

HDSL2 Tutorial: Spectral Compatibility and Real-World Performance Advances

Dr. George A. Zimmerman
Chief Scientist, PairGain Technologies

(714) 481-4549 george_zimmerman@pairgain.com

June 25, 1998

Abstract

This tutorial is intended to provide the background for understanding OPTIS HDSL2 and a summary of the new HDSL2 technology. There are three key themes:

1.) PairGain's vast field experience with DSL systems, far outweighing our nearest competitors, gives us a unique, practical perspective on DSL technology advancement. Field experience with DSL systems translates into principles which produce robust, reliable systems which add value to the customer, not just laboratory curiosities or specsmanship. In particular, deployed, real-world gains are often different from those shown in small-scale lab exercises. These result in a need to evaluate new technologies in a variety of noise environments, where the noise environments are often particular to the technology being proposed.

2.) Our real-world experience has been focused the development of extended reach technologies. It has prevented potential deployment disasters when solutions focused merely on the standards groups' test environment.

3.) PairGain has been and is leading the T1E1 standards development of HDSL2. We developed the OPTIS (Overlapped PAM Transmission with Interlocking Spectra) system which is the agreed line code for ANSI HDSL2. PairGain formed the standards coalitions necessary to bring the OPTIS proposal into agreement for the ANSI HDSL2 standard. The solution is the result of trading off the requirements of spectral compatibility with robust performance. The key characteristics of OPTIS are present in a most alternatives presented to date. In limiting environments it represents near-optimal performance, bringing significant advancement to symmetric transmission in real-world operating environments.

PAIRGAIN
THE COMMUNICATIONS COMPANY

Outline

- ◆ Overview of ANSI HDSL2 Problem & Background
 - Crosstalk noise: differences between ANSI & ETSI
 - (Spectrally) Symmetric Echo-Cancelled Transmission:
 - Effects of coding on usable bandwidth & latency
 - Spectral Compatibility: limits on transmitted energy
- ◆ Characteristics of the ANSI HDSL2 solution (OPTIS):
 - Spectrally Asymmetric Solutions
 - Shaping PSDs for Compatibility & Performance
 - Robustness in Mixed Noise
 - “Interlocking Spectra”: Spectral folding in DFES
 - 1-D Trellis coding with a flexible encoder
- ◆ Current issues & Extension to ETSI & multi-rate operation

HDSL2 Tutorial - June 1998

2

1. Introduction & Problem Overview

Problem Overview

PAIRGAIN
THE COMMUNICATIONS COMPANY

Technology Advancement Principles:

- ◆ Observable value added to the customer
- ◆ Robust, real-world, reliable performance gains
 - Not just taking advantage of “noiseless” spec's or lab environments
- ◆ Cost effective, deployable solutions
 - Must be implementable, from the silicon to systems levels
- ◆ Compatibility with existing services
 - Maintain performance of new technology and existing

HDSL2 Tutorial - June 1998

3

Instead of simply adding to the confusion currently in the xDSL market by offering marginally better solutions, it has been our goal to develop HDSL2 which meets certain principles of in advancing technology. These are:

- Observable value added: in order to be worthy of consideration, performance gains (reach or bandwidth) must be significant. For both the equipment vendor & the service provider it is not worthwhile to constantly change technology for small gains in bandwidth or reach (10-20%). An

example of this is the constant changes in ADSL & RADSL data rates.

- Robust real-world performance gains: Pairgain's experience in deploying HDSL systems had made us extremely sensitive to the issues involved with real-world performance. Our advancements must hold up in the real world and not require our customer to manage disturbers in their loops. We could claim near 18 kft (24 AWG) 2-pair HDSL performance with 2B1Q today, but this would be by re-specifying loops. There are many examples of such "specsmanship" in the industry today where the technology hasn't really been advanced, the burden has simply been passed to the customer to reduce the noise environment in their loops.
- Compatibility with existing services: At Pairgain, we realize that the installation of new systems must not interfere with the existing services or important planned services in the network. As a result, we carefully evaluate spectral compatibility of our services and present them for peer-review in the T1E1standards forum.
- Cost-effective deployability: Pairgain's HDSL success was based on coming directly to market with a deployable solution. Our philosophy is to do the same with new services. In some instances this may not make us first to market, but it avoids wasting the customers' time & money on evaluating development-stage or prototype systems.

Problem Overview



HDSL2 - Future Symmetric Transmission

- ◆ CSA-reach single-pair T1
 - Provides same performance and spectral compatibility for T1 service on 1 pair that existing HDSL provides on 2 pairs
 - Works even in the presence of repeated T1 and harsh real-world crosstalk mixtures
 - Low latency coding required to meet T1 application
- ◆ Advanced transmission technology
 - Beyond simple 2B1Q, CAP, DMT or other echo-cancelled symmetric transmission schemes.
 - High-gain trellis-coded 16-level PAM with advanced spectral shaping
 - Beyond common trellis codes and transmit spectra

HDSL2 Tutorial - June 1998

HDSL2 is NOT second-generation HDSL. It is NOT a replacement to 2B1Q HDSL. What HDSL2 IS is a complement to existing HDSL. It fills the niches for single-pair T1, and some reach-extended 2-pair HDSL, beyond CSA, without doublers. When really long-reach (beyond 18 kft 24 AWG or 15 kft 26AWG) becomes

important, however, use of existing 2B1Q HDSL with doublers is probably the way to go.

HDSL2 is responsive to customers. Carriers requested that ANSI HDSL2 have the following operational characteristics.

- An advanced single-pair T1 system should go CSA reach to have real value
- It should tolerate all the normal loop disturbances (crosstalk, bridged taps, etc.) that HDSL tolerates
- It should not interfere more with existing services than existing HDSL does.
- It should be as cost-effective as conventional HDSL
- It needs to work in the real-world!

Through PairGain's leadership, true HDSL2 technology will truly advance the transmission technology. It goes beyond simple line codes like 2B1Q, CAP, or DMT, and uses advanced concepts in spectral shaping and error-correction codes to achieve performance near to the theoretical limit - Shannon Capacity.

Problem Overview



HDSL2: Why is it hard?

- ◆ HDSL already performs near to optimum levels
 - Large effort required for significant performance improvement
- ◆ HDSL2 requirements are near fundamental limits
 - Latency, power, interference to and from other services
- ◆ Powering capabilities limited by number of pairs
 - Single-pair HDSL2 line-powered doublers unlikely, limits HDSL2 to niche
- ◆ Increasingly diverse services in loop plant complicates spectral issues

HDSL2 Tutorial - June 1998

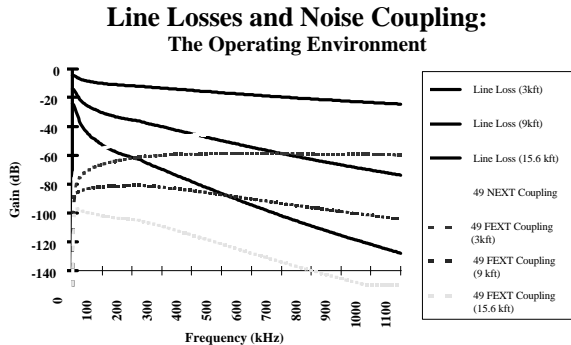
HDSL2 is technically a very challenging problem. HDSL2 is more than simply "single-pair T1". While it is easy to create single-pair T1 & E1 systems by increasing the rate of existing HDSL technologies such as CAP or 2B1Q, giving these systems the robust performance that will hold up under the real-world challenges in the local loop requires careful testing and experience. Multiple vendors have claimed single-pair T1 systems as "HDSL2 compatible". Often these systems claim similar reach to HDSL2. In most cases, however, the reaches claimed are in minor or minimal crosstalk situations, or under particular laboratory test scenarios. Deployment at these reaches in the real world makes the user of the technology vulnerable to interference from other circuits either at the time of installation or at some future date.

2. Background

2.1. Noise Environment

Background: Noise Environment

PAIRGAIN
TECHNOLOGIES, INC.



HDSL2 Tutorial - June 1998

In the North American model, there are three distances of interest: “Revised Resistance Design” (RRD) loops (< 1300 ohms, typically less than 18,000 ft (18 kft)), which describes more than 85% of all residential loops, “Carrier Serving Area” (CSA) loops (< 9000 feet, 26 AWG, or < 12,000 feet 24 AWG), which describes 50% of all loops, and 75% of all business district loops, and “Distribution Area” (DA) loops (approx. 3000-4000 feet, 26 AWG). HDSL2 development was targeted for 1.544 Mbps rates on CSA distance loops. For reliable deployment, this reach must be achieved even in the presence of heavy crosstalk.

Because crosstalk in the real world is a result of the signals deployed in the local loop, it is important that the noise environments used for testing be sensitive to new technologies being deployed. Noise must be sensitive not only to measurements made of the existing loop, but also to the effect of the new signals in the loop.

This is a fundamental difference between ANSI and ETSI noise models, and is important to the proper evaluation of HDSL2 systems:

ANSI uses a model sensitive to the transmit PSD proposed with a non-periodic (or very long period) Gaussian noise model. This model, while accurate to the local loop environment, makes design of repeatable test equipment difficult.

ETSI has traditionally used a fixed noise model with a specified method of generation for evaluating HDSL technologies. This greatly facilitates design of test equipment and repeatability of results. However, it gives an unrealistic advantage to systems with higher PSDs in the high frequency end of the passband, and its periodicity gives an unrealistic bias to performance measurements and masks the advantages of error correcting codes.

Background: Noise Environment

PAIRGAIN
TECHNOLOGIES, INC.

Key Differences: ANSI & ETSI

- ◆ ANSI:
 - Noise environment varies with signals in loop
 - Allows for accurate modeling of effect of new signals on themselves
 - Complicates test equipment & analysis greatly
 - Does not specify generation method
- ◆ ETSI:
 - Noise environment fixed based on measurements
 - Allows for predictable, repeatable testing & test equipment
 - Specifies generation method (periodic synthesis by tones)
 - Does not generally model effect of new signals on themselves

HDSL2 Tutorial - June 1998

Crosstalk noise couples to the line either at the receiver end as Near End Crosstalk (NEXT), or along the far end of the line as Far-End Crosstalk (FEXT). Existing and proposed signal power spectral densities are coupled through the appropriate NEXT & FEXT models. **Cooperation with service providers is essential in determining the appropriate disturber sets to use**, and this has been important to the specification of ANSI HDSL2.

The simplified (ANSI) model for NEXT coupling is: $|HNEXT(f)|^2 = X49 * f^{1.5} * (N/49)^{0.6}$, where X49 is the crosstalk coupling constant for 49 disturbers in a 50 pair binder, f is the frequency in Hz, and N is the number of disturbers in the binder group. X49 is approximately 8.82E-14. (ANSI) FEXT coupling is typically modeled as: $|HFEXT(f)|^2 = k * l * f^2 * (N/49)^{0.6}$, where N is the number of disturbers, l is the coupling length, f is the frequency of transmission, |Hchannel|² is the channel frequency response, and k is the FEXT coupling constant, approximately 8.00E-20 / (ft*Hz²), for the 1% worst-case coupling of 49 disturbers in a 50 pair binder.

It is important to note that these crosstalk models represent the 1% worst-case coupling scenario for all the packings of N disturbers into a 50-pair binder group. It is fairly common to see less crosstalk than this (the other 99%), especially if there are fewer than half of the pairs in the binder group carrying the disturber in question; however, field measurements have shown that the model is accurate for the 1% worst case. This fact is important when one considers differing cable binder sizes (say 25 pairs, in residential loops or 600 pairs near a telco central office). The worst-case crosstalk rarely gets worse than the 1% worst-case for 49 disturbers in a 50-pair binder, and is often better. The important consideration is not the number of disturbers, but rather the proportion of the cable which is full. Basically, most of the crosstalk between pairs comes from about 6 pairs.

Thus, it is possible, but extremely unlikely (1 in about 14 million) that 6 disturbers in a 50 pair binder may be as bad as all 49. Likewise, 100 disturbers in a 600 pair cable are in the 1% worst case likely to look more like 8 disturbers in a 50 pair binder (and are 99% likely to be better). However, 24 disturbers in a 25 pair binder are very likely to be just as bad as 49 pairs in a 50 pair binder.

It is natural to consider generating many such noise sequences and evaluating the results in the laboratory. We show in this section that the results will be highly repeatable, but, unless the number of noise samples is very large, they will test systems in a biased, optimistic fashion, especially for Gaussian noise. This bias (roughly 2 dB) is evident in the ETSI HDSL noise model. The reason the bias occurs is that for large sample sets the distribution of the maximum value of a set of random variates taken as a random variable in itself (over different sequence trials) will be approximately Gaussian, with mean and variance given below:

$$m = F^{-1}(1/N)$$


$$s^2 = \frac{N-1}{N^3 f(m)^2}$$

where $F(x)$ is the cumulative distribution function, and $f(x)$ is the probability density function from which N samples are drawn. It is readily apparent that for large N the distribution of the maximum value will be constrained due to the $1/N$ dependence of the standard deviation. For 8192 baud sampled Gaussians, the mean value of the peak-to-rms ratio would be 3.85, and the variance would be 0.256, or 0.5 standard deviations.

Background: Noise Environment



ETSI Noise Model: Effect of Periodic Noise Sequences



HDSL2 Tutorial - June 1998

Thus the peak-to-rms ratio only varies more than 26% (1 dB) in less than 2% of the sequences. Typically the mean value of the expected noise peak is short of the $1e-7$ BER point by 2.4dB, and thus produces repeatable,


overly optimistic margin results. The figure in the slide above, shows the probability density of the maximum (absolute) value of a set of N Gaussian random variables for various values of N . Testing with an artificial noise source, such as this can lead to perceived performance benefits which, while repeatable, do not hold up in the real world. Real HDSL2 technology & PairGain's existing 2B1Q HDSL are tested for reach in environments of Gaussian noise with very long periods to avoid this problem.

2.2. Symmetric EC Transmission

Background: Symmetric EC Transmission



Symmetric Echo Cancelled Transmission: Received SNR



HDSL2 Tutorial - June 1998

Widely deployed xDSL systems like ISDN and HDSL transmit the same spectrum in the upstream and downstream direction. The bi-directional, "spectrally symmetric" transmission is accomplished by echo-cancellation. This greatly simplifies the worst-case performance prediction. Given that the transmit power level is high enough (and it is, by specification), self-crosstalk will limit performance. Self-crosstalk is crosstalk from the transmitters of other ISDN or HDSL lines into the receiver of the desired ISDN or HDSL line. ISDN and HDSL are both specified to operate in a reliable, worst-case environment, known as 1% worst-case 49 self-NEXT (Near-End Crosstalk) and FEXT in a 50 pair binder. Aslanis and Cioffi (IEEE Trans. on Comm. 40:2) examined the achievable rates for symmetric echo-cancelled NEXT limited transmission on CSA loops, showing that approximately 400 kHz of useful bandwidth was available (if high coding gain were used). Figure 3 shows the received SNR for symmetric echo-cancelled transmission in 49 self-NEXT+FEXT for the three classes of loops considered.

These SNR plots are directly related to the carrying capacity at each frequency.

$$C(f) = \log_2 \left(1 + \frac{S(f)|H(f)|^2}{N(f)\Gamma} \right)$$

where $S(f)$ is the received signal power spectral density at frequency f , $|H(f)|^2$ is the magnitude squared of the wireline loop transfer function, $N(f)$ is the noise power spectral density at the receiver, as before, and is the effective SNR gap. For coded systems, SNR gap is defined as (9.75 - (effective coding gain)) dB. For transmission where SNR margin is required, the effective SNR gap is increased by the desired margin, and is defined as = 9.75 - (effective coding gain) + Margin (dB).

Background: Symmetric EC Transmission



Symmetric Echo-Cancelled Transmission

Table 1: Usable Bandwidth for Symmetric Echo-Cancelled Transmission (49 NEXT)

Loop Length 26 AWG	Uncoded (SNR > 10 dB)		6 dB Coded (SNR > 4 dB)	
	3000 ft (DA)	1430 kHz	1888 kHz	
9000 ft (CSA)	272 kHz	356 kHz		
13500 ft (RRD)	98 kHz	138 kHz		

Table 2: Achievable rates for Symmetric Echo-Cancelled Transmission (49 NEXT case)

Loop Length 26 AWG	Uncoded Transmission		6 dB Coded Transmission	
	3000 ft (DA)	5574 kbps	8195 kbps	
9000 ft (CSA)	931 kbps	1401 kbps		
15600 ft (RRD)	356 kbps	536 kbps		

HDSL2 Tutorial - June 1998

The inclusion or exclusion of coding gain modifies the useful bandwidth accordingly. HDSL2 uses powerful codes with more than 5 dB gain. If we consider the useful bandwidth as the region where greater than 1 bit per second per Hz is achieved, uncoded systems can generally make use of regions with SNR greater than 10 dB (due to the 10 dB SNR gap), whereas systems employing can often use SNRs down to 4 dB. Extending the transmission bandwidth beyond these bounds adds minimally to the carrying capacity of the line. Approximate usable bandwidths are shown in Table 1. Achievable rates can be directly computed from the SNR curves, by integrating the Shannon-Hartley capacity shown above. Table 2 shows these results when including the 6 dB margin usually used in telephone circuits.

Background: Symmetric EC Transmission



Performance Gain Required: Symmetric Echo-Cancelled Transmission (9kft 26 AWG, 1.5Mbps)

Bandwidth (kHz)	Margin at Capacity		dB From Capacity Required for x dB Margin on CSA 6 w/ 49 NEXT							
	Simplified		6		5		4.5			
	Next	Unger	Simp.	Unger	Simp.	Unger	Simp.	Unger		
400	7.65	8.4	1.65	2.4	2.65	3.4	3.15	3.9	3.65	4.4
355	7.45	8.15	1.45	2.15	2.45	3.15	2.95	3.65	3.45	4.15
300	6.95	7.65	0.95	1.65	1.95	2.65	2.45	3.15	2.95	3.65
295	6.9	7.6	0.9	1.6	1.9	2.6	2.4	3.1	2.9	3.6
260	6.25	6.95	0.25	0.95	1.25	1.95	1.75	2.45	2.25	2.95

Bandwidth (kHz)	Gain Required for x dB Margin on CSA 6 w/ 49 NEXT							
	6		5		4.5		4	
	Simp.	Unger	Simp.	Unger	Simp.	Unger	Simp.	Unger
400	8.09	7.34	7.09	6.34	6.59	5.84	6.09	5.34
355	8.29	7.59	7.29	6.59	6.79	6.09	6.29	5.59
300	8.79	8.09	7.79	7.09	7.29	6.59	6.79	6.09
295	8.84	8.14	7.84	7.14	7.34	6.64	6.84	6.14
260	9.49	8.79	8.49	7.79	7.99	7.29	7.49	6.79

HDSL2 Tutorial - June 1998

Studies such as the one shown above demonstrate that the use of coded modulation gain alone would require realized performance gains of 7 to 8 dB.

While asymptotic gains are common in evaluating codes, HDSL2 performance required the actual measured gain of the code to be high. As a result, most common coding strategies topped out around 5 dB.

Many coding schemes were proposed showing measured gains of 5 to 5.5 dB.

Among the schemes considered were:


- 1) Trellis coded modulation with:
 - Long constraint length (K=11-19) convolutional codes with sequential or other suboptimal decoding (speed / latency constrained to 5.1-5.5 dB gain)
 - Long constraint length (K=9) convolutional codes with Viterbi decoding (performance constrained to 5.1 dB)
- 2) Turbo trellis-coded modulation (latency constrained to approximately 5 dB)
- 3) Multi-Stage coded systems:
 - Using Conventional Convolutional codes (5 dB)
 - Using Turbo Codes (not fully examined)

Additional performance gain required moving away from Spectrally Symmetric Echo-Cancelled transmission.

2.3. Spectral Compatibility

Background: Spectral Compatibility

PAIRGAIN
THE CUMULATIVE GAIN



HDSL2 Tutorial - June 1998

This chart shows just some of the existing and planned services in the (North American) local loop. With such a wide range of frequencies, (everything from 40 kHz for ISDN to 772 kHz T1), it is easy and tempting to consider specifying systems to have reduced performance with one or another signal as a crosstalk disturber. However, useful, practical, deployable cost-effective systems must answer to the challenge of coexistence with this complicated “soup” of signals which currently inhabit the local loop.

Most importantly notice that NEXT & FEXT from a new signal will impede transmission of these existing signals. If the new signal extends in full-power bandwidth substantially beyond 400 kHz at CSA reach, it will cause the same type of SNR pinch-off in other signals that it might in itself under the symmetric self-NEXT condition described earlier. Particularly important was the relation of the upstream of HDSL2 to the downstream of ADSL.

Echo Cancelled ADSL’s downstream provides a near-end crosstalk source across the entire bandwidth from 25 kHz to 1.1 MHz. For PSD levels near -40 dBm-Hz, this limits the useful bandwidth for systems like HDSL2 to less than approximately 280 kHz. Non-interference into existing services requires HDSL2 to not degrade ADSL service any more than ADSL is currently affected by HDSL (2B1Q) service. The consideration of interference into Echo-cancelled ADSL also limits the usable bandwidth of the HDSL2 upstream to approximately 280 kHz, depending slightly on the exact shape of the HDSL2 upstream spectrum. In order to maintain this level of spectral compatibility, out-of-band energy must be severely limited (-40 dB or lower). This limitation implies the use of oversampling in the HDSL2 analog front ends.

Similarly, interference from AMI T1 into the downstream limited the usefulness of transmission

above approximately 600 kHz, where the T1 signal reaches its full bandwidth.

3. Key Characteristics:


Given all the considerations of the previous discussion, it is clear that in order to truly perform, HDSL2 must go beyond the technology employed in simple-coded symmetric transmission systems. This section describes some of characteristics required for true HDSL2 performance.

3.1. Spectral Asymmetry

Characteristics: Spectral Asymmetry

PAIRGAIN
THE CUMULATIVE GAIN

Spectrally Asymmetric Signals: FEXT-limited Transmission



HDSL2 Tutorial - June 1998

It has long been known that if one divides the upstream and downstream in either the frequency or time domain, one will be subject to the much less harsh noise environment of FEXT. This principle has been used for ADSL and VDSL systems extensively; however, FEXT-based systems are typically limited by interference from other services. In practice, analysis of spectrally asymmetric signals such as those used for HDSL2 FEXT and NEXT limitations needed to be considered jointly. It was essential to define the signals that HDSL2 must coexist with disturbance from. Accuracy in this definition was extremely important, and at one point an incorrect assumption (that in the worst case, crosstalk would be of the same type of disturber) nearly led to disaster. It is imperative that designers of the new systems keep real environments found in the local loop in mind, lest they be substantially worse disturbers than ordinarily considered, and not make too many assumptions on deployment practices.

3.2. Robustness in Mixed Noise

Characteristics: Robustness In Mixed Noise



xDSL Limiting Environments

- ◆ Unknown Limitations by mixtures of other services cause:
 - Unexpected trouble calls (service interruptions)
 - Unnecessary loop management
 - Failed installation attempts
 - Additional spares
 - Special engineering rules
- ◆ Deployed xDSL systems are limited by interference from similar systems (self-NEXT)
 - This is what makes HDSL robust and reliable

HDSL2 Tutorial - June 1998

14


Unless properly considered, unknown limitations caused by interference from mixtures of crosstalk (e.g., self crosstalk + T1 crosstalk) might cause HDSL2 T1 circuits to go down, or cause the need for additional spares and waste money due to failed installation attempts. When such limitations become apparent the only solution, other than switching technologies, would be to employ special engineering rules or to practice loop management which would have been unnecessary if the system had been designed properly in the first place. These actions cost service providers time & money, and, as a result, we strive to provide robust, reliable solutions and adequately specified xDSL systems.

Going from analysis (or test) to a deployable system requires careful consideration as to the appropriateness and completeness of test environments. Pairgain's experience in deploying xDSL has taught that systems such as HDSL are extremely robust simply because they are self-NEXT limited. One can test them in noise from like sources and get a good worst-case performance measurement. Newer systems, such as ADSL, RADSL, and HDSL2 are not typically self-NEXT limited. Therefore one must consider many test environments in determining the worst case environments. Because mixtures are often worst-case, the number of such combinations can be extremely large and is not always obvious.

Although it complicates testing considerably, mixed crosstalk is often much more common in the loop environment than the homogeneous environments traditionally used for disturbance planning. As a result, robustness in mixed crosstalk is very important. There are several approximate models for mixed crosstalk. They vary from pessimistic models (1.2 dB too high in limiting cases) where the 1% worst-case power for each disturber to optimistic models (1.2 dB too low in limiting cases) where the disturber power is added prior

to passing through the noise coupling model. Accurate prediction of worst-case mixed crosstalk is dependent on the frequency content of the signal being disturbed, and hence becomes an iterative process in design; however, because the variation from pessimistic to optimistic is low enough (2.4 dB), we can safely use one of these simple models to avoid any potential catastrophic failures in mixed crosstalk.

Characteristics: Robustness In Mixed Noise



HDSL2 Tutorial - June 1998

This chart graphically depicts the failure of POET, one spectrally asymmetric proposal for HDSL2. Each chart shows received signal-to-noise ratio of the POET downstream as a function of frequency. In order to properly operate, the average of the SNR across approximately the bottom 600 kHz of the band must be at the level of the dotted line. Note that POET appears to work well in the two specified environments shown at top. POET power levels and spectral shaping were optimized for these environments, and the results suggested that good performance and spectral compatibility could be achieved in both self-disturbance and in T1 disturbance. However, notice that in self-disturbance a significant amount of the (downstream) SNR comes from frequencies greater than 300 kHz (where the POET downstream does not overlap the upstream). Because of spectral compatibility constraints in the design of the POET system, the received downstream power above 300 kHz is very small, but this is tolerable because in self-NEXT only there is little noise there. As might be expected, in T1 disturbance, there is little SNR above 300 kHz, but SNR is adequate below 300 kHz, and hence good margins can be obtained in pure-T1 environments when no other signals are present.

When as little as 1 T1 signal is added to the self-NEXT mixture (and an optimistic model is applied for crosstalk mixtures), the noise due to the 1 T1 signal deflates the good SNR region above 300 kHz, sinking margins in the POET system by as much as 9 dB. This

behavior makes POET unsuitable for deployment in loops containing T1, since even the slightest mixtures of POET and T1 could cause the HDSL2-based T1 lines to go down. Note that there are other cases of crosstalk mixtures which are damaging to POET, however, T1 mixtures are most dramatic. Note also that coexistence of T1 and HDSL terminations at the customer premises regularly occur in the loop environment. In fact, PairGain has both HDSL and AMI T1 signals terminating within our own building.

4. OPTIS HDSL2 Solution

OPTIS HDSL2 Solution



OPTIS HDSL2 Solution

- ◆ OPTIS agreed as basis for ANSI Standard (P.A. Dec 97)
 - Overlapped PAM Transmission with Interlocking Spectra
 - Spread downstream avoids noise (similar to DMT ADSL)
 - High gain coded modulation and precoding for robust transmission
 - Low-latency 1-D trellis-coded modulation
- ◆ Robust even in crosstalk mixtures
- ◆ Shaped for spectral compatibility
 - Upstream PSD "bumped" to improve performance without increasing bandwidth

HDSL2 Tutorial - June 1998

PairGain's OPTIS HDSL2 solution uses advanced spectral placement and frequency planning techniques learned from ADSL to provide robust performance in mixtures. In June 1997, PairGain introduced OPTIS to the ANSI T1E1.4 solution, after working carefully behind the scenes to put together a partnership with Level One Communications & ADC Telecommunications PAM to form the new HDSL2 standard. These three then formed the nucleus of a larger group (including, at the time, Globespan & Paradyne) supporting the OPTIS proposal in the ANSI T1E1.4 committee, and meeting the real-world HDSL2 requirement without much of the traditional DSL line-code infighting.

PairGain synthesized multiple ideas generated during the HDSL2 effort into the OPTIS proposal, and further refined the baud rates and modulation techniques to produce robust systems capable of good performance even in mixed crosstalk. In earlier contributions PairGain presented overlapped spectra based on oversampled CAP/QAM and symmetric symbol rates (OverCAPPED transmission); Adtran presented results on performance and spectral compatibility of a partially overlapped echo-canceled transmission using PAM and asymmetric symbol rates (POET PAM); PairGain and


Level One demonstrated that both the POET system and the OverCAPPED system were subject to substantial degradation in mixed crosstalk environments; ADC and Level One separately showed how boosts could be added to the upstream PSD to improve margins in the presence of EC-ADSL crosstalk without significantly impacting spectral compatibility.

4.1. Asymmetric Spectra

OPTIS HDSL2 Solution: Asymmetric Spectra



ADSL, HDSL, T1, HDSL2



HDSL2 Tutorial - June 1998

OPTIS HDSL2 uses a unique spectrally-shaped waveform which overlaps the upstream and the downstream in an interlocking fashion. As seen on this chart, this allows both directions of transmission to make use of the sweet spot in the transmission characteristics of the line occupying frequencies between normal HDSL and AMI T1. At the same time OPTIS avoids self-interference by using a new concept called "interlocking spectra". The next slide shows how OPTIS folds the received signal to "interlock" the high-SNR region of the downstream to mask the affect of interference in the middle of the band.


4.2. "Interlocking Spectra" & Spectral Folding in DFEs

OPTIS HDSL2: Spectral Folding



PairGain OPTIS HDSL2 in Harsh Real-World Mixtures

Received SNR for PairGain HDSL2 in Mixed Crosstalk:
39 Self NEXT/FEXT +10 T1 NEXT, CSA4



HDSL2 Tutorial - June 1998


This slide shows the SNR profile for operation of the initial OPTIS proposal in a far more severe mixed crosstalk environment than that which killed the POET system. In this example, crosstalk comes from a combination of 39 HDSL2 and 10 T1 signals in the same binder. Note that the downstream SNR folds to fill in a hole in the SNR spectrum below 300 kHz. What this means in simple terms is that the OPTIS HDSL2 system uses different regions of the frequency spectrum to support areas where the SNR may be weak. In this example the high SNR at 300 kHz would fold in to support the low SNR at 220 kHz. This gives the robust performance which, combined with the high-performance error-correction coding, allows for CSA reach even in realistic mixed crosstalk environments.

OPTIS HDSL2: Spectral Folding



SNR Folding in DFEs

- DFEs with fractionally-spaced feed-forward equalizers inherently have a baud sampling at the slicer
 - The feedback transfer function aliases as a result
 - Complex (QAM) and real (PAM) systems alias differently
 - Folding vs. Stacking



HDSL2 Tutorial - June 1998

20

For PAM signals, SNR is computed using an optimal DFE calculation for PAM, generally due to J. Salz:

$$SNR_{out} = 10 * \log_{10} \left(\exp \left(\frac{2}{f_{baud}} \int_0^{f_{baud}/2} \ln(1 + f_{SNR}(f)) df \right) \right) \text{dB}$$

where $f_{SNR}(f)$ is the folded received signal-to-noise ratio, defined as:

$$f_{SNR}(f) = \sum_{n=-\infty}^{\infty} \frac{S(f + f_{baud} \times n) |H(f + f_{baud} \times n)|^2}{N(f + f_{baud} \times n)}$$

and $S(f)$ is the desired signal's (e.g., ISDN or HDSL) transmit power spectral density, $|H(f)|^2$ is the magnitude squared of the wireline loop transfer function, and $N(f)$ is the total noise power spectral density (crosstalk plus background noise) computed as described above. Nominally, SNR folding calculated out to 3 times the Nyquist rate, which has been sufficient for the xDSL signals used to date.

It is important to consider the nature of the folded SNR. As stated in Salz's work (BSTJ), this optimal relation holds for an optimized receive filter. One can show that the folded SNR relation is the result of a multi-source optimization, where each of the aliased terms are separately weighted.


CAP/QAM or PAM?

- Differences in spectral folding separated the PAM & CAP solutions for OPTIS dramatically (note chart and margins are for the original OPTIS)

Table 1: Uncoded Theoretical Margins, 30kb per dimension, CAP/QAM & PAM

Interferers	CSM - Upstream		CSM - Downstream	
	HDSL2 CAP/QAM	HDSL2 PAM	HDSL2 CAP/QAM	HDSL2 PAM
39 Self NEXT/FEXT	21 dB	20 dB	92 dB	135 dB
49 HDSL2 NEXT/FEXT	23 dB	22 dB	99 dB	128 dB
25 T1 NEXT	195 dB	194 dB	151 dB	160 dB
49 FDM ADSL NEXT/FEXT	68 dB	68 dB	161 dB	150 dB
39 EC ADSL NEXT/FEXT	17 dB	16 dB	166 dB	156 dB
24 T1 NEXT + 24 HDSL2 NEXT/FEXT	43 dB	42 dB	36 dB	11 dB
24 FDM ADSL NEXT/FEXT + 24 HDSL2 NEXT/FEXT	16 dB	16 dB	55 dB	123 dB

OPTIS HDSL2 Mixed Crosstalk SNRs: CSA, 24 T1 NEXT + 24 Self NEXT/FEXT



HDSL2 Tutorial - June 1998

21


Quite notable is the difference between the CAP/QAM and PAM margins, more than 4.5 dB in the case of the original OPTIS. This is due to the folding of the downstream SNR in the interlocking spectra. The chart shows the received SNR in the downstream direction as a function of frequency, along with the PAM and CAP/QAM folded SNRs. Note that since PAM DFEs alias SNR reflectively back at half the PAM baud rate, the high downstream SNR region beyond the upstream-downstream overlap folds in to fill the area of poor downstream SNR due to NEXT from the boosted upstream region. The use of this "interlocking" PAM folding thus provides substantial benefit to OPTIS HDSL2.

CAP/QAM and PAM differ in their folding, because CAP/QAM is processed at the slicer as an analytic signal. As a result, CAP/QAM does not fold reflectively, since the reflection is due to the "negative" frequency components of a real-only signal. CAP/QAM instead aliases only by translation (as opposed to reflection) about the CAP/QAM baud rate, equal to half the PAM baud rate for the same number of bits per dimension. Aliasing doesn't benefit the CAP/QAM signal as much as the PAM signal because the frequencies which fold to the SNR-poor high end of the passband are substantially higher in frequency. In fact, the portion of the CAP/QAM SNR which would fold into the SNR deficit region does not begin until 460 kHz, which is beyond the edge of the downstream passband in this case. One can extrapolate that even if the passband were extended for CAP/QAM the benefit would be minimal since the SNR folded in would be below 10 dB. It is noteworthy that this effect is not usually observed in CAP/QAM systems because they tend to use spectra with little energy above the folding frequency.

4.3. Trellis Coded Modulation

1-D PAM Trellis-Coded Modulation

- ◆ 1-D, rate 1/2, TCM used with Tomlinson precoding
- ◆ Powerful 5.1 dB trellis codes necessary to gain full margin
 - Several codes of equivalent performance & complexity presented
 - Example: 512-state PAM code: 5.10 dB gain with 216 usec latency



- ◆ Programmable encoder was proposed as a standards compromise
 - Allows for future performance enhancement & backwards compatibility

The study of trellis code candidates for HDSL2 would be a tutorial unto itself. Ultimately agreement centered around a traditional form of trellis-coded modulation:

1-dimensional, rate 1/2 trellis coded modulation. In this scheme, one of the 3 bits/symbol is protected with a powerful convolutional code. The remaining bits are transmitted unprotected. The result is a 16-level PAM signal conveying 3 bits of information per symbol. The two least significant bits (finest levels) are formed by the output of the rate 1/2 encoder, while the two most significant bits are raw transmit bits. A one-dimensional code was chosen for its superior latency characteristics, while giving equivalent performance to longer-latency multidimensional codes.

The use of a programmable encoder was part of a standards compromise to break the logjam surrounding several technically equivalent 1-D coding techniques. HDSL2 transceivers are required to meet a set of stringent performance requirements to guarantee performance. In the programmable encoder structure, the coefficients of the code would be exchanged during startup, determining which taps of the shift register would be XOR'ed together to form the outputs. The code shown above will provide 5.1 dB coding gain, sufficient to meet the performance requirements. The programmable encoder structure also allows for future performance enhancement with backwards compatibility; however, the uncoded bits in the fixed 1-D encoding scheme limits such enhancements to fractions of a dB.

4.4. Spectral Compatibility with Existing HDSL

Spectral Compatibility with Existing HDSL

- ◆ Measurements of the original OPTIS spectrum showed that existing HDSL units might suffer some margin degradation
 - Heavy HDSL2 NEXT into HDSL downstream
 - Difference in Baud-Sampled vs. Fractional-rate adaptive FFEs
- ◆ Subsequent minor re-shaping of boosted region improved compatibility
- ◆ Adaptations to ETSI rate will need to scale frequency of "bump"
 - Compatibility with single-pair 2B1Q is an open question

Because existing HDSL units were designed with fixed receive filters and baud-spaced feed-forward equalizers, they were unable to adapt to provide optimal DFE performance in the presence of the new HDSL2 signal. Laboratory measurements with simulated HDSL2 NEXT into deployed HDSL units demonstrated performance degradations from 1 to 3 dB. The degradation was predominantly due to the raised portion of the HDSL2 upstream PSD in the region of 200 kHz. Minor modification of the "bump", slightly increasing its starting frequency, corrected the issue, and again multiple vendors verified this with laboratory measurements.

This susceptibility is an important consideration for ETSI, as ETSI will need to scale the OPTIS PSD in frequency, including the starting frequency of the "bump" in order to maintain compatibility with 2-pair HDSL. Unlike ANSI, ETSI has an existing single-pair HDSL recommendation, and careful consideration needs to be given to spectral compatibility between HDSL2 and the existing 1168 kbaud 2B1Q signal.

5. Current Issues in HDSL2

Current Issues

- ◆ ANSI work continues on details
 - Activation sequence, Line Powering, EOC, Framing, Options
- ◆ Line powering & HDSL2 repeaters
 - Problems associated with single-pair line powering
 - Extensibility of technology to low performance solutions
 - Transmit Power Control
- ◆ Extensions to ETSI E1 & rate-adaptive transmission
 - Spectral Shape:
 - Modify for ETSI & Rate Adaption?
 - Bits/symbol:
 - Use to scale rate adaptation?

Since December 1997, the T1E1.4 committee has been busy with the details of turning agreement on the fundamental transmission for HDSL2 into an interoperable standard, hopefully by the end of 1998, but likely by mid 1999.

One of the particularly difficult issues has been how to reconcile the desire for doublers (repeaters) for HDSL2 with the inherent limitations of a single-pair T1 technology.

Use of doublers will be more limited with HDSL2 than with HDSL for several reasons. Chief among these is a simple consequence of Ohm's law. Because power for the HDSL2 doubler must go over a single-pair instead of 2 pairs, span powering losses are higher in 1-pair HDSL2 than 2-pair HDSL. As a result, one-pair HDSL2 has much less powering capability than 2-pair HDSL. It is often thought that HDSL2 might be able to avoid the doubler issue merely by reducing the data rate and using more pairs. Unfortunately, halving bit rates does not double the distance, but rather achieves only about 30% reach extension.

There are currently several proposals on the table at T1E1 for line powering of HDSL2 units. Most of these allow high voltage powering with some form of current limiting. Even so, only 1 or maybe (in the future) 2 doublers could be powered, compared to the current 4 to 5 span powering capability of 2-pair 2B1Q HDSL. The issue of rate adaptation is also still open. However, discussions with service planners suggest that a rate-selectable version of HDSL2, similar to PairGain's 2B1Q HDSL-based Campus-HRS product would be preferable to automatic rate adaptation.

6. Some References

Apologizes to authors, individuals, and companies who made contributions to HDSL2 development that are not reflected in this reference list. This list is intended

to be informative to the reader, not necessarily definitive of contribution to the final result.

6.1. Background:

1. J. Aslanis & J. Cioffi, "Achievable Information Rates on Digital Subscriber Loops: Limiting Information Rates with Crosstalk Noise", IEEE Trans. Comm., 40:2, February 1992, pp. 361-372.
2. ANSI Standard T1.413-1995, "Network and Customer Installation Interfaces - Asymmetric Digital Subscriber Line (ADSL) Metallic Interface".
3. George Zimmerman, "Approaches to CSA-reach Single-pair HDSL Transmission", PairGain Contribution T1E1.4/96-063, April 22-26, 1996.

6.1.1. Background: Coding

1. E. Biglieri, D. Divsalar, P. J. McLane, and M. K. Simon, Introduction to trellis-coded modulation with applications, Chapter 4, Macmillan, 1991.
2. G.D. Forney, "Coset Codes - Part I: Introduction and Geometrical Classification", IEEE Trans. Info. Theory, 34:5, September 1988, pp. 1123-1151.
3. Simon Lin, "Relative Complexity of VLSI Viterbi Decoders in HDSL Transceivers", PairGain Contribution T1E1.4/96-064, 22-26 April 1996, Colorado Springs, CO.
4. B. Grayson, "Complexity and Achievable Gain of Sequential Decoding for Large Trellis Codes", PairGain Contribution T1E1.4/96-065, 22-26 April 1996, Colorado Springs, CO.
5. G. Zimmerman, "Turbo Codes for Bandlimited Channels: Potential and State-of-the-Technology", PairGain Contribution T1E1.4/96-067 22-26 April 1996, Colorado Springs, CO.
6. H. Herzberg "Powerful Multilevel Coding Scheme for Single Pair HDSL" GlobeSpan contribution, T1E1.4/97-176, May 1997.
7. George Zimmerman, "On Rate Loss in "Pinched Off" or Spectrally Constrained DFE-based Coded Transmission," PairGain Contribution T1E1.4/97-241, June 30, 1997
8. Chris Heegard and Stan Ling, "Multi-Stage Trellis Coded Modulation for HDSL2," Level One contribution T1E1.4/97-056, February 1997
9. Mike Tu, "A comparison of MS/TCM coding schemes for HDSL2", PairGain Contribution, T1E1.4/97-299, Sep. 1997.

6.1.2. Background: Spectral Compatibility

1. George Zimmerman, "Normative text for Spectral Compatibility Evaluations", PairGain Contribution, T1E1.4/97-180, May 12, 1997. (revised to 97-180r1, June 30, 1997).
2. George Zimmerman, "On the Importance of Crosstalk from Mixed Sources", PairGain Contribution, T1E1.4/97-181, May 12, 1997
3. George Zimmerman, "Test environments for HDSL2", PairGain Contribution T1E1.4/97-182, May 15, 1997.
4. Hiroshi Takatori, "Spectrum Compatibility from HDSL2 to T1", Level One, T1E1.4/96-345, November 11-14, 1996.

6.2. Key Characteristics of OPTIS HDSL2 Modulation

1. George Zimmerman, "A Modulation Strategy for HDSL," PairGain Contribution T1E1.4/96-340 Nov. 1996
2. Kevin Schneider and Richard Goodson, "A Modulation Technique for CSA Range HDSL2," ADTRAN Inc. T1E1.4/97-073
3. Jack Liu, "HDSL2 Transmit Spectra", PairGain Contribution, T1E1.4/97-178, May 15, 1997.
4. Mike Rude, "Refined HDSL2 Transmission Masks: Performance & Compatibility", ADC Telecommunications Contribution, T1E1.4/97-170, May 15, 1997.
5. Hiroshi Takatori, "Evaluation and Performance of the POET System", Level One Contribution, T1E1.4/97-191, May 15, 1997.
6. Hiroshi Takatori, "Spectral Shaped Transmission for HDSL2", Level One Contribution, T1E1.4/97-192, May 15, 1997.

6.3. OPTIS HDSL2 Solution

1. George Zimmerman, "Performance and Spectral Compatibility of OPTIS HDSL2," PairGain Contribution T1E1.4/97-237r1, June 30, 1997
2. Mike Rude, "OPTIS Performance and Spectral Compatibility," ADC Telecommunications Contribution T1E1.4/97-239, June 30, 1997
3. Hiroshi Takatori, "Performance Evaluation of OPTIS for HDSL2," Level One Contribution, T1E1.4/97-240, June 30, 1997

4. Mike Rude, Hiroshi Takatori, George Zimmerman, "Updated OPTIS PSD Mask and Power Specification for HDSL2", ADC Telecommunications, Level One Communications, PairGain Technologies, T1E1.4/97-435, December 1997.
5. Mike Rude, Peter Brackett, Hiroshi Takatori, George Zimmerman, Gunter Niemeyer, "Basis for HDSL2 Standard," ADC Telecommunications, Adtran, LevelOne Communications, PairGain Technologies, Siemens EZM, T1E1.4/97-471, Sep. 1997.
6. "Proposal to Break the FEC Logjam for HDSL2," Adtran, Cicada, Siemens, Tellabs, Westell , T1E1.4/97-443,, Dec. 1997.

6.4. Folded Spectra in DFEs

1. G.J. Pottie & M.V. Eyuboglu, "Combined Coding and Precoding for PAM and QAM HDSL Systems", IEEE JSAC, August 1991, pp. 861-870.

6.5. Coding

1. Mike Tu, "A 512-State PAM TCM for HDSL2," PairGain contribution T1E1.4/97-300 September 1997
2. Chris Heegard and S. Ling, "Simulation Results of 1-D TCM for HDSL2," Heegard Consulting & Level One Contribution, T1E1.4/97-xxx, Sept 1997.
3. Jack Liu, Stanley Ling, Mike Rude, Chris Heegard, "Proposal for HDSL2 Transmission: FEC Scheme", PairGain, Level One, ADC Telecommunications, Heegard Consulting Contribution, T1E1.4/97-301, 25 September 1997.
4. Richard Goodson, "Performance and Characteristics of One-Dimensional Codes for HDSL," ADTRAN contribution, T1E1.4/97-337, 25 September 1997.
5. Richard Goodson, "Proposal for a Programmable Trellis Encoder for HDSL2," ADTRAN contribution, T1E1.4/97-337, 25 September 1997.

6.6. Spectral Compatibility with Existing HDSL

1. George Zimmerman, "On Spectral Compatibility with DFE-Based Systems using Baud-Rate Sampling and Fixed Receive Filters", PairGain Contribution T1E1.4/97-271, September 1997.
2. Peter Brackett & Kevin Schneider, "The Importance of Testing Widely-Deployed Equipment in the Presence of Interference from New Line Codes", Adtran Contribution T1E1.4/97-339, September 1997.