Multi-user Techniques and MIMO Transmission for VDSL

Raphael Cendrillon and Marc Moonen
SISTA Research Group
ESAT Department, K.U. Leuven
in co-operation with
Tom Bostoen, Piet Vandaele and Katleen Van Acker
Alcatel, Antwerp

April 1, 2002

Contents

1 Data Model **6**

1.1 Data Model . . . 6
1.2 Upstream Cancellation . . . 7
1.3 Downstream Pre-compensation . . . 7
1.4 Cyclic Prefix Duration . . . 8

2 Crosstalk Mitigation **9**

2.1 Crosstalk Cancellation (Upstream) . . . 9
2.1.1 Zero Forcing . . . 9
2.1.2 MMSE . . . 9
2.1.3 Vectored Receiver (QR) . . . 10
2.1.4 Generalized Decision Feedback Equalizer (GDFE) . . . 11
2.2 Crosstalk Pre-compensation (Downstream) . . . 13
2.2.1 Zero Forcing . . . 13
2.2.2 Vectored Transmitter . . . 14
2.3 Simulations . . . 15
2.3.1 Simulation Environment . . . 15
2.3.2 Theoretical Capacity Limit . . . 16
2.3.3 Pre and Post-Cancellation (Full Vectoring) . . . 17
2.3.4 Crosstalk Free . . . 18
2.3.5 Performance Results . . . 18
2.4 Conclusion . . . 19

3 Complexity Study **22**

3.1 Memory . . . 22
3.1.1 Crosstalk Reciprocity . . . 22
3.1.2 Tone Grouping . . . 23
3.2 Complexity . . . 26
3.2.1 Limited Frequency Band . . . 26
3.2.2 Dominant Interferers . . . 26
3.3 Conclusion . . . 30

4 Summary - Crosstalk Cancellation **32**

5 Near-Far Problem **34**

5.1 Conventional PBO . . . 35

5.1.1 Reference Frequency Method . . . 35

5.1.2 Reference Length Method . . . 35

5.1.3 Equalized FEXT Method . . . 35

5.1.4 Reference Noise Method . . . 36

5.1.5 Summary . . . 36

5.2 Rate Regions . . . 36

5.2.1 Static Power Allocation . . . 37

5.2.2 Relationship to Interference Cancellation . . . 37

5.2.3 Dynamic Power Allocation . . . 38

5.2.4 Summary . . . 40

5.3 Empirical Channel Models . . . 41

6 Power Allocation with Crosstalk Cancellation **42**

6.1 Optimal Power Allocation: Simultaneous Vector Waterfilling . . . 42

6.2 Optimal Power Allocation: Iterative Vector Waterfilling . . . 44

6.3 Optimum Receiver Structure: Generalized DFE . . . 45

6.4 Rate Region of the Optimal Scheme . . . 46

6.5 Rate Region of the Near-Optimal Scheme . . . 47

6.6 Summary . . . 49

7 Power Allocation without Crosstalk Cancellation **51**

7.1 Rate Region . . . 51

7.2 Iterative Scalar Waterfilling . . . 52

7.3 Noise-Only Scalar Waterfilling with PBO . . . 54

7.4 Summary . . . 54

8 Power Allocation and Partial Cancellation **56**

8.1 Resource Allocation . . . 56

8.2 (Tone, Interferer) Selection . . . 57

8.3 Crosstalk Canceller Design . . . 58

8.4 Performance . . . 58

8.5 Summary . . . 59

9 Future Work **61**

9.1 Downstream Power Back-off 61

9.2 Optimal Power Allocation without Crosstalk Cancellation 61

9.3 Optimal Partial Cancellation 62

9.4 Partial Unbundling and Spectral Masks 62

9.5 Flexible FDD Bandplans 62

9.6 Adaptive Issues 62

9.7 Bonding and the Multi-Access Channel 63

9.8 Crosstalk Cancellation in Unbundled and Asynchronous Environments 63

10 Summary - Multi-user Power Allocation **64**

Introduction

This report details the progress to date of our investigations into crosstalk cancellation techniques for VDSL. VDSL (Very high bit rate Digital Subscriber Line) is an emerging standard which will support high speed multimedia services over the existing telephone network. Frequencies up to 12 MHz are used to result in downstream data rates up to 52 Mbps depending on loop conditions.

Operating in such high frequency ranges introduces considerable crosstalk into the received signal both in the form of Near End Crosstalk (NEXT) and Far End Crosstalk (FEXT). Put simply, FEXT corresponds to electromagnetic coupling between two signals traveling in the same direction, whilst NEXT corresponds to coupling between two signals traveling in opposite directions.

Crosstalk pre-compensation and multi-dimensional crosstalk-cancellation are two novel techniques that have been proposed to reduce crosstalk in xDSL systems. These approaches are based on the MIMO (Multiple Input Multiple Output) system model. Instead of modulating/demodulating the signals on each twisted pair separately, MIMO algorithms rely on a joint processing of all signals leading to increased performance.

This report is organized as follows.

In Part 1 we investigate crosstalk cancellation techniques in VDSL. Chapter 1 describes our adopted data model which is based on a per-tone analysis. In Chapter 2 several crosstalk cancellation techniques are presented based on both linear and decision feedback operations. Here we see the considerable rate gains that are achievable through crosstalk cancellation. This chapter also describes the simulation environment and gives a performance comparison (based on simulation) of the different techniques. Chapter 3 discusses reduction of memory requirements and computational complexity through techniques such as tone grouping, limited frequency bands and partial crosstalk cancellation.

In Part 2 we investigate the near-far problem in VDSL and look at different methods for multi-user power allocation. In Chapter 5 we introduce the near-far problem and give an overview of conventional power back-off schemes. We also give an overview of the information theory of multi-access channels which forms a theoretical basis for further analysis. Chapter 6 investigates optimal power allocation when crosstalk cancellation is employed whilst in Chapter 7 we look at power allocation in the absence of crosstalk cancellation. In both chapters we present near-optimal schemes aswell as several low complexity, sub-optimal schemes.

In Chapter 8 we investigate the combination of partial crosstalk cancellation with multi-user power allocation. We develop a general partial cancellation scheme based on the selection of dominant (tone, interferer) pairs. This is effectively an extension to the simple techniques presented in Chapter 3. Finally Chapter 9 examines extensions to the work done to date and open questions.

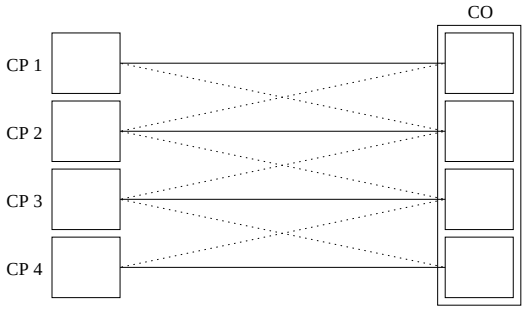


Figure 1: VDSL Environment

Chapter 1

Data Model

1.1 Data Model

The data model for our analysis assumes perfect synchronization at the central office (CO). Crosstalk cancellation in the asynchronous case will be the subject of future studies. As such, there is no inter-carrier interference (ICI) and we can model the system on a per-tone basis.

For one tone $k \in \{1 \dots K\}$ and N subscribers

$$\underbrace{\begin{bmatrix} y_1(k) \\ \vdots \\ y_N(k) \end{bmatrix}}_{\mathbf{y}(k)} = \underbrace{\begin{bmatrix} h_{11}(k) & \dots & h_{1N}(k) \\ \vdots & \vdots & \vdots \\ h_{N1}(k) & \dots & h_{NN}(k) \end{bmatrix}}_{\mathbf{H}(k)} \underbrace{\begin{bmatrix} x_1(k) \\ \vdots \\ x_N(k) \end{bmatrix}}_{\mathbf{x}(k)} + \underbrace{\begin{bmatrix} n_1(k) \\ \vdots \\ n_N(k) \end{bmatrix}}_{\mathbf{n}(k)}$$

where $x_i(k)$ is the DMT sub-symbol transmitted by user i on tone k (for a particular DMT symbol) and $y_i(k)$ is the corresponding received sub-symbol. $h_{ij}(k)$ represents the crosstalk coupling coefficient between lines i and j whilst $h_{ii}(k)$ represents the direct transfer function for line i. Note that $\begin{bmatrix} h_{ij}(1) & \dots & h_{ij}(K) \end{bmatrix}$ is related to the impulse response of the channel between lines i and j by the discrete Fourier transform (DFT). $n_i(k)$ is an additive noise term for user i and absorbs all alien crosstalk, thermal noise etc. which we do not care to cancel.

In terms of Power Spectral Densities (PSDs) (dropping k for clarity)

$$\underbrace{\begin{bmatrix} \sigma^2_{y_1} \\ \vdots \\ \sigma^2_{y_N} \end{bmatrix}}_{\mathbf{S}_y} = \underbrace{\begin{bmatrix} |h_{11}|^2 & \dots & |h_{1N}|^2 \\ \vdots & \vdots & \vdots \\ |h_{N1}|^2 & \dots & |h_{NN}|^2 \end{bmatrix}}_{|\mathbf{H}|^2_{\text{elem}}} \underbrace{\begin{bmatrix} \sigma^2_{x_1} \\ \vdots \\ \sigma^2_{x_N} \end{bmatrix}}_{\mathbf{S}_x} + \underbrace{\begin{bmatrix} \sigma^2_{n_1} \\ \vdots \\ \sigma^2_{n_N} \end{bmatrix}}_{\mathbf{S}_n}$$

Note that $|.|^2_{\text{elem}}$ represents an element-wise absolute value and squaring operation. We can split $\mathbf{S}_y$ into signal, interference and noise terms

$$\mathbf{S}_y = \mathbf{S}_{\text{signal}} + \mathbf{S}_{\text{interference}} + \mathbf{S}_{\text{noise}}$$

where

$$\mathbf{S}_{\text{signal}} = \underbrace{\begin{bmatrix} |h_{11}|^2 & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & |h_{NN}|^2 \end{bmatrix}}_{\text{diag}\{|\mathbf{H}|^2_{\text{elem}}\}} \mathbf{S}_x$$

$$\mathbf{S}_{\text{interference}} = \underbrace{\begin{bmatrix} 0 & |h_{12}|^2 & \dots & |h_{1N}|^2 \\ |h_{21}|^2 & 0 & \ddots & \vdots \\ \vdots & \ddots & 0 & \left|h_{(N-1),N}\right|^2 \\ |h_{N1}|^2 & \dots & \left|h_{N,(N-1)}\right|^2 & 0 \end{bmatrix}}_{\text{offdiag}\{|\mathbf{H}|^2_{\text{elem}}\}} \mathbf{S}_x$$

$$\mathbf{S}_{\text{noise}} = \mathbf{S}_n$$

Hence we can write the signal to interference plus noise ratio (SINR)

$$\mathbf{SINR} = \mathbf{S}_{\text{signal}} ./ \left(\mathbf{S}_{\text{interference}} + \mathbf{S}_{\text{noise}}\right)$$

where $./$ represents element-wise division of two vectors.

1.2 Upstream Cancellation

If we apply some linear operation to the received vector $\mathbf{y}$ to try and reduce interference[1]

$$\begin{aligned} \mathbf{z} &= \mathbf{W}\mathbf{y} \\ &= \widetilde{\mathbf{H}}\mathbf{x} + \mathbf{W}\mathbf{n} \end{aligned} \tag{1.1}$$

where $\widetilde{\mathbf{H}} = \mathbf{W}\mathbf{H}$. Again we can split the PSD of $\mathbf{y}$ into signal, interference and noise terms

$$\begin{aligned} \mathbf{S}_{\text{signal}} &= \text{diag}\left\{\left|\widetilde{\mathbf{H}}\right|^2_{\text{elem}}\right\}\mathbf{S}_x \\ \mathbf{S}_{\text{interference}} &= \text{offdiag}\left\{\left|\widetilde{\mathbf{H}}\right|^2_{\text{elem}}\right\}\mathbf{S}_x \\ \mathbf{S}_{\text{noise}} &= |\mathbf{W}|^2_{\text{elem}}\,\mathbf{S}_n \end{aligned}$$

1.3 Downstream Pre-compensation

In the downstream case we wish to apply some linear operation to the vector $\mathbf{x}$ prior to transmission. Ideally this pre-distortion should destructively interfere with the introduced interference. We can write this operation

$$\mathbf{y} = \mathbf{H}\mathbf{P}\mathbf{x} + \mathbf{n} \tag{1.2}$$

Again we can split the PSD of $\mathbf{y}$ into signal, interference and noise terms

$$\begin{aligned} \mathbf{S}_{\text{signal}} &= \text{diag}\left\{\left|\overline{\mathbf{H}}\right|^2_{\text{elem}}\right\}\mathbf{S}_x \\ \mathbf{S}_{\text{interference}} &= \text{offdiag}\left\{\left|\overline{\mathbf{H}}\right|^2_{\text{elem}}\right\}\mathbf{S}_x \\ \mathbf{S}_{\text{noise}} &= \mathbf{S}_n \end{aligned}$$

where $\overline{\mathbf{H}} = \mathbf{H}\mathbf{P}$.

[1] not all cancellation techniques use linear operations, however this analysis can be extended to the decision feedback techniques as shown in Chapter 2.

1.4 Cyclic Prefix Duration

Our assumption of no ICI only holds if all modems in a binder are synchronized, and if the cyclic prefix is longer than the largest channel impulse response. If this is not the case, we lose tone orthogonality and the techniques described here will fail.

Here we show that indeed, the cyclic prefix (CP) is of sufficient duration to prevent ICI.

$$\begin{aligned} \text{CP}_{\text{duration}} &= \text{CP}_{\text{length}}/\text{Sampling Rate} \\ &= 150/17.664 \times 10^6 \\ &= 15.625\mu s \end{aligned}$$

where $\text{CP}_{\text{length}}$ is the CP length in DMT sub-symbols.

Typically $\geq 99\%$ of the impulse response energy resides within $\text{CP}_{\text{duration}}$.

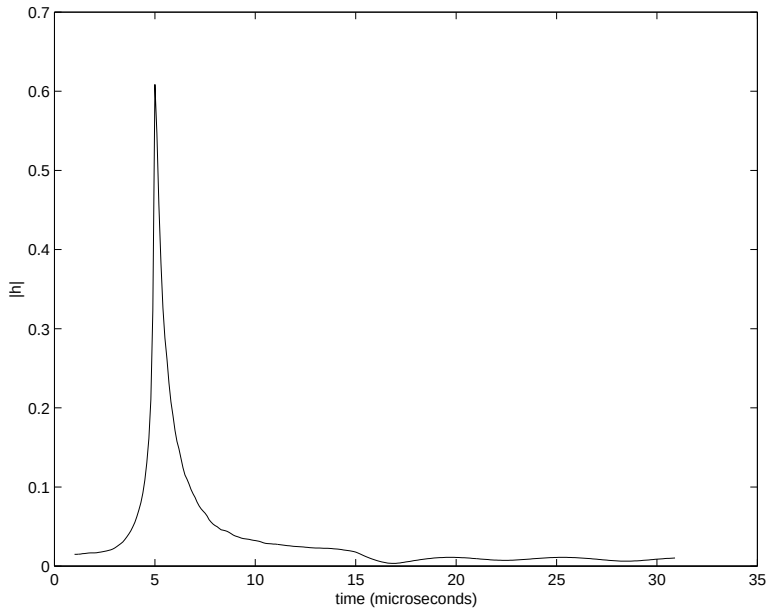


Figure 1.1: Example Channel Impulse Response

Chapter 2

Crosstalk Mitigation

In this chapter we present techniques for both crosstalk cancellation (upstream) and crosstalk pre-compensation (downstream). We also describe the simulation environment and performance of different crosstalk mitigation techniques.

2.1 Crosstalk Cancellation (Upstream)

2.1.1 Zero Forcing

Zero forcing (ZF) is the simplest solution and is optimal in the absence of noise since it removes all interference. $\mathbf{W}$ is calculated by simply inverting the channel

$$\mathbf{W}_{zf} = \mathbf{H}^{-1}$$

As a result

$$\mathbf{y} = \mathbf{x} + \mathbf{H}^{-1}\mathbf{n}$$

In terms of PSDs

$$\begin{aligned} \mathbf{S}_{\text{signal}} &= \mathbf{S}_x \\ \mathbf{S}_{\text{interference}} &= \mathbf{0} \\ \mathbf{S}_{\text{noise}} &= \left|\mathbf{H}^{-1}\right|^2_{\text{elem}} \mathbf{S}_n \end{aligned}$$

Zero forcing can suffer from significant performance degradation when $\mathbf{H}$ is badly conditioned. In this case multiplication of $\mathbf{y}$ with $\mathbf{H}^{-1}$ leads to significant noise enhancement. A good review of ZF and other linear multi-user detection techniques can be found in [1].

2.1.2 MMSE

Minimum Mean Squared Error (MMSE) crosstalk cancellers try to achieve some balance between crosstalk cancellation and noise filtering. They overcome the noise enhancement problems of ZF equalizers by focusing on minimizing the SINR rather than the SIR.

The cost function for an MMSE crosstalk canceller is

$$\begin{aligned} J &= \mathcal{E}\left\{\|\mathbf{x} - \mathbf{W}\mathbf{y}\|_2^2\right\} \\ \mathbf{W}_{\text{mmse}} &= \arg\min_{\mathbf{W}} J \\ &= \text{diag}\left\{\mathbf{S}_x\right\}\mathbf{H}^H\left(\mathbf{H}\text{diag}\left\{\mathbf{S}_x\right\}\mathbf{H}^H + \text{diag}\left\{\mathbf{S}_n\right\}\right)^{-1} \end{aligned}$$

where $(.)^H$ represents the complex conjugation and transpose operation. In terms of PSDs

$$\begin{aligned}
\mathbf{S}_{\text{signal}} &= \textbf{diag}\left\{\left|\widetilde{\mathbf{H}}\right|^2_{\text{elem}}\right\}\mathbf{S}_x \\
\mathbf{S}_{\text{interference}} &= \textbf{offdiag}\left\{\left|\widetilde{\mathbf{H}}\right|^2_{\text{elem}}\right\}\mathbf{S}_x \\
\mathbf{S}_{\text{noise}} &= |\mathbf{W}_{\text{mmse}}|^2_{\text{elem}}\,\mathbf{S}_n
\end{aligned}$$

where $\widetilde{\mathbf{H}} = \mathbf{W}_{\text{mmse}}\mathbf{H}$.

2.1.3 Vectored Receiver (QR)

Vectored techniques attempt to remove interference using decision feedback[2] and are similar to the classical decision feedback (DF) equalizers. Other interference cancellation techniques adopted in the wireless field like Successive Interference Cancellation (SIC) and BLAST[3] have similar structures.

Taking the QR decomposition of the channel matrix

$$\begin{aligned}
\mathbf{H} &= \mathbf{QR} \\
\mathbf{y} &= \mathbf{QRx} + \mathbf{n}
\end{aligned}$$

where $\mathbf{Q}$ is a unitary matrix and $\mathbf{R}$ upper triangular.

We first process the received vector $\mathbf{y}$ with a linear (unitary) transformation

$$\begin{aligned}
\mathbf{z} &= \mathbf{Q}^H\mathbf{y} \\
&= \mathbf{Rx} + \widetilde{\mathbf{n}}
\end{aligned}$$

If $\mathbf{n}$ has independent, identically distributed (i.i.d.) components then so does $\widetilde{\mathbf{n}}$. Writing $\mathbf{z}$ explicitly

$$\begin{bmatrix} z_1 \\ z_2 \\ \vdots \\ z_N \end{bmatrix} = \begin{bmatrix} r_{11} & r_{12} & \dots & r_{1N} \\ 0 & r_{22} & \dots & r_{2N} \\ 0 & 0 & \ddots & \vdots \\ 0 & 0 & 0 & r_{NN} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_N \end{bmatrix} + \widetilde{\mathbf{n}} \tag{2.1}$$

Examining (2.1) it becomes clear that we can extract $x_N \dots x_1$ through an iterative procedure consisting of hard decisions and back-substitution

$$\widehat{x}_i = \text{dec}\left[\tfrac{1}{r_{ii}} z_i - \textstyle\sum_{j=i+1}^{N} \tfrac{r_{ij}}{r_{ii}}\widehat{x}_j\right] \quad i = N,\, N-1,\, \dots,\, 1$$

We can also express this in vector notation. The structure of the vectored receiver is shown in figure 2.1. It consists of a feedfoward filter $\mathbf{F} = \text{diag}\left\{\mathbf{R}\right\}^{-1}\mathbf{Q}^H$ and a feedback filter $\mathbf{B} = \mathbf{I} - \text{diag}\left\{\mathbf{R}\right\}^{-1}\mathbf{R}$. The causal nature of the decision feedback operation requires that the feedback filter $\mathbf{B}$ be strictly upper triangular.

Notice that if you ignore the non-linear decision operation, then the feedback section effectively implements

$$(\mathbf{I} - \mathbf{B})^{-1} = \mathbf{R}^{-1}\text{diag}\left\{\mathbf{R}\right\}$$

This means that the entire vectored receiver implements

$$\begin{aligned}
(\mathbf{I} - \mathbf{B})^{-1}\mathbf{F} &= \mathbf{R}^{-1}\mathbf{Q}^H \\
&= \mathbf{H}^{-1}
\end{aligned}$$

which is simply the ZF crosstalk canceller. For this reason, we refer to the vectored receiver as the Zero Forcing-Decision Feedback Equalizer (ZF-DFE).

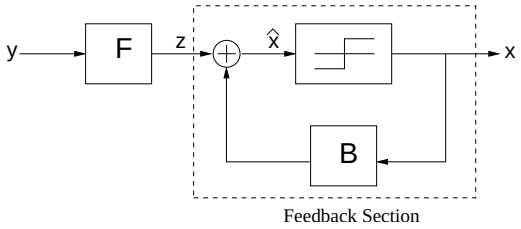


Figure 2.1: Decision Feedback Equalizer Structure

Assuming previous decisions are correct gives an (optimistic) estimate of SINR

$$\begin{aligned}
\mathbf{S}_{\text{signal}} &= \mathbf{diag}\left\{|\mathbf{R}|^2_{\text{elem}}\right\}\mathbf{S}_{\mathbf{x}} \\
\mathbf{S}_{\text{interference}} &= \mathbf{0} \\
\mathbf{S}_{\text{noise}} &= \mathbf{S}_{\tilde{\mathbf{n}}} = \mathbf{S}_{\mathbf{n}} \\
\mathbf{SINR}_{\text{vec rx}} &= \mathbf{diag}\left\{|\mathbf{R}|^2_{\text{elem}}\right\}(\mathbf{S}_x ./ \mathbf{S}_n)
\end{aligned}$$

One potential issue in the use of vectoring is the error propagation. This can be reduced through the use of Soft Interference Cancellation techniques [4][5]. These techniques also look promising in the asynchronous scenario.

2.1.4 Generalized Decision Feedback Equalizer (GDFE)

The generalized decision feedback equalizer (GDFE) is another receiver structure which attempts to remove interference through decision feedback. If we ignore the non-linear operations, the GDFE implements an MMSE crosstalk canceller, in contrast to the Vectored Receiver which implements a ZF crosstalk canceller[6].

One nice aspect of the GDFE is that it can achieve the full capacity of the multi-user channel and we show how this is achieved further on. To begin with, we describe the structure of the GDFE in more detail.

We begin by assuming that the noise power does not vary between users and only changes with frequency. Hence

$$\mathbf{S}_n = \sigma_n^2 \mathbf{1}_{N\times 1}$$

Under this assumption, we can rewrite the expression for the MMSE crosstalk canceller as follows

$$\begin{aligned}
\mathbf{W}_{\text{mmse}} &= \text{diag}\left\{\mathbf{S}_x\right\}\mathbf{H}^H\left(\mathbf{H}\text{diag}\left\{\mathbf{S}_x\right\}\mathbf{H}^H + \sigma_n^2\mathbf{I}\right)^{-1} \\
&= \left(\mathbf{H}^H\mathbf{H} + \sigma_n^2\text{diag}\left\{\mathbf{S}_x\right\}^{-1}\right)^{-1}\mathbf{H}^H
\end{aligned} \tag{2.2}$$

We now break the MMSE crosstalk canceller into two parts. The first part is simply the matched filter $\mathbf{H}^H$. The second part contains the rest of the filter and we denote this as $\mathbf{R}_b$

$$\mathbf{R}_b \triangleq \left(\mathbf{H}^H\mathbf{H} + \sigma_n^2\text{diag}\left\{\mathbf{S}_x\right\}^{-1}\right)^{-1}$$

$\mathbf{R}_b$ can be further decomposed using a Cholesky decomposition. Let

$$\mathbf{U}^{-1}\mathbf{D}^{-1}\mathbf{U}^{-H} \overset{\text{chol}}{=} \mathbf{R}_b$$

such that $\mathbf{U}$ is an upper triangular matrix with 1's on it's main diagonal, and $\mathbf{D}$ is a diagonal matrix. With these definitions, we can define the GDFE. Like the vectored receiver, it has the feedback structure as shown in figure 2.1. However now the feedforward filter $\mathbf{F} = \mathbf{D}^{-1}\mathbf{U}^{-H}\mathbf{H}^H$, whilst the feedback filter $\mathbf{B} = \mathbf{I} - \mathbf{U}$.

Notice first of all that if we ignore the non-linear decision operation, that the feedback section of the GDFE is equivalent to multiplication with the matrix $\mathbf{U}^{-1}$. This means that the entire GDFE implements

$$\begin{aligned}\mathbf{U}^{-1}\mathbf{D}^{-1}\mathbf{U}^{-H}\mathbf{H}^{H} &= \mathbf{R}_b\mathbf{H}^{H} \\ &= \mathbf{W}_{\text{mmse}}\end{aligned}$$

This is simply the MMSE crosstalk canceller as defined in (2.2).

We now examine the GDFE structure and show that it achieves the full capacity of the multi-user channel. This derivation comes directly from [6]. Define the variance of $\mathbf{x}$ as $\sum_{xx} \triangleq \mathcal{E}\left\{\mathbf{x}\mathbf{x}^H\right\}$ and the variance of $\mathbf{y}$ as $\Sigma_{yy} \triangleq \mathcal{E}\left\{\mathbf{y}\mathbf{y}^H\right\}$. Note that $\Sigma_{xx} = \text{diag}\left\{\mathbf{S}_x\right\}$. Furthermore, we also define the variance of $\mathbf{x}$ when $\mathbf{y}$ is given $\Sigma_{x|y} \triangleq \mathcal{E}\left\{\mathbf{x}\mathbf{x}^H|\mathbf{y}\right\}$.

$\Sigma_{x|y}$ can be evaluated as follows. We first define the optimal linear estimate of $\mathbf{x}$ given $\mathbf{y}$ using the Weiner equation

$$\begin{aligned}\overline{\mathbf{x}} &= \Sigma_{xy}\Sigma_{yy}^{-1}\mathbf{y} \\ &= \mathbf{W}_{\text{mmse}}\mathbf{y}\end{aligned}$$

The estimate $\overline{\mathbf{x}}$ will differ from the true $\mathbf{x}$, and we call this difference the estimation error $\mathbf{e}$. Hence

$$\mathbf{x} = \overline{\mathbf{x}} + \mathbf{e}$$

We are now in a position to write $\mathbf{x}$ as a function of $\mathbf{y}$

$$\mathbf{x} = \Sigma_{xy}\Sigma_{yy}^{-1}\mathbf{y} + \mathbf{e}$$

If $\mathbf{y}$ is given, then the variance of $\mathbf{x}$ is simply the variance of $\mathbf{e}$. Hence

$$\begin{aligned}\Sigma_{x|y} &= \mathcal{E}\left\{\mathbf{e}\mathbf{e}^H\right\} \\ &= \mathcal{E}\left\{\left(\mathbf{x} - \Sigma_{xy}\Sigma_{yy}^{-1}\mathbf{y}\right)\left(\mathbf{x} - \Sigma_{xy}\Sigma_{yy}^{-1}\mathbf{y}\right)^H\right\} \\ &= \Sigma_{xx} - \Sigma_{xx}\mathbf{H}^H\left(\mathbf{H}\Sigma_{xx}\mathbf{H}^H + \sigma_n^2\mathbf{I}\right)^{-1}\mathbf{H}\Sigma_{xx} \\ &= \mathbf{A}^{-1} - \mathbf{A}^{-1}\mathbf{B}\left(\mathbf{D}\mathbf{A}^{-1}\mathbf{B} + \mathbf{C}^{-1}\right)^{-1}\mathbf{D}\mathbf{A}^{-1} \\ &= (\mathbf{A} + \mathbf{B}\mathbf{C}\mathbf{D})^{-1} \\ &= \left(\Sigma_{xx}^{-1} + \sigma_n^{-2}\mathbf{H}^H\mathbf{H}\right)^{-1} \\ &= \mathbf{R}_b\sigma_n^2\end{aligned}$$

where we use the matrix inversion lemma to get from line 4 to 5.

With these definitions, let us examine the capacity of the multi-user channel. This can be written in two ways $I(\mathbf{x};\mathbf{y})$ and $I(\mathbf{y};\mathbf{x})$ and both are equivalent. Hence

$$\begin{aligned}C &= I(\mathbf{x};\mathbf{y}) \\ &= H(\mathbf{x}) - H(\mathbf{x}|\mathbf{y}) \\ &= \frac{1}{2}\log\frac{|\Sigma_{xx}|}{|\Sigma_{x|y}|}\end{aligned}$$

Now notice that $\mathbf{z}$ from figure 2.1 can be written

$$\begin{aligned}\mathbf{z} &= \mathbf{F}\mathbf{y} \\ &= \mathbf{D}^{-1}\mathbf{U}^{-H}\mathbf{H}^H\mathbf{y} \\ &= \mathbf{U}\mathbf{U}^{-1}\mathbf{D}^{-1}\mathbf{U}^{-H}\mathbf{H}^H\mathbf{y} \\ &= \mathbf{U}\mathbf{W}_{\text{mmse}}\mathbf{y} \\ &= \mathbf{U}\mathbf{x} - \mathbf{U}\mathbf{e}\end{aligned}$$

Furthermore, the signal fed into the decision device

$$\begin{aligned}\widehat{\mathbf{x}} &= \mathbf{z} + (\mathbf{I} - \mathbf{U})\,\mathbf{x} \\ &= \mathbf{x} - \mathbf{U}\mathbf{e}\end{aligned}$$

So the decision device detects each user in the presence of the noise term $\mathbf{U}\mathbf{e}$. Examining the variance of the noise term $\mathbf{U}\mathbf{e}$ we find

$$\begin{aligned}\mathcal{E}\left\{(\mathbf{U}\mathbf{e})\,(\mathbf{U}\mathbf{e})^H\right\} &= \sigma_n^2 \mathbf{U}\mathbf{R}_b\mathbf{U}^H \\ &= \sigma_n^2 \mathbf{D}^{-1}\end{aligned}$$

Notice that $\mathbf{D}$ is diagonal, hence the noise experienced by each user is decorellated. This allows us to do detection of each user independently, reducing the need for a highly complex ML detector. With independent detection, the rate achieved by each user is now

$$\begin{aligned}c_i &= I(x_i;\widehat{x}_i) \\ &= H(x_i) - H(x_i|\widehat{x}_i) \\ &= H(x_i) - H(\mathbf{U}\mathbf{e}) \\ &= \frac{1}{2}\log\left(\sigma_{x_i}^2\right) - \frac{1}{2}\log\left(\sigma_n^2 d_i^{-1}\right) \\ &= \frac{1}{2}\log\left(d_i\frac{\sigma_{x_i}^2}{\sigma_n^2}\right)\end{aligned}$$

and the total rate of the system

$$C_{\text{gdfe}} = \frac{1}{2}\log\left(\sigma_n^{-2}\left|\mathbf{D}\Sigma_{xx}\right|\right)$$

Note that since $\mathbf{U}$ is upper triangular

$$|\mathbf{U}| = \prod_i u_{ii}$$

However $u_{ii} = 1\,\forall i$ since $\mathbf{U}$ has 1's on the main diagonal. Hence $|\mathbf{U}| = 1$. Using this observation we proceed

$$\begin{aligned}\sigma_n^{-2}\left|\mathbf{D}\right| &= \sigma_n^{-2}\left|\mathbf{U}^H\right|\left|\mathbf{D}\right|\left|\mathbf{U}\right| \\ &= \sigma_n^{-2}\left|\mathbf{R}_b\right|^{-1} \\ &= \left|\Sigma_{x|y}\right|^{-1}\end{aligned}$$

Hence

$$\begin{aligned}C_{\text{gdfe}} &= \frac{1}{2}\log\frac{\left|\Sigma_{xx}\right|}{\left|\Sigma_{x|y}\right|} \\ &= C\end{aligned}$$

and we see that the GDFE achieves the full capacity of the multi-user channel.

2.2 Crosstalk Pre-compensation (Downstream)

2.2.1 Zero Forcing

In downstream transmissions we are concerned with the application of some transformation $\mathbf{P}$ to $\mathbf{x}$ prior to transmission. As a result noise enhancement is no longer an issue and the ZF and MMSE solutions coincide.

The ZF solution prevents any crosstalk by setting

$$\begin{aligned}\mathbf{P}_{\text{zf}} &= \frac{1}{\alpha}\mathbf{H}^{-1} \\ \alpha &= \max_i \left\| \left[\mathbf{H}^{-1}\right]_{\text{row } i} \right\|\end{aligned}$$

Multiplication with $\mathbf{H}^{-1}$ can lead to a significant increase in the transmit power of a pair. This is why we introduce the scaling term α which ensures that the transmit power constraint for each pair is maintained. As a result

$$\begin{aligned}\mathbf{y} &= \frac{1}{\alpha}\mathbf{x} + \mathbf{n} \\ \mathbf{SINR}_{\text{zf}} &= \frac{1}{\alpha^2}\mathbf{S}_{\mathbf{x}\cdot}/\mathbf{S}_{\mathbf{n}}\end{aligned}$$

In the case when the constituent pairs of a binder have significantly different gains (eg. for pairs with different lengths or when one pair has a bridged tap) α must be large to satisfy the power constraint for the weakest pair. This can unnecessarily reduce the transmit power on the strong pair and lead to a solution which is quite sub-optimal.

2.2.2 Vectored Transmitter

Vectoring in the downstream addresses the issue of power amplification which occurs in the ZF solution. A modulus operation, coupled with feedback is used to ensure that the transmit power is unaffected by the interference pre-compensation operation[2].

Similar to the vectored receiver we do a QR decomposition of the channel matrix

$$\begin{aligned}\mathbf{H}^T &= \mathbf{QR} \\ \mathbf{H} &= \mathbf{R}^T\mathbf{Q}^T \\ \mathbf{P}_{\text{vec}} &= (\mathbf{Q}^H)^T \\ \mathbf{y} &= \mathbf{H}\mathbf{P}_{\text{vec}}\mathbf{x} + \mathbf{n} \\ &= \mathbf{R}^T\mathbf{x} + \mathbf{n}\end{aligned} \tag{2.3}$$

where $(.)^T$ represents the transpose operation. Note that pre-multiplication of the vector $\mathbf{x}$ by $\mathbf{P}_{\text{vec}}$ does not alter the transmit power since $\mathbf{P}_{\text{vec}}$ is a unitary matrix.

Writing $\mathbf{y}$ explicitly

$$\begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_N \end{bmatrix} = \begin{bmatrix} r_{11} & 0 & 0 & 0 \\ r_{12} & r_{22} & 0 & 0 \\ \vdots & \vdots & \ddots & 0 \\ r_{1N} & r_{2N} & \dots & r_{NN} \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_N \end{bmatrix} + \mathbf{n}$$

At this point we could apply a simple pre-compensation operation on the elements of $\mathbf{x}$ to prevent interference from occurring. This would correspond to multiplying the vector of true data symbols $\mathbf{u} = [u_1 \dots u_N]^T$ by the matrix $\mathbf{R}^{-T}$ to form $\mathbf{x}$. Such an operation would remove all interference however it would also lead to a power increase in the transmitted symbols $\mathbf{x}$. Instead we use a modulus operation to maintain constant transmit power

$$x_i = \text{mod}_{2M_i}\left[u_i - \sum_{j=1}^{i-1}\frac{r_{ji}}{r_{ii}}x_j\right] \tag{2.4}$$

which is in essence the same operation as used in the Tomlinson-Harashima pre-coder (see [7] for a good overview). Here $\text{mod}_{2M_i}[.]$ represents the modulus operation centered at 0 which insures that x_i is always in the range $(-M_i \dots M_i)$

$$\text{mod}_{2M_i}[g] = g - 2M_i\left\lfloor\frac{g+M_i}{2M_i}\right\rfloor$$

Here M_i is the constellation size for user i. In other words u_i has a constellation formed by the odd integers in the range $(-M_i \ldots M_i)$. The extension of this operation to complex constellations is straightforward.

Since all operations are done at the transmitter error propagation is not an issue in the vectored transmitter (unlike the vectored receiver). A modulus operation must also be applied at the receiver before we make a hard decision

$$\widehat{u}_i = \text{dec}\left[\text{mod}_{2M_i}\left[\frac{y_i}{r_{ii}}\right]\right] \tag{2.5}$$

We can write (2.4) in vector form as

$$\mathbf{x} = \text{mod}_{2M_i}\left[\underbrace{\mathbf{u} - \text{diag}\left\{\mathbf{R}\right\}^{-1}\text{offdiag}\left\{\mathbf{R}^T\right\}\mathbf{x}}_{\mathbf{a}}\right] \tag{2.6}$$

and using (2.5) and (2.3) we can show

$$\begin{aligned}
\widehat{\mathbf{u}} &= \text{dec}\left[\text{mod}_{2M}\left[\text{diag}\left\{\mathbf{R}\right\}^{-1}\mathbf{y}\right]\right] \\
&= \text{dec}\left[\text{mod}_{2M}\left[\text{diag}\left\{\mathbf{R}\right\}^{-1}\left(\text{diag}\left\{\mathbf{R}^T\right\}\mathbf{x} + \text{offdiag}\left\{\mathbf{R}^T\right\}\mathbf{x} + \mathbf{n}\right)\right]\right] \\
&= \text{dec}\left[\text{mod}_{2M}\left[\mathbf{x} + \underbrace{\text{diag}\left\{\mathbf{R}\right\}^{-1}\text{offdiag}\left\{\mathbf{R}^T\right\}\mathbf{x} + \text{diag}\left\{\mathbf{R}\right\}^{-1}\mathbf{n}}_{\mathbf{b}}\right]\right]
\end{aligned}$$

Upon substitution of (2.6) we find

$$\begin{aligned}
\widehat{\mathbf{u}} &= \text{dec}\left[\text{mod}_{2M}\left[\text{mod}_{2M}\left[\mathbf{a}\right] + \mathbf{b}\right]\right] \\
&= \text{dec}\left[\text{mod}_{2M}\left[\mathbf{a} + \mathbf{b}\right]\right] \\
&= \text{dec}\left[\text{mod}_{2M}\left[\mathbf{u} + \text{diag}\left\{\mathbf{R}\right\}^{-1}\mathbf{n}\right]\right]
\end{aligned}$$

Examining this one can see that the probability of making an error is almost the same as for a conventional system without the mod operation. For symbols not on the edge of the constellation the number of nearest neighbours (which dominates the probability of error) is still 4. For symbols on the edge of the constellation the number of nearest neighbours increases from 3 (or 2) to 4 due to the wrap around effect of the mod operation. This will increase the probability of error slightly however for large constellation sizes M_i the increase will be negligible.

Hence use of the vectored transmitter results in the following SINR

$$\mathbf{SINR}_{\text{vec tx}} = \text{diag}\left\{|\mathbf{R}|^2_{\text{elem}}\right\}(\mathbf{S}_x ./ \mathbf{S}_n)$$

2.3 Simulations

2.3.1 Simulation Environment

All simulations are done using a modified version of the Alcatel Matlab simulation environment for VDSL. Simulations are done in the frequency domain using PSDs and we ignore the effects of ISI and ICI.

The scenario used is as follows:

Channel Transfer Functions

- Measured British Telecom cable Transfer Function (TF), spline interpolated to frequencies of interest

- length - 1200 m
- 8× 0.5mm pairs

Echo and NEXT Transfer Functions

- Empirical TFs for AWG 24

Bridged Taps

- None
  OR
- Bridged tap of 40m @ 100m from TX on pair 5

Transmit PSDs

- Normal (-60 dBm/Hz flat)
- Uses 998 FDD Band Plan
- Frequency Range: 10kHz-8MHz

Noise

- Analog Front End (AFE) + Echo + NEXT + AWGN
- Alien Crosstalk Type A
- RFI: none

Plotted in Figure 2.2 and 2.3 are some sample TFs from the measurement set. Figure 2.2 (a) contains the direct TF for pair 1, Figure 2.2 (b) contains the direct TF for pair 5 (with bridged tap) and Figure 2.3 contains the crosstalk TF from pair 1 to 3.

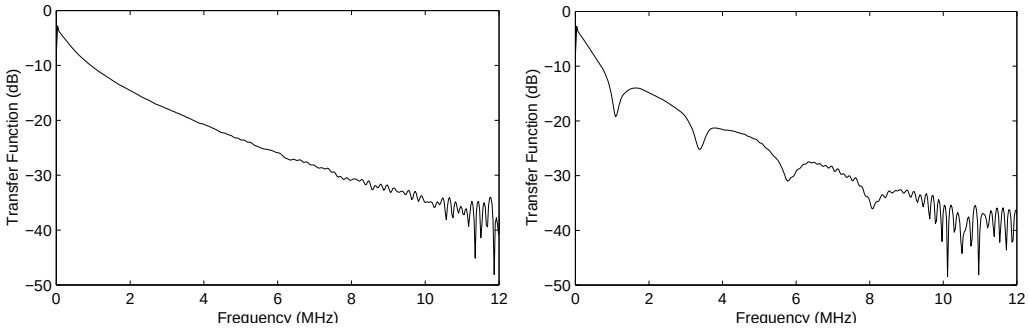
Notice the large fluctuations of the TFs in the 10-12 MHz frequency range. This is due to inaccuracies in the measurement equipment and is not related to the characteristics of the line itself. Due to this, we limit the range of frequencies in our simulations to 8 Mhz.

2.3.2 Theoretical Capacity Limit

A useful measure for evaluation of crosstalk cancellation schemes is the theoretical capacity of a MIMO channel[8]

$$C = \frac{1}{2N}\log_2 \left|\mathbf{I}_N + \mathbf{H}\text{diag}\left\{\mathbf{S}_x ./ \mathbf{S}_n\right\}\mathbf{H}^H\right| \quad \text{bps/Hz/user}$$

Here $|.|$ represents the determinant of a matrix.



(a) Normal Pair (b) Bridged Tap Pair

Figure 2.2: Direct Transfer Function for normal pair (pair 1) and bridged tap pair (pair 5)

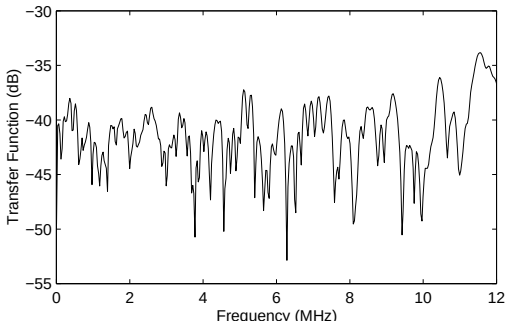


Figure 2.3: Crosstalk Transfer Function from pair 1 to 3

2.3.3 Pre and Post-Cancellation (Full Vectoring)

One way of achieving the Theoretical Capacity of a MIMO channel is to use pre and post-cancellation through co-operation at the transmitter and receiver. This technique is discussed here only for comparison. It is not realizable in the VDSL context since it requires knowledge of the entire vector $\mathbf{x}$ at each network terminal (NT) (ie. co-location of modems at the CP side).

Taking the SVD of the channel matrix[9]

$$\mathbf{H} = \mathbf{U}\Sigma\mathbf{V}^H$$

where $\mathbf{U}$ and $\mathbf{V}$ are unitary matrices and Σ is diagonal. We use a pre-compensation at the LT $\mathbf{x} = \mathbf{V}\mathbf{u}$ and cancellation at the LT $\mathbf{z} = \mathbf{U}^H\mathbf{y}$. Here $\mathbf{u}$ represents the vector of true symbols. Hence

$$\begin{aligned} \mathbf{y} &= \mathbf{H}\mathbf{V}\mathbf{u} + \mathbf{n} \\ \mathbf{z} &= \mathbf{U}^H\mathbf{y} \\ &= \mathbf{U}^H\mathbf{H}\mathbf{x} + \mathbf{U}^H\mathbf{n} \\ &= \Sigma\mathbf{u} + \widetilde{\mathbf{n}} \end{aligned}$$

If $\mathbf{n}$ is i.i.d. then so is $\widetilde{\mathbf{n}}$

$$\begin{aligned} \mathbf{S}_{\text{signal}} &= |\Sigma|^2 \mathbf{S}_{\mathbf{x}} \\ \mathbf{S}_{\text{interference}} &= \mathbf{0} \\ \mathbf{S}_{\text{noise}} &= \mathbf{S}_{\widetilde{\mathbf{n}}} = \mathbf{S}_{\mathbf{n}} \\ \mathbf{SINR}_{\text{svd}} &= |\Sigma|^2 (\mathbf{S}_{\mathbf{x}}./\mathbf{S}_{\mathbf{n}}) \end{aligned}$$

This technique will sometimes simply be referred to as the SVD technique.

2.3.4 Crosstalk Free

In our simulations we also compare performance of the crosstalk cancellers to the crosstalk free scenario. In this case, the crosstalk channels are set to zero whilst the direct channels remain unchanged

$$h_{ij}^{(\mathrm{cf})} = \begin{cases} 0 & i \neq j \\ h_{ij} & i = j \end{cases}$$

Hence

$$\begin{aligned} \mathbf{S}_{\mathrm{signal}} &= \mathrm{diag}\left\{ |\mathbf{H}|^2_{\mathrm{elem}} \right\} \mathbf{S}_x \\ \mathbf{S}_{\mathrm{interference}} &= \mathbf{0} \\ \mathbf{S}_{\mathrm{noise}} &= \mathbf{S}_n \\ \mathbf{SINR}_{\mathrm{cf}} &= \mathrm{diag}\left\{ |\mathbf{H}|^2_{\mathrm{elem}} \right\} \mathbf{S}_x \,./\, \mathbf{S}_n \end{aligned}$$

2.3.5 Performance Results

Average Data Rates

Contained within tables 2.1 and 2.2 are the average data rates (per user) under the various crosstalk cancellation schemes

Bitloading	continuous		integer	
Bridged Taps	no	yes	no	yes
Crosstalk Free	9.81	9.75	9.89	9.83
No Cancellation	7.21	6.39	7.33	6.45
ZF	9.81	9.74	9.89	9.82
MMSE	9.81	9.74	9.89	9.82
Vectored Rx	9.81	9.75	9.89	9.83
Pre and Post	9.81	9.75	9.89	9.83
Theoretical Limit	9.81	9.75	-	-

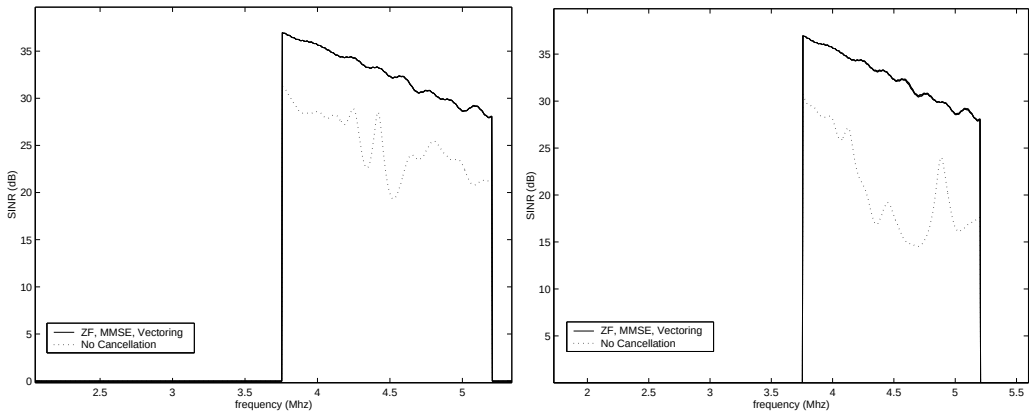
Table 2.1: US Average User Data Rate (Mbps/user)

Bitloading	continuous		integer	
Bridged Taps	no	yes	no	yes
Crosstalk Free	41.82	40.89	41.32	40.44
No Cancellation	34.03	32.60	34.27	32.92
ZF	39.59	34.70	39.28	34.83
Vectored Tx	41.82	40.89	41.31	40.44
Pre and Post	41.85	40.93	41.33	40.47
Theoretical Limit	41.85	40.93	-	-

Table 2.2: DS Average User Data Rate (Mbps/user)

Signal to Interference + Noise Ratio

Plotted in Figure 2.4 are the US SINRs for user 3 in both the normal channels and the channels with a bridged tap. Figure 2.5 shows the same for the DS case. In Figure 2.6 the SINRs for user 5 (the user who's pair contains the bridged tap) is plotted for the bridged tap channel.



(a) Normal channel (b) Bridged tap

Figure 2.4: User 3 Upstream SINR

Data Rates vs. User

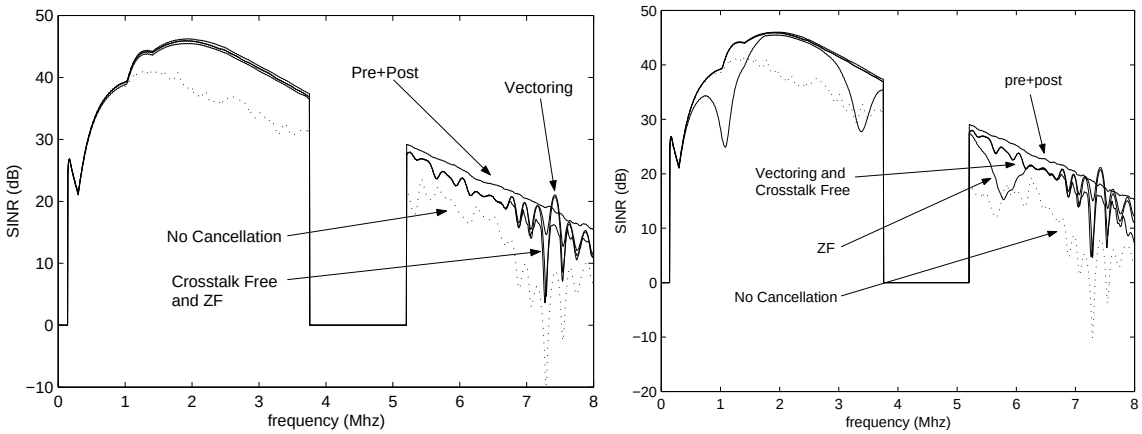
Contained in Figures 2.7 and 2.8 is the data rate across users for the different crosstalk cancellation schemes in the US and DS directions respectively.

2.4 Conclusion

Significant gains are achievable through crosstalk mitigation. Looking at table 2.1 we see that typical data rate gains are around 7 Mbps in the DS direction and 2.5 Mbps in the US. Such performance improvements justify further investigation in crosstalk mitigation.

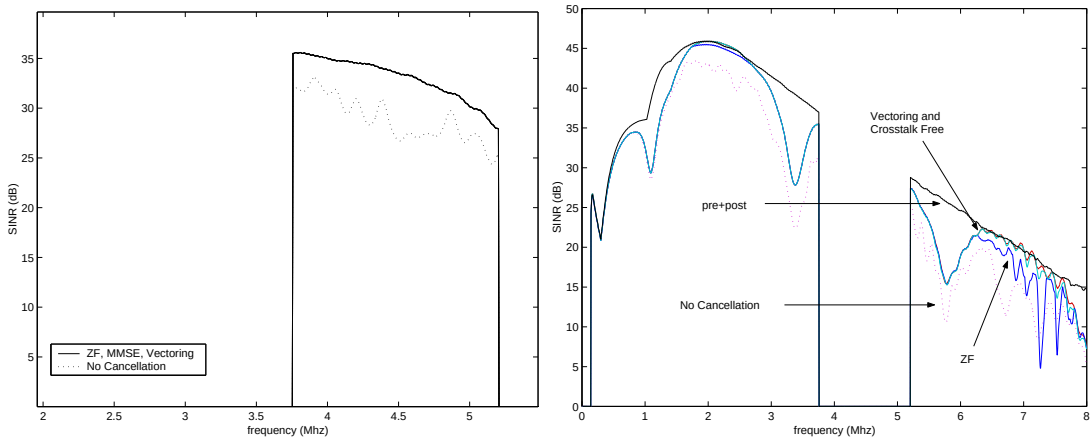
As can be seen in table 2.1 there is essentially little difference between the cancellation schemes in the US direction even in the presence of bridged taps. For this reason it is preferable to choose a linear crosstalk cancellation technique like ZF or MMSE since these schemes do not suffer from error propagation.

In the DS direction however it is quite important to choose a good pre-compensation scheme and we see a significant difference between the vectored transmitter and ZF pre-compensator (table 2.2). This difference is exaggerated in the bridged tap channel and is a result of the power amplification problem. As such it is recommended to use the vectored transmitter for DS pre-compensation. Additionally, error propagation is not a problem in the DS case since we have full knowledge of all transmitted symbols.



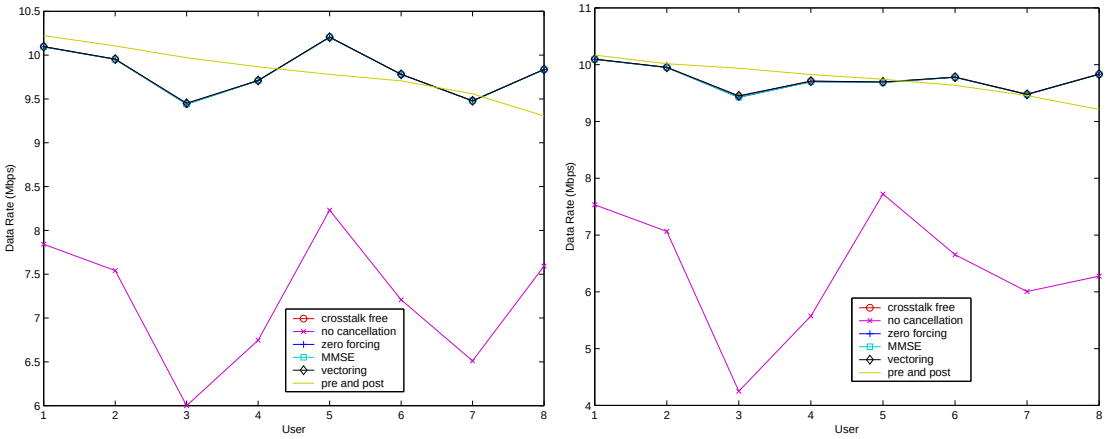
(a) Normal channel (b) Bridged tap

Figure 2.5: User 3 Downstream SINR



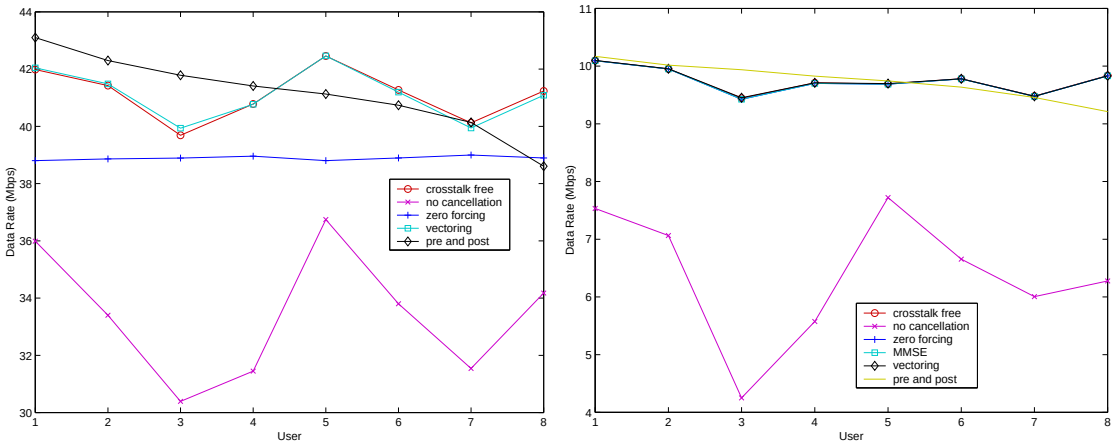
(a) Upstream (b) Downstream

Figure 2.6: User 5 SINR in bridged tap channel



(a) Normal channel (b) Bridged tap

Figure 2.7: Upstream Data Rates Across Users



(a) Normal channel (b) Bridged tap

Figure 2.8: Downstream Data Rates Across Users

Chapter 3

Complexity Study

In terms of memory requirements and computational complexity, simultaneous cancellation of the many users constituting a binder can impose a significant burden on the Central Office (CO) hardware. As a result, we are interested in ways of reducing complexity whilst maintaining close to the full performance provided by the crosstalk cancellation schemes of Chapter 2.

This chapter contains the results of an initial study into complexity reduction methods. Section 3.1 focuses on techniques for memory reduction through exploitation of crosstalk reciprocity and tone grouping. Section 3.2 addresses the complexity issue through techniques which only do crosstalk cancellation in frequency bands that suffer significantly from crosstalk effects and by only cancelling dominant interferers.

Other methods for complexity reduction include replacing full crosstalk cancellation with Dynamic Spectrum Management (water filling across users and frequency). This technique attempts to mitigate crosstalk by intelligent allocation of power resources instead of actual operations on the signal level. Dynamic spectrum management and related techniques are discussed further in Chapters 5-8.

3.1 Memory

3.1.1 Crosstalk Reciprocity

One proposed way of reducing complexity is to exploit the assumed reciprocity between crosstalk coefficients. In other words, it is assumed that crosstalk coupling coefficients between lines are symmetrical

$$h_{ij} = h_{ji}$$

This allows us to halve the required memory coefficients since we only need to store h_{ij} for $i \geq j$. To evaluate the validity of this assumption, we have constructed a set of matrices $\mathbf{H}_{\text{recip}}(k)$ which are formed by knowledge of the upper triangular part of $\mathbf{H}(k)$ and the reciprocity assumption. Mathematically the relationship is

$$h_{ij}^{(\text{recip})} = \begin{cases} h_{ij} & i \leq j \\ h_{ji} & i > j \end{cases}$$

where $\mathbf{H}_{\text{recip}}(k) = \left[h_{ij}^{(\text{recip})}\right]$ and we recall that $\mathbf{H}(k) = [h_{ij}]$.

Using this formulation we can evaluate a kind of percentage error ϵ in the channel matrix estimate $\mathbf{H}_{\text{recip}}(k)$, as a result of the reciprocity assumption

$$\epsilon(k) = \frac{\|\mathbf{H}(k) - \mathbf{H}_{\text{recip}}(k)\|_{\text{fro}}}{\|\mathbf{H}(k)\|_{\text{fro}}}$$

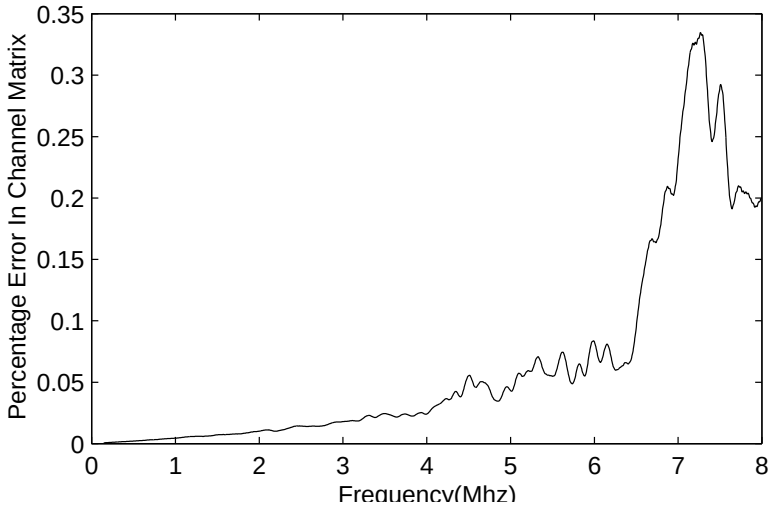


Figure 3.1: Percentage Error in channel matrix estimate $\mathbf{H}_{\text{recip}}(k)$ due to reciprocity assumption

Equalizer	US		DS	
	Bitrate	Degradation	Bitrate	Degradation
Crosstalk Free	9.81	-	41.82	-
No Equalization	7.21	-	34.03	-
ZF	7.50	2.31	33.12	6.47
MMSE	7.53	2.28	-	-
Vectoring	7.53	2.28	34.30	7.52
Pre and Post	7.42	2.39	34.18	7.67

Table 3.1: Average Data Rate (Mbps) and Performance Degradation under Reciprocity Assumption

where $\|.\|_{\text{fro}}$ represents the Frobenius norm. The resultant error is plotted in Figure 3.1 for the normal (non bridged tap) channel. Using $\mathbf{H}_{\text{recip}}(k)$ we can also design crosstalk cancellation filters (as shown in Chapter 2) and evaluate data rate degradation as a result of the reciprocity assumption. The results of this is shown in table 3.1. We see that the crosstalk reciprocity assumption is quite inaccurate at high frequencies and this leads to significant performance degradation. Hence the reciprocity assumption is not a wise choice.

3.1.2 Tone Grouping

Under the assumption that there is strong correlation between neighboring (in frequency) elements of the crosstalk transfer functions (ie. $h_{ij}(k)$ correlated with $h_{ij}(k+1)$) it is possible to reduce memory requirements through the use of tone grouping. Tone grouping essentially replaces the crosstalk cancellation parameters of a group of tones with some single set of shared parameters. This results in a reduction in memory usage by a factor M when M is the group size.

In our simulations the crosstalk cancellation parameters for a group are defined by the parameters of the group's central tone. Other ways of calculating the shared parameters are possible however this choice results in the simplest implementation. Furthermore, the group size is constant with frequency. It may be more optimal to have large group sizes at lower frequencies (where the crosstalk transfer functions are highly correlated) and small group sizes at higher frequencies.

Analysis US

In the linear crosstalk cancellation schemes (ZF, MMSE) analysis of the performance degradation in the US due to tone grouping is quite straightforward. We can express the received signal vector after crosstalk cancellation as

$$\mathbf{z}(k) = \mathbf{W}(c)\mathbf{y}(k) = \mathbf{W}(c)\mathbf{H}(k)\mathbf{x}(k) + \mathbf{W}(c)\mathbf{n}(k)$$

Note that for tone k we use the crosstalk cancellation parameters from tone c; the central tone for the group of tone k

$$c = m + \left\lfloor \tfrac{M}{2} \right\rfloor \quad \forall\, k \in \{m \ldots m + M - 1\}$$

Contrast this with the normal scenario in which $c = k$.

For any linear scheme we can simply determine the PSDs of the signal, interference and noise terms in the usual fashion

$$\begin{aligned} \mathbf{S}_{\text{signal}} &= \text{diag}\left\{ |\mathbf{W}(c)\mathbf{H}(k)|^2_{\text{elem}} \right\} \mathbf{S}_x \\ \mathbf{S}_{\text{interference}} &= \text{offdiag}\left\{ |\mathbf{W}(c)\mathbf{H}(k)|^2_{\text{elem}} \right\} \mathbf{S}_x \\ \mathbf{S}_{\text{noise}} &= |\mathbf{W}(c)|^2_{\text{elem}} \, \mathbf{S}_n \end{aligned}$$

For the vectored receiver, analysis proceeds as follows. Again writing out the received vector

$$\begin{aligned} \mathbf{H}(c) &= \mathbf{Q}(c)\mathbf{R}(c) \\ \mathbf{z}(k) &= \mathbf{Q}^H(c)\mathbf{H}(k)\mathbf{x}(k) + \mathbf{Q}^H(c)\mathbf{n}(k) \\ &= \mathbf{G}(k)\mathbf{x}(k) + \widetilde{\mathbf{n}}(k) \\ &= \begin{bmatrix} g_{11}(k) & \ldots & g_{1N}(k) \\ \vdots & \ddots & \vdots \\ g_{N1}(k) & \ldots & g_{NN}(k) \end{bmatrix} \begin{bmatrix} x_1(k) \\ \vdots \\ x_N(k) \end{bmatrix} + \widetilde{\mathbf{n}}(k) \end{aligned}$$

where $\mathbf{G}(k) = \mathbf{Q}^H(c)\mathbf{H}(k)$. The decision feedback part of the receiver now proceeds as follows

$$\widehat{x}_i(k) = \text{dec}\left[\frac{z_i(k)}{r_{ii}(c)} - \sum_{j=i+1}^{N} \frac{r_{ij}(c)}{r_{ii}(c)} \widehat{x}_j \right]$$

Assuming previous decisions are correct

$$\widehat{x}_i(k) = \text{dec}\left[\frac{g_{ii}(k)}{r_{ii}(c)} x_i + \sum_{j=1}^{i-1} \frac{g_{ij}(k)}{r_{ii}(c)} x_j + \frac{1}{r_{ii}(c)} \sum_{j=i+1}^{N} \left(g_{ij}(k) - r_{ij}(c) \right) x_j + \frac{1}{r_{ii}(c)} \widetilde{n}_i \right] \tag{3.1}$$

Examining (3.1) we can see three terms corresponding to the scaled signal from the user of interest and two interference terms. The first interference term arises from the fact that $\mathbf{Q}^H(c)$ cannot perfectly upper triangularize $\mathbf{H}(k)$. The second term results mostly from the discrepancy between $\mathbf{R}(c)$ and $\mathbf{R}(k)$.

If we ignore the scaling ($\frac{g_{ii}(k)}{r_{ii}(c)}$) on the signal of interest we can determine the PSDs of the signal, interference and noise terms as follows

$$\begin{aligned} \mathbf{S}_{\text{signal}} &= \text{diag}\left\{ |\mathbf{G}(k)|^2_{\text{elem}} \right\} \mathbf{S}_x \\ \mathbf{S}_{\text{interference}} &= \text{offdiag}\left\{ |\mathbf{G}(k) - \mathbf{R}(c)|^2_{\text{elem}} \right\} \mathbf{S}_x \\ \mathbf{S}_{\text{noise}} &= \mathbf{S}_n \end{aligned}$$

Analysis DS

Turning our attention to the DS case, analysis for ZF crosstalk pre-compensation proceeds as follows

$$\mathbf{y}(k) = \mathbf{H}(k)\mathbf{P}(c)\mathbf{x}(k) + \mathbf{n}(k)$$

and hence the PSDs are

$$\begin{aligned}\mathbf{S}_{\text{signal}} &= \text{diag}\left\{|\mathbf{H}(k)\mathbf{P}(c)|^2_{\text{elem}}\right\}\mathbf{S}_x \\ \mathbf{S}_{\text{interference}} &= \text{offdiag}\left\{|\mathbf{H}(k)\mathbf{P}(c)|^2_{\text{elem}}\right\}\mathbf{S}_x \\ \mathbf{S}_{\text{noise}} &= \mathbf{S}_n\end{aligned}$$

In the vectored transmitter we proceed as follows. The channel is well conditioned and diagonally dominant hence back-substitution (with the $\mathbf{R}^T$ matrix) leads to little power change. As a result we ignore the $\text{mod}_{2M_i}[.]$ operations to simplify analysis. The vectored transmitter now corresponds to a ZF pre-compensator with power scaling by $\text{diag}\{\mathbf{R}\}$. Writing out the received vector

$$\mathbf{y}(k) = \mathbf{H}(k)\left(\mathbf{Q}(c)^T\right)^H\mathbf{x}(k) + \mathbf{n}(k)$$

where $\mathbf{H}(c) = \mathbf{R}(c)^T\mathbf{Q}(c)^T$. Using the definition of the vectored transmitter (2.4) and ignoring the $\text{mod}_{2M_i}[.]$ operations

$$\mathbf{x}(k) = \mathbf{R}(c)^{-T}\text{diag}\left\{\mathbf{R}(c)\right\}\mathbf{u}(k)$$

hence

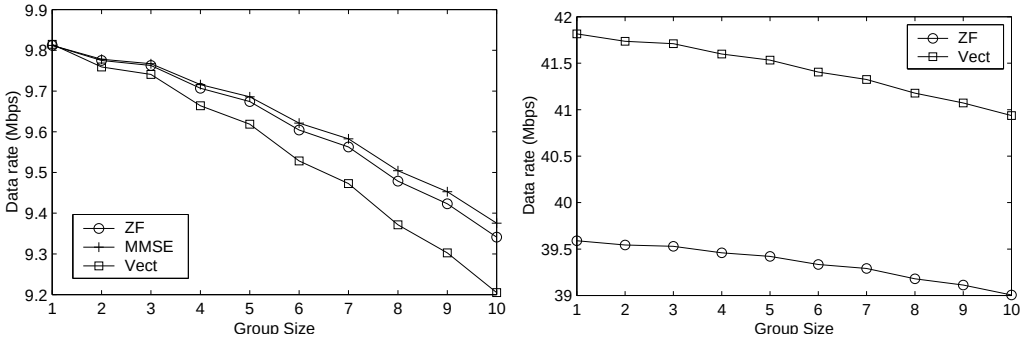
$$\mathbf{y}(k) = \mathbf{H}(k)\underbrace{\left(\mathbf{Q}(c)^T\right)^H\mathbf{R}(c)^{-T}\text{diag}\left\{\mathbf{R}(c)\right\}}_{\mathbf{P}(c)}\mathbf{u}(k) + \mathbf{n}(k)$$

In terms of PSDs

$$\begin{aligned}\mathbf{S}_{\text{signal}} &= \text{diag}\left\{|\mathbf{H}(k)\mathbf{P}(c)|^2_{\text{elem}}\right\}\mathbf{S}_x \\ \mathbf{S}_{\text{interference}} &= \text{offdiag}\left\{|\mathbf{H}(k)\mathbf{P}(c)|^2_{\text{elem}}\right\}\mathbf{S}_x \\ \mathbf{S}_{\text{noise}} &= \mathbf{S}_n\end{aligned}$$

Performance

Using the analysis described above we have evaluated performance degradation for the various crosstalk cancellation schemes when tone grouping is used. Shown in Figure 3.2 is the average data rate per user vs. group size M for the different schemes. As can be seen it is possible to use a group size of up to 6 (which results in a memory reduction by a



(a) Upstream (b) Downstream

Figure 3.2: Average Data Rate vs. Group Size M

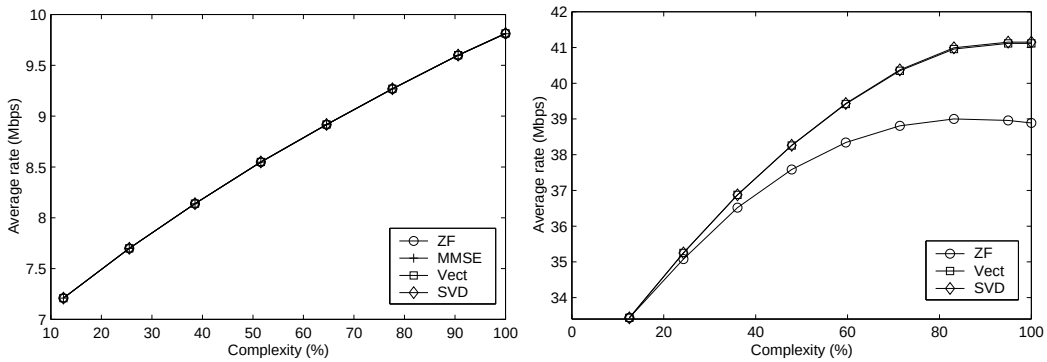
factor of 6) with only minor (<500kbps) decreases in data rate.

3.2 Complexity

3.2.1 Limited Frequency Band

One way of reducing complexity is to cancel crosstalk only on a limited selection of tones. We select the tones that exhibit the largest data rate increase through crosstalk cancellation and in this way minimize the data rate loss as a result of the neglected tones.

Plotted in Figure 3.3 are the results of simulations done where we perform crosstalk cancellation in the limited frequency band as described. Unfortunately the reduction in complexity is virtually linear with reduction in data rate which makes this scheme unattractive.



(a) Upstream

(b) Downstream

Figure 3.3: Average Data Rate vs. Tones Cancelled

3.2.2 Dominant Interferers

It is also possible to reduce complexity by only cancelling the dominant interferers of each user. Viewed in a MIMO (rather than multi-access) framework, we desire to detect each user by observing only a sub-set of the signals on the CO lines.

For the linear (ZF, MMSE) techniques this leads to a reduction in complexity by a factor $\frac{n}{N}$ where n is the number of CO signals observed and N is the total number of lines. For vectored techniques the reduction in complexity is similar although the expression is more complex as we shall see.

Linear Techniques

Crosstalker Selection

To begin with we need a method for selecting which subset of CO lines to observe when detecting a user. Intuitively, we might try to cancel the dominant crosstalkers, ie. those that contribute the most interference energy to the user of interest's line. Hence we wish to observe the lines corresponding to the dominant interferer's direct channels. For the i'th user this corresponds to the largest elements of row i of the channel matrix $\mathbf{H}$.

It turns out that this method for selecting the observation lines is optimal when the channel matrix $\mathbf{H}$ is diagonally dominant.

Upstream - ZF

In the upstream direction recall that (1.1)

$$\mathbf{z} = \mathbf{W}(\mathbf{H}\mathbf{x} + \mathbf{n})$$

Let us define $\mathbf{D}_i$ as the union of i and the set of dominant interferers for user i. Then by only observing a subset of the CO lines in the detection of user i we impose the following constraint

$$[\mathbf{W}]_{ij} = 0 \text{ if } j \notin \mathbf{D}_i$$

For convenience we define

$$\begin{aligned} \widetilde{\mathbf{W}}_{\text{row } i} &= \text{non} - \text{zero elements of } \mathbf{W}_{\text{row } i} \\ \widetilde{\mathbf{H}}_i &= \text{corresponding rows of } \mathbf{H} \\ &= \mathbf{H}_{\text{rows } \mathbf{D}_i} \end{aligned}$$

We can now define the zero forcing equalizer through the LS criterion

$$\begin{aligned} \mathbf{W}^{\text{opt}}_{\text{row } i} &= \arg \min_{\mathbf{W}_{\text{row } i}} \mathcal{E} \left| x_i - \mathbf{W}_{\text{row } i}\mathbf{H}\mathbf{x} \right|^2 \\ \widetilde{\mathbf{W}}^{\text{opt}}_{\text{row } i} &= \arg \min_{\widetilde{\mathbf{W}}_{\text{row } i}} \mathcal{E} \left| x_i - \widetilde{\mathbf{W}}_{\text{row } i}\widetilde{\mathbf{H}}_i\mathbf{x} \right|^2 \\ &= \arg \min_{\widetilde{\mathbf{W}}_{\text{row } i}} \mathbf{J} \end{aligned}$$

Solving $\nabla_{\widetilde{\mathbf{W}}_{\text{row } i}} \mathbf{J} = 0$ yields

$$\begin{aligned} \widetilde{\mathbf{W}}^{\text{opt}}_{\text{row } i} &= \left[\begin{array}{ccc} \mathbf{0}_{1\times i-1} & 1 & \mathbf{0}_{1\times N-i} \end{array} \right] \widetilde{\mathbf{H}}_i^H \left(\widetilde{\mathbf{H}}_i \widetilde{\mathbf{H}}_i^H \right)^{-1} \\ \left(\widetilde{\mathbf{W}}^{\text{opt}}_{\text{row } i} \right)^T &= \left(\widetilde{\mathbf{H}}_i^T \right)^{\dagger} \left[\begin{array}{c} \mathbf{0}_{i-1\times 1} \\ 1 \\ \mathbf{0}_{N-i\times 1} \end{array} \right] \end{aligned}$$

where $(.)^{\dagger}$ represents the pseudo-inverse operation.

To help clarify things we give an example of ZF crosstalk canceller design. Given the US channel matrix

$$\mathbf{H} = \left[\begin{array}{cccc} 40 & 1 & 3 & 2 \\ 0.5 & 50 & 4 & 3 \\ 5 & 7 & 20 & 2 \\ 1 & 3 & 8 & 10 \end{array} \right]$$

Lets assume we desire to cancel 2 interferers per user which will lead to a complexity of $\frac{3}{4}$ the full complexity. Thus

$$\mathbf{D}_1 = \{1, 3, 4\}$$

So $\mathbf{W}$ must have the form

$$\mathbf{W} = \left[\begin{array}{cccc} w_{11} & 0 & w_{13} & w_{14} \\ 0 & w_{22} & w_{23} & w_{24} \\ w_{31} & w_{32} & w_{33} & 0 \\ 0 & w_{42} & w_{43} & w_{44} \end{array} \right]$$

$$\widetilde{\mathbf{W}}_{\text{row } 1} = \left[\begin{array}{ccc} w_{11} & w_{13} & w_{14} \end{array} \right]$$

$$\widetilde{\mathbf{H}}_1 = \left[\begin{array}{cccc} 40 & 1 & 3 & 2 \\ 5 & 7 & 20 & 2 \\ 1 & 3 & 8 & 10 \end{array} \right]$$

$$\begin{aligned} \widetilde{\mathbf{W}}^{\text{opt}}_{\text{row } 1} &= \left[\begin{array}{cccc} 1 & 0 & 0 & 0 \end{array} \right] \widetilde{\mathbf{H}}_1^H \left(\widetilde{\mathbf{H}}_1 \widetilde{\mathbf{H}}_1^H \right)^{-1} \\ &= \left[\begin{array}{ccc} 0.0254 & -0.0019 & -0.0047 \end{array} \right] \\ \mathbf{W}^{\text{opt}}_{\text{row } 1} &= \left[\begin{array}{cccc} 0.0254 & 0 & -0.0019 & -0.0047 \end{array} \right] \end{aligned}$$

A similar procedure is used to obtain rows $2 \dots 4$ of $\mathbf{W}^{\text{opt}}$.

Upstream - MMSE

We follow a similar derivation as before except now

$$\begin{aligned}
\mathbf{W}_{\text{row}\,i}^{\text{opt}} &= \arg\min_{\mathbf{W}_{\text{row}\,i}} \mathcal{E}\left|x_i - \mathbf{W}_{\text{row}\,i}\left(\mathbf{H}\mathbf{x} + \mathbf{n}\right)\right|^2 \\
\widetilde{\mathbf{W}}_{\text{row}\,i}^{\text{opt}} &= \arg\min_{\widetilde{\mathbf{W}}_{\text{row}\,i}} \mathcal{E}\left|x_i - \widetilde{\mathbf{W}}_{\text{row}\,i}\widetilde{\mathbf{H}}_i\mathbf{x} - \widetilde{\mathbf{W}}_{\text{row}\,i}\widetilde{\mathbf{n}}\right|^2 \\
&= \arg\min_{\widetilde{\mathbf{W}}_{\text{row}\,i}} \mathbf{J}
\end{aligned}$$

where $\widetilde{\mathbf{n}}$ is defined similar to $\widetilde{\mathbf{H}}_i$ as the rows corresponding to the non-zero elements of $\widetilde{\mathbf{W}}_{\text{row}\,i}$. Again solving $\nabla_{\widetilde{\mathbf{W}}_{\text{row}\,i}}\mathbf{J} = 0$ we find

$$\widetilde{\mathbf{W}}_{\text{row}\,i}^{\text{opt}} = \left[\begin{array}{ccc} \mathbf{0}_{1\times i-1} & 1 & \mathbf{0}_{1\times N-i} \end{array}\right] \widetilde{\mathbf{H}}_i^H \left(\widetilde{\mathbf{H}}_i\widetilde{\mathbf{H}}_i^H + \frac{\sigma_n^2}{\sigma_x^2}\mathbf{I}\right)^{-1} \tag{3.2}$$

where $\mathbf{I}_n$ is the $n \times n$ identity matrix.

Downstream - ZF

In the downstream we recall (1.2)

$$\mathbf{y} = \mathbf{H}\mathbf{P}\mathbf{x} + \mathbf{n}$$

Define $\mathbf{V}_i$ as the union of i and the set of users who consider user i as a dominant interferer. Now since we only wish to do operations on a subset of the CO lines in the pre-compensation for user i we impose the following constraint

$$[\mathbf{P}]_{ji} = 0 \text{ if } j \notin \mathbf{V}_i$$

For convenience we define

$$\begin{aligned}
\overline{\mathbf{P}}_{\text{col}\,i} &= \text{non} - \text{zero elements of } \mathbf{P}_{\text{col}\,i} \\
\overline{\mathbf{H}}_i &= \text{corresponding columns of } \mathbf{H} \\
&= \mathbf{H}_{\text{columns}\,\mathbf{v}_i}
\end{aligned}$$

We can now define the zero forcing equalizer through the LS criterion

$$\mathbf{x} \overset{\text{LS}}{=} \mathbf{H}\mathbf{P}^{\text{opt}}\mathbf{x} + \mathbf{n}$$

Hence

$$\begin{aligned}
\mathbf{H}\mathbf{P}^{\text{opt}} &\overset{\text{LS}}{=} \mathbf{I} \\
\overline{\mathbf{H}}_i\overline{\mathbf{P}}_{\text{col}\,i}^{\text{opt}} &\overset{\text{LS}}{=} \left[\begin{array}{c} \mathbf{0}_{i-1\times 1} \\ 1 \\ \mathbf{0}_{N-i\times 1} \end{array}\right] \\
\overline{\mathbf{P}}_{\text{col}\,i}^{\text{opt}} &= \left(\overline{\mathbf{H}}_i\right)^{\dagger} \left[\begin{array}{c} \mathbf{0}_{i-1\times 1} \\ 1 \\ \mathbf{0}_{N-i\times 1} \end{array}\right]
\end{aligned}$$

Returning to our previous example but assuming now that $\mathbf{H}$ is the DS channel matrix

$$\mathbf{H} = \left[\begin{array}{cccc} 40 & 1 & 3 & 2 \\ 0.5 & 50 & 4 & 3 \\ 5 & 7 & 20 & 2 \\ 1 & 3 & 8 & 10 \end{array}\right]$$

Lets assume we desire to do pre-compensation operations using 2 dominant interferers per user which leads to a complexity reduction by $\frac{3}{4}$. Thus

$$\mathbf{V}_1 = \{1, 3\}$$

So $\mathbf{P}$ must have the form

$$\mathbf{P} = \begin{bmatrix} p_{11} & 0 & p_{13} & p_{14} \\ 0 & p_{22} & p_{23} & p_{24} \\ p_{31} & p_{32} & p_{33} & 0 \\ 0 & p_{42} & p_{43} & p_{44} \end{bmatrix}$$

$$\overline{\mathbf{P}}_{\mathrm{col}\,1} = \begin{bmatrix} p_{11} \\ p_{31} \end{bmatrix}$$

$$\overline{\mathbf{H}}_1 = \begin{bmatrix} 40 & 3 \\ 0.5 & 4 \\ 5 & 20 \\ 1 & 8 \end{bmatrix}$$

$$\begin{aligned} \overline{\mathbf{P}}^{\mathrm{opt}}_{\mathrm{col}\,1} &= \left(\overline{\mathbf{H}}_1\right)^{\dagger} \begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \end{bmatrix} \\ &= \begin{bmatrix} 0.0254 \\ -0.0058 \end{bmatrix} \\ \mathbf{P}^{\mathrm{opt}}_{\mathrm{col}\,1} &= \begin{bmatrix} 0.0254 \\ 0 \\ -0.0058 \\ 0 \end{bmatrix} \end{aligned}$$

A similar procedure is used to obtain columns $2 \ldots 4$ of $\mathbf{P}^{\mathrm{opt}}$.

Vectored Techniques

In vectoring we rotate ($\mathbf{Q}$) and iteratively detect users in a coupled fashion. As such, we cannot simply detect each user in the isolated way that we could with linear equalizers. The result of this is that each user can no longer have their own set of dominant interferers. If user a is a dominant interferer of user b then both users must be detected together.

Effectively, we need to break the pool of users into non-interfering sets and to do this we employ sparse-matrix re-ordering techniques. This is quite similar to a technique employed in [10] although the context is different.

Non-Interfering Sets

To divide our user pool into non-interfering sets we proceed as follows. We first normalize the channel matrix $\mathbf{H}$ by it's diagonal elements, and force the resulting matrix to be symmetric. Symmetry is required for the sparse-matrix re-ordering algorithm and is intuitively satisfying since it implies that if user $a \in$ set of user b then user $b \in$ set of user a. Mathematically

$$\begin{aligned} \mathbf{A} &= \mathrm{diag}\left\{\mathbf{H}\right\}^{-1}\mathbf{H} \\ \mathbf{B} &= \mathrm{threshold}\left[\mathbf{A} + \mathbf{A}^{T},\, \alpha\right] \\ \mathbf{p} &= \mathrm{symrcm}\left[\mathbf{B}\right] \end{aligned}$$

Here $\mathrm{threshold}\left[\mathbf{G},\, \alpha\right]$ denotes a simple thresholding operation that sets all elements of its argument matrix $\mathbf{G}$ that are below the threshold α to 0 and all other elements to 1. $\mathrm{symrcm}\left[\mathbf{B}\right]$ returns the permutation vector $\mathbf{p}$ (by the Symmetric Reverse Cuthill-McKee method) such that $\mathbf{B}(\mathbf{p}, \mathbf{p})$ tends to have it's non-zero elements closer to the diagonal than $\mathbf{B}$. Note that $\left[\mathbf{B}(\mathbf{p}, \mathbf{p})\right]_{i,j} \triangleq \left[\mathbf{B}\right]_{\mathbf{p}(i),\mathbf{p}(j)}$. For an overview of sparse-matrix re-ordering techniques see [11].

Depicted in Figure 3.4 are the results of a matrix re-ordering. The plot on the left contains $\mathbf{B}$ before re-ordering and each point represents what is considered to be a significant coupling between users. Note that the matrix is forced to be symmetrical about the diagonal. The plot on the right contains $\mathbf{B}(\mathbf{p}, \mathbf{p})$ and here we see the users divided into three decoupled sets.

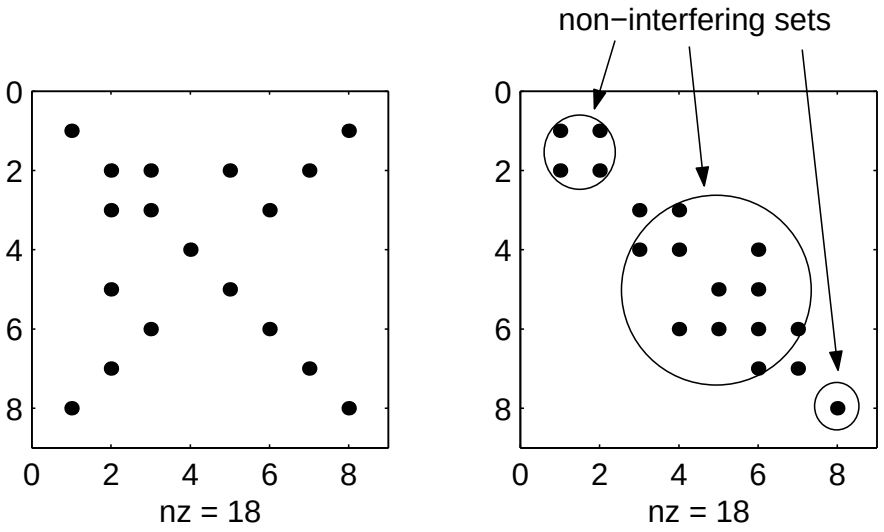


Figure 3.4: The user coupling matrix $\mathbf{B}$ before and after re-ordering into non-interfering sets

Using the re-ordering technique described, the channel matrix $\mathbf{H}$ can be broken into smaller and smaller sub-matrices (each corresponding to a set) by increasing the threshold parameter α. Each set now has it's own $\mathbf{Q}$ and $\mathbf{R}$ matrix which can be used to detect the symbols of that set's members, independently of the other users.

We essentially ignore the coupling parameters *between* sets in detection. As a result decreasing set size reduces complexity but results in a performance penalty aswell.

Complexity Reduction

We now turn our attention to the complexity reduction achieved through the use of non-interfering sets. The complexity for vectored detection of a set of size n is

$$\text{complexity}(n) = \frac{n^2}{2} - \frac{n}{2} + 1 \text{ multiplications/received vector}$$

Hence if we split a channel with N users into l sets with n_i users in the ith set, the complexity is reduced by a factor

$$\frac{\sum_{i=1}^{l} n_i^2 - n_i + 2}{N^2 - N + 2}$$

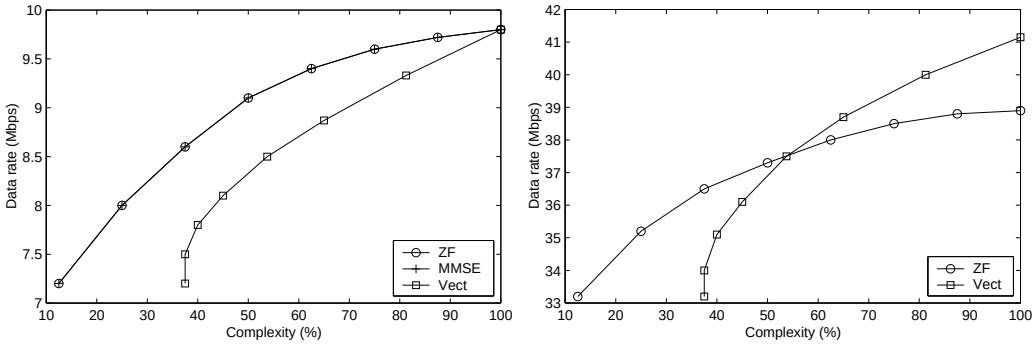
Performance

Figure 3.5 contains the performance versus complexity for the different crosstalk cancellation schemes. As can be seen for the linear schemes it is possible to obtain 91% of the performance gains achieved using full crosstalk cancellation with only $\frac{3}{4}$ of the complexity (5 interferers cancelled) . Reducing the complexity to $\frac{1}{2}$ we still retain 70% of the full performance.

Vectored techniques suffer significantly larger performance degradation and achieve only 66% of the performance gain when the complexity is reduced to $\frac{3}{4}$.

3.3 Conclusion

In this chapter we have reviewed a number of methods to reduce both the memory overhead and complexity of the crosstalk cancellation techniques presented in Chapter 2. It was seen in Section 3.1 that crosstalk reciprocity is not a valid



(a) Upstream (b) Downstream

Figure 3.5: Average Data Rate vs. Interferers Cancelled

assumption and leads to significant performance degradation. Tone grouping turned out to be more successful and it is possible to reduce memory overhead by a factor of 6 whilst suffering minimal (<500kbps) performance degradation.

In Section 3.2 techniques for reducing complexity were investigated and we saw that cancelling crosstalk in a limited frequency band gave only linear complexity reduction versus performance degradation. Hence this technique is not particularly attractive. Cancelling only dominant interferers on the other hand allows complexity to be reduced by a factor of $\frac{3}{4}$ or $\frac{1}{2}$ while maintaining 91% and 70% respectively of the performance gains. It is important to note that this was the case with linear techniques (ZF, MMSE) and that vectored methods show significantly larger performance degradation.

Chapter 4

Summary - Crosstalk Cancellation

The first part of this report (Chapters 1-3) gave an overview of crosstalk pre-compensation and cancellation in VDSL. In Chapter 1 we developed a data model and standard method for analysis of crosstalk mitigation techniques. Chapter 2 described different crosstalk mitigation methods both linear and with decision/non-linear feedback. These schemes were evaluated through simulation and we found the following:

- Significant performance gains are possible in both upstream and downstream transmissions. Typically data rate gains are around 7 Mbps in DS and 2.5 Mbps in the US.
- In upstream operation there is little difference between the various schemes hence it is preferable to choose a simple scheme such as MMSE crosstalk cancellation.
- In downstream operation power amplification causes significant performance degradation in the linear schemes. This problem is even more pronounced when the channel contains bridged taps. As a result it is recommended that Vectored Transmission be used for crosstalk pre-compensation.

Due to the high computational requirements of the crosstalk mitigation schemes, in Chapter 3 we investigated techniques for complexity reduction. It was concluded that:

- For the channels measurements used, crosstalk reciprocity is not a valid assumption and should not be used as a means for memory reduction.
- Tone grouping can be used to reduce memory requirements by a factor of up to 6 with little (<500kbps) performance loss.
- Cancelling crosstalk on a limited selection of tones leads to an almost linear decrease in complexity with performance degradation. As a result, this technique is not attractive for complexity reduction.
- Only cancelling the dominant interferers allows complexity to be reduced by a factor of $\frac{3}{4}$ or $\frac{1}{2}$ whilst maintaining 91% and 70% respectively of the performance gains obtained through full crosstalk cancellation. This is an attractive option for complexity reduction of linear crosstalk mitigation. Note that performance degradation was considerably greater when using vectored crosstalk mitigation.

It should be noted that these results where obtained by deactivating tones above 8 MHz (due to inaccuracies in channel measurements above this frequency). This may give an unfair disadvantage to the partial crosstalk cancellation schemes described above since the crosstalk coupling is fairly consistent in the lower frequencies. In the high frequencies, on the other hand, we see a large difference between crosstalk couplings and hence benefit most by cancelling only dominant interferers/tones. This point will be explored further in the following chapters where we replace the measured channel models by empirical ones, allowing us to run simulations in the full frequency range (up to 12 MHz).

Overall, significant gains are achievable using crosstalk cancellation and through partial cancellation we can achieve this will a relatively small (3-4x) increase in complexity. In the following chapters we explore the multi-user power allocation problem and it's combination with crosstalk cancellation. As we shall see, the use of power allocation significantly increases the benefits of crosstalk cancellation allowing us to achieve large gains (e.g. $20 \rightarrow 65$ Mbps) with low complexity.

Chapter 5

Near-Far Problem

xDSL is a continually evolving technology. Beginning with ISDN in the late 80's, we have seen a plethora of DSL standards develop like HDSL, SDSL, ADSL, VDSL and so on. As these standards evolve, we have seen a continuous roll-out of systems. Older, low-speed systems like HDSL and ADSL operate over long distances from the CO. Modern, high-speed systems and systems proposed for the future will operate on much shorter loops, typically from ONUs at the end of each subscriber's street.NTs

What this continuous roll-out means, in terms of the network structure, is that we now have a mixture of many different services across a whole range of loop lengths. LTs are now spread between the CO and the end of the street, and NTs can be anything between 300m (VDSL) and several kilometers (ADSL) from a common termination point.

The result of this is the 'near-far effect', a problem well known in mobile communications and one which is becoming increasingly important in DSL. The near-far effect results when the crosstalk signal of a user on a short line (near-end user) overwhelms the direct signal of a user on a long line (far-end user).

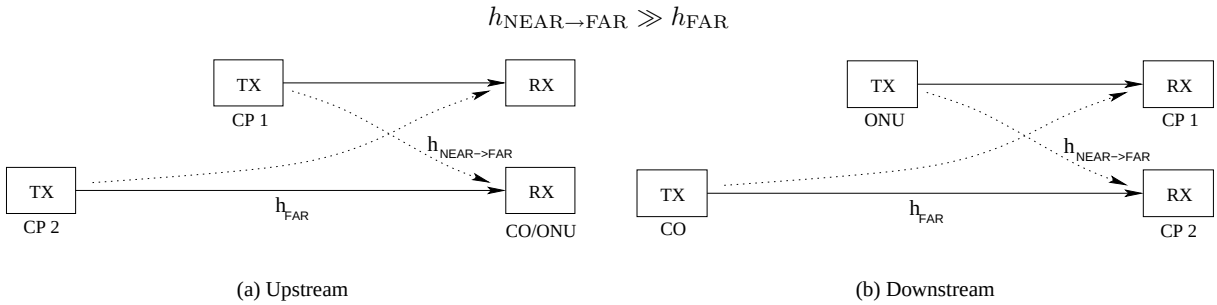
The near-far effect can occur in both up- and down-stream communications. In the up-stream it occurs when a user close to the CO/ONU overwhelms the signal of a far-end user. In the down-stream it occurs when the signal from an ONU overwhelms the signal from the CO. Both of these cases are illustrated in Figure 5.1.

The near-far problem has been well studied in wireless fields and is typically addressed using power back-off (PBO), or power control. With this technique, the base-station commands near-end users to reduce their transmit power so that the received power of all users is approximately equal.

Similar techniques have been proposed for DSL applications, which we refer to collectively as 'conventional PBO'. While these techniques achieve some success, there are significant differences between the wireline and wireless channels which must be taken into account. These include both the relatively stationary nature of the DSL channel and it's high frequency selectivity. To address these issues conventional PBO schemes are usually adjusted using some heuristic motivation. Whilst this allows for reasonable performance, many of these techniques lose their robustness as a result. To allow them to work well in most situations, conventional PBO schemes are typically designed using worst case scenarios. Whilst this ensures a high QOS for most of the Carrier Serving Area (CSA) it can be overly restrictive on near-end users, since crosstalk is rarely as bad as is predicted.

In this part of the report we investigate PBO in the DSL environment and it's relationship with crosstalk cancellation. Th rest of the report is structured as follows. In the next chapter we give an overview of conventional PBO techniques. We also introduce the information theory of multi-access channels (MAC). This describes the rates that can theoretically be achieved in a multi-user DSL system.

This sets the stage for consideration of more theoretically based PBO schemes, and in the following chapters we investigate optimal multi-user power allocation. As we will see, the optimal solution varies depending on whether we employ full, partial or no crosstalk cancellation and all these cases are investigated in turn.



(a) Upstream (b) Downstream

Figure 5.1: Near-Far Problem

5.1 Conventional PBO

Conventional PBO schemes are heuristically motivated ways of reducing power on near-end users so that they do not significantly degrade the performance of far-end users. In this section we give an overview of the conventional PBO schemes, specifically the: reference frequency method, reference length method, equalized FEXT method and the reference noise method. For a good overview in literature (similar to the one given here) see [12].

5.1.1 Reference Frequency Method

Under the reference frequency method, the transmit PSD of a user is decreased by a constant factor. The constant factor varies from user to user and is determined as follows. The constraint placed on the transmit PSD of a user, is that it's received power at a predefined frequency (the reference frequency) should match the received power of a 'reference loop' (which is typically the longest loop in the binder). In equations we can say

$$S_i(f) = \frac{|h_{rr}(f_r)|^2}{|h_{ii}(f_r)|^2} S_r(f_r)$$

where r denotes the reference user, and f_r the reference frequency.

5.1.2 Reference Length Method

In the reference length method, we require the received PSD of a user to match that of the 'reference loop' not only at a single (reference) frequency, but at all frequencies. This results in the following expression for the transmit PSD of user i

$$S_i(f) = \frac{|h_{rr}(f)|^2}{|h_{ii}(f)|^2} S_r(f)$$

The reference length method, as stated, leads to unnecessary performance degradation since it forces near-end users to decrease their PSDs at high frequencies to match the received PSDs of far-end users. Since far-end users do not use the higher tones (due to large attenuation) this is an unnecessary restriction. This issue is addressed in the multiple reference length method where we define several reference loops, each one corresponding to a particular frequency range.

5.1.3 Equalized FEXT Method

Instead of focusing on the received PSDs, the equalized FEXT method focuses on the FEXT that users induce into each other. It forces the FEXT that all users induce into each other to be equal to the FEXT induced by the reference loop. This results in the following expression for the transmit PSD

$$S_i(f) = \frac{L_r}{L_i} \frac{|h_{rr}(f)|^2}{|h_{ii}(f)|^2} S_r(f)$$

Here L_i is the length of user i's line, whilst L_r is the length of the reference line .

5.1.4 Reference Noise Method

The reference noise method is a generalization of the equalized FEXT method to any target FEXT. The FEXT that users induced into each other is made equal to some 'reference noise' which can be chosen by the designer. This results in the following expression for the transmit PSD

$$S_i(f) = \frac{\eta(f)}{L_i \, K_{\mathrm{FEXT}} \, f^2 \, |h_{ii}(f)|^2}$$

Here $\eta(f)$ is the reference noise, and K_{FEXT} a constant representing the coupling between lines (see [13] for details). If f is given in MHz and L_i in km's then $K_{\mathrm{FEXT}} = 10^{(-90/20)}$ in the ETSI models.

5.1.5 Summary

As you might have noticed, the conventional techniques for PBO are generally based on heuristical arguments and are not directly related to the data rates of the various users. Furthermore, the performance of each method varies dramatically with the situation (line lengths of users, crosstalk transfer functions etc.). There is no single scheme which performs well in all conditions. Since conventional PBO is based on worst case channel models, it can be overly restrictive and unnecessarily reduce near-end user rates.

This is one of the main motivations for looking at the multiple-access channel (MAC) from an information theory perspective. We can then design PBO schemes perform well in all situations and this is investigated further in chapter 6.

5.2 Rate Regions

A rate region is simply a graph representing the possible rate trade-offs that can be achieved in a MAC. By definition, any point inside the graph is achievable. The graph has N dimensions where N is the number of co-channel users. In this section we give expressions for the rate region of the MAC under both fixed (static) and freely chosen (dynamic) transmit power spectra.

We also examine successive interference cancellation (SIC) and show that this leads to the optimal receiver structure in the sense that, with proper coding, it can achieve any point on the boundary of the rate region[14].

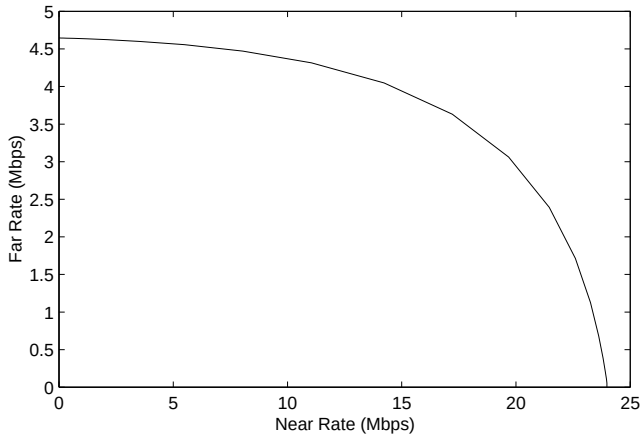


Figure 5.2: Example Rate Region

5.2.1 Static Power Allocation

To aid explanation we begin with a 2 user example. This is extended to an arbitrary number of users further on. Let us define our channel model as follows

$$\begin{aligned}\begin{bmatrix}\mathbf{y}_1\\ \mathbf{y}_2\end{bmatrix} &= \begin{bmatrix}\mathbf{H}_{11} & \mathbf{H}_{12}\\ \mathbf{H}_{21} & \mathbf{H}_{22}\end{bmatrix}\begin{bmatrix}\mathbf{x}_1\\ \mathbf{x}_2\end{bmatrix}+\begin{bmatrix}\mathbf{n}_1\\ \mathbf{n}_2\end{bmatrix}\\ &= \begin{bmatrix}\mathbf{H}_1 & \mathbf{H}_2\end{bmatrix}\begin{bmatrix}\mathbf{x}_1\\ \mathbf{x}_2\end{bmatrix}+\mathbf{n}\end{aligned} \tag{5.1}$$

Here $\mathbf{x}_i$ is a vector containing the transmitted (frequency domain) symbols for user i from all tones, $\mathbf{y}_i$ is the corresponding received vector and $\mathbf{H}_{ij}$ represents the transfer function from user j to user i. $\mathbf{H}_i \triangleq \begin{bmatrix}\mathbf{H}_{1i}\\ \mathbf{H}_{2i}\end{bmatrix}$ and contains both the direct and crosstalk channels from the transmitter of user i to all receivers. We use this later since, in the upstream, we have co-ordination between receivers and this leads to a form of receive diversity. Co-ordination between the receivers is possible since they are co-located in the CO.

In our case the users are synchronized so the $\mathbf{H}_{ij}$'s are diagonal matrices (no inter-carrier interference). The results shown below have very similar counterparts in the asynchronous case, but as one would expect the receiver structure becomes more complex.

Just to re-cap, here we are only concerned with the upstream communication channel. This corresponds to the MAC since co-ordination is possible at the receiver side. Whilst similar techniques can be applied to downstream communication, this corresponds to a broadcast rather than multiple access channel since co-ordination is possible between the transmitters only. The rate regions of broadcast channels are still being explored in literature and are currently unknown.

Rate regions for very simple MACs where first explored in [15] as early as 1971. This work was then extended to include frequency selective channels [16], channels with vector outputs [17] and channels with more than 2 users [18].

In our example we have a 2 user, vector MAC. The rate region is then defined as[17]

$$\mathbf{C}_{\text{static}} = \left\{ (R_1,R_2): \begin{array}{rl} R_1 \le I(\mathbf{x}_1;\mathbf{y}|\mathbf{x}_2) & = \frac{1}{2}\log_2\left|\mathbf{I}+\mathbf{H}_1\mathbf{S}_1\mathbf{H}_1^H\mathbf{N}^{-1}\right| \\ R_2 \le I(\mathbf{x}_2;\mathbf{y}|\mathbf{x}_1) & = \frac{1}{2}\log_2\left|\mathbf{I}+\mathbf{H}_2\mathbf{S}_2\mathbf{H}_2^H\mathbf{N}^{-1}\right| \\ R_1+R_2 \le I(\mathbf{x}_1,\mathbf{x}_2;\mathbf{y}) & = \frac{1}{2}\log_2\left|\mathbf{I}+\left(\mathbf{H}_1\mathbf{S}_1\mathbf{H}_1^H+\mathbf{H}_2\mathbf{S}_2\mathbf{H}_2^H\right)\mathbf{N}^{-1}\right| \end{array} \right\} \tag{5.2}$$

where $\mathbf{S}_i = \mathcal{E}\left\{\mathbf{x}_i\mathbf{x}_i^H\right\}$ (the auto-correlation of user i) and $\mathbf{N} = \mathcal{E}\left\{\mathbf{n}\mathbf{n}^H\right\}$ (the noise auto-correlation). $I(\mathbf{a};\mathbf{b})$ is the mutual information between vectors $\mathbf{a}$ and $\mathbf{b}$ and is defined as

$$I(\mathbf{a};\mathbf{b}) = H(\mathbf{a}) - H(\mathbf{a}|\mathbf{b})$$

The rate region is plotted in figure 5.3.

Note that this is simply a polyhedron (a pentagon in this case) which results from the application of the 3 constraints in (5.2) plus the requirement that R_1 and R_2 are both non-negative.

The extension of (5.2) to an arbitrary number of users results in the following expression for the rate region[18][19] which is also a polyhedron under static power allocations

$$\mathbf{C} = \left\{ (R_1,\ldots,R_N): \sum_{i\in\mathcal{S}} R_i \le \frac{1}{2}\log_2\frac{\left|\sum_{i\in\mathcal{S}}\mathbf{H}_i\mathbf{S}_i\mathbf{H}_i^H+\mathbf{N}\right|}{|\mathbf{N}|} \forall \mathcal{S}\subseteq\{1,\ldots,N\} \right\}$$

5.2.2 Relationship to Interference Cancellation

One nice aspect of the rate region is that it has an intuitive interpretation related to the optimal receiver structure for the MAC. It turns out that Successive Interference Cancellation (SIC) allows us to achieve points A and B in figure 5.3[14].

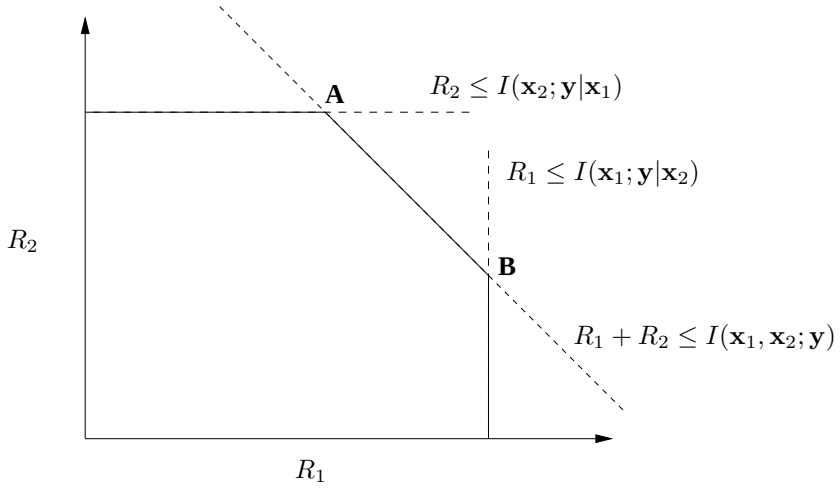


Figure 5.3: Two User Static Rate Region

Point A corresponds to detecting user 1 first, subtracting the interference from the received signal and then detecting user 2. Point B corresponds to detection in the reverse order.

To see this we examine the rates achieved by SIC when we detect user 1 first. User one is detected in the presence of the interference of user 2 and background noise.

$$\begin{aligned} R_1 &= \frac{1}{2}\log_2\left|\mathbf{I}+\mathbf{H}_1\mathbf{S}_1\mathbf{H}_1^H\left(\mathbf{N}+\mathbf{H}_2\mathbf{S}_2\mathbf{H}_2^H\right)^{-1}\right| \\ &= \frac{1}{2}\log_2\left|\mathbf{I}+\left(\mathbf{H}_1\mathbf{S}_1\mathbf{H}_1^H+\mathbf{H}_2\mathbf{S}_2\mathbf{H}_2^H\right)\mathbf{N}^{-1}\right| - \frac{1}{2}\log_2\left|\mathbf{I}+\mathbf{H}_2\mathbf{S}_2\mathbf{H}_2^H\mathbf{N}^{-1}\right| \qquad (5.3) \\ &= I(\mathbf{x}_1,\mathbf{x}_2;\mathbf{y}) - I(\mathbf{x}_2;\mathbf{y}|\mathbf{x}_1) \\ &= I(\mathbf{x}_1;\mathbf{y}) \qquad (5.4) \end{aligned}$$

Following this the signal from user 1 is subtracted from the received signal and we detect user 2 in the presence of background noise alone. Hence

$$\begin{aligned} R_2 &= \frac{1}{2}\log_2\left|\mathbf{I}+\mathbf{H}_2\mathbf{S}_2\mathbf{H}_2^H\mathbf{N}^{-1}\right| \qquad (5.5) \\ &= I(\mathbf{x}_2;\mathbf{y}|\mathbf{x}_1) \end{aligned}$$

Here $\mathbf{S}_1$ and $\mathbf{S}_2$ have Gaussian distributions (since this maximizes the mutual information) and we assume an arbitrarily small probability of error in the detection of both users (this is possible with appropriate coding since we operate at Shannon Capacity). Note that we are discussing the theoretically achievable rate region so we don't concern ourselves with practical issues like decision feedback errors in the SIC.

As can be seen from (5.3) and (5.5) using this detection order with SIC results in the rates of point A. Following a similar derivation, detecting the users in the reverse order results in the rates at point B. By time-sharing between the 2 detection orders we can achieve any point on the line A-B.

It should now be clear that under any power distribution $(\mathbf{S}_1, \mathbf{S}_2)$ the optimal receiver structure consists of SIC. This is also true with an arbitrary (>2) number of users[14].

5.2.3 Dynamic Power Allocation

In the previous section we characterized the rate region under static power allocation and showed that SIC is the optimal receiver structure allowing us to achieve any point on the boundary of the rate region. We now define the rate region

under dynamic power allocation, where we are free to set the transmit spectra of each user, provided we stay within some constraints on the total transmit power of each user.

The rate region under dynamic power allocation is simply the union of static rate regions for all valid power allocations. By 'valid power allocation' we refer to one that satisfies the total power constraint on each transmitter ie. $\mathrm{tr}(\mathbf{S}_i) \leq P_i \, \forall i$. Here $\mathrm{tr}(.)$ denotes the trace operator.

Using 5.2 we can define the dynamic rate region as

$$\mathbf{C} = \bigcup_{\substack{\mathbf{S}_1, \mathbf{S}_2 \\ \mathrm{tr}(\mathbf{S}_1) \leq P_1 \\ \mathrm{tr}(\mathbf{S}_2) \leq P_2}} \left\{ (R_1, R_2) : \begin{array}{c} R_1 \leq \frac{1}{2} \log_2 \left| \mathbf{I} + \mathbf{H}_1 \mathbf{S}_1 \mathbf{H}_1^H \mathbf{N}^{-1} \right| \\ R_2 \leq \frac{1}{2} \log_2 \left| \mathbf{I} + \mathbf{H}_2 \mathbf{S}_2 \mathbf{H}_2^H \mathbf{N}^{-1} \right| \\ R_1 + R_2 \leq \frac{1}{2} \log_2 \left| \mathbf{I} + \left(\mathbf{H}_1 \mathbf{S}_1 \mathbf{H}_1^H + \mathbf{H}_2 \mathbf{S}_2 \mathbf{H}_2^H \right) \mathbf{N}^{-1} \right| \end{array} \right\} \tag{5.6}$$

Many valid power allocations will not form part of the boundary of the dynamic rate region since they are extremely sub-optimal. As such, the union in (5.6) can be made more precise as we now show.

Each point on the boundary of the rate region corresponds to an optimisation. The cost function of this optimisation (to be maximised) is the weighted sum of the user's rates[16]. Changing the weights shifts the priority from one user to another and this allows us to traverse the boundary of the rate region. Using this we can express the rate region as follows

$$\mathbf{C} = \bigcup_{\alpha \in [0,1]} \left\{ (R_1, R_2) : \begin{array}{c} R_1 \leq \frac{1}{2} \log_2 \left| \mathbf{I} + \mathbf{H}_1 \mathbf{S}_1^*(\alpha) \mathbf{H}_1^H \mathbf{N}^{-1} \right| \\ R_2 \leq \frac{1}{2} \log_2 \left| \mathbf{I} + \mathbf{H}_2 \mathbf{S}_2^*(\alpha) \mathbf{H}_2^H \mathbf{N}^{-1} \right| \\ R_1 + R_2 \leq \frac{1}{2} \log_2 \left| \mathbf{I} + \left(\mathbf{H}_1 \mathbf{S}_1^*(\alpha) \mathbf{H}_1^H + \mathbf{H}_2 \mathbf{S}_2^*(\alpha) \mathbf{H}_2^H \right) \mathbf{N}^{-1} \right| \end{array} \right\} \tag{5.7}$$

where $\mathbf{S}_i^*(\alpha)$ is the optimal power allocation for user i using a weighting between users of α.

$$\begin{aligned} (\mathbf{S}_1^*(\alpha), \mathbf{S}_2^*(\alpha)) \quad &= \quad \arg \max_{\mathbf{S}_1, \mathbf{S}_2} \alpha R_1 + (1-\alpha) R_2 \\ &\text{s.t.} \quad \mathrm{tr}(\mathbf{S}_1) \leq P_1 \\ &\text{s.t.} \quad \mathrm{tr}(\mathbf{S}_2) \leq P_2 \end{aligned} \tag{5.8}$$

The resulting rate region is depicted in figure 5.4. As can be seen, the dynamic rate region is formed by a union of static rate regions. Furthermore, each static rate region typically intercepts the boundary of the dynamic rate region at one point. For $\alpha > 0.5$ priority is given to user 1 and this point corresponds to detecting user 1 last. For $\alpha < 0.5$ priority is given to user 2 and this point corresponds to detecting user 2 last. For $\alpha = 0.5$, both users are given equal priority, and both detection orders form part of the boundary.

Note that using (5.7) requires the rate region to be convex. Proof of the convexity of the rate region is given in [16].

Another way of visualising the rate region is as a polygon bounded by an infinite set of inequality constraints. Each inequality corresponds to a single value of α. Mathematically we can express this as

$$\mathbf{C} = \bigcap_{\alpha \in [0,1]} \left\{ (R_1, R_2) : \alpha R_1 + (1-\alpha) R_2 \leq \begin{array}{ll} \alpha I(\mathbf{x}_1; \mathbf{y}) + (1-\alpha) I(\mathbf{x}_2; \mathbf{y} | \mathbf{x}_1) & \alpha \leq 0.5 \\ \alpha I(\mathbf{x}_1; \mathbf{y} | \mathbf{x}_2) + (1-\alpha) I(\mathbf{x}_2; \mathbf{y}) & \alpha \geq 0.5 \end{array} \right\} \tag{5.9}$$

This is plotted in Figure 5.5. Each dashed line corresponds to an inequality. Notice that almost all inequalities (excluding the one corresponding to $\alpha = 0.5$) become tangent to the rate region boundary at exactly one point. In solving the optimisation problem (5.8), we effectively move these constraints as far out as possible. This in turn maximises our rate region.

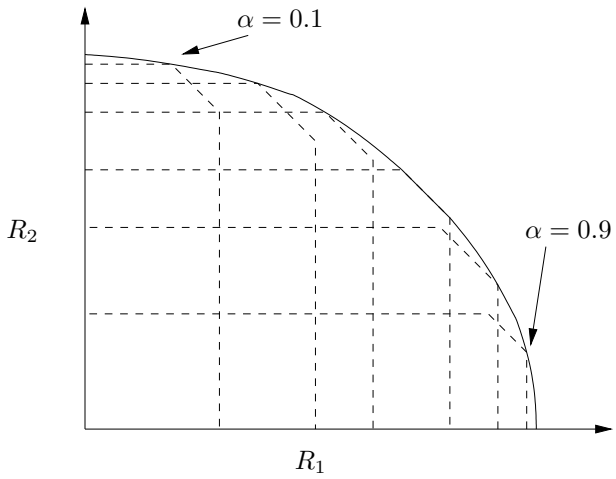


Figure 5.4: Dynamic Rate Region as an Optimisation of Weighted Rate-Sums

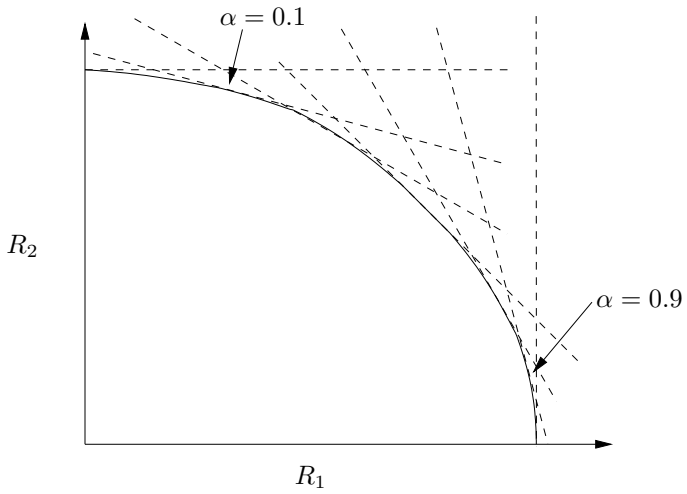


Figure 5.5: Dynamic Rate Region as a Polygon Bounded by Infinite Constraints

5.2.4 Summary

In summary, we have expressed the rate region in several ways, as the union of static rate regions (for all valid power allocations), as the solution to a weighted optimisation, and as a polyhedron bounded by infinite inequality constraints. We have presented these different expressions for the rate region in order to give some intuition behind it's construction. This will become important in the following chapters as we develop more optimal PBO schemes.

In section 5.2.2 we showed the relationship between successive interference cancellation (SIC) and the rate region. It was seen that SIC is the optimal receiver structure for the MAC and allows us to operate at any point on the boundary of the rate region.

One important aspect that we have not yet discussed is the solution to (5.8) which yields the optimal transmit PSDs for the MAC. In certain special cases (e.g. $\alpha = 0.5$ in frequency selective channels, arbitrary α in frequency flat channels) this optimisation is convex and has a known solution. It must be stressed however that in it's general form (5.8) has no closed form solution and is an open research problem. This issue is discussed further in following chapters.

5.3 Empirical Channel Models

In chapters 1 - 3 of this report we used transfer functions measured from actual cables. Due to the accuracy of measuring equipment, channel measurements were not reliable above 8 MHz and so we were forced to restrict our simulations to frequencies below this. Unfortunately, the effect this has on our results is quite considerable, since the majority of the gains derived from crosstalk cancellation, complexity reduction (ie. partial cancellation) and the more advanced PBO schemes occurs in the higher frequency ranges of the channel.

To overcome this, all simulations in the following chapters are done using empirical channel models. These models represent the worst 1% of cases and the performance of a system in this environment serves as a lower bound for more realistic scenarios.

The empirical models for direct channel and crosstalk transfer functions come from the ETSI standards [13]. All other simulation parameters are as in section 2.3.1 except that we now use tones up to 12 MHz and alien crosstalk model D.

Simulations use a network configuration of 4 near-end users and 4 far-end users at 300 m and 1200 m from the CO respectively.

Chapter 6

Power Allocation with Crosstalk Cancellation

We developed the rate region, in the previous chapter, to help us visualise the rate-tuples that are achievable in a multi-user DSL system. As we saw, operating on the boundary of this region corresponds to an optimisation of the weighted sum of user rates.

In this chapter, we investigate this optimisation problem in further detail.

In section 6.1 we show that solving the optimisation with non-equal weights translates to a non-convex problem. Thus we are forced to restrict our attention to the equal weight case, which corresponds to the section of the rate region boundary where the gradient is -1.

In section 6.1-6.3 we develop a power allocation scheme and receiver structure which together maximise the (unweighted) sum of rates.

The power allocation scheme, which is discussed in section 6.1, is a form of 'simultaneous waterfilling' and has a low-complexity iterative implementation as shown in section 6.2. The optimal receiver structure is the generalised decision feedback equaliser (GDFE) and achieves channel capacity with reasonable complexity as is described in section 6.3.

Due to the nature of the DSL channel, it is possible to achieve near-perfect crosstalk cancellation. We show in section 6.4 that this causes the rate-region to simplify to a rectangle in which the optimal operating point is the same, regardless of the weights assigned to each user's rate. This allows us to avoid the non-convex optimisation problem and we use knowledge of this near-perfect crosstalk cancellation in section 6.5 to formulate some very simple power allocation schemes and receiver structures which exhibit near optimal performance.

6.1 Optimal Power Allocation: Simultaneous Vector Waterfilling

As mentioned in section 5.2.3, operating on the border of the rate region requires us to solve a number of optimisation problems. These optimisation problems maximise a weighted sum of rates. Different weightings correspond to different points on the boundary, and each point has a corresponding power allocation. In this section we describe the solution to this optimisation problem.

When successive interference cancellation is employed, simultaneous vector waterfilling is the optimal power allocation scheme in the sense that it maximises the (unweighted) rate-sum[20]. The proof for this is contained in Appendix A.

This proof relies on the fact that the unweighted-rate sum is a concave function, and thus can be maximized through convex optimisation. Unfortunately, when $\alpha \neq 0.5$ this is not the case and we find that the optimisation function is non-convex.

To see this more clearly we first examine the unweighted cost function (this is equivalent to $\alpha = 0.5$)

$$J = R_1 + R_2 \tag{6.1}$$

This optimisation is constrained by the maximum transmit power of each user. However these power constraints are linear and hence can be included into the cost function without altering convexity[21]. This is done using Lagrange multipliers and is commonly known as the Kuhn-Tucker technique.

If we detect user 1 first

$$\begin{aligned} R_1 &= I(\mathbf{y};\mathbf{x}_1) \\ &= H(\mathbf{y}) - H(\mathbf{y}|\mathbf{x}_1) \\ R_2 &= I(\mathbf{y};\mathbf{x}_2|\mathbf{x}_1) \\ &= H(\mathbf{y}|\mathbf{x}_1) - H(\mathbf{y}|\mathbf{x}_1,\mathbf{x}_2) \end{aligned}$$

hence

$$\begin{aligned} J &= H(\mathbf{y}) - H(\mathbf{y}|\mathbf{x}_1,\mathbf{x}_2) \\ &= I(\mathbf{x}_1,\mathbf{x}_2;\mathbf{y}) \\ &= \frac{1}{2}\log\frac{\left|\mathbf{H}_1\mathbf{S}_1\mathbf{H}_1^H + \mathbf{H}_2\mathbf{S}_2\mathbf{H}_2^H + \mathbf{N}\right|}{|\mathbf{N}|} \end{aligned} \tag{6.2}$$

Note that equation (6.2) does not vary if we change the detection order.

Since the $\log|.|$ function is concave in the space of positive semi-definite matrices[21] we can use convex optimisation to find the optimum power allocations under the power constraint. Hence with equal weight given to the rates of all users ($\alpha = 0.5$) the optimum power allocation can be found.

When the rates of the users are un-evenly weighted we can no longer ensure the convexity of the cost function. The weighted cost function with two users can be expressed as

$$\begin{aligned} J &= \alpha R_1 + (1-\alpha)R_2 \\ &= \alpha H(\mathbf{y}) + (1-2\alpha)H(\mathbf{y}|\mathbf{x}_1) - (1-\alpha)H(\mathbf{y}|\mathbf{x}_1,\mathbf{x}_2) \\ &= \frac{\alpha}{2}\log\left|\mathbf{H}_1\mathbf{S}_1\mathbf{H}_1^H + \mathbf{H}_2\mathbf{S}_2\mathbf{H}_2^H + \mathbf{N}\right| + \frac{1-2\alpha}{2}\log\left|\mathbf{H}_2\mathbf{S}_2\mathbf{H}_2^H + \mathbf{N}\right| + \frac{1-\alpha}{2}\log|\mathbf{N}| \end{aligned} \tag{6.3}$$

Except in special cases ($\alpha \in \{0, 0.5, 1\}$) this is non-convex and extremely difficult to solve, especially when the number of dimensions (ie. tones) is high.

As shown in Appendix A, the solution to the unweighted optimisation problem (6.1) is simultaneous vector waterfilling. The word 'vector' implies that each user waterfills not against their direct channel alone, but against the composite channel consisting of every path (both direct and through crosstalk coupling) from their NT to the base station. In (5.1) we defined the direct channel for user i as $\mathbf{H}_{ii}$ and the composite channel as $\mathbf{H}_i$.

Thus vector waterfilling exploits the ability to transmit on both the direct and crosstalk channels. This is intuitively satisfying since both of these channels will be used by the crosstalk canceller in detecting the user's signal.

Single-user vector waterfilling is the optimal power allocation in a single-user channel with receive diversity. In this case the optimal transmit PSD satisfies

$$S^*(k) + \frac{1}{\mathbf{h}^H(k)\mathbf{N}(k)^{-1}\mathbf{h}(k)} = \frac{1}{\lambda} \tag{6.4}$$

Where $\frac{1}{\lambda}$ is the waterfilling level and is related to the total power constraint of the user.

In the multi-user scenario, waterfilling must occur against both the background noise and interference from other users. This results in a set of coupled equations, since the power allocation of each user depends on the interference experienced from, and hence the power allocation of, other users. The optimal transmit PSD for user i in the multi-user case is

$$S_i^*(k) + \frac{1}{\mathbf{h}_i^H(k)\left[\mathbf{N}(k) + \sum_{j\neq i}\mathbf{h}_j(k)\mathbf{h}_j^H(k)S_j^*(k)\right]^{-1}\mathbf{h}_i(k)} = \frac{1}{\lambda_i} \tag{6.5}$$

Here $\frac{1}{\lambda_i}$ is the waterfilling level for user i and varies with each user.

We must now solve this set of equations simultaneously for all users. A closed form expression for this would be highly complex, but fortunately there are simple, iterative solutions which can guarantee convergence to the global optimum. These iterative methods also exhibit rapid convergence and we discuss them further in the next section.

6.2 Optimal Power Allocation: Iterative Vector Waterfilling

In the previous section we showed that simultaneous vector waterfilling is the optimal power allocation when we desire to maximise the (unweighted) rate-sum. We also assumed that interference cancellation was used.

As we noted, the simultaneous waterfilling equations of (6.5) are quite complicated to solve. In this section we present a low-complexity, iterative algorithm which converges to simultaneous waterfilling[19]. Furthermore, the scheme is guaranteed to converge to the global optimum and does so rapidly[19].

Algorithm 1 Iterative Vector Waterfilling[19]

initialise $S_i(k) = 0 \ \forall i, k$
repeat
  for $i = 1 \, to \, N$

$$S_i(k) = \left[\frac{1}{\lambda_i} - \frac{1}{\mathbf{h}_i^H(k) \left[\mathbf{N}(k) + \sum_{j \neq i} \mathbf{h}_j(k) \mathbf{h}_j^H(k) S_j(k) \right]^{-1} \mathbf{h}_i(k)} \right]^+ \forall k$$

  end
until convergence

In algorithm 1 $[x]^+ \stackrel{\triangle}{=} \max(0, x)$. Recall that $\frac{1}{\lambda_i}$ is the waterfilling level for user i and is chosen such that $\sum_{\mathrm{k}} S_i(k) = P_i$.

The proof of convergence for this method proceeds as follows[19]. To begin with, iterative vector waterfilling is a non-decreasing operation with respect to the rate-sum. To see this examine Figure 6.1.

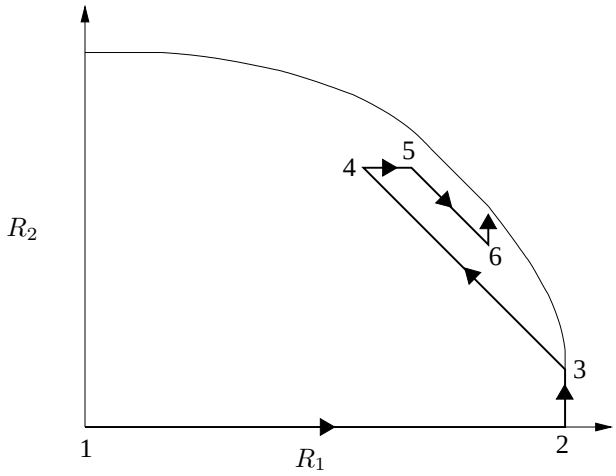


Figure 6.1: Convergence of Iterative Waterfilling

As the algorithm iterates we move through the following steps which are all non-decreasing w.r.t. the rate-sum.

1. Initially, both users are switched off and we have a rate-sum of zero.
2. Waterfill user 1 against the background noise.
3. User 2 is allocated power by waterfilling against the background noise and the interference of user 1. At this point we use interference cancellation with user 2 detected first. As a result, activating user 2 has no impact on the rate of user 1 (since we subtract the interference from user 2 before detecting user 1).
4. Change detection orders. This has no effect on the rate-sum since changing detection orders simply moves us along the 45 degree section of the pentagon defined in 5.2, that is we change vertices on the static rate region for the current power allocation.

5. Waterfill user 1 against the background noise and interference from user 2. This can only increase the rate of user 1 and has no effect on the rate of user 2 due to our new detection order.
6. Change detection orders.
7. Repeat 3-6.

This concludes the proof that iterative vector waterfilling is a non-decreasing operation. Now since the rate-sum is upper bounded (by the capacity of the MAC), the algorithm must converge. At convergence iterative waterfilling is equivalent to simultaneous waterfilling and is thus optimal.

6.3 Optimum Receiver Structure: Generalized DFE

In this section we describe the generalised decision feedback equaliser (GDFE) which achieves the theoretical MAC capacity while allowing low-complexity detection on a per-user basis.

As is frequently stated, the maximum-likelihood detector is the optimal receiver structure and allows us to achieve the full capacity (mutual information) of the MAC. Unfortunately such a structure is extremely complex since it must evaluate the likelihood of all possible transmit vector combinations. For a given tone, the vector $\mathbf{x}$ contains the transmitted symbols of every user. As a result we must evaluate the likelihood function over the $\prod_j 2^{b_j}$ different values that $\mathbf{x}$ can take, where b_j is the bitloading for user j on the tone of interest. This must be repeated on every tone which results in a completely unrealisable complexity.

The alternative to this is to use some form of crosstalk cancellation which removes the coupling between the received signals of different users. This allows us to perform detection on an individual basis and we only need to evaluate $\sum_j 2^{b_j}$ different possibilities using a conventional slicer.

Unfortunately, while many simple crosstalk cancellation structures (e.g. zero forcing) can remove the coupling between received signals, they reduce the capacity of the channel. So while we decrease complexity, we also decrease the data rate by doing processing on the received signal. This is known in information theory as the Data Processing Lemma. The entropy of a signal (and hence the mutual information between two signals) always decreases with deterministic processing. In equations

$$\begin{aligned} H(f(\mathbf{x})) &\leq H(\mathbf{x}) \\ I(f(\mathbf{x});\mathbf{y}) &\leq I(\mathbf{x};\mathbf{y}) \end{aligned}$$

where $f(\mathbf{x})$ is any deterministic function.

One structure which decouples the received signals, and hence allows for simple detection, without decreasing the capacity of the channel, is the generalised decision feedback (DF) equaliser (GDFE)[6]. Similar MMSE-DF schemes have also been proposed and shown to be optimal in terms of capacity[14]. For comparison, the 'Vectored Receiver' presented in section 2.1.3 is a zero forcing (ZF) version of the DF crosstalk canceller. The details of the implementation of the GDFE are given in section 2.1.4. A derivation is also given that shows that the GDFE achieves the capacity of the channel, whilst allowing independent detection of each user's signal.

The GDFE consists of a feed-forward (FF) and feed-back (FB) section. In the FF section the interference and noise prior to the decision is whitened (allowing us to achieve the capacity of the channel). In the FB section, the interference of users who have already been detected is subtracted (ie. SIC). The FB section allows us to implement detection on a user by user basis; avoiding the high complexity of a ML detector.

Note that the GDFE will achieve the full capacity of the channel, and is optimal for any choice of α in (6.4). The structure of the receiver does not change based on our weighting of the different users rates. As we change weightings (and priorities) between users we simply change the power allocation, and possibly the detection order, to achieve our desired operation point on a vertex of a particular static rate region. The optimal receiver structure, the GDFE, remains unchanged.

6.4 Rate Region of the Optimal Scheme

Using the optimal power allocation and receiver structure described in section 6.2 and 6.3 we can achieve a single section on the rate region where $\alpha = 0.5$ (ie. where the gradient of the rate region boundary is -1). This corresponds to maximising the unweighted rate-sum.

We can also achieve two other points, corresponding to switching one of the users off and applying single-user vector waterfilling (see eq. (6.4)) to the other user.

Since (6.4) is non-convex, we cannot obtain optimal power allocations for weighted rate-sums when $\alpha \notin \{0, 0.5, 1\}$. This means that we cannot operate at all points on the rate region boundary.

In this section we formulate an upper and lower bound on the rate region of the DSL MAC. By showing that these bounds are virtually equal, we see that the rate region is virtually rectangular. This indicates that nearly perfect crosstalk cancellation is achieved by the GDFE. The most important aspect of this result, is that a single point on the rate region (the upper right hand corner) is near-optimal regardless of the weights used in the optimisation function. This allows us to avoid the non-convex optimisation and simply maximise the unweighted sum of rates regardless of our desired rate distribution.

We begin with the development of the upper bound. Exploiting the convexity of the rate region defined in (5.9) (all polygons are indeed convex[21]), we can define an upper bound on the rate region

$$\mathbf{C} \subseteq \mathbf{C}_{\text{upper bound}} = \left\{ (R_1, R_2) : \begin{array}{c} R_1 \leq \frac{1}{2}\log_2 \left|\mathbf{I} + \mathbf{H}_1 \mathbf{S}_1^{\text{single}} \mathbf{H}_1^H \mathbf{N}^{-1}\right| \\ R_2 \leq \frac{1}{2}\log_2 \left|\mathbf{I} + \mathbf{H}_2 \mathbf{S}_2^{\text{single}} \mathbf{H}_2^H \mathbf{N}^{-1}\right| \\ R_1 + R_2 \leq \frac{1}{2}\log_2 \left|\mathbf{I} + \left(\mathbf{H}_1 \mathbf{S}_1^{\text{sim}} \mathbf{H}_1^H + \mathbf{H}_2 \mathbf{S}_2^{\text{sim}} \mathbf{H}_2^H\right) \mathbf{N}^{-1}\right| \end{array} \right\} \tag{6.6}$$

where $\mathbf{S}_i^{\text{single}}$ is the power allocation for user i with single-user vector waterfilling as defined in (6.4). $\mathbf{S}_i^{\text{sim}}$ is the power allocation for user i with simultaneous vector waterfilling as defined in (6.5). This can be calculated efficiently using the iterative waterfilling detailed in algorithm 1. Note that the upper bound in (6.6) is just (5.9) with the intersection limited to $\alpha \in \{0,\ 0.5,\ 1\}$.

We now substitute single-user waterfilling with single class waterfilling. Thus (6.6) can be replaced with

$$\mathbf{C} \subseteq \mathbf{C}_{\text{upper bound}} = \left\{ (R_1, R_2) : \begin{array}{c} R_{\text{near}} \leq \frac{1}{2}\log_2 \left|\mathbf{I} + \mathbf{H}_{\text{near}} \mathbf{S}_{\text{near}}^{\text{single}} \mathbf{H}_{\text{near}}^H \mathbf{N}^{-1}\right| \\ R_{\text{far}} \leq \frac{1}{2}\log_2 \left|\mathbf{I} + \mathbf{H}_{\text{far}} \mathbf{S}_{\text{far}}^{\text{single}} \mathbf{H}_{\text{far}}^H \mathbf{N}^{-1}\right| \\ R_{\text{near}} + R_{\text{far}} \leq \frac{1}{2}\log_2 \left|\mathbf{I} + \left(\mathbf{H}_{\text{near}} \mathbf{S}_{\text{near}}^{\text{sim}} \mathbf{H}_{\text{near}}^H + \mathbf{H}_{\text{far}} \mathbf{S}_{\text{far}}^{\text{sim}} \mathbf{H}_{\text{far}}^H\right) \mathbf{N}^{-1}\right| \end{array} \right\} \tag{6.7}$$

Here $\mathbf{S}_{\text{near}}^{\text{single}}$ denotes the power allocation for near-end users, when we de-activate far-end users and do simultaneous waterfilling on the near-end users only. A similar definition applies for $\mathbf{S}_{\text{far}}^{\text{single}}$ except in this case we de-activate the near-end users and do simultaneous waterfilling on the far-end users only.

R_{near} and R_{far} denote the average rate achieved by the near-end and far-end users respectively. So we can replace the 2 user rate region with a 2 class rate region. Note that maximizing the average rate of a set of identical users (a single class), corresponds to maximising the unweighted sum of rates. Hence simultaneous vector waterfilling is optimal in the sense that it will maximise the average rate of a single class.

To form an upper bound we now proceed as follows:

- Do iterative vector waterfilling on far-end users with near-end users de-activated. Denote resulting operating point (in the rate region) as $A \triangleq (0, r_{\text{far}}^a)$
- Do iterative vector waterfilling with all users activated. Denote resulting operating point when the near-end users are detected first as $C \triangleq (r_{\text{near}}^c, r_{\text{far}}^c)$. Denote the resulting operating point whn the far-end users are detected first as $D \triangleq \left(r_{\text{near}}^d, r_{\text{far}}^d\right)$.

- Do iterative vector waterfilling on near-end users with far-end users de-activated. Denote resulting operating point as $F \triangleq \left(r^f_{\text{near}}, 0\right)$
- Now construct the following points $B \triangleq (r^c_{\text{near}}, r^a_{\text{far}})$ and $E \triangleq \left(r^f_{\text{near}}, r^d_{\text{far}}\right)$
- Since the true rate region passes through points A and C, and since it is convex, it must lie below the line $\overline{AB}$ and above the line $\overline{AC}$.
- Since the true rate region passes through points D and F, and since it is convex, it must lie inside the line $\overline{EF}$ and outside the line $\overline{DF}$.
- Hence we can form an upper-bound on the rate region by connecting points $ABCDEF$.
- Furthermore, we can form a lower-bound on the rate region by connecting points $ACDF$.

This development is depicted in figure 6.2.

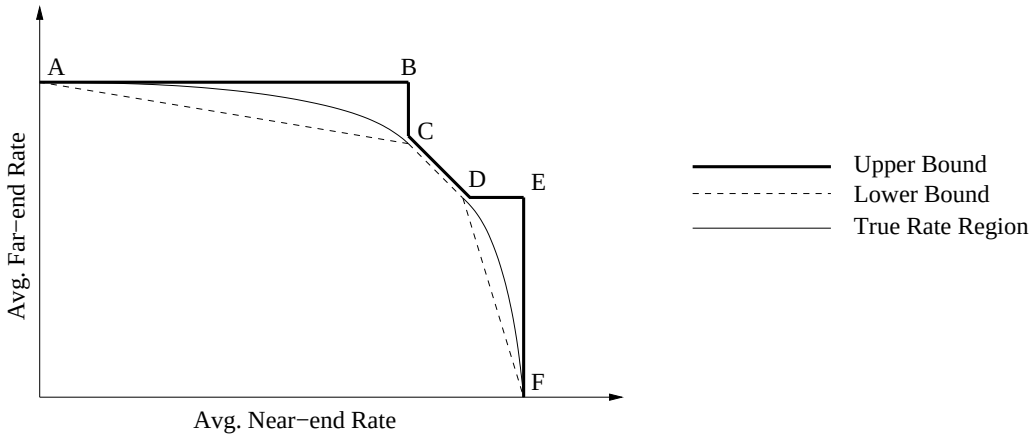


Figure 6.2: Upper and Lower Bounds of the Rate Region

The upper and lower bounds just described where formed using the VDSL simulation environment. The upper-bound is depicted in figure 6.3. As can be seen, due to the properties of the VDSL channel, points B, C, D, and E virtually co-incide. This means that our upper and lower-bound are extremely tight (they also virtually co-incide). Hence the true rate region can be approximated extremely well by the lower-bound $ACDF$.

Using this observation we conclude that the true rate region is virtually rectangular. This indicates that our crosstalk canceller can successfully remove virtually all of the crosstalk. Increasing the rates of the near-end users has a negligible effect on the rates of the far-end users and vice-versa.

Thanks to this, we only need to concern ourselves with operating at a single point on the rate region, the upper right hand corner. This point is near-optimum for any positive choice of weights. As such, we can use iterative waterfilling to achieve near-optimum performance. Note that we are not inhibited by our inability to solve the weighted optimisation when $\alpha \notin \{0, 0.5, 1\}$.

6.5 Rate Region of the Near-Optimal Scheme

In the previous section we noticed that crosstalk cancellation (in a DSL environment) achieves near perfect results. This suggests that we can replace iterative waterfilling with single-user (noise-only) waterfilling and achieve near-optimal results. Furthermore, we can replace the GDFE with a much simpler linear MMSE crosstalk canceller with negligible impact.

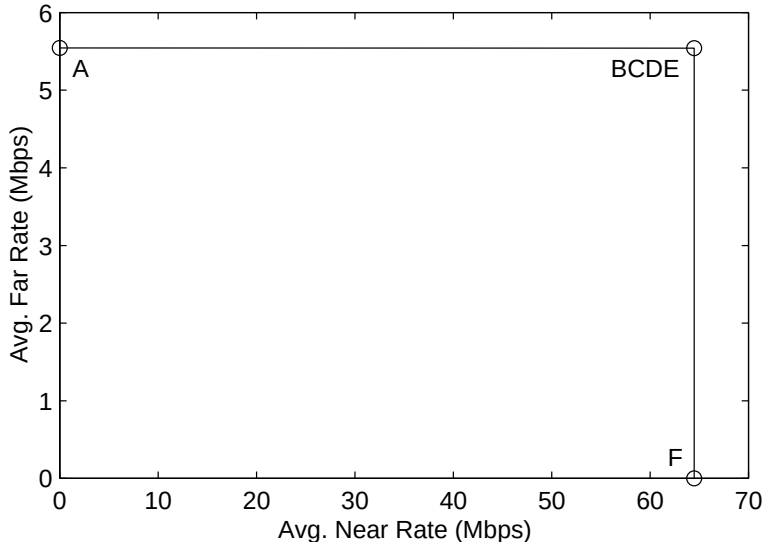


Figure 6.3: Upper Bound on Rate Region

This fact was first noted in [2]. As a result, we can considerably reduce the complexity of both our power allocation scheme and receiver structure. In a sense the crosstalk cancellation is so perfect in the DSL environment, that we do not need to perform PBO on the near-end users.

Another aspect of the DSL environment is that a user's direct channel always exceeds their crosstalk channels that couple *into* other user's lines. This means that we can neglect the vector aspect of the channel when performing waterfilling, again with minimal performance reduction. In terms of the composite channel we defined in (5.1), we can say that the difference between the Frobenius norm of the composite channel $\mathbf{H}_i$ and the direct channel $\mathbf{H}_{ii}$ is negligible.

Using these observations, we replace the vector waterfilling in (6.4) with noise-only scalar waterfilling, defined as

$$S_i^*(k) = \left[\frac{1}{\lambda_i} - \frac{\sigma_{n_i}^2(k)}{|h_{ii}(k)|^2}\right]^+$$

Here $\sigma_{n_i}^2(k)$ is the noise power on user i's line at tone k. Note that in our definition of noise-only waterfilling, we explicitly ignore the interference term and waterfill only against the background noise.

Shown in figure 6.4 is the rate region achieved when using noise-only scalar waterfilling and a linear MMSE crosstalk canceller. As can be seen, the scheme provides near-optimal performance with considerably reduced complexity. Also included is the performance of a conventional PBO scheme, in this case the equalised FEXT method (with and without crosstalk cancellation) for comparison.

Listed in table 6.1 are the average rates of near-end and far-end users with the optimal (iterative vector waterfilling + GDFE) and near optimal (conv. scalar waterfilling + linear MMSE) schemes. The upper bounds on the near-end/far-end rates (iterative water-filling on near-end/far-end only + GDFE) are also shown.

To give an overview we compare the rates achieved by the different schemes in table 6.2. As can be seen, without PBO the far-end users' lines become unusable. The equalised FEXT method improves the situation slightly, but has to reduce the rates of the near-end users to do this. Crosstalk cancellation yields relatively minor improvement to the performance of the equalised FEXT method. Notice the considerable gains (5.9->64.4 Mbps on near-end users) that can be achieved using the proposed power allocation scheme with crosstalk cancellation.

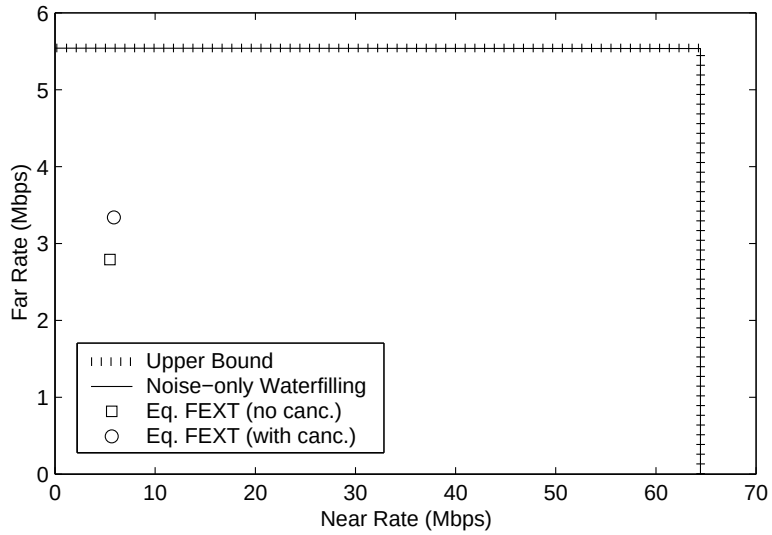


Figure 6.4: Rate Region using Near-Optimal Scheme

Users Active	Power Allocation	Crosstalk Canc.	Avg. Rate (Mbps)	
			Far-End	Near-End
Far-End Only	It. Vector Water.	GDFE	5.544	0
Near-End Only	It. Vector Water.	GDFE	0	64.467
Upper Bound			5.544	64.467
All	It. Vector Water.	GDFE	5.541	64.466
All	Noise-only Scalar Water.	Linear MMSE	5.537	64.448

Table 6.1: Comparison of Upper Bound, Optimal and Near-Optimal Schemes

6.6 Summary

In this section we examined in more depth the role of the cost function in defining the rate region for a vector MAC. The cost function itself took two forms, an unweighted- ($\alpha = 0.5$) and a weighted rate-sum.

We showed in section 6.1 that for the unweighted rate-sum, the optimal power allocation is simultaneous vector waterfilling. Furthermore, we showed that iterative vector waterfilling is a low-complexity iterative method which is guaranteed to converge to the simultaneous solution. Unfortunately, when the cost function is a weighted rate-sum, the optimisation becomes non-convex and solutions are highly complex.

The optimum receiver structure is the generalized decision feedback equalizer. When combined with optimal power allocation this has been shown[6] to achieve the mutual information of the vector MAC.

In sections 6.4 and 6.5 we evaluated the rate regions with the optimal power allocation scheme and receiver structure. We saw that, due to the nature of the DSL channel, near-perfect crosstalk cancellation can be achieved. Hence very simple power-allocation and crosstalk cancellation schemes achieve near-optimum performance. Additionally we noticed the considerable gains in rates that can be achieved (5.9->64.4 Mbps on near-end users) by using waterfilling based schemes instead of conventional PBO.

Power Allocation	Crosstalk Canc.	Avg. Rate (Mbps)	
		Far-End	Near-End
None	None	0.0551	23.982
Equalised FEXT	None	2.792	5.486
Equalised FEXT	Linear MMSE	3.339	5.880
Noise-only Scalar Water.	Linear MMSE	5.537	64.448

Table 6.2: Comparison of Conventional PBO with Near-Optimal Scheme

Chapter 7

Power Allocation without Crosstalk Cancellation

In chapter 6 we examined optimal power allocation schemes when crosstalk cancellation is used at the receiver. Whilst this led to simple optimisation problems and low complexity receivers, we frequently encounter situations where we cannot do crosstalk cancellation. This is the case when we cannot afford to increase run-time complexity. It also occurs when we operate in an unbundled environment.

For these reasons we now investigate power allocation in the absence of crosstalk cancellation.

This chapter is organised as follows. First we re-evaluate the rate region in the absence of crosstalk cancellation and it's relationship to a weighted optimisation. This is similar to the definition in section 5.2.3, however now the optimisation becomes non-convex, even when the rates are weighted equally. As a result of this we focus on sub-optimal power allocation schemes in the rest of the chapter. We examine two sub-optimal schemes for power allocation without cancellation.

The first was proposed by Cioffi et. al[22] and consists of iterative scalar waterfilling. Although this scheme performs well without crosstalk cancellation, in the presence of crosstalk cancellation it is considerably sub-optimal.

The second scheme is our suggestion and offers similar performance (in the absence of cancellation) with reduced complexity. Additionally, our scheme has near-optimal performance when crosstalk cancellation is present. This makes it general enough to be used in partially unbundled scenarios or with partial crosstalk cancellation.

7.1 Rate Region

The rate region in the absence of crosstalk cancellation can be defined similar to that in (5.2). Since we only use each user's direct line during detection, we can no longer exploit the receive diversity and hence the upper bound on the rate for a user becomes $I(\mathbf{x}_i;\mathbf{y}_i)$ instead of $I(\mathbf{x}_i;\mathbf{y})$. Furthermore, since we do not perform interference subtraction, no user can be detected with knowledge of another user's signal. This means that terms such as $I(\mathbf{x}_1;\mathbf{y}_1|\mathbf{x}_2)$ become $I(\mathbf{x}_1;\mathbf{y}_1)$.

Keeping these points in mind, the rate region is now defined as

$$\mathbf{C} = \bigcup_{\substack{S_1(k),S_2(k) \\ \sum_k S_1(k) \le P_1 \\ \sum_k S_1(k) \le P_2}} \left\{ (R_1,R_2): \begin{array}{l} R_1 \le I(\mathbf{x}_1;\mathbf{y}_1) = \frac{1}{2}\sum_{k=1}^{K} \log_2\left(1+\frac{|h_{11}(k)|^2 S_1(k)}{\sigma^2_{n_1}(k)+|h_{12}(k)|^2 S_2(k)}\right) \\ R_2 \le I(\mathbf{x}_2;\mathbf{y}_2) = \frac{1}{2}\sum_{k=1}^{K} \log_2\left(1+\frac{|h_{22}(k)|^2 S_2(k)}{\sigma^2_{n_2}(k)+|h_{21}(k)|^2 S_1(k)}\right) \end{array} \right\}$$

As in section 5.2.3, the boundary points of this rate-region can be achieved by allocating power according to a weighted rate-sum optimisation problem. The cost function for this optimisation is now

$$\begin{aligned} J &= \alpha R_1 + (1-\alpha) R_2 \\ &= \frac{1}{2}\sum_{k=1}^{K} \alpha \log_2 \left(1 + \frac{|h_{11}(k)|^2 S_1(k)}{\sigma_{n_1}^2(k) + |h_{12}(k)|^2 S_2(k)}\right) + (1-\alpha)\log_2\left(1 + \frac{|h_{22}(k)|^2 S_2(k)}{\sigma_{n_2}^2(k) + |h_{21}(k)|^2 S_1(k)}\right) \end{aligned}$$

Unfortunately, unlike (6.4), this function is non-convex even when the rates are weighted equally. Furthermore, since we now longer have crosstalk cancellation, the rate region is not rectangular and hence the problem must be solved for all values of $\alpha \in [0, 1]$.

Due to the non-convexity of this optimisation problem we restrict our attention to sub-optimal, low-complexity solutions in the following sections.

7.2 Iterative Scalar Waterfilling

To address the issue of power allocation in the absence of crosstalk cancellation, Cioffi et. al[22] proposed iterative scalar waterfilling. The advantages of this scheme are it's ability to be implemented in a distributed manner, it's simplicity and the considerable performance gains it yields over conventional PBO methods. The ability to be implemented in an, at least partially distributed manner makes this scheme particularly suited to unbundled environments where there is no centralised control.

Since iterative waterfilling yields only one point on the rate region, this algorithm was modified to include an outer loop which allows us to traverse the rate region, trading-off rates between near-end and far-end users. Recall that due to the absence of crosstalk cancellation the rate region is not rectangular. As such, we desire to characterise all operating points on the region boundary in addition to their corresponding power allocations.

The outer loop works by modifying the total power constraints on each user and operates as follows[22]. If a user is found to significantly exceed their (desired) target rate, their total power constraint is reduced. If a user operates below their target rate then their total power constraint is increased. There is also a set of absolute power constraints for each modem which may never be exceeded. These reflect the maximum output power that each modem itself can support.

The outer loop of the algorithm continues until all target rates are achieved. Naturally, it must be known a-priori which target rates are achievable, or the algorithm will never converge. This must be decided by some centralised agent, and hence the algorithm cannot be implemented in a completely de-centralised manner.

Define T_i as the target rate for user i and P as the total power constraint for each modem. The algorithm is listed in algorithm 2.

Although this algorithm performs reasonably well in the absence of crosstalk cancellation, when we apply it in an environment with crosstalk cancellation it is significantly sub-optimal. When crosstalk cancellation is present, we compare this scheme with the optimal, iterative vector waterfilling defined in algorithm 1 to determine the degree of sub-optimality. The results are shown in table 7.1 and the performance of noise-only scalar waterfilling, where we waterfill against the background noise only, is included for comparison.

Power Allocation	Avg. Rate (Mbps)	
	Far-End	Near-End
Equalized FEXT	3.339	5.880
It. Scalar Water.	1.669	64.449
It. Vector Water.	5.537	64.448
Noise-only Scalar Water.	5.537	64.448

Table 7.1: Comparison of Power Allocation Schemes with Crosstalk Cancellation (Linear MMSE)

It is interesting to note that whilst iterative vector waterfilling (against the background noise and interference) and noise-only scalar waterfilling perform quite well, iterative scalar waterfilling, where we waterfill against the interference of other

Algorithm 2 Iterative Scalar Waterfilling**[22]**

initialise $P_i = P \quad \forall i$
repeat
  initialise $S_i(k) = 0 \quad \forall i, k$
  repeat
    for $i = 1\,to\,N$
      $S_i(k) = \left[\frac{1}{\lambda_i} - \frac{\sum_{j \neq i} |h_{ij}(k)|^2 S_j(k) + \sigma^2_{n_i}(k)}{|h_{ii}(k)|^2}\right]^+ \quad \forall k$
      set R_i *to be the resulting data rate*
    end
  until convergence
  for $i = 1\,to\,N$
    if $R_i > T_i + \epsilon$, *set* $P_i = P_i - \delta$
    if $R_i < T_i$, *set* $P_i = P_i + \delta$
    if $P_i > P$, *set* $P_i = P$
  end
until $R_i > T_i$ *for all* i

users but against the direct channel only, exhibits significantly sub-optimal performance when crosstalk cancellation is present.

In the absence of crosstalk cancellation, all schemes have quite similar performance as shown in table 7.2.

These results where generated using the scenario described in section 5.3. With crosstalk cancellation all users were given full power. In the absence of crosstalk cancellation the near-end users had their total power constraint decreased by 23 dB.

Power Allocation	Avg. Rate (Mbps)	
	Far-End	Near-End
Equalized FEXT	2.792	5.486
It. Scalar Water.	2.774	20.379
It. Vector Water.	2.801	20.390
Noise-only Scalar Water.	2.800	20.390

Table 7.2: Comparison of Power Allocation Schemes in Absence of Crosstalk Cancellation

Whilst the performance of these schemes with full crosstalk cancellation might appear to be of little consequence, this is not the case. For example, in a partially unbundled environment, an operator may have crosstalk cancellation on their in-domain lines, but cannot perform crosstalk cancellation on out-of-domain lines where another competitive carrier (CLEC) is operating. This means we need a general power allocation and PBO scheme which is applicable to environments both with and without crosstalk cancellation. A similar arguement applies to scenarios where, to reduce computational complexity, we only cancel dominant crosstalkers. We explore partial cancellation further in chapter 8.

As can be seen from the results in tables 7.1 and 7.2, iterative vector waterfilling overcomes the problems of iterative scalar waterfilling and is applicable in environments with or without crosstalk cancellation. Unfortunately, iterative vector waterfilling has an increased complexity and cannot be applied in a distributed manner.

On the other hand, the noise-only scalar waterfilling scheme which we introduced in section 6.5 has lower complexity than the iterative scalar waterfilling proposed in [22], only requires centralised power control of in-domain lines, and performs well with or without crosstalk cancellation. This makes it applicable to the partially unbundled and partial cancellation scenarios.

In the following section we give a more detailed look at the noise-only scalar waterfilling scheme and it's combination with PBO.

7.3 Noise-Only Scalar Waterfilling with PBO

In section 6.5 we saw that with full crosstalk cancellation, noise-only scalar waterfilling was near-optimal. Furthermore, since the crosstalk is almost completely removed, PBO was unnecessary on the near-end users.

In the absence of crosstalk cancellation this is no longer the case and it is still necessary to perform PBO on near-end users. The degree of PBO determines the trade-off between near-end and far-end rates and this forms a rate region which, in contrast to the case with crosstalk cancellation, is no longer rectangular.

Despite the necessity for PBO, we will see that a simple scheme, consisting of constant PBO on each user gives virtually identical performance to iterative scalar waterfilling. This performance is considerably better than conventional PBO schemes. Furthermore, the scheme we propose here has a much lower complexity than iterative scalar waterfilling and is applicable to partially unbundled and partial cancellation scenarios as mentioned in the previous section.

The algorithm is listed in algorithm 3.

Algorithm 3 Noise-only Scalar Waterfilling

set $S_i^{\text{noise only}}(k) = \left[\frac{1}{\lambda_i} - \frac{\sigma_{n_i}^2(k)}{|h_{ii}(k)|^2}\right]^+ \quad \forall i,k$
initialise $\kappa_i = 1 \ \forall i$
repeat
  set $S_i(k) = \kappa_i S_i^{\text{noise only}}(k) \ \forall i,k$
  set R_i *to be the resulting data rate*
  for $i = 1 \, to \, N$
    if $R_i > T_i + \epsilon$, *set* $\kappa_i = \kappa_i - \delta$
    if $R_i < T_i$, *set* $\kappa_i = \kappa_i + \delta$
    if $\kappa_i > 1$, *set* $\kappa_i = 1$
    if $\kappa_i < 0$, *set* $\kappa_i = 0$
  end
until $R_i > T_i$ *for all* i

As can be seen the scheme is quite simple. Each user's power allocation is simply noise-only scalar waterfilling, with a constant PBO factor κ_i . Varying the κ_i's allows us to traverse the rate region, trading-off near-end rates against far-end rates. This scheme is approximately equivalent to taking convex combinations between the power allocations for single-user scalar waterfilling, and iterative scalar waterfilling.

Shown in figure 7.1 are the rate regions for iterative scalar waterfilling and the noise-only waterfilling with constant PBO scheme we propose here. As can be seen the performance of both schemes is virtually identical.

7.4 Summary

In this chapter we examined power allocation schemes in the absence of crosstalk cancellation. In section 7.1 we discussed the rate region in the absence of crosstalk cancellation, and showed that defining it's boundary results in a non-convex optimisation. Due to this, we restricted our attention to low-complexity, sub-optimal schemes in sections 7.2 and 7.3.

In section 7.2 we gave an overview of a scheme proposed by Cioffi et. al. in [22]. Whilst this scheme has a simple implementation, and gives large performance gains over conventional PBO methods ($5.5 \rightarrow 20.4$ Mbps on near-end users), we found it to be considerably sub-optimal when applied to environments with crosstalk cancellation. The inability of this power allocation scheme to be applied to both scenarios with and without crosstalk cancellation limits it's applicability to unbundled environments, and scenarios where we only cancel dominant interferers to reduce implementational complexity.

To address these issues, in section 7.3 we presented a scheme based on noise-only scalar waterfilling with constant PBO. We saw earlier, in section 6.5 that this scheme is near-optimal in scenarios with full crosstalk cancellation. Furthermore, it exhibits virtually identical performance to iterative scalar waterfilling in the absence of crosstalk cancellation, with considerably reduced complexity.

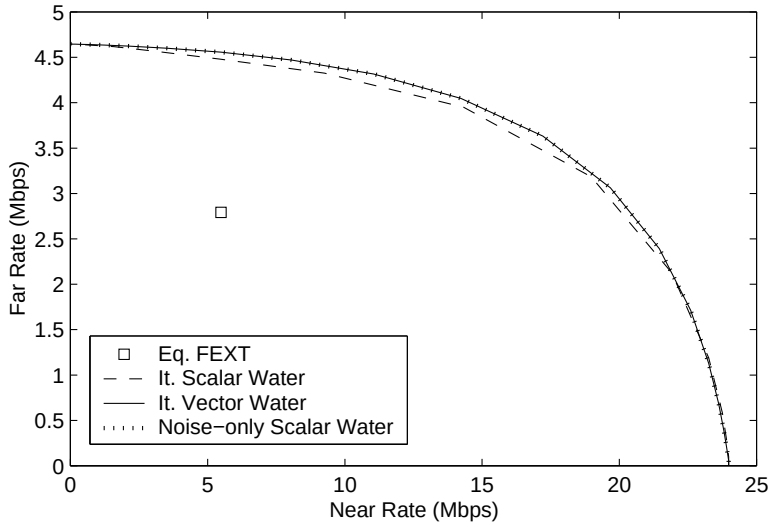


Figure 7.1: Rate Regions of Power Allocation Schemes in the Absence of Crosstalk Cancellation

Chapter 8

Power Allocation and Partial Cancellation

In the previous chapters we focused our attention on two problems, that is, power allocation with full crosstalk cancellation and without crosstalk cancellation. These situations are in fact, extremes of a more general case where we perform partial crosstalk cancellation.

Partial crosstalk cancellation can occur in a number of ways. We may wish to reduce the computational complexity at the receiver. By only cancelling dominant interferers, we can gain most of the benefits of crosstalk cancellation, whilst significantly reducing complexity. Due to local loop unbundling it is also possible that several competing operators may use different lines within the same binder. This means that co-ordination is only possible between lines belonging to the same operator (in-domain) and not between lines of different operators (out-of-domain).

Unfortunately, the optimal power allocation in a partial cancellation scenario is coupled with the choice of which crosstalkers to be considered as dominant. This coupling is extremely complicated and makes the optimal choice of power allocation and dominant crosstalkers difficult.

For this reason, we focus on sub-optimal methods in this section. We present a general scheme for allocation of computing resources across users, tones and interferers. By treating computing power as a resource to be allocated (in a similar manner to transmit power) we show that this scheme achieves high performance while considerably reducing the run-time complexity.

8.1 Resource Allocation

In chapter 3 we made some initial investigations into partial cancellation schemes. The schemes where rather arbitrary however, in that we either cancelled all interferers on a sub-set of tones (sec. 3.2.1) or a sub-set of interferers on all tones (sec. 3.2.2).

We now present a scheme which allows the choice of dominant (tone, interferer) pairs to be made freely on a per user basis. The division of available computing resources between users can also be adjusted to give more priority to users who suffer the most from crosstalk.

Using a linear MMSE crosstalk canceller (see sec. 2.1.2), cancelling one interferer, from a single user's signal, at one tone corresponds to a single multiplication. Put another way, when we perform partial crosstalk cancellation, the crosstalk cancelling filter $\mathbf{W}(k)$ is a sparse matrix. The number of non-zero elements in $\mathbf{W}(k)$ corresponds to the number of multiplications per received DMT-block for tone k.

Let us arbitrarily set the total number of available multiplications (per received DMT-block) to

$$cMN$$

where M is the number of tones, N the number of users in a common binder and c a freely chosen parameter. Using this definition, we can limit the complexity of crosstalk cancellation to c times that of a standard single-user frequency domain equaliser (FDE).

To allow the division of computing resources between users to be tuned, we define a parameter μ_i which determines the proportion of total multiplications allocated to user i. The number of multiplications for user i is

$$\text{mults}_i = \mu_i cMN \quad \textstyle\sum_i \mu_i = 1$$

At this point we have divided multiplications between users and we now face the following problem, given that we have a fixed number of multiplications available for a certain user, how do we allocate these multiplications across tones and interferers.

8.2 (Tone, Interferer) Selection

As we mentioned, the optimal solution to the computing resource allocation problem would be highly complex and this arises from the high dimensional, combinatorial nature of the problem. So in considering the allocation of resources across tones and interferers we opt for a heuristic choice.

Let us define the usefulness or utility of cancelling a given interferer on a given tone, as the increase in rate it would cause to the user, if all other interferers had already been cancelled. Hence we can define a utility measure as a function of user, tone and interferer

$$\begin{aligned} \partial R^i_{k,j} &= \log\left(1+\frac{\text{signal}_i(k)}{\text{noise}_i(k)}\right) - \log\left(1+\frac{\text{signal}_i(k)}{\text{interference}_{i,j}(k)+\text{noise}_i(k)}\right) \\ &= \log\left(\frac{(\text{signal}_i(k)+\text{noise}_i(k))\,(\text{interference}_{i,j}(k)+\text{noise}_i(k))}{\text{noise}_i(k)\,(\text{signal}_i(k)+\text{interference}_{i,j}(k)+\text{noise}_i(k))}\right) \end{aligned}$$

where i is the user of interest, j the interferer, and k the tone. We also define the following

$$\begin{aligned} \text{signal}_i(k) &= |h_{ii}(k)|^2 S_i(k) \\ \text{interference}_{i,j}(k) &= \Gamma\,|h_{ij}(k)|^2 S_j(k) \\ \text{noise}_i(k) &= \Gamma\sigma^2_{n_i}(k) \end{aligned}$$

Here Γ represents the SNR gap to capacity and is a function of the target BER, desired noise margin and coding gain.

Using this definition we can use $\exp\left(\partial R^i_{k,j}\right)$ as our metric for dominant (tone, interferer) pair selection

$$\triangle^i_{k,j} = \exp\left(\partial R^i_{k,j}\right) = \begin{cases} \frac{(\text{signal}_i(k)+\text{noise}_i(k))(\text{interference}_{i,j}(k)+\text{noise}_i(k))}{\text{noise}_i(k)(\text{signal}_i(k)+\text{interference}_{i,j}(k)+\text{noise}_i(k))} & i \neq j \\ \infty & i = j \end{cases}$$

Here we have defined $\triangle^i_{k,j}$ as infinity when $i = j$ to force the direct lines to always be equalised. Another way of thinking of this is that we always include user i's direct line during the detection of user i.

Let us define a vector which contains the metrics for user i sorted in descending order

$$\text{t}^i = \text{sort}\left(\left[\begin{array}{cccccccc} \triangle^i_{1,1} & \dots & \triangle^i_{1,N} & \dots & \triangle^i_{K,1} & \dots & \triangle^i_{K,N} \end{array}\right]\right)$$

We use this as a threshold in the choice of dominant interferers. We can now define the set of dominant interferers for user i as

$$\mathbf{D}^i = \left\{ \; (k,j) \quad : \quad \triangle^i_{j,k} \geq \text{t}^i(\mu_i cMN) \; \right\}$$

where $t^i(a)$ denotes the a'th element of vector t^i. Notice that $\left|\mathbf{D}^i\right| = \mu_i cMN$ where $|.|$ denotes cardinality.

Further to this, the set of dominant interferers for user i at a single tone k can be defined as

$$\mathbf{D}^i_k = \left\{ \; j \quad : \quad (k,j) \in D^i \; \right\}$$

8.3 Crosstalk Canceller Design

In the previous section we devised a heuristic method for selecting the dominant (tone, interferer) pairs of a given user. In this section we describe the design of a partial crosstalk canceller based on the selection of dominant (tone, interferer) pairs. The derivation here is similar to that presented in section 3.2.2.

First recall our system model. We define the received symbols on a per-tone basis

$$\mathbf{y}(k) = \mathbf{H}(k)\mathbf{x}(k) + \mathbf{n}(k)$$

where $\mathbf{x}(k)$ is the transmitted signal from all users at tone k and similar definitions apply for $\mathbf{y}(k)$, $\mathbf{H}(k)$ and $\mathbf{n}(k)$.

We wish to design some crosstalk cancelling matrix $\mathbf{W}(k)$ such that we minimize

$$J(k) = \mathcal{E}\left|\mathbf{x}(k) - \mathbf{W}(k)\mathbf{y}(k)\right|^2$$

Element by element we can express this as

$$J_i(k) = \mathcal{E}\left|x_i(k) - \mathbf{W}_{\mathrm{row}\,i}(k)\mathbf{y}(k)\right|^2$$

Since we perform partial crosstalk cancellation, $\mathbf{W}_{\mathrm{row}\,i}(k)$ will contain only a few non-zero entries which correspond to the dominant interferers in $\mathbf{D}_k^i$. Let us define

$$\begin{aligned}
\widetilde{\mathbf{W}}_{\mathrm{row}\,i}(k) &= \text{non}-\text{zero elements of } \mathbf{W}_{\mathrm{row}\,i}(k) \\
\widetilde{\mathbf{H}}_i(k) &= \text{corresponding rows of } \mathbf{H}(k) \\
&= \mathbf{H}(k)_{\mathrm{rows}\,\mathbf{D}_k^i}
\end{aligned}$$

The optimal crosstalk canceller in the MMSE sense is now

$$\begin{aligned}
\mathbf{W}^*_{\mathrm{row}\,i}(k) &= \arg\min_{\mathbf{W}_{\mathrm{row}\,i}(k)} J_i(k) \\
&= \arg\min_{\mathbf{W}_{\mathrm{row}\,i}(k)} \mathcal{E}\left|x_i(k) - \mathbf{W}_{\mathrm{row}\,i}(k)\left(\mathbf{H}(k)\mathbf{x}(k) + \mathbf{n}(k)\right)\right|^2 \\
\widetilde{\mathbf{W}}^*_{\mathrm{row}\,i}(k) &= \arg\min_{\widetilde{\mathbf{W}}_{\mathrm{row}\,i(k)}} \mathcal{E}\left|x_i(k) - \widetilde{\mathbf{W}}_{\mathrm{row}\,i}(k)\widetilde{\mathbf{H}}_i(k)\mathbf{x}(k) - \widetilde{\mathbf{W}}_{\mathrm{row}\,i}(k)\widetilde{\mathbf{n}}(k)\right|^2 \\
&= \arg\min_{\widetilde{\mathbf{W}}_{\mathrm{row}\,i}} \widetilde{J}_i(k)
\end{aligned}$$

where $\widetilde{\mathbf{n}}(k)$ is defined similar to $\widetilde{\mathbf{H}}_i(k)$ as the rows corresponding to the non-zero elements of $\widetilde{\mathbf{W}}_{\mathrm{row}\,i}(k)$. Solving $\nabla_{\widetilde{\mathbf{W}}_{\mathrm{row}\,i}} \widetilde{J}_i(k) = 0$ we find

$$\widetilde{\mathbf{W}}^*_{\mathrm{row}\,i} = S_i(k)\ \left[\ \mathbf{0}_{1\times i-1}\quad 1\quad \mathbf{0}_{1\times N-i}\ \right]\ \widetilde{\mathbf{H}}_i^H(k)\ \left(\widetilde{\mathbf{H}}_i(k)\mathbf{S}_x(k)\widetilde{\mathbf{H}}_i^H(k) + \mathbf{S}_n(k)\right)^{-1}$$

where $\mathbf{S}_x(k)$ contains the PSD of all users at tone k. Naturally this depends on the power allocation method employed.

Just to re-cap, we now have a way of designing the partial crosstalk canceller $\mathbf{W}(k)$ based on the choice of dominant (tone, interferer) pairs discussed in section 8.2. Effectively we move through each (tone, user) pair, and design an MMSE crosstalk canceller which only uses knowledge of the signals of dominant interferers at that tone, for that user.

8.4 Performance

In the previous two sections we proposed a partial crosstalk cancellation scheme which selects dominant (tone, interferer) pairs for each user, and then designs MMSE crosstalk cancellers based on these selections.

To build a complete system, we combine this with the noise-only waterfilling + constant PBO power allocation scheme that was discussed in section 7.3. As mentioned previously, this power allocation scheme performs well with or without crosstalk cancellation and as a result is well suited to the partial cancellation scenario.

To evaluate performance we ran simulations with $c = 3$ which limits the number of multiplications to 3 times the complexity of a standard FDE. The simulation scenario is as described in section 5.3.

For a fixed choice of complexity allocation (μ_i's) varying the degree of power backoff (κ_i's) produces a single rate region. Furthermore, sweeping through the different values of μ_i produces a family of rate regions. The union of all of these regions is the total achievable region with this scheme.

Shown in figure 8.1 are the rate regions of our simulation. When μ_i is set to 0 for the near-end users and $\frac{1}{4}$ for the far end users the region is virtually rectangular. This is because the far-end users cancel crosstalk almost perfectly and experience no degradation in rate as the power of the near-end users is increased. On the other hand, the near-end users rate is significantly below the rate achieved by full crosstalk cancellation since the crosstalk between near-end users is not removed.

In the other extreme we set μ_i to $\frac{1}{4}$ for the near-end users and 0 for the far-end users, we see the far-end user's rate degrades rapidly as the near-end user's power is increased. We also see that since near-end users can now cancel the crosstalk they induce into one another we come very close to the full cancellation rate for near-end users.

The most important thing to notice is that with significantly reduced complexity ($\frac{3}{8}$ of full cancellation or 3 times the complexity of a standard FDE) we can achieve significant gains in performance ($20 \rightarrow 64$ Mbps on near-end users).

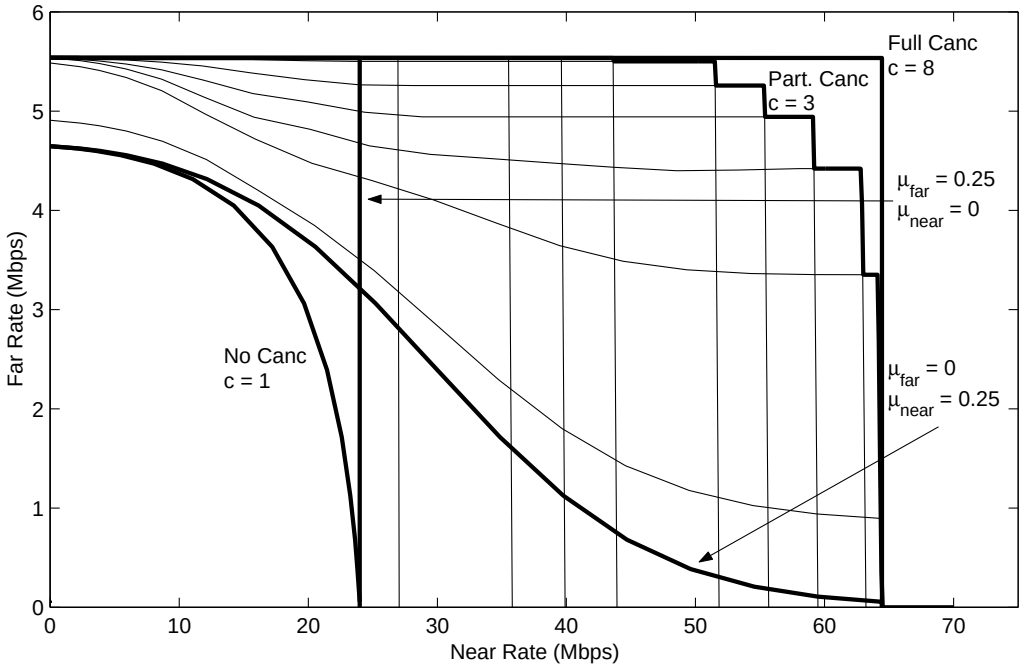


Figure 8.1: Rate Regions for Partial Crosstalk Cancellation with $c = 3$ and varying values for μ_i

Shown in figure 8.2 are the rate regions that are achieved under different complexities. These rate regions are formed by taking the union across all values of μ_i for a given choice of c. Notice that with $c = 4$ (4 times the complexity of FDE, half the complexity of full cancellation) we can almost achieve the performance of full crosstalk cancellation.

8.5 Summary

In this chapter we presented a partial crosstalk cancellation scheme which is combined with multi-user power allocation. The scheme consists of

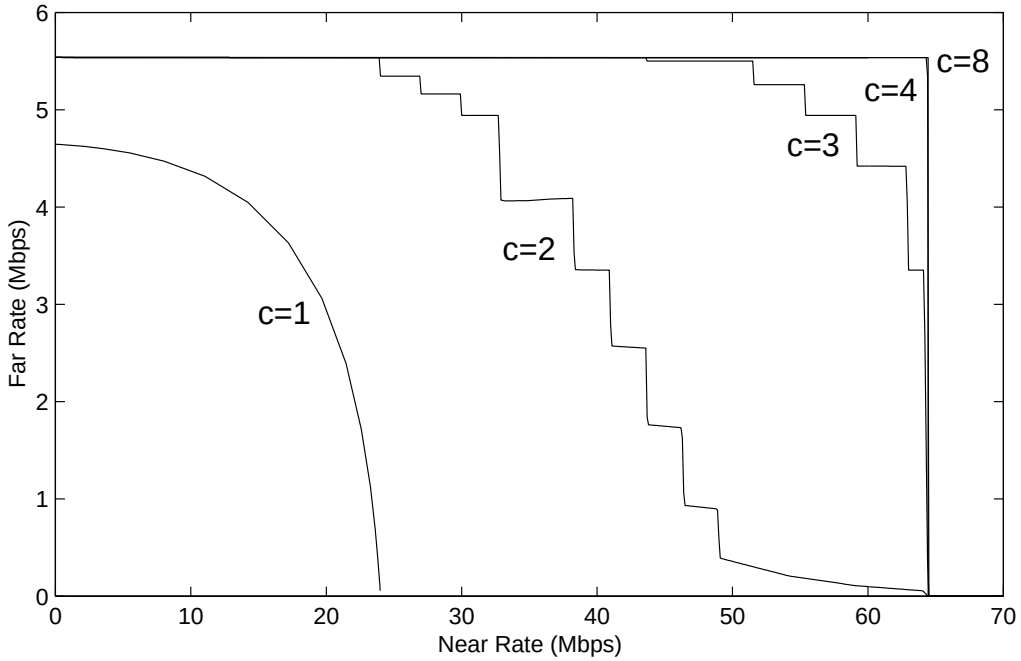


Figure 8.2: Rate Regions for Partial Crosstalk Cancellation with varying c

- A simple method for selecting dominant (tone, interferer) pairs
- Noise-only waterfilling with constant PBO to trade-off near-end/far-end rates

By treating computational complexity as a resource, and allocating it across users in a similar fashion to the allocation of transmit power, our scheme achieves significant performance gains with low complexity. With a complexity increase 3 times that of a standard 1-tap FDE, near-end rates go from 20 to 64 Mbps.

Chapter 9

Future Work

Our investigations in Crosstalk Cancellation and Multi-user Power Allocation have left a number of questions unanswered and exposed directions for future research. In this chapter we give a brief overview of some possible extensions to the work done to date.

9.1 Downstream Power Back-off

The power back-off schemes discussed in Chapters 5-8 where all applied to US communication. As mentioned earlier, the near-far effect also occurs in the DS direction when crosstalk from an ONU overwhelms the direct signal from a CO located further away from the CP end. Currently this occurs less frequently than the near-far effect in the US direction. However with the continual roll-out of VDSL, the near-far effect in the DS direction will become increasingly common.

In the DS direction we have co-operation at the transmitter side and this is referred to in information theory as the 'broadcast channel'. Unfortunately, the rate region of the broadcast channel is still unknown and significant research is being done in this area[23][24][25].

It is also possible to apply simple schemes such as iterative waterfilling or conventional PBO to the DS PBO problem.

9.2 Optimal Power Allocation without Crosstalk Cancellation

We investigated, in Chapter 7, multi-user power allocation in the absence of crosstalk cancellation. We saw that determining the optimal power allocation results in a non-convex optimisation which we could not solve. Indeed the rate region of interference channels is still unknown[26].

The solution to this problem is interesting for two reasons. On the one hand, it allows us to find an upper limit on the performance of a system without crosstalk cancellation. With this we can determine the sub-optimality of any scheme. In addition to this, knowing the form of the optimal power allocation gives us insight into the system and may hint at low complexity algorithms which are near-optimal.

The interaction between power allocation and partial crosstalk cancellation (in particular the choice of dominant (tone, interferer) pairs) is particularly interesting.

Unbundled environments are one situation where we are unable to do crosstalk cancellation unless we resort to blind techniques which have extremely high complexity. Furthermore, there is no centralized control in the unbundled system and this means that any practical power allocation strategy has to be implemented in a distributed fashion.

Such distributed systems lend themselves to analysis through game theory and the study of power control from a game theory perspective is receiving considerable attention[22][27][28] and may deserve further investigation.

9.3 Optimal Partial Cancellation

In section 8.2 we showed a simple scheme for choosing dominant (tone, interferer) pairs in a partial crosstalk canceller. While this scheme led to good performance, it is interesting to investigate more optimal schemes. Knowledge of the optimal scheme also allows us to quantify the sub-optimality of simpler schemes.

9.4 Partial Unbundling and Spectral Masks

Conventional power back-off and the use of spectral masks is currently proposed in VDSL standards as a way of addressing crosstalk and the near-far effect. These spectral masks are designed using worst 1% case channel models and as a result they can be overly restrictive on near-end users.

As we discussed, the main concept behind dynamic spectrum management and multi-user power allocation is that by exploiting knowledge of the direct and crosstalk channels, we no longer need to use worst case designs. We can adjust power allocations specifically to suite the channel at hand and the desired rates for the users in a binder. We saw in Chapters 6-8 that taking this approach leads to significant gains over conventional PBO.

Unfortunately it may not always be possible to violate spectral mask constraints, particularly when we operate in an unbundled environment. In this case, it may be interesting to investigate power allocation in an optimisation framework (as before), but with the addition of constraints on maximum PSDs relating the the relevant spectral masks.

Since the spectral masks form a set of linear constraints on the transmit spectra of the users, they do not alter the convexity of the optimisation problem. Hence with crosstalk cancellation we can still formulate power allocation as a convex optimisation[2].

It may be interesting to investigate the gains we can achieve by exploiting channel knowledge whilst still obeying the spectral mask constraints. Naturally the gains will be reduced since the spectral masks decrease the size of the valid ‘solution space’, however the gains may still be considerable.

9.5 Flexible FDD Bandplans

It has been suggested in literature that considerable gains are achievable by allowing not just the transmit power to be dynamically allocated, but by allowing full freedom in design of the FDD bandplan aswell[29][2].

In this scheme, all users in a binder are forced to use the same bandplan (to prevent intra-binder NEXT), however the bandplan can be designed on a per-binder basis, exploiting full knowledge of the crosstalk channels.

An added advantage of this is that it allows a trade-off between upstream and downstream rates aswell as between users. This means that the carrier could adjust bandplans on the basis of the applications being carrier within a binder group.

9.6 Adaptive Issues

In the work presented in this report, we always assumed perfect channel knowledge. Whilst this is a reasonable assumption given the static nature of the DSL channel; channel identification and tracking is an important issue which remains to be investigated.

Other important issues include the effect of new users entering the bundle on system performance. Investigation of adaption schemes for crosstalk cancellation parameters when a new user enters the system which do not disrupt existing services is also an important area.

9.7 Bonding and the Multi-Access Channel

Bonding is a method commonly used for providing high-speed access over a set of co-ordinated DSL lines. In principle, bonding simply multiplexes a data stream across multiple-lines giving the appearance of a high-rate 'virtual' connection.

Co-ordination between bonded lines can on occur different levels. The most simple and common-place implementation involves co-ordination on the transport layer. There have been a number of suggestions that significant gains can be achieved through co-ordination on the physical layer through signal processing techniques.

The advantages of this are twofold. Bonding can use signal processing to:

- Remove inter-line interference from within the bonded set of lines and
- Exploit full spatial diversity

In a multi-user scenario bonding can exploit the correlation between external interference/noise seen at different bonded receivers. This could lead to improved interference cancellation and noise filtering. The essential gains in bonding come from allowing a co-ordinated transceiver at the CP in addition to the CO.

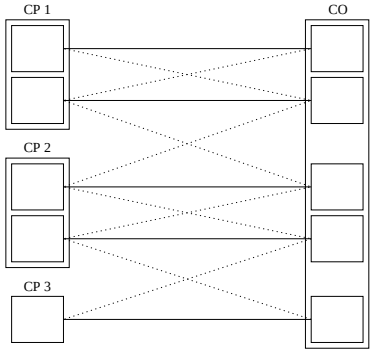


Figure 9.1: Bonding

9.8 Crosstalk Cancellation in Unbundled and Asynchronous Environments

In an asynchronous environment a single operator occupies a bundle but runs unsynchronised VDSL systems across the different twisted pairs. In the unbundled case, multiple operators will occupy a bundle and it is highly unlikely that they will share a common clock.

The result of this is inter-carrier interference (ICI) between the different VDSL systems in the bundle. ICI does not change the structure of optimal power allocation schemes however the MUD becomes considerably more complex. Instead of doing crosstalk cancellation on a per-tone basis, we are now forced to cancel crosstalk jointly across all tones. This can lead to matrix inversions of e.g. 20 users and 1000 tones: 20000×20000 matrices during initialisation and 20000^2 multiplications per DMT-block during run-time.

Clearly this kind of complexity is unrealisable. In the unbundled case, crosstalk cancellation must be done in a somewhat blind fashion, since we do not have access to the crosstalker's direct channels. This increases complexity even further.

The search for asynchronous/unbundled crosstalk cancellation strategies with realisable complexity and good performance is a challenging task. There have been a number of contribution in literature on soft interference cancellation which is applicable to unbundled environments[4][30]. Due to the sparseness of the crosstalk coupling matrix, even in the asynchronous case, the combination of partial crosstalk cancellation with soft cancellation techniques looks promising.

Chapter 10

Summary - Multi-user Power Allocation

In the second part of this report we investigated the near-far effect and multi-user power allocation in VDSL. In Chapter 5 we introduced the near-far effect which occurs when the crosstalk signal of a near-end user overwhelms the direct signal of a far-end user.

This problem is traditionally addressed using conventional PBO, an overview of which was given in Chapter 5. Conventional PBO is designed using worst 1% case channel models which can lead to overly conservative PBO on the near-end users. Furthermore, since conventional PBO is not in any way related to the desired rates of the various users in a binder, it gives no mechanism for trading off near/far rates according to the desired services.

Given knowledge of the crosstalk channels we can do significantly better than conventional PBO and this was the subject of Chapters 5-8. In Chapter 5 we gave an overview of the information theory of the MAC. This allows us to formulate multi-user power allocation as an optimisation problem.

In Chapter 6 we discussed multi-user power allocation when crosstalk cancellation is used at the receiver. In this case the optimisation becomes convex and leads to 'simultaneous vector waterfilling'. A simple, near-optimal power allocation scheme called 'noise-only waterfilling' was also presented. These schemes gave significant gains ($5.9 \rightarrow 64$ Mbps on near-end users) over conventional PBO.

Power allocation without crosstalk cancellation was investigated in Chapter 7 and shown to lead to a non-convex optimisation problem. The unrealisable complexity of solving this forced us to focus on sub-optimal schemes. We presented two schemes, 'iterative scalar waterfilling' as suggested by Cioffi et. al.[22] and one based on a combination of noise-only waterfilling with constant PBO. Both schemes gave significant gains ($5.5 \rightarrow 20$ Mbps on near-end users) over conventional PBO while the noise-only waterfilling scheme was significantly simpler.

The concept of partial cancellation was introduced in the first part of this report and allows us to achieve the majority of gains from crosstalk cancellation with a fraction of the complexity. This is done by cancelling crosstalk only on the tones and interferers where it causes the largest performance degradation.

In Chapter 8 we presented a scheme which combines multi-user power allocation and partial crosstalk cancellation. The scheme combines noise-only waterfilling with a simplified method for selecting dominant (tone, interferer) pairs. Complexity in the signal processing equipment at the CO is treated as a resource to be allocated across tones and users, in a similar manner to the allocation of transmit power. Using the partial cancellation scheme it is possible to increase near-end user rates from $20 \rightarrow 64$ Mbps with just 3 times the complexity of a standard 1-tap frequency domain equaliser (FDE).

The investigations to date have left several open questions and in Chapter 9 we suggested several areas for future research. Some particularly relevant topics include downstream PBO, bonded VDSL and cancellation of asynchronous crosstalk.

Appendix A
Proof on Optimality of Simultaneous Waterfilling

We now present a light-weight version of the derivation for optimal power allocation in the vector MAC where the optimisation is done with equal weights ($\alpha = 0.5$). This is based on the work in [20] but considerably simplified, and applied to the DSL context. We also omit the proof of convexity of the optimisation function.

To begin with, let us define our cost function as the unweighted rate-sum under the power constraints for each user

$$J = C_{\text{sum}} + \lambda_1 \left(P_1 - \sum_k S_1(k) \right) + \lambda_2 \left(P_2 - \sum_k S_2(k) \right) \tag{10.1}$$

Here λ_1 and λ_2 are the Lagrange multipliers corresponding to the power constraints (the existence of which is proven in [20]), P_i is the total power constraint for user i, and $S_i(k)$ is the PSD of user i at tone k. The sum of rates $C_{\text{sum}} = R_1 + R_2$

First note that the sum of rates is always $I(\mathbf{x}_1, \mathbf{x}_2; \mathbf{y})$ regardless of detection order. This can be seen by examining (5.3) and (5.5). Hence

$$\begin{aligned} C_{\text{sum}} &= \frac{1}{2} \log_2 \left| \mathbf{I}_{KN} + \left(\mathbf{H}_1 \mathbf{S}_1 \mathbf{H}_1^H + \mathbf{H}_2 \mathbf{S}_2 \mathbf{H}_2^H \right) \mathbf{N}^{-1} \right| \\ &= \frac{1}{2} \sum_{k=1}^{K} \log_2 \left| \mathbf{I}_N + \left(\mathbf{h}_1(k) \mathbf{h}_1^H(k) S_1(k) + \mathbf{h}_2(k) \mathbf{h}_2^H(k) S_2(k) \right) \mathbf{N}(k)^{-1} \right| \end{aligned}$$

Recall that N is the number of users (2 in this case) and K the number of tones. We have essentially split our multi-carrier system into K single carrier systems. $\mathbf{h}_i(k)$ is the composite (direct and crosstalk) channel for user i on tone k (see (5.1)) and $S_i(k)$ is the corresponding transmit PSD. $\mathbf{N}(k)$ is the noise covariance matrix for tone k.

Using the definition of the derivative

$$\frac{\partial f(x)}{\partial x} = \lim_{\triangle x \to 0} \frac{f(x + \triangle x) - f(x)}{\triangle x}$$

$$\begin{aligned} \frac{\partial C_{\text{sum}}}{\partial S_i(k)} &= \lim_{\triangle S_i(k) \to 0} \frac{\left[\log \left| \mathbf{I} + \left(\sum_{j=1}^{2} \mathbf{h}_j(k) \mathbf{h}_j^H(k) S_j(k) + \mathbf{h}_i(k) \mathbf{h}_i^H(k) \triangle S_i(k) \right) \mathbf{N}(k)^{-1} \right| - \log \left| \mathbf{I} + \sum_{j=1}^{2} \mathbf{h}_j(k) \mathbf{h}_j^H(k) S_j(k) \mathbf{N}(k)^{-1} \right| \right]}{\triangle S_i(k)} \\ &= \lim_{\triangle S_i(k) \to 0} \frac{\log \left| \mathbf{I} + \mathbf{h}_i(k) \mathbf{h}_i^H(k) \triangle S_i(k) \left[\mathbf{N}(k) + \sum_{j=1}^{2} \mathbf{h}_j(k) \mathbf{h}_j^H(k) S_j(k) \right]^{-1} \right|}{\triangle S_i(k)} \end{aligned}$$

Since the cost function (10.1) is convex, at the optimum value for $S_i(k)$

$$\begin{aligned} \frac{\partial J}{\partial S_i(k)}\left(S_i^*(k)\right) &= 0 \\ \frac{\partial}{\partial S_i(k)} C_{\text{sum}}\left(S_i^*(k)\right) - \lambda_i &= 0 \end{aligned}$$

hence

$$\frac{\partial}{\partial S_i(k)} C_{\text{sum}}\left(S_i^*(k)\right) = \lambda_i$$

$$\lim_{\triangle S_i(k)\to 0} \frac{\log\left|\mathbf{I} + \overbrace{\mathbf{h}_i(k)}^{\mathbf{A}}\overbrace{\mathbf{h}_i^H(k)}^{\mathbf{B}}\overbrace{\triangle S_i(k)}^{\mathbf{C}}\overbrace{\left[\mathbf{N}(k) + \sum_{j\neq i}\mathbf{h}_j(k)\mathbf{h}_j^H(k)S_j(k) + \mathbf{h}_i(k)\mathbf{h}_i^H(k)S_i^*(k)\right]^{-1}}^{\mathbf{D}}\right|}{\triangle S_i(k)} = \lambda_i$$

Since $|\mathbf{I} + \mathbf{ABCD}| = |\mathbf{I} + \mathbf{CBDA}|$

$$\lim_{\triangle S_i(k)\to 0} \frac{\log\left|1 + \triangle S_i(k)\,\mathbf{h}_i^H(k)\overbrace{\left[\mathbf{N}(k) + \sum_{j\neq i}\mathbf{h}_j(k)\mathbf{h}_j^H(k)S_j(k) + \mathbf{h}_i(k)\mathbf{h}_i^H(k)S_i^*(k)\right]^{-1}}^{k}\mathbf{h}_i(k)\right|}{\triangle S_i(k)} = \lambda_i \qquad (10.2)$$

Using a Taylor expansion for (10.2)

$$\begin{aligned} \log(1+\triangle x k) &= k\triangle x + O(\triangle x^2) + \ldots \\ \lim_{\triangle x\to 0}\frac{\log(1+\triangle x k)}{\triangle x} &= k \end{aligned}$$

hence

$$\mathbf{h}_i^H(k)\left[\overbrace{\mathbf{N}(k) + \sum_{j\neq i}\mathbf{h}_j(k)\mathbf{h}_j^H(k)S_j(k)}^{\mathbf{Z}} + \overbrace{\mathbf{h}_i(k)S_i^*(k)^{0.5}}^{\mathbf{x}}\overbrace{S_i^*(k)^{0.5}\mathbf{h}_i^H(k)}^{\mathbf{x}^H}\right]^{-1}\mathbf{h}_i(k) = \lambda_i$$

Using the matrix inversion lemma

$$\begin{aligned} \lambda_i &= \frac{1}{S_i^*(k)}\mathbf{x}^H\left[\mathbf{Z} - \mathbf{x}\mathbf{x}^H\right]^{-1}\mathbf{x} \\ &= \frac{1}{S_i^*(k)}\mathbf{x}^H\left[\mathbf{Z}^{-1} - \frac{\mathbf{Z}^{-1}\mathbf{x}\mathbf{x}^H\mathbf{Z}^{-1}}{1+\mathbf{x}^H\mathbf{Z}^{-1}\mathbf{x}}\right]\mathbf{x} \\ &= \frac{1}{S_i^*(k)}\mathbf{x}^H\mathbf{Z}^{-1}\mathbf{x}\left[1 - \frac{\mathbf{x}^H\mathbf{Z}^{-1}\mathbf{x}}{1+\mathbf{x}^H\mathbf{Z}^{-1}\mathbf{x}}\right] \\ &= \frac{1}{S_i^*(k)}\frac{1}{1+\left(\mathbf{x}^H\mathbf{Z}^{-1}\mathbf{x}\right)^{-1}} \end{aligned}$$

Hence

$$S_i^*(k)\left[1 + \left(\mathbf{x}^H\mathbf{Z}^{-1}\mathbf{x}\right)^{-1}\right] = \frac{1}{\lambda_i}$$

and

$$S_i^*(k) + \frac{1}{\mathbf{h}_i^H(k)\left[\mathbf{N}(k) + \sum_{j \neq i} \mathbf{h}_j(k)\mathbf{h}_j^H(k)S_j(k)\right]^{-1}\mathbf{h}_i(k)} = \frac{1}{\lambda_i} \tag{10.3}$$

As can be seen, the optimal power allocation is a form of simultaneous waterfilling against the vector channel. The matrix $\mathbf{Z}$ represents the covariance of the background noise and interference from other users, whilst $\mathbf{x}\mathbf{x}^H$ is the covariance of user i's signal after it has passed through the channel. So the sum of a user's power, plus the noise and interference to channel ratio should equal a constant waterfilling level ($\frac{1}{\lambda_i}$) at all frequencies.

This expression is exactly the same for more than 2 users. We limited the derivation here to 2 users only to simplify notation. The extension is trivial.

A closed form expression which solves (10.3) simultaneously for all users is highly complex. Fortunately, a simple iterative method exists which exhibits rapid and well-behaved convergence to the simultaneous solution[19]. This is referred to as 'iterative waterfilling' and is discussed further in chapter 6.

Bibliography

[1] S. Verdu, *Code Division Multiple Access Communications*, ch. Adaptive Multiuser Detection, pp. 97–116. Kluwer Academic Publishers, 1995.

[2] G. Ginis and J. Cioffi, "Vectored Transmission for Digital Subscriber Line Systems," *IEEE J. Select. Areas Commun.*, 2001. submitted.

[3] G. Foschini, "Layered Space-Time Architecture for Wireless Communication in a Fading Environment When Using Multi-Element Antennas," *Bell Labs Technical Journal*, pp. 41–59, Autumn 1996.

[4] K. Cheong, W. Chio, and J. Cioffi, "Multiuser Soft Interference Canceler via Iterative Decoding for DSL Applications," *IEEE J. Select. Areas Commun.*, vol. 20, pp. 363–371, Feb. 2002.

[5] M. Sellathurai and S. Haykin, "A Simplified Diagonal BLAST Architecture with Iterative Parallel-Interference Cancelation Receivers," in *Proc. of the Int. Conf. on Commun.*, (Helsinki, Finland), June 2001.

[6] 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*, (Beijing), Aug. 2000.

[7] E. Lee and D. Messerschmitt, *Digital Communications*, ch. 10, pp. 460–464. Kluwer Academic, 2 ed., 1994.

[8] G. Foschini and M. Gans, "On Limits of Wireless Communications in a Fading Environment when Using Multiple Antennas," *Wireless Personal Communications*, vol. 6, pp. 311–335, 1998.

[9] G. Taubock and W. Henkel, "MIMO Systems in the Subscriber-Line Network," in *Proc. of the 5th Int. OFDM-Workshop*, (Hamburg), pp. 18,1–18,3, 2000.

[10] J. Hench *et al.*, "Method and Apparatus for Impairment Diagnosis in Communication Systems." International Patent, May 2001. International Publication Number: WO 01/35609 A1.

[11] I. Duff, A. Erisman, and J. Reid, *Direct Methods for Sparse Matrices*. Oxford Science Publications, 1986.

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

[13] ETSI, "TS 101 270-1 Standard, Transmission and Mutliplexing (TM); Access transmission systems on metallic access cables; VDSL; Part 1: Functional Requirements," 1999. Version 1.2.1 pg. 32.

[14] M. Varanasi and T. Guess, "Optimum Decision Feedback Multiuser Equalization with Successive Decoding Achieves the Total Capacity of the Gaussian Multiple-Access Channel," in *Proc. of Asilomar Conf. on Signals, Syst. and Comput.*, pp. 1405–1409, 1997.

[15] R. Ahlswede, "Multi-way communication channels," in *Proc. 2nd Int. Symp. on Inform. Theory*, (Aermenia, U.S.S.R), pp. 23–52, Sept. 1971.

[16] S. Verdu, "Multiple Access Channels with Memory with and without Frame Synchronism," *IEEE Trans. on Inform. Theory*, vol. 35, pp. 605–619, May 1989.

[17] R. Cheng and S. Verdu, "Gaussian Multiaccess Channels with ISI: Capacity Region and Multiuser Water-Filling," *IEEE Trans. on Inform. Theory*, vol. 39, pp. 773–785, May 1993.

[18] D. Tse and S. Hanly, "Multiaccess Fading Channels-Part I: Polymatroid Structure, Optimal Resources Allocation and Throughput Capacities," *IEEE Trans. on Inform. Theory*, vol. 44, pp. 2796–2815, Nov. 1998.

[19] W. Yu, W. Rhee, S. Boyd, and J. Cioffi, "Iterative Water-filling for Gaussian Vector Multiple Access Channels," *IEEE Trans. Inform. Theory*, 2001. Submitted 6 November 2001.

[20] P. Viswanath, D. Tse, and V. Anantharam, "Asymptotically Optimal Water-Filling in Vector Multiple-Access Channels," *IEEE Trans. on Inform. Theory*, vol. 47, pp. 241–267, Jan. 2001.

[21] S. Boyd and L. Vandenberghe, *Convex Optimization*. http://www.stanford.edu/class/ee364/: Stanford University, 2001. Textbook Draft.

[22] W. Yu, G. Ginis, and J. Cioffi, "Distributed Multiuser Power Control for Digital Subscriber Lines," *IEEE J. Select. Areas Commun.*, vol. 20, June 2002. Pre-print available at http://www.stanford.edu/ weiyu.

[23] T. Cover, "Comments on broadcast channels," *IEEE Trans. Inform. Theory*, vol. 44, pp. 2524–2530, Oct. 1998.

[24] A. Goldsmith and M. Effros, "The capacity region of broadcast channels with intersymbol interference and colored Gaussian noise," *IEEE Trans. Inform. Theory*, vol. 47, pp. 211–219, Jan. 2001.

[25] W. Yu and J. Cioffi, "Sum Capacity of Gaussian Vector Broadcast Channels," *IEEE Trans. Inform. Theory*, 2001. submitted.

[26] G. Kramer, "Genie-aided outer bounds on the capacity of interference channels," in *Proc. IEEE Int. Symp. on Inform. Theory*, p. 103, 2001.

[27] G. Foschini and Z. Miljanic, "A simple distributed autonomous power control algorithm and its convergence," *IEEE. Trans. Veh. Tech.*, vol. 42, pp. 641–646, Nov. 1993.

[28] A. MacKenzie and S. Wicker, "Game theory and the design of self-configuring, adaptive wireless networks," *IEEE Commun. Mag.*, vol. 39, pp. 126–131, Nov. 2001.

[29] W. Yu and J. Cioffi, "eFDMA Capacitty of Gaussian Multiple-Access Channels With ISI," *IEEE. Trans. Comms.*, vol. 50, pp. 102–111, Jan. 2002.

[30] C. Zeng and J. Cioffi, "Crosstalk Cancellation in xDSL Systems," *IEEE J. Select. Areas Commun.*, 2001. Submitted.