An Adaptive Multiuser Power Control Algorithm for VDSL

Wei Yu, George Ginis and John Cioffi *

March 20, 2001

Abstract

This paper considers the optimal power control problem in a multiuser interference network. The frequency selective interference channel is modeled as a non-cooperative game. We establish the existence and uniqueness of a Nash equilibrium in the game, and propose an iterative water-filling algorithm to efficiently reach the Nash equilibrium. The Nash equilibrium points correspond to competitively optimal power allocation schemes in the interference network. An adaptive power control algorithm for upstream VDSL power back-off is developed based on competitive optimality. 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 dramatic performance improvement when compared to existing methods.

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 to 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 data rate requirements for earlier systems are considerably lower than the system capacity.

*Address: Packard Room 360, 350 Serra Mall, Stanford University, Stanford, CA 94305-9515, USA. e-mails:{weiyu,gginis,cioffi}@stanford.edu

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 a central issue. This is especially true for the emerging VDSL standard, where frequency up to 17MHz 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 aspects. Although the DSL environment varies from loops to loops, 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 frequency 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 uses frequency range up to 17.6MHz, which is divided into 4096 tones, with each tone 4.3125kHz wide.

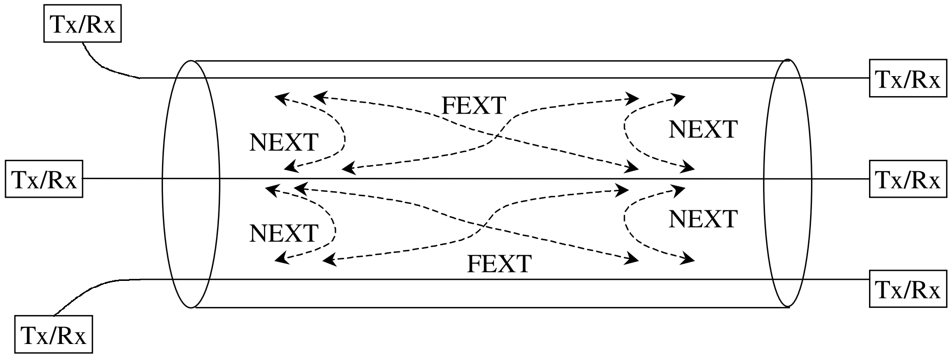

Figure 1: The DSL crosstalk environment.

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 crosstalks created by transmitters located on the same side as the receiver. Far-end crosstalk (FEXT) refers to crosstalks 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 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

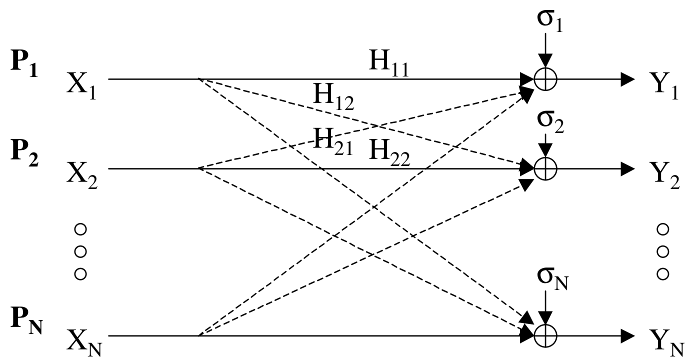

Figure 2: A Gaussian interference network.

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 $0 \le f \le F_s$, $F_s = \frac{1}{2T_s}$, 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:

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

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

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

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 $\{R_1, \cdots R_N\}$ 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, we 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:

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

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. We circumvent the difficulty here 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, 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 especially in view of the near-far problem mentioned before. Figure 3 illustrates a configuration in which two VDSL loops in the same binder emanate from a central site to the customer premises (CP). The central site may be a central office (CO), optical network unit (ONU), or local exchange (LEx). The performance of each VDSL loop, both downstream and upstream is limited by VDSL FEXT. When all transceivers at the central site transmit at the same power spectral density (PSD), then at any distance from the central site, the downstream signal levels on all lines are approximately equal. As a result, at any distance from the central site, FEXT contributions to other lines in the binder are equal. In contrast, if all transmit PSDs in the upstream direction are the same, the level of FEXT at the central site is not the same from line to line due to the variability in the lengths of each loop. In particular, the upstream performance on longer loops is severely compromised by upstream transmission on shorter loops. To remedy this intra-VDSL spectral compatibility problem, upstream transmitters on shorter loops must reduce their transmit power spectral densities (PSDs) such that they do not unfairly compromise the upstream data rates on longer loops. The process of reducing the upstream transmit PSD to improve

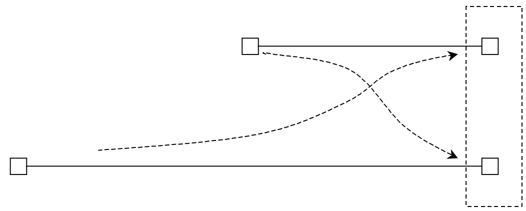

Figure 3: A situation requiring upstream power back-off.

spectral compatibility is known as upstream power back-off (UPBO) [8] [9], and it has been under intensive study in the VDSL standardization bodies.

Several power back-off algorithms have been proposed in VDSL, and an extensive review of these methods can be found in [9]. 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 equalizes the PSD level of a shorter loop with the PSD level of some longer reference loops. 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 [10], or all PSDs are in some sense symmetrical [11]. The first attempt in finding the true global optimum is due to Cherubini [12], where an approach based on quantum annealing is used to minimize the total energy subject to rate constraints on each user. We will take a different approach here. 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, which 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 fixes 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 [13]. In the following, we 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. If we make 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:

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

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

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. Here, we will offer 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 $\epsilon_1 = \sup\{\alpha_1(f)\alpha_2(f)\}$, $\epsilon_2 = \sup\{\alpha_1(f)\}\frac{1}{F_s}\int_0^{F_s}\alpha_2(f)df$, and $\epsilon_3 = \sup\{\alpha_2(f)\}\frac{1}{F_s}\int_0^{F_s}\alpha_1(f)df$. If either $\epsilon_1 < 1$, or $\epsilon_1 + \epsilon_2 < \frac{1}{2}$, or $\epsilon_1 + \epsilon_3 < \frac{1}{2}$, 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 $\sup \alpha_1(f) \sup \alpha_2(f) = \epsilon_1 < 1$. 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:

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

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

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

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

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 need 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, if we concentrate on the competitive optimal power allocations only, and if 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 optimize 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.

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 [14] 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. However, in this example, we do not have $\sup \alpha_1(f) \sup \alpha_2(f) < 1$, $\forall f$, since $\sup\{\alpha_1(f)\}$ is about 40dB$-\Gamma$, and $\sup\{\alpha_2(f)\} < 1$ is about -65dB$-\Gamma$. But, both the second and the third conditions in 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 [14] [15]. The usual power spectrum density constraint is not applied, except for below 1.1MHz where the protection of ADSL

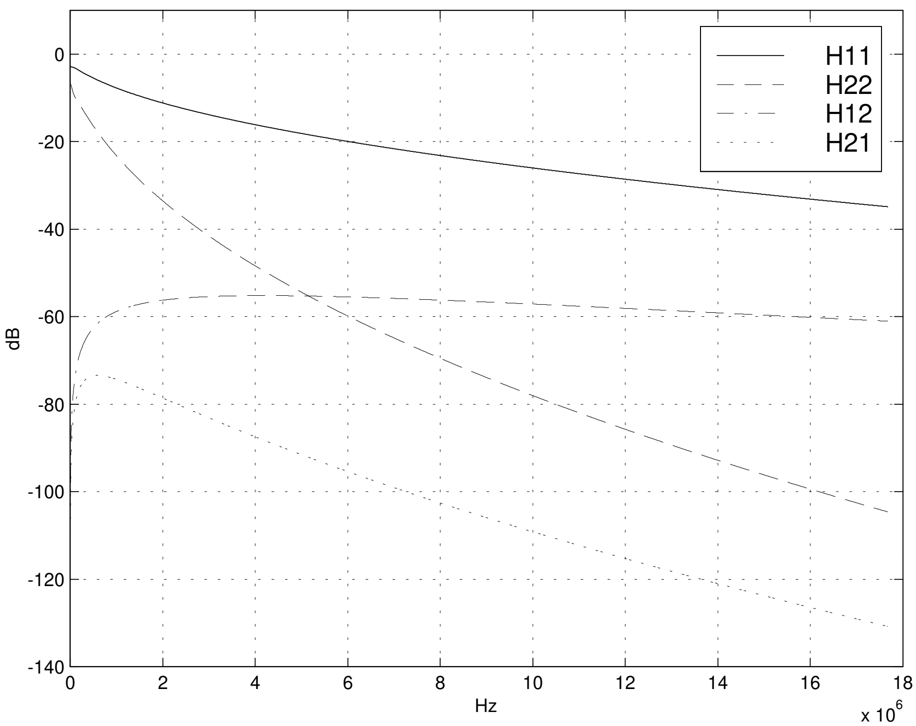

Figure 4: Typical loop coupling and transfer functions: 3000ft vs 1000ft.

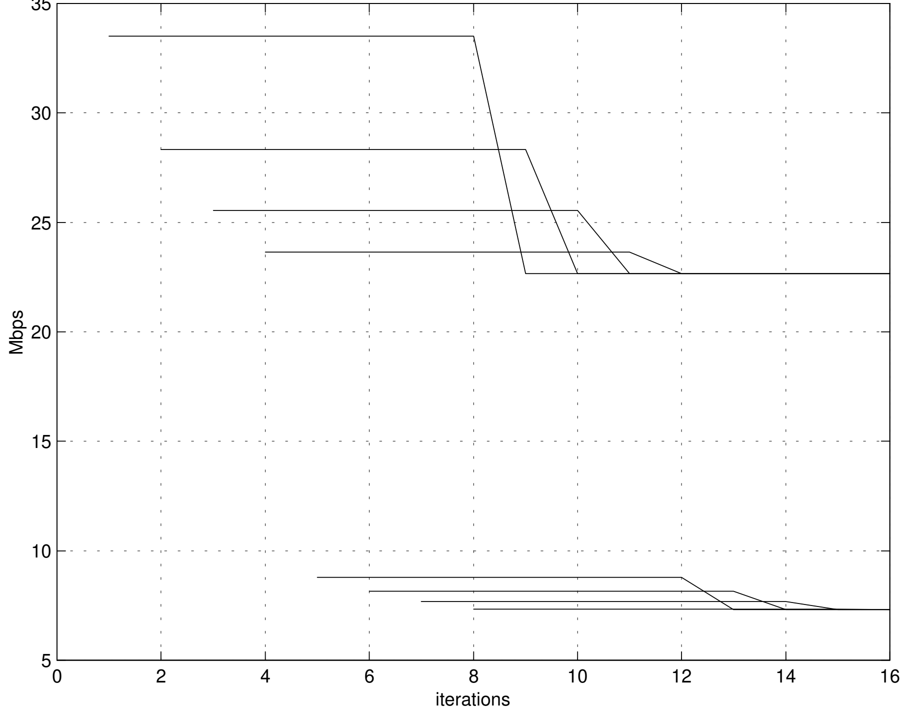

Figure 5: Convergence of iterative water-filling algorithm: four 3000ft loops at -15.5dBm vs four 1000ft loops at 11.5dBm.

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 [16]. 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) [17] 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 [14] 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 34Mbps 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

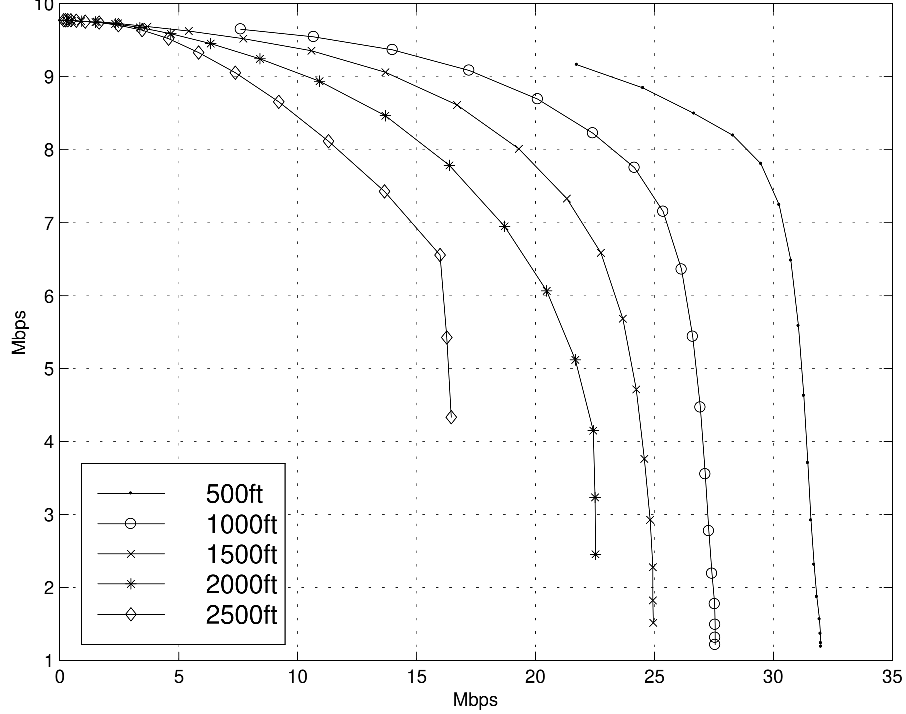

Figure 6: Competitive optimal upstream rate regions in VDSL: 3000ft vs various lengths.

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 9Mbps in 3000ft loops and 22Mbps in 500ft loops, or 7Mbps on 3000ft loops and 30Mbps 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 7.8Mbps 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 dramatic 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

loop length (ft)	reference noise (Mbps)	competitive optimum (Mbps)
500	14.0	28.5
1000	11.5	24.0
1500	10.0	19.5
2000	9.0	15.5
2500	8.5	11.0

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.

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.

6 Conclusion

In this paper, we have considered the problem of optimal power control in a frequency selective multiuser interference network. We defined a notion of competitive optimality. The competitive optimal point is the Nash equilibrium in the interference channel game, and we showed that under a set of sufficient conditions, the Nash equilibrium exists, is unique, and is stable. We then proposed an iterative water-filling algorithm, and demonstrated its convergence to the competitively optimal point under a power constraint. This iterative algorithm is used as the core of an adaptive power control scheme. The new power control scheme considers all loop transfer functions and cross-couplings, 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. This suggests that loop transfer functions and coupling functions are extremely valuable information in VDSL deployment.

References

[1] John Cioffi, Vladimir Oksman, Jean-Jacques Werner, Thierry Pollet, Paul M.P. Spruyt, Jacky S. Chow, and Krista S. Jacobsen, "Very-high-speed digital subscriber lines," *IEEE Communications Magazine*, vol. 37, no. 4, pp. 72–79, April 1999.

[2] Gerard J. Foschini and Zoran 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] Wei Yu and John 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] Max H. M. Costa, "On the Gaussian interference channel," *IEEE Transactions on Information Theory*, vol. IT-31, pp. 751–760, November 1985.

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

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

[11] Andrew Sendonaris, Venugopal Veeravalli, and Behnaam 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.

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

[13] Drew Fudenberg and Jean Tirole, *Game Theory*, MIT Press, 1991.

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

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

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

[17] Kent McCamman, "G.vdsl: VDSL band plan for North America," Geneva, 2000, ITU Contribution D.715.