Vectored DSLs with DSM:

The road to Ubiquitous Gigabit DSLs

J. Cioffi, M. Brady, V. Pourahmad, S. Jagannathan, W. Lee, Y. Kim, C. Chen, K. Seong, D. Yu
Stanford University
Stanford, CA 94305-9515 USA
cioffi@stanford.edu

M. Ouzzif, H. Mariotte, R. Tarafi
France Telecom R&D
Lannion, FRANCE
meryem.ouzzif@francetelecom.com

G. Ginis, B. Lee, T. Chung, P. Silverman
ASSIA Inc.
Redwood City, CA 94065-1198 USA
gginis@assia-inc.com

Abstract- **This paper investigates the Digital Subscriber Line (DSL) data-rate increases possible with the Dynamic Spectrum Management (DSM) methods known as Level 3 Vectoring and Level 2 Band Preference. A suggested sequence of increasingly binder-adaptive DSM steps appears herein to assist and motivate DSL service providers and equipment vendors to progress in use of DSM. In particular, studies of bounds of spectral balancing find Band Preference as a practical method that provides the highest possible Level 2 performance in bundled or unbundled DSL environments. Investigations of vectoring begin with differential vectoring and show very high DSL data rates. These data rates increase further through the use of full-binder vectoring, leading to projections of feasible DSM vectored implementations of multi-100Mbps DSLs.**

Keywords: **DSM, spectrum balancing, vectoring.**

I. INTRODUCTION

Dynamic Spectrum Management (DSM) provides an evolutionary path towards a goal of ubiquitous single-line 500 Mbps/customer Digital Subscriber Line (DSL) service. This paper suggests a series of steps from present DSM use in DSL to these very highest DSL speeds attainable with future ***vectored*** DSM. Vectored-DSM DSLs coordinate simultaneous binder transmissions to reduce and to exploit crosstalk, while also removing most other noises. Such vectored DSLs promise to best use standardized VDSL systems to achieve 500 Mbps per-line data rates on lines of up to 400 meters in length. A neighborhood DSL binder of 200 lines could then carry 100 Gbps of potentially shared bandwidth. While such binder speeds exceed the speeds of present PON [1] or DSL deployments, this paper explores various DSM technologies that would lead towards these eventual binder capacities.

DSM methods are increasingly gaining acceptance for use in DSL networks. DSM methods (see [1] for an overview) have recently appeared in various standards, most notably the North American DSM Report [2] to be released in 2006. DSM elements also appear in the "PLOAM" [3], ADSL[4],[5] and "VDSL2" [6] efforts of the ITU, and in various evolving reports and documents of the DSL Forum [7],[8]. Many service providers and vendors currently evaluate the introduction of early DSM methods in their networks and products. The DSM Report provides a widely used characterization of DSM levels described in Table 1 that will prove convenient for describing DSM evolution in this paper.

TABLE I
DYNAMIC SPECTRUM MANAGEMENT LEVELS [2].

DSM Level	Description
0	No DSM
1	Single-line Politeness and Impulse Control
2	Multiple-line Spectrum Balancing (spectra controls)
3	Multiple-line Vectored Coordinated LT-side downstream transmission and upstream reception

Level 0 (no DSM) systems still represent the majority of DSL services in present use. However, Level 1 systems have been deployed in some DSL networks and are in trial use in many other DSL networks. Level 1 DSM systems make use of the data- and control-management interfaces defined in [2]-[6] to monitor, and to adjust when appropriate, the spectra levels and degree of forward-error-correction used on individual DSLs without specific knowledge of other DSLs. Such DSM systems can be viewed as components of operational support systems known as Spectrum Maintenance Centers (SMC) and administered by service providers. The basic reference

[1] Passive Optical Networks (PONs) share a bandwidth of one of roughly 1 Gbps (APONs are 622 Mbps, EPONs are 1 Gbps, and new GPONs can be as high as 2.4 Gbps) over 32 to 256 users. No equivalent binder sharing yet occurs in deployed DSL systems.

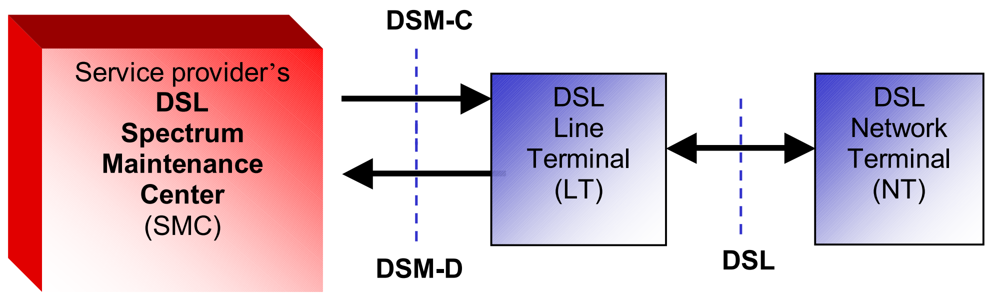

Fig. 1. DSM Reference Diagram.

diagram appears in Fig. 1 where the DSM-D interface carries data from vendor equipment through an Element Management System (EMS) for LTs (line terminals or "DSLAMs") or an Auto-Configuration Server (ACS) for NTs (network terminals or CPE modems) [7]. The DSM-C or control interface carries commands usually known as "profile settings" to the DSL lines. The DSM-C and DSM-D interfaces are known as a Management Information Base (MIB) in some standards and systems. References [2],[3], and [7] specify these interfaces in greater detail.

Many examples of productive example DSM use occur in [1] and [2], particularly the informative annex C of [2]. The interface to the SMC is standardized so that equipment vendors can support the service provider's management[2]. Such early DSM systems can provide very large improvements in the economics of DSL service provisioning and maintenance and thereby motivate progress in DSM use. This paper focuses upon the multiple-line DSM levels of 2 and 3, leaving the reader interested in Level 1 to review [1] and [2] and the references therein.

Section 2 investigates Level 2 DSM systems for both bundled and unbundled environments[3]. Level 2 systems make use of binder characterization, which is essentially the knowledge of binder topology or of the crosstalk power-spectra coupling between some or all the binder's lines. Section 2 overviews the theoretical bounds of Optimal Spectrum Balancing [9] that applies only to bundled networks. Section 2 then proceeds to practical open-interface implementations known as "band preference" in [2]; the latter of which uses open standardized interfaces and can apply equally well in bundled or unbundled situations. In particular the use of mutually compatible adaptive spectra within a binder is illustrated in terms of DSL rate/range improvements in Section 2. These methods provide gain and motivate vectored DSM systems in Section 3.

Section 3 proceeds to Level 3 vectored systems that imply line-terminal coordination of the transmitted downstream and received upstream signals within those DSLs controlled by any single service provider. Such systems have largest gains when all lines are vectored, but also offer significant improvements in either bundled or unbundled environments, in which latter case some band preference may also be very helpful. Section 3 also explores the fascinating concept of the full capacity of the cable of twisted pairs (which is hundreds of times the data rates of passive optical networks in use today) and explains how a "roving capacity" of the binder can be dynamically allocated to selected lines that need higher data rate. Section 4 concludes with the suggestion that 500 Mbps per-line DSLs may indeed be feasible with the combined efforts of vendors, service providers, and an open embracement of DSM technology.

II. Level 2 DSM: Band Preference

Fig. 2 illustrates the basic concept of a rate region in spectrum balancing. Two mutually crosstalking DSL services share a binder. Each acts as noise to the other. The short line has generally higher data rates, but can create a substantial noise to the long line, especially when the short line's transmitter is closer to the longer line's receiver.[4] Each point in the rate region is achievable with a different pair of spectra for the users. The differing spectral-pair choices offer a multitude of DSL speeds. Static or "fixed" choice of spectra on all lines not surprisingly causes a reduction in achievable rates to the lower left region, often a much smaller space of possible rates than with dynamic spectra. DMT DSLs already ubiquitously use dynamic spectra and this existing facility (often known as "bit swapping") can be configured (with complete compliance to all standards) to enlarge the rate region with respect to fixed use. In fact, careful use of the widely understood and practiced "water-fill" loading with ample but polite fixed margins can effect significant data rate increases as

[2] One or two vendors may initially attempt to incorporate the management into their products through proprietary interfaces, but such interfaces may have negative impact on open management, costs, unbundling regulation, and consequently DSM practice and standardization. However, most vendors now support DSM interface standardization.

[3] Level 1 DSM systems apply in the same manner to bundled or unbundled DSL networks.

[4] Such situations occur when a fiber-fed DSL terminal has lines that share a binder with existing central-office fed ADSLs, an increasingly common configuration in most networks; another example is the upstream transmissions of any mixture of VDSL line lengths.

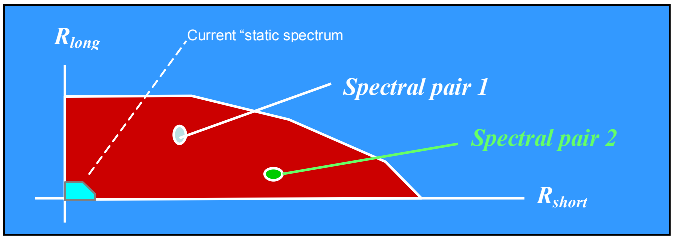

Fig. 2. 2-user rate region for spectrum balancing.

in the iterative water-filling of [1] and the informative appendices of [2]. These Level 1 methods essentially are single-line managed and do not centrally use or need to know the binder topology, allowing a fully distributed implementation of the choice of spectra within the rate region to the individual DSL lines. Service provider management methods within the SMC can impose basic margin limits and impulse control settings for customer lines according to line and network history, experience, and observed rate/range footprint and/or maintenance cost reductions.

Level 2 DSM makes use of knowledge of the binder topology. Binder identification can occur through a number of mechanisms within the SMC. Mutual interference between lines can be inferred via sophisticated correlation methods based upon reported standardized DSM-D data from the lines that do report to the SMC, telephone-number/address correlation to public information already on the worldwide web, and other published information about topology. Inference of suspected crosstalkers and types can be calculated for same-provider DSLs as well as competitive-provider DSLs without violation of regulatory policy. Binder identification is a topic beyond the scope of this paper, but the sequel presumes it can be estimated with some reasonable tolerance and probability that is a function of the service-provider operational practice and the regulatory situation.

Subsection 2.1 investigates the bounds of the rate region in Fig. 1 for a bundled-only Optimal Spectrum Management (or Optimal Spectrum Balancing) situation in which all users know all other users' practices and exact crosstalk coupling (without vectoring). The OSM data-rate bounds obtained are often significantly beyond the current practice of DSL. Subsection 2.2 then proceeds with the standardized practice of "band preference" – a single bit indication given to DSL line modems that an SMC is present and attempting to estimate binder topology and set the lines that it controls to profiles that better serve all lines within a binder. DSM-capable modems that observe band preference can then act to enhance their own performance by adjusting their internal algorithms accordingly. Essentially all the gain possible with the proprietary OSM can be achieved with an open, standardized, and distributed Level 2 DSM implementation by cooperation between service providers and their equipment. Subsection 2.3 returns to some of the regulatory advantages with an open implementation.

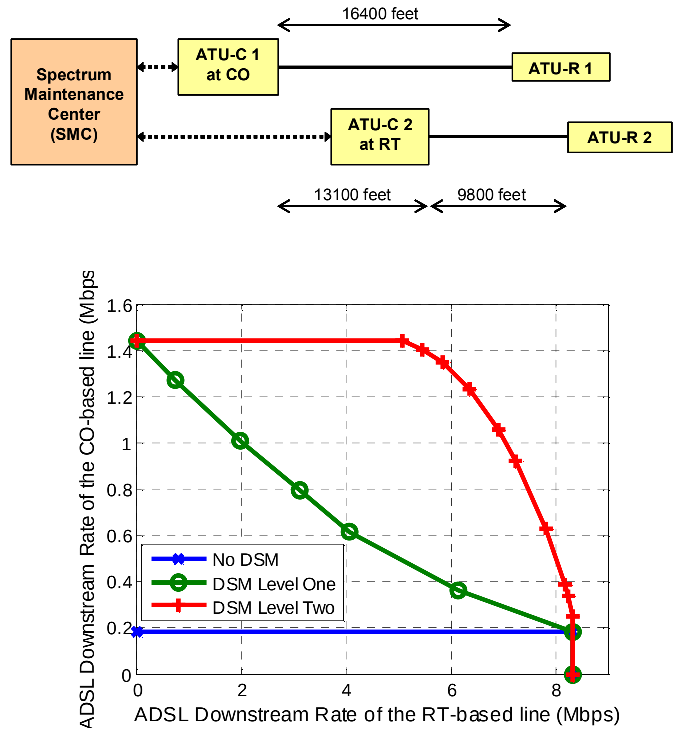

Fig. 3. Illustration of level 2 OSB gains over level 1 polite water-filling use.

A. *Theoretical Bounds of Centrally Controlled Bit-Swapping*

Optimal Spectrum Management (OSM) bounds were investigated by Cendrillon in [9]. OSM is called OSB (optimal spectrum balancing) in Annex A of [2]. The OSB technique essentially describes a procedure to maximize a weighted sum of data rates for all DSL subscribers within a binder. The OSB procedure requires central knowledge of the allowable spectral masks for all users and the crosstalk power coupling transfer functions between all pairs of users. The optimization procedure makes the assumption that crosstalk occurs only on the same tone between users (which requires Level 3 synchronization, see Section 3 on Level 3) and proceeds to allocate all the tone-energies (equivalently power spectra) and the number of bits/tone for Discrete Multi-Tone (DMT) standardized ADSLs and VDSLs from the SMC of Fig. 1. By varying the relative weighting of all the users (so that these non-negative weights all sum to 1), the outer limit of the rate region possible in Fig. 2 can be constructed for any binder. Such a result is useful in that it allows determination of the limits of the various user data rate trade-offs.

Level 2 OSB methods provide a large gain over Level 1 methods in situations where crosstalk is strong, particularly mixtures of long and short lines. Fig. 3 from [2] illustrates a situation with 4 long ADSL lines of 16,400 ft (6km) and 4 short lines of 9800 ft (3.6 km).[5] The largest achievable region is Level 2 DSM while the static Level 0 DSM has very poor performance, which is classic for long-short-DSL mixture.

[5] Coding gain of 3 dB with margin of 6dB, Gap of 9.8 dB with ADSL1 spectral and standard compliance.

Both Levels 1 and 2 use 6 dB margin and 3 dB of coding gain in the simulations. Level 1 has intermediate performance and uses iterative water-filling with a maximum allowed margin of 6 dB.

The algorithm of OSB is very computationally complex, essentially having the number of calculations that rises exponentially in the number of users for each and every tone, and so a few approximations have been derived in [10a,b] and [11] to simplify the OSB calculations and essentially get the same performance. All these OSB approximations presume a central implementation of bit-swapping by the SMC (with some interesting message-passing options in the SCALE algorithm of Papandriopoulos in [11]). As such, unfortunately, these methods all then require a bundled implementation that is almost always contrary to regulatory practice. Furthermore, the delays involved with central control of swapping with change of noise conditions force a need for rapid SMC response and high-bandwidth control/data flows from both DSL modems to the SMC. So even if bundling is allowed in certain areas of the world, the high-speed central reaction to noise changes is not desirable and may not be feasible in most situations. Additionally, there is no facility in current standards for such centrally controlled bit swapping. Thus, the introduction of band preference in [12],[13] allows a mechanism to achieve the OSB gains with a decentralized implementation that is of essentially no complexity increase with respect to present Level 1 implementations and can work in bundled or unbundled situations.

B. Band Preference Distributed Loading

Band Preference (BP) as standardized in [2] is a simple one-bit indicator from an SMC to all DSLs controlled by that SMC. BP signals to the DSL (over the DSM-C) to use one of two loading algorithms. The band preference is either "on" or "off." Different service provider's SMCs can make independent band-preference decisions if the binder or cable is unbundled. When BP is off, the DSL lines operate as they would otherwise (presumably Level 1 or Level 0 DSM). When BP is on, DSL transceivers each independently and locally run an algorithm that is a modified version of the water-filling they would run otherwise. In standardized DSLs, there is a quantity known as the PSDMASK[n] that is distributed by an SMC (or set somehow by an operator or standard) to DSL modems. The PSDMASK represents an upper limit on power spectra at each tone when BP is off or on. When BP is on, the SMC uses the PSDMASK to provide an indication of a band-preference weighting factor to the loading algorithm (so the PSDMASK setting when BP was off remains the power-spectral density limit if the DSL line modems decide to stop using BP because of a noise change). This BP-on weighting factor is typically 2 to 32 values in different bands that amplify the loading algorithm's internal use of energy in applying water-filling (or equivalently discrete loading) algorithms. These weighting factors allow the distributed modified water-filling algorithms to achieve OSB performance. The modems continue to load, and if noise changes occur that exceed a threshold, then the algorithm reverts to normal (BP off) water-fill loading (because the line's noise changed too much) and then waits for any future subsequent control from the SMC to use a new set of BP weighting factors with BP reset to on.

To understand the simplicity of the BP loading algorithm, it helps to restate the basic water-fill loading algorithm that assigns an energy E_n to each tone as a function of a measured channel to noise gain ratio g_n on each tone and a constant gap Γ such that the total energy for the user does not violate total power, maximum margin, or power-spectra density constraints. The basic equation is

$$E_n + \frac{\Gamma}{g_n} = \text{constant on all tones} \; , \qquad (1)$$

Normal (Level 1 or 0 DSM) water-filling

subject to the energy being non-negative and satisfying the PSDMASK constraints. Discrete implementations are discussed shortly. This simple water-filling concept leads either to maximum data rate at some given margin and allowed power (rate-adaptive DSLs), or to minimum power at some given data rate and power (fixed-rate fixed-margin polite DSLs).

With band preference on and the supplied set of weighting factors $\alpha_n \geq 0$ (essentially linear equivalents of the supplied PSDMASK[n] with some interpretation and scaling), the water-filling in Equation (1) is modified to

$$\alpha_n \cdot E_n + \frac{\Gamma}{g_n} = \text{constant on all tones} \; . \qquad (2)$$

Band-Preference (Level 2 DSM) water-filling

(2) is also subject to the energy being non-negative and satisfying the PSDMASK and power constraints. Rate adaptive or fixed-rate/margin water-filling can also be implemented within the BP water-fill just as in the normal case in (1). The only difference is the use of the energy scaling factor α_n that is supplied by the SMC. Usually there are less than 16 values of α_n for any user's bands. The value of α_n is computed as the linear-gain factor corresponding to the ratio of PSDMASK[n] when the setting for BP changes from 0 to 1. If the noise changes substantially on a line so that $g_{n,new}/g_{n,old}<$ threshold, the modem resets $\alpha_n = 1$ (BP is then effectively off for this modem and the scale factor is reset so normal water-filling in (1) is again used with the PSDMASK supplied when BP was last off or the nominal default). The reported standardized PSD[n] and SNR[n] in the DSM-D are more than sufficient for the SMC to see that BP has been reset by the particular modem. Alternately, a modem that could or would not implement BP would simply observe the supplied mask, but would be less resilient to large noise changes.

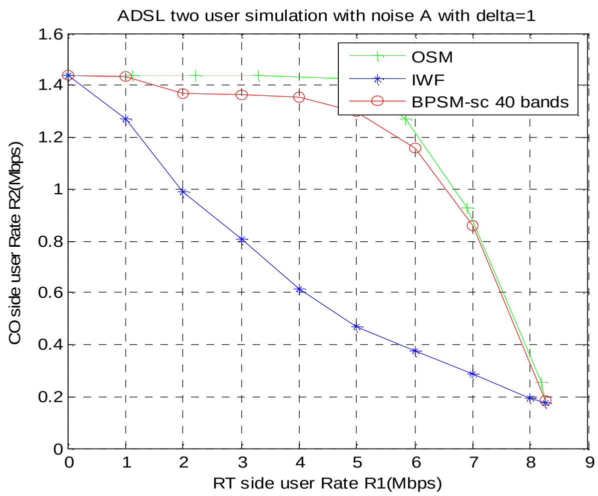

(a)

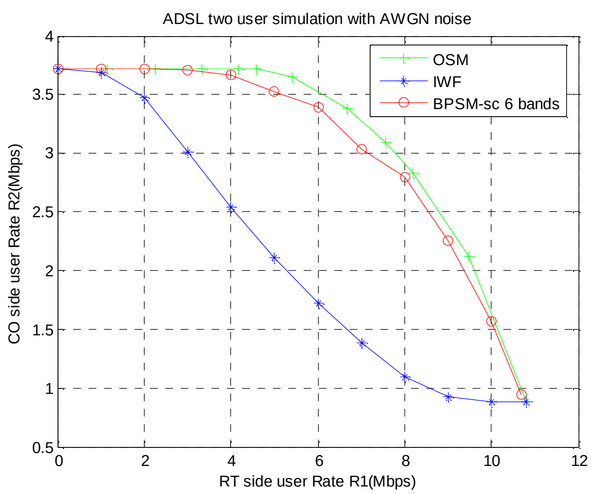

(b)

Fig. 4. BP example for (a) situation in Fig. 3 and (b) with just AWGN.

Example results of BP use appear in Fig. 4 for an ADSL mixed-binder situation with and without other non-adaptive noises. Fig. 4(a) is for exactly the same situation as in Fig. 3, while in Fig. 4(b), the noise is reduced to just the -140 dBm/Hz AWGN floor. Fig. 4(a) is roughly the same for 20 or more bands in band preference and degrades by about 10% (in data rate) for smaller number of bands down to 6[6]. In Fig. 4(b) with the flat AWGN, only 6 bands were necessary. In each case, OSB is also shown and the BP clearly comes very close in terms of rate-region limits.

The SMC can determine the factors α_n by various algorithms For instance, an algorithm of the type in [9]-[11] could be executed in the SMC and then the resulting PSDs limited or approximated over a few bands and the scale factors computed and distributed. Better mechanisms instead may approximate centrally exactly the situation ongoing in the distributed environment of the binder to determine appropriate values of α_n. The authors prefer the latter implementations for reasons of complexity and the more accurate central replication of the actual distributed resultant implementation. Such better algorithms may to date rest proprietary.

Most DSL loading algorithms are not quite "water-filling" per se. Instead, they are approximated by discrete (integer bits per tone) methods like the Levin-Campello procedure described in [14]. Such practical loading algorithms typically require an incremental energy table for the specific cost using the applied coding method to add each additional bit to a constellation ΔE_n (there may be many efficient ways to implement such a system without storing a full table of incremental energies for each tone). For BP water-filling, the table of incremental energies is scaled by α_n and then loading proceeds as in the normal discrete-loading situation. The presence of unequal α_n will tend to prefer loading in some bands over others, whence the name "band preference."

C. *Unbundled or Unregulated Use*

Band preference enables the use of Level 2 DSM in unbundled environments where binders of DSLs may be excited by the equipment of different service providers. Each service provider's SMC can independently make a decision on whether to set BP=1 and consequently supply a set of α_n (via a new PSDMASK on the DSM-D) to any subset of or all its controlled lines. If there is no SMC for any service provider, then $\alpha_n \equiv 1$ for all that service provider's presumed DSLs. Such BP=off use can lead to a presumption of worst-case spectra use by other service providers' SMCs, or those SMCs can be more adaptive to the observation of crosstalk as reported in the general already standardized data that is today supplied by standard-compliant DSL modems, each to their respective SMCs. If multiple service providers are using BP, then they all may experience yet larger average DSL data rates, but none may exceed existing applicable standardized (or regulator imposed) power spectra masks, nor may they exceed the programmed PSDMASK when BP is off. Thus, band preference represents a significant and practical improvement from a regulatory and implementation viewpoint.

III. Level 3 DSM: Vectoring

Level 3 DSM or "vectoring" enables the highest speeds known for DSLs and applies to single-line DSLs (so no bonding is necessary). With vectoring, a common DSLAM (more than likely a common multi-line card in the DSLAM or LT) synchronizes downstream DSL transmissions to a common DMT symbol clock (either the 4.3125 kHz clock of ADSL and VDSL or the 8.625 kHz double-width clock option of high-speed VDSLs). Digitally duplexed loop-timed circuits will then provide the same common symbol clock in the upstream direction. The term "vector" applies to the co-generation of a synchronized vector of simultaneously transmitted user signals in the downstream direction. These signals can go to different users in different customer locations

[6] Fewer than 6 rapidly degrades performance to the level of Iterative Water-Filling (IWF) that is shown for reference.

but are launched in cooperation into the binder. In a dual fashion, a vector of upstream signals is simultaneously sampled and provided to the same common DSLAM/card. Such vectored upstream systems can essentially eliminate all hostile crosstalk and indeed can even exploit friendly reinforcing crosstalk, thus leading to very high speed DSLs. Downstream vectored systems can eliminate all FEXT or can even use FEXT for diversity reinforcement of downstream transmissions.

MIMO (multiple-input-multiple-output) [15] is not vectoring. While MIMO requires bonding, vectoring does not require bonding. MIMO systems have been shown to improve more than "N times" the data rates because of mutual crosstalk cancellation when both transmitter and receiver are attached to all the lines. Vectoring can get most of the MIMO gains despite the fact that vectored systems must additionally satisfy the restriction that only one-side (either the transmitter for downstream OR the receiver for upstream) is allowed to subtract (or pre-subtract) crosstalk. Vectoring systems are detailed in [14],[16], and further detailed in [17] and can increase complexity moderately in DSL, leading to a question of the amount of gain versus the cost increase. Such complexity evaluations are addressed in the references, and instead this paper provides some basic characteristics of different types of vectoring as well as the possible data rates for good vectored designs are addressed, leaving the level of vectoring to DSLAM vendors.

A. Differential Vectoring

Ginis' differential vectoring [16] uses only differential excitation of each twisted pair (thus, phantoms and split-pair circuits are ignored). Differential vectoring is the earliest form of vectoring to be used. A classic result from that pioneering work is that all co-generated or co-received FEXT crosstalk in either direction can be eliminated (assuming different frequency bands for upstream and downstream, so no NEXT). Furthermore, upstream crosstalk and noises (like radio noise or impulse noise) can be suppressed even if this noise's source is not in the vectored group (this spatial cancellation performs better if the number of vectored lines exceeds the number of sources at any DMT frequency). The upstream ability to cancel out-of-vector-group noises is calibrated by the spatial correlation of the noise. When the correlation is high, then the out-of-group noise is largely canceled in upstream vector reception. High correlation means there is a single or small number of noise sources. Low correlation means there are a large number of sources each contributing a small part of the noise. AM radio noise and impulse noise have very high spatial correlation so they can be suppressed in the upstream by vectoring. Unbundled crosstalk noise from many unbundled sources can only be partially reduced in vectoring (so some level of band preference may need to accompany vectoring in unbundled environments with vectoring). For downstream transmission, all noise appears spatially uncorrelated because of the absence of common vectored receiver (whether or not the noise actually is correlated, since no cancellation can occur at the receiver[7]). When all lines of a binder are differentially vectored and active simultaneously for a single service provider, the rate region is aberrantly a hyper-rectangular box and each user has a single maximum achievable rate independent of what all others do.

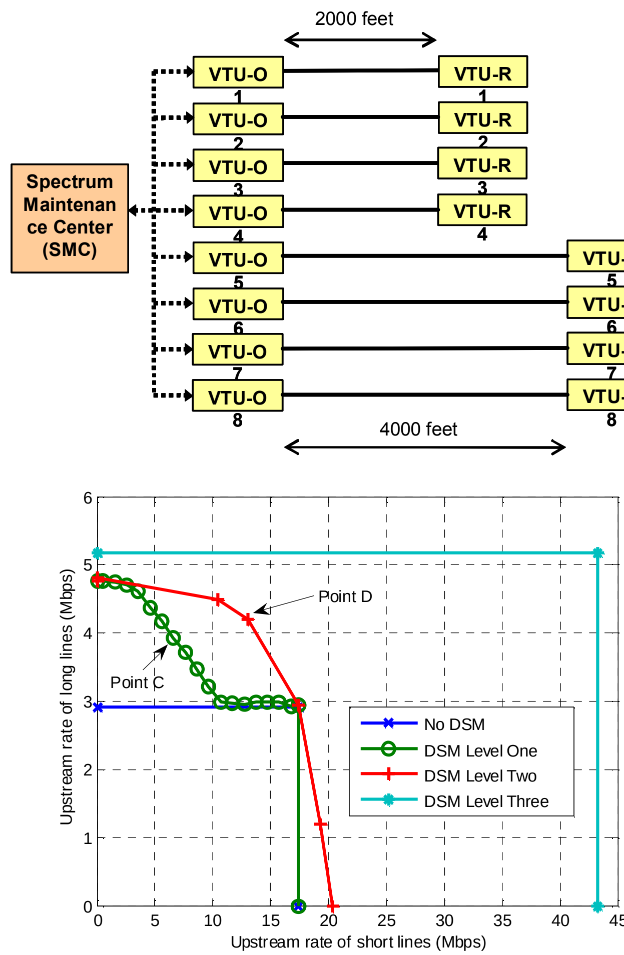

Fig. 5. VDSL DSM Level comparison

A number of plots from North America and Germany appear here for such vectoring. Fig. 5 (from [2]) illustrates the further substantial gain of Level 3 differential vectoring on the line configuration also shown. The parameters for this simulation are (coding gain 3 dB, margin 6 dB, Gap 9.5 dB, 14.5 dBm, 998 bandplan, bit cap 15, see [6]).

Fig. 5 illustrates the classic rectangular (or box in more dimensions) Level 3 rate region for differential vectoring. Fig. 5 is for upstream transmission. Fig. 6(a) illustrates a situation in Germany with a very high-power crosstalker from what is known as an HDB3. Standard VDSL [6] with the 997 bandplan was used and both upstream and downstream data rates appear for all lines of the same length with and without differential vectoring (the upper red curves use vectoring). The

[7] An exception is perfect space-time correlation where for instance a known sinusoidal signal could be subtracted, but then a sinusoid is really not noise.

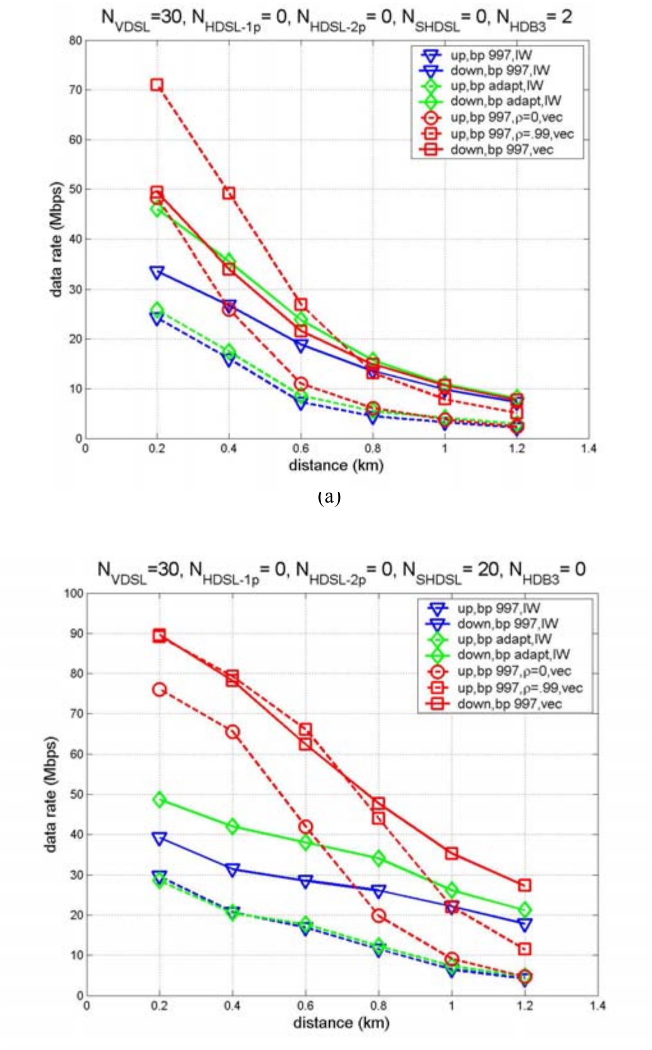

Fig. 6. Illustration of (a) HDB3 crosstalk and of (b) SHDSL crosstalk for .4mm lines.

results shown in Fig. 6 are also transferable to the 998 bandplan, which in Germany seems to become more relevant. A correlation coefficient for upstream spatial noise correlation is listed with 0 meaning no correlation and 1 meaning largest spatial correlation physically possible on the HDB3 noise. The latter is the realistic value because there are only 2 HDB3 sources (and 30 vectored lines). Fig. 6(b) also appears for the case of the more modern but equally invasive SHDSL crosstalk, where spatial correlation could be expected to be intermediate to 0 and 1.

Fig. 7 illustrates a situation proposed by North American operators for fiber-to-the-curb standardization when differential vectoring is applied. In these plots, VDSL2 with a 7-band North American service provider proposal is used [20]. The upstream and downstream rate goals are illustrated (and vary with length from 150 Mbps symmetric to 50-down/30-up at longer lengths) also and in all cases. The data rates are very

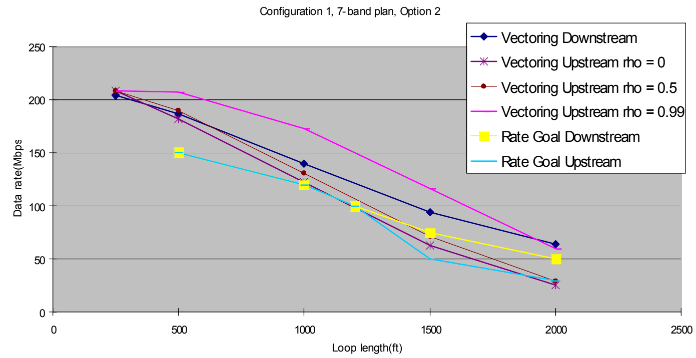

Fig. 7. North American fiber-to-the-curb high-speed VDSL data rates with differential vectoring.

high again for differential vectoring. A range of spatial correlation is used on the upstream -125 dBm/Hz noise model suggested by the service providers.

B. Full Vectored Binder Capacity

Full binder capacity exploits split-pair transfers and common-mode crosstalk within the cable of twisted pairs. This subsection reviews both effects, starting with the split-pair transfers and finishing with common-mode crosstalk. Full vectored binder capacity investigates a binder of *U* DSL lines that contains *2U* wires. There are many transfer modes of energy in such a binder of wires that are left dormant in differential vectoring. When these modes are all considered, the full capacity of the binder can be computed. The possibility of exploitation of such modes is clear in MIMO systems (such as in [18],[19] where symmetric Gigabit DSLs are observed on 4 bonded 300-meter-plus telephone lines by using all 8 wires). The extra capacity is reduced for single-line systems, but is still substantial.

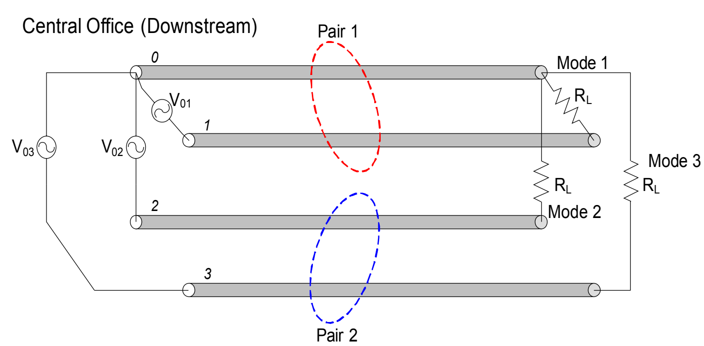

Fig. 8. Split-pair excitation of two twisted pair.

Split-pair Transfers

Fig. 8 illustrates split-pair transfers between two bonded twisted pairs or a quad in the downstream direction with three loads. The concept easily generalizes to more than two twisted pairs. With split-pair excitation, one wire (0) in one of the pairs is viewed as a reference for 3 sources, V_{01}, V_{02}, and

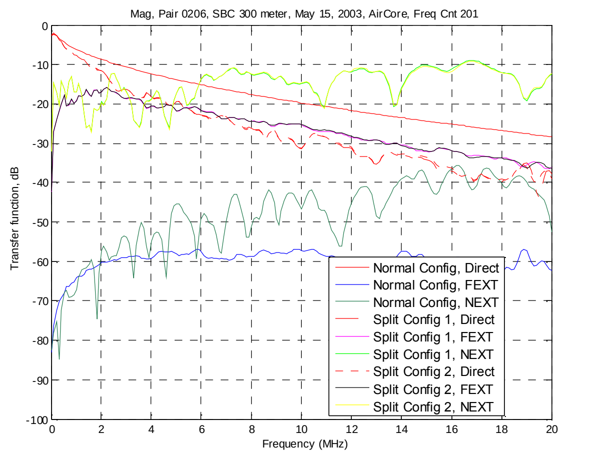

(a)

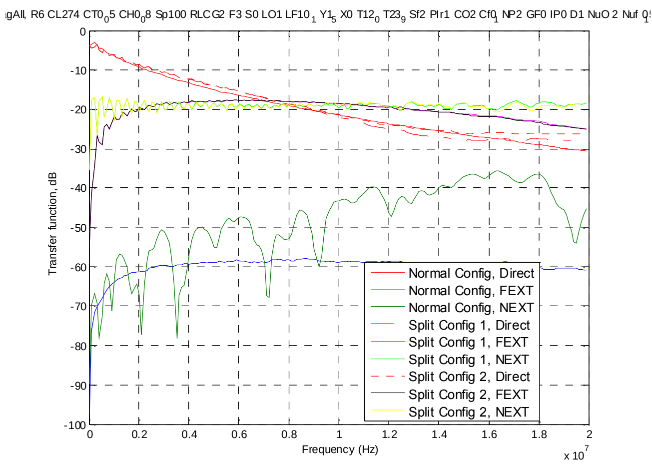

(b)

Fig. 9. Comparison of (a) measured split-pair transfers as well as (b) computed using the theory in Section 5 of the US DSM Report [2].

V_{03}. In normal differential operation, the second subscriber's excitation voltage is direct and is $V_2 = V_{02} - V_{03}$ and only two excitations appear. However, 3 sources are possible more generally and in fact can all be assigned to a single subscriber with bonding of the pairs. More generally with *U*>2 twisted pairs, then up to *2U-1* sources can be assigned to the *U* DSL subscribers when all excited pairs of wires are properly terminated. Some typical split-pair transfers appear in Fig. 9. In the plots, normal configuration refers to the usual differential excitation of a twisted pair (for example, the sources are placed between 0,1 and 2,3 respectively in Fig. 8). Split configuration 1 refers to excitation of the pairs by placing the sources between 0,2 and 1,3 respectively while split configuration 2 refers to placing the sources between 0,3 and 1,2 respectively. In all cases, the loads are placed in the same configuration as the sources.

Fig. 10 illustrates some data rates computed for split-pair excitation of quads with 3 sources as illustrated in Fig. 8. Fig.

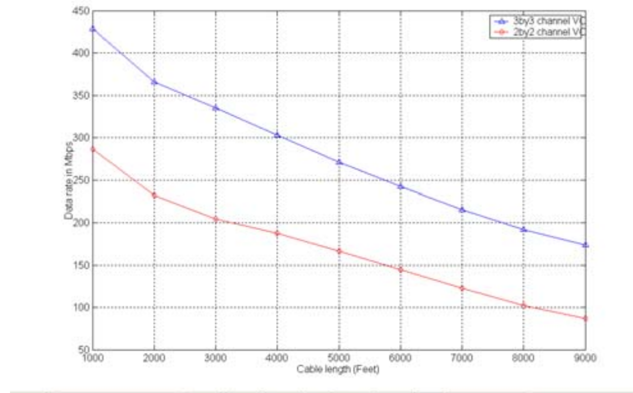

(a)

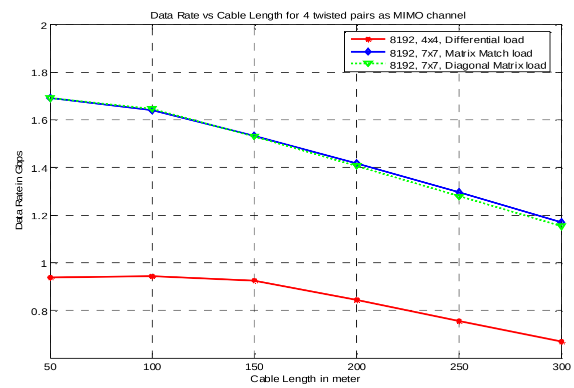

(b)

Fig. 10. Symmetric upstream and downstream data rates with and without split-pair use for (a) two pair (3 modes) and (b) four pair (7 modes).

10(a) is for 0.4mm European quads and Fig. 10(b) is for a bundle of 4 American twisted pairs. The only other noise illustrated is that of -140 dBm/Hz, so the plots represent an upper limit. The plots also represent the symmetric upstream and downstream data rates. So, while differential vectoring alone (see 2x2 result in Fig. 10(a)) could achieve as much as 200 Mbps at 1 km, the use of split pairs increases that upper bound on rate by about 50% to 300 Mbps symmetric.

Common Crosstalk Modes

Fig. 11 alternately illustrates the possibility of common-mode transfers within a binder where the sheath has been well-connected to ground at all junctions. The configuration of Fig. 11 was simulated using the same Multi-conductor Transmission Line (MTL) model as in Chapter 5 of the DSM Report [2] with some additions to include the sheath as further detailed in this section. The additions allow calculation of the common-mode direct and common-mode FEXT transfer functions with respect to a common surrounding and grounded sheath as in Fig. 11.

Fig. 11 shows two twisted pairs explicitly, but easily generalizes to *U* twisted pairs. Fig. 11 shows the downstream

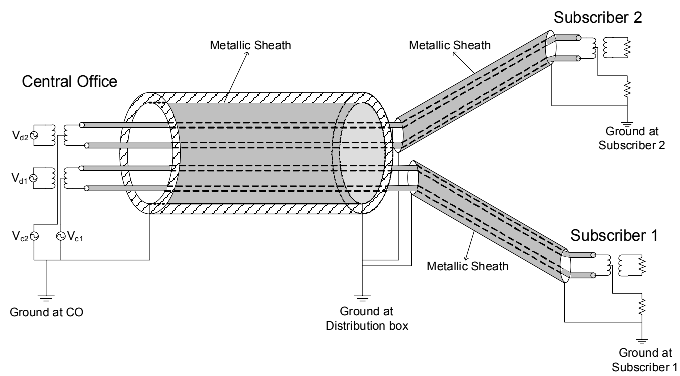

Fig. 11. Binder with grounded sheath and full common- and differential-mode excitation.

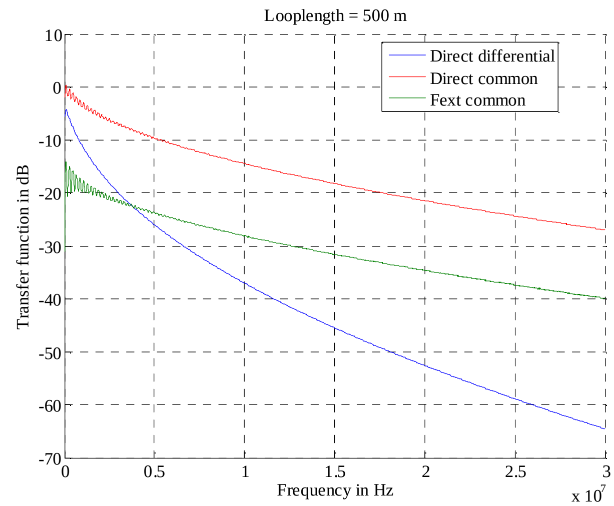

Fig. 12. Illustration of common transfers, direct and crosstalk, as well as nominal differential transfer.

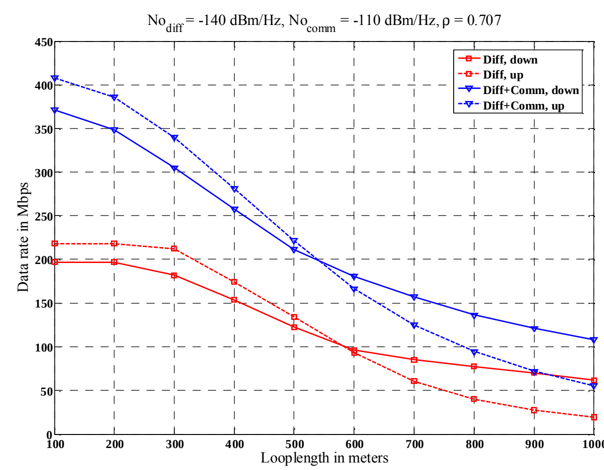

Fig. 13. Illustration of the additional data rate available from full binder use for a single line, upstream and downstream shown for a popular proposed band plan for VDSL2 [23].

configuration. Common mode excitation downstream would likely be achieved by driving the line-side center tap of the central-office (or line-terminal) transformer with respect to ground. The central office ground is presumed attached to the sheath also (which is the specified practice of most telco's, although such grounding is sometimes accidentally missing). Each subscriber presumably attaches also to the sheath ground of the drop cable and has the connection from the network interface to the receiving modem box (this requires some care in installation and has implication for splitters and micro-filter circuits that are not addressed here). This contribution presumes the integrity of the sheath ground connection and models, then exploits it, leaving debate on viability to future discussion. The intent then is to evaluate the data rate improvements that would occur if the effort were to be (or has been) made to ensure sheath grounding at all junctions as in Fig. 11. Note the subscriber side also extracts the signal from the line-side center tap of the subscriber transformer. The upstream configuration is a dual that is easily derived by reversing the source and loads in the diagram and presuming reasonable hybrid coupling circuits to separate up and downstream transmissions in all transfer modes. More detail on the calculation of the transfer modes in Fig. 11 is provided in [22] and involves use of the well-known method of images in multi-line transmission analysis. However, Fig. 12 shows the transfer functions for a 500 m segment of .5mm twisted pairs.

Fig. 13 is a plot of data rate for upstream (dashed curves) and downstream (solid curves) with full vectoring (blue curves) versus the same situation with differential vectoring (red curves). At shorter line lengths (all lines in the binder are the same length for this simulation, even though they go to different users), upstream actually exceeds downstream data rate because the OSFA band plan [23] used has significant upstream bandwidth allocated at high frequencies. Note the background noise in the common receivers is set 1000 times larger (at -110 dBm/Hz) than the level of differential noise; nonetheless, the vectoring still significantly increases rates in the presence of the large noise. The reason for setting the common noise so high is that receiver circuits that sense common mode signals will be far more sensitive to environmental noise. Some additional similar simulations show that for lengths greater than 500 meters that additional common-mode noise causes a loss of about 50 Mbps for each 10 dB. The loss is smaller than 50 Mbps for shorter lengths and almost negligible for lengths under 200 meters, but appears to be limited by bit-cap effects (without bit caps at 15, the data rate can increase a few 100 Mbps more).

Reference [22] provides greater detail on reinforcement and reassignment of data rates from one line to another with full vectoring. For this higher-level description, a summary of findings is that particularly for upstream transmission, there appears to be a significant trade-off between the data rates of different upstream users even with Level 3. This differs from Level 3 differential vectoring, where each line essentially

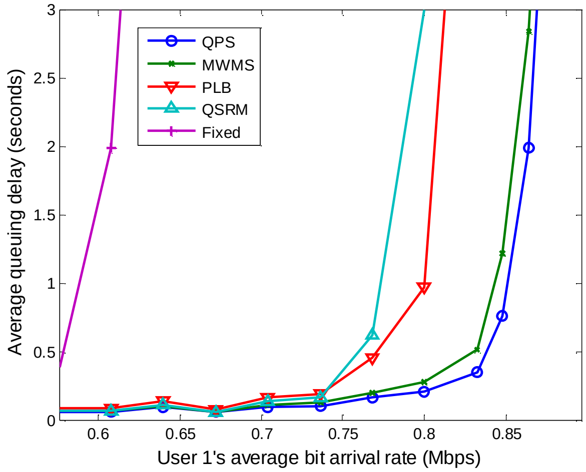

Fig. 14. Illustration of advantage of dynamic rate reassignment, exploiting the rate region of the situation in Fig. 3.

operates free of crosstalk from other differentially excited users. Thus, the trade-offs present at Level 1 and Level 2 between different lines' data rates return, as does a non-rectangular rate region looking more like Fig. 2 again with full binder vectoring.

C. *The time dimension*

A final area of improvement promised by any non-rectangular rate region such as that of Fig. 2 is in the area of cross-layer optimization of performance, particularly dynamic reassignment of data rates from one user to another at Levels 1, 2, or 3 (with full vectoring). One promising method for such reassignment and the possible gains is what is known as "queue-proportional scheduling," where each user's data rate capability is selected at a boundary point on the rate region such that each user's rate is proportional to the number of packets or bits in that user's queue, see [21]. Fig. 14 illustrates for the conditions of Fig. 3 the advantage of dynamic rate reassignment at every 10 ms (40 DMT symbols in ADSL or VDSL). That is, every 10 ms, the data rates are reassigned depending on the current queue depth where the packet arrivals are Poisson distributed and the length of packets is exponentially distributed with the average of 4 Kbytes.

Except for the fixed rate assignment labeled "Fixed," all the other methods plotted are dynamic and use various assumptions on the criteria for rate reassignment. Basically, one observes that for reasonable delay, approximately a 30-50% increase in data rate is available by dynamically exploiting the rate region with time. Compared to queue proportional scheduling, other dynamic reassignment methods provide slightly less improvement as described in [21].

IV. Conclusion

Dynamic Spectrum Management (DSM) is already in use while opening new vistas of heretofore unforeseen high data rates for DSL use. Unbundling does not prevent the use of DSM and indeed many Level 2 and Level 3 DSM gains are possible within any regulatory situation if the equipment and management are well conceived and designed to open standardized interfaces with service-provider management. Speeds to as high as 500 Mbps (35 MHz times x 15 bits/Hz) per user line become foreseeably feasible with vectored DSLs (and up to twice this rate or 1 Gbps using split-pair cross-overs on up to ½ the lines in a binder) at lengths of 300-400 meters. Such speeds within a cable of 200 lines add to 100 Gbps of possibly shared capacity for fiber-to-the-curb DSL systems, eliminating the need for fiber to traverse the last, most costly segment to each customer and generally well-exceeding the offered and projected speeds supplied by the much more expensive passive optical approaches. Thus, DSL capabilities, through the grace of DSM, may well extend throughout this new millennium.

Acknowledgements

The authors would like to acknowledge support from NSF FAST project, contract CNS-0427677; the project website is http://www.princeton.edu/fastcopper/. The authors would also like to thank Eva Berndt and Peter Muggenthaler, T-Systems, Darmstadt, for supporting in defining scenarios which could be relevant for the situation in Germany.

References

[1] J. M. Cioffi and M. Mohseni, "Dynamic Spectrum Mangement," *ISSLS 3/04*, Edinburgh, SCOTLAND.

[2] DSM Draft Report, *ANSI NIPP-NAI Contribution 2005-031R5,* October 2005, Las Vegas, NV.

[3] ITU-T Recommendation G.997.1, *Physical layer management for digital subscriber line (DSL) transceivers*, November 2005, Geneva

[4] ITU-T Rec. G.992.1, *Asymmetrical digital sub. line (ADSL) transceivers*, July 1999, Geneva.

[5] ITU-T Recommendation G.992.3 *Asymmetric digital subscriber line transceivers 2 (ADSL2),* January 2001, Geneva and ITU-T Recommendation G.992.5 *Asymmetric Digital Subscriber Line (ADSL) transceivers - Extended bandwidth ADSL2 (ADSL2+)*, January 2005,Geneva.

[6] ITU-T Rec. G.993.2 *Very high speed digital subscriber line 2,* February 2006, Geneva.

[7] DSL Forum Tech Report TR-069, *CPE WAN Management Protocol,* May 2004, Fremont, CA.

[8] H. Kerkskey, W. Lupton, A. Okiamianski (editors), DSL Forum Working Text *WT-131 ACS Northbound Interface Requirements* (Revision 2), September, 2005.

[9] R. Cendrillon, W. Yu, M. Moonen, J. Verlinden, T. Bostoen, "Optimal Spectrum Management," *ANSI T1E1.4 Working Group (DSL Access) Meeting*, contrib. 2004-459, North Carolina, May 2004. (Now part of the DSM Report T1E1.4/2003-018RE).

[10a] R. Lui and W. Yu, "Low complexity near optimal spectrum balancing for digital subscriber lines," *IEEE International Conf. on Communications. (ICC)*, Seoul, Korea, 2005.

[10b] R. Cendrillon, M. Moonen, "Iterative Spectrum Balancing for Digital Subscriber Lines," *IEEE International Communications Conference (ICC)*, Seoul, May, 2005.

[11] J. Papandriopoulos and J. Evans, "Low-Complexity Distributed Algorithms for Spectrum Balancing in Multi-User DSL Networks" (with minor corrections), *IEEE International Conference on Communications*, Istanbul, Turkey, to appear in June 2006.

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

[13] J. M. Cioffi, W. Rhee, M. Mohseni, and M.H. Brady, “Band Preference in Dynamic Spectrum Management,” *Eurasip Conference on Signal Processing*, Vienna, Austria, Sept 2004.

[14] T. Starr, M. Sorbara, J. M. Cioffi, and P. Silverman, *DSL Advances* (Chapter 11), Prentice-Hall: Upper Saddle River, NJ, 2003.

[15] M. Tsatsanis, “VDSL2 for Business Services: A new Frontier for DMT Technology,” 2005, http://www.convergedigest.com/blueprints/ttp03 /bp1.asp?ID=234&ctgy=Loop.

[16] G. Ginis and J. M. Cioffi, “Vectored transmission for digital subscriber line systems,” *IEEE J. Select. Areas. Comm*., Vol. 20, no. 5, pp. 1085-11-4, 2002.

[17] J. M. Cioffi, EE379C, EE479 class notes, http://www.stanford.edu/class/ ee379c/, http://eeclass.stanford.edu/ee479/

[18] B. Lee, J. M. Cioffi and M. Mohseni, “Gigabit DSL,” submitted to *IEEE Transactions on Comm*.

[19] J. M. Cioffi, B.Lee, M. Mohseni, A. Leshem, and Y. Li, “GDSL (Gigabit DSL),” *ANSI T1E1.4 contribution 2004-487R1*, August 2004, Washington, DC.

[20] J. M. Cioffi, C.Y. Chen, et al, “Update of Achievable Data Rates for Vectored VDSLs,” *ANSI Contribution NIPP-NAI-2006-040*, Las Vegas, NV: March 20, 2006.

[21] K. Seong, R. Narasimhan, and J. M. Cioffi, “Queue proportional scheduling via geometric programming in fading broadcast channels,” *IEEE J. Select. Areas Comm*., vol. 24, no. 8, Aug. 2006.

[22] J. M. Cioffi, S. Jagannathan, et al., “Full Binder Level 3 DSM Capacity and Vectored DSL Reinforcement,” *ANSI Contribution NIPP-NAI-2006-041*, Las Vegas, NV, March 20, 2006.

[23] V. Oksman and R. Stolle, “Proposal for NA VDSL2 band-plan for Profiles 17a and 30a,’’ *ANSI Contribution NIPP-NAI-2006-011R1*, Dana Point, CA, January 18, 2006.