Partial Crosstalk Precompensation in Downstream VDSL ⋆

Raphael Cendrillon [a,*], George Ginis [b], Marc Moonen [a], Katleen Van Acker [c]

[a] *Katholieke Universiteit Leuven, Kasteelpark Arenberg 10, Heverlee 3001, Belgium*

[b] *Texas Instruments, 2043 Samaritan Drive, San Jose, CA 95124, USA*

[c] *Alcatel Bell, Francis Wellesplein 1, Antwerp 2018, Belgium*

Abstract

Very high bit-rate digital subscriber line (VDSL) is the latest generation in the ongoing evolution of DSL standards. VDSL aims at bringing truly broadband access, greater than 52 Mbps in the downstream, to the mass consumer market. This is achieved by transmitting in frequencies up to 12 MHz. Operating at such high frequencies gives rise to crosstalk between the DSL systems in a binder, limiting achievable data-rates. Crosstalk is typically 10-15 dB larger than other noise sources and is *the* primary limitation on performance in VDSL. In downstream transmission several crosstalk precompensation schemes have been proposed to address this issue. Whilst these schemes lead to large performance gains, they also have extremely high complexities, beyond the scope of current implementation.

In this paper we develop the concept of *partial crosstalk precompensation*. The majority of the crosstalk experienced in a DSL system comes from only a few other lines within the binder. Furthermore its effects are limited to a small subset of tones. Partial precompensation exploits this by limiting precompensation to the tones and lines where it gives maximum benefit. As a result, these schemes achieve the majority of the gains of full crosstalk precompensation at a fraction of the run-time complexity. In this paper we develop several partial precompensation schemes. We show that with only 20% of the run-time complexity of full precompensation it is possible to achieve 80% of the performance gains.

Key words:
bonding, crosstalk cancellation, crosstalk precompensation, crosstalk selectivity, reduced complexity, vectoring, very-high bit-rate digital subscriber line (VDSL)

1 Introduction

Very high bit-rate digital subscriber line (VDSL) is the latest generation in the ongoing evolution of DSL standards. VDSL aims at bringing truly broadband access, greater than 52 Mbps in the downstream, to the mass consumer market. In practice this is achieved by operating over short loop lengths, less than 1.2 km, and transmitting over a wide frequency range up to 12 MHz.

Unfortunately, operating at such high frequencies in a medium originally designed for voice-band transmission (< 4 kHz) leads to its own problems. The biggest of these is electromagnetic coupling which arises between nearby pairs within cable binders. This electromagnetic coupling gives rise to interference or *crosstalk* between the DSL systems in a binder, limiting achievable data-rates.

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 *frequency division duplexing* (FDD). FEXT on the other hand is still a major problem. FEXT is typically 10-15 dB larger than other noise sources and is *the* primary limitation on performance in VDSL.

In *upstream* (US) transmission FEXT can be removed through the use of crosstalk cancellation techniques[1,2,3]. These techniques typically rely on the fact that *receiving* (RX) modems are co-located at the CO. This allows reception to be done in a joint fashion, and crosstalk to be filtered out.

Unfortunately in *downstream* (DS) transmission crosstalk cancellation is not possible since the RX modems are located at different *customer premises* (CPs). However TX modems are co-located hence we can do *crosstalk precompensation*[4,5]. Effectively, crosstalk precompensation involves the distortion of each modem's signal prior to transmission. This distortion destructively interferes with the crosstalk introduced during transmission. As such the RX modems receive a crosstalk free

⋆ 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*; IWT SOLIDT Project, *Solutions for xDSL Interoperability, Deployment and New Technologies*; 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.

* Corresponding author.
Email addresses: cendrillon@ieee.org (Raphael Cendrillon), gginis@dsl.stanford.edu (George Ginis), moonen@esat.kuleuven.ac.be (Marc Moonen), katleen.van_acker@alcatel.be (Katleen Van Acker).

signal.

Whilst crosstalk precompensation leads to large performance gains, it also has a high complexity. For example, in a binder of 25 users linear crosstalk precompensation requires several billion multiplications/second. Non-linear schemes add even more complexity. This is beyond the scope of current DSL platforms.

To address this problem we develop the concept of *partial crosstalk precompensation*. We will show that by limiting precompensation to the largest crosstalkers, and the tones worst affected by crosstalk the majority of the data-rate gains of full precompensation can be achieved at a fraction of the run-time complexity.

Previous work investigated a similar concept namely *partial crosstalk cancellation*, which can only be applied in US transmission when RX modems are co-located[6]. Here we extend this work to the design of partial precompensators for DS transmission. Exploiting the properties of the DSL channel allows us to simplify the design of the crosstalk precompensator. We first propose a novel linear precompensator based on a channel diagonalization which achieves near-optimal performance. In contrast to previously proposed precompensators, e.g. [4], no modification of *CP equipment* (CPE) is required. Note that this is extremely desirable due to the difficulty of modifying the millions of CPEs already in place, which are all owned and operated by different customers.

The rest of the paper is organized as follows: In Sec. 2 we describe the system model for crosstalk coupling in downstream transmission. We introduce the concept of crosstalk precompensation in Sec. 3, and describe the *diagonalizing precompensator.* Due to the high complexity of full precompensation, we develop the partial precompensator in Sec. 4. We note that the majority of crosstalk typically comes from only a few other lines within the binder. So by only precompensating for these we can reduce run-time complexity substantially. Furthermore, crosstalk coupling varies dramatically with frequency, the worst effects being limited to a small subset of tones. We discuss this fact, and schemes which can exploit it in Sec. 5. As we show, achieving maximum reduction in run-time complexity requires us to exploit both the *space* and *frequency-selectivity* of the crosstalk channel. In Sec. 6 we consider the distribution of complexity between users. In Sec. 7 we evaluate the performance of the different partial precompensation schemes and finally conclusions are drawn in Sec. 8.

2 Downstream System Model

We assume that all transmitting modems are co-located at the CO/ONU. This is a prerequisite for crosstalk precompensation since transmission of all modems must be co-ordinated on a signal level. If transmission of the different modems is syn-

chronized, the cyclic structure of the DMT blocks allows us to model crosstalk coupling independently on each tone. We assume there are N users within the binder-group. Transmission of a single DMT block can be modeled as

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

The vector $\mathbf{x}_k \triangleq \left[x_k^1, \cdots, x_k^N\right]$ contains transmitted signals on tone k, where x_k^n denotes the signal transmitted onto line n at tone 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 is assumed to be spatially white and Gaussian such that $\mathcal{E}\left\{\mathbf{z}_k \mathbf{z}_k^H\right\} = \sigma_k^2 \mathbf{I}_N$. $k = 1, \ldots, K$ where K is the number of tones in the DMT system (e.g. for VDSL $K = 4096$). $\mathbf{H}_k$ is the $N \times N$ 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 $\mathbf{S}_k \triangleq \mathcal{E}\left\{\mathbf{x}_k \mathbf{x}_k^H\right\}$ with the transmit PSD of user n denoted as $s_k^n \triangleq [\mathbf{S}_k]_{n,n}$. As is common practice[7], we assume a spectral mask limits the maximum transmit PSD such that

$$s_k^n \leq s_k^{\max}, \ \forall\, n \tag{1}$$

In DSL channels with co-located TXs the channel matrix $\mathbf{H}_k$ is said to be *row-wise diagonal dominant* (RWDD) since it satisfies the following property

$$|h_k^{n,n}| \gg |h_k^{n,m}|\,, \forall n \neq m \tag{2}$$

In other words the direct channel of any user always has a larger gain than the crosstalk channel from any other user's TX into that user's RX. This is in contrast to the (similarly defined) *column-wise diagonal dominant* property which is seen in DSL channels with co-located *receivers* (RXs). These channels were studied in the context of crosstalk cancellation for upstream transmission in previous work[6].

We quantify the degree of row-wise diagonal dominance using the parameter α_k

$$|h_k^{n,m}| \leq \alpha_k \, |h_k^{n,n}|\,, \ \forall m \neq n \tag{3}$$

The row-wise diagonal dominance property has been verified through extensive cable measurements, see e.g. the semi-empirical crosstalk channel models in [8]. It will be exploited in the design of our crosstalk precompensators in the remaining sections.

3 Crosstalk Precompensation

3.1 Optimal Crosstalk Precompensation/Cancellation

It is instructive to first consider the case where *both* the TXs and RXs of the modems within a binder are co-located, such that both transmission and reception can be co-ordinated between modems on a signal level. This allows channel capacity to be achieved in a simple fashion[9]. 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 $\mathbf{U}_k$ and $\mathbf{V}_k$ are matrices containing the left and right singular vectors of $\mathbf{H}_k$ as columns. The diagonal matrix Λ_k contains the singular values $\Lambda_k \triangleq \text{diag}\{\lambda_k^1, \ldots, \lambda_k^N\}$. It is assumed that $\mathbf{H}_k$ is non-singular which is ensured by (2) provided that $h_k^{n,n} \neq 0,\ \forall n$.

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

$$\mathbf{P}_k = \mathbf{V}_k \tag{5}$$

such that $\mathbf{x}_k = \mathbf{P}_k \widetilde{\mathbf{x}}_k$. Note that since $\mathbf{V}_k$ is unitary, provided $s_k^n \leq s_k^{\text{mask}},\ \forall n$ then $\widetilde{s}_k^n \leq s_k^{\text{mask}},\ \forall n$ also. So the pre-filtering operation preserves compliance with the spectral masks (1).

At the RX we apply the filter

$$\mathbf{W}_k = \Lambda_k^{-1} \mathbf{U}_k^H \tag{6}$$

to generate our estimate of the transmitted symbol

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

where $\mathbf{e}_n \triangleq [\mathbf{I}_N]_{\text{col}\, n}$, $\mathbf{I}_N$ is the $N \times N$ 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 $\mathcal{E}\left\{|\widetilde{z}_k^n|^2\right\} = \sigma_k^2 \left(\lambda_k^n\right)^{-2}$. So the post-filtering operation removes crosstalk perfectly without causing noise enhancement.

Applying a conventional slicer to $\widehat{x}_k^n$ achieves the following rate for user n on tone

k

$$c_k^n = \log\left(1 + \frac{1}{\Gamma}\sigma_k^{-2}\left(\lambda_k^n\right)^2 \widetilde{s}_k^n\right) \tag{7}$$

Γ represents the SNR-gap to capacity and is a function of the target BER, coding gain and noise margin[10]. The maximum achievable rate of the multi-line DSL channel is

$$C = \sum_k \log\left|\mathbf{I}_N + \frac{1}{\Gamma}\sigma_k^{-2}\mathbf{H}_k\mathbf{S}_k\mathbf{H}_k^H\right| \tag{8}$$

It is straight-forward to show that $\sum_n \sum_k c_k^n = C$ which confirms that formulas (7) and (8) are consistent. 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 post-filter requires the receiving modems to be co-located. In downstream DSL this is typically not the case since receiving modems are located at different CPs.

3.2 Near-Optimal Diagonalizing Precompensation

We can exploit the row-wise diagonal dominance of $\mathbf{H}_k$ in the downstream transmission case to simplify the optimal crosstalk cancellation/precompensation structure considerably. As we will show, near-optimal rates can be achieved with co-ordination only on the TX (CO) side of the binder. Co-ordination between RXs (CPs) is not required.

The SVD of a RWDD matrix $\mathbf{H}_k$ (2) can be well approximated

$$\mathbf{U}_k \simeq \widetilde{\mathbf{U}}_k \triangleq \mathbf{I}_N \tag{9}$$

$$\Lambda_k \simeq \widetilde{\Lambda}_k \triangleq \mathrm{diag}\{|h_k^{1,1}|, \ldots, |h_k^{N,N}|\} \tag{10}$$

$$\mathbf{V}_k \simeq \widetilde{\mathbf{V}}_k \triangleq \mathbf{H}_k^H \mathrm{diag}\{|h_k^{1,1}|, \ldots, |h_k^{N,N}|\}^{-1} \tag{11}$$

by which we mean that $\mathbf{H}_k = \widetilde{\mathbf{U}}_k\widetilde{\Lambda}_k\widetilde{\mathbf{V}}_k^H$ with $\widetilde{\mathbf{U}}_k$ unitary, $\widetilde{\Lambda}_k$ diagonal and $\widetilde{\mathbf{V}}_k$ *approximately* unitary. The approximation becomes exact as $\alpha_k \to 0$.

Define

$$\mathbf{G}_k \triangleq \widetilde{\mathbf{V}}_k^H\widetilde{\mathbf{V}}_k$$

and $g_k^{n,m} \triangleq [\mathbf{G}_k]_{n,m}$. Now we can upper bound

$$\begin{aligned} g_k^{n,n} &= |h_k^{n,n}|^{-2}\textstyle\sum_m |h_k^{n,m}|^2 \\ &\leq 1 + (N-1)\alpha_k^2 \end{aligned}$$

where we use (3) in the second line. Similarly we can lower bound $g_k^{n,n} \geq 1$ so

$$1 \leq g_k^{n,n} \leq 1 + (N-1)\alpha_k^2 \tag{12}$$

Also

$$
\begin{aligned}
|g_k^{n,m}| &= \left| |h_k^{n,n}|^{-1} |h_k^{m,m}|^{-1} \sum_l h_k^{n,l} h_k^{m,l^*} \right| \\
&\leq |h_k^{n,n}|^{-1} |h_k^{m,m}|^{-1} \sum_l \left| h_k^{n,l} \right| \left| h_k^{m,l} \right| \\
&= |h_k^{m,n}| \, |h_k^{m,m}|^{-1} + |h_k^{n,m}| \, |h_k^{n,n}|^{-1} + \sum_{l \neq n,m} \left| h_k^{n,l} \right| |h_k^{n,n}|^{-1} \left| h_k^{m,l} \right| |h_k^{m,m}|^{-1} \\
&\leq 2\alpha_k + (N-2)\alpha_k^2 \qquad (13)
\end{aligned}
$$

where we use (3) in the last line. Combining (12) and (13) implies

$$\widetilde{\mathbf{V}}_{\mathbf{k}}^H \widetilde{\mathbf{V}}_k \simeq \mathbf{I}_N$$

with the approximation becoming exact as $\alpha_k \to 0$.

It is remarked that $\widetilde{\mathbf{U}}_k \triangleq \mathbf{H}_k \mathrm{diag}\{|h_k^{1,1}|, \ldots, |h_k^{N,N}|\}^{-1}$, $\widetilde{\Lambda}_k \triangleq \mathrm{diag}\{|h_k^{1,1}|, \ldots, |h_k^{N,N}|\}$, $\widetilde{\mathbf{V}}_k \triangleq \mathbf{I}_N$ does not promote a similar SVD approximation as e.g. for $\mathbf{F}_k \triangleq \widetilde{\mathbf{U}}_k^H \widetilde{\mathbf{U}}_k$ it is not possible to attain bounds similar to (12) and (13).

From this a near-optimal crosstalk cancellation/precompensation structure of section 3.1 can be postulated as

$$\mathbf{W}_k = \widetilde{\mathbf{U}}_k \; = \; \mathbf{I}_N \qquad (14)$$

$$\mathbf{P}_k = \widetilde{\mathbf{V}}_k^{-H} \; = \; \frac{1}{\beta_k} \mathbf{H}_k^{-1} \mathrm{diag}\{|h_k^{1,1}|, \ldots, |h_k^{N,N}|\} \qquad (15)$$

The term β_k ensures that compliance with the spectral masks (1) is maintained after application of the precompensator and is defined

$$\beta_k \triangleq \max_n \left\| \left[\mathbf{H}_k^{-1} \mathrm{diag}\{|h_k^{1,1}|, \ldots, |h_k^{N,N}|\} \right]_{\mathrm{row}n} \right\| \qquad (16)$$

Now $\mathbf{H}_k^{-1} \mathrm{diag}\{|h_k^{1,1}|, \ldots, |h_k^{N,N}|\} \simeq \mathbf{V}_k$ and $\mathbf{V}_k$ is unitary, so in practice $\beta_k \simeq 1$.

In a companion paper[11] it will be shown that the application of (14) - (16) instead of (5) - (6), incurs a capacity loss which is $\mathcal{O}\left(\log_2(1 + N\alpha_k^2) - \log_2(1 - N\alpha_k^2)\right)$. This is small in practice. In this paper we are going to verify the performance for the diagonalizing precompensator based on real-life scenarios and simulations.

Since $\mathbf{W}_k = \mathbf{I}_N$ RX post-filtering is not required. This is important since in downstream DSL receiving modems are not co-located. Furthermore it is observed that the optimal transmitter structure is well approximated by a *channel diagonalizing design*. That is, application of the precompensator (15) diagonalizes the channel matrix $\mathbf{H}_k$. Each modem observes their original direct channel (scaled by β_k^{-1})

with crosstalk perfectly removed. For this reason we term this the *diagonalizing precompensator.*

One of the major benefits of this novel design in addition to its simplicity, is that no modification of *CP equipment* (CPE) is required. This is in contrast to, for example, the Tomlinson-Harashima based crosstalk precompensators which require a modulo operation to be applied at the RX[4]. This is highly undesirable since there are already millions of CPEs in place all owned and operated by different customers. Replacing *CO equipment* (COE) is much easier since it is typically managed by a single operator. In addition, COE and CPE are typically manufactured by different hardware vendors, which makes joint design more difficult.

A drawback, however, of the diagonalizing precompensator is that it still has a high run-time complexity. Define $\mathbf{p}_{k,m} \triangleq [\mathbf{P}_k]_{\text{col}\,m}$. The transmitted vector can be written as

$$\begin{aligned}\mathbf{x}_k &= \mathbf{P}_k \widetilde{\mathbf{x}}_k \\ &= \textstyle\sum_m \mathbf{p}_{k,m} \widetilde{x}_k^m\end{aligned}$$

The term $\mathbf{p}_{k,m}\widetilde{x}_k^m$ corresponds to the contribution that user m makes to the transmitted vector. $\mathbf{p}_{k,m}$ is a length N vector so precompensating for the crosstalk of one user at one tone requires N multiplications per DMT block. Precompensation for N users on K tones at a block-rate b (DMT blocks/second) requires N^2Kb multiplications per second. Thus the complexity rapidly grows with the number of users in a binder. For example, in a 25 user system with 4096 tones and a block rate of 4000 (typical VDSL settings) the complexity is 10 billion multiplications per second. Certainly large performance gains can be achieved with crosstalk precompensation. However it can be extremely complex, certainly beyond the complexity available in current-day systems. This is the motivation behind partial crosstalk precompensation.

4 Partial Crosstalk Precompensation

In Fig. 2 we have plotted the crosstalk channels from a set of measurements of a 24 AWG cable. As can be seen, the severity of crosstalk varies significantly with both frequency and space. We make two observations:

First, since electromagnetic transmission follows a distance squared law, the majority of the crosstalk that a line experiences comes from the 4 or 5 surrounding lines within a binder. We refer to this as the *space selectivity* of crosstalk. This is illustrated in Fig. 5. To illustrate this further we evaluated the proportion of crosstalk caused by the i largest crosstalkers into user n on tone k. All users were set to have identical transmit PSDs, so crosstalker m on tone k is said to be larger than crosstalker m' if $|h_k^{n,m}| > |h_k^{n,m'}|$. We averaged this calculation across all users n

and all tones k. The result is shown in Fig. 3. As can be seen close to 80% of the crosstalk power is caused by the 3 largest crosstalkers.

Second, the crosstalk coupling varies significantly with frequency. Electromagnetic coupling increases with frequency and reflections within the binder can lead to nulls in the transfer function. We refer to this as the *frequency selectivity* of crosstalk. To illustrate this further we evaluated the proportion of crosstalk contained within the i strongest tones between TX n and RX m. Tone k is said to be stronger than tone k' if $|h_k^{n,m}| > |h_{k'}^{n,m}|$. We averaged this calculation across all TXs n and RXs m. The result is shown in Fig. 4. As can be seen almost 80% of the crosstalk power is contained within half of the tones.

As we saw crosstalk coupling varies significantly with space. This suggests that the majority of the benefits of crosstalk precompensation can be realized with a significant reduction in run-time complexity by only precompensating for the largest crosstalkers of each user on each tone. Furthermore, the effects of crosstalk vary significantly with frequency. We can vary the degree of crosstalk precompensation between none and full, or anything in between, to match the severity of crosstalk experienced on a tone. This can lead to even further reductions in run-time complexity.

4.1 Principle

We now describe the design of the partial crosstalk precompensator in more detail. From the perspective of RX n on tone k, define the indices of the crosstalkers sorted in order of crosstalk strength

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

Based on the potential benefit of removing crosstalk on tone k, user n decides how many crosstalkers should be *precompensated out* of its received signal. We denote the number of crosstalkers that user n would like removed from tone k as $r_{k,n}$ and the corresponding set of crosstalkers to be removed

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

We describe how to chose $r_{k,n}$ later in section 5. For now let us just assume that all users have defined the set of crosstalkers they would like precompensated out of their RX signals. Let us define the set of RXs who would like TX m to be

precompensated out of their received signals, on tone k

$$\mathbb{N}_k^m \triangleq \{n \,|m \in \mathbb{M}_k^n\} = \{n_{k,m}(1), \ldots, n_{k,m}(t_{k,n})\}$$

where $t_{k,n}$ is the number of such RXs. Our goal is to design a precompensation filter $\mathbf{P}_k$ that satisfies

$$[\mathbf{H}_k\mathbf{P}_k]_{n,m} = \begin{cases} |h_k^{n,n}| & n = m \\ 0 & n \in \mathbb{N}_k^m \end{cases}, \ \forall n, m \tag{17}$$

That is, all RXs should see their original direct channel, and TX m should cause no crosstalk to the RXs in the set $\mathbb{N}_k^m$. Equivalently, RX n should experience no crosstalk from TXs in the set $\mathbb{M}_k^n$. Since our goal is to reduce run-time complexity, $\mathbf{P}_k$ should also have a sparse structure

$$[\mathbf{P}_k]_{n,m} = 0, \ \forall\, n \notin \{\mathbb{N}_k^m, \, m\} \tag{18}$$

Define the reduced channel matrix for TX m

$$\overline{\mathbf{H}}_{k,m} \triangleq \begin{bmatrix} h_k^{(m,m)} & [\mathbf{H}_k]_{\text{row } m, \text{ cols } \mathbb{N}_k^m} \\ [\mathbf{H}_k]_{\text{rows } \mathbb{N}_k^m, \text{ col } m} & [\mathbf{H}_k]_{\text{rows } \mathbb{N}_k^m, \text{ cols } \mathbb{N}_k^m} \end{bmatrix} \tag{19}$$

Also define column m of the precompensation matrix $\mathbf{p}_{k,m} \triangleq [\mathbf{P}_k]_{\text{col } m}$ and its reduced version

$$\overline{\mathbf{p}}_{k,m} \triangleq [\mathbf{p}_{k,m}]_{\text{rows } \{m, \, \mathbb{N}_k^m\}} = \left[\overline{p}_{k,m}(1) \ \ldots \ \overline{p}_{k,m}(t_{k,n}+1)\right]^T$$

Combining (18) with the constraint (17) leads to

$$\overline{\mathbf{H}}_{k,m}\overline{\mathbf{p}}_{k,m} = |h_k^{m,m}|\, \overline{\mathbf{e}}_1$$

where $\overline{\mathbf{e}}_m$ is the mth column of the $(t_{k,n}+1) \times (t_{k,n}+1)$ identity matrix. Hence the constraints (17) and (18) can both be satisfied by choosing

$$\overline{\mathbf{p}}_{k,m} = |h_k^{m,m}| \left[\overline{\mathbf{H}}_{k,m}^{-1}\right]_{\text{col } 1} \tag{20}$$

and

$$\mathbf{p}_{k,m} = \begin{cases} \overline{p}_{k,m}(1) & n = m \\ \overline{p}_{k,m}(i+1) & n = n_{k,m}(i) \\ 0 & \text{otherwise} \end{cases} \tag{21}$$

So to summarize, our partial precompensator is designed as follows. First each RX selects the set of crosstalkers that they would like precompensated out of their received signal. Based on this each TX determines which RXs it must do precompensation for. We design the matrix $\mathbf{P}_k$ in a column-wise fashion. Each column

corresponds to a particular TX. For column m, we find the crosstalk coupling submatrix $\overline{\overline{\mathbf{H}}}_{k,m}$ which models the coupling between TX m and the RXs it must do precompensation for. Then (20) and (21) show how to design column m of $\mathbf{P}_k$. We must also apply a scaling β_k to ensure that the signal after precompensation does not violate the spectral mask constraints (1).

The RWDD of $\mathbf{H}_k$ ensures that for partial precompensators $\beta_k \simeq 1$. The proof is straightforward but rather lengthy so we exclude it here. It is based on the observation that from (19), RWDD in $\mathbf{H}_k$ ensures RWDD in $\overline{\overline{\mathbf{H}}}_{k,m}$.

Note that precompensation of the selected crosstalkers of RX n at tone k now requires only $r_{k,n}$ multiplications per DMT block in contrast to the N multiplications required for full crosstalk precompensation. This technique has many similarities to hybrid selection/combining from the wireless field[12]. There selection is also used between transmit and/or receive antennas to reduce run-time complexity and reduce the number of analog front-ends (AFE) required.

4.2 Achievable Data-rate

We first examine the case when all modems generate signals $\widetilde{x}_k^n$ (prior to precompensation) that have equal PSDs on tone k. Then

$$\mathcal{E}\left\{\widetilde{\mathbf{x}}_k \widetilde{\mathbf{x}}_k^H\right\} = \widetilde{s}_k \mathbf{I}_N$$

Now the use of β_k (16) ensures that application of the precompensation matrix will not increase the transmit powers

$$\begin{aligned} \mathcal{E}\left\{|x_k^n|^2\right\} &= \mathcal{E}\left\{[\mathbf{P}_k]_{\text{row } n}\, \widetilde{\mathbf{x}}_k \widetilde{\mathbf{x}}_k^H\, [\mathbf{P}_k]_{\text{row } n}^H\right\} \\ &= \widetilde{s}_k \left\|[\mathbf{P}_k]_{\text{row } n}\right\|^2 \\ &\leq \widetilde{s}_k \end{aligned} \tag{22}$$

where we use (16) in the last line. Furthermore, the fact that $\mathbf{P}_k$ is almost unitary ensures that the components of $\mathbf{x}_k$ will be approximately uncorrelated. Now the received signal on line n

$$\begin{aligned} y_k^n &= \mathbf{h}_{k,n}\mathbf{x}_k + z_k^n \\ &= \mathbf{h}_{k,n}\mathbf{P}_k\widetilde{\mathbf{x}}_k + z_k^n \\ &= \beta_k^{-1}\left|h_k^{n,n}\right| \widetilde{x}_k^n + \textstyle\sum_{m \notin \mathbb{M}_k^n} h_k^{n,m} x_k^m + z_k^n \end{aligned}$$

where $\mathbf{h}_{k,n} \triangleq [\mathbf{H}_k]_{\text{row } n}$ and we use (17) in the last line. The first term corresponds to the signal, the second to the non-precompensated crosstalk and the third to the

Algorithm 1 Line Selection Only

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

noise. So, using (22), the SINR on line n at tone k with partial pre-compensation is

$$SINR_k^n = \frac{\beta_k^{-2} \left|h_k^{n,n}\right|^2 \widetilde{s}_k}{\sum_{m \notin \mathbb{M}_k^n} \left|h_k^{n,m}\right|^2 \widetilde{s}_k + \sigma_k^2}$$

When the users adopt different transmit powers a bound can be made

$$SINR_k^n \geq \frac{\beta_k^{-2} \left|h_k^{n,n}\right|^2 \widetilde{s}_k^n}{\sum_{m \notin \mathbb{M}_k^n} \left|h_k^{n,m}\right|^2 s_k^{\text{mask}} + \sigma_k^2}$$

The resulting achievable data rate of user n on tone k is thus

$$b_{k,n} = \log_2 \left(1 + \frac{1}{\Gamma} SINR_k^n\right)$$

It is now clear how to design a partial crosstalk precompensator for a particular tone, assuming that the number of multiplications to be spent on each RX $r_{k,n}$ is specified. With the partial precompensator we precompensate only the largest crosstalkers out of each RX's signal. This is how we exploit the space-selectivity of crosstalk to reduce run-time complexity, as will be further detailed in the next section.

4.3 Line Selection

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

$$\sum_k r_{k,n} \leq cK,\ \forall n$$

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

The reduction in run-time complexity from this algorithm comes from space-selectivity only. Since the degree of partial precompensation 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 an algorithm which exploits both space and frequency-selectivity. The advantage of algorithm 1 is its simplicity. The algorithm requires only $O(KN)$ multiplications and K sorting operations of N values to initialize the partial crosstalk precompensator for one user. Here we define initialization complexity as the complexity

of determining $r_{k,n}$, $\forall k$. Initialization complexity does not include actual calculation of the crosstalk precompensation parameters $\mathbf{P}_k$ for each tone. This requires $\mathcal{O}(\sum_k (t_{k,n}+1)^3)$ multiplications for user n regardless of the partial precompensation algorithm employed. We assume that the direct and crosstalk channel gains $|h_k^{n,m}|^2$, $\forall n, m, k$ are available and do not need to be calculated.

The initialization complexities (in terms of multiplications and logarithm operations per user) of the different partial precompensation algorithms are listed in Tab. 1. The required number of sort operations is listed in Tab. 2. All algorithms have equal run-time complexity.

5 Complexity Distribution across Frequency

5.1 Tone Selection

In the previous section we presented Alg. 1 for partial crosstalk precompensation. 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 as well.

In low frequencies crosstalk coupling is minimal so we would expect minimal gains from precompensation. In high frequencies on the other hand crosstalk coupling can be severe. However in high frequencies the direct channel attenuation is often so large that the channel can support only minimal bitloading even in the absence of crosstalk. This limits the potential gains of crosstalk precompensation. The largest gains from crosstalk precompensation 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 $r_{k,n}$ largest crosstalkers are precompensated out of its received signal

$$b_{k,n}(r_{k,n}) \triangleq \log\left(1+\frac{1}{\Gamma}\frac{|h_k^{n,n}|^2\,\tilde{s}_k^n}{\sum_{i=r_{k,n}+1}^{N}\left|h_k^{n,m_{k,n}(i)}\right|^2\tilde{s}_k^{m_{k,n}(i)}+\sigma_k^2}\right) \tag{23}$$

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

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

and the indices of the tones ordered by this gain

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

Algorithm 2 Tone Selection Only

$$r_{k,n} = \begin{cases} N & k \in \{k_n(1), \ldots, k_n(cK/N)\} \\ 0 & \text{otherwise} \end{cases}$$

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 precompensation algorithm: Alg. 2. This algorithm simply employs full crosstalk precompensation for user n on the cK/N tones with the largest gain and no precompensation on all other tones. This leads to a run-time complexity of cK multiplications/DMT-block/user.

Note that in this algorithm $r_{k,n}$ is restricted to take only the values 0 or N. As a result it is not possible to only precompensate for 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.

5.2 Joint Line-Tone Selection

In Sec. 4.3 and 5.1 we described partial precompensation 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 precompensation done on each tone $r_{k,n}$ to match the potential gains. In practice this means that we allow $r_{k,n}$ to take on values other than 0 and N (unlike algorithm 2) whilst also allowing $r_{k,n}$ to vary from tone to tone (unlike algorithm 1).

At this point, it is interesting to evaluate the sub-optimality of the algorithms we described so far through comparison with a truly optimal partial precompensation algorithm. The problem of partial precompensation 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

$$\max_{\{r_{k,n}\}_{k=1,\ldots,K}} \sum_k b_{k,n}(r_{k,n}) \quad \text{s.t.} \sum_k r_{k,n} \leq cK$$

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. 3 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

Algorithm 3 Joint Line-Tone Selection

init $v_{k,n}(r) = \left(b_{k,n}(r) - b_{k,n}(0)\right)/r \quad \forall\, k,\, r > 0$

repeat

$(k_s, r_s) = \arg\max_{(k,r)} v_{k,n}(r)$

$r_{k_s,n} = r_s$

$v_{k_s,n}(r) = 0 \qquad r = 1, \ldots, r_s$

$v_{k_s,n}(r) = \left(b_{k_s,n}(r) - b_{k_s,n}(r_s)\right) / \left(r - r_s\right) \qquad r = r_s + 1, \ldots, N$

while $\sum_k r_{k,n} < cK$

practical problem. Define the value of precompensating for the r largest crosstalkers of user n on tone k as

$$v_{k,n}(r) = \frac{b_{k,n}(r) - b_{k,n}(0)}{r}$$

Recall that $b_{k,n}(r)$ is the rate achieved by user n on tone k when the r largest crosstalkers are precompensated and is evaluated using (23). 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 r multiplications are spent on tone k.

To find the optimal distribution of available complexity $[r_{1,n}, \ldots, r_{K,n}]$ for user n the algorithm begins by initializing $v_{k,n}(r)$ for all values of r and k. It then proceeds as follows

(1) Find choice of tone k and number of precompensated crosstalkers r with largest value $v_{k,n}(r)$. Store this in (k_s, r_s)
(2) Set the number of lines to be observed on tone k_s to r_s
(3) Set value of precompensating r_s or less crosstalkers on tone k_s to zero. This prevents re-selection of previously selected pairs.
(4) Update value of precompensating $r_s + 1$ or more crosstalkers on tone k_s. The rate increase and cost should be relative to the currently selected number of crosstalkers.

The algorithm iterates through steps 1-4 until the allocated complexity exceeds cK. It then takes the solution of the previous iteration. Since the algorithm allocates at most N multiplications in each iteration, the allocated complexity from the previous iteration will be at the least $cK - N$. With $K = 4096$ typically $cK \gg N$. Hence the difference between the desired run-time complexity cK and that of the solution provided by the algorithm is minimal.

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. 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. The algorithm can require up to KN sort operations which can have sizes as large as KN. Although this algorithm is more complex than the previous algorithms, this is not such an issue since the DSL channel is quite static in time. Due to this, updates to the crosstalk precompensator will only be required every few hours. To give a feeling for complexity, this algorithm typically takes less than 10 seconds to run on a standard PC.

6 Complexity Distribution between Users

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

$$\sum_k r_{k,n} \leq cK, \, \forall n$$

However since crosstalk precompensation of all lines in a binder is integrated into a single processing module at the CO, the multiplications can be shared between users. That is, the true constraint is on the total complexity of crosstalk precompensation for *all users*

$$\sum_n \sum_k r_{k,n} \leq cKN$$

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

$$\kappa_n = \mu_n cKN \quad \text{s.t.} \sum_n \mu_n = 1$$

Here μ_n is a parameter which determines the proportion of computing resources allocated to user n. This allows us to view partial precompensation 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. [13,14], 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 inside a binder.

As we will show, limiting crosstalk precompensation on each tone to the users who benefit the most leads to further reductions in run-time complexity with minimal performance loss.

7 Performance

We now compare the performance of the partial crosstalk precompensation algorithms described in sections 4.3, 5.1 and 5.2. As we show, the ability to exploit

both space and frequency-selectivity is essential for achieving the lowest possible run-time complexity.

We use a set of measured crosstalk channel transfer functions from a 0.5 mm (24 AWG) cable. This contains 8 pairs. The first 4 pairs are 900 m. long and whilst the last 4 are 1200m. The direct and crosstalk channels are depicted in Fig. 1 and Fig. 2 respectively. We assume there are $N = 8$ modems operating out of a common CO/*optical network unit* (ONU) as depicted in Fig. 6. Other simulation parameters are listed in Tab. 3.

We examine the distribution of run-time complexity between users as described in Sec. 6. Fig. 7 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 900 m. and 1200 m. users. Users of one class receive an equal number of multiplications; $2\mu cK$ and $2(1-\mu)cK$ multiplications per DMT-block for the 1200 m. and 900 m. users respectively. By varying the parameter μ we can trace out the boundary of the rate region. We see in Fig. 7 that with 30% of the run-time complexity ($c = 2.4$) of full crosstalk precompensation we can achieve the majority of the operating points within the rate region.

In Fig. 8 the achievable rate regions of the different partial precompensation algorithms are compared for 20% complexity ($c = 1.6$). Note the considerably larger rate region which is achieved by exploiting both space and frequency-selectivity in Alg. 3.

To give an example of the possible gains we consider the case when we have a desired service of 20 Mbps on the 1200 m. lines. Under this constraint Tab. 4 shows the rates that can be achieved on the 900 m. lines. The allocation of complexity between the users is shown. Also included is the rate gain as a proportion of the total possible rate gain that can be achieved with full crosstalk cancellation. So by definition the rate gain of no cancellation is 0%, and the rate gain of full cancellation is 100%.

Essentially we allocate just enough complexity to the 1200 m. lines such that they achieve 20 Mbps. This corresponds to finding the smallest possible μ, that still achieves the 1200 m. target rate. Once this is done, any left over complexity is allocated to the 900 m. lines. The better a partial precompensation algorithm is, the smaller the value of μ it will be able to reach whilst still achieving the 1200 m. target rate.

With tone selection we see that $\mu = 0.8$ is required achieve the target rate on the 1200 m. lines. This allocates 80% of the available complexity to the 1200 m. lines. With the remaining 20% the rate on the 900 m. lines can be increased to 26.4 Mbps. This corresponds to 23% of the achievable rate gain.

Using line selection gives better performance. Less complexity needs to be allo-

cated to the 1200 m. lines, and they achieve their target with $\mu = 0.4$. This leaves 60% of the available complexity to the 1200 m. lines, allowing them to achieve 42% of the potential gains.

Joint selection gives a much higher performance than either line or tone selection alone. The 1200 m. line target rate is achieved with only $\mu = 0.2$ and the 900 m. lines can increase their rates to 35.9 Mbps, which is 80% of the achievable gain. This underscores the importance of exploiting both space and frequency selectivity when designing partial precompensators.

So using joint selection we can achieve 70% and 80% of the achievable gains on the 1200 m. and 900 m. lines respectively. This is done with only 20% of the run-time complexity of full precompensation.

8 Conclusions

Crosstalk is *the* dominant source of performance degradation in modern DSL systems. In downstream transmission, several crosstalk precompensation schemes have been proposed to address this. Whilst the schemes lead to large performance gains, they have high run-time complexities, typically beyond the scope of implementation for current systems.

Crosstalk channels in DSL are space and frequency selective. That is, the majority of crosstalk comes from a few users and its effects are limited to a subset of tones. Partial precompensation exploits this by limiting precompensation to the tones and lines where it gives the maximum benefit. As a result, these schemes can achieve the majority of the gains of full crosstalk precompensation at a fraction of the run-time complexity.

In this paper we presented several crosstalk precompensation algorithms. *Line Selection* precompensates only the largest crosstalkers of each user. This allows it to exploit the space selectivity of crosstalk, however since the number of precompensated crosstalkers is the same on each tone, frequency selectivity cannot be exploited. *Tone Selection* runs full precompensation on the tones which benefit most, however since it is an 'all or nothing' approach it cannot exploit space selectivity by canceling just the largest crosstalkers. *Joint Line-Tone Selection* gives the best performance, limiting precompensation to the lines and tones which benefit most.

With the *Joint Line-Tone Selection* algorithm it is possible to achieve 80% of the performance gains of full crosstalk precompensation with only 20% of the run-time complexity.

We considered the allocation of run-time complexity between users. This allows

complexity to be distributed to the users who benefit most, leading to further reductions in complexity. In a similar fashion to the work done in multi-user power allocation[14,13] this led to the development of rate regions. However here we consider the allocation of computing resources rather than transmit power.

In this paper we have restricted our attention to VDSL however the partial precompensation technique can also be applied to other wireline systems such as other DSLs, and high-speed multi-pair Ethernets.

References

[1] G. Ginis, J. Cioffi, Vectored Transmission for Digital Subscriber Line Systems, IEEE J. Select. Areas Commun. 20 (5) (2002) 1085–1104.

[2] C. Zeng, J. Cioffi, Crosstalk cancellation in ADSL systems, in: Proc. IEEE Globecom, 2001, pp. 344–348.

[3] K. Cheong, W. Choi, J. Cioffi, Multiuser Soft Interference Canceler via Iterative Decoding for DSL Applications, IEEE J. Select. Areas Commun. 20 (2) (2002) 363–371.

[4] G. Ginis, J. Cioffi, A Multi-user Precoding Scheme achieving Crosstalk Cancellation with Application to DSL Systems, in: Proc. 34th Asilomar Conference, Pacific Grove, CA, 2000, pp. 1627–1631.

[5] W. Yu, Spatial Multiplex in Downlink Multiuser Multiple-Antenna Wireless Environments, in: Proc. of IEEE Global Communications Conf. (GLOBECOM), San Francisco, 2003.

[6] R. Cendrillon, M. Moonen, G. Ginis, K. Van Acker, T. Bostoen, P. Vandaele, Partial Crosstalk Cancellation Exploiting Line and Tone Selection in Upstream DMT-VDSL, accepted for publication in EURASIP Journal on Applied Signal Processing.

[7] Very-high bit-rate Digital Subscriber Lines (VDSL) Metallic Interface (2003).

[8] Transmission and Multiplexing (TM); Access transmission systems on metallic access cables; VDSL; Functional Requirements (2003).

[9] G. Taubock, W. Henkel, MIMO Systems in the Subscriber-Line Network, in: Proc. of the 5th Int. OFDM-Workshop, 2000, pp. 18.1–18.3.

[10] T. Starr, J. Cioffi, P. Silverman, Understanding Digital Subscriber Line Technology, Prentice Hall, 1999.

[11] R. Cendrillon, M. Moonen, Linear Crosstalk Precompensation and Cancellation with Near-optimal Performance, in preparation.

[12] D. Gore, A. Paulraj, Space-time block coding with optimal antenna selection, in: Proc. of the Int. Conf. on Acoustics, Speech and Sig. Processing (ICASSP), 2001, pp. 2441–2444.

[13] R. Cheng, S. Verdu, Gaussian Multiaccess Channels with ISI: Capacity Region and Multiuser Water-Filling, IEEE Trans. Inform. Theory 39 (3) (1993) 773–785.

[14] W. Yu, G. Ginis, J. Cioffi, Distributed Multiuser Power Control for Digital Subscriber Lines, IEEE J. Select. Areas Commun. 20 (5) (2002) 1105–1115.

Table 1
Initialization complexity - Multiplications and log operations per user

Scheme	Initialization Complexity		$N = 8$, $K = 4096$	
	Mults.	Logs	Mults.	Logs
Line Selection Only	KN	0	33×10^3	0
Tone Selection Only	$K(N+5)$	0	53×10^3	0
Joint Selection	$K(0.5N^2+2.5N+3)$	$K(N+1)$	225×10^3	37×10^3

Table 2
Initialization complexity - Sort operations 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
Joint Selection	0	0	$\leq KN$

Table 3
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[8]

Table 4
Achievable data-rates with different algorithms

Precompensation Technique	μ	Complexity			Rate Gain		Rate (Mbps)	
		Total	900m	1200m	900m	1200m	900m	1200m
None	-	0%	0%	0%	0%	0%	22.7	17.2
Tone Selection	0.8	20%	8%	32%	23%	70%	26.4	20.0
Line Selection	0.6	20%	16%	24%	41%	70%	29.5	20.0
Joint Selection	0.2	20%	32%	8%	80%	70%	35.9	20.0
Full	-	100%	100%	100%	100%	100%	39.1	21.2

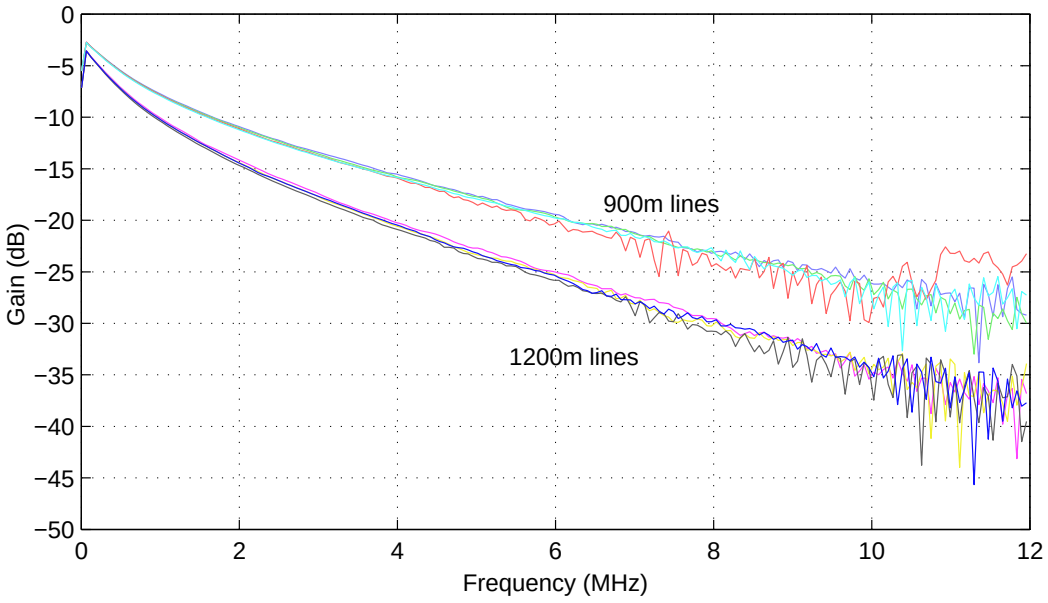


Figure 1. Direct channel transfer functions of measured binder

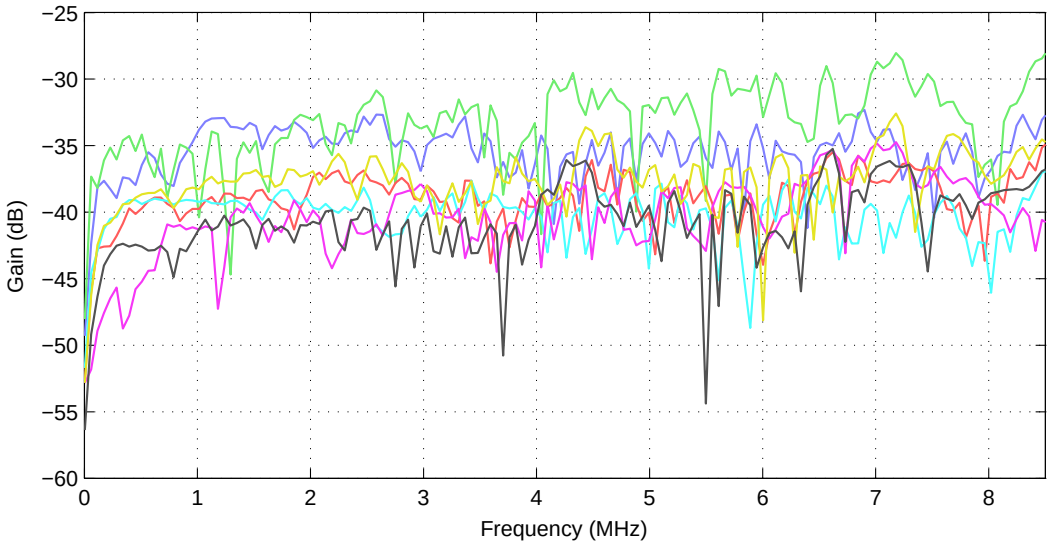


Figure 2. FEXT transfer functions of measured binder

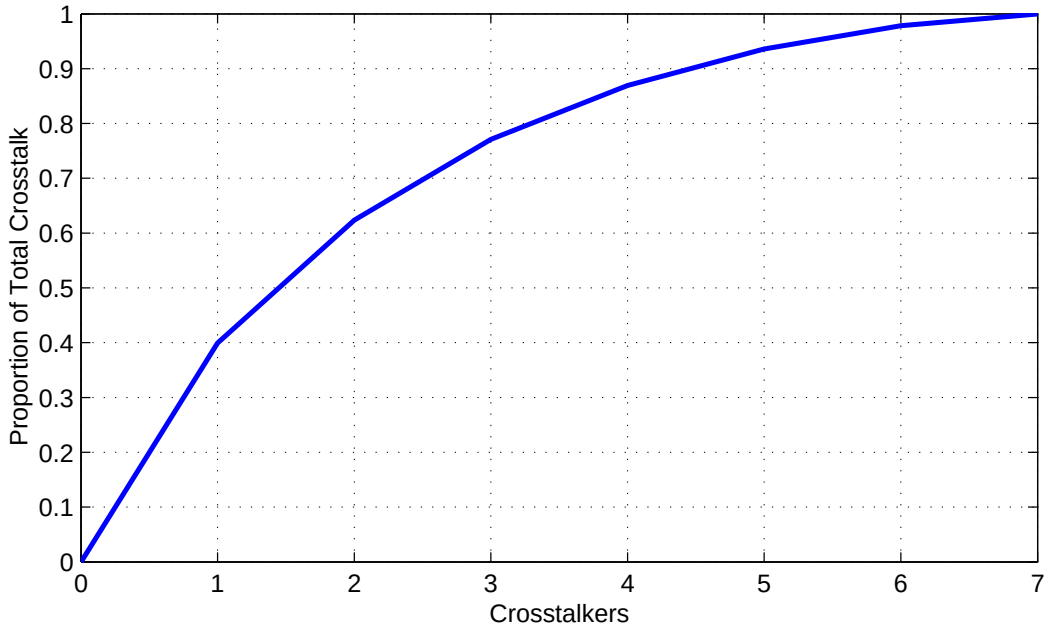


Figure 3. Proportion of crosstalk caused by largest crosstalkers

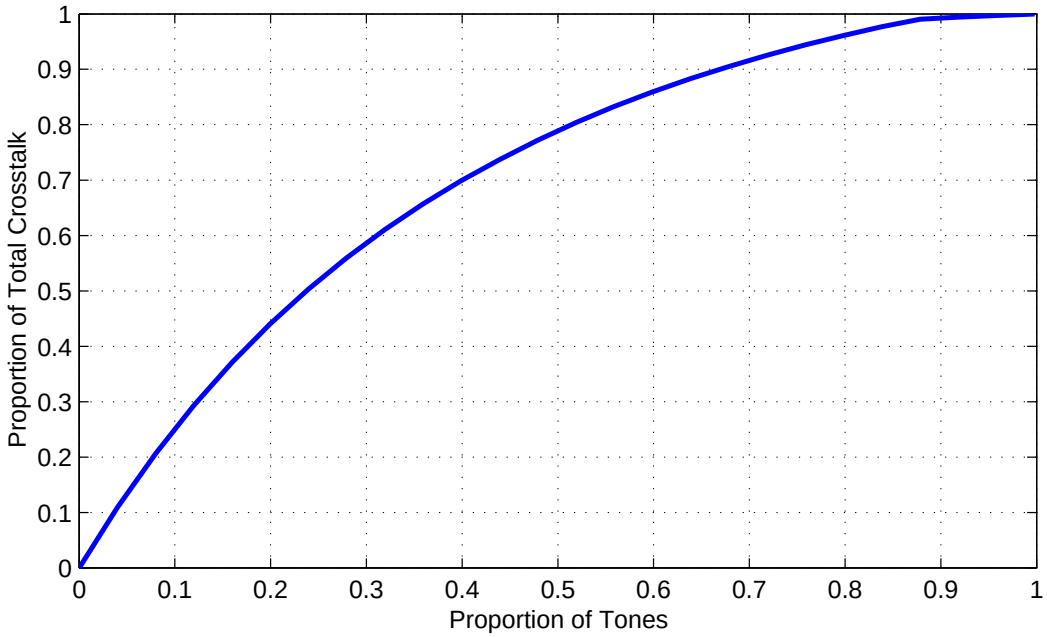


Figure 4. Proportion of crosstalk contained within worst tones

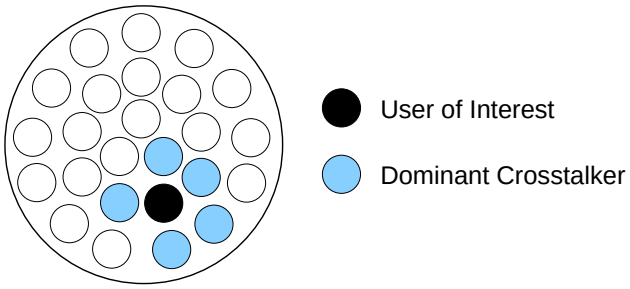


Figure 5. Geometry of a 25-pair binder

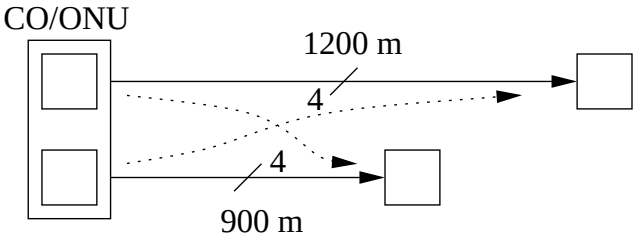


Figure 6. Downstream transmission

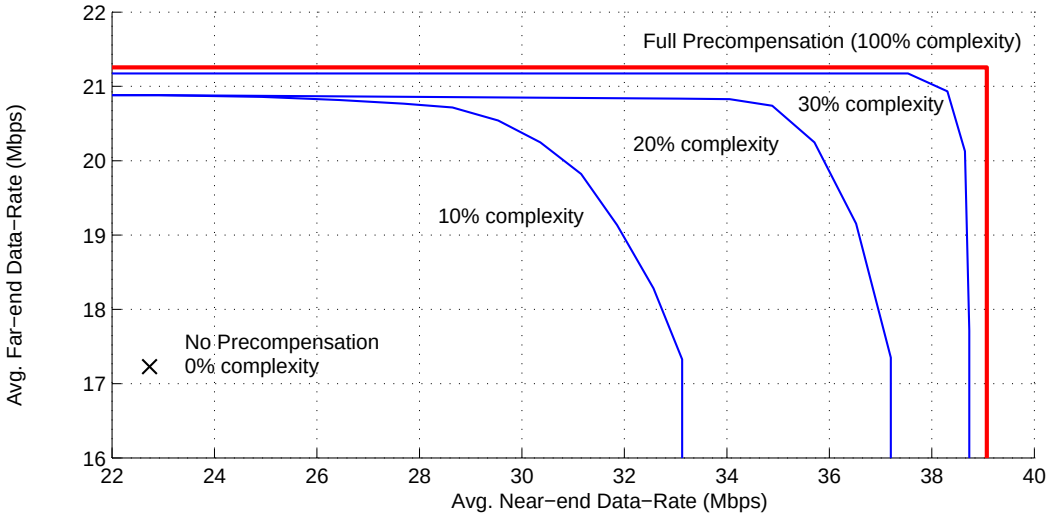


Figure 7. Achievable rate regions vs. complexity (Joint selection algorithm)

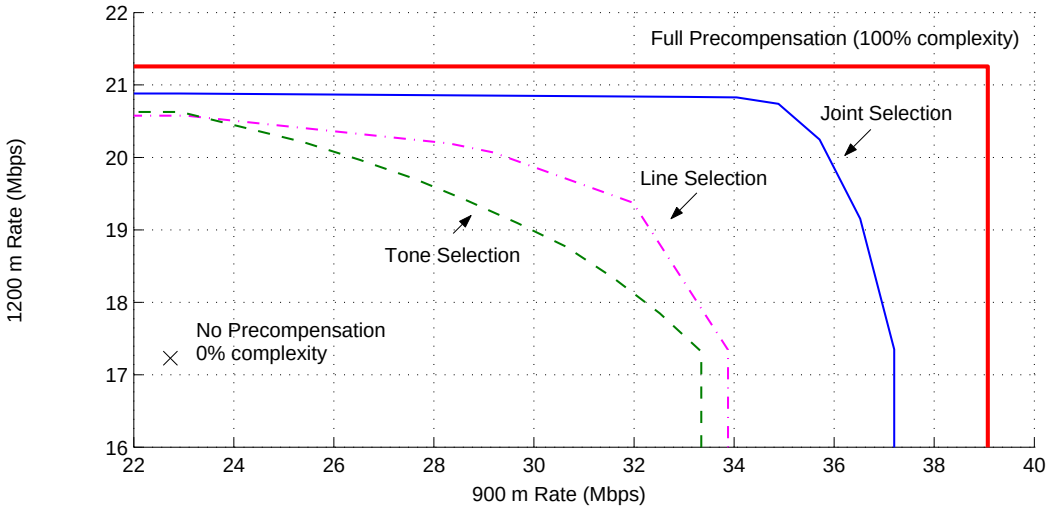


Figure 8. Achievable rate regions of different algorithms (20% complexity)