ATIS Committee NIPP-NAI
Working Group NAI (DSL Access) **NIPP-NAI-2006-087R1**
Savannah, GA: June 26, 2006

CONTRIBUTION

TITLE: **DSM is for Unbundled DSL**

SOURCE:

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

PROJECT: **T1E1.4, DSM**

ABSTRACT

This contribution reinforces the understanding that Dynamic Spectrum Management (DSM) is very much so intended for use in unbundled networks; that is those networks served by more than one co-located, and absolutely uncoordinated, service providers. All levels of DSM have significant performance improvements possible with unbundled networks (as well as with single-provider networks). Description of various unbundled uses and advantages of DSM at each of the 3 levels 1, 2, and 3 appear throughout. This contribution is for information only.

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

DSM is for Unbundled DSL

J. Cioffi and 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 M. Brady

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

ABSTRACT

This contribution reinforces the understanding that Dynamic Spectrum Management (DSM) is very much so intended for use in unbundled networks; that is those networks served by more than one co-located, and absolutely uncoordinated, service providers. All levels of DSM have significant performance improvements possible with unbundled networks (as well as with single-provider networks). Description of various unbundled uses and advantages of DSM at each of the 3 levels 1, 2, and 3 appear throughout. This contribution is for information only.

1. Introduction

Dynamic Spectrum Management (DSM) well addresses unbundled environments where multiple service providers may co-locate network DSL equipment without any coordination. Specifically line-parameter adaptation can be autonomous to each line, shared within groups of lines for a single service provider, or autonomous to independent uncoordinated groups of lines. Each service provider independently controls its own lines with its own spectrum maintenance center (SMC). At any level of DSM, prudent internal operation of the SMC by its service provider (independent of other service providers) leads to greater stability of all service providers' lines, resulting in trouble reduction and consequent range expansion. Good vendor compliance to DSM capabilities in the DSM report assists the SMC's ability to stabilize and improve the network[1]. The goal of this contribution is to further the understanding of DSM use in

[1] And conversely, poor vendor compliance (which unfortunately abounds today) or attempts to control locally or in unknown proprietary manner inhibits network stability.

unbundled networks and to allay any conjured fear, uncertainty, and doubt about its use in unbundled situations.

The 3 levels of increasing DSM capability in the currently in-ballot DSM Report [1] are sequentially reviewed in the 3 main sections of this document, each considering use of the associated level of DSM in an unbundled situation.

2. Level 1 Unbundled – Single-Line Stability and Footprint enlargement

Impulse and intermittent noises or effects (which can be other DSL's energizing or de-energizing) are of paramount consequence to DSL service stability. Customers experiencing such noise are more likely to generate various maintenance actions such as trouble calls or truck rolls, as well as to drop the service altogether (sometimes measured by the "churn rate"). Actual deployment qualification range is determined by such "impulse" noises and unfortunately almost never (yet) by the well-defined and implemented 1% crosstalk noise tests upon which all standards heretofore have focused. Location and separation of stable and unstable lines allows the stable lines to offer higher rates, improving average deployment rate/range footprints. The unstable lines can be remedied by several strategies in any individual service provider's SMC, separate from any other SMC. These are discussed in subsections to follow.

2.1 Forward Error Correction

The most basic remedy for impulse and intermittent noise, well-known and specifically installed in ADSL1 and all subsequent DSL standards, is the forward-error-correction (FEC) of byte-wide Reed Solomon codes. The corrective capability of such codes is today measured by the INP (impulse noise protection) and DELAY parameters of the various standards. The increase of the INP parameter (at any fixed delay) tends to increase the fraction of parity, thus allowing more byte errors to be corrected and stabilizing the line (perhaps reducing raw data rates but, in prudent SMC use, INCREASING throughput as well as stabilizing the customer experience). Increasing DELAY typically allows higher values of INP, but larger delay may not be tolerable for certain applications. Both INP and/or DELAY can be changed for each line by that line's SMC. No coordination between lines is necessary for such changes.

There are a number of advantages of using the standardized FEC for its original intent, mitigation of impulse noise. First, for a parity overhead increase, errors of any type can be corrected as long as the total number of byte errors is less than half (or all with erasures) the number of parity bytes. Some of the authors were directly responsible for the introduction of FEC in ADSL1 and the programmability of N (block length) down to very short values was indeed very much so intentional. These low N values were mandated by the standard so that a maximum parity of 16 bytes could represent a significant portion of the codeword. Subsequent vendor INP (and/or "S" = DMT-symbols/codeword parameter) definitions unfortunately re-interpreted the use of FEC so that it was essentially useful for dealing with only stationary Gaussian noise, and not for intermittent nor impulse effects. The FEC re-interpretation reduced the number of DMT symbols per burst-length-corrected to two or less and focused instead on total coding gain for only this stationary Gaussian noise[2]. This re-interpretation was belatedly realized as a significant

[2] This mistake manifests itself in very large block lengths like N=240 (with 16 parity bytes) for what are called "interleave" or "medium" settings. The reduction to 3 delay settings and essentially minimal values of INP is not as

mistake and has been somewhat corrected in later ITU-T Recommendations where higher values for INP are supported as options. The DSM report draft suggests that an INP as large as 16 may be supported for all modes of operation. No other impulse mitigation method is as effective in improving DSL stability, but some alternatives are discussed in the subsequent subsections of this section. These alternatives are a second choice when reasonable INP values are not available, either because of standards interpretation or vendor fault/blocking of the capability. The appropriate use of FEC is effective on any line bundled or unbundled..

FEC with INP has a number of additional advantages from an unbundled perspective: The foremost advantage is that any PSD settings or MAXSNRM settings are respected with FEC use. This respect helps ensure that politeness is maintained between lines. Thus, there is an overall reduction in crosstalk that occurs, assisting all lines, even when impulse noise is present. Even when only one service provider -- among many -- uses politeness, overall crosstalk is reduced for all service providers. Thus, all service providers' customers benefit. We caution the reader that strict interpretation of crosstalk models for inter-line transfers is virtually assured to be incorrect in any specific situation. Therefore any reduction in overall noise should be viewed a statistical benefit to the set of all customers in terms of trouble tickets, truck rolls, churns or more generally operational efficiency.

FEC also applies whether or not the intermittent noise is ``periodic.'' Various equipment-qualification test noises may need to be periodic for ease of test implementation. These noises provide some indication of FEC strength and effectives as well as functionality. However, DSL equipment designed to exploit the test-lab noise periodicity, rather than more robustly designed to mitigate the nonstationary fluctuating noise, has dubious value.

2.2 Polite Adaptive Profiling

When vendors have blocked proper INP settings, as is often unfortunately the case, or when standardized INP ranges are interpreted to limit such values, polite profiling is an alternative that can be effective. Polite profiling is however less desirable than an increase in INP. Polite profiling can also be combined with higher INP values to mitigate further impulse noise. Polite adaptive profiling essentially narrows the range of rate adaptation for a modem based on a history of line observation. Target margins are maintained at low (typically 6 dB) levels, but the maximum and minimum rates are selected to keep the customer operating with reduced (ideally no) retrains while avoiding wide user-data-rate swings, the latter of which can increase churn rates for a service as well as increase trouble tickets. While the service is still rate adaptive, the range of rate adaptation is adaptively tuned to each line at a fixed low margin. By setting the MAXSNRM to a reasonably low value, then politeness is enforced[3]. Such polite adaptive profiling improves the stability of an SMC's controlled lines and still maintains a low crosstalk into other services. It is thus beneficial in unbundled environments.

2.3 Poor SMC practice that can lead to instability

---

effective as was necessary or intended for impulse noise in the original addition of FEC to ADSL1. Long block length is not very helpful for intermittent impulse noises.

[3] Some DSL equipment may ignore or incorrectly implement the MAXSNRM parameter and thus does not conform to standards, perhaps reducing (but not eliminating) the gains attainable by polite profiling, despite the equipment bug.

There are some SMC practices that are ill-advised for that SMC's lines and also are ill-advised in terms of these practices' effects on other lines. Two poor practices are:

(1) increasing Target SNR margins
(2) imposition of fake large noises on DSL loading.

Large target margins imposed on a line cause it to use more power than necessary to achieve the desired data rate. Crosstalk is thus increased. Very often, intermittent or impulse noises are so large that they overwhelm any feasibly large margin so the line is not stabilized anyway and then simply becomes a crosstalking beacon into other lines while still exhibiting instability itself. Most service providers who do increase target margins are aware of the detrimental effects on other lines (both their own and those that may be served by another service provider), but feel they have no choice because the INP settings they would have preferred are blocked by vendor equipment. Polite adaptive profiles is a preferred alternative, but can also be frustrated by artificial limitations on the number of profiles that certain types of DSLAM equipment will accept.

Rather than provide proper INP or allow sufficient numbers of profiles, some vendors have proposed forcing modems to train as if they see a large "phantom" noise that is actually not present or may occur infrequently. Thus, instead of using INP as a fix (perhaps with polite profiling), the modems are forced to artificially enlarge transmited signal spectra, and thus consequently to radiate larger crosstalk. Enlarged crosstalk then leads to other lines in turn imposing even larger phantom noises on their lines. The whole situation is destructive in any environment, however in an unbundled environment it can be especially destructive to the competing carriers' services, therefore these authors recommend it be avoided. However, it is an example of poor SMC practice that could lead to unbundling problems. We cite such effects in this section so that the reader understands that good SMC design is very much compatible and indeed overall positive for unbundled operation.

3. Level 2 Unbundled – Spectrum Balancing and Politeness

Level 2 DSM may also be practiced in unbundled environments. While the DSM Report's [1, Annex A] theoretical bound known as Optimal Spectral Balancing is derived only for situations of a single service provider, this derivational facility does not prohibit the good performance from being achieved or approached in unbundled environments with various approximations as outlined in [3],[4].

3.1 Band Preference and Distributed Implementations

A practical way in the DSM Report to use Level 2 DSM in unbundled networks is for the SMC to set the band preference bit on some line and consequently for DSM Level-2-capable equipment to observe the rule that margins at all frequencies must be less than the MAXSNRM parameter. Such observation prevents any system from "overloading" any frequency band with high margins simply because the supplied PSDMASK is lower in another band, a good overall practice for any service provider. Further agreement or understanding of supplied PSDMASKs, margins (frequency-dependent or not), and other parameters (for instance frequency-dependent bit-caps are possible on some systems and of interest to some service providers and might be wise to standardize so all could benefit) can then be interpreted by any line in connection with ONLY its own SMC.

However, even without any change in the regulatory environment, the more focused use of MAXSNRM by use of Band Preference enhances the overall politeness of the lines in either the unbundled or bundled environment. Again it is observed that appropriate use of DSM techniques enhances the environment for all users of a binder, even if only some of the lines are controlled by an SMC.

4. Level 3 Unbundled – Cancellation of the Other Provider's Xtalk

Level 3 DSM is often called vectoring. Figure 1 illustrates the use of vectoring in unbundled networks. It is not necessarily MIMO (multiple-input multiple-output) nor is vectoring the same as bonding. Vectoring is the co-generation of a downstream vector of transmitted signals for a subset of a single service provider's DMT lines (typically VDSL with digital duplexing is required). Vectoring is also the co-reception of an upstream vector of received signals for those same lines[4]. Vectored systems cancel FEXT and are typically used with frequency-division multiplexing (as in VDSL). At any DMT frequency, noise from outside the group of vectored lines – THAT IS IN PARTICULAR ANOTHER **UNBUNDLED** SERVICE PROVIDER's CROSSTALK – can also be removed in the upstream direction. The condition for such removal is that the vectored lines all see correlated noise (or more practically that the number of sources of that other noise at any frequency is smaller than the number of vectored lines). When this later condition is met, vectoring helps the service provider that uses this vectoring with a large data rate gain. Simultaneously, the other service providers see certainly no more noise and may see less noise (because the first service provider needs less power since they see little noise and observe a MAXSNRM simultaneously). The other service providers, of course, can also use vectoring, independently as their choice. Near/far effects for different service providers are largest in the upstream direction for VDSL because of the different positions of customers, and thus vectoring very much assists unbundled situations.

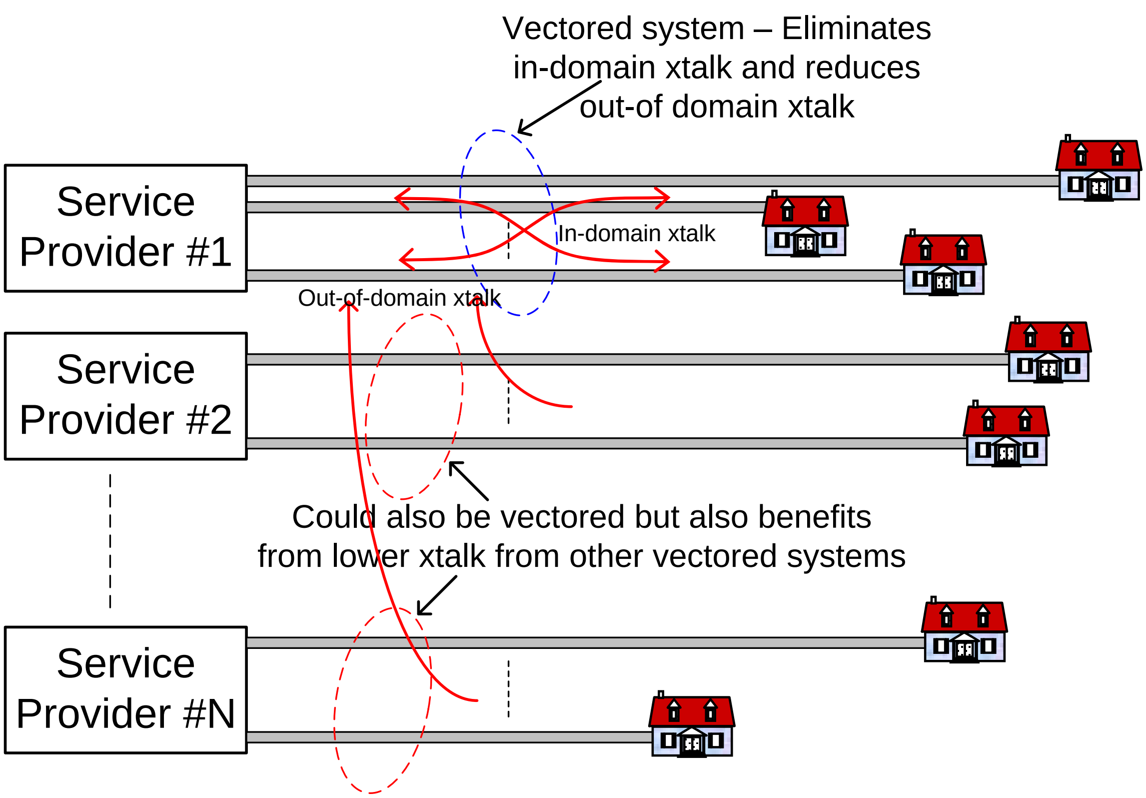

Figure 1: Illustration of the use of vectoring in unbundled networks

Downstream "out-of-domain" noise (see [2]) cannot be cancelled with vectoring, but near/far effects are much less in downstream VDSL. If there are such near/far effects, the downstream situation is still also better than no vectoring because vectored lines can use lower power to meet the same performance/data-rate requirements and thus overall binder crosstalk is reduced, just as in Levels 1 and 2.

[4] Bonding can be used with or with or without vectoring, and vice-versa, see [2].

As an example, the following simulations for VDSL2 show the benefit of reduced power in vectored systems for a given data rate and SNR margin requirement, even in unbundled networks. The reduced power on each line of the vectored system translates into lower crosstalk within the network thereby improving the performance of all the lines. In the simulations, 25 lines are considered in an unbundled network and the number of vectored lines is varied. If M lines are vectored, then the remaining (25-M) lines (which may correspond to other vectored or non-vectored systems in the network) are considered as out-of-domain crosstalk or noise. Therefore M=1 corresponds to using no vectoring and M=25 corresponds to Full-vectoring.

For all plots in this contribution, the following parameters are used:
gap-coding gain+margin = 9.5-4.5+6=12dB
bit cap = 15
VDSL2 12 MHz, North American bandplan is used
US0 is not used
Line type: 24 AWG
All lines are of the same loop length

The out-of-domain crosstalk observed at each of the lines in the vectored system will be correlated and this correlation can be exploited in the upstream direction by the lines that are vectored (due to receive coordination). This noise correlation among the vectored lines is captured using the correlation coefficient ρ which is varied from 0 to 0.99. Figures 2 and 3 show the power reduction per vectored line that arises from vectoring in the upstream VDSL2 direction. The target data rate was set to be 3Mbps for both the plots. Figure 2 corresponds to a loop length of 2Kft for all lines whereas Figure 3 corresponds to 3.8Kft.

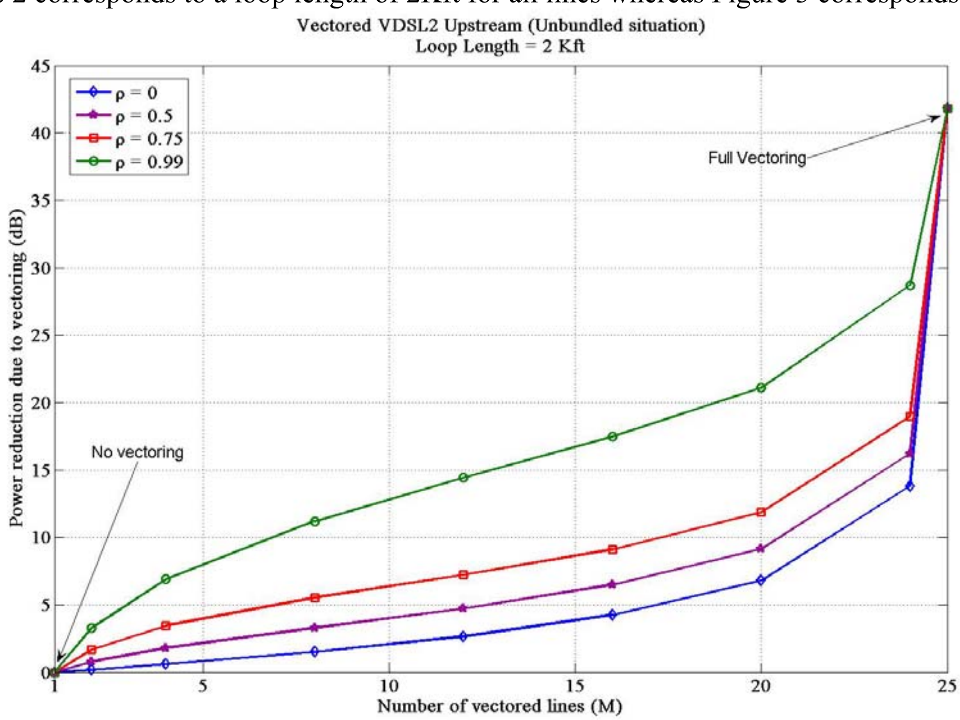

Figure 2: Power reduction per vectored line in upstream VDSL2 with vectoring. Loop length = 2000 feet, Data rate = 3Mbps, Margin = 6dB.

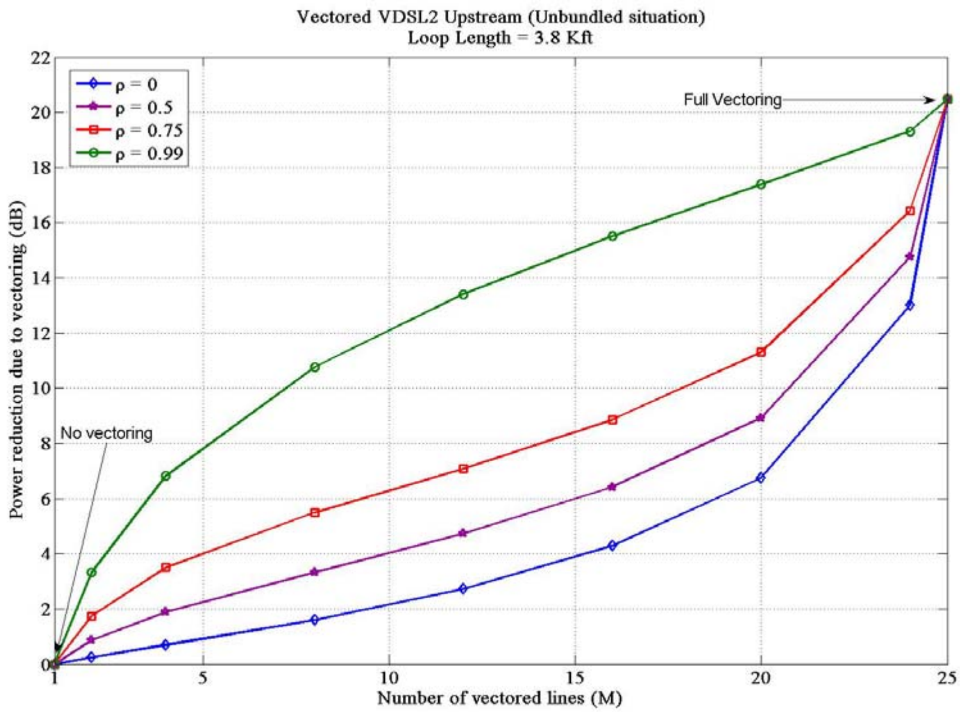

Figure 3: Power reduction per vectored line in upstream VDSL2 with vectoring. Loop length = 3800 feet, Data rate = 3 Mbps, Margin = 6 dB.

Figures 2 and 3 show that vectoring leads to significant power reduction (and hence lower crosstalk) even if all the lines are not vectored (i.e. $1<M<25$) and is hence very useful in unbundled networks. For $\rho = 0$, the coordinated receiver for the vectored system can not (or does not) exploit the correlation in the out-of-domain crosstalk. Even in this extreme, the elimination of the in-domain crosstalk (i.e. the cross talk between the vectored lines), leads to good power reductions that help the entire network (compared to the non-vectored M=1 case). As the noise correlation ρ increases (and is exploited at the receiver), the benefit of vectoring increases in the upstream. Even if half of the lines in the network are vectored (M = 12), the power can be reduced by close to 15 dB for ρ=0.99 and a loop length of 2 Kft (correspondingly, an improvement of 13.5 dB for a loop length of 3.8 Kft). Therefore, the benefits of full vectoring (M=25) can be quickly approached in an unbundled network with an increase in the number of vectored lines. The consequent power reduction significantly reduces crosstalk to other lines in the network.

Figure 4 shows the power reduction per vectored line that can be obtained in the downstream. Unlike the upstream, the correlation in the out-of-domain crosstalk experienced by the different lines of the vectored system cannot be exploited since there is no receive coordination. However, the in-domain crosstalk can still be eliminated using coordinated transmission. This leads to useful power gains and hence reduced crosstalk in the network. The target data rate for the plot was 16 Mbps with an SNR margin of 6 dB.

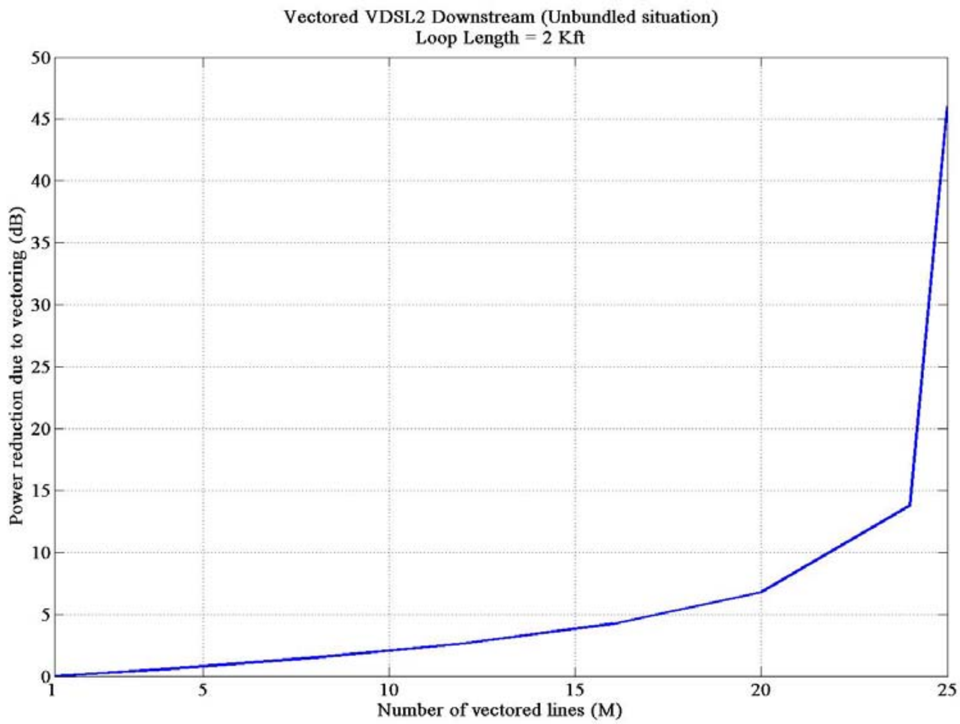

Figure 4: Power reduction per vectored line in the VDSL2 downstream due to vectoring. Loop length = 2000 feet, Data rate = 16 Mbps, Margin = 6 dB.

A companion contribution [5] investigates symmetric transmission with vectoring.

5. Conclusions

This information-only contribution explains how all 3 levels of DSM have good benefits and use in unbundled environments, where no sharing nor coordination of any type, nor even any regulation, is necessary to effect significant gains. Thus, the intent is to reduce confusion that might have previously associated DSM with single-service-provider networks.

6. References

[1] DSM (Draft) Technical Report, *ATIS NIPP-NAI Contribution 2006-028R2,* June 2006, Savannah, GA.

[2] M. Brady, J. Cioffi, G. Ginis, and P. Silverman, "Overview of Technologies for Vectored VDSL," ANSI Contribution *NIPP-NAI-2006-086,* June 26, 2006, Savannah, GA.

[3] J.Cioffi, M. Brady, W. Lee, G. Ginis, and P. Silverman, "An Explanation of Band Preference Use," ANSI Contribution *NIPP-NAI-2006-036,* March 20, 2006, Las Vegas, NV.

[4] M. Brady, W. Lee, J.Cioffi, G. Ginis, and P. Silverman, "Examples of Band Preference Gains in VDSL," ANSI Contribution *NIPP-NAI-2006-037,* March 20, 2006, Las Vegas, NV.

[5] J. Cioffi, S. Jagannathan, G. Ginis, and M. Brady, "DSM Upstream Data Rate Revival: the Restoration of Symmetry," ANSI Contribution *NIPP-NAI-2006-092,* June 26, 2006, Savannah, GA.