Project: **T1E1.4: Dynamic Spectrum Management**

Title: **Iterative Water-filling for DSM (200R4)**

Contact: **J.M. Cioffi, W. Yu, S.T. Chung, and G. Ginis**

Dept of EE, Stanford University
363 Packard EE, 350 Serra Mall
Stanford, CA 94305-9515
cioffi@stanford.edu , 1-650-723-2150, F: 1-650-724-3652

Date: **November 5, 2001 (Greensboro, NC)**

Dist'n: **T1E1.4**

Abstract:

This contribution introduces and investigates "iterative waterfilling" for the situations of no coordination between DSL lines (autonomous DSM or Level 0 coordination) and of Level 1 coordination. In the case of no coordination, DMT modems are shown to autonomously implement dynamic spectra that converge to a globally good solution for each and every line that has performance often well in excess of that projected by static-spectrum-management (T1.417) calculations. For Level 1 coordination, the performance improves further in that service rate can be traded among customers to allow a greater possibility of functioning service.

A reference document [1] is appended to this contribution that details the mathematics behind iterative water-filling for the very interested reader. The early part of this contribution simply summarizes results and attempts to motivate interest in iterative water-filling.

NOTICE

This contribution has been prepared to assist Standards Committee T1 - Telecommunications. This document is offered to the Committee as a basis for discussion and is not a binding on any of the companies listed as authors. The requirements are subject to change after further study. The authors specifically reserve the right to add to, amend, or withdraw the statements contained herein.

Iterative Water-filling for DSM (200R4)

J.M. Cioffi, W. Yu, S.T. Chung, and G. Ginis
Department of Electrical Engineering
Stanford University, Stanford, CA 94305
Phone: 650-723-2525 Fax: 650-724-3652

Abstract:

This contribution introduces and investigates "iterative waterfilling" for the situations of no coordination between DSL lines (autonomous DSM or Level 0 coordination) and of Level 1 coordination. In the case of no coordination, DMT modems are shown to autonomously implement dynamic spectra that converge to a globally good solution for each and every line that has performance often well in excess of that projected by static-spectrum-management (T1.417) calculations. For Level 1 coordination, the performance improves further in that service rate can be traded among customers to allow a greater possibility of functioning service.

A reference document [1] is appended to this contribution that details the mathematics behind iterative water-filling for the very interested reader. The early part of this contribution simply summarizes results and attempts to motivate interest in iterative water-filling.

1. Introduction

A surprising result in the optimization of a multi-user DSL channel is that a simple extension of the spectrum-optimizing "water-filling" procedure for a single-user channel provides a nearly globally optimum sharing of spectra between the all the lines in a binder. This procedure is called "iterative water-filling." Surprisingly iterative water-filling is already approximated serendipitously by some standardized DMT modems. This contribution investigates the performance improvements of iterative water-filling under two situations:

1. No coordination (Level 0 coordination) between DSL lines
2. Level 1 coordination (sharing data rate and power levels at a network maintenance center).

In both cases, the improvements are large – allowing performance well beyond that anticipated for DSL or previously computed according to the methods of T1.417 (static spectrum management).

Section 2 reviews binder characterization with rate regions and line coordination in DSM. Section 3 overviews the basic concept of iterative waterfilling and discusses Level Zero and Level One coordination. Section 4 then analyzes three examples for Level 0 and Level 1 coordination with iterative water-filling: (1) a solution to the RT-located-ADSL FEXT into CO-located ADSL problem, (2) a large improvement in the well-known VDSL-FEXT into ADSL crosstalk problem and (3) a large improvement in the allowable mixture of asymmetric and symmetric VDSL services within the same binder. An attachment [1] provides detailed mathematical description of iterative water-filling.

2. Rate Regions and Binder Characterization

Figure 1 illustrates the basic crosstalking binder. Consistent with present DSL standards, this paper maintains the presumption that DSL systems in the same binder have upstream and downstream transmissions separated (at least at frequencies above 200 kHz) by frequency-division multiplexing. Thus, the dominant crosstalk in the binder is FEXT in all the examples. The FEXT can have a variety of power-

transfer models, but this paper uses exactly those standardized in the T1.417 standard. Figure 2 illustrates the **rate region** that might occur for two active DSLs in the same binder: a short line and a long line in the (for example) downstream direction. The short line's data rate is on the horizontal axis, while the long line's data rate is on the vertical axis. Any combination of data rates in the shaded region of Figure 2 is achievable. The data rates for the shorter line are generally higher than those for the long line, so the region is not symmetric. Each point in this achievable rate region is characterized by different spectral pairs for the two individual DSLs. The two DSLs can jointly achieve any point in the region, but each such point has its own spectra for each line. For 3 or more DSLs, the concept of the rate region easily generalizes, but is not easily plotted. Thus, a multidimensional region could be computed for all the possible *L*-tuples of data rates that are achievable on *L* active DSL lines in a binder. Each such achievable point in the rate region would correspond to a distinct *L*-tuple of spectra for the lines. Each different binder situation has a different rate region. Static spectrum management tries to find a single spectrum for each type of DSL, which is equivalent to finding a single L-tuple of rates that is common to all possible rate regions for all possible binders. Dynamic spectrum management need only find the corresponding rate region for each binder situation, and then select a desirable operation point within that rate region.

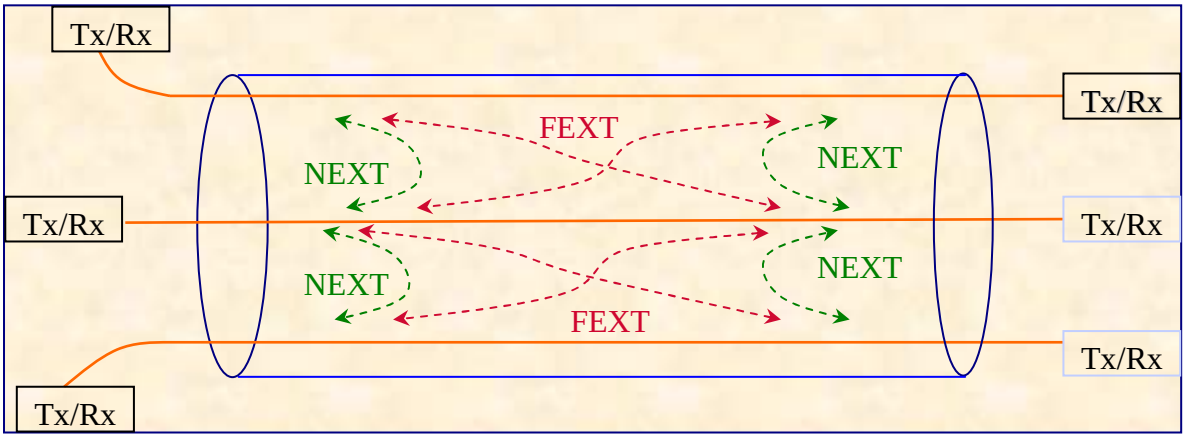


Figure 1 – Illustration of crosstalk between lines in a binder.

Initially, one suspects that the various choices of spectral combinations for the lines would require some level of coordination between the lines, perhaps at a central maintenance facility. Actually, as long as the *L*-tuple of attempted data rates on the lines is within the rate region for a binder, no coordination is necessary. In fact, the optimization can be achieved autonomously (Level 0 coordination) by certain types of normally used spectrum procedures (i.e., bit-swapping in ADSL). Coordination can help identify (and thus repair) the situation when the *L*-tuple attempted (or to be attempted) is not in the achievable rate region.

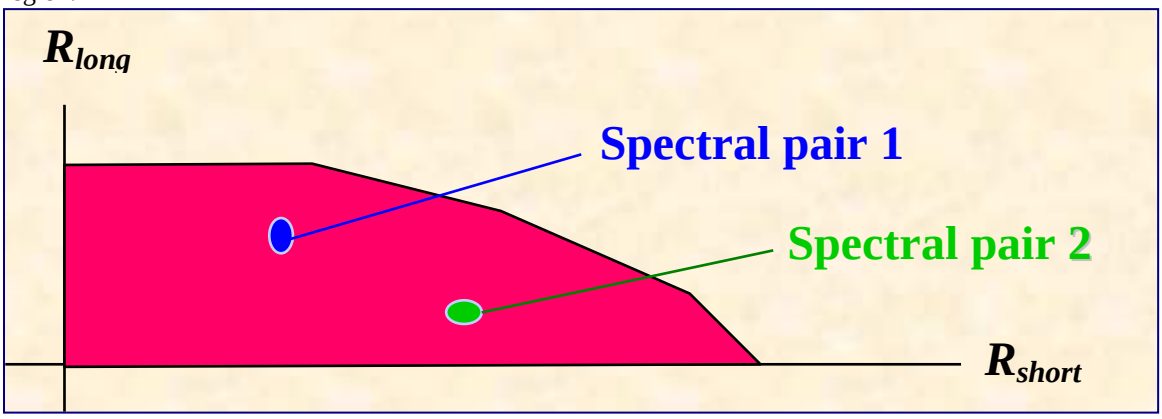


Figure 2 – Illustration of data rate region for two active DSLs, one short and one long

DSM mission statements and early drafts [2] have identified 3 levels of coordination, 0, 1, and 2. Level 0 means no coordination, an existing state of the art in many DSL deployments. Level 1 allows line-characterization data (typically the data rate, used power level, corresponding margin/performance level, DSL type, and SNR, possibly frequency dependent) to be shared for fault isolation and possibly line-spectra control purposes. Level 2 allows co-generation of signals for different DSL pairs and is not investigated in this contribution. This contribution investigates rate regions that would be possible using Level 0 and Level 1 coordination (actually these rate regions are the same for both Level 0 and Level 1, with Level 1 allowing a greater flexibility in choice of point within the rate region for a given group of binder-sharing DSLs).

3. Iterative Water-filling Procedure

Iterative water-filling makes use of the basic water-filling procedure [3] reviewed by Figure 3. This particular procedure provides the optimum spectrum for transmission on any channel characterized by a frequency-dependent SNR, $SNR(f)$. Conceptually, inverse of the SNR, $(NSR(f))$ is plotted (the yellow curve in Figure 3). A given amount of power is available leading to the energy (shown in red) distribution that occurs as if the energy were "water" poured into the $NSR(f)$ curve, which will naturally fill to a flat level. A number of practical procedures (see [1]) can be used to approximate the implementation of water-filling in practice, including the so-called Levin-Campello algorithm used in DMT ADSL and VDSL bit-swapping [3]. These procedures very slowly implement or maintain the water-filling procedure and at each step move in a convex-convergent like manner towards a better solution. Alternatively, during initial training, the entire bit/energy distribution can be calculated from the total power and NSR information.

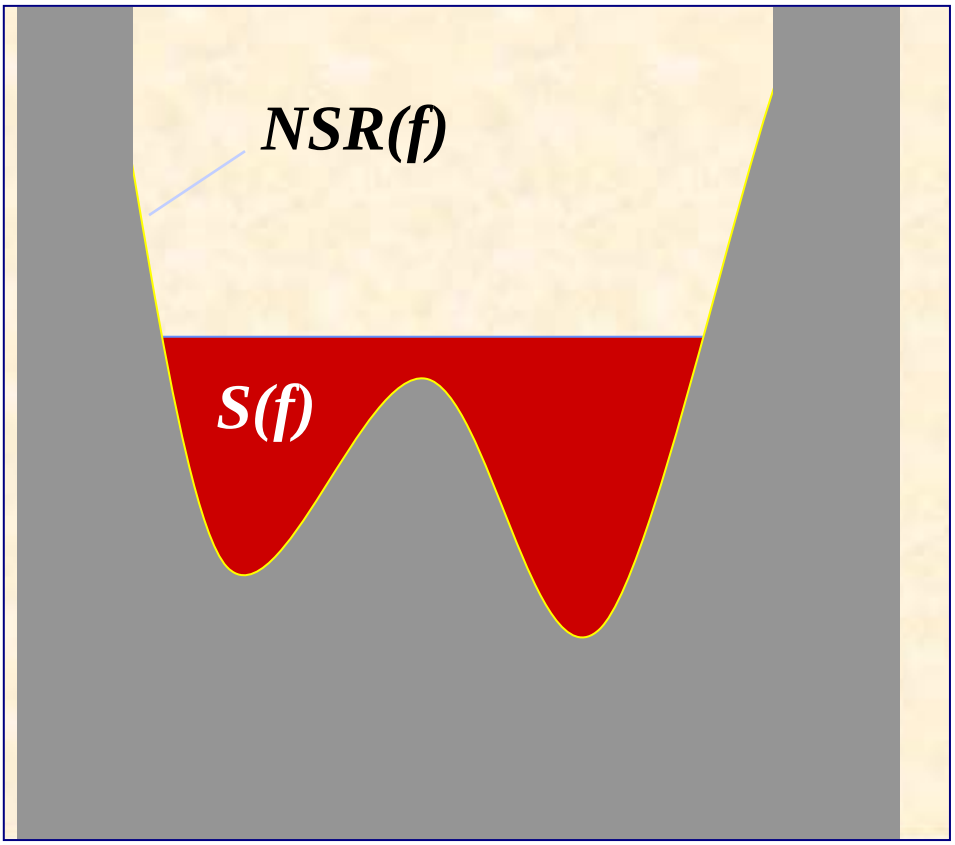


Figure 3 – Basic Water-Filling for single channel spectrum optimization.

The iterative water-filling procedure makes successive or continuous use of water-filling on each line in a globally beneficial way for all lines mutually crosstalking within a binder. Its basic derivation in the

attached paper derives from the obvious (in retrospect) observation that if all lines are optimum, then each must be a water-filling solution in view of the spectra and consequent crosstalk of all other lines [1]. This is sometimes called a "stationary point" in optimization theory. A margin-maximizing version of single-line water-filling (which actually minimizes power at a fixed margin level, see [3]) is typically used in practice and can also be used in iterative water-filling, in which case the data rate to be attempted on each line is known to the modems on that line. While it is possible to have multiple stationary points and consequently multiple minima – a case of such has yet to be uncovered over a wide range of DSL lines considered. In fact, one can prove multiple minima will not occur under a mild set of assumptions easily satisfied by DSL lines [1]. Figure 4 thus illustrates the iterative water-filling procedure in a flow diagram. The various lines may start with any initial spectra, and water-filling is successively applied to each line while the spectra of the others are held constant. After 2 or 3 passes, the procedure will converge to a stationary point in practice.

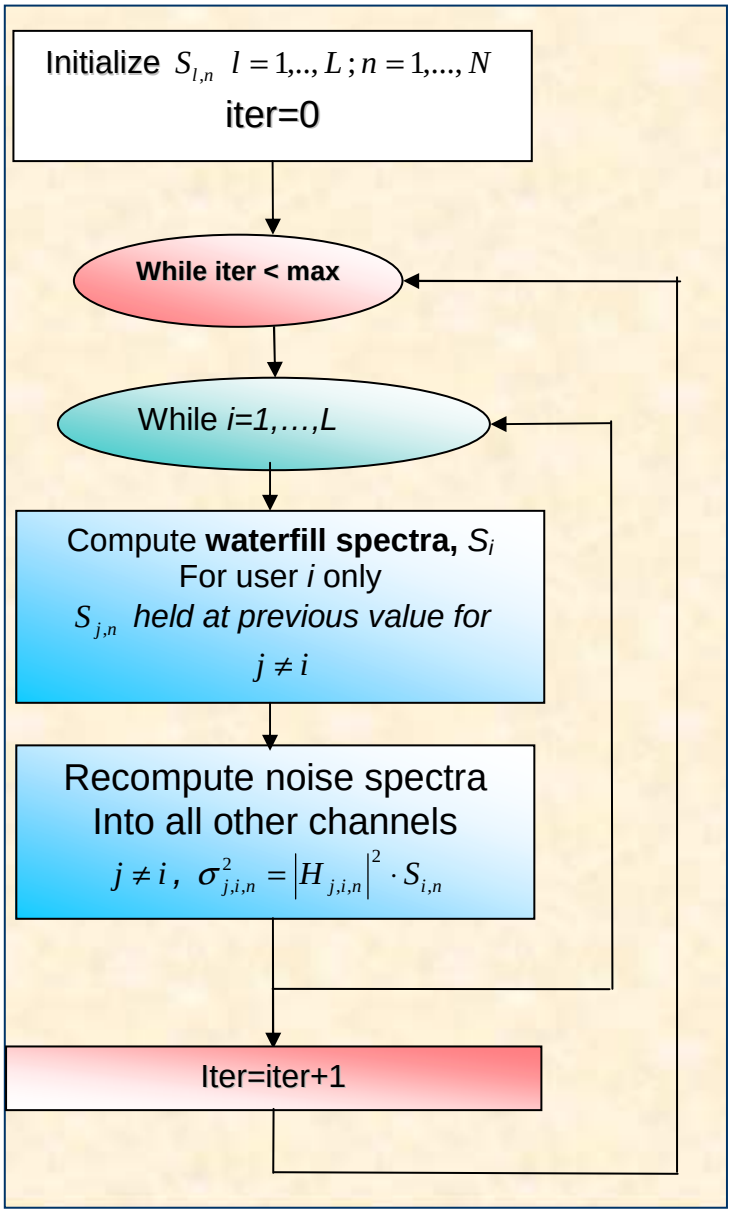


Figure 4 – Flow Chart for Iterative Water-filling

In Figure 4, L is the number of lines involved in the procedure, n is a frequency or tone index, j and i are user/line indices that go from 1 to L. Also, $\left|H_{j,i,n}{}^{2}\right|$ is the transfer function magnitude from line i into line j on tone n, while $S_{i,n}$ is the power spectal density at tone n of line i, and $\sigma^2_{j,i,n}$ is the crosstalk noise on line j from line i on tone n. Any line that has a fixed spectrum (for instance an HDSL, SDSL or perhaps QAM VDSL) simply remains as it is, and all other lines move to the best global position with respect to this fixed spectrum The ultimate result is all lines work better. If all implement iterative waterfilling, the overall solution is yet better than if only some of the lines implement waterfilling. This provides a very acceptable migration strategy from the fixed DSLs of today to the variable DSL spectra of ADSL and future VDSL, both symmetric and asymmetric.

3.1 Distributed or Autonomous Water-filling

Iterative water-filling converges because at each step the solution can be shown to be moving closer to a global optimum. One can envision that each time a line senses some increase in crosstalk in a given region of spectra, it will move information and energy away from that region to a region of comparatively less crosstalk. Another line or lines sensing that reduction will correspondingly move energy toward that region. Each step is an improvement, and eventually all the lines will converge. The procedure in Figure 4 assumes one line at a time changes its spectrum. However, the same convergence applies even if each of the lines autonomously (Level 0 – no coordination in DSM) implement water-filling as long as each step improves the line's spectra with respect to the current spectra. With discrete quantization of information distribution and energy in practice, the gradual convergence occurs until the very last movements of bits where a small cycle of energy/bit moves of negligible consequence on performance can occur.

The distributed implementation of iterative water-filling is a very important result in that it implies that existing DMT systems (ADSL and VDSL) are already unwittingly implementing iterative water-filling. Section 4 illustrates some unexpected improvements that may occur in practice that were not previously noted nor appreciated. Also illustrated are the resultant spectra for various binder topologies.

3.2 Computing the Rate Region

An outer loop can be applied in iterative water-filling in Figure 4 to generate the rate region. The outer loop can either drop the data rate attempted (margin-adaptive water-filling) or drop the total power used (rate-adaptive water-filling) to locate different points in the rate region. Each such reduction on any given DSL line leads to a different operating point for all lines in the mutually crosstalking binder, and thus a different point in the rate region. By performing such training for several different power/rate levels on each of the lines, the achievable rate region for all the lines can be constructed and stored.

3.2.1 Level Zero (No Coordination) Implementation

In this case, the rate region can be only computed somewhat manually by each modem simply training to a given fixed margin level (say 6 dB) for a set of data rates it attempts to achieve. While this clearly has some level of improvement with respect to current spectrum-management practice, some level of data reporting and consequent control, even if manual, leads to an improvement.

3.2.2 Level One

In level one coordination, data can be reported to a maintenance facility. Such reporting may only be the achieved data rate and or margin and the associated transmit power level for each reporting modem. The maintenance center could on a basis of acceptable margin (working modem), then map the possible jointly achieved rate points and thus trace a version of the rate region. That would require only Level One Data coordination. However, with Level One Control, the maintenance center could instruct the various modems at which rates and power levels they can operate (assuming the modem then implements water-filling if so instructed or simply just changes its data rate and spectrum use according to some known standardized pattern such as in SDSL or QAM VDSL). Since each combination of possible data rates

corresponds to a different spectrum combination, the maintenance center could choose an L -tupple that is achievable and then recommend the corresponding rate to each participating line. The water-filling procedure implemented at that rate would then ensure all lines converge to a set of spectra that enables the safe implementation of recommended data rates.

4. Results of Iterative Water-Filling

This section provides three examples of the Level-0 and also Level-One use of iterative waterfilling: (1) RT-located-ADSL FEXT into CO-located ADSL (2) VDSL FEXT into ADSL and (3) mixture of asymmetric and symmetric VDSL service in the same binder.

4.1 RT-located-ADSL FEXT into CO-located ADSL

The problem considered here is similar to one considered by Voyan in [5] of the 3 ADSL systems shown in Figure 5. The ADSL receiver on Line 1 will sense large FEXT from the RT-located ADSL downstream transmitters, both of which are only 5 kft away from the ADSL receiver. Under nominal T1.417 assumptions, line 1 basically does not work – that is the achievable data rate is 300 kbps downstream. Voyan finds a frequency-division-multiplexed Level-One coordination DSM solution to the problem where basically Line 2 uses only higher frequencies and Lines 1 and 3 use other frequencies in the lower range so that each l. They achieve in excess of 1 Mbps all all 3 lines. In our experiments, Line 2 quickly migrates to higher frequency use as Voyan also notes (except here with no coorination nor FDM imposed) and need not be further considered (as they also noted).

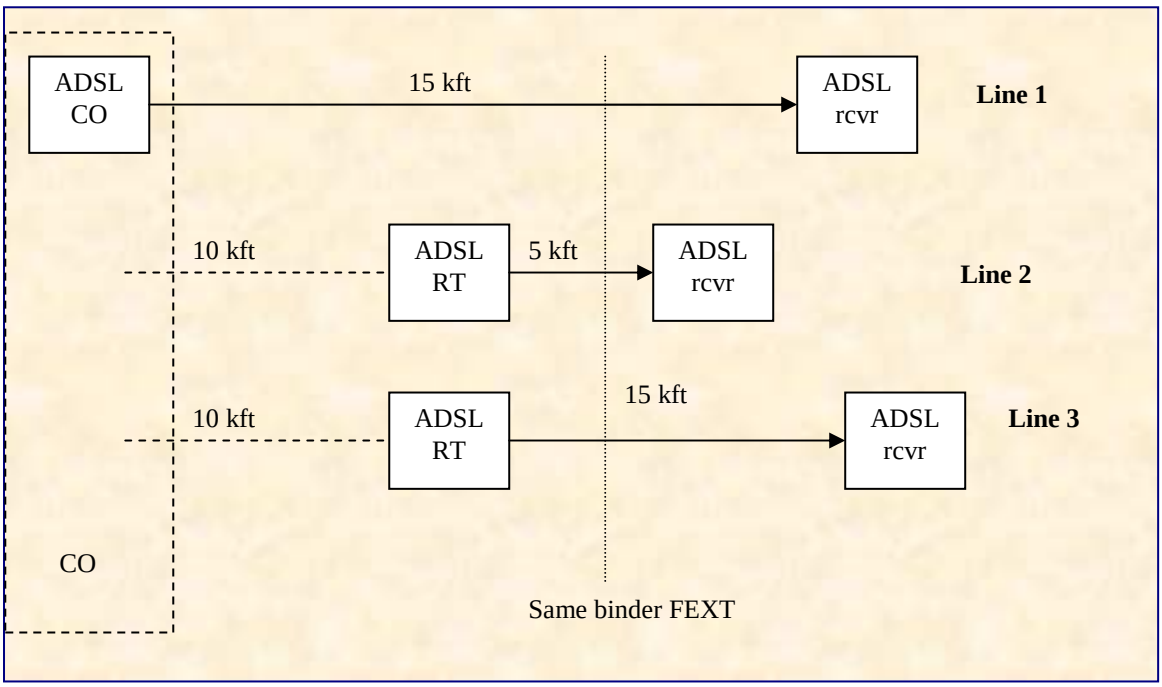


Figure 5 – Illustration of RT-located-ADSL FEXT into CO-located ADSL.

Instead of using FDM, we use here Level Zero (i.e., no coordination) knowing by allowing each line to train at 1 Mbps with margin-adaptive (power-minimizing for given rate) water-filling. In fact it is possible to get 1 Mbps on line 1, 1.6 Mbps on Line 2, and 1.2 Mbps on Line 3. The rate region for Lines 1 and 3 are is illustrated in Figure 6 – Line 2 migrated to higher frequencies and had little further affect and so is not shown in the rate region, nor discussed further. The solution for (1.0, 1.6, 1.2) is not FDM, but actually has an overlap of spectra – quite different than the solution in [5], but still attaining the same data rates with no

coordination and distributed iterative water-filling. Line 2 can run as high as 10 Mbps VDSL without interfering significantly if one desires to try to think of the 3-dimensional rate region. Line one used tones 35-130 at about -43 dBm/Hz. Line 2 used tones 35 to 220 at -90 dBm/Hz. Line 3 used tones 35-100 at -51 dBm/Hz.

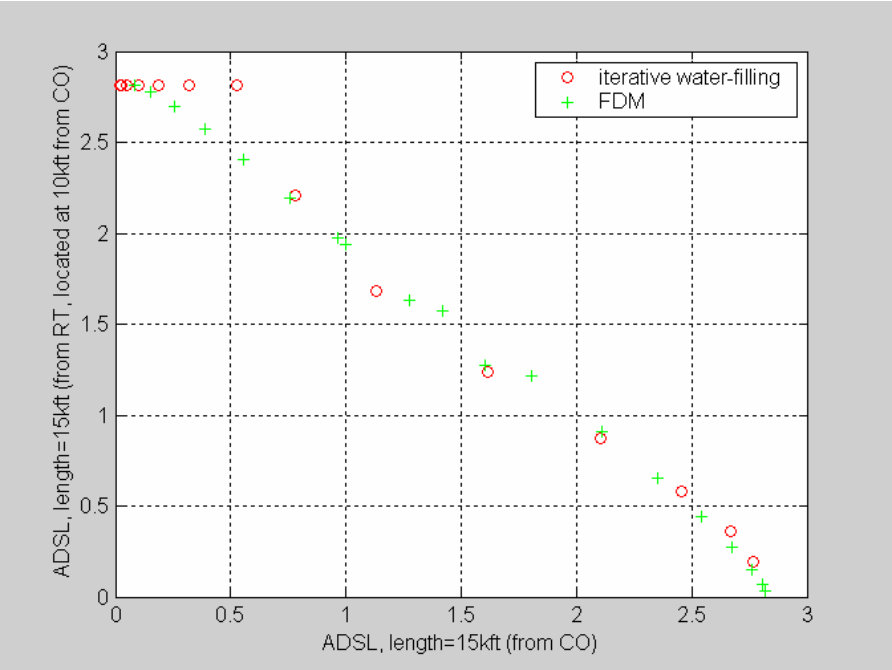


Figure 6 – Lines 1 and 3 rate regions

4.2 VDSL FEXT into ADSL

A problem feared by service providers in potential deployment of VDSL is the downstream FEXT that VDSL ONU emits into the ADSL downstream receivers attached to lines in the same binder. For instance, the situation investigated here has a VDSL ONU at 3000' from its customer, and ADSL central office at 9000' from its customer. Both customers are at approximately the same location. VDSL downstream signals below 1.1 MHz create a large FEXT for ADSL, severely reducing its performance under static spectrum management. Some studies [4] have found the interaction to be so detrimental to the performance of both ADSL and VDSL, leaving static spectrum management with a difficult trade-off that discouraged VDSL deployment.

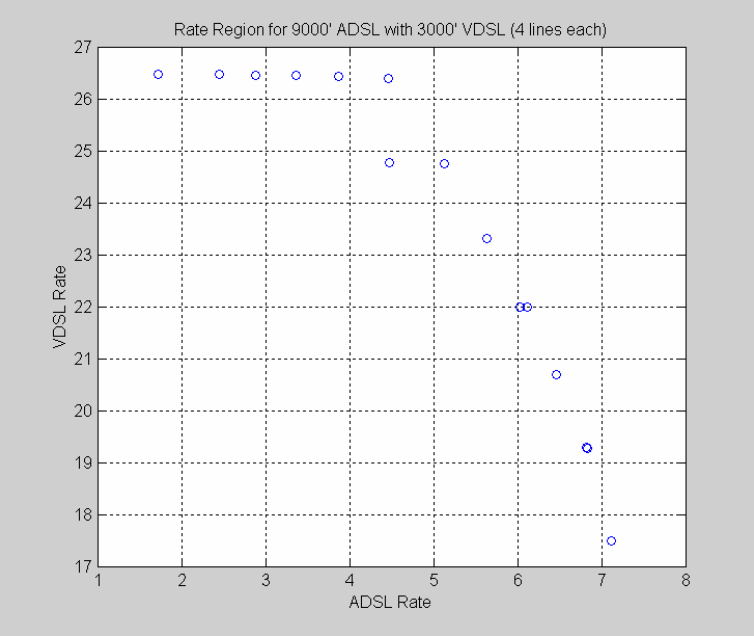


Figure 7 – Rate region for ADSL and VDSL example.

Iterative waterfilling, even with no coordination (i.e. Level Zero coordination) is automatically implemented between DMT ADSL modems and DMT VDSL modems. Under this new presumption of

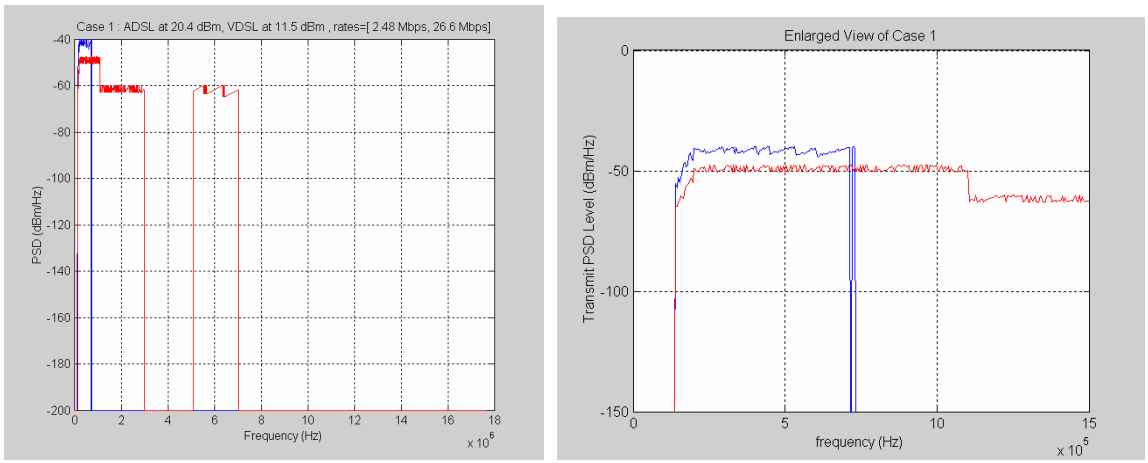


Figure 8 – ADSL and VDSL spectra when at full power [2.48 Mbps, 26.6 Mbps]

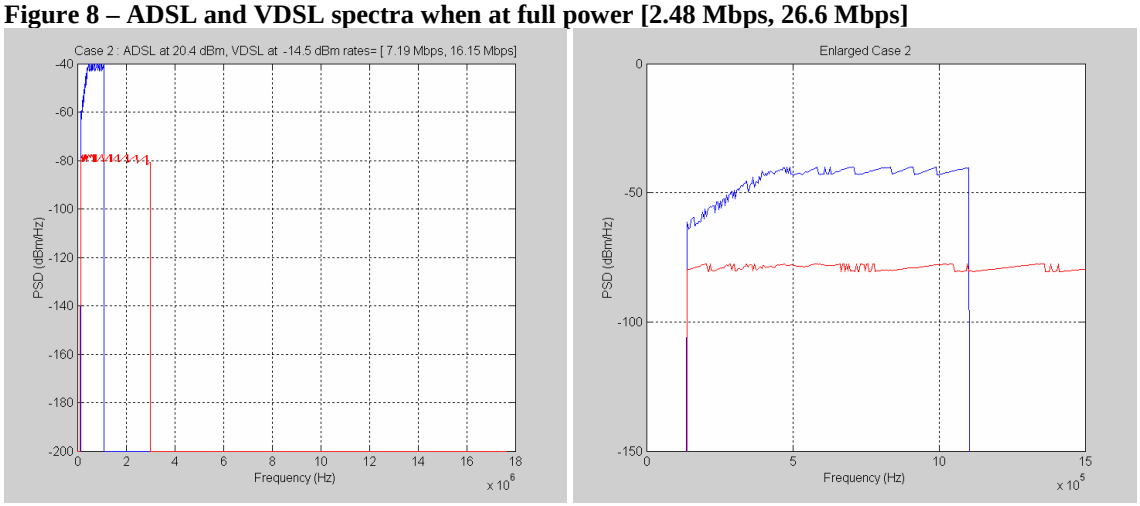


Figure 9 – ADSL and VDSL spectra with full ADSL power, but low VDSL power [7.19,16.2] Mbps

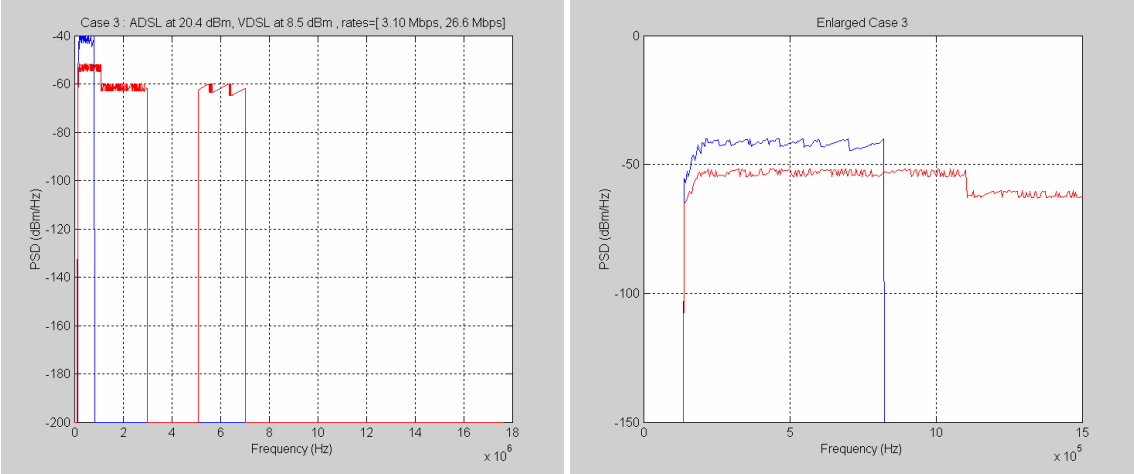


Figure 10 – ADSL and VDSL spectra at intermediate condition [3.1 , 26.6] Mbps

both modems converging to the iterative-water-filling solution creates the rate region of Figure 7. Figure 7 illustrates downstream rate region for 4 VDSL modems at 3000' and 4 ADSL modem at 9000' (both on 26-gauge). All 4 ADSL modems have the same spectrum and so do all 4 VDSL modems in this case because the line lengths and impairments are the same (it would be hard to draw the 8-dimensional region, but we

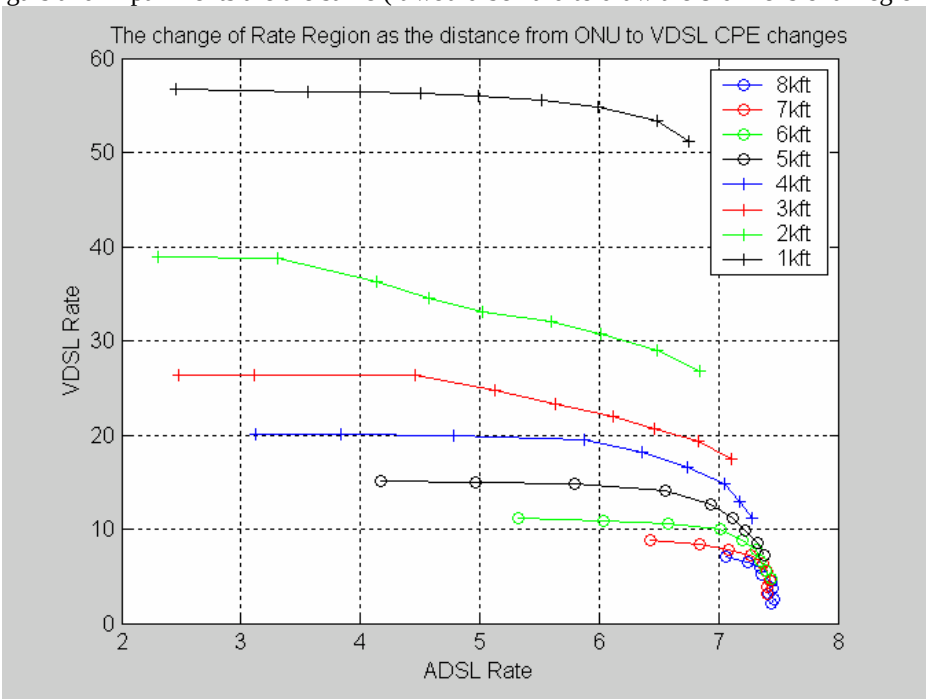


Figure 11(a) VDSL/ADSL with varying VDSL line length.

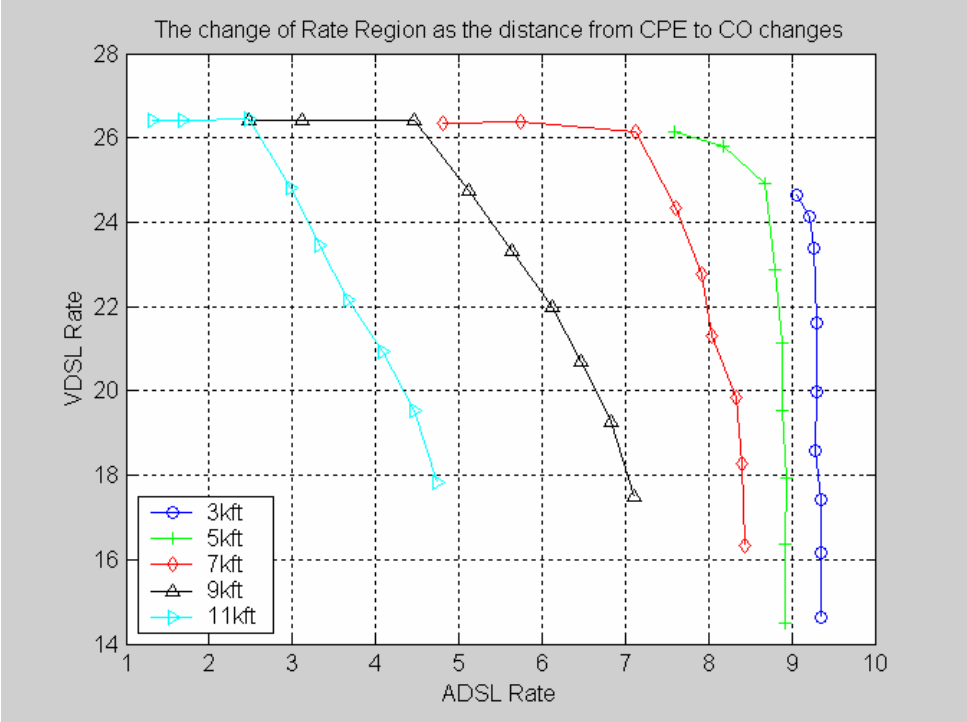


Figure 11(b) VDSL/ADSL with varying ADSL line length.

also wanted to ensure more than just one crosstalker for each of ADSL and VDSL as that leads to very large data rates). The upstream data rates are not shown, but are 3 Mbps for VDSL and 500 kbps for ADSL. In the nominal situation of T1.417 using the 998-standardized spectrum for VDSL of -60 dBm/Hz causes the downstream rate for ADSL to be 1.5 Mbps best case while VDSL hopes to achieve 18 Mbps downstream in recent standards. Clearly, the rate regions here indicate that the VDSL systems can operate at 25-26 Mbps while the ADSL system operates somewhere between 4 and 5 Mbps. The spectra of the two signals is shown in Figures 8-10 for various power levels of the ADSL and VDSL. Note that the ADSL signals tend to operate a full power, but over very different bandwidths as the iterative water-filling determines a cut-off frequency that is a function of the data rate attempted on the ADSL line. The VDSL signal does not vary so much in used band, but does vary power level automatically with margin-adaptive water-filling used (which can be implemented to implement the lowest power to achieve a given margin of 6 dB at a given data rate as in our examples). The VDSL PSD level varies over a wide range, taking only as much power as is necessary to implement a target bit rate (say 26.6 Mbps) while then reducing crosstalk into the ADSL so its performance then increases.

Figures 11(a) and 11(b) show the rate regions as the length of the VDSL or ADSL lines are varied.

4.3 VDSL Examples

As an alternative, this subsection investigates some trade-offs in VDSL when iterative water-filling is used for a 1000' loop and a 3000' foot loop. The objective here is to determine what symmetric services might be offered on the 1000' loop while the 3000' foot loop maintains a 26/3 Mbps asymmetric VDSL service.

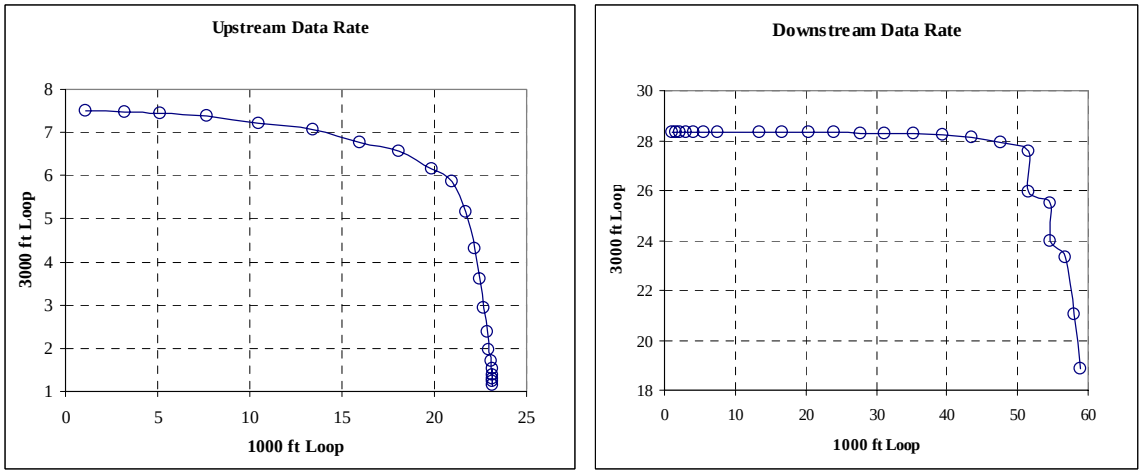


Figure 12 – VDSL 998 Upstream and Downstream rate regions for 1000 and 3000 ft loop in same binder with distributed Level 0 iterative water-filling

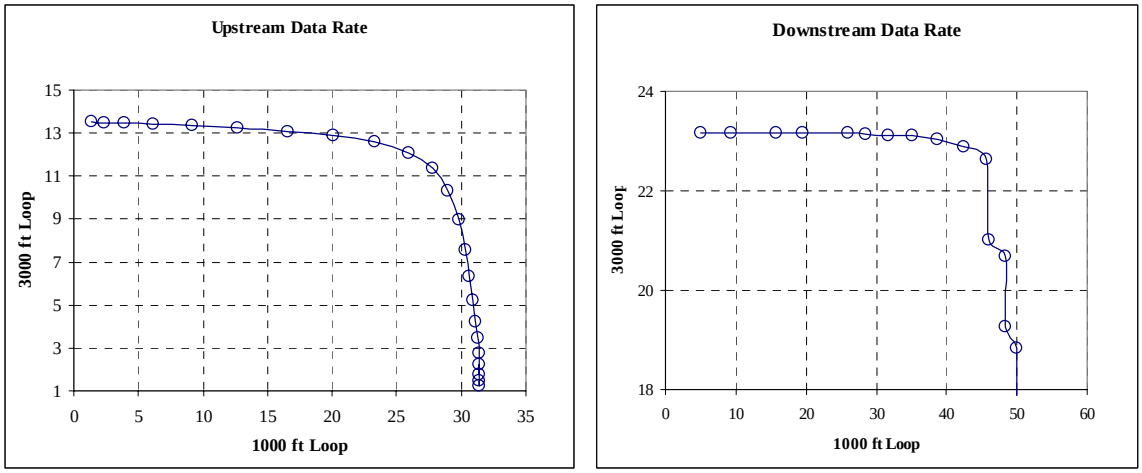


Figure 13 – VDSL Flexible Upstream and Downstream rate regions for 1000 and 3000 ft loop in same binder with Level 0 distributed water-filling, but Level 1 control of up/down cut-off frequencies

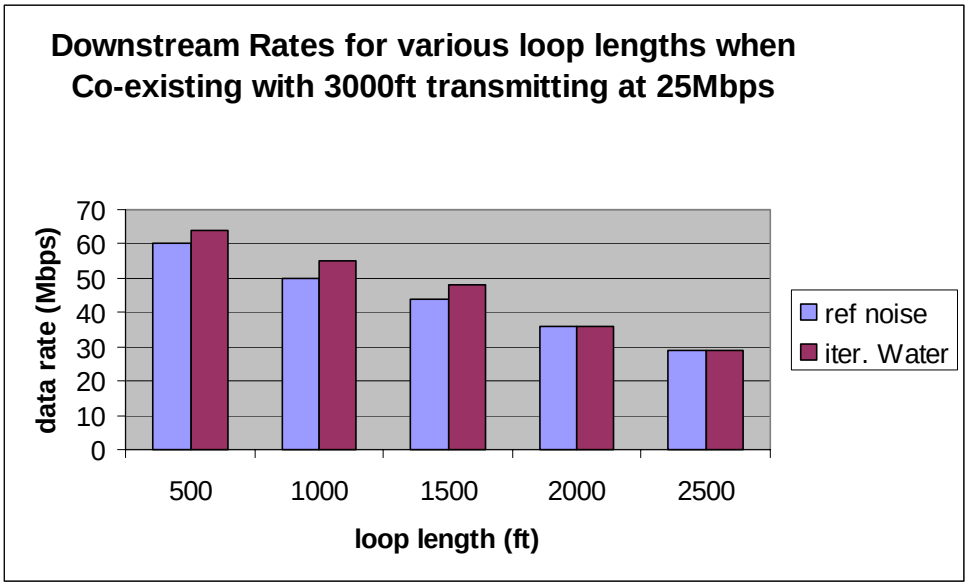


Figure 14(a) – Downstream data rate comparisons with 4 lines at 3000' at 25 Mbps.

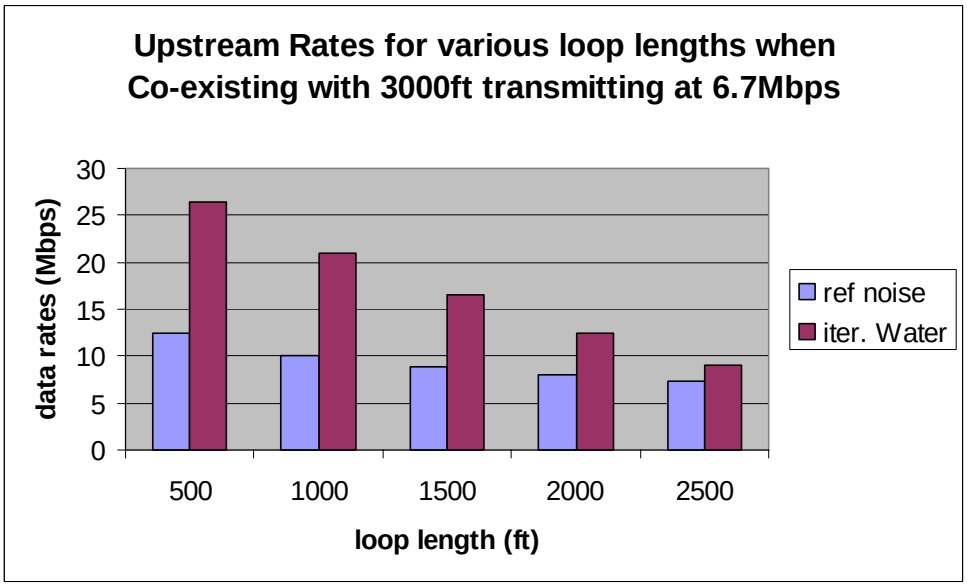


Figure14(b) – Upstream data rate comparisons with 4 lines at 3000' at 6.7 Mbps.

Any point in the upstream rate region in Figure 12 (or Figure 13) can be paired with any point in the downstream rate region of Figure 12 (13). Readers may recall how various service providers argued as to the relative merits of asymmetric versus symmetric service, eventually reaching a compromise where no group saw a data rate it liked with 18 Mbps/1.5 Mbps downstream service being argued by one group of service providers as perhaps the bare minimum for video service, while 6 Mbps symmetric was argued by another group to be below what they wanted but perhaps interesting. Figure 9 shows that even within the 998 plan of VDSL, DMT water-filling VDSL modems avoid each other within the common band in a mutually beneficial way autonomously, leading to 26/6 Mbps on a worst-case 3000 foot VDSL loop and at least 20 Mbps symmetric on the 1000 ft loop. Such a situation might well occur in practice with residential loops being further than small business loops. In any case, the data rates significantly exceed those that were projected by static spectrum management. Figure 13 alters the cut-off frequencies of up and down, given that with DSM, the 998 plan may not have provided the best trade off. Indeed, that leads to 26/13

Mbps on the asymmetric VDSL loop at 3000 ft while the 1000 ft loop has at least 30 Mbps symmetric service. These cut-off frequencies for up/down need not be fixed if DSM is used.

Figure 14 illustrates the trade-off versus line length as the loop length of VDSL instead varies with respect to the best known (and not used because it cannot be implemented in QAM VDSL) reference-length spectrally shaped power-back-off methods of VDSL and the corresponding ADSL data rates.

5. Conclusion

This paper suggests that significant improvements can occur with no coordination if some or all DSL modems in a binder implement water-filling, particularly if the rates attempted are within the achievable rate regions. An early deployment would emphasize margin-adaptive loading, which minimizes power consumption at a fixed 6 dB margin and data rate for various systems. If the data rate set for the binder is not achievable, Level 1 data collection can be used to compute what rate regions are achievable at a central maintenance facility and thereby recommend known data rates to each of the DSL modems to attempt that can be implemented (then with distributed iterative water-filling or "bit-swapping" as it is sometimes called).

The improvements are sufficiently large to consider at least increasing deployment of modems that may already implement water-filling or approximate it rather than modems that use some fixed less desirable spectra that imposes a more severe crosstalk on other lines in the same binder.

6. References

[1] See the attached copy of the document submitted also elsewhere for publication as Journal paper.

[2] J. Cioffi et al., "Proposed Scope and Mission for Dynamic Spectrum Management," *ANSI Contribution T1E1.4/2001-188R4,* November 5, 2001, Greensboro, NC.

[3] T. Starr, J. Cioffi, and P. Silverman, Understanding DSL Technology. Prentice Hall: Upper Saddle River, NJ, 1999.

[4] K. Jacobsen, "Spectral Compatibility of ADSL and VDSL: Parts 1 and 2," *T1E1.4/97-404*, Dec 8, 1997, Sacramento, CA – see also ETSI TM6 TD05, 1997, Verona, Italy.

[5] M. Tsatsanis and I. Kanellakopoulis, "Identification of Crosstalk using MIB-reported data," *ANSI Contribution T1E1.4/2001-278*, Greensboro, NC, November 5, 2001.

COMMITTEE T1 - TELECOMMUNICATIONS **T1E1.4/2001-200R4**

Working Group T1E1.4 (DSL Access)
Greensboro, NC; November 05 - 09, 2001

CONTRIBUTION [1]

Title : **AN ADAPTIVE MULTIUSER POWER CONTROL ALGORITHM**
Source : **Stanford University**
Project : **Dynamic Spectrum Management**

ABSTRACT

This paper considers the multiuser power control problem in an interference network. We introduce the notion of competitive optimality for the ISI interference channel, and propose an iterative water-filling algorithm to reach such competitively optimal power allocations. An adaptive multiuser power control algorithm based on iterative water-filling is illustrated for the upstream VDSL power back-off problem as an example of dynamic spectrum management. By taking into account all loop transfer functions and cross-couplings, and by offering an opportunity for loops to negotiate the best use of power and frequency, the new power control algorithm is found to have a substantial performance improvement when compared to existing methods.

[1] Contacts: Wei Yu (email: weiyu@dsl.stanford.edu), Seong Taek Chung (email: stchung@dsl.stanford.edu), George Ginis (email: gginis@dsl.stanford.edu), and John M. Cioffi (email: cioffi@stanford.edu) Phone: 1-650-723-2150, FAX: 1-650-723-9251. Address: 350 Serra Mall, Packard Building Room 363, Stanford, CA 94305-9515, U.S.A.

NOTICE

This contribution has been prepared to assist Standards Committee T1-Telecommunications. This document is offered to the committee as a basis for discussion and is not a binding proposal. The requirements are subject to change in form and numerical value after more study. Specifically, the right to add to, or amend, the statements contained herein is reserved.

1 Introduction

Optimal power control is a central problem in the design of multiuser communication systems Typical communication systems are interference-limited. In these systems, each user's performance depends not only on its own power allocation, but also on the power allocation of all other users. Consequently, the system design generally involves performance trade-off among different users. Such trade-off is an important consideration and this is the subject of current investigation. This paper considers the digital subscriber line (DSL) environment as a multiuser system. The aim is to provide an optimal power allocation scheme that maximizes the achievable data rates for multiple DSL modems in the presence of mutual interference.

The digital subscriber line technology brings high speed data services to home via ordinary telephone copper pairs [1]. The DSL environment is considered as a multiuser environment because telephone lines from different households are bundled together on the way from the central office, and different lines in the bundle create crosstalk into each other. Such crosstalk could be the dominant noise source in a loop. However, early DSL systems such as ADSL and HDSL are designed as single-user systems. This is because single-user systems are considerable easier to design, and the target data rate requirements for earlier systems are considerably lower than the system multiuser capacity. As the demand of higher data rates increases and the evolution of DSL systems moves toward higher frequency bands, where the crosstalk problem is more pronounced, spectral compatibility and power control have emerged as central issues. This is especially true for the emerging VDSL standard, where frequencies up to 20MHz can be used. This paper will use VDSL as an example, and show that a multiuser system design with an optimal power allocation scheme could result in a large system performance improvement compared to a single-user design that does not take the multiuser aspect into account.

The power control problem in DSL systems differs from the more widely-studied power control problem in wireless systems in the following key aspect: although the DSL environment varies from loop to loop, it does not vary over time. Fading and mobility are not issues, and consequently, the assumption of perfect channel knowledge is realistic, and it will be assumed here. This allows the possibility of implementing sophisticated centralized power control schemes. On the other hand, unlike the wireless situation where flat fading can often be assumed, the DSL loops are severely frequency selective. Thus, the optimal power allocation scheme needs to consider not only the total amount of power allocated for each user, but also the allocation of power in each frequency. Nevertheless, power control schemes designed for wireless systems [2] [3] still provide us with considerable insight. In particular, the VDSL systems suffer from a near-far problem similar to that in CDMA systems. The near-far problem occurs when two transmitters located in different distances both attempt to communicate with the central office. When one transmitter is much closer to the central office than the other, the interference due to the closer transmitter overwhelms the signal from the farther transmitter. This paper will propose a power allocation scheme that overcomes this problem. The proposed solution takes advantage of the channel and crosstalk coupling functions that are available, and allocates power across the frequencies for each user based on the idea of competitive optimality [4]. It turns out that the channel information is

crucial in providing a dramatic data rate improvement over existing power control methods.

The rest of the paper is organized as follows. Section 2 reviews the DSL environment, and models a typical DSL loop as an interference network. Section 3 defines and characterizes the competitive equilibrium in such a network, and devises an iterative method to achieve the equilibrium. An adaptive power allocation method for VDSL will then be proposed in section 4 based on the idea of competitive equilibrium. System performance are characterized in section 5 and conclusions are drawn in section 6.

2 The VDSL Environment

DSL modems use frequencies above the traditional voice band to carry high-speed data. To combat intersymbol interference in the severely frequency selective telephone channel, DSL transmission uses Discrete Multitone (DMT) modulation, which divides the frequency band into a large number of sub-channels and lets each sub-channel carry a separate data stream. The use of DMT modulation allows easy implementation of arbitrary power allocation in each frequency tone, thus making spectral shaping easy to realize. The VDSL transmission studied in this paper has been standardized. It uses frequencies range from 0 to 17.6MHz, and the frequency band is divided into 4096 tones, with each tone 4.3125kHz wide.

A DSL bundle can consist of up to 100 subscriber lines bundled together. Because of their close proximity, each line creates electromagnetic interference into each other, thus causing crosstalk noise (see Figure 1). Near-end crosstalk (NEXT) refers to crosstalk created by transmitters located on the same side as the receiver. Far-end crosstalk (FEXT) refers to crosstalk created by transmitters located on the opposite side. NEXT is typically much larger than FEXT. xDSL standards in North American and Europe use frequency division duplexing to divide upstream and downstream transmission, while time-division duplexing is used in Japan. We will mainly concentrate on frequency duplexed systems in this paper. Because NEXT is typically much larger than FEXT, the VDSL standard requires all loops to transmit in the same direction in each frequency band, thus avoiding NEXT completely. The entire frequency band is divided into a number of upstream and downstream bands.

The DSL environment is a multiuser environment because the background noise in the loop is typically small, and the system performance is limited by crosstalk. The DSL environment consists of multiple transmitters and multiple receivers interfering into each other, so it can be modeled as an interference channel. This is different from the much-studied multiple access or broadcast scenarios where many-to-one or one-to-many situations are considered. Multiaccess and broadcast models are not suitable for the DSL environment because data coordination among the transmitters or the receivers are generally not possible. Coordination cannot be done on the remote terminals because they are distributed in different locations. Coordination at the central office side can be done only if all loops are controlled by the same service provider, which may not be the case under the current regulatory condition. Thus the DSL environment is most appropriately modeled as an interference channel where many-to-many transmissions take place.

From a theoretical point of view, the interference channel is little understood. Even for the

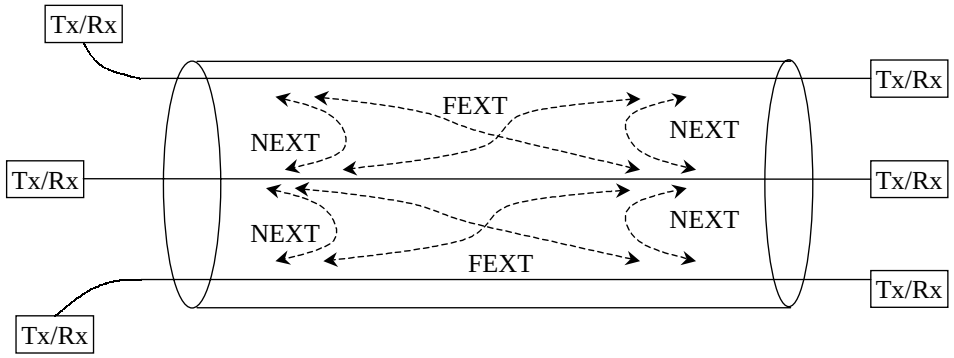


Figure 1: The DSL crosstalk environment.

simplest two-user additive white Gaussian interference channel, only partial achievable regions and outer bounds are available [5]. The optimal transmission in an interference channel, as commonly believed, probably involves some aspects of multiuser detection. In fact, if the interference level is very high, interference subtraction can achieve a data rate as if all interference can be completely removed [6] [7]. For this to happen, however, the interfering channel must be stronger than the direct channel, which does not occur in a realistic DSL environment. Unfortunately, when the interference level is low, the capacity problem is largely unsolved. The best achievable region uses partial superposition coding at the transmitter and interference cancelation ideas in the receiver [5] but its optimality has not been shown. In this light, we will restrict our attention to transmission techniques where no multiuser detection takes place, and focus solely on the problem of optimal power allocation for each user in the network.

Consider the interference channel model depicted in Figure 2. There are N transmitters and N receivers in the network. The channel from user i to user j is modeled as an ISI channel, whose transfer function in frequency domain is denoted as $H_{ij}(f)$, where $0 \leq f \leq F_s$, $F_s = \frac{1}{2T_s}$, and T_s is the sampling rate. In addition to the interference noise, each receiver also sees a background noise whose power spectrum density is denoted as $\sigma_i(f)$. Denote the power allocation for each transmitter as $P_i(f)$, which must satisfy a power constraint:

$$\int_0^{F_s} P_i(f)df \leq \mathbf{P_i}. \tag{1}$$

The achievable data rate for each user while treating all interference as noise is:

$$R_i = \int_0^{F_s} \log_2 \left(1 + \frac{P_i(f)|H_{ii}(f)|^2}{\Gamma \left(\sigma_i(f) + \sum_{j \neq i} P_j(f)|H_{ji}(f)|^2 \right)} \right) df, \tag{2}$$

where Γ denotes the SNR-gap with depends on the probability of error, the modulation scheme and the coding applied. A coding and modulation scheme that approaches the information theoretical capacity has $\Gamma = 0$dB. The objective of the system design is to maximize the set of rates

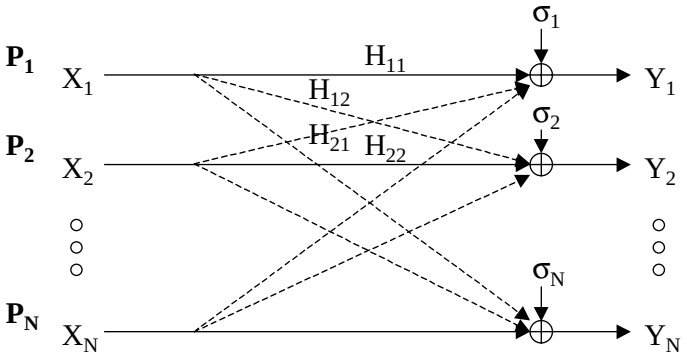


Figure 2: A Gaussian interference network.

$\{R_1, \cdots R_N\}$ subject to the power constraints (1). Notice that for each transmitter, increasing its power at any frequency band would increase its own data rate, but it also causes increased interference to other users, and is therefore detrimental to other user's transmission. Thus, the design must consider the trade-off among the data rates of all users. For example, if the objective is to maximize the sum capacity, then there is no guarantee of a minimal data rate to any one user. The sum-rate optimal power allocation is often the one that assigns high data rates to users closer to the central office, and low data rates to users farther away. Realistic DSL deployment often requires multiple service rates be supported for all users, and the required level of service of each user could be arbitrary. Thus a single figure of merit is often inadequate to represent the system performance. A convenient way to fully characterize the trade-off among the users is through the notion of capacity region, which is defined as:

$$\mathcal{R} = \{(R_1, \cdots, R_N) : \exists (P_1(f), \cdots P_N(f)) \text{ satisfying (1) and (2)}\}. \tag{3}$$

The capacity region characterizes all possible data rate combinations among all users. Although in theory, the capacity region can be found by an exhaustive search through all possible power allocations, or by a series of optimizations involving weighted sums of data rates, computational complexities of these approaches are prohibitively high. This is because the capacity formula is a non-convex function of power allocations, and consequently, the usual numerical algorithms are able to find only the local maxima, and not the global maximum. The solution proposed here circumvents the difficulty by defining a different notion of competitive optimality. Although the proposed method does not achieve all points in the capacity region as defined above, it will nevertheless perform much better than the current DSL systems.

The current DSL systems are designed as single-user systems. In addition to the power constraint, each user is also subject to a power spectrum density constraint. The power spectrum density constraint limits the worst-case interference level from any modem, thus each modem can be designed so that it withstands the worst-case noise. Such design is conservative in the sense that realistic deployment scenarios often have interference much smaller than the worst-case noise,

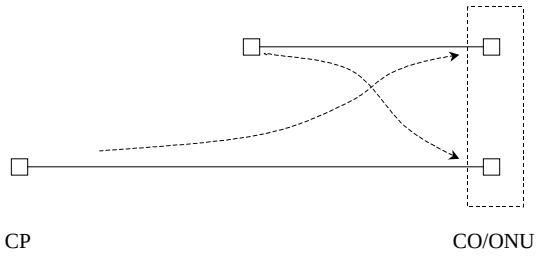


Figure 3: A situation requiring upstream power back-off.

and current systems are not designed to take advantage of this fact. In addition, the same power spectrum density constraint is applied to all modems uniformly regardless of their geographic location. The lack of a mechanism to allocate different amount of power to different users in different locations is problematic because of the near-far problem mentioned before. Figure 3 illustrates a configuration in which two VDSL loops in the same binder emanate from the central office (CO) to the customer premises (CP). When both transmitters at the CP-side transmit at the same power spectral density, because of the difference in line attenuation, the FEXT caused by the short line can overwhelm the data signal in the long line. The upstream performance of the long line is therefore severely affected by the upstream transmission of the short line. To remedy this spectral compatibility problem between short and long lines, the short lines must reduce their upstream power spectral densities so that they do not cause unfair interference onto the long lines [8] [9] [10]. This reduction of upstream transmit power spectral density is known as upstream power back-off (UPBO) and it has been under intensive study in the VDSL standardization bodies. Note that the downstream direction does not suffer from a similar problem because although all transmitters at the CO-side also transmit at the same power spectral density, the FEXT they cause to each other is identical at any fixed distance from CO [9]. This downstream FEXT level is always much smaller than the data signals, so it does not pose a serious problem to downstream transmission.

Several upstream power back-off algorithms have been proposed in VDSL. An extensive review of these methods can be found in [10]. All current power back-off methods attempts to reduce the interference emission by shorter loops by forcing it to emulate the behavior of a longer loop. For example, in the constant power back-off method, a constant factor is reduced across the frequency in upstream transmission bands, so that at a particular reference frequency the received PSD level from shorter loops is the same as the received PSD level from a longer reference loop. A generalization of this method is called the reference length method where variable amount of back-off is implemented across the frequency so that the received PSD for a shorter loop is the same as some longer reference loop at all frequencies. Imposing the same-PSD criterion for shorter loops across the entire frequency band may sometime be too restrictive. This is because high frequency bands usually have too much attenuation to be useful by long loops, thus short loops should be able to transmit at high frequency bands without worrying about its interference. This observation leads

to the multiple reference length method, which sets a different reference length at each upstream frequency band. All three methods mentioned so far equalize the PSD level of a shorter loop to the PSD level of some longer reference loop. While this may be easy to implement, better performance can be obtained if the interference levels themselves are equalized instead. Examples of such approaches are the equalized-FEXT method, which forces the FEXT emission by shorter loops to be equal to the FEXT from a longer reference loop, and the reference noise method which forces the FEXT emission to equal to a more general reference noise. Although currently, there has not been a consensus on a single favorite method in the standardization community, it is clear however, a versatile method such as reference noise that allows shaped spectrum is most likely to provide optimal performance.

All previously proposed power back-off methods require the power or noise spectrum by the short loops to comply with a reference loop or a reference noise. Although the optimal choice of the reference is far from being clear, these approaches are simple to implement because they only require each loop to adjust its power spectrum according to a reference, and thus does not require any knowledge of the network configuration. If, however, in addition, all loop and coupling characteristics in the network are known to either the loops themselves, or a centralized third party, adaptive adjustment of power spectrum level can then be implemented, and much better performance is potentially possible. However, the optimization problem involved here is complex, because the number of variables is large, and due to the non-convex nature of the problem, many local minima exist. Early attempts at solving this problem often resorted to added constraints such as all transmitter power spectrum densities are the same [11], or all PSDs are in some sense symmetrical [12]. The first attempt in finding the true global optimum is due to Cherubini [13], where an approach based on quantum annealing is used to minimize the total energy subject to rate constraints on each user. This paper will take a different approach. Instead of finding the global maximum, we concentrate on a notion of competitive optimality, which has the intuitive appeal of being the locally optimal solution that all users have the incentive to move toward. These competitive optimal points are easy to characterize, and they lead to a power control method that offers the following advantages compared to the previous methods:

- Unlike previous methods that set some PSD level for each VDSL transmitter based solely on its interference emission level, the new power allocation method strikes a balance between maximizing each user's own data rate and minimizing its interference emission. In particular, the frequency selective nature of the channel is dealt with explicitly.

- By taking into account all loop transfer functions and cross-couplings, the new method offers the loops an opportunity to negotiate the best use of power and frequency with each other.

- The usual PSD constraint, which is in place for the purpose of controlling interference, is no longer needed. Only total power constraints apply.

- Unlike previous methods, which fix a data rate for each loop regardless of actual service requirement, the new method naturally supports multiple service requirements in different loops.

- The proposed method does not involve arbitrary decisions on the reference noise or reference length.
- Much better performance can be achieved.

These benefits are possible because the loop and interference coupling functions are assumed to be available to each loop or to a centralized agent. The performance improvement of this competitive optimal power allocation will suggest that this channel information is extremely valuable in loop planning and service deployment.

3 Competitive Optimality

The traditional information-theoretic view of an interference channel allows the different transmitters, while sending independent data streams, to be cooperative in their respective coding strategies, so that interference cancelation can take place in the receivers. If such cooperation cannot be assumed, the interference channel can be better modeled as a non-cooperative game. Under this viewpoint, each user competes for data rates with the sole objective of maximizing its *own* performance regardless of all other users. This scenario is particularly realistic in the current unbundled environment where different loops in the same binder could belong to different service providers, and they indeed compete in the local access market. Now, because each modem has a fixed power budget, each user should adjust its power allocation to maximize its own data rate, while regarding all other interference as noise. If such power adjustment is done continuously for all users at the same time, it is natural to ask, will they eventually reach an equilibrium? Such an equilibrium will be a desired system operating point, since at the equilibrium, all users have reached its own local maximum, and nobody has an incentive to move away from its current power allocation. From a game theory perspective, this equilibrium point is called the Nash equilibrium. A Nash equilibrium is defined as a strategy profile in which each player's strategy is an optimal response to the other player's strategy [14]. The following will characterize the Nash equilibrium in the Gaussian interference channel game, and determine its existence and uniqueness in realistic channels.

Let's focus on a two-user interference channel, and consider the following simplified model:

$$\mathbf{y}_1 = \mathbf{x}_1 + \mathcal{A}_2\mathbf{x}_2 + \mathbf{n}_1 \tag{4}$$

$$\mathbf{y}_2 = \mathbf{x}_2 + \mathcal{A}_1\mathbf{x}_1 + \mathbf{n}_2, \tag{5}$$

where the channel transfer functions are normalized to unity. The square magnitude of the interference transfer functions $\mathcal{A}_1$ and $\mathcal{A}_2$ are denoted as $\alpha_1(f)$ and $\alpha_2(f)$, respectively. Let $N_1(f)$ and $N_2(f)$ denote noise power spectrum densities. The two senders are considered as two players in a game. The structure of the game, i.e., the interference coupling functions and noise power, are assumed to be common knowledge to both players. A strategy for each player is its transmit power spectrum, $P_1(f)$ and $P_2(f)$, subject to the power constraints $\int_0^{F_s} P_1(f)df \leq \mathbf{P}_1$, and $\int_0^{F_s} P_2(f)df \leq \mathbf{P}_2$, respectively. (We will only consider deterministic, or pure strategy here.) The

payoff for each user is its respective data rate. Under the simplifying assumption that no interference subtraction is performed regardless of interference strength, the data rates are:

$$R_1 = \int_0^{F_s} \log\left(1 + \frac{P_1(f)}{N_1(f) + \alpha_2(f)P_2(f)}\right) df, \tag{6}$$

$$R_2 = \int_0^{F_s} \log\left(1 + \frac{P_2(f)}{N_2(f) + \alpha_1(f)P_1(f)}\right) df. \tag{7}$$

If we compare the above expression with equation (2), we can identify:

$$N_1(f) = \frac{\Gamma\sigma_1(f)}{|H_{11}(f)|^2} \tag{8}$$

$$\alpha_2(f) = \frac{|H_{21}(f)|^2}{\Gamma|H_{11}(f)|^2}, \tag{9}$$

and similarly for $N_2(f)$ and $\alpha_1(f)$. Thus, the simplified model incurs no loss of generality. The data rate game considered here is not a zero-sum game, i.e. one player's loss is not equal to the other player's gain. Our interest is in characterizing all the pure strategy Nash equilibria. Since at a Nash equilibrium, each user's strategy is the optimal response to the other player's strategy, and for each user, the optimal power allocation given other player's power level is the water-filling of the power against the combined noise and interference, a Nash equilibrium is reached if water-filling is simultaneously achieved for all users,

A complete characterization of the simultaneous water-filling point in the interference channel is surprisingly difficult to do. This paper offers several sufficient conditions for the existence and uniqueness of the Nash equilibrium in the two-user case.

Theorem 1 *Suppose that $\alpha_1(f)\alpha_2(f) < 1$, $\forall f$, then at least one pure strategy Nash equilibrium in the Gaussian interference game exists. Further, let $\epsilon_0 = \sup\{\alpha_1(f)\}\sup\{\alpha_2(f)\}$, $\epsilon_1 = \sup\{\alpha_1(f)\alpha_2(f)\}$, $\epsilon_2 = \sup\{\alpha_1(f)\}\frac{1}{F_s}\int_0^{F_s}\alpha_2(f)df$, and $\epsilon_3 = \sup\{\alpha_2(f)\}\frac{1}{F_s}\int_0^{F_s}\alpha_1(f)df$. If any of the following conditions, $\epsilon_0 < 1$, $\epsilon_1 + \epsilon_2 < \frac{1}{2}$, or $\epsilon_1 + \epsilon_3 < \frac{1}{2}$ is satisfied, then the Nash equilibrium is unique, and is stable.*

Proof: At a Nash equilibrium, both users have a water-filling level, call them (K_1, K_2). The first idea in proving the existence of a Nash equilibrium under a power constraint $(\mathbf{P}_1, \mathbf{P}_2)$ is to establish the existence of a Nash equilibrium under a fixed water level. Assume $\alpha_1(f)\alpha_2(f) < 1\ \forall f$, and fix (K_1, K_2). The Nash equilibrium power allocation $(P_1(f), P_2(f))$ can be found by simultaneously solving the following water-filling conditions at each frequency f: when $P_1(f)$ (or $P_2(f)$) is zero, the combined interference and noise must be greater than or equal to K_1 (or K_2); when $P_1(f)$ and $P_2(f)$ are positive, we should have:

$$P_1(f) + \alpha_2(f)P_2(f) + N_1(f) = K_1, \tag{10}$$

$$P_2(f) + \alpha_1(f)P_1(f) + N_2(f) = K_2. \tag{11}$$

Now, if either $K_1 < N_1(f)$, or $K_2 < N_2(f)$, then trivially $P_1(f) = 0$ or $P_2(f) = 0$ satisfies the condition. So, without loss of generality, assume that $K_1 > N_1(f)$ and $K_2 > N_2(f)$. If

$\alpha_1(f) > \frac{K_2 - N_2(f)}{K_1 - N_1(f)}$, setting $P_1(f) = 0$, and $P_2(f) = K_2 - N_2(f)$ satisfies the condition. Also, if $\alpha_2(f) > \frac{K_1 - N_1(f)}{K_2 - N_2(f)}$, setting $P_2(f) = 0$, and $P_1(f) = K_1 - N_1(f)$ satisfies the condition. The above two conditions on α_1 and α_2 cannot be both true at the same time because $\alpha_1(f)\alpha_2(f) < 1$. If neither is true, then equations (10) and (11) have a positive solution, which is the water-filling power allocation. Thus, under all cases, the simultaneous water-filling condition has a solution, and the solution must be a Nash equilibrium.

Next, we establish that for a given power constraint $(\mathbf{P}_1, \mathbf{P}_2)$, there exists (K_1, K_2) whose Nash equilibrium has exactly this power. For each (K_1, K_2), denote the total power level at the corresponding Nash equilibrium as $(\mathbf{P}_{K_1}, \mathbf{P}_{K_2})$. Observe that when $\alpha_1(f)\alpha_2(f) < 1$, if $K_1 < K_1'$ and $K_2 = K_2'$, then $\mathbf{P}_{K_1} \leq \mathbf{P}_{K_1'}$ and $\mathbf{P}_{K_2} \geq \mathbf{P}_{K_2'}$. This can be verified by working through the simultaneous water-filling condition. Now, start with $K_1 = K_2 = 0$. Increase K_1 until $\mathbf{P}_{K_1} = \mathbf{P}_1$, then increase K_2 until $\mathbf{P}_{K_2} = \mathbf{P}_2$. But then, we have $\mathbf{P}_{K_1} \leq \mathbf{P_1}$ by observation. So, we can increase K_1 again, until $\mathbf{P}_{K_1} = \mathbf{P_1}$, then increase K_2, etc. The increasing sequences of K_1's and K_2's converge because they cannot go to infinity with finite power constraints. The limit point is a Nash equilibrium corresponding to $(\mathbf{P}_1, \mathbf{P}_2)$. Thus, we have established the existence of a Nash equilibrium under the power constraint.

To prove uniqueness, let $(P_1^N(f), P_2^N(f))$ be the power distribution at a Nash equilibrium, whose existence we already established. Start with any power distribution $P_1^{(0)}(f)$ that satisfies the power constraint. Water-fill for $P_2^{(0)}(f)$, assuming $P_1^{(0)}(f)$ as interference. Then, water-fill for $P_1^{(1)}(f)$, assuming $P_2^{(0)}(f)$ as interference. Then, $P_2^{(1)}(f)$, etc. We will show that this iterative water-filling process converges in $\mathbf{L}_1$-norm, $\frac{1}{F_s}\int_0^{F_s} |P_1^{(k)}(f) - P_1^N(f)|df$. Denote $(\cdot)_+$ as the positive part of a function, and $(\cdot)_-$ as the negative part. Then,

$$\begin{aligned}
&\max\left\{\int_0^{F_s}\left(P_1^{(k+1)}(f) - P_1^N(f)\right)_+ df, \int_0^{F_s}\left(P_1^{(k+1)} - P_1^N\right)_- df\right\} \\
&\leq \sup\alpha_2(f)\max\left\{\int_0^{F_s}\left(P_2^{(k)} - P_2^N\right)_+ df, \int_0^{F_s}\left(P_2^{(k)} - P_2^N\right)_- df\right\} \\
&\leq \sup\alpha_2(f)\sup\alpha_1(f)\max\left\{\int_0^{F_s}\left(P_1^{(k)} - P_1^N\right)_+ df, \int_0^{F_s}\left(P_1^{(k)} - P_1^N\right)_- df\right\},
\end{aligned}$$

which is a contraction if $\sup\alpha_1(f)\sup\alpha_2(f) = \epsilon_0 < 1$. So, $P_1^{(k)} \to P_1^N$ in $\mathbf{L}_1$-norm as $k \to \infty$.

The above condition may be too restrictive in certain cases. To derive the second and third sufficient conditions, let $\Delta^{(k)}P_1(f) = P_1^{(k)}(f) - P_1^N(f)$. The difference in interference from the interference that would have been caused by the Nash equilibrium power allocation is $\alpha_1(f)\Delta^{(k)}P_1(f)$. This difference in interference would cause a difference in user 2's power allocation by at most $\alpha_1(f)\Delta^{(k)}P_1(f) - \frac{1}{F_s}\int_0^{F_s}\alpha_1(f)\Delta^{(k)}P_1(f)df$. (We are subtracting the mean here because the water-filling process is insensitive to the absolute interference level change, and only to the relative interference level change.) This difference in user 2's power allocation would in turn cause an interference level difference in user 1: $\alpha_2(f)\alpha_1(f)\Delta^{(k)}P_1(f) - \alpha_2(f)\frac{1}{F_s}\int_0^{F_s}\alpha_1(f)\Delta^{(k)}P_1(f)df$. Now, this difference in interference would cause difference in user 1's power allocation by at most an

amount:

$$\begin{aligned}\Delta^{(k+1)}P_1(f) &\leq \alpha_2(f)\alpha_1(f)\Delta^{(k)}P_1(f) - \alpha_2(f)\frac{1}{F_s}\int_0^{F_s}\alpha_1(f)\Delta^{(k)}P_1(f)df - \\ &\quad \frac{1}{F_s}\int_0^{F_s}\alpha_2(f)\alpha_1(f)\Delta^{(k)}P_1(f)df - \frac{1}{F_s}\int_0^{F_s}\alpha_2(f)df\frac{1}{F_s}\int_0^{F_s}\alpha_1(f)\Delta^{(k)}P_1(f)df\end{aligned}$$

The $\mathbf{L}_1$-norm of $\Delta^{(k+1)}P_1(f)$ the above can be bounded by triangular inequality as below:

$$\begin{aligned}\frac{1}{F_s}\int_0^{F_s}|\Delta^{(k+1)}P_1(f)|df &\leq \frac{1}{F_s}\int_0^{F_s}|\alpha_2(f)\alpha_1(f)\Delta^{(k)}P_1(f)|df + \\ &\quad \frac{1}{F_s}\int_0^{F_s}|\alpha_2(f)|df\frac{1}{F_s}\int_0^{F_s}|\alpha_1(f)\Delta^{(k)}P_1(f)|df + \\ &\quad \frac{1}{F_s}\int_0^{F_s}|\alpha_2(f)\alpha_1(f)\Delta^{(k)}P_1(f)|df + \\ &\quad \frac{1}{F_s}\int_0^{F_s}\alpha_2(f)df\frac{1}{F_s}\int_0^{F_s}|\alpha_1(f)\Delta^{(k)}P_1(f)|df \\ &\leq \sup\{\alpha_2(f)\alpha_1(f)\}\frac{1}{F_s}\int_0^{F_s}|\Delta^{(k)}P_1(f)|df + \\ &\quad \sup\{\alpha_1(f)\}\frac{1}{F_s}\int_0^{F_s}\alpha_2(f)df\frac{1}{F_s}\int_0^{F_s}|\Delta^{(k)}P_1(f)|df + \\ &\quad \sup\{\alpha_2(f)\alpha_1(f)\}\frac{1}{F_s}\int_0^{F_s}|\Delta^{(k)}P_1(f)|df + \\ &\quad \sup\{\alpha_1(f)\}\frac{1}{F_s}\int_0^{F_s}\alpha_2(f)df\frac{1}{F_s}\int_0^{F_s}|\Delta^{(k)}P_1(f)|df.\end{aligned}$$

Thus, if

$$\sup\{\alpha_2(f)\alpha_1(f)\} + \sup\{\alpha_1(f)\}\frac{1}{F_s}\int_0^{F_s}\alpha_2(f)df = \epsilon_1 + \epsilon_3 < \frac{1}{2}, \tag{12}$$

the iterative water-filling algorithm becomes a contraction, and $P_1^{(k)}(f) \to P_1^N(f)$ in $\mathbf{L}_1$-norm as $k \to \infty$. The same analysis can be applied to $P_2(f)$, which yields the third condition.

The convergence of the iterative water-filling process implies that the Nash equilibrium is unique. This is because the starting point is arbitrary so in particular, it could be another Nash equilibrium if it were not unique. But each Nash equilibrium is its own fixed point. So, this cannot happen. The stability of the Nash equilibrium also follows from the convergence of the iterative procedure. □

Corollary 1 *If the condition for existence and uniqueness of the Nash Equilibrium is satisfied, then an iterative water-filling algorithm, where in every step, each modem updates its power spectrum density regarding all interference as noise, converges, and it converges to the unique Nash equilibrium from any starting point.*

Proof: This is a direct consequence of the iterative water-filling procedure in the proof for the theorem above. □

The condition of the theorem 1 is not a mere technicality. The following is a simple example where the condition is not satisfied, and iterative water-filling does not converge. Consider a two-user case where there are only two frequencies of concern. Let $\alpha_1(f_1) = \alpha_1(f_2) = \alpha_2(f_1) = \alpha_2(f_2) = 2$. Let power constraints and background noise all be 1. Start with the power allocation $P_1(f_1) = 1$, $P_1(f_2) = 0$. Then successive water-filling of user 2 and user 1 produces a sequence of oscillation that will never converge.

4 Adaptive Power Control

Because of the frequency-selective nature of the DSL channel and the DSL crosstalk coupling, power control algorithms in the DSL environment not only needs to optimally allocate power among different users, they also need to allocate power optimally in the frequency domain. This requirement brings in many extra variables, and makes the design of optimal power control for DSL difficult. However, by concentrating on the competitive optimal power allocations, and assuming that the existence and uniqueness conditions under a total power constraint for Nash equilibrium are satisfied, then total power alone is sufficient to represent all competitive power allocations. Consequently, power control can be done based solely on total power despite the frequency selectivity. This simplifies computation tremendously. Although the competitive optimal solutions are in general not globally optimal, as we will soon show, it still gives impressive improvement compared to the current power back-off algorithms.

We propose the following power control algorithm. The goal is to achieve certain target rates for each user. The adaptive algorithm runs in two stages. The inner stage takes some power constraints for each user as the input, and derives the competitive optimal power allocation and data rates as output. This is accomplished by the iteration water-filling procedure. The procedure works as follows. With a fixed total power constraint for each user, the first user updates its power allocation as the water-filling spectrum of its channel regarding all other user's crosstalk as noise. Water-filling is then successively applied to the second user, the third user, $\cdots$, then again to the first user, second user, etc, until it converges.

The outer stage finds the optimal total power constraint for each user. The outer procedure adjust each user's power based on the outcome of the inner iterative water-filling. If a user's data rate is below its target rate, its power will be increased, unless it is already at the modem power limit in which case its power stays the same. If a user's data rate is much above its target rate, its power will be decreased. If the data rate is only slightly above the target rate, its power will be unchanged. The outer procedure converges when the set of target rates is achieved.

We summarize the algorithm in the following.

Algorithm 1 *Consider a K-user system. Let $\mathbf{P}$ be the modem power limit. Let T_i be the target rate. The power control works as follows:*

Initialize $\mathbf{P}_i = \mathbf{P}$, $i = 1, \ldots K$.
repeat
  repeat
    for $i = 1$ to K
      $N(f) = \sum_{j=1, j\neq i}^{K} |H_{ji}(f)|^2 P_j(f) + \sigma_i(f)$;
      $P_i(f)$ = water-filling spectrum with channel $|H_{ii}(f)|^2$, noise $N(f)$,
        and power constraint $\mathbf{P}_i$
      R_i = data rate on channel $|H_{ii}(f)|^2$ with the power allocation $P_i(f)$,
        and noise $N(f)$
    end
  until the desired accuracy is reached.
  for $i = 1$ to K
    If $R_i > T_i + \epsilon$, set $\mathbf{P}_i = \mathbf{P}_i - \delta$.
    If $R_i < T_i$, set $\mathbf{P}_i = \mathbf{P}_i + \delta$.
    If $\mathbf{P}_i > \mathbf{P}$, set $\mathbf{P}_i = \mathbf{P}$.
  end
until $R_i > T_i$ for all i.

The algorithm is found to work well with the parameters $\delta = 3$dB and ϵ equal to roughly 10% of the target rate. The outer iteration converges only if the set of target rates is achievable. Unfortunately, which set of target rates is achievable cannot be known *a priori*. Thus, a centralized agent with the full knowledge of all channel and interference transfer functions must decide, by running through all possible total power constraints, which sets of target rates can be deployed in a DSL bundle. In effect, we have separated the power allocation problem into two parts. A centralized agent needs to decide on a target rate and a power constraint for each loop in the bundle. Then the loops themselves can undergo the iterative water-filling procedure to arrive at the desired rates without the need of centralized control. The amount of information that needs to be passed by the central control to each loop is small.

Comparing to conventional power control methods, this new method offers two key advantages. First, because the interference is controlled systematically, no power spectral density constraints are needed, thus, allowing a more efficient use of total power by all users. Secondly, because at each stage, a single-user water-filling is performed, which optimizes each user's data rate regarding all other users as noise, the iterative water-filling algorithm offers an opportunity for different loops in a binder to negotiate the use of frequency. Thus, each loop has an incentive to move away from frequency bands when interference is strong, and concentrate on the frequency bands that it can most efficiently utilize.

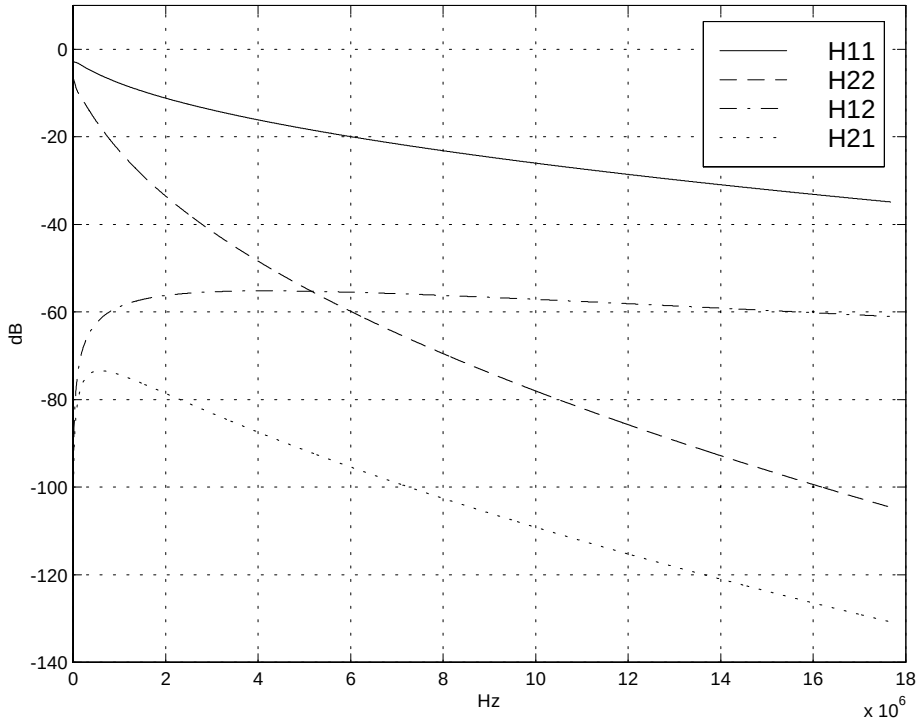


Figure 4: Typical loop coupling and transfer functions: 3000ft vs 1000ft.

5 Performance

A realistic VDSL environment is simulated here. In Figure 4, sample channel and crosstalk transfer functions are plotted for two modems located at 3000 feet and 1000 feet from the central office, respectively. H_{ij} refers to the transfer function from user i to user j. The crosstalk transfer function is computed using the FEXT crosstalk models [15] where cross-coupling increases with frequency as $f^{3/2}$. The 26 gauge, or .4mm lines are modeled here. Observe that in this case, for all $j \neq i$, $\forall f$, $H_{ii}(f) \gg H_{ij}(f)$ by at least 20dB. So, $\forall f$,

$$\alpha_1(f)\alpha_2(f) = \frac{|H_{12}(f)|^2}{\Gamma \cdot |H_{22}(f)|^2} \cdot \frac{|H_{21}(f)|^2}{\Gamma \cdot |H_{11}(f)|^2} < 1, \tag{13}$$

since Γ is about 16dB for an uncoded QAM transmission with 6dB margin. Thus, a Nash equilibrium exists. Further, all three conditions of Theorem 1 turn out to be satisfied. So the Nash equilibrium is unique, and the iterative water-filling procedure converges. In realistic simulations, we have not encountered a case where iterative water-filling does not converge.

A maximum modem power of 11.5dBm is applied to each modem [15] [16]. The usual power spectrum density constraint is not applied, except for below 1.1MHz where the protection of ADSL and other services is guaranteed. A number of non-VDSL crosstalkers are also included. This

includes 10 ADSL, 4 HDSL, and 10 ISDN crosstalks, comprising the so-called noise A model [17]. We evaluated the performance of a binder with 8 VDSL lines, 4 of which are at the same distance 3000 feet from the central office, and the other 4 are also of the same distance L from the central office, with L varies from 500 feet to 2500 feet. The North American frequency plan (so-called Plan 998) [18] is used to separate upstream and downstream. The 998 plan uses the $3.75-5.2$MHz band, $8.5-12.0$MHz band plus an optional $30-138$kHz band for upstream transmission. Frequency bands corresponding to the amateur radio frequencies [15] are notched off.

The iterative water-filling algorithm is applied to the 8-user scenario. Figure 5 shows the convergence speed for the case where two sets of loops are of lengths 1000ft and 3000ft respectively. A total power constraint of -15.5dBm is set for the 1000ft loops, and 11.5dBm is set for the 3000ft loops. The iterative algorithm successively performs the water-filling power allocation for each loop while holding the other 7 loops fixed. As the figure shows, after the first water-filling, the 1000ft loop achieves a rate of 32Mbps because it sees no interference at that point. The subsequent loops get smaller data rate because of the interference from loops that are previously water-filled. However, at 9th water-fill, where the first loop is re-visited, it also drops its data rate in response to other loops power allocations. The algorithm converges after just two water-fillings for each user.

With different total power constraints, many different rate-tuples are achievable. The data rates for each set of 4 users of the same length are the same, so the rate region is two-dimensional. The upstream rate regions are plotted in Figure 6. For example, when 4 lines are at 3000ft, and the other 4 lines are at 500ft, we can achieve a data rate of 7.8Mbps in 3000ft loops and 18Mbps in 500ft loops, or 7Mbps on 3000ft loops and 26Mbps on 500ft loops. The data rate trade-offs are clearly illustrated, and thus, different classes of services can be naturally supported in the same loop configuration depending on the service requirements.

We now compare the performance with the achievable rates using a conventional power back-off method. We chose the reference noise method, with a reference noise equal to the FEXT caused by a 3000ft loop. (This is in effect the same as the equalized-FEXT method.) This way, all loops are forced to emit the same amount of interference as a 3000ft loop, regardless of its actual length. With this power back-off method, the four 3000ft loops always have a capacity of 6.7Mbps each. The performance of the other four loops are tabulated in Table 1.

As the results show, the competitively optimal power allocation method offers a large improvement in performance. This improvement is possible because the new power control algorithm considers all loops in the binder as a whole, taking into account all interactions and globally allocating power to each user. Although the competitively optimal operating points are not necessarily globally optimal, it already offers drastic improvement over current power back-off methods which consider only each loop by itself. These competitively optimal points are easy to find because iterative water-filling converges very fast.

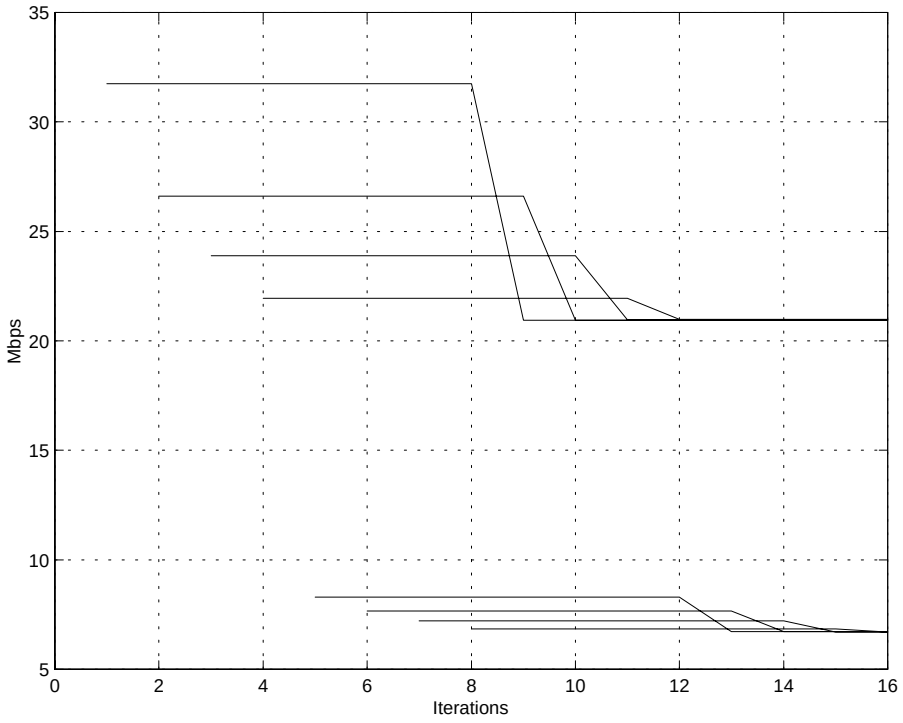


Figure 5: Convergence of iterative water-filling algorithm: four 1000ft loops at -15.5dBm vs four 3000ft loops at 11.5dBm.

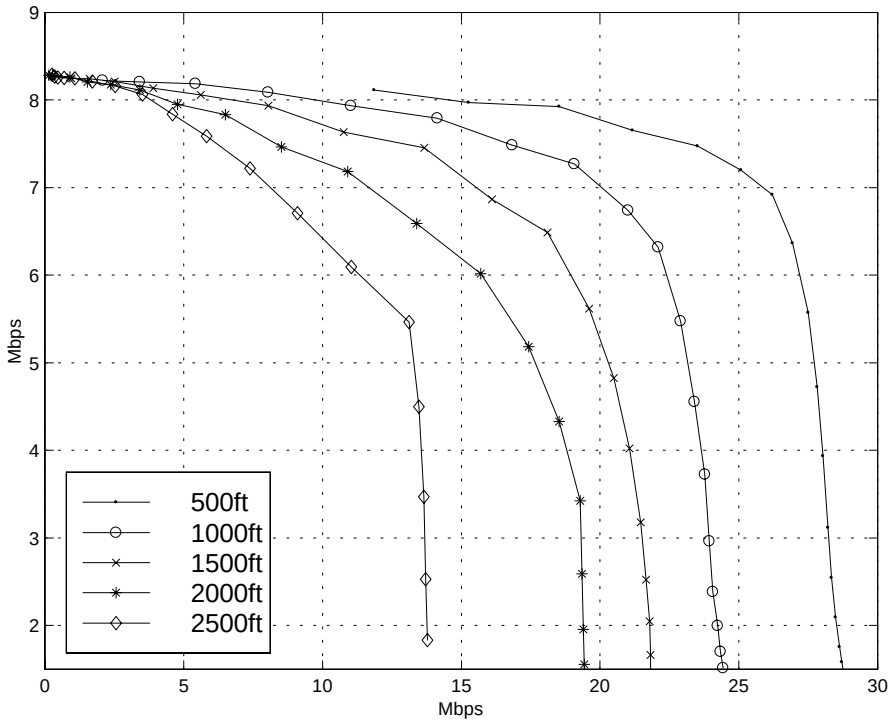


Figure 6: Competitive optimal upstream rate regions in VDSL: 3000ft vs various lengths.

loop length (ft)	reference noise (Mbps)	competitive optimum (Mbps)
500	12.5	26.5
1000	10.1	21.0
1500	8.9	16.5
2000	8.0	12.5
2500	7.3	9.0

Table 1: Upstream data rates of reference noise power back-off vs. competitively optimal power back-off. The reference length is at 3000ft, with data rate 7.8Mbps.

6 Conclusion

This paper considered the problem of optimal power control in a frequency selective multiuser interference network. A notion of competitive optimality is defined in the interference channel, and it represents the Nash equilibrium in the game theoretical sense. Under a set of sufficient conditions, this Nash equilibrium exists, is unique, and is stable. An iterative water-filling algorithm which converges to the competitively optimal point is then proposed. This iterative algorithm is used as the core of an adaptive power control scheme. The new power control scheme uses loop transfer functions and cross-couplings to derive a target data rate for each user, and it offers the loops to negotiate the best use of power and frequency with each other. When applied to the VDSL upstream power back-off problem, this adaptive power control is found to outperform current power back-off schemes such as the reference noise method dramatically.

References

[1] J. Cioffi, V. Oksman, J.-J. Werner, T. Pollet, P. M.P. Spruyt, J. S. Chow, and K. S. Jacobsen, "Very-high-speed digital subscriber lines," *IEEE Communications Magazine*, vol. 37, no. 4, pp. 72–79, April 1999.

[2] G. J. Foschini and Z. Miljanic, "A simple distributed automonous power control algorithm and its convergence," *IEEE Transactions on Vehicular Technology*, vol. 42, no. 4, pp. 641–6, November 1993.

[3] N. Bambos, "Toward power-sensitive network architectures in wireless communications: concepts, issues, and design aspects," *IEEE Personal Communications Magazine*, vol. 5(3), pp. 50–59, June 1998.

[4] W. Yu and J. Cioffi, "Competitive equilibrium in Gaussian interference channels," in *Proceedings International Symposium on Information Theory (ISIT)*, 2000.

[5] T.S. Han and K. Kobayashi, "A new achievable rate region for the interference channel," *IEEE Transactions on Information Theory*, vol. 27, pp. 49–60, January 1981.

[6] M. H. M. Costa, "On the Gaussian interference channel," *IEEE Transactions on Information Theory*, vol. IT-31, pp. 751–760, November 1985.

[7] H. Sato, "The capacity of the Gaussian interference channel under strong interference," *IEEE Transactions on Information Theory*, vol. IT-27, pp. 786–788, November 1981.

[8] K. S. Jacobsen, "Text proposal for upstream power back-off," Amsterdam, Holland, November 1999, ETSI TM6, Contribution TD15.

[9] K. S. Jacobsen, "Upstream power back-off for VDSL - practical considerations," Fiji, January 1999, ITU SG 15, Contribution FI-074.

[10] K. S. Jacobsen, "Methods of upstream power backoff on very high-speed digital subscriber lines," *IEEE Communications Magazine*, pp. 210–6, March 2001.

[11] J. T. Aslanis Jr., *Coding for Communication Channels with Memory*, Ph.D. thesis, Stanford University, 1990.

[12] A. Sendonaris, V. Veeravalli, and B. Aazhang, "Joint signaling strategies for approaching the capacity of twisted-pair channels," *IEEE Transactions on Communications*, vol. 46, no. 5, pp. 673–685, May 1998.

[13] G. Cherubini, E. Eleftheriou, and S. Olcer, "On the optimality of power back-off methods," Vancouver, Canada, August 2000, ANSI T1E1.4, Contribution 235.

[14] D. Fudenberg and J. Tirole, *Game Theory*, MIT Press, 1991.

[15] "Very-high-speed digital subscriber lines: System requirements," November 1998, J. Cioffi (editor), ANSI T1E1.4/98-043R8.

[16] "Very-high-speed digital subscriber line, part 1: Functional requirements," June 1999, ETSI TM6, TS 101 270-1, v.1.1.5.

[17] V. Oksman and J. Cioffi, "Noise models for VDSL performance verification," Clearwater, Florida, December 1999, ANSI T1E1.4, Contribution 99-438R2.

[18] K. McCammon, "G.vdsl: VDSL band plan for North America," Geneva, 2000, ITU Contribution D.715.