ATIS Committee NIPP-NAI
Working Group NAI (DSL Access) **NIPP-NAI-2007-129R1**
Miami, FL: September 24, 2007

CONTRIBUTION

TITLE: **The Inherent Simplicity of Distributed Band Preference**

SOURCE:

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

PROJECT: **T1E1.4, DSM, Information only**

ABSTRACT

This new contribution investigates the simplicity of Distributed Band Preference [1]. Distributed Band Preference is found to be very simple in terms of the implementation burden upon any component. The component implementation supports two loading algorithms that can share the same memory and signal-processing implementation. Only one is used at any given time, depending upon whether the optional maximum-SNR-margin-mode bit is set or not set (the default). The simplicity comes at very little (most cases none) loss with respect to the best Level 2 DSM performance. Simultaneously, the modems themselves retain a full ability to adapt to time-varying conditions, thus avoiding a latency associated with a more centralized Level 2 control.

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

The Inherent Simplicity of Distributed Band Preference (129R1)

J. Cioffi, W. Lee, S. Jagannathan

Dept of EE, Stanford University
363 Packard EE, Stanford, CA 94305-9515
Cioffi@stanford.edu,
P: +1-650-723-2150 ; Fax: +1-650-724-3652

G. Ginis

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

ABSTRACT

This new contribution investigates the simplicity of Distributed Band Preference [1]. Distributed Band Preference is found to be very simple in terms of the implementation burden upon any component. The component implementation supports two loading algorithms that can share the same memory and signal-processing implementation. Only one is used at any given time, depending upon whether the optional maximum-SNR-margin-mode bit is set or not set (the default). The simplicity comes at very little (most cases none) loss with respect to the best Level 2 DSM performance. Simultaneously, the modems themselves retain a full ability to adapt to time-varying conditions, thus avoiding a latency associated with a more centralized Level 2 control.

1. Introduction

Level 2 DSM's spectrum balancing can effect large data rate gains through a Spectrum Management Center's (SMC's) central imposition of politeness upon users that otherwise transmit too much power in their victim's essential spectrum band. Distributed Band Preference [1] has been suggested as fully distributed system that uses only the control of a single bit (maximum SNR-margin-mode) to achieve the full practical gain of Level 2 DSM, essentially achieving the theoretical bound on Level 2 performance in [2]. This present contribution attempts to answer the question

> ***What implementation burden does Distributed Band-Preference place on the DSL component to achieve this gain?"***

The answer is very little[1]. This contribution reviews the gains and then studies the complexity imposed upon a (presumably chip) implementation of the level-2-DSM-capable DSL modem. A short Section 2 reviews basics, while Section 3 attends to the simplicity of implementation.

2. Band-Preference Basics

Figure 1 illustrates the basic band preference Level 2 DSM concept. The SMC marks users/customers as either "strong" (Margin-cap = ON) or "weak" (Margin-cap = OFF, the default). In Figure 1, Maximum-SNR-margin-mode is simplified to "Margin-cap." Figure 1 shows an upstream-VDSL example of a short line interfering with a long line. If both lines have about the same noise, user 1 is a strong user. User 2 is instead a weak user and may need some help from user 2 in the form of increased politeness. This is essentially a preference for user 1 to use higher frequencies over lower frequencies, leaving user 1 to transmit unencumbered at the lower frequencies, whence the name "band preference."

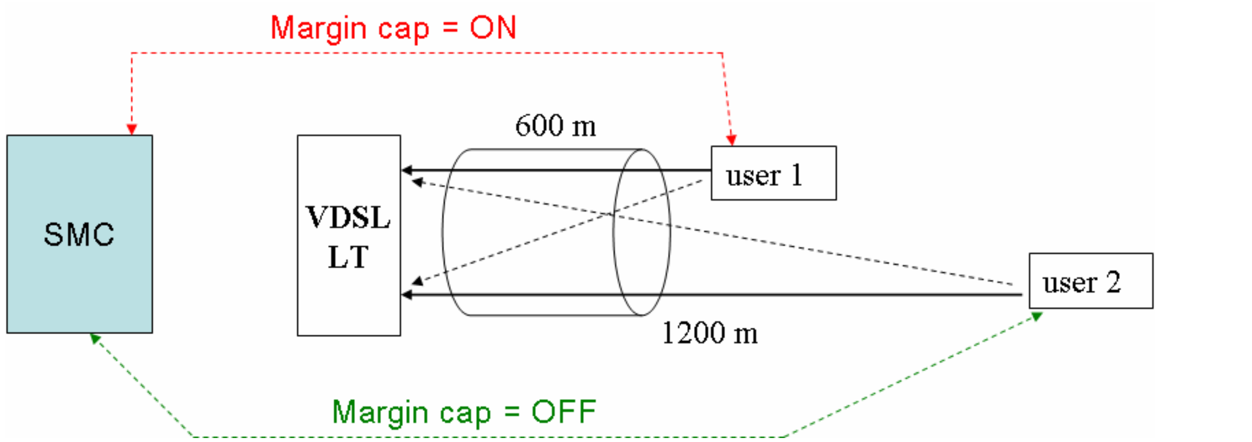

Figure 1 – Illustration of SMC setting Maximum-SNR-Margin-Mode (abbreviated "Margin-cap") on for extra politeness from user 1.

Various studies (referenced in [1]) have found such spectrum balancing to provide a large gain for both users particularly in situations of the strong/weak (or "near/far") crosstalk. Figure 2 illustrates band-preference at work for the situation of Figure 1. There, as is the case for practical situations of interest in Level 2 DSM, band preference gets very, very close to the theoretical optimum performance.

The difference between a weak and a strong user depends on the noise and the line length/configuration. An SMC assesses the reported noise and line information and decides whether to request ultra-polite behavior (that is set the margin-cap=ON) or to leave the modem operating in the usual manner (which is less polite and based on getting maximum performance without regard to radiated crosstalk into other users). Various "power back-off" methods often make erroneous assumptions about reference lengths or noises to which the strong user must train instead of allowing the modems to adapt. Power back-off in the field has been largely thus ineffective and often disabled. Band preference allows the SMC to decide for a situation where essentially any power reduction or skewing towards less favorable frequencies by a strong user would be beneficial. The strong modem then decides how polite to be (without sacrificing its own data rate) according to the algorithm in Section 3.

[1] While a number of proposals have been made for new MIB elements and controls by various parties, including many unrelated to the authors of this contribution.

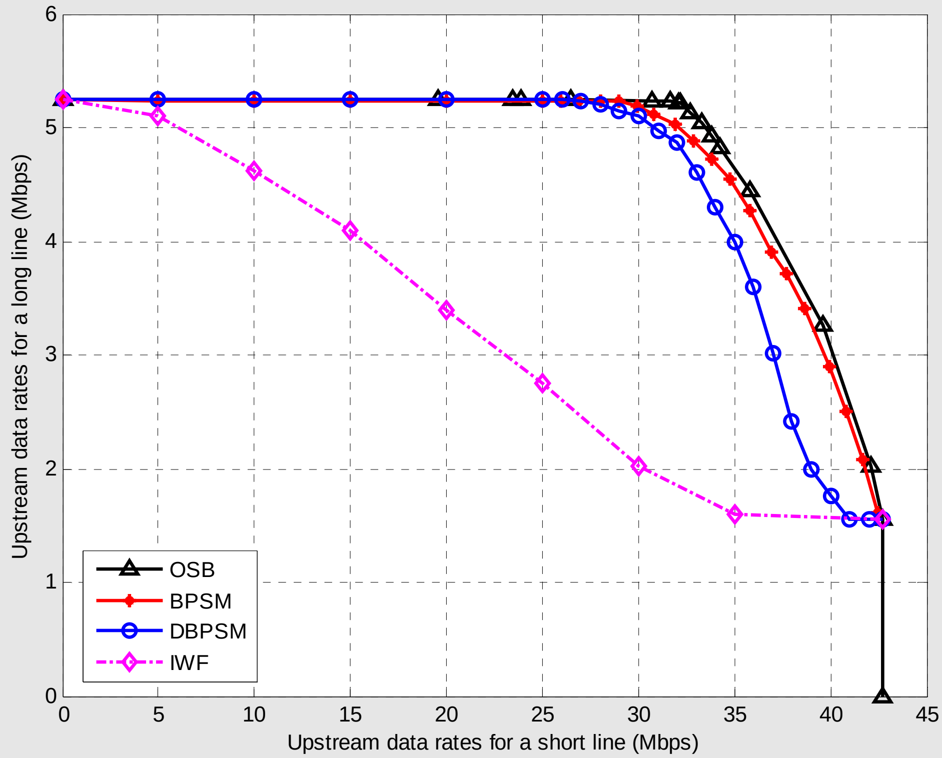

Figure 2 – Performance of various methods for the situation of Figure 1.

3. The Strong and Weak Loading Algorithms

There are two algorithms described here, which are almost identical in terms of complexity. Only one runs in any given modem according to the margin-cap bit setting.

3.1 The Weak-User Algorithm

The basic water-filling or "weak-user" algorithm is:

$$Power(n) = \max\left\{0, K - \frac{\Gamma \cdot Noise(n)}{|Channel(n)|^2}\right\} \quad , \qquad (1)$$

where the power is for all tones n such that it is either zero or the difference between a frequency-independent constant K and the scaled (by the gap Γ) channel noise-to-signal ratio. This method is best for single-user DSL and well-known throughout the DSL industry. There are various practical approximations to it, but the modems' receivers compute a basic spectrum that looks close to Equation (1) when margin-cap is OFF. The constant *K* is the "water level."

Figure 3 illustrates a so-called "Greedy" or "Levin-Campello" algorithm used to approximate water-filling (in fact this algorithm is optimum for the situation where an integer number of bits must be assigned to each tone, see [3]). A table of energies is maintained for each tone (typically this is a single column of energies in memory that is scaled by a factor related to the inverse gain-to-noise ratio on each tone to generate efficiently a table column for each tone in implementation, thus saving memory at the expense of a single multiply). This table contains the incremental energy required to add each successive bit to the tone. Bits are assigned to tones in a greedy manner by placing the next bit to be loaded on the tone requiring least incremental energy to carry that bit. It is a very simple and widely used algorithm. Bit swapping usually continuously compares the tone currently requiring the most energy for its highest-order bit transmitted to all next-incremental energies on other tones. If line or noise changes make another incremental energy less (more attractive), the bit-swapping control moves (swaps) the bit in to the lower energy position.

The weak algorithm thus tends to send bits in the best (often lower frequency) tones first.

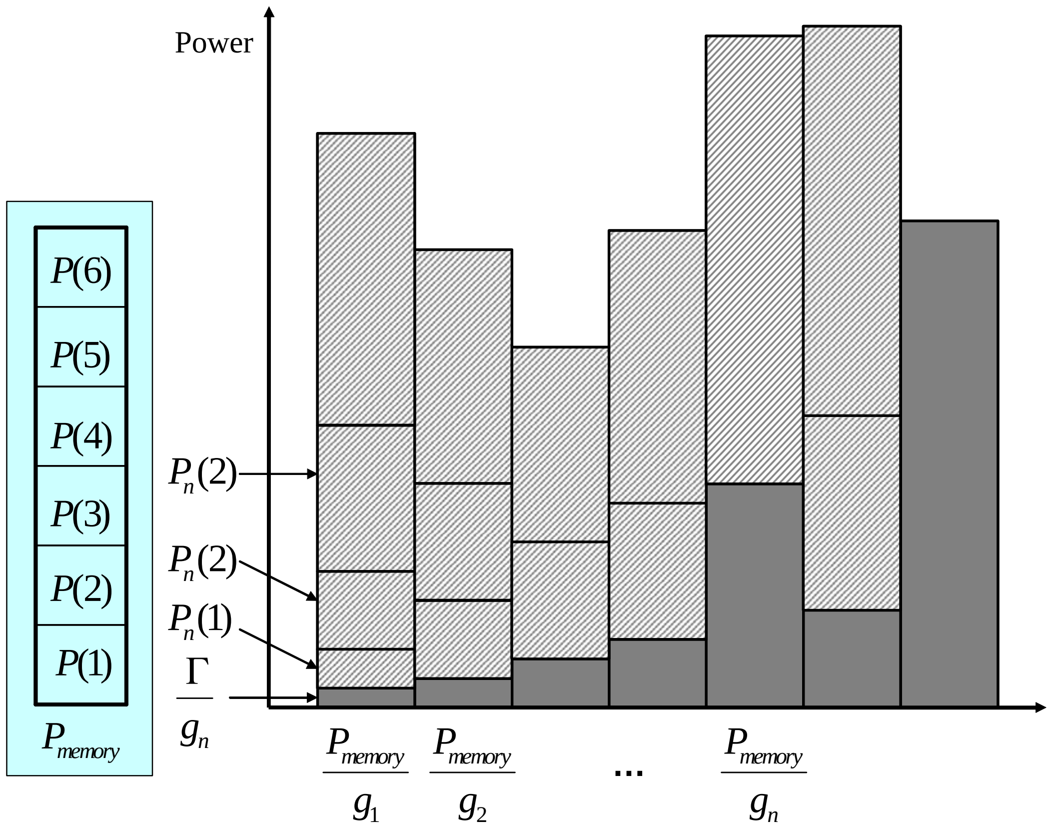

Figure 3 – Weak loading Algorithm's use of bit-loading tables (sometimes called the "Levin-Campello" method, see [3]).

3.2 The Strong-User Algorithm

If the weak-user algorithm can be described as "greedy" by giving to the rich tones and not to the poor tones, the strong algorithm might be described as "Robin Hood", i.e., robbing the rich to give to the poor. The strong-user algorithm initially runs a water-filling algorithm, thus same complexity so far. However, when the strong algorithm is completed, the modem looks at evacuating the better tones one-by-one, by giving to poor tones as long as it is possible for a poor tone to carry an extra bit below the given power-spectral density and total power constraints as described in more detail below. Figure 4 also contrasts the spectra produced by the weak and strong algorithms.

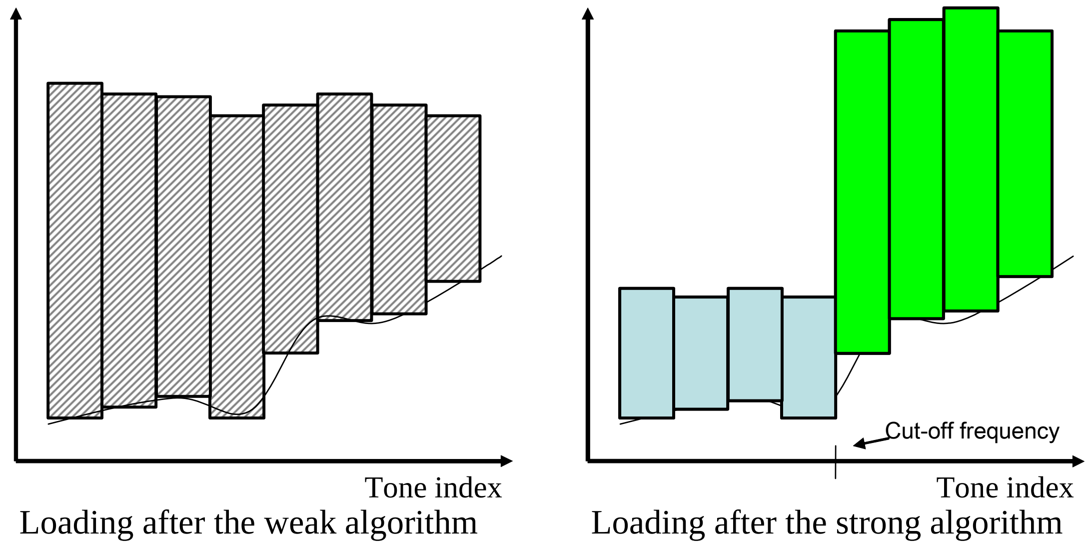

Figure 4 – Comparison of the Weak-User and Strong-User Loading Algorithms' consequent spectra for a near/far (short/long) upstream VDSL situation.

Thus, the steps are

1. Run the greedy algorithm
2. two bands of tones (higher-SNR and lower-SNR) are determined such that by removing bits from the higher-SNR band and placing them into the positions of least incremental energy on the tones of the lower-SNR band, water-filling is approximately and INDEPENDENTLY maintained in both bands.
3. Repeat Step 2 until no more bits can be removed from the high-SNR band and placed in the low-SNR band without violating a power-spectral-density or power constraint

It is well-known that water-filling on practical channels often leads to a nearly flat energy/power-spectral-density distribution in the used band (that is an ON/OFF band with the ON and OFF tones carefully selected will perform nearly identical to water-filling in practice). Essentially then there are two flat energy bands after completion of Steps 2 and 3.

One immediately notes that if noises were equal on the strong and weak users of Figure 1, the transmitted energy of the short-line (strong) users would move to higher-frequency tones where they would not affect the long-line (weak) users as much. The strong users give as much as they can in terms of frequency use and power. If conditions change, that strong users continue to maintain the best "giving" position they can while still maintaining their data rate (and margin, say 6 dB typically) and satisfying Step 2 above. Thus all dynamics are retained, even if the strong user eventually returns to the water-filling distribution because that is the only one that can achieve the given data rate and margin. Therefore, the "strong-user" algorithm does not result in inferior performance compared to basic water-filling when the noise varies. The two bands of higher-SNR and lower-SNR tones would be identified by a single "cut-off" frequency between them. This cut-off frequency is determined to maximize the separation between the flat energy/power-spectral-density in each of the bands. This maximum (and the consequent cut-off frequency) is determined by swapping bits from the higher-SNR band's positions of most incremental energy reduction to the low-SNR band's positions of least incremental additional energy.

There are other methods for the 2nd loading algorithm that may be of similar good performance so it is not necessarily unique.

The complexity of the "strong-user" algorithm is the same as the usual bit-swapping and the same energy table as in the greedy algorithm is used. Just the criterion for moving bits prohibits use of the best tones if they can be vacated. Steps 2 and 3 of the "strong-user" algorithm should not be confused with communicating bit-swaps with the DSLAM. The bits and gains are conveyed to the DSLAM only after the entire loading is completed. The SMC plays a crucial role in determining whether a user should use the "strong-user" algorithm or continue with the weak-user algorithm. An SMC/service-provider should not demand the strong algorithm unless there is a strong possibility of overall binder improvement. The decision to use weak or strong may demand sophisticated signal processing, but that processing is not in the modem chips or DSLAMs.

4. References

[1] J.Cioffi, W. Lee and M. Brady, "DSL Level 2 Distributed Band Preference," *ATIS Contribution NIPP-NAI 2007-056R1*,Minearpolis, MN, May 14, 2007.

[2] 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).

[3] T. Starr, M. Sorbara, J. Cioffi and Peter Silverman, <u>DSL Advances,</u> Prentice Hall, 2003.