Introduction to ADSL

Christian Truppel

ADSL Explained

The costs of exchanging large numbers of bits and bytes used to be equally high for all available methods. Whilst large companies used microwave and satellite networks to span long distances, and rented expensive T1 and E1 lines for dedicated telephone and data communications between facilities, such technologies were not often available to small and medium enterprises (SMEs), not to mention home users. The consumer-end copper wire was designed and installed to carry analogue voice or modem signals at low frequency, which enabled the wire to transport the signals to an exchange up to 18,000 feet (5,500 meters) away without attenuation or degradation of sound quality.

Nowadays, however, superior data capacities can be achieved at a relatively low cost using Asymmetric Digital Subscriber Line (ADSL) technology. Asymmetric, in this case stands for a higher downstream data rate (6.144 mbps) than upstream (640 kbps). Digital means that the information, whatever kind it may be,is transferred digitally and ‘Subscriber Line’ is to say, that home or end users of this digital data transfer service have to subscribe themselves, using the existing telephone lines. Modern telephone networks start at the consumer end with twisted-pair copper cable that runs to a switching centre or exchange, where it connects into the telephone network. The backbone that connects these exchanges is now mainly fibre optic cable, which provides the high-capacity necessary for providing high quality transmission to multiple households.

ADSL Basics

ADSL uses discrete multitone (DMT) modulation to convert a single wideband communication channel into multiple narrowband subchannels. This is done by segmenting the data into blocks and using an IFFT operation at the transmitter and a Fast Fourier Transform (FFT) operation at the receiver. Bits and power are allocated to individual subchannels (tones of the IFFT/FFT) to maximize the data rate for a fixed margin or to maximize the margin for a fixed data rate. The process of allocating bits and power to individual subchannels is referred to as bit loading.

The use of a cyclic prefix (preceding the tail of the signal after the IFFT to the block to be transmitted) allows for simplified equalization at the receiver if the channel memory is not larger than the length of the cyclic prefix. In this case, the effects of the channel are removed by equalization, which simply reduces to multiplication by a complex number on a subchannel by subchannel basis. This can be shown by noting that the combination of the IFFT, channel, and FFT results in a diagonal matrix relating the input block of data to the received block of data, with the channel response as the elements of the diagonal.

For cases where the channel memory is larger than the length of the cyclic prefix, it is common to use a channel shortening filter in the receiver (referred to as a time domain equalizer or TEQ) which is designed so that the cascade of the TEQ and channel has the majority of it's power limited to a window of the length of the cyclic prefix + 1. This is in contrast to the single channel case, where a linear equalizer is typically designed to invert the channel in the minimum mean square error (MMSE) sense.

While data and other items, which are periodic with respect to the subchannel spacing only affect a single subchannel, typical noise sources are not periodic with respect to the subchannel spacing and thus tend to affect multiple subchannels. The mechanism through which this occurs is often referred to as FFT spreading. Viewing the blocking operation before the FFT at the receiver as windowing in time with a rectangular window, noise is spread across subchannels at the rate of a sinc() function (the FFT of a rectangular window). This has implications when there are nulls in the channel response, or the noise near a subchannel is much larger than the noise affecting any of the neighbouring subchannels.

Typical ADSL deployments are frequency division duplex (FDD), meaning that the upstream (consumer to service provider) and downstream (service provider to consumer) signals occupy separate frequency bands separated by a transition band. A single channel (twisted pair wire) is used to carry both signals, and is coupled to the transmitter and receiver of each modem through a hybrid circuit. In broad terms, the hybrid performs four-to-two wire conversion using an isolation transformer. More specifically, the hybrid performs analogue echo cancellation to reduce the reflection of the transmitted signal (or echo) in the receiver. The hybrid echo rejection depends on the reflected line impedance (through the transformer), which in turn varies for different loop topologies.

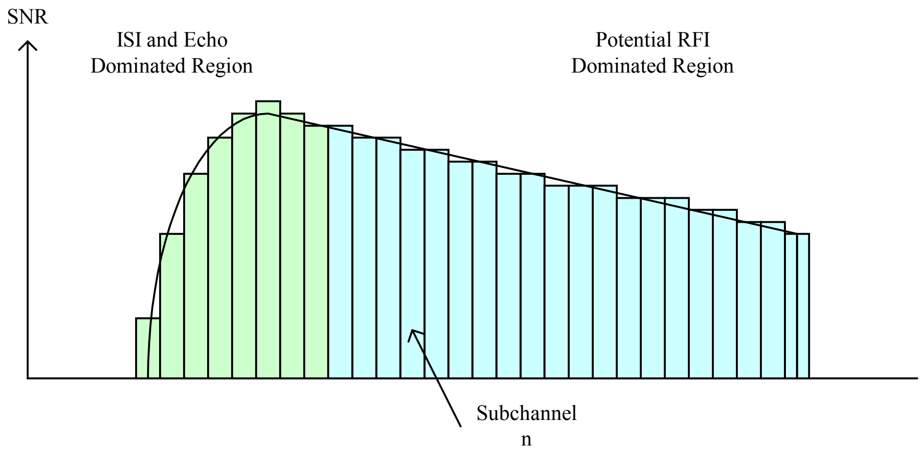

Figure 2: An example of the downstream PSD illustrating how DMT divides a single channel into multiple subchannels. ISI, echo and potential RFI dominated regions are also shown.

ADSL technology and improvements

The ADSL technology is the most frequently used DSL technology from a worldwide perspective. As of today, an ADSL modem is used for the data as well as the Plain Old Telephone System (POTS) for the voice transmission. The splitter separates the corresponding frequency bands. This allows a single twisted-pair copper wire to carry normal telephone communications in the 0 to 3.4 kHz range (the narrow band), data upload from the consumer in the 30 kHz to 138kHz range and data download to the consumer at up to 1104kHz (the wide band).

Modulation is the method of converting digital signals into waveforms and transmitting them over a wire. For the construction of these waveforms in an efficient way, 'modulation transform' employs the use of an encoder and a modulator. ADSL uses discrete multitone (DMT) modulation to convert a single wideband communication channel into multiple narrowband subchannels. This is done by segmenting the data into blocks. Bits and power are allocated to individual subchannels to maximize the data rate for a fixed margin or to maximize the margin for a fixed data rate. The process of allocating bits and power to individual subchannels is referred to as bit loading.

Each of these channels is then encoded (using a QAM encoder) with data and transmitted in parallel. Each subchannel has limited bandwidth, and thus uses longer symbol rates and slower baud-rate when tranceiving to reach high-bandwidth levels. DMT is the choice for ADSL and the encoder efficiency has been improved in the latest ADSL developments, such as ADSL2/2+, where new algorithms and various enhancements to QAM and the encoding process is greatly improving the data rates and overall performance of ADSL technology.

ADSL (G.992.1)

The original ADSL standard G.992.1 is well known and widely deployed worldwide. The standard allows for up to 8Mbps downstream and up to 1Mbps upstream with a reach of about 15,000 feet.

ADSL2 (G.992.3)

The ADSL2 standard contains numerous improvements compared to the previous standards. A majority of the improvements are based on experiences and feedback from the field. Also known as G.dmt.bis, the standard specifies 800Kbps to 1.5mbps upstream bandwidth and 8 to 16 mbps downstream with the same 5km reach as ADSL. The frequency spectrum used by ADSL2 is still in the 1.1 MHz limit.

In addition to faster start-up times (from ~10s to ~3s) and native support for voice with VoDSL, the ADSL2 standard contains a number of important developments.

Line diagnostics

Built-in diagnostic functions in ADSL2 transceivers allow for more efficient troubleshooting of errors on subscriber services. Items that are measurable are interference, loop attenuation and noise ratios (signal/noise ratio) at both ends, giving the network operator real-time information on the quality of an individual ADSL2 line. The diagnostics data is an excellent tool used to determine whether a customer has a good-enough line to handle certain types of services, for example high-bandwidth Internet or video services.

Improved modulation

The improved modulation in ADSL2 gives the ability for a line to hold a higher data rate on long lines, compared to classic ADSL. The bandwidth drop-off is not as sharp with ADSL2 as in classic ADSL. This is accomplished by a greatly enhanced modulation scheme with higher gain rates, reduced overhead, better algorithms for signal processing and enhanced efficiency in the actual modulation.

Power management

During periods of inactivity and low traffic levels, an ADSL2 line can enter certain power-saving modes. Instead of always being on full power, the ADSL2 device reduces the power level, which in turn reduces the data rate significantly while giving a dynamic approach to power management. In addition to saving electricity at both ends, this feature will enhance the line quality of 'surrounding' lines, as the probability of crosstalk and near-end echo is reduced. This will effectively give users in the same area a better service. Keeping electricity consumption low will also have a positive effect on the cost of ownership and will save natural resources. Saving power is an especially important signal for companies with an environmental profile.

Improved rate adaptation

The aggregation points for ADSL lines are basically massive bundles of copper wire, terminated at a central location. Electromagnetically, interference between individual wires in these bundles is inevitable, and has a negative effect on the line quality - this is known as crosstalk. Other elements of interference, such as bad copper, water or other RF transmitters, will also negatively affect the quality of a line. The improved rate adoption in ADSL2 enables the rate to be dynamically changed, as the quality of a given line changes. In times of heavy interference for certain lines the data rate will be adapted transparently to match the quality of the line, this will avoid the loss of service at the expense of bandwidth.

ADSL2+ (G.992.5)

By using a larger frequency spectrum, ADSL2+ gives a great increase in data rates while still maintaining the characteristics and features of its parent, ADSL2. The total reach is slightly increased compared to ADSL2. The G.992.5 standard defines a reach of 18Kft, or around 6Km. ADSL2+ can reach bandwidths of up to 26Mbps downstream and around 1-1,5Mbps upstream.

Bandwidth doubling

ADSL2+ uses a frequency range all the way up to 2.2Mhz, essentially doubling the frequency spectrum. The result is simply doubled bandwidth on shorter distances, generally this only effective on distances below 3000m; longer lines will not be able to enjoy the doubled data rates. The use of frequencies above 1.1Mhz means that ADSL2+ cuts in on the frequencies defined for VDSL, as VDSL has a defined range of 1.1 to 12Mhz. Therefore it is not recommended to run ADSL2+ and VDSL on the same wiring.

Transmit Spectrum Based Performance Improvements: ADSL 2 & ADSL2+

Because ADSL is based on DMT modulation, it has a great deal of flexibility in shaping its transmit spectrum. This flexibility can be used to improve the reach of the ADSL system, for handling mixed CO and RT deployments, and for crosstalk minimization.

Spectral Shaping for Extended Reach

The general shape of the DSL channel is such that higher frequencies are attenuated more than lower frequencies. Additionally, channel attenuation increase as the loop length increases. Because FDD ADSL systems allocate higher frequencies to the downstream, to improve the performance of the ADSL on long loops, it is typically necessary to improve the downstream data rate.

ADSL2 approaches this problem with a specific annex (reach extended ADSL2), which uses spectral shaping for putting power where the channel is better, either by shrinking the range of the downstream frequencies and boosting the power or, by overlapping the downstream with the upstream. Furthermore, the upstream power can be moved lower in frequency to avoid crosstalk and reduce the echo into the downstream.

ADSL2+ Handling of the Mixed CO and RT Deployment Scenario

The unbundling of the local loop has resulted in the potential for one operator to serve a particular area from a CO while another operator services the same area from a RT. Because the RT is potentially much closer to the end user than the CO, the result is that crosstalk from the RT can seriously degrade the performance of the ADSL system operating out of the CO. The degradation is dependent on the distances between the CO, RT and end user, and the coupling between the separate lines and could be negligible or severe.

ADSL2+, a high rate version of the new ADSL standard where the downstream bandwidth is doubled, offers a potential solution to the mixed CO and RT deployment scenario. The basic idea is to shape the ADSL2+ spectrum (potentially through simply turning off subchannels) such that the crosstalk in the lower frequencies is minimized. Because the loop is short, the RT deployed ADSL2+ system will still be able to achieve reasonable rates only using the higher subchannels. At the same time, the CO deployed ADSL system, which if on a long loop is constrained to the lower subchannels by loop attenuation, will see less crosstalk from the RT deployed ADSL2+ system and thus still achieve reasonable rates.

Crosstalk Minimization

Following the lines of being a good neighbour to existing systems (of which the mixed CO and RT deployment is an example), ADSL2 allows for some additional methods of crosstalk minimization. These include using a ceiling based power cutback mechanism, which removes power while tending to maintain the same data rate, L2 mode which reduces the transmit power when less data is present, and the ability to do iterative waterfilling for crosstalk minimization with full control over bit loading changes during showtime.

Compensation of Impairments

The ADSL channel is affected by a number of impairments which result in a reduction of the achievable data rate, sometimes to the point of making ADSL unusable from a business perspective, for example, it is difficult to market 32 kbps as broadband. As such, it is important to consider techniques to mitigate the impact of these impairments on the data rate during modem design. This section considers ISI, bridged taps and RFI and shows how appropriate modem design can be used to limit their impact on the ADSL system. Figure 2 illustrates the parts of the channel that are affected by each of these impairments.

Intersymbol Interference (ISI) and Programmable Tx and Rx Filters

ISI in the ADSL channel is a result of both the twisted pair medium and FDD filters. It has been observed that TEQs which perform well with the strong ISI around the transition band tend to produce notches in frequency. Because of noise spreading from the FFT, notches result in a loss in SNR and consequently data rate. So, the symbols passing through those transmit filters are effectively blended together in time. This interference typically extends beyond the cyclic prefix length in ADSL.

In addition, every CPE modem encounters different channel and system environments, so each one requires unique filter settings to attain optimum performance. These adaptive settings are essential in dealing with varying channel conditions, as well as the different DSLAM requirements for the transmit and receive channels. A technical solution is represented by features such as programmable coefficients for digital pre-compensation, band-split and power spectral density (PSD) shaping filters, as well as multiple corner frequencies for analogue filters – resulting in ideal system interoperability and greatly enhanced data rates.

Transmit Filter Profiling

The modem, or better the silicon embedded determines the transfer filter profile best suited to the characteristics of the involved DSLAM:

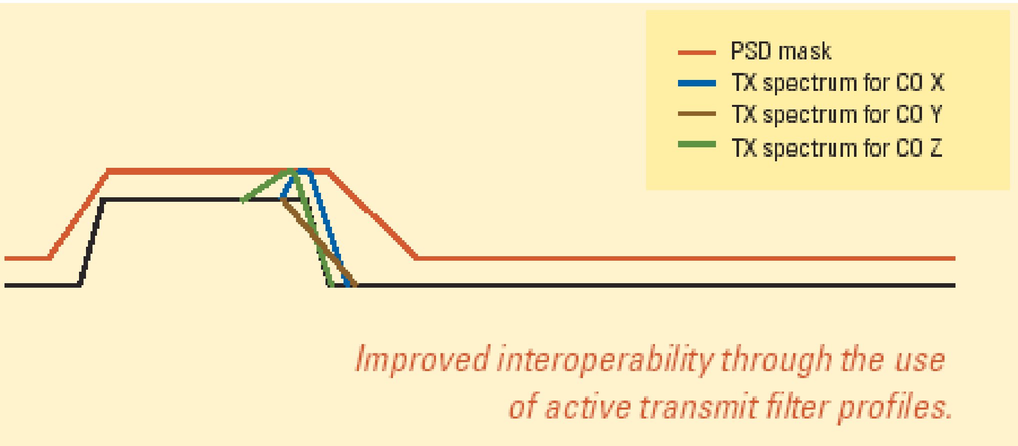

PSP Compliance

The silicon accommodates a wide range of PSD masks for Annex A, B, C, ADSL2/ ADSL2+, All Digital Loop (ADL), Reach-Extended ADSL (READSL) and more:

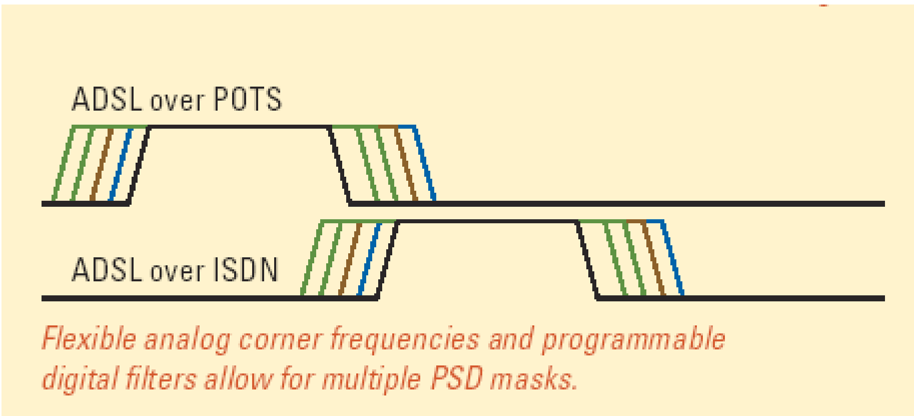

Digital Pre-Compensation

In addition to accommodating analogue process variations, programmable filters enable digital pre-compensation for transmit transfer function variations caused by impedance mismatch:

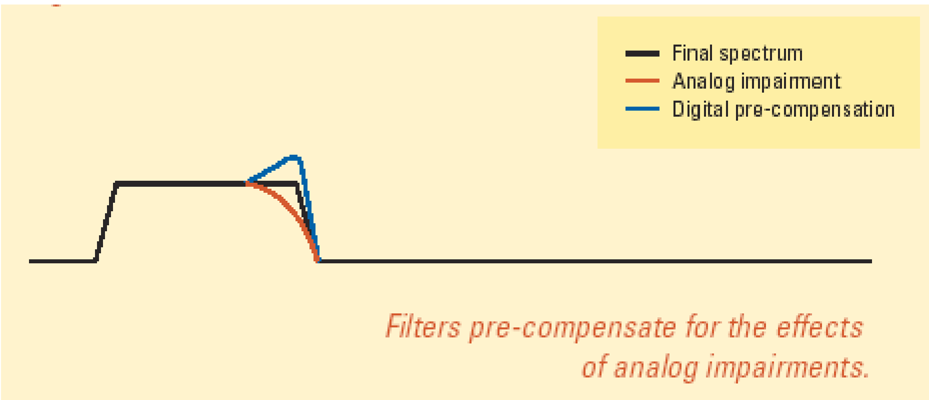

Echo and Bridged-Taps

ADSL is a full-duplex communication system, meaning that both upstream and downstream signals are transmitted simultaneously through the same transfer medium. It is the task of the receiver to separate the locally transmitted signal (echo) from the far-end received signal.

Bridged taps refer to multiple distribution cables connected to a single feeder cable. Only one of the distribution cables is connected and the others are left open. While this architecture provides flexibility for the operator in allocating lines, bridged taps tend to create nulls in the channel and cause problems with impedance matching. Depending on the architecture, home wiring can have a similar impact.

The reflections of the transmitted signal caused by bridged taps result in a large echo component in the received signal. Even if the ADSL system is operating in a FDD configuration, the increased echo (if not compensated for) results in a reduction in the data rate. This is because on longer loops the echo power can be larger than the received signal power and effectively limit the gain settings of the receiver and thus increase the modem's effective noise floor. Additionally, spreading from the FFT allows one band to bleed over into the other band and thus look like an additional noise source.

While using sharp bandsplit, filters can help reduce the amount of echo available for spreading, the drawback is that they can create equalization problems for the other receiver. Additionally, they do not address the modem noise floor issue. As such, in dealing with the additional echo caused by bridged taps it is reasonable to use a two part approach.

First, to optimize the receiver dynamic range, the hybrid must adapt to the varying reflected line impedance caused by varying loop topologies. In the simplest implementation this can be achieved by using multiple hybrids tuned for different loop topologies.

For the component of the echo not removed by better hybrid matching, an echo canceller (EC) can be used to remove the residual echo signal. The basic EC structure assumes that the echo is a filtered version of the transmitted signal. Echo cancellation is then performed by subtracting the output of a filter (identified during training to approximate the echo channel) fed by the transmitted signal from the received signal. ADSL systems can be designed such that a traditional EC is used in the time domain, or the echo cancellation can be carried out in the frequency domain (using a form of cyclic echo synthesis).

Switchable/ Adaptive Hybrid

A hybrid acts as an analogue echo canceller that subtracts the transmitted signal from the receive path once the signal has been filtered by the transmit echo transfer function. In practical systems, this transfer function can only be approximated since it is highly dependent on the line impedance, which varies greatly depending on the loop characteristics. Loop characteristics are conditioned by changes in line diameter, RF interference, in-house wiring and many more.

For example, as defined in the ETSI standard, " ETSI has normalised the value of the reference impedance of the real-world line at 135 Ω for a wide range of xDSL performance and conformance tests, including ADSL tests. This value is considered as being a reasonable average of characteristic impedances (Z_0) observed for a wide range of commonly used European distribution cables." This definition exists worldwide for the different lines deployed, such as i.e. the 26 AWG (American Wire Gauge Standard for ADSL lines) in USA.

It is out of date to say that this reference impedance does not really match the "real-world" line, which is always affected by its loop characteristics and depending on the of the region in which the CPE modem will be deployed. *Switchable* Hybrids are designed specifically for the loop characteristics shown in the following figure.

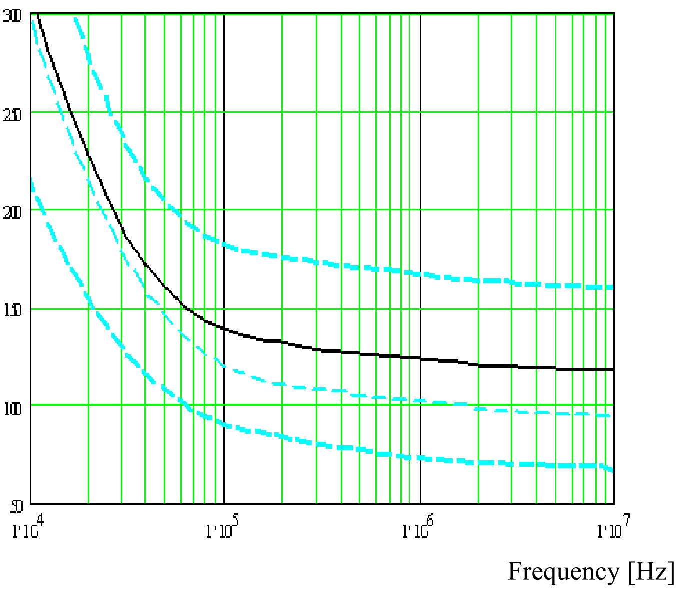

Figure: Theblack line represents the ETSI reference impedance and the blue lines represent the added impedance options (source: Texas Instruments Inc.).

The moment a consumer switches on an ADSL modem, it chooses the best impedance out of the available options during the training phase. In this way, it takes care of impedance variations thereby minimizing the potential existing mismatch with the ETSI standard loop in the real world. Unlike single fixed hybrids, which ultimately sacrifice data rate and loop plant coverage, the switchable hybrid solution offers an adaptive flexibility that greatly enhances the modem's performance.

The following are examples for the use of the switchable or adaptive hybrid:

Bridged-Tap Loop Impedance

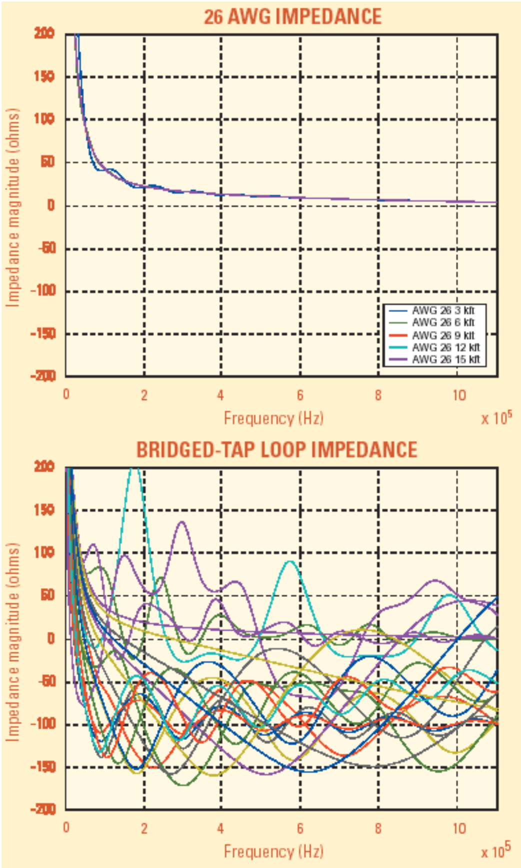

Switchable hybrids address a broad range of impedence scenarios, including straight (top) and bridged-tap (bottom) loops.

Performance testing with a fixed hybrid versus a switched hybrid yields results that prove significant improvements in down stream data rate when employing switchable hybrid technology as shown below.

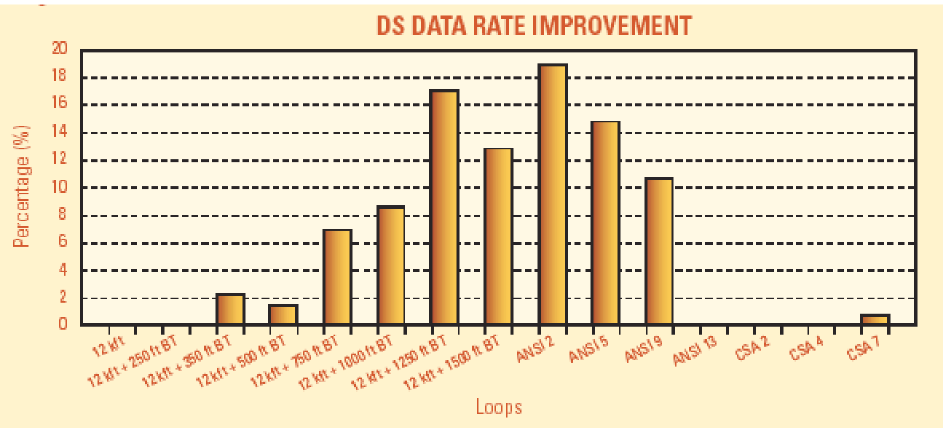

In-house wiring (SH = Switchable Hybrid function 'On')

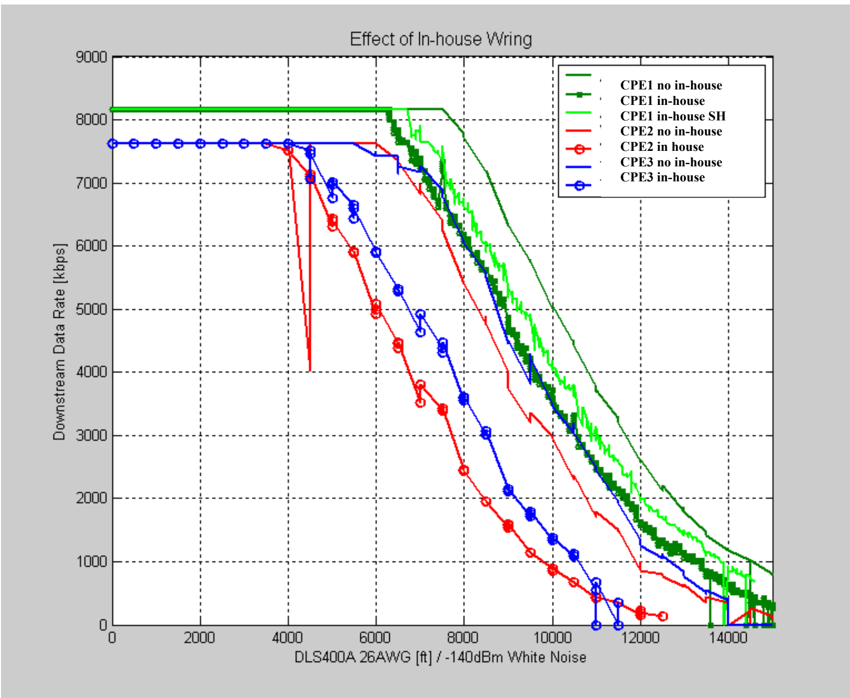

Echo Cancellation

Echo cancellation is used to eliminate the echo signals that were not removed by hybrid matching. The basic EC structure assumes that the echo is a filtered version of the transmitted signal. Echo cancellation is then performed by subtracting the output of a filter fed by the transmitted signal from the received signal. This filter is identified during training to approximate the echo channel.

Radio Frequency Interference (RFI)

RFI describes the intrusion of foreign signals, such as AM radio into the ADSL band, and is caused by an imperfect balance on the twisted-pair wire and modem front end. Spreading of the sinusoidal-like interferer from the FFT results in a potential data rate reduction across a large number of subchannels. Because of this, a number of algorithms have been developed for dealing with RFI.

While the primary criterion of the TEQ design is channel shortening, simple MMSE-based approaches for TEQ design also tend to put nulls at the locations of strong RFI sources. While there is some reduction in rate as a consequence of the null, in general the noise spreading is much reduced and the rate trade-off is appropriate. As such, if the RFI source is present during TEQ training, the TEQ can be used to compensate for the RFI.

Receiver windowing is a second approach that can be used for RFI compensation. Receiver windowing exploits the information in the cyclic prefix to form a window, which affects the noise but not the signal, provided that the channel memory is shortened to the length of the cyclic prefix minus the length of the window. The result is a window with sidelobes, which decay faster than those of the rectangular window. Therefore, even if RFI appears after modem training, the modem still exhibits a greater immunity to its detrimental effects. The trade-off is the additional constraints (or fewer degrees of freedom) in the channel shortening. An additional change to the new ADSL2 standard, allowing more receiver flexibility in the tone ordering, can further reduce the effects of RFI.

Nevertheless, typically no single equalization structure offers optimum performance for all possible channel conditions encountered in real DSL environments.

Configurable Equalization Structure

A configurable architecture is given that offers a clean implementation and high system cost utilization for four different structures using the same basic time domain equalizer and Fast Fourier Transform components, but with the addition of simple delays, downsampling and routing. The four possible equalization structures are:

> single-path > dual path > oversampled > double-rate

Use of these structures improves the equalization performance, enabling further fine-tuning using digital signal processor (DSP) algorithms.

ADSL systems use the cyclic prefix (CP) to minimize the effects of ISI. If the impulse response of the effective channel (including transmit and receive filters) is shorter than the cyclic prefix, no ISI occurs. However, this is not the case in most practical systems. Therefore, a TEQ is used to shorten the effective channel length.

Dual-Path TEQ

A dual-path equalization architecture allows the equalization structure to be optimized for different parts of the communication channel. For multicarrier systems such as ADSL, this allows one equalizer to be optimized for the transition band (where high ISI and echo are strong) and a second equalizer for the remainder of the band (where mild ISI, low echo, and possible RFI dominate the performance). Afterwards, the outputs of the two paths are compared, and the better of the two is selected as the output for the subchannel.

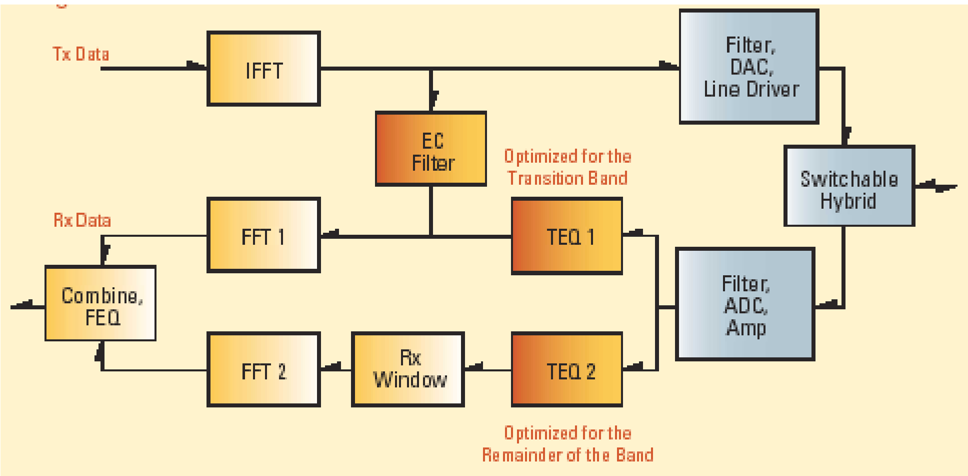

Dual-path receiver configuration.

Receiver Windowing

Receiver Windowing picks up after the TEQ in compensating for RFI. Receiver windowing exploits the information in the cyclic prefix to form a window with sidelobes that decay faster than those of the rectangular window. Therefore even if RFI appears after training, the modem still exhibits a greater immunity to its detrimental effects.

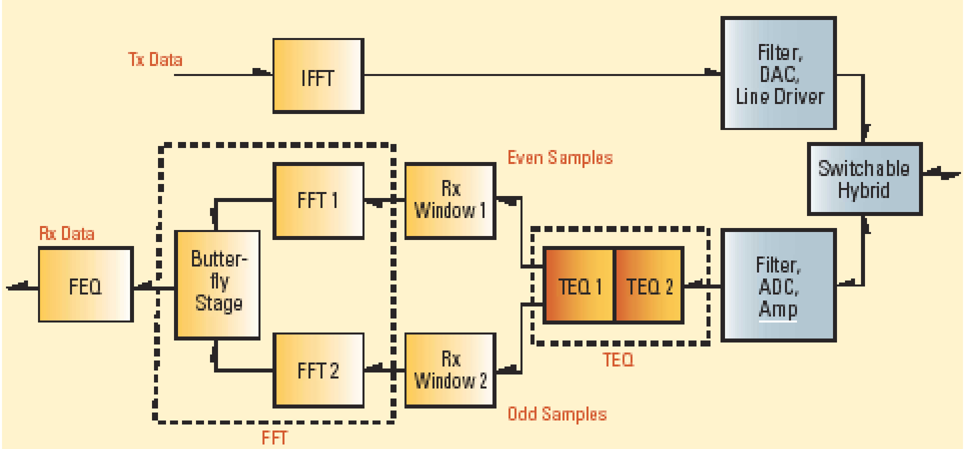

Double-rate receiver configuration for ADSL2+.

Summary

There are several techniques such as echo cancellation, time domain windowing, configurable time domain equalization (TEQ) and others to compensate for several of the line and spectral conditions that have prevented operators from providing service to all households. These techniques will not compensate for all line conditions but for items like bridge taps and radio frequency interference, the sum of these techniques will effectively increase the carrier service area (CSA).

Within the DSL spectrum there are lower and upper frequency tones. Normally there is a buffer area between these tones to ensure that there is no echo. However there can be improved performance with overlapping since you can take full advantage of the lower tones, which provided a longer reach and increased rates, this does impact echo and requires enhanced echo cancellation to compensate.

Most DSL solutions also use time domain equalization (TEQ) filtering techniques for the upper tones to compensate for line disturbers. There are also some that have configurable TEQs. In dual TEQ configuration, one path is assigned to the upper tones and one for the transitional tones where there is overlap to increase the use of the lower tones in the spectrum.

There are other anomalies out in the network such as RF interference (RFI) that can hit the tones. Some modems today address this problem through the use of time domain windowing. If the RFI hits exactly on a tone, it will wipe that one out and there is limited impact. However, if the RFI hits slightly off a tone it will spread out and take out portions of neighbouring tones, causing transmission problems. Time domain windowing limits the number of tones affected by RFI.

There are also hybrids included to detect other line conditions such as bridgetaps (most commonly wire slices near or in the home) and adjust the transmission to get the best performance out of the line.

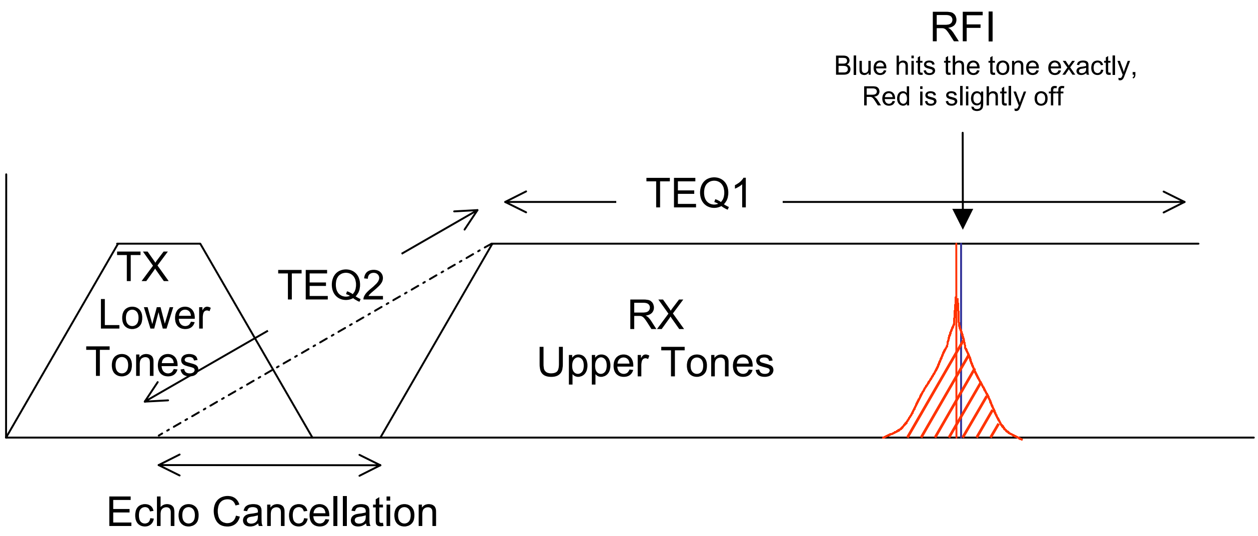

Conclusions

Increasing the coverage area of ADSL modems can be achieved by a combination of flexible modem design and new ADSL standards. This article summarized some methods to compensate for common impairments, which along with the flexibility in the ADSL transmit spectra can be used to improve the data rates at all loop lengths. An efficient design for an ADSL2+ system based on a modem designed for flexible impairment compensation was provided.

References

[1] Arthur Redfern, Fernando Mujica and Murtaza Ali, “Expanding the DSL Service Area”

[2] Dynamic Adaptive Equalization™ – DAE Tech Brief, Texas Instruments Incorporated.