Partial Crosstalk Cancellation Exploiting Line and Tone Selection in Upstream DMT-VDSL

Raphael Cendrillon and Marc Moonen
Katholieke Universiteit Leuven - ESAT/SCD
Kasteelpark Arenberg 10, Leuven-Heverlee 3001, Belgium
{raphael.cendrillon,marc.moonen}@esat.kuleuven.ac.be
Tel. +32 16 32 1788 Fax +32 16 32 1970

George Ginis
Texas Instruments
2043 Samaritan Drive, San Jose, CA 95124, USA
gginis@dsl.stanford.edu

Katleen Van Acker, Tom Bostoen and Piet Vandaele
Alcatel Bell
Francis Wellesplein 1, Antwerp 2018, Belgium
{katleen.van_acker,tom.bostoen,piet.vandaele}@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. We present both the optimal algorithm for partial crosstalk cancellation which has a high initialization complexity and a number of low-complexity algorithms which achieve near-optimal performance in a wide range of scenarios.

Index Terms

DSL, interference cancellation, reduced complexity, partial crosstalk cancellation, crosstalk selectivity, hybrid selection/combining

This work was carried out in the frame of IUAP P5/22, *Dynamical Systems and Control: Computation, Identification and Modeling* 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.

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-group*s 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.

There are two types of crosstalk, near-end crosstalk (NEXT) and far-end crosstalk (FEXT). NEXT occurs when the upstream signal of one modem couples into the downstream signal of another or vice versa. FEXT occurs when two signals traveling in the same direction couple. In VDSL NEXT is avoided through the use of FDD. FEXT on the other hand is still present. 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 based on linear pre- and post-filtering[1], successive interference cancellation[2], [3] and turbo coding[4]. These schemes are applicable to upstream (US) transmission where the receiving modems are co-located. In downstream transmission it is also possible to pre-compensate for crosstalk since the transmitters are then co-located at the central office (CO)[2]. Cancellation of crosstalk from alien systems like HPNA and HDSL has also been investigated[5], [6].

Since crosstalk is the dominant source of performance degradation in VDSL removing it leads to spectacular performance gains, e.g. $50 \rightarrow 130$ Mbps in the upstream direction[2]. 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 6.5 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.

This paper is organised as follows: In Sec. II we describe the system model for the crosstalk environment. Sec. III describes crosstalk cancellation, its performance and complexity. Due to the high complexity of full crosstalk cancellation in Sec. IV and V we introduce the concept of partial crosstalk cancellation which exploits both the space and frequency-selectivity of the crosstalk channel. This takes advantage of the fact that the majority of the crosstalk experienced by a modem comes from only a few other crosstalkers in the binder. Furthermore since crosstalk coupling varies dramatically with frequency, the worst effects of crosstalk are limited to a small selection of tones. Exploiting these two properties leads to significant reductions in complexity. In Sec. VI we describe a partial cancellation algorithm which exploits space-selectivity. An algorithm which exploits frequency-selectivity only is described in Sec. VII. As we will see, achieving the largest possible reduction in run-time complexity requires algorithms to exploit both forms of selectivity and in Sec. VIII we describe such algorithms. The performance of the algorithms is compared in Sec. IX and conclusions are drawn in Sec. X.

II. Upstream System Model

We begin by assuming that all receiving modems are co-located at the 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

$$\begin{bmatrix} y_k^1 \\ \vdots \\ y_k^{N+1} \end{bmatrix} = \begin{bmatrix} h_k^{(1,1)} & \cdots & h_k^{(1,N+1)} \\ \vdots & \ddots & \vdots \\ h_k^{(N+1,1)} & \cdots & h_k^{(N+1,N+1)} \end{bmatrix} \begin{bmatrix} x_k^1 \\ \vdots \\ x_k^{N+1} \end{bmatrix} + \begin{bmatrix} z_k^1 \\ \vdots \\ z_k^{N+1} \end{bmatrix}$$
$$\mathbf{y}_k = \mathbf{H}_k \mathbf{x}_k + \mathbf{z}_k$$

Here x_k^n and y_k^n denote the symbols transmitted and received by user n on tone k respectively. The tone k is in the range $1, \ldots, K$ where K is the number of tones in the DMT system (e.g. for VDSL $K = 4096$). $h_k^{(n,n)}$ is the direct channel of user n at tone k, and $h_k^{(n,m)}$ is the crosstalk channel from user m into user n. z_k^n represents the additive noise experienced by user n on tone k and is assumed to be spatially white and Gaussian such that $\mathcal{E}\left\{\mathbf{z}_k \mathbf{z}_k^H\right\} = \sigma_k^2 \mathbf{I}_N$. We denote the transmit auto-correlation on tone k as $\mathcal{E}\left\{\mathbf{x}_k \mathbf{x}_k^H\right\} = \mathbf{S}_k$ 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.

A matrix $\mathbf{A}$ is said to be *column-wise diagonal dominant* if it satisfies

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

where $a^{(n,m)} \triangleq [\mathbf{A}]_{n,m}$. Whilst $\mathbf{A}$ is said to be *row-wise diagonal dominant* if it satisfies

$$\left|a^{(n,n)}\right| \gg \left|a^{(n,m)}\right|, \forall n \neq m \tag{2}$$

If $\mathbf{A}$ satisfies both (1) and (2) it is said to be *strictly diagonal dominant.*

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

$$\left|h_k^{(m,m)}\right| \gg \left|h_k^{(n,m)}\right|, \forall n \neq m \tag{3}$$

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 [7]). 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

$$\mathbf{H}_k \stackrel{\text{svd}}{=} \mathbf{U}_k \Lambda_k \mathbf{V}_k^H \tag{4}$$

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 $\Lambda_k \triangleq \text{diag}\{\lambda_k^1, \ldots, \lambda_k^{N+1}\}$. It is assumed that $\mathbf{H}_k$ is non-singular which is ensured by (3) provided that $h_k^{(n,n)} \neq 0, \forall n$.

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

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

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

$$\begin{aligned} \widehat{x}_k^n &= \mathbf{w}_k^n \mathbf{y}_k \\ &= \mathbf{w}_k^n \left(\mathbf{H}_k \mathbf{P}_k \widetilde{\mathbf{x}}_k + \mathbf{z}_k\right) \\ &= \widetilde{x}_k^n + \widetilde{z}_k^n \end{aligned}$$

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 $\widetilde{z}_k^n \triangleq \mathbf{e}_n^H \Lambda_k^{-1} \mathbf{U}_k^H \mathbf{z}_k$. Here we use $[\mathbf{A}]_{\text{row}\,n}$ and $[\mathbf{A}]_{\text{col}\,n}$ to denote the nth row and column of matrix $\mathbf{A}$ respectively. Note that $\mathcal{E}\left\{|\widetilde{z}_k^n|^2\right\} = \sigma_k^2 (\lambda_k^n)^{-2}$. The pre and post-filtering operations remove crosstalk without causing noise enhancement. Applying a conventional slicer to $\widehat{x}_k^n$ achieves the following rate for user n on tone k

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

Γ represents the SNR-gap to capacity and is a function of the target BER, coding gain and noise margin[8]. The maximum achievable rate of the multi-line DSL channel is

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

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 $\widetilde{\mathbf{S}}_k$. 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 column-wise diagonal dominance 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.

Theorem 1: Any column-wise diagonal dominant matrix $\mathbf{H}_k$ which satisfies (16) can be decomposed

$$\mathbf{H}_k = \mathbf{Q}_k \mathbf{\Sigma}_k$$

such that $\mathbf{Q}_k$ is unitary and $\mathbf{\Sigma}_k$ is strictly diagonal dominant with positive diagonal elements. Furthermore, the off-diagonal elements of $\mathbf{\Sigma}_k$ can be bounded using (26) and (28).

Proof: See Appendix I. ■

The strict diagonal dominance of $\mathbf{\Sigma}_k$ allows us to make the approximations

$$\begin{aligned} \mathbf{\Sigma}_k &\simeq \text{diag}\left\{\mathbf{\Sigma}_k\right\} \\ \mathbf{\Sigma}_k^{-1} &\simeq \text{diag}\left\{\mathbf{\Sigma}_k\right\}^{-1} \end{aligned} \tag{5}$$

Hence

$$\mathbf{H}_k \simeq \mathbf{Q}_k \text{diag}\left\{\mathbf{\Sigma}_k\right\} \mathbf{I}_N \tag{6}$$

Comparison with (4) yields $\mathbf{U}_k \simeq \mathbf{Q}_k$, $\Lambda_k \simeq \text{diag}\left\{\mathbf{\Sigma}_k\right\}$ and $\mathbf{V}_k^H \simeq \mathbf{I}_N$. So the optimal transmit/receive structure of the previous section is well approximated by

$$\begin{aligned} \mathbf{P}_k &\simeq \mathbf{I}_N \\ \mathbf{w}_k^n &\simeq \mathbf{e}_n^H \text{diag}\left\{\mathbf{\Sigma}_k\right\}^{-1} \mathbf{Q}_k^H \\ &\simeq \mathbf{e}_n^H \mathbf{\Sigma}_k^{-1} \mathbf{Q}_k^H \\ &= \mathbf{e}_n^H \mathbf{H}_k^{-1} \end{aligned}$$

where we use (5) to go from line 2 to 3. In [9] an upper bound is proposed for the capacity loss incurred due to the above approximation. This is shown to be minimal for all practical DSL channels.

Since $\mathbf{P}_k = \mathbf{I}_N$ *pre-filtering is not required*. This is important since in upstream DSL transmitting modems are not co-located. Furthermore 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

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

Note that noise enhancement is not a problem since $\mathbf{H}_k^{-1} \simeq \mathbf{\Sigma}_k^{-1} \mathbf{Q}_k^H$. $\mathbf{Q}_k^H$ is unitary hence it does not alter the statistics of the noise. $\mathbf{\Sigma}_k^{-1}$ is approximately diagonal hence it scales the signal and noise equally.

Using this scheme crosstalk cancellation of one user at one tone requires N multiplications per DMT block. So crosstalk cancellation for $N+1$ users on K tones at a block-rate b (DMT blocks/second) requires $(N^2+N)Kb$ multiplications per second. Thus the 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 6.5 billion multiplications 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 cable consisting of 8×0.5mm 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. A 25-pair binder is depicted in Fig. 4. As can be seen, each pair is typically surrounded by 4-5 neighbours. Since electromagnetic coupling decreases rapidly with distance, each pair will experience significant crosstalk from only a few other surrounding pairs within the binder. Naturally twisted-pairs which are nearby within a binder-group will cause each other more crosstalk. 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.

To illustrate the space-selectivity of crosstalk we calculated the proportion of total crosstalk energy that is caused by the i largest crosstalkers of user n on tone k. All users have identical transmit PSDs hence from the perspective of user n, crosstalker m is said to be larger than crosstalker q at tone k if $|h_k^{(n,m)}| > |h_k^{(n,q)}|$. The result was averaged across all tones k and every line n within the binder. The measurements were done using the British Telecom cable and the result is shown in Fig. 2. As can be seen on average approximately 90% of crosstalk energy is caused by the 4 largest crosstalkers.

Second, crosstalk channels vary significantly with frequency. So whilst a user may experience significant crosstalk on one tone, weak crosstalk may be experienced on other tones. We refer to this as the *frequency-selectivity* of crosstalk which arises naturally from the frequency dependent nature of electromagnetic coupling.

To illustrate the frequency-selectivity of crosstalk we calculate the proportion of total crosstalk energy contained within the i worst tones. From the perspective of user n and crosstalker m, tone k is said to be worse than tone l if $|h_k^{(n,m)}| > |h_l^{(n,m)}|$. The result is shown in Fig 3. Approximately 90% of the crosstalk is contained within 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. These observations suggest that we can achieve the majority of the performance gains of crosstalk cancellation by canceling only the largest crosstalkers on each tone and we refer to this as partial crosstalk cancellation.

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 the design of partial crosstalk cancellation 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

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

and the corresponding received signals

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

We also define the set of lines which are not observed in the detection of user n on tone k

$$\begin{aligned}\underline{\mathbb{M}}_k^n &\triangleq \{1,\, \ldots,\, n-1,\; n+1,\, \ldots,\, N+1\} \setminus \mathbb{M}_k^n \\ &= \left\{\underline{m}_{k,n}(1),\, \ldots,\, \underline{m}_{k,n}(N-p_{k,n})\right\}\end{aligned}$$

where $\mathbb{A} \setminus \mathbb{B}$ denotes the elements contained in set $\mathbb{A}$ and not in set $\mathbb{B}$. We form an estimate of the transmitted symbol using a linear combination of the received signals on the *observation lines only*

$$\widehat{x}_k^n = \overline{\mathbf{w}}_k^n \overline{\mathbf{y}}_k^n$$

Note that crosstalk cancellation for user n at tone k now requires only $p_{k,n}$ multiplications per DMT block in contrast to the N multiplications required for full crosstalk cancellation. This technique has many similarities to hybrid selection/combining from the wireless field[10]. 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 cancellation coefficients $\overline{\mathbf{w}}_k^n$. We begin with a reduced system model which only contains the signals observed in the detection of user n at tone k

$$\overline{\mathbf{y}}_k^n = \overline{\mathbf{H}}_k^n \overline{\mathbf{x}}_k^n + \underline{\mathbf{H}}_k^n \underline{\mathbf{x}}_k^n + \overline{\mathbf{z}}_k^n$$

$\overline{\mathbf{x}}_k^n$ contains the signals transmitted onto the set of observed lines $\{n,\, \mathbb{M}_k^n\}$

$$\overline{\mathbf{x}}_k^n \triangleq \left[x_k^n,\, x_k^{m_{k,n}(1)}\, \cdots\, x_k^{m_{k,n}(p_{k,n})}\right]^T$$

and $\overline{\mathbf{H}}_k^n$ contains the corresponding channels

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

where $[\mathbf{A}]_{\text{rows}\, \mathbb{A},\, \text{cols}\, \mathbb{B}}$ denotes the sub-matrix formed from the rows $\mathbb{A}$ and columns $\mathbb{B}$ of matrix $\mathbf{A}$. $\underline{\mathbf{x}}_k^n$ contains the signals transmitted onto the set of non-observed lines $\underline{\mathbb{M}}_k^n$

$$\underline{\mathbf{x}}_k^n \triangleq \left[x_k^{\underline{m}_{k,n}(1)}\, \cdots\, x_k^{\underline{m}_{k,n}(N-p_{k,n})}\right]^T$$

and $\underline{\mathbf{H}}_k^n$ contains the corresponding channels

$$\underline{\mathbf{H}}_k^n \triangleq \begin{bmatrix} [\mathbf{H}_k^n]_{\text{row } n,\, \text{cols}\, \underline{\mathbb{M}}_k^n} \\ [\mathbf{H}_k^n]_{\text{rows}\, \mathbb{M}_k^n,\, \text{cols}\, \underline{\mathbb{M}}_k^n} \end{bmatrix}$$

$\overline{\mathbf{z}}_k^n$ contains the noise seen on the observed lines

$$\overline{\mathbf{z}}_k^n \triangleq \left[z_k^n,\, z_k^{m_{k,n}(1)}\, \cdots\, z_k^{m_{k,n}(p_{k,n})}\right]^T$$

We choose a ZF design which was shown in Sec. III-B to be a near-optimal transmit/receive structure. The partial cancellation filter is designed to remove all crosstalk from crosstalkers in the set $\mathbb{M}_k^n$

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

where $\overline{\mathbf{e}}_n \triangleq \left[\mathbf{I}_{p_{k,n}+1}\right]_{\text{col } n}$. Hence

$$\widehat{x}_k^n = x_k^n + \overline{\mathbf{w}}_k^n \underline{\mathbf{H}}_k^n \underline{\mathbf{x}}_k^n + \overline{\mathbf{w}}_k^n \overline{\mathbf{z}}_k^n$$

The first term is the transmitted signal whilst the second and third terms are the residual crosstalk and filtered noise respectively.

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

$$\max_{\mathbb{M}_k^n} c_k^n \text{ s.t. } |\mathbb{M}_k^n| \leq p_{k,n} \tag{7}$$

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

A. Residual Interference

Column-wise diagonal dominance in $\mathbf{H}_k$ implies the same in $\overline{\mathbf{H}}_k^n$. Hence we can use the decomposition defined in Theorem 1

$$\overline{\mathbf{H}}_k^n = \overline{\mathbf{Q}}_k^n \overline{\mathbf{\Sigma}}_k^n$$

where $\overline{\mathbf{Q}}_k^n$ is unitary and $\overline{\mathbf{\Sigma}}_k^n$ strictly diagonal dominant. Hence

$$\overline{\mathbf{\Sigma}}_k^n \simeq \operatorname{diag}\left\{\overline{\mathbf{\Sigma}}_k^n\right\}$$

We define $\overline{\rho}_{k,n}^{(i,j)} \triangleq \left[\overline{\mathbf{\Sigma}}_k^n\right]_{i,j}$. Since $\overline{\mathbf{\Sigma}}_k^n$ is column-wise diagonal dominant

$$\begin{aligned} \left[\overline{\mathbf{H}}_k^n\right]_{\text{col}\,1} &= \sum_i \overline{\rho}_{k,n}^{(i,1)} \left[\overline{\mathbf{Q}}_k^n\right]_{\text{col}\,i} \\ &\simeq \overline{\rho}_{k,n}^{(1,1)} \left[\overline{\mathbf{Q}}_k^n\right]_{\text{col}\,1} \end{aligned} \tag{8}$$

Now since the diagonal elements of $\overline{\mathbf{\Sigma}}_k^n$ are positive taking the norm of both sides of (8) yields

$$\begin{aligned} \overline{\rho}_{k,n}^{(1,1)} &\simeq \left\|\left[\overline{\mathbf{H}}_k^n\right]_{\text{col}\,1}\right\|_2 \left\|\left[\overline{\mathbf{Q}}_k^n\right]_{\text{col}\,1}\right\|_2^{-1} \\ &\simeq \left|h_k^{(n,n)}\right| \end{aligned}$$

where we use the column-wise diagonal dominance of $\overline{\mathbf{H}}_k^n$ and the observation $\left[\overline{\mathbf{H}}_k^n\right]_{1,1} = h_k^{(n,n)}$. Hence

$$\begin{aligned} \overline{\mathbf{w}}_k^n &= \overline{\mathbf{e}}_1^H \left(\overline{\mathbf{H}}_k^n\right)^{-1} \\ &\simeq \overline{\mathbf{e}}_1^H \operatorname{diag}\left\{\overline{\mathbf{\Sigma}}_k^n\right\}^{-1} \left(\overline{\mathbf{Q}}_k^n\right)^H \\ &\simeq \left(\overline{\rho}_{k,n}^{(1,1)}\right)^{-1} \left[\overline{\mathbf{Q}}_k^n\right]_{\text{col}\,1}^H \\ &\simeq \left|h_k^{(n,n)}\right|^{-1} \left[\overline{\mathbf{Q}}_k^n\right]_{\text{col}\,1}^H \end{aligned}$$

From (8)

$$\begin{aligned} \left[\overline{\mathbf{Q}}_k^n\right]_{\text{col}\,1}^H &\simeq \left(\overline{\rho}_{k,n}^{(1,1)}\right)^{-1} \left[\overline{\mathbf{H}}_k^n\right]_{\text{col}\,1}^H \\ &\simeq \left|h_k^{(n,n)}\right|^{-1} \left[\begin{array}{cccc} \left(h_k^{(n,n)}\right)^* & \left(h_k^{(m_{k,n}(1),n)}\right)^* & \cdots & \left(h_k^{(m_{k,n}(p_{k,n}),n)}\right)^* \end{array}\right] \end{aligned} \tag{9}$$

Thus we find

$$\overline{\mathbf{w}}_k^n \simeq \left|h_k^{(n,n)}\right|^{-2} \left[\begin{array}{cccc} \left(h_k^{(n,n)}\right)^* & \left(h_k^{(m_{k,n}(1),n)}\right)^* & \cdots & \left(h_k^{(m_{k,n}(p_{k,n}),n)}\right)^* \end{array}\right]$$

Using (3) we can make the approximation

$$\overline{\mathbf{w}}_k^n \underline{\mathbf{H}}_k^n \simeq \left(h_k^{(n,n)}\right)^* \left|h_k^{(n,n)}\right|^{-2} \left[\underline{\mathbf{H}}_k^n\right]_{\text{row}\,1}$$

hence the residual interference

$$\overline{\mathbf{w}}_k^n \underline{\underline{\mathbf{H}}}_k^n \underline{\mathbf{x}}_k^n \simeq \left(h_k^{(n,n)}\right)^* \left|h_k^{(n,n)}\right|^{-2} \sum_{m \in \underline{\mathbb{M}}_k^n} h_k^{(n,m)} x_k^m$$

The power of the residual interference is thus

$$\mathcal{E}\left\{\overline{\mathbf{w}}_k^n \underline{\underline{\mathbf{H}}}_k^n \underline{\mathbf{x}}_k^n \left(\overline{\mathbf{w}}_k^n \underline{\underline{\mathbf{H}}}_k^n \underline{\mathbf{x}}_k^n\right)^H\right\} \simeq \left|h_k^{(n,n)}\right|^{-2} \sum_{m \in \underline{\mathbb{M}}_k^n} \left|h_k^{(n,m)}\right|^2 s_k^m$$

B. Filtered Noise

Using (3) and (9) we can make the approximation

$$\overline{\mathbf{w}}_k^n \overline{\mathbf{z}}_k^n \simeq \left(h_k^{(n,n)}\right)^* \left|h_k^{(n,n)}\right|^{-2} z_k^n$$

The power of the filtered noise is thus

$$\mathcal{E}\left\{\overline{\mathbf{w}}_k^n \overline{\mathbf{z}}_k^n \left(\overline{\mathbf{w}}_k^n \overline{\mathbf{z}}_k^n\right)^H\right\} \simeq \left|h_k^{(n,n)}\right|^{-2} \sigma_k^2$$

C. SINR after Partial Crosstalk Cancellation

After crosstalk cancellation we have the following estimate of the transmitted signal

$$\widehat{x}_k^n = x_k^n + \overline{\mathbf{w}}_k^n \underline{\underline{\mathbf{H}}}_k^n \underline{\mathbf{x}}_k^n + \overline{\mathbf{w}}_k^n \overline{\mathbf{z}}_k^n$$

The Signal to Interference plus Noise Ratio (SINR) at the input of the decision device is thus

$$\mathrm{SINR}_k^n \simeq \frac{\left|h_k^{(n,n)}\right|^2 s_k^n}{\sum_{m \in \underline{\mathbb{M}}_k^n} \left|h_k^{(n,m)}\right|^2 s_k^m + \sigma_k^2} \tag{10}$$

with the approximation becoming exact in strongly column-wise diagonal dominant channels.

There are two interesting observations to make at this point. First, as we expected the ZF crosstalk canceller removes crosstalk caused by the modems in the set $\mathbb{M}_k^n$ perfectly. Second, more surprisingly the ZF crosstalk canceller does not change the statistics of the crosstalk caused by modems outside of the set $\mathbb{M}_k^n$. It also does not change the statistics of the noise. So *the column-wise diagonal dominant property of* $\mathbf{H}_k$ *ensures us that a ZF partial crosstalk canceller will not cause enhancement of the crosstalk caused by modems outside* $\mathbb{M}_k^n$ *or of the noise.*

D. Line Selection Algorithm

Maximizing SINR_k^n and thus rate c_k^n corresponds to minimizing the amount of interference in the set $\underline{\mathbb{M}}_k^n$. 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

$$\{q_{k,n}(1), \ldots, q_{k,n}(N)\} \quad \text{s.t} \quad \left|h_k^{(n,q_{k,n}(i))}\right|^2 s_k^{q_{k,n}(i)} \geq \left|h_k^{(n,q_{k,n}(i+1))}\right|^2 s_k^{q_{k,n}(i+1)}, \forall i$$
$$q_{k,n}(i) \neq n, \forall i$$

Remark 1: Optimal Line Selection

In column-wise diagonal dominant channels the set $\mathbb{M}_k^n$ which maximizes the rate of user n on tone k subject to a complexity constraint of $p_{k,n}$ multiplications/DMT-block (see optimisation in (7)) is

$$\mathbb{M}_k^n = \{q_{k,n}(1), \ldots, q_{k,n}(p_{k,n})\}$$

Proof: Follows from examination of (10). ■

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

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

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

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

Algorithm 1 Line Selection Only

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

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 $O(KN)$ multiplications 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 $\overline{\mathbf{w}}_k^n$ for each tone. This requires $\mathcal{O}(\sum_k (p_{k,n}+1)^3)$ multiplications for user n regardless of the partial cancellation algorithm employed. We assume that the direct and crosstalk channel gains $\left|h_k^{(n,m)}\right|^2$, $\forall n, m, k$ are available and do not need to be calculated.

The initialization complexity (in terms of multiplications and logarithm operations per user) of the different partial cancellation algorithms is listed in Tab. I. The required number of sort operations of each size is listed in Tab. II. 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 high so 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 cancelled

$$r_{k,n}(p_{k,n}) \triangleq \log\left(1 + \frac{1}{\Gamma}\frac{\left|h_k^{(n,n)}\right|^2 s_k^n}{\sum_{i=p_{k,n}+1}^{N}\left|h_k^{(n,q_{k,n}(i))}\right|^2 s_k^{q_{k,n}(i)} + \sigma_k^2}\right) \tag{11}$$

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

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

and the indices of the tones ordered by this gain

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

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 $r_{k,n}(0)$.

Algorithm 2 Tone Selection Only

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

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 multiplications/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 $\mathcal{O}(KN)$ multiplications 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.

A. Simple Joint Tone-Line Selection

As we saw in Sec. VI-C 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 cancelled 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, as will be demonstrated in Sec. IX.

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

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

Note that if we work in a logarithmic scale then $\overline{g}_{k,n}(m)$ 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 $\overline{g}_n(d_n(i)) \triangleq \overline{g}_{k_n(i),n}(m_n(i))$. This allows us to define the indices of (crosstalker, tone) pairs ordered by gain

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

We can now define our simplified 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

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

Algorithm 3 Simple Line-Tone Selection

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

This leads to a run-time complexity of cK multiplications per DMT-block per user. The benefit of this algorithm is its low complexity. Pair selection for one user has a complexity of $\mathcal{O}(KN)$ multiplications 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. In Sec. IX we will see that this algorithm leads to near-optimal performance.

B. Optimum Joint Tone-Line Selection

It is interesting to evaluate the sub-optimality of the algorithms we described so far through an upper bound achieved by a truly optimal partial cancellation algorithm. The problem of partial cancellation is effectively a *resource allocation problem*. Given cK multiplications per user we need to distribute these across tones such that the largest rate is achieved

$$\max_{\{\mathbb{M}_k^n\}_{k=1,\ldots,K}} \sum_k c_k^n \quad \text{s.t.} \sum_k |\mathbb{M}_k^n| \leq cK$$

Since the channel is column-wise diagonal dominant Remark 1 allows us to determine in a simple fashion the best set of lines to observe in the detection of user n. Hence our problem simplifies to

$$\max_{\{p_{k,n}\}_{k=1,\ldots,K}} \sum_k c_k^n \quad \text{s.t.} \sum_k p_{k,n} \leq cK$$

An exhaustive search could require us to evaluate up to N^K different allocations. In VDSL $K = 4096$ which makes any such search numerically intractable.

Due to the structure of the problem it is possible to come up with a greedy algorithm, Alg. 4 which will iteratively find the optimal allocation for some values of c. The algorithm cannot find a solution for any arbitrary value of c however the range of values of c generated by the algorithm are so closely spaced that this is not a practical problem. Define the value of canceling p crosstalkers on tone k as

$$v_{k,n}(p) = \frac{r_{k,n}(p) - r_{k,n}(0)}{p}$$

Recall that $r_{k,n}(p)$ is the rate achieved by user n on tone k when the p largest crosstalkers are canceled and is evaluated using (11). Value is the increase in rate (benefit) divided by the increase in run-time complexity (cost). It measures increase in bit-rate per multiplication when p multiplications are spent on tone k. The algorithm begins by initializing $v_{k,n}(p)$ for all values of p and k. It then proceeds as follows

1) Find choice of tone k and cancelled crosstalkers p with largest value $v_{k,n}(p)$. Store this in (k_s, p_s)
2) Set lines to be observed on tone k_s to $\mathbb{M}_{k_s}^n = \{q_{k_s,n}(1), \ldots, q_{k_s,n}(p_s)\}$
3) Set value of canceling p_s or less crosstalkers on tone k_s to zero. This prevents re-selection of previously selected pairs.
4) Update value of canceling p_s+1 or more crosstalkers on tone k_s. The rate increase and cost should be relative to the currently selected number of crosstalkers.

Algorithm 4 Optimal Line-Tone Selection

init $v_{k,n}(p) = (r_{k,n}(p) - r_{k,n}(0))\,/p \quad \forall\, k,\, p > 0$

repeat

$(k_s, p_s) = \arg\max_{(k,p)} v_{k,n}(p)$

$\mathbb{M}_{k_s}^n = \{q_{k_s,n}(1), \ldots, q_{k_s,n}(p_s)\}$

$v_{k_s,n}(p) = 0 \qquad p = 1, \ldots, p_s$

$v_{k_s,n}(p) = (r_{k_s,n}(p) - r_{k_s,n}(p_s))\,/\,(p - p_s) \qquad p = p_s + 1, \ldots, N$

while $\sum_k |\mathbb{M}_k^n| < cK$

The algorithm iterates through steps 1-4 until the allocated complexity exceeds cK. This yields an upper bound on the partial crosstalk cancellation performance for a given complexity. Since the algorithm allocates at most

N multiplications in each iteration, the total allocated complexity will be at the most $cK + N$. With $K = 4096$ typically $cK \gg N$. Hence the difference between the desired run-time complexity and that of the solution provided by the algorithm is minimal. The upper bound is thus tight.

Like Alg. 3, this algorithm can exploit both the space and frequency-selectivity of crosstalk to reduce run-time complexity. This algorithm generates a resource allocation at the end of each iteration which is optimal. That is, of all the resource allocations of equal run-time complexity the one generated by this algorithm achieves the highest rate. Unfortunately this algorithm is considerably more complex than Alg. 3. Pair selection for a single user requires $\mathcal{O}(KN^2)$ multiplications and $O(KN)$ logarithm operations. It is hard to define the exact sorting complexity since it varies significantly with the scenario. Sorting complexity is typically much higher than any of the other algorithms and can require up to KN sort operations which can have sizes as large as KN.

C. 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 multiplications can be shared between users. That is, the true constraint is on the total complexity of crosstalk cancellation for *all users*

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

The available complexity can be divided between users based on our desired rates for each. Denote the number of multiplications/DMT-block allocated to user n as κ_n, then

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

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. [11], [12]) we can define a rate region as the set of all achievable rate-tuples under a given total complexity constraint. This allows us to visualise the different trade-offs that can be achieved between the rates of different users inside a binder.

Limiting crosstalk cancellation on each tone to the users who benefit the most leads to further reductions in run-time complexity with minimal performance loss. This is demonstrated in Sec. IX-B.

IX. Performance

We now compare the performance of the partial crosstalk cancellation algorithms described in sections VI, VII and VIII. Performance is compared over a range of scenarios with crosstalk channels which exhibit both space and frequency-selectivity. As we show, the ability to exploit both space and frequency-selectivity is essential for achieving low run-time complexity in all scenarios.

We use semi-empirical transfer functions from the ETSI VDSL standards[7]. 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) 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. So 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. The number of lines in the binder is always 8, so $N = 7$. Other simulation parameters are listed in Tab. III.

A. Equidistant Lines $(8 \times 1000m)$

In the first scenario the binder contains 8×1000m lines. Since the lines are of equal length the crosstalk channels exhibit frequency-selectivity only; no space-selectivity is present. Shown in Fig. 5 are the rates achieved by each of the algorithms versus run-time complexity. Complexity is shown as a percentage relative to full crosstalk cancellation ($c = N$).

Alg. 1 can only exploit space-selectivity. There is no space-selectivity in this scenario so this algorithm gives extremely poor performance. Worst of all, we actually see a non-convex rate vs. run-time complexity curve. So doing partial crosstalk cancellation gives worse performance than time-sharing. In other words we could do full crosstalk cancellation for some fraction of the time, and none for the rest and this would lead to better performance than Alg. 1 with the same run-time complexity. The reason for this is as follows: As we increase the number of crosstalkers canceled $p_{k,n}$ the increase in signal-to-interference ratio (SIR) grows rapidly.

We illustrate this with the following example. Consider a binder with 7 crosstalkers. Let us assume that the crosstalkers all have identical crosstalk channels $\chi_k^{(n)}$ to user n as is the case in our simulation. Cancelling the first crosstalker causes the SIR to increase from $\frac{1}{7}|h_k^{(n,n)}|^2|\chi_k^{(n)}|^{-2}$ to $\frac{1}{6}|h_k^{(n,n)}|^2|\chi_k^{(n)}|^{-2}$. Cancelling the sixth crosstalker gives a much larger SIR increase from $\frac{1}{2}|h_k^{(n,n)}|^2|\chi_k^{(n)}|^{-2}$ to $|h_k^{(n,n)}|^2|\chi_k^{(n)}|^{-2}$. In general cancelling the pth crosstalker leads to an SIR increase of $(N-p+1)^{-1}(N-p)^{-1}|h_k^{(n,n)}|^2|\chi_k^{(n)}|^{-2}$. So the increase in SIR grows with rapidly with p as $p \to N$. Recall that $c_k^n = \log(1 + SINR_k^n) \simeq SINR_k^n$ for low $SINR_k^n$. So when crosstalkers have equal strength and the SINR is low, data-rate gain will grow rapidly with the number of crosstalkers cancelled p. This is why cancelling N crosstalkers typically gives greater than N times the data-rate gain of canceling one crosstalker. This leads to the non-convex rate-complexity curve of Fig. 5.

When the channel exhibits space-selectivity the first crosstalker causes much more interference than the second and so on. This effect counter-acts the rapid growth of SIR with p. As a result the best trade-off between performance and complexity usually occurs somewhere between no and full crosstalk cancellation.

Alg. 2 cannot exploit space-selectivity. In this scenario this is not a problem since all crosstalkers have equal strength. Alg. 2 can implement a form of frequency-sharing. This is analogous to the time-sharing just discussed and allows this algorithm to cancel e.g. 6 crosstalkers on half of the tones instead of 3 crosstalkers on all of the tones. For this reason Alg. 2 will always give a convex rate vs. complexity curve. Comparing the performance of Alg. 2 to the optimal algorithm Alg. 4 we see that it gives near-optimal performance in this scenario.

Alg. 3 also gives near-optimal performance. Note that with 29% of the complexity of full crosstalk cancellation we can achieve 89% of the performance gains.

B. Near-Far Scenario $(4 \times 300\text{m}, 4 \times 1200\text{m})$

We now evaluate the selection algorithms in a binder consisting of 4×300m loops and 4×1200m loops. In this configuration the lines suffer the near-far effect causing all users to see the 300m near-end lines as the largest sources of crosstalk. This space-selectivity assists the partial cancellation algorithms in reducing run-time complexity.

Frequency-selectivity is present in this scenario and is most pronounced on far-end lines. Near-end lines have relatively flat channels and benefit less from algorithms which exploit frequency-selectivity alone.

Fig. 6 contains the rates of the 300m near-end users versus complexity under the different algorithms. Fig. 7 contains the same for the 1200m far-end users.

Alg. 1 cannot exploit frequency-selectivity. On near-end lines frequency-selectivity is minimal and reasonable performance is still achieved. Again we see a non-convex rate-complexity curve however above 43% complexity Alg. 1 gives near-optimal performance. On far-end users frequency-selectivity is pronounced and Alg. 1 gives poor performance.

Alg. 2 cannot exploit space-selectivity and on near-end users this leads to poor performance which is virtually identical to time sharing. On far-end users frequency-selectivity is pronounced and this algorithm still achieves reasonable performance despite its inability to exploit space-selectivity.

Alg. 3 can exploit both space and frequency-selectivity. As a result it gives near-optimal performance for both near and far-end users. With 43% complexity this algorithm can achieve 99% of the performance gains on near-end users. On far-end users 29% complexity achieves 97% of the performance gains.

We now examine the distribution of run-time complexity between users as described in Sec. VIII-C. Fig. 8 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 $2\mu_{\text{far}}cK$ multiplications per 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. 8 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. 9 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 in Alg. 3 and Alg. 4.

C. Distributed Scenario $(300 : 100 : 1000m)$

Simulations were run in a distributed scenario consisting of 8 lines ranging from 300m to 1000m in 100m increments. Alg. 3 exhibited near-optimal performance and could increase the average rate from 9.7 Mbps to 23.7 Mbps with only 29% of the complexity of full crosstalk cancellation. This is equivalent to 2 times the complexity of a conventional FEQ.

We have seen that the performance of algorithms which exploit only one type of selectivity such as Alg. 1 and Alg. 2 varies considerably with the scenario. By exploiting both space and frequency-selectivity Alg. 3 consistently gave near-optimal performance. This algorithm is also considerably less complex than the optimal algorithm, Alg. 4.

X. Conclusions

Crosstalk is *the* limiting factor in VDSL performance. Many crosstalk cancellation techniques have been proposed and these lead to significant performance gains. Unfortunately 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. It was seen that designing a partial crosstalk canceller requires us to choose which lines to observe when detecting each user on each tone. This is equivalent to choosing the (crosstalker, tone) pairs to cancel in the detection of each user. We described different algorithms for choosing pairs. These included simplistic algorithms such as Alg. 1 which exploits space-selectivity only, and Alg. 2 which exploits frequency-selectivity only. In Sec. IX we saw that the performance of these two algorithms varies greatly depending on the scenario. Robust performance requires us to exploit both space and frequency-selectivity together.

We presented an optimal algorithm (Alg. 4) for partial crosstalk cancellation. Whilst this algorithm is highly complex its ability to exploit both space and frequency-selectivity led to good performance in all scenarios. Partial crosstalk canceller initialization for one user in this algorithm requires $\mathcal{O}(KN^2)$ multiplications and $\mathcal{O}(KN)$ logarithms.

A simple joint selection algorithm (Alg. 3) was described which decouples the problem of (crosstalker, tone) pair selection thereby reducing initialization complexity significantly. This algorithm gave near-optimal performance in all of the scenarios we evaluated and has an initialization complexity of only $\mathcal{O}(KN)$ multiplications per user.

With Alg. 3 it is possible to increase the average rate from 9.7 to 23.7 Mbps using only 2 times the run-time complexity of a conventional single-user detector (SUD) ie. frequency domain equalizer (FEQ), as is currently implemented in VDSL modems. With this complexity the algorithm achieves 89% of the performance gains of full crosstalk cancellation.

By treating computational complexity as a resource to be divided across tones and users we developed rate regions in Sec. IX. These allow us to visualize all of the achievable rate-tuples under a certain run-time complexity constraint. This is quite similar to work done in the areas of multi-user power allocation (see e.g. [11], [12]) however here we consider the allocation of computing resources rather than transmit power.

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[13]. Hence the MIMO-CDMA channel is column-wise diagonal dominant and the partial crosstalk cancellation techniques developed here can be directly applied.

In this work we have considered crosstalk cancellation, which is applicable only to upstream DSL where receivers are co-located at the CO. In downstream DSL it is also possible to mitigate the effects of crosstalk through crosstalk pre-compensation[2]. The development of partial crosstalk pre-compensation algorithms with reduced run-time complexity is the subject of ongoing research.

The simulations done here neglected the problem of power loading and assumed flat transmit PSDs. The use of non-flat PSDs through multi-user water-filling or power back-off is currently the subject of much activity in the research community (see e.g. [12], [14]). The use of non-flat PSDs increases space and frequency-selectivity and would allow partial cancellation to achieve even greater run-time complexity reductions whilst maintaining similar performance. The combination of multi-user power allocation and partial cancellation will lead to even larger achievable rates with implementable run-time complexities. This is an important area for future work.

APPENDIX I
PROOF OF THEOREM 1

Sort the diagonal elements of $\mathbf{H}_k$ in decreasing order

$$\{t_k(1), \ldots, t_k(N+1)\} \quad \text{s.t.}\, h_k^{(t_k(i),t_k(i))} \geq h_k^{(t_k(i+1),t_k(i+1))}, \forall i \tag{12}$$

We define the permutation matrix

$$\mathbf{\Pi}_k \triangleq \left[\begin{array}{ccc} \mathbf{e}_{t_k(1)} & \cdots & \mathbf{e}_{t_k(N+1)} \end{array} \right]^T$$

We use $\mathbf{\Pi}_k$ to re-order the rows and columns of $\mathbf{H}_k$

$$\widetilde{\mathbf{H}}_k = \mathbf{\Pi}_k^H \mathbf{H}_k \mathbf{\Pi}_k \tag{13}$$

Define $\widetilde{h}_k^{(n,m)} \triangleq \left[\widetilde{\mathbf{H}}_k\right]_{n,m}$. From (13)

$$\widetilde{h}_k^{(n,m)} = h_k^{(t_k(n),t_k(m))}$$

Using (12) yields

$$\widetilde{h}_k^{(n,n)} \geq \widetilde{h}_k^{(m,m)}, \forall m > n \tag{14}$$

So application of $\mathbf{\Pi}_k$ re-orders the rows and columns of $\mathbf{H}_k$ such that its diagonal elements are in decreasing order. Define α which measures the degree of column-wise diagonal dominance of the channel

$$\alpha \triangleq \max_{\substack{n,m \\ n \neq m}} \arctan \frac{\left|h_k^{(n,m)}\right|}{\left|h_k^{(m,m)}\right|} = \max_{\substack{n,m \\ n \neq m}} \arctan \frac{\left|\widetilde{h}_k^{(n,m)}\right|}{\left|\widetilde{h}_k^{(m,m)}\right|} \tag{15}$$

such that

$$\left|h_k^{(n,m)}\right|^2 \leq \tan^2 \alpha \left|h_k^{(m,m)}\right|^2, \forall n \neq m \tag{16}$$

Using the QR decomposition

$$\widetilde{\mathbf{H}}_k = \widetilde{\mathbf{Q}}_k \widetilde{\mathbf{R}}_k$$

We define the QR decomposition such that $\widetilde{\mathbf{R}}_k$ has positive values on the diagonal. This is without loss of generality. From [2]

$$\left|\widetilde{h}_k^{(n,n)}\right| \sqrt{1-4\alpha^2} \leq \left|\widetilde{r}_k^{(n,n)}\right| \leq \left|\widetilde{h}_k^{(n,n)}\right| \sqrt{1+N\tan^2\alpha} \tag{17}$$

where $\widetilde{r}_k^{(n,m)} \triangleq \left[\widetilde{\mathbf{R}}_k\right]_{n,m}$.

Now

$$\left\| \left[\widetilde{\mathbf{R}}_k\right]_{\text{col}\, m} \right\|_2^2 = \left\| \left[\widetilde{\mathbf{H}}_k\right]_{\text{col}\, m} \right\|_2^2, \forall m$$

implies

$$\begin{aligned}
\left|\widetilde{r}_k^{(n,m)}\right|^2 &= \left\|\left[\widetilde{\mathbf{H}}_k\right]_{\mathrm{col}\,m}\right\|_2^2 - \left|\widetilde{r}_k^{(m,m)}\right|^2 - \sum_{i\notin\{n,m\}}\left|\widetilde{r}_k^{(i,m)}\right|^2, \forall n\neq m \\
&\leq \left|\widetilde{h}_k^{(m,m)}\right|^2 + \sum_{i\neq m}\left|\widetilde{h}_k^{(i,m)}\right|^2 - \left|\widetilde{r}_k^{(m,m)}\right|^2, \forall n\neq m \\
&\leq \left|\widetilde{h}_k^{(m,m)}\right|^2\left[1+N\tan^2\alpha\right] - \left|\widetilde{r}_k^{(m,m)}\right|^2, \forall n\neq m \qquad (18)\\
&\leq \left|\widetilde{r}_k^{(m,m)}\right|^2\left[\frac{1+N\tan^2\alpha}{1-4\alpha^2}-1\right], \forall n\neq m \\
&= \left|\widetilde{r}_k^{(m,m)}\right|^2\left[\frac{4\alpha^2+N\tan^2\alpha}{1-4\alpha^2}\right], \forall n\neq m
\end{aligned}$$

where we use (15) to get from line 2 to 3 and the lower bound in (17) to get from line 3 to 4. Hence

$$\left|\widetilde{r}_k^{(n,m)}\right|^2 \leq f_1(\alpha)\left|\widetilde{r}_k^{(m,m)}\right|^2, \forall n\neq m \qquad (19)$$

where

$$f_1(\alpha) \triangleq \frac{4\alpha^2+N\tan^2\alpha}{1-4\alpha^2}$$

From (18)

$$\begin{aligned}
\left|\widetilde{h}_k^{(m,m)}\right|^2\left[1+N\tan^2\alpha\right] &\geq \left|\widetilde{r}_k^{(m,m)}\right|^2 + \left|\widetilde{r}_k^{(n,m)}\right|^2, \forall n\neq m \\
&\geq \left|\widetilde{r}_k^{(m,m)}\right|^2
\end{aligned}$$

Hence

$$\left|\widetilde{h}_k^{(m,m)}\right|^2 \geq \frac{\left|\widetilde{r}_k^{(m,m)}\right|^2}{\left[1+N\tan^2\alpha\right]} \qquad (20)$$

Now

$$\left\|\left[\widetilde{\mathbf{R}}_k\right]_{\mathrm{col}\,n}\right\|_2^2 = \left\|\left[\widetilde{\mathbf{H}}_k\right]_{\mathrm{col}\,n}\right\|_2^2$$

implies

$$\begin{aligned}
\left|\widetilde{r}_k^{(n,n)}\right|^2 &= \left\|\left[\widetilde{\mathbf{H}}_k\right]_{\mathrm{col}\,n}\right\|_2^2 - \sum_{m\neq n}\left|\widetilde{r}_k^{(m,n)}\right|^2 \\
&\geq \left|\widetilde{h}_k^{(n,n)}\right|^2 - \sum_{m\neq n}\left|\widetilde{r}_k^{(m,n)}\right|^2 \\
&\geq \left|\widetilde{h}_k^{(n,n)}\right|^2 - \left|\widetilde{r}_k^{(n,n)}\right|^2 N f_1(\alpha)
\end{aligned}$$

where we use (19) to get from line 2 to 3. So

$$\begin{aligned}
\left|\widetilde{r}_k^{(n,n)}\right|^2 &\geq \frac{\left|\widetilde{h}_k^{(n,n)}\right|^2}{1+Nf_1(\alpha)} \\
&\geq \frac{\left|\widetilde{h}_k^{(m,m)}\right|^2}{1+Nf_1(\alpha)}, \forall m>n \\
&\geq \frac{\left|\widetilde{r}_k^{(m,m)}\right|^2}{\left[1+Nf_1(\alpha)\right]\left[1+N\tan^2\alpha\right]}, \forall m>n \\
&\geq \frac{\left|\widetilde{r}_k^{(n,m)}\right|^2}{f_1(\alpha)\left[1+Nf_1(\alpha)\right]\left[1+N\tan^2\alpha\right]}, \forall m>n \qquad (21)
\end{aligned}$$

where we use (14) to get from line 1 to 2, (20) from 2 to 3 and (19) from 3 to 4. Since $\widetilde{\mathbf{R}}_k$ is upper triangular $\widetilde{r}_k^{(n,m)} = 0, \forall m < n$ which implies

$$\left|\widetilde{r}_k^{(n,n)}\right|^2 \geq \left|\widetilde{r}_k^{(n,m)}\right|^2 = 0, \forall m < n \tag{22}$$

Combining (21) and (22)

$$\left|\widetilde{r}_k^{(n,m)}\right|^2 \leq f_2(\alpha)\left|\widetilde{r}_k^{(n,n)}\right|^2, \forall n \neq m \tag{23}$$

where

$$f_2(\alpha) \triangleq f_1(\alpha)\left[1 + N f_1(\alpha)\right]\left[1 + N\tan^2\alpha\right]$$

From (13)

$$\begin{aligned}\mathbf{H}_k &= \mathbf{\Pi}_k\widetilde{\mathbf{H}}_k\mathbf{\Pi}_k^H \\ &= \left(\mathbf{\Pi}_k\widetilde{\mathbf{Q}}_k\mathbf{\Pi}_k^H\right)\left(\mathbf{\Pi}_k\widetilde{\mathbf{R}}_k\mathbf{\Pi}_k^H\right) \\ &= \mathbf{Q}_k\mathbf{\Sigma}_k\end{aligned}$$

where we exploit the fact that $\mathbf{\Pi}_k$ is unitary. We have defined

$$\begin{aligned}\mathbf{Q}_k &\triangleq \mathbf{\Pi}_k\widetilde{\mathbf{Q}}_k\mathbf{\Pi}_k^H \\ \mathbf{\Sigma}_k &\triangleq \mathbf{\Pi}_k\widetilde{\mathbf{R}}_k\mathbf{\Pi}_k^H\end{aligned} \tag{24}$$

Since $\mathbf{\Pi}_k$ is unitary $\mathbf{Q}_k\mathbf{Q}_k^H = \mathbf{I}_{N+1}$ and hence $\mathbf{Q}_k$ is unitary. Note that unlike $\widetilde{\mathbf{R}}_k$, $\mathbf{\Sigma}_k$ is not strictly diagonal dominant. Define the inverse permutation order

$$\{v_k(1), \ldots, v_k(N+1)\} \quad \text{s.t.}\, h_k^{(n,m)} = \widetilde{h}_k^{(v_k(n),v_k(m))} \tag{25}$$

Define $\rho_k^{(n,m)} \triangleq [\mathbf{\Sigma}_k]_{n,m}$. Compare (13) and (24). Thus (25) implies

$$\rho_k^{(n,m)} = \widetilde{r}_k^{(v_k(n),v_k(m))}$$

Using (19)

$$\begin{aligned}\left|\rho_k^{(n,m)}\right|^2 &= \left|\widetilde{r}_k^{(v_k(n),v_k(m))}\right|^2 \\ &\leq f_1(\alpha)\left|\widetilde{r}_k^{(v_k(m),v_k(m))}\right|^2, \forall v_k(n) \neq v_k(m)\end{aligned}$$

Thus

$$\left|\rho_k^{(n,m)}\right|^2 \leq f_1(\alpha)\left|\rho_k^{(m,m)}\right|^2, \forall n \neq m \tag{26}$$

For small α, $f_1(\alpha) \ll 1$ hence

$$\left|\rho_k^{(m,m)}\right| \gg \left|\rho_k^{(n,m)}\right|, \forall n \neq m \tag{27}$$

So *column-wise diagonal dominance in* $\mathbf{H}_k$ *implies column-wise diagonal dominance in* $\mathbf{\Sigma}_k$. Similarly using (23)

$$\begin{aligned}\left|\rho_k^{(n,m)}\right|^2 &= \left|\widetilde{r}_k^{(v_k(n),v_k(m))}\right|^2 \\ &\leq f_2(\alpha)\left|\widetilde{r}_k^{(v_k(n),v_k(n))}\right|^2, \forall v_k(n) \neq v_k(m)\end{aligned}$$

Thus

$$\left|\rho_k^{(n,m)}\right|^2 \leq f_2(\alpha)\left|\rho_k^{(n,n)}\right|^2, \forall n \neq m \tag{28}$$

For small α, $f_2(\alpha) \ll 1$ hence

$$\left|\rho_k^{(n,n)}\right| \gg \left|\rho_k^{(n,m)}\right|, \forall n \neq m \tag{29}$$

So *column-wise diagonal dominance in* $\mathbf{H}_k$ *implies row-wise diagonal dominance in* $\mathbf{\Sigma}_k$. Combining (27) and (27) leads to strict diagonal dominance in $\mathbf{\Sigma}_k$. Since the diagonal elements of $\widetilde{\mathbf{R}}_k$ are positive, the diagonal elements of $\mathbf{\Sigma}_k$ are also positive. Thus $\mathbf{H}_k$ can be decomposed

$$\mathbf{H}_k = \mathbf{Q}_k\mathbf{\Sigma}_k$$

where $\mathbf{Q}_k$ is unitary and $\mathbf{\Sigma}_k$ is strictly diagonal dominant with positive diagonal elements.

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] H. Dai and V. Poor, "Turbo multiuser detection for coded DMT VDSL systems," *IEEE J. Select. Areas Commun.*, vol. 20, no. 2, pp. 351–362, Feb. 2002.

[5] C. Zeng and J. Cioffi, "Crosstalk cancellation in ADSL systems," in *Proc. Global Commun. Conf., Globecom '01*, 2001, pp. 344–348.

[6] K. Cheong, W. Choi, and J. Cioffi, "Multiuser soft interference canceler via iterative decoding for DSL," *IEEE J. Select. Areas Commun.*, vol. 20, no. 2, pp. 363–371, Feb. 2002.

[7] *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.

[8] G. Forney and M. Eyuboglu, "Combined Equalization and Coding Using Precoding," *IEEE Commun. Mag.*, vol. 29, no. 12, pp. 25–34, Dec. 1991.

[9] R. Cendrillon and M. Moonen, "Simplified TX/RX Structure and Power Allocation for Co-ordinated DSL Systems," in preparation.

[10] 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.

[11] R. Cheng and S. Verdu, "Gaussian Multiaccess Channels with ISI: Capacity Region and Multiuser Water-Filling," *IEEE Trans. Inform. Theory*, vol. 39, no. 3, pp. 773–785, May 1993.

[12] 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.

[13] S. Chung, "Digital Transmission Techniques for Frequency Selective Gaussian Interference Channels," Ph.D. dissertation, Stanford University, 2003.

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

LIST OF FIGURES

1 FEXT Transfer Functions for 0.5 mm British Telecom Cable 20
2 Proportion of Crosstalk caused by i largest crosstalkers 21
3 Proportion of Crosstalk contained within i worst tones 21
4 Geometry of a 25-pair Bundle 21
5 Data-rate vs. Run-time Complexity (Equidistant Lines) 22
6 Near-end Data-rate vs. Run-time Complexity 22
7 Far-end Data-rate vs. Run-time Complexity 23
8 Achievable Rate Regions vs. Complexity (Simple Joint Selection Algorithm) 23
9 Achievable Rate Regions of Different Algorithms ($c = 2$) 24

LIST OF TABLES

I Initialization Complexity (Mults. and Log Operations) of Partial Crosstalk Cancellation Algorithms (per user) 20
II Initialization Complexity (Sort Operations) of Partial Cancellation Algorithms (per user) 20
III Simulation Parameters 20

TABLE I

INITIALIZATION COMPLEXITY (MULTS. AND LOG OPERATIONS) OF PARTIAL CROSSTALK CANCELLATION ALGORITHMS (PER USER)

Scheme	Initialization Complexity		$N = 7$, $K = 4096$		$N = 99$, $K = 4096$	
	Mults.	Logs	Mults.	Logs	Mults.	Logs
Line Selection Only	KN	0	$29{\times}10^3$	0	$0.4{\times}10^6$	0
Tone Selection Only	$K(N+5)$	0	$49{\times}10^3$	0	$0.4{\times}10^6$	0
Simple Joint Selection	$3K(N+1)$	0	$98{\times}10^3$	0	$1.2{\times}10^6$	0
Optimal Joint Selection	$K(0.5N^2+2.5N+3)$	$K(N+1)$	$184{\times}10^3$	$33{\times}10^3$	$21.1{\times}10^6$	$0.4{\times}10^6$

TABLE II

INITIALIZATION COMPLEXITY (SORT OPERATIONS) OF PARTIAL CANCELLATION ALGORITHMS (PER USER)

Scheme	Sort Operations		
	Sort Size N	Sort Size K	Sort Size KN
Line Selection Only	K	0	0
Tone Selection Only	0	1	0
Simple Joint Selection	0	0	1
Optimal Joint Selection	0	0	KN

TABLE III

SIMULATION PARAMETERS

Number of DMT Tones	4096
Tone Width	4.3125 kHz
Symbol Rate	4 kHz
Coding Gain	3 dB
Noise Margin	6 dB
Symbol Error Probability	$< 10^{-7}$
Transmit PSD	Flat -60 dBm/Hz
FDD Band Plan	998
Cable Type	0.5 mm (24-Gauge)
Source/Load Resistance	135 Ohm
Alien Crosstalk	ETSI Type A[7]

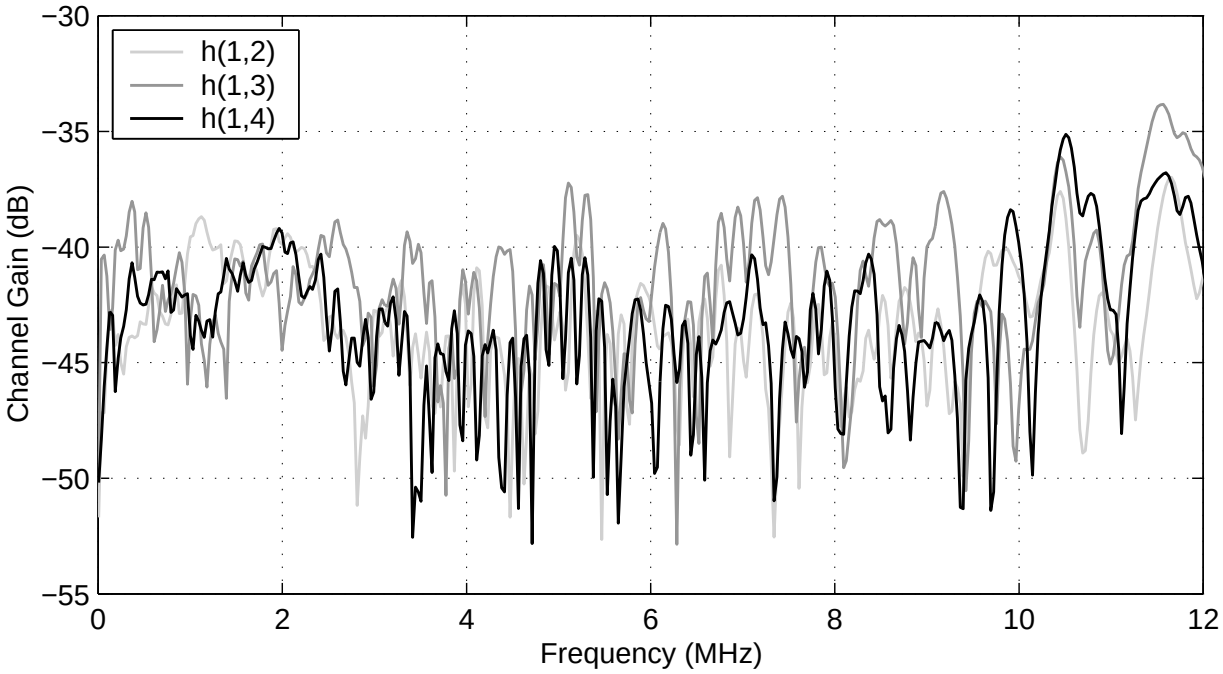


Fig. 1. FEXT Transfer Functions for 0.5 mm British Telecom Cable

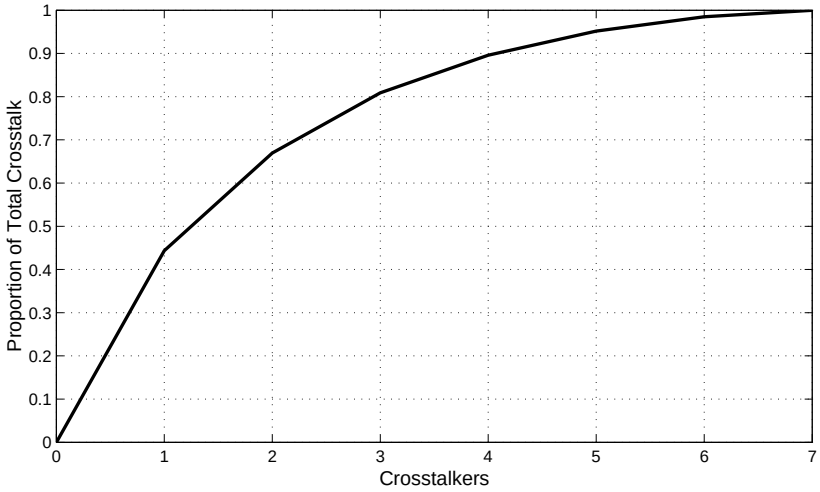


Fig. 2. Proportion of Crosstalk caused by i largest crosstalkers

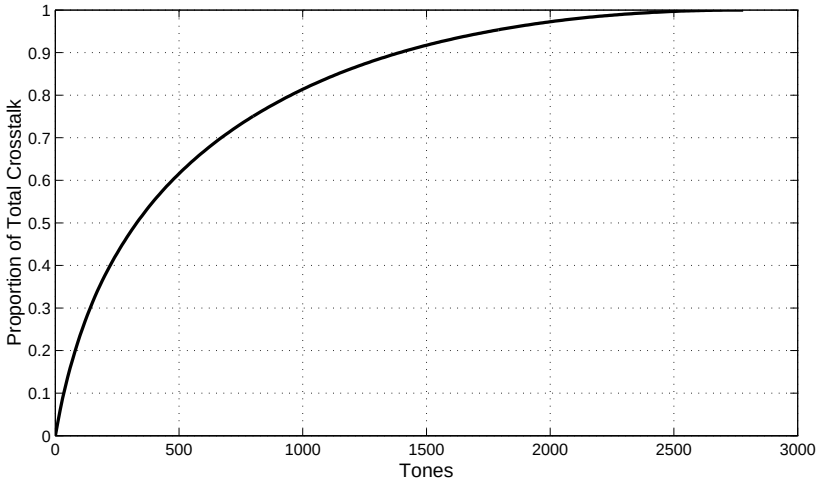


Fig. 3. Proportion of Crosstalk contained within i worst tones

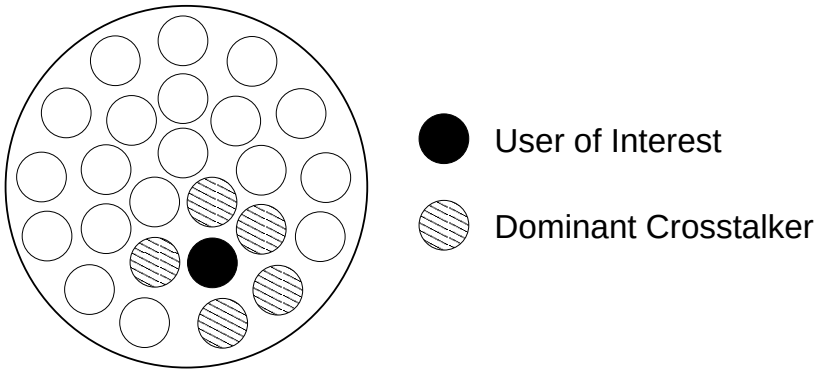


Fig. 4. Geometry of a 25-pair Bundle

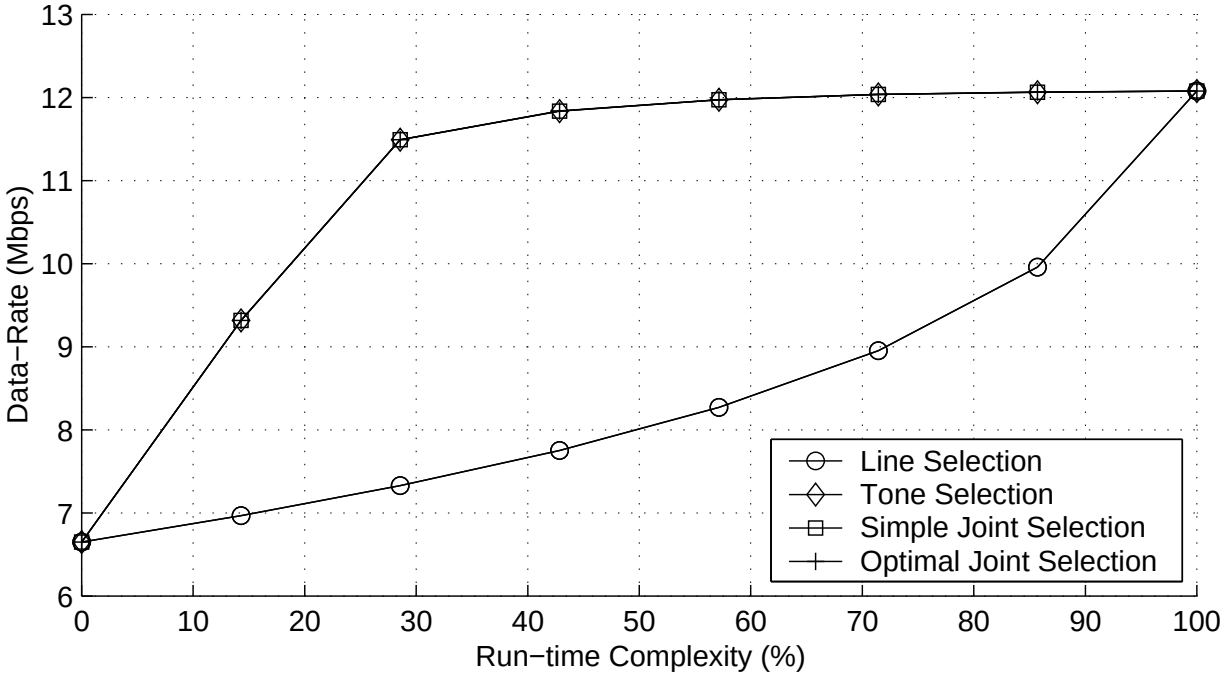


Fig. 5. Data-rate vs. Run-time Complexity (Equidistant Lines)

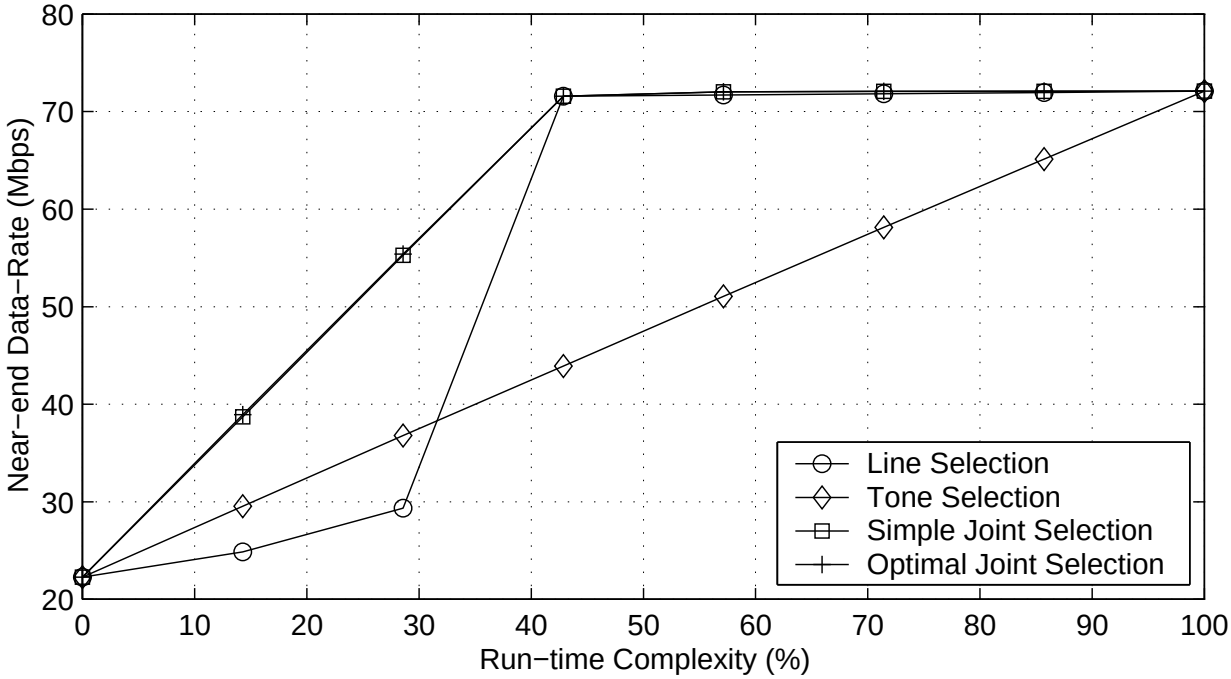


Fig. 6. Near-end Data-rate vs. Run-time Complexity

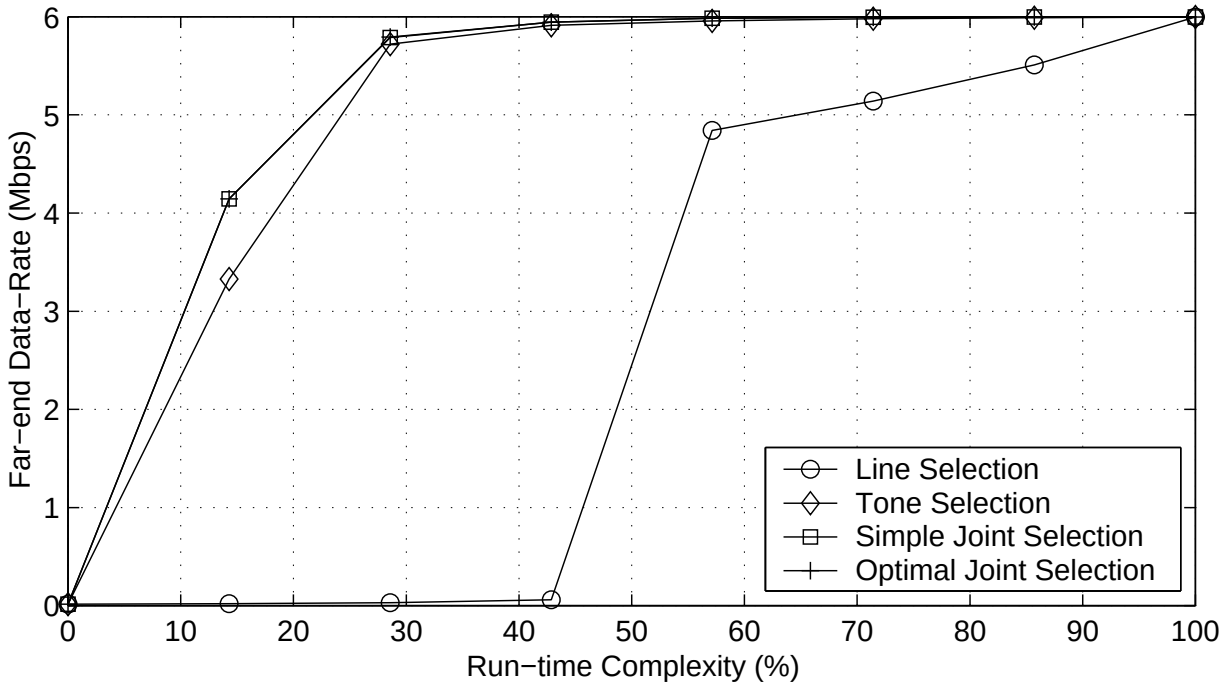


Fig. 7. Far-end Data-rate vs. Run-time Complexity

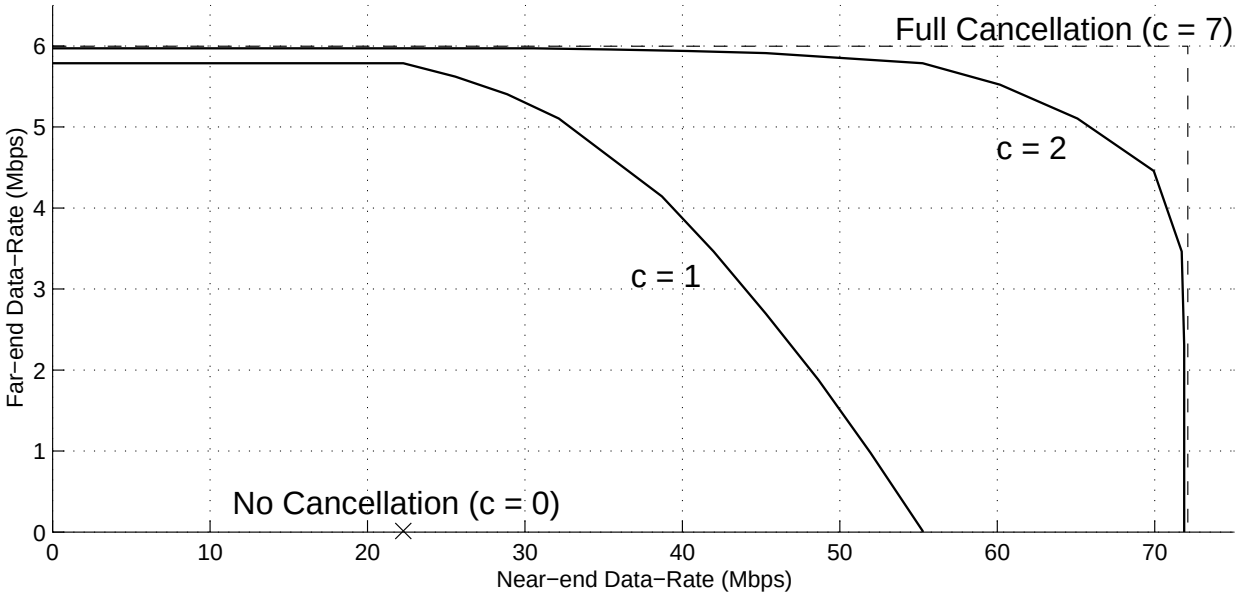


Fig. 8. Achievable Rate Regions vs. Complexity (Simple Joint Selection Algorithm)

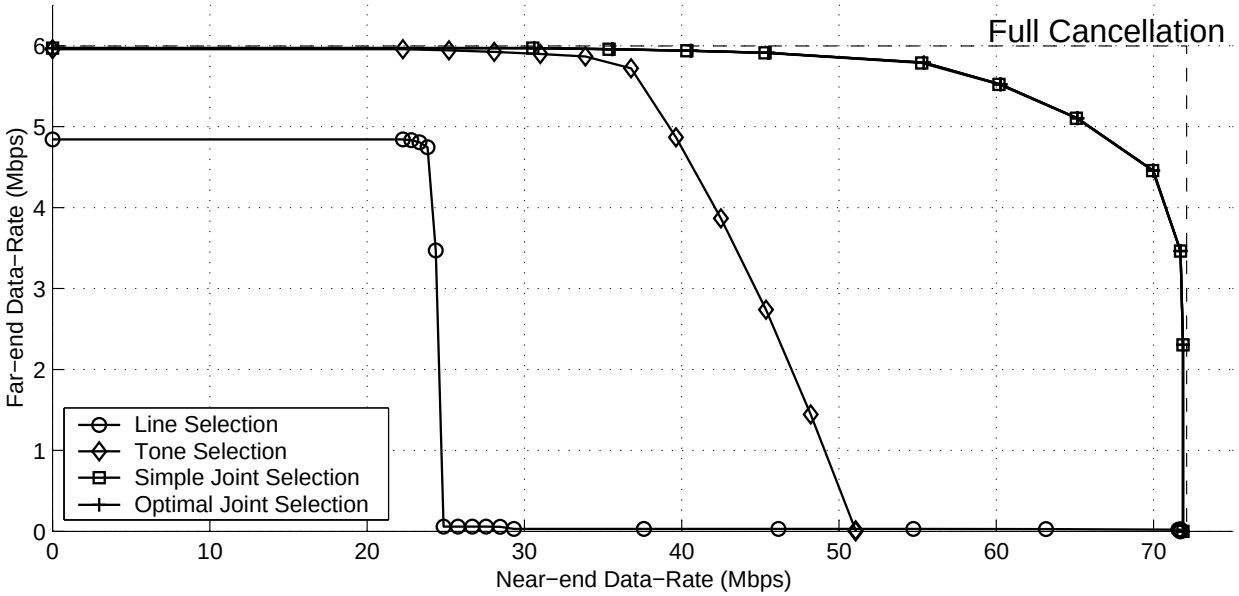


Fig. 9. Achievable Rate Regions of Different Algorithms ($c = 2$)