BAND PREFERENCE IN DYNAMIC SPECTRUM MANAGEMENT

John M. Cioffi[1,2], *W. Rhee*[2], *M. Mohseni*[1,2], *and M.H. Brady*[1,2]

(1) Dept of Electrical Engineering, Stanford University 363 Packard EE Bldg., Stanford, CA 94305-9515 phone: +1 650-723-2150, fax: +1-650-724-3652, email: cioffi@stanford.edu (2) ASSIA-Inc., www.assia-inc.com .

ABSTRACT

The use of band preference in DSL systems with dynamic spectrum management (DSM) is investigated. Specifically, band preference is a DSM center's allocation of approximate PSD masks for subsequent use in the commonly encountered fixed-margin or "iterative" water-filling DMT DSL systems is found to improve data rates when DSL line lengths differ. The results of fully optimized spectrum management are shown to be achieved, but without the need for full central control of spectra.

1. INTRODUCTION

Optimal Spectrum Management (OSM) was introduced by and studied by Cendrillon, Moonen, Verlinden, and Bostoen et. al. in [1] and shows some nice DSM gains in Level-2-coordinated [2] spectrum management by a DSM center. In some cases, the OSM gains were significant over uncoordinated DSM (see [3],[4] for overview of DSM techniques), particularly when line lengths differ greatly in the same binder. Reference [1] concludes (and we agree) that the central coordination however renders the gains difficult to achieve in a practical system because of the need in [1] to control spectra centrally through a complex dual-decomposition optimization algorithm. This paper instead proposes use of the PSDMASK of ITU DSL standard G.997.1 [5] and observes that the usual discrete-integer water-filling can thereby obtain essentially the same performance in [1]. This paper elaborates on the early observations in [6].

Section 2 discusses band preference, which is an alternative to achieve the gains in OSM. Band preference is a loose emphasis of a band in the uncoordinated loading normally used by DSL modems in practice. Band preference maintains the usual bit-swapping and gain-swapping procedures that are so very necessary today in practical systems without loss in performance[1]. Band preference essentially tells the receiving modem to respect and observe a power-spectral density mask specified by the PSDMASK (or by the "tss" of G.992.3/5's "spectral toolbox" [7],[8]) to "load with preference" as in Section 2, allowing a priority to some bands in water-filling in general and preventing loading to levels that might interfere with other DSLs. Section 2 describes why an indication (called the "prefband" bit in the DSM report of [2]) is necessary to ensure that loading algorithms know how to observe simultaneous imposition of maximum margin constraints when band preference is used (and why claims the bit is not necessary are incorrect). Section 3 discusses some results from the use PSDMASKs to obtain with discrete water-filling the best results in [1].

Figure 1 illustrates the basic concept of a DSM Center that is presumed throughout (see, [2-5]). In such a system, typically data rates, suggested PSD bounds in perhaps a few bands, and maximum margin levels can be easily recommended to lines. These decisions are made depending on data reported to the center from the lines, like margins, powers, rates, and insertion-loss/noise measurements. Subsequently the lines continue to adapt somewhat autonomously with potential changes in rate or power spectrum levels suggested at infrequent intervals. Such DSM Centers then allow the DSL line to adaptive to any rapidly changing environmental situations (noises turning on or off). Central continuous coordination of the power spectrum level in detail is at least in present DSL systems not considered feasible, thus preventing a complex centralized control of spectra no matter how well it might perform. This paper investigates the use of the simple infrequent PSDMASK-level setting that is possible through the DSM Center.

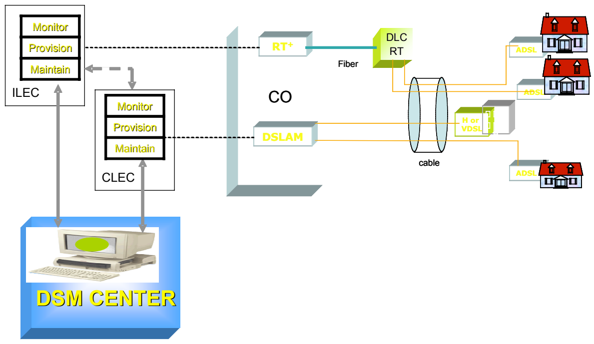

Figure 1 – DSM Center.

This paper applies to DMT ADSL and VDSL systems and presumes their structure as in [2],[7], and [8].

2. BAND PREFERENCE BASICS

The theoretical water-filling procedure solves the equations

---

[1] A centralized controller or DSM Center could not expected to be able to react and change bit distributions fast enough and certainly would be slower than the modems reactions themselves, which is one of the reasons that [1] concludes that water-filling or approximations done by modems in distributed fashion (sometimes called "iterative water-filling") is good enough.

$$E_n + \frac{\Gamma}{|H_n|^2} \cdot \sigma_n^2 = \text{constant} \qquad n=1,\ldots,\text{NSC} \qquad (1)$$

for non-negative energies E_n on each of the tones (NSC is the maximum number of tones). The gap Γ is a constant determined by code choices and desired margins at bit error rate 1e-7 in DSLs. The channel attenuation on each frequency is specified by $|H_n|^2$ and the noise energy on each tone is specified by σ_n^2, both of which are measured (or their ratio measured directly) during training and updated during "showtime" operation of DMT DSL modems. This procedure is viewed as running continuously in time with updates at periodic or on-channel-change intervals. The updating is distributed or autonomous in that each modem individually executes the procedure. Fixed-margin water-filling allows only the minimum sufficient energy to be allocated in water-filling so that the specified MAXSNRM (maximum margin) is achieved.

This theoretical water-filling procedure is well-known in DSL and can be approximated in a number of ways including various greedy algorithms[2] for discrete integer bit restrictions where successive bits are loaded in the least energy-consumptive bit positions on all the tones until the desired max bit rate limit has been attained with no more than a maximum specified margin (often known as MAXSNRM in various DSL standards) and no less than a minimum margin (often known as TARSNRM or TSNRM in various DSL standards). In greedy algorithms, tones that have already been loading to a maximum bit cap have infinite (large) cost to add additional bits, thus preventing bits/tone in excess of the bit-cap. This paper observes that tones for which the addition of a bit would cause the PSDMASK bound specified to be exceeded should also have infinite (large) cost in those positions, which may or may not be implemented by various manufacturers. Changes in channel and or noise are monitored and the algorithms are run continuously allowing the DMT transmission facility to move bits so that good energy use and margin is maintained. The infinite cost associated with exceeding an imposed spectral mask is maintained in showtime operation so that bits do not reallocate to a PSDMASK-limited band even if that band would be more attractive than other bands in theoretical water-filling.

The PSDMASK's imposition of essentially an infinite cost to adding bits on a particular tone when the existing energy of that tone is already at/near the mask level is used in the proposed **band preference**. Essentially, it is too hard to add a bit in that band because of the PSDMASK constraint, thus forcing the discrete water-filling algorithm (that is the greedy procedure in the previous paragraph) to place bits in another "preferred" band. The PSDMASK may have thus been set well below allowed masks in an effort to prefer the use of other bands, presumably because the DSM Center has determined such band preference is valuable to the lines. A central control of the bit distribution would likely be impractical (even if regulatory positions might sometimes allow it) because of the speed of response in changing the bit distribution as needed or as appropriate for time-varying channel effects (like crosstalk changes, etc.). Band preference is instead specified at time of initialization by the DSM Center, presumably through the wise choice of the PSDMASK levels (or alternatively possibly through tss_n levels) that are already used in advancing ADSL systems. One could also implement band preference through a frequency-dependent "bit-cap" (maximum number of bits per tone) in loading.

The energy on a particular tone, E_n, is determined by 3 components

$$E_n = E_{0,n} \cdot g_n^2 \cdot tss_n^2 \quad (2)$$

where $E_{0,n}$ is a nominal initial transmit energy level that is determined by standardized fixed or standardized MIB-specified (i.e., PSDMASK) spectral masks possibly offset by flat "power-back-off" amounts on those tones that are allowed for use (the tones allowed for use are often specified by the CARMASK parameters of various DSL standards). Those familiar with standards would call $E_{0,n}$ the REFPSD. For instance, an ADSL modem with no power back off and no use of PSDMASK would have an $E_{0,n}$ that corresponds to -40 dBm/Hz (or some other value determined in the various annexes of standards) that would be known to both transmitter and receiver. The quantity g_n^2 specifies a receiver-computed gain that is usually between -14.5 dB and +2.5 dB for ADSL standards and in theory could be any non-negative (linear) value. This "gain" is passed to the transmitter through a reverse control channel in DMT DSL either at the exchange of initialization or during "bit-swapping" in live operation of the modem. The tss_n^2 parameter is fixed for any use of the modem by the MIB and can be between 0 and 1. One immediately notes with theoretical water-filling, the *tss* is almost useless as the gain parameter could undo any *tss* effect and set the energy levels to the desired water-filling levels (if gains g_n were not upper bounded). Of course, a *tss*=0 value would prevent use of the tone and could not be inverted. In practice the upper bounding of the possible gain choices does allow the *tss* to impact the loading algorithm's limits and is thus a useful tool. In particular an upper limit of gain at +2.5 dB prevents significant inversion (if this limit were raised for instance in ADSL2 by what is called EXTGI[3], then more inversion is possible). Nonetheless, the unmodified water-filling procedure might indeed try to restore a band by a positive value of gains when tss is low.

[2] These are often called Levin-Campello procedures although others lay claim to the title also, see [4],[9] for instance.

[3] EXTGI is a transmitting-modem-supplied parameter (to receiver) that tells how much additional positive gain is possible so that gains can be as high as 2.5+EXTGI where $0 \leq$ EXTGI<25.6 dB. EXTGI specifically may be zero or selected to be zero by a modem manufacturer, but is not within the control of the service provider.

The situation of the $E_0 + 2.5$ dB being the maximum in (1) corresponds directly to an infinite (large in finite precision) cost associated with adding an additional bit to that tone in discrete loading. This nonlinearity of discrete loading is important to band preference.

In a system where PSDMASKs have been set intentionally lower in an otherwise good band, the infinite cost associated with loading more bits beyond those levels in that band forces loading into other available bands of transmission not yet at their PSDMASK levels, essentially then preferring those other bands. For instance, in the ADSL CO/RT mix situations of Section 2.2 of [1] and VDSL upstream example of Section 2.3 of [1], theoretical water-filling is ineffective because it attempts to continue to load into the lower frequency band that looks more attractive (but then limited by the crosstalk created into the 2nd user on the longer line). For those same situations, *if the receiver also knows the PSDMASK setting*, discrete water-filling would sense the infinite cost of exceeding the PSDMASK if it had been set at the appropriately low level to avoid crosstalk into the longer line. Thus, the discrete water-filling would then begin adding bits to the 2nd higher-frequency band on the shorter line in both examples and get the same results as OSM. Figure 2 illustrates this simple effect.

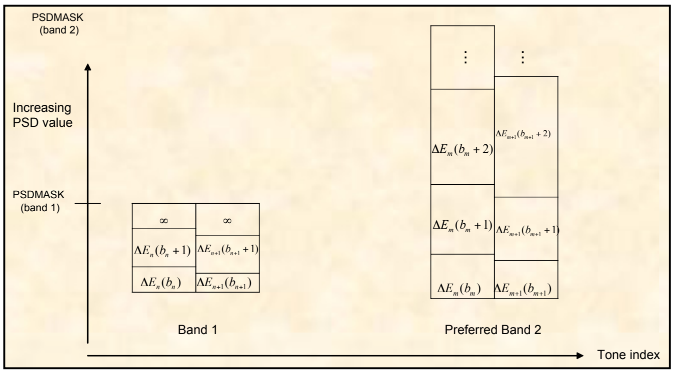

Figure 2 – Illustration of the preferred band 2 that has a much higher PSDMASK limit, even though the incre mental energies for additional bits (indicated here by the height of a box) are smaller up to the point where they become infinite.

Again, to implement this band-preference with discrete water-filling by computing at which bit-loadings to place infinite cost in the discrete water-filling algorithm, the receiving modem must know the PSDMASK settings. The emerging DSM standard [2] has a one-bit PREFBAND indicator be sent to the receiving modem as a MIB parameter. Some think this bit is unnecessary because a careful reading of the ADSL2 standards will find a footnote several pages beyond the PSDMASK section that states that the EXTGI parameter cannot be exploited to the extent that a PSDMASK is violated. This is true for the careful compliant designer. However, there is also a MAXSNRM (maximum SNR margin) parameter. Margin in modems with discrete water-filling is defined as the worst margin on any of the tones (continuous formulae otherwise are difficult to use). The authors' experience has shown that many loading methods of some vendors modems will -- if a band is de-emphasized as in band preference (for instance, the lower band on the left in Figure 2) – have a low margin on tones in this band that is below MAXSNRM. Thus, the MAXSNRM is satisfied even if the preferred band (tones on the right in Figure 2) have margins well in excess of MAXSNRM. This unhappy situation does satisfy the letter of the ADSL2 standards and is acceptable and unfortunately practiced heavily. However, a modem directed by bandpref=1 (meaning band preference is on) is being told by the service provider to not only observe the PSDMASK but also to hold the largest (not the smallest) of the margins in any band or on any tone to be less than MAXSNRM. While good loading practice could lessen the need for the bandpref bit, no specific loading algorithm is mandated in standards. Thus, an interpretation of band lowering that it is for band preference (and not for some other purpose where high margin in preferred bands would not be offensive) is indeed necessary, and the reason the DSM Report in the USA has mandated this extra bit in the MIB.

3. SOME RESULTS

This section illustrates how the PREFBAND option can be used to obtain results such as those in [1] using OSM.

As an illustration, we repeat the tables from [1] with the use of the PREFBAND indication and discrete water-filling with infinite costs here. The first situation is similar to that shown in Figure 1 where an RT and CO DSLAM share the same binder corresponding to perhaps a newer DSL customer served from the RT and an older one already in existence from before the RT installation. The new DSL can dramatically decrease through crosstalk noise the data rate of the older CO-based DSL. In this simple example of the old customer at 4 km and the new customer, the customer would drop from 1 Mbps speed to less than 300 kbps without any DSM or fixed-margin (iterative) water-filling as is noted in several other places (see the references of [3] for example). Water-filling. Theoretical water-filling preserves both customers at 1 Mbps. However, the newer customer can actually obtain a much higher data rate with either OSM or the preferred-band discrete-water-filling implementations.

Table 1 – CO/RT ADSL mix from [1]

	4 km data rate (Mbps)	3 km data rate (Mbps)
Theoretical water-filling	1.0	1.0
Discrete water-fill PREBAND=1	1.0 Mbps	3.3 Mbps
OSM	1.0 Mbps	3.3-3.7 Mbps

The 2nd row of Table 1 has PSDs of -40 dBm/Hz to 250 kHz and -60 dBm/Hz 250 kHz to 450 kHz, -40 dBm/Hz 500 to 1.108 MHz on the 3km loop and -40 on 4 km loop. Table 2 illustrates another example for upstream VDSL where essentially the same band preference equivalence to OSM is noted. Figure 3 illustrates the upstream VDSL rate region increase from the use of band preference for two lines of lengths 600 meters and 900 meters in upstream VDSL for “998-Plan” [1] iterative water-filling with noise A ([1]) and for water-filling using band preference with PSDMASK levels of -72 dBm/Hz below 5 MHz and -55 dBm/Hz above 5 MHz on the 600 meter line, and -50 dBm/Hz on the 900 meter line. The data rate on the shorter line is dramatically improved with band preference.

To the extent that Noise A is based on static ADSL models, Figure 3 represents worst-case performance because the ADSLs then are not modelled correctly (they water-fill also, so model “A” noise in standards is grossly incorrect, but used anyway here). Note that 6 Mbps on the 900m loop up-stream can be achieved while nearly 20 Mbps upstream oc-curs on the 600m loop. These rates are considerably higher than what would be achieved if fixed spectra were used on the VDSL lines and more than double what are achievable with fixed spectra today.

Table 2 – Upstream VDSL from [1]

	900m data rate (Mbps)	600m data rate (Mbps)
Theoretical water-filling	6.0	4.5
Discrete water-fill PREBAND=1	6.0	14.0
OSM	6.0 Mbps	14.0 Mbps

The 2nd row of Table 2 has PSD’s of -72 dBm/Hz to 5.5 MHz and -55 dBm/Hz above 8 MHz on the 600m loop and -50 on 900m loop.

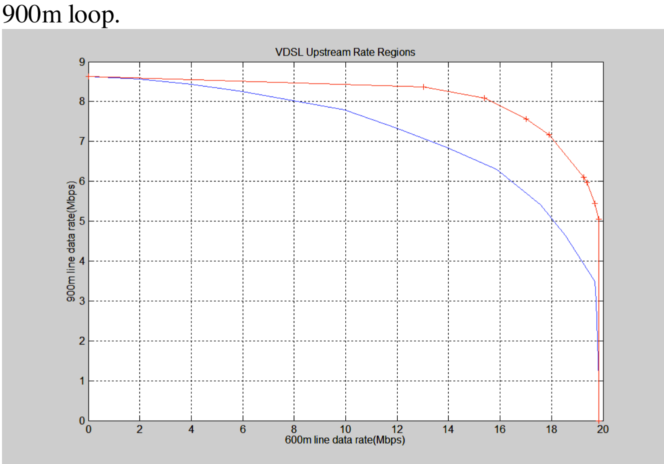

Figure 3 – Upstream VDSL rate region for band prefer-ence (upper curve has band preference, lower has not) – 4 VDSL lines at 600m and 4 VDSL lines at 900m.

Figure 4 illustrates that band preference is provides most of its gain when mutually crosstalking loops have very different lengths. As the lengths approach the same, no band prefer-ence is necessary and all iterative-water-filling loops can use the same PSD mask levels. A single 900m or longer loop was held at 10 Mbps (80% of its maximum rate when no other loop is present) while a single other loop was varied in length between 600m and 900m. The vertical axis plots the fraction of the maximum short-loop rate that is achieved with and without band preference.

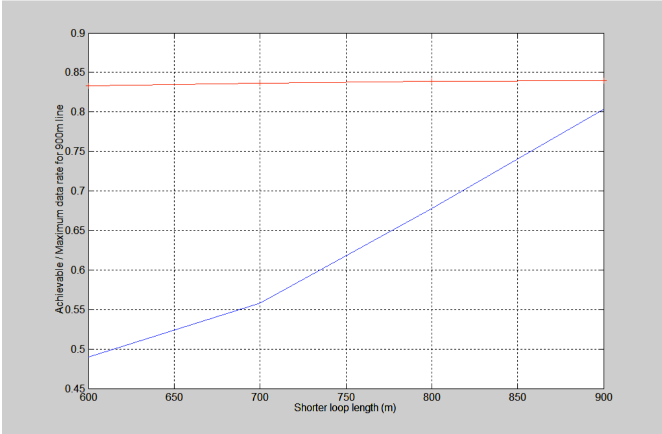

Figure 4 – Band prefererence (upper curve) versus no preference when the shorter loop length is varied.

4. CONCLUSION

Band preference is a practical method for achieving highest performance in mixed binder DSLs with a minimum of co-ordination between the lines. Such band preference is feasi-ble with existing DSL standards if the full suite of mainte-nance parameters available in those standards are used well by a DSM Center. Band preference assists DSM systems to offer DSL customers with the highest mutually compatible data rates and represents a very large gain over existing static spectrum management systems.

REFERENCES

[1] R. Cendrillon, M. Moonen, J. Verlinden, T. Bostoen, “On the optimality of iterative water-filling,” T1E1.4 Contribution 2003-325, December 8, 2003, San Diego, CA.

[2] DSM Draft Report, *ANSI T1E1.4/2003-018R9,* De-cember 8, 2003, San Diego, CA.

[3] J. Cioffi and M. Mohseni, Dynamic Spectrum Man-agement, ISSLS 2004, Edinburgh, Scotland, March 2004.

[4] T. Starr, M. Sorbara, J. Cioffi, and P. Silverman, DSL Advances, Prentice-Hall, 2003, Chapter 11.

[5] *PLOAM Standard G.997.1,* International Telecom-munications Union, 2002.

[6] J. Cioffi and M. Mohseni, et al., “Band Preference with Water-Filling in DSM,” *ANSI Contribution T1E1.4/2003-321R1,* December 8, 2003, San Diego, CA 2003.

[7] *ADSL2 Standard G.992.3,* International Tele-communications Union, 2003.

[8] *ADSL2+ Standard G.992.5,* International Tele-communications Union, 2003.

[9] T. Starr, J. Cioffi, and P. Silverman, Understanding DSL , Prentice-Hall, 1999.