Discrete Multi-Tone Asymmetric Digital Subscriber Line

Kevin Baker

Abstract

This paper describes discrete multi-tone (DMT) line coding, which has become the standard asymmetric digital subscriber line (ADSL) modulation technique as selected by the American National Standards Institute (ANSI) T1E1 committee. Presented here is an overview of DMT ADSL including quadrature amplitude modulation (QAM) encoding, rate-adaptation, bit-swapping as well as some of the overall advantages to using DMT for ADSL modulation.

Introduction

Asymmetric digital subscriber line (ADSL) is widely hailed as the technology of choice for telecommunications companies to provide high-speed Internet access to their customers. ADSL has allowed the service providers to expand on the unshielded twisted pairs 1 MHz capacity by providing a downstream bandwidth (towards the customer premise) of up to 6.144 Mb/s and an upstream bandwidth (towards the network) of up to 640 kb/s. The essential principle of ADSL is that by incorporating frequency division multiplexing (FDM) plain old telephone service (POTS) is transmitted in the baseband while the upstream and downstream data rates are modulated at higher frequencies. Perhaps most importantly, ADSL is cost effective to both the service provider and the customer by utilizing the existing resources of the unshielded twisted pair.

Discrete multi-tone (DMT) has become the accepted standard for ADSL modulation. The reason for this is that DMT line coding has several advantages that lend exceptionally well to ADSL. These advantages include rate adaptation as a function of the signal-to-noise ratio (SNR), adjustable bit rates and an inherent immunity to impulse noise and radio frequency interference (RFI).

ADSL Overview

There are a number of key components that comprise the overall ADSL architecture. These components are the ADSL transceiver unit - central office (ATU-C), the ADSL transceiver unit - remote (ATU-R), the plain old telephone service (POTS) splitter and the digital subscriber line access multiplexer (DSLAM). Figure 1 below helps to illustrate the fundamental architecture of ADSL and these key components.

In the downstream direction POTS and ADSL signals are transmitted from the central office to the customer premise. At the customer premise, a POTS splitter is utilized to demultiplex the POTS and ADSL signals. This is accomplished by implementing a passive low-pass filter (LPF) with a cutoff frequency of approximately 4 kHz (the frequency spectrum of POTS). The splitter is also used to guarantee POTS service even in the event that ADSL service should fail. There are two combined outputs of the splitter, one that transmits a POTS-only signal to the telephones on-site and one that transmits both the POTS and ADSL signal to the ATU-R, which is also commonly referred to as an ADSL modem.

The ATU-R houses a high-pass filter (HPF), which is implemented to block the POTS signal and POTS associated DC voltages and spikes due to on and off-hook conditions. The resulting data signal is transmitted via Ethernet (10/100BaseT) to an Ethernet interface card contained in the customer's personal computer (PC).

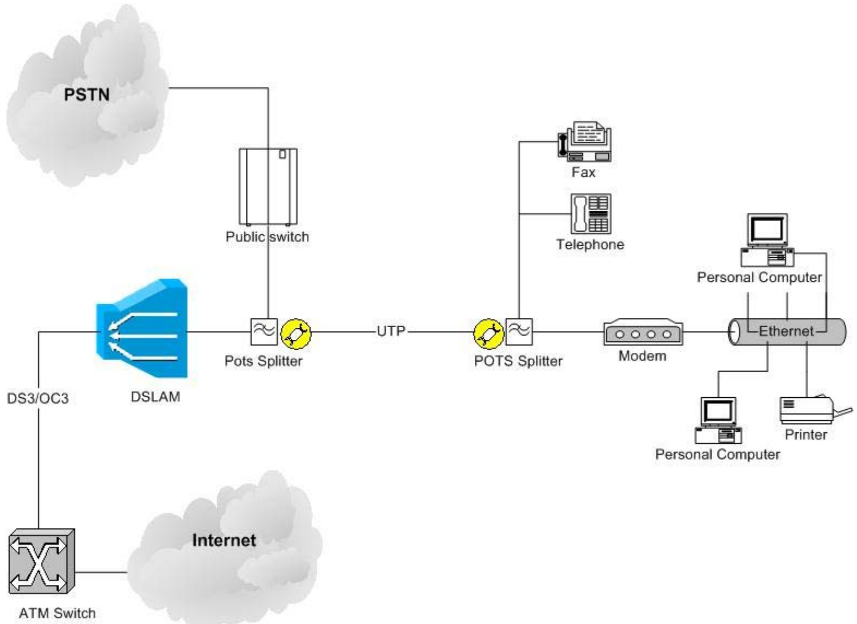

Figure 1. ADSL Architecture

In the upstream direction (towards the central office), the POTS and ADSL signals are multiplexed by the POTS splitter for transmission over the unshielded twisted pair to the central office. At the central office, the combined POTS and ADSL signals leave the main distribution frame (MDF) and are sent to a DSLAM. The DSLAM is used to house several ATU-Cs and respective POTS splitters for multiple subscribers. Again the POTS splitter is incorporated as an LPF to demultiplex the POTS signals (DC to 4kHz) and transmits these signals off towards the public switched telephone network (PSTN) and ultimately a POTS switch.

The ADSL signal is transmitted to the respective ATU-C card housed in the DSLAM, which also incorporates an HPF to block DC voltages and spikes. The ADSL signals from the individual ATU-Cs within the DSLAM are then multiplexed onto a high speed bus in the neighborhood of 155 Mb/s and transmitted as an OC3 or DS3 signal off to the data (usually ATM) network via what is known as a network termination (NT) card in the DSLAM.

DMT ADSL Obstacles

Theoretically and ideally, on any given twisted pair an ATU-C and an ATU-R should achieve the exact same throughput. However, numerous differences between transceiver designs affect performance in application. It is important to note however that design conflicts are far more rare and by implementing an ATU-C and an ATU-R manufactured by the same vendor can in many cases rectify numerous performance issues.

At least in theory, the optimum transmission technique is to "match" the signal to the transmission medium, obviously an extremely difficult task. In practice, distributing the transmitted power in the area of the frequency spectrum that will be best received at the far end of the medium is something to strive for. In other words, the transmission medium is an enormous factor to account for when designing transceivers. This why it is so important ADSL transceivers be able to adapt quickly to varying line conditions. With the unshielded twisted pair medium, higher frequencies are attenuated considerably more than lower frequencies. Therefore, the inherent characteristics of the medium affect different frequencies quite differently. As will soon be explained, a DMT transceiver can monitor as well as adapt to UTP varying conditions and in addition, continually update to maintain optimum performance. This is how DMT systems transmit the "best" possible signal and data rate.

For ADSL systems and like most communications systems, the primary obstacle is noise. Noise presents itself in a variety of ways such as thermal noise, impulse noise and radio frequency interference (RFI).

Impulse noise such as that from household appliances, weather conditions or that from a phone ringing is in fact broadband with respect to frequency. However impulse noise is averaged across many what are called DMT sub-channels or tones. Therefore DMT copes well when faced with the adverse conditions of impulse noise.

Radio stations are the most significant source of radio frequency interference (RFI). Particularly true when noting that the ADSL band of 1 MHz is well within the AM band. DMT systems are inherently smart though to place signal power where it will be most advantageous – at the receiver and therefore not be subjected to elimination by interference. As will soon be explained further, DMT systems place energy into frequency areas that can optimize it and not waste excess energy where it will not be optimized.

ADSL Signal Encoding

Even with today's most sophisticated modulation techniques, the narrow 4 kHz baseband frequency of POTS can only achieve a throughput of up to 56 kb/s. This can be seen by referring back to simple DS0 mathematics:

8 bits X 8 kHz sampling rate = 64 kb/s

Although this simple equation offers a throughput of 64 kb/s, 56 kb/s is the theoretical limit. In actuality, the downstream data rate is transmitted at 64 kb/s, however, there is a lower SNR due to line impairments and non-linearity. The local loop has a digital-to-analog converter (DAC) in the central office that converts the downstream data into an analog waveform that represents the data. The DAC runs at 64 kb/s but due to noise, loop length, and other network impairments, speeds are reduced. In addition to this, Federal and other regulatory agencies, in order to minimize crosstalk between adjacent lines have signal power level requirements. To meet these requirements for maximum signal strength levels on the telephone lines, a 56 kb/s limit is the theoretical maximum.

As initially stated, the ANSI T1E1 committee has standardized DMT as the line code to be used in ADSL transmission systems, due to its unique ability to overcome the severe distortion of the unshielded twisted pair at this frequency band. DMT divides the channel into a number of sub-channels, commonly referred to as tones, each of which is quadrature amplitude modulated (QAM) on a separated carrier. These carrier frequencies are multiples of one basic frequency. The subsequent available spectrum ranges from approximately 20 kHz to 1.104 MHz, while the low 20 kHz are reserved for POTS, which are more realistically capped closer to 4 kHz. Noise and line conditions are constantly measured for each tone separately, in order to optimize transmission at any given time. This is exemplified in Figure 2 below.

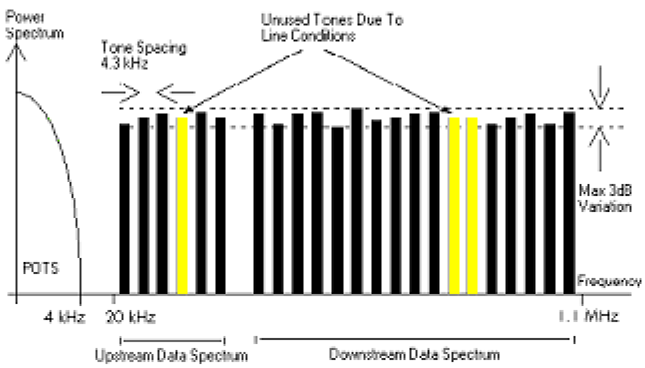

Figure 2. DMT ADSL Spectrum

To support bi-directional channels, ADSL systems divide the available bandwidth by frequency division multiplexing (FDM), where non-overlapping bands are assigned for the downstream and upstream data. DMT modulation is actually a form of FDM (see Figure 3). The input data stream is split into N channels having the same bandwidth but a different centre frequency. Using many channels with a very narrow bandwidth results in several advantages. The most significant of these advantages is that all channels become independent regardless of line characteristics and therefore can be individually encoded and decoded.

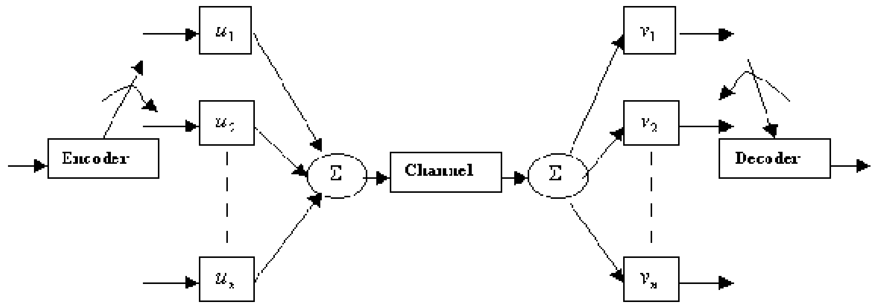

Figure 3. Block Diagram of a DMT Transceiver

The standard (ANSI) ADSL system uses 256 frequency sub-channels or tones for the downstream data and 32 tones for the upstream data. All channels have bandwidth of 4.3 kHz and the frequency difference between two successive channels is also 4.3 kHz. Each tone supports a maximum of 15 bits, which of course is subject to the signal-to-noise ratio (SNR). Higher frequency tones particularly in the downstream data spectrum are prone to higher attenuation and noise and are thus often allocated a fewer number of bits per sub-channel than the lower frequencies within the spectrum.

Quadrature Amplitude Modulation

In order for ADSL to achieve the 15-bit maximum that any sub-channel or tone can accommodate, quadrature amplitude modulation (QAM) is employed. In its essence, QAM is simply a refined combination of amplitude modulation (AM) and phase shift keying (PSK). Using a 3-bit QAM constellation example, we can see that eight binary combinations are required to represent the signal to be transmitted. The example presented here assumes two possible measures of amplitude (1 and 2V) combined with four possible phase shift combinations.

Utilizing the QAM and PSK combination, a stream of 3-bit words can be extracted from a relatively large bit stream. Figure 4 depicts the QAM-encoded signals of the bit stream relative to Table 1, where each wave period is phase shifted in relation to the wave of the period immediately preceding it. It is important to note however that this is only a minute example of QAM in an ADSL system. 32,768 different combinations of phase shifts and amplitudes would be required to achieve the maximum allocation of 15-bits of data to an individual sub-channel or tone.

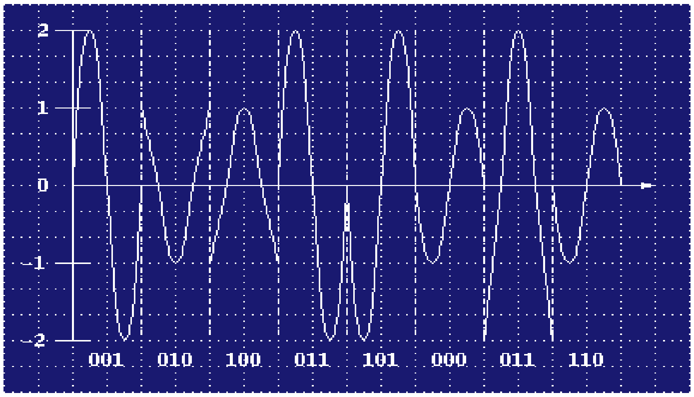

Figure 4. QAM Encoding

Rate Adaptation in ADSL Systems

As stated earlier, the ADSL signal that is transmitted along an unshielded twisted pair is subject to numerous forms of impairments due to inherent line conditions. These impairments vary from factors such as the gauge of the UTP to the lines proximity to sources of radio frequency interference (RFI). Compensation for these conditions is achieved by incorporating rate-adaptive ADSL (RADSL) systems to achieve the maximum desired throughput. There are primarily three RADSL systems that are employed by DSLAM vendors today, which are, manual, initialization and dynamic RADSL, all of which are usually a component of the vendors network management system (NMS).

Manual rate adaptation is where the NMS specifies a desired bit rate that must be supported by the ATUs. Synchronization will not occur between the ATUs if the line conditions are such that the specified bit rate will not be achieved. While both ATUs are synchronized and the connection is in what is commonly referred to as the "showtime" state, no further rate adaptation will occur.

With the rate adaptation being performed at initialization, a minimum and maximum bit rate is specified by the NMS for the ATUs to support. The ATUs will attempt to maximize their throughput by achieving the maximum bit rate possible given their respective line conditions. Again if the minimum bit rate is not achievable the ATUs will not be able to achieve synchronization. As well, during showtime, no further rate adaptation takes place.

As the name implies, dynamic rate adaptation monitors the line conditions during showtime and adapts to any variances in the line conditions in order to achieve the maximum bit rate and throughput possible between the ATUs as defined by the NMS.

Bit Swapping

With all three rate-adaptive modes, a process known as bit-swapping occurs during showtime. As the ADSL system operates, each 4 kHz tone is constantly monitored for the quality of the line. Optimum performance is maintained by adjustments by the system to the bits-per-tone distribution. If the noise margin for a particular tone is degraded below the specified SNR, as many bits as required are allocated to other tones that can accommodate additional bits. For clarification purposes, this procedure is exemplified in a simplistic form below in Figure 5.

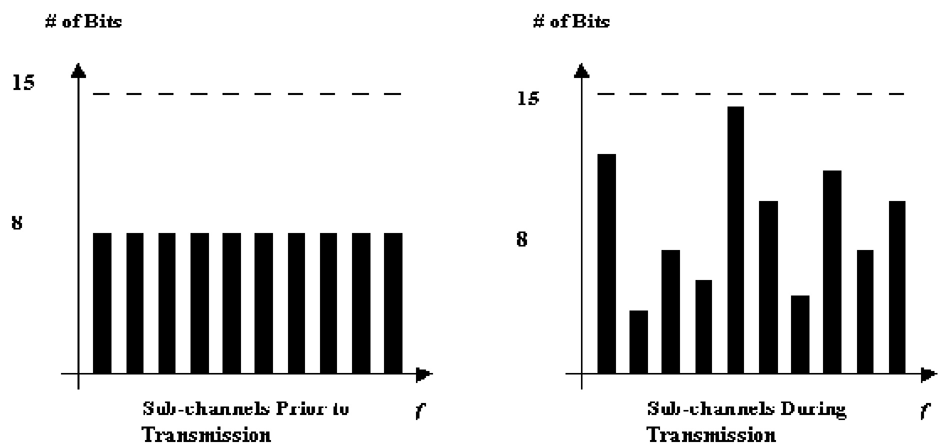

Figure 5. Bit-Swapping

Therefore, as stated, the encoder receives the data bit stream and encodes it into N QAM constellation points. This encoding is done according to the bit-loading table, which defines the number of bits carried by each tone. Clearly high signal-to-noise ratio carriers can carry more bits than low SNR carriers so the bit-loading table reflects the variation of the SNR over frequency.

Conclusion

Discrete multi-tone asymmetric digital subscriber line (DMT ADSL) incorporates frequency division multiplexing (FDM) to transmit plain old telephone service (POTS) signals in the baseband, while modulating the ADSL upstream and downstream data rates at higher frequencies. DMT further divides these bi-directional channels into sub-channels or tones, which are quadrature amplitude modulated (QAM) on a separated carrier. QAM also enables the system to achieve the 15-bit maximum that any tone can accommodate. Compensation for the impairments in the twisted pair line conditions is achieved through rate-adaptive ADSL (RADSL), which consequently results in maximum throughput. Finally, a process known as bit-swapping is employed to further provide optimum performance by adjusting the bits-per-tone distribution as specified by the signal-to-noise ratio (SNR).