
ADSL Technology Explained

Autor:

Data de publicació: 20-09-2013

ADSL Technology Explained, Part 1: The Physical Layer

ADSL Technology Explained, Part 1: The Physical Layer

Louis Litwin, Michael Pugel, Rob Rhodes, and John Richardson

3/1/2001 04:35 PM EST

Post a comment

in Share

For residential and commercial users with an ongoing need for broadband data access, but who do not send out correspondingly large data streams, asymmetric digital subscriber line (ADSL) services work well. This service is so named because the data rate sent to the user (downstream) is much greater than the data rate sent from the user (upstream). This asymmetric model is based on typical Internet usage patterns.

For example, a user sends a Web page request (small amount of upstream data) and receives the HTML for the Web page with graphics and sound (large amount of downstream data). Various services, such as cable modems, satellite services, and DSL exist to provide such access.

An ADSL system uses existing telephone wire to allow bidirectional data communications between a user and the telephone company's central office (CO). Some other popular services, such as an ISDN line or a standard dial-up modem, also use the phone lines to communicate. However, those services prevent the simultaneous operation of standard analog phone service on the same phone line. An important advantage of ADSL is that it allows the plain old telephone system (POTS) signal to co-exist with the ADSL data signal.

We begin our tour of the ADSL system architecture with a look at the physical layer (PHY). Topics covered at the PHY will include ADSL's multicarrier modulation technique, common impairments, and phone-line characteristics.

Spectrum allocation

The ADSL PHY was designed so that it could peacefully co-exist with the standard POTS spectrum. The two services can co-exist because the ADSL spectrum only uses the frequencies above POTS. The POTS spectrum goes from near DC to approximately 4 kHz. A frequency guard band is placed between the POTS spectrum and the ADSL spectrum to help avoid interference. The ADSL spectrum starts above the POTS band and extends up to approximately 1.1 MHz. The lower part of the ADSL spectrum is for upstream transmission (from the customer to the CO) and the upper part of the spectrum is for downstream transmission. There are actually two different ways that the upstream and downstream spectra can be arranged (see Figure 1).

In a frequency division multiplexed (FDM) system, the upstream and downstream spectra use separate frequency ranges. They can vary for different implementations, but typically the upstream band is from 25 to 200 kHz and the downstream band is from 200 kHz to 1.1 MHz. Other divisions are also permitted within the ADSL standard. This system is free from the occurrence of a type of interference called self-crosstalk. One drawback, however, is that the downstream bandwidth is reduced in comparison to an echo-cancelled system.

An echo-cancelled system allows the downstream band to overlap with the upstream band. The upstream band still uses the frequencies from 25 to 200kHz, but the downstream band can now extend over the upstream band. The main advantage of this system is that it significantly extends the available downstream bandwidth. However, it does require echo-canceling circuitry due to the full-duplex transmission. In addition, the presence of self-crosstalk causes additional interference.

The DMT approach

The PHY of ADSL uses a multicarrier modulation technique known as discrete multitone (DMT). A DMT system transmits data on multiple subcarriers in a manner very similar to the orthogonal FDM (OFDM) technique that is used in many wireless applications. A DMT modulator takes in N data symbols in parallel and transmits the symbols on N subcarriers. The data rate on each subcarrier is $1/N$ the original data rate.

Reducing the data rate results in a DMT symbol period that is N times as long as the original symbol period. Increasing the symbol period can make the symbol longer than the time span of the channel. This situation can make it easier to combat the effects of intersymbol interference.

The DMT signal is formed by using an Inverse Fast Fourier Transform (IFFT) to generate orthogonal subcarriers at the transmitter. The data symbols at the transmitter are treated as being in the frequency domain and act as complex weights for the basis functions (orthogonal sinusoids at different frequencies) of the IFFT. The IFFT then converts the data symbols into a time-domain "sum of sinusoids" signal.

The block of IFFT output samples is known as a DMT symbol. This time-domain signal is transmitted across the channel, and an FFT is used at the receiver to bring the signal back into the frequency domain. A block diagram of a typical ADSL transmitter/receiver pair is shown in Figure 2.

A $2N$ -point IFFT is used to generate the DMT symbol, and the N negative-frequency IFFT bins are the complex conjugate of the N positive-frequency bins. This symmetric spectrum results in a real time-domain signal. The DMT signal is centered at DC with the subcarriers around DC zeroed out (not used) to create a hole in the DMT spectrum in order to make room for the POTS spectrum. DMT is thus a true baseband system.

DMT supports inclusion of a cyclic prefix. A cyclic prefix is a block of samples with a length, LP , that is a replica of the last LP samples of the DMT symbol. The prefix is then transmitted first, followed by the $2N$ samples of the DMT symbol. The length LP is chosen such that it will be longer than the length of the channel response. The cyclic prefix contains redundant information. However, the DMT receiver exploits the presence of the prefix in order to mitigate the effects of the channel. The use of the cyclic prefix will be described in further detail in the Impairments section.

The dynamic bit allocation technique allows DMT to make efficient use of the available channel capacity. This technique enables the system to vary the number of bits per symbol for each subcarrier based on the subcarrier's signal-to-noise ratio (SNR). Subcarriers with a low SNR transmit binary phase-shift keying (BPSK) or quadrature PSK (QPSK) because they are robust modulation formats. If the subcarrier's SNR is very low, that subcarrier will not be used to transmit data at all. Subcarriers with a higher SNR transmit higher-order quadrature amplitude modulation (QAM) in order to achieve an increased throughput.

Impairments

There are several significant types of impairments encountered in an ADSL system: additive white Gaussian noise (AWGN), crosstalk, impulse noise, bridged taps, and radio noise.

AWGN is the thermal noise that is common to all communication systems. In a digital system such as ADSL, AWGN can cause symbol errors to occur at the receiver when noise pushes the received sample beyond a decision boundary. Like many other digital communication systems, ADSL employs error-control coding to help mitigate the effect of AWGN. Coding adds redundancy to the transmitted signal and exploits the redundancy at the receiver to detect and correct errors.

ADSL uses three layers of coding. The innermost code in the PHY is a convolutional code. These codes get their name because the encoding process can be viewed as the convolution of the message with the code's impulse response. The Viterbi algorithm is used at the receiver to decode the received sequence.

Convolutional codes

Convolutional codes are good at correcting random errors. However, the nature of the decoding algorithm is such that the decoder can cause burst errors to occur if errors are made during the decoding process.

A Reed-Solomon block code is used on top of the convolutional code. Reed-Solomon codes are powerful codes that are good at detecting and correcting burst errors, such as those generated by the Viterbi decoder. The ADSL specification allows Reed-Solomon code-word lengths of up to 255 bytes with the addition of up to 16 parity bytes for each code word. The outermost code is a cyclic redundancy check (CRC) code. The CRC can detect errors, but it cannot correct them. The CRC code is used as a top-level error-detection mechanism in order to detect any errors that remain after Viterbi and Reed-Solomon decoding.

Because bundled telephone cable contains many wires for many different users, crosstalk is a common impairment. These wires radiate electromagnetically and can induce currents in other wires in the cable. This interference effect is known as crosstalk. There are two basic types of crosstalk and they both appear at the receiver as additive noise.

Near-end crosstalk (NEXT) occurs when a transmitter interferes with a receiver located on the same end of the cable. Far-end crosstalk (FEXT) occurs when the transmitter interferes with a receiver on the opposite end of the cable. The effect of NEXT is more severe than FEXT since the FEXT interference travels the entire length of the cable and is attenuated by the time it reaches the receiver.

Crosstalk can be further subdivided into self-crosstalk and foreign crosstalk. Self-crosstalk is interference from another ADSL system using the same spectrum allocation. Foreign crosstalk is interference from an ADSL system using a different spectrum allocation or from a completely different type of system (such as ISDN). One way to reduce the effects of crosstalk is with spectrum allocation.

In echo-cancelled ADSL, the upstream and downstream channels overlap (see Figure 1). Since the same frequency band is being used for transmission and reception, the system will suffer from self- and foreign crosstalk. However, in FDM ADSL the upstream and downstream channels use separate frequency bands. This system will not suffer from self-crosstalk, although foreign crosstalk will still be present.

Interference that is short in duration but of a large magnitude is known as impulse noise. Impulse noise can be caused by lightning or by a motor turning on and creating a power surge. ADSL systems use a combination of interleaving and coding to correct the errors caused by impulse noise. The interleaving process rearranges data so that those samples that were located contiguously in time are spaced far apart. Impulse noise can cause a burst of errors that is hard for the decoders to correct. The use of interleaving combined with coding spreads out these errors in time to improve decoding performance.

An additional impairment

Bridged taps are an additional impairment type found in ADSL systems. A bridged tap is a section of wire connected to the loop at one end and unterminated at the other end. Examples of bridged taps are unterminated wires that are laid out in areas where housing is still being built. When a transmitted signal arrives at a bridged tap, the signal divides. Part of the energy continues to the receiver and the rest of the energy reflects off of the unterminated end. This reflection causes delayed versions of the signal to arrive at the receiver, and these reflections distort the received signal. The effect is very similar to the interference that occurs with a multipath channel in a wireless communication system.

Bridged taps cause two problems. The first is intersymbol interference where a received DMT symbol is distorted due to delayed versions of the previous DMT symbol. The effect of intersymbol interference is removed by discarding the cyclic prefix at the receiver. The cyclic prefix contains redundant information, and so it is not needed at the receiver. The length of the prefix is chosen such that the delayed versions of the previous symbol only distort the cyclic prefix and not the actual data part of the DMT symbol. The ability to remove the effect of intersymbol interference by discarding the cyclic prefix is one of the advantages of having a long symbol period.

The second problem is the intrasymbol interference that is caused when delayed versions of a DMT symbol cause the symbol to interfere with itself. The effect of intrasymbol interference in the frequency domain is a shaping of the received-signal spectrum. The received spectrum is essentially a multiplication of the transmitted spectrum and the channel's frequency response. The distortion due to intrasymbol interference is removed by using a frequency-domain equalizer.

And finally, radio noise is interference due to a wireless source. The copper phone lines act as antennae and pick up this interference. The most common source of radio noise comes from AM radio since its spectrum overlaps with the ADSL spectrum. Coding can help correct the errors caused by radio noise, and adaptive RF cancellation filters can also be used. The previously mentioned technique of dynamic bit allocation can be used to turn off subcarriers near

frequencies of interference.

ADSL in the physical plant

ADSL operates within the existing POTS plant structure. The phone company has been in existence for nearly 100 years and despite modernization, the structure today remains much as it did years ago. There are more than 700 million phone lines in the world with two-thirds in the US alone. Figure 3 illustrates the basic structure that is utilized in many US phone-service deployments.

Mainline routes shown on the left side of Figure 3 are connected by fiber between toll offices (TOs). These TOs provide the interface between local exchanges and long-haul fiber runs for long-distance service. The TOs also connect COs that control local exchange service. The CO houses the main switching equipment for the home and also serves as the location for the head-end DSL equipment.

Multiple interconnects are often made between several COs and TOs for operational reliability. From the CO, the network is further subdivided into customer service areas (CSAs) that are serviced by remote terminals (RTs). This region represents the last several miles before the home, such as a large neighborhood or small city.

The CSA is then further subdivided into a distribution area (DA) that is initiated at a feeder distribution interface (FDI) serving up to 500 phone lines. The FDI represents the last point where the phone lines are still bundled together. The phone wires take separate routes to each home and typically cover the last mile of service routing.

Standard copper twisted pair is the primary medium for routing between the CO and RT, and then on to the home residence. Fiber deployment through these regions is technically possible but economically unjustifiable, especially for points past the FDI. Fiber routing will eventually be used for routes between the CO and RT. This upgrade will support improvements in ADSL service as well.

The ability to supply phone service over a very long distance of wire, often called reach, is limited by the ability of the phone switching equipment at the CO to function correctly. The network switch can operate with a maximum DC load resistance of approximately 1,500 . Most of this resistance is found in the copper lines running between the CO and home.

The distance at which this resistance is reached is known as the revised resistance design (RRD) distance. Voice quality over the phone often suffers prior to reaching the RRD distance, which can extend to more than 3 miles. The phone company employs tricks to improve quality, including thicker wire and load coils. ADSL performance is also a function of the distance from the CO and service is often extremely limited or impossible on these long reaches.

Plain old copper

The copper pair wire that runs between the CO and the home residence has the greatest impact on ADSL system

performance. The wire provides simultaneous signal routing of bidirectional voice communications, which will occupy the spectrum up to approximately 4 kHz. Above this frequency, the ADSL signals are inserted. The wire is best modeled as a lossy transmission line. Most plants employ one of four copper wire gauges ranging from 19 to 26 gauge.

Thinner wire (higher gauge number) is typically used close to the CO to allow smaller bundles to be formed. Thicker wire is used close to the residence to extend the reach from the CO to home while keeping within the resistance limits required for proper voice switching.

The transmission-line model contains a shunt capacitance in parallel with a shunt resistance (due to dielectric losses) along with a series resistance and series inductance. The cable specifications dictate that the capacitance (per mile) is constant for all wire gauges. As a result, the nominal line impedance is typically around 120 Ω . The series resistance changes by approximately 20% over the possible wire gauges and increases logarithmically with frequency. Typical loss through the wire is 3 to 6 dB per mile in the voice band, depending on wire gauge, and also increases logarithmically with frequency.

Most system impairments are created by, or enter the system through, the copper wire. The wire, through its lossy nature, accounts for degradation in SNR. Reflections are created in the system due to unterminated lines such as bridged taps and wire splices. These reflections create self-interference. Reflections cause energy cancellation at certain frequencies as discussed in the Impairments section. The wire is typically unshielded from the FDI (point of fanout) to the home location and is susceptible to ingress (pickup of external signals in the frequency range of interest).

Part two of this article will complete the discussion of PHY equipment and move on to discuss the link layer. Other elements of the PHY plant will be introduced with a focus on the equipment at the home. Link layer topics include ADSL modem initialization and packet framing. ATM and IP will be covered in the network-layer section. Finally, the application layer will address some new services such as voice over DSL (VoDSL) and video streaming.

The main physical (PHY) layer structure of asymmetric digital subscriber line (ADSL) technology, as explored in part one of this article, involves a modulation scheme known as discrete multitone (DMT). In this second part, we examine other PHY layer components, move on to the handshake and initialization process of ADSL, and finally end with a look at the network and application layer protocols.

The phone line interface in ADSL is an analog circuit, but the signal processing for a DMT signal is done digitally, so an interface is needed between these two circuits. An analog-to-digital converter (ADC) is connected in the receive path of the hybrid to interface downstream signals at the home, while a digital-to-analog converter (DAC) is used to connect into the transmit path for upstream data.

Multicarrier systems contain many simultaneous signals. If each signal's peak amplitude is represented by x , and all signals simultaneously reach peak signal level, the resulting level would be $x \cdot 20 \log N$, where N is the number of signals. The signals in a DMT system have a statistical nature - they can be considered uncorrelated random processes (that is, their cross-correlation is equal to zero). The possible peak amplitude may be large, but the probability of this level occurring is low.

The peak-to-average ratio (PAR) is used to define the ratio between a signal's peak level and its average level over time. Most multicarrier systems use a modified definition for PAR that is based on the statistical likelihood of exceeding

a certain peak level (such as the probability of clipping in the DAC output). For 256 subcarriers and a clipping probability of 10^{-7} , the PAR value is around 5 (14 dB). The PAR value partially determines the operating parameters for the ADC and DAC.

The important parameters associated with the design of the converter for ADSL include the PAR factor mentioned above, the number of bits per subcarrier, and the required signal-to-noise ratio (SNR). Typical DACs use 10-b resolution, a level considered acceptable for up to 8 b per subcarrier.

The ADC must take into account all of these parameters, plus additional bits of resolution for input noise (receiver SNR is lower than transmitter SNR) and for echo energy. Typically one to two extra bits are employed in the ADC.

Other ADSL components

The POTS splitter is a passive three-terminal circuit which is key to the coexistence of the existing phone service and ADSL service. The splitter contains a common terminal (from the phone plant end of the service), a low-pass filter to the POTS side, and a high-pass filter to the ADSL side. The filters are designed to diplex the signals onto the outside phone line and have stopband impedance characteristics that minimize the effect of changing from "on hook" to "off hook" condition.

The residence modem also contains several other blocks. A hybrid simultaneously connects the transmitter and receiver to the same copper wire. The hybrid contains the line conditioning and gain control for proper operation over a range of signal levels.

The digital circuits connected to the DAC and ADC converters contain the signal processing and memory necessary to perform the demodulation and data conversion to get the information from the phone wire to the home device.

The equipment in the central office (CO) is similar to the equipment used in the home except it is structured as a bank of modems. These modems, one for each home in use, along with some network and phone interfacing equipment, comprise the device known as a digital subscriber loop access multiplexer (DSLAM).

Phone systems employ tricks to extend the usable range of their system. One such trick is to employ load coils in loops that potentially extend far from the CO (more than 3 miles or 16 kft). The copper wire will attenuate frequencies in the voice band at these distances, degrading operation. However, the revised resistance design (RRD) distance may not have been reached yet, so switching operation is still possible.

In order to maintain voice service operation, load coils are added in series with the line at periodic distance intervals. These coils are used to compensate for the effect of the cable capacitance through the voice band region on the line. As a result, the frequency response above the voice band (ADSL spectrum) rolls off at an accelerated rate. Load coils typically need to be removed from phone lines that will carry ADSL services. In some cases the plant must be re-engineered to compensate for the missing load coils.

ADSL has evolved to the point of receiving attention at a standards level. In 1998, agreement was reached on a set of standards for ADSL. G.992.1 is part of a suite of standards (the G.99x.x series) covering several DSL systems as well as protocols and tests. Key PHY layer specifications are outlined in Table 1.

Table 1: G.992.1 PHY Layer Parameters

Downstream

Overall symbol rate
4 kHz

Number of carriers per DMT
symbol 256

Subcarrier spacing
4.3125 kHz

Cyclic prefix length
32 samples

Operational modes
FDM or echo cancelled

FDM mode frequency range
64 to 1100 kHz

Echo cancelled mode frequency range
13 to 1100 kHz

Number of bits assigned per subcarrier
0 to 15 (no bits assigned to 64k QAM)

Synchronization
Pilot tone at subcarrier 64, $f = 276$ kHz

Upstream

Number of subcarriers per DMT symbol
32

Cyclic prefix length
4 samples

FDM mode frequency range
11 to 43 kHz

Echo cancelled mode frequency range
11 to 275 kHz

Synchronization Pilot tone at subcarrier
16, $f = 69$ kHz

Handshake/initialization
Per G.994.1

The PHY layer specification outlines power levels and spectral masks that must be maintained.

The lower three to six subcarriers are set to a gain of "0" (turned off) to permit operation of POTS, with the addition of a splitter at the home entry point of the phone line.

It is important to note that these standards encompass only a framework for operation. Individual networks and providers are free to adapt their system within that framework. The standards provide a set of boundaries for equipment manufacturers.

In the beginning

When an ADSL termination unit-residence (ATU-R), or modem, is first connected to a DSL network, it goes through a fairly extensive initialization process. This process identifies and qualifies both the capabilities of the network equipment and of the underlying physical infrastructure.

The initialization process consists of four major phases. The first phase is a handshake using the G.994.1 or G.hs protocol. (G.hs is a precursor to the G.992.1 specific initialization and is used by other DSL and telecommunication devices.). The remaining three phases - transceiver training, channel analysis, and exchange - are covered directly in the G.992 standard and apply specifically to standards-based ADSL networks.

The G.994.1 handshake is used to determine the nature and capabilities of the endpoints (such as an ADSL modem) and to indicate which protocol will be used for the remainder of the initialization. The signaling method used for the handshake interchange is designed to be robust to address channel characteristics that could be atrocious. Biphase shift keying (BPSK) modulation is used to modulate multiple single-tone subcarriers, all carrying the same data.

The subcarriers used are selected based on the typical impairments likely to be present in a given global region. The handshake has several possible variants, but, fundamentally, the two endpoints exchange a message which contains information about the endpoint type, and a number of related subparameters such as the frequency range and number of DMT subcarriers supported.

The second phase of initialization is transceiver training. Receivers at each end of the line acquire the DMT symbol stream, adjust receiver gain, perform symbol timing recovery, and train any equalizers. There is an optional echo cancellation training step that can also be performed during this phase, but the specification does not define the training signal to be used.

Characterization

The transmitter power at each end of the line is set to a predetermined level at this phase, allowing a preliminary estimate of loop attenuation by the receivers. The received upstream power level is reported back to the ATU-R transmitter to allow limited power level adjustment (attenuation), if needed, to meet spectral mask requirements. The training phase is conducted with all available upstream and downstream subcarriers modulated, using two of the four constellation points of a QPSK constellation.

In the third phase, the transceivers exchange capability information and perform detailed channel characterization. For example, the ADSL termination unit-CO (ATU-C) specifies the minimum SNR margin for the system and whether it can support functions such as trellis coding and echo cancellation. Similar information is exchanged about the ATU-R. Although some of these same parameters were exchanged during the G.994 handshake, the handshake message parameters are used to gather information only and are not necessarily used for the connection.

During this third phase, both transceivers attempt to measure specific channel characteristics such as unusable subcarriers, loop attenuation on a per subcarrier basis, SNRs, and any other channel impairments that would affect the potential transmitted bit rates. Based on the discovered channel characteristics, the ATU-C makes the first offer of the overall bit rates and coding overhead that will be used for the connection.

Four possible rates are offered, in decreasing order of preference. In the current release of the ADSL standard, the ATU-C completely controls the final bit rate. All subcarriers are modulated simultaneously with the same information. The primary tool for channel measurement is a pseudo-random bit sequence.

Setting the rates

The last phase of the initialization sets the final overall transmission rates in both the upstream and downstream directions for the connection. These final rates are determined based on calculated channel parameters measured during the channel analysis phase, and are not necessarily the same as the preliminary rates offered during that phase.

As the ATU-C controls data rates, if the ATU-R cannot support any of the offered rates, both terminals will return to the

beginning of the initialization process. Otherwise the ATU-R responds with the rate it can support.

Since ADSL uses multiple orthogonal subcarriers, each subcarrier can be assigned a modulation format (number of bits per subcarrier) and relative gain independently. The ATU-C assigns bits and gains for the downstream direction while the ATU-R assigns the upstream parameters.

The resulting assignment maximizes the amount of traffic that can be carried over the PHY layer. The last part of the exchange phase is a synchronized transition from the highly robust BPSK and QPSK modulations used during the initialization to the full traffic rate modulations (such as higher-order QAM) assigned during the exchange phase. At the conclusion of the initialization steps, the system is ready to pass higher-layer traffic.

ADSL and ATM

Along with TCP/IP, ATM, with its key layers and its relationship with ADSL, is one of the key enabling technologies for broadband applications using ADSL.

ADSL uses a framed transport structure in which frames are encoded and modulated into DMT symbols. The ADSL frames can be further grouped into superframes, each consisting of 68 frames. An ADSL frame is sent every 250 microseconds, therefore a superframe is sent every 17 ms.

In full-rate ADSL, a frame can be broken down further into two parts, each being 125 microseconds. These two parts can be classified as the fast data path and the interleaved data path. The fast data path could have a higher bit error rate (BER) because the interleaver is not used to mitigate the effects of impulse noise. However, the removal of the interleaver significantly reduces the latency of this data path (which is well suited for time-sensitive information such as interactive audio and video).

ATM, which allows data to be sent asynchronously, uses cells consisting of 53 B of information. The cells' small size allows the efficient multiplexing of data from multiple sources. ATM is connection-oriented - once a connection is established, the connection carries traffic that meets the quality of service (QoS) requirements requested by the source and destination.

The ATM and ATM adaptation layer (AAL) are the two most important of the several layers in the ATM protocol stack. The ATM layer is responsible for the definition of logical connections through the network. Logical connections in ATM are known as virtual circuits (VCs).

VCs represent fundamental ways of switching in an ATM network; a VC is established between two end users on the network. Bundles of VCs are called virtual paths (VPs) and they share the same end-point. A single VP carries the cells from multiple VCs, and the cells are subsequently switched together.

The AAL is responsible for inserting higher-layer information into cells to be transported over the network. It consists of

two sublayers called the segmentation and re-assembly (SAR) sublayer and the common part convergence sublayer (CPCS). The SAR sublayer segments the upper-layer protocol data units (PDUs) into 48-B cell payloads (SAR PDUs). The SAR PDUs are then passed to the ATM layer to form a complete cell.

The CPCS is responsible for performing functions for different classes of service. These are referred to as AAL1-5. A summary of the types of services is shown in Figure 1.

AAL5 is most often used for connectionless Internet traffic, as it allows the entire 48 B of cell payload to transport data with minimal overhead. A 10% overhead is typical when transporting Internet traffic over ATM.

Figure 2 illustrates a basic ADSL network architecture representing the connection between the service provider and the customer. An ATM connection is set up between the ADSL modem and a termination point in the back-end network. This connection is referred to as a permanent virtual circuit (PVC).

The deployment model for ADSL is based on the current dial-up system, which uses point-to-point protocol (PPP) to support network services such as authentication and client addressing. The ADSL Forum recommends this PPP-over-ATM-over-ADSL model, and it has become the standardized method for accessing data networks over ADSL. For the ever-popular encapsulation method of transmission, the Internet Engineering Task Force (IETF) has defined a method called RFC-2364 for PPP encapsulation over AAL5.

A DSLAM is used in all instances of ADSL deployments to aggregate traffic. With this device, the data from the ADSL modems is statistically multiplexed onto a common upstream link that interfaces to an ATM network. The ATM switch then routes the cells to their destination based on the cell header information. This destination is an IP router that reassembles the data cells into packets for transmission across the Internet.

Advanced services

ADSL is primarily deployed for Internet connectivity. Recently, other services have become available. For example, ADSL allows the service provider to offer multiple voice lines over a single copper pair. Instead of the voice being carried by POTS, it is now packetized directly in ATM and carried over the ADSL link. The voice packets are routed to a point in the service provider's network that interfaces through a gateway into the POTS network.

ATM's QoS capabilities allow the network to deliver acceptable voice-quality services. High-quality audio and video streaming are other potential services that can be offered using ADSL.

The demand for bandwidth and high-value content and services is soaring in both the home and office, and ADSL provides a robust and cost-effective mechanism to meet the demand. Existing infrastructures can often be reused, providing for fast and economical deployment.

Because ADSL was designed to be consumer-friendly, it can coexist with POTS. By using existing voice circuits instead

of a shared broadcast medium, ADSL can better individualize services. The roll-out of ADSL will continue to increase as business models are created to exploit the capabilities, both business and technical, that are available.

ADSL Technology Explained, Part 2: Getting to the Application Layer
ADSL Technology Explained, Part 2: Getting to the Application Layer

Louis Litwin, Michael Pugel, Rob Rhodes, and John Richardson

4/2/2001 12:11 PM EDT

[Post a comment](#)

NO RATINGS

[Login to Rate](#)

The main physical (PHY) layer structure of asymmetric digital subscriber line (ADSL) technology, as explored in part one of this article, involves a modulation scheme known as discrete multitone (DMT). In this second part, we examine other PHY layer components, move on to the handshake and initialization process of ADSL, and finally end with a look at the network and application layer protocols.

The phone line interface in ADSL is an analog circuit, but the signal processing for a DMT signal is done digitally, so an interface is needed between these two circuits. An analog-to-digital converter (ADC) is connected in the receive path of the hybrid to interface downstream signals at the home, while a digital-to-analog converter (DAC) is used to connect into the transmit path for upstream data.

Multicarrier systems contain many simultaneous signals. If each signal's peak amplitude is represented by x , and all signals simultaneously reach peak signal level, the resulting level would be $x \cdot 20 \log N$, where N is the number of signals. The signals in a DMT system have a statistical nature - they can be considered uncorrelated random processes (that is, their cross-correlation is equal to zero). The possible peak amplitude may be large, but the probability of this level occurring is low.

The peak-to-average ratio (PAR) is used to define the ratio between a signal's peak level and its average level over time. Most multicarrier systems use a modified definition for PAR that is based on the statistical likelihood of exceeding a certain peak level (such as the probability of clipping in the DAC output). For 256 subcarriers and a clipping probability of 10^{-7} , the PAR value is around 5 (14 dB). The PAR value partially determines the operating parameters for the ADC and DAC.

The important parameters associated with the design of the converter for ADSL include the PAR factor mentioned above, the number of bits per subcarrier, and the required signal-to-noise ratio (SNR). Typical DACs use 10-b resolution, a level considered acceptable for up to 8 b per subcarrier.

The ADC must take into account all of these parameters, plus additional bits of resolution for input noise (receiver SNR is lower than transmitter SNR) and for echo energy. Typically one to two extra bits are employed in the ADC.

Other ADSL components

The POTS splitter is a passive three-terminal circuit which is key to the coexistence of the existing phone service and ADSL service. The splitter contains a common terminal (from the phone plant end of the service), a low-pass filter to the POTS side, and a high-pass filter to the ADSL side. The filters are designed to diplex the signals onto the outside phone line and have stopband impedance characteristics that minimize the effect of changing from "on hook" to "off hook" condition.

The residence modem also contains several other blocks. A hybrid simultaneously connects the transmitter and receiver to the same copper wire. The hybrid contains the line conditioning and gain control for proper operation over a range of signal levels.

The digital circuits connected to the DAC and ADC converters contain the signal processing and memory necessary to perform the demodulation and data conversion to get the information from the phone wire to the home device.

The equipment in the central office (CO) is similar to the equipment used in the home except it is structured as a bank of modems. These modems, one for each home in use, along with some network and phone interfacing equipment, comprise the device known as a digital subscriber loop access multiplexer (DSLAM).

Phone systems employ tricks to extend the usable range of their system. One such trick is to employ load coils in loops that potentially extend far from the CO (more than 3 miles or 16 kft). The copper wire will attenuate frequencies in the voice band at these distances, degrading operation. However, the revised resistance design (RRD) distance may not have been reached yet, so switching operation is still possible.

In order to maintain voice service operation, load coils are added in series with the line at periodic distance intervals. These coils are used to compensate for the effect of the cable capacitance through the voice band region on the line. As a result, the frequency response above the voice band (ADSL spectrum) rolls off at an accelerated rate. Load coils typically need to be removed from phone lines that will carry ADSL services. In some cases the plant must be re-engineered to compensate for the missing load coils.

ADSL has evolved to the point of receiving attention at a standards level. In 1998, agreement was reached on a set of standards for ADSL. G.992.1 is part of a suite of standards (the G.99x.x series) covering several DSL systems as well as protocols and tests. Key PHY layer specifications are outlined in Table 1

Table 1: G.992.1 PHY Layer Parameters

Downstream
Overall symbol rate
4 kHz
Number of carriers per DMT
symbol256
Subcarrier spacing
4.3125 kHz

Cyclic prefix length
32 samples

Operational modes
FDM or echo cancelled

FDM mode frequency range
64 to 1100 kHz

Echo cancelled mode frequency range
13 to 1100 kHz

Number of bits assigned per subcarrier
0 to 15 (no bits assigned to 64k QAM)

Synchronization
Pilot tone at subcarrier 64, $f = 276$ kHz

Upstream

Number of subcarriers per DMT symbol
32

Cyclic prefix length
4 samples

FDM mode frequency range
11 to 43 kHz

Echo cancelled mode frequency range
11 to 275 kHz

Synchronization Pilot tone at subcarrier
16, $f = 69$ kHz

Handshake/initialization
Per G.994.1

The PHY layer specification outlines power levels and spectral masks that must be maintained.

The lower three to six subcarriers are set to a gain of "0" (turned off) to permit operation of POTS, with the addition of a splitter at the home entry point of the phone line.

It is important to note that these standards encompass only a framework for operation. Individual networks and providers are free to adapt their system within that framework. The standards provide a set of boundaries for equipment manufacturers.

In the beginning

When an ADSL termination unit-residence (ATU-R), or modem, is first connected to a DSL network, it goes through a fairly extensive initialization process. This process identifies and qualifies both the capabilities of the network equipment and of the underlying physical infrastructure.

The initialization process consists of four major phases. The first phase is a handshake using the G.994.1 or G.hs protocol. (G.hs is a precursor to the G.992.1 specific initialization and is used by other DSL and telecommunication devices.). The remaining three phases - transceiver training, channel analysis, and exchange - are covered directly in the G.992 standard and apply specifically to standards-based ADSL networks.

The G.994.1 handshake is used to determine the nature and capabilities of the endpoints (such as an ADSL modem) and to indicate which protocol will be used for the remainder of the initialization. The signaling method used for the handshake interchange is designed to be robust to address channel characteristics that could be atrocious. Biphase shift keying (BPSK) modulation is used to modulate multiple single-tone subcarriers, all carrying the same data.

The subcarriers used are selected based on the typical impairments likely to be present in a given global region. The handshake has several possible variants, but, fundamentally, the two endpoints exchange a message which contains information about the endpoint type, and a number of related subparameters such as the frequency range and number of DMT subcarriers supported.

The second phase of initialization is transceiver training. Receivers at each end of the line acquire the DMT symbol stream, adjust receiver gain, perform symbol timing recovery, and train any equalizers. There is an optional echo cancellation training step that can also be performed during this phase, but the specification does not define the training signal to be used.

Characterization

The transmitter power at each end of the line is set to a predetermined level at this phase, allowing a preliminary estimate of loop attenuation by the receivers. The received upstream power level is reported back to the ATU-R transmitter to allow limited power level adjustment (attenuation), if needed, to meet spectral mask requirements. The training phase is conducted with all available upstream and downstream subcarriers modulated, using two of the four constellation points of a QPSK constellation.

In the third phase, the transceivers exchange capability information and perform detailed channel characterization. For example, the ADSL termination unit-CO (ATU-C) specifies the minimum SNR margin for the system and whether it can support functions such as trellis coding and echo cancellation. Similar information is exchanged about the ATU-R. Although some of these same parameters were exchanged during the G.994 handshake, the handshake message parameters are used to gather information only and are not necessarily used for the connection.

During this third phase, both transceivers attempt to measure specific channel characteristics such as unusable subcarriers, loop attenuation on a per subcarrier basis, SNRs, and any other channel impairments that would affect the potential transmitted bit rates. Based on the discovered channel characteristics, the ATU-C makes the first offer of the overall bit rates and coding overhead that will be used for the connection.

Four possible rates are offered, in decreasing order of preference. In the current release of the ADSL standard, the ATU-C completely controls the final bit rate. All subcarriers are modulated simultaneously with the same information. The primary tool for channel measurement is a pseudo-random bit sequence.

Setting the rates

The last phase of the initialization sets the final overall transmission rates in both the upstream and downstream directions for the connection. These final rates are determined based on calculated channel parameters measured during the channel analysis phase, and are not necessarily the same as the preliminary rates offered during that phase.

As the ATU-C controls data rates, if the ATU-R cannot support any of the offered rates, both terminals will return to the beginning of the initialization process. Otherwise the ATU-R responds with the rate it can support.

Since ADSL uses multiple orthogonal subcarriers, each subcarrier can be assigned a modulation format (number of bits per subcarrier) and relative gain independently. The ATU-C assigns bits and gains for the downstream direction while the ATU-R assigns the upstream parameters.

The resulting assignment maximizes the amount of traffic that can be carried over the PHY layer. The last part of the exchange phase is a synchronized transition from the highly robust BPSK and QPSK modulations used during the initialization to the full traffic rate modulations (such as higher-order QAM) assigned during the exchange phase. At the conclusion of the initialization steps, the system is ready to pass higher-layer traffic.

ADSL and ATM

Along with TCP/IP, ATM, with its key layers and its relationship with ADSL, is one of the key enabling technologies for broadband applications using ADSL.

ADSL uses a framed transport structure in which frames are encoded and modulated into DMT symbols. The ADSL frames can be further grouped into superframes, each consisting of 68 frames. An ADSL frame is sent every 250 microseconds, therefore a superframe is sent every 17 ms.

In full-rate ADSL, a frame can be broken down further into two parts, each being 125 microseconds. These two parts can be classified as the fast data path and the interleaved data path. The fast data path could have a higher bit error rate (BER) because the interleaver is not used to mitigate the effects of impulse noise. However, the removal of the interleaver significantly reduces the latency of this data path (which is well suited for time-sensitive information such as interactive audio and video).

ATM, which allows data to be sent asynchronously, uses cells consisting of 53 B of information. The cells' small size allows the efficient multiplexing of data from multiple sources. ATM is connection-oriented - once a connection is established, the connection carries traffic that meets the quality of service (QoS) requirements requested by the source and destination.

The ATM and ATM adaptation layer (AAL) are the two most important of the several layers in the ATM protocol stack. The ATM layer is responsible for the definition of logical connections through the network. Logical connections in ATM are known as virtual circuits (VCs).

VCs represent fundamental ways of switching in an ATM network; a VC is established between two end users on the network. Bundles of VCs are called virtual paths (VPs) and they share the same end-point. A single VP carries the cells from multiple VCs, and the cells are subsequently switched together.

The AAL is responsible for inserting higher-layer information into cells to be transported over the network. It consists of two sublayers called the segmentation and re-assembly (SAR) sublayer and the common part convergence sublayer (CPCS). The SAR sublayer segments the upper-layer protocol data units (PDUs) into 48-B cell payloads (SAR PDUs). The SAR PDUs are then passed to the ATM layer to form a complete cell.

The CPCS is responsible for performing functions for different classes of service. These are referred to as AAL1-5. A summary of the types of services is shown in Figure 1.

AAL5 is most often used for connectionless Internet traffic, as it allows the entire 48 B of cell payload to transport data with minimal overhead. A 10% overhead is typical when transporting Internet traffic over ATM.

Figure 2 illustrates a basic ADSL network architecture representing the connection between the service provider and the customer. An ATM connection is set up between the ADSL modem and a termination point in the back-end network. This connection is referred to as a permanent virtual circuit (PVC).

The deployment model for ADSL is based on the current dial-up system, which uses point-to-point protocol (PPP) to support network services such as authentication and client addressing. The ADSL Forum recommends this PPP-over-ATM-over-ADSL model, and it has become the standardized method for accessing data networks over ADSL. For the ever-popular encapsulation method of transmission, the Internet Engineering Task Force (IETF) has defined a method

A DSLAM is used in all instances of ADSL deployments to aggregate traffic. With this device, the data from the ADSL modems is statistically multiplexed onto a common upstream link that interfaces to an ATM network. The ATM switch then routes the cells to their destination based on the cell header information. This destination is an IP router that reassembles the data cells into packets for transmission across the Internet.

Advanced services

ADSL is primarily deployed for Internet connectivity. Recently, other services have become available. For example, ADSL allows the service provider to offer multiple voice lines over a single copper pair. Instead of the voice being carried by POTS, it is now packetized directly in ATM and carried over the ADSL link. The voice packets are routed to a point in the service provider's network that interfaces through a gateway into the POTS network.

ATM's QoS capabilities allow the network to deliver acceptable voice-quality services. High-quality audio and video streaming are other potential services that can be offered using ADSL.

The demand for bandwidth and high-value content and services is soaring in both the home and office, and ADSL provides a robust and cost-effective mechanism to meet the demand. Existing infrastructures can often be reused, providing for fast and economical deployment.

Because ADSL was designed to be consumer-friendly, it can coexist with POTS. By using existing voice circuits instead of a shared broadcast medium, ADSL can better individualize services. The roll-out of ADSL will continue to increase as business models are created to exploit the capabilities, both business and technical, that are available.

Rob Rhodes is manager of the Communications Design group at Thomson Multimedia (Indianapolis, IN). The group is responsible for the development of the PHY interfaces for satellite, cable, and DTV communications. He can be reached at rhodesr@tce.com.

Mike Pugel is a principal member of the technical staff at Thomson Multimedia (Indianapolis, IN). He is currently working on advanced communication receiver concepts. He can be reached at pugelm@tce.com.

Louis Litwin is a member of the technical staff at Thomson Multimedia (Princeton, NJ). His focus is the development of wireless communication devices used for digital home networking and mobility applications. He can be reached at litwinl@tce.com.

John Richardson is a member of the technical staff at Thomson Multimedia (Princeton, NJ). He is primarily involved in the areas of digital home networking and ADSL system development. He can be reached at richardsonj@tce.com.

Resources

Busby, M., Demystifying ATM/ADSL, Worldware Publishing, TX, 1998.

Chen, W., DSL: Simulation Techniques and Standards Development for Digital Subscriber Line Systems, Macmillian, IN, 1998.

Cioffi, J., Silverman, P., and Starr, T., Understanding Digital Subscriber Line Technology, Prentice Hall, NJ, 1999.

Goralski, W., ADSL and DSL Technologies, McGraw-Hill, NY, 1998.

Rhodes, R., Pugel, M., Litwin, L., and Richardson, J., "Digital Subscriber Line Technology Tutorial," International Conference on Consumer Electronics, Los Angeles, CA, June 11, 2000 (invited tutorial).

Rob Rhodes is manager of the Communications Design group at Thomson Multimedia (Indianapolis, IN). The group is responsible for the development of the PHY interfaces for satellite, cable, and DTV communications. He can be reached at rhodesr@tce.com.

Mike Pugel is a principal member of the technical staff at Thomson Multimedia (Indianapolis, IN). He is currently working on advanced communication receiver concepts. He can be reached at pugelm@tce.com.

Louis Litwin is a member of the technical staff at Thomson Multimedia (Princeton, NJ). His focus is the development of wireless communication devices used for digital home networking and mobility applications. He can be reached at litwinl@tce.com.

John Richardson is a member of the technical staff at Thomson Multimedia (Princeton, NJ). He is primarily involved in the areas of digital home networking and ADSL system development. He can be reached at richardsonj@tce.com.