ATIS Committee NIPP-NAI
Working Group NAI (DSL Access) **NIPP-NAI-2007-159**
Vancouver, Canada: November 12, 2007

CONTRIBUTION

TITLE: **DSM Level 2: Polite and Efficient Access-Network Throughput Increase**

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: **DSM, Information only**

ABSTRACT

Level 2 Dynamic Spectrum Management (DSM) addresses spectrum balancing across multiple DSL Lines in a cable/binder. DMS Level 2 analyzes simultaneous Level 2 DSM-Data from multiple DSLs, and increases data rates via Spectrum Management Center (SMC) re-profiling of single DSL lines in the cable/binder through the DSM-Control interface. Level 2 DSM re-profiling considers crosstalk between lines and can use programmable power spectral densities and/or other inputs to DMT-based DSL loading algorithms. Neighborhood-level diagnostic reports may also be provided by the SMC, possibly facilitating service-provider actions for further DSL access-network improvement. This information-only contribution suggests Level 2 methods for using true noise margins and Level 2 strong/weak user classification to increase speeds.

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

DSM Level 2:

DSM Level 2: Polite and Efficient Access-Network Throughput Increase

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 2 Dynamic Spectrum Management (DSM) addresses spectrum balancing across multiple DSL Lines in a cable/binder. DMS Level 2 analyzes simultaneous Level 2 DSM-Data from multiple DSLs, and increases data rates via Spectrum Management Center (SMC) re-profiling of single DSL lines in the cable/binder through the DSM-Control interface. Level 2 DSM re-profiling considers crosstalk between lines and can use programmable power spectral densities and/or other inputs to DMT-based DSL loading algorithms. Neighborhood-level diagnostic reports may also be provided by the SMC, possibly facilitating service-provider actions for further DSL access-network improvement. This information-only contribution suggests Level 2 methods for using true noise margins and Level 2 strong/weak user classification to increase speeds.

1. Introduction

Level 2 DSM balances customer data-rate and power-consumption interests dynamically. With a stable binder (presumably stabilized concurrently or previously through the use of Level 1 DSM, see[1]), crosstalk between users of the same service provider can be balanced (while any other service provider's users are treated as noise via Level 1 DSM). Level 2 DSM collects data from all the SMC's subtended lines in a binder, but applies single-line control mechanisms (to change margins and other parameters to the individual DSL modems' loading algorithms). This contribution describes some basic physical-layer DSM actions that lead to significant data-rate improvements.

The concept of a rate region is fundamental to Level 2 DSM. Rate regions, such as Figure 1, illustrate the possible allowed combinations of data rates, sometimes called "rate tuples." For most cross-talking DSL binders, the rates of the individual customers are not independent. The rate region illustrates the trade-offs between the various combinations. A level 2 DSM SMC can conceptually compute the rate region and store it for the users in a particular cable or binder of DSLs[1]. No individual DSL can compute or know this region because it has no access to the cross-talking transfers known as Xlog in the DSM Report. If individual DSLs acted independently, then some internal rectangular region of Figure 1 would represent the possible data rates instead of the full region, consequently reducing the data rates and possible revenue for the service provider. A combination of desired customer data rates is attempted if it is safely within the rate region for the binder. A rate tuple that is within the region is feasible. The rate region is a function of crosstalk and other noises. Other noises can include non-DSL noises like AM noise or background noise, but also includes crosstalk from any lines that may not be subtended (controlled) by the SMC – for instance, the crosstalk from another service providers DSLs would be treated as uncontrollable noise. Level 2 DSM is thus completely compatible with unbundling. The size of the rate region may be smaller for each and every service provider when unbundling is present. Nonetheless, this smaller unbundled rate region can still be considerably larger than what is possible if there is no SMC. Level 2 DSM increases this rate region.

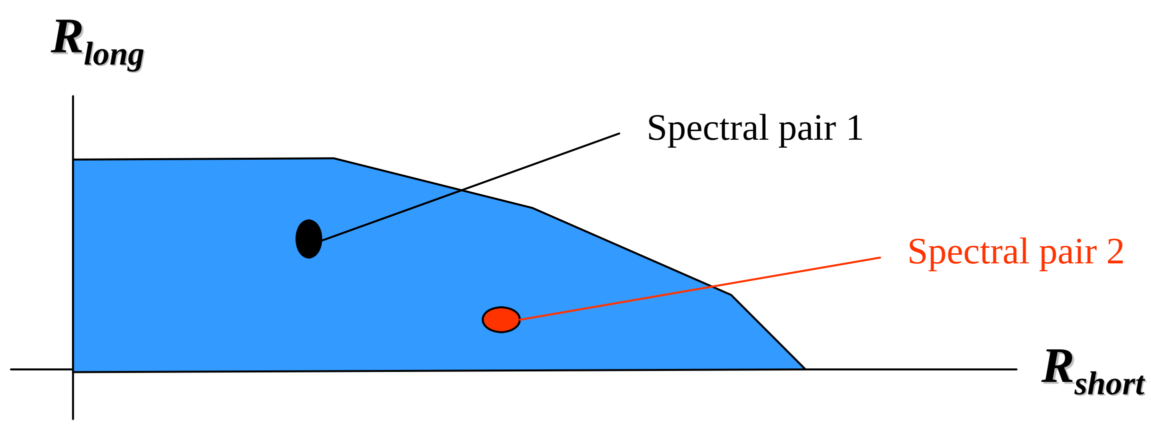

Figure 1 – basic rate region.

DSL binders can be highly dynamic. Typically, time variation of crosstalk transfer functions and line insertion losses is small (at most a few dB over wide temperature ranges). However, the actual DSL signals themselves can be time-varying much more rapidly. Energization or de-energization of DSLs can occur and create significant changes in crosstalk noise, both from same-SMC and uncontrollable unbundled DSLs. A DSL spectrum will also change significantly from the use of power-saving modes, sometimes called "L2" modes. Bit-swapping and gain-swapping algorithms can also introduce modest changes in DSL spectra with time. Non DSL noises can also change. An SMC is responsible for surveying the binder's lines over a period of time to observe the dynamics of the lines and the noises (both those it controls and those noises the SMC can only observe in reported signal-to-noise ratios). Once observed, sanguine SMC decisions can increase speeds with acceptably small probability of an attempted rate tuple outside the rate region or equivalently acceptably low probability of a fault. Any reasonable SMC's Level 2 DSM strategy will consider time variation for continued line stability (as with Level 1 DSM [1]), but also to exploit opportunity in the enlarged Level 2 DSM rate regions.

[1] Actual calculation and storage of the rate region is conceptual. An actual SMC would instead test desired user rate tuples to see if they are in the region rather than compute and store the entire region. Such tests are often called "admission tests."

Section 2 discusses DSM Level 2's use of a bi-lateral loading algorithm politeness strategy to improve rate regions. Section 3 discusses further improvement of rate regions through improved use of user-specific margins. Section 4 suggests some areas for standards focus that assist Level 2 DSM progress.

2. Politeness: Strong versus Weak users

A simple Level 2 DSM strategy uses a Level 2 control bit to alert individual DSLs to use one of two loading strategies (called 'strong" and "weak" in [2],[3]). These two loading strategies dynamically allocate energy and the number of bits to each of the tones in a DMT DSL modem. Each strategy makes use of the same well-understood and standardized bit-swapping protocols. One strategy is the usual DMT "water-filling" or "greedy" method that rewards the tones with best SNR with most information/bits. This algorithm is called the "weak user" strategy and is typically used by DSL customers who have limited usable line spectrum available. This loading strategy continues to maintain politeness through maximum SNR margins. The other alternative strategy re-allocates energy from best-SNR tones below some cut-off frequency in an attempt to reduce crosstalk into other lines. This second strategy is known as the "strong-user" strategy and is exceptionally polite. Either loading algorithm continues to observe binder dynamics and retains the DSLs essential ability to react to changes in noise, thus guarding against instability. Further, such dynamic reaction substantially reduces power consumption and their emitted crosstalk into other lines.

This contribution, like [1], will assume that the sum of transmit powers is best minimized for a stable system with any given data rate. Thus, sum of transmit power will be used as a criterion for consumed power, because consumed power typically scales with transmit power. Sometimes some users may be more offensive than others in terms of emitted crosstalk. In those cases, those users may be more strongly weighted in consideration of their politeness.

2.1 Band Preference

Originally the band-preference concept [4] recognized that short lines, or more generally strong users[2], best de-emphasize a low-frequency band. The basic concept is that a strong line may have other victims in the binder who must share the use of this low-frequency band. Thus, if possible, the strong user should politely reduce its use of the lower frequencies. Such frequency-selective politeness augments the well-known maximum margin concept that should be used on all lines. This concept is most simply used by an SMC partitioning its users within a binder into two classes "strong" and "weak." The SMC then sends an indication (typically a bit) to the strong users to use the exceptionally polite algorithm. Further, when this weak/strong bit is set to strong, it is useful for the SMC to also send a single cut-off frequency to the modem to strengthen the performance of the strong algorithm. This frequency suggests to the DSL's strong loading algorithm to extract as many bits as possible from frequencies below the cut-off and place them in a water-fill fashion that applies only to frequencies above the cut-off frequency. If the cut-off frequency is not supplied, the DSL modem can attempt to approximate it by various methods described in [3]. Frequencies below the cut-off can be partially (or completely in extreme cases) vacated. A typical strategy would be to move the bits requiring most energy on tones below the cut-off frequency to tones of least incremental energy above the cut-off (after initially running the normal water-fill algorithm on all tones). Such a strategy leaves a second water-fill energy distribution below the cut-off frequency, but the

[2] It may be that a long line with very low noise is a strong user and a short line with very large noise is a weak user, so simply saying "long/short" or "near/far" can be slightly inaccurate. "Strong/Weak" more accurately conveys the concept with respect to the desired data rate that the DSL customer should receive.

water-level is different than the water-level above the cut-off. Figure [2] illustrates the results of such a situation described in more detail in [3].. Generally, one cut-off frequency for all the strong users is good enough to achieve near-optimal performance. If one strong user starts moving bits to the tones over its cut-off frequency, then it is reasonable to allow other strong users to use that frequency region because weak users will suffer from strong interferences in that region if they tried to load bits anyway. Therefore, the cut-off frequency is determined as the minimum of all cut-off frequencies which may be required by strong users to achieve rate targets.

2.2 Rate Region Enlargement

This subsection briefly illustrates an example better detailed in [3]. Essentially 10 lines of lengths uniformly distributed between 1600' and 4500' are used for upstream VDSL2 within a binder. The crosstalk within the binder uses the accepted 1% models in the DSM Report and other standards. Figure 2 attempts to summarize the 10-dimensional rate region in two dimensions by holding 4 strong users, designated by the SMC, at high fixed rates of [10 8 6 5] Mbps and plotting the sum of the 5 longest-line (weak) users versus the data rate on the 2nd-shortest line. The Level 2 DSM rate region is enlarged with respect to the Level 1 DSM rate regions through the use of the strong algorithm on the 5 shortest of the lines. Both weak and strong users achieve higher data rates through politeness than is possible with Level 1 alone. Stability is maintained through Level 1 DSM methods that manage FEC and maximum margin parameters on all lines. Data rate gains are roughly 10 to 50% depending on which user and which operating point is selected. All users on all curves use a 6 dB margin.

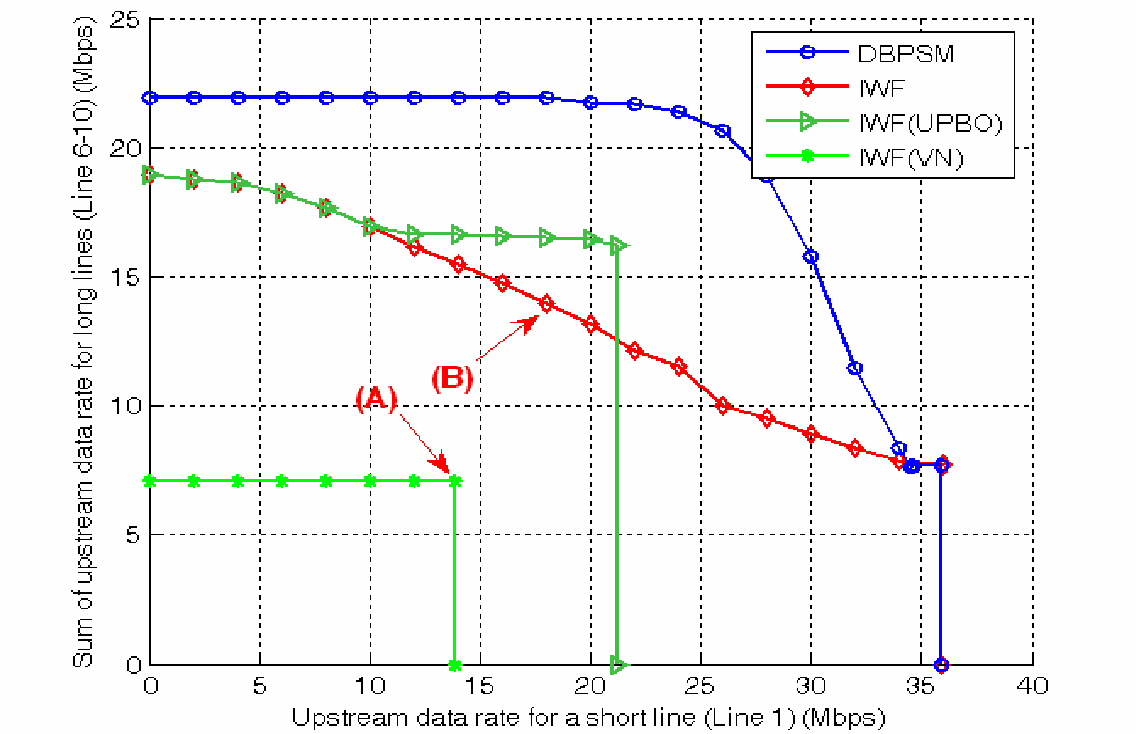

Figure 2 – Rate regions for 10 upstream VDSL users

The curves in Figure 2 also show performance for a particular choice of "virtual noise" suggested in the literature [7], [8] where the choice of VN is the worst-case level of the crosstalker for that line (so thus, the situation corresponding to the line training without crosstalk and then the crosstalking line suddenly energizing). Other VN choices are considered in [3] without change in the basic observation that the

DBPSM and politeness is better – there is no VN choice that can lead to the DBPSM performance level generally. The data rates are lower for all users because the spectra cannot adapt to the best balance between cross-talkers.

2.3 Power Reduction

Figure 3 instead plots the sum of the power transmitted by all 10 lines (an indicator of the power consumed as in [1]) for the different methods shown. A fair power-consumption comparison selects a point corresponding to the same data rates on all lines. Figure 3 attempts to illustrate such fair comparisons by indexing the horizontal axis to a VN reference, which is equivalent to picking a point on the VN curves (since all methods can achieve these lower-rate points) and then summing the transmit power of all users. As the VN index approaches 1, the methods of choosing the worst-case crosstalker turning on in [7] and [8] is increasingly used at VN=1 on the right. VN=0 corresponds to no use of virtual noise. Clearly Level 2 DSM can gain about 3 dB (or reduce power consumption by a factor of 2) for the same data rates as Level 1 DSM. Furthermore, the VN, since it causes victim modems to transmit at high levels and thus increases crosstalk, does the worst and uses almost a factor of 10 more total power to achieve the same rate point. If the power saved by DBPSM were instead used to maximize rate/reach the data rates would approximately double.

The increased power usage by VN is more prominent when dynamics of lines are also considered, as in Figure 3b. As the probability of being in the on-state approaches 1, all lines are in the on-state most of the time, while the probability of 0 means no line is in active (on-) state. With this dynamic configuration of lines, the spectrum balancing method allows the use of appropriate power levels. However, VN does not profit from this possible power-reduction opportunity because of static setting of worst-case noise.

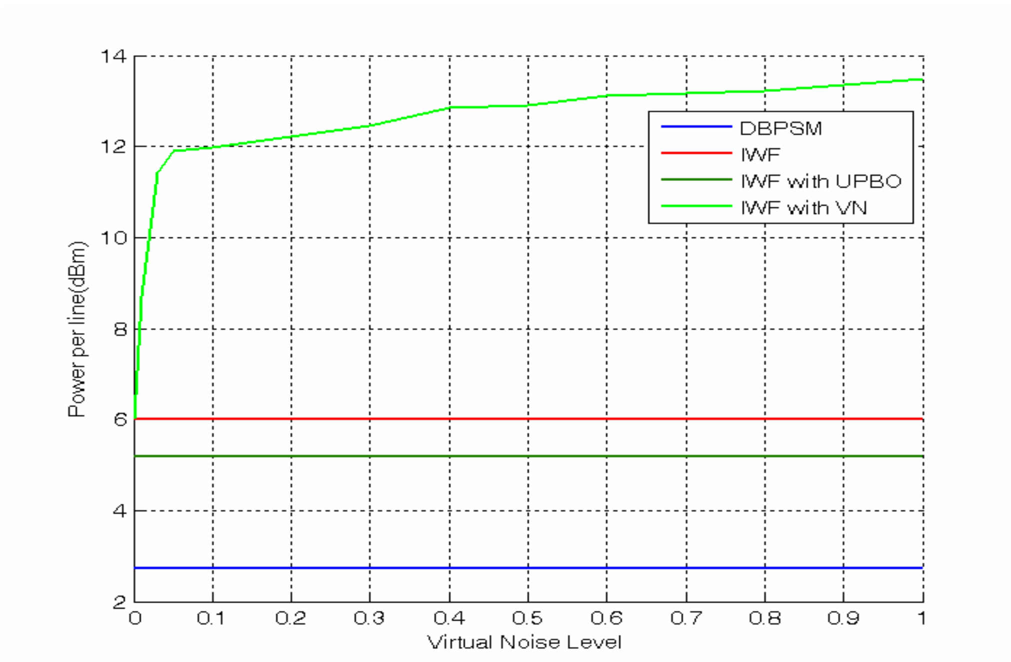

Figure 3 (a) – Power reduction for 10 upstream VDSL users

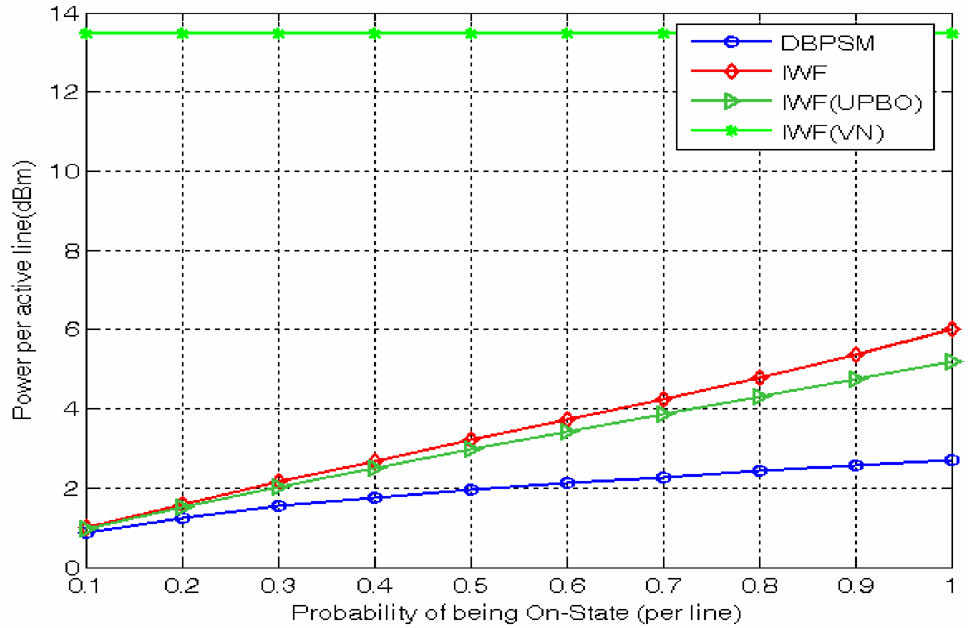

Figure 3(b) – Power reduction for 10 upstream VDSL users with dynamics of users

3. True Margins

An additional and compatible Level 2 strategy uses true margins [6] internal to the SMC in determining the consequent DSM-C interface parameters. True margins follow the original definition of margin in digital communications generally" In DSL, margin measures the amount by which a random (and uncontrollable) noise can be increased before unacceptable line performance is exhibited. This unacceptable performance is measured by an outage probability, P_{out} as in [6]. In Level 2 DSM, such a true margin should not include crosstalk noise that is controlled. However, the practice of DSL understandably uses Signal-to-Interference-and-Noise Ratios (SINR). Fortunately true margins can be translated into equivalent SINR margins so no change to existing DSL modem standards or practice is necessary. The SMC can internally effect these translations and thus can instruct different lines to have different margins, rather than to have a practice of a uniform margin (say 6 dB target margin and 16 dB maximum margin) on all lines.

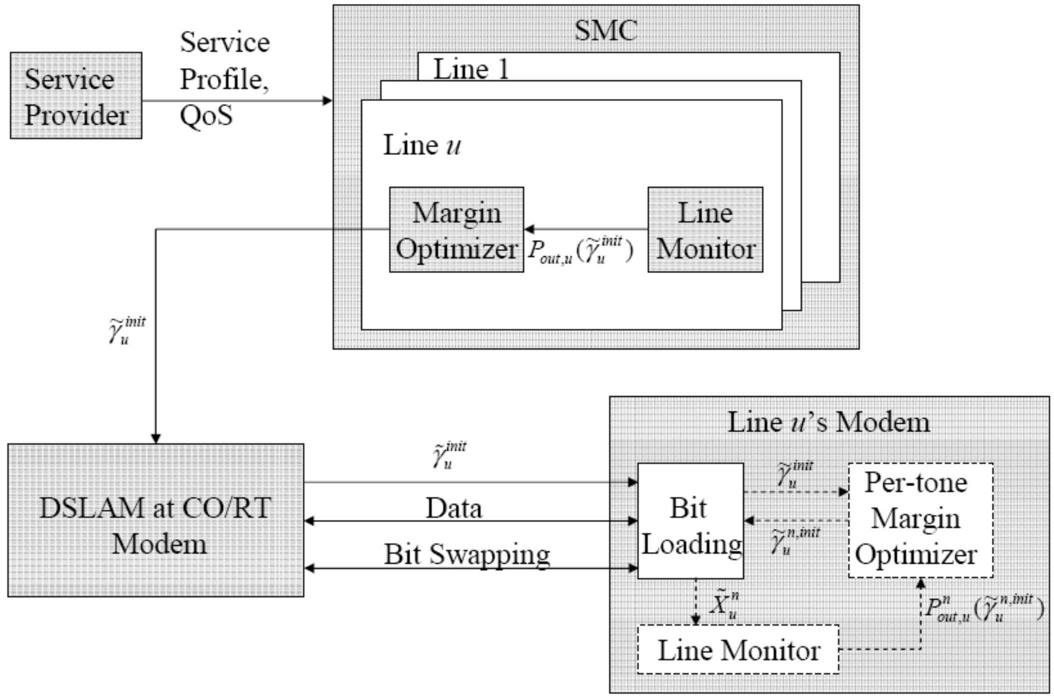

Figure 4 – Margin optimization

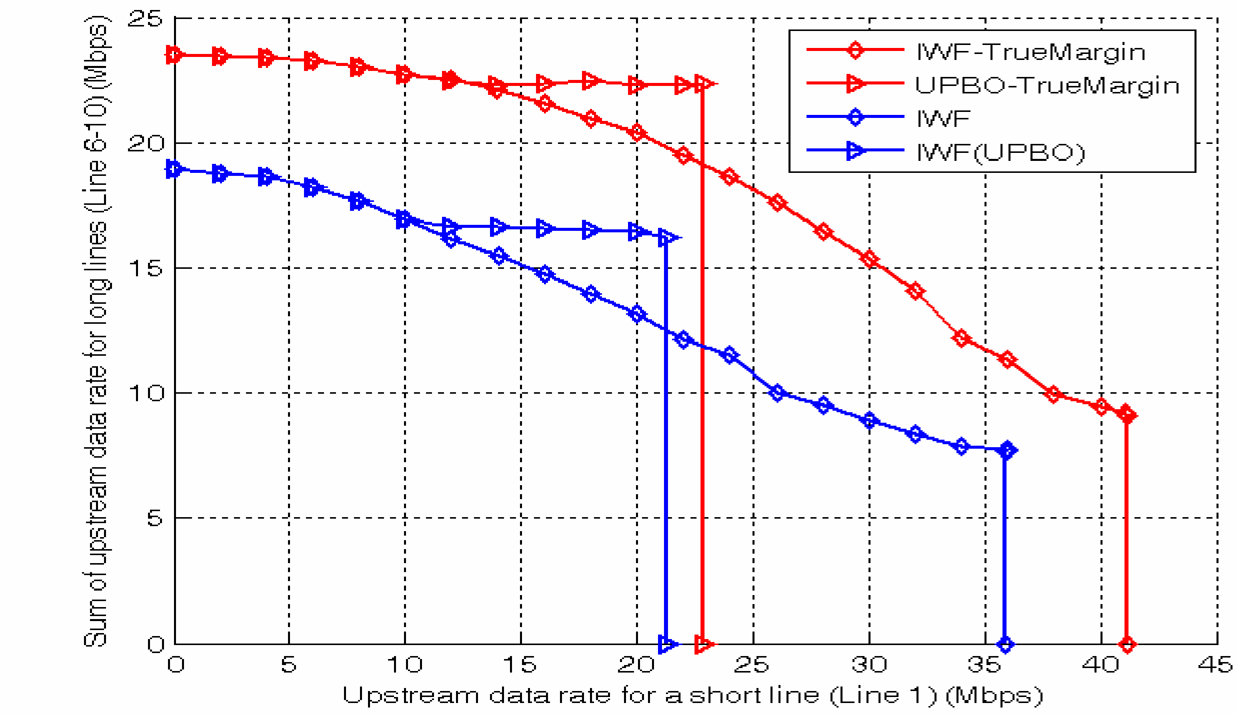

Figure 5– Rate regions enlargement by using 6 dB true margin for IWF (10 users)

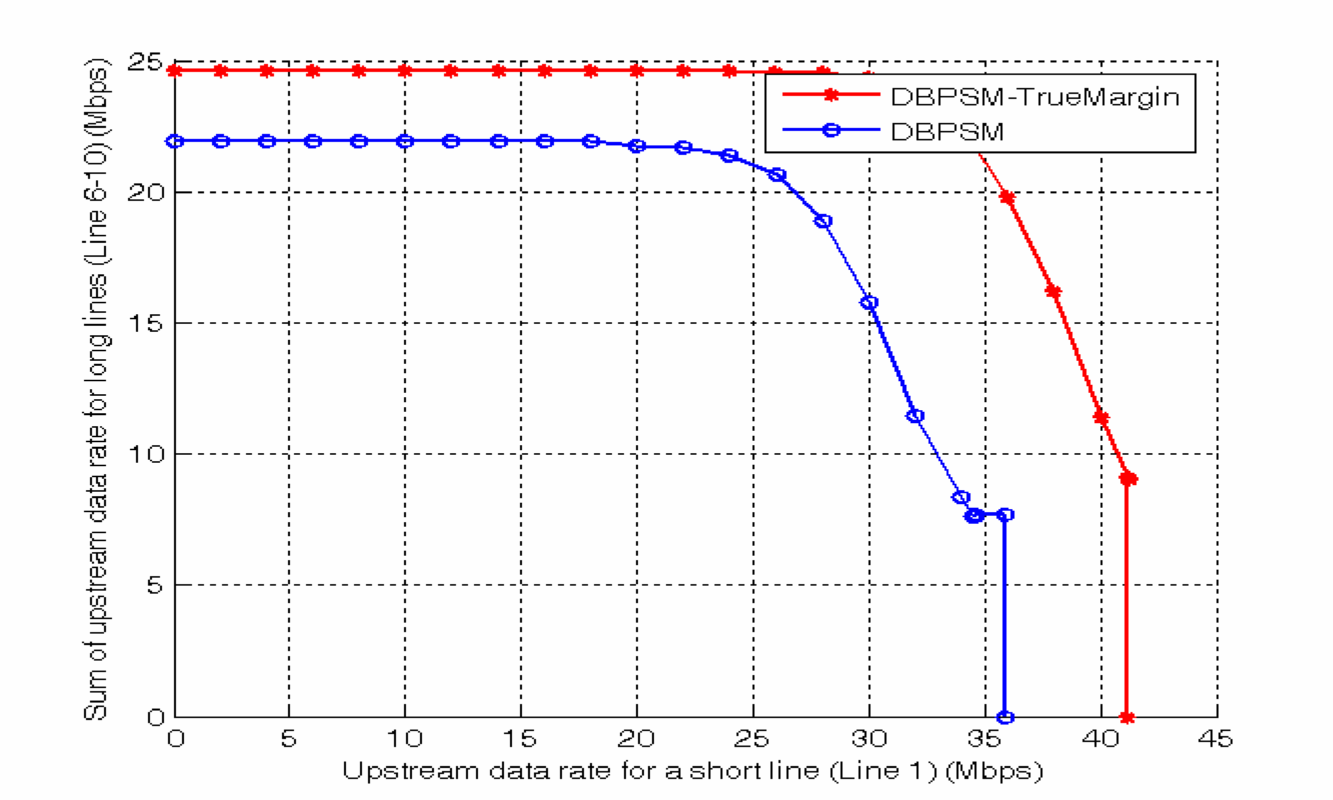

Figure 6– Rate regions enlargement by using 6 dB true margin for DBPSM (10 users)

4. Recommendations for Level 2 DSM

4.1 DSM 2 controls

DSM Level 2 support by an SMC and capability by DSL modems should be signaled through a simple handshake procedure that confirms the compatibility.

Additionally, for the strong-algorithm use, there should be a DSM-C bit indication that suggests use of the strong algorithm by the DSL modems (separate for upstream and downstream). This bit when set should be accompanied by a cut-off frequency DSM-C parameter to assist the algorithms.

Finally, with or without strong-algorithm use, the max-SNR-margin mode should be cleanly defined as a control parameter to signal to Level 2 DSM-capable modems that the margin on all tones should satisfy the MAXSNRM parameter.

4.2 Reporting Accuracy

The accuracy of Hlog, Xlog, SNR, and QLN parameters (in addition to accurate margins in Level 1 DSM data reporting) limits the SMC's ability to assist best performance. Improving accuracy should be seen as an opportunity for increases in overall DSL deployment and not as a backward threat to existing equipment designs.

An important point here is that a typical SMC feature is to report on DSL rates and performance, specific to various regions and/or equipment types. Deviations in those reports between different types of equipment caused by inaccurate values in reported parameters might incorrectly inform a network operator regarding the capabilities of equipment. One purpose of standardization is of course to try to minimize such incorrect views by having multiple-vendors consistently support the standard management interfaces.

The exact sensitivities today are large, and it would benefit standards groups to understand such sensitivities so that proper accuracy tests can thus best ensure the full Level 2 DSM gains are attained.

A component of accuracy is also the consistency of the data itself. For instance SNR's reported at one point in time, while margins and data rates reported for another instant in time, and perhaps bit distributions for a 3rd point in time causes an uncertainty in operating conditions. The uncertainty can lead to the SMC being more conservative in terms of rate-region enlargement for equipment with inconsistent data reporting. The concept of SCPG's (self-consistent parameter groups) was introduced[9] to try to assist vendors in achieving best Level 2 DSM compatibility and also deserves serious consideration.

4.3 Control Accuracy

Presently, there is considerable variation in the programmed PSDMASK[n] parameter for SMC programming or even for quasi-static settings in, for instance, standardized power "back-off" schemes. Such inconsistency might be addressed by tests. Often PSDMASK[n] and even the simpler CARMASK[n] are not implemented at all, which means a simple test could identify such non-compliance.

The MAXSNRM should be also tested to see if actual measured margin exceeds the specified maximum, as mentioned in Section 4.1 and [1].

Also, if true margin translation is used, this depends upon correct implementation of target margins in DSL equipment. A sampling of several vendors' equipment shows gross errors in the actual measured margins for highly static well-understood laboratory tests. The misses in field use are even larger. Thus tests to ensure that reported/implemented margin in loading algorithms is correct would be of value.

3.4 Keep the Dynamics

As in [1], faster swapping speed would be of assistance to both strong and weak algorithms and should be investigated.

Possibly, even dynamic provisioning of the weak/strong algorithm selection and cut-off frequency might be considered (that is provisioning during showtime), although not absolutely necessary. A fast SMC might be able to use such dynamic capability to its advantage.

4. References

[1] J. Cioffi, W. Lee, S. Jagannathan, G. Ginis, and P. Silverman, "DSM Level 1: Polite Low-Power Single-DSL Stability," *ATIS Contribution NIPP-NAI 2007-158,*Vancouver, Canada, November 12, 2007

[2] J.Cioffi, W. Lee, S. Jagannathan, and G. Ginis, "The Inherent Simplicity of Distributed Band Preference," *ATIS Contribution NIPP-NAI 2007-129R2,*Miami, FL, September 24, 2007.

[3] 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.

[4] J. Cioffi and M. Mohseni, "Preference in Water-filling with DSM,'' *ANSI contribution T1E1.4-2003-321R1*, San Diego, CA, December 8, 2003.

[5] Dynamic Spectrum Management Technical Report, "ATIS Committee NIPP Pre-published document ATIS-PP-0600007," May 2007.

[6] S. Jagannathan, C. S. Hwang, and J.Cioffi, "Margin Optimization in DSM Levels 1 and 2," *ATIS Contribution NIPP-NAI 2007-135R1,*Miami, FL, September 24, 2007.

[7] P. Spruyt, "Decreasing the OpEz of Remote DSL," IEC Broadband World Forum, October 8, 2007, Berlin, Germany.

[8] C. Posthuma, "DSM Implementation and Benefits," DSL Forum Birds of a Feather Session, August 28, 2007, Nashville, TN.

[9] Mehdi Mohseni, George Ginis, *G.ploam, G.vdsl: Self-consistent Parameter Groups*, ITU-T SG 15/Q4 Temporary Document RJ-087R1, Redbank NJ, October 2007.