Greener Copper with Dynamic Spectrum Management

(Invited Paper)

John Cioffi[1,2], Sumanth Jagannathan[1], W. Lee[1], H. Zou[1], A. Chowdhery[1],
W. Rhee[2], G. Ginis[2], and P. Silverman[2]

[1] Stanford University, Stanford, CA 94305-9515 USA
`cioffi@stanford.edu`
[2] ASSIA, Inc., Redwood City, CA 94065-1317 USA
`gginis@assia-inc.com`

Abstract. This paper investigates the benefits of Dynamic Spectrum Management (DSM) in terms of reducing the power consumption and improving the data rates in digital-subscriber-line (DSL) networks. The proposed techniques at the three different DSM levels simultaneously also provide a significant improvement in the stability of DSLs. The proposed DSM methods are compared with other non-DSM solutions, which sacrifice power and/or data-rate in order to improve the stability, while also harming other DSLs through impolite power usage. Various examples are presented showing that the proposed DSM methods can avoid such unnecessary impoliteness, and that stability and politeness can be simultaneously achieved in a DSL network that is efficiently managed using DSM.

Keywords: DSM, spectrum balancing, power savings.

1 Introduction

Dynamic Spectrum Management (DSM) methods stabilize and increase Digital Subscriber Line (DSL) data rates through service-provider surveillance and management of various DSL physical-layer parameters. Correct use of DSM methods encourages “politeness” or power-control of DSL transmitters so that they radiate less crosstalk into other DSLs. This process of statistical noise reduction also lowers DSL transmit power. Since the power consumed by DSL modems is often dominated by circuits used for transmitting power, there is a consequent large secondary DSM benefit of reduced power consumption by DSL systems. This reduction in consumed power can be beneficial to DSL networks and is consistent with growing worldwide pressure for telecommunication networks to reduce their power consumption.

DSL management of many service providers often leads to a power increase because service providers may simply (but in retrospect, incorrectly) attempt to increase the margin on any observed unstable DSL. The margin is a controlled parameter dating to the earliest DSL standards' specifications. Specifically, the

C. Wang (Ed.): AccessNets, LNICST 6, pp. 62–77, 2009.

margin specifies how much larger (in dB) the line noise may rise before DSL performance (that is bit error rate) becomes unacceptable. A management system's increase of margin on an unstable DSL line means either or both of a power increase and a rate reduction, and is largely ineffective against the enormous intermittent noises that often lead to field-DSL instability[1]. This power increase may also induce a crosstalk increase, which is generally considered as impolite behavior. Section 2 discusses margin and the fallacy of its increase for unstable lines, and instead shows that DSM Level 1 methods of politeness, often called "tiered rate adaptation" (TRA), instead do stabilize the DSL and reduce average power consumption of DSLs simultaneously - DSM's first example of greening of the copper. DSM's power reductions can be very large when TRA is used correctly compared to any impolite management as Section 2 shows.

Section 3 advances to a frequency-dependent form of politeness known as spectrum balancing or DSM Level 2. A particular version of this politeness, known as "band-preference" in North American DSM standards documents [1], is shown to effect additional greening of the copper (that is, additional power savings), while caution is urged with respect to other impolite methods sometimes called "virtual noise," which are here shown to substantially increase power consumption if used. Thus, Section 3 illustrates a correct, polite, and low-power way of implementing DSM, further greening the copper.

Section 4 then explores the limits of DSM Level 3 (the highest DSM level in standards), also known as "vectoring," for further reduction of transmit power, which should ultimately lead to the smallest power consumption and thus, the best greening of the copper.

2 DSM Level 1: Dynamics and Politeness

DSL modems are dynamic by standard. The long-mandated DSL-standardized "bit-swapping" procedure [2] allows the transmitter and receiver to react simultaneously to changes in the line's noise by redistributing energy from regions of relatively high noise increase to those of noise decrease (or of relatively low noise increase). This capability to react to noise changes helps maintain an optimized information (or bit) distribution on the DSL through a hand-shaking procedure, where the DSL's transmitter and receiver use a control channel to reallocate information synchronously. This "bit-swapping" is the key difference between the multi-carrier DSL systems and the Orthogonal Frequency Division Multiplexing (OFDM) systems used successfully in wireless transmission, where the information distribution is not changed and there is no bit-swapping. DSL systems thus call the highly adaptive transmission format "Discrete Multi-Tone" or DMT to distinguish it from the non-swapping wireless OFDM systems. The dynamic bit-swapping capability improves line stability and reduces the number of retrains or restarts of DSL modems when used with effective DSM methods as in this section. While standards mandate bit-swapping, some DSLs are nevertheless non-compliant and

---

[1] Simply stated, the intermittent noise is so large that no practical value of margin is large enough to accommodate this transient noise.

require retraining upon significant noise changes. DSM can also help such systems through TRA, although not as much as bit-swapping systems.

This section begins with a discussion of DSL stability and intermittent noises, and then progresses to the appropriate polite and power-saving solutions that use DSM Level 1.

2.1 The DSL Environment and Stability

Instabilities limit today's DSL deployment range[2]. Such deployment range is typically well below the DSL range exhibited in carefully crafted and standardized lab/interoperability tests, which prior to the advent of DSM were thought to have tested "worst-case" conditions. Thus, range of field-deployed DSL is often much less on average than found in such "worst-case" testing unless DSM is correctly used. Instability is typically measured by 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. These large intermittent noises often easily exceed any DSL management system's "margin" against noise, thus often rendering a management system's margin increase ineffective.

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],[3]-[10]. The main causes are large intermittent noises that occur when the DSL is in operation, which are so large that bursts of errors or even modem retrains occur.

Video quality is strongly reduced by DSL instability, but data and voice-over-internet services also exhibit strong correlation between stability and customer satisfaction. Level 1 DSM has several effective mechanisms to stabilize DSLs and thereby to improve deployment range.

Level 1 DSM[3] [1] focuses largely upon the management of a single line. A Level 1 DSM system-controller collects DSL performance data from the DSL equipment's maintenance interface. The Level 1 Spectrum Management Center (SMC) then assesses DSL stability. The Level 1 DSM system may consequently re-profile[4] the DSL through the maintenance interface. Each DSL service provider can have their own DSM system - there is no central controller for

---

[2] "Range" means the curve of data rate versus line length. DSLs typically offer lower data rates on longer lines; however, DSM methods can statistically improve the average "range," thus raising the curve of data rate vs. length.

[3] Level 1 DSM is sometimes also called "Dynamic Line Management" or DLM, although not in standards.

[4] "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."

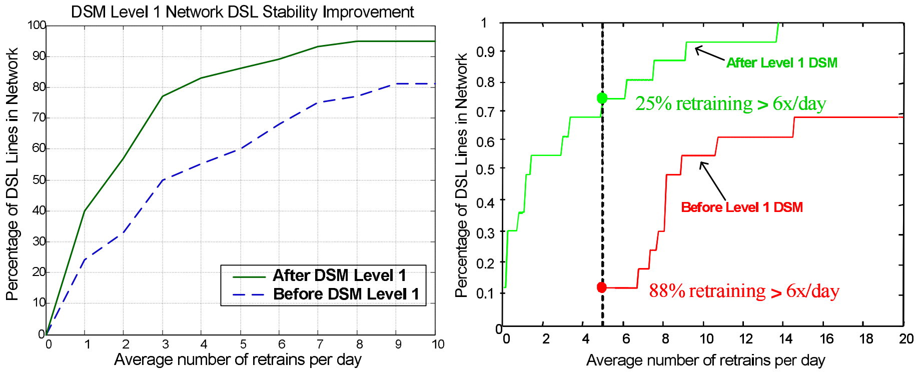

(a) (b)

Fig. 1. Stability improvement with DSM Level 1 (a) bit-swapping DSLs (b) without bit-swapping

all service providers in DSM. However, spectrum management standards may impose some guidelines for polite management in various situations.

Figures 1(a) and 1(b) reflect the average stability experience of field data taken from several DSL service providers around the world on bit-swapping and non-bit-swapping equipment, respectively. Figure 1(a) illustrates the percentage of DSL customers on a typical network that have fewer retrains per day than the abscissa value (horizontal axis). The figure illustrates DSL stability before and after the use of DSM Level 1 management known as TRA (see Section 2.3). If 5 retrains per day is the threshold for unacceptable stability,the stable-DSL-customer improvement is roughly 25% in Fig. 1(a)). Experience shows that such a stability improvement typically extends actual deployment range by at least 20% (which could mean 30-50% more customers possible for a higher-speed DSL service).

Figure 1(b) illustrates the stability improvement without bit-swapping. All improvements here come through DSM Level 1 management that tries to leave sufficient tolerance in the DSL's programmed parameters to improve stability. A large stability improvement is possible with Level 1 DSM in Fig. 1(b) for non-swapping modems, but the total number of stable customers in Fig. 1(b) with no swapping is understandably smaller than with swapping in Fig. 1(a). For instance, at 5 retrains/day, the improvement in stability is almost 60% of the customers in the non-swapping composite networks of Fig. 1(b), increasing the number of stable customers to 75%. The bit-swapping system can increase this same point to almost 90% in Fig. 1(a). With or without swapping, the number of stable customers is not 100%, even after DSM Level 1 correction - but the customer-satisfaction gains are very large, rendering proper DSM-use a major initiative by most DSL service providers today. However, there are green and good methods to achieve such gains, and then some not-so-green methods to effect a smaller gain in stability. The next subsection reviews margin and its use/abuse.

2.2 Margins as Control Parameters

The DSL management parameter called "margin" measures the amount by which noise can increase before the DSL link performance is unacceptably degraded. Margin thus helps guard against unforeseen noise increases. Most DSL systems today use a default margin value of 6dB, which means the noise can increase in power by a factor of 4 before the DSL starts making appreciable errors. Management of an unstable line thus might increase margin, a process here called "automatic margin adaptation" or AMA. Typically, retraining the DSL with higher margin increases power and/or reduces data rate. It is impolite because the DSL will use maximum power levels if the margin is set sufficiently high. However, margin levels are often not sufficiently high to offset large transient noises or intermittent/impulse noise. Furthermore, such large noise only occurs for a fraction of the time. Thus, increasing the margin simply increases the power (and/or reduces the rate) all the time, while typically having little or no effect on a very intermittent noise. AMA is thus not power efficient. DSL modems have long used 3 margin parameters:

1. Target SNR margin (TSNRM)
2. Minimum SNR margin (MINSNRM)
3. Maximum SNR margin (MAXSNRM).

The target margin is used during training of the modem and is the margin used to counter a future noise-increase as stated above, typically 6 dB. The DSL modems continue to measure and update the margin of the link as the noise (and thus margin) changes. MINSNRM is a margin level (typically 0 dB) below which the DSL link will retrain. MAXSNRM is important and is the maximum margin allowed in bit-swapping. If the noise reduces, bit-swapping reduces power to keep the consequent margin below the MAXSNRM in standard-compliant modems. Setting the MAXSNRM adaptively is important in politeness, reducing transmitted power, emitted crosstalk, and also consumed power when intermittent noise is not present. Thus, wise Level 1 DSM might better set the MAXSNRM low than set the TSNRM high. The next subsection explores this effect with TRA.

2.3 Tiered Rate Adaptation (TRA)

TRA maintains a low TSNRM (typically 6 dB) even on unstable lines. The MAXSNRM is also kept relatively low, creating a low used power and a low emitted crosstalk when intermittent noise is not present. TRA tiers the allowed data-rate range with several (typically overlapping) ranges within which the modem is allowed to train as shown in Fig. 2(a).

Instead of increasing the TSNRM, TRA instead caps the data rate at a level for which instability has acceptably low probability of occurrence. It is thus polite because power is maintained at a low level when noise is normal (not increased). A minimum rate is also imposed (below which the DSL will retrain) such that a retrain during a very unlikely (but possible) large noise event will not cause

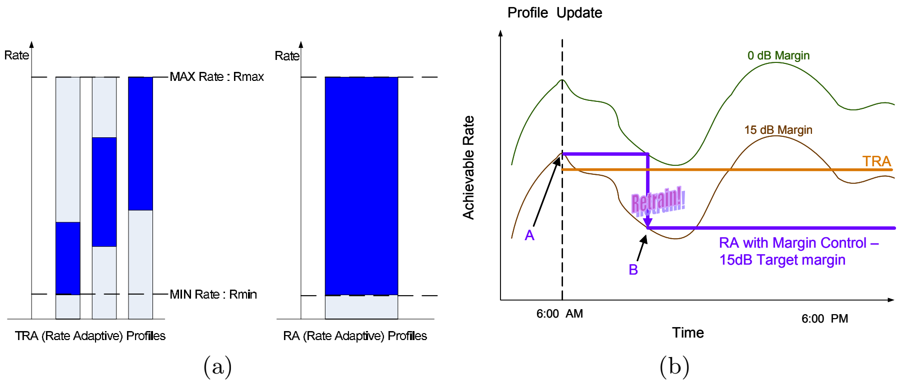

Fig. 2. Illustrations of (a) the tiered rate-ranges of TRA as opposed to the usual RA profile with a single, wide rate-range (b) The "stuck-at" low-rate problem of AMA (blue curve with retrain) and aversion of the problem by TRA

the modem to stay at a very low data rate[5]. AMA systems will retrain anyway upon a large such noise event, thus training at an even lower data rate because of the large margin, leading to a very poor eternally maintained low data rate. Figure 2(b) illustrates the "stuck-at-low-rate" problem of AMA.

Section 3 proceeds to show that there is a rate/reach loss, along with an increased power-usage for a target rate, from AMA because of the additional crosstalk it creates.

3 DSM Level 2: Frequency-Dependent Politeness

Level 2 DSM goes beyond TRA stabilization and considers the frequency-dependent effect that short lines' crosstalk may have on the longer victim lines. The first two subsections of this section look at spectrum balancing and its practical implementation using band-preference algorithms. The last subsection looks at a frequency-dependent extension of the non-DSM AMA known as "virtual noise," a pre-programmed large noise that the DSL is told to pretend to be always present, stripping it of dynamics and forcing a maximum power transmission, crosstalk, and power consumption. Several examples are provided to show the power savings of polite Level 2 DSM with respect to virtual-noise use.

3.1 Spectrum Balancing Concept

Spectrum balancing theory projects bounds on the best trade-offs between the data rates of different mutually crosstalking DSL customers. Because of crosstalk, those rates are not independent, so a region of trade-off can be plotted as shown in Fig. 3. Optimal Spectrum Balancing (OSB) [12] (with considerable complexity) computes this largest region. Gains in data rate are sometimes large when there are long-line DSLs as victims in the same binder with short-line users.

---

[5] Unless the customer intentionally resets the modem by power-cycling it.

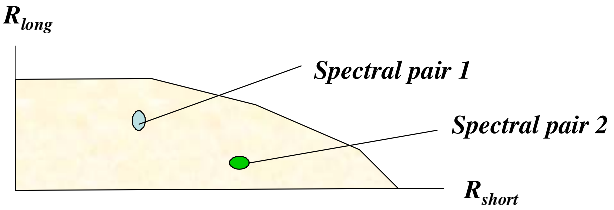

Fig. 3. A rate region for two users with asymmetric line length

The rate region contains all the possible pairs (generally U-tuples, where U is the number of users) that can be achieved by balancing the spectra of the different users. The data rate of one DSL customer can often depend upon the practices of another when the crosstalk between their lines is strong.

3.2 Band-Preference

Band-preference observes that shorter-line users (or more generally those with better channels) may often choose several different energy allocations to reduce crosstalk into other users. Often a short line's use of higher frequencies, largely above the lower frequency band that long lines can only typically use, can lead to much less crosstalk to long lines. These higher frequencies on one line may require somewhat more energy for transmission of a given data rate on that line, but through crosstalk reduction, the longer line uses much less power. The general trade-off is often very good for overall DSL data rates and power consumption. The practice of using only the higher frequencies above some management-specified cut-off frequency is known in DSM Level 2 as band-preference; that is, a preference applied to the use of higher frequencies.

This simple band-preference concept is easily managed in a distributed manner by letting each service provider's management system (SMC) decide if a line is strong (that is, will use band-preference on higher frequencies) or weak (the line needs to use all frequencies available to achieve its target minimum data rate). The strong lines then use a more polite form of bit-swapping that is described elsewhere [13][14][15]. The weak lines use normal swapping. Each line then does its best (still also retaining TRA) to minimize its own power consumption while also meeting the politeness and data-rate constraints. Figure 4 provides a representative example using band-preference, where this highly distributed (and thus usable in unbundled or bundled environments) method performs very close to the actual bounds of a theoretical optimum spectrum balanced system (that would require central control of all lines and their swapping) [15][16]. The proximity is so good that band-preference essentially allows optimal performance with very practical implementation by each DSL circuit of their own bit-swapping procedures.

3.3 Zap Your Neighbor First: Virtual Noise Power Increases

A non-DSM future proposal is Virtual Noise (VN) [17]. Management by VN mandates a DSL circuit to pretend that a pre-programmed large noise is always

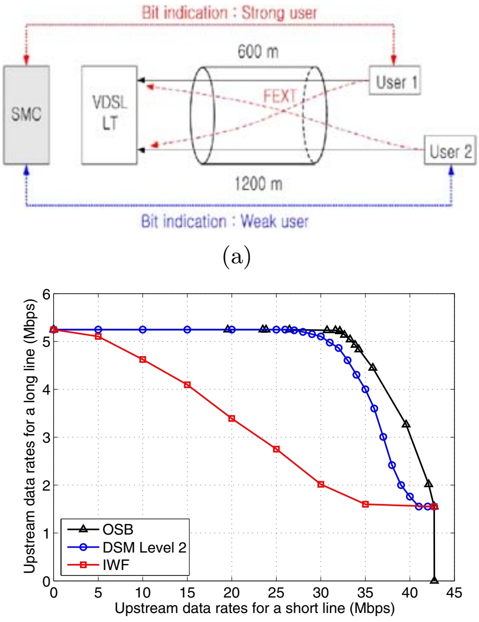

(a)

(b)

Fig. 4. (a)Upstream VDSL band-preference example. (b) Rate regions corresponding to Fig. 4(a) for the DSM Level 2, OSB, and iterative water-filling.

present as a threat to it. By pretending the noise is always present, and if the noise is the worst that can occur, then the DSL is always ready for this worst case. Effectively bit-swapping is disabled (unless the actual noise exceeds the programmed virtual-noise level). VN is similar to AMA, but is frequency-selective so that AMA's all-frequency abusive use of spectrum is narrowed to some frequencies in VN. Both VN and AMA essentially keep transmit power at the highest possible level. VN can also be more impolite than AMA since the amount of effective margin increase caused by the pretended programmed virtual noise can be larger than an AMA system's highest margin. Thus, the most abusive of AMA and VN actually depends on the management practice, but both attempt to address instability by reducing data rate and increasing power (and thus, crosstalk) without allowing dynamics of modems. Typically, VN proponents cite crosstalk transients as the reason for VN use. Crosstalk transients are caused by transition of a DSL from a dormant/off state to an "on" state. However, TRA is here shown to be a more appropriate remedy than these static and worst-case methods.

The non-DSM virtual noise presumes the use of a central control system for each service provider to determine and pre-program a large fake noise into each DSL's modems. A VN quandary thus occurs in an environment where each service provider subsequently sets their VN increasingly higher to protect against their competitors' most recent increase in crosstalk, leading each in turn to send at higher power levels. Thus, each user creates more crosstalk into their neighbors, leading to a loss in the data rates of all lines, and essentially larger

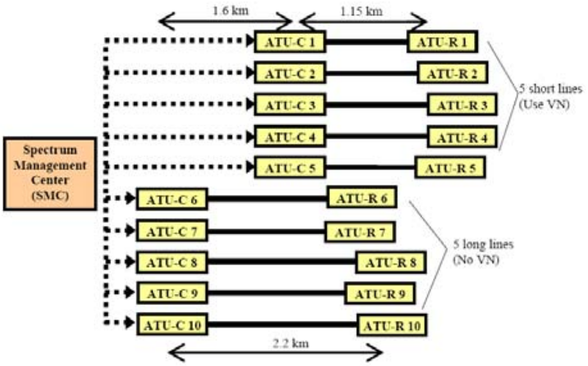

Fig. 5. Illustration of mixed-binder ADSL2+ lines. (RA = rate adaptive)

or maximum crosstalk between all lines. In an unbundled environment, there is no ability to set VN on different circuits with knowledge of the levels on other service providers' circuits. VN is thus a largely static operation of the DSL as if the large noise were always present (if the actual noise exceeds the virtual noise, then the DSL can adapt, but only in this case). All lines essentially prepare for the worst reaction of the others' possible VN (and thus power) settings.

The band-preference method of the previous subsection (a frequency-dependent DSM method) is usually implemented by classifying DSLs into two subsets, strong and weak. Power is set at maximum levels only when it must be so set because the largest noise is actually present. This subsection's simulations emulate all possible settings of virtual noise from none to all users using the largest-possible feared noise A level of 0 means that the virtual noise is less than or equal to the background noise and a level of 1 means all lines use a virtual noise that is equal to the sum of all others' crosstalk. Levels between 0 and 1 thus attempt to emulate various intermediate strategies. AMA is presumed to be equivalent to at least one of the possible choices for VN within the range considered, and is thus listed on the same curve as VN in the figures that follow. More detailed simulation settings can be found in [16].

A first example of mixed binder ADSL2+ is described in Fig. 5 and Table 1. Ten users are located roughly 1.15 km to 2-3 km from a central office or a remote terminal with a target rate of 10 Mbps for video service on those served from a service provider's fiber-fed RT (at 1.6 km) in the same binder. The lines emanating from the central office are all allowed to rate-adapt to the best service possible. One operator, serving lines 1-5, uses virtual noise. The other CO-based operator does not use VN. The 2nd service provider's lines are harmed by the VN use: Figure 6(a) shows the data rate loss of the 2nd service provider's lines with respect to TRA, and further, the larger loss with respect to band-preference.

Figure 6(b) shows the average power per line required to achieve [10, 10, 10, 10, 10, 5, 5, 5, 5, 5] (Mps), respectively. Figure 6(b) shows large power savings. More examples can be provided for a variety of situations and the losses can be

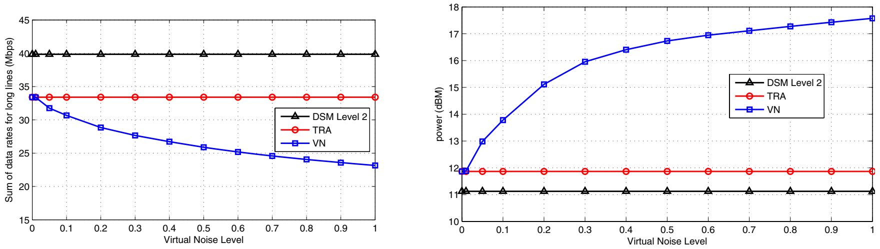

(a) (b)

Fig. 6. (a) Sum of data rates on lines 6-10 for downstream ADSL2+ with DSM Level 2 Band-Preference, DSM Level 1 TRA, and Virtual-Noise/AMA. (b) Average power use for the configuration of Fig. 5.

Table 1. Line lengths and target rates of Fig. 5

User Index	1	2	3	4	5	6	7	8	9	10
CO/RT Location (km)	1.6	1.6	1.6	1.6	1.6	0	0	0	0	0
CPE Location (km)	2.75	3.0	3.2	3.2	3.0	1.8	1.8	2.3	2.0	2.2
Target Rate (Mbps)	10	10	10	10	10	RA	RA	RA	RA	RA

larger in more boundary situations. The example provided is mid-range in terms of losses to be expected.

A second situation occurs in Figures 7, 8(a) and 8(b) for upstream VDSL. Gains are larger because higher-speed DSLs use wider bandwidth and thus, experience larger amounts of crosstalk, and consequently, benefit even more from politeness. Figure 7 illustrates a line-length distribution for upstream VDSL. Many other line length distributions could be used without changing the conclusions. This particular set is reasonable. In these simulations, the operators for odd-numbered lines will use various strategies for VN, while the operator for the even numbered lines use DSM Level 1 or 2 without VN. The target data rates appear in Table 2.

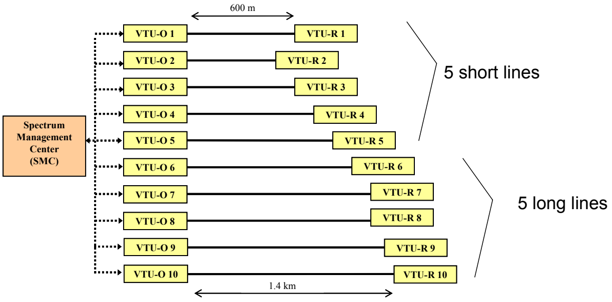

Fig. 7. Line-length distribution of 10 lines used in the upstream VDSL study of Level 1 and 2 DSM vs VN. (RA = rate adaptive). VN is used on odd-numbered lines in simulations.

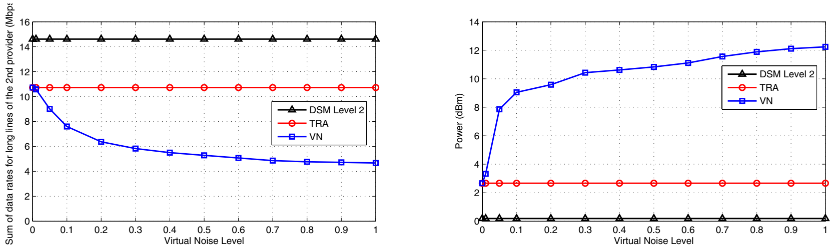

(a) (b)

Fig. 8. (a) Upstream VDSL VN's loss of data rate by a 2nd service provider when a first provider uses VN. (b) Average power use for the configuration of Fig. 7.

Table 2. Line lengths and target rates of Fig. 7

User Index	1	2	3	4	5	6	7	8	9	10
CO/RT Location (km)	0	0	0	0	0	0	0	0	0	0
CPE Location (km)	0.6	0.52	0.58	0.76	0.98	1.1	1.2	1.2	1.25	1.4
Target Rate (Mbps)	13	10	8	6	2	RA	1.5	RA	1.5	RA

Figure 8(b) illustrates, in addition to any power savings, the effect of virtual noise on other customers by again sweeping a virtual noise level from 0, for no VN use, to 1. The sum of the 2nd operators' data rates for lines 6, 8, and 10 is plotted both with and without virtual noise.

A final example in Figures 9(a) and 9(b) illustrates a similar situation for 1.5 km ADSL2+ single loop with AM radio noise varying by as much as 10 dB in amplitude assuming that tones from 550KHz to 1.7MHz are affected by AM noise and that AM noise source could exist at every 10 KHz radio channel in this frequency band.

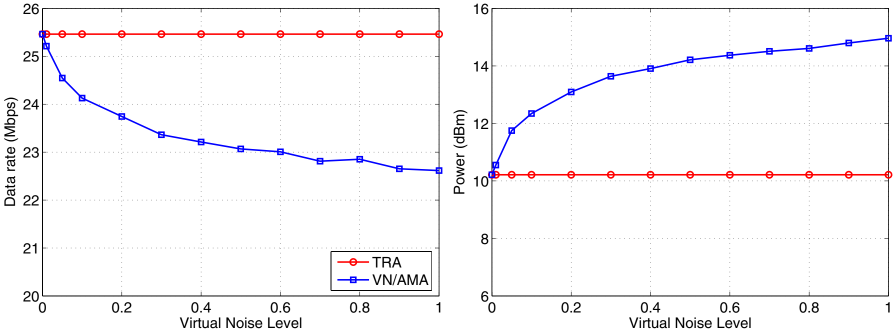

(a) (b)

Fig. 9. (a) ADSL2+ data rates for variable level AM noise. (b) Average power use for variable level AM noise.

4 DSM Level 3: Vectored Power Reduction

Level 3 DSM attempts cancellation of crosstalk. If this is possible, then transmitted power levels can be further reduced. This section quantifies such potential power savings.

4.1 Noise Cancellation Basics

DSM Level 3 vectoring [18] can cancel nearly all upstream crosstalk and also downstream far-end crosstalk. Such "Vectored VDSL" data rates can increase substantially over non-vectored environments[6]. Specific to management interfaces, vectoring naturally introduces a concept of user/customer order or priority. Users with higher priority have the crosstalk from their neighbors pre-removed before making their decisions, and thus, get a higher data-rate increase. The assignment of decoding order can significantly impact the different users' data rates and the total power consumption of the users' transceivers. This section evaluates the magnitude of data rate gains (and/or power consumptions) under a specific type of management interface that effectively determines the appropriate order for the service providers subtended by a common Level 3 vectored VDSL DSLAM.

Figure 10 illustrates the basic vector-cancelling situation: A common outside noise impinges upon two DSL lines, and the cancellation reconstructs the noise from line 2's receiver and subtracts it from line 1. In digitally-duplexed VDSL2, this cancellation can occur independently on every tone of the synchronized DMT line signals, thus simplifying implementation. Essentially, on any tone during every symbol period, line 2's data is detected in the presence of the noise, and the resultant noise-only signal is reconstructed from the decision and channel knowledge[7]. The consequent noise is then filtered to align its gain and phase with the noise on line 1 and then subtracted from line 1. Any far-end crosstalk (FEXT) from line 2 into line 1 can also be removed by filtering the decision-device output and subtracting that FEXT noise also. Line 1 then has lower noise prior to its detection and thus, can carry a higher data rate in the absence of the common external noise and FEXT[8]. Such common noise can be crosstalk from other lines out of the vectored group. This entire process applies equally with reversal of the terms "user 1" and "user 2," thus emphasizing the trade-off in decoding order on each tone. For 3 or more users, the concept readily generalizes

---

[6] A concatenation of Level 3 DSM and VDSL is sometimes called VDSL3 (which also could eventually be the name for the current G.vector effort within the ITU, which focuses on Level 3 DSM vectoring use in VDSL).

[7] Such reconstructed noise is the difference between the tone decision devices' input and output.

[8] It is well established now [19] that the diagonal dominance of common-vector-group FEXT allows a linear pre-decision filtering to remove this FEXT only. Thus, order for such diagonally dominated systems' FEXT is usually of little consequence. However, this is not true for the other common noises, nor is it true for shared line use as in Section 4.3.

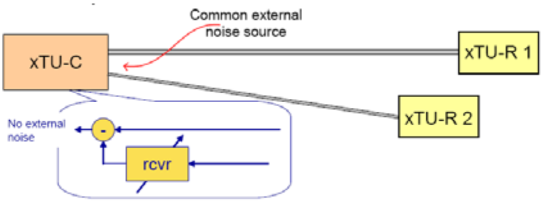

Fig. 10. Simple illustration of vectoring

with a highest priority (benefit) for the last user in the order, decreasing to the first user whose benefit is least.

4.2 Reducing Power by Cancelling Noise

Figure 11(a) illustrates the savings in power for a 3 Mbps upstream vectored VDSL vs. line length with simulation settings found in [20]. The power savings can be substantial for this case of 4 vectored VDSLs and 4 other VDSLs outside the vectored group. For such a system, with a good analog front-end design, the spatial correlation of the noises should approach 1 (that is, .99 is most realistic). Typically, when the number of vectored lines equals or exceeds the number of out-of-vector group lines, the spatial noise correlation is high. This effect can be made pronounced by using split-pair sensors in the analog front-end thus augmenting the number of dimensions used. Even with low spatial correlation, the savings is substantial. A factor of 10 or more savings essentially causes the transmitted power to be negligible. At this point, signal processing costs would limit further power reduction.

Figure 11(b) instead illustrates the power savings as a function of data rate, and of course the power increases as the data rate increases. However, the gap

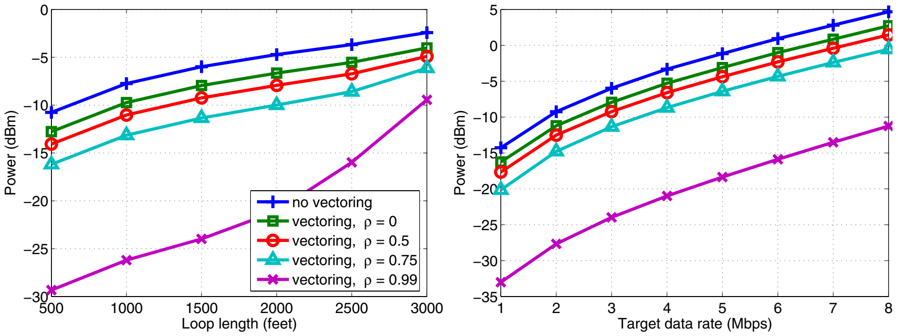

(a) (b)

Fig. 11. (a) Power savings vs. loop length for 3 Mbps upstream vectored VDSL. (b) Power savings versus data rate for upstream vectored VDSL.

(savings) becomes relatively constant for these 1500ft lines. Again there are 4 vectored lines and 4 lines with FEXT from outside the vectored group. Additional configurations appear in [20][21].

4.3 CuPONs

CuPONs [22] copy the passive-optical-network (PON) architecture of shared line bandwidth for up to 2-4 customers, using the multiple-pair drops as shown in Fig. 12(a). Such multi-drop sharing allows much higher bandwidths - for instance, 4 line drops can have over 1 Gbps of symmetric bandwidth as first noted in [23]. 1 Gbps shared over 4 users is 250 Mbps each, while it is 500 Mbps symmetric for each of two users. Such data rates are very feasible at line lengths of 300 to 500 meters (1000-1500 feet). The CuPON makes use of all the copper and in particular, allows vectored receivers over the 4 bonded pairs. Then crosstalk from outside the vectored group can again be cancelled. The concept of order between the users sharing common lines is again important in CuPON.

The downstream problem is known in multi-user information theory as the broadcast channel. This channel requires a duality transformation [24] of the measured FEXT and noise correlation matrices to use the algorithm discussed in Section 4.2 directly. Also, the spatial correlation between noises needs to be reported for the (2 to 4) common pairs at each customer location. Technically, only the noise-whitened equivalent insertion loss (Xlog's) could be reported, but this would be an abuse of the definition of insertion loss as currently defined in standards. Thus, a separate reporting of the spatial noise correlation matrix would be desirable.

Figure 12(b) illustrates the further potential power reduction and rate increase using the CuPON architecture with respect to best power and rates for no vectoring and vectoring.

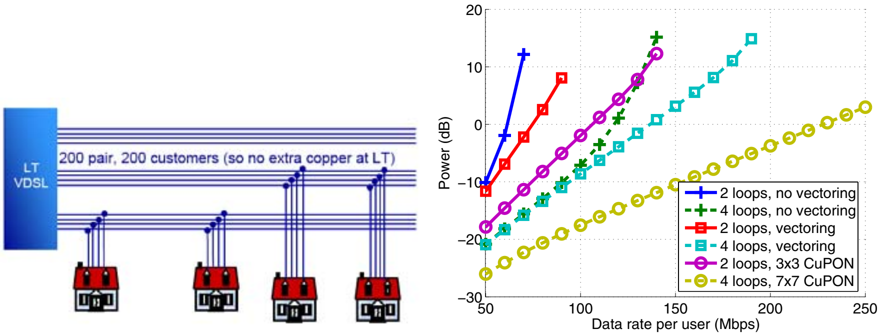

(a) (b)

Fig. 12. (a) Basic CuPON access-network architecture. (b) Power savings with shared CuPON architecture versus data rate.

5 Conclusion

Dynamic Spectrum Management offers many successively larger mechanisms to reduce consumed power in DSL systems, as well as to improve data rates. Such power savings could be of serious consequence to large DSL deployments that routinely consume megawatts of power today. The consequent savings provide a way to make copper greener in an evolving and increasingly energy-efficient world.

References

1. Dynamic Spectrum Management Technical Report, ATIS Committee NIPP Prepublished document ATIS-PP-0600007 (2007)
2. Starr, T., Sorbara, M., Cioffi, J.M., Silverman, P.J.: DSL Advances. Prentice-Hall, Englewood Cliffs (2003)
3. Masson, J. L., France Telecom R&D: DSL access and advantages from DSM in France. In: IEEE Globecom Access, Washington, DC (2007)
4. Foster, K., BT Design UK: DSL & Dynamic Line Management. In: IEEE Globecom Access, Washington, DC (2007)
5. Starr, T.: AT&T: Experiences from Dynamic Spectrum Management. In: IEEE Globecom Access, Washington, DC (2007)
6. Cioffi, J.M.: Dynamic Spectrum Management (DSM): 3 Steps to Ubiquitous High-Speed DSL. In: IEEE Globecom Access, Washington, DC, ASSIA, Inc. (2007)
7. Kerpez, K.: Telcordia: Operationalizing DSM into DSL Test and Maintenance. In: IEEE Globecom Access, Washington, DC (2007)
8. Cioffi, J.M.: ASSIA, Inc.: Dynamic Spectrum Management (DSM): 3 Steps to Ubiquitous High-Speed DSL. In: IEC Broadband World Forum, Berlin (2007)
9. Polano, M.: Telecom Italia: xDSL stability and performance: current issues and future solutions. In: IEC Broadband World Forum, Berlin (2007)
10. Berndt, E.: Goals for the Future Access Network and Economic Aspects. In: IEC Broadband World Forum, Berlin (2007)
11. Cook, J.: BTexact UK: Simulation parameters for discussion in the NICC-DSL TG. NICC Contribution PNO-DSLTG/CP38 (04)2 (2004)
12. Cendrillon, R., Yu, W., Moonen, M., Verlinden, J., Bostoen, T.: Optimal Multiuser Spectrum Balancing for Digital Subscriber Lines. IEEE Trans. Commun. 54(5), 922–933 (2006)
13. Cioffi, J., Lee, W., Jagannathan, S., Ginis, G.: The Inherent Simplicity of Distributed Band Preference. ATIS Contribution NIPP-NAI-2007-129R2, Miami, FL (2007)
14. Jagannathan, S., Cioffi, J.M.: Distributed adaptive bit-loading for spectrum optimization in multi-user multicarrier systems. Elsevier Physical Communication 1(1), 40–59 (2008)
15. Lee, W., Kim, Y., Brady, M., Cioffi, J.: Distributed Band-Preference Dynamic Spectrum Management for Digital Subscriber Lines. Submitted to IEEE Trans. Commun.
16. Lee, W., Jagannathan, S., Cioffi, J., Ginis, G., Silverman, P.: Higher-Rate Level 2 DSM Power-Saving Examples. ATIS Contribution NIPP-NAI-2007-160R2, Vancouver, Canada (2007)
17. Verlinden, J., Bruyssel, D.V.: Virtual Noise Mechanism. ATIS Contribution NIPP-NAI-2005-049, San Francisco, CA (2005)

18. Ginis, G., Cioffi, J.M.: Vectored Transmission for Digital Subscriber Line Systems. IEEE J. Select. Areas Commun. 20(5), 1085–1104 (2002)
19. Cendrillon, R., Ginis, G., Bogaert, E.V., Moonen, M.: A Near-Optimal Linear Crosstalk Precoder for Downstream VDSL. IEEE Trans. Commun. 55(5), 860–863 (2007)
20. Lee, W., Chen, C., Jagannathan, S., Cioffi, J., Ginis, G., Silverman, P.: Power and Rate Management for Level 3 DSM Vectored DSLs. ATIS Contribution NIPP-NAI-2007-162R2, Vancouver, Canada (2007)
21. Cioffi, J., Jagannathan, S., Ginis, G., Brady, M.: DSM is for Unbundled DSL. ATIS Contribution NIPP-NAI-2006-087R1, Savannah, GA (2006)
22. Cioffi, J.M., Jagannathan, S., Mohseni, M., Ginis, G.: CuPON: The Copper Alternative to PON - 100 Gb/s DSL Networks. IEEE Commun. Mag. 45(6), 132–139 (2007)
23. Lee, B., Cioffi, J.M., Jagannathan, S., Mohseni, M.: Gigabit DSL. IEEE Trans. Commun. 9(55), 1689–1692 (2007)
24. Cioffi, J.M.: EE479 course reader, `http://eeclass.stanford.edu/ee479`