

Partial Crosstalk Cancellation Exploiting Line and Tone Selection in VDSL

R. Cendrillon, M. Moonen
SCD, Electrical Engineering Dept.
K.U. Leuven, Belgium
raphael.cendrillon@esat.kuleuven.ac.be

G. Ginis
Texas Instruments
San Jose, USA
gginis@dsl.stanford.edu

K. Van Acker, T. Bostoen, P. Vandaele
Alcatel Bell
Antwerp, Belgium
katleen.van_acker@alcatel.be

Abstract—VDSL is the next step in an on-going evolution of DSL systems. In VDSL downstream data rates of up to 52 Mbps are supported by operating over short loop lengths and using frequencies up to 12 MHz. Unfortunately such high frequencies result in crosstalk between pairs within a binder-group. Many crosstalk cancellation techniques have been proposed to address this. Whilst these schemes lead to impressive performance gains their complexity grows with the square of the number of lines within a binder. In binder-groups which can carry up to hundreds of lines this complexity is outside the scope of current implementation.

In this paper we investigate partial crosstalk cancellation for upstream VDSL. The majority of the detrimental effects of crosstalk are typically limited to a small sub-set of lines and tones. Furthermore significant crosstalk is often only seen from neighbouring pairs within the binder configuration. We present a number of algorithms which exploit these properties to reduce the complexity of crosstalk cancellation. These algorithms are shown to achieve the majority of the performance gains of full crosstalk cancellation with significantly reduced run-time complexity.

I. INTRODUCTION

VDSL is the next step in the on-going evolution of DSL systems. Supporting data rates up to 52 Mbps in the downstream, VDSL offers the potential of bringing truly broadband access to the consumer market. VDSL supports such high data rates by operating over short line lengths and transmitting in frequencies up to 12 MHz.

The twisted pairs in the access network are distributed within large *binder-groups* which typically contain anything from 20-100 individual pairs. As a result of the close distance between twisted-pairs within binders and the high frequencies used in VDSL transmission there is significant electromagnetic coupling between near-by pairs. This electromagnetic coupling leads to interference or *crosstalk* between the different systems operating within a binder. Far-end Crosstalk (FEXT) is typically 10-15 dB larger than the background noise and is the dominant source of performance degradation in VDSL.

Many crosstalk cancellation schemes have been proposed for VDSL see e.g. [1], [2], [3]. These schemes are applicable to upstream (US) transmission where the receiving modems are co-located.

Whilst the benefits of crosstalk cancellation are large, complexity can be extremely high. For example in a bundle with 20 users all transmitting on 4096 tones, and operating at a block rate of 4000 blocks/second the complexity of linear crosstalk

cancellation exceeds 6.2 billion multiplications per second. This is outside the scope of present-day implementation and may remain infeasible economically for several years. Other techniques such as soft-interference cancellation and non-linear crosstalk cancellation add even more complexity.

What is required is a crosstalk cancellation scheme with scalable complexity. It should support both conventional single-user detection (SUD) and full crosstalk cancellation. Furthermore it should exhibit graceful performance degradation as complexity is reduced. We present an upstream crosstalk cancellation scheme which exhibits these properties. It is shown that by exploiting the space and frequency-selective nature of crosstalk channels this crosstalk cancellation scheme can achieve the majority of the performance gains of full crosstalk cancellation with a fraction of the run-time complexity.

II. UPSTREAM SYSTEM MODEL

We begin by assuming that all receiving modems are co-located at the central office (CO) as is the case in upstream transmission. This is a prerequisite for crosstalk cancellation since signal level co-ordination is required between receivers. Through synchronized transmission and the cyclic structure of DMT blocks crosstalk can be modeled independently on each tone. We assume there are $N+1$ users within the binder-group so that each user has N interferers. Transmission of a single DMT block can be modeled as

$$\mathbf{y}_k = \mathbf{H}_k \mathbf{x}_k + \mathbf{z}_k$$

Here $\mathbf{x}_k \triangleq [x_k^1, \dots, x_k^{N+1}]$ and is the vector of transmitted signals on tone k . x_k^n is the signal transmitted onto line n at tone k . The tone k is in the range $1, \dots, K$. $\mathbf{y}_k$ and $\mathbf{z}_k$ have similar structures. $\mathbf{y}_k$ is the vector of received signals on tone k . $\mathbf{z}_k$ is the vector of additive noise on tone k and contains thermal noise, alien crosstalk, RFI etc. We assume $\mathcal{E}\{\mathbf{z}_k \mathbf{z}_k^H\} = \sigma_k^2 \mathbf{I}_{N+1}$. This is without loss of generality since a noise-whitening pre-filter can be applied at the RXs prior to crosstalk cancellation.

$\mathbf{H}_k$ is the $(N+1) \times (N+1)$ channel transfer matrix on tone k . $h_k^{n,m} \triangleq [\mathbf{H}_k]_{n,m}$ is the channel from TX m to RX n on tone k . The diagonal elements of $\mathbf{H}_k$ contain the direct-channels whilst the off-diagonal elements contain the crosstalk channels. We denote the transmit auto-correlation on tone k as $\mathcal{E}\{\mathbf{x}_k \mathbf{x}_k^H\} = \mathbf{S}_k$ with $s_k^m \triangleq [\mathbf{S}_k]_{m,m}$. Note that $\mathbf{S}_k$ is a diagonal matrix since co-ordination is not available between the different customer premises (CP) transmitters.

In DSL channels with co-located receivers the channel matrix $\mathbf{H}_k$ is column-wise diagonal dominant (CWDD) and

This work was carried out in the frame of IUAP P5/22, *Dynamical Systems and Control: Computation, Identification and Modelling* and P5/11, *Mobile multimedia communication systems and networks*; the Concerted Research Action GOA-MEFISTO-666, *Mathematical Engineering for Information and Communication Systems Technology*; FWO Project G.0196.02, *Design of efficient communication techniques for wireless time-dispersive multi-user MIMO systems* and was partially sponsored by Alcatel-Bell.

satisfies the following property

$$\left| h_k^{(m,m)} \right| \gg \left| h_k^{(n,m)} \right|, \forall n \neq m \quad (1)$$

In other words the direct channel of any user always has a larger gain than the channel from that user's transmitter into any other user's receiver. This property has been verified through extensive cable measurements (see the semi-empirical crosstalk channel models in [4]). It will be exploited in the remaining sections.

III. CROSSTALK CANCELLATION

A. Optimal Crosstalk Cancellation

When both the transmitters and receivers of the modems within a binder are co-located, channel capacity can be achieved in a simple fashion[1]. Using the singular value decomposition (SVD) define

$$\mathbf{H}_k \stackrel{\text{svd}}{=} \mathbf{U}_k \mathbf{\Lambda}_k \mathbf{V}_k^H \quad (2)$$

where the columns of $\mathbf{U}_k$ and $\mathbf{V}_k$ are the left and right singular vectors of $\mathbf{H}_k$ respectively and the singular values $\mathbf{\Lambda}_k \triangleq \text{diag}\{\lambda_k^1, \dots, \lambda_k^{N+1}\}$. It is assumed that $\mathbf{H}_k$ is non-singular which is ensured by (1) provided that $h_k^{(n,n)} \neq 0, \forall n$.

Define the true set of symbols $\tilde{\mathbf{x}}_k \triangleq [\tilde{x}_k^1 \dots \tilde{x}_k^{N+1}]^T$ which are generated by the QAM encoders. Define $\mathcal{E}\{\tilde{\mathbf{x}}_k \tilde{\mathbf{x}}_k^H\} \triangleq \tilde{\mathbf{S}}_k = \text{diag}\{\tilde{s}_k^1, \dots, \tilde{s}_k^{N+1}\}$. For a given $\tilde{\mathbf{S}}_k$ the optimal transmitter structure pre-filters $\tilde{\mathbf{x}}_k$ with the matrix

$$\mathbf{P}_k = \mathbf{V}_k$$

such that $\mathbf{x}_k = \mathbf{P}_k \tilde{\mathbf{x}}_k$. At the receiver we apply the filter $\mathbf{w}_k^n = \mathbf{e}_n^H \mathbf{\Lambda}_k^{-1} \mathbf{U}_k^H$ to generate our estimate of the transmitted symbol

$$\begin{aligned} \hat{x}_k^n &= \mathbf{w}_k^n \mathbf{y}_k \\ &= \mathbf{w}_k^n (\mathbf{H}_k \mathbf{P}_k \tilde{\mathbf{x}}_k + \mathbf{z}_k) \\ &= \tilde{x}_k^n + \tilde{z}_k^n \end{aligned}$$

where $\mathbf{e}_n \triangleq [\mathbf{I}_{N+1}]_{\text{col } n}$, $\mathbf{I}_{N+1}$ is the $(N+1) \times (N+1)$ identity matrix, and $\tilde{z}_k^n \triangleq \mathbf{e}_n^H \mathbf{\Lambda}_k^{-1} \mathbf{U}_k^H \mathbf{z}_k$. Here we use $[\mathbf{A}]_{\text{row } n}$ and $[\mathbf{A}]_{\text{col } n}$ to denote the n th row and column of matrix $\mathbf{A}$ respectively. Note that $\mathcal{E}\{|z_k^n|^2\} = \sigma_k^2 (\lambda_k^n)^{-2}$. The pre and post-filtering operations remove crosstalk without causing noise enhancement. Applying a conventional slicer to $\hat{x}_k^n$ allows the following rate to be achieved for user n on tone k

$$c_k^n = \log \left(1 + \frac{1}{\Gamma} \sigma_k^{-2} (\lambda_k^n)^2 \tilde{s}_k^n \right)$$

Γ represents the SNR-gap to capacity and is a function of the target BER, coding gain and noise margin[5]. The maximum achievable rate of the multi-line DSL channel is

$$C = \sum_k \log \left| \mathbf{I}_{N+1} + \frac{1}{\Gamma} \sigma_k^{-2} \mathbf{H}_k \mathbf{S}_k \mathbf{H}_k^H \right|$$

It is straight-forward to show $\sum_n \sum_k c_k^n = C$. So through the application of a simple linear pre and post-filter, and a conventional slicer it is possible to operate at the maximum achievable rate of the DSL channel for the given $\tilde{\mathbf{S}}_k$. Unfortunately application of a pre-filter requires the transmitting modems to be co-located. In upstream DSL this is typically not the case since transmitting modems are located at different CPs.

B. Simplified, Near-Optimal Crosstalk Cancellation

As a result of the CWDD of $\mathbf{H}_k$ rates close to the maximum can be achieved with a very simple receiver structure. Furthermore pre-filtering is not required so such rates can be achieved without co-located transmitting modems. We now show why this is true.

The CWDD of $\mathbf{H}_k$ ensures that its columns are approximately orthogonal. Hence the right singular vectors of $\mathbf{H}_k$ can be approximated

$$\begin{aligned} \mathbf{P}_k &= \mathbf{V}_k \\ &\simeq \mathbf{I}_{N+1} \end{aligned}$$

So pre-filtering is not required. This is important since in upstream DSL transmitting modems are not co-located. At the receivers we apply the filter

$$\begin{aligned} \mathbf{w}_k^n &= \mathbf{e}_n^H \mathbf{\Lambda}_k^{-1} \mathbf{U}_k^H \\ &\simeq \mathbf{e}_n^H \mathbf{V}_k \mathbf{\Lambda}_k^{-1} \mathbf{U}_k^H \\ &\simeq \mathbf{e}_n^H \mathbf{H}_k^{-1} \end{aligned}$$

Hence the optimal receiver structure is well approximated by a *linear zero-forcing (ZF)* design. Thus we can achieve close to maximum-rate using the following estimate

$$\hat{x}_k^n = \mathbf{e}_n^H \mathbf{H}_k^{-1} \mathbf{y}_k$$

Note that noise enhancement is not a problem since $\mathbf{H}_k^{-1} \simeq \mathbf{\Lambda}_k^{-1} \mathbf{U}_k^H$. $\mathbf{U}_k^H$ is unitary hence it does not alter the statistics of the noise. $\mathbf{\Lambda}_k^{-1}$ is diagonal hence it scales the signal and noise equally. In [6] an upper bound is proposed for the capacity loss incurred due to the above approximation. Using worst case models, this is shown to be less than 1% for 99% of DSL channels.

Using this scheme crosstalk cancellation of one user at one tone requires N mults./DMT block. So crosstalk cancellation for $N+1$ users on K tones at a block-rate b (DMT blocks/second) requires $(N^2 + N)Kb$ mults. per second. This complexity rapidly grows with the number of users in a bundle. For example, in a 20 user system with 4096 tones and a block rate of 4000 the complexity is 6.2 billion mults. per second. So whilst crosstalk cancellation leads to significant performance gains it can be extremely complex, certainly beyond the complexity available in current-day systems. This is the motivation behind partial crosstalk cancellation.

IV. CROSSTALK SELECTIVITY

In Fig. 1 some crosstalk transfer functions are plotted from a set of measurements of a British Telecom (BT) cable consisting of $8 \times 0.5\text{mm}$ pairs. Examining this plot we can make two observations:

First, from a particular user's perspective, some crosstalkers cause significant amounts of interference, whilst others cause little interference at all. We refer to this as the *space-selectivity* of crosstalk since the crosstalk channels vary significantly between lines. Space-selectivity arises naturally due to the physical layout of binders. Since electromagnetic coupling decreases rapidly with distance, each pair will experience significant crosstalk from only a few other surrounding pairs within the binder. The near-far effect also gives rise to space-selectivity. In upstream transmission, modems which are located closer to the CO will cause more crosstalk than

those located further away. Using the channel measurements provided by BT it was seen that 90% of crosstalk energy is caused by the 4 largest crosstalkers.

Second, crosstalk channels vary significantly with frequency. We refer to this as the *frequency-selectivity* of crosstalk which arises naturally from the frequency dependent nature of electromagnetic coupling. Using the BT channel measurements it was seen that 90% of the crosstalk is contained within less than half of the tones.

So the effects of crosstalk vary considerably with both space and frequency. Furthermore, the majority of its effects are contained within a relatively small subset of tones and crosstalkers. Some tones will see more significant crosstalkers than others and we can scale between conventional single-user detection (SUD) and full crosstalk cancellation on a tone-by-tone basis. On each tone we choose the degree of crosstalk cancellation based on the severity of crosstalk experienced. By only canceling the largest crosstalkers and by varying the degree of crosstalk cancellation on each tone, partial crosstalk cancellation can approach the performance of full crosstalk cancellation with a fraction of the run-time complexity.

V. PARTIAL CROSSTALK CANCELLATION

A. Partial Crosstalk Canceller Structure

We now describe partial crosstalk cancellation design in more detail. In the detection of user n we observe the direct line of user n (to recover the signal) and $p_{k,n}$ additional lines (to enable crosstalk cancellation). $p_{k,n}$ varies both with the tone k and user n to match the severity of crosstalk seen by that user on that tone. Note that $p_{k,n} = N$ corresponds to full crosstalk cancellation whilst $p_{k,n} = 0$ corresponds to none (ie. SUD). Define the set of extra observation lines

$$\mathbb{M}_k^n \triangleq \{m_{k,n}(1), \dots, m_{k,n}(p_{k,n})\}$$

and the corresponding received signals

$$\bar{\mathbf{y}}_k^n \triangleq \left[y_k^n, y_k^{m_{k,n}(1)} \dots y_k^{m_{k,n}(p_{k,n})} \right]^T$$

We form an estimate of the transmitted symbol using a linear combination of the received signals on the *observation lines* only

$$\hat{x}_k^n = \bar{\mathbf{w}}_k^n \bar{\mathbf{y}}_k^n$$

Note that crosstalk cancellation for user n at tone k now requires only $p_{k,n}$ mults./DMT block in contrast to the N mults. required for full crosstalk cancellation. This technique has many similarities to hybrid selection/combining from the wireless field[7]. There selection is also used between receive antennas to reduce run-time complexity and reduce the number of analog front-ends (AFE) required.

B. Partial Crosstalk Canceller Design

We now describe the design of the partial crosstalk canceller. We begin with a reduced channel matrix which contains only the paths between user n and users in the set $\mathbb{M}_k^n$

$$\bar{\mathbf{H}}_k^n \triangleq \begin{bmatrix} h_k^{(n,n)} & [\mathbf{H}_k^n]_{\text{row } n, \text{cols } \mathbb{M}_k^n} \\ [\mathbf{H}_k^n]_{\text{rows } \mathbb{M}_k^n, \text{col } n} & [\mathbf{H}_k^n]_{\text{rows } \mathbb{M}_k^n, \text{cols } \mathbb{M}_k^n} \end{bmatrix}$$

The linear ZF partial canceller is designed

$$\bar{\mathbf{w}}_k^n \triangleq \bar{\mathbf{e}}_1^H \left(\bar{\mathbf{H}}_k^n \right)^{-1}$$

where $\bar{\mathbf{e}}_1^H$ is the first column of $\mathbf{I}_{p_{k,n}+1}$ and $[\mathbf{A}]_{\text{rows } \mathbb{A}, \text{cols } \mathbb{B}}$ denotes the sub-matrix formed from the rows $\mathbb{A}$ and columns $\mathbb{B}$ of matrix $\mathbf{A}$.

VI. LINE SELECTION

In DSL the majority of the crosstalk that a particular user experiences comes from only a few of the other users within the system. We have referred to this effect as the *space-selectivity* of the crosstalk channel and we exploit it to reduce the complexity of crosstalk cancellation. In practice this corresponds to observing only the subset $\mathbb{M}_k^n$ of the lines at the CO when detecting user n .

In this section we investigate the optimal choice for the subset $\mathbb{M}_k^n$. Our problem is thus

$$\max_{\mathbb{M}_k^n} c_k^n \text{ s.t. } |\mathbb{M}_k^n| \leq p_{k,n} \quad (3)$$

where $|\mathbb{A}|$ denotes the cardinality of set $\mathbb{A}$ and c_k^n is the rate of user n on tone k .

Define the SVD of the reduced channel matrix

$$\bar{\mathbf{H}}_k \stackrel{\text{svd}}{=} \bar{\mathbf{U}}_k \bar{\mathbf{\Lambda}}_k \bar{\mathbf{V}}_k^H$$

CWDD in $\bar{\mathbf{H}}_k$ implies the same in $\bar{\mathbf{H}}_k^n$. Hence we can approximate $\bar{\mathbf{V}}_k \simeq \mathbf{I}_{p_{k,n}+1}$. Under this approximation

$$\begin{aligned} \bar{\mathbf{w}}_k^n &= \bar{\mathbf{e}}_1^H \bar{\mathbf{H}}_k^{-1} \\ &\simeq \mathbf{e}_n^H \bar{\mathbf{\Lambda}}_k^{-1} \bar{\mathbf{U}}_k^H \end{aligned}$$

Since $\bar{\mathbf{U}}_k$ is unitary, and $\bar{\mathbf{\Lambda}}_k^{-1}$ is diagonal, application of $\bar{\mathbf{w}}_k^n$ does not cause noise enhancement. It is straight-forward to show that, after application of the partial crosstalk canceller the Signal to Interference plus Noise Ratio (SINR) is closely approximated by

$$\text{SINR}_k^n \simeq \frac{|h_k^{(n,n)}|^2 s_k^n}{\sum_{m \notin \{n, \mathbb{M}_k^n\}} |h_k^{(n,m)}|^2 s_k^m + \sigma_k^2} \quad (4)$$

with the approximation becoming exact in strongly CWDD channels. See [8] for a more thorough proof.

Maximizing SINR_k^n and thus rate c_k^n corresponds to maximizing the amount of interference captured in the set $\mathbb{M}_k^n$. Note that we assume a sufficient number of noise sources and crosstalkers such that the background noise and residual interference are approximately Gaussian. So to maximize rate c_k^n we simply choose $\mathbb{M}_k^n$ to contain the largest crosstalkers of user n on tone k . Define the indices of the crosstalkers of user n on tone k sorted in order of crosstalk strength

$$\{q_{k,n}(1), \dots, q_{k,n}(N)\}$$

such that $q_{k,n}(i) \neq n, \forall i$ and

$$|h_k^{(n,q_{k,n}(i))}|^2 s_k^{q_{k,n}(i)} \geq |h_k^{(n,q_{k,n}(i+1))}|^2 s_k^{q_{k,n}(i+1)}, \forall i$$

At this point we can propose a simple approach to partial crosstalk cancellation: Alg. 1. Assume we operate under a complexity limit of cK mults./DMT-block/user

$$\sum_k |\mathbb{M}_k^n| \leq cK, \forall n$$

Algorithm 1 Line Selection Only

$$\mathbb{M}_k^n = \{q_{k,n}(1), \dots, q_{k,n}(c)\}, \forall n, k$$

This corresponds to c times the complexity of a conventional frequency domain equalizer (FEQ) as is currently implemented in VDSL modems. In this algorithm we simply cancel the c largest crosstalkers on each tone, hence

$$p_{k,n} = c, \forall n, k$$

The reduction in run-time complexity from this algorithm comes from space-selectivity only. Since the degree of partial cancellation stays constant across all tones this algorithm cannot exploit the frequency-selectivity of the crosstalk channel. As we will see, this leads to sub-optimal performance when compared to algorithms which exploit both space and frequency-selectivity. The advantage of this algorithm is its simplicity. The algorithm requires only $O(KN)$ mults. and K sorting operations of N values to initialize the partial crosstalk canceller for one user. Here we define initialization complexity as the complexity of determining $\mathbb{M}_k^n, \forall k$. Initialization complexity does not include actual calculation of the crosstalk cancellation parameters $\bar{\mathbf{w}}_k^n$ for each tone. This requires $O(\sum_k (p_{k,n} + 1)^3)$ mults. for user n regardless of the partial cancellation algorithm employed. The initialization complexity (in terms of mults./user) of the different partial cancellation algorithms is listed in Tab. I. All algorithms have equal run-time complexity.

VII. TONE SELECTION

In the previous section we presented Alg. 1 for partial crosstalk cancellation. This algorithm exploits the space-selectivity of the crosstalk channel, ie. the fact that crosstalk varies significantly between different lines. Crosstalk coupling also varies significantly with frequency and this can also be exploited to reduce run-time complexity.

In low frequencies crosstalk coupling is minimal so we would expect minimal gains from crosstalk cancellation. In higher frequencies on the other hand crosstalk coupling can be severe. However in high frequencies the direct channel attenuation is so high that the channel can only support minimal bitloading even in the absence of crosstalk. This limits the potential gains of crosstalk cancellation. The largest gains from crosstalk cancellation will be experienced in intermediate frequencies and this is where most of the run-time complexity should be allocated. Define the rate achieved by user n on tone k when the $p_{k,n}$ largest crosstalkers are canceled

$$r_{k,n}(p_{k,n}) \triangleq \log \left(1 + \frac{1}{\Gamma} \frac{|h_k^{(n,n)}|^2 s_k^n}{\Psi_{k,n}(p_{k,n}) + \sigma_k^2} \right) \quad (5)$$

where

$$\Psi_{k,n}(p_{k,n}) \triangleq \sum_{i=p_{k,n}+1}^N |h_k^{(n,q_{k,n}(i))}|^2 s_k^{q_{k,n}(i)}$$

Define the gain of full crosstalk cancellation ($p_{k,n} = N$)

$$g_{k,n} \triangleq r_{k,n}(N) - r_{k,n}(0)$$

Algorithm 2 Tone Selection Only

$$\mathbb{M}_k^n = \begin{cases} \{1, \dots, n-1, n+1, \dots, N+1\}, \\ \quad k \in \{k_n(1), \dots, k_n(cK/N)\} \\ \emptyset, \text{ otherwise} \end{cases}$$

and the indices of the tones ordered by this gain

$$\{k_n(1), \dots, k_n(K)\} \quad \text{s.t. } g_{k_n(i),n} \geq g_{k_n(i+1),n}, \forall i$$

Note that by operating on a logarithmic scale $g_{k,n}$ can be calculated by dividing the arguments of the logarithms in $r_{k,n}(N)$ and $r_{k,n}(0)$.

We can now define another partial crosstalk cancellation algorithm: Alg. 2. This algorithm simply employs full crosstalk cancellation on the cK/N tones with the largest gain and no cancellation on all other tones. This leads to a run-time complexity of cK mults./DMT-block/user.

Note that in this algorithm $p_{k,n}$ is restricted to take only the values 0 or N . As a result it is not possible to only cancel the largest crosstalkers and this algorithm cannot exploit space-selectivity. The initialization complexity of this algorithm is $O(KN)$ mults. and one sort of size K , per user.

VIII. JOINT TONE-LINE SELECTION

In Sec. VI and VII we described partial cancellation algorithms which exploit only one form of selectivity in the crosstalk channel. To achieve maximum reduction in run-time complexity it is necessary to exploit both space and frequency-selectivity. We should adapt the degree of crosstalk cancellation done on each tone $p_{k,n}$ to match the potential gains. In practice this means that we allow $p_{k,n}$ to take on values other than 0 and N whilst also allowing $p_{k,n}$ to vary from tone to tone.

As we saw in Sec. VI observing the direct line of a crosstalker allows us to remove the crosstalk it causes to the user being detected. Hence line selection is equivalent to choosing which subset of crosstalkers we desire to cancel. When combined with tone selection our problem is effectively to choose which (crosstalker, tone) pairs to cancel in the detection of a certain user.

The rate improvement from canceling a particular crosstalker on a particular tone is dependent on the other crosstalkers that will be canceled on that tone. As such there is an inherent coupling in crosstalker selection which greatly complicates matters. In this algorithm we remove this coupling by ignoring the effect of other crosstalkers in the system. This greatly simplifies (crosstalker, tone) pair selection with only a small performance penalty. A discussion of the near-optimality of this algorithm and performance comparisons with the truly optimal partial cancellation scheme (which is based on a greedy algorithm) is given in [8].

Define the gain of canceling crosstalker m on tone k in the detection of user n , and in the absence of all other crosstalkers

Algorithm 3 Simple Line-Tone Selection

$$\mathbb{M}_k^n = \{m : (m, k) \in \{d_n(1), \dots, d_n(cK)\}\}$$

$$\bar{g}_{k,n}(m) \triangleq \log \left(1 + \frac{|h_k^{(n,n)}|^2 s_k^n}{\Gamma \sigma_k^2} \right) - \log \left(1 + \frac{1}{\Gamma} \frac{|h_k^{(n,n)}|^2 s_k^n}{|h_k^{(n,m)}|^2 s_k^m + \sigma_k^2} \right)$$

Note that if we work in a logarithmic scale then $\bar{g}_{k,n}(m)$ can be calculated by simply dividing the arguments of each log function. Define (crosstalker, tone) pair $d_n(i) \triangleq (m_n(i), k_n(i))$ and its corresponding gain $\bar{g}_n(d_n(i)) \triangleq \bar{g}_{k_n(i),n}(m_n(i))$. This allows us to define the indices of (crosstalker, tone) pairs ordered by gain

$$\{d_n(1), \dots, d_n(KN)\} \text{ s.t. } \bar{g}_n(d_n(i)) \geq \bar{g}_n(d_n(i+1)), \forall i$$

We can now define our joint tone-line selection algorithm: Alg. 3. In the detection of user n we observe the direct line of crosstalker m on tone k if the pair

$$(m, k) \in \{d_n(1), \dots, d_n(cK)\}$$

This leads to a run-time complexity of cK mults./DMT-block/user. The benefit of this algorithm is its low complexity. Pair selection for one user has a complexity of $\mathcal{O}(KN)$ mults. and one sort of size KN . Furthermore, this algorithm exploits both the space and frequency-selectivity of the crosstalk channel, allowing it to cancel the largest crosstalkers on the tones where they do the most harm.

IX. COMPLEXITY DISTRIBUTION BETWEEN USERS

So far we have limited the run-time complexity of detecting each user to cK such that

$$\sum_k |\mathbb{M}_k^n| \leq cK, \forall n$$

If crosstalk cancellation of all lines in a binder is integrated into a single processing module at the CO, then mults. can be shared between users. That is, the true constraint is on the total complexity of crosstalk cancellation for *all users*

$$\sum_n \sum_k |\mathbb{M}_k^n| \leq cK (N+1)$$

This would typically be the case when DSL services are deployed from an optical network unit (ONU). Then all DSL lines will be served from a single DSLAM which may centralize crosstalk cancellation into a single module. The available complexity can then be divided between users based on our desired rates for each. Denote the number of mults./DMT-block allocated to user n as κ_n , then

$$\kappa_n = \mu_n cK (N+1) \quad \text{s.t.} \quad \sum_n \mu_n = 1$$

Here μ_n is a parameter which determines the proportion of computing resources allocated to user n . This allows us to

view partial cancellation as a resource allocation problem not just across tones, but users as well. Given a fixed number of multiplications we must divide them between users based on the desired rate of each user. In a similar fashion to work done in multi-user power allocation (see e.g. [9]) we can define a rate region as the set of all achievable rate-tuples under a given total complexity constraint. This allows us to visualize the different trade-offs that can be achieved between the rates of different users within a binder.

X. PERFORMANCE

We now compare the performance of the partial cancellation algorithms previously described.

We use semi-empirical transfer functions from the ETSI VDSL standards for 0.5mm (24-Gauge) lines[4]. Note that in these channel models each user sees identical crosstalk channels to all crosstalkers of equal line length. That is, the variation of crosstalk channel attenuation with the distance between lines within the binder is not modeled. When a binder consists of lines of varying length the model does capture the near-far effect. All users will see the modems located closest to the CO (near-end lines) as the largest sources of crosstalk. On the other hand when a binder consists of lines of equal length all users will see equal crosstalk from all other users. In that case there will be no space-selectivity in the crosstalk channel model. In reality we would expect more space-selectivity than is contained within these channel models. Hence we can expect the reduction in run-time complexity to be even larger than that shown here.

In our simulations the binder consisted of 8 lines, so $N = 7$. There are 4 near-end users located 300 m. from the CO and 4 far-end users located 1200 m. from the CO. We used $K = 4096$ tones, a coding gain of 3 dB, noise margin 6 dB, target error probability 10^{-7} , 998 Bandplan with flat transmit PSD at -60 dBm/Hz, and noise model A from the ETSI standards[4].

We distribute run-time complexity between users as described in Sec. IX. Fig. 2 contains the achievable rate regions under varying complexities c using Alg. 3. The rate region was constructed by dividing multiplications between the two classes of near-end and far-end users. Users of one class receive an equal number of multiplications; $2\mu_{\text{near}}cK$ and $2\mu_{\text{far}}cK$ mults./DMT-block for the near-end and far-end users respectively. By varying the parameter μ_{far} we can trace out the boundary of the rate region. Note that $\mu_{\text{near}} = 1 - \mu_{\text{far}}$. We see in Fig. 2 that with $c = 2$ (29% of the run-time complexity of full crosstalk cancellation) we can achieve the majority of the operating points within the rate region.

In Fig. 3 the achievable rate regions of the different partial cancellation algorithms are compared for $c = 2$. Note the considerably larger rate region which is achieved by exploiting both space and frequency-selectivity with Alg. 3.

To give an example of the performance gains achieved using *joint tone-line selection* imagine that we have a desired rate of 4 Mbps for the far-end (1200m) lines. The rates that can be achieved for the near-end lines under the different schemes are then listed in Table II. Note that by using *joint line-tone selection* we can achieve 99% of the data-rate achieved using Full Crosstalk Cancellation with only 29% of the run-time complexity. Also note the significant gains of *joint line-tone selection* over *line selection only* or *tone selection only*. This demonstrates the importance of exploiting both the space and frequency selectivity of the crosstalk channels.

Scheme	Init. Complexity (Mults.)	$N = 7, K = 4096$
Line Selection Only	KN	29×10^3
Tone Selection Only	$K(N + 5)$	49×10^3
Joint Selection	$3K(N + 1)$	98×10^3

TABLE I

INITIALIZATION COMPLEXITY OF PARTIAL CANCELLATION ALGORITHMS

XI. CONCLUSIONS

Crosstalk is the limiting factor in VDSL performance. Many crosstalk cancellation techniques have been proposed and these lead to significant performance gains. Unfortunately full crosstalk cancellation has a high run-time complexity and this grows rapidly with the number of users in a binder.

Crosstalk channels in the DSL environment exhibit both space and frequency-selectivity. The majority of the effects of crosstalk are limited to a small number of crosstalkers and tones. Partial crosstalk cancellation exploits this by only performing crosstalk cancellation on the tones and lines where it gives the most benefit. This allows it to give close to the performance of full crosstalk cancellation with considerably reduced run-time complexity.

In this paper we presented several partial crosstalk cancellation algorithms for upstream transmission. A simple joint line-tone selection algorithm (Alg. 3) was presented which can achieve 99% of the performance of full cancellation with only 29% of the run-time complexity.

Whilst this paper has focused on crosstalk cancellation in VDSL the techniques here are also applicable to MIMO-CDMA systems. Taking into account the processing gain, the interference path typically has 15-20 dB more attenuation than the main path. Hence the MIMO-CDMA channel is column-wise diagonal dominant and the partial crosstalk cancellation techniques developed here can be directly applied.

Interesting areas for future work include partial crosstalk pre-compensation (for downstream transmission), and combinations of partial cancellation with dynamic spectrum management.

REFERENCES

- [1] G. Taubock and W. Henkel, "MIMO Systems in the Subscriber-Line Network," in *Proc. of the 5th Int. OFDM-Workshop*, 2000, pp. 18.1–18.3.
- [2] G. Ginis and J. Cioffi, "Vectored Transmission for Digital Subscriber Line Systems," *IEEE J. Select. Areas Commun.*, vol. 20, no. 5, pp. 1085–1104, June 2002.
- [3] W. Yu and J. Cioffi, "Multiuser Detection in Vector Multiple Access Channels using Generalized Decision Feedback Equalization," in *Proc. 5th Int. Conf. on Signal Processing, World Computer Congress*, 2000.
- [4] D. C. R. and M. Moonen, "Transmission and Multiplexing (TM); Access transmission systems on metallic access cables; VDSL; Functional Requirements, ETSI Std. TS 101 270-1/1, Rev. V.1.2.1, 1999.
- [5] G. Forney and M. Eyuboglu, "Combined Equalization and Coding Using Precoding," *IEEE Commun. Mag.*, vol. 29, no. 12, pp. 25–34, Dec. 1991.
- [6] R. Cendrillon and M. Moonen, "An Analysis of Linear Crosstalk Cancellation and Pre-compensation Techniques which Achieve Near Capacity in DSL Channels," ESAT, Katholieke Universiteit Leuven, Belgium, Tech. Rep., June 2003, available at <http://www.esat.kuleuven.ac.be/~rcendrillon>.
- [7] D. Gore and A. Paulraj, "Space-time block coding with optimal antenna selection," in *Proc. of the Int. Conf. on Acoustics, Speech and Sig. Processing*, 2001, pp. 2441–2444.
- [8] R. Cendrillon, M. Moonen, et al., "Partial Crosstalk Cancellation Exploiting Line and Tone Selection in Upstream DMT-VDSL," *Submitted to EURASIP Journal on Applied Signal Processing*, Feb 2003, available at <http://www.esat.kuleuven.ac.be/~rcendrillon>.
- [9] W. Yu, G. Ginis, and J. Cioffi, "Distributed Multiuser Power Control for Digital Subscriber Lines," *IEEE J. Select. Areas Commun.*, vol. 20, no. 5, pp. 1105–1115, June 2002.

Scheme	Far-end Rate	Near-end Rate
No Cancellation	0.1 Mbps	22 Mbps
Line Selection Only	4.0 Mbps	24 Mbps
Tone Selection Only	4.0 Mbps	42 Mbps
Joint Line-Tone Selection	4.0 Mbps	71 Mbps
Full Cancellation	4.0 Mbps	72 Mbps

TABLE II

RATES ACHIEVED UNDER DIFFERENT PARTIAL CANCELLATION SCHEMES




Fig. 1. FEXT Transfer Functions for 0.5 mm British Telecom Cable




Fig. 2. Achievable Rate Regions vs. Complexity (Joint Selection Algorithm)




Fig. 3. Achievable Rate Regions of Different Algorithms ($c = 2$)