upATIS Committee NIPP-NAI
Working Group NAI (DSL Access) **NIPP-NAI-2007-158**
Vancouver, Canada: November 12, 2007

CONTRIBUTION

TITLE: **DSM Level 1: Polite Low-Power Single-DSL Stability**

SOURCE:

J. Cioffi	**Stanford**	+1-650-723-2150	**cioffi@stanford.edu**
W. Lee			
S. Jagannathan			
G. Ginis	**ASSIA**	+1-650-654-3400	gginis@assia-inc.com
P. Silverman			

PROJECT: **T1E1.4, DSM, Information only**

ABSTRACT

Level 1 Dynamic Spectrum Management (DSM) addresses individual DSL stability, diagnosing Level 1 DSM Data from an individual line and (if possible) stabilizing the line through Level 1 DSM-Control line re-profiling. Individual line diagnostic reports may also be provided by the Spectrum Maintenance Center (SMC), possibly facilitating service-provider actions for further DSL stabilization. This information-only contribution suggests best Level 1 stabilization methods using existing politeness capabilities while reducing power consumption. This contribution also cautions against some non-DSM methods that increase power consumption excessively.

- CONTACT: J.M. Cioffi, Cioffi@stanford.edu, Tel: +1-650-723-2150 ; Fax: +1-650-724-3652

DSM Level 1:
Polite Low-Power Single-DSL Stability (158)

J. Cioffi, W. Lee, S. Jagannathan

Dept of EE, Stanford University
363 Packard EE, Stanford, CA 94305-9515
Cioffi@stanford.edu,
P: +1-650-723-2150 ; Fax: +1-650-724-3652

G. Ginis and P. Silverman

ASSIA Inc.
203 Redwood Shores Parkway , Suite 250
Redwood City, CA 94065-1198
gginis@assia-inc.com ,
+1-650-654-3400 , Fax: +1-650-654-3404

ABSTRACT

Level 1 Dynamic Spectrum Management (DSM) addresses individual DSL stability, diagnosing Level 1 DSM Data from an individual line and (if possible) stabilizing the line through Level 1 DSM-Control line re-profiling. Individual line diagnostic reports may also be provided by the Spectrum Maintenance Center (SMC), possibly facilitating service-provider actions for further DSL stabilization. This information-only contribution suggests best Level 1 stabilization methods using existing politeness capabilities while reducing power consumption. This contribution also cautions against some non-DSM methods that increase power consumption excessively.

1. Introduction

Instability limits today's xDSL deployment range. Such deployment range is typically well below the DSL range exhibited in carefully crafted and standardized lab/interoperability tests, which were previously thought to have tested "worst-case" conditions. Instability is typically measured in retrain counts, packet-error counts, or combinations of both. If such measures are unacceptably high, the DSL is labeled “unstable.” Statistically, both instability measures (retrains, packet errors) increase with loop length, but many short loops can have high instability caused by large intermittent noises. Customer satisfaction, measured in complaint-call volume, number of dispatches (“truck rolls”), and/or customer turnover (“churn”), decreases with increasing instability; DSL service providers determine deployment range to maintain an acceptable level of customer satisfaction. Thus, deployment range is today almost always limited by such instability. Examples of such instability and its importance appear in references [1]-[5].

Video services are strongly affected by instability, but data and voice-over services also exhibit strong correlation between stability and customer satisfaction. Level 1 Dynamic Spectrum Management (DSM) has several effective mechanisms to stabilize DSLs and thereby to improve deployment range. ***These methods merit better understanding, use, and DSM-capability support through correct DSM-capability implementation, which is the purpose of this contribution.***

Level 1 DSM[1] [6] focuses largely upon the management of a single line. The Spectrum Management Center (SMC) collects line performance data through the DSM-D interface from DSM-capable equipment. The Level 1 SMC then assesses DSL stability. The SMC may consequently re-profile[2] the DSL through the DSM-C interface. Higher Level DSM systems may use all the data from a neighborhood or binder to reprofile single lines (Level 2 DSM) or even may use all such neighborhood/binder data for coordinated reprofiling of multiple lines (Level 3 DSM “vectoring”). The gains of higher-level DSM crosstalk and spectrum management are meaningful only when stability of lines has been first established, since such stability otherwise limits deployment range.

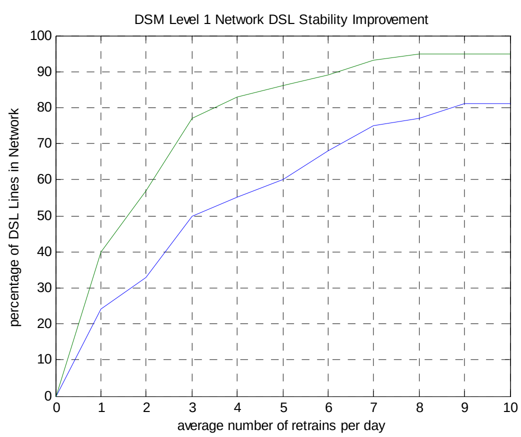

Figure 1(a) stability with bit swapping

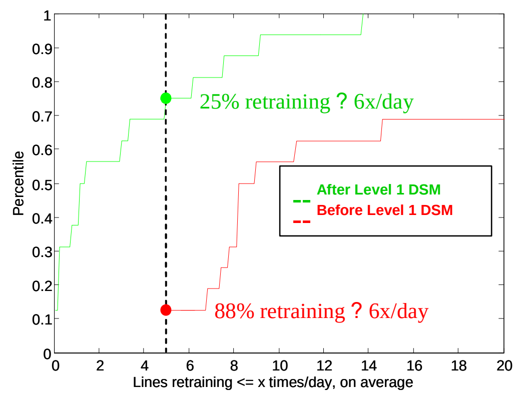

Figure 1(b) stability without bit swapping

Figures 1(a) and 1(b) respectively reflect average stability experience of field data taken from several DSL service providers around the world on bit-swapping and non-bit-swapping equipment. The plots do not reflect any specific service provider’s data, but rather a composite to guard the specific interests of any particular service provider. Figure 1(a) on the left illustrates the percentage of an average network that have fewer retrains per day than the abscissa value (horizontal axis). Figure 1(a) illustrates DSL stability before and after use of an SMC supporting DSM Level 1 for equipment that correctly bit swaps. If 5 retrains per day were used as a threshold for unacceptable stability, the stable-customers improvement is roughly 25%. Experience shows that such a stability improvement typically extends actual deployment range by at least 20%.

[1] Level 1 DSM is sometimes also abbreviated by “Dynamic Line Management” or DLM, although not in standards.
[2] “Re-profiling” is a term commonly used to mean changing a file of DSM-control parameters imposed through the DSM-C interface. Such a file is sometimes called a profile, whence the term “re-profiling.”

Figure 1(b) illustrates stability without bit-swapping. A larger stability improvement is possible with Level 1 DSM, but the total number of stable customers is understandably smaller. For instance at 5 retrains/day, the improvement is almost 60% of the customers in these non-swapping composite networks. With and without swapping, the number of stable customers is not 100%, even after DSM Level 1 correction – but the customer-satisfaction gains are very large (thus increasing deployment range 20% or more). Clearly, better equipment support of Level 1 capability (that is correct implementation of DSM-D and DSM-C interfaces) improves the percentages.

The SMC statistics on rate-range and stability for individual DSL lines, and overall as a function of equipment type also have value for manual corrective actions like targeted customer dispatch (presumably on the smaller number of unstable lines because of the SMC's automated corrective re-profiling of many troubled lines). These targeted manual actions provide additional stability gains above the levels in Figures 1a and 1b. Such statistics also may guide future equipment-type choice for new services or existing services' equipment replacement. Thus Level 1 DSM can substantially reduce instability, if used correctly. And consequently, the DSL deployment range increases.

Section 2 of this informative contribution examines polite Level 1 mechanisms for stabilization, including forward-error correction (Reed Solomon code), use and speed of bit swapping, and profile tiering. These methods are found to be robust to noise changes (induced by any mechanism including L2-mode-induced crosstalk variations) while holding power consumption to a minimum. Section 2 also observes that mechanisms that increase margins are less effective and simultaneously increase power consumption excessively compared to polite Level 1 mechanisms - specifically including but not limited to the static spectrum management method known in the ITU as "virtual noise" [7].

Section 3 concludes by suggesting some improvements in Level 1 DSM-capable equipment that would further enhance stability.

2. Polite Level 1 and Power Consumption Basics

This section enumerates some basic issues in Level 1 DSM, illustrating some existing field improvements and making recommendations on directions to pursue further in standards. Cautions also appear for some methods that reduce rates and increase power consumption.

2.1 DSL Power Consumption

DSL modems increasingly have power consumption determined by their analog driver circuits. While the power consumption of complex digital signal processing reduces to low levels with silicon-manufacturing advances (that is, reductions to 65 nm feature size and beyond), consumed analog driver-circuit power does not scale. Instead consumed analog power consistently increases linearly with transmit power independent of manufacturing feature size. By contrast, a digital integrated circuit in today's 65nm technology consumes roughly ¼ the power it once did in 130nm technology, and advances will continue to drive digital power usage to smaller levels yet. Thus, the power of a 100 mW line driver (20 dBm) remains proportional to transmit power and is roughly 500 mW [8]. This paper's analysis will use a factor of 5 throughout to relate consumed power to transmit power. Thus, a 17 dBm transmit driver would consume 250 mW, while a 10 dBm transmit driver would consume about 50mW. Presuming other digital processing scales down in pro-forma good DSL chip designs, the majority of power-reduction opportunity is in this paper assumed to be in transmit-power reduction.

Certainly L2 mode can also reduce both analog and digital component power roughly proportional to the de-energized/total time ratio. However, L2 mode can introduce more time-varying noise, simultaneously thus contributing to instability. Thus, use of L2 mode additionally motivates use of Level 1 DSM stabilization capability. If an average video (increasingly a driving application for DSL) customer observes their television for 3-6 hours per day, the opportunity for actual L2 power reduction may be roughly up to 9-10 dB or a factor of 10. Clearly, such a factor of 10 should not consequently lead to more than a factor of 10 increase in L2-on-time transmit power to overcome time-varying crosstalk, or L2 mode's intent would thus be defeated. Some non-DSM approaches that increase margins by 10 dB or more, or other approaches that insert large virtual noises will thus defeat L2 mode's power savings. Thus, the Level 1 DSM stabilization should be polite, particularly when L2 mode is used.

As a simple example, suppose a good level 1 DSM system reduces transmit power by a factor of 7 dB with respect to a non-DSM system. Then the maximum power of 500 mW is reduced to 100 mW, a savings on average of 400 mW. Over a 5-million-DSL network, this is a savings of 2 Megawatts.

2.2 DSM Coding

DSL standards require use of Forward Error Correcting (FEC) codes, which were specifically proposed (see [9]) to mitigate time-varying noises of all types to which the bit-swapping mechanism may not be able to adapt. An FEC code will augment $R \leq 16$ parity bytes within a codeword of $N \leq 255$ bytes and can thus correct up to 8 bytes in error within that codeword[3]. A single isolated noise "strike", or one repeated at periodic intervals, is one type of intermittent noise known as "impulse noise". Such a strike can render many bytes in error, so interleaving was introduced to distribute the effects of such an isolated impulse so that it is corrected by the FEC [9]. However, periodic impulses or those that occur often but irregularly can easily overwhelm interleaving – indeed interleaving depths of more than 8-10 ms (rectifing noise at the 50-60 Hz power-line repetition rate) often provide no additional stability for periodic impulses. Thus, the ratio of R/N needs to be increased to a level sufficient to correct nearly all errors. The original ADSL1 standard ANSI T.413-1995 did not limit the R/N ratio (see [9], which was eventually adopted in ADSL1). Higher R/N ratios imply shorter N since $R \leq 16$. Shorter N means the FEC encoder and decoder must run faster (at any given data rate). Extremely short DSL N choices as low as 24 can be useful on some DSLs, particularly short lines with data rates that would otherwise be very high except for intermittent noises. More typically, values of N=128, 64, 48, and even 32 can be useful.

Trellis codes are not forward-error-correcting in a strict sense and instead are used to improve performance with stationary noises (see the original trellis code proposal in [10] and the arguments made there). On a line with stationary noise, a trellis code helps improve the SNR margin by up to 4.2 dB in the absence of FEC[4]. However, such gain is of little value on a chronic unstable line where the intermittent noise dominates. Impulse noise with a trellis code often instead leads to error propagation that increases the length of the resultant error burst. Indeed SMCs for such lines might best turn off the trellis code and increase the R/N ratio to a level where customer satisfaction (that is packet errors) is acceptable. SNR margin, which was not limiting in this intermittent-noise case, will reduce (because the R/N exceeds 8%), but SNR margin has less meaning on such a line. Furthermore, the consequent higher-R/N error rate on the line will become acceptable and SNR margin, an exclusively stationary-noise approximation, is no longer the most important chronic-line characterization for customer satisfaction. If the error rate is low or zero with and without the intermittent noise present, the SNR margin need not be

[3] Erasure decoding can increase this number of correctable bytes up to 16.

[4] FEC codes with R/N approximately at 8% can also provide about 3.8 dB of coding gain against stationary noise in DSL. The use of both together can improve SNR margin to about 6 dB with sophisticated iterative decoding methods on lines without impulse noise, with R/N closer to 4%. This use of FEC, however, is also based on stationary noise – such coding gains are largely meaningless on lines with intermittent/impulse noises.

high. Indeed, target and minimum SNR margins can often be reduced on such lines to reduce unnecessary retraining – that is, if there are zero errors, there is no need to trigger a retrain no matter what the margin level.

Figure 2's example considers an R/N ratio of 1/3, a powerful setting for correcting intermittent-induced error rates as high as 17% of the bytes. Such a setting may see a coding gain that is only about 1 dB instead of 3.8 dB (trellis coding presumed off – trellis coding on such a line can introduce negative coding gain because of error propagation in the Viterbi decoder). Thus, the system needs to use 3 dB more power, or better yet, now accept 3 dB less margin (for no power increase). The lower (6-3 = 3 dB) margin is not important because the SMC has determined the major problem on the line is intermittent noise by observing the service over long periods of time. Indeed, the margin might even be set lower than 3 dB because this line is stable and the underlying DMT modem will maintain synchronization easily even at 0 margin where the FEC would still eliminate all 1e-7 stationary noise errors). Thus, this solution is polite – power consumption has not been increased (and even possibly decreased) and the chronic line's customer is now satisfied with near error-free performance. As long as the line was not already at maximum length for the stationary noise, additional parity can be added. Deployment range consequently statistically increases because more lines are stable and the service provider can extend deployment range.

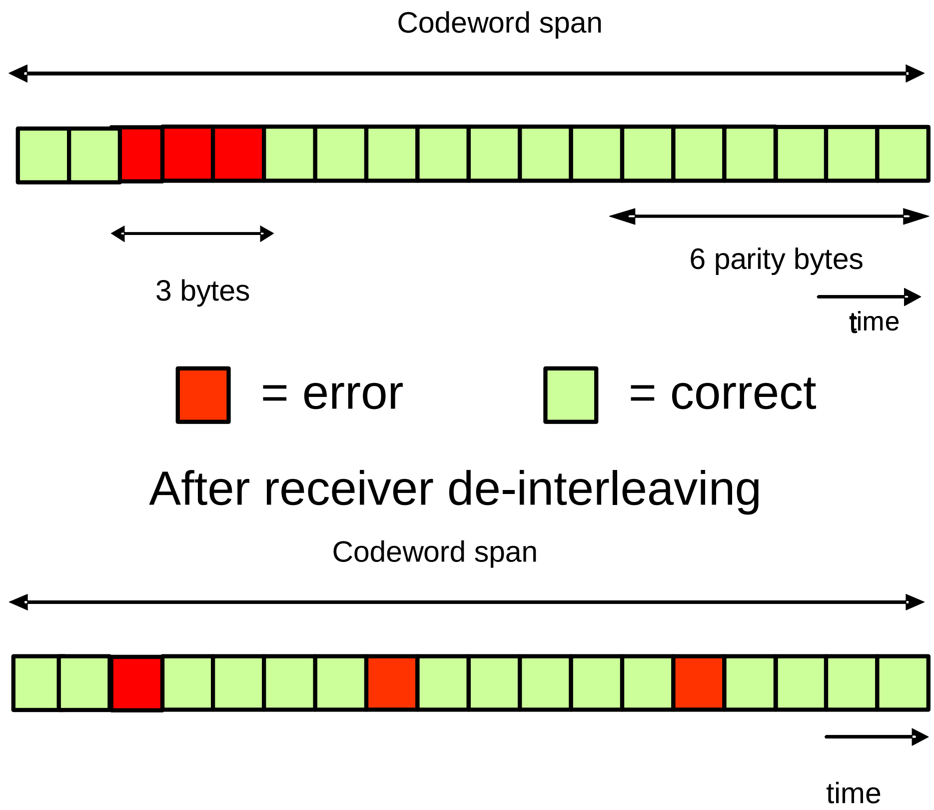

Figure 2 – Illustration of 3/18 bytes in error that R/N = 1/3 can completely correct.

By contrast, an often-implemented non-DSM alternative is considered. To understand the alternative, its origin is first considered in the indented section immediately to follow in this paragraph:

> *Somehow confusion slipped into pre-DSM equipment-provider simplifications; for instance some vendors introduced a Fast, Interleave, and Medium interleave simplification that unfortunately keeps the R/N ratio at about 4-8% for all data rates*

and maximizes SNR margin. The interleave setting is typically 16 ms or more in Interleave Mode (Values of interleave depth above 8ms only help on isolated impulses and usually not on 60Hz repetitive impulses, so 8ms might have been just as good for interleaving – however the argument here is valid for either setting). A better solution is to return to the originally specified and intended higher R/N possibilities (see [9]). Later ADSL2/2+ and VDSL2 standards introduced a concept of INP (impulse noise protection) parameter that was defined as "how long of an isolated impulse" can be corrected in units of 250 μ*s (DMT symbol periods) so that consequent interleave depths and R, N settings would be determined. Typically, vendors interpreted higher INP as "increase the interleave memory/depth/delay" because the INP definition was for an isolated impulse. A second parameter DELAY has to be set low to force higher R/N ratios. Indeed an INP/DELAY ratio close to 1 can force R/N =1/2 (50%), a good setting for a line with chronic intermittent noise. However, the settings in the ADSL2/2+ and VDSL2 standards largely restrict INP and delay value choices so that this high R/N ratio is not achievable (thus in some sense going backwards from the ADSL1 standard that originally allowed such ratios – even if vendors may not have implemented them). Thus an improvement would be to allow INP/DELAY ratios as low as 1 at all speeds, which implies neither large memory nor high interleave delay (but does imply a fast decoder).*

Thus, with some vendor equipment's R/N ratio restrictions, the service provider then feels forced to increase the SNR margin variable from 6 dB to 16 dB (a factor of 10 increase) or more. Because the intermittent noise is so large, this margin increase is often still insufficient to eliminate the customer-satisfaction problem. However, now the modem uses 10 times more analog power to get the same data rate that could have been achieved with fewer errors by a higher R/N ratio at much less power.

Thus margin increase is not polite – and indeed tends to defeat any power savings that might have otherwise been reduced by L2 mode.

Further, wide L2 mode use would introduce more intermittent noise than already present (typically already 10-30% of the customers), thus increasing overall power use if the "margin-increase" approach were maintained in an (largely unsuccessful) attempt to reduce power. Clearly this was not the intent of L2 mode. A better solution would be use of L2 mode with better FEC R/N ratios selectively applied to any line affected by L2 time-varying crosstalk.

Thus, FEC offers a very good power-consumption and customer-satisfaction solution, allowing better stability and better use of L2 mode to amplify power savings. Margin increases instead increase consumed power, and have relatively little effect on stability with intermittent noise.

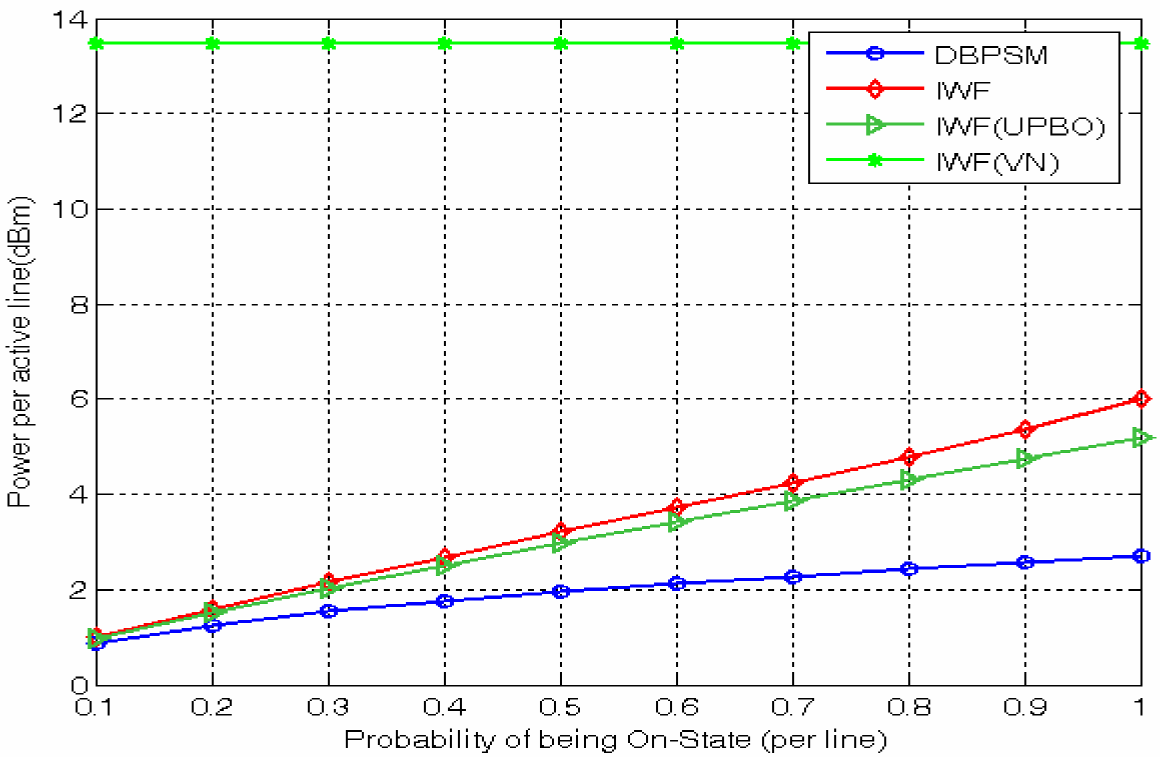

Figure 3 – Comparison of power consumption

A third possibility has been introduced in ITU standards under the name of "virtual noise." This method is basically equivalent to a frequency-dependent margin with respected to a presumed reference noise. The virtual noise is interpreted as a large (possibly maximum) noise that the line experiences intermittently. Similar to margins, this virtual noise may have to be set very high for large intermittent noises, causing a higher transmit power to be constantly used when the modem is on. Higher transmit power means higher consumed power, and of course a simultaneous increase crosstalk into victim lines. When the victim line senses the higher crosstalk, the victim line, then consequently also at larger power could increase its crosstalk into other lines including the original offending line, forcing those lines to use virtual noises of higher level, thus increasing their power and crosstalk. Figure 3 shows one example of the increased overall power by this approach compared to polite spectrum balancing methods for 10 VDSL upstream users with fixed rate targets [14].

Virtual noise proponents would argue that virtual noises for all users could be simultaneously set to cause power spectra usage that effects spectrum balancing – however, such a mechanism loses all dynamics and ability to reduce power consumption when worst-case situations are not present. However, such cases do not correspond to OSB generally, and the PSDMASK instead would be the preferred stationary-binder solution. For a dynamic situation, again such a central imposition of spectra is unnecessarily complex as in [14]. The overall situation then is all lines using higher power with greater crosstalk from other lines, so all then exhibit lower data rates and/or higher power consumption. Thus, any stability comes at a premium in an increase in consumed power and loss of data rate when compared to a prudent use of FEC. This method should be avoided – indeed it poses a regulatory issue of severe consequences in that a user of virtual noise harms other lines unnecessarily (FEC or other polite methods could have been used and introduced less harm, so a network operator sharing the cable plant with a user of virtual noise would have a legitimate argument of having experienced undue harm).

2.3 Bit and Gain Swapping

Bit-swapping is essential to DMT transmission [11] and is mandatory in DMT-DSL standards in use today. However, dysfunctional or non-functional bit-swapping equipment appears in some networks (see Figure 1b for the average retraining counts across such networks). Some intermittent noises will remain sufficiently long (that is, be "quasi stationary") for bit and gain swapping to move bits and energy among tones to maintain the desired data rate (or maintain the best specified margin for fixed-rate service). There is a transient when swapping occurs. Typically such transients are protected by the FEC discussed in Subsection 2.2. The noise need not be identical in spectrum each time it occurs, so swapping then allows dynamics of the modems themselves to act locally to do the best they can. If they fail, a retrain occurs. However, clearly from Figures 1a and 1b that represent experiences across huge samples of customers, such retraining occurs less often when bit-swapping is used. Use of high margins does not "turn off" bit swapping but does increase power consumption as in Subsection 2.2. Thus, use of these options is increasingly unfavorable when bit-swapping is used.

There are some basic parameters that guide bit and gain swapping in Table 1:

Table 1 – Some swapping-relevant control parameters

Parameter	Explanation
TSNRM	**Target SNR margin** The margin that the modem tries to attain and maintain in rate-adaptive DSLs.
MINSRM	**Minimum SNR margin** The margin below which the DSL modem will force a retrain
MAXSNRM	**Maximum SNR margin** The largest margin the modem's swapping should allow (sometimes excess above this level is instead minimized if the modem's internal gain settings are at standardized limits)
Gains range	**Range of allowed gains** Typically +2.5 dB to -14.5 dB
MAXNOMPSD and/or PSDMASK	**PSD Limit** The level below which gain-swapping must maintain the PSD.
minRate	The **minimum data rate** below which the modem should retrain (or report failure if not achievable)
maxRate	The **maximum data rate** at which the modem should operate
PCB	**Power "Cut Back"** The amount of power reduction (sometimes frequency dependent) that a receiving modem suggests is acceptable to a transmit modem during training.

Typically, these parameters have been well-specified in DSL standards. An important trade-off is between the range of gains and the PCB parameters. Initially a receiving modem should be able to determine if power use is excessive so it can request lower power and adjust its analog gain-control settings accordingly (so that best ADC range is maintained). The fine-gains adjustment of over 17 dB can then be used to maintain margin at or below MAXSNRM. PCB should not be set too low so that the modem would not be able to recover from an intermittent noise. Level 1 DSM uses a MAXSNOMPSD parameter (generalized to a frequency-dependent mask in Level 2) that an SMC can use to help the modems with initial prudent power level setting. If the SMC has seen large noises, the PSD level requested is an indication to the modem that excessively large PCB is unwise, thus maintaining the 17 dB gain-swapping range.

Understanding of this trade-off between SMC and vendor choices would be useful in further level 1 DSM specification (as well as higher level DSM specification). Field experience suggests that line noise is highly dynamic, causing variations in data rate by a factor of 3 or more over the course of a year when swapping is not used well. The best solution is to guide the modem's dynamics effectively with the value selection in Table 1 from an SMC so that stability is maintained, but not to attempt precise limitations on the modem (like trying to guess the exact worst-case noise and thus freeze the modem).

2.3.1 Excessive field margins

Presently, some equipment does not adhere to standardized implementation of MAXSNRM. (The contribution from ASSIA and Stanford, NIPP-NAI-2006-114 in August 2006, showed a number of examples.) With such equipment, most of the stable lines have margins exceeding 20 dB in a typical deployment, even if the MAXSNRM is set at lower values than 20 dB – the number of lines with at least 20 dB (100 times more power than they need if stable) often exceeds 80% in even advanced DSL networks today. This is a standards compliance violation, but such violations are common. Clearly, with or without L2 mode, such non-compliance causes a large overuse of power. Secondly, it causes excessive crosstalk between lines, thus reducing the achievable rates of any longer lines that need full power, and thus contribute to instability.

Current tests of MAXSNRM are grossly insufficient in that they look only for a small reduction in power when the noise is reduced, but do not register that the margin may be greater than the MAXSNRM. Where possible, an SMC could observe such non-compliance and if the MAXNOMPSD function is working (or some equivalent mechanism for reducing PSD), then this can be set and thus emulate the desired function. Often the manufacturers who have not complied with standardized MAXSNRM also prevent change of the PSD, a second standards-compliance violation, so large crosstalk may simply be a consequence of serious multiple simultaneous non-compliances. DSL stability and range would thus improve if adequate tests and a clear shared understanding within the industry of the use of MAXSNRM were to be developed promptly.

2.3.2 Swap Speed

Presently, swap speed must be at least one bit swapped every 800ms. This is excessively slow. Numerous investigations with faster swapping ([16]) have found that the processing of swapping to approximate a water-fill energy use is convergent at any speed. Thus, the main limitation is ability of the DSL modem control electronics to implement a swap. VDSL1 [12] introduced an optional express swapping that allowed large numbers of bits to be simultaneously swapped. The SOS [13] protocol currently under investigation for VDSL2 would restore that express-swap capability in some situations of interest.

However, a simple improvement could be based on the observation that swaps occur on super-frame boundaries, which occur at least every 17ms in DSLs (and faster in new DSLs). Swapping one bit on each such boundary corresponds to more than 50 swaps per second, more than an order of magnitude improvement in swapping speed (and thus reducing any transient's duration from an abrupt noise change). Swapping of 10 bits per superframe would cover a wide range of abrupt noise changes. Thus, somewhat like FEC, an existing mechanism's limits might much more easily be enlarged than specification of entirely new approaches that may be less effective.

2.4 Profile Tiers

Profile tiers are a highly effective Level 1 DSM mechanism implemented in the SMC that leads to the large gains in Figures 1a and 1b. Unfortunately modem FEC implementations prevented wide use of the methods in Subsection 2.2. No margin increases (nor virtual noises) were used in Figures 1a and 1b. The basic idea is the SMC observes the line performance over an interval in time, including retrains and DSL speed variations. The minRate and maxRate of Table 1 are thus set to ensure that the modem does not retrain often. Typically, a wide range of data rates offered for a "rate-adaptive" DSL service are tiered into overlapping narrow ranges for each customer (for which all have a minRate that exceeds the minimum for which the customer paid). Each customer's minRate and maxRate is thus adaptively determined so that the customer sees a stable low-power-consumed data rate with little errors or retrains. Such tiers leads to at least 20% gains in deployment range for all speeds, and thus increases the actual customers' speeds.

DSL Customers who experience a frequent day-to-day variation in data rates often complain or drop service. Thus, "seamless" rate adaptation may not improve customer satisfaction – lower retrain counts and few packet errors do improve customer satisfaction, but variation in rate is another indication of instability. However, to date, the authors are not aware of any functional seamless rate adaptive installations, so there has been no opportunity to correlate that rapid rate change with customer satisfaction. However, slower rate variation (of course accompanied by retrains) has high correlation with unsatisfied customers – thus, it is reasonable to believe that even faster rate variation might also be viewed as unstable by a customer. And again, instability limits deployment range. The goal of Tiered Rate Adaptation (TRA) is to avoid this instability by providing the customer with data rates that are either stable.

Tiered Rate Adaption (TRA), as described above, can improve deployment range by more than 20% in Level 1 DSM. As FEC R/N range is improved and use/speed of swapping improve, even yet larger improvements occur in Level 1 DSM at power savings in the range of megawatts for a large DSL network. Large margins and/or virtual noise increase power consumption.

3. Some Suggestions for Polite low-power stability (and effective L2 use)

This section concludes with suggestions for serious consideration in Issue 2 of the DSM Report and in DSL Standards. All would help L2 modes achieve their intent.

3.1 Return to original FEC R/N wide range

The current fast, interleave, medium delay settings, as well as the range of INP and DELAY parameters is not enough. Ratios of R/N = ½ at all speeds should be allowed and be programmable by an SMC.

It should also be possible for an SMC to independently turn on and off both trellis codes and FEC.

3.2 Increase Swapping Speed

Swapping speeds of 50 to 500 swaps per second should be considered, particularly in higher-speed DSLs.

3.3 Enforce MAXSNRM

Tests that actually measure the margin and compare it to the requirement should be devised and used in compliance and interoperability testing.

3.4 Tiers

Service providers might want to better understand profile tiers and their use in conjunction with all the recommendations in this section, and in particular when FEC, swapping, or margin-based solutions cannot be implemented fully because of equipment conformance violations.

3.5 Avoid solutions that force high-power-protection use

High-margin-setting solutions increase power consumption and reduce rates, but do retain some dynamics. They should be used only when all other options are exhausted. Use of virtual noise causes even higher power consumption and should not be encouraged for use anywhere. Use of virtual noise also can lead to unfair harm between co-located service providers.

4. References

[1] T. Starr (ATT), "Experiences from Dynamic Spectrum Management," IEEE Globecom, Access 2007, Washington, DC, November 29, 2007.

[2] J. Le Masson, R. Le Meur, and H. Mariotte (France Telecom), "DSL Access Advantages from DSM in France," IEEE Globecom, Access 2007, Washington, DC, November 29, 2007.

[3] K. Foster, (British Telecom), "DSL Access and Gains from DSM," IEEE Globecom, Access 2007, Washington, DC, November 29, 2007.

[4] E. Berndt (Deutsche Telekom – T-Systems), "Goals for Future Access Networks and Economic Propspects," Broadband World Forum, Berlin, Germany, October 8, 2007.

[5] M. Polano (Telecom Italia), "Next Generation DSL: Pushing the Frontiers of Broadband," Broadband World Forum, Berlin, Germany, October 8, 2007.

[6] ATIS Dynamic Spectrum Management Report – Issue 1, May 2007.

[7] D. van Bruyssel, "G.vdsl: VDSL2 Virtual Noise Mechanism: Proposed Text," (Alcatel), ITU SG15/Q4 contribution HA-073, March 21, 2005, Huntsville, AL.

[8] M. Mohseni and G. Ginis, "G.gen: G.vdsl: G.adsl: Margin Cap Mode achieves power savings," (ASSIA), ITU SG15/Q4 Contribution RJ-075, October, 8, 2007, Red Bank, NJ.

[9] J.T. Aslanis, P. Tong, and N. Zogakis (Amati/Stanford), "An ADSL Proposal for Selectable Forward Error Correction with Convolutional Interleaving," ATIS Contribution T1E1.4-1991-180, August 20, 1992, Portland, Oregon.

[10] P. Tong, J. Aslanis, and J. Cioffi (Amati/Stanford), “A Trellis-Code Proposal for ADSL Line Signals,” ATIS Contribution T1E1.4-1991- 176, August 20, 1992, Portland, Oregon (delayed from May 1991 meeting).

[11] J. Cioffi, “The Essential Merit of Bit Swapping,” ATIS Contribution T1E1.4/1998-318, Dallas, TX.

[12] ITU Standard G.993.1 (VDSL1), 2003.

[13] Sigurd Schelstraete, et. al, *G.vdsl: SOS working text proposal*, ITU-T SG15/Q4 – Temporary Document, RJ-084, Redbank NJ, October 2007

[14] W. Lee, S. Jagannathan, J.Cioffi, G. Ginis, and P. Silverman “Higher-Rate Level 2 DSM Power-Saving Examples,” *ATIS Contribution NIPP-NAI 2007-160,*Vancouver, Canada, November 12, 2007.

[15] S. Jagannathan, ``The need for Margin Cap Mode Lab Test Results,”: ATIS contribution NIPP-NAI-2006-114, August 30, 2006, conference call.

[16] S.T. Chung, ``Transmission Schemes for Frequency Selective Guassian Interference Channels,“ Ph.D. Dissertation, Stanford Univerisity, November 2003.