

Dynamic Spectrum Management for Next-Generation DSL Systems

Kee Bong Song and Seong Taek Chung, Stanford University

George Ginis, Texas Instruments

John M. Cioffi, Stanford University

ABSTRACT

The performance of DSL systems is severely constrained by crosstalk due to the electromagnetic coupling among the multiple twisted pairs making up a phone cable. In order to reduce performance loss arising from crosstalk, DSL systems are currently designed under the assumption of worst-case crosstalk scenarios leading to overly conservative DSL deployments. This article presents a new paradigm for DSL system design, which takes into account the multi-user aspects of the DSL transmission environment. Dynamic spectrum management (DSM) departs from the current design philosophy by enabling transceivers to autonomously and dynamically optimize their communication settings with respect to both the channel and the transmissions of neighboring systems. Along with this distributed optimization, when an additional degree of coordination becomes available for future DSL deployment, DSM will allow even greater improvement in DSL performance. Implementations are readily applicable without causing any performance degradation to the existing DSLs under static spectrum management. After providing an overview of the DSM concept, this article reviews two practical DSM methods: iterative water-filling, an autonomous distributed power control method enabling great improvement in performance, which can be implemented through software options in some existing ADSL and VDSL systems; and vectored-DMT, a coordinated transmission/reception technique achieving crosstalk-free communication for DSL systems, which brings within reach the dream of providing universal Internet access at speeds close to 100 Mb/s to 500 m on 1–2 lines and beyond 1 km on 2–4 lines. DSM-capable DSL thus enables the broadband age.

INTRODUCTION

During the last decade, digital subscriber line (DSL) technology has been praised as an attractive broadband access service for residential and business areas [1]. Several DSL standards have

been successfully established with tens of millions of customers throughout the world. Currently, the most popular asymmetric DSL (ADSL) service reliably delivers up to 8 Mb/s depending on the distance of the customer from the telephone office. However, the need for even higher DSL speeds is already evident due to the increasing demands of the customer imposed by the proliferation of high-bandwidth multimedia services. The question is, then, “Is there a realm of even higher speeds reachable by the current DSL system?”

Dynamic spectrum management (DSM), the evolutionary step in current DSL spectrum management strategies, addresses this question. By inherently accommodating features like automatic fault detection and maintenance, and by guaranteeing spectral compatibility among different services, DSM promises faster and more reliable information delivery.

At present, DSL modems of the various standards operate independently; the high transmit energy of one line may cause severe interference (crosstalk) in neighboring lines, thus preventing the normal operation of DSL modems. The existing spectrum management approach in DSL mitigates this “DSL hogging” problem by enforcing power spectral density (PSD) masks for all the modems, in which the PSD mask designs were based on the worst case crosstalk emission scenario. The resulting conservative rules are described by the American National Standard Institute (ANSI) spectrum management standard [2] (affecting both the system specifications and the deployment guidelines).

Clearly such “static” spectrum management is effective only for the specific transmission scenario presumed in the design stage, which is hardly the case in practice. As a result, the current design approach necessarily leads to DSL deployment that is overly restrictive in terms of both coverage and speed. DSM allows for adaptive allocation of the available resources among multiple lines, depending on the channel characteristics and the crosstalk activity on the line.

Two principles will guide improvements in DSM:

The extension of optical fiber towards the customer side implies higher DSL data rates thanks to smaller signal attenuation in the shorter lines. A potentially more important observation, however, is that with an ONU a single DSLAM may attach to all the DSL modems.

- DSL modems should not transmit more power than necessary to achieve their target data rates with good quality of service.
- DSL modems should not transmit more bandwidth than useful for communication.

These two simple principles will guide enormous improvement. In particular, this article eventually shows a series of basic improvements that lead to 100 Mb/s symmetric service on single or multiple DSL lines, enabling broadband communications. A 500 m cable of 50 twisted pairs, then, can carry 5 Gb/s of symmetric data rate to 50 customers, more than has been contemplated in fiber or coaxial systems to date, illustrating that the copper connection to the customer need not be an impediment to the implementation of broadband services throughout the world.

At the time of this writing, DSM is an ongoing project of the T1E1.4 DSL Access Working Group [3]. This article overviews the basic concepts of DSM, presents the technical issues involved, and introduces some of the proposed DSM methods for spectra balancing and interference cancellation. The fundamental idea in DSM is the balanced control among different DSL modems and possibly among different DSL service providers. This concept of “harmonized” system operation is motivated by the multi-user characteristics of DSL systems, which are illustrated in the next section.

DSM WITH THE EVOLUTION OF DSL

DSL AS A MULTI-USER COMMUNICATION ENVIRONMENT

In the current deployment (shown in Fig. 1a), DSL transmission occurs between the modems at the customer side and the modems at the central office (CO) side, which are connected through twisted pair copper lines. Several twisted pairs are physically bundled together in the same cable, resulting in electromagnetic coupling, which induces interference known as crosstalk in the neighboring lines. Near-end crosstalk (NEXT) refers to the interfering signals at the receiver originating from transmitters at the same side of the affected receiver. Far-end crosstalk (FEXT) signals originate from transmitters at the other side of the affected receiver.

Since in the current DSL topology the spectral incompatibility among different DSL services and different service providers can aggravate the crosstalk problem, the existing static spectrum management enforces overly conservative PSD constraints for all DSL services, thereby severely limiting the achievable transmission performance. Clearly better spectra allocation must be addressed. DSM, even in this noncoordinated situation, immediately achieves remarkable improvement in both DSL service coverage and transmission speed by enabling DSL modems to “autonomously” adapt their own spectra, without causing harm to existing DSL systems.

The increasing investment in fiber to the curb (FTTC) installations is gradually resulting in a topology where the optical fiber is deployed between the CO and an optical network unit (ONU) located closer to the customer side, as

shown in Fig. 1a. Then twisted pairs emanate from the ONU providing the last mile connection to the customer. The ONU equipment includes a DSL access multiplexer (DSLAM) performing the appropriate multiplexing/demultiplexing operations between the fiber link and the DSL lines. The corresponding inverse operations are done at the CO. Very-high-speed DSL (VDSL) will mainly be deployed in such “fiber-assisted” environments.

The extension of optical fiber towards the customer side implies higher DSL data rates thanks to smaller signal attenuation in the shorter lines. A potentially more important observation, however, is that with an ONU a single DSLAM may attach to all the DSL modems. This change of paradigm allows the possibility to coordinate the DSL modems at the DSLAM side, enabling additional significant performance gain on top of the autonomous DSM [4].

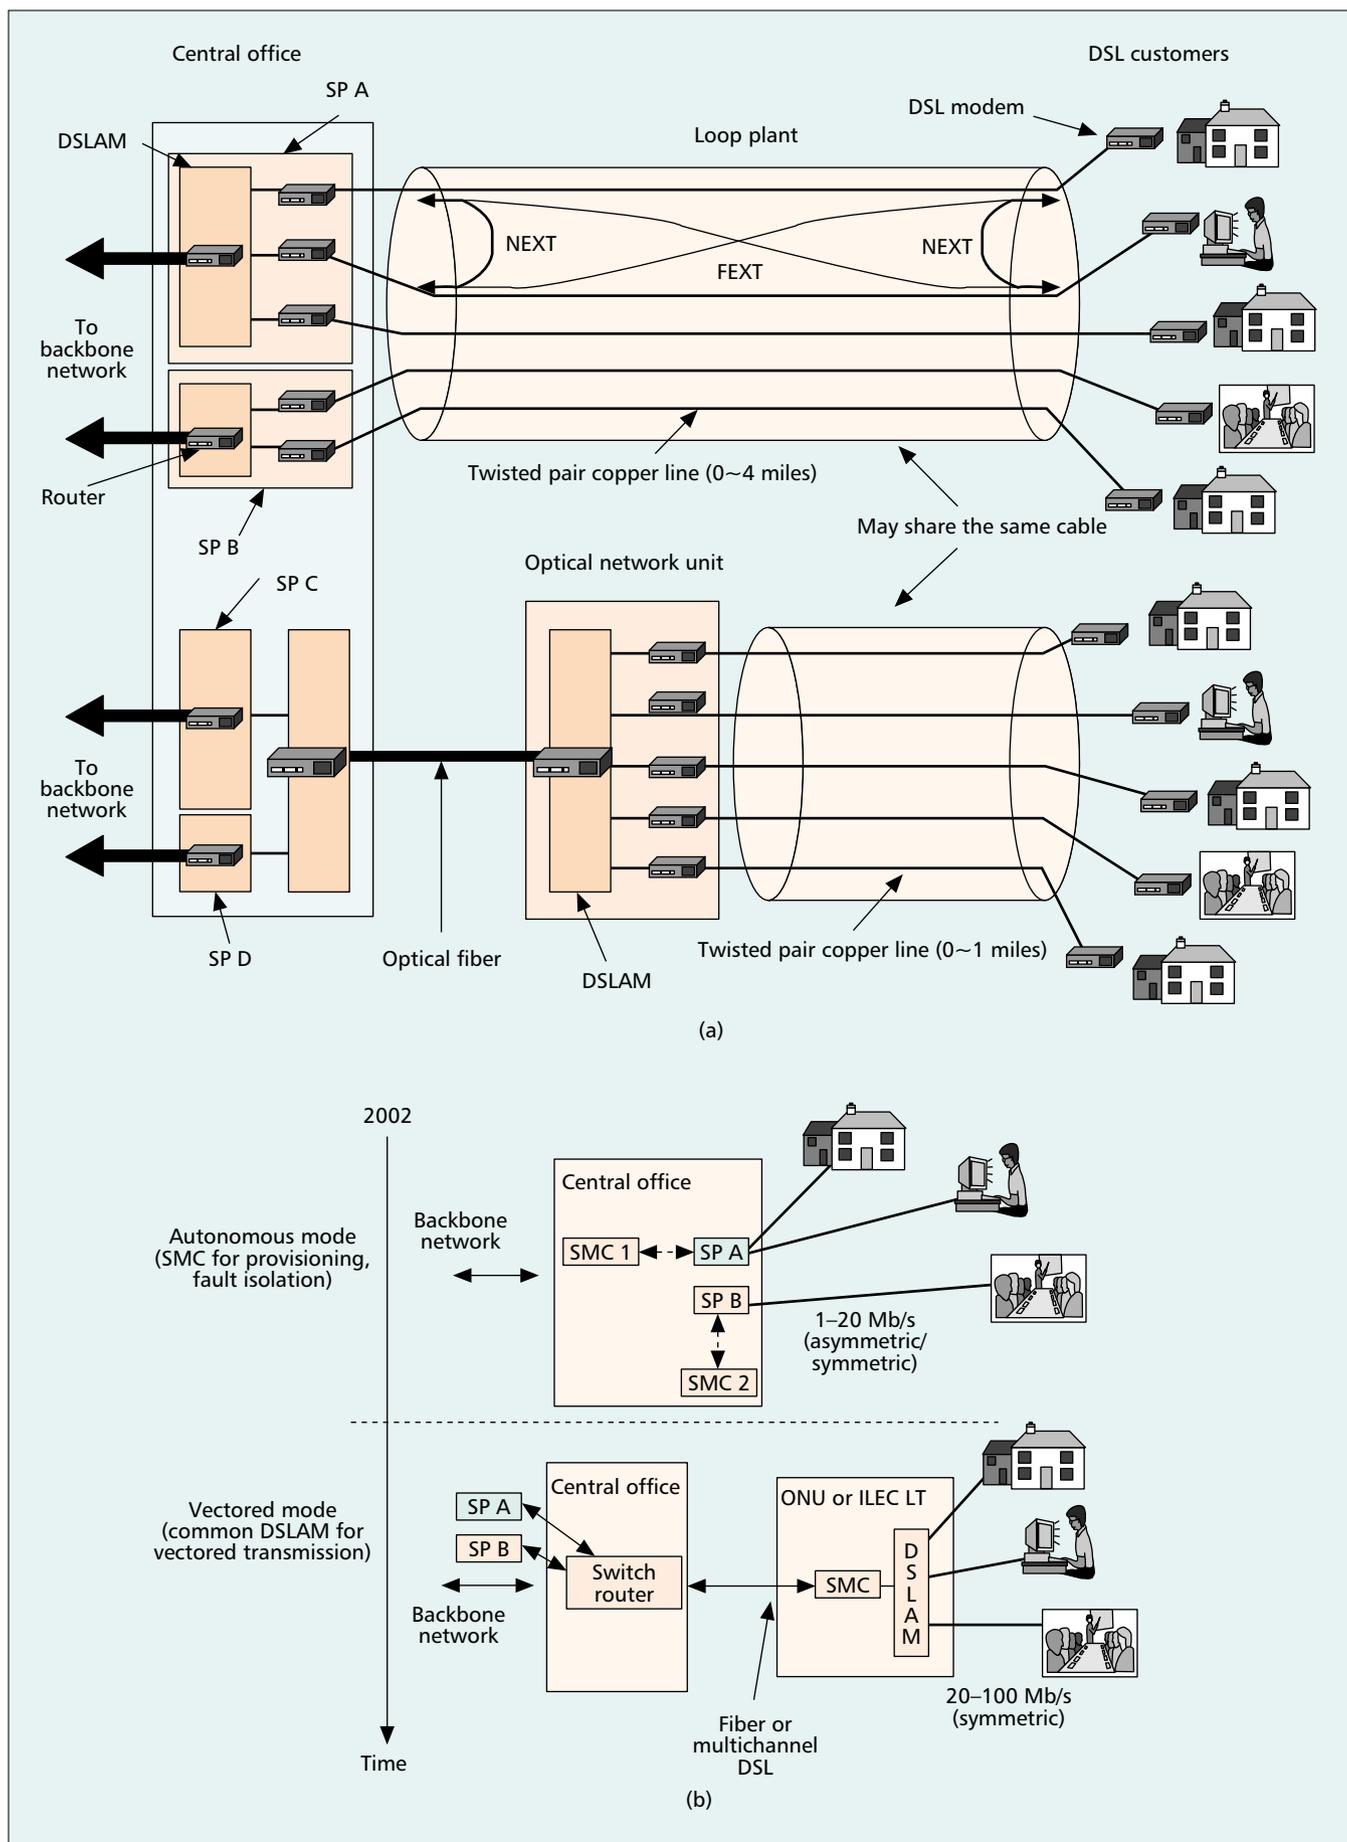
DSM CHARACTERIZATION BY MODE OF OPERATION

Figure 1b shows the timeline of the DSL/DSM evolution characterized by its operational modes. Although signal coordination enables the ultimate highest rates and performance in DSL, it is possible to obtain substantial gains even without any coordination [5]. Autonomous mode is essentially the present DSL environment where no coordination is allowed. In addition, in time each service provider may have its own spectrum management center (SMC) that reports line information and spectrum control information to provide automatic or manual repair and fault prevention. The most advanced spectrum utilization is achieved in vectored mode, where synchronized signals can be co-transmitted and co-received at one side (i.e., the ONU side). In this mode the SMC is directly located next to the DSLAM, and controls the transmission and reception of signals.

A SYSTEM PERSPECTIVE OF DSM

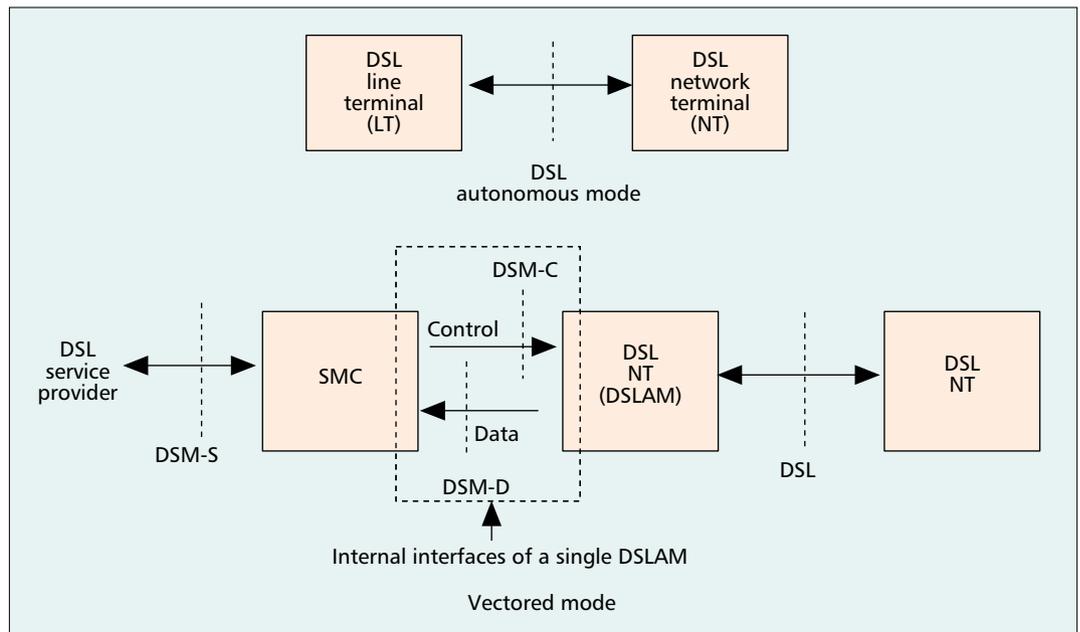
The basic DSM reference model is shown in Fig. 2. The DSL system can be any of the existing DSL standards. In an autonomous operation, there is no SMC control or SMC data reporting; furthermore, some or all DSL modems use autonomous DSM training algorithms that enhance the performance of every line in the binder. Significant gain is possible in this immediately practical and deployable mode. Further improvement is possible with some level of coordination in the vectored mode. In this case, line and crosstalk information is reported by the DSL modems to the SMC via the interface DSM-D, which is typically internal to a DSLAM at the ONU location. The SMC receives those data and provides control information through DSM-C, where the degree of control may depend on the level of coordination. The SMC also provides processed data results across the interface DSM-S for fault isolation, and maintenance and provisioning purposes.

Coordinated advanced DSM systems essentially execute the following steps, perhaps periodically, depending on the specific transmission



■ **Figure 1.** a) The DSL environment; b) a timeline of DSL/DSM.

At the core of DSM are advanced network optimization techniques to provide optimal resource allocation and transmission performance.



■ Figure 2. DSM reference model.

scenario: the network information acquisition phase, where loop characteristics, transmission parameters, and traffic information as well as external user commands are collected by a single service provider; the negotiation and optimization phase, in which all the information acquired in the previous step is used to distribute the network resources among the service provider's customers and the communication parameters are determined to best serve those customers; and the coordinated operation phase, in which the optimized communication and network parameters are provided to the modems within a single service provider's DSLAM to achieve "cooperative" data transmission and reception.

Several DSM system implementations are possible. For instance, when the modems are all collocated at an ONU cabinet serviced by a single provider, an independent entity may collect all the network information during the information acquisition phase. Of more immediate interest, each modem may autonomously react to line conditions in a way that is globally beneficial to all the DSL lines in the binder so that no system will perform any worse than if no DSM were used, and some or all may perform better. This autonomous step has very high early gains and creates momentum toward increasing coordination when lines get shorter (where coordination offers additional gains).

PHYSICAL LAYER TECHNIQUES IN DSM

At the core of DSM are advanced network optimization techniques to provide optimal resource allocation and transmission performance. The key elements of the DSM implementation include iterative water-filling (in the autonomous mode), the coordinated transmission method (in the vectored mode), the multi-input multi-output (MIMO)

interference identification [6] (applicable to the vectored mode), and rate and power control.

The following subsections introduce two promising physical layer techniques that can be used in the DSM optimization process: iterative water-filling [7, 8] for the autonomous mode and vectored DMT [4] for the vectored mode. Each algorithm balances multiple DSL modems either from the view of competition or from the view of cooperation. These algorithms are not only highly compatible with existing DSL modem technologies, but also result in impressive performance enhancements.

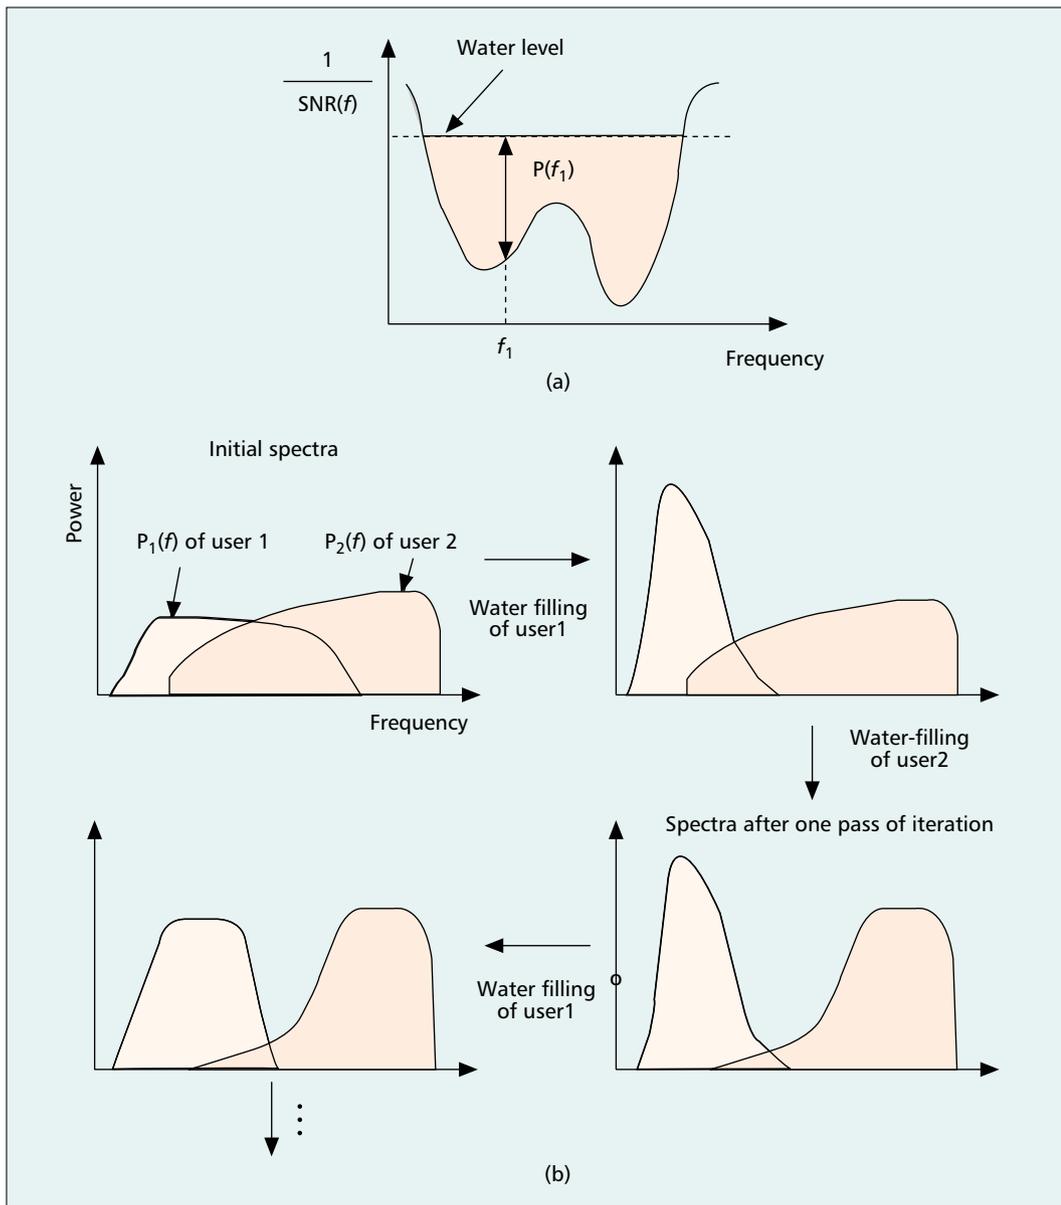
AUTONOMOUS MODE: ITERATIVE WATER-FILLING

Water-filling is a well known optimal-power distribution algorithm for the single-user communication channel [1] and provides the basis for the power and bit allocation schemes in most DMT¹-based modems. The single-user water-filling concept is illustrated in Fig. 3a. Given the channel signal-to-noise ratio (SNR) information in the frequency domain,² $SNR(f)$, the optimal spectrum ($P(f)$) maximizing the data rate is obtained by allocating more power to frequency bands with higher channel SNR. The strategy is pictorially the same as "pouring" the total power in the bowl of the inverse $SNR(f)$ curve, hence the name water-filling. A power-minimizing or margin-fixing version of water-filling may be more commonly deployed to practice DSM, in which case the target data rate is fixed and the objective is to minimize the total power used. The optimal allocation scheme is still interpreted as pouring the power in the inverse $SNR(f)$ curve; in this case, however, the water level is dictated by the target data rate rather than by the total power constraint.

Iterative water-filling can be viewed as an extension of the water-filling process for a multi-user communication environment [7, 8]. This algorithm is based on formulating power alloca-

¹ Discrete multitone (DMT) is a multicarrier modulation method adopted in a number of existing DSL standards. In DMT, the information bits are mapped into a set of quadrature amplitude modulated (QAM) signals, where these QAM signals are transmitted through independent sub-channels in the frequency domain. A detailed description of this method can be found in [1].

² In general, SNR is defined as signal power divided by noise power at a specific frequency tone or band. Channel SNR refers to SNR with unit signal power across the entire frequency band, which concisely characterizes the communication channel. For notational simplicity, $SNR(f)$ denotes the channel SNR at frequency f .



For ADSL channels, the data rates achieved by iterative water-filling are very close to those achieved by the optimal rate and power allocation. Indeed no better spectra have been found in stationary binders.

■ Figure 3. a) Water-filling for a single user; b) iterative water-filling for two users.

tion in the multi-user interference network as a competitive game, in which each user tries to maximize his own data rate regarding crosstalk interference from other users as noise. Starting from any initial spectra, the water-filling procedure on each line is performed independently across all DSL lines within a binder. It can be shown that this distributed iteration process converges to a competitively optimal equilibrium point, sometimes called a *stationary point* in optimization theory. While multiple stationary points may exist in the multi-user interference channel, such a case has not yet been found over a wide range of DSL lines tested, and indeed can be mathematically proven not to exist in frequency bands up to 2 MHz. Furthermore, the convergence and uniqueness of an iterative water-filling solution from any initial spectra are guaranteed for all DSL lines [7]. Power-minimizing iterative water-filling is also possible, in which case the target data rate on each line must be known to

the modems in advance. Maximum data rates are established today even in static spectrum management by provisioning rules known as profiles; thus, DSM's need to impose a target data rate maintains the existing autonomous nature of DSL deployment today.

For ADSL channels, the data rates achieved by iterative water-filling are very close to those achieved by the optimal rate and power allocation [9]. Indeed, no better spectra have been found in stationary binders. This result is highly encouraging; specifically, it implies that iterative water-filling, in which rate and power for each frequency tone of the user is adapted autonomously with no central control, performs essentially the same as when the spectra of all the users are controlled by a centralized network entity.

Figure 3b illustrates the iterative water-filling process for two users. As can be seen, at each iteration step, each user's spectrum "moves away" from the frequency region characterized

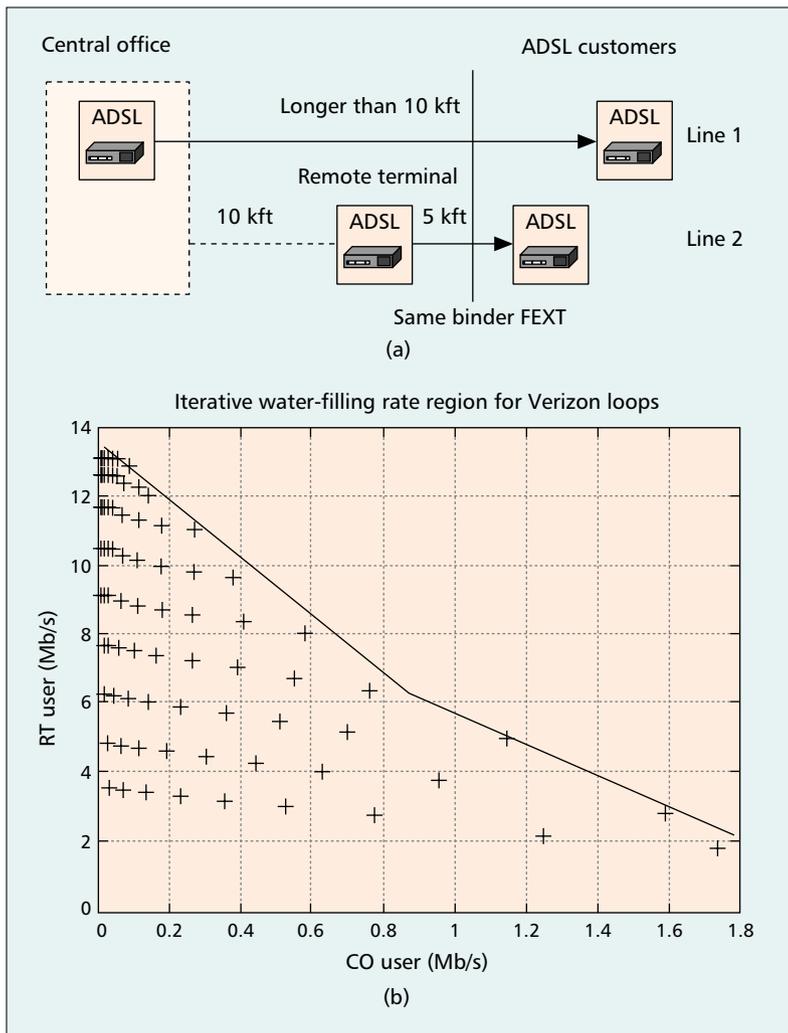


Figure 4. a) The ADSL transmission scenario; b) the achievable rate region by iterative water-filling.

by strong interference thanks to the water-filling process explained above, thereby approaching better performance step by step. For most of the DSL lines tested, the algorithm converged within two or three iterations.

The iterative water-filling algorithm can be summarized for offline analysis as follows:

- Initialize iter_num=0
- Repeat
 - For $i = 1, \dots, L$
 - Compute water-filling spectrum, S_i , for user i
 - Compute crosstalk spectra into all other channels

$$\sigma_{j,n}^2 = \sum_{i \neq j} |H_{ji,n}|^2 S_{i,n}$$

until iter_num < max

Here, L is the number of users or DSL lines; max is the maximum number of iterations allowed; $S_{i,n}$ is the i th user's spectrum at the n th frequency tone; $\sigma_{j,n}^2$ is the crosstalk variance of the j th user at the n th frequency tone; and $H_{ji,n}$ is the crosstalk transfer function from user i into user j at the n th frequency tone.

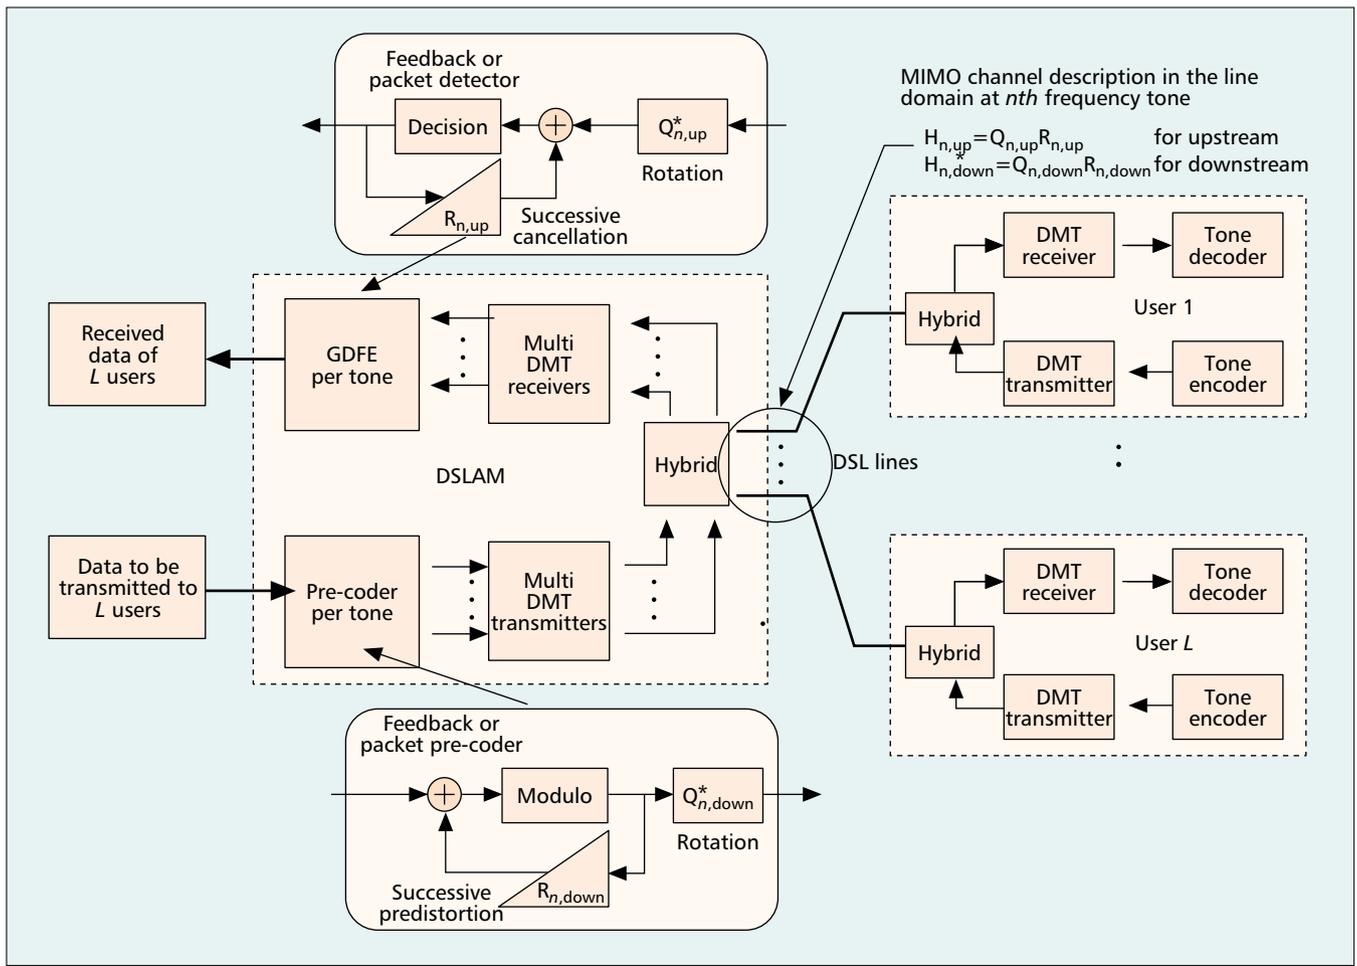
The above procedure implicitly assumes that the user water fillings occur in a specific order. However, it should be emphasized that for DSL channels, the same convergence result is obtained whether or not the user water fillings are ordered or successive. In other words, each modem only needs to autonomously execute its own water-filling process.

Fortunately, this iterative water-filling algorithm has already been implemented “unwittingly” in several existing DMT-based DSL systems. The well-known bit-swapping technique, which serves to adapt modem operation to slow channel changes, implicitly executes iterative water-filling. Actually, for such systems the benefits of iterative water-filling would be much more pronounced if the conservative rules of the current spectrum management standards were lifted. Such lifting combined with iterative water-filling will not lead to performance degradation on any existing line that continues to use static spectrum management. However, the more lines that use DSM, the greater the advantage. A “single” static line, such as a 998 VDSL system with a fixed PSD mask [2], can unnecessarily limit the performance of all future lines.

DSM is permitted in the existing standard [2] under the definition of “new technology.” This accommodation essentially paves the way for future DSM use by simply allowing new systems to be deployed as long as they do not harm any existing systems more than those systems are allowed to harm themselves. In essence, DSM-enabled DSLs will gradually “clean up” the loop plant in a graceful and consistently performance-improving manner. Service providers will be motivated to use DSM-enabled modems rather than static modems as they deploy more line terminals and more DSLs, since line maintenance will decrease and data rates will dramatically increase.

Figure 4 illustrates a classic field problem in current ADSL deployments, for which iterative water-filling can achieve dramatic data rate increases. In Fig. 4a the ADSL receiver on line 1 experiences large FEXT from the remote-terminal (RT)-located ADSL downstream transmitter which is only 5 kft away from the ADSL receiver. Under the existing static spectrum management rules, line 1 simply does not work; the achievable data rate is only 300 kb/s for downstream, and field personnel must be dispatched.

Crosstalk measurements for a typical configuration as in Fig. 4a were reported by the Verizon Broadband Integration Lab [10]. The worst case crosstalk coupling pair is considered here. Using static rate adaptive training, the Verizon Broadband Integration Lab was able to obtain 9 Mb/s in the short line and 100 kb/s in the long line. Figure 4b shows the achievable rates by allowing both lines to use power-minimizing (minimum power for a given rate) iterative water-filling. For example, rate pairs (2,1.8), (4,1.4), (6,0.9) and (8,0.6) Mb/s are achievable for the short and long lines, respectively. At the expense of reduced rates on the short line, enormous improvements are possible in the long line. The spectrum of the short line is seen to “migrate” to higher frequency tones, yielding to the long line.

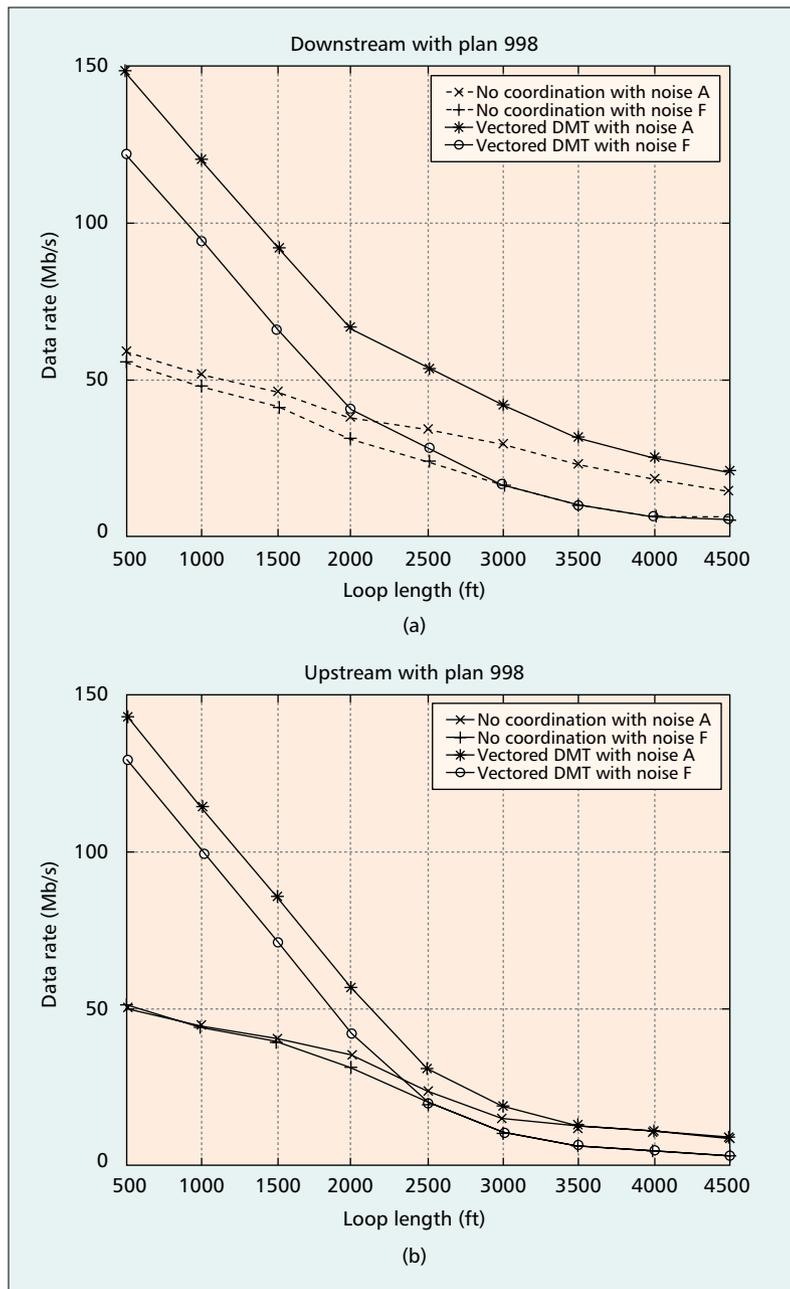


■ Figure 5. A vectored DMT diagram.

The iterative water-filling algorithm can “optionally” be augmented by an outer control loop where centralized resource allocation is performed, thus providing more flexible services to customers in the form of potentially higher data rate choices. Clearly, this would require a higher degree of coordination, where each modem reports the relevant DSM data, such as the achieved data rate or the power margin and the associated transmit power level to the SMC. Based on this reporting, the SMC would determine the achievable data rate pairs or, more generally, the achievable data region, from which it can command each modem to perform distributed iterative water-filling with different transmit power level and target data rate.

Even this small amount of cooperation among the modems offers a great deal of flexibility. For example, upon the request of a customer or a higher level of network hierarchy, the SMC can command a modem operating in an excess power margin to increase the data rate, thereby providing a higher-speed service, or to decrease its transmit power, resulting in less crosstalk interference with other modems. This concept of flexible service provisioning is not feasible in the static spectrum management environment, under which each modem is forced to use the fixed PSD mask regardless of where it is located and “accept” the resulting data rate whether or not it is satisfactory.

Although the well-known water-filling algorithm results in a bit distribution with continuous values, all practical DMT systems require that the number of bits assigned to a frequency tone is an integer or the fraction of integers. Methods to perform discrete bit allocation have been developed in the past for single-user systems. For multi-user DMT systems, a discrete bit-loading and power control algorithm was recently proposed to provide some improvement in bit and power allocation [11]. In this algorithm, a single network entity calculates and updates the incremental power table that specifies the amount of power necessary to assign one more information bit to the current bit assignment over all frequency tones and all users. Based on a greedy algorithm, a particular user and a specific frequency tone corresponding to the minimum entry in the table is selected and gets a one-bit increase. This bit allocation update process continues until the transmit power reaches the maximum allowed, or the data rate meets the specified target. Although the complexity for the update of the incremental table at each bit allocation step can be high, the performance of multi-user discrete bit loading leads to reduced total transmit power for a given data rate with respect to the one obtained through iterative water-filling for the same target data rate. However, it is unlikely that such a loading method would be used, because when coordination is



■ **Figure 6.** VDSL data rates: a) downstream; b) upstream.

allowed, the better and easier choice is the vectored mode described below.

VECTORED MODE: VECTORED DMT

Vectored DMT can achieve FEXT-free transmission in a coordinated environment where multiple DSL lines “share” their physical signal information [4]. One-sided coordination can occur when a fiber-assisted DSL system (shown in Fig. 1a) is deployed. In this case, the DSLAM is capable of generating and receiving synchronized DSL signals from multiple lines. By employing frequency-division duplexing (FDD), which separates the upstream and downstream frequency bands, NEXT is easily mitigated; the major interference, then, comes from FEXT. The performance loss caused by FEXT becomes much more severe when the lengths of the DSL lines differ signifi-

cantly. For instance, if all the modems at the customer side transmit the same power level in the same frequency band, FEXT caused by short lines at the upstream receiver of a longer line may be even higher than the direct signals; thus, the upstream transmission is severely affected.

Vectored DMT exploits the optimum MIMO signal processing structure, the generalized decision feedback equalizer (GDFE). GDFE offers a unified view of signal detection in the interference channel. In the DSL environment, major interference sources include intersymbol interference (ISI), a signal distortion caused by the frequency selectivity of the line, and crosstalk interference, more generally referred to as interuser interference (IUI). In its most general form, GDFE thus views ISI and IUI in the same domain and tries to eliminate both simultaneously. DMT is a special case of GDFE where only ISI is eliminated.

When DMT is employed as a modulation technique, the signals from multiple lines can be processed separately at each frequency tone in the user or line domain. The GDFE structure then can be derived to eliminate the crosstalk. Figure 5 illustrates the block diagram of a vectored DMT system. One should note that all DSL modems at the customer side are essentially the same as before, except that the transmission and reception now occurs synchronously, which is feasible in a one-sided coordinated situation with a common DSLAM, and indeed is included in the most recent VDSL-DMT standard under the name *zippering* or *digital duplexing* [12]. The DSLAM processes the signals to be transmitted on each line or to be received from each line. It is assumed that all the channel information, such as signal attenuation and crosstalk characteristics on each line, is reported to the SMC and is known to the DSLAM in advance, which will be subsequently expressed as $H_{n,up}$ and $H_{n,down}$, as in Fig. 5.

For the upstream case, the DSLAM becomes the cooperative receiver. The received QAM signals from the lines are first processed with the usual DMT technique. After the normal DMT processing stage, the signal at each line contains crosstalk interference from other lines. For each frequency tone, all the signals from multiple lines are then collected, and interference cancellation is performed in the line domain: a rotation operator $Q_{n,up}^*$ in Fig. 5) is first applied such that the signal to be decoded in the current line contains interference only from the previously decoded symbols in other lines; finally, decision-assisted successive cancellation ($R_{n,up}$ in Fig. 5) is performed.

On the other hand, the DSLAM becomes the cooperative transmitter in the downstream direction, in which case the corresponding GDFE has a *pre-coder* structure. The multi-line signals to be transmitted are first collected for each frequency tone, and appropriate predistortions or precodings (characterized by $Q_{n,down}$ and $R_{n,down}$ in Fig. 5) are successively introduced across all the lines. These distortions are such that, after passing through the channel, the received signal at the modem on the customer side becomes crosstalk-free.

While the concept of the GDFE is applicable to any system, the use of DMT as a modulation technique greatly reduces the computational com-

plexity of transmitter/receiver processing since each frequency tone can be separately processed in the line domain. Moreover, the major crosstalk sources for each line are usually due to at most 2–4 neighboring lines in the cable. Therefore, the complexity increase at the DSLAM is the additional GDFE/pre-coder processing block associated with at most 4×4 matrix operations/tone for each line. There is essentially no increase in complexity on the customer side.

Figure 6 compares the performance for the VDSL system, where one-sided vectored DMT is compared with noncoordinated transmission. The assumed FDD plan is the one for VDSL proposed for North America (also known as 998). Clearly, the gains are spectacular, especially for short loops where transmission is FEXT-limited. For example, a downstream data rate of 50 Mb/s limits noncoordinated systems to a reach of 1150 ft with a type A noise model, whereas vectored DMT extends the reach to 2650 ft.

Two-sided coordination is not possible in DSL practice since modems at the customer side are geographically separate; however, in some private network environments utilizing multiple lines, modems at both the transmitter and receiver sides could cooperate. In such cases, instead of deploying FDD, NEXT cancellation methods can be used, thus allowing use of the entire frequency band in both the upstream and downstream directions.

SUMMARY AND FUTURE DIRECTIONS OF DSM

DSM offers the most efficient data transmission on the conventional twisted pair line telephone network. Globally beneficial dynamic spectra variation among DSL modems of different lines is the key idea in DSM. The earliest autonomous modes of DSM are immediately deployable and offer significant performance improvement in any environment. The highest performing vectored modes build on the autonomous mode to add coordination where economically prudent, and are motivated from a regulatory perspective by the increasing uses of optical fibers to replace copper in the feeder loop.

Iterative water-filling can achieve much better spectra allocation compared to the current static management rules, and vectored DMT provides essentially FEXT-free transmission. DSM allows more flexible provisioning to customers, and enables service providers to optimize their resource management. Such DSM will eventually allow a binder of 50 twisted pairs at 500 m to carry 5 Gb/s, more than any broadband delivery system of similar range has yet to contemplate.

REFERENCES

- [1] T. Starr, J. M. Cioffi, and P. J. Silverman, *Understanding Digital Subscriber Line Technology*, Prentice Hall, 1999.
- [2] "Spectrum Management For Loop Transmission Systems," ANSI Std. T1. pp. 417–2001.

- [3] "Proposed Scope and Mission for Dynamic Spectrum Management," ANSI Cont. T1E1.4/2001-188R4, Greensboro, NC, Nov. 2001; more documents can be found at <http://www-isl.stanford.edu/~cioffi/dsm/>
- [4] G. Ginis and J. M. Cioffi, "Vectored Transmission for Digital Subscriber Line Systems," *IEEE JSAC*, vol. 20, no. 5, June 2002, pp. 1085–1104.
- [5] J. M. Cioffi et al., "New Technology in Spectrum Management," ANSI Contribution T1E1.4/2002-128, Atlanta, GA, Apr. 8, 2002.
- [6] J. M. Cioffi et al., "Channel Identification with Dynamic Spectrum Management," ANSI Cont. T1E1.4/2001-147R1, Greensboro, NC, Nov. 5, 2001.
- [7] W. Yu, G. Ginis, and J. M. Cioffi, "An Adaptive Multiuser Power Control Algorithm for VDSL," *IEEE JSAC*, vol. 20, no. 5, June 2002, pp. 1105–15.
- [8] W. Yu and J. M. Cioffi, "Competitive Equilibrium in the Gaussian Interference Channel," *IEEE ISIT 2000*, p. 431.
- [9] S. T. Chung and J. M. Cioffi, "Rate and Power Control in a Two-User Multicarrier Channel with No Coordination: The Optimal Scheme vs. Suboptimal Methods," to appear, *IEEE VTC 2002 Fall*, Vancouver, BC, Canada.
- [10] J. M. Cioffi, S. T. Chung, and J. Lee, "To 2001-273R1 Using Measured Verizon DSL SNRs," ANSI Cont. T1E1.4/2002-069, Vancouver, BC, Canada, Feb. 18, 2002.
- [11] J. Lee, R. V. Sonalkar, and J. M. Cioffi, "Multiuser Discrete Bit Loading for DMT-Based DSL Systems," to appear, *IEEE GLOBECOM 2002*.
- [12] "Transmission and Multiplexing (TM); Access Transmission Systems on Metallic Access Cables; Very High Speed Digital Subscriber Line (VDSL); Part 2: Transceiver Specification," ETSI TS 101 270-2, 2001.

BIOGRAPHIES

KEE BONG SONG [S '02] (kbsong@stanford.edu) received his B.S.E.E in 2000 from the Korea Advanced Institute of Science and Technology (KAIST), Taejon, and his M.S.E.E in 2002 from Stanford University. He is currently working toward his Ph.D. degree at Stanford University. During the summer of 2002 he worked as a member of technical staff at Agere Systems where he was involved in next-generation high-speed wireless LAN system design. His specific research interests include the coding and detection algorithms for wireless and wireline communication systems.

SEONG TAEK CHUNG [S '99] (stchung@dsl.stanford.edu) received his B.S.E.E. in 1998 from Seoul National University, Korea, and his M.S.E.E. in 2000 from Stanford University. He is now pursuing his Ph.D. degree at Stanford University. He worked in the Wireless Research Department at Lucent Technologies Bell Labs as a member of technical staff during the summer of 2000. His research area includes multi-user communication theory as well as wireless communication theory.

GEORGE GINIS [S '97] (gginis@ti.com) received his diploma in electrical and computer engineering from the National Technical University of Athens, Greece, in 1997, and M.S. and Ph.D. degrees in electrical engineering from Stanford University, California, in 1998 and 2002, respectively. He is currently with the Broadband Communications Group of TI, San Jose, California. His research interests include multi-user transmission theory, interference mitigation, and their application to wireline and wireless communications.

JOHN M. CIOFFI [F '96] (cioffi@stanford.edu) received his B.S.E.E. in 1978 and his Ph.D. in electrical engineering in 1984, from Stanford University. He worked for Bell Laboratories from 1978 to 1984 and IBM Research from 1984 to 1986, and is currently an electrical engineering professor at Stanford, 1986–present. He founded Amati Com. Corp in 1991 (purchased by TI in 1997) and was officer/director from 1991 to 1997. He currently is on the Board of Directors of Marvell, IteX, Teknovus, and Multicom. He is on the advisory boards of Actelis, GoDigital, Ikanos, Iospan, Kestrel, Charter Ventures, and Portview Ventures, and a member of the U.S. National Research Council's CSTB. His specific interests are in the area of high-performance digital transmission. He has published over 200 papers and holds over 40 patents.

The iterative water-filling algorithm can "optionally" be augmented by an outer control loop where centralized resource allocation is performed, thus providing more flexible services to customers in the form of potentially higher data rate choices.