equipment had reached the point where HDSL was economically preferred over traditional T1/E1 transmission equipment for nearly all new installations. Traditional T1/E1 equipment is still used for very short lines (less than 3 kft) that require no repeater and for very long lines (more than 30 kft) that would require more than two HDSL repeaters.

The annual maintenance costs of HDSL lines are lower than for T1/E1 lines because HDSL lines have fewer repeaters to fail, superior transmission robustness, and improved diagnostic capabilities. However, existing T1/E1 lines are rarely replaced by new HDSL lines due to the cost of installing the new line.

Although HDSL is most used by local exchange carriers (telephone companies), there is some use of HDSL in private networks to provide high-speed links between buildings within a campus.

2.4.3 HDSL Transmission

Echo-canceled hybrid dual-duplex 2B1Q transmission is used for nearly all HDSL systems worldwide, with some discrete multitone (DMT) and carrierless AM/PM (CAP) systems used in parts of Europe. For 1.544 Mb/s transport, dual-duplex transmission uses each pair of wires to convey one-half of the two-way payload (768 kb/s) plus framing and embedded operations channel (eoc) overhead of 16 kb/s for a total of 784 kb/s transmission. Two pairs of wires make up the total 1.544 Mb/s HDSL transmission system. Since the same overhead information is conveyed on both wire pairs, the receiver selects one wire pair for the overhead information. Usually, the receiver selects the wire pair with the better signal-to-noise ratio (SNR).

Several alternatives were considered for the original HDSL systems: single duplex, dual simplex, and dual duplex.

Single duplex provides the benefits of using only one pair of wires and requiring only one transmitter-receiver pair at each end of the line. See Figure 2.7. The two directions of transmission may be separated by frequency division multiplexing (FDM) or by echo-canceled hybrid (ECH) transmission. However, transmitting the full payload rate over most loops was beyond the

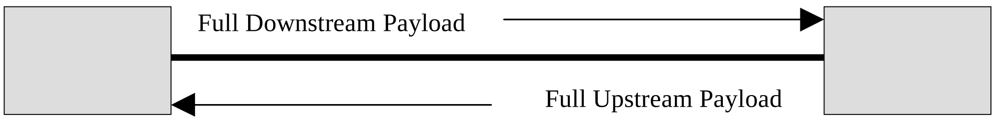

Figure 2.7 Single-duplex HDSL.

abilities of the technology in the early 1990s. Furthermore, the large bandwidth needed presented concerns for spectral compatibility with other types of transmission systems. The single-pair 1.544 Mb/s HDSL systems (sometimes called SDSL) developed in the early 1990s had a loop reach of less than 6 kft on 26 AWG wire; this short reach greatly limited their utility. Only with the most advanced technology available in the late 1990s does it appear that single-duplex 1.544 Mb/s transport may become practical for full carrier serving area (CSA) loop reach. HDSL2, described in Section 2.4.4, employs single-duplex transmission.

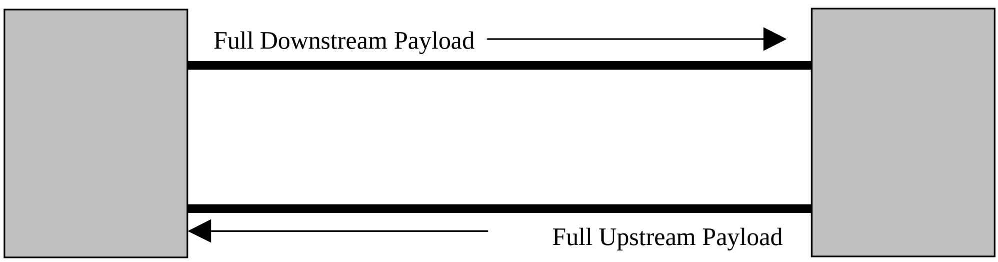

Figure 2.8 Dual-simplex HDSL.

Dual-simplex transmission uses two pairs of wires, with one wire pair carrying the full payload in one direction and the second wire pair carrying the full rate transmission in the opposite direction. See Figure 2.8. This provides a very simple means for the separation of the signals in the two different directions of transmission. Traditional T1 carrier uses dual-simplex transmission. Dual-simplex transmission has the disadvantage of transmitting a signal with a wide frequency bandwidth, which is subject to great loss and crosstalk at the higher frequencies. Due to crosstalk, the signals sent on the two pairs of wires are not fully segregated. Thus, the dual-simplex transceivers may be simpler, but the resulting performance is inferior to dual duplex.

Dual-duplex transmission improves the achievable loop reach and spectral compatibility by sending only one-half of the total information on each wire pair. See Figure 2.9. HDSL further reduces the bandwidth of the transmitted signal by using ECH transmission to send the two directions of transmission in the same frequency band. The dual-duplex HDSL transmitted signal power is progressively less for frequencies above 196 kHz. As a result, the signal crosstalk and attenuation is reduced. Another advantage of dual-duplex transmission is that using one pair of wires can easily provide a half-rate transmission system.

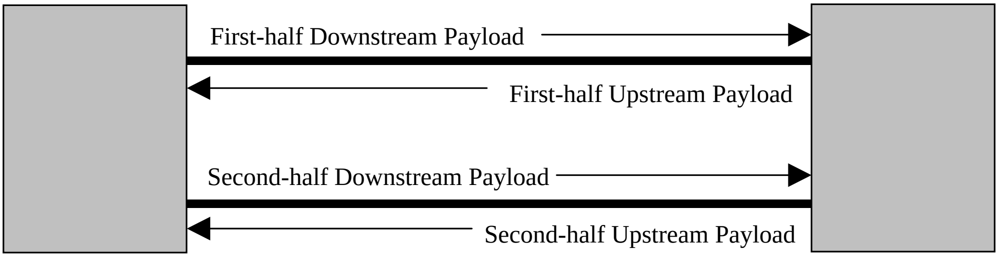

Figure 2.9 Dual-duplex HDSL.

One-pair, fractional-rate HDSL systems are used for the transport of fractional-rate private-line services of 768 kb/s and below and also for small loop carrier systems supporting 12 or fewer voice channels. Fractional-rate HDSL plug-ins for the D4 channel bank permit up to 12 DS0s of HDSL transport information to be multiplexed with information from other channel units in the same D4 channel bank.

Identical maintenance information (indicator bits and eoc) is conveyed on each wire pair of the dual-duplex HDSL system. This redundant transport of the overhead permits the use of the same transceiver components for one-, two-, and three-wire pair HDSL systems. Furthermore, the redundant overhead information ensures reliable operation of the maintenance functions even if transmission has failed or is impaired on one of the loops.

2.4.3.1 Timing

HDSL framing contains positions for stuff quats (quaternary symbols representing two bits). The stuff quats are added to frames as needed to synchronize the T1/E1 payload bit rate to the HDSL line transmission rate. To permit effective echo-canceled operation, the upstream and downstream HDSL symbol rates must be exactly the same. There are some situations where the upstream T1/E1 payload bit rate may slightly differ from the downstream payload bit rate. The stuff quats, together with a little buffering, permit the payload rate to differ slightly from the HDSL line rate. Many public network T1/E1 circuits are loop timed, which means that the upstream timing is derived from the downstream bit clock. Loop-timed circuits do not require stuff quats. However, this feature is provided on all HDSLs in the event that a circuit is not loop timed.

2.4.3.2 Delay (Latency)

Traditional T1 transmission systems have an end-to-end signal transfer delay of less than 100 microseconds. Due to the digital signal processing, HDSL circuits typically have approximately 400 microseconds of signal transfer delay as measured one way between the DSX-1 interface and the T1.403 interface. The additional delay found in HDSL systems rarely poses a problem, but there have been a few cases where upper-layer protocol handshakes have timed-out due to the total end-to-end delay. For this reason, HDSL systems are designed to assure that the one-way signal transfer delay for a nonrepeatered HDSL line is less than 500 microseconds.

HDSL lines with one mid-span repeater have twice this delay. Other network elements including SONET terminals and digital cross-connect systems (DCS) may have delays in excess of 500 microseconds. Thus, end systems should allow for a few milliseconds of network delay regardless of the presence of HDSL.

2.4.3.3 Bit Error Rate

HDSL systems, like BRI and ADSL, are designed to assure better than 10^{-7} bit error rate (BER) on worst-case loops with crosstalk noise power 6 dB greater than the theoretical worst-case crosstalk model. This design criteria is based on the engineering judgment and agreements among leading experts in the T1E1.4 standards working group. A decade of field experience has proved this design criteria to be a good compromise between over-engineering (underuse due to overly conservative design) and under-engineering (poor reliability due to the lack of robustness).

Nonetheless, there are two prevalent misconceptions regarding the BER design of HDSLs and other DSLs. The first misconception is that most HDSLs operate at 10^{-7} BER. The 10^{-7} BER value is for a *worst-worst*-case situation, which is rarely seen in the field. Approximately 99% of HDSLs in the field operate at a BER better than 10^{-9}. When errors do occur, they tend to appear in short bursts. This bursty characteristic is more benign than random bit errors. The second misconception is that HDSLs are grossly over-engineered. Considering the design with 6 dB of margin beyond a worst-worst-case model, it is easy to see why some people have this opinion. However, the seemingly overconservative design is justified for two reasons. HDSLs are required to operate reliably all the time for all qualified loops. Unlike voice-band modems used on switched circuits, one can not “hang up” and dial up again in hopes of obtaining a better connection. Furthermore, the real-world environment contains many impairments that can consume the 6 dB design margin (e.g., water in cable, bad splices, poor-quality inside wire, or a line longer than indicated in cable records).

2.4.4 Second-Generation HDSL

Standards development for a second-generation HDSL technology (HDSL2) began in 1995 to provide the same bit rate and loop reach as first-generation HDSL but using one pair of wires instead of two. This reduction of wire pairs is important because many LECs have a shortage of spare wire pairs in some areas. HDSL2 uses more sophisticated modulation and powerful coding techniques. Carefully selected offset frequency placement of the up- and downstream directions is used for HDSL2 to help combat crosstalk. The newer versions of HDSL borrow many ideas from ADSL. A rate-adaptive version of HDSL is likely to appear. There has also been consideration of HDSL placed in a frequency band above baseband analog voice or above basic rate ISDN. The term SDSL (symmetric, or single-pair DSL) has also been used to describe later versions of HDSL.

2.4.4.1 Performance Requirements

Although several suggestions for line codes were made to T1E1.4 following a request in 1995 (T1E1.4/95-044), progress was slow until detailed requirements were established. These requirements, specified primarily by the operating companies, were first proposed in March 1996 (T1E1.4/96-094 and T1E1.4/96-095) and revised since that time (T1E1.4/97-180, 180R1, 181, 469). They currently are as follows:

Loop Reach: CSA coverage (same as two-pair ANSI HDSL):

9000 ft (2.7 km) of 26 AWG (0.4 mm)

12000 ft (3.6 km) of 24 AWG (0.5 mm)

Bridged taps limited to 2.5 kft total, 2 kft per tap

Cable parameters as specified in T1.601

Impairments/Performance: minimum of 5 dB of performance margin with 1% worst-case crosstalk from the following interfering services:

49-disturber HDSL

39-disturber HDSL2

39-disturber EC-ADSL

49-disturber FDM-ADSL

25-disturber T1

24 T1 + 24 HDSL2

24 FDM-ADSL + 24 HDSL

Spectral Compatibility: To all existing services, no more impairment than the services tolerate today, with the following exceptions: shall not degrade HDSL by more that 2 dB (T1E1.4/97-434, 440R1) and ADSL by more than 1 dB (T1E1.4/97-444). These services include the following customer interface specifications: T1.413 (ADSL), TR-28 (HDSL), ANSI T1.403 (DS1), and T1.601 (ISDN-BRA).

Latency: The maximum latency for HDSL2 is to be no more than for HDSL (500 µs).

2.4.4.2 Impairments

The impairments were selected as typical of severe-case crosstalk combinations that HDSL2 may encounter. Of the CSA test loops in ANSI TR-28, it was found that CSA 4 represented the limiting case. Near-end crosstalk coupling is modeled using the Unger model, as specified in T1E1.4/96-036, and far-end crosstalk coupling is modeled as specified in ANSI T1.413 Annex B. Models for T1.601, TR-28, and T1.403 transmitters were taken from T1.413 Annex B. A variety of models for the echo-canceled (EC) and frequency division multiplexed

(FDM) version of ADSL were used. Most of the latest work incorporated modified versions of the PSDs from Annex B.4 and B.5 of T1.413. Most of the variation dealt with split points for the FDM, roll-off of the upstream PSD, and low frequency roll-off of the downstream EC PSD. It was generally accepted that the Sinc term from B.4 and B.5 should not be used. The mixed crosstalk cases were added to the requirements (T1E1.4/97-180, 181) after it was found that they were worse than homogeneous crosstalk for non–self-NEXT limited modulation techniques.

Impulsive noise has not been considered to be a significant impairment in the T1E1.4 deliberations. Also, all calculations regarding spectral compatibility are with respect to other ANSI DSLs. No calculations/measurements have been published with respect to their ETSI or ITU counterparts.

2.4.4.3 Spectral Compatibility

Determining spectral compatibility between new and existing service proved to be a significant challenge. For ISDN-BRA it was easily shown that the proposed line codes were definitely less of an impairment than self-NEXT. The other listed services were not so easy. For T1.403, (DS1/T1), the initial technique involved measuring the total amount of NEXT power present at the T1 receiver. This was compared to the power from T1 self-crosstalk to see if a problem resulted. In several contributions, the crosstalk was weighted by a measured (T1E1.4/97-071) or calculated T1 receive filter. Later, it was noticed that spectral compatibility with T1 was eased since the first segment from the CO only has 15 dB of loss and not the 30 dB that the other segments must operate over.

With ADSL, spectral compatibility was defined by ideal margin calculations. It was found that slight changes in the assumed noise floor, transmit PSDs, and minimum carrier number (for the FDM case) could have a significant impact on the performance estimate. Most calculations have found that the agreed PSD would degrade ADSL (T1.413) margins by less than 1 dB for the worst-case standard interferer combination.

With HDSL, the initial compatibility work was done using theoretical calculations, but later testing (T1E1.4/97-339) showed that for some modulation formats this was insufficient. (This is addressed more completely in the next section.)

2.4.4.4 Modulation Format

Early on, both symmetric echo-canceled transmission (SET) and frequency division multiplexed transmission (FDM) approaches were considered. SET proved to have a self-crosstalk limitation 2 to 3 dB short of the requirements. In contrast, FDM transmission is not limited by self-NEXT, but by (ingress) crosstalk from other services. It is also limited by (egress) crosstalk into other services due to the higher transmit frequencies involved with transmitting a symmetric payload in this manner. The ingress and egress crosstalk make the FDM solution even less desirable than SET. A "staggered FDM" scheme (T1E1.4/96-340) was proposed in an attempt to limit these undesirable effects.

In T1E1.4/97-073, partially overlapped echo-canceled transmission (POET) was proposed. POET involves overlapping, but not identical, spectra in the two transmit directions.

These spectra are carefully shaped to provide maximum performance in the presence of self- and foreign crosstalk while causing minimal degradation of other services due to POET crosstalk into other services. Various versions of this approach were proposed in the standards process, all incorporating the same basic concept (POET-PAM (97-073), OverCAPped (97-179), OPTIS (97-237,320), MONET (97-307,412)).

One characteristic that all these POET modulation schemes exhibit is the effect of heterogeneous crosstalk on performance. For SET, performance in homogeneous and heterogeneous crosstalk is quite similar. However, with POET modulation it is possible to have performance in the presence of heterogeneous crosstalk that is significantly worse than performance in the presence of homogeneous crosstalk. The actual performance of these systems also varies with the symbol rate and modulation type. With digital oversampled transceivers, it is possible to decouple the transmit PSD from the actual symbol rate. (This uses principles similar to those used in a traditional CAP transceiver.) This property was first exploited in a CAP version of POET (T1E1.4/97-170), but ultimately it was found that with the impairment set for HDSL2, PAM modulation reaps even larger benefits from this decoupling (T1E1.4/97-237). For each unique crosstalk PSD, there is a particular symbol rate that gives maximum performance. For ease of implementation, a single symbol rate that offers performance at near the optimal level over a wide variety of crosstalk PSDs is desirable.

Most of the later modulation scheme proposals have PSDs where some of the upper frequencies are boosted above the nominal value. These “boosted” portions of the PSD are also above the level of any other DSL that operates at those frequencies. This boost was first introduced in T1E1.4/97-170 and incorporated in a pronounced way in T1E1.4/97-273. After this concept was introduced, it was discovered that, when transmitting such signals, theoretical calculations alone were not sufficient to predict spectral compatibility with existing services. Testing of deployed HDSL systems (T1E1.4/97-339) revealed a significant difference between theoretical calculations and measured performance in the presence of OPTIS crosstalk. As a result, modifications were made to the proposed HDSL2 PSD to reduce this degradation (T1E1.4/97-435). Final measurements after modification showed this degradation to be 2 dB or less (T1E1.4/97-434, 440R1).

The current agreed-upon modulation format incorporates the key elements proposed in T1E1.4/97-257:

- The upstream and downstream transmitters will each have a unique spectral shape.
- The upstream and downstream transmitter spectra will be partially overlapped in frequency.
- The shape of the transmit spectrum will be decoupled from the symbol rate to allow for flexible use of excess bandwidth.
- The transmit modulation used will be pulse amplitude modulation (PAM).
- Coded modulation will be used.

The result (T1E1.4/97-435) is a POET system using a modification of the OPTIS PSD. This modulation format uses PAM with 3 information bits per symbol and a 16-level coded constellation. A symbol rate of one-third the payload rate in both the NT-to-LT and LT-to-NT directions was chosen as a good compromise symbol rate. Advantage is achieved through the use of excess bandwidth in the LT-to-NT direction and a high degree of spectral shaping in both directions. The transmit power is approximately 16.5 dBm in each direction. This modulation technique has been shown (via optimal DFE calculations) to have a minimum theoretical uncoded margin on the worst-case required loop of 1.0 dB. Realized performance near the theoretical values is only possible through the use of a fractionally spaced equalizer.

The spectral shaping employed in the agreed PSDs was designed specifically for this application: to maximize the folded SNR at the HDSL2 receiver (with a symbol rate of 517.33 kHz) in the presence of the specific crosstalk mixes listed in the requirements, while simultaneously minimizing the impact of egress crosstalk into the ANSI DSLs. See Figure 2.10. Not only are maximizing HDSL2 performance while minimizing egress crosstalk conflicting goals, simply trying to simultaneously minimize egress crosstalk into two different DSLs may result in a conflict. For instance, to minimize impact into ANSI HDSL, it is desirable to use lower and wider PSDs, whereas to minimize crosstalk into ANSI T1.413 ADSL, a narrower and higher PSD is preferable. Since this optimization was done specifically for the set of ANSI DSLs in the requirements, one should not extrapolate that a frequency scaled version of the same shaping filter would be the best solution for the ITU version of HDSL2. Further work on this issue will be required before conclusions can be drawn.

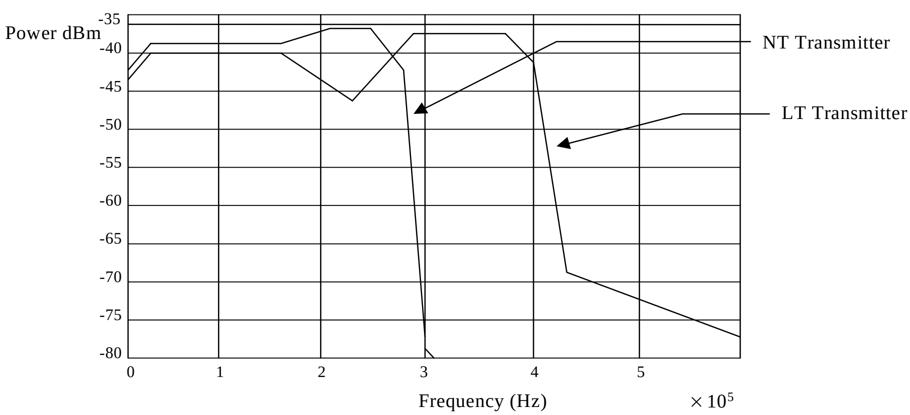

Figure 2.10 Nominal transmit PSDs for the HDSL2 modulation format.

2.4.4.5 Trellis Code Structure

To meet the difficult requirements, coded modulation must be used to increase the crosstalk limited performance of HDSL2 over that possible using only uncoded modulation. With the latency limitation, interleaved concatenated coding and turbocoding techniques proved to be impossible. This left traditional trellis-coded modulation in combination with channel equalizing precoding (such as Tomlinson-Harashima precoding). Although multidimension and multilevel coding approaches were examined, a simple one-dimensional trellis-coded PAM approach has proved to be the best at achieving high coding gains with minimal latency (T1E1.4/97-337).

With the agreed modulation format, 4.0 dB of coding gain is needed to achieve the requirements. For one-dimensional codes with Viterbi decoding, 32 states are needed to achieve over 4 dB of realized BER coding gain. However, the 5.0 dB margin budget must include some non–coding-related implementation loss which must be made up for in coding gain. (This implementation loss, can also affect the realized coding gain.) Thus, a variable amount of coding gain may be needed, based on other losses in the design of the system. The agreement includes a programmable rate ½ one-dimensional trellis (T1E1.4/97-443). See Figure 2.11. This structure allows receivers the flexibility to trade off complexity of the trellis-decoder with complexity in the remainder of the transceiver. This programmable structure also allows for alternative decoding techniques to be used (e.g., sequential decoding), which require substantially different codes than those used for Viterbi decoding.

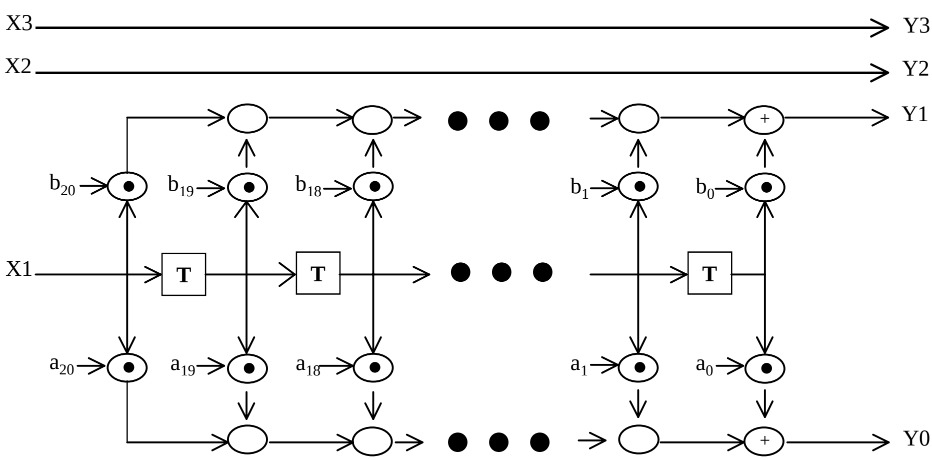

Figure 2.11 Block diagram of the programmable encoder (from T1E1.4/97-443).

2.4.4.6 Complexity Differences Relative to HDSL

Since the requirements for ANSI HDSL2 were quite challenging, a significant complexity increase as compared to HDSL is required to meet them. In this section, we briefly examine the complexity differences.

- The transmit power of HDSL2 is 3 dB higher than that of HDSL. Furthermore, the use of precoding and the spectral shaping together cause the peak-to-rms ratio to be larger than that of 2B1Q-based HDSL. Higher peak voltage levels will increase the power consumption of the line driving circuitry.
- The channel equalizing precoder has a function that is similar to the feedback filter of the decision feedback equalizer used in HDSL. However, the data in the precoder are many bits (12–16) wide instead of the 2 bits for 2B1Q DFE. This increases the complexity. Furthermore, the presence of the precoder increases the complexity of the echo canceler in the same manner.
- To meet the performance requirements through the use of Viterbi decoding, it appears that trellis codes on the order of 512 states will be necessary. The Viterbi decoder for such a code represents an enormous increase in both processing power and memory compared to uncoded HDSL systems.
- To obtain adequate performance, HDSL2 requires a fractionally spaced equalizer and echo canceler, both which are of significantly higher complexity than the baud-spaced equivalents typically employed in HDSL transceivers.

2.5 ADSL

2.5.1 ADSL Definition and Reference Model

Asymmetric digital subscriber line (ADSL) is a local loop transmission technology that simultaneously transports the following via one pair of wires:

- Downstream (towards customer) bit rates of up to about 9 Mb/s
- Upstream (towards network) bit rates of up to 1 Mb/s
- Plain old telephone service (POTS, i.e., analog voice)

The bit rate towards the customer is much greater than from the customer, hence the term *asymmetric*. Analog voice is transmitted at baseband frequencies and combined with the passband data transmission via a low-pass filter (LPF) that is commonly called a “splitter.” In addition to the splitters, the ADSL consists of an ADSL transmission unit at the Central Office side (ATU-C), a local loop, and an ADSL transmission unit at the remote side (ATU-R).

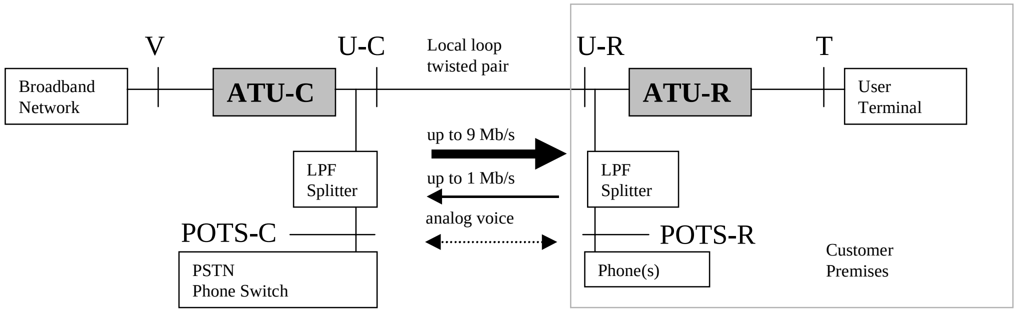

Figure 2.12 ADSL reference model.

2.5.2 ADSL Origins

The early conceptual definition of ADSL began in 1989, primarily by the work of J. W. Lechleider and others at Bellcore. Early ADSL development began at Stanford University and AT&T Bell Labs in 1990. ADSL prototypes arrived at telephone companies and Bellcore Laboratories in 1992, and early ADSL products moved into field technology trials in 1995. ADSL drew from the earlier work on voice-band modems, ISDN, and HDSL.

In October 1998, the ITU gave preliminary approval to (“determination” in ITU language) a set of ADSL Recommendations. Recommendation G.922.1 specifies full-rate ADSL. This is nearly identical to ANSI T1.413 Issue 2 with two major exceptions:

1. the T1.413 tone-based initialization sequence is replaced by a message-based process described in G.994.1, and
2. a special mode has been added to improve performance in the presence of crosstalk from TCM-type ISDN used in Japan.

Recommendation G.992.2 (previously known as G.lite) specifies ADSL for use without a POTS splitter. G.922.2 is based on G.992.1 with the following major differences:

1. added provisions for power saving modes at the ATU-C and ATU-R,
2. the addition of a fast retrain mechanism to permit rapid recovery from on/off-hook events,
3. the number of tones is reduced from 256 to 128, and
4. the number of bits per tone is reduced from 15 to 8.

Recommendation G.994.1 (previously known as G.hs) specifies a message-based initialization handshake to allow multimode DSL transceivers to negotiate a common operating mode. Recommendation G. 995.1 provides an overview of the family of DSL recommendations. Recommendation G.996.1 specifies methods for measuring the performance of DSL equipment.

Recommendation G.997.1 specifies the physical layer operations, administration, and maintenance provisions for ADSL. This includes the ADSL embedded operations, channel (eoc), and Management Information Bases (MIBs).

2.5.3 ADSL Capabilities and Application

2.5.3.1 2.5.3.1ADSL1, ADSL2, and ADSL3

The ADSL concept evolved during the early 1990s. At first, ADSL was considered at a fixed rate 1.5 Mb/s downstream and 16 kb/s upstream for video dial tone (VDT) MPEG-1 applications. Some members of the industry refer to this as ADSL1. Later, it became clear that some applications would require higher speeds and that more advanced transmission techniques would enable the higher speeds. Three Mb/s downstream and 16 kb/s upstream ("ADSL2") was briefly considered to enable two simultaneous MPEG-1 streams. In 1993, interest shifted to ADSL3 with 6 Mb/s downstream and at least 64 kb/s upstream to support MPEG2 video. The Issue 1 ANSI T1.413 ADSL standard grew out of the ADSL3 concept. The terms ADSL1, ADSL2, and ADSL3 have seen little use after approval of the ANSI T1.413 standard.

2.5.3.2 RADSL

Rate-adaptive digital subscriber line (RADSL) is a term that applies to ADSL systems capable of automatically determining the transport capacity of the individual local loop and then operating at the highest rate suitable for that local loop. The ANSI T1.413 standard provides the capability for rate-adaptive operation. The rate adaptation occurs upon line start-up, with an adequate signal quality margin to assure that the start-up line rate can be maintained during nominal changes in the line transmission characteristics. Thus, RADSL will automatically provide higher bit rates on loops with better transmission characteristics (less loss or less noise). RADSL implementations have supported maximum downstream rates in the range of 7 to 10 Mb/s and maximum upstream rates in the range of 512 to 900 kb/s. On long loops (5.5 km/18 kft or more), RADSL may operate at rates of about 512 k/s downstream and 128 kb/s upstream.

RADSL borrowed the concept of rate adaptation from voice-band modems. RADSL provides the benefit of one version of equipment that assures the highest possible transmission rate for each local loop and also permits operation on long loops at lower rates.

2.5.4 ADSL Transmission

The ADSL concept contains two fundamental parts: (1) near-end crosstalk is reduced by having the upstream bit rate and bandwidth much less than the downstream bit rate, and (2) simultaneous transport of POTS and data by transmitting data in a frequency band above voice telephony. Two-way transmission of multimegabit rates is not on most telephone lines due to the combined effect of loop loss and crosstalk. As shown in Figure 2.13, received signal power diminishes in proportion with frequency, and received crosstalk noise increases with frequency. Thus, two-way transmission is not possible at frequencies where the crosstalk noise overwhelms the received signal.

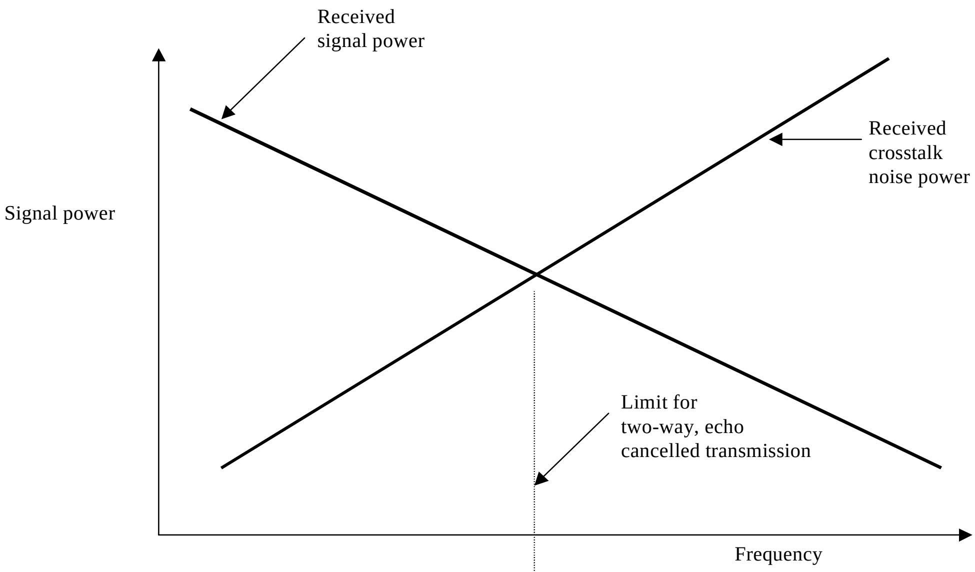

Figure 2.13 Two-way transmission is limited to lower frequencies.
This figure is based upon one created by Kim Maxwell of Independent Editions.

ADSL performs two-way transmission where possible: below the two-way cut-off frequency. The upper frequencies that are unsuitable for two-way transmission are used for one-way transmission. This permits downstream transmission rates far in excess of those possible for two-way transmission.

Many ADSL systems use a frequency division multiplexed (FDM) transmission technique, which places upstream transmission in a frequency band separate from the downstream band to prevent self-crosstalk. The guard band is necessary to facilitate filters that prevent POTS noise from interfering with the digital transmission. See Figure 2.14.

Some ADSL systems use an ECH transmission technique where the upstream frequency band resides within the downstream band. See Figure 2.15. By overlapping the bands, the total transmitted bandwidth may be reduced. However, the ECH is subject to self-crosstalk, and its implementation involves more complex digital signal processing. There is some debate as to whether the digital complexity is offset by simplification of the analog front end.

Due to the lack of self-crosstalk at the CO end, FDM ADSL offers much better upstream performance than ECH ADSL. However, the wider downstream bandwidth of ECH ADSL permits better downstream performance, especially for shorter loops.

The performance of symmetric DSL is primarily limited by self–near-end-crosstalk (self-NEXT). ADSL overcomes self NEXT at the customer end simply by reducing the source of the self-NEXT. By reducing the upstream bit rate, the upstream channel may be positioned to mini-

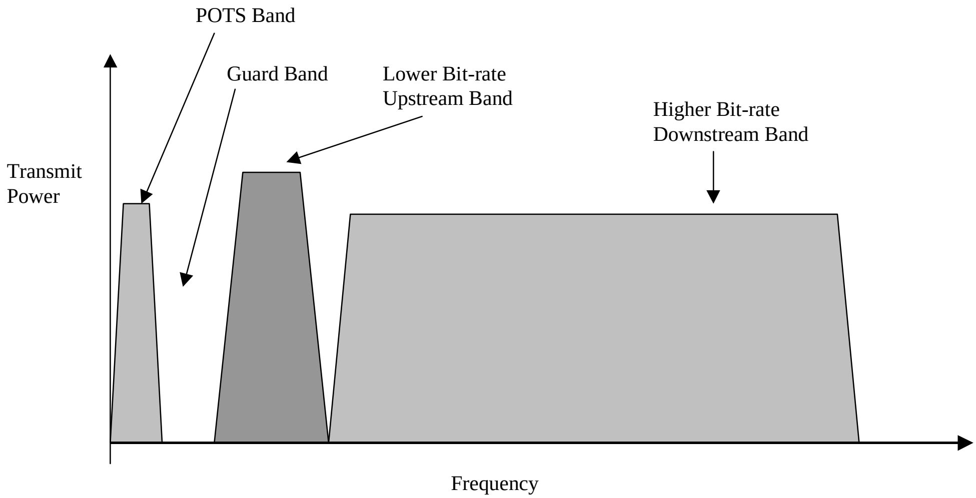

Figure 2.14 FDM ADSL.

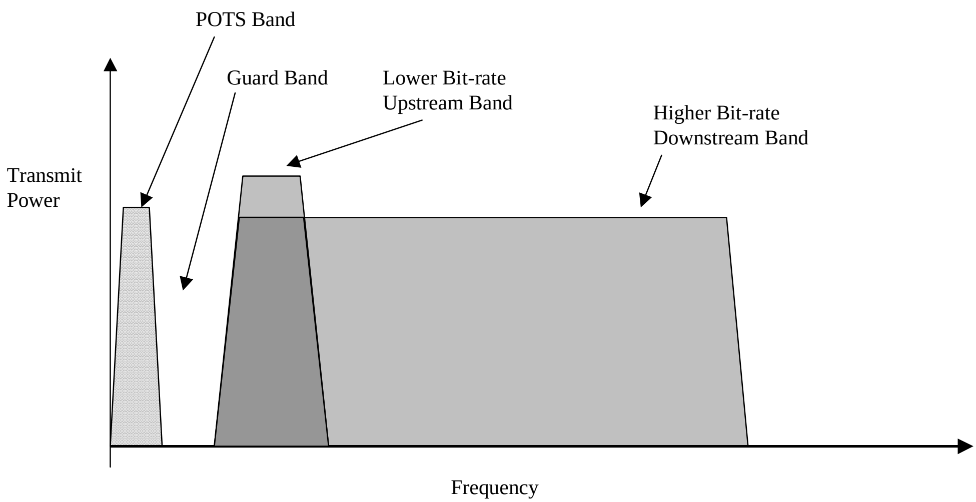

Figure 2.15 ECH ADSL.

mize crosstalk into the downstream transmission. For ADSL, reception of the upstream channel is made easier by placing it at lower frequencies where loop loss is less and crosstalk noise is less.

ADSL systems employ advanced digital transmission techniques to enhance performance. The modulation and frequency placement of the transmitted signal dynamically adapts to achieve optimal performance from the characteristics unique to the subscriber line being used. Trellis codes are used to reduce the effect of steady-state wide-band noise. Adaptive equalizers protect against narrowband noise such as radio frequency interference (RFI). Forward error control (FEC) codes and interleaving protect against noise impulses. Interleaving protects against error bursts by shuffling blocks of data so that a long burst of errors results in a few (correctable) errors in each block, rather than a large (uncorrectable) number of errors in one block. An interleaving depth of 20 ms will protect against noise bursts up to 500 μs in duration. Surveys of loop impulse events suggest that a vast majority of impulses are less than 500 μs in duration. However, this degree of interleaving causes an additional transport delay of 20 ms, which can slow down the throughput of protocols such as TCP/IP, which require packets to be acknowledged before more data are sent.

ADSL loops may have bridged taps, but no loading coils. An in-depth discussion of ADSL transmission techniques is provided in Chapter 6.

2.5.5 ADSL's Future

ADSL will be integrated into fiber-fed digital loop carrier (DLC) systems to address those loops that are not served directly from a CO. ADSL is well suited to providing high bit rates over the DLC-fed loops, which are rarely longer than 3.7 km (12 kft). Despite an industry standard for ADSL (ANSI T1.413), early ADSL systems did not interoperate. Efforts by ADSL manufacturers and standards committees are expected to achieve multivendor interoperability for future ADSL systems. In addition to the physical layer, full interoperability requires compatibility at all layers of the protocol stack.

It has become clear that ADSL is the access technology that asynchronous transfer mode (ATM) needed to open the door to home and small office. Before ADSL, ATM appeared to be restricted to only those who could afford the premium prices of links at 45 Mb/s and above: big business and the backbone network. Work is underway to deal with ATM transport over ADSL's unique characteristics: error rates, latency, asymmetry, and dynamic rate change.

For a time it appeared that the focus for ADSL evolution was high speeds such as 10 Mb/s downstream and 1.5 Mb/s upstream. However, this direction has faded due to overlap with VDSL, spectral compatibility concerns, and doubts regarding the need for these speeds. Instead, the focus now is on improved loop reach at more modest rates near 1 Mb/s, and lower-cost, lower-power, and reduced crosstalk implementations. ADSL systems are being developed to convey multiple digital derived voice circuits in addition to high-speed data.

2.5.5.1 ADSL + ISDN

Some vendors are introducing a version of ADSL where the upstream and downstream frequency band have been placed above the 0 to 80 kHz ANSI T1.601 basic rate ISDN transmission band. For BRI using the 4B3T line code, the BRI frequency band is 0 to 120 kHz. This substantially reduces the ADSL bit rates but does permit simultaneous ISDN and ADSL service on one loop. The ADSL + ISDN configuration is unlikely to provide the full 5.5 km (18 kft) reach normally provided by ISDN. ADSL + ISDN is of particular interest in Germany and France, where ISDN service is particularly widespread. This configuration may also be used to provide two voice circuits and moderately high data rates.

2.5.5.2 Splitterless ADSL

The customer premises installation of ADSL service can require new or modified inside telecommunications wiring. For a conventional ADSL configuration, the ADSL terminates at the network interface device (NID), where a low-pass filter (the *splitter*) extracts the voice-band signals that are attached to the red and green inside wires to the telephones, and the wideband signals are connected to the yellow and black inside wires to the customer's ADSL modem. This requires the installation of the splitter and also requires the use of the yellow and black inside wires, which are not found in some premises or may already be used for second-line voice service. Furthermore, in some cases, substandard inside wire has been used, which will impair ADSL (and even ISDN) operation. As a result, new inside wire will often be required from the NID to the customer's ADSL modem.

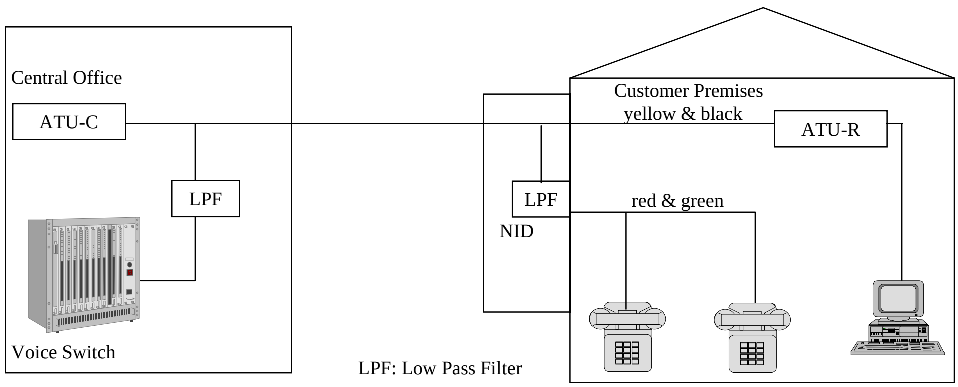

Figure 2.16 Conventional ADSL configuration with splitter.

The most common ADSL POTS splitter configuration (shown in Figure 2.16) places a low-pass filter (LPF) for the voice wiring at or near the NID and a high-pass filter (HPF) located within the ATU-R. Alternatively, the splitter (LPF and HPF) may be integrated within the ATU-R. The splitter within the ATU-R has the disadvantages of the possible loss of POTS service when the ATU-R is removed and possibly excessive crosstalk when using existing premises wiring.

The splitterless ADSL concept eliminates the splitter filter at the customer end of the line. Many other terms have been used to describe this concept: *ADSL Lite*, *Consumer DSL (CDSL)*, or *Universal ADSL (UADSL)*. Splitterless ADSL is defined in ITU Recommendation G.992.2. The ADSL modem and phones are all directly connected to the existing red and green wires within the premises. Simultaneous data and voice operation are supported. ADSL installation is easily performed by plugging the ADSL modem into any phone jack in the premises. Neither new inside wiring nor splitter installation is needed. See Figure 2.17.

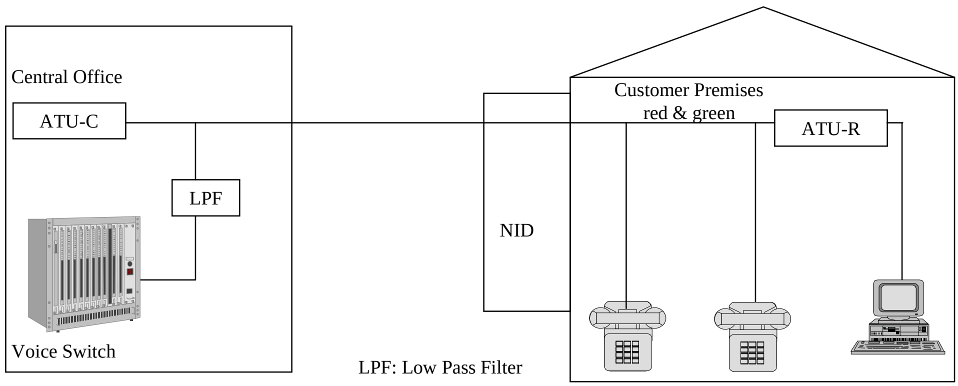

Figure 2.17 ADSL configuration without splitter at customer premises.

The concept is wonderful; the practical implementation is being investigated and will present compromises. The ADSL splitter has two purposes: (1) the splitter attenuates POTS signaling noise, which could corrupt ADSL data transmission, and (2) the splitter attenuates the ADSL signals to prevent audible noise on the telephones. Due to the nonlinear impedance of some telephones, ADSL-transmitted energy at frequencies well above the audible band can be modulated into the voice band. Shifting the ADSL transmission bands to higher frequencies can reduce these problems. However, this will reduce the ADSL data rates and the loop reach. Brief error bursts are likely to occur when the phone rings, and surely during the *ring trip* transient when a ringing phone is taken off-hook. An objectionable hissing sound may be heard in phones with poor wideband characteristics.

One solution to these problems is to place a low-pass filter in series with each phone. See Figure 2.18. This filter would be inexpensive and have modular connectors so that an untrained customer could install it in seconds. Installation for a wall phone would not be so easy, but at worst the customer could buy a new “ADSL-compatible” wall phone with a low-pass filter built inside the phone. This configuration should prevent POTS noise from impairing ADSL transmission and ADSL noise from being heard on the telephones. No new inside wire is needed, and no splitter installation is needed at the NID. The customers could plug their ADSL modem into any telephone wall jack. The ADSL data rate would be somewhat less than the conventional ADSL configuration. There would be some reduction in data rates due to other noise and the loading effect of many filters and wiring stubs. The ITU Recommendation G.992.2 (“G.lite”) has less performance than the full-rate Recommendation G.992.1 (“G.DMT”), due to a reduced number of DMT tones, and fewer bits-per-tone. Reduced voice-band transmission quality may also result from many low-pass filters being placed in parallel. Another concern is trouble resulting from customers who forget to place a LPF in the line to one of their phones.

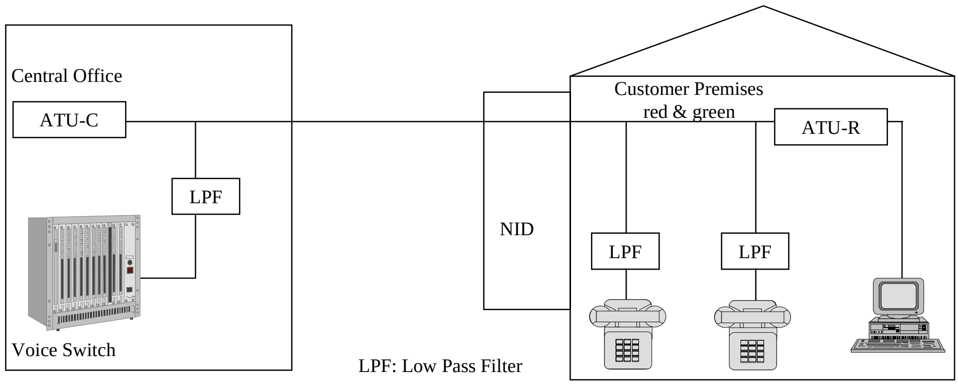

Figure 2.18 ADSL configuration with low-pass filter by each phone.

If the technical and operational hurdles are overcome, splitterless ADSL may ultimately become the dominant type of ADSL. For the near term, most ADSLs are being installed with a splitter at both ends of the line. The use of a splitter at the customer site may see continued use for the installation of service to customers who require higher-bit-rate service. Some ADSL service providers have suggested that their ADSL service should work with both the splitterless and splittered customer configurations, while using the same type of ATU-C at the CO. The progress of splitterless ADSL was accellerated by the technical and marketing activities of the Universal ADSL Working Group (UAWG), a collection of leading telephone companies and computer companies.

2.6 VDSL

Very-high-bit-rate digital subscriber line (VDSL) is an extension of ADSL technology to higher rates, up to 52 Mb/s downstream. At such high bit rates, the loops must be so short that optical fiber will be used for all but the last few thousand feet.

2.6.1 VDSL Definition and Reference Model

Most DSLs are primarily intended for use on loops from a CO to a customer premises and secondarily for use from fiber-fed distribution multiplexers. The opposite is true for VDSL. VDSL will primarily be used for loops fed from an optical network unit (ONU), which is typically located less than a kilometer from the customer. Few VDSL loops will be served directly from a CO.

Optical fiber connects the ONU to the CO. VDSL transmission over a twisted wire pair is used for the few thousand feet from the ONU to the customer premises. See Figure 2.19. VDSL requirements developed by the T1E1.4 standards working group specified the following objectives for rates and distances from the ONU to the customer site:

Downstream rate (Mb/s)	Upstream rate (Mb/s)	Distance (kft - m)
52	6.4	1,000 – 300
26	3.2	2,500 – 800
26	26	1,000 – 300
13	13	1,800 – 600
13	1.6	3,750 – 1,200

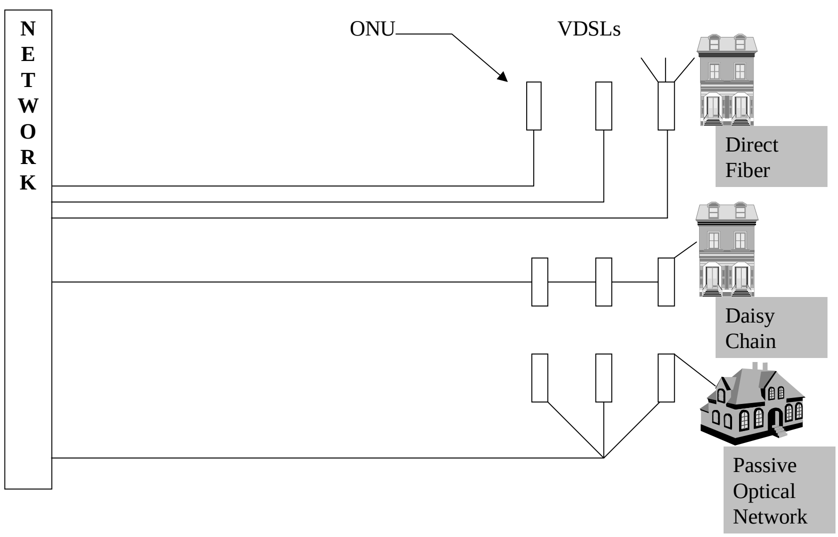

Figure 2.19 VDSL architectures.

The fiber from the network to the ONUs may be connected directly to the ONU, daisy chained, or via a passive optical splitter.

2.6.2 VDSL Origins

Discussion of the VDSL concept began in standards committees in late 1994, with the definition of VDSL system requirements in ETSI TM6 and T1E1.4. Several proposals are currently be studied by these groups.

2.6.3 VDSL Capabilities and Applications

VDSL, as part of a full-service network (FSN), is intended to support all applications simultaneously: voice, data, and video. Ultimately, VDSL would support high-definition television (HDTV) and high-performance computing applications. Symmetric application of VDSL will provide two-way data rates up to 26 Mb/s that will be attractive for business sites where fiber-to-the-building is not justified.

The DAVIC VDSL type specification employs carrierless amplitude phase (CAP) modulation for rates of 13, 25.92, and 51 Mb/s downstream and 1.6 Mb/s upstream via an unshielded twisted wire pair. The DAVIC VDSL specification is based on a passive NT architecture, which permits direct connection of multiple VDSL transceivers at the customer end of the line. See Figure 2.20. Typically, the passive NT architecture requires the ONU to be less than 100 meters from the customer VDSL units, thus making it more suitable for fiber-to-the-pedestal and in-premises applications.

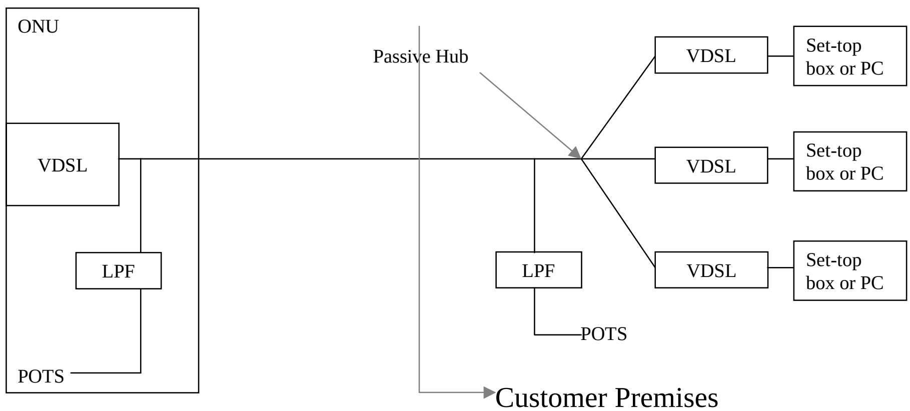

Figure 2.20 VDSL passive hub architecture.

The VDSL Active hub architecture shown in Figure 2.21 permits a greater (rate) × (reach) product by using a point-to-point configuration for loop transmission. The active hub consists of a single VDSL transceiver, and a separate short-haul link within the premises to each terminal (shown), or a short-haul bus within the premises (not shown).

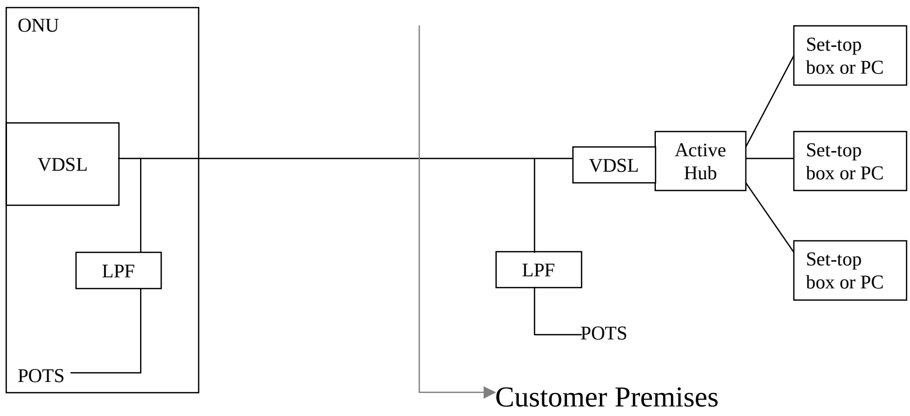

Figure 2.21 VDSL active hub architecture.