

13

ADSL Transmission

New lamps for old.
—Aladdin

Asymmetrical Digital Subscriber Line (ADSL) technology allows old ways of using telephone lines to be swapped for new ways that give much higher data rates. In fact ADSL goes one better because narrowband services do not have to be removed when ADSL is introduced. It looks like an offer you can't refuse.

ADSL challenges the Hybrid Fiber Coax (HFC) technology that cable operators can use. HFC supports bi-directional transmission with individual end-users over legacy plant that was designed for broadcast cable television. ADSL in contrast, and later perhaps High-speed Digital Subscriber Line (HDSL) or Very-high-speed Digital Subscriber Line (VDSL) technology, allows higher data rate communication over legacy plant that was designed for bi-directional voice communications with end-users (see Figure 13.1). These other technologies are described in more detail elsewhere (see Chapter 18 and reference [1] of this chapter).

Another metaphor for ADSL is that of putting new wine in old skins, since it is the way that old telephone lines are used that is changed and not the lines themselves. This may serve as a warning. The danger of putting new wine in old skins is that it may ferment and split the skins. Likewise putting

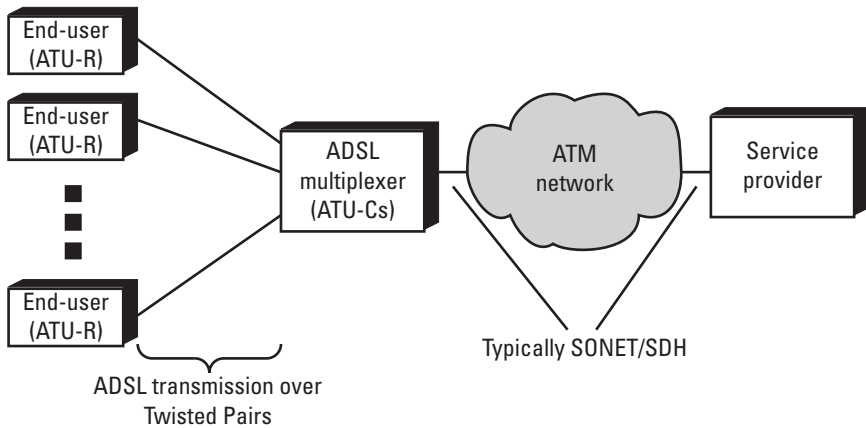


Figure 13.1 Typical ADSL Application

adaptive ADSL transmissions in a cable bundle may cause a reaction like fermentation and produce unpredicted emergent behavior.

Despite the danger of putting new wine in old skins, the quotation from “Aladdin” may be the more apt since ADSL will release the genie of the Internet from the lamp of narrowband access.

13.1 Tones, Modulation, and Coding

ADSL transmission uses a large number of discrete tones that are modulated to carry data and it decreases the sensitivity to noise by Trellis coding.

13.1.1 Tones

Standard ADSL [2] uses Discrete Multi-Tone (DMT) transmission with tones that start at 4.3125 kHz and are spaced 4.3125 kHz apart. There are 255 possible tones downstream, from an ADSL Termination Unit - Central (ATU-C) to an ADSL Termination Unit - Remote (ATU-R), and 31 possible tones upstream. If ADSL operates on the same line as PSTN or narrowband ISDN the lowest frequency tones are not used. For PSTN the reduction in bandwidth is acceptable since PSTN only needs 4 kHz, but it can be easier to treat ISDN as an incompatible service because it needs an order of magnitude more.

The particular tones used, the energy per tone, and the number of bits per tone are agreed during initialization. These values can also be adjusted

during normal operation. Echo-cancellation [1] can be used to allow the same tone to be used in both directions. In simpler implementations, uni-directional tones alone are used. This reduces the data rate downstream, because upstream transmission must use the lower frequency tones that are the least attenuated.

13.1.2 Modulation

Each DMT tone is modulated to a point in an amplitude-phase constellation. These points represent the amplitude and phase of the DTMF tone. The cosine direction (see Figure 13.2) represents the component that is in phase. The sine direction represents the component that is 90 degrees out of phase.

All constellations have the same nominal average power, although the actual average power may differ from this because of adaptation during operation. The level of noise at the tone frequency determines the constellation used, because the lower the noise the more points the receiver can detect.

The number of bits carried by a constellation determines its form. Constellations carrying an even number of bits are square. These constellations are generated recursively from constellations carrying two less bits by splitting each point into four (see Figures 13.2 and 13.3).

The constellation carrying three bits is a special case because it has four arms and a rotational symmetry (see Figure 13.2). Constellations carrying a higher number of odd bits are generated from the square constellation carrying one bit less by extending the edges (see Figure 13.3)

These agreed ADSL constellations are not optimal because they do not stack points in a hexagonal array. Like packing fruit in a tray, there can be more points in the same space if they are packed hexagonally rather than




Figure 13.2 Growth of constellations from 2 Bits to 4 Bits.




Figure 13.3 Growth of constellations from 4 Bits.

octagonally. The agreed constellations also do not take account of the cost of the points in power, which depends on the distance from the origin.

An optimal constellation has the greatest number of points for a specified average power, with the minimum spacing between points determined by the level of noise. It is more efficient to encode bits across multiple constellations because the number of points in an optimal constellation is rarely a power of two.

13.3.3 Trellis Coding

Trellis coding [3] may be applied before data is mapped onto constellation points to reduce the effect of noise. Trellis coding uses more symbols than needed. The choice of symbol provides additional information about previous data and so reduces the sensitivity to noise.

Trellis coding can have a negligible effect on signal bandwidth if amplitude-phase modulation is used. The cost of Trellis coding is an increase in the complexity of the decoder needed to recover the data. The decoder is more complex because it needs to take account of every subsequent symbol that would be affected by the transmitted data before it makes a decision. Typically the Trellis decoder uses the Viterbi algorithm [3], which eliminates possible sequences of symbols as they become unlikely. Decoding is simplified if dummy data is sent periodically to reset the decoder.

In a simple example of Trellis coding, the value of each bit is represented by one of two possible pairs of orthogonal symbols. The choice of symbol pair depends on the value of the previous bit. The four orthogonal symbols are (+,+), (-,+), (-,-) and (+,-). The second half of the symbol is determined by the current bit, but the first half of the symbol is flipped or not flipped depending on the value of the previous bit.

In this example, if noise causes one half of a symbol to change during transmission then the bit can still be recovered because the information is

spread over three half symbols periods (see Figure 13.4). If the four symbols correspond to the four constellation points of phase modulation, the coding gain costs almost no power or bandwidth when the data rate is much less than the carrier frequency.

In general, a Trellis encoder can have a number of states with the new state and the symbol to be transmitted determined by the input data and the previous state. It is called Trellis coding because a trellis pattern is created by the transitions between the states of the encoder (see Figure 13.5).

ADSL data is mapped onto ordered pairs of DMT tones when Trellis coding is used. If there are an odd number of active tones then an inactive tone that carries no bits is used in the first ordered pair. Each active tone carries at least two bits since phase modulation allows two bits to be carried for the same energy as one and the Trellis coding more than compensates for the reduced distance between adjacent points.

Normally the number of constellation points in the second tone of each tone pair is twice what is needed if there was no Trellis coding. The constellation point in the second tone is determined by the least significant bits that are carried by the pair and by the choice made for previous pairs.

Two dummy bits are added to the last two-tone pairs for the ADSL frame to reset the decoder for the next frame. This is done to simplify the decoder as these bits ensure that each transmitted frame has no residual impact on the next frame.

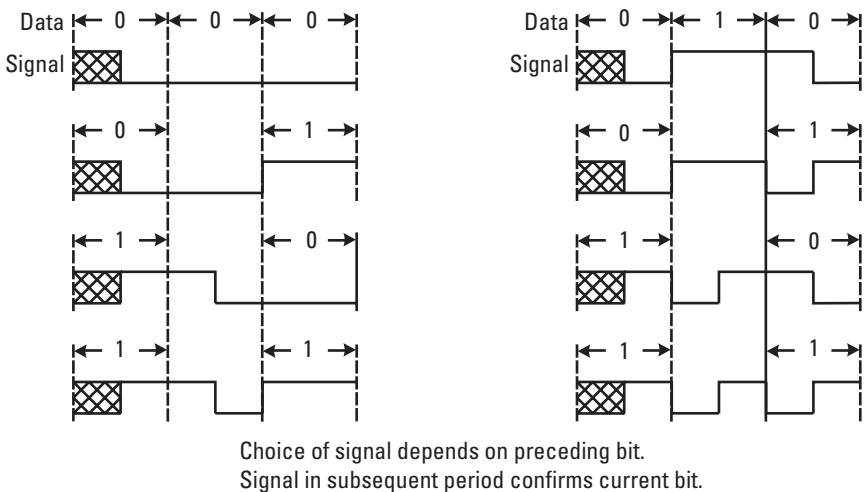


Figure 13.4 A simple example of trellis coding.




Figure 13.5 Generation of the trellis pattern.

The Trellis coding for ADSL is less than optimal when it is implemented by detecting constellation points and then calculating the transmitted data. This is because information is lost when signals are mapped onto constellation points.

13.2 Frames, Superframes, and Symbols

Tones are ordered according to the number of bits that they carry and then allocated to the transmitted ADSL frames.

13.2.1 Frames, Superframes, and Tone Allocation

ADSL frames are grouped into superframes that are 17 ms long and contain 69 symbol periods (see Figure 13.6). The first 68 of these symbol periods are allocated to the payload frames and the final symbol period is allocated to superframe synchronization.

If 250 tones are used in the downstream direction then payload data would be carried in one million tones every second (see Figure 13.7). In the upstream direction, the data rate is about an order of magnitude less. Tones are ordered according to the number of bits that they carry. The bits at the beginning of the payload frame are assigned to the tones that carry the least bits.

13.2.2 Payload Frames and ADSL Channels

Payload frames contain either the fast channel or the interleaved channel, or both. The interleaved channel has a greater delay because data is held back for transmission on later frames. The number of bytes in a frame and the allocation of bytes to the channels are flexible.




Figure 13.6 Frames, Superframes, and Symbols.




Figure 13.7 Upstream and downstream capacities.

The interleaved channel is allocated to the second half of the frame because tone allocation places the tones that carry the most bits there, and these tones are susceptible to corruption if the transmitted signal is clipped. The interleaved channel is better protected against this corruption because the interleaving makes the forward error correction more effective. Tomlinson precoding [1], which avoids clipping, is not used for ADSL.

13.3 Forward Error Correction

Forward error correction is added to each ADSL channel to reduce the need for retransmission. The interleaved ADSL channel has the more effective forward error correction but the cost is yet further delay.

13.3.1 Error Detection and Error Correction

Error detection and forward error correction work by transmitting the information as a codeword that contains additional information. The redundant information in the codeword allows errors to be detected. In forward error correction there is sufficient redundancy for a limited number of errors to be corrected.

ADSL uses the Reed-Solomon method [3], where the codeword is created by adding additional bytes. These bytes are generated using polynomial division with modulo arithmetic of polynomial coefficients.

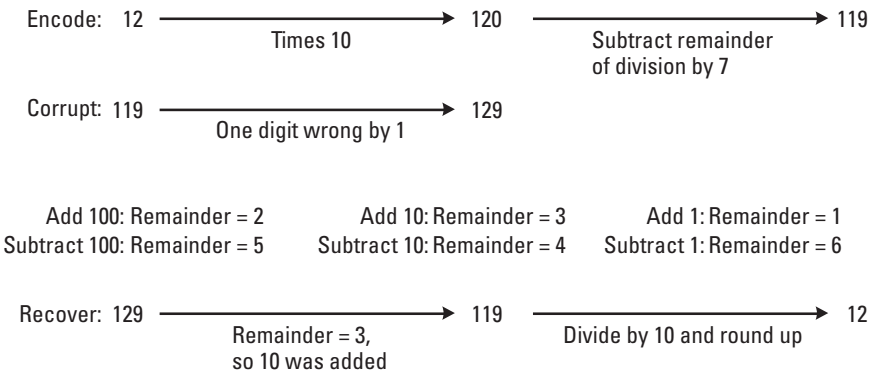


Figure 13.8 Equivalent of Reed-Solomon error correction in simple arithmetic.

The original message bytes are used as the coefficients of a polynomial expression, the message polynomial, and the additional bytes are the remainder when this is divided by the generating polynomial of the Reed-Solomon code. Modulo arithmetic ensures that the coefficients of the remainder are also equivalent to bytes, and these are appended to create the codeword. The generating polynomial is mathematically similar to a prime number in ordinary arithmetic.

In ordinary arithmetic errors can be detected if the only numbers that are used are numbers that are perfectly divisible by another number. For example, numbers that are divisible by seven can be generated by multiplying an initial number by ten and then subtracting the remainder. This allows all errors in arithmetic to be detected except for those errors that introduce multiples of seven.

Errors in arithmetic can also be corrected because different errors produce different remainders when the numbers are divided by seven (see Figure 13.8). There are six ways that a single digit in a three-digit number can be out by one, e.g. a three being written as a two or a four. These six ways correspond exactly to the six possible non-zero remainders when a three-digit number is divided by seven.

This simple example of error correction works because all possible remainders of division by seven are generated by powers of the number ten, which is the base used in ordinary arithmetic. This relationship between the prime number seven and the base ten corresponds to the relationship between the generating polynomial and its length. The effect of noise can be assessed because errors are similar to simple corruption of digits.

13.3.2 Forward Error Correction in ADSL

The bit streams of both the interleaved and the fast channel are scrambled before the Reed-Solomon codewords are generated. For the fast channel a Reed-Solomon codeword is constructed by adding error correction bytes in each frame. Forward error correction in the fast channel is effective if the errors are distributed evenly in time because the chance of overloading a frame with errors is low. Unfortunately errors are often concentrated in bursts that occur in a minority of the frames.

Interleaving the data from different frames makes the error correction more effective because the frames that are error free help to correct the frames that contain many errors. The interleaved transmission also increases the effect of error correction by having a longer Reed-Solomon codeword. Bytes in the interleaved channel are transmitted sequentially but with spaces between them that are filled with the bytes from previous frames. Only the first byte in each frame is not delayed.

13.4 Channels, Ports, and Framing

Payload frames have different modes depending on the number of ADSL channels and whether or not rate adaptation is applied to the bearer channels within them. Rate adaptation is controlled by frame overhead bytes that also perform conventional overhead functions.

13.4.1 Bearer Channels, ADSL Channels, and ATM Ports

Both the fast and the interleaved ADSL channels can support up to seven bearer channels (see Figure 13.9). These form two groups, a simplex (unidirectional) group of four bearer channels (AS0 to AS3) and a duplex (bidirectional) group of three channels (LS0 to LS2). The simplex channels only carry data downstream.

The simplex group normally has a higher data rate than the duplex group because the duplex group is limited by the data rate for upstream transmission. Each bearer channel can operate at a multiple of 32 Kbps. LS0 can also operate at 16 Kbps if it is used to carry signaling for the other bearer channels.

The bearer channels can either be used separately for synchronous transmission or they can be combined to form two ATM ports, one fast and the other interleaved.

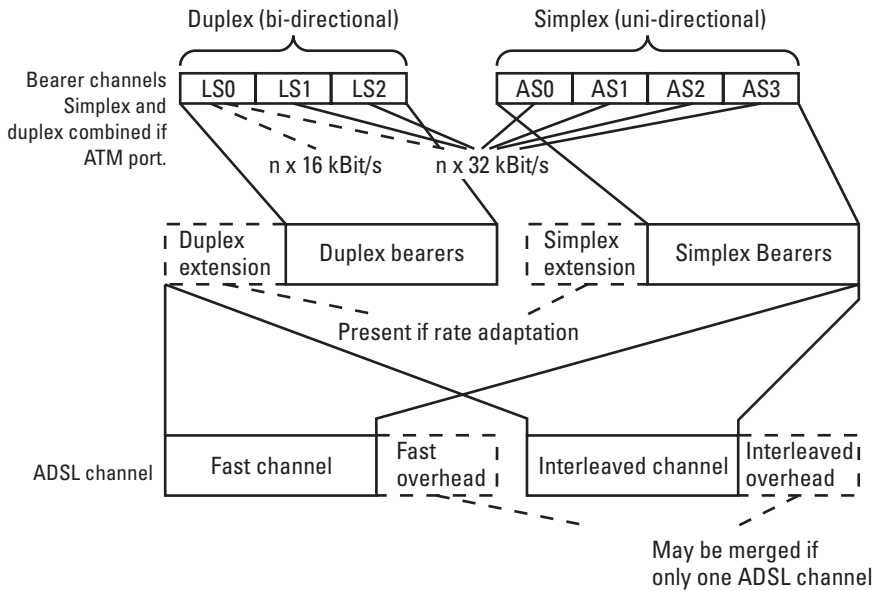


Figure 13.9 Channels, ports, and framing.

13.4.2 Framing and Overheads

The bearer channels are carried within the ADSL channels in one of four possible framing modes. The mode depends on whether the ADSL channels include extension bytes for rate adaptation, whether these extension bytes are used, and whether the ADSL channel overheads are merged if only one channel is present.

In practice this flexibility is not needed because Service Providers obtain greater flexibility at the ATM layer with the simplest ADSL configuration (i.e. one interleaved channel and a single overhead). The only concern is that for delay sensitive applications the fast channel may be necessary instead.

The simplex and duplex bearer channels are multiplexed together onto their ADSL channel and prefixed with an overhead byte. The fast and interleaved ADSL channels carry different information in their overhead bytes and the bytes carry different information in the different frames of the superframe.

The information in the overhead bytes also varies with the framing mode. In the full framing mode, simplex and duplex extension bytes are added after the multiplexed bearer channels. For upstream transmission there

is no simplex extension byte because no simplex channels are carried upstream.

The overhead bytes control the use of the extension bytes for rate adaptation by indicating the bearer channels that they are allocated to. They also indicate whether rate adaptation is necessary and whether there is an underflow or an overflow. If there is an overflow then the extension byte is added to the normal bytes. If there is an underflow then the final normal byte of the identified bearer channel is ignored.

The overhead bytes in frame zero of the superframe always carry the CRCs of their ADSL channels. These CRCs are calculated for the previous frame. The overhead bytes also carry 24 indicator bits, which carry alarms and can provide a timing reference, and the Embedded Operations Channel (EOC) and ADSL Overhead control Channel (AOC).

In frames other than 0, 1, 34 and 35, the least significant bit of the fast overhead byte determines whether it controls rate adaptation or whether the overhead carries the EOC instead. In frames other than 0 the interleaved overhead is always allocated to rate adaptation unless the interleaved channel is present but does not carry any bearer channels. If there are interleaved bearer channels and there is no merged overhead then the interleaved extension byte for the duplex bearer channel group is used for AOC when it is not needed to carry an overflow for a bearer channel. The interleaved overhead byte is also used for the AOC if the interleaved channel carries no bearer channels. Unlike the EOC, the AOC is carried in whole bytes.

13.5 Summary

Asymmetrical Digital Subscriber Line (ADSL) transmission uses a very large number of discrete tones that can each be modulated to carry a number of bits. The immunity of these bits to noise can also be increased through Trellis coding, which introduces redundancy to increase the effect that the transmitted data has on the transmitted signal. Tones are ordered according to the number of bits that they carry and are allocated, tones with fewest bits first, to the ADSL frames.

Payload data is carried in 68 of the 69 frames that make up a 17 ms superframe. Payload frames contain either a fast ADSL channel or an interleaved ADSL channel, or both. Forward error correction is added to each ADSL channel payload to reduce the need for retransmission. The interleaved ADSL channel has the greater delay but the more effective forward error correction.

One of four payload-framing modes is used depending on the number of ADSL channels and whether or not there is rate adaptation. Rate adaptation is controlled by the frame overhead bytes, which also carry error detection for the superframes and indication bits for alarms and for network timing. The frame overhead bytes are also used for the Embedded Operations Channel (EOC) and the ADSL Overhead control Channel (AOC).

Each ADSL channel in the downstream direction, from an ADSL Termination Unit-Central (ATU-C) to an ADSL Termination Unit-Remote (ATU-R), can carry a group of bi-directional (duplex) channels and a group of unidirectional (simplex) channels. Only the duplex group of bearer channels can be carried in the upstream direction. The bearer channels on each ADSL channel can be used separately for synchronous channels or combined to create an ATM port.

References

- [1] Gillespie, A., *Access Networks: Technology and V5 Interfacing*, ISBN 089-006-928-X, Norwood, MA: Artech House, 1997.
- [2] *Asymmetrical Digital Subscriber Line (ADSL) Transceiver*, ITU-T Recommendation G.992.1, 1998.
- [3] Clark, G. C., Jr, and Cain, J. B., *Error-Correction Coding for Digital Communications*, ISBN 0-306-40615-2, Plenum Press (New York), 1981.