Dynamic Spectrum Management
– A methodology for providing significantly higher broadband capacity to the users

JOHN M. CIOFFI[1] AND MEHDI MOHSENI

Dynamic Spectrum Management (DSM) methods are overviewed and analyzed in terms of benefit to DSL performance, rates, ranges and operational costs. Two DSM steps are recommended for near-term and longer-term improvements in DSL: (1) transmit power-minimization at fixed good quality of service by all or many DSLs in service, and (2) cooperative physical-layer signal vectoring at line-terminal access multiplexers. These steps are show to make 100 Mb/s symmetric DSL services practically viable everywhere DSL subscribers demand service.

John M. Cioffi is a Professor at Stanford University, California, USA

Mehdi Mohseni is a PhD student at Stanford University, California, USA

1 Introduction

The successful installation of a largely fiber-based network over the past two decades has encouraged higher access network speeds, a trend expected to continue. While high capacity can be delivered to individual customers over leased fiber facilities, lower cost access alternatives have become increasingly interesting. In particular, the world-wide success of the digital subscriber line (DSL) service, which can deliver increasingly high speeds on the existing copper twisted pair at a fraction of the cost of leased facilities, intensifies interest to best use this copper.

Cable television facilities have been upgraded to include a portion of fiber, enlarging the bandwidth to each customer and to provide two-way digital communication allowing fast internet and digital voice services to be provided on that facility, thus creating competition for incumbent telephone-company service providers. Thus, telephone companies have then studied the business case to enlarge the fiber network closer to, but not all the way to the customer with connections to the households that utilise the already existing copper lines. During the last years there has been an interesting evolution of the DSL technology that provides higher and higher capacity. The increase in capacity depends on a set of factors like length of the copper line, diameter on the twisted pairs, degree of symmetrical and asymmetrical traffic, the coding, etc. To be able to make good broadband demand forecasts, it is of crucial importance to understand the technological evolution.

This paper gives perspectives of the new Dynamic Spectrum Management (DSM) methods, which have the potential to increase the data rate capacity for broadband significantly compared with the DSL methods used today. Dynamic Spectrum Management (DSM) automates the provisioning, maintenance, and operations of DSL. Figure 1 decomposes a typical service provider's DSL costs into equipment and operations. The fraction of cost attributed to operations is already more than half for most service providers (and as high as 70 – 80 % in countries with higher labor costs) and increasing relative to equipment costs.

Customer service (truck rolls) visits contribute most of the "operations" cost. Higher DSL service rates at any given range (loop length) increase the probability of a customer service visit and consequently increase operations costs. The exact DSL data rate and range and consequent cost depend on the service provider, local practices and demand, and the cost of labor. Some service providers with longer loops already incur very large operational costs while others with largely urban populations and short loops (and perhaps lower local labor costs) have not yet experienced operations costs as high, but will eventually.

This overview of DSM suggests and details two steps in DSL evolution that allow significant operational-cost decrease each while progressively increasing the data rates at those lower operational costs: The first step is most simply and colloquially described as "stop hogging". This step introduces a politeness in DSL transmission spectra that allows each user to get excellent quality of service without overwhelming (via crosstalk) other DSL signals in the same trans-

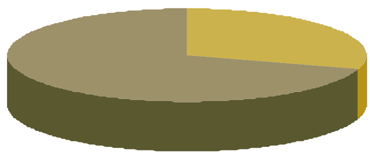

Figure 1 DSL cost decomposition

[1] *J. Cioffi is also with ASSIA Inc and SBC and many of the results in this paper merit credit to ASSIA and SBC.*

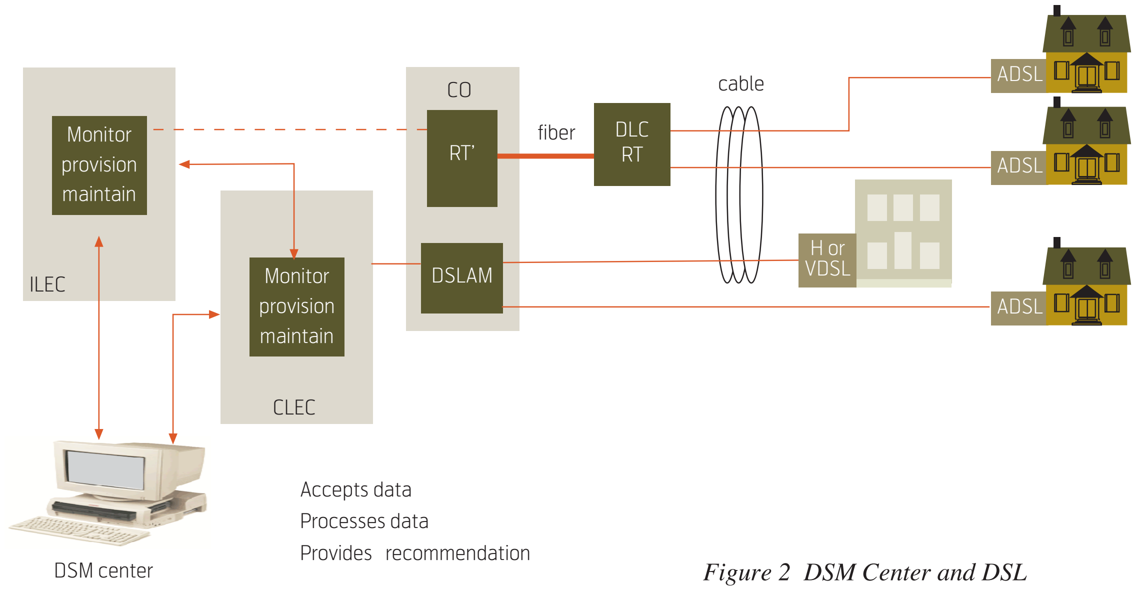

Figure 2 DSM Center and DSL

mission binder. This step is more formally described as "adaptive spectrum" control. Section 2 describes the adaptive spectrum concept and provides some DSM example improvements from DSL customers in service – this first step can be implemented largely on existing DSL service platforms with increasingly automated (software) use of information and controls available in those platforms. The second step is more long-term and involves conscious signal alignment of different transmissions through vectored DSLAMS that attach to many or all of a line terminal's distribution lines. Section 3 projects the further improvements of this second future 3rd-Generation-DSLAM DSM step. DSM is presently a standards project in the T1E1.4 American group and the current draft is at [1], while a complete list of many contributions to that DSM effort is in [2]. Other overview references occur in [3] – [5].

2 Adaptive spectra and the DSM center

Figure 2 illustrates the generic DSM Center that accepts and processes data before providing recommendations to a service provider's maintenance and provisioning tools. The types of data that can be forwarded to the DSM Center include line margins, transmit power levels used, bits/tone tables, insertion-loss per tone, noise per tone, actual power-spectral-density levels/tone, errored seconds, and known loop conditions like bridged taps, loop lengths, and binder taper-code/service-area (allowing the knowledge of other lines in the same binder).

The types of recommendations that are returned after processing can be the data rates (and associated reliabilities of achievement), allowed maximum margins, forward-error-correction choices, and power-spectral densities. Such recommendations can change daily or even more often, especially as DSL speeds increase and approach levels where a service-visit without DSM would be too likely.

2.1 The ADSL RT and DSM

Figure 2 also illustrates a classic problem for the growing DSL service provider, mixture of an existing DSL service and a new DSL service from a fiber-fed terminal in the same binder. Movement of the existing subscriber to the fiber terminal may not occur immediately (or ever) and increases operational costs. The fiber fed lines are usually short. Practice today (despite standards indications to the contrary) is for the fiber line to play at full power (say 20 dBm in ADSL) and the long line to play at reduced power (say 12 dBm) because of ill-conceived spectrum masks in current static spectrum management standards. Actually, the short line creates a very large crosstalk into the long line, reducing the rate of or disrupting the existing long-line customer. The short line meanwhile plays with 100 -100,000 times the power it needs to achieve a high data rate with no errors.

The situation in Figure 2 can be described by the rate region shown for two users in Figure 3. The rate

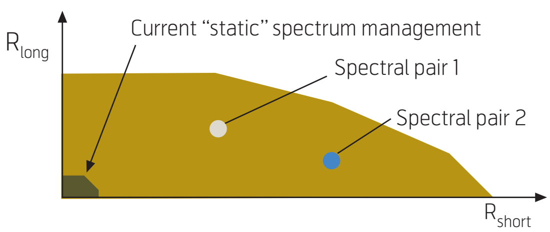

Figure 3 2-user rate region

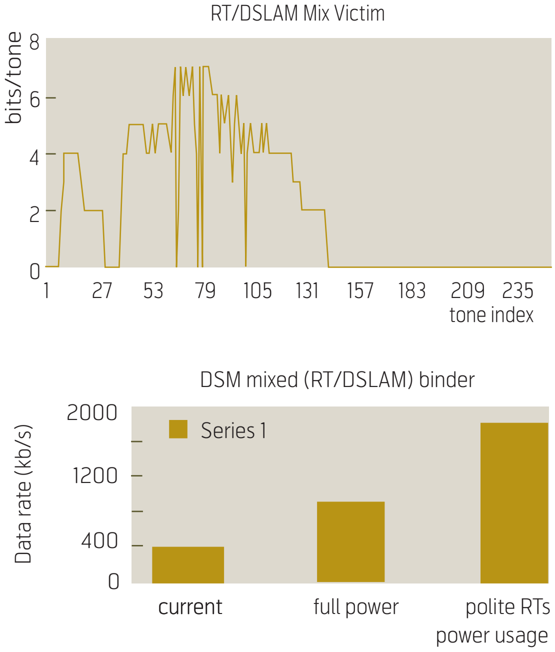

Figure 4 Victim ADSL bit distribution and rates

region in Figure 3 illustrates that different combinations of data rates are possible but each such combination requires a different pair of spectra. Static spectrum management's use of worst-case spectra causes a much smaller region to be achievable. The rate region is particularly appropriate for the situation of ADSL remote-terminal/central-office mixture.

Figure 4 illustrates at its top the bit distribution of a customer in service on a 13.6 kft loop in the USA and various rates below. Several newer RT ADSL circuits are crosstalking in the same binder and cause the rapid bits/tone drop in spectrum use for this medium-length line, rendering a service data rate of only 384 kb/s. This line is also using only 14 dBm of transmit power (instead of the full 20 dBm) because of a concern that such transmit power would cause an excess power spectral density of more than –40 dBm/Hz (an optional limit suggested in some current static spectrum management [6]). This limit is intended to protect HDSL circuits from down-to-up ADSL NEXT. The spectrum in Figure 4 also suggests that HDSLs are present in the binder (because it slopes up in downstream spectrum instead of immediately rising in the 140–250 kHz range). However, the HDSLs in this binder are all operating with highly excessive margins, so an increase of ADSL spectra could be possible (high HDSL margins are common – cursory calculations can show the likelihood of long-length HDSL in the same binder as a very long ADSL, where both are at margin limits, occurs in about one binder out of 100,000 or more – the cost of repairing such a rare event might well be less than the service revenue lost by lowering the rates of millions of ADSL customers). Thus, simply increasing the ADSL power to the allowed maximum of 20.5 dBm increases the data rate to 904 kb/s. Alternatively, if all the RT lines in this binder operate with minimum power for 16 dB of margin at 1.536 Mb/s (and even at rates as high as 3–6 Mb/s), the 13.6 kft loop sees much less crosstalk and actually operates itself in excess of 1.536 Mb/s (the same speed as before the RT is installed and still with 14 dBm of transmit power). Thus DSM provides two solutions for this line, either one of which would lead to very large data rate increase. Similar rate increases have been independently noted in [5], [7] and [8].

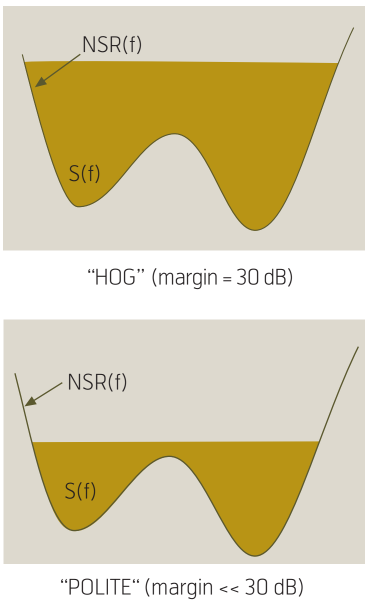

Figure 5 Illustration of water-filling for both hog and polite DSL modems

2.2 Iterative water-filling

The analysis of the situation shown in Figures 2 – 4 is often called "iterative water-filling". Water-filling is a term (see [6]) used to describe the calculation of the best spectrum for a transmission line, in particular a DSL loop. The water-filling procedure is illustrated in Figure 5 for both a "hogging" modem that uses too much power and a polite modem. In both cases transmitted power is viewed as "water" poured from above into the noise-to-signal-ratio frequency curve *NSF(f)*. The dark region is the best power spectral density that can be used on this line viewed by itself without knowledge of other lines. The impolite or "hog" modem continues to pour energy in until the source is exhausted, while the polite modem only uses enough energy to ensure no errors (or high quality of service).

Water-filling is typically approximated within various practical constraints by the DMT DSL modems that are used in most DSL services (see [9] for a discussion of how simple PSD constraints can be used with water-filling to implement "optimum spectrum management."). Polite water-filling modems transmit at no greater than some service-provider-specified maximum rate (usually the maximum rate the customer was offered or purchased). At this rate, good DSL modems then minimize the power they need to achieve this rate with some maximum margin, typically called MAXSNRM in DSL standards. In particular, polite-water-filling spectrum use also avoids or reduces noise/crosstalk in bands where crosstalk is damaging, more important yet than simple power reduction. When this is implemented correctly, "hogging" DSL lines that use too much power are eliminated, greatly reducing crosstalk. The DSM center recommends data rates and the MAXSNRM to the various DSL lines of the service provider. There is no coordination of the modems other than the usual service provider's specification of maximum data rates and maximum margin to be attempted by any given customer. The results can be simulated in a computer simulation by simply running the water-filling procedure for each line successively with all other spectra (and thus consequent crosstalk) viewed as noise. Repeatedly iterating this process imitates the actual binder action, allowing the "iterative water-filling" to project the performance of the system. Note this simulation is very different from the incorrect static-spectrum-management assumption that each line uses a fixed spectrum – the results are also often very different. (The assumption is incorrect because ADSL modems do not have fixed spectra.)

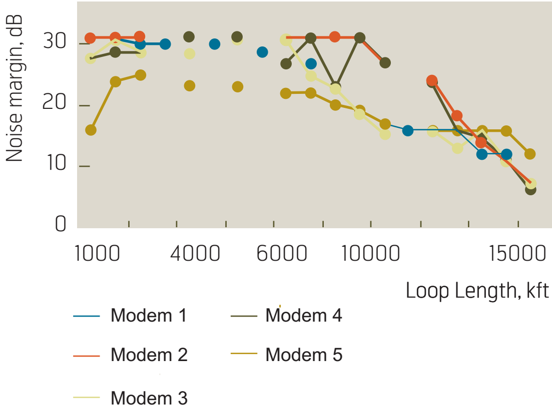

Figure 6 Margins for 1.536 Mb/s DSL

An unfortunate note:

Figure 6 illustrates the measured and reported margins for several different manufacturers' modems (the names have been removed and simply referred to as "modem1," "modem2," etc.). All the manufacturers here exceeded an already high 16 dB MAXSNRM required by the service provider by a significant amount to as much as 13,000 ft (at 1.536 Mb/s). The excessive margins are an artefact of early ADSL standards limiting the amount of power back-off that can be imposed by a receiving modem running water filling to 14.5 dB, which is clearly not enough in Figure 6. This high-margin problem can be repaired by DSLAM software releases that allow the DSM Center to specify the initial transmit spectrum/power level as a function of line history (so a line with 20 dB more margin than the maximum required would be trained on its next initialization at a power level 20 dB lower for example).

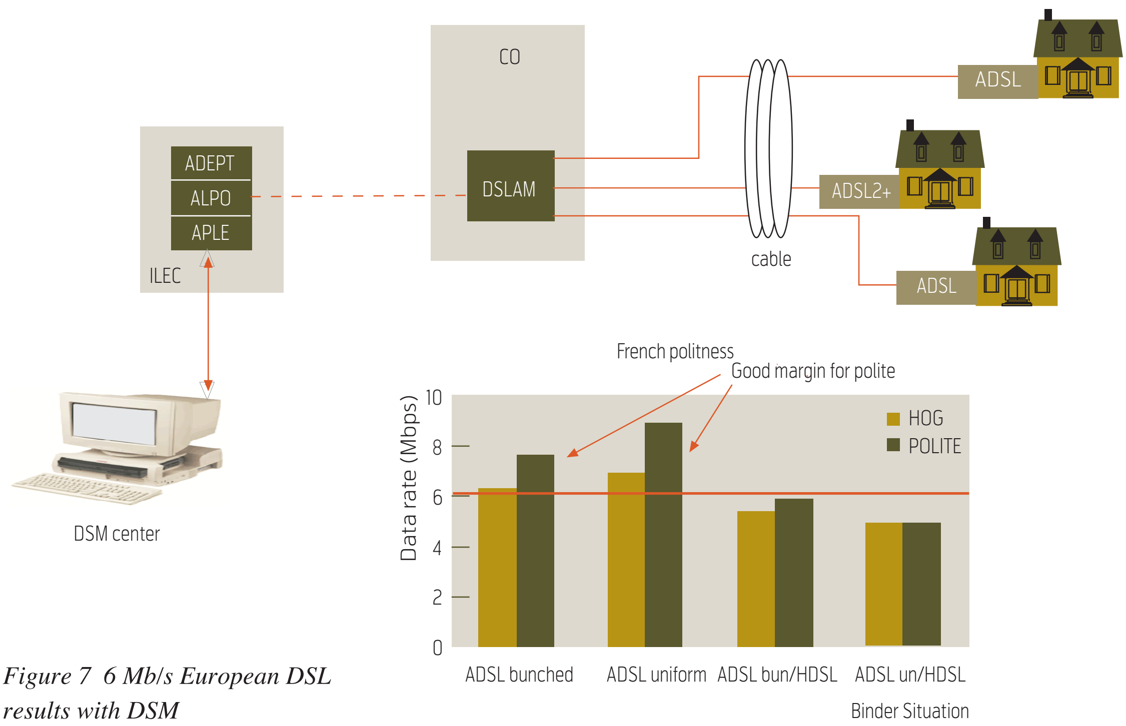

Figure 7 6 Mb/s European DSL results with DSM

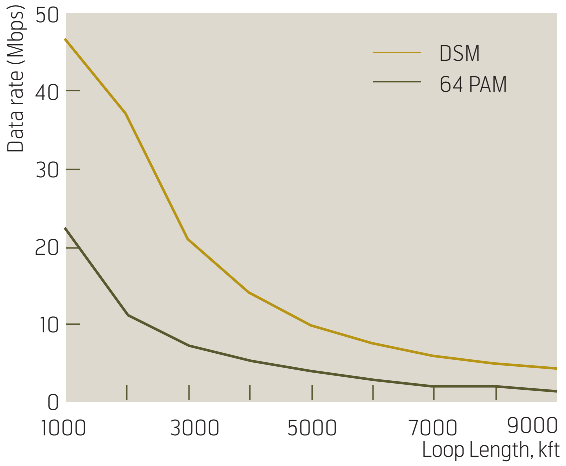

Figure 8 Symmetric rate limits of DSM without band preference

2.3 DSM for higher asymmetric speeds without RT mixture

Many service providers may not have RTs or mixtures of fiber-fed and central office loops in the same binder. Nonetheless, the simple "iterative water-filling" DSM procedure of Section 2.2 still leads to large improvements, especially at higher data rates because shorter lines do not excessively transmit power. Thus, the dominant impairment at higher speeds, crosstalk, is reduced substantially from "shorter lines." An example appears in Figure 7.

In this DSM simulation, a service provider desired 6 Mb/s at 3 km for video expansion of their DSL services. Static spectrum management with assumptions of fixed spectra and "hogging" DSL modems leads to the results on the left (ochre) in each pair of columns where 6 Mb/s is not achievable or just barely achievable. However, with simple limiting of MAXSNRM to 6 dB and aversion of HDSL in the same binder, the DSM results (on the right [brown] in each pair of columns) show plenty of extra data rate (or extra safety for ensuring the rate). The result is true whether the lines are "bunched" (all in the last 2.5 to 3 km) or "uniform" (uniformly distributed between 0 and 3 km). In particular, with HDSL in other binders and polite operation, the DSM situation shows 8 – 9 Mb/s, leaving some room for non-ideal effects. Operation without politeness clearly will not be possible.

2.4 Ultimate limits of balanced adaptive spectra

Figure 8 illustrates the ultimate limits of non-vectored uncoordinated transmission systems using iterative water-filling. In this figure, no spectral limits are placed on any DSL system downstream or upstream and each adapts for best symmetric transmission rate. All lines are the same length, which represents a worst-case situation (mixtures of shorter lines and longer lines allows the shorter lines to use less power, use higher frequencies, and thereby yet less crosstalk into the longest lines). The entire binder of 25 lines was active and 1 % worst-case crosstalk coupling between all pairs of lines was used.

Also shown in Figure 8 is a best-case fixed-spectra (PAM or "SHDSL-like") symmetric data rate. The range is doubled for both 10 Mb/s symmetric service and 5 Mb/s symmetric service by iterative water-filling (two bonded 5 Mb/s make one 10 Mb/s symmetric service). It can also be shown that this symmetric performance with iterative water-filling will not significantly reduce the data rate of any existing ADSL service if instead some of the lines in the binder were asymmetric instead of symmetric. However, the fixed spectra SHDSL-like choice essentially annihilates ADSL performance if in the same binder.

2.5 Band preference

Band preference in water-filling can essentially achieve optimum (largest) rate regions (such as in [8], [9]). In band preference, the DSM Center is presumed to know all the crosstalking paths and line lengths (while in Subsections 2.1–2.4, just worst-case assumptions were made and no line knew the length of any other and adapted only locally). This central knowledge is called "Level Two" coordination in the DSM standard [1]. In this case a coordinated allocation of energy can be determined by jointly optimizing all spectra used at the DSM Center. Such best spectra can be imposed by setting PSDMASK parameters (for instance see ADSL2 standards G.992.3 [10] or G.992.5 [11] – so-called "spectrum toolbox" or tssi parameters) for each line on a basis of the DSM Center predictions of binder performance. Such "band preference" can force a short line to use higher frequencies even though lower frequencies would have been better in iterative water-filling with no spectral masks imposed. Figure 9 illustrates an upstream VDSL rate region increase from the use of band preference for two lines of lengths 600 meters and 900 meters in upstream VDSL for "998-Plan" [13] iterative water-filling with noise A ([13]) and for water-filling using band preference with PSDMASK levels of –72 dBm/Hz below 5 MHz and –55 dBm/Hz above 5 MHz on the 600 meter line, and –50 dBm/Hz on the 900 meter line. The data rate on the shorter line is dramatically improved with band preference.

To the extent that Noise A is based on static ADSL models, Figure 9 represents worst-case performance because the ADSLs then are not modelled correctly (they water-fill also, so model "A" noise in standards is grossly incorrect, but used anyway here). Note that 6 Mb/s on the 900 m loop upstream can be achieved while nearly 20 Mb/s upstream occurs on the 600 m loop. These rates are considerably higher than what would be achieved if fixed spectra were used on the VDSL lines and more than double what is achievable with fixed spectra today.

Figure 10 illustrates that band preference provides most of its gain when mutually crosstalking loops have very different lengths. As the lengths approach the same, no band preference is necessary and all iterative water-filling loops can use the same PSD mask levels. A single 900 m or longer loop was held at 10 Mb/s (80 % of its maximum rate when no other loop is present) while a single other loop was varied in length between 600 m and 900 m. The vertical axis plots the fraction of the maximum short-loop rate that is achieved with and without band preference.

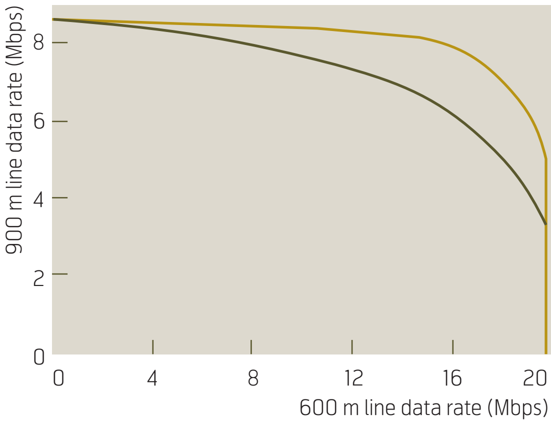

Figure 9 Upstream VDSL rate region for band preference (upper curve has band preference, lower has not) – four VDSL lines at 600 m and four VDSL lines at 900 m

2.6 Intermittent and impulse noise

Impulse noise is by its very nature not well characterized in DSL. A better term would be "intermittent noise" because customer premises (and sometimes other locations) noise can last much longer than 1 ms and can be repetitive and frequency selective. When it occurs in many cases, even interleaved operation of a DSL modem is insufficient to prevent severe erroneous seconds. Not all lines are "chronic" as such and have this problem. In fact, only a small percentage may be so adversely affected – but those lines can dominate operations costs through truck rolls to rewire the customer's premises. Increase of transmit power is of very little benefit for this type of noise. Proprietary "impulse-skewed-loading" assumes some consistency to noises that just may not often be valid.

The practical solution is forward error correction with sufficiently high percentage of parity to eliminate the errors on a line. Interleaving will not help remove the effect of an intermittent noise that is continuing to occur several times within an interleave depth. One might define a chronic line as

> **Chronic line**: *A line for which even interleaved forward-error correction is not sufficient to produce very low or zero errors.*

In ADSL and VDSL, Reed-Solomon (RS) codes are used and are characterized by N bytes in a codeword, of which P are parity bytes, and the remaining $K = N - P$ carry all other information. In an RS code, $t = P$ bytes in error can be corrected if erasures are used, which is common with DSL systems and intermittent noise. If no erasures are used, then $t = P / 2$ erroneous bytes can be corrected.

> **Correct-Chronic-Line Simple Rule**: *If the fraction of erroneous seconds (or time intervals) is denoted f_{errors}, then the selection of $t / N > f_{errors}$ will correct the errors and the line will no longer be chronic.*

The above rule essentially presumes any interleaving uniformly distributes errors, so is somewhat of a best case, but could be applied in an ADSL system. Usually $P = 16$ is the maximum value in DSL, so the fraction t / N is increased by decreasing N. The codeword length N is supposed to be programmable in ADSL standards, but it is not a parameter included in present management interfaces to the telco. It should be. Only the telco can make reliable assessment of the correct N value by evaluation of error performance and customer-service need. Decrease of N reduces the data rate on the line. However, chronic lines have poor or no data rate because of all the bit errors, so reduction in data rate using a lower N value that increases the effective throughput is an improvement.

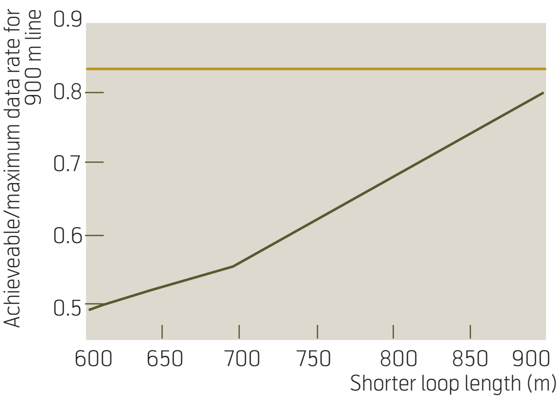

Figure 10 Band preference (upper curve) versus no preference when the shorter loop length is varied

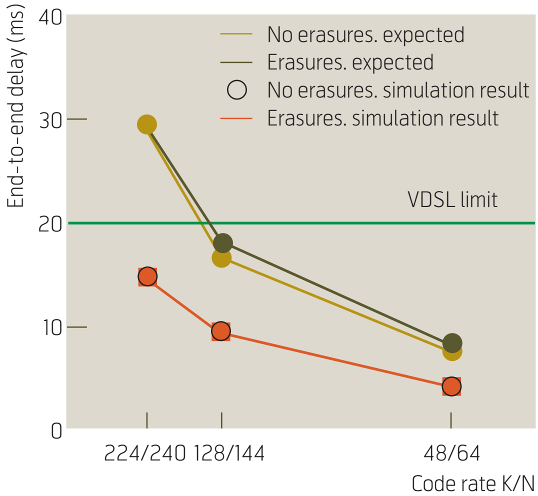

Figure 11 Delay versus code parameters (K = N – P) for FT "impulses"

ADSL1 [12] and VDSL1 [13] standards allow *N* to be set by the ATU-C. Thus, motivated DSLAM (or RT) manufacturers who want their customers to have better ability to correct intermittent noise will add the ability to change the default DSLAM *N* value on chronic lines (note the default *N* is often very good for lines without impulse problems or lines for which the current interleaving is sufficient to eliminate impulse errors, but not for chronic lines).

ADSL2 [10], [11] standards instead allow the ATU-R to select the *N* parameter according to two supplied parameters: Impulse "length" (INP) and maximum interleaving delay (DELAYMAX). The impulse length in ms is computed by (INP/4) ms. The DELAYMAX is given in ms and must be at least 4 ms. Unfortunately, current versions of the standards (which need an FEC amendment) only allow an impulse "length" of up to 2 DMT symbols (or 500 microseconds). A loading-algorithm engineer could use the following formulas to convert the supplied maximum delay and now renamed "impulse strength indication INP" into the usual FEC parameters (where $t = P$, the number of parity bytes chosen by the receiver if erasures are used, and $t = P / 2$ if no erasures are used and rate includes information and synch/control bits)

$$D' = depth\ in\ bytes = \frac{(INP/4) \times \left(\frac{rate\ in\ kbps}{8}\right)}{t} \quad (1)$$

$$N = \frac{(DELAYMAX) \times \left(\frac{rate\ in\ kbps}{8}\right)}{(D' - 1)} + 1 \approx \frac{DELAYMAX}{burst\ length} \cdot t \quad (2)$$

where burst length (*INP* / 4) and DELAYMAX again are in ms. Then the remaining G.992.3/5 [10], [11] DMT/framing parameters (*M*, *T*, *L*, see [yy]) can be computed and *S* (number of symbols/codeword) inferred for all latency paths as well as the values of *B* (the rate) for each frame bearer and latency path. Also once *N* is known, then $D = ND'$, the interleave depth in codewords. For current largest value of $INP = 2$ and smallest value of $DELAYMAX = 4$, one finds that $N \geq 8 \cdot t$, which with strong error correction so $t = 16$ leads to a smallest blocklength of $N = 128$, much larger than would occur on chornic lines according to (2).

An unused bit in the NPAR(3) of G.994.1 [14] used to convey INP will now allow an increase in INP to 4, 8, and 16 in the addendum to the ITU ADSL2 and ADSL2+ standards about to be released in 2004. With such values at lower data rates, all practical values of *N* could be attained, thus allowing ADSL2 to have almost as good error protection on chronic lines as ADSL1 where the DSLAM controls instead directly the *N* value. While maintaining the specified (DELAYMAX/burstlength) ratio, the ATU-R in ADSL2 should not pick a small *t* parameter just to decrease *N*. Such a poor strategy would increase the depth but not assist the chronic line – the choices should be made to obey the chronic line rule, which forces a larger value of *t*. Even approximate adherence to the above mathematical rules should assist on chronic lines. Service providers with DSM capability may well train the modem several times, each time increasing the INP value in ADSL2 until zero (or decreasing the *N* value in ADSL1 while holding *t* or parity constant) until small numbers of code violations and/or errored seconds are observed over a time period determined by the service provider.

An example of the delay necessary to eliminate all impulses measured (about 33,000) in France Telecom's network appears in Figure 11. Even in this study, many of the most severe impulses (those persisting much longer than 1 ms in intermittent patterns) could not be recorded, but the trend in the graph to the right as *P* / *N* increases would also be such that eventually as long as a few bits make it through the channel, the line can function (albeit at a lower data rate). Figures 12(a) and 12(b) illustrate a severe chronic American customer in service (with complaints to provider).

This chronic customer sees about 34,000 code violations in a 15-minute interval, about 60 % of the bytes are in error. Thus, with erasures then a code with $N = 24$ and $P = 16$ would correct all errors. This line otherwise has a maximum attainable rate in excess of 5 Mb/s and would function at 1.5 Mb/s with such error correction with no customer complaint at this maximum rate for which the customer paid. Figure 13 illustrates the concept of the DSM Center in this case, monitoring perhaps over a significant period of time (could be days or weeks) for code violations and then training the customer at an appropriate level. If the customer later desired a 5 Mb/s service, then the telco would know before provisioning that service that customer-premises wiring needs replacement (or a splitter perhaps located at the entry point) to reduce or eliminate the problem. The service provider might so notify the customer and advise of a charge covering the 5 Mb/s-upgrade service visit.

Both the ADSL1 and ADSL2 methods can be executed with a MAXSNRM also implied as in Subsections 2.1–2.5. The power used for a lower *N* will be higher, but still the minimum necessary to guard against stationary and crosstalk noise.

3 Signal alignment and DSM vectoring

Figure 14 illustrates the concept of a vectored DSLAM. All lines emanating from an LT are coordinated so that transmitted downstream signals are aligned or co-generated. Similarly received upstream signals are synchronized and aligned also. Such systems can eliminate ALL upstream NEXT and FEXT theoretically (even if the NEXT comes from signals that are not coordinated). Such systems can also eliminate all downstream FEXT (but not NEXT). Upstream NEXT is cancelled on each tone by subtracting constructed crosstalk of other users.

FEXT is eliminated in both directions via triangularization of the matrix binder response, followed by a successive cancellation on each tone of the effects of previously decoded (or previously precoded downstream) signals on other users on that same tone. These

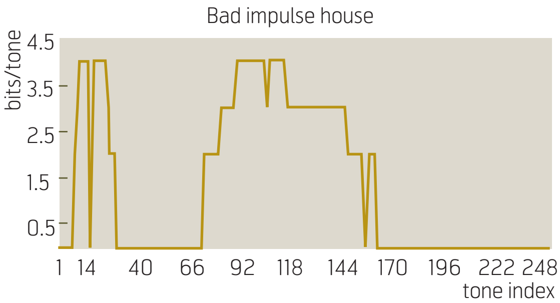

Figure 12(a) Bit distribution for American DSL customer with intermittent noise

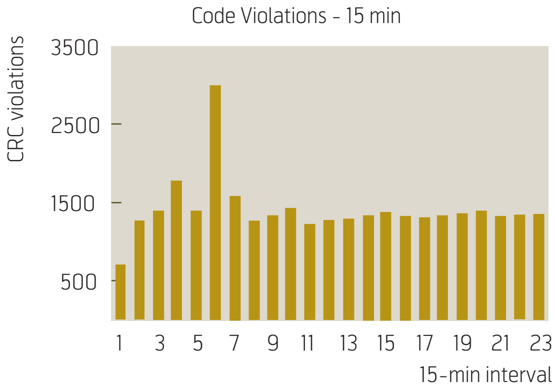

Figure 12(b) Code violations (with 24 ms interleave enabled) for same customer as in Figure 12(a)

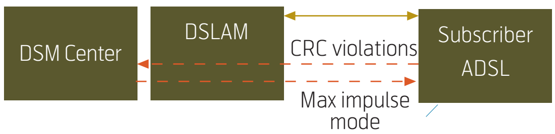

Figure 13 DSM Center with impulse/code adjustments

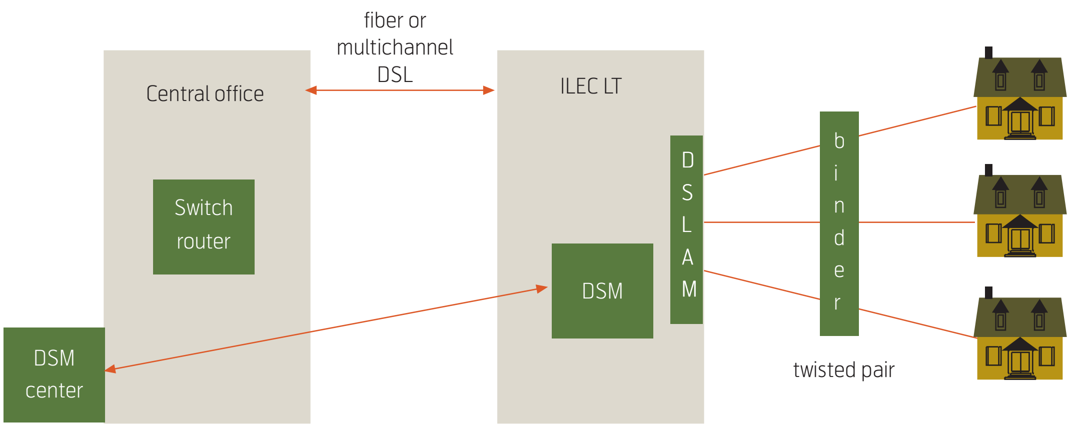

Figure 14 Basic signal alignment with vectored DSLAM

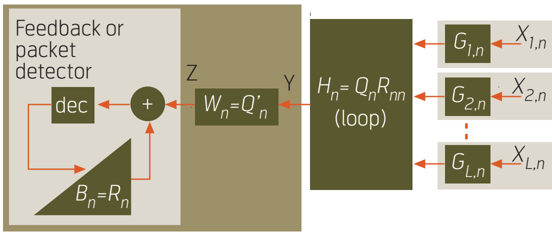

Figure 15 Upstream vectored-receiver coordination

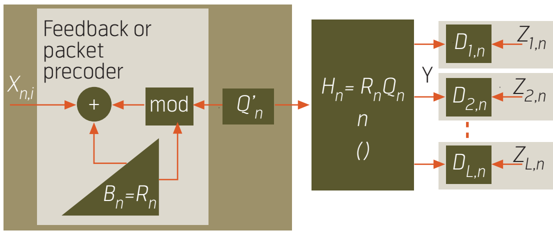

Figure 16 Downstream vectored-transmitter coordination

results are independent of the frequency band planning that is used. Allocating more downstream frequencies favors downstream rates, but does not change the elimination of crosstalk fundamentals of FEXT and NEXT upstream and only FEXT downstream.

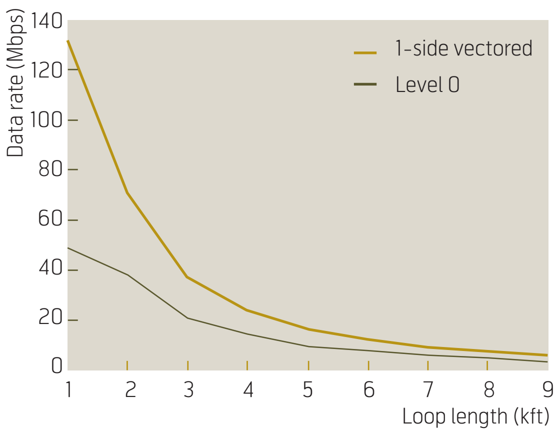

Figure 17(a) Signal Alignment, lengths to 9000 feet

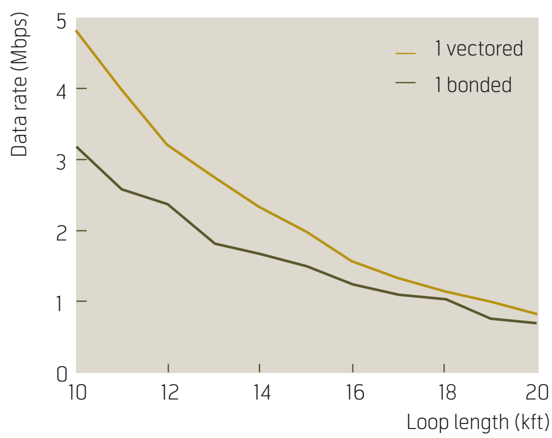

Figure 17(b) Signal Alignment, lengths 9–20 kft

3.1 Bonding versus vectoring

Bonding is a link-layer (TPS-PMS) multiplexing and de-multiplexing of data rates on several lines used to service a single DSL customer. Vectoring is the co-generation of physical-layer (PMD) transmitted signals, co-processing of physical-layer received signals, or both. Bonded lines may optionally also be vectored for the receivers on both sides (in which case NEXT and FEXT both can be eliminated if all the lines in a binder are coordinated and bonded, so there is then no crosstalk of any type upstream or downstream). More realistically, when only a few lines (say 2 or 4) are bonded, there is no advantage in terms of data rate per line with respect to just vectoring and no bonding. Bonding increases the speed of a service by using more lines. Vectoring instead increases the speed of a DSL service by better coordinated signal processing. Either can be used without the other or together.

3.2 Vectored signal processing structures for DSM

Figures 15 and 16 illustrate the basic processing structures for upstream and downstream (DMT-based) DSL systems with presumed digital duplexing on all lines [13] and all such lines using the same DMT symbol phase. For upstream, signals are co-received. Any locally generated downstream signals can be removed through crosstalk cancellation, which is essentially a multi-dimensional echo canceller, and is not shown. On each tone the remaining structure in Figure 15 first conditions the signal to a triangular interference structure through a matrix filter. One of the users will have no crosstalk whatsoever on each tone and is decoded first. Then that user's influence on the next user is constructed and removed from all subsequent user's decoding, one by one, in a successive process independently of each DMT tone. The matrix filter and triangular feed-back structure depend on knowledge of the exact binder crosstalk matrix and noise autocorrelation matrix. The matrix filter also rejects all NEXT from uncoordinated crosstalk signals.

Figure 16 illustrates the dual downstream (per tone) transmit structure that uses a matrix triangular precoder, followed by a matrix pre-filter to co-generate all signals.

3.3 Ultimate limits of signal alignment DSM

Figures 17(a) and 17(b) illustrate the increase in data rates over those of Subsection 2.4 and Figure 8 over iterative water-filling alone. Note that 100 Mb/s symmetric speeds (thus 200 Mb/s asymmetric) is possible at 1500' or 500 m while 10 Mb/s symmetric is possible to 7000' or 2 km. 100 Mb/s on two lines is about 900 m in range.

3.4 Finding MIMO (binder identification)

The identification of the crosstalk transfer functions between lines, both gain and phase, is necessary for vectored transmission. This is an extension of the insertion loss gain/phase presently reported by ADSL2 and ADSL2+ modems to inter-wire transfer functions. There are several sophisticated methods for identifying such crosstalk coupling, and only the simplest is described here for two cases: (1) synchronized start and (2) unsynchronized start.

Synchronized start:

Synchronization of the lines upon start-up means that they all train at the same time. This is feasible with bonded lines and in certain situations with no bonding also. Figure 18 illustrates the basic concept, which is that each line transmits independently (or a variant is all transmit at the same time, but with clearly distinct robust, i.e. "white" training sequences) at one time only during training a sequence that is used not only to determine insertion loss gain/phase on that line, but also to determine the FEXT transfer function from/to (upstream/downstream) all other lines in the binder. All XTU-R modems know the time slot of their line as well as all other lines through the single-sided coordination at the line terminal (or fiber-fed DSLAM). The learned transfer functions are combined into a matrix of transfer functions. QR factorization (see Figures 15 and 16, which is a factorization of the binder matrix transfer function at each tone into the product of an orthogonal matrix Q and a triangular matrix R) is then executed upon the resulting matrix to obtain the settings for the matrix filters and triangular processing of Figures 15 and 16. With digital duplexing on all lines to the same symbol clock, this process is independent for each tone of a DMT system. An overall bit/data assignment for all users (picking an acceptable point in the vectored rate region) can then be made. The learned downstream transfer functions must be reported back to the coordinated side via a reverse or secure channel (in much the same way that ADSL2 returns the insertion loss HLIN[n]).

For systems with overlapping spectra, the NEXT transfer functions can be simultaneously learned for downstream into upstream (and for upstream into downstream where desired) at the same time. Thus, only one transmitter is active at each time (or orthogonal to signals sent by all other transmitters).

Asynchronous training:

Asynchronous systems will be slaved to the same master symbol clock with digital duplexing on all lines in the same line terminal as with all vectoring systems. However, the individual customers may train at different times (representing time of energy excitation, offering of service, retraining for any reason, etc.). In this case upstream, a modem beginning its own training needs to energize at a low but detectable energy level. The upstream co-reception process will detect the new user and subtract all other users' signals (since they are in operation and detectable – a line for which signals are not detectable most of the time would retrain anyway, see synchronous training). In this case, the new row of the channel matrix would be determined (again independently for each tone) with the on-line transfer and all the FEXT coupling then known and used in subsequent calculation. For downstream asynchronous start, for full performance gain, all remotes have to be capable of either (1) determining downstream FEXT from a new subscriber themselves or (2) collecting packets of channel outputs periodically or on command that are returned to the central processor. This is called "Level 3" coordination in the DSM standard and preferred (so that client modems are otherwise not unduly pressed in terms of processing capability).

3.5 Phantom of the DSL

Figure 19 illustrates the basic concept of a phantom signal for two twisted pairs (the use of the term "phantom" is somewhat misguided as the signal is very real and exists – indeed the phantom is what creates what we normally call crosstalk, which exists). For quad situations, the phantom component can be high. Typically, the phantom is defined between

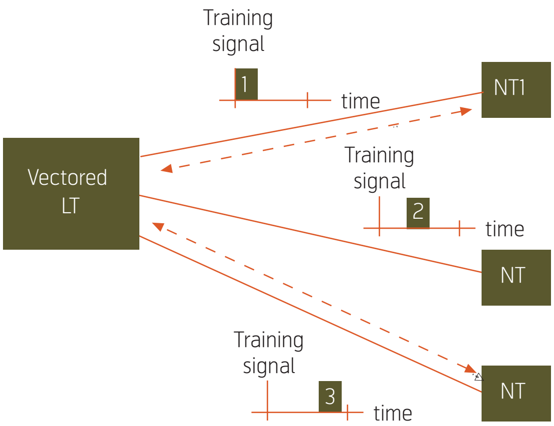

Figure 18 Synchronous training illustration

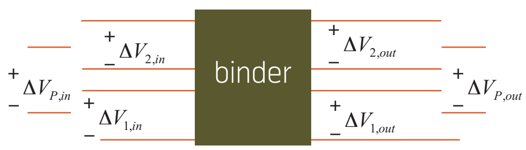

Figure 19 Illustration of phantoms (see [1])

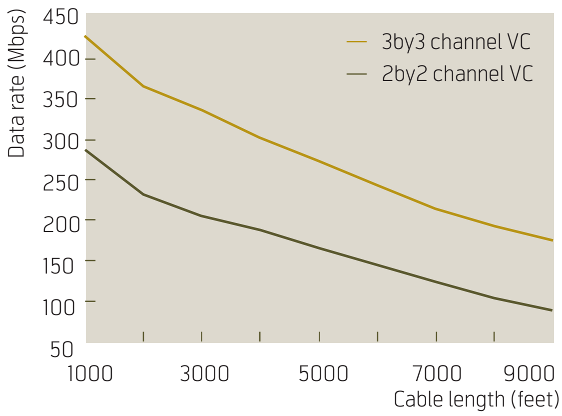

Figure 20 Illustration of quad capacity (2 loops or used with phantoms so 3 loops)

hypotentical "center tap voltages" of the two pairs, denoted ΔV_p in Figure 19. In reality, there are four conductors that create three closed loops, not two, if appropriately terminated and excited as a group. The third or extra loop can have very real and significant data-carrying capacity. It can thus be exploited in bonded and vectored schemes. The gain can be more than 50 % or less depending on loop geometries, imperfections, and coupling. Bonding lines with significant coupling may indeed be physically and practically difficult, except for the case of so-called "quads" where the four wires of a single quad twisted as an ensemble (so a twisted quad) may in many cases be maintained to a customer. The concept of a quad generalizes to any number of conductors but would then require yet further bonded and vectored use of multiple pairs. In reality, this is a random possibility for transmission that is very real (i.e. not a phantom) but would not be reliable unless pairs were selected carefully for customers. The data rates for vectored quads (based on French cable) are shown in Figure 20.

Such auto-selection and routing of pairs in a cable in a distribution frame and at customer premises may be possible in multiple-dwelling units or elsewhere and likely represents the last known physical capacity in the copper loop plant that can be exploited to squeeze the very last highest rates from copper. It would be highly adaptive and capacity would need to be assessed by a DSM center on a binder-to-binder basis for the various customers. Given the high costs of fiber in the very last segments of the network, this last frontier of bandwidth in phantoms may need to be someday exploited. Fortunately, DSM offers significant bandwidth increase of very large factor even without exploitation of this last opportunity.

3.6 Data-dependent DSM

The exploitation of silent intervals that abound in data transmission (and even in video when video is off) create a very interesting possibility for DSM that has not yet been exploited in a centrally coordinated manner. The use of vectoring allows several opportunities for such systems. Silent signals on other lines reduce crosstalk for the period of silence, meaning another victim line could increase its rate during that time frame. Futhermore, a vector transmitter could actually exploit crosstalk from other lines to reinforce a signal when its nominal signal is silent, allowing a statistical boost in data rate over the theoretical level of "no FEXT and no NEXT" so actually not only is noise low, but signal received is higher also. The shortest lines in a binder would benefit most when longer lines are silent and can temporarily boost the short-line's rate.

4 Conclusion

Dynamic Spectrum Management through automated intelligent control of basic DSL line parameters can significantly reduce operations costs for DSL service providers. Simultaneously, data rates can be substantially increased, readily to 100 Mb/s symmetric services on no more than two lines to 1 km range and beyond. The ultimate limits of DSM may require new chips and vectored equipment, but a great deal of improvement is possible with existing systems through the use of basic MIB parameters and procedures already inherent in existing DSL systems like ADSL.

DSL methods through the use of DSM can resurrect a viable broadband strategy with good operational cost and excellent service providers, enabling a real worldwide and sustainable broadband revolution. No other solution has such realistic promise. These higher speeds will establish a continuing improved business case for wider bandwidth telephone company networks via the higher speed lower-cost copper connection to the customer.

References

1 *DSM Draft Report*. December 8, 2003, San Diego, CA. ANSI T1E1.4/2003-018R9.

2 *DSM Living List*. December 8, 2003, San Diego, CA. ANSI T1E1.4/2003-017R4.

3 Song, K B, Chung, S T, Ginis, G, Cioffi, J. Dynamic spectrum management for next-generation DSL systems. *IEEE Communications Magazine*, 40 (10), 101–109, 2002.

4 Starr, T, Sorbara, M, Cioffi, J, Silverman, P. *DSL Advances*. Upper Saddle River, NJ, Prentice-Hall, 2003. (Specifically Chapter 11 by J. Cioffi.)

5 Bostoen, T (Alcatel), Cioffi, J (Stanford). Dynamic Spectrum Mangement. *DSL Forum Birds of a Feather Special Seminar*, Paris Hilton, Nov 11, 2003.

6 T1.417-2002 ANSI Spectrum Management Standard.

7 Kerpez, K, Cioffi, J et al. Response to 2001-273R1 using Telcordia DSL Analysis. *ANSI Contribution T1E1.4/2002-063R1*, Vancouver, BC, February 18, 2002.

8 Cendrillon, R, Moonen, M, Yu, W, Verlinden, J, Bostoen, T. On the Optimality of Iterative Water-filling. *ANSI Contribution T1E1.4/2003-325*, San Diego, CA, December 8, 2003.

9 Cioffi, J, Mohseni, M et al. Band Preference with Water-Filling in DSM. *ANSI Contribution T1E1.4/2003-321R1*, San Diego, CA, December 8, 2003.

10 ITU. *Asymmetric digital subscriber line transceivers 2 (ADSL2)*. Geneva, International Telecommunications Union, 2003. (Recommendation G.992.3 (07/02.)

11 ITU. *Asymmetrical Digital Subscriber Line (ADSL) transceivers – Extended bandwidth ADSL2 (ADSL2+)*. Geneva, International Telecommunications Union, 2003. (Recommendation G.992.5 (05/03).)

12 ITU. *Asymmetrical digital subscriber line (ADSL) transceivers*. Geneva, International Telecommunications Union, 2000. (Recommendation G.992.1.)

13 VDSL1 DMT Standard T1.424-2004. American National Standards Institute, 2004 (see draft T1 letter ballot 1184).

14 ITU. *Handshake procedures for digital subscriber line (DSL) transceivers*. Geneva, International Telecommunications Union, 2002. (Recommendation G.994.1.)

15 ITU. *Physical layer management for digital subscriber line (DSL) transceivers*. Geneva, International Telecommunications Union, 2002. (Recommendation G.997.1.)

John M. Cioffi – BSEE, 1978, Illinois; PhDEE, 1984, Stanford; Bell Laboratories, 1978–1984; IBM Research, 1984–1986; EE Prof., Stanford, 1986–present. Cioffi founded Amati Com. Corp in 1991 (purchased by TI in 1997) and was officer/director from 1991 to 1997. He is currently on the Board of Directors of Marvell, Teknovus, ASSIA (Chair), Teranetics, and ClariPhy. He is on the advisory board of Portview Ventures and Wavion. Cioffi's specific interests are in the area of high-performance digital transmission. Various Awards: Holder of Hitachi America Professorship in Electrical Engineering at Stanford (2002); Member, National Academy of Engineering (2001); IEEE Kobayashi Medal (2001); IEEE Millennium Medal (2000); IEEE Fellow (1996); IEE JJ Tomson Medal (2000); 1999 U. of Illinois Outstanding Alumnus, 1991 IEEE Comm. Mag. best paper; 1995 ANSI T1 Outstanding Achievement Award; NSF Presidential Investigator (1987–1992), ISSLS 2004 Outstanding Paper award. Cioffi has published over 250 papers and holds over 40 patents.

email: Cioffi@stanford.edu

Mehdi Mohseni received the BS degree in electrical engineering from Sharif University of Technology, Tehran, Iran, in 2001 and the MS degree in electrical engineering from Stanford University, California, in 2003. Since 2003 he is working towards a PhD degree at Stanford University. His research interests include multi-user information theory, wireless and broadband communications.

email: mmohseni@stanford.edu