Vectored Transmission for Digital Subsriber Line Systems

George Ginis and John M. Cioffi
Department of Electrical Engineering
Stanford University

This work was supported by gifts from Samsung, Intel, AMD, TI, France Telecom, Telcordia, SK Telecom and CommTech.

Parts of this work have been presented at the 34th Asilomar Conference on Signals, Systems, and Computers, and will be presented at ICC 2001.

Abstract

This paper describes the "vectored" transmission technique for Digital Subscriber Line (DSL) systems, which utilizes user coordination at the Central Office (CO) or Optical Network Unit (ONU). This method exploits the co-location of the downstream transmitters and of the upstream receivers, in order to achieve Far-End Crosstalk (FEXT) cancellation and perform multi-user transmission optimization. Computer simulations demonstrate that the performance improvements are particularly pronounced in environments with strong FEXT such as in Very-high-speed DSL (VDSL).

Discrete Multi-Tone (DMT) is employed for each user with additional constraints on the Cyclic Prefix (CP) length, and with the assumption of block-synchronized transmission and reception for downstream and upstream transmission correspondingly. Within each tone, upstream crosstalk is removed by Multiple-Input-Multiple-Output (MIMO) decision feedback at the receiving side, while downstream crosstalk is eliminated by analogous pre-processing at the transmitting side. Additionally, the issue of transmission energy allocation in frequency and among users is addressed. Assuming Frequency Division Duplexing (FDD), the corresponding optimization problem is formulated and solved via convex programming both for a fixed upstream-downstream band plan, and for a dynamically programmable band plan. The case of power back-off as a means to reduce the impact of crosstalk on alien systems is also treated.

Keywords

Digital subscriber lines, interference cancellation, multiuser detection, transmission optimization, energy allocation, power back-off, spectrum management, QR decomposition, convex optimization.

I. Introduction

DSL systems are rapidly gaining popularity as a broadband access technology capable of reliably delivering high data rates over telephone subscriber lines. The successful deployment of Asymmetric DSL (ADSL) systems has helped reveal the potential of this technology, and current standardization efforts focus on VDSL, which allows the use of bandwidth up to 20 MHz ([1], [2], [3], [4], [5]). ADSL can reach downstream rates up to 6 Mbps, while VDSL aims to deliver asymmetric service with downstream rates up to 52 Mbps, and symmetric service with rates up to 13 Mbps. Yet, several results in this paper advocate that DSL communication is still far from reaching its full potential, and that the gradual "shortening" of the loop combined with the application of advanced algorithms can realize the deployment of symmetric services with rates approaching 80 Mbps.

The DSL system topology is shown in Fig. 1, where twisted-pairs emanate from the CO, and reach out to the Customer Premises Equipment (CPE) of the DSL customers (see [6] for more details). The installation of ONU's (at points between the CO and the CPE's) aims to shorten the loop lengths, so that the reach and performance of DSL service are improved. Typically, ONU's are connected to the CO through fiber.

The dominant impairment in DSL systems is crosstalk arising from electromagnetic coupling between neighboring twisted-pairs. FEXT describes the coupled signals, which originate from the same side as the affected signal, while Near-End Crosstalk (NEXT) describes the coupled signals, which originate from the opposite side of the affected signal. (Fig. 1 shows both NEXT and FEXT signals induced by upstream transmission of user 1.) In VDSL, the impact of NEXT from same systems is suppressed by employing FDD to separate downstream and upstream transmission. Then, FEXT from other VDSL systems (often called "self-FEXT") is the major performance-constraining factor, especially as the loop length becomes shorter. It should not be overlooked that non-VDSL systems induce both NEXT and FEXT, but their effect is not as severe as self-FEXT, due to the fact that they occupy a much narrower bandwidth.

The objective of *vectored transmission* is to eliminate self-FEXT by treating the channel as a MIMO system and employing joint signal processing of all signals at either the receiver side (for upstream transmission) or at the transmitter side (for downstream transmission). DMT modulation allows the partition of the channel into orthogonal frequency bins, where crosstalk can be cancelled independently. For upstream, crosstalk is removed in a decision feedback manner, while for downstream, crosstalk is eliminated by pre-distorting the transmitted symbols. Thus, transmission is made essentially interference-free for each user.

Crosstalk coupling is strongest among the twisted-pairs in a binder group, therefore mitigating self-FEXT within a binder group has the biggest performance benefit. Thus, in order to obtain the maximum advantage from vectored transmission, the CO/ONU transceivers corresponding to a binder group must be co-located. Vectoring can also be applied to a subset of a binder group, however the benefits will be reduced.

It is a fact that in the current "unbundled" environment different service providers may actually share a binder group, which invalidates the assumption of co-location of the CO transceiver equipment. However, there are several indications that ONU deployment will lead to a different architecture, which will increasingly allow for vectored transmission ([7]). More specifically, the current architecture of "line unbundling" becomes impractical with the installation of ONU's, since it implies that each service provider uses an individual fiber to provide a proprietary connection to the ONU, and that the ONU must be large enough to accomodate a shelf or rack for each service provider. These difficulties (combined with the fact that vectored transmission can realize large performance gains) may likely cause the evolution to "packet unbundling", where service bandwidth is leased at the transport layer, instead of the physical layer.

This paper outlines the vectored transmission technique, and attempts to address some of the practical issues related to its implementation. Section II gives the channel model, and defines coordinated DMT transmission, which makes it sufficient to perform crosstalk mitigation independently in each tone. Section III presents the crosstalk cancelling structures for upstream and downstream communication, and discusses their relation to other known methods, while addressing some relevant practical issues. In Section IV, the issue of transmission optimization is treated, where the objective is the optimal allocation of energy across frequency and among users given a set of energy and induced crosstalk constraints. A solution to the question of joint energy and frequency planning is also presented. Finally, Section V shows simulation results.

The crosstalk problem has been addressed before in [8], [9], where MIMO Minimum-Mean-Square-Error (MMSE) linear equalizers were derived, as opposed to the decision-directed ones adopted in this paper. Reference [10] employs the singular value decomposition (similarly to [11] applied in wireless) to achieve crosstalk cancellation assuming co-location of both transmitters and receivers. Other related work includes [12], [13], where "wider than Nyquist" transmitters were shown to provide performance advantages compared to "Nyquist-limited" ones, and [14], where crosstalk cyclostationarity (induced by transmitter synchronization) combined with over-sampling were shown to result in higher SNR values.

The following notation is used throughout the paper: Upper-case letters denote matrices, and bold letters denote vectors. The superscripts ^{T}, and * denote the transpose, and conjugate transpose operations correspondingly. The symbols $(.)_k$, and $||.||$ denote correspondingly the k^{th} element of a vector, and the square norm of a vector. The operator $diag(.)$ has a different meaning depending on its argument: If the argument is a matrix, then its off-diagonal elements are zeroed. If the argument is a list of elements (or matrices), then a diagonal (or block diagonal) matrix is produced from those elements (or matrices).

II. Channel Model and DMT Transmission

The DSL channel model for the architecture of Fig. 1 is now presented. The L users are assumed to correspond to a subset of the twisted-pairs of a binder group. The sampled output for a specific user for either upstream or downstream transmission depends on the present and past input symbols of both the intended user and the other crosstalking users. A block of N output samples for user i satisfies:

$$\mathbf{y}_i = H^c_{i,1}\mathbf{x}^p_1 + \cdots + H^c_{i,i}\mathbf{x}^p_i + \cdots + H^c_{i,L}\mathbf{x}^p_L + \mathbf{n}_i \quad (1)$$

where $H^c_{i,1}, \ldots, H^c_{i,L}$ are convolution matrices derived from the channel impulse response matrix, $\mathbf{y}_i$ is the vector of N output samples of receiver i, $\mathbf{x}^p_k$ is the vector of $N + \nu$ input symbols of user k, and $\mathbf{n}_i$ is the vector of N noise samples of receiver i. Evidently, ν represents the maximum memory of the transfer and crosstalk coupling functions expressed in number of samples. The noise samples represent the superposition of several noise sources such as crosstalk from neighboring DSL systems, radio frequency ingress, and impulse noise. In the following, $\mathbf{n}_i$ is considered to be white and gaussian, and, without loss of generality, has unit variance.

Two fundamental assumptions are required in order to proceed:

- All users must employ block transmission with a CP of at least length ν (as outlined in [15]).
- Block transmission and reception at the CO/ONU must be synchronized as illustrated in the timing diagram of Fig. 2.

Given the co-location assumption, synchronized block transmission is relatively straightforward to implement, but synchronized block reception needs some more attention. The block boundaries for upstream transmission must be aligned in such a way, so that the blocks of all users arrive simultaneously at the CO/ONU. This block-level synchronization must be performed during initialization, and it is analogous to the problem of synchronized uplink transmission in a wireless environment.

Synchronization at the CO/ONU is automatically achieved when "Zipper" FDD is used (see [16]). According to this technique, a Cyclic Suffix (CS) is included in addition to the CP, with size larger than the channel propagation delay. "Zippering" offers the benefit of eliminating residual NEXT and near-echo resulting from "spectral leakage" at frequencies close to the upstream/downstream band edges. Nevertheless, in the remainder of this paper, the less stringent assumptions of the previous paragraph will be made, with the tacit understanding that residual NEXT and near-echo are mitigated by transmitter pulse-shaping and receiver windowing as in [17].

Taking the above into account, (1) becomes:

$$\mathbf{y}_i = H_{i,1}\mathbf{x}_1 + \cdots + H_{i,i}\mathbf{x}_i + \cdots + H_{i,L}\mathbf{x}_L + \mathbf{n}_i \quad (2)$$

where $\mathbf{x}_k$ is a vector of N input symbols of user k, and $H_{i,j}$, $i, j = 1, \ldots, L$ are circulant matrices. Combining the L users, (2) becomes:

$$\mathbf{y} = H\mathbf{x} + \mathbf{n} \tag{3}$$

where $\mathbf{y} = \left[\mathbf{y}_1^T\, \mathbf{y}_2^T\, \ldots\, \mathbf{y}_L^T\right]^T$, $\mathbf{x} = \left[\mathbf{x}_1^T\, \mathbf{x}_2^T\, \ldots\, \mathbf{x}_L^T\right]^T$, $\mathbf{n} = \left[\mathbf{n}_1^T\, \mathbf{n}_2^T\, \ldots\, \mathbf{n}_L^T\right]^T$, and H is a matrix whose (i, j) block is $H_{i,j}$. The noise covariance matrix is assumed to be $R_{\mathbf{nn}} = I$.

Applying the Discrete Fourier Transform (DFT) modulation principle of [18], [15], an Inverse Discrete Fourier Transform (IDFT) operation is performed on each transmitted data block (prior to appending the CP), and a DFT operation is performed on each received data block (after discarding the CP):

$$\begin{aligned} \mathbf{Y} &= Q_{mDFT} H Q_{mIDFT} \mathbf{X} + Q_{mDFT} \mathbf{n} \\ &= \Lambda \mathbf{X} + \mathbf{N}' \end{aligned} \tag{4}$$

where:

$$\mathbf{Y} = \left[\mathbf{Y}_1^T\, \mathbf{Y}_2^T\, \ldots\, \mathbf{Y}_L^T\right]^T = Q_{mDFT}\mathbf{y} \tag{5}$$

$$\mathbf{x} = Q_{mIDFT}\left[\mathbf{X}_1^T\, \mathbf{X}_2^T\, \ldots\, \mathbf{X}_L^T\right]^T = Q_{mIDFT}\mathbf{X} \tag{6}$$

$$\mathbf{N}' = \left[\mathbf{N}_1'^T\, \mathbf{N}_2'^T\, \ldots\, \mathbf{N}_L'^T\right]^T = Q_{mDFT}\mathbf{n} \tag{7}$$

$$Q_{mIDFT} = diag\left(\overbrace{Q_{IDFT}, Q_{IDFT}, \ldots, Q_{IDFT}}^{L}\right) \tag{8}$$

with Q_{IDFT} the $N \times N$ IDFT matrix, $Q_{DFT} = Q_{IDFT}^*$ the corresponding DFT matrix, and $Q_{mDFT} = Q_{mIDFT}^*$. Also, $\mathbf{Y}_k$, $\mathbf{X}_k$, $\mathbf{N}_k'$ each contain N samples corresponding to user k, $R_{\mathbf{N}'\mathbf{N}'} = I$, and Λ is a matrix whose (i,j) block is the diagonal matrix $\Lambda_{i,j} = Q_{DFT} H_{i,j} Q_{IDFT}$.

Therefore, (4) gives a channel description, where the samples are stacked in groups corresponding to users, and each of the groups contains samples corresponding to tones. It is desirable to reorganize these samples, so that they are stacked in groups corresponding to tones, and each of these groups contains samples corresponding to different users. To this end, a permutation matrix P having NL rows and NL columns is defined, which is composed of the $N \times N$ blocks $P_{i,j}$ where $i, j = 1, \ldots, L$. The block $P_{i,j}$ contains all zeros, except for a one at position (j, i). Some inspection reveals that, when matrix P is right-multiplied with a vector of size NL, it essentially reorders its elements from L groups of N components into N groups of L components.

Also, note that $P^{-1} = P^* = P$. Applying this reordering operation to both the transmitter and the receiver samples, yields:

$$\begin{aligned} \mathbf{Z} &= P\Lambda PP\mathbf{X} + P\mathbf{N}' \\ &= P\Lambda P\mathbf{U} + \mathbf{N} \\ &= T\mathbf{U} + \mathbf{N} \end{aligned} \tag{9}$$

where:

$$\mathbf{Z} = \left[\mathbf{Z}_1^T\,\mathbf{Z}_2^T\,\ldots\,\mathbf{Z}_N^T\right]^T = P\mathbf{Y} \tag{10}$$

$$\mathbf{U} = \left[\mathbf{U}_1^T\,\mathbf{U}_2^T\,\ldots\,\mathbf{U}_N^T\right]^T = P\mathbf{X} \tag{11}$$

$$\mathbf{N} = \left[\mathbf{N}_1^T\,\mathbf{N}_2^T\,\ldots\,\mathbf{N}_N^T\right]^T = P\mathbf{N}' \tag{12}$$

$$T = P\Lambda P = diag\left(\overbrace{T_1, T_2, \ldots, T_N}^{N}\right) \tag{13}$$

with $R_{\mathbf{NN}} = I$. Finally, (9) becomes:

$$\mathbf{Z}_i = T_i\mathbf{U}_i + \mathbf{N}_i\ ,\ i = 1, \ldots, N \tag{14}$$

Therefore, $\mathbf{Z}_i$, $\mathbf{U}_i$, and $\mathbf{N}_i$ contain the received samples, transmitted symbols and noise samples of all users corresponding to tone i, and T_i fully characterizes MIMO transmission within tone i. In the following, a distinction between upstream and downstream will be made by adopting the notation $T_{i,up}$ and $T_{i,down}$.

Eq. (14) indicates that crosstalk cancellation can be performed independently in each tone. Therefore, as shown in the next section, an array of canceler blocks can be employed at the CO/ONU to remove crosstalk within each tone for upstream communication. Similarly, precoder blocks can be used at the CO/ONU to pre-distort the transmitted signals within each tone, so that the received signals are crosstalk-free. However, determining the parameters of the canceler/precoder blocks relies on perfect channel matrix and noise covariance matrix knowledge at the CO/ONU. This is fairly reasonable for DSL, since the twisted pair channels are stationary, and systems can afford the luxury of training-based channel identification during initialization.

Practical Issues for the VDSL environment

A few comments are now made about the CP and synchronization assumptions with regard to the VDSL environment. The additional requirement of having a CP longer than the memory of both the transfer and the crosstalk coupling functions can be satisfied without suffering an excessive loss. Fig. 3 shows FEXT coupling measurements for loops with length of 1640 feet. Since only magnitude data is provided, linear phase is assumed in order to derive the impulse responses, and it is found that 99.9% of the signal energy is contained within 9 μsec. With a DMT block size of 4096 samples and sampling rate of 17.664 MHz, this corresponds to 159 samples. Therefore, the currently proposed CP length of 320 samples (corresponding to a 7.8% loss) is more than adequate, and it is expected that this will remain true with any typical FEXT coupling function.

The average delay of a typical twisted pair is approximately 1.5 μs/kft (see [19]). Given that VDSL loops usually have lengths shorter than 6000 ft, and with the previous DMT assumptions, the propagation delay corresponds to fewer than 160 samples. Therefore, even if "Zippering" is used, the length of the CP plus the CS does not exceed the proposed 320 samples. In cases where the channel has unusually long memory, a MIMO Time-Domain-Equalizer may be used at the CO/ONU as described in [20]. Although this technique applies to upstream communication only, a MIMO extension of the precoder proposed in [21] may be utilized for downstream communication.

III. Crosstalk Cancellation via QR Decomposition

Starting from (14), the methods to remove the crosstalk within each tone are now described first for upstream and then for downstream communication. In the following, the matrices $T_{i,up}$, $T_{i,down}$ are assumed to be non-singular (the validity of this claim and the consequences of ill-conditioning will be examined later).

A. Upstream

For upstream transmission, the co-location of the CO/ONU transceiver equipment gives the opportunity to perform joint signal processing of the received samples. The computation of the QR decomposition of matrix $T_{i,up}$ yields:

$$T_{i,up} = Q_i R_i \tag{15}$$

where Q_i is a unitary matrix and R_i is an upper triangular matrix. If the received samples are "rotated/reflected" by Q_i^*, then (14) becomes:

$$\tilde{\mathbf{Z}}_i = Q_i^* \mathbf{Z}_i \tag{16}$$

$$= R_i \mathbf{U}_i + \tilde{\mathbf{N}}_i \tag{17}$$

where $\tilde{\mathbf{N}}_i = Q_i^* \mathbf{N}_i$ has an identity covariance matrix. Since R_i is upper triangular and $\tilde{\mathbf{N}}_i$ has uncorrelated components, the input $\mathbf{U}_i$ can be recovered by back-substitution combined with symbol-by-symbol detection. Thus, a decision feedback structure is derived with the feedforward matrix being Q_i^*, and the feedback matrix being $I - R_i$. The detection of the k^{th} element of $\mathbf{U}_i$ is expressed as:

$$\left(\hat{\mathbf{U}}_i\right)_k = dec\left[\frac{1}{r_{k,k}^i}\left(\tilde{\mathbf{Z}}_i\right)_k - \sum_{j=k+1}^{L} \frac{r_{k,j}^i}{r_{k,k}^i}\left(\hat{\mathbf{U}}_i\right)_j\right] \ , \ k = L, L-1, \dots, 1 \tag{18}$$

where $r_{k,j}^i$ is the (k, j) element of R_i. Assuming that the previous decisions are correct, crosstalk is completely cancelled, and L "parallel" channels are created within each tone. The detection SNR of user k on tone i then equals:

$$SNR_k^i = |r_{k,k}^i|^2 \mathcal{E}_{k,up}^i \tag{19}$$

where $\mathcal{E}_{k,up}^i$ is the energy of $(\mathbf{U}_i)_k$.

The operations described above define the canceler block corresponding to a single tone, which is shown in Fig. 4. Combining the canceler blocks of all tones, and taking into account DMT transmission, the system for upstream Vectored-DMT transmission is obtained as in Fig. 5.

B. Downstream

For downstream transmission, joint signal processing of the transmitted symbols is allowed. The QR decomposition of $T_{i,down}^T$ results in:

$$T_{i,down}^T = Q_i R_i \tag{20}$$

where again Q_i is a unitary matrix and R_i is an upper triangular matrix. Assuming that the symbols are "rotated/reflected" by Q_i^{T*} prior to being transmitted:

$$\mathbf{U}_i = Q_i^{T*} \mathbf{U}_i' \tag{21}$$

So, with the choice of:

$$\mathbf{U}_i' = R_i^{-T} diag(R_i^T)\tilde{\mathbf{U}}_i \tag{22}$$

crosstalk-free reception is achieved, where the transmitted symbols in tone i are the elements of $\tilde{\mathbf{U}}_i$. However, this operation may cause undesired energy increase, so, utilizing the concept of Tomlinson-Harashima precoding (see [23], [24]), it is replaced by:

$$\left(\mathbf{U}_i'\right)_k = \Gamma_{M_{i,k}}\left[\left(\tilde{\mathbf{U}}_i\right)_k - \sum_{j=1}^{k-1}\frac{r_{j,k}^i}{r_{k,k}^i}\left(\mathbf{U}_i'\right)_j\right],\ k = 1, 2, \ldots, L \tag{23}$$

The following operation is performed at the receiver:

$$\left(\hat{\mathbf{Z}}_i\right)_k = \Gamma_{M_{i,k}}\left[\frac{(\mathbf{Z}_i)_k}{r_{k,k}^i}\right],\ k = 1, 2, \ldots, L \tag{24}$$

where $\Gamma_{M_{i,k}}$ is defined as:

$$\Gamma_{M_{i,k}}[x] = x - M_{i,k}d\left\lfloor\frac{x + \frac{M_{i,k}d}{2}}{M_{i,k}d}\right\rfloor \tag{25}$$

and $M_{i,k}$ is the constellation size of user k on tone i, while d is the constellation point spacing. (If x is complex, then $\Gamma_{M_{i,k}}[x] = \Gamma_{\sqrt{M_{i,k}}}[\Re(x)] + \jmath\Gamma_{\sqrt{M_{i,k}}}[\Im(x)]$.) It is shown in [22] that these operations result in:

$$\hat{\mathbf{Z}}_i = \tilde{\mathbf{U}}_{\mathbf{i}} + \left[diag(R_i^T)\right]^{-1}\mathbf{N}_i \tag{26}$$

which implies crosstalk-free reception. The detection SNR of user k on tone i equals:

$$SNR_k^i = |r_{k,k}^i|^2\mathcal{E}_{k,down}^i \tag{27}$$

where $\mathcal{E}_{k,down}^i = \tilde{\mathcal{E}}_k^i$ is the energy of $\left(\tilde{\mathbf{U}}_i\right)_k$.

The MIMO precoder described above corresponds to a single tone and is shown in Fig. 6. Combining the precoders of all tones and including the DMT transmitters and receivers, the Vectored-DMT system for downstream transmission is obtained as in Fig. 7.

C. Some Comments about Cancellation via QR

An alternative derivation of the crosstalk cancellation structures described above is given in [25], where it is proved that these are special cases of the Zero-Forcing (ZF) version of the

Generalized Decision-Feedback-Equalizer (GDFE) (see [26]). The GDFE is a MIMO extension of the well-known DFE employed for Inter-Symbol-Interference (ISI) equalization.

It can be observed that the upstream crosstalk cancellation scheme is similar to the V-BLAST architecture proposed for wireless communication with multiple antennas (see [27], [28]). Actually, based on the equivalency between V-BLAST and GDFE (shown in [29]), it is deduced that the upstream QR cancellation method is identical to the processing steps of BLAST. The main difference lies in the fact that, instead of twisted pairs suffering from crosstalk, there are multiple antennas receiving the signals of multiple transmitting antennas. Also, it is worth noting that the decision-feedback multiuser detectors described in [30] share several similarities with the QR cancellation scheme.

Under the assumption that transmit and receive filtering at the CO/ONU and at the CPE's is identical, and noise within a tone has the same statistics for all users, the reciprocity property for twisted pair transmission implies that $T_{i,up} = T_{i,down}^T$. In that case, (15) and (20) give the QR decomposition of the same matrix.

For the upstream channel, it is *always* true regardless of the loop topology, that the diagonal element of a column of T_i is larger in magnitude than the off-diagonal elements of the same column. This occurs because in upstream transmission the crosstalk coupled signal originating from a specific transmitter can never exceed the "directly" received signal of the same transmitter, and typically the magnitude difference is more than 20 dB. This can be verified in Fig. 8, which refers to an extreme case, where a 500 feet pair neighbors a 6000 feet pair (and both consist of AWG-26 copper wire). The insertion losses are shown together with the FEXT coupling losses, assuming upstream transmission and co-location at the CO/ONU. The data were obtained using the standard models adopted by the T1E1.4 comittee ([1]). Clearly, the insertion loss of a signal is always smaller than the coupling loss that it experiences when it propagates into a neighboring pair.

Visualizing the columns of T_i in vector space, it is seen that the columns are almost orthogonal to each other, which implies that Q_i is close to being an identity matrix. Thus, the magnitudes of the diagonal elements of R_i do not differ significantly from those of the diagonal elements of T_i, which indicates that QR cancellation performs almost as good as perfect crosstalk removal. This is illustrated in Fig. 9 for a two user case, where $\mathbf{t}_1$, $\mathbf{t}_2$ are the columns of T_i, $\mathbf{q}_1$, $\mathbf{q}_2$ are the columns of Q_i, r_{mn} is the (m,n) element of R, and $r_{11} = ||\mathbf{t}_1||$. The fact that $\mathbf{t}_1$, $\mathbf{t}_2$ are close to orthogonal to the axes implies that $r_{22} \approx ||\mathbf{t}_2||$. As shown in the figure, this holds for both possible detection orderings. Appendix A gives a more mathematical argument for this result.

Fig. 10 shows the upstream *channel* SNR's for the same loop topology as in Fig. 8, either assuming that no crosstalk exists, or that crosstalk is cancelled by vectored transmission. (Channel SNR is defined as equal to the SNR with unit transmission energy.) The background noise is white with Power Spectral Density (PSD) equal to -140 dBm/Hz. Since the QR decompositions differ when the columns of T_i are swapped, the two possible ordering are displayed. Clearly, the plots corresponding to the same loop overlap almost perfectly, which means that QR cancellation is as good as having no crosstalk, and that the impact of ordering can be ignored. Also, this serves as evidence that the crosstalk ZF criterion is fairly adequate, and employing an MMSE crterion will not yield noticeable improvement.

All the previous arguments were made for upstream transmission, but they can easily be extended to downstream transmission by starting with the observation that the crosstalk signals at a specific receiver can never exceed in magnitude the "directly" received signal. Alternatively, the same conclusions can be reached by using the transpose relationship between the upstream and downstream channel matrices, which in most practical situations is approximately true.

D. Practical Considerations

The computational cost incurred by the QR cancellation is decomposed into the cost of the QR decompositions and the cost associated with signal processing. DSL channels are stationary, so the QR decompositions need to be computed infrequently (during initialization). Generally, this requires $\mathcal{O}(L^3)$ flops per tone (e.g. using the Householder transform as in [32]), but it can be greatly reduced by taking advantage of the crosstalk environment characteristics. It is known that the crosstalk noise in a pair originates mostly from just three or four neighboring pairs, which implies that a typical T_i matrix is almost sparse with only three or four relatively large off-diagonal elements per row. Therefore, approximating T_i as a sparse matrix, Givens rotations can be employed to triangularize T_i with a reduced cost of $\mathcal{O}(L^2)$ flops. On the other hand, the real-time computational burden due to the canceler and precoder blocks cannot be reduced. In a straightforward implementation, the operations dominating the total cost are those of (16) and (21), and require $\mathcal{O}(L^2)$ flops per tone.

Although the assumption of perfect channel matrix knowledge is fairly reasonable in the given environment, it is still worth examining the effects of channel estimation errors. Let the estimated upstream channel matrix for tone i be:

$$\hat{T}_{i,up} = T_{i,up} + E_i \tag{28}$$

where E_i is the channel estimation error. Then, performing the QR decomposition with the reciprocity assumption gives:

$$\hat{T}_{i,up} = \hat{Q}_i \hat{R}_i - E_i \tag{29}$$

$$\hat{T}_{i,down} = \hat{R}_i^T \hat{Q}_i^T - E_i^T \tag{30}$$

where $\hat{Q}_i$, $\hat{R}_i$ are the QR factor estimates. Starting from (16), the effect on upstream communication can be computed:

$$\begin{aligned}
\tilde{\mathbf{Z}}_i &= \left(\hat{R}_i - \hat{Q}_i^* E_i\right)\mathbf{U}_i + \hat{Q}_i^* \mathbf{N}_i \Rightarrow \\
\hat{\mathbf{U}}_i &= dec\left[\mathbf{U}_i - \left[diag\left(\hat{R}_i\right)\right]^{-1} \hat{Q}_i^* E_i \mathbf{U}_i + \left[diag\left(\hat{R}_i\right)\right]^{-1} \hat{Q}_i^* \mathbf{N}_i\right] \\
&= dec\left[\left\{1 - \left[diag\left(\hat{R}_i\right)\right]^{-1} diag\left(Q_i^* E_i\right)\right\}\mathbf{U}_i - \right. \\
&\qquad \left. \left[diag\left(\hat{R}_i\right)\right]^{-1}\left[Q_i^* E_i - diag\left(Q_i^* E_i\right)\right]\mathbf{U}_i + \left[diag\left(\hat{R}_i\right)\right]^{-1} \hat{Q}_i^* \mathbf{N}_i\right]
\end{aligned} \tag{31}$$

where $dec\,[.]$ acts on the elements of a vector. So, the estimation errors impact transmission by introducing a "bias" in the detection (first term in (31)), and also by resulting in some residual crosstalk (second term in (31)). A similar analysis can be applied for downstream communication, but modulo arithmetic complicates the expressions. Ignoring the modulo operations and employing (21), (22) gives:

$$\begin{aligned}
\mathbf{Z}_i &= \left(\hat{R}_i^T \hat{Q}_i^T - E_i^T\right)\mathbf{U}_i + \mathbf{N}_i \\
&= diag\left(\hat{R}_i^T\right)\tilde{\mathbf{U}}_i - E_i^T \hat{Q}_i^{T*} \mathbf{U}_i' + \mathbf{N}_i \Rightarrow \\
\hat{\mathbf{Z}}_i &= \tilde{\mathbf{U}}_i - \left[diag\left(\hat{R}_i^T\right)\right]^{-1} E_i^T \hat{Q}_i^{T*} \hat{R}_i^{-T} diag\left(\hat{R}_i^T\right)\tilde{\mathbf{U}}_i + \left[diag\left(\hat{R}_i^T\right)\right]^{-1}\mathbf{N}_i
\end{aligned} \tag{32}$$

Just as before, the effects of the estimation errors can be separated into a detection bias term and a residual crosstalk term.

Equations (31) and (32) reveal that the impact of channel estimation errors is aggravated, when any of the diagonal elements of $\hat{R}_i$ is small. Although channel matrix singularity is almost impossible in the DSL environment, an ill-conditioned channel (implying small diagonal elements) cannot be ruled out, thus increasing the impact of channel estimation errors and posing several computational problems. Such cases arise in high frequencies (e.g. in loop topologies similar to that of Fig. 8), or in the presence of bridged taps. Nevertheless, the energy allocation algorithms proposed in Section IV prevent the occurence of such phenomena by not allowing transmission in frequencies where the diagonal elements of R_i are small.

Evidently, Vectored-DMT is successful in cancelling self-FEXT. One could argue that self-NEXT at the CO/ONU side can also be cancelled through the use of "NEXT-cancelers" operating similarly to echo-cancelers. In that case, upstream transmission is essentially crosstalk-free, which hints that downstream transmission may occupy the whole band without any impact on upstream. However, downstream transmission in the bands employed for upstream still suffers from self-NEXT. Simulations have shown that (unless the self-NEXT coupling is very weak) the performance benefits of allowing downstream transmission in the whole band are minuscule. In the following, the use of FDD is assumed for the separation of downstream and upstream transmission, thus, no NEXT-cancelers are needed.

IV. Transmission Optimization

Methods to optimize transmission are now presented, where the objective is the maximization of a weighted data rate sum. These methods refer to the problems of energy allocation in frequency, energy allocation in frequency while observing constraints on induced crosstalk, and energy allocation combined with upstream/downstream frequency selection.

A. Energy Allocation

The optimization objective is the maximization of the weighted sum of the data rates of all users:

$$\max \sum_{k=1}^{L} a_k R_k \tag{33}$$

where $a_k \geq 0$ is the weight assigned to the k^{th} user , and R_k is the achievable data rate of the k^{th} user, which may refer to either the upstream or the downstream direction. In order to compute the data rate, the gap approximation is employed (see [33]). Taking into account the fact that vectoring essentially "diagonalizes" the channel (and assuming no error propagation in upstream), the upstream and downstream achievable rates are obtained:

$$R_{k,up} = \sum_{i \in \mathcal{N}_{up}} \frac{1}{2} \log_2 \left(1 + \frac{|r^i_{k,k}|^2 \mathcal{E}^i_{k,up}}{\Gamma} \right) \tag{34}$$

$$R_{k,down} = \sum_{i \in \mathcal{N}_{down}} \frac{1}{2} \log_2 \left(1 + \frac{|r^i_{k,k}|^2 \mathcal{E}^i_{k,down}}{\Gamma} \right) \tag{35}$$

where Γ is defined as the transmission gap, and depends on the probability of error requirement, the coding gain, and the required margin. Also, $\mathcal{N}_{up}$ and $\mathcal{N}_{down}$ are the sets of upstream and downstream tone indices correspondingly, which depend on the FDD plan. The assumption of no error propagation is justified given that errors in DSL are rather infrequent.

The parameters with respect to which optimization takes place are $\mathcal{E}^i_{k,up}$ for upstream, and $\mathcal{E}^i_{k,down}$ for downstream transmission, and these are constrained by limits on the transmitted energy. In upstream transmission, the total transmit energy is constrained by:

$$\sum_{i \in \mathcal{N}_{up}} \mathcal{E}^i_k \leq \mathcal{E}_{k,up} \tag{36}$$

where $\mathcal{E}^i_k$ is the energy of $(\mathbf{U}_i)_k$ in (17), and $\mathcal{E}_{k,up}$ is the maximum allowed upstream transmitted energy of user k. Since $\mathcal{E}^i_{k,up} = \mathcal{E}^i_k$, it is deduced that:

$$\sum_{i \in \mathcal{N}_{up}} \mathcal{E}^i_{k,up} \leq \mathcal{E}_{k,up} \tag{37}$$

In downstream transmission, the total transmit energy constraint is expressed as:

$$\sum_{i \in \mathcal{N}_{down}} \mathcal{E}^i_k \leq \mathcal{E}_{k,down} \tag{38}$$

where $\mathcal{E}^i_k$ is the energy of $(\mathbf{U}_i)_k$ in (21), and $\mathcal{E}_{k,down}$ is the maximum allowed downstream transmitted energy of user k. Unfortunately, this constraint does not translate directly to a constraint for $\tilde{\mathcal{E}}^i_k = \mathcal{E}^i_{k,down}$, due to the non-linear precoding operation of (23).

However, simulation results for the previously defined extreme loop topology indicate that the covariance matrix of $\mathbf{U}_i$ can be approximated fairly accurately by the covariance matrix of $\tilde{\mathbf{U}}_i$ at all tones. For simulation purposes it was assumed that the transmission PSD is flat at -60 dBm/Hz, the noise PSD is flat at -140 dBm/Hz, and the transmission gap is 12 dB. Thus, an integer bit distribution was derived, where the maximum number of bits per tone was 12. The precoder was simulated in the time domain, and the sample covariance matrices of $\tilde{\mathbf{U}}_i$, $\mathbf{U}'_i$ and

$\mathbf{U}_i$ were computed with 10000 iterations for a number of tones. Fig. 11 shows the elements (1,1), (2,2) and (1,2) of these covariance matrices. It is seen that user 1 transmits data on all frequencies, while user 2 does not actually transmit any data at frequencies above 2 MHz. At frequencies below 2 MHz, the correlation between the users is very small. At frequencies higher than 2 MHz, the rotation/reflection operation means that the precoder output of user 2 has some energy, however the correlation still remains small. At all frequencies, the energy increase due to precoding is indiscernible.

Therefore, it is seen that the precoder does not result in considerable correlation between the transmitted signals of different users. It is reasonable to assume that this result holds generally, since the simulated loops correspond to a worst-case situation with regard to the crosstalk coupling.

Therefore, the approximation $\mathcal{E}_k^i \approx \tilde{\mathcal{E}}_k^i = \mathcal{E}_{k,down}^i$ is made, and (38) for downstream becomes:

$$\sum_{k=1}^{L} \sum_{i \in \mathcal{N}_{down}} \mathcal{E}_{k,down}^i \leq \mathcal{E}_{k,down} \tag{39}$$

With this in mind, it is seen that the energy allocation problem of (33) becomes independent for each user, and thus the a_k weights are irrelevant in this scenario. The optimization problem for each transmission direction is broken into k waterfilling problems expressed by:

$$\max \quad \sum_{i \in \mathcal{N}_{up}} \frac{1}{2} \log_2 \left(1 + \frac{|r_{k,k}^i|^2 \mathcal{E}_{k,up}^i}{\Gamma} \right) \tag{40}$$

$$\text{subject to} \quad \sum_{i \in \mathcal{N}_{up}} \mathcal{E}_{k,up}^i \leq \mathcal{E}_{k,up} \tag{41}$$

and by:

$$\max \quad \sum_{i \in \mathcal{N}_{down}} \frac{1}{2} \log_2 \left(1 + \frac{|r_{k,k}^i|^2 \mathcal{E}_{k,down}^i}{\Gamma} \right) \tag{42}$$

$$\text{subject to} \quad \sum_{i \in \mathcal{N}_{down}} \mathcal{E}_{k,down}^i \leq \mathcal{E}_{k,down} \tag{43}$$

Solutions to these problems are easy to derive and simulation results are shown in Section V-A.

Specifically for downstream transmission, a constraint on the total transmitted energy of all users may also be relevant:

$$\sum_{k=1}^{L} \sum_{i \in \mathcal{N}_{down}} \mathcal{E}_{k,down}^i \leq \mathcal{E} \tag{44}$$

but will not be pursued any further.

It is worth reminding that the resulting transmission spectra are optimal only in the context of Vectored-DMT. Reference [36] describes a strategy to obtain the optimal spectra of the vector multiple access channel, of which the upstream DSL channel is a special case.

B. Energy Allocation with Power Backoff

As explained previously, all users in vectored transmission correspond to a group of neighboring twisted pairs. This does not preclude the operation of other "alien" DSL systems in neighboring twisted pairs, which on one hand cause crosstalk into the vectored systems, and on the other hand suffer from crosstalk originating from the vectored systems. The current approach in dealing with this problem is to impose limits on the transmitted PSD's, so that the performance of systems is not excessively affected by crosstalk.

Additionally, VDSL systems suffer from the fact that upstream signals on short lines detrimentally affect upstream performance on long lines (similarly to the near-far situation in wireless communications). In order to avoid imposing an overly restrictive universal PSD mask, power back-off methods have been proposed (see [34]), which effectively make the PSD mask dependent on the loop length of the specific user. A similar scenario, where the downstream communication of neighboring DSL systems may suffer considerably is shown in Fig. 12. Note that this situation will occur increasingly often as the installation of ONU's progresses, while twisted pair connections to the CO's remain.

Vectoring combined with full channel matrix knowledge can prove effective in limiting the crosstalk induced by vectored systems, without resorting to the introduction of a universal PSD mask, or the use of power back-off methods (which do not take into account knowledge about crosstalk coupling resulting from matrix channel identification).

Eq. (9) can be augmented to include the received samples of alien systems:

$$\begin{bmatrix} \mathbf{Z} \\ \mathbf{Z}_n \end{bmatrix} = \begin{bmatrix} T & C_n \\ C & T_n \end{bmatrix} \begin{bmatrix} \mathbf{U} \\ \mathbf{U}_n \end{bmatrix} + \begin{bmatrix} \mathbf{N} \\ \mathbf{N}_n \end{bmatrix} \tag{45}$$

where $\mathbf{Z}_n$, $\mathbf{U}_n$, and $\mathbf{N}_n$ are vectors of the received samples, of the transmitted symbols and of the noise samples respectively of the alien systems. The definitions of the block matrices C, C_n and T_n depend on both the channel and the characteristics of the alien DSL systems, and although T is block diagonal, this property will not generally hold for the other matrices.

In the case where $\mathbf{Z}$ and $\mathbf{Z}_n$ correspond to systems belonging to different service providers, it is currently very difficult to identify the crosstalk coupling matrices C and C_n. This occurs, because in the current unbundling framework there is no way for the first provider to obtain access to either $\mathbf{Z}_n$ or $\mathbf{U}_n$, and similarly for the second provider. However, the "third party" concept (see [35]) overcomes this difficulty by introducing an impartial third-party site, which captures all transmitted and received data, and is capable of producing estimates of the crosstalk coupling matrices.

With this in mind, the knowledge of C can be taken for granted, and a number of criteria can be devised to mitigate the effect of crosstalk into the alien systems. A possible solution would be to transmit in the null-space of C, thus causing theoretically zero crosstalk. Of course, this implies that some dimensions of $\mathbf{U}$ must be "sacrificed", so that transmission in the null-space is achieved. Although one could envisage scenarios where application of this method may be promising, this will not be further pursued here.

Limiting the FEXT in the mean square sense, results in the following conditions:

$$\begin{aligned}\sum_{k=1}^{L}\sum_{i\in\mathcal{N}_{up}}|c_{j,(k-1)N+i}|^2\mathcal{E}_{k,up}^{i} &\leq \mathcal{E}_{j,up}^{c}\ ,\ j=1,\ldots,MN_N &(46)\\ \sum_{k=1}^{L}\sum_{i\in\mathcal{N}_{down}}|c_{j,(k-1)N+i}|^2\mathcal{E}_{k,down}^{i} &\leq \mathcal{E}_{j,down}^{c}\ ,\ j=1,\ldots,MN_N &(47)\end{aligned}$$

where M is the number of neighboring systems, N_N is the number of "dimensions" (e.g. number of tones) per neighboring system, $\mathcal{E}_{j,up}^{c}$, $\mathcal{E}_{j,down}^{c}$ are the maximum allowable crosstalk energies in sample j of the neighboring systems for upstream and downstream, and $c_{j,l}$ is the (j,l) element of the $MN_N \times LN$ matrix C. Note that this approach can be generalized, so that both FEXT and NEXT are restricted.

The set of inequalities in (46) and (47), combined correspondingly with those of (37) and (39), form a set of linear inequality constraints. Including the rate maximization objective of (33) yields the following optimization problems:

$$\begin{aligned}\max\quad & \sum_{k=1}^{L}a_k\sum_{i\in\mathcal{N}_{up}}\frac{1}{2}\log_2\left(1+\frac{|r_{k,k}^{i}|^2\mathcal{E}_{k,up}^{i}}{\Gamma}\right) &(48)\\ \text{subject to}\quad & \sum_{i\in\mathcal{N}_{up}}\mathcal{E}_{k,up}^{i}\leq\mathcal{E}_{k,up}\ ,\ k=1,\ldots,L &(49)\\ & \sum_{k=1}^{L}\sum_{i\in\mathcal{N}_{up}}|c_{j,(k-1)N+i}|^2\mathcal{E}_{k,up}^{i}\leq\mathcal{E}_{j,up}^{c}\ ,\ j=1,\ldots,MN_N &(50)\end{aligned}$$

and

$$\max \quad \sum_{k=1}^{L} a_k \sum_{i \in \mathcal{N}_{down}} \frac{1}{2} \log_2 \left(1 + \frac{|r_{k,k}^i|^2 \mathcal{E}_{k,down}^i}{\Gamma} \right) \tag{51}$$

$$\text{subject to} \quad \sum_{i \in \mathcal{N}_{down}} \mathcal{E}_{k,down}^i \leq \mathcal{E}_{k,down} \ , \ k = 1, \ldots, L \tag{52}$$

$$\sum_{k=1}^{L} \sum_{i \in \mathcal{N}_{down}} |c_{j,(k-1)N+i}|^2 \mathcal{E}_{k,down}^i \leq \mathcal{E}_{j,down}^c \ , \ j = 1, \ldots, MN_N \tag{53}$$

The objective functions are convex (since they are sums of log functions), and the constraints form convex sets (because they are linear inequalities). Thus, solutions can be efficiently produced using convex programming techniques, and such an example is given in Section V-B. It is worth observing that other restrictions (such as PSD masks or bit caps) are easy to include in the above optimization problems, since they only require the introduction of linear inequality constraints, which preserve the convexity of the problem.

C. Energy Allocation and Upstream/Downstream Frequency Selection

Although all existing DSL systems employing FDD have a fixed upstream/downstream frequency duplexing band plan, there is evidence that a dynamically configured band plan offers significant advantages (see [37]). Such a plan is common for all users, but is determined during modem initialization, depending on the specific transmission environment, as well as on the user requirements. The disadvantages of a fixed band plan are exemplified in the presence of bridged taps, where transmission in one direction may face a disproportionate degradation, while transmission in the opposite direction may remain unscathed. On the other hand, adopting a dynamic plan in such a case enables a fairer distribution of the impact on both upstream and downstream.

The optimization objective is now expressed by:

$$\max \sum_{k=1}^{L} (a_{k,up} R_{k,up} + a_{k,down} R_{k,down}) \tag{54}$$

where $a_{k,up}, a_{k,down} \geq 0$ are the weights assigned to upstream and downstream transmission for user k, and $R_{k,up}$, $R_{k,down}$ are the achievable upstream and downstream rates of user k. Here, the optimization parameters involve not just the energies assigned but also the selection of upstream/downstream tones. However, if (34) and (35) are used, the partition of the set of tones into $\mathcal{N}_{up}$ and $\mathcal{N}_{down}$ is a binary constrained problem, whose solution has very high complexity.

Instead, the binary constraints can be relaxed, which greatly simplifies the computations. This idea has previously been exploited in [38] for the computation of the FDMA capacity of the Gaussian multiple access channel in the presence of ISI. In more detail, it is initially assumed that each tone is time-shared between upstream and downstream, thus obtaining the following achievable rates:

$$R_{k,up} = \sum_{i=1}^{N} t_{i,up} \frac{1}{2} \log_2 \left(1 + \frac{|r_{k,k}^i|^2 \mathcal{E}_{k,up}^i}{t_{i,up}\Gamma}\right) \tag{55}$$

$$R_{k,down} = \sum_{i=1}^{N} t_{i,down} \frac{1}{2} \log_2 \left(1 + \frac{|r_{k,k}^i|^2 \mathcal{E}_{k,down}^i}{t_{i,down}\Gamma}\right) \tag{56}$$

where $t_{i,up}$, $t_{i,down}$ describe the fraction of time in tone i used for upstream and downstream transmission correspondingly, and $t_{i,up} + t_{i,down} = 1$, $t_{i,up}, t_{i,down} \geq 0$. The existence of $t_{i,up}$ and $t_{i,down}$ in the denominators inside the log expressions implies that the assigned energy is "boosted", since transmission takes place over only some fraction of time. The energy constraints for user k are:

$$\sum_{i=1}^{N} \mathcal{E}_{k,up}^i \leq \mathcal{E}_{k,up} \tag{57}$$

$$\sum_{i=1}^{N} \mathcal{E}_{k,down}^i \leq \mathcal{E}_{k,down} \tag{58}$$

Therefore, the optimization problem has the following form:

$$\max \quad \sum_{k=1}^{L} \left[a_{k,up} \sum_{i=1}^{N} t_{i,up} \frac{1}{2} \log_2 \left(1 + \frac{|r_{k,k}^i|^2 \mathcal{E}_{k,up}^i}{t_{i,up}\Gamma}\right) + a_{k,up} \sum_{i=1}^{N} t_{i,down} \frac{1}{2} \log_2 \left(1 + \frac{|r_{k,k}^i|^2 \mathcal{E}_{k,down}^i}{t_{i,down}\Gamma}\right)\right] \tag{59}$$

$$\text{subject to} \quad \sum_{i=1}^{N} \mathcal{E}_{k,up}^i \leq \mathcal{E}_{k,up} \ , \ k = 1, \dots, L \tag{60}$$

$$\sum_{i=1}^{N} \mathcal{E}_{k,down}^i \leq \mathcal{E}_{k,down} \ , \ k = 1, \dots, L \tag{61}$$

$$t_{i,up} + t_{i,down} = 1 \ , \ i = 1, \dots, N \tag{62}$$

The objective function is convex, because it is a sum of functions of the form $x \log(1 + y/x)$, which are shown to be convex in $x, y \geq 0$ in [38]. The constraint sets are clearly convex, since they are defined by linear inequalities. Therefore, the problem is convex and a variety of methods can be used to efficiently derive a solution.

Still, such a solution would actually yield a hybrid between an FDD and a Time Division Duplexing (TDD) implementation. Since an FDD implementation is required, an approximate such solution is obtained by rounding $t_{i,up}$ and $t_{i,down}$. Naturally, this is suboptimal, but when the number of tones is fairly large, it will be adequately close to the optimal solution. Simulation results in Section V-C validate this claim. Note that the power-backoff constraints of the previous subsection can also be included in the problem formulation without considerably affecting the difficulty of obtaining a solution.

In all the previous problems, the objective has been the maximization of a weighted data rate sum. It should be evident that by adjusting the weights, one obtains different surface points of the data rate region achievable by vectored transmission, and thus the whole multi-dimensional surface can be traced. However, visualizing the inherent tradeoffs becomes difficult when the weighted sums include more than three terms. The practical question, which will likely be posed to a service provider, is whether the vectored system can support a set of rate requirements, and if so, what is the energy allocation achieving the requirements. Appendix B shows that this problem actually has a duality relationship with the weighted data rate sum problem, and thus the weighted sum problem provides an alternative method to solve the "feasibility" problem.

V. Simulation Results

The simulation results presented next refer to the VDSL environment. The general parameters and the DMT specifications (shown in Table I) are from [1], [4], and [3], [5] correspondingly. The specified coding gain, noise margin and probability of error target lead to a transmission gap of 12 dB. In the following, infinite granularity for the number of bits on each tone is assumed, and error propagation effects in upstream are ignored. Deliberately, no bit cap was taken into account, since the purpose of the simulations is to demonstrate the achievable upper bounds.

A. *Results for Energy Allocation Optimization and Channel Estimation Errors*

The energy allocation method of Subsection IV-A amounts to performing water-filling independently for each user. The upstream/downstream frequency division is chosen as shown in Table II, and the noise models used are specified in [39]. A "uniform" loop topology with 20 users is assumed, meaning that all loops have the same length. Figures 13 and 14 compare the achievable data rates of "conventional" DMT and Vectored-DMT versus the loop length with either noise A or noise F, for upstream and downstream transmission respectively. Although no PSD constraint was enforced, the resulting PSD's were observed to be typically below -52 dBmW/Hz, and the peak never exceeded the value of -45 dBmW/Hz.

Evidently, vectoring improves performance significantly, and this can be viewed in two ways: For a given loop length, VDMT allows the achievement of much higher data rates. The rate increases are considerable for lengths shorter than 4500 feet in noise A, or 3500 feet in noise F. The gains become spectacular in short loops, where obviously transmission is FEXT-limited. Also, it is seen that VDMT extends the maximum loop length given a data rate requirement. As an example, in Fig. 14, a downstream rate requirement of 50 Mbps limits DMT systems with noise of type A to loops shorter than 1150 feet, while employing VDMT extends the reach to 2650 feet. Simulations not shown here indicate that including a bit cap of 11 bits per tone has noticeable effects only in loops shorter than 1500 ft, where upstream rates are restricted to 80 Mbps, while downstream rates are restricted to 90 Mbps.

Next, the impact of channel estimation errors in a worst-case scenario is assessed. Equations (31) and (32) indicate that the effects of such errors are residual crosstalk, which may be treated as noise, and detection biasing, which can be modeled as a decrease in the minimum distance of the constellation points (e.g. see [40]). Here, the detection bias is ignored, and simulation results are shown, where the energy allocation algorithm takes into account estimation errors by considering them as extra noise. By making the approximations $Q_i \approx I$ and that E_i has zero diagonal elements, the covariance matrix of the extra noise term in tone i is $\left[diag\left(\hat{R}_i\right)\right]^{-1} E_i E\left(\mathbf{U}_i\mathbf{U}_i^*\right) E_i \left[diag\left(\hat{R}_i\right)\right]^{-*}$ for upstream, and $\left[diag\left(\hat{R}_i^T\right)\right]^{-1} E_i^T E\left(\mathbf{U}_i'\mathbf{U}_i'^*\right) E_i^{T*} \left[diag\left(\hat{R}_i^T\right)\right]^{-*}$ for downstream. Assuming a worst-case situation, E_i has all its off-diagonal elements equal to ε and the variance of $\mathbf{U}_i$ and $\mathbf{U}_i'$ is σ_u^2. With the additional approximation of uncorrelated components of $\mathbf{U}_i'$, it is found that the energy of the noise contribution due to channel estimation errors in tone i for user k is:

$$\sigma_{i,k}^2 = (L-1)|r_{k,k}^i|^{-2}|\varepsilon|^2\sigma_u^2 \tag{63}$$

holding for both upstream and downstream transmission.

Making the pessimistic assumption that $\sigma_u^2 = -50$ dBm, Figures 15 and 16 show the impact of channel estimation errors on the achievable upstream and downstream data rates for the noise A environment. The curves are parameterized by the *relative* estimation error $|\varepsilon|/|r_{k,k}^i|$, and it is seen that as channel identification becomes less accurate, Vectored-DMT performance diminishes to that of DMT.

B. Example of Energy Allocation with Power Back-Off Optimization

Next, an example is given where energy allocation is further constrained by the requirement that some "alien" DSL system must also be protected against crosstalk from the vectored system. Assume that there are two vectored users and a single affected CO-based system as in Fig. 12, where obviously downstream transmission from the CO is severely impaired. The loop lengths of the vectored users are 2500 ft and 2000 ft, the loop length of the affected user is 3000 ft, and the crosstalk coupling length between the vectored and the alien system is 1500 ft.

Assuming the existence of a VDSL system on the affected line, a possible strategy to mitigate the impact of crosstalk is to postulate that the induced crosstalk must not exceed the crosstalk PSD caused by a certain number of VDSL disturbers with loop lengths equal to the length of the affected line (this is analogous to the power back-off method of [41]). Employing the well-known T1E1.4 FEXT model, (47) must then be satisfied with:

$$\mathcal{E}^{c}_{j,down} = PSD(f_j) k_{FEXT} l f_j^2 |H(f_j)|^2 \left(\frac{N_d}{49}\right)^{0.6}, \ j = 1, \ldots, N \tag{64}$$

where f_j is the center frequency of tone j, $H(f)$ and l are the transfer function and the length of the affected line (in feet), $PSD(f)$ is the VDSL PSD mask, N_d is the number of disturbers, and k_{FEXT} is a constant equal to $7.999 \cdot 10^{-20}$. In the following, the coefficients $c_{j,l}$ of (47) were computed by making the approximation that signals transmitted in a tone do not cause crosstalk in neighboring tones, and N_d was chosen to be 20.

Figure 17 shows the achievable downstream rate curve, where the different points are obtained by adjusting the weights in (33). The nonlinear programming package MINOS (see [42]) is used, with all transmission parameters set as before, and assuming noise of type A. The dashed curve indicates the achievable rate bounds in the absence of crosstalk constraints. The practical use of this result is that a service provider may perform power-backoff in a "selective" way, so that the performance impact can be distributed according to the service priorities.

C. Example of Joint Energy Allocation and Frequency Planning

Finally, it is assumed that the frequency plan for upstream/downstream duplexing is allowed to vary in each vectored bundle depending on the loop characteristics and the service requirements. Results for two users are shown, where user 1 and user 2 have loop lengths of 1500 ft and 3000 ft correspondingly. Noise A is considered, and all transmission parameters are the same as before, with the difference being that the radio bands are now ignored and the number of tones is set to 128 in order to reduce the computational burden. The MINOS package is again used.

By having $\alpha_{1,up} = \alpha_{2,up}$, $\alpha_{1,down} = \alpha_{2,down}$, and adjusting the upstream/downstream weights, the rate curves of Fig. 18 are produced. The rates achieved with the currently proposed VDSL band plans 997 and 998 (see [43]) are also shown. The optimal frequency allocation for almost symmetric transmission is shown in Fig. 19. The corresponding achievable rates are $R_{1,up} = 89.7$ Mbps, $R_{1,down} = 90.8$, $R_{2,up} = 32.0$ Mbps, and $R_{2,down} = 29.3$ Mbps, which are obtained with $\alpha_{1,up} = \alpha_{2,up} = 0.4998$.

VI. Conclusions

This paper presented the concept of vectored transmission applied to DSL systems. The method can be summarized as employing joint signal processing techniques at the CO/ONU site, and thus taking advantage of user coordination to mitigate crosstalk. Structures were presented, where crosstalk is removed individually in each tone, and it was demonstrated that these perform very close to perfect crosstalk cancellation. Furthermore, adopting transmission optimization techniques allowed extra performance benefits, while at the same time transmission constraints (such as emission constraints) were taken into account. In these cases, weighted data rate sum maximization criteria were used.

The implementation of vectored systems holds the promise to greatly improve DSL systems, offering advantages to both customers and service providers. At the cost of synchronization of all users, channel matrix identification, and extra signal processing at the CO/ONU, considerable data rate increases can be achieved for both upstream and downstream transmission. Universal PSD mask constraints can be replaced by flexible power back-off strategies adaptable to the specific crosstalk environment. Additionally, fixed frequency duplexing plans can be substituted by dynamically determined plans tailored to the specific channel and user requirements. In conclusion, service providers are given the freedom to define different classes of service and to allocate the available resources according to the requirements of each class.

Appendix

I. Perturbation Bounds for the QR Decomposition

With the presumption of weak coupling between neighboring pairs, a bound can be obtained regarding how close the QR approach performs to perfect cancellation. If T_i is considered as a perturbation of the matrix corresponding to crosstalk-free transmission, then:

$$T_i = diag(T_i) + E_i^p \tag{65}$$

where E_i^p is the perturbation term, which consists of the off-diagonal elements of T_i. Obtaining the QR decompositions of T_i and $diag(T_i)$ yields:

$$diag(T_i) = \tilde{Q}_i \tilde{R}_i = \tilde{R}_i \tag{66}$$

$$T_i = Q_i R_i \tag{67}$$

where obviously $\tilde{Q}_i = I$. Making use of bounds derived in [31] (holding to the first order), one finds:

$$||R_i - diag(T_i)||_F \leq \sqrt{2}\frac{|t^i_{max}|}{|t^i_{min}|}\sqrt{\sum_{k\neq j}|t^i_{k,j}|^2} \Rightarrow \tag{68}$$

$$||diag(R_i) - diag(T_i)||_F \leq \sqrt{2}\frac{|t^i_{max}|}{|t^i_{min}|}\sqrt{\sum_{k\neq j}|t^i_{k,j}|^2} \tag{69}$$

where $||.||_F$ denotes the Frobenius norm, t^i_{max} is the diagonal element of T_i with the maximum magnitude, t^i_{min} is the diagonal element of T_i with the minimum magnitude, and $t^i_{k,j}$ is the (k, j) element of T_i. Since the diagonal elements of R_i determine the detection SNR's, it is deduced from (69) that the QR cancellation technique performs very close to perfect cancellation in cases where the coupling is weak and when the loop lengths are rather uniform.

II. Duality between the Feasibility Problem and the Weighted Data Rate Sum Maximization Problem

The rate feasibility problem is expressed by the following requirements:

$$R_k(\mathbf{E}) \geq r_k \ , \ k = 1, \ldots, L \tag{70}$$

where r_k is the rate requirement, $R_k(\mathbf{E})$ is the rate of user k, and $\mathbf{E}$ represents the energy allocation. Here, the requirements refer to either upstream or downstream, but (70) (as well as the conclusions that follow) can be extended to a joint downstream/upstream feasibility problem. Additionally, the domain of $\mathbf{E}$ is defined, such that the relevant energy constraints (for instance those of (37), (39), (46), or (47)) are satisfied.

The functions $R_k(\mathbf{E})$ are concave in $\mathbf{E}$, since their domains are convex sets, and they are defined as sums of log functions. Thus, the feasibility problem may be posed as the following convex problem:

$$\min \quad 0 \tag{71}$$

$$\text{subject to} \quad -R_k(\mathbf{E}) + r_k \leq 0 \ , \ k = 1, \ldots, L \tag{72}$$

where any feasible solution is actually acceptable. Reference [44] describes the so-called Phase-I method, which can be employed here to provide a solution. An alternative formulation is:

$$\min \quad -R_L(\mathbf{E}) + r_L \tag{73}$$

$$\text{subject to} \quad -R_k(\mathbf{E}) + r_k \leq 0 \ , \ k = 1, \ldots, L-1 \tag{74}$$

Here, the problem is feasible, if and only if the optimal value is found to be smaller or equal to zero. Interior point methods may be applied, where the algorithm may terminate as soon as some value smaller or equal to zero is achieved. In that case, the problem is declared feasible and the corresponding $\mathbf{E}$ is produced. If the optimum value is found to be larger than zero, then the problem is declared infeasible.

Forming the Lagrangian gives:

$$L(\mathbf{E}, a_1, a_2, \ldots, a_L) = \sum_{k=1}^{L-1} a_k \left(-R_k(\mathbf{E}) + r_k\right) - R_L(\mathbf{E}) + r_L \tag{75}$$

and the dual function is:

$$g(a_1, a_2, \ldots, a_L) = \min_{\mathbf{E}} L(\mathbf{E}, a_1, a_2, \ldots, a_L) = \sum_{k=1}^{L-1} a_k r_k + r_L - \max_{\mathbf{E}} \left[\sum_{k=1}^{L-1} a_k R_k(\mathbf{E}) + R_L(\mathbf{E})\right] \tag{76}$$

The function $g(a_1, a_2, \ldots, a_L)$ is concave as the minimum of a set of affine functions. The dual problem then is:

$$\max_{a_1, a_2, \ldots, a_L} \quad g(a_1, a_2, \ldots, a_L) = \tag{77}$$

$$\max_{a_1, a_2, \ldots, a_L} \quad \left[\sum_{k=1}^{L-1} a_k r_k + r_L - \max_{\mathbf{E}} \left[\sum_{k=1}^{L-1} a_k R_k(\mathbf{E}) + R_L(\mathbf{E})\right]\right] \tag{78}$$

where $a_k \geq 0, k = 1, \ldots, L$. In the above, the weighted data rate sum maximization problem appears as a subproblem.

Let d^* be the optimal value of (78) and p^* be the optimal value of (73). Slater's condition (see [44]) states that if (74) are strictly feasible (i.e. there exists an $\mathbf{E}$, such that $-R_k(\mathbf{E}) + r_k < 0$ for $k = 1, \ldots, L-1$), then $d^* = p^*$. Therefore, solving the dual problem provides an alternative way to answer the feasibility question. Again, interior point methods may be employed, which may terminate as soon as a positive value is achieved, and in which case the problem is declared infeasible. On the other hand, if the optimal value is found to be negative, then the problem is declared feasible and the corresponding $\mathbf{E}$ is produced.

References

[1] "Very-high-speed digital subscriber lines (VDSL) metallic interface, Part 1: Functional requirements and common specification," *T1E1.4 Contribution 2000-009R3,* Editor: Q. Wang (available at http://www.t1.org/).

[2] "Very-high-speed digital subscriber lines (VDSL) metallic interface, Part 2: Technical specification of a single-carrier modulation (SCM) transceiver," *T1E1.4 Contribution 2000-011R3,* Editor: Vladimir Oksman (available at http://www.t1.org/).

[3] "Very-high-speed digital subscriber lines (VDSL) metallic interface, Part 3: Technical specification of a multi-carrier modulation transceiver," *T1E1.4 Contribution 2000-013R4,* Editor: Sigurd Schelstraete (available at http://www.t1.org/).

[4] "Transmission and multiplexing (TM); Access transmission systems on metallic access cables; Very high speed digital subscriber line (VDSL); Part 1: Functional requirements," *ETSI Doc. Nb. TS 101 270-1, Ver. 1.2.1* (available at http://www.etsi.org/).

[5] "Transmission and multiplexing (TM); Access transmission systems on metallic access cables; Very high speed digital subscriber line (VDSL); Part 2: Transceiver specification," *ETSI Doc. Nb. TS 101 270-2, Ver. 1.1.1* (available at http://www.etsi.org/).

[6] T. Starr, J.M. Cioffi and P.J. Silverman, "Twisted pair transmission," Chapter 3 of *Understanding Digital Subscriber Line Technology,* Upper Saddle River NJ: Prentice Hall, 1999.

[7] J.M. Cioffi, "Unbundled DSL evolution," *T1E1.4 Contribution 2001-088,* Los Angeles, CA, Feb. 2001 (also available at http://www.t1.org/).

[8] M.L. Honig, K. Steiglitz and B. Gopinath, "Multichannel signal processing for data communications in the presence of crosstalk," *IEEE Trans. Commun.,* vol. 38, no. 4, pp. 551-558, Apr. 1990.

[9] M.L. Honig, P. Crespo and K. Steiglitz, "Suppression of near- and far-end crosstalk by linear pre- and post-Filtering," *IEEE J. on Select. Areas Commun.,* vol. 10, no. 3, pp. 614-629, Apr. 1992.

[10] G. Tauböck and W. Henkel, "MIMO systems in the subscriber-line network," *5th International OFDM Workshop,* Hamburg, Germany, pp. 18.1-18.3, Sep. 2000.

[11] G.G. Raleigh and J.M. Cioffi, "Spatio-temporal coding for wireless communication," *IEEE Trans. Commun.,* vol. 46, no. 3, pp. 357-366, Mar. 1998.

[12] G.D. Golden, J.E. Mazo and J. Salz, "Transmitter design for data transmission in the presence of a data-like interferer," *IEEE Trans. Commun.,* Vol. 43, No. 2/3/4, pp. 837-850, Feb./Mar./Apr. 1995.

[13] P.S. Kumar and S. Roy, "Optimization for crosstalk suppression with noncoordinating users," *IEEE Trans. Commun.,* Vol. 44, No. 7, pp. 894-905, Jul. 1996.

[14] M. Abdulrahman and D.D. Falconer, "Cyclostationary crosstalk suppression by decision feedback equalization on digital subscriber loops," *IEEE J. Select. Areas Commun.,* vol. 10, no. 3, pp. 640-649, Apr. 1992.

[15] A. Peled and A. Ruiz, "Frequency domain data transmission using reduced computational complexity algorithms," *IEEE ICASSP,* Denver, CO, pp. 964-967, Apr. 1980.

[16] F. Sjöberg, M. Isaksson, R. Nilson, P. Ödling, S.K. Wilson and P. O. Börjesson, "Zipper: A duplex method for VDSL based on DMT," *IEEE Trans. Commun.,* vol. 47, no. 8, pp. 1245-1252, Aug. 1999.

[17] F. Sjöberg, R. Nilson, M. Isaksson, P. Ödling and P. O. Börjesson, "Asynchronous zipper," *Proc. of ICC '99,* Vancouver, Canada, pp. 231-235, 1999.

[18] S.B. Weinstein and P.M. Ebert, "Data transmission by frequency-division multiplexing using the discrete fourier transform," *IEEE Trans. Commun. Technology,* vol. 19, pp. 626-634, Oct. 1971.

[19] J.A.C. Bingham, "The DSL as a medium for high-speed data", Chapter 3 of *Multicarrier Modulation and xDSL,* New York: Wiley, 2000.

[20] N. Al-Dhahir, "FIR channel-shortening equalizers for MIMO ISI channels," *IEEE Trans. Commun.,* to appear.

[21] K.W. Cheong, J.M. Cioffi, "Precoder for DMT with insufficient cyclic prefix," *Proc. of ICC '98,* Atlanta, GA, pp. 339-343, 1998.

[22] G. Ginis and J.M. Cioffi, "A multi-user precoding scheme achieving crosstalk cancellation with applications to DSL systems," *Proc. 34th Asilomar Conf. on Signals, Systems, and Computers,* Pacific Grove, CA, pp. 1627-1631, Oct. 2000.

[23] M. Tomlinson, "New automatic equaliser employing modulo arithmetic," *Electr. Letters,* Vol. 7, pp. 138-139, Mar. 1971.

[24] H. Harashima and H. Miyakawa, "Matched-transmission technique for channels with intersymbol interference," *IEEE Trans. Commun.,* Vol. COM-20, No. 4, pp. 774-780, Aug. 1972.

[25] G. Ginis and J. M. Cioffi, "Vectored-DMT: A FEXT canceling modulation scheme for coordinating users," *ICC 2001,* Helsinki, Finland, accepted for presentation.

[26] J. M. Cioffi and G. D. Forney, "Generalized decision-feedback equalization for packet transmission with ISI and gaussian noise," Chapter 4, *Communication, Computation, Control and Signal Processing (A tribute to Thomas Kailath),* Editors: A. Paulraj, V. Roychowdhury, and C. Schaper. Boston MA: Kluwer, 1997, pp. 79-127.

[27] P.W. Wolniansky, G.J. Foschini, G.D. Golden, R.A. Valenzuela, "V-BLAST: An architecture for realizing very high data rates over the rich-scattering wireless channel," *URSI ISSSE '98*, Pisa, Italy, pp. 295-300, 1998.

[28] G. J. Foschini, G. D. Golden, R. A. Valenzuela and P. W. Wolniansky, "Simplified processing for high spectral efficiency wireless communication employing multi-element arrays," *IEEE JSAC,* vol. 17, no. 11, pp. 1841-1852, Nov. 1999.

[29] G. Ginis and J. M. Cioffi, "On the relation between BLAST and the GDFE," *submitted to IEEE Commun. Letters.*

[30] S. Verdu, *Multiuser Detection.* Cambridge University Press, 1998.

[31] G.W. Stewart, "On the perturbation of LU, Cholesky, and QR factorizations," SIAM J. Matrix Anal. Appl., vol.14, no. 4, pp. 1141-1145, Oct. 1993.

[32] G.H. Golub and C.F. Van Loan, "Orthogonalization and least squares," Chapter 5 of *Matrix Computations,* Third Edition, Baltimore MD: Johns Hopkins University Press, 1996.

[33] J.M. Cioffi, "Course notes for EE379A - Digital Communication," Stanford University (available at http://www.stanford.edu/class/ee379a/reader.html).

[34] K.S. Jacobsen, "Methods of upstream power backoff on very high-speed digital subscriber lines," *IEEE Commun. Mag.,* pp. 210-216, Mar. 2001.

[35] C. Zeng, Carlos Aldana, Atul Salvekar and J.M. Cioffi, "Crosstalk identification in xDSL systems," *submitted to IEEE JSAC special issue on Multiuser Detection.*

[36] W. Yu, W. Rhee, J.M. Cioffi and S. Boyd, "Iterative waterfilling for vector multiple access channels", *accepted for presentation at Int. Symp. Information Theory (ISIT),* 2001.

[37] J.M. Cioffi, "Proposal for study of dynamic spectrum management for the evolving unbundling architecture of DSL," *T1E1.4 Contribution 2001-090,* Los Angeles, CA, Feb. 2001 (available at http://www.t1.org/).

[38] W. Yu and J.M. Cioffi, "FDMA capacity of the gaussian multiple access channel with ISI," *Proc. of ICC 2000,* New Orleans, LA, pp. 1365-1369, 2000.

[39] V. Oksman and J.M. Cioffi, "Noise models for VDSL performance verification," *T1E1.4 Contribution 99-438R2,* Clearwater, Florida, December 1999 (available at http://www.t1.org/).

[40] A. Leke and J.M.Cioffi, "Impact of Imperfect Channel Knowledge on the Performance of Multicarrier Systems," *Proc. Globecom '98*, Sydney, Australia, pp. 951-955, 1998.

[41] K.S. Jacobsen, "The equalized-FEXT upstream power cutback method to mitigate the near-far problem in VDSL," *ETSI TM6 Contribution TD05,* Sophia Antipolis, France, Nov. 1998 (available at http://www.etsi.org/).

[42] "SOL optimization software," *Stanford Business Software Inc,* (http://www.sbsi-sol-optimize.com/index.html).

[43] J. Cook, "Proposed VDSL band plans," *ETSI TM6 Contribution TD13,* Montreux, Switzerland, Feb. 2000 (available at http://www.etsi.org/).

[44] S. Boyd and L. Vandenberghe, "Course notes for EE364 - Convex Optimization," Stanford University (available at http://www.stanford.edu/class/ee364/).

List of Tables

TABLE I

Simulation parameters.

Number of DMT Tones	4096
Tone Width	4.3125 KHz
Symbol Rate	4 KHz
Coding Gain	3.5 dB
Noise Margin	6 dB
Symbol Error Probability	$< 10^{-7}$
Maximum Power	14.5 dBmW
Cable Type	26-Gauge (0.4 mm)
Source/Load Resistance	100 Ohm
Amateur Radio Bands	as specified in [1]

TABLE II

Upstream-downstream frequency assignment.

Upstream Bands (MHz)	Downstream Bands (MHz)
0.03 – 0.138	0.138 – 2.5
2.5 – 3.75	3.75 – 7.5
7.5 – 14.5	14.5 – 17.7

List of Figures

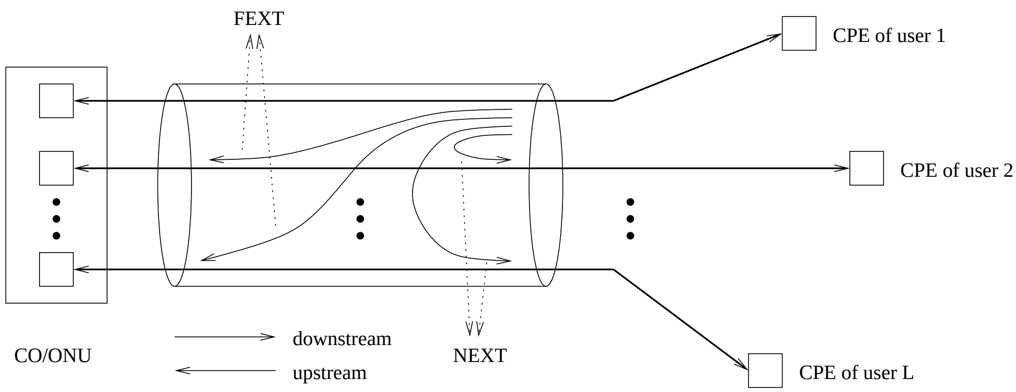

Fig. 1. DSL system topology.

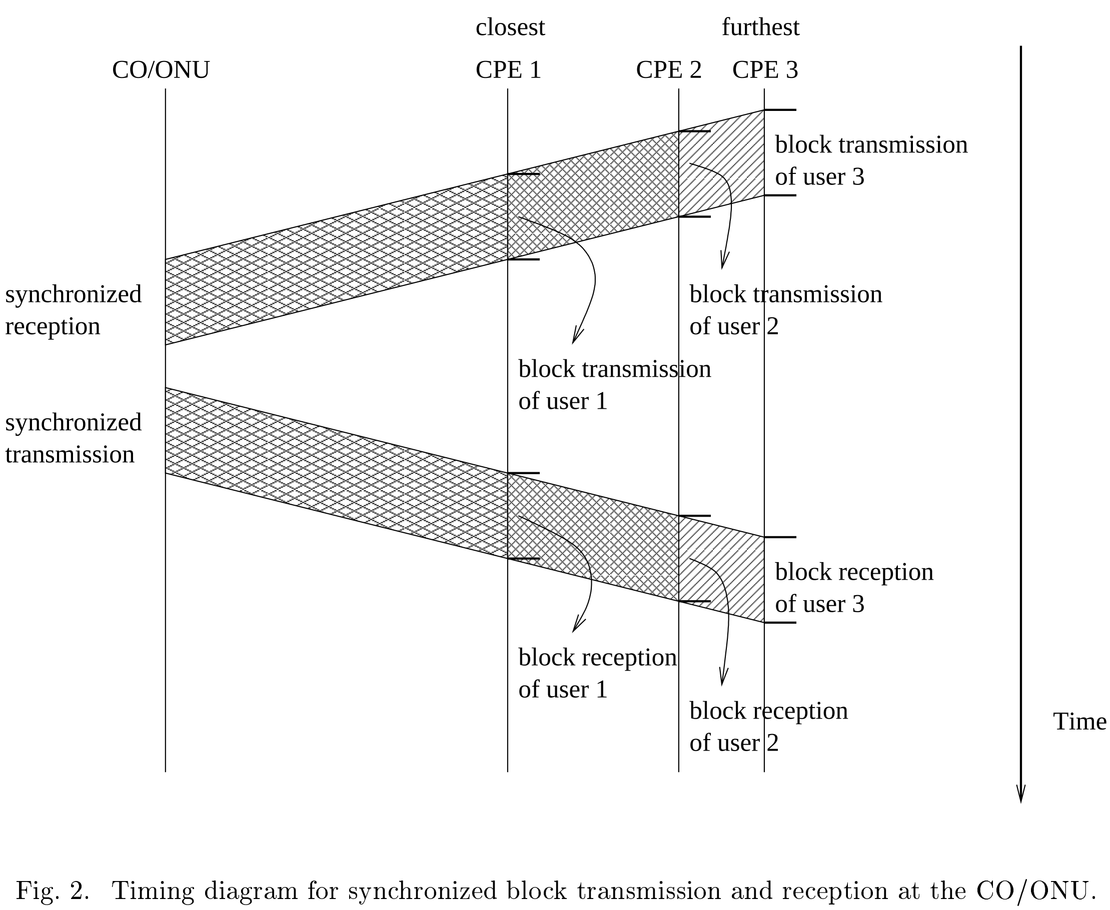

Fig. 2. Timing diagram for synchronized block transmission and reception at the CO/ONU.

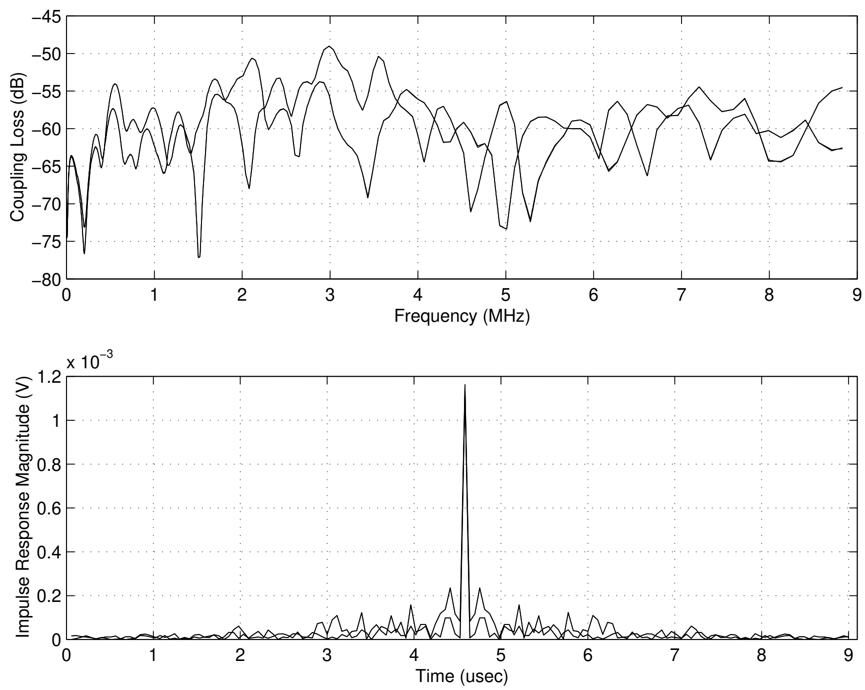

Fig. 3. FEXT frequency response measurements and corresponding impulse responses.

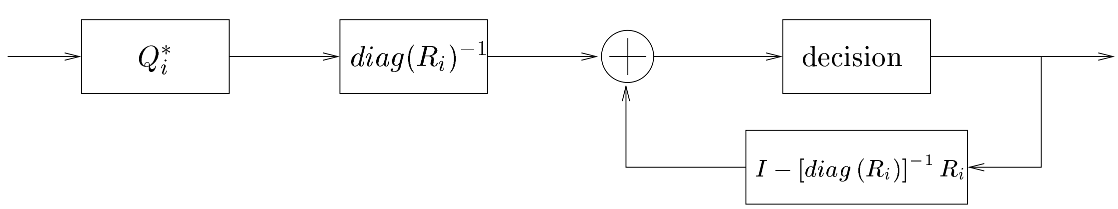

Fig. 4. Canceler block.

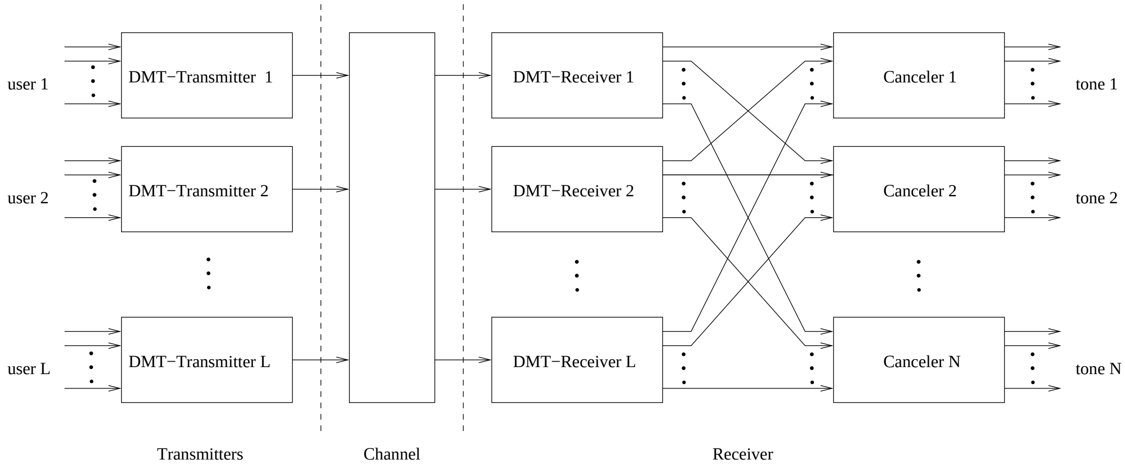

Fig. 5. Vectored-DMT for upstream communication.

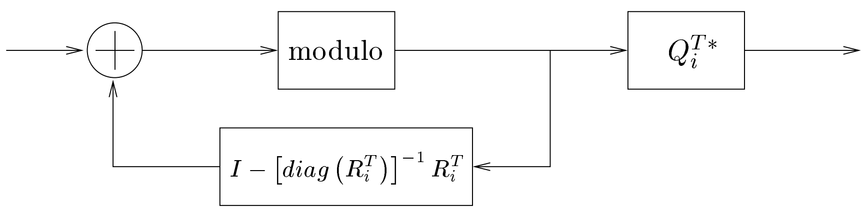

Fig. 6. Precoder block.

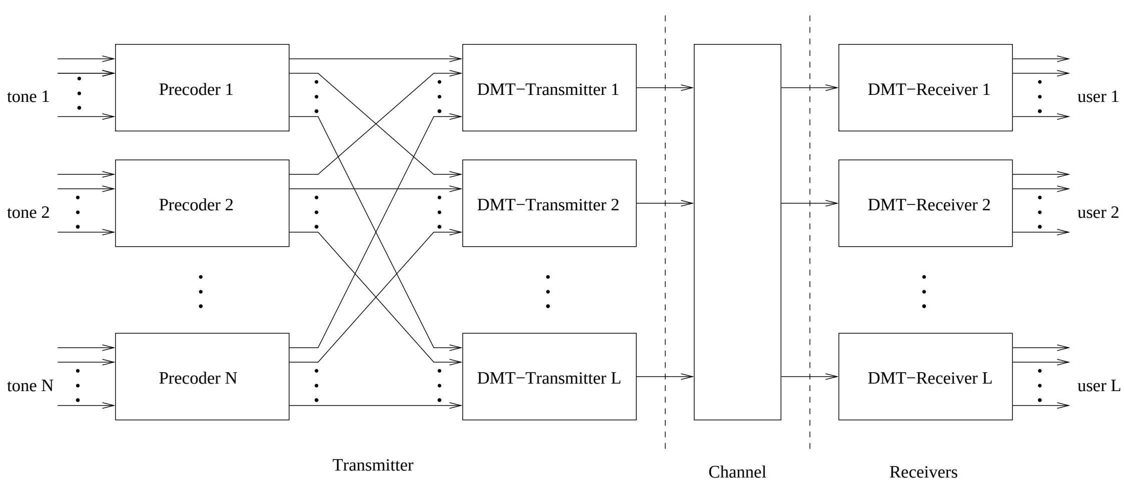

Fig. 7. Vectored-DMT system for downstream communication.

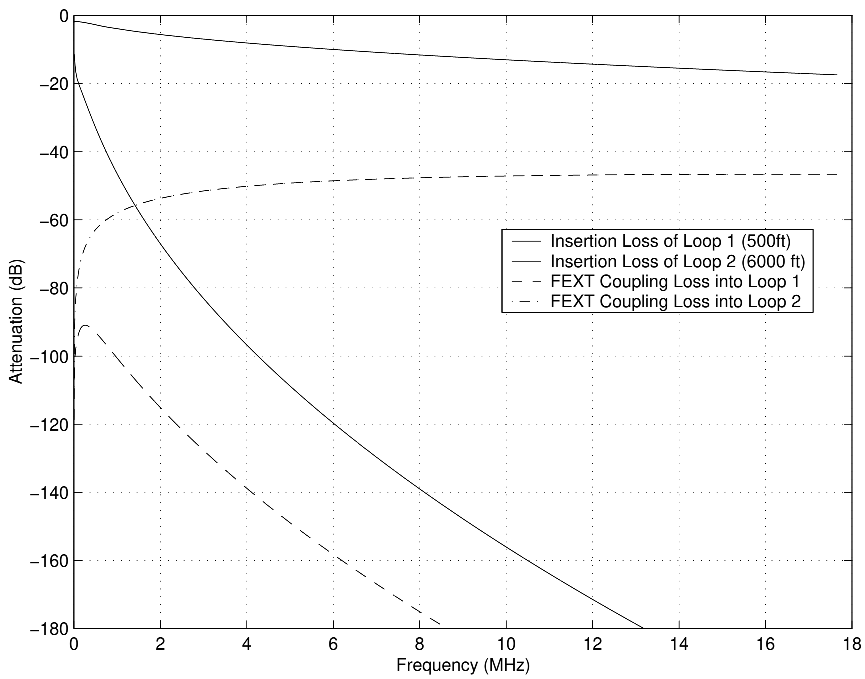

Fig. 8. Insertion loss and FEXT coupling loss for an extreme loop topology.

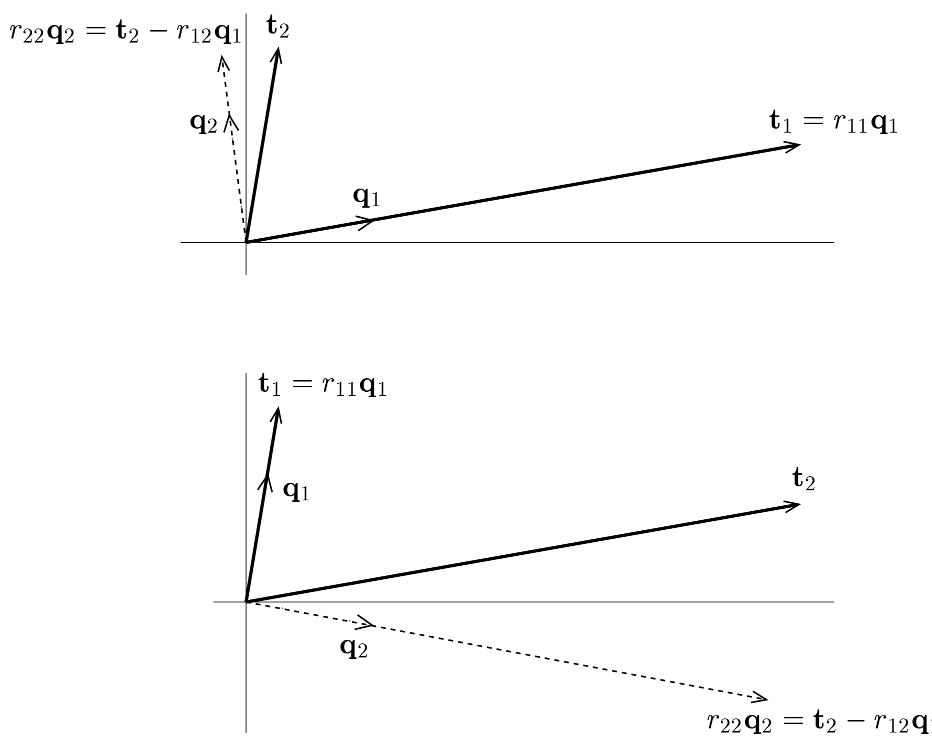

Fig. 9. Illustration of QR decomposition for the two possible orderings.

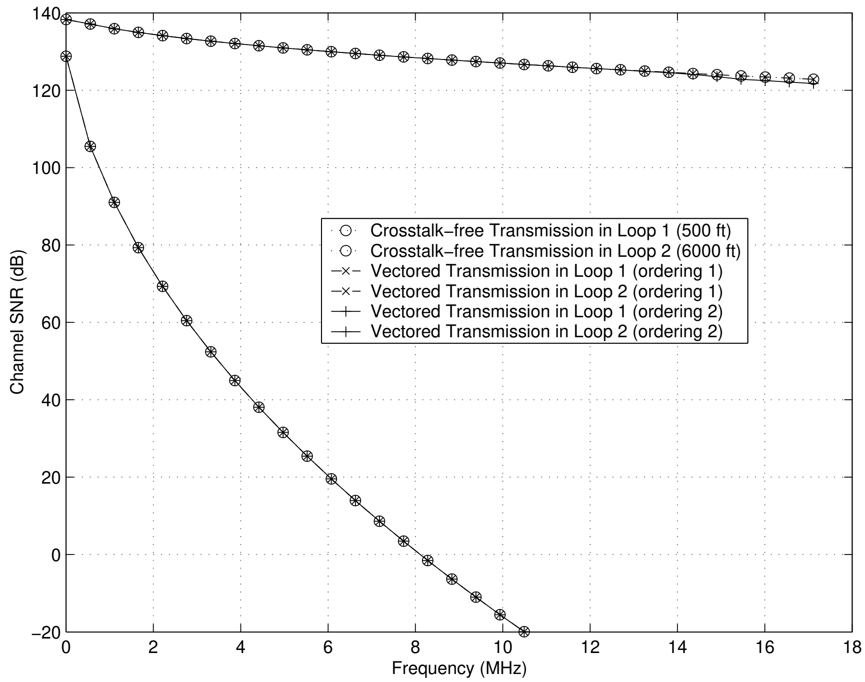

Fig. 10. Channel SNR's in crosstalk-free environment and in crosstalk environment with vectoring (shown for the two possible orderings).

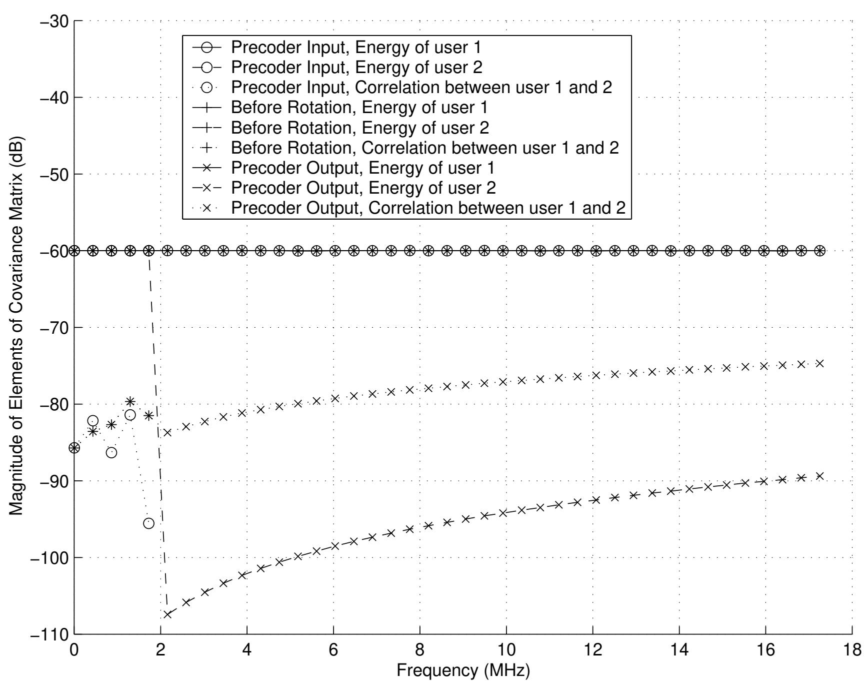

Fig. 11. Magnitude of elements of covariance matrices of precoder samples.

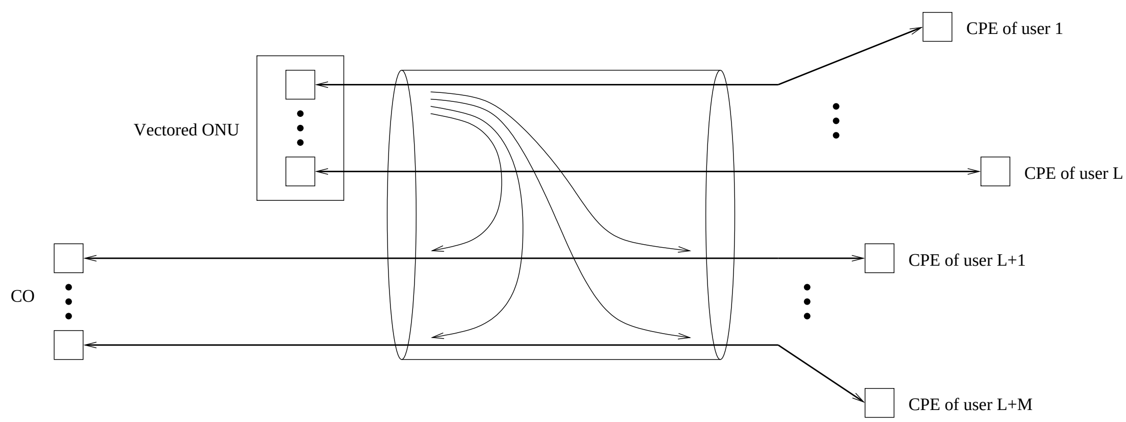

Fig. 12. DSL system topology with co-existence of CO and ONU.

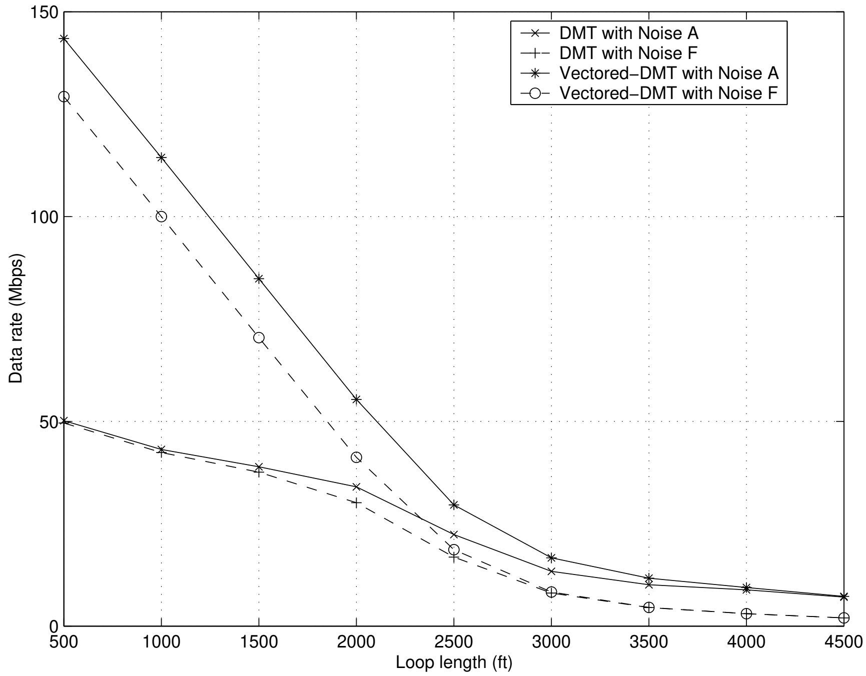

Fig. 13. Upstream VDSL data rates.

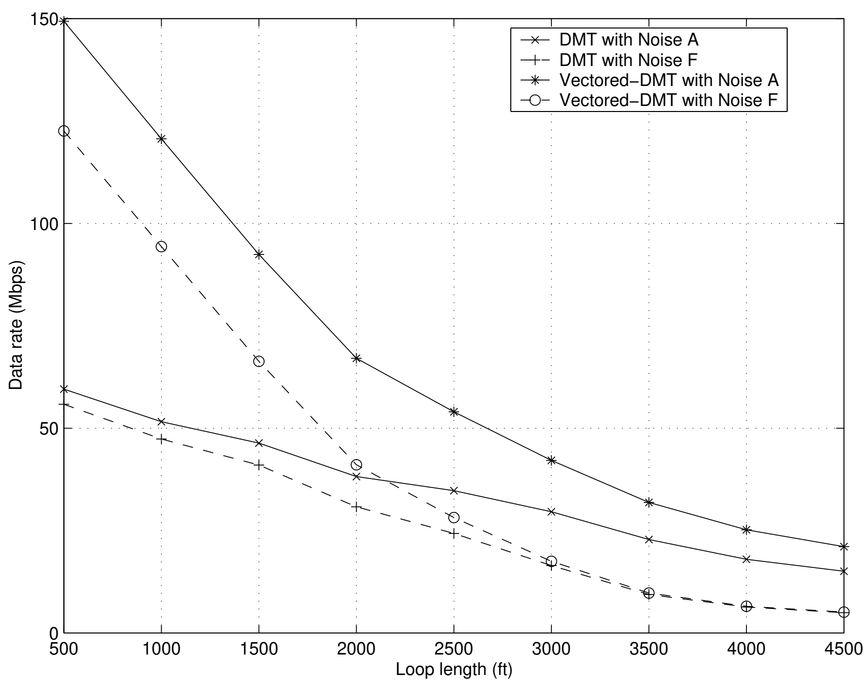

Fig. 14. Downstream VDSL data rates.

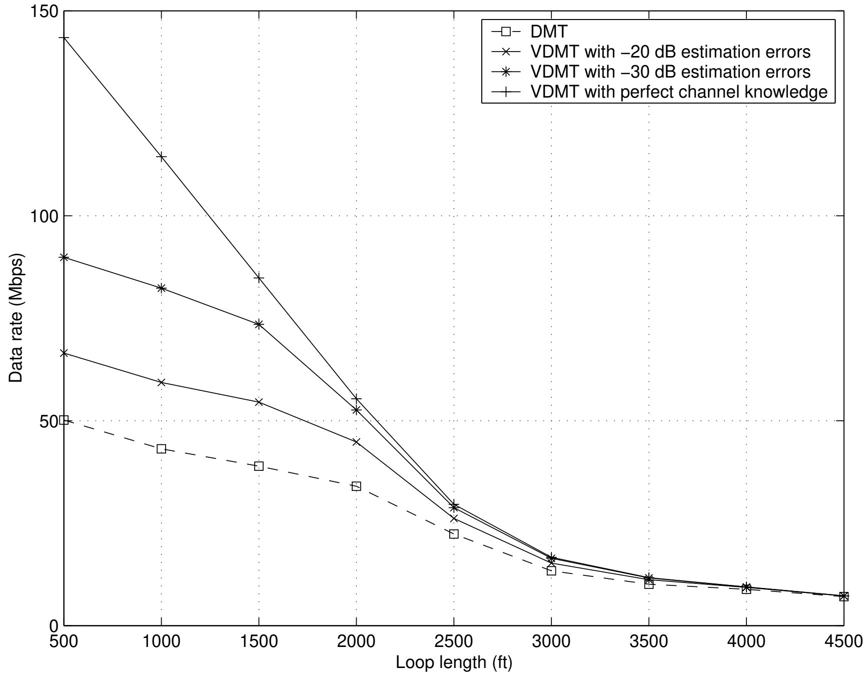

Fig. 15. Upstream VDSL data rates (noise A) with channel estimation errors.

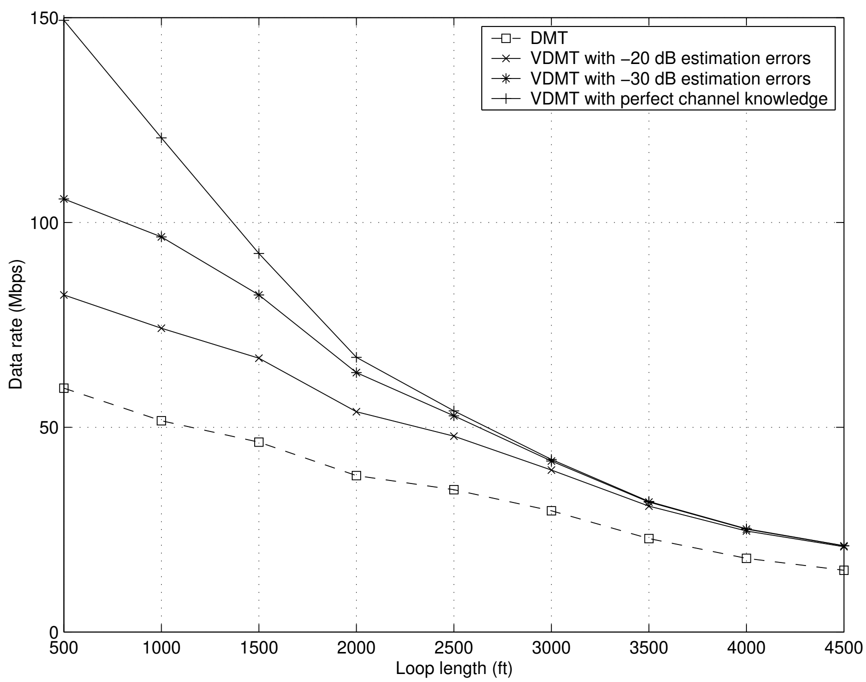

Fig. 16. Downstream VDSL data rates (noise A) with channel estimation errors.

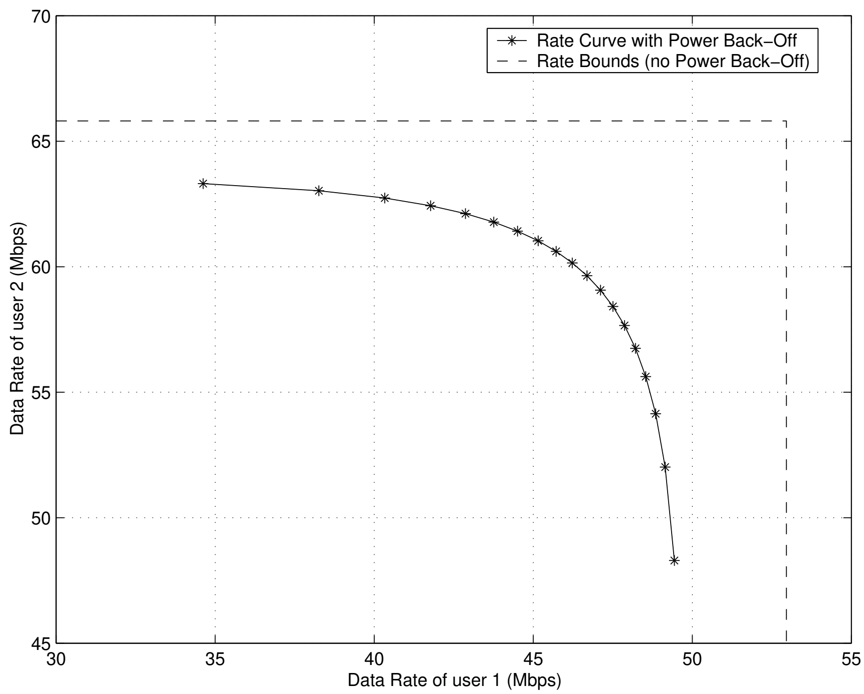

Fig. 17. Rate region for 2 users with power back-off.

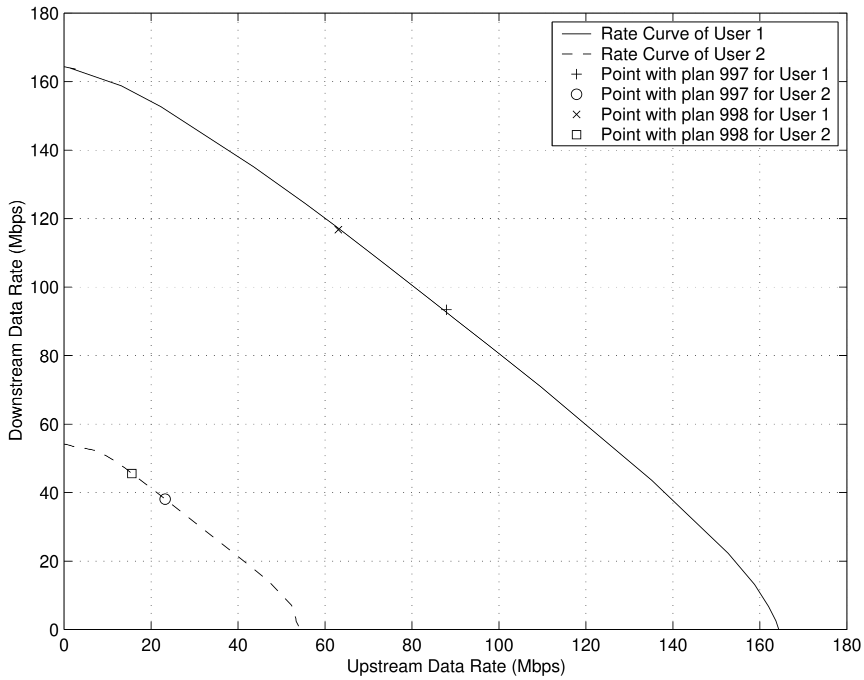

Fig. 18. Upstream/downstream rate regions for 2 users with optimum frequency duplexing.

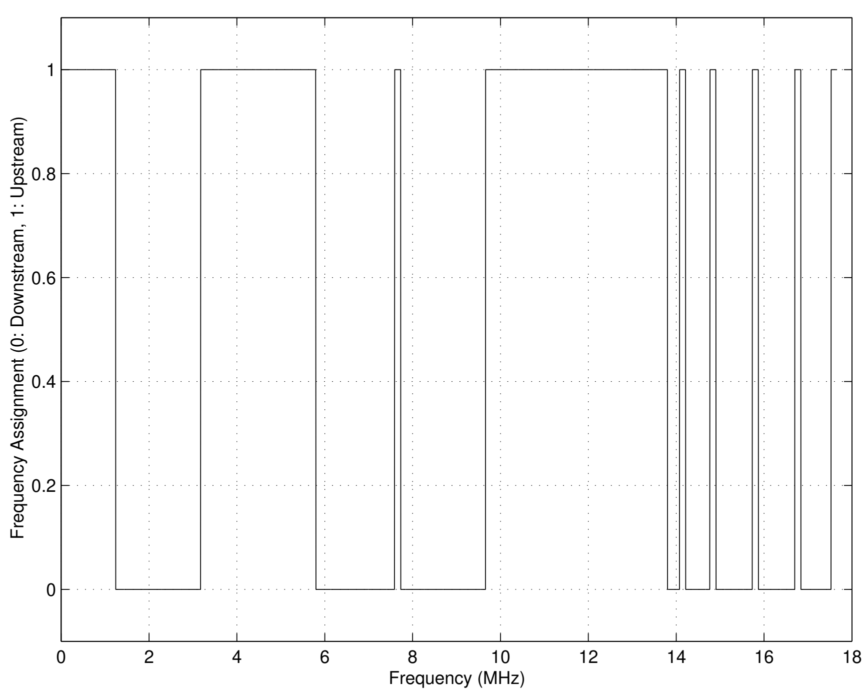

Fig. 19. Optimum frequency duplexing plan with symmetric service for upstream and downstream.