Distributed Multiuser Power Control for Digital Subscriber Lines

Wei Yu, George Ginis, and John M. Cioffi, *Fellow, IEEE*

***Abstract*—This paper considers the multiuser power control problem in a frequency-selective interference channel. The interference channel is modeled as a noncooperative game, and the existence and uniqueness of a Nash equilibrium are established for a two-player version of the game. An iterative water-filling algorithm is proposed to efficiently reach the Nash equilibrium. The iterative water-filling algorithm can be implemented distributively without the need for centralized control. It implicitly takes into account the loop transfer functions and cross couplings, and it reaches a competitively optimal power allocation by offering an opportunity for loops to negotiate the best use of power and frequency with each other. When applied to the upstream power backoff problem in very-high bit-rate digital subscriber lines and the downstream spectral compatibility problem in asymmetric digital subscriber lines, the new power control algorithm is found to give a significant performance improvement when compared with existing methods.**

***Index Terms*—Asymmetric digital subscriber line (ADSL), dynamic spectral management, game theory, interference channel, iterative water-filling, Nash equilibrium, power control, spectral compatibility, upstream power backoff, very-high bit-rate digital subscriber line (VDSL).**

I. Introduction

POWER control is a central issue in the design of interference-limited multiuser communication systems. 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 involves a performance tradeoff among the different users. Such a tradeoff is the subject of current investigation. This paper considers the digital subscriber line (DSL) system as a multiuser environment. The aim is to develop a power allocation scheme that is able to jointly optimize the performance of 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 can be considered as a multiuser environment because telephone lines from different customer premise sites are bundled together on the way to the central office, and different lines in the bundle create crosstalk interference into each other. Such crosstalk can be the dominant noise source in a loop. However, early DSL systems such as asymmetric digital subscriber line (ADSL) and high-speed digital subscriber line (HDSL) are designed as single-user systems in the sense that the transmit power-spectral-densities for all modems are fixed regardless of the loop environment. This is because single-user systems are considerably easier to design, and the target data rates for earlier systems are well below the system multiuser capacity. As the demand for higher data rates increases, spectral compatibility and power control emerge as central issues for the following two reasons: first, high-speed DSL systems such as very-high bit-rate digital subscriber lines (VDSL) are evolving toward higher frequency bands, where the crosstalk problem is more pronounced. Second, optical network units (ONU) are increasingly being deployed closer to customer premises, and they can potentially emit strong crosstalk into neighboring lines. The goal of this paper is to show that in many cases, a multiuser system design with an optimal power allocation scheme could significantly improve the system performance when compared with a single-user design. Further, simple distributed power allocation schemes that are implementable in existing modems can be used to realize much of the gain.

The power control problem in DSL systems differs from the more widely studied power control problem in wireless systems (e.g., [2]–[5]) in two key aspects. First, although the DSL transmission 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 made here. 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 to each user, but also the allocation of power over frequencies. Nevertheless, power control schemes designed for wireless systems [2], [6], [4] still provide us with considerable insight. In particular, the DSL systems suffer from a near–far problem similar to that in code division multiple acfess (CDMA) systems. The near–far problem arises when two transmitters located in different distances attempt to communicate with the same central office at the same time. When one transmitter is much closer to the central office than the other, the interference coming from the closer transmitter overwhelms the signal from the farther transmitter. The power control algorithm proposed in this paper is capable of overcoming this problem.

The proposed power control algorithm is based on the formulation of the multiuser environment as a noncooperative game. This game-theory point of view has appeared in several recent works that studied the power control problem for wireless networks [7]–[10]. These existing works typically focus on the

Manuscript received April 14, 2001; revised December 18, 2001. This work was supported in part by a Stanford Graduate Fellowship and in part by Alcatel, Fujitsu, Samsung, France Telecom, IBM, Voyan, Sony, and Telcordia. This work was presented in part at Globecom 2001, San Antonio, TX.

The authors are with the Electrical Engineering Department, Stanford University, Stanford, CA 94305-9515 USA (e-mail: weiyu@dsl.stanford.edu; gginis@dsl.stanford.edu; cioffi@dsl.stanford.edu).

Publisher Item Identifier S 0733-8716(02)05384-2.

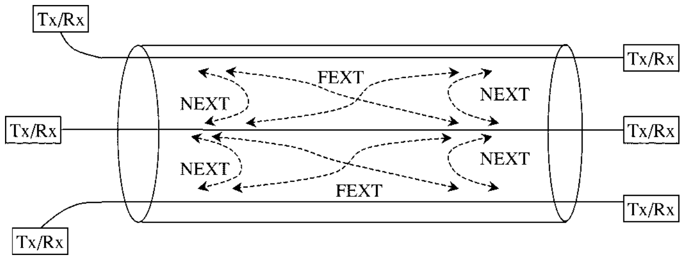

Fig. 1. The DSL crosstalk environment.

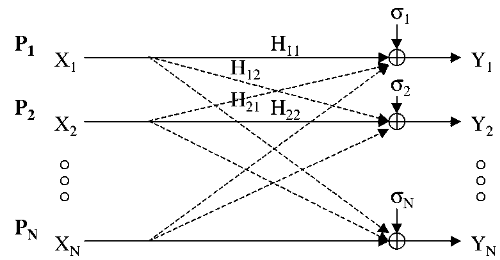

Fig. 2. A Gaussian interference network.

CDMA system, and the power control algorithms studied there considered only flat-fading channels. In a DSL environment, the frequency-selective nature of the channel is crucial and it must be dealt with explicitly. The main result of this paper is that under a wide range of conditions, the frequency-selective interference channel game has a unique Nash equilibrium. This result leads to a power control algorithm based on the concept of competitive optimality, and it further suggests that power control can be implemented distributively and asynchronously with minimal centralized control. Distributed power control schemes have important advantages over centralized schemes, especially as the local access market moves toward "unbundling" where different service providers can potentially share the same binder.

The rest of the paper is organized as follows. Section II reviews the DSL environment and models it as an interference channel. Section III defines and characterizes the competitive equilibrium in such a network and devises an iterative method to achieve the equilibrium. In Section IV, a distributed power allocation method based on the idea of competitive equilibrium is proposed for DSL. System performance is characterized in Section V and conclusions are drawn in Section VI.

II. The DSL Environment

DSL modems use frequencies above the traditional voice band to carry high-speed data. The telephone channels are severely frequency selective. One way to combat intersymbol interference (ISI) is to use discrete multitone (DMT) modulation, which divides the frequency band into a large number of ISI-free subchannels and lets each subchannel carry a separate data stream. This paper considers the DMT modulation scheme as standardized for ADSL in [11] and VDSL in [12]. The use of DMT modulation allows arbitrary power assignment in each frequency tone, thus making spectral shaping easy to realize.

A DSL binder can consist of up to 100 subscriber lines bundled together. Because of their close proximity, the lines create electromagnetic interference into each other, thus causing crosstalk noise (see Fig. 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 end of the line. NEXT is usually much stronger than FEXT. To avoid NEXT, DSL transmission uses either frequency-division duplex (FDD), where all loops transmit in the same direction in every frequency band, or time-division duplex (TDD), where all loops transmit in the same direction in every time slot. North American and European standards use FDD, while TDD is used in Japan. This paper mainly considers frequency-division duplex systems such as the one in [12], where the entire frequency band is divided into four to six 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 as shown in Fig. 2. This model is usually referred to as an interference channel, and it is different from a multiple access or a broadcast channel which models many-to-one or one-to-many transmissions. The multiple access and broadcast models are not suitable for the DSL environment considered here because data coordination among the transmitters or the receivers is generally not possible. Coordination is not possible at the remote terminals because they are located in different geographical locations. Coordination at the central office side is feasible only if all loops are controlled by the same service provider, which might not be the case in an unbundled environment.

From a theoretical point of view, the capacity of the interference channel is still an unsolved problem. Even for the simplest two-user additive white Gaussian interference channel, only partial achievable regions are known [13]. 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 [14], [15]. For this to happen, however, the interference coupling must be stronger than the direct channel, which does not typically occur in a realistic DSL environment. When the interference level is low, multiuser detection becomes difficult and the capacity region is unknown. In this light, this paper restricts attention to transmission techniques where no interference subtraction takes place and focuses solely on the problem of optimal power allocation for each user in the network.

Consider the interference channel model depicted in Fig. 2. There are N transmitters and N receivers in the network. The channel from user i to user j is modeled as a frequency-selective channel, whose transfer function in the frequency domain is denoted as $H_{ij}(f)$, where $0 \le f \le F_s$, $F_s = 1/2T_s$, and T_s is the sampling rate. In addition to the interference, each receiver also experiences background noise, whose power-spectral-density (PSD) 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

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

where Γ denotes the SNR-gap. The SNR-gap defines the gap between a practical coding and modulation scheme and the channel capacity. The SNR-gap depends on the coding and modulation scheme used and also on the target probability of error. (At theoretical capacity, $\Gamma = 0$ dB.) The objective of the system design is to maximize "jointly" the set of rates $\{R_1, \ldots R_N\}$ subject to the power constraints (1). Notice that for each transmitter, increasing its power at any frequency increases its own data rate, but this also increases its interference into other users and is, therefore, detrimental to other users transmissions. Thus, the system design must consider the tradeoff among the data rates of all users, and a single figure of merit is often inadequate to represent the system performance. For example, it is not enough to consider just the maximization of the sum rate, because it does not guarantee a minimal data rate for 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, creating inherent unfairness. As realistic DSL deployments may require an arbitrary level of service for each user, it is necessary to fully characterize the performance tradeoff among all users. A convenient way to fully represent such a tradeoff is to use the concept of a rate region defined as

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

The rate region characterizes all possible data rate combinations among all users subject to the power constraints.

Despite its attractiveness, the above rate region turns out to be not so easy to compute. Although in theory, the rate 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, the computational complexity of either approach is prohibitively high. This is because the capacity expression is a nonconvex function of power allocations, and, consequently, usual numerical algorithms are capable of finding local maxima only, and not the global maximum. This paper circumvents this difficulty by viewing the interference channel as a noncooperative game. The focus is shifted to a different notion of competitive optimality, which, although strictly smaller than the rate region defined above, nevertheless gives substantial improvement in performance over the current DSL systems.

As mentioned earlier, the current DSL systems are designed as single-user systems. In addition to the power constraint, each user is subject to a PSD constraint. The PSD constraint limits the worst case interference level from any modem, and each modem is designed to withstand the worst case noise. Such a design is conservative in the sense that realistic deployment scenarios often have interference much smaller than the worst case noise. Moreover, the same PSD constraint is applied to all modems regardless of their geographic locations. 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. Fig. 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 with the same PSD, due to 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, the short line must reduce its upstream PSD so that it does not cause excessive interference into the long line [16]–[18]. This reduction of upstream transmit PSD is known as upstream power backoff (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 PSD, the FEXT they create into each other is identical at any fixed distance from CO [17]. This downstream FEXT level is always much weaker than the data signals, so it does not pose a serious problem to downstream transmission.

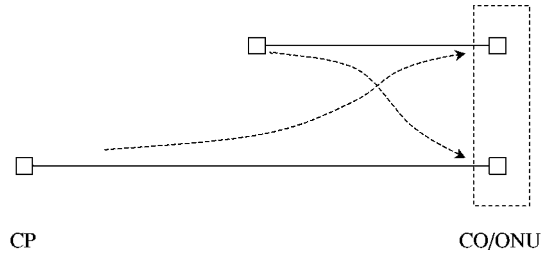

Fig. 3. A situation requiring upstream power backoff.

Several upstream power backoff algorithms have been proposed for VDSL. An extensive review of these methods can be found in [18]. All current power backoff methods attempt to reduce the interference emission of shorter loops by forcing them to emulate the behavior of a longer loop. For example, in the *constant power backoff* method, the PSD is reduced by a constant factor across the 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 the *reference length* method, where variable amount of backoff is implemented across 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 sometimes be too restrictive. This is because in long loops, high-frequency bands usually suffer from much attenuation and are not used, thus, short loops should be able to transmit in those bands without worrying about interference. This observation leads to the *multiple reference length* method, which sets a different reference length for 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 be equal to a more general reference noise.

All previously discussed power backoff methods require the power or noise spectrum of the short loops to comply with a reference loop or a reference noise. Although the optimal choice of reference may not be easy to make, once the references are standardized, these approaches are simple to implement, because they only require each loop to adjust its power spectrum according to a predetermined reference, and thus do not require any knowledge of the network configuration. If, however, the loop characteristics of the network are known in addition, more sophisticated adaptive power control methods can be implemented, and much better performance is possible. For example, if all loop and coupling transfer functions are known to a centralized agent, the optimal power allocation can be precomputed according to some desired criterion, and an appropriate PSD can be assigned to each user. However, finding the optimal power spectrum may be computationally complex, since the optimization problem involves a large number of variables, and due to the nonconvex nature of the problem, many local minima exist. Early attempts in solving this problem have resorted to additional constraints on the PSD [19], [20], and more recent work has focused on advanced techniques such as simulated annealing [21]. Moreover, existing approaches typically require a centralized agent, and they can only be implemented when a single service provider controls the entire bundle. In an unbundled environment where loops within the same bundle may be operated by competing service providers, loop information is typically not shared among the service providers, and it is impractical for a centralized control agent to enforce spectral compatibility.

For the reasons outlined above, this paper mainly focuses on distributed power control algorithms that do not require centralized control. Each line is assumed to have the knowledge of its own channel transfer function and noise profile only, and each DSL modem is allowed to locally optimize its own performance. This locally optimized power control scheme leads to a game-theory interpretation of the interference channel. The locally optimized power allocation is a Nash equilibrium in the game, and it has the intuitive appeal of being the operating point where all users have an incentive to move toward. The Nash equilibrium is computationally easy to characterize, and the power control algorithm offers the following advantages when compared with previous methods.

- The power control algorithm can be implemented distributively without the need of a centralized agent.
- Unlike previous methods that set a PSD level for each 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, it deals with the frequency-selective nature of the channel explicitly.
- The loop transfer functions and cross couplings are implicitly taken into account, and 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 reference noise or reference length.
- Although not globally optimum, the proposed method performs much better than existing methods.

III. Competitive Optimality

The traditional information-theoretic view of an interference channel allows the transmitters, while sending independent data streams, to be cooperative in their respective coding strategies so that interference cancellation may take place in the receivers. If such cooperation cannot be assumed, the interference channel should be alternatively modeled as a noncooperative game. Under this viewpoint, each user competes for data rates with the sole objective of maximizing its *own* performance regardless of all other users. Since each modem has a fixed power budget, the data rate maximization is done by adjusting the power allocation over frequencies. If such power adjustment is done continuously for all users at the same time, it is natural to ask the following question: Can an equilibrium be reached eventually? Such an equilibrium is called a Nash equilibrium, and it is defined as a strategy profile in which each player's strategy is an optimal response to the other player's strategy [22]. The goal of this section is to characterize the Nash equilibrium in the Gaussian interference channel game and to determine conditions for its existence and uniqueness in realistic channels.

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

$$\mathbf{y}_1 = \mathbf{x}_1 + \mathcal{A}_2\mathbf{x}_2 + \mathbf{n}_1 \quad (4)$$

$$\mathbf{y}_2 = \mathbf{x}_2 + \mathcal{A}_1\mathbf{x}_1 + \mathbf{n}_2 \quad (5)$$

where $\mathbf{x}_i$ and $\mathbf{y}_i$ are user i's input and output signals, respectively. The channel transfer functions are normalized to unity. $\mathcal{A}_1$ and $\mathcal{A}_2$ are the interference channels, and their transfer functions are denoted as $\alpha_1(f)$ and $\alpha_2(f)$, respectively. $\mathbf{n}_1$ and $\mathbf{n}_2$ are additive noise with power-spectral-densities $N_1(f)$ and $N_2(f)$, respectively. The two senders are considered as two players in a game. The structure of the game, i.e., the interference coupling functions and the noise power, are assumed to be common knowledge to both players. The strategies for the two players are the transmit power spectra $P_1(f)$ and $P_2(f)$, with power constraints $\int_0^{F_s} P_1(f)\,df \le \mathbf{P}_1$, and $\int_0^{F_s} P_2(f)\,df \le \mathbf{P}_2$, respectively. (Only deterministic or pure-strategies are considered here.) The payoffs for the players are the respective data rates. Under the simplifying assumption that no interference subtraction is performed regardless of interference strength, the data rates are

$$R_1 = \int_0^{F_s} \log_2 \left(1 + \frac{P_1(f)}{N_1(f) + \alpha_2(f)P_2(f)}\right) df \quad (6)$$

$$R_2 = \int_0^{F_s} \log_2 \left(1 + \frac{P_2(f)}{N_2(f) + \alpha_1(f)P_1(f)}\right) df. \quad (7)$$

Comparing the above expression with (2), it is easy to identify

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

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

and similarly for $N_2(f)$ and $\alpha_1(f)$. Thus, the simplified model incurs no loss of generality. The interference channel game considered here is not a zero-sum game, i.e., one player's loss is not equal to the other player's gain.

The main objective here is to characterize all pure-strategy Nash equilibria in an interference channel game. At a Nash equilibrium, each user's strategy is the optimal response to the other player's strategy. So fixing $P_2(f)$, the optimal $P_1(f)$ must be the solution to the following optimization problem:

$$\begin{aligned} \max \quad & \int_0^{F_s} \log_2 \left(1 + \frac{P_1(f)}{N_1(f) + \alpha_2(f)P_2(f)}\right) df \\ \text{s.t.} \quad & \int_0^{F_s} P_1(f)\, df \leq \mathbf{P_1} \\ & P_1(f) \geq 0, \qquad \forall f. \end{aligned} \quad (10)$$

The solution to this problem is the well-known water-filling power allocation. More precisely, let $\tilde{N}(f) = N_1(f) + \alpha_2(f)P_2(f)$. Then, the water-filling power allocation is

$$P_1(f) = \begin{cases} 0, & \text{if } \tilde{N}(f) \geq K_1 \\ K_1 - \tilde{N}(f), & \text{if } \tilde{N}(f) \leq K_1 \end{cases} \quad (11)$$

where K_1 is a constant chosen so that the power constraint is met. Likewise, fixing $P_1(f)$, the optimal $P_2(f)$ should also be a water-filling power allocation against the combined interference from $P_1(f)$ and the noise. Thus, a Nash equilibrium is reached if and only if the water-filling condition is simultaneously achieved for both users. The characterization of Nash equilibria is therefore equivalent to a characterization of "simultaneous water-filling" points. The idea of simultaneous water-filling is illustrated in Fig. 4. The following theorem 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 two-user Gaussian interference game exists. Further, let $\epsilon_0 = \sup\{\alpha_1(f)\}\sup\{\alpha_2(f)\}$, $\epsilon_1 = \sup\{\alpha_1(f)\alpha_2(f)\}$, $\epsilon_2 = \sup\{\alpha_1(f)\}(1/F_s)\int_0^{F_s} \alpha_2(f)\, df$, and $\epsilon_3 = \sup\{\alpha_2(f)\}(1/F_s)\int_0^{F_s} \alpha_1(f)\, df$. If any of the following conditions, $\epsilon_0 < 1$, $\epsilon_1 + \epsilon_2 < 1/2$, or $\epsilon_1 + \epsilon_3 < 1/2$ is satisfied, then the Nash equilibrium is unique and is stable.

The proof of Theorem 1 is lengthy and it is included in the Appendix. The basic idea is that under suitable conditions, the Nash equilibrium can be reached by an iterative water-filling procedure, where each user successively optimizes his power spectrum while regarding other users' interference as noise. The main purpose of Theorem 1 is to characterize conditions under which such an iterative water-filling procedure converges. The following corollary is a direct consequence of the proof.

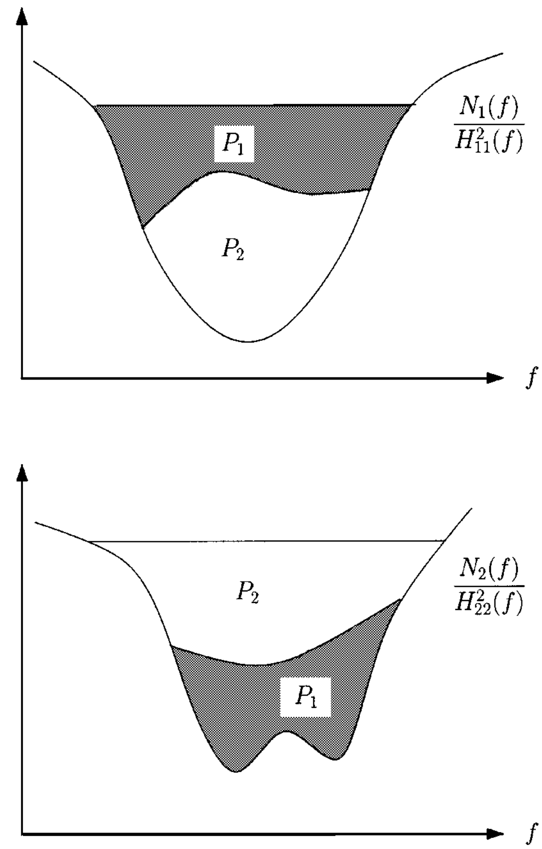

Fig. 4. Simultaneous water-filling.

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

The condition of Theorem 1 is not a mere technicality. The following simple example illustrates a case where the Nash equilibrium is not unique. 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. The set of power allocations $P_1(f_1) = P_2(f_2) = 1$ and $P_1(f_2) = P_2(f_1) = 0$ is one Nash equilibrium, and the set of power allocations $P_1(f_1) = P_2(f_2) = 0$ and $P_1(f_2) = P_2(f_1) = 1$ is a different Nash equilibrium.

IV. DISTRIBUTED POWER CONTROL

Because of the frequency-selective nature of the DSL channel, power control algorithms for DSL applications need to allocate power optimally not only among different users, but also in the frequency domain. This requirement brings in many extra

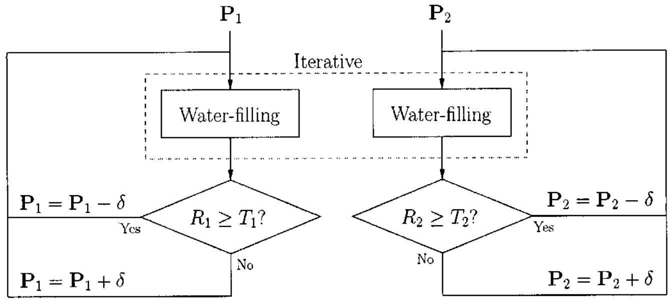

Fig. 5. A simplified illustration of the distributed power control algorithm based on iterative water filling.

variables and makes the design of optimal power control for DSL challenging. However, if only competitively optimal power allocations are considered, the total power alone is sufficient for power control purposes. Assuming that the condition for Theorem 1 is satisfied, then for each set of power constraints, a unique Nash equilibrium exists. Thus, it is possible to reach all possible competitive equilibria by adjusting the total power constraints for each user in spite of the frequency selectivity. This simplifies the computational complexity tremendously. Although the competitively optimal solutions are in general not globally optimal, it will soon be shown that they give significant improvements over current methods.

The following is a description of the proposed power control algorithm. The goal is to achieve a set of target rates for all the users. The adaptive algorithm runs in two stages. The inner stage takes a set of power constraints for each user as the input and derives the competitively optimal power allocation and its associated data rates as output. This is accomplished by the iterative water-filling procedure, which works as follows: with a fixed total power constraint for each user, the first user updates its power allocation by deriving a water-filling spectrum while regarding all other user's crosstalk as noise. Water-filling is then successively applied to the second user, the third user, etc., then applied again to the first user, second user, etc., until the process converges.

The outer stage finds the optimal total power constraint for each user. The outer procedure adjusts 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 is increased, unless this exceeds the power constraint. If a user's data rate is much above its target rate, its power is decreased. If the data rate is only slightly above the target rate, its power remains unchanged. The outer procedure converges when the set of target rates is achieved.

The algorithm is summarized in the following, and a simplified illustration is shown in Fig. 5:

Algorithm 1 Consider a K-user system. Assume a common power constraint $\mathbf{P}$ for all users. Let T_i be the target rate for user i. The power control algorithm works as follows:

Initialize $\mathbf{P}_i = \mathbf{P}$, $P_i(f) = 0$, $i = 1, \ldots K$.
repeat
  repeat
    for $i = 1$ to K,
      $N(f) = \sum_{j=1,\, j\neq i}^{K} |H_{ji}(f)|^2 P_j(f) + \sigma_i(f)$;
      Set $P_i(f)$ to be the water-filling spectrum with noise $N(f)$ and total power $\mathbf{P}_i$.
      Set R_i to be the resulting data rate.
    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 inner loop is the iterative water-filling procedure. Although Theorem 1 only gives the condition for its convergence for the two-user case, it is observed in practice that iterative water-filling converges for arbitrary number of users. The outer iteration is an *ad hoc* method to find the appropriate power constraint for each user so that the competitive equilibrium corresponding to the power constraint satisfies the target rate requirement. Since data rates are monotonic functions of total power, the linear adjustment used in the above algorithm converges as long as the set of target rates is achievable. The algorithm is found to work well with parameters $\delta = 3$ dB and ϵ equal to roughly 10% of the target rate.

The outer loop of the power control algorithm essentially attempts to find the minimum amount of power that is needed to support the target data rate. In fact, the inner and outer loops can be combined. The usual water-filling maximizes the achievable data rate under a fixed power constraint. This is referred to as a "rate-adaptive" water-filling. On the other hand, a "margin-adaptive" water-filling minimizes the total transmission power subject to a fixed rate constraint. The proposed algorithm can be alternatively thought of as each user doing "margin-adaptive" water-filling against each other. Most ADSL modems deployed today already have the capability to do margin adaptation. As long as the set of target rates are achievable, the set of desirable competitively optimal power allocations can be reached using margin adaptation without the need for centralized control. In fact, the iterative procedure can even be done asynchronously, because the Nash equilibrium points are stable. Thus, the proposed power control algorithm is remarkably easy to implement in practice.

To truly implement the proposed power control algorithm distributively, each user must know its target data rate *a priori*. It is important for the target rates to be within the achievable

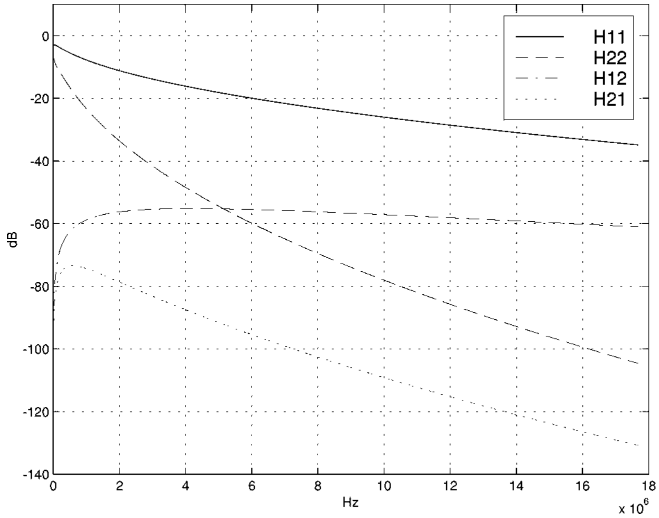

Fig. 6. Typical loop coupling and transfer functions: 3000 versus 1000 ft.

rate region, as otherwise some or all of the users would operate with negative margin. Unfortunately, the set of achievable target rates cannot be determined distributively. Some 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. However, this usually occurs in the loop planning stage, and it needs to be done only once.

Compared with conventional methods, the key advantage of this new power control algorithm is the following: the iterative water-filling algorithm offers an opportunity for different loops in a binder to negotiate the best use of frequency with each other. Thus, each loop has an incentive to move away from frequency bands where interference is strong and concentrate on the frequency bands that it can most efficiently utilize. This method of controlling the interference removes the arbitrary PSD constraint, and it is able to bring a large overall improvement in system performance.

V. Performance

The performance of the distributed power control algorithm is examined in this section. The first example is taken from the upstream power backoff problem in VDSL. A realistic VDSL environment is simulated here. Fig. 6 shows an example of channel and crosstalk transfer functions for two modems located at 3000 and 1000 ft from the central office. H_{ij} refers to the transfer function from user i to user j. The crosstalk transfer function is computed using the FEXT crosstalk models [12] where cross coupling increases with frequency as $f^{3/2}$. The 26-gauge, or 0.4-mm lines, are modeled here. Observe that at high frequencies, the crosstalk transfer function is actually larger than the direct channel. However, it is always true that

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

where Γ is about 16 dB for an uncoded QAM transmission with 6-dB margin. Further, in the frequency range of interest (up to 12 MHz), all three conditions of Theorem 1 are satisfied. So, if the binder consists of two users only, it would have a unique Nash equilibrium, and the Nash equilibrium could be reached by the iterative water-filling process. Although this example considers a realistic deployment scenario with eight VDSL lines, our experience has been that iterative water-filling always converges in DSL channels regardless of the number of lines.

A maximum transmit power of 11.5 dBm is applied to each modem [12]. The usual PSD constraint is not applied, except for below 1.1 MHz where the protection of ADSL and other services is guaranteed. A number of non-VDSL disturbers are also included. This includes 10 ADSL, 4 HDSL, and 10 ISDN disturbers, comprising the so-called noise A model [23]. The loop topology is shown in Fig. 7. It consists of eight VDSL lines, four of which are at the same distance 3000 ft from the central office, and the other four are also at the same distance L from the central office, with L varying from 500 to 2500 ft. The North American frequency plan (so-called Plan 998) [24] 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 [12] are notched off.

The iterative water-filling algorithm is applied to the eight-user scenario. Fig. 8 shows the convergence behavior for the case where two sets of loops are at lengths 1000 and 3000 ft, respectively. A total power constraint of -15.5 dBm is set for the 1000-ft loops, and 11.5 dBm is set for the 3000-ft loops. The iterative algorithm successively performs the water-filling power allocation for each loop, while holding the other seven loops fixed. As the figure shows, after the first water-filling, a 1000-ft loop achieves a data rate of 32 Mbps as it sees no interference at this point. The subsequent loops achieve smaller data rates due to the interference coming from loops that were previously water filled. At the ninth water-fill, the first loop is revisited. 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 power constraints, many different rate-tuples are achievable. The set of all possible rate-tuples is the rate region. Fig. 9 shows the upstream rate regions for the eight VDSL loops. The data rates within each set of four users of the same length are the same, so the rate region can be plotted in a two-dimensional graph. Each curve in the figure corresponds to a different loop topology. The outermost curve corresponds to the topology with four lines at 500 ft and four lines at 3000 ft. The next curve corresponds to the topology with four lines at 1000 ft and four lines at 3000 ft, etc. The rate region illustrates the data rate tradeoff among the users. For example, with four lines at 500 ft and the other four lines at 3000 ft, it is possible to achieve 18 Mbps in 500-ft loops and 7.8 Mbps in 3000-ft loops, or 26 Mbps on 500-ft loops and 7 Mbps on 3000-ft loops, etc. This ability to provide many classes of different service levels is inherent in the proposed power control method.

The proposed power control algorithm also compares favorably with existing power backoff methods. As an example, consider the reference noise method, where the reference noise is chosen to be equal to the FEXT caused by a 3000-ft loop. This choice of the reference noise forces all loops to emit the same

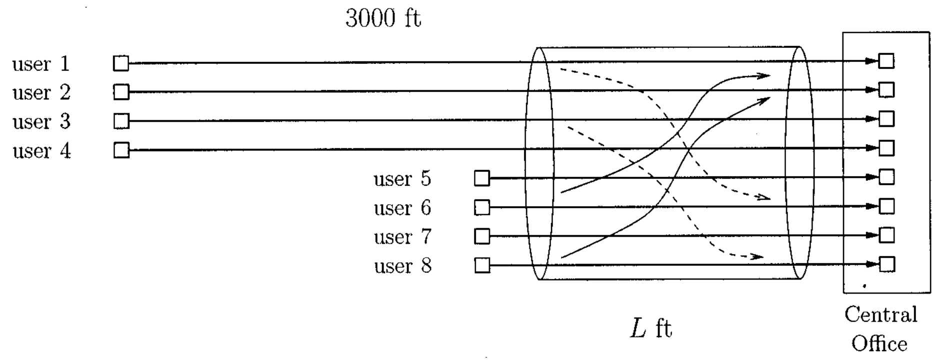

Fig. 7. Loop topology for VDSL upstream power backoff. (L = 500–2500 ft.)

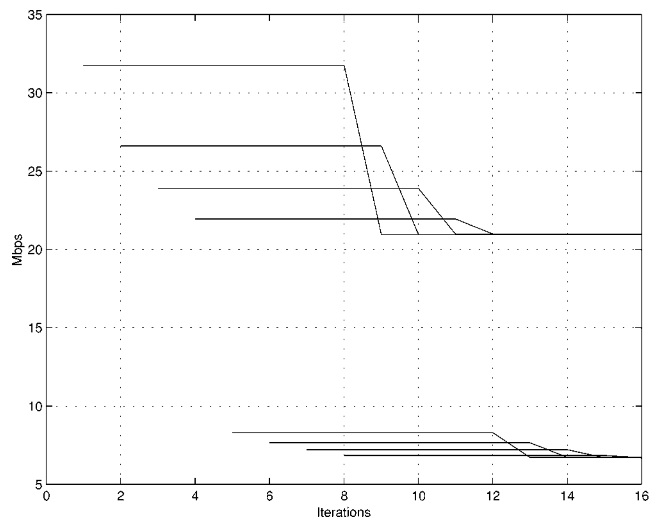

Fig. 8. Convergence of iterative water-filling algorithm. The upper four lines correspond to four 1000-ft loops at −15.5 dBm. The lower four lines correspond to four 3000-ft loops at 11.5 dBm.

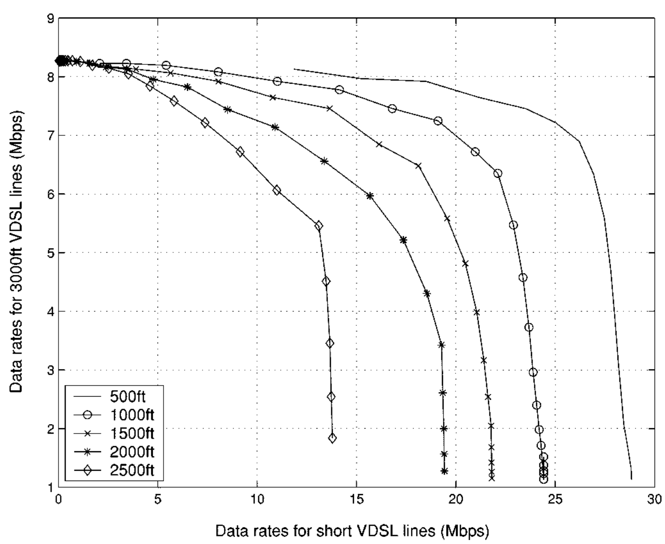

Fig. 9. Competitive optimal upstream rate regions in VDSL: 3000 versus various lengths.

amount of interference as a 3000-ft loop, regardless of their actual lengths. (This is also called the equalized-FEXT method.)

TABLE I
COMPARISON OF UPSTREAM DATA RATES BETWEEN REFERENCE NOISE POWER BACKOFF AND COMPETITIVELY OPTIMAL POWER CONTROL. THE REFERENCE LENGTH IS AT 3000 ft, WITH A DATA RATE OF 6.7 Mbps

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

Using this reference noise, the four 3000-ft loops always achieve a data rate of 6.7 Mbps each. Table I tabulates the performance of the other four loops. As the results show, the competitively optimal power allocation, although not globally optimal, nevertheless offers a significant improvement in performance over current static spectrum management methods. This improvement is possible because the iterative water-filling algorithm implicitly takes into account the interaction among the loops.

Fig. 10 shows a second example where the iterative water-filling algorithm is used in an ADSL scenario. This deployment configuration consists of a central office (CO)-based ADSL line residing in the same binder as a remote terminal (RT)-based ADSL line, and it is increasingly common as service providers install ONU to be located close to customer homes in order to enlarge the service area. However, as downstream transmitters are now located in geographically different places, downstream power control becomes necessary. This is true in the example illustrated, as the RT-based ADSL emits strong downstream interference into the CO-based ADSL. In fact, without power control, the CO-based ADSL does not function at all. Fig. 11 illustrates the rate region for the two ADSL lines when the power control algorithm based on iterative water-filling is used. Again, a graceful tradeoff between the two lines is possible, and a data rate of 1.4 Mbps can be supported in both lines.

VI. CONCLUSION

This paper considers the power control problem in a frequency-selective multiuser interference channel. The interference channel is viewed as a noncooperative game, and the Nash equilibrium of the game is characterized under a set of

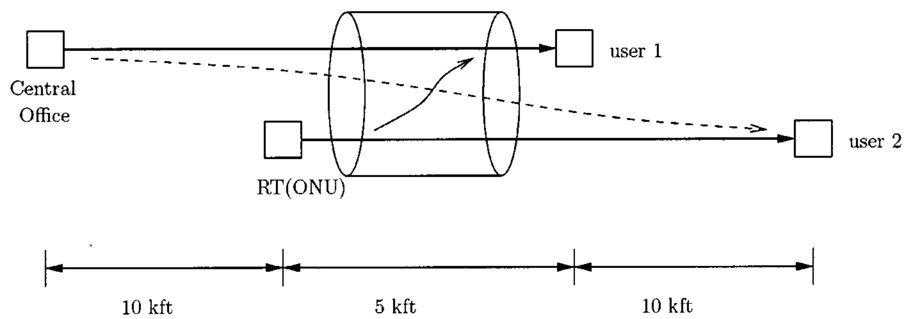

Fig. 10. CO-based ADSL versus RT-based ADSL.

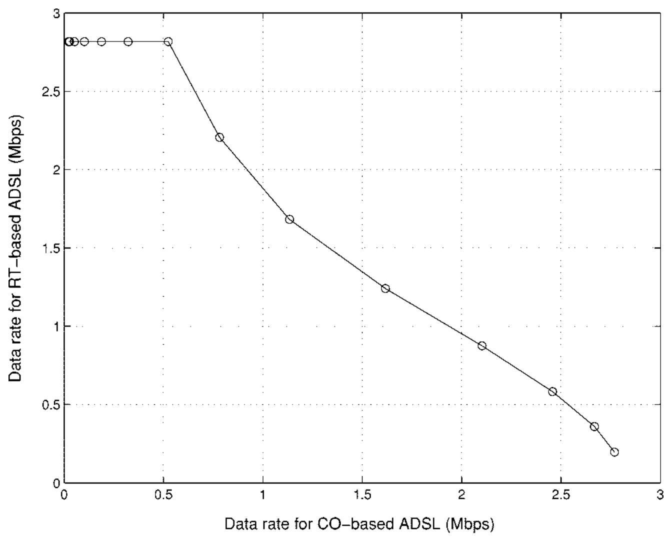

Fig. 11. Competitive optimal downstream rate regions for two ADSL loops.

sufficient conditions. The Nash equilibrium corresponds to a set of competitively optimal power allocations, and it leads to a distributed power control algorithm based on iterative water-filling. The iterative water-filling algorithm implicitly takes into account the loop transfer functions and cross couplings, and it allows the loops to negotiate the best use of power and frequency with each other. The new power control algorithm does not require centralized control, and it is implementable in today's existing modems. When applied to the VDSL upstream power backoff problem and the ADSL spectral compatibility problem, the new method is shown to provide a significant performance improvement when compared with existing methods.

APPENDIX
PROOF OF THEOREM 1

The Nash equilibrium points correspond to power allocations where each user's power spectrum is a water-filling against the combined interference and noise. Call the water-filling level at the Nash equilibrium (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, the following must be true:

$$P_1(f) + \alpha_2(f)P_2(f) + N_1(f) = K_1 \tag{13}$$

$$P_2(f) + \alpha_1(f)P_1(f) + N_2(f) = K_2. \tag{14}$$

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 $\alpha_1(f) > (K_2 - N_2(f))/(K_1 - N_1(f))$, setting $P_1(f) = 0$, and $P_2(f) = K_2 - N_2(f)$ satisfies the condition. Also, if $\alpha_2(f) > (K_1 - N_1(f))/(K_2 - N_2(f))$, 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 (13) and (14) will have a positive solution

$$P_1(f) = \frac{(K_2 - N_2(f)) - \alpha_1(K_1 - N_1(f))}{1 - \alpha_1\alpha_2}$$

$$P_2(f) = \frac{(K_1 - N_1(f)) - \alpha_2(K_2 - N_2(f))}{1 - \alpha_1\alpha_2}. \tag{15}$$

Thus, under all cases, the simultaneous water-filling condition has a solution, and the solution is 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 $\mathbf{P}_{K_1} \le \mathbf{P}_{K_1'}$ and $\mathbf{P}_{K_2} \ge \mathbf{P}_{K_2'}$. 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 $\mathbf{P}_{K_1} \le \mathbf{P_1}$ by the previous 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_1s and K_2s cannot go to infinity with finite power constraints, so they must converge. 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 was 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. Continue iteratively to obtain $P_2^{(1)}(f) \rightarrow P_1^{(2)}(f) \rightarrow P_2^{(2)}(f) \rightarrow \cdots$. We show next that this iterative water-filling process converges in $\mathbf{L}_1$-norm, $(1/F_s)\int_0^{F_s} |P_1^{(k)}(f) - P_1^N(f)|\,df$. Define $(\cdot)^+ = \max(0, \cdot)$, and $(\cdot)^- = -\min(0, \cdot)$. Then

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

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

The above condition may be too restrictive in certain cases. To derive the second and third sufficient conditions, let $\Delta^{(k)}P_1(f) = P_1^{(k)}(f) - P_1^N(f)$. Comparing with the interference emitted by $P_1^N(f)$, the power distribution $P_1^{(k)}(f)$ causes a difference in the interference level that is equal to $\alpha_1(f)\Delta^{(k)}P_1(f)$. This difference in interference would cause a difference in user 2's power allocation by at most $\alpha_1(f)\Delta^{(k)}P_1(f) - (1/F_s)\int_0^{F_s}\alpha_1(f)\Delta^{(k)}P_1(f)\,df$. (The mean is subtracted here, because the water-filling process is insensitive to the absolute interference level change and is only affected by the relative interference level change.) This difference in user 2's power allocation would in turn cause an interference level difference in user 1: $\alpha_2(f)\alpha_1(f)\Delta^{(k)}P_1(f) - \alpha_2(f)(1/F_s)\int_0^{F_s}\alpha_1(f)\Delta^{(k)}P_1(f)\,df$. Now, this difference in interference would cause difference in user 1's power allocation by at most an amount

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

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

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

Thus, if

$$\begin{aligned}
&\sup\{\alpha_2(f)\alpha_1(f)\} + \sup\{\alpha_1(f)\}\frac{1}{F_s}\int_0^{F_s}\alpha_2(f)\,df\\
&\qquad\qquad = \epsilon_1 + \epsilon_3 < \tfrac{1}{2}
\end{aligned}\qquad (16)$$

the iterative water-filling algorithm is a contraction, and $P_1^{(k)}(f) \rightarrow P_1^N(f)$ in $\mathbf{L}_1$-norm as $k \rightarrow \infty$. 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 iterative water-filling process converges to the same Nash equilibrium from any starting point. But each Nash equilibrium is its own fixed point. So, there could not have been more than one Nash equilibria. The stability of the Nash equilibrium also follows from the convergence of the iterative procedure. □

REFERENCES

[1] T. Starr, J. M. Cioffi, and P. J. Silverman, *Understanding Digital Subscriber Line Technology*. Englewood Cliffs, NJ: Prentice-Hall, 1999.

[2] G. J. Foschini and Z. Miljanic, "A simple distributed automonous power control algorithm and its convergence," *IEEE Trans. Veh. Technol.*, vol. 42, pp. 641–646, Nov. 1993.

[3] R. D. Yates, "A framework for uplink power control in cellular radio systems," *IEEE J. Select. Areas Commun.*, vol. 13, pp. 1341–1347, Sept. 1995.

[4] S. V. Hanly, "An algorithm for combined cell-site selection and power control to maximize cellular spread spectrum capacity," *IEEE J. Select. Areas Commun.*, vol. 13, no. 7, pp. 1332–1340, Sept. 1995.

[5] N. Bambos, "Toward power-sensitive network architectures in wireless communications: Concepts, issues, and design aspects," *IEEE Personal Commun. Mag.*, vol. 5, pp. 50–59, June 1998.

[6] R. D. Yates and C.-Y. Huang, "Integrated power control and base station assignment," *IEEE Trans. Veh. Technol.*, vol. 44, pp. 638–644, Apr. 1995.

[7] H. Ji and C.-Y. Huang, "Non-cooperative uplink power control in cellular radio systems," *Wireless Networks*, vol. 4, no. 3, pp. 233–240, 1998.

[8] S.-J. Oh, T. L. Olsen, and K. M. Wasserman, "Distributed power control and spreading gain allocation in CDMA data networks," in *Proc. IEEE Conf. Comp. Communications (INFOCOM)*, 2000, pp. 379–385.
[9] N. Feng, N. B. Mandayam, and D. J. Goodman, "Joint power and rate optimization for wireless data services based on utility functions," presented at the Proc. Conf. Information Science Systems (CISS), 1999.
[10] A. B. MacKenzie and S. B. Wicker, "Game theory and the design of self-configuring, adaptive wireless networks," *IEEE Commun. Mag.*, vol. 39, pp. 126–131, Nov. 2001.
[11] "Asymmetrical digital subscriber line (ADSL)," ITU, ITU-G.992.1, July 1999.
[12] "Very-high bit-rate digital subscriber lines (VDSL): Part 1: Functional requirements and common specification. Part 3: Technical specification of a multi-carrier modulation transceiver," ANSI, ANSI - T1E1.4/2002-031R1, 2001-013R2.
[13] T. S. Han and K. Kobayashi, "A new achievable rate region for the interference channel," *IEEE Trans. Inform. Theory*, vol. 27, pp. 49–60, Jan. 1981.
[14] A. Carleial, "Interference channels," *IEEE Trans. Inform. Theory*, vol. 24, pp. 60–70, Jan. 1978.
[15] H. Sato, "The capacity of the Gaussian interference channel under strong interference," *IEEE Trans. Inform. Theory*, vol. IT–27, pp. 786–788, Nov. 1981.
[16] K. S. Jacobsen, "Text proposal for upstream power back-off," ETSI, ETSI-TM6/TD15, Nov. 1999.
[17] ——, "Upstream power back-off for VDSL—Practical considerations," ITU, ITU SG-15/FI-074, Jan. 1999.
[18] ——, "Methods of upstream power backoff on very high-speed digital subscriber lines," *IEEE Commun. Mag.*, pp. 210–216, Mar. 2001.
[19] J. T. Aslanis, Jr., "Coding for communication channels with memory," Ph.D. dissertation, Stanford Univ., Stanford, CA, 1990.
[20] A. Sendonaris, V. Veeravalli, and B. Aazhang, "Joint signaling strategies for approaching the capacity of twisted-pair channels," *IEEE Trans. Commun.*, vol. 46, pp. 673–685, May 1998.
[21] G. Cherubini, E. Eleftheriou, and S. Olcer, "On the optimality of power back-off methods," ANSI, ANSI - T1E1.4/235, Aug. 2000.
[22] D. Fudenberg and J. Tirole, *Game Theory*. Cambridge, MA: MIT Press, 1991.
[23] V. Oksman and J. M. Cioffi, "Noise models for VDSL performance verification," ANSI, ANSI - T1E1.4/99-438R2, Dec. 1999.
[24] K. McCammon, "G.vdsl: VDSL band plan for North America," ITU, ITU Contribution D.715, 2000.

Wei Yu received the B.A.Sc. degree in computer engineering and mathematics from the University of Waterloo, Waterloo, ON, Canada, in 1997, and the M.S. in electrical engineering from Stanford University, Stanford, CA, in 1998. He is currently pursuing the Ph.D. degree at Stanford University.

His research interests include communication theory, multiuser information theory, OFDM systems, and DSL systems.

George Ginis (S'97) received the Diploma in electrical and computer engineering from the National Technical University of Athens, Athens, Greece, in 1997, and the M.S. and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, in 1998 and 2002, respectively.

His current research interests include multiuser transmission theory, interference mitigation, and their applications to wireline and wireless communications.

John M. Cioffi (S'77–M'78–SM'90–F'96) received the B.S.E.E. degree from the University of Illinois, Urbana-Champaign, in 1978, and the Ph.D.E.E. degree from Stanford University, Stanford, CA, in 1984.

He was with Bell Laboratories, Holmdel, NJ, from 1978 to 1984 and IBM Research, San Jose, CA, from 1984 to 1986. He has been with Stanford University as an Electrical Engineering Professor from 1986 to the present. He founded Amati Communications Corporation, Palto Alto, CA, in 1991 (it was purchased by Texas Instruments, Incorporated in 1997) and was Officer/Director from 1991 to 1997. He is currently on the boards or advisory boards of BigBand Networks, Coppercom, GoDigital, Ikanos, Ionospan, Ishoni, IteX, Marvell, Kestrel, Charter Ventures, and Portview Ventures. He is a Member of the U.S. National Research Council's CSTB. His specific interests are in the area of high-performance digital transmission.

Dr. Cioffi has received the following awards: the National Academy of Engineering in 2001, IEEE Kobayashi Medal in 2001, IEEE Millennium Medal in 2000, IEE JJ Tomson Medal in 2000, 1999 University of Illinois Outstanding Alumnus, 1991 IEEE COMMUNICATIONS MAGAZINE best paper; 1995 ANSI T1 Outstanding Achievement Award, and NSF Presidential Investigator from 1987 to 1992. He became a member of the National Academy of Engineering in 2001. He has published over 200 papers and holds over 40 patents, most of which are widely licensed, including basic patents on DMT, VDSL, and V-OFDM.