

INFORMATION TO USERS

This manuscript has been reproduced from the microfilm master. UMI films the text directly from the original or copy submitted. Thus, some thesis and dissertation copies are in typewriter face, while others may be from any type of computer printer.

The quality of this reproduction is dependent upon the quality of the copy submitted. Broken or indistinct print, colored or poor quality illustrations and photographs, print bleedthrough, substandard margins, and improper alignment can adversely affect reproduction.

In the unlikely event that the author did not send UMI a complete manuscript and there are missing pages, these will be noted. Also, if unauthorized copyright material had to be removed, a note will indicate the deletion.

Oversize materials (e.g., maps, drawings, charts) are reproduced by sectioning the original, beginning at the upper left-hand corner and continuing from left to right in equal sections with small overlaps.

ProQuest Information and Learning
300 North Zeeb Road, Ann Arbor, MI 48106-1346 USA
800-521-0600

UMI[®]

Characteristics and Performance of Various VDSL RFI Suppression Techniques

By: Richard Abela, B.Eng.

A thesis submitted to the
Faculty of Graduate Studies and Research
In partial fulfillment of the requirements for the degree of

Master of Applied Science

Ottawa-Carleton Institute for Electrical and Computer Engineering
Department of Systems and Computer Engineering
Carleton University
Ottawa, Ontario

© Copyright 2005, Richard Abela



Library and
Archives Canada

Bibliothèque et
Archives Canada

0-494-08362-X

Published Heritage
Branch

Direction du
Patrimoine de l'édition

395 Wellington Street
Ottawa ON K1A 0N4
Canada

395, rue Wellington
Ottawa ON K1A 0N4
Canada

Your file *Votre référence*

ISBN:

Our file *Notre référence*

ISBN:

NOTICE:

The author has granted a non-exclusive license allowing Library and Archives Canada to reproduce, publish, archive, preserve, conserve, communicate to the public by telecommunication or on the Internet, loan, distribute and sell theses worldwide, for commercial or non-commercial purposes, in microform, paper, electronic and/or any other formats.

The author retains copyright ownership and moral rights in this thesis. Neither the thesis nor substantial extracts from it may be printed or otherwise reproduced without the author's permission.

AVIS:

L'auteur a accordé une licence non exclusive permettant à la Bibliothèque et Archives Canada de reproduire, publier, archiver, sauvegarder, conserver, transmettre au public par télécommunication ou par l'Internet, prêter, distribuer et vendre des thèses partout dans le monde, à des fins commerciales ou autres, sur support microforme, papier, électronique et/ou autres formats.

L'auteur conserve la propriété du droit d'auteur et des droits moraux qui protègent cette thèse. Ni la thèse ni des extraits substantiels de celle-ci ne doivent être imprimés ou autrement reproduits sans son autorisation.

In compliance with the Canadian Privacy Act some supporting forms may have been removed from this thesis.

Conformément à la loi canadienne sur la protection de la vie privée, quelques formulaires secondaires ont été enlevés de cette thèse.

While these forms may be included in the document page count, their removal does not represent any loss of content from the thesis.

Bien que ces formulaires aient inclus dans la pagination, il n'y aura aucun contenu manquant.


Canada

The undersigned recommend to
the Faculty of Graduate Studies and Research
acceptance of the thesis

**Characteristics and Performance of Various
VDSL RFI Suppression Techniques**

Submitted by
Richard Abela, B.Eng.

In partial fulfillment of the requirements for the degree of
Master of Applied Science

Professor Mohammed El-Tanany, Thesis Supervisor

Professor Rafik Goubran, Chair

May 2005, Carleton University

Abstract

The past few years have seen ADSL widely deployed as the telephone companies attempted to grow their market share in the consumer service industry of high speed internet access. Bundling a digital video service offering together with data and phone services is seen by many telephone companies as the next step for protecting, if not growing, their market share. VDSL is one technology that promises to provide the required rates over existing twisted pairs for supporting these services.

VDSL employs the frequency band spanning from 25 kHz to 12 MHz, a significantly larger band than used by ADSL. As a consequence, the VDSL band is adjacent to or overlaps with other radio signalling bands such as Amateur Radio (HAM) and AM broadcasts. These signals may introduce radio frequency interference (RFI) in VDSL signals, which will degrade VDSL performance.

Many approaches can be found in the literature for reducing the impact of RFI on VDSL communication. These methods however may not always be as practical or efficient as predicted when applied to real systems adhering to a common standard. This thesis attempts to quantify the impact of RFI on typical multi-carrier modulation (MCM) based-VDSL systems, also known as discrete multi-tone (DMT) VDSL, and researches some of the most promising suppression techniques for mitigating the effect of RFI in such systems. Elaborate end-to-end simulations are carried out to investigate the effect of these suppression techniques on system performance, both with and without RFI present. The thesis concludes by highlighting the performance and complexity characteristics of the various RFI suppression techniques considered.

Acknowledgements

First, I would like to thank my supervisor Professor Mohamed El-Tanany, whose understanding, guidance and encouragement were crucial in allowing me to complete this thesis on a part time basis.

I also want to thank all of my DSP colleagues at Catena Networks whom I credit for creating a stimulating and supportive work environment in which to grow. A special thanks goes to Eric Verreault for all his help, and in particular, for patiently reading the drafts of this thesis and providing valuable feedback. I have truly enjoyed the many though-provoking discussions we have had throughout the years.

I want to express by deepest gratitude towards my parents, Suzanne and Bernard, who have nurtured and fostered my desire for knowledge and my ability to learn. They have given me a wonderful gift, which I hope to bestow to my own children.

My final and most important thanks goes to my wife Christine whom I love very much. She has supported and encouraged me in my Masters studies from day one and this accomplishment would not have been possible without her care and affection. She is the inspiration of my life.

Table of Contents

Abstract	i
Acknowledgements	ii
Table of Contents	iii
List of Figures	vi
Acronyms and Abbreviations	ix
List of Symbols	xi
Chapter 1 Introduction	1
1.1 Thesis Objectives	2
1.2 Thesis organisation	2
Chapter 2 Background Information	4
2.1 Introduction to Discrete Multi-Carrier Modulation (MCM).....	4
2.2 The Transceiver and its Physical Layer Characteristics	5
2.2.1 Physical Medium Dependent (PMD) Functional Model	6
2.2.2 Multi-Carrier Modulation Specification	8
2.2.3 Mitigating Intersymbol Interference and Near-end Echo Transients	10
2.2.4 Cyclic Extension and Symbol Rate.....	13
2.2.5 Synchronisation of Transmitters and Receivers	15
2.2.6 Constellation Encoding, Channel Capacity and Bit Swapping.....	17
2.2.7 Transmit signal characteristics.....	20
2.2.8 Egress control.....	23
2.3 Received Signal Characteristics and Noise Impairments	25
2.3.1 Time-Domain ISI.....	26
2.3.2 Near-end Echo	26
2.3.3 Noise Impairments	28
2.3.3.1 Additive Noise.....	28
2.3.3.2 Impulse Noise.....	28
2.3.3.3 Crosstalk.....	28
2.3.3.4 Ingress Radio Frequency Interference.....	29

Chapter 3 Literature Review	32
3.1 Analog RFI Cancellation	32
3.2 Time-Domain Windowing	32
3.3 Adaptive Digital Notch Filtering	34
3.4 Digital RFI Cancellation	35
Chapter 4 Problem Statement	37
4.1 Unresolved Issues	37
4.2 Motives for Further Research and Investigation.....	41
Chapter 5 System Model	43
5.1 VDSL System Parameters.....	45
5.2 Transmitter.....	46
5.3 Channel Model.....	47
5.3.1 Loop Model.....	47
5.3.2 Near-end Echo	49
5.3.3 Ingress RFI Modelling	49
5.3.4 Other Noise Models	50
5.4 VTU-O Receiver Design.....	51
5.4.1 Analog RFI Cancellation	53
5.4.2 Receive Symbol Alignment	54
5.4.3 Time-Domain Equalization.....	56
5.5 Analysis of Loop Models in the Context of System Performance.....	58
5.5.1 Impulse Response of Loop Models.....	58
5.5.2 Transfer Function Magnitude of Loop Models.....	60
5.5.3 Input Impedance of Loop Models.....	62
5.5.4 Propagation Delay of Loop Model VDSL1-x.....	63
5.5.5 Dispersion Time of Loop Model VDSL1-x.....	64
5.5.6 Echo Path	65
5.5.7 Cyclic Extensions Requirements	69
Chapter 6 Digital RFI Suppression Techniques	71
6.1 Receiver Windowing	71
6.1.1 Raised Cosine Windowing Extending into Cyclic Extension.....	73

6.1.1.1	Raised Cosine Windowing Computational Requirements	78
6.1.2	Frequency-Domain implementation of Time Windowing	80
6.1.3	Constrained Windowing and Controlled ICI	80
6.1.3.1	Constrained Windowing Computational Requirements.....	88
6.2	RFI Suppression using an adaptive notch filter	89
6.2.1.1	Notch Filtering Computational Requirements	94
6.3	RFI Cancellation using Time Domain Modelling	95
Chapter 7	Simulation Results	97
7.1	Performance Benchmarks	97
7.1.1	Impact of Near-end Echo on VDSL Performance	98
7.1.2	VDSL Performance Sensitivity to Receiver Alignment	99
7.1.2.1	Nominal Benchmark	105
7.1.3	RFI impact on VDSL Performance.....	105
7.2	Windowing Performance	109
7.2.1	Nominal Windowing Performance	110
7.2.2	Windowing Performance with Symbol Misalignments.....	113
7.2.3	Windowing Performance in Additive White Gaussian Noise	119
7.2.4	Ability of Windowing to Combat RFI.....	122
7.3	Notch Filtering Performance Analysis.....	129
7.4	RFI Cancellation Performance Analysis.....	135
7.5	Summary	139
Chapter 8	Conclusions and Recommendations.....	142
8.1	Thesis Conclusions	142
8.2	Recommendations for further study.....	145
References		146

List of Figures

Figure 2-1: PMD Functional Model Diagram	7
Figure 2-2: VDSL Band Allocation.....	9
Figure 2-3: Cyclic Extensions and Symbol Overlap.....	14
Figure 2-4: Symbol Alignment	16
Figure 2-5: Upstream M1 PSD Template	21
Figure 2-6: Transmit Time-Domain Window.....	24
Figure 2-7: Transfer Function Magnitude of Transmit Window.....	24
Figure 2-8: First 14 Spectral Sidelobes of Transmit Window	25
Figure 2-9: Transmit Signal at Line Interface.....	26
Figure 4-1: Cyclic Suffix Requirement with Echo Path Propagation Delay	39
Figure 5-1: Simulation System Model.....	44
Figure 5-3: Receiver Model.....	52
Figure 5-4: Alignment of 5000-Foot Loop Impulse Response.....	55
Figure 5-5: Nominal Receiver Alignment	56
Figure 5-6: Impulse Response of Loop Model VDSL1-3000ft (26 AWG).....	58
Figure 5-7: Impulse Response of Loop Model VDSL5.....	59
Figure 5-8: Transfer Functions of Loop Models VDSL1-x of Various Length	60
Figure 5-9: Transfer Functions of Loop Models VDSL4 through VDSL7	61
Figure 5-10: Input Impedance Magnitude of Loop Model VDSL1-3000ft.....	62
Figure 5-11: VDSL1-x Propagation Delay versus Loop Length.....	63
Figure 5-12: VDSL1-x Dispersion Time versus Loop Length	64
Figure 5-13: Echo Path Impulse Response with VDSL1-3000ft.....	66
Figure 5-14: Echo Path Impulse Response with VDSL5.....	66
Figure 5-15: Transfer Function of Echo Path with VDSL1-3000ft.....	67
Figure 5-16: Transfer Function of Echo Path with VDSL5.....	67
Figure 5-17: Duration of Echo Path Impulse Response versus Loop Length (VDSL1-x).....	68
Figure 6-1: Raised Cosine Receiver Time-Domain Windows	73
Figure 6-2: Transfer Function Magnitude of Raised Cosine Windows.....	74

Figure 6-3: Bins 0 through 20 of Raised Cosine Window Transfer Functions	75
Figure 6-4: Receive Signal PSD with Raised Cosine Windowing	77
Figure 6-5: Receive SNR with and without Raised Cosine Windowing	77
Figure 6-6: Windowing and Folding of Received Symbol.....	79
Figure 6-7 : Constrained Windowing Original Architecture	81
Figure 6-8: Proposed Constrained Windowing Architecture.....	82
Figure 6-9: Magnitude of Constrained Time-Domain Windows.....	85
Figure 6-10: Transfer Function Magnitude of Constrained Windows.....	85
Figure 6-11: Bins 0 through 20 of Constrained Window Transfer Functions	86
Figure 6-12: Receive Signal PSD with Asymmetric {1, -1} Windowing	87
Figure 6-13: Receive SNR with and without Asymmetric {1,-1} Windowing	88
Figure 6-14: Transfer Function of Notch Filter ($r = 0.96$).....	91
Figure 6-15: Impulse Response of Single Notch Filter ($r = 0.96$)	92
Figure 6-16: Impulse Response of Two Consecutive Notch Filters ($r = 0.96$).....	92
Figure 6-17: Receive Signal PSD with Notch Filtering.....	93
Figure 6-18: Receive SNR with and without Notch Filtering	94
Figure 7-1: Rate Reach with no Echo for Various Receiver Alignments.....	101
Figure 7-2: Rate Reach without Transmit Shaping for Various Receiver Alignments ..	104
Figure 7-3: Rate Reach with Transmit Shaping for Various Receiver Alignments	104
Figure 7-4: Transfer Function of the Demodulator's Nominal Window	106
Figure 7-5: RFI and Upstream PSD.....	106
Figure 7-6: Receive SNR With and Without RFI Present	107
Figure 7-7: RFI Impact on VDSL Rates	108
Figure 7-8: Raised Cosine Windowing Impact in RFI-Free Environment	112
Figure 7-9: Constrained Windowing Impact in RFI-Free Environment.....	112
Figure 7-10: Raised Cosine Windowing Impact with Symbol Misalignments	113
Figure 7-11: Constrained Windowing Impact with Symbol Misalignments	114
Figure 7-12: Raised Cosine Windowing Impact in Different Noise Environment.....	121
Figure 7-13: Constrained Windowing Impact in Different Noise Environments.....	122
Figure 7-14: Raised Cosine Windowing Impact with RFI at bin 870.5	124
Figure 7-15: Windowing Impact with RFI at bin 870.5	124

Figure 7-16: Raised Cosine Windowing Impact with RFIs at bins 870.5 and 2343.5 ...	126
Figure 7-17: Windowing Impact with RFIs at bins 870.5 and 2343.5	126
Figure 7-18: Raised Cosine Windowing Impact with RFI at bin 2343.5	127
Figure 7-19: Windowing Impact with RFI at bin 2343.5	128
Figure 7-20: Notch Filtering Impact in RFI-Free Environment	131
Figure 7-21: Notch Filter and R-C Windowing Impact with Single RFI	134
Figure 7-22: Notch Filter and R-C Windowing Impact with Dual RFI.....	134
Figure 7-23: RFI Cancellation of 4 kHz RFI at bin 870.5 with 2000 foot Loop.....	136
Figure 7-24: RFI Cancellation of 400 Hz RFI at bin 870.5 with 1000-foot loop.....	138

Acronyms and Abbreviations

AC	Alternating current
ADC	Analog-to-digital converter
ADSL	Asymmetric digital subscriber line
AFE	Analog front end
AGC	Automatic gain control
AM	Amplitude modulation
ANSI	American National Standards Institute
AWG	American wire gauge
AWGN	Additive white Gaussian noise
CIR	Channel impulse response
CO	Central office (or local exchange)
CP	Cyclic prefix
CPE	Customer premise equipment
CS	Cyclic suffix
DAC	Digital-to-analog converter
dBm	Logarithmic unit of power with respect to reference impedance
DC	Direct current
DFE	Decision feedback equalizer
DFT	Discrete Fourier transform
DMT	Discrete multi-tone
DS	Downstream
DSL	Data subscriber line
DSP	Digital signal processing
EMC	Electromagnetic compatibility
EMI	Electromagnetic interference
ERL	Echo return loss
FEXT	Far-end crosstalk
FFT	Fast Fourier transform

ft	Feet
HAM	Amateur Radio
IDFT	Inverse discrete Fourier transform
IFFT	Inverse fast Fourier transform
ICI	Inter-carrier interference
ISI	Inter-symbol interference
ITU	International Telecommunications Union
kbps	Kilo bits per second
kft	kilofeet, a thousand feet
Mbps	Mega bits per second
MCM	Multi-carrier modulation
NEXT	Near-end crosstalk
ONU	Optical network unit
PAR	Peak-to-Average power ratio
PBO	Power back-off
POTS	Plain old telephone service
PSD	Power spectral density
QAM	Quadrature amplitude modulation
RF	Radio-frequency
RFI	Radio-frequency interference
Rx	Receive
SNR	Signal-to-noise ratio
Tx	Transmit
US	Upstream
VDSL	Very-high-bit-rate digital subscriber line
VTU	VDSL transceiver unit
VTU-O	VDSL transceiver unit at the ONU (service-provider end of the line)
VTU-R	VDSL transceiver unit at the remote customer site

List of Symbols

a, b	RFI cancellation model parameters
A, B, C, D	two port network parameters
b	channel loading measured in number of bits per signaling dimension
b_{bin}	bin capacity measured in bits
β	number of samples during which consecutive symbols overlap
C	channel capacity measured in number of bits per real dimension
Δf	frequency spacing between bins
f_o	carrier frequency estimate
f_s	sampling frequency
Γ	gap, ratio between required SNR and theoretical SNR to attain channel capacity
Γ^*	gap incorporating SNR margin and coding gain
$h(n)$	discrete time channel impulse response
H	channel transfer function DFT
L_{CE}	total cyclic extension length
L_{CP}	length of cyclic prefix
L_{CS}	length of cyclic suffix
M	length of discrete time channel impulse response
N_{sc}	number of defined tones/bins
ω_k	notch frequency
r	pole contraction factor
$r(n)$	received signal in discrete time
$rect_{2N_{sc}}(n)$	rectangular window of size $2N_{sc}$
$RLCG$	per-unit-length two port line model parameters
$s_1(t), s_2(t)$	signal contents of two different bins
SNR_i	signal to noise ratio in the i^{th} bin
θ	phase delay due to the time shifting of the receiver window
U_i	window-induced ICI experienced by the i^{th} bin

$v(n)$	received windowed signal in discrete time
$v_p(n)$	signal obtained by the periodic repetition of $v(n)$
$V(\omega)$	Fourier transform of the received windowed symbol
V_{echo}	echo signal represented as a single-ended voltage
V_{out}	VDSL driver output signal represented as a single-ended voltage
V_{tx}	transmitted signal represented as a single-ended voltage
$w(n)$	window in discrete time
W_i	window transfer function DFT value at the i^{th} bin
Φ	carrier phase estimate
x_k	transmitted sample sequence of k^{th} symbol
$x(n)$	transmitted signal in discrete time
Z_i	complex constellation value transmitted on the i^{th} tone
\tilde{Z}_i	decoded (sliced) constellation value of the i^{th} tone
Z_L	impedance looking into the line
Z_{ref}	VDSL reference impedance
Z_S	source impedance at the transmitter
Z_T	termination impedance at the receiver

Chapter 1

Introduction

The business boundaries that once clearly segmented wireline and wireless voice, data and video services offerings have all but disappeared. Indeed, it is not uncommon today for a telephone or cable company to offer a host of these services via their proprietary infrastructure. In the case of telephone companies, offering video services to customers over their twisted pair infrastructure remains a challenge due to the channel's impairments and limited bandwidth. Many of the telephone companies are now considering a Multi-Carrier Modulation (MCM) technology called Very-high-bit-rate Digital Subscriber Line (VDSL) as a solution for transporting the high data rates required for video transmission over short distances. The idea is to re-use the existing copper pair infrastructure routed to the customer's premise, but terminate the connection fairly close to the customer, say at the curb, as to achieve the high data rates provided by VDSL. From there, multiple customer connections can be aggregated and fed back to the central office via a high bandwidth media such as optical fibre.

A specific impairment with potentially severe repercussions on VDSL communication is ingress Radio Frequency Interference (RFI). This impairment has random characteristics in terms of presence, frequency location, and power level, making it difficult to qualify and mitigate its impact on system performance. Many approaches have been proposed to combat RFI in VDSL systems, each with their own set of promises relating to performance and complexity, and conversely, their own set of drawbacks.

1.1 Thesis Objectives

The exercise of identifying an effective and economical RFI suppression technique is a non-trivial task. Indeed, the analysis of various RFI suppression methods found in the open literature often lack the necessary thoroughness to properly back up the claimed benefits in the context of a VDSL system. The analysis deficiencies vary from false assumptions to a failure to take into account certain VDSL system features and scenarios. These inconsistencies thus make it difficult to compare the performance value of different methods.

The objective of this thesis is to identify, investigate, evaluate and compare various digital approaches to mitigating ingress RFI in the context of MCM-based VDSL systems. Specifically, the following two goals are pursued:

1. Compare the RFI suppression effectiveness and complexity of the identified techniques.
2. Investigate possible non-RFI benefits and drawback of these same techniques.

This thesis focuses on RFI suppression applications located at the service provider's end of the line. The nature and intensity of ingress RFI at this location are more confined than at the customer premise equipment (CPE) location as will be explained in section 2.3.3.4. This reality needs to be considered when selecting an RFI suppression solution at the service provider's end of the line.

1.2 Thesis organisation

This thesis is organised as follows. A brief introduction to multi-carrier modulation is presented in chapter 2, followed by a description of the transceiver characteristics as defined by the ANSI VDSL trial standard [4] and similar ITU draft standard [11], as well as a description of channel induced impairments. Next, chapter 3 presents a review of the literature pertaining to RFI suppression in DSL systems. The problem statement along

with supporting arguments is formulated in chapter 4. A complete end to end VDSL system model, representative of practical systems, is developed and described in chapter 5. This chapter also takes a closer look at the important characteristics of the loop models and their potential impact on system performance. Chapter 6 is dedicated to the study of the digital RFI suppression techniques under consideration. Chapter 7 presents various simulation results along with corresponding analysis. Conclusions and recommendations for further study are given in chapter 8.

Chapter 2

Background Information

2.1 Introduction to Discrete Multi-Carrier Modulation (MCM)

The ANSI VDSL standard [4] specifies two versions of transceivers. The one of interest for the purpose of this thesis is based on multi-carrier modulation (MCM). This technique partitions a given channel into subchannels by using a set of orthonormal basis functions to perform modulation and demodulation. If these functions are selected such that they remain orthogonal after transmission through the known channel, the resulting subchannels will be independent and ideal¹, thus providing optimum transmission performance. The sets of functions that achieve such partitioning are the eigen-functions of the channel. They can be very difficult to determine and may require infinite complexity and delay for realisation with most DSL channels (see Proakis [20] and Starr et al. [26] for a more in-depth analysis of multi-carrier communication principles).

A more practical form of MCM called discrete multi-tone (DMT) is used by VDSL as well as Asymmetric Digital Subscriber Line (ADSL) transmission (see Refs. [3] and [12]) to maximise the use of the channel's capacity. DMT relies on the Discrete Fourier Transform (DFT) to subdivide the bandwidth of a channel into a number N of subchannels, referred to as bins, each one containing a carrier modulated by a data stream. By subdividing the channel, each subchannel can be approximated as ideal, meaning that the frequency response and noise spectral density over the subchannel can

¹ A channel is said to be ideal if its frequency response has a constant amplitude and a phase that is a linear function of frequency over the channel's bandwidth.

be considered constant, thus avoiding the need for complicated equalization. This allows for simple capacity analysis of each subchannel with the objective of optimal distribution of transmitted power and data over all subcarriers.

In MCM/DMT based VDSL, Quadrature Amplitude Modulation (QAM) is applied to each subcarrier using a constellation (complex-valued), and accordingly a bit allocation, determined on a per subchannel basis during an initialisation and channel estimation procedure. All subcarriers however operate at a common symbol rate. Such a modulator/demodulator system is elegantly implemented with the Fast Fourier Transform (FFT) algorithm and its inverse counterpart (IFFT). The FFT and IFFT are efficient algorithms to implement the DFT and inverse DFT (IDFT). As the symbol length goes to infinity, it can be shown that DMT converges to the optimum MCM that uses the channel's eigen-functions as modulation basis functions ([26]).

In this thesis, the terms “subchannel” and “bin” are used interchangeably as is the case for the terms “tone”, “carrier” and “subcarrier”. Also, any reference made to the VDSL standard in the following sections essentially pertains to the Multi-Carrier Modulation specification of the standard.

2.2 The Transceiver and its Physical Layer Characteristics

A few high-level concepts and system definitions need to first be established before taking a closer look at the details of the MCM VDSL transceiver. To start, the VDSL standard specifies a point to point transmission link. Although bi-directional by nature, VDSL treats separately the flow of data in each direction, which is frequency multiplexed.

The VDSL standard differentiates between the VDSL transceiver units (VTU) at each end of the transmission link. The VTU-O is the unit at the service-provider end of the line and may be located at a central office (CO) or in an external cabinet when VDSL is deployed from a mid-loop location. The VTU-R is the transceiver unit at the remote

customer location, also referred to as the customer premise equipment (CPE). The flow of data from the VTU-O to the VTU-R is referred to as the downstream, while the flow of data in the opposite direction is referred to as the upstream. The channel bandwidth is thus divided and allocated to either the upstream or downstream according to predefined band plans [11].

As is detailed in the following sections, the VDSL standard allows for flexibility in setting many of the system parameters for the purpose of optimising the transmission protocol based on loop characteristics as well as features supported by the modems. Settings for these parameters are often negotiated by the two modems during an initialisation procedure prior to establishing data communication.

2.2.1 Physical Medium Dependent (PMD) Functional Model

The ANSI VDSL standard [4] and its ITU counterpart [11] define the functional model of the PMD sublayer as presented in Figure 2-1. Note that details relating to the receiver path are not provided in this model. It is left up to the designer to specify the processing required for efficient and reliable demodulation of the received signal.

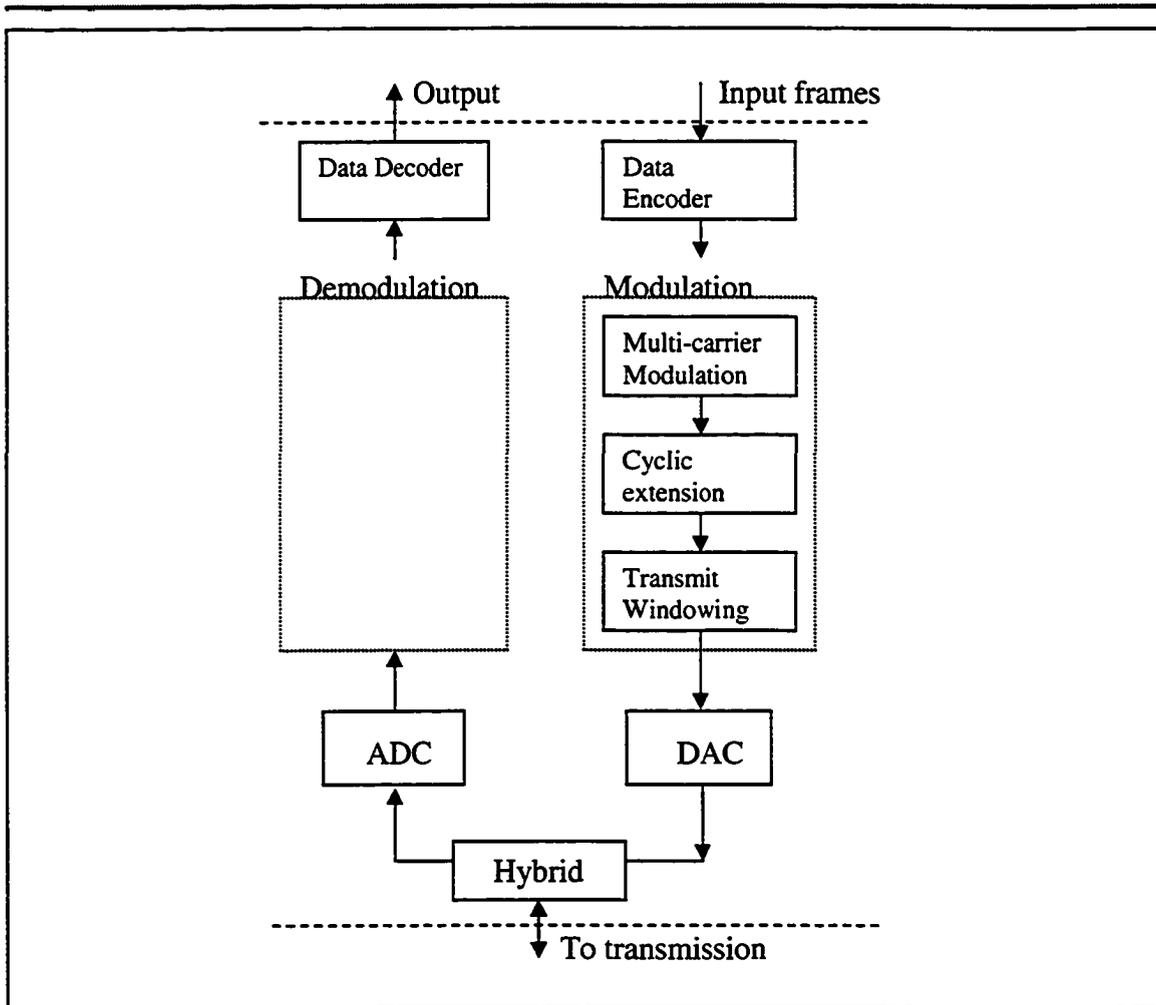


Figure 2-1: PMD Functional Model Diagram

In the transmit direction, the model assumes the reception of input frames containing the bytes which are to be modulated onto a DMT symbol. A number of bits is pre-assigned to each tone and encoded into a complex-valued constellation point that is used to modulate the tone for the duration of the symbol period. Modulation and summation of the tones are achieved by applying an IFFT on the block of data consisting of all assigned constellation points. The output of the IFFT is then cyclically extended and shaped in time before being converted to a serial stream and fed to a Digital to Analog Converter (DAC) for transmission over the channel. Further details on the above processing are presented in the following sections.

In the other direction, the received signal is assumed to be sampled and passed through an Analog to Digital Converter (ADC). The digital signal is then demodulated by computing its FFT on a symbol by symbol basis (block processing). The demodulated signal now represents multiple constellation points (one for every bin), and is sent to the constellation decoder for data frames extraction. Although not defined by the standard, extra processing is required prior and after the FFT computation to compensate for the non ideal characteristics of the loop (transmission medium) as well as for other noise impairment such as crosstalk and additive noise. A receiver design is proposed in section 5.4.

2.2.2 Multi-Carrier Modulation Specification

The number of carriers N_{sc} allowed by the VDSL standard varies in increments of power of 2, ranging from 2^8 (256) to 2^{12} (4096). The frequency spacing Δf between bins is fixed at 4.3125 kHz, the same as defined for ADSL. As the VDSL transmission spectrum is limited to below 12 MHz, the number of sub-carriers that can be used is less than the maximum number allowed by the partitioning.

Carriers can be assigned to transport data in the upstream direction, in the downstream direction, or to not transport data at all; in which case the carrier is usually turned off. For example, carriers must be turned off in frequency bands used by Plain Old Telephone Sets (POTS) and amateur radio (HAM). As a minimum, the carrier in bin 0 (direct current – DC bin) must be turned off. Active bins (subchannels with active carriers) are assigned the upstream or downstream according to frequency band plans. The duplexing band allocation specified by the VDSL standard [4] is presented in Figure 2-2.

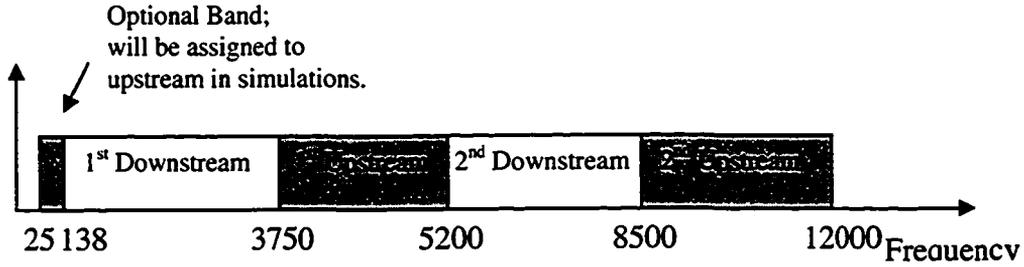


Figure 2-2: VDSL Band Allocation

As previously mentioned, each carrier is QAM modulated using N_{sc} complex Z_i values generated by the constellation encoder. Note that Z_0 corresponding to the DC bin shall always be set to zero. Since an IFFT is used to perform the modulation, the N_{sc} length modulating vector consisting of the Z_i values needs to be augmented by its conjugate counterpart such that the new $2N_{sc}$ length modulating vector produces a real time-domain sequence when transformed by the IFFT. In other words, the augmented modulating vector is now Hermitian, such that:

$$Z_i = \text{conj}(Z_{2N_{sc}-i}), i = N_{sc}, \dots, 2N_{sc}-1 \quad (2-1)$$

The DC and Nyquist frequency carriers shall not be modulated, therefore:

$$Z_i = 0 \text{ for } i = 0 \text{ and } i = N_{sc} \quad (2-2)$$

The IFFT implements the following modulating transform:

$$x_k = \sum_{i=0}^{2N_{sc}-1} Z_i e^{j\frac{2\pi ki}{2N_{sc}}}, k = 0, \dots, 2N_{sc}-1 \quad (2-3)$$

The sequence x_k consists of samples associated to a symbol with a sampling frequency given by:

$$f_s = 2N_{sc} \times \Delta f \quad (2-4)$$

The IFFT block processing is performed on a symbol per symbol basis.

The sequence Z_i is recovered by performing the inverse of the modulating transform, the FFT that is, on associated x_k sequence. Of course, in real applications, the transmitted x_k sequence would never be available to the VDSL receiver due to impairments introduced by the communication channel, including the Analog Front End (AFE) of the VDSL transceivers.

2.2.3 Mitigating Intersymbol Interference and Near-end Echo Transients

A DMT signal transmitted over a real channel is subject to distortion and various types of noise. Here, the channel includes the loop as well as the transmitter and receiver Analog Front End (AFE). The AFE implements the hybrid function, which performs two wire to four wire signal conversion, some analog processing such as filtering, and the digital to analog or analog to digital conversion which will introduce noise and some non-linearities. Given a discrete time signal $x(n)$ transmitted over the channel, the received signal after sampling can be modelled at a high level as:

$$r(n) = x(n) * h(n) + noise(n) \quad (2-5)$$

where $h(n)$ is the discrete time channel impulse response (CIR) of length M , $*$ denotes the linear convolution operation, and $noise(n)$ includes various impairments that are common in VDSL systems (see section 2.3.3) as well as channel non-linearities which are dependent on $x(n)$. The channel impulse $h(n)$ is generally non-zero for more than one sample due to the delay and attenuation characteristics of the channel which vary with frequency [26]. Assuming the transmitted signal $x(n)$ consisted of back to back symbols generated as per equation (2-3), time domain intersymbol interference (TD-ISI) is introduced in the received signal $r(n)$, as any consecutive $2N_{sc}$ samples of $r(n)$ contains contributions from $2N_{sc} + M - 1$ samples of the signal $x(n)$ due to the convolution by $h(n)$. The severity of the TD-ISI is linked to the dispersion time of the channel, which is the time interval over which a “significant” percentage of the channel impulse response (CIR) energy is contained. This definition is further elaborated in section 5.5.1.

It can be shown that the presence of TD-ISI in the received symbols translates into carrier interference at the output of the FFT demodulator (see Van Acker [28]). This frequency-domain interference is generally categorised in three types:

1. Intercarrier Interference of type 1 (ICI1): Interference from one bin to another bin during the same symbol period.
2. Intersymbol Interference (ISI): Interference from one bin in one symbol period to the same bin in a different symbol period.
3. Intercarrier Interference of type 2 (ICI2): Interference from one bin in one symbol period to a different bin in a different symbol period.

These frequency-domain interferences are associated with a loss of orthogonality over the demodulation time interval between the received signal contents, referred to loosely as “tones” or “carriers”, of different bins. The signal contents, or tones, of two different bins, say $s_1(t)$ and $s_2(t)$, are said to be orthogonal to each other if and only if their inner product is zero as given by equation (2-6)

$$\int_a^b s_1(t) s_2^*(t) dt = 0 \quad (2-6)$$

where the integration interval is set to the demodulation time interval. If during this time interval the signal in a given bin contains information from two or more different symbols, either due to channel dispersion or symbol misalignment, it will not possess a pure periodic behaviour over the interval and will therefore not be orthogonal to other tones. The frequency content of such a signal will exhibit spreading about the centre frequency of the bin, which in turn will cause frequency-domain interference due to the rectangular window operation that is intrinsic to the FFT demodulator. To elaborate, the time windowing operation corresponds in the frequency domain to the convolution of the received signal segment’s frequency content with the transfer function of the rectangular

window, which has zeroes at all carrier locations relative to its main lobe:

$$r(n) \times \text{rect}_{2N_{sc}}(n) \xrightarrow{F} R(f) * \frac{\sin(2N_{sc}\pi f / f_s)}{\sin(\pi f / f_s)} e^{-j2\pi f \theta} \quad (2-7)$$

where $\text{rect}_{2N_{sc}}(n)$ is a rectangular window of size $2N_{sc}$ shifted appropriately in time to capture a given demodulation time interval and θ is the phase delay due to the time shifting of the window. If the segment of the received signal to be demodulated contains frequency components other than those at the defined tone locations, the leaked frequency content is picked-up by the spectral sidelobes of demodulator's time window, and injected as interference in each of the tones as per equation (2-7). This phenomenon is similar to the smoothing effect introduced by windowing when used for performing spectrum estimation [21].

In the case of ICI1, tone interference is due to frequency leakage from other tones of a same symbol, while for ISI and IC2, interference is due to frequency leakage derived from tones of other symbols. In any event, the frequency-domain interference leads to an increase in data error probability, or equivalently, a decrease in data transmission rate.

To reduce the negative impact of TD-ISI and ensuing frequency-domain interference, the VDSL standard specifies a symbol to be an extended version of a given time domain data sequence x_k produced by the IFFT modulator (see section 2.2.4 for more details). The reasoning behind this symbol construction is to try to contain the occurrence of TD-ISI at the receiver within the extension, and discard the corresponding samples prior to demodulation. The remaining data sequence is associated in theory to tones that are orthogonal to each other over the demodulation time interval, making it suitable for demodulation by the FFT transform. Proakis [20] emphasises that a DFT-based multi-carrier modem is susceptible to ICI unless an adequate symbol extension is used to contain TD-ISI. This is due to the first spectral sidelobe inherent to demodulator's rectangular windowing operation, which is only 13 dB down from the peak at the desired subcarrier.

A receive noise impairment which has as an effect similar to TD-ISI is the near-end echo (see section 2.3.2). This impairment arises due to the non-ideal characteristics of the hybrid, whose function is to remove the locally transmitted signal from the receive path of the transceiver. In reality, hybrids always provide some transfer function path between the transceiver's transmitter and receiver. Thus, the signal at the VDSL receiver contains a near-end echo component corresponding to DMT symbols generated by the local transmitter. Although the carriers of near-end echo occupy different frequency bands than those occupied by the far end transmitted carriers (frequency multiplexing of the upstream and downstream), these two sets of carriers are not usually symbol aligned at the receiver, and therefore not orthogonal to each other. In fact, the propagation delay introduced by the channel makes it generally impossible to achieve symbol alignment of these carriers at both ends of the line. The propagation delay refers here to the minimum time required for any "significant" energy component of a transmitted signal to propagate through the channel and is further defined in section 5.4.2. Thus, aligning the streams at one end of the line would lead to a misalignment at the other end of the line equivalent to twice the channel's propagation delay. Should an echo symbol transition occur in the demodulation time interval, the echo segment delimited by this interval will contain frequency components outside of the tone locations. These frequency components will get picked-up by the spectral sidelobes of demodulator's time window and injected as interference in each of the received tones as per equation (2-7). On the other hand, if it was possible to avoid demodulating sample blocks containing echo symbol transitions, the echo tones would then be orthogonal to the received tones of interest over the demodulation time interval, and no interference (i.e. ICI) would ensue. This suggests that the symbol extension be also used to guard against demodulation of samples containing echo symbol transitions.

2.2.4 Cyclic Extension and Symbol Rate

As mentioned in the previous section, a VDSL time-domain symbol is constructed by concatenating the output of the IFFT modulator with an extension. In actuality, the total extension is specified as two parts, a cyclic prefix (CP) and a cyclic suffix (CS), which

consist of the last L_{CP} and first L_{CS} samples of the IFFT modulator's block output and which are concatenated at opposite ends of the block. Thus, for every IFFT output sequence x_k generated, its last L_{CP} and first L_{CS} samples are prepended (cyclic prefix) and appended (cyclic suffix) respectively to itself to form a symbol of length $2N_{SC} + L_{CP} + L_{CS}$. Prior to a symbol being transmitted, its first and last β samples undergo envelope shaping, also referred to as transmit windowing. The VDSL parameter β is negotiated at initialisation, with a maximum limit value of $16 \times 2n$, where:

$$n = \log_2(N_{SC}) - 8 \quad (2-8)$$

The windowed parts of consecutive symbols are to be overlapped during transmission. Beside allowing symbol overlapping, the time-domain windowing of transmitted symbols contributes to reduce Radio Frequency Interference (RFI) emissions, or egress RFI, as explained in section 2.2.8. The extended symbol construction and transmission overlap is illustrated in Figure 2-3.

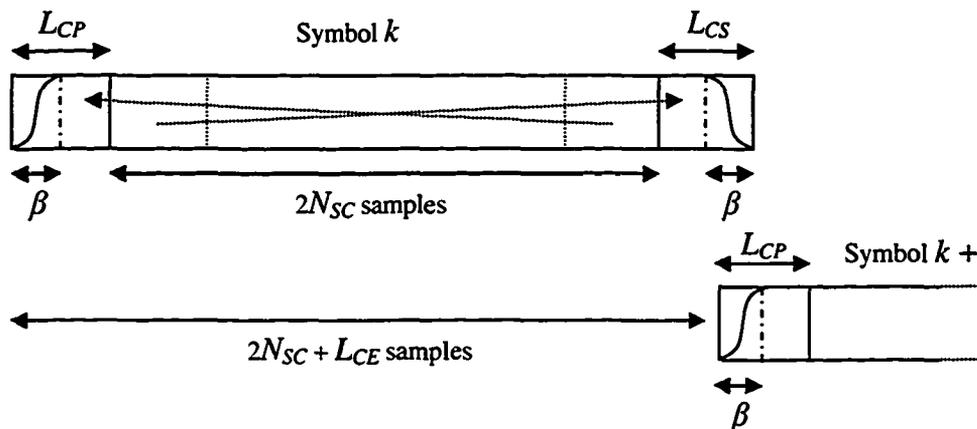


Figure 2-3: Cyclic Extensions and Symbol Overlap

The total cyclic extension length is defined as $L_{CE} = L_{CP} + L_{CS} - \beta$. The values of L_{CP} , L_{CS} , and β are to be selected during initialisation under the restrictions that $\beta < L_{CP}$, $\beta < L_{CS}$ and $L_{CE} = m \times 2n+1$, where m is an integer value. The symbol transmission rate is thus given by:

$$\text{symbol rate} = f_s / (2N_{sc} + L_{CP} + L_{CS} - \beta), \quad (2-9)$$

where the sampling frequency $f_s = 2N_{sc} \times \Delta f$. A VDSL transceiver must support as a minimum a cyclic extension length L_{CE} of value $40 \times 2n$, which results in a symbol transmission rate of 4 kHz. This is the same symbol rate as used by ADSL systems.

2.2.5 Synchronisation of Transmitters and Receivers

In the absence of TD-ISI and echo symbol transitions (i.e. ideal channel conditions), performing FFT demodulation on any non-shaped $2N_{sc}$ samples of a symbol results in transforms with identical magnitudes but different phases, which is due to the cyclic nature of the VDSL symbol. This would allow for demodulation of an arbitrary selection of $2N_{sc}$ samples within the non-shaped part of a received symbol, as long as phase equalization was applied after the demodulator to compensate for differences in timing alignment. Thus, with TD-ISI and/or echo tones present, the $2N_{sc}$ samples to be demodulated can be judiciously selected from the received symbol as to avoid the samples containing TD-ISI and echo symbol transitions, as long as the cyclic prefix/suffix combination is of sufficient length.

It turns out that the requirement on the cyclic prefix/suffix combination length is minimised when transmitters at both end of the line (i.e. VTU-O and VTU-R) start symbol transmission in opposite directions simultaneously, which is a feature required by the VDSL standard. The stream symbol misalignments at both ends of the lines are then made equal to the channel propagation delay. In this case, frequency interference due to TD-ISI and near-end echo can be avoided with proper receiver alignment as long as the cyclic prefix/suffix combination length (measured in number of samples) is equal or greater than the sum of the channel dispersion time and the channel propagation delay.

This condition is best understood by examining the scenario depicted in Figure 2-4. The requirements on the cyclic prefix and cyclic suffix to maintain signal orthogonality for a given channel are further described in [24].

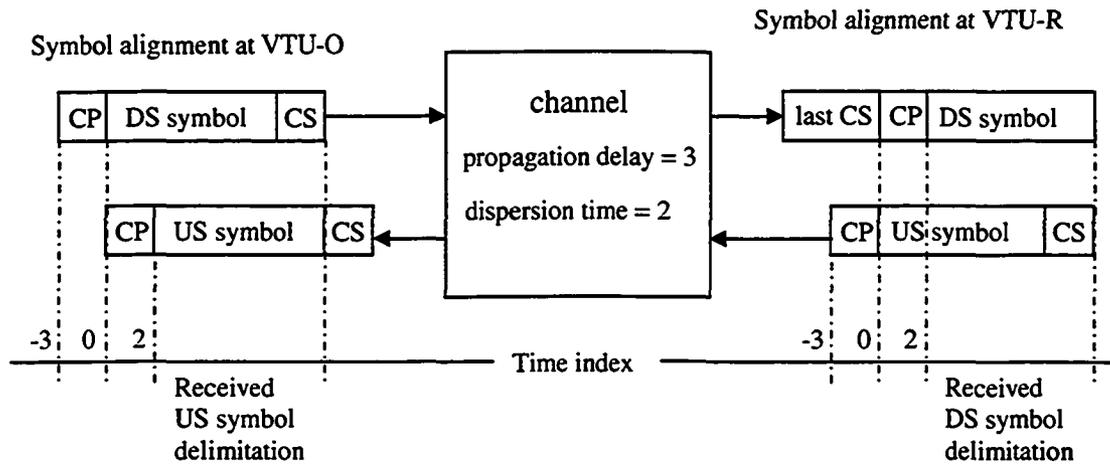


Figure 2-4: Symbol Alignment

Another potential timing related impairment pertains to near-end crosstalk (NEXT) caused by other VDSL systems, referred to as self-NEXT (see section 2.3.3.3). Like near-end echo, self-NEXT originates from carriers that occupy different frequency bands than those occupied by the received carriers, which were transmitted by the far-end. If the VDSL transceivers sharing the same binder group start transmission of the DMT symbols at different times, the self-NEXT is not orthogonal with the received far end signal, unless the symbol transition of each self-NEXT contributor happens to fall within the cyclic prefix/suffix combination of the received symbols. Indeed, when the self-NEXT exhibits symbol transitions within the receiver's demodulating period, its frequency content during this same period is non-zero between tone locations, and this content is picked-up as interference by the spectral sidelobes of the demodulator's time window. To avoid this frequency interference, the VDSL standard defines an optional synchronous mode of operation, for which all transceivers sharing the same binder group start transmission of the DMT symbols at the same time. The idea is to contain the self-NEXT symbol transitions to within the cyclic suffix of the received symbols, along with the echo symbol transitions. The NEXT due to other VDSL systems is thus made orthogonal to the received tones of interest. Synchronous mode requires however that the

cyclic extension size as well as the number of samples to undergo transmit envelope shaping be the same for all modems sharing a binder. This may unfortunately reduce transmission efficiency for some systems, as data rates can only be fully optimised if the cyclic extension size as well as timing synchronisation between the VTU-O and VTU-R is determined on a line-by-line basis.

2.2.6 Constellation Encoding, Channel Capacity and Bit Swapping

In DMT, each bin's capacity (subchannel capacity) is analysed individually at initialisation to assess the size of QAM constellation it can support. QAM constellations supporting 1 to 15 bits per bin are pre-defined in the VDSL standard. However, the maximum size of constellation to be used during any given communication session is negotiated during initialisation, under the condition that this number lie between 8 and 15.

The theoretical capacity for an Additive White Gaussian Noise (AWGN) channel as proven by Shannon is expressed in [26] as:

$$C = \frac{1}{2} \log_2(1 + \text{SNR}) \quad (2-10)$$

where C is the capacity measured in number of bits per real dimension and SNR (signal-to-noise ratio) is the ratio of transmitted energy per symbol to the noise power spectral density. Practical line codes however cannot achieve this theoretical transmission efficiency. For a chosen probability of symbol error, channel loading per signalling dimension must be reduced to (see [26]):

$$b = \frac{1}{2} \log_2(1 + \text{SNR} / \Gamma) \quad (2-11)$$

where Γ , called the gap, is the ratio between the required SNR to achieve a capacity with a given line code and probability of symbol error and the minimum theoretical SNR to achieve this same capacity. Subcarriers in VDSL are modulated using QAM, which achieves a gap Γ of 9.759 dB at a probability of subsymbol error 10^{-7} . This probability of subsymbol error is generally accepted in the industry as a worst-case target for

DMT-based systems [26] and is the upper bound specified by the VDSL standard [4]. Furthermore, since QAM has 2 dimensions, the number of bits that can be allocated on a given sub-carrier is actually twice the number specified by equation (2-11).

In an effort to compensate for the performance inefficiency of a line code, which is characterised by the gap Γ , error correction coding is widely implemented to increase achievable bit rates. On the other hand, it is a good design practice to allow for SNR margin, which is an amount of additional signal-to-noise ratio used to calculate channel capacity. This excess SNR compensates for some channel capacity estimation error, but reduces the achievable bit rates. The SNR margin also allows for the system to tolerate an increase in noise without exceeding the allowed symbol error rate. It is not uncommon to allow for a margin of 6 dB, as specified by the VDSL standard [4], and a coding gain of 3.5 dB when determining bit loading [26]. Thus, a modified gap Γ^* can be constructed to take into account coding gain and SNR margin:

$$\begin{aligned}\Gamma^* &= \Gamma + \text{Margin} - \text{Coding Gain} & (2-12) \\ &= 9.759 \text{ dB} + 6 \text{ dB} - 3.5 \text{ dB} \\ &= 12.259 \text{ dB}\end{aligned}$$

The capacity of a given DMT bin measured in bits then becomes:

$$b_{\text{bin}} = \log_2(1 + \text{SNR}_{\text{bin}} / \Gamma^*) \quad (2-13)$$

where SNR_{bin} is the signal to noise ratio in that specific bin. In practice, the result of equation (2-13) can be rounded to the nearest integer if the SNR_{bin} is compensated accordingly. Adjustment of the SNR_{bin} is achieved by performing fine adjustments to the transmit energy setting of the bin in question, as allowed by the VDSL standard [4]. The capacity for the system is simply the sum of the all bin capacities. It is interesting to note the following in a VDSL system: 1) an SNR_{bin} increase of approximately 3 dB is required to increase the bin bit allocation by one; 2) a minimum SNR_{bin} of 14 dB is required to carry any data at all if the fine adjustment of the transmit energy is ignored; and 3) that the maximum number of bits per sub-carrier is achieved with an SNR_{bin} of 57.5 dB, again ignoring fine adjustment of the transmit energy.

The initialisation procedure for determining bit allocation has one major drawback: it is based on the channel carrying capacity at one instant in time. Should the channel transmission characteristics or the noise level change with time, the ability of individual bins to carry the bit load established at initialisation while respecting a prescribed bit rate may be compromised. This is one of the reasons for making an SNR margin allowance when initially allocating bits to bin. Thus, the SNR on any given bin can degrade by an amount equivalent to the margin before the bit error rate degrades beyond the established target. Should the margin be insufficient to compensate for the channel degradation, the number of bits allocated to the problematic bins needs to be reduced dynamically. Since the VDSL standard does not allow for dynamic rate adaptation, any reduction in bits allocated to a given bin must be compensated by an equivalent, but possibly distributed, increase in bits allocated to other bins. For example, a bin whose SNR margin is of 6 dB can be allocated an extra bit and have its margin reduced to 3 dB. Furthermore, transmit energy can be changed in the appropriate bins to equalise margins over all bins. This procedure is referred to as bit swapping and is essentially a mechanism to dynamically redistribute the SNR margin equally over all bins when the transmission medium characteristics change. It is important to note that bit swapping relies on the assumption that the rate of SNR degradation be not equal over all bins. Under certain assumptions, a VDSL system can compensate over time for an overall average SNR degradation equal to the SNR margin set at initialisation, without having the bit error rate exceed the established target. This assumes of course that all bins were initially allocated a number of bits inferior to the negotiated upper limit. Conversely, if all bins were already allocated the maximum number of bits, a reduction in margin would neither increase rates nor compensate for a change in line condition. Another fact to consider is that bit swapping cannot occur instantly. When an impairment suddenly appears on the line, sub-optimum performance is always observed for the duration of time it takes to perform bit swapping.

More information on bit loading strategies for DMT systems can be found in [26] and [28].

2.2.7 Transmit signal characteristics

The VDSL standard [4] specifies a purely resistive source and load termination impedance of 100 Ohm over the entire VDSL frequency band. The maximum allowed average power transmitted by a VTU-O in the downstream is 14.5 dBm (power in dB relative to 1 mW) when deployed from a central office (CO) and 11.5 dBm when deployed from a cabinet. More stringent limits are required for lines deployed from a cabinet because they may share the same binder as lines deployed from a further distance at the CO and whose downstream signals will necessarily be more attenuated, and thus more susceptible to NEXT (see section 2.3.3.3) generated at the cabinet. The upstream average power transmitted by a VTU-R is limited to 14.5 dBm, regardless of the deployment scenario. The VTU-R however must apply power back off (PBO) as a function of the estimated loop length. The purpose of PBO is to compensate for the upstream attenuation differences in lines sharing the same binder due to the different loop lengths. This promotes equalised FEXT over all lines in the binder.

Several power spectral density (PSD) templates are provided in the VDSL standard [4] for the different stream directions and deployment scenario. They represent the maximum allowed PSD for the signal transmitted by a VTU-O or VTU-R transceiver. The PSD limits are specified for both in-band and out-of-band regions, the latter being the frequency bands not associated with transmission of the VDSL stream in question. Figure 2-5 shows as an example the upstream M1 PSD template.

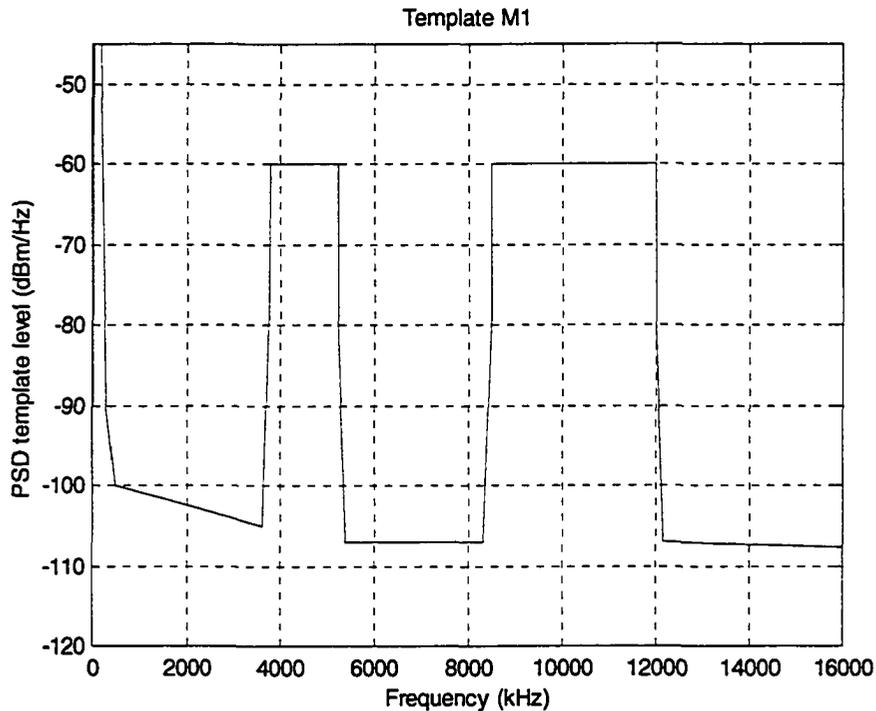


Figure 2-5: Upstream M1 PSD Template

A 175 kHz wide transition band is also specified for each of the in-band to out-of-band transitions. It extends from a given transition frequency into the adjacent out-of-band region. The transmitted PSD is allowed to roll-off over this region. The roll-off upper limit is defined by a straight line going from -80 dB/Hz at the transition frequency, down to the specified out-of-band PSD limit at the other end of the transition band.

Of special concern to VDSL are the spectral compatibility issues with amateur radio service. Amateur radio and VDSL are susceptible to radiated emissions from each other. The amateur frequency bands, also referred to as HAM bands, are presented in Table 2-1.

Table 2-1: Amateur radio bands recognised by ANSI

Start Frequency (MHz)	End Frequency (MHz)
1.810	2.000
3.500	4.000
7.000	7.300
10.100	10.150
14.000	14.350
18.068	18.168
21.000	21.450
24.890	24.990
28.000	29.700

Radiated emissions from VDSL systems are referred to as egress, while radiated emissions from other sources that interfere with VDSL are referred to as ingress. To avoid potential egress interference to amateur radio services, the VDSL transceiver must reduce its transmitted PSD within the amateur radio bands to below -80 dBm/Hz [4]. This is generally achieved by turning off selected interfering bins and performing transmit windowing to reduce the spectral sidelobe leakage inherent to IFFT modulation (see [21] for more information on the frequency response effect of windowing). As for dealing with ingress interference, commonly referred to as RFI, this is the central topic of this thesis as it pertains to upstream transmission.

2.2.8 Egress control

As mentioned in the previous section, it is required that the VDSL PSD within the HAM bands be inferior to -80 dBm/Hz. This can be partially achieved by turning off the sub-carriers within those bands and time windowing the transmitted DMT symbol as to shape and reduce the sidelobe levels of its corresponding spectrum. The raised cosine window specified by equation (2-14) is typically used for transmit shaping as it possesses desirable low sidelobe features and provides good control over the number of samples to be shaped:

$$\begin{aligned} Window_{rc}(n) &= 1, \quad \left| n - \frac{M-1}{2} \right| \leq \alpha \frac{M-1}{2} \quad 0 < \alpha < 1 \\ &= \frac{1}{2} \left[1 + \cos \left(\frac{n - (1+\alpha)(M-1)/2}{(1-\alpha)(M-1)/2} \pi \right) \right], \quad \alpha \frac{M-1}{2} \leq \left| n - \frac{M-1}{2} \right| \leq \frac{M-1}{2} \end{aligned} \quad (2-14)$$

Figure 2-6 depicts in the time domain the raised cosine transmit window with a beta factor of 0.029, corresponding to 512 samples out of 9088 being shaped, the maximum allowed by the VDSL standard. The window's transfer function magnitude is depicted in Figure 2-7. Assuming the worst case scenario of a -53 dBm/Hz in-band PSD (maximum allowed by the VDSL standard) adjacent to the HAM bands, the sidelobes of the transmitted PSD would need to be 27 dB less than the peak PSD. Taking a closer look at the first few spectral sidelobes of the transmit window (Figure 2-8), it is apparent that this upper limit is only achieved as of the 7th bin away from the main lobe. This leads to the requirement that a minimum of 7 in-band tones, adjacent to each side of the HAM band, be turned off in order to meet the egress attenuation requirements within the HAM bands. These tones are referred to as RFI guard bins. The number of RFI guard bins could also be increased to 10 to provide 30 dB of egress attenuation with respect to the peak PSD, and thus allow for some margin.

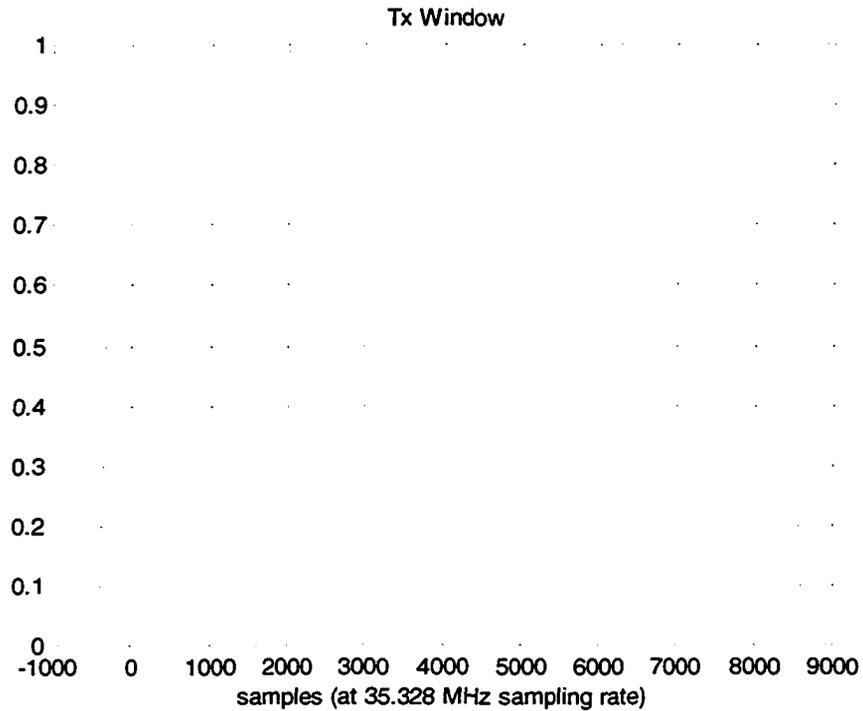


Figure 2-6: Transmit Time-Domain Window

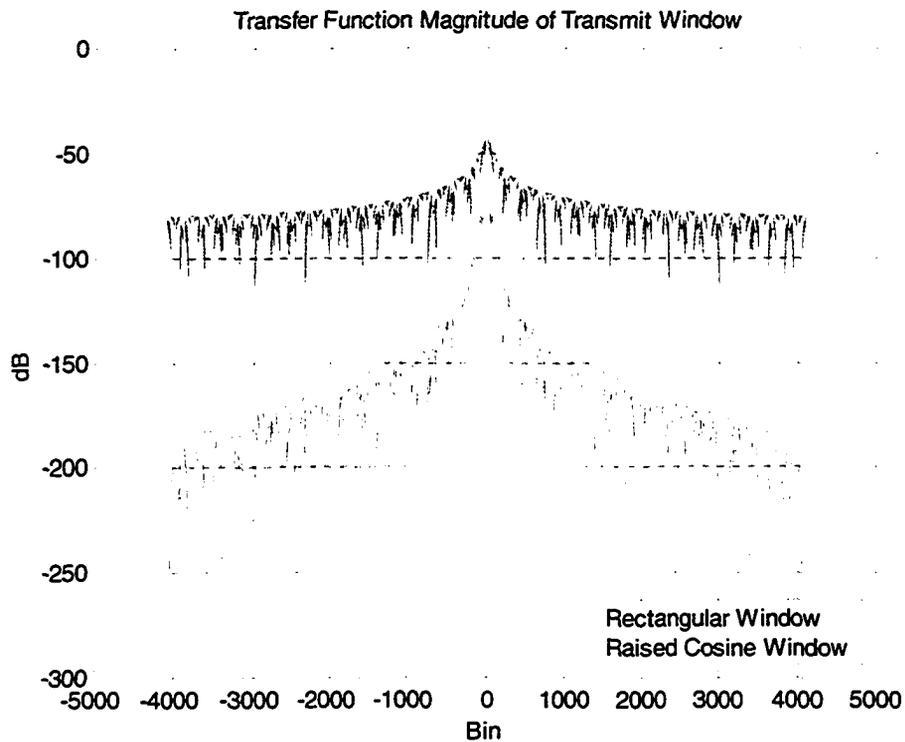


Figure 2-7: Transfer Function Magnitude of Transmit Window

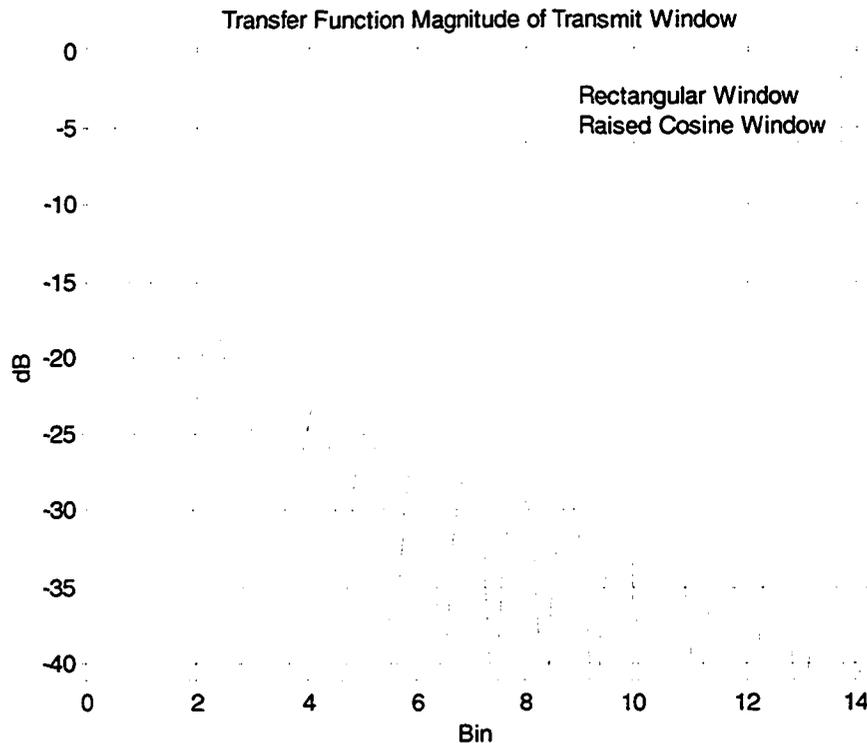


Figure 2-8: First 14 Spectral Sidelobes of Transmit Window

2.3 Received Signal Characteristics and Noise Impairments

As mentioned in section 2.2.3, a transmitted VDSL signal $x(n)$ consisting of back to back cyclically extended symbols will undergo channel distortion before being received, along with noise and near-end echo, at the other end of the channel. The received signal is modelled as:

$$r(n) = x(n) * h(n) + echo(n) + noise(n) \quad (2-15)$$

where $h(n)$ is the discrete-time channel impulse response of length M , and $*$ denotes the linear convolution operation.

2.3.1 Time-Domain ISI

As explained in section 2.2.3, the non-zero impulse response of the overall transmission channel will cause time domain ISI to be present in the received VDSL signal.

Specifically, any given received signal sample will contain a contribution from M individual signal samples transmitted by the far end, where M is the length of the impulse response for the discrete-time model of the channel. To what degree this ISI translates into an impairment will depend on how well it can be contained within a symbol's cyclic prefix as delimited by the receiver.

2.3.2 Near-end Echo

In two-wire transmission, transmit and receive signals are both present at the interface between the transceiver and the line. The transmit signal at this interface can be represented as a single-ended voltage V_{tx} as in Figure 2-9, where V_{out} is the VDSL signal at the output of the driver, Z_S is the source impedance and Z_L is the impedance looking into the line, including the transformer when present. The transmit signal at the two-wire interface is thus given as:

$$V_{tx} = \frac{Z_L}{Z_L + Z_S} * V_{out} \quad (2-16)$$

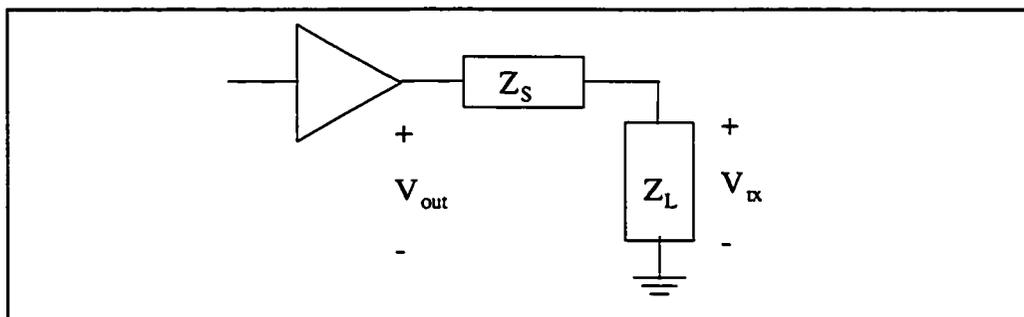


Figure 2-9: Transmit Signal at Line Interface

As previously mentioned, the hybrid circuitry removes an estimate of the transmit signal from the received signal. This estimate is based on a reference impedance Z_{ref} that models the line impedance. However, a mismatch between the model impedance Z_{ref} and the actual line impedance Z_L will cause a portion of the transmit signal, referred to as echo, to make its way to the receive path. The echo signal is given by:

$$V_{echo} = \frac{Z_L}{Z_L + Z_S} * V_{out} - \frac{Z_{ref}}{Z_{ref} + Z_T} * V_{out} \quad (2-17)$$

where Z_T is the termination impedance at the receiver. For the case of $Z_{ref} = Z_T$, the echo voltage becomes:

$$V_{echo} = \frac{Z_L - Z_S}{Z_L + Z_S} * V_{out} * \frac{1}{2} \quad (2-18)$$

The echo return loss is given by

$$ERL_{dB} = 20 \log \left(\frac{V_{echo}}{V_{tx}} \right) \quad (2-19)$$

The time-domain echo signal can be calculated by simply calculating the IFFT of V_{echo} .

The signal seen by the VDSL receiver will thus consist of a component corresponding to the signal transmitted by the far end transceiver as well as a near-end echo. In order to completely avoid the introduction of echo related ICI in the demodulated received signal, the echo must be made orthogonal to the far-end transmitted signal by ensuring that no echo symbol transitions occur in the selected block of received samples to be demodulated. In other words, the receive symbol delimitation should be aligned in time such that the signal samples containing echo symbol transients are contained within the cyclic suffix, which would ensure orthogonality between upstream and downstream signals.

2.3.3 Noise Impairments

Various types of noise are frequently present in the VDSL receive signal. The most common ones can be classified in one of four categories:

1. Additive Noise
2. Impulse Noise
3. Crosstalk
4. Radio Frequency Interference

The following sections briefly highlight the most common noises. More detailed descriptions can be found in section 3.6 of [26].

2.3.3.1 Additive Noise

Additive noise is random in nature. It is often considered to be Additive White Gaussian Noise (AWGN), but can also have a coloured spectrum.

2.3.3.2 Impulse Noise

Impulse noise is a temporary impairment induced on the phone lines by nearby electromagnetic events. Due to the random nature of the noise source, impulse noise cannot be readily modelled as its characteristics vary widely. A simple example of impulse noise would be a burst of Additive White Gaussian Noise of sufficient power relative to the signal of interest.

2.3.3.3 Crosstalk

Crosstalk arises due to signal coupling between different wire pairs bundled together into a binder. The coupled signal may originate from other VDSL systems, in which case it is considered self-crosstalk, or it may originate from other types of systems (ex. ADSL, T1, ISDN, etc.) and is considered an alien crosstalk. Crosstalk is further categorised as near-end Crosstalk (NEXT) when the coupled signal is transmitted from the receiver's end of the binder (i.e. near-end transmitter), or Far-End Crosstalk (FEXT) when the

coupled signal is transmitted from the receiver's opposite end of the binder (i.e. far-end transmitter).

The negative impact of NEXT self-crosstalk on VDSL transmission can be avoided if the NEXT is made orthogonal to the received signal of interest. The VDSL synchronous mode of operation attempts to do just that, as was explained in section 2.2.4.

2.3.3.4 Ingress Radio Frequency Interference

Various types of radio transmissions also share parts of the frequency range used by VDSL transmission, the most common ones being amplitude modulated (AM) broadcasts and amateur radio (HAM) service. Due to cable imbalance, the transmission medium (twisted pair) used by VDSL systems will act as both a receiver and transmitter antenna. VDSL and radio frequency transmissions therefore have the potential to interfere with each other in the shared frequency bands, which is to be avoided. Control of VDSL egress emissions was briefly touched on in section 2.2.7. This section and much of the remainder of this thesis focuses on the impact that ingress RFI has on VDSL transmission and ways to combat it.

The types of interfering radio-frequency (RF) wireless signals and the manner in which they are coupled onto phone lines are well documented in [5], [23], [23] and [26]. Of all ingress RFI sources, amateur radio (HAM) is commonly considered as the most significant one in VDSL system for two basic reasons. First, the highest anticipated RFI level on a phone line is a result of having a HAM transmitter located in close physical proximity to the line and VDSL transceivers, thus producing a RF path to the VDSL channel with little attenuation. This scenario is most likely to occur at the customer's premise location. AM broadcast and other distress bands transmitters sharing the VDSL spectrum will typically be located further away from a VDSL transceiver, thus increasing path loss and reducing RFI. Secondly, the presence and operating frequency of HAM radio signals are random in nature. HAM transmission is intermittent and bursty in nature, and can thus suddenly appear and disappear. Their transmit frequencies can also change regularly within the amateur bands, typically every few minutes. Consequently, a

VDSL transceiver initialising a communication session will not be able to properly assess HAM RFI for the purpose of ascertaining channel capacity. Since the VDSL standard does not allow for dynamic rate adaptation, a sudden change in the nature of a HAM interferer could force an established VDSL session to be dropped. It should also be noted that RFI generated by AM broadcasts is less an issue for VTU-O receivers than for VTU-R since AM radio bands are confined to the downstream VDSL bands.

HAM communication uses single-sideband (SSB) modulation with a bandwidth 4 kHz (speech bandwidth). The modulating signal has On-Off activities in the order of 50-100 ms (syllabic rate) and of 5-10 seconds (speech activities). HAM related ingress RFI levels on twisted pairs are known to vary and to be dependent on network topology, installation practices and cable types. Aerial dropwires, vertical cables in high-rise buildings, and in-house wires are generally considered the most vulnerable channels to ingress RFI. De Clercq and Spruyt [6] report RF ingress differential power measurements as high as -5 dBm (calculated in 100 Ohm) for the case of a measurement performed in a house with the radio transmit antenna located at a distance of approximately 24 m. Foster and Cook [8] raise the maximum estimated worst case RFI power through analytic means up to approximately 0 dBm. These levels however refer to customer premises scenarios. Ingress RFI at the more isolated and protected VTU-O receiver location is expected to be limited to lower levels, although the difference is difficult to quantify. As an example, it is assumed in [24] that the RFI power levels at the central office are 10 dB lower than at the customer side. This seems to be in agreement with the ANSI test procedure [4], which specifies an amateur radio interference noise with a differential mode power of -10 dBm.

A DMT system such as VDSL is very susceptible to RF interferers due to the FFT demodulating process, which inherently performs a rectangular window operation in the time-domain. In the frequency domain, this operation corresponds to a convolution with the $\sin(x)/x$ shaped transfer function of the rectangular window, leading to the known sidelobe leakage effect explained in [21]. Thus, a relatively narrowband yet powerful RF noise at the receiver will have its energy spread over many bins after demodulation.

Consequently, even though VDSL does not transmit data over the HAM bands, the spectral sidelobes resulting from the finite-point FFT of the HAM RFI will interfere with the tones carrying data.

Chapter 3

Literature Review

3.1 Analog RFI Cancellation

An adaptive analog RFI cancellation technique is described by Ödling et al. [19] and Sehlstedt [23]. This method assumes a strong correlation between the common mode signal on the twisted pair wires and the ingress RFI, thus permitting the use of the common mode signal as a reference RFI signal. This reference is used to create an in-phase component and a quadrature component which are adaptively scaled to model and cancel the RFI, as well as to determine the residual RFI after cancellation. A quadrature modulator is used for constructing the RFI model while a quadrature demodulator is used to determine the residual RFI, which is fed back to the adaptive RFI modelling algorithm. The motive behind analog RFI cancellation is to avoid saturating the analog front end of a VDSL receiver with the RFI signal. The specific concern is the possibility of overloading the analog to digital converter, which would result in the received signal being clipped. RFI suppressions of about 35 dB have been reported with the use of this method. Further suppression would be achieved with the use of appropriate digital algorithm in following processing stages.

3.2 Time-Domain Windowing

There are many documented digital domain approaches and algorithms for combating ingress RFI in VDSL systems. A simple and common technique is to apply receiver time-domain windowing to reduce sidelobe leakage of the RFI into active bins, as described by Proakis and Manolakis [21]. In order to preserve inter-bin orthogonality,

the window, usually of raised cosine type, must extend into the cyclic extensions in order to avoid generating intercarrier interference (ICI). This method is explored in [5], [10] and [28]. In particular, de Clercq et al. [5] and Isaksson et al. [10] claim that the head and tail ends of the extended block of samples can be folded back onto the symbol after windowing, but prior to demodulation, as to limit the required size of the fast Fourier transform (FFT) to the number of carriers. This is also the basis for the technique elaborated in [7] and [28], which essentially reproduces the time-windowing operation in the frequency domain, allowing the windowing operation to be combined with a frequency domain equalizer.

One potential drawback of raised cosine windowing is that it could increase frequency interference related to TD-ISI and echo. The reason for this is that portions of the cyclic extensions are now demodulated to recover data instead of acting as guard time to confine TD-ISI and echo transients. On the other hand, the use of an optimal receiver windowing scheme with lower spectral sidelobes should generally improve system immunity to various noise sources non-orthogonal to the received symbol, not just RFI, leading to reduced interference and noise power leakage into the signal of interest. These potential noise sources include additive channel noise, crosstalk, ingress RFI, and the portions of the TD-ISI and echo symbol transitions that are not confined to the cyclic extensions, either due to symbol misalignment or because the combined length of the channel dispersion time and propagation delay is greater than can be accommodated by the cyclic extensions.

An alternative constrained windowing scheme is presented in [14], where the windowing operation is limited to the symbol interval (excluding cyclic extensions) and allowed to generate controlled ICI. This ICI is controlled, in that any given bin is only subject to window induced interference originating from one or two bins at known positions. A frequency-domain decision feedback equalizer (DFE) can then be used to cancel-out the portion of the ICI that is due to the windowing of the data signal. The DFE however cannot remove the window-induced interference originating from the noise signal as it is uncorrelated to the recovered data. This method therefore suffers from an SNR

degradation in an additive noise environment. The proposed scheme has two receiver data paths, one conventional and the other implementing the windowing function followed by the frequency DFE. This allows to only apply the windowing operation on selected bins, thus keeping computations to a minimum and avoiding the SNR degradation due to additive noise on bins free of RFI. The time windowing function is performed in the frequency domain, which is more computationally intensive than its time-domain counterpart, but allows for the output of a single FFT demodulator to be used by both paths. A bin selection logic block routes the output of one of the two paths to the slicer² input on a bin by bin basis. The actual selections are performed during the communication initialisation process, and remain fixed thereafter. Such a scheme however does not handle intermittent and frequency varying RFI sources unless the windowing operation is applied to all bins, which would render the conventional path of the architecture superfluous.

3.3 Adaptive Digital Notch Filtering

Another common approach to combating ingress RFI in VDSL systems is to employ an adaptive digital notch filter to track and suppress the radio interferer as documented in [9], [15], [16] and [22]. These algorithms however were all developed and evaluated in the context of the single carrier version of the VDSL standard. The reason this approach is typically avoided for DMT-based systems is that the notch filtering operation increases TD-ISI, which is unacceptable according to de Clercq et al. [5], since time-domain equalization is to be avoided as explained in section 5.4.3. Furthermore, as the algorithm relies on the use of one second order IIR notch filter per RFI source, the more there are RFI sources to suppress, the longer the impulse response of the total RFI filter will be, and the more TD-ISI will be introduced. The extent of the TD-ISI increase that can be expected in a DMT-based VDSL system due to notch filtering is unfortunately not readily available in the surveyed literature. It is generally assumed however that the notch filter's width must be severely constrained as to avoid interfering with VDSL

² Decision device which slices the signal into predefined partitions.

transmission in adjacent bins, resulting in a filter with a long impulse response. However, such a constraint is unwarranted, as the notch filter would operate in the HAM bands where no VDSL transmission occurs, and in the worst case, could be allowed to extend into the RFI guard bins (section 2.2.8) without affecting data rates. The bandwidth requirements on notch filter could then be relaxed to the equivalent of twice the number of RFI guard bins, which would shorten the filter's impulse response and reduce its TD-ISI contribution. Any undesirable frequency attenuation or phase distortion associated such relaxed notch filtering could be easily compensated for with the frequency domain equalizer.

3.4 Digital RFI Cancellation

Finally, another class of interferer suppression algorithm widely referred to in the literature is known as digital RFI cancellation. This class of algorithms is based on modelling the disturber's noise contributions on selected DMT bins for the purpose of subtracting them from the corresponding tones. To elaborate, digital RFI cancellation consists of modelling the narrowband disturber in either the time or frequency domains, followed by the cancellation operation, which is almost always performed in the frequency domain due to the nature of the DMT signal. The parameters required to model the RFI can be obtained by observing the frequency-domain signal in the reserved bins lying in the HAM bands where VDSL transmission does not occur (see section 2.2.7). The FFT demodulator output essentially acts as an antenna in those reserved bins, allowing for RFI presence and measurement of the interferer's characteristics including frequency position.

The algorithms presented in [13], [17], [18] and [29] estimate the RFI carrier frequency and build a zero or first order time domain model of the RFI envelope on a symbol by symbol basis. The model is used to estimate the frequency domain RFI contribution on selected bins, which is then subtracted from the corresponding tones. To illustrate, [17], [18] and [29] model the RFI envelope as a first-order polynomial for the duration of the DMT symbol:

$$RFI(t) = (a + bt) * \cos(2\pi f_o t + \phi) \quad (3-1)$$

where a and b are model parameters to be determined and f_o and ϕ are the carrier frequency and phase to be estimated. The RFI model is actually solved in the frequency domain, where cancellation occurs. It is then multiplied by the frequency representation of the receiver window before being subtracted from the tones of interest.

It is questionable if a first order model is adequate to approximate the RF interferer over a VDSL symbol period for the purpose of performing RFI cancellation. To elaborate, Wiese and Bingham acknowledge in [29] that such a model is valid as long as the bandwidth of the interferer is much less than the DMT symbol frequency. This condition turns out not to be met for HAM interference since the VDSL symbol transmission rate is 4 kHz, which is approximately equal to the assumed RFI bandwidth. The RFI will thus contain frequency components that complete a full cycle during a single VDSL symbol. One possible reason for the above premise to have been considered is that some early DMT based VDSL proposals prescribed a tone spacing of 43 kHz. In this case, RFI due to amateur radio could indeed be considered as a narrowband noise, and consequently, reasonably approximated by a first order model. Unfortunately, some of the later publications ([17] and [18]) seem to be unaware of the more recent developments in the VDSL standard.

Chapter 4

Problem Statement

With the variety of approaches proposed in the literature to combat ingress RFI in VDSL systems (see Chapter 3), selection of an appropriate algorithm in designing a VDSL system is a non-trivial task. The documented evaluations of these algorithms each have their own set of assumptions and performance criteria, which may or may not be relevant, or even valid, in the context of a real VDSL system. At the very least, the different assumptions and criteria used by different authors make it difficult to carry out a fair assessment and comparison of the various algorithms. For example, the performance provided by a certain RFI suppression algorithm could be adequate for a VTU-O application, but not for a VTU-R one, due to the different received VDSL frequency bands and expected ingress RFI levels at opposite ends of the line. Also, the different implementation complexities, computational or otherwise, associated with the different algorithms need to be considered as they typically influence the selection of a solution when engineering a real product.

This thesis focuses on the analysis and development of algorithms to control HAM related ingress RFI in VTU-O receivers for the reasons given in section 1.1 and section 2.3.3.4.

4.1 Unresolved Issues

The published evaluations of the RFI suppression techniques often fail to take into account many subtleties of a VDSL system. These oversights can lead to a misleading assessment of an algorithm's advantages and drawbacks. For example, the use of adaptive digital notch filters in DMT based systems is usually avoided as it is known to

lengthen the impulse response of the overall communication channel, leading to increase TD-ISI. Although the notch filtering algorithms may be adapted for DMT based VDSL systems, it is not clear to what extent such filtering would increase TD-ISI in a VDSL system for reasons given in section 3.3.

The effect of receiver windowing on the system's ability to deal with TD-ISI and near-end echo/NEXT is another issue requiring clarification. As described in section 3.2, it is conceivable for a receiver windowing operation to increase or decrease frequency interference, which depends on the VDSL system's design and environment. Of particular interest is the prospect of using raised cosine windowing to reduce the negative impact of non-orthogonal echo and TD-ISI, as this technique, in contrast to constrained windowing, diminishes the guard time provided by the cyclic extensions. Indeed, by extending the time-windowing operation into the cyclic extensions, the guard time, whose purpose is to confine TD-ISI and echo symbol transitions, is effectively reduced, which in turn could degrade the system performance on certain loops. This leads to the question of receiver alignment and its impact on data rates. Recall from section 2.2.3 that receiver symbol alignment offers, generally speaking, a compromise between containing TD-ISI and echo/NEXT symbol transitions in the cyclic extensions. Because channel induced TD-ISI tends to taper off in time as per the channel's impulse response, one may at first think that priority should be given to avoid having echo/NEXT symbol transitions occur within the $2N_{SC}$ samples to be demodulated. However, transmit pulse shaping, as well as receive pulse shaping when properly implemented, may well reduce the negative impact of echo/NEXT symbol transitions to the extent that it may be more beneficial to avoid TD-ISI instead. Another point is that echo and NEXT symbol transitions are usually considered instantaneous, i.e. not spread in time, to simplify the alignment analysis. This allows for the cyclic suffix length requirement to just be set equal to the channel propagation delay as was assumed so far. In reality, both echo and NEXT paths have an impulse response length greater than one, thus effectively spreading across many time samples, and perhaps even delaying, the transient energy corresponding to the echo and NEXT symbol transitions. This could in turn reduce the required cyclic suffix length, as the interfering transient energy at the receiver would be pushed out towards the

boundary of the received symbol as illustrated in Figure 4-1. Such potential gain could be diminished however if the receiver window extends into the cyclic suffix.

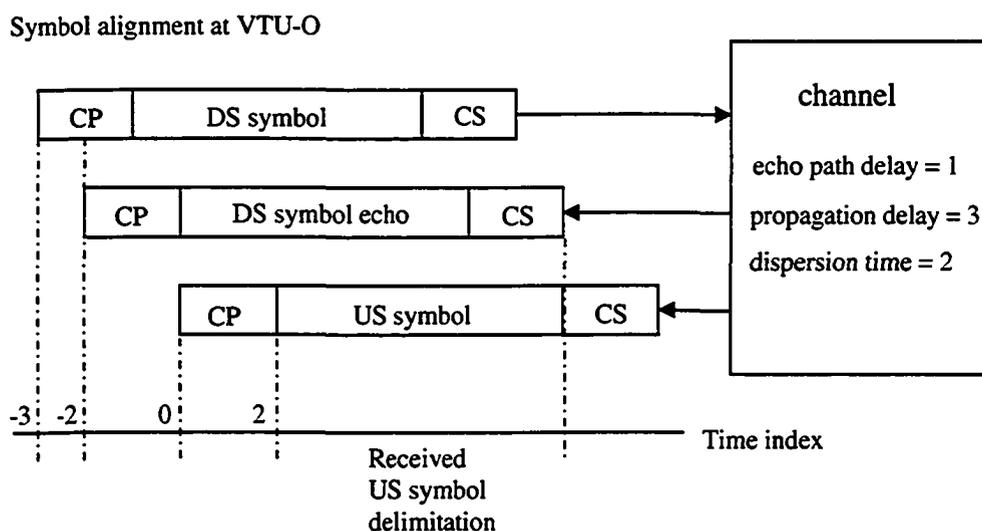


Figure 4-1: Cyclic Suffix Requirement with Echo Path Propagation Delay

As the constrained windowing technique does not extend into the cyclic extensions, it is possible that it would be more effective than its raised cosine window counterpart at dealing with channels exhibiting a long dispersion time or a long propagation delay. The drawback of this approach is a reduction of achievable data rates in a white noise environment, as well as the requirement for a DFE and a dual receiver path. However, as mentioned in section 3.2, the windowing algorithm needs to operate on all bins in order to properly address the RFI with its expected power level and random nature. As a consequence, the dual receiver path architecture can be dropped as a requirement and the algorithm's performance should be re-evaluated in this context.

The development of the RFI cancellation methods documented in [13], [17], [18], and [29] are based on a hypothesis that appears to be false in the context of VDSL. Specifically, it is assumed in these publications that the radio frequency interferer can be adequately approximated by a zero or first order model over the duration of the DMT symbol for the purpose of cancelling RFI. As explained in section 3.4, the validity of this

assumption is questionable. Further investigation is therefore required to assess the merits of this RFI cancellation approach.

The following illustrates the difficulty of comparing the various RFI suppression methods for VDSL. Jeong and Yoo [13] state that time-domain windowing is effective for minimising interference on tones far from the centre of the RFI frequency, but transmission on the tones close to the RFI carrier frequency is said to remain very vulnerable. The significance of this statement in the context of a VDSL system is unclear as no data or analysis was presented to support the claim. Also of note is that the reference in question was published prior to the release of the current VDSL standard, which is also the case for most of the literature surveyed. In this particular case, it is possible that the tones deemed closed to the RFI carrier frequency by Jeong and Yoo would in practice be required as RFI guard bins, and therefore not be used for data transmission in the first place.

With regards to using available literature to compare RFI suppression techniques, special consideration must be given when interpreting the published evaluations that are based on different sets of conditions, scenarios, and simulation models. In some cases, certain simulation parameters are not even defined. The contested parameters include, but are not limited to, RFI power levels and frequencies, RFI bandwidth and modulating schemes, channel models, and other noise impairment specifics such as crosstalk models and noise floor.

Some RFI suppression methods are also more sensitive to specific conditions than others. For example, windowing methods do not perform RFI detection or RFI frequency estimation, and are thus immune to detection and frequency estimate errors, whether they be due to noise or otherwise. Furthermore, the windowing operation remains fixed regardless of the number of RFI sources present, contrary to adaptive notch filtering and RFI cancellation. The above aspects could potentially make windowing more robust and effective than other type of RFI suppression techniques, but such benefits need to be qualified through appropriate simulations

Finally, the following are some of the inconsistencies or discrepancies pertaining to the analysis and simulation results presented in the reviewed literature:

1. Failure to perform end to end simulation (transmitter – channel – receiver) to assess the proposed algorithm performance.
2. Use of frequency domain simulation instead of time domain simulation, neglecting the effect of TD-ISI, which might interact with the features of various proposed algorithms in different ways (e.g. decision feedback equalization, time modelling, frequency estimation, windowing that extends into the cyclic extensions, etc.).
3. Failure to take into account echo interference generated by the data stream in the opposite direction, which could affect windowing performance.
4. Approximation of VDSL signal with a Gaussian noise as in [23] instead of a DMT symbols.
5. Use of different system parameter settings among the various publications (e.g. number of carriers, size of cyclic prefix and cyclic suffix, size of receiver windowing, etc.).
6. Use of different channel features in simulation such as number and type of modelled FEXT, loop model, AWGN level, RFI model, number of RFI sources, etc.

4.2 Motives for Further Research and Investigation

Whether RFI suppression is required or not in a VTU-O receiver is a valid question in itself. At the very least, it would seem worthwhile to investigate the impact of RFI on system performance, if only for informative purposes. However, even if the rate degradation due to RFI may be considered acceptable, some sort of RFI suppression strategy is most likely required to deal with the random presence of RFI as the VDSL

standard [4] does not allow for dynamic rate adaptation³. In any case, an appropriate RFI suppression technique would serve to maximise performance in the presence of an RFI impairment.

There are many aspects to consider in selecting an RFI suppression strategy. In the end, the adopted solution is usually some sort of compromise between feasibility, product cost, and performance. In this case, cost can be measured as a function of computational requirements and development effort, while performance can be measured as a function of robustness and achievable transmission rates. As will be shown in Chapter 7, RFI suppression techniques affect system performance in various ways and to different extent, with and without RFI present. As a consequence, proper evaluation of these techniques requires that their performance be assessed in a broad system context rather than in an isolated environment with ingress RFI present. To this end, a system level architecture needs to be properly developed and modelled for the purpose of evaluating and simulating the RFI suppression algorithms in different scenarios. In particular, special attention should be given to the following aspects when investigating the RFI suppression techniques:

1. Ability to minimise the impact of ingress RFI, with one or more RF sources, on system performance.
2. Ability to minimise the impact of other non-RFI related impairments on system performance.
3. Impact on system performance in an RFI free environment.
4. Resource requirements.

³ Dynamic rate adaptation is being considered for a future standard.

Chapter 5

System Model

A simulation system model capable of accurately representing a real VDSL system under various scenarios is crucial for evaluating the several RFI suppression techniques studied in this thesis. Not only does it provide a level playing field to compare the various algorithms but more importantly, it is a powerful tool for investigating and exposing the strengths and weaknesses of each technique under consideration. A high level view of the proposed system model is shown in Figure 5-1. This model is used in simulation to determine the impact of RFI and related suppression techniques, in terms of achievable data rates, in various scenarios (see Chapter 7). The achievable data rate is determined by calculating the DMT capacity of the channel as elaborated in section 2.2.6. The receiver SNR profile required for this calculation is established through the simulation of 4-QAM symbol transmission.

Modelling of a complete VDSL system is begun by implementing a transmitter and receiver that adhere to the VDSL standard [4]. The standard however does not specify many of the required or desirable features for the receiver design, such as the procedures for assessing the channel's propagation delay and for performing receiver symbol alignment, as well as other mechanisms to handle common channel impairments. As these features often impact the evaluation of the RFI suppression techniques, they need to be identified, developed, and implemented in the simulation environment. The other major component of the system requiring modelling is the channel. First, various loop models are elaborated in order to properly characterise the corresponding channel impulse responses and input impedances. The channel impulse response dictates the attenuation, TD-ISI, and propagation delay, while the loop input impedance is required to

calculated near-end echo. Secondly, a set of impairment signals, such as additive noise, echo, crosstalk, and especially RFI, are generated for injection into the channel model. Unfortunately, the many types of noise impairments along with their varying parameters make it impractical to perform exhaustive simulations of all noise scenarios. The noises and associated settings combinations are therefore judiciously selected to allow useful conclusions to be made from simulation results.

The simulation environment used for this thesis was developed with the Matlab tool. As the focus of this thesis is the investigation of RFI suppression techniques for the VTU-O receiver, only the upstream transmission is modelled end to end. The downstream transmitted signal is only modelled for the purpose of generating near-end echo at the VTU-O receiver.

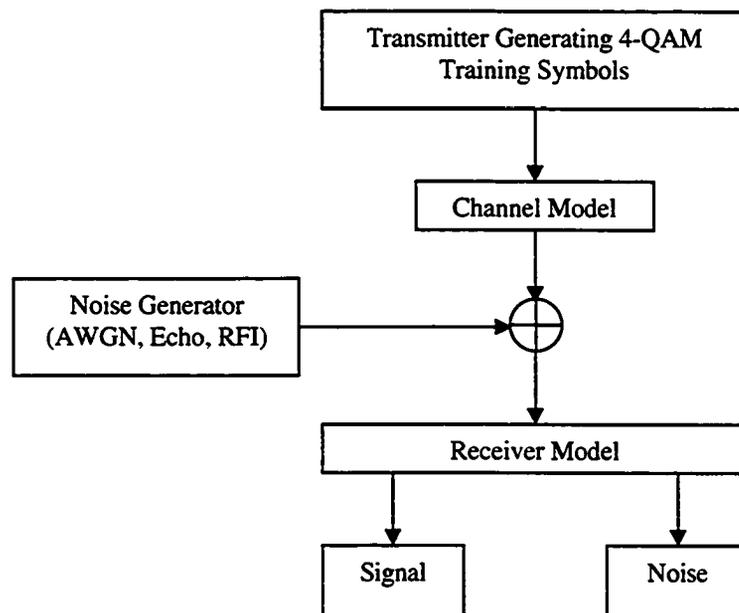


Figure 5-1: Simulation System Model

5.1 VDSL System Parameters

The VDSL standard specifies a set of system parameters. In some cases the parameter values are fixed while in others, ranges of permitted values are specified. For reference, these parameters are listed below along with values and names of variables (in parentheses) as used in the simulation environment. For parameters that are variable, boundary values are selected as they represent limit cases, which tend to best highlight the characteristics of the algorithms under analysis:

1. Subcarrier spacing fixed at 4312.5 Hz (DF).
2. Number of subcarriers set at its maximum value of 4096 (N_{sc}). In this way, potential VDSL vulnerabilities can be exposed throughout the useable bandwidth.
3. Number of samples during which consecutive transmitted symbols overlap set at its maximum value of 256 ($sym_overlap$).
4. Maximum number of bits per carrier set to its highest value of 15 ($bmax_us$) to maximise the overall data rate sensitivity to SNR.

From the above, the following VDSL parameter values are derived:

1. Sampling frequency = $2*N_{sc}*DF = 35.328$ MHz (fs)
2. L_CE set to 640 samples in order to achieve a 4 kHz symbol rate with 4096 subcarriers
3. Cyclic prefix length arbitrarily set equal to the cyclic suffix length, therefore taking on the value of $(L_CE + sym_overlap)/2 = 448$ samples

The following parameters/estimates are used for calculating achievable rates and are assigned typical values [26]:

1. Coding Gain assumed to be 3.5 dB (*cg_dB*)
2. Target system margin set to typical 6 dB (*margin_dB*)
3. Target symbol error rate of 10^{-7}

5.2 Transmitter

The model of the transmitter is defined in the VDSL standard [4]. Its characteristics were summarised in section 2.2. In particular, the simulation environment allows control over whether or not shaping is applied to the envelope of the transmitted signal. The characteristics of the transmit shaping window are defined in section 2.2.8.

Upstream and downstream power spectral densities settings are based on the central office deployment M1 templates presented in [4] (see Figure 2-5 for example). Transmit power is set to -60 dBm/Hz across all bands for simplicity. Note that this level is inferior to the -38 dBm/Hz limit permitted in the first upstream band (25-138 kHz), which is also considered out-of-band for the downstream. The bins in the HAM bands (Table 2-1) as well as the 10 RFI guard bins on each side of the bands have their carriers turned off.

Data frames fed to the IFFT modulators (upstream and downstream) in the simulations are constructed by randomly generating two bits for each active bin, and encoding them using a four point constellation. This scheme is sufficient for performing SNR estimation at the receiver. Furthermore, the upstream and downstream transmitters (at opposite ends of the loop) have their symbols period synchronised. In so doing, it is assumed that TD-ISI is dealt with the cyclic prefix and that the cyclic suffix is used to compensate for the channel propagation delay and contain the modelled near-end echo (see Figure 2-4)

5.3 Channel Model

5.3.1 Loop Model

The VDSL standard [4] defines a set of loops to be used in testing procedures, of which the subset listed in Table 5-1 have been modelled for analysis and simulations purposes. This subset was chosen as it covers both loop reach and the limit cases well.

Table 5-1: Modelled Loops

Loop ID	Description ⁴
VDSL1-x ⁵	26 AWG underground cable of length x varying from 500 ft to 6000 ft in 500 ft increments
VDSL4	Horizontal aerial 26 AWG cable of medium length (3300 ft total), with 300 foot 24 AWG bridged tap ⁶ at 3000 feet and 150 foot 24 AWG bridged tap ⁶ at 3150 feet.
VDSL5 ⁷	650 foot underground 24 AWG cable, followed by 250 foot overhead aerial 24 AWG cable, followed by two 50 foot horizontal bridged tap ⁶ at 90 degrees to each other, one being made of 24 AWG cable, the other made of DW 10 ⁸ cable.
VDSL6 ⁷	1650 feet of underground 26 AWG cable, followed by 650 feet of underground 24 AWG cable, followed by VDSL5
VDSL7 ⁷	1650 feet of underground 26 AWG cable, followed by 2300 feet of underground 24 AWG cable, followed by VDSL5

Each segment of these loops can be characterised as a two port network (“ABDC” matrix modelling), that is derived from the per-unit-length two port model as described in sections 3.5.2 through 3.5.4 of [26]. The values used for the *RLCG* model parameters are those defined in [1]. The ABCD matrix characterising the overall loop is computed by multiplying in cascade the ABCD matrices corresponding to each of the segment. The

⁴ The ANSI VDSL standard [4] specifies cable lengths for loop models in feet, not meters.

⁵ Loop model VDSL1-x is typically used for characterizing performance with respect to loop reach.

⁶ Length of wire pair that is connected to the loop at one end and is unterminated at the other end. These wire “stubs” allow for easy spliced branching connections to be made in the future.

⁷ Loop models VDSL5 through VDSL7 are meant to be used with various crosstalk injections, but will be used in simulation solely for their channel impulse response and input impedance.

⁸ Drop-Wire 10” (British Reinforced 0.5 mm copper PVC-insulated conductors, PVC-insulated steel strength member, Polyethylene sheath)

loop's model parameters are then used to calculate the loop's transfer function (insertion loss) as per equation 3.48 in [26], repeated here for convenience:

$$H_{Loop} = \frac{Z_S + Z_T}{A \cdot Z_T + B + C \cdot Z_S \cdot Z_T + D \cdot Z_S} \quad (5-1)$$

Note that the loop's transfer function is dependent on the transmitter's source impedance Z_S and the receiver termination impedance Z_T . The loop transfer functions used in this thesis are all calculated with these impedances set to the reference impedance value specified by the VDSL standard.

The loop impulse response is determined by simply performing an IFFT on the corresponding transfer function. A 2500 sample truncated version of the loop impulse responses is used in simulation.

The input impedance Z_L looking into the loop is calculated using the ABCD parameters as follows:

$$Z_L = \frac{A \cdot Z_T + B}{C \cdot Z_T + D} \quad (5-2)$$

Again, receiver termination impedance Z_T is set to the nominal VDSL impedance for modelling purposes. The impedance looking into the loop is used to calculate the echo signal voltage as per equation (2-18).

The characteristics of the loops specified by the VDSL standard were investigated and results are presented in section 5.5.

5.3.2 Near-end Echo

Near-end echo is generated using equation (2-18) with Z_L as calculated by equation (5-2), Z_S set to 100 Ohm, and V_{out} set to the downstream signal:

$$V_{echo} = \frac{1}{2} * \frac{Z_L - 100}{Z_L + 100} * V_{tx_downstream} \quad (5-3)$$

Since both upstream and downstream transmitters have their symbol period synchronised, the received upstream symbols will be delayed with respect to the transmitted downstream symbol by an amount equivalent to the channel propagation delay.

The simulation environment allows control over whether or not near-end echo is generated.

5.3.3 Ingress RFI Modelling

The RFI impairment is generated by applying a single-sideband-modulation algorithm [20] to a baseband signal. The carrier frequency is selected from within the HAM bands presented in Table 2-1. Two RFI specific frequencies are used for simulation, one at 3.754 MHz, which is the equivalent of bin 870.5, and the other at 10.106 MHz which is the equivalent of bin 2343.5. These locations correspond to the beginning of the two regions where the HAM bands intersect with the upstream bands, and represent worse case scenarios due to the large RFI sidelobes which interfere with some of the active bins. Gain scaling is used to control differential mode RFI power, which is set to – 10 dBm in all simulations.

Unless otherwise noted, a reference RFI baseband signal with a bandwidth of 4 kHz is used for all simulations. It and its Hilbert transform are constructed as follows. A random white noise sequence is first generated and then processed through two consecutive Butterworth filters, one a tenth order low pass, the other a fifth order high pass. The filtering operation is actually performed at $1/128^{\text{th}}$ the system sampling

frequency. This allows for the use of relatively simple filters while ensuring that the output signal is band limited to within the 300 to 4000 Hz region. The resulting signal is then used to construct its Hilbert transform, followed by up sampling of the signal pair by a factor of 128.

The above base band signal represents a worst case scenario where speech is continuous. In real life, the speech signal would exhibit periods of inactivity as well as interruptions due to its syllabic nature. This would substantially decrease the average RFI power.

It is worth mentioning that the above RFI model is typically more elaborate than the ones referenced by the publications presented in Chapter 3.

5.3.4 Other Noise Models

A host of noise impairments have been known to occur on telephone lines and models for them are readily available (see [4] and [26]). The Additive White Gaussian Noise model is probably the most versatile of them all as any given system will often exhibit a significant noise element well represented by this model. AWGN has the further advantages that it is well understood, simple to simulate, and easy to account for in performance calculation. The simulation environment allows for randomly generated AWGN with a specified PSD to be injected into the channel. An AWGN PSD of -140 dBm/Hz is specified in the test procedures of the VDSL standard [4] and commonly used in simulations, as reported in many of the reference publications ([13], [14], [22], [24], [29]).

The other prevalent noise class found in VDSL systems is crosstalk. The various types of crosstalk impairments are commonly modelled in the frequency domain by filtering with a given crosstalk transfer function a sequence whose spectrum is representative of the signal from which originates the crosstalk. This modelling technique is described by Starr et al. in [26] and is specified in the test procedures of the VDSL standard [4].

Unfortunately, the models of the crosstalk transfer function are only defined in

magnitude, making any time domain analysis inaccurate. Furthermore, the VDSL standard test procedures call for crosstalk impairment to be generated by using the model to filter a random coloured signal, which is not orthogonal to the received symbols. There are no obvious reasons to believe that a noise generated in such a way would expose the characteristics of the RFI suppression techniques any better than simple AWGN. Moreover, this crosstalk modelling approach overlooks the ability of the VDSL synchronous mode of operation to combat NEXT self-crosstalk, which is based on ensuring orthogonality between the self-NEXT and the received symbols (see section 2.3.3.3). It is also conceivable for the impulse response of the near-end crosstalk path to exhibit a significant propagation delay, or a spreading in time. This would further increase the effectiveness of the VDSL synchronous mode of operation, as self-NEXT symbol transitions in this case are more easily confined to the cyclic suffix of the received symbols (see Figure 2-4). For the above reasons, crosstalk impairments are not modelled in the simulation environment.

5.4 VTU-O Receiver Design

The proposed design model for the PMD sublayer of a VDSL receiver is presented in Figure 5-3. The key functions of a VDSL receiver are as follows. First, a hybrid circuit is used to remove the transmit signal from the receive path while at the same time connecting the transmit and receive paths to the line. The hybrid circuitry uses a model of the impedance looking into the line to estimate the transmit signal present in the receive path, and then removes this estimate from the received signal. Next, the received signal undergoes analog processing by the Analog Front End (AFE), such as anti-aliasing filtering and analog RFI cancellation (optional), before being sampled and digitised by the Analog to Digital Converter (ADC). The digital samples are then aligned (see section 5.4.2 for details), equalised in the time-domain (optional), and grouped into symbols (serial to parallel conversion) such that any time domain ISI (TD-ISI) and near-end echo is concentrated in the cyclic prefix and cyclic suffix respectively. At this point, time-domain receiver windowing and RFI cancellation can be optionally applied to each symbol, followed by the removal of the symbol's cyclic extensions and demodulation by

the FFT transform. The frequency data is then processed with an optional frequency domain RFI suppression algorithm and passed through a frequency domain equalizer (FDEQ), which compensates each bin individually for the attenuation and phase shift incurred during channel transmission. Finally, the equalised frequency samples are sent to a constellation decoder (slicer) to extract the bits for the corresponding frame.

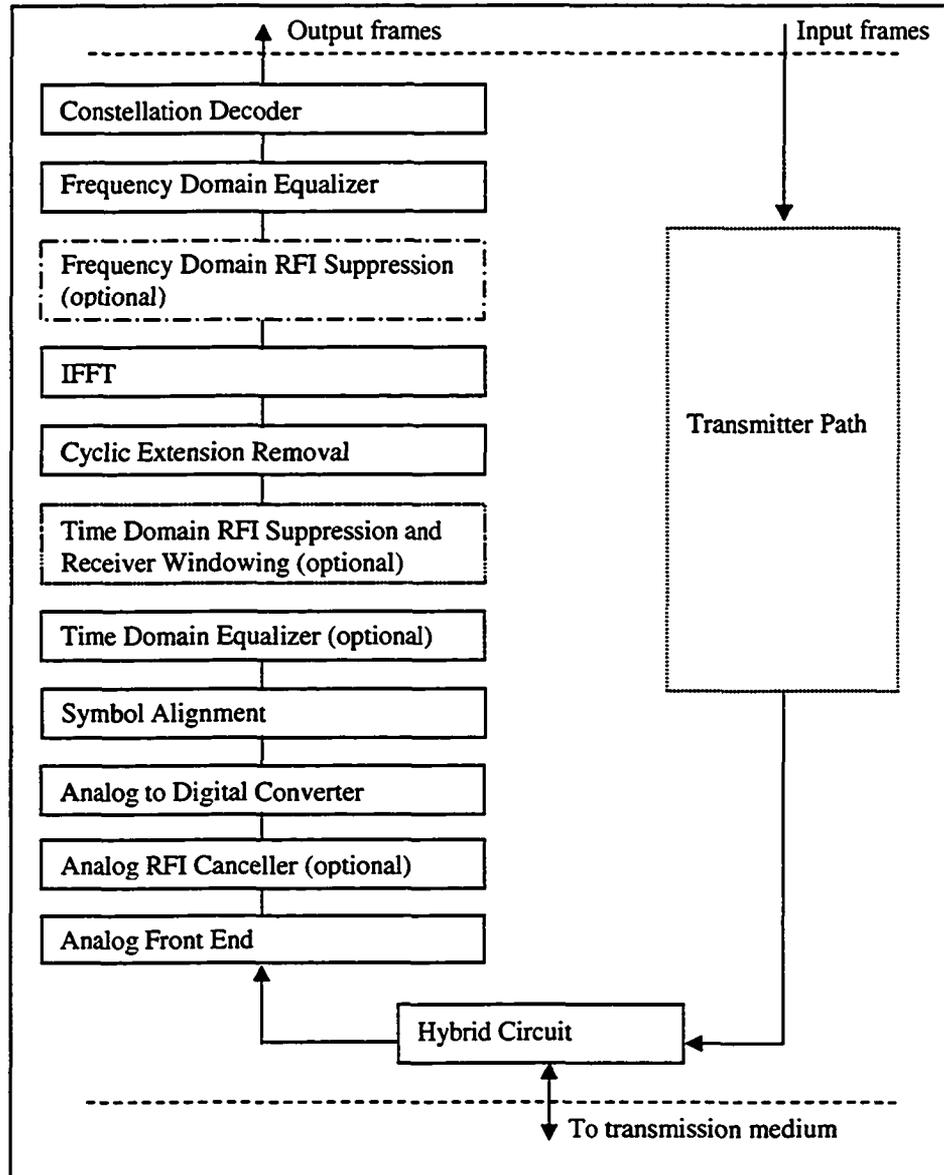


Figure 5-3: Receiver Model

The following sections provide details of the receiver as well as other related engineering insights. The various RFI suppression techniques, including receiving windowing, are described in detail in Chapter 6.

5.4.1 Analog RFI Cancellation

An Analog RFI canceller is primarily used to ensure that the received RFI corrupted signal is compatible with the Analog Front End (AFE) of the VDSL transceiver, and more specifically, the analog-to-digital converter (ADC). Even when such a device is implemented, a digital RFI suppression technique remains desirable for dealing with residual RFI at the output of the analog RFI canceller (see Ödling [19]).

According to de Clercq [5], the level of ingress RFI noise at the receiver can be significantly higher than the received VDSL signal. The concern is that it may saturate an ADC dimensioned for the received VDSL signal, or conversely, impose the use of a more elaborate and expensive ADC to handle the increased dynamic range of the signal at the receiver. However, although the RFI noise level may well be greater at the receiver than the far-end transmitted VDSL signal, other noise contributors may have levels just as great. Indeed, considering an echo return loss (ERL) of 15 dB and a signal transmitted by the VDSL transceiver at its maximum allowed power level of 14.5 dBm, the echo seen by the receiver will have a power level of -0.5 dBm. Assuming that the ADC is designed to have a dynamic range capable of handling the echo, it should be able therefore to handle ingress RFI without requiring analog RFI cancellation. Furthermore, should the RFI power ever be greater than the combined echo and received VDSL signal power, the input signal to the ADC would no longer have the Gaussian distribution associated with a DMT signal (see Starr et al. [26]). The Peak to Average Ratio (PAR) would then tend to decrease, from a ratio of 15 dB for the case of a Gaussian distributed signal with a clip probability of about 10^{-8} , towards the PAR of the RFI signal which is typically around the 3 dB PAR of a pure sine wave. This allows for the design requirements of the AFE to be relaxed.

Considering the likely RFI power levels to be seen at the VTU-O as well as the other design requirements for the receiver, it is assumed that RFI does not cause the analog-front end to saturate. Analog RFI cancellation is therefore not deemed necessary and the RFI suppression strategy can be implemented entirely in the digital domain.

5.4.2 Receive Symbol Alignment

Receiver symbol alignment delineates the symbol segment of the received signal to be demodulated by the FFT. It has a significant impact on system performance as it essentially determines the portion of the total cyclic extension to be used for containing TD-ISI, with the remainder being allocated for containing symbol transitions related to near-end echo as well as NEXT when synchronous mode is implemented.

Assuming that transmitters at both ends of the loop are indeed synchronised, receiver symbol alignment needs to be delayed with respect to the transmitter symbol timing to compensate for the channel dispersion time and propagation delay, while taking into account transmit pulse shaping (see Figure 2-4). There are loops however whose channel impulse response (CIR) exhibit a combined channel dispersion time and propagation delay greater than the sum of the maximum allowed cyclic prefix and suffix (see Van Acker [28] and section 5.5). For these cases, a compromise must be made between containing TD-ISI and echo/NEXT symbol transitions, which are both sources of frequency domain interference.

The symbol alignment algorithm used by the modelled receiver, performs a sliding biased window summation of length $2*N_{sc}$ over the squared channel impulse response. The interval offset corresponding to the largest sum is identified as the channel propagation delay estimate, and is bounded in theory by the channel's minimum and maximum group delays over the transmission frequencies used. The sliding window is biased from 1 to 0.9 over $2*N_{sc}$ points to give more weight to earlier samples. This ensures that the selected summation used for estimating the channel propagation delay corresponds to the window having the channel impulse response energy located at the beginning of its

interval, consistent with the definition of propagation delay. As an example, Figure 5-4 shows the squared impulse response of a modelled 5000-foot loop, normalised with respect to its peak magnitude, as well as the biased window positioned at the location that maximises the alignment criterion.

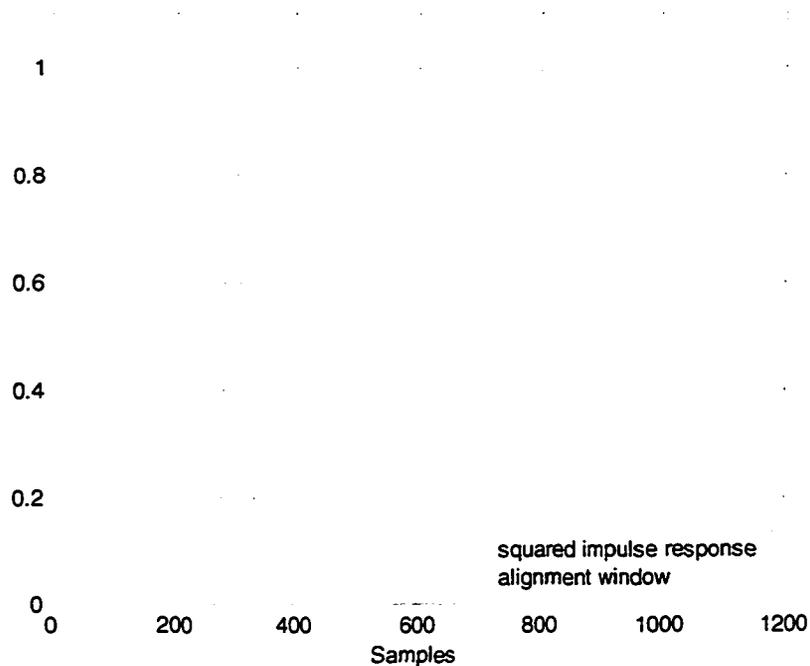


Figure 5-4: Alignment of 5000-Foot Loop Impulse Response

Setting the extended symbol alignment (start of received cyclic prefix) to the propagation delay essentially allocates the cyclic prefix for TD-ISI containment and the cyclic suffix for near-end echo containment as illustrated in Figure 5-5. This is the nominal setting for the simulation environment. In doing so, it is implicitly assumed that the channel propagation delay is similar in value to the channel dispersion time since the cyclic prefix and cyclic suffix lengths are defined to be equal at the transmitter.

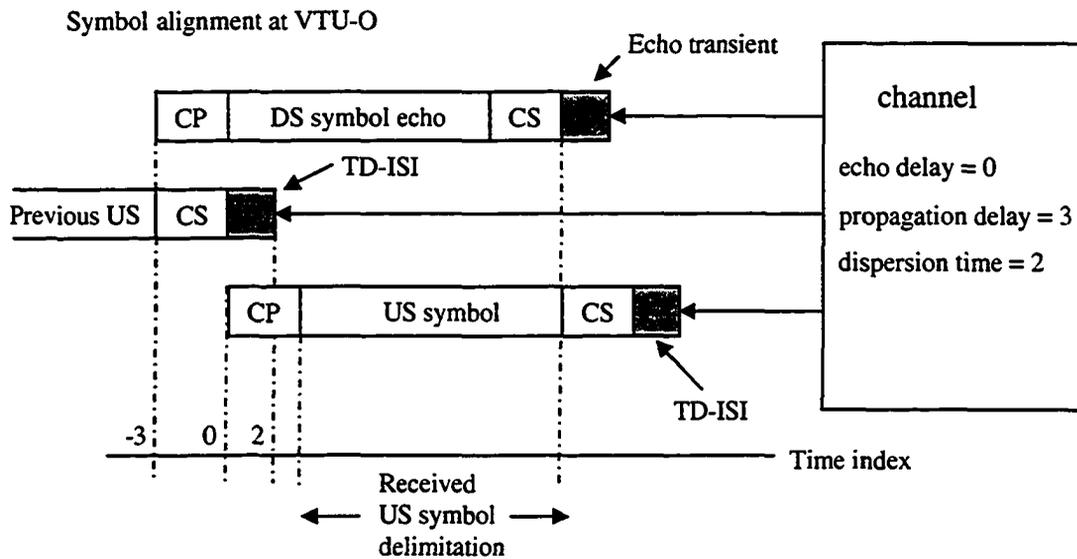


Figure 5-5: Nominal Receiver Alignment

Adding a timing advance to the nominal alignment serves to increase echo containment at the expense of TD-ISI containment, while the opposite is true when adding a timing delay. Such a timing adjustment allows for a trade-off to be made between using the cyclic extensions for containing TD-ISI versus containing near-end echo symbol transitions, regardless of the specific cyclic prefix and cyclic suffix lengths used by the transmitter. Note that a change in receiver alignment introduces the equivalent of a channel linear phase distortion, which is easily compensated for by the FDEQ. The above timing adjustment process is controlled in the simulation environment via an alignment offset variable. This allows for investigating the performance consequences of not having TD-ISI and/or echo symbol transitions properly contained.

5.4.3 Time-Domain Equalization

Recall from section 2.2.3 that frequency interference due to TD-ISI is avoided if the TD-ISI is confined to the cyclic prefix, as these corrupted samples are disposed of prior to demodulation.

A time domain equalizer (TDEQ) is often used in DMT receivers to address excessive TD-ISI that cannot be contained in the cyclic prefix. This would occur in a VDSL system when the channel dispersion time is greater than time spanned by the allocated cyclic prefix, which is meant to contain TD-ISI. The purpose of the equalizer is to reduce the dispersion time of the overall channel such that the TD-ISI generated by the equalised channel can be confined to the cyclic prefix. This approach is widely used in ADSL systems.

Time domain equalization has two major drawbacks. First, the processing power required to implement an effective TDEQ is significant and increases with the sampling rate. Since the VDSL receiver operates on up to 128 times more carriers than the ADSL receiver (i.e. upstream receiver), its TDEQ sampling rate must increase accordingly, which can become very expensive to implement. Secondly, the optimisation of the TDEQ is conventionally based on a simple time domain criterion (e.g. minimum mean square error with a specified target impulse response [1]) for which a solution is readily available. Al-Dhahir and Cioffi [2] demonstrate however that such a TDEQ does not typically maximise the achievable bit rate in a DMT system.

Since VDSL was conceived for use on relatively short loops with little receiving filtering, the corresponding channel dispersion times are shorter than those seen in ADSL deployments (see section 5.5 for impulse responses of targeted loops). It then becomes feasible for a VDSL receiver to contain the TD-ISI within the cyclic prefix with no further processing required. This suggests a different approach in designing a VDSL receiver compared to the ADSL receiver design. Instead of performing equalization to reduce the dispersion of the channel impulse response, a VDSL transceiver should simply aim to avoid contributing to the channel impulse response. Thus, time domain transmit and receive filtering, including the TDEQ, should be avoided as much as possible.

5.5 Analysis of Loop Models in the Context of System Performance

This section analyses the characteristics of the loop models used for simulation, which are defined in Table 5-1. The attributes of the echo paths corresponding to these loops are also investigated. A good understanding of these characteristics is essential in order to determine the cyclic extensions requirements, evaluate the impact of loop impairments on system performance, and make correct interpretation of simulation results.

5.5.1 Impulse Response of Loop Models

The impulse response of a loop determines the dispersion time and the propagation delay of the loop, which in turn characterize the loop's contributions to TD-ISI and non-orthogonal echo. Figure 5-6 and Figure 5-7 depict the impulse responses of loop models VDSL1-3000ft and VDSL5 as examples.

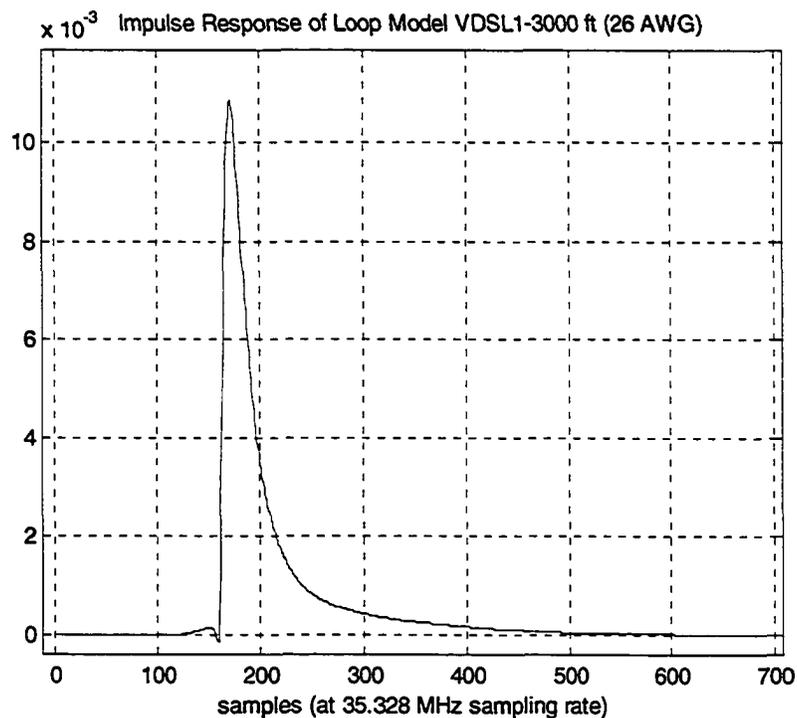


Figure 5-6: Impulse Response of Loop Model VDSL1-3000ft (26 AWG)

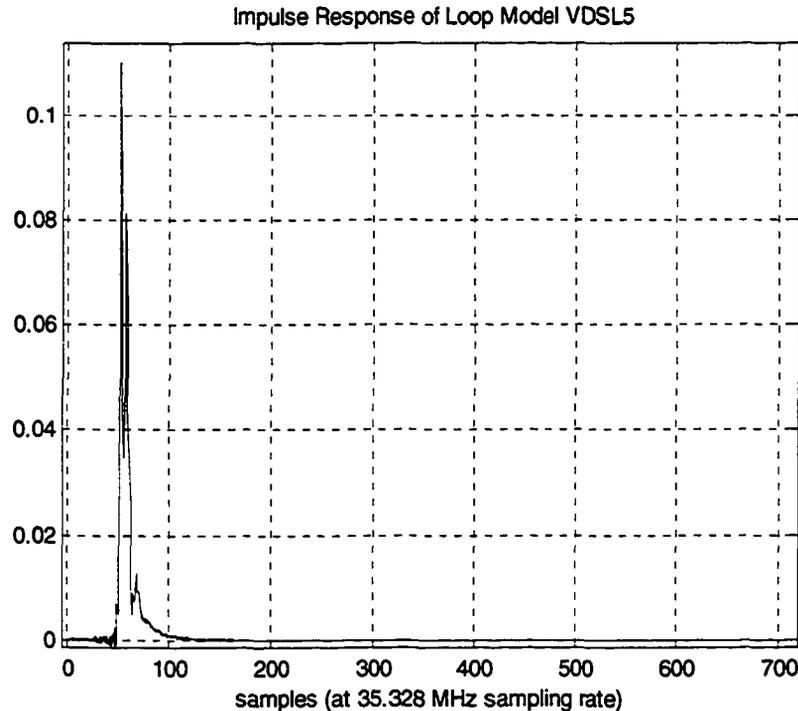


Figure 5-7: Impulse Response of Loop Model VDSL5

The propagation delay is often associated with the length, measured in samples, from time zero to the response's peak. The energy alignment measurement described in section 5.4.2 however is a more accurate interpretation of propagation delay. It essentially identifies the start of a region containing the maximum impulse energy as the propagation delay. This definition is consistent with the receiver alignment algorithm proposed in this thesis (see section 5.4.2) and will be used herein after.

The spreading of the impulse response energy, or dispersion time, gives an indication of the duration and severity of the TD-ISI to expect. The dispersion time, measured in samples, will refer herein after to the window length required to contain 99.5 % of the impulse energy, assuming the start of the window is offset by an amount equivalent to the propagation delay. It is to be interpreted as the duration of the impulse response minus the propagation delay. The specified energy threshold was to a large extent chosen arbitrarily. Nevertheless, the dispersion time measurements calculated with this threshold correlate well with the simulation results presented in Chapter 7.

5.5.2 Transfer Function Magnitude of Loop Models

The transfer function magnitude curves of various length VDSL1-x loop models are presented in Figure 5-8. These curves along with the noise energy at the receiver provide insight into the channel capacity of the corresponding loops. Recall from section 2.2.6 that the capacity of a bin is basically a log function of the bin's SNR. A minimum 14 dB SNR is required to allocate 1 bit to a bin, while the maximum 15 bit allocation requires 57.5 dB of SNR. Assuming a transmit PSD of -60 dBm/Hz and a noise PSD at the receiver of -140 dBm/Hz, a bin can have 22.5 dB of channel attenuation and still carry the maximum bit allotment. On the other hand, a bin that is subject to channel attenuation greater than 66 dB will not be allocated any bits at all. As an example under this scenario, a 3000 foot 26 AWG loop modelled as VDSL1-3000ft could only carry data on bins below approximately 1500 (out of 4096).

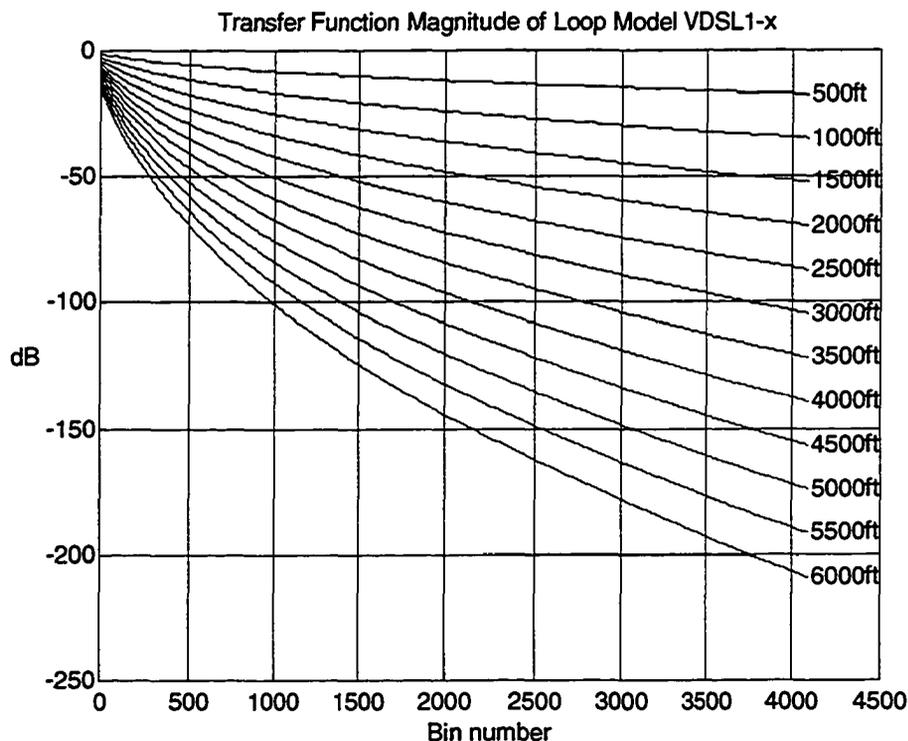


Figure 5-8: Transfer Functions of Loop Models VDSL1-x of Various Length

The transfer function magnitude curves of VDSL loops 4 through 7 are presented in Figure 5-9.

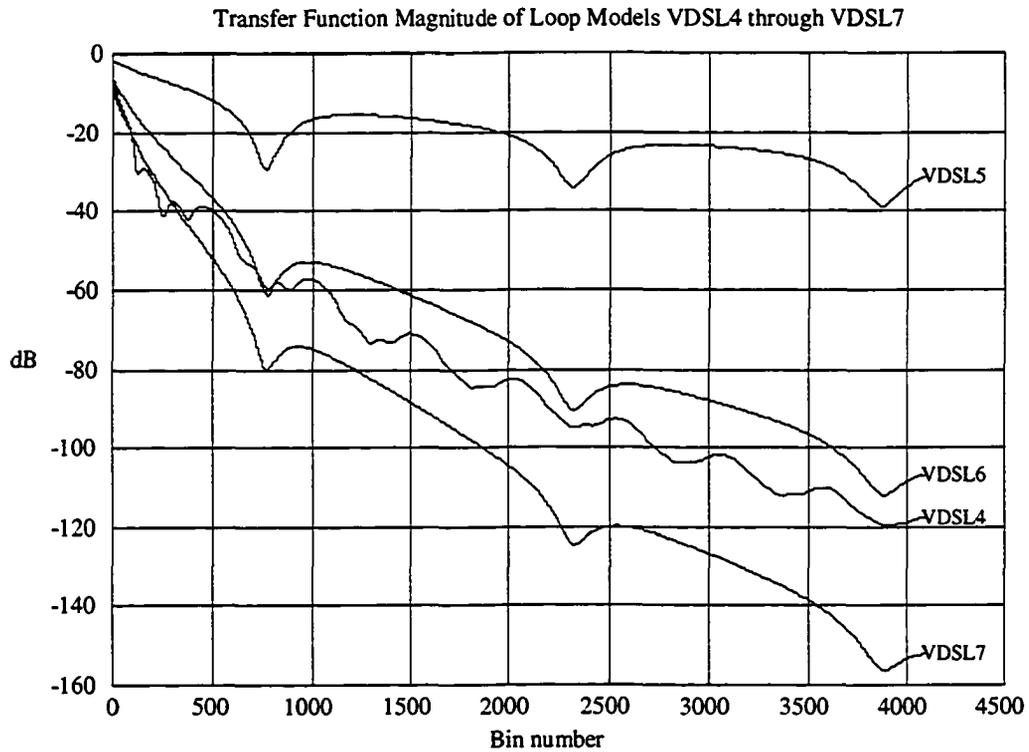


Figure 5-9: Transfer Functions of Loop Models VDSL4 through VDSL7

5.5.3 Input Impedance of Loop Models

Recall that the modelled hybrid has its reference input loop impedance set equal to the 100 Ohm VDSL termination impedance. Consequently, an actual input loop impedance different from this value will lead to near-end echo as per equation (2-18). As an example, Figure 5-10 depicts the input impedance magnitude of the loop model VDSL1-3000ft over the first 500 bins. Note that the impedance mismatch, compared to 100 Ohm, is significant at the lower frequencies. It is thus expected that the echo will be most severe in the lower bins. The near-end echo paths for the various loops are investigated in section 5.5.6.

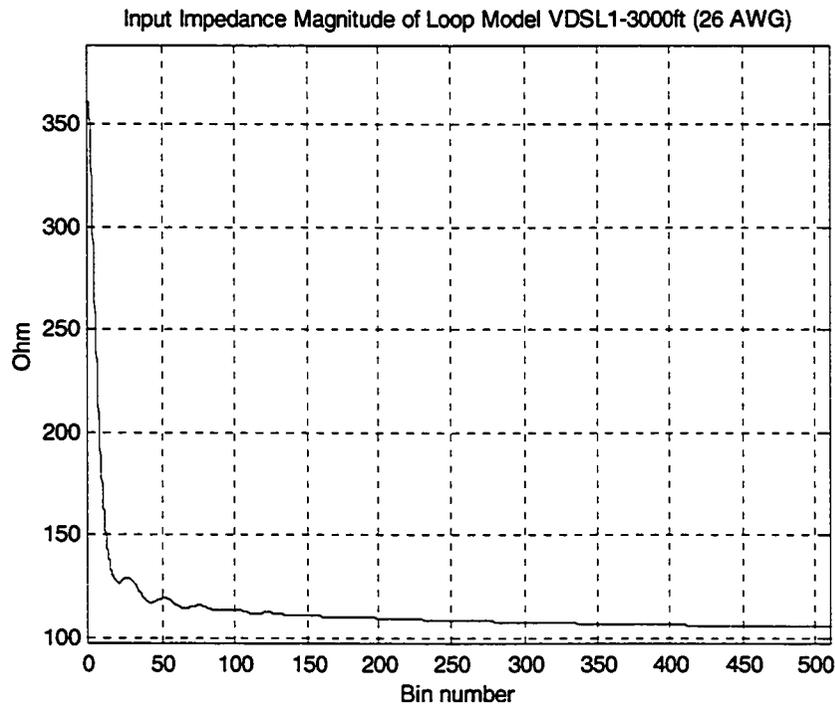


Figure 5-10: Input Impedance Magnitude of Loop Model VDSL1-3000ft

5.5.4 Propagation Delay of Loop Model VDSL1-x

Figure 5-11 depicts the propagation delay introduced by the loop model VDSL1-x as a function of loop length. The propagation delay is calculated according to the energy alignment procedure of section 5.4.2.

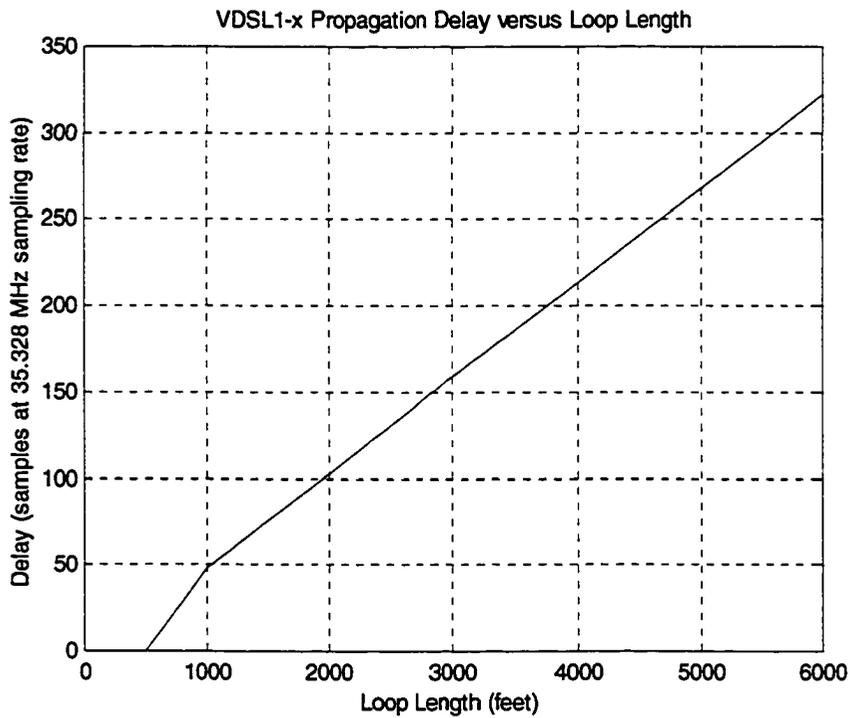


Figure 5-11: VDSL1-x Propagation Delay versus Loop Length

5.5.5 Dispersion Time of Loop Model VDSL1-x

Figure 5-12 depicts the dispersion time of the loop model VDSL1-x as a function of loop length. The dispersion time is calculated according to the definition of section 5.5.1.

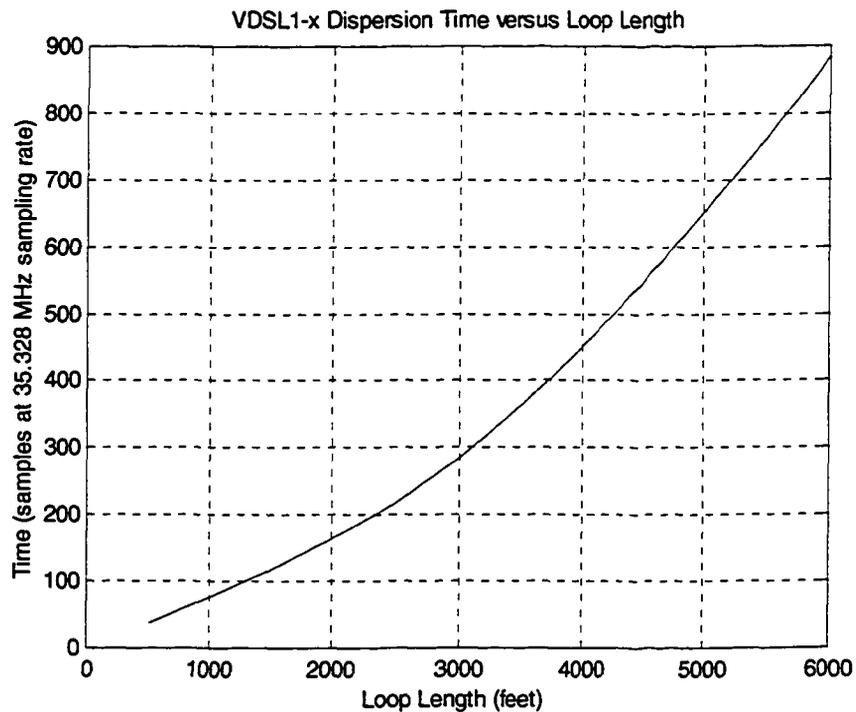


Figure 5-12: VDSL1-x Dispersion Time versus Loop Length

5.5.6 Echo Path

The echo path impulse response determines the severity of the near-end echo in terms of amplitude, delay (propagation delay) and energy spread (dispersion time). The more the energy of an echo symbol transition is delayed and spread in time, the less stringent are the requirements on the receiver alignment and the combined cyclic prefix/suffix length. To see this, recall from section 5.4.2 that confining the transient energy of an echo symbol transition to the cyclic suffix of the received symbol requires accurate symbol alignment. This assumes that the cyclic prefix/suffix combination is of sufficient length to both contain the TD-ISI in the cyclic prefix and the echo transient in the cyclic suffix. Now, an echo path impulse response exhibiting a significant propagation delay or dispersion time will push out the echo transient energy towards the boundary of the received extended symbol, reducing the impact of a symbol misalignment or an insufficient cyclic suffix length.

Figure 5-13 and Figure 5-14 show two examples of echo path impulse responses where the energy spread can be clearly observed. The transfer functions corresponding to these impulse responses are shown in Figure 5-15 and Figure 5-16. They confirm that lower bins will indeed be more prone to echo than higher ones.

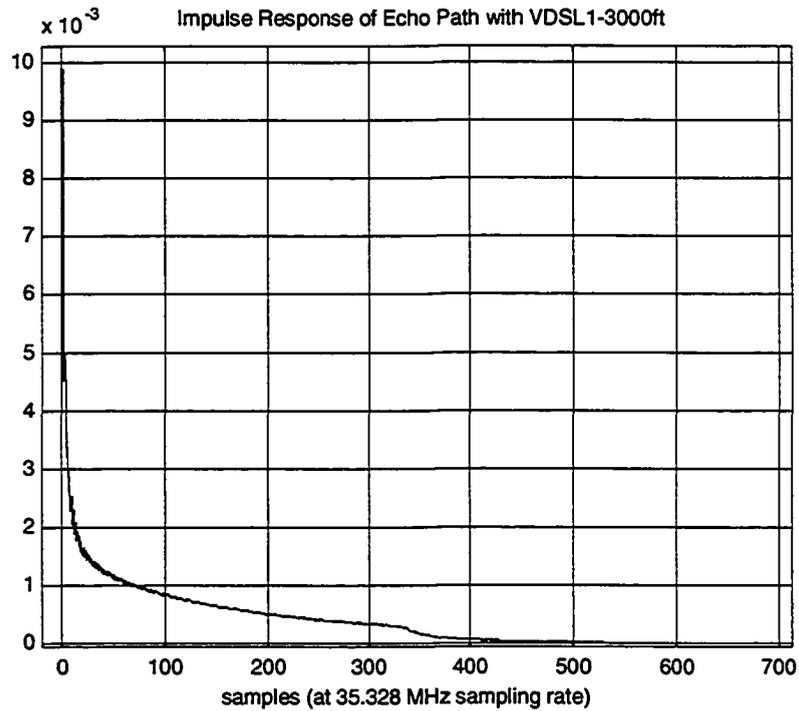


Figure 5-13: Echo Path Impulse Response with VDSL1-3000ft

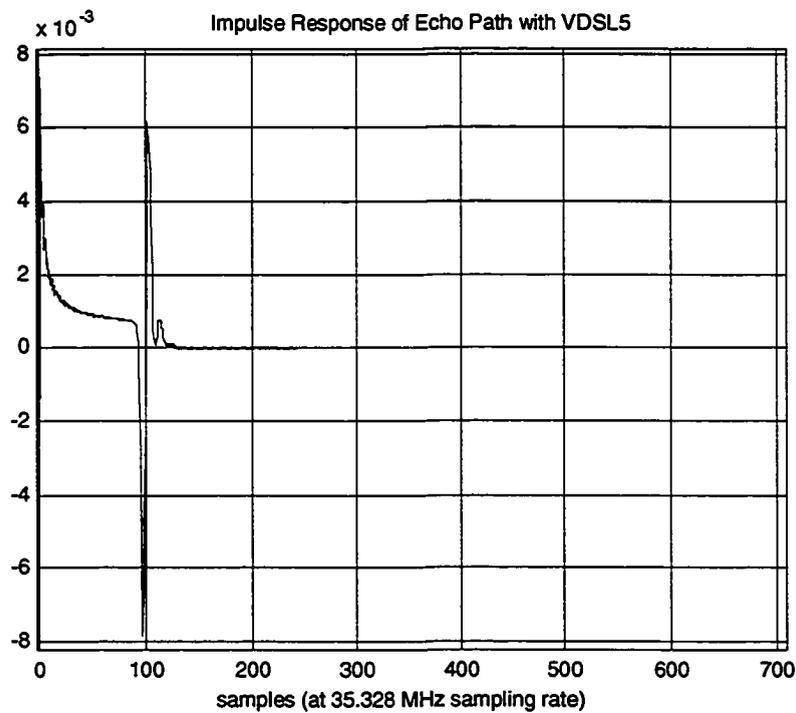


Figure 5-14: Echo Path Impulse Response with VDSL5

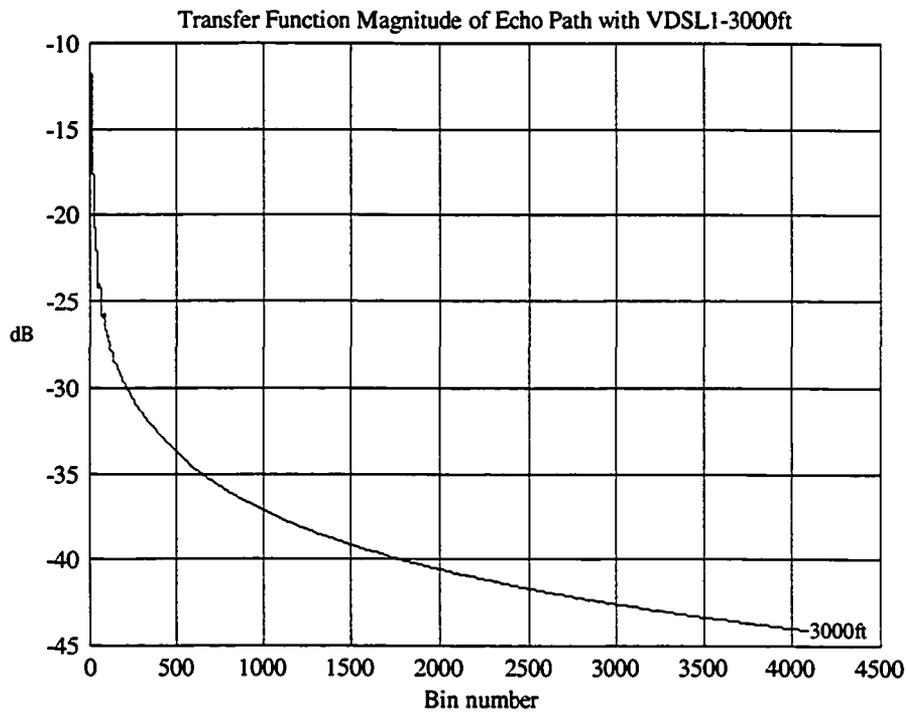


Figure 5-15: Transfer Function of Echo Path with VDSL1-3000ft

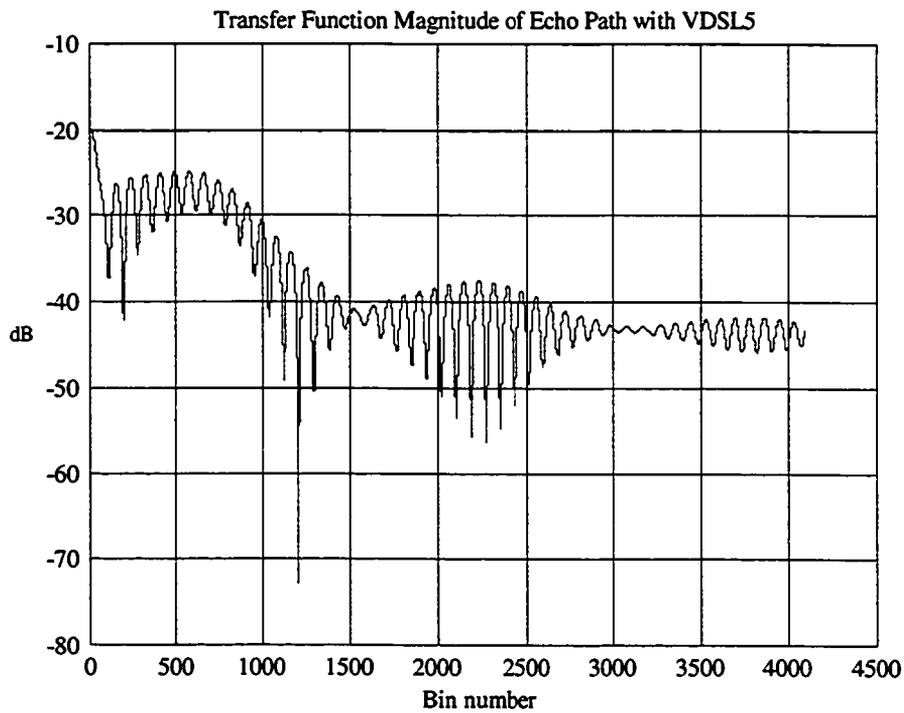


Figure 5-16: Transfer Function of Echo Path with VDSL5

Figure 5-17 displays the duration of the echo path impulse response as a function of loop length with the loop model VDSL1-x. The impulse response duration is calculated according to the dispersion time definition given in section 5.5.1, assuming a propagation delay of zero.

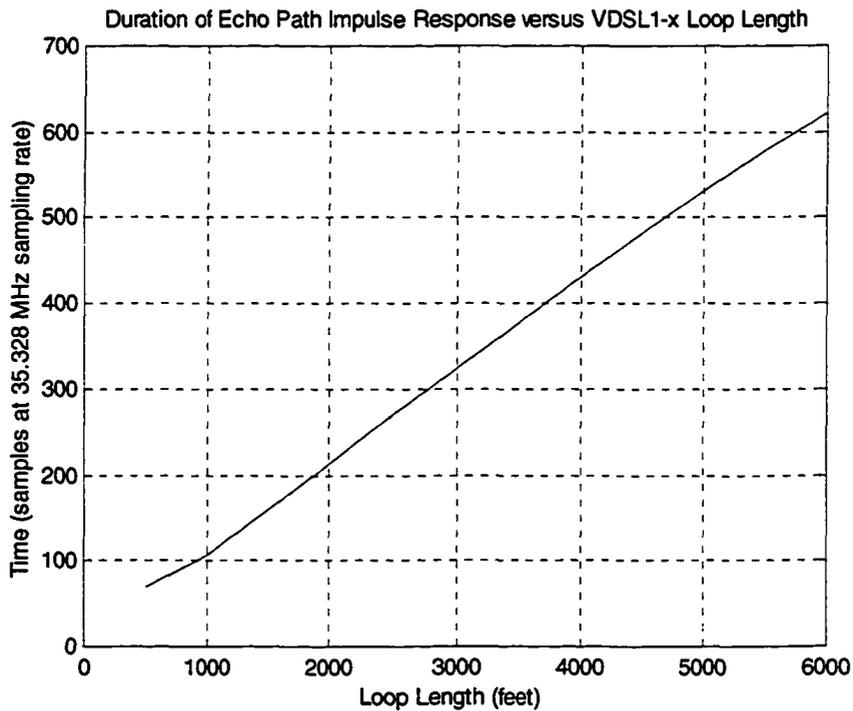


Figure 5-17: Duration of Echo Path Impulse Response versus Loop Length (VDSL1-x)

5.5.7 Cyclic Extensions Requirements

The following tables present a summary of the loop model characteristics that impact on cyclic extensions requirements, as well as the extension margins assuming the maximum permitted cyclic prefix/suffix combination length of 896 samples. The characteristics are measured in samples at a sampling frequency of 35.328 MHz.

Table 5-2: Time Domain Characteristics of Various Length VDSL1 Loops

Loop Characteristics Measured in Samples	500 ft	1000 ft	1500 ft	2000 ft	2500 ft	3000 ft	3500 ft	4000 ft	4500 ft	5000 ft	5500 ft	6000 ft
Loop Propagation Delay	0	48	76	103	130	159	186	213	241	268	295	322
Loop Dispersion Time	38	76	117	163	217	285	364	452	547	651	762	883
Echo Path Impulse Response Duration	69	107	160	214	269	324	378	431	483	532	578	622
Loop Delay + Dispersion	38	124	193	266	347	444	550	665	788	919	1057	1205
Cyclic Prefix/Suffix Combined Margins	858	772	703	630	549	452	346	231	108	-23	-161	-309

Table 5-3: Time Domain Characteristics of VDSL4 through VDSL7 Loops

Loop Characteristics Measured in Samples	VDSL4 (3300 ft)	VDSL5 (950 ft)	VDSL6 (3250 ft)	VDSL7 (4900 ft)
Loop Propagation Delay	175	39	171	245
Loop Dispersion Time	469	72	293	542
Echo Path Impulse Response Duration	428	107	321	479
Loop Delay + Dispersion	644	111	464	787
Cyclic Prefix/Suffix Combined Margins	252	785	432	109

As can be observed from the above tables, many loop models exhibit dispersion times and propagation delays that cannot be properly dealt with by the permitted cyclic extensions. Specifically, the expected TD-ISI and echo transients occurring with

VDSL1-x loop models of lengths greater than 5000 feet cannot be entirely contained in the cyclic extensions. Containment is also marginal with the loop model VDSL1-4500ft as well as with the loop model VDSL7. On the other hand, these same loop models have significant dispersion times as well as significant echo path impulse response durations, alleviating the shortcomings of the cyclic extensions. For example, the energy contained in the tail end of the impulse response of a loop with a large dispersion time should be relatively little and will consequently contribute little TD-ISI. The same principle applies for the echo path impulse response as explained in section 5.5.6. The relations between the loop model characteristics and system performance are explored through simulation, with results presented in section 7.1.

Chapter 6

Digital RFI Suppression Techniques

The various digital RFI suppression techniques identified in the literature review are studied in this chapter. The following aspects are given special attention where warranted:

- expected performance, with and without RFI present
- computation requirements
- ability to handle more than one RFI source
- identification of estimation, adaptation, and modelling requirements
- other drawbacks and benefits

It is convenient to group the algorithms under consideration into one of three classes of RFI suppression techniques: receiver windowing, adaptive notch filtering, and digital RFI cancellation based on modelling.

6.1 Receiver Windowing

The amount of RFI power leaked into neighbouring bins is directly related to the transfer function of the time domain window used by the demodulator, and more specifically, to the magnitude and location of the sidelobes. Traditionally, DMT demodulators have used a rectangular window, resulting in sidelobe levels characterised by $\sin(x)/x$. The transfer function magnitude of this receiver rectangular window (boxcar) is shown in Figure 6-2 and Figure 6-3. By performing non-rectangular windowing on the time-domain received data samples prior to the demodulator, the sidelobe levels can be substantially reduced,

although the main lobe bandwidth actually increases (see Manolakis and Proakis [21]). This operation can be considered the most natural approach to combating RFI as it addresses the very source of the problem: noise corruption of active bins due to the frequency sidelobe leakage inherent to the demodulator's rectangular window. The intent is thus to minimise the sidelobe power leakage that is introduced by the demodulation process. This approach has two major advantages over other RFI suppression techniques. First no interferer-specific information is required to suppress the RFI. In fact, windowing may actually increase VDSL immunity to other types of impairments aside from ingress RFI, such as non-orthogonal echo and other out-of-band noise sources. It may even help confine to a few bins the interference due to an in-band narrowband noise source. Second, the windowing operation remains constant whether one or more noise sources are present. Complexity and robustness of the algorithm are thus fixed and well defined. It should also be noted that the model-based RFI cancellation algorithm presented later on in section 6.3 requires non-rectangular receiver windowing to achieve the RFI suppression levels reported in [29]. Analysis of various types of receiver windowing operation and implementation is presented in the following sections.

Another impairment that is particularly well suited for suppression by receiver windowing is asynchronous self-NEXT. Recall that this impairment is caused by other near-end VDSL transceivers sharing the same binder group but starting transmission of symbols at different times (asynchronous mode). The main benefit of the asynchronous scenario is that the cyclic extension size as well as timing synchronisation between the VTU-O and VTU-R is determined on a line-by-line basis to maximise transmission efficiency. The downside is that asynchronous self-NEXT is typically not orthogonal to the received symbols, even though both signals operate in different frequency bands. The reason for this is that the symbol transitions of asynchronous self-NEXT are not controlled and therefore cannot be contained within the cyclic extensions of the received symbols. This will translate into some level of frequency interference due to self-NEXT energy being introduced in the received data. Applying non-rectangular receiver windowing to the received signal however diminishes this interference significantly. This has for effect of smoothing out the edges of the frames, thus reducing the spectral

sidelobes leakage of NEXT and the ensuing interference. Simulation results highlighting the above are presented in [10].

6.1.1 Raised Cosine Windowing Extending into Cyclic Extension

An important objective of VDSL receiver processing is to maintain the ideal characteristics of the subchannels as to avoid the need for a costly equalizer. As a consequence, receiver processing should keep the introduction of intercarrier interference (ICI) to a minimum. To avoid having the receiver time windowing operation generate ICI, the window's transfer function must have zero crossings at the bin frequencies. For a non-rectangular window however to have this property, it must extend beyond the orthogonality interval into the cyclic extensions (see [5] and [29]). This is indeed the case for the proposed family of raised cosine windows displayed in Figure 6-1. These windows extended into the cyclic extensions such that half of the pulse shaping occurs there. The demodulation interval must also be extended accordingly.

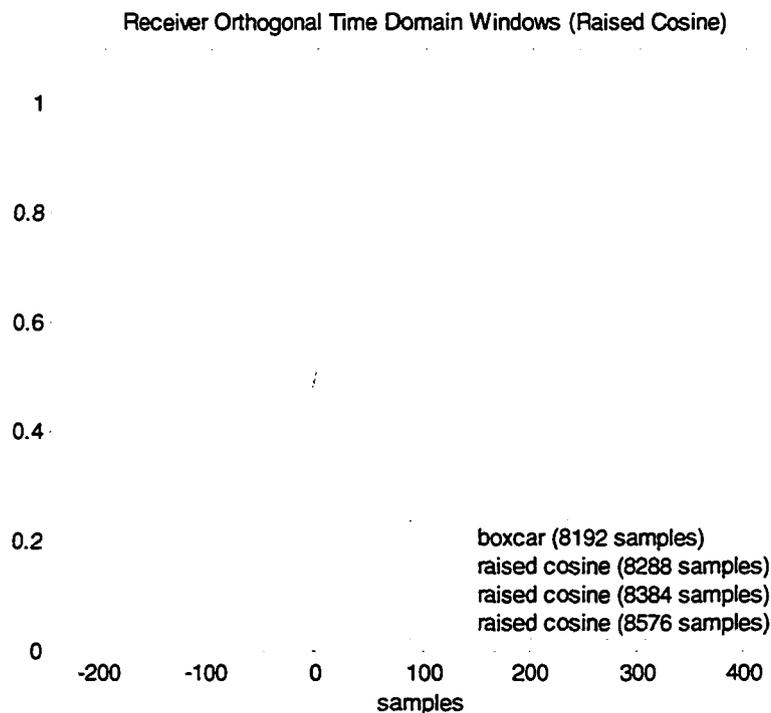


Figure 6-1: Raised Cosine Receiver Time-Domain Windows

The transfer functions of raised cosine windows are presented in Figure 6-2 and Figure 6-3. Note that they do indeed exhibit the appropriate nulls at the bin frequencies.

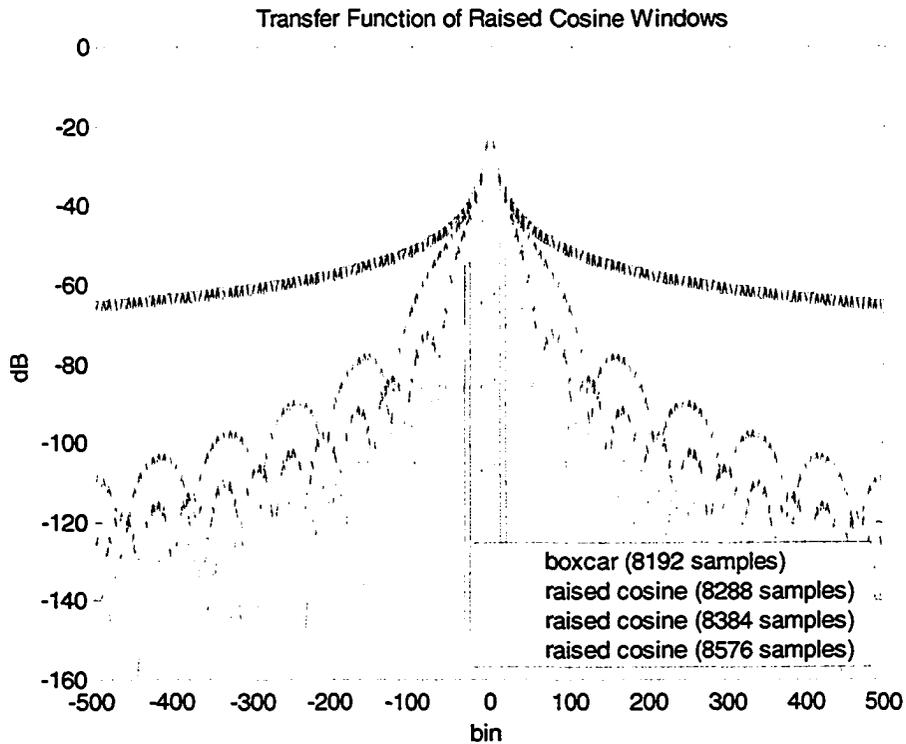


Figure 6-2: Transfer Function Magnitude of Raised Cosine Windows

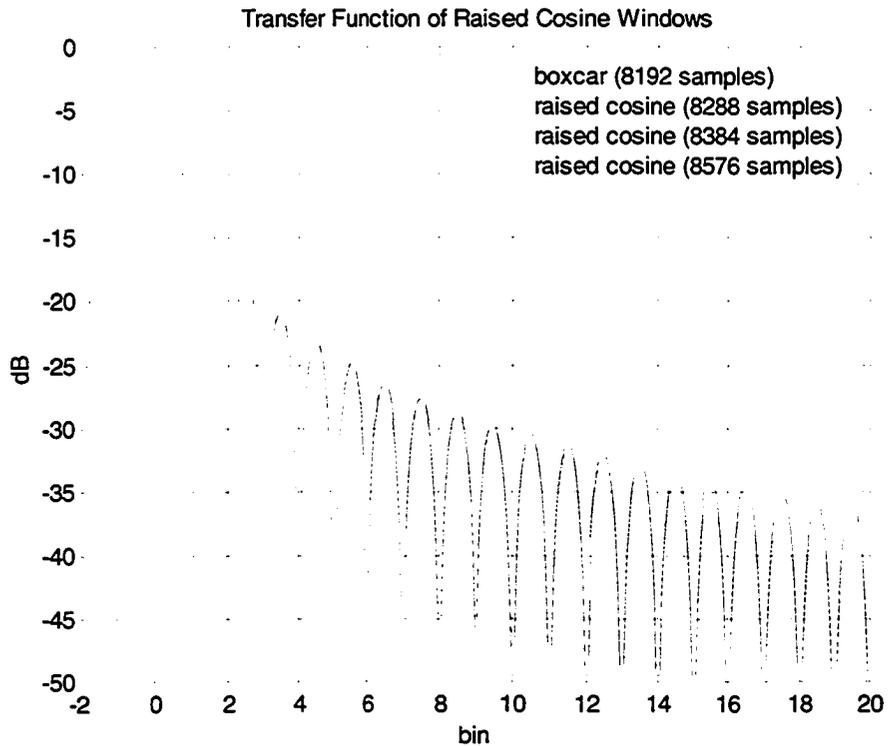


Figure 6-3: Bins 0 through 20 of Raised Cosine Window Transfer Functions

Table 6-1 lists the magnitudes of the first few dominant sidelobes of various receiver windows, including those studied in later sections. The sidelobes that are within 10 bins from the main lobe are of no concern in the context of RFI, as noise leakage attributed to them will, in the worst case, fall within the 10 RFI guard bins. The enhanced sidelobe attenuation provided by raised cosine windowing is rather little in the immediate vicinity of the main lobe, but becomes substantial at further distances. For example, the 8576 sample raised cosine window provides approximately 8 dB of extra attenuation at a distance of 20 bins, and 40 dB of extra attenuation at a distance of 100 bins. Keep in mind that every extra 3 dB of attenuation allows potentially for an extra bit to be carried on a given bin. Ideally, the maximum level of sidelobes would be -130 dB with respect to the main lobe to ensure that the effect of a worst case -10 dBm RFI source be buried in the noise floor, which is assumed here to be -140 dBm/Hz.

Table 6-1: Sidelobe Magnitude of Receiver Window Transfer Functions

Window	Sidelobe Magnitude with respect to Peak (dB)			
	9.5 bin away	10.5 bin away	11.5 bin away	19.5 bin away
Rectangular (8192 samples)	-29.5 dB	-30.4 dB	-31.2 dB	-35.7 dB
Raised Cosine 8288 samples	-29.6 dB	-30.5 dB	-31.3 dB	-36.2 dB
Raised Cosine 8384 samples	-29.9 dB	-30.9 dB	-31.8 dB	-37.5 dB
Raised Cosine 8576 samples	-31.2 dB	-32.4 dB	-33.6 dB	-43.5 dB
Asymmetric 1, -0.5	-36.6 dB	-37.5 dB	-38.0 dB	-42.2 dB
Asymmetric 1, -1	-48.1 dB	-49.9 dB	-51.6 dB	-61.1 dB
Hanning	-59.6 dB	-62.5 dB	-65.1 dB	-79.9 dB

The increased demodulation time interval associated with raised cosine windowing potentially reduces the system's ability to handle TD-ISI and near-end echo. The reason for this is that part of the cyclic extensions is now demodulated for data recovery instead of acting as a guard time against TD-ISI and echo transients. Any time domain ISI or echo symbol transition present in the windowed portion of the cyclic extensions consequently introduces interference in the frequency domain data. On the other hand, the raised cosine windowing operation should generally improve the demodulator's immunity to out-of-band noise, including near-end echo transients and NEXT, since it reduces frequency sidelobe leakage.

The raised cosine windowing and demodulation operations should not extend into the shaped portions of the transmitted signal, as this would introduce ICI due to the samples being shaped as well as ISI due to symbol overlap in this region. With a cyclic extension of 448 samples and a symbol overlap of 256 samples, the windowing incursion into the cyclic extensions should be limited to within 192 samples. This corresponds to a maximum raised cosine window width of 8576 samples.

Figure 6-4 shows components of the receive signal PSD before and after raised cosine windowing for the case of a 2000 foot loop and one RFI located at bin 870.5. For clarity, the near-end echo component at the receiver is not shown. Note the RF noise contribution on each tone, before and after windowing, relative to the upstream signal.

Receive SNR, with and without RFI present, is plotted in Figure 6-5, emphasising the effectiveness of raised cosine windowing at combating RFI.

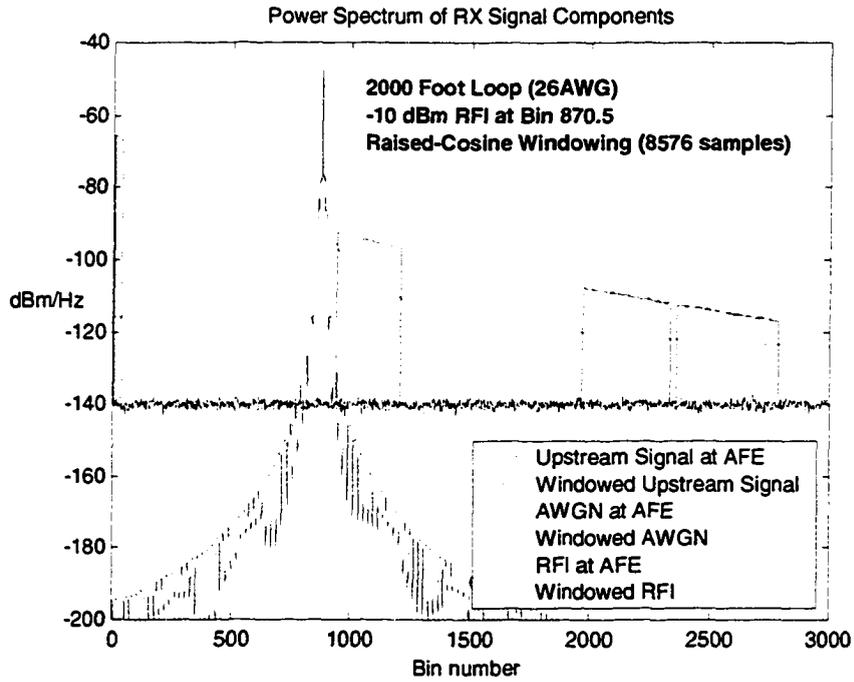


Figure 6-4: Receive Signal PSD with Raised Cosine Windowing

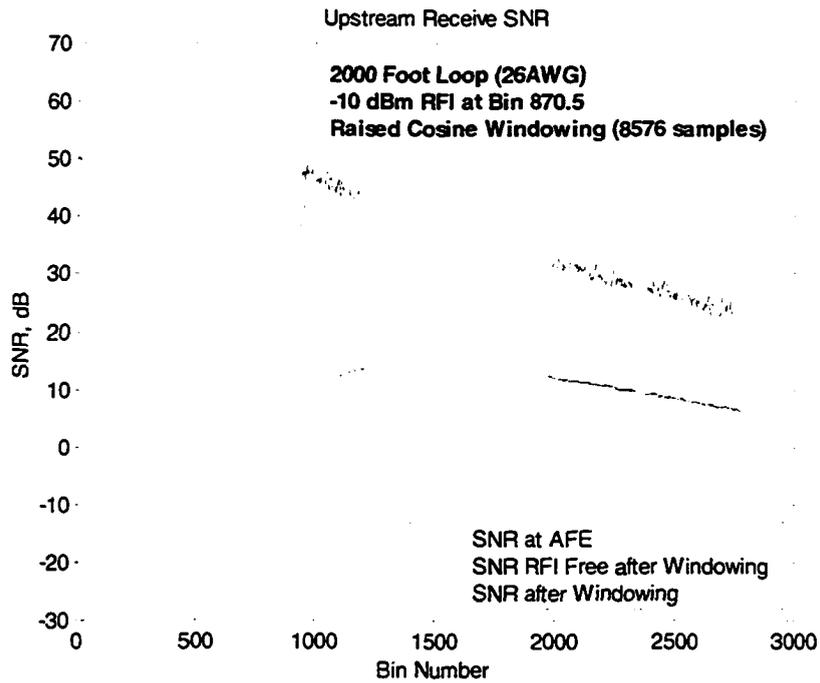


Figure 6-5: Receive SNR with and without Raised Cosine Windowing

6.1.1.1 Raised Cosine Windowing Computational Requirements

The computation requirement for executing the raised cosine windowing algorithm arises from two processes. First, one real multiplication is required for each sample to undergo shaping. If windowing is to be performed in real time, these multiplications must be executed at the sampling rate, even though only part of the symbol requires to be shaped. This translates into a computation requirement of 35.3×10^6 real multiplications per second, which is equivalent to $2 \cdot N_{SC} + L_{CE} = 8832$ real multiplications per symbol period. On the other hand, if system performance can accommodate an extra symbol delay, the windowing multiplications can be spread out across this symbol delay, resulting in a requirement of multiplications per symbol period equivalent to the total number of samples to be shaped in a symbol. For the 8576 samples raised cosine window, this number is 768.

The other aspect adding to the algorithm's computational complexity is the need to demodulate part of the cyclic extensions, which at first look requires an FFT size increase. However, [5] and [10] suggest folding the shaped portions of the cyclic extensions back into the symbol boundaries prior to demodulation, thus avoiding the need for a larger FFT. The folding operation is performed by taking the shaped portions of the cyclic prefix and suffix and adding them to the outermost and opposite ends of the elementary symbol (without the cyclic extensions), as illustrated in Figure 6-6. This is essentially the inverse of the procedure used for cyclically extending a symbol. It thus follows that the DMT signal contained in the $2N_{SC}$ samples resulting from the windowing and folding operations is the same as the DMT signal contained in the original $2N_{SC}$ samples due to the symmetry properties of the window and the cyclical periodicity of the DMT symbol.

What is less obvious is the frequency domain consequence of performing the folding operation on a general signal or noise impairment, such as RFI, which does not share the cyclical periodicity of the DMT symbol. It turns out that the DFT of the windowed and folded $2N_{SC}$ samples is simply a phased shifted version of the frequency samples obtained

by directly performing a double sized FFT on the windowed samples followed by a decimation of two. To see this, consider the Fourier transform $V(\omega)$ of the windowed symbol defined as $v(n)$ over an extended symbol interval and zero elsewhere. The sampled version of $V(\omega)$ evaluated at $\omega = 2\pi k/N$ corresponds to the DFT of a signal $v_p(n)$ obtained by the periodic repetition of $v(n)$ every N samples as shown in [21]. If N is equal or greater than the number of non-zero samples in $v(n)$, then $v(n) = v_p(n)$ for $0 \leq n \leq N - 1$. This corresponds to the case of the double size FFT where $N = 4N_{SC}$ which is greater than the number of windowed samples, including the shaped portions of the cyclic extensions. On the other hand, a nominal sized FFT corresponds to $N = 2N_{SC}$, which is less than the number of windowed samples. In this case, the signal $v_p(n)$ corresponding to the IDFT of $V(\omega)$ evaluated at $\omega = 2\pi k/2N_{SC}$ consists of periodic repetitions of $v(n)$ every $2N_{SC}$ samples which overlap with each other leading to time-domain aliasing. But this aliasing is exactly what is reproduced when performing the folding operation on the windowed symbol samples, besides from an added time shift. It therefore follows that the $2N_{SC}$ FFT of the windowed and folded symbol corresponds to $V(\pi k/2N_{SC})$ times a phase shift factor which can be compensated for in the FDEQ in order to recover the required signal.

The folding operation requires a number of real additions equivalent to the number of windowed samples extending beyond the nominal rectangular window.

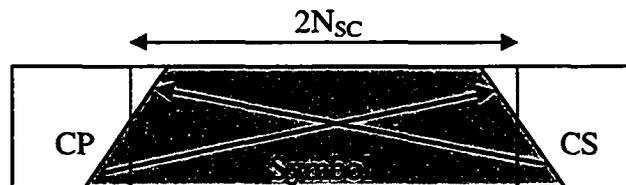


Figure 6-6: Windowing and Folding of Received Symbol

6.1.2 Frequency-Domain implementation of Time Windowing

An implementation of the frequency-domain operation corresponding to time-windowing is proposed by [28] in the context of combining receiver windowing and per-tone equalization. Even though equalization is not required, the frequency-domain version of time-windowing at first glance has the merit of not requiring an extended sized FFT demodulator to handle the extra shaped samples of the cyclic extension. However, a closer look at the algorithm development reveals that the windowing procedure is based on the same folding operation as presented in [28], and therefore does not present any advantage over the previously explored methods. This approach will therefore not be investigated any further.

6.1.3 Constrained Windowing and Controlled ICI.

The major drawbacks of raised cosine windowing are that it extends into the cyclic extensions and that the sidelobe attenuation it provides may in some cases be insufficient to properly suppress RFI interference. The constrained windowing algorithm on the other hand potentially offers greater sidelobe attenuation and operates only over the DMT symbol interval. The full cyclic extensions are rejected prior to demodulation, leaving them entirely for their original purposes.

As mentioned in section 3.2, constraining the receiver windowing operation to the DMT symbol interval induces ICI in the recovered data. It is possible however to specify constrained windows that produce a controlled amount of ICI, where any given tone is subject to window-induced ICI originating from only a few particular bins. In other words, the transfer functions of these windows have zero crossings aligned with all but a few of the bin frequencies. As the resulting ICI is limited and deterministic in nature, it is feasible to use a decision feedback equalizer (DFE) in the frequency domain to cancel-out the window-induced ICI as proposed by Kapoor and Nedic [14]. This technique is similar in principal to the partial response signalling (PRS) method outlined by Proakis [20], except that in this case, the controlled interference is introduced at the receiver

instead of at the transmitter. As a consequence, the DFE needs to take into account the channel's frequency response when generating the feedback signal. Furthermore, any noise present in the signal also contributes to window induced ICI, but since the resulting impairment is uncorrelated with the recovered data, it cannot be cancelled out by the DFE. This leads to a noise floor increase as documented later on in Table 6-2. This method therefore suffers SNR degradation in an additive noise environment.

The architecture offered by Kapoor and Nedic [14] is shown in Figure 6-7. It makes use of two receiver data paths, one conventional and the other implementing the windowing function accompanied by a frequency-domain DFE. The time windowing function is implemented in the frequency domain as a convolution operation to allow for a single FFT demodulator to be used by both paths. A bin selection logic block routes the output of one of the two paths to the slicer input on a bin by bin basis. Selection is determined according to a performance criterion evaluated based on channel conditions at initialisation, including noise sources. As a consequence, this algorithm will be of no avail should the presence or frequency location of the RFI change after initialisation, as is known to occur.

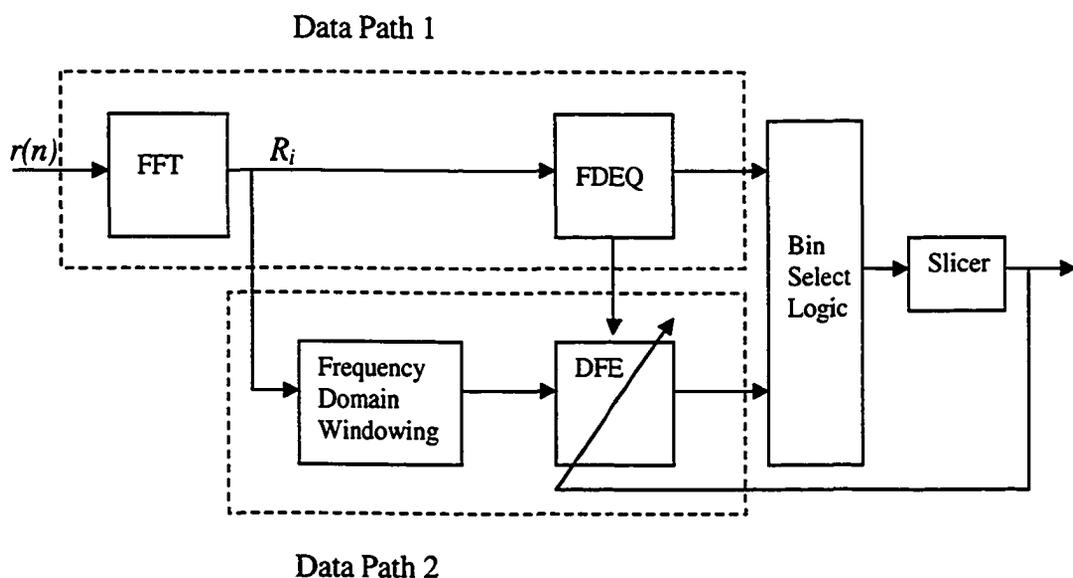


Figure 6-7 : Constrained Windowing Original Architecture

In order to handle intermittent and frequency varying RFI sources, the scope of the windowing algorithm can be extended to apply to all bins. This is a reasonable approach considering that an RFI source might well interfere with most active bins as illustrated in Figure 7-5 of section 7.1.3. As a consequence, Data Path 1 and the bin selection logic in Figure 6-7 are no longer required and can be eliminated. In this case, there are no longer any reasons for implementing the windowing function in the frequency domain, and a more efficient time-domain windowing operation can thus be performed. The resulting architecture is presented in Figure 6-8.

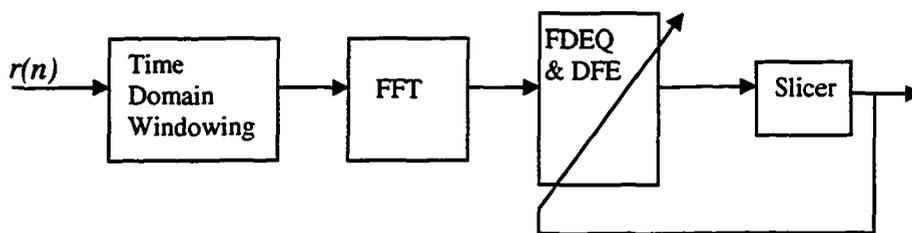


Figure 6-8: Proposed Constrained Windowing Architecture

One possible motive for not pursuing the described architecture simplification would be if it were beneficial for certain bins to never undergo the windowing operation. This could be the case for bins sufficiently far-away from the HAM bands as to never be significantly affected by the expected RFI. It would be advantageous in this scenario to not apply constrained windowing to these bins in order to avoid increasing their noise floor. However, simulation results show that with an assumed worst case RFI level of -10 dBm, even a tone located 1000 bins away from the RFI carrier is subject to -80 dBm of interference (see section 7.1.2.1), which would severely impair the bin's ability to carry data.

The following equation development, as given by Kapoor and Nedic [14], describes the constrained windowing operation along with the resulting controlled ICI. With nominal rectangular window demodulation, the recovered constellation points are given as:

$$\begin{aligned} R_i &= \frac{1}{\sqrt{2N_{sc}}} \sum_{n=0}^{2N_{sc}-1} r(n) e^{-j\frac{2\pi ni}{2N_{sc}}}, \quad i = 0, \dots, 2N_{sc} - 1 \\ &= H_i Z_i + Noise_i, \quad i = 0, \dots, 2N_{sc} - 1 \end{aligned} \quad (6-1)$$

where, $r(n)$ are the received samples, H is the channel transfer function DFT, Z_i are transmitted constellation points, and $Noise_i$ are the DFT points corresponding to the received noise component samples, including echo and RFI, given by:

$$Noise_i = \frac{1}{\sqrt{2N_{sc}}} \sum_{n=0}^{2N_{sc}-1} noise(n) e^{-j\frac{2\pi ni}{2N_{sc}}}, \quad i = 0, \dots, 2N_{sc} - 1 \quad (6-2)$$

Consider the window whose transfer function is given by:

$$W_i = \frac{1}{\sqrt{2N_{sc}}} \sum_{n=0}^{2N_{sc}-1} w(n) e^{-j\frac{2\pi ni}{2N_{sc}}}, \quad i = 0, \dots, 2N_{sc} - 1 \quad (6-3)$$

Then, the recovered constellation points with windowing applied is given by the following equations:

$$\begin{aligned} \hat{R}_i &= \frac{1}{\sqrt{2N_{sc}}} \sum_{n=0}^{2N_{sc}-1} r(n) w(n) e^{-j\frac{2\pi ni}{2N_{sc}}}, \quad i = 0, \dots, 2N_{sc} - 1 \\ &= \beta_i Z_i + V_i + U_i, \quad i = 0, \dots, 2N_{sc} - 1 \end{aligned} \quad (6-4)$$

where:

$$\begin{aligned} \beta_i &= W_0 H_i, \quad i = 0, \dots, 2N_{sc} - 1 \\ V_i &= \frac{1}{\sqrt{2N_{sc}}} \sum_{n=0}^{2N_{sc}-1} noise(n) w(n) e^{-j\frac{2\pi ni}{2N_{sc}}}, \quad i = 0, \dots, 2N_{sc} - 1 \\ U_i &= \sum_{l \neq i} Z_l H_l W_{l-i}, \quad i = 0, \dots, 2N_{sc} - 1 \end{aligned} \quad (6-5)$$

In the above, U_i represents the window-induced ICI experienced by the i^{th} bin from all other active bins. By requiring the window's DFT to be zero for negative indices, U_i will only contain ICI contributions from bins with indices higher than i . This requirement leads to complex-valued windows in the time domain (i.e. with an imaginary component). Since ICI induced on any given bin will only originate from higher active tones, decoding of tones can proceed sequentially, in a descending order, by first cancelling the interference which is solely inflicted by higher tones that have already been decoded. The cancellation is performed on a tone by tone basis through the use of a Decision Feedback Equalizer as given by equation (6-6):

$$\bar{R}_i = [1, -W_1, -W_2] \times [\hat{R}_i, \tilde{Z}_{i-1}H_{i-1}, \tilde{Z}_{i-2}H_{i-2}] \quad (6-6)$$

where \tilde{Z}_i is a decoded tone value. For the asymmetric windows, $W_2 = 0$.

Further details on the algorithm and its properties are presented in [14]. Table 6-2 lists the proposed constrained windows along with the values of their non-zero bins and their associated noise floor increase penalty.

Table 6-2: Constrained Windows Features

Proposed Window Name	Non-Zero Transfer Function Coefficients	Noise Floor Increase
Hanning	$W_0 = 1, W_1 = -2, W_2 = 1$	7.78 dB
Asymmetric{1, -1}	$W_0 = 1, W_1 = -1$	3 dB
Asymmetric{1, -0.5}	$W_0 = 1, W_1 = -0.5$	0.97 dB

The time-domain magnitudes of the proposed windows are presented in Figure 6-9. Note that these windows have indeed the same width as the elementary VDSL symbol (without the cyclic extensions). The transfer functions of these windows are shown in Figure 6-10 and Figure 6-11. As expected, the transfer functions have non-zero values at the bin locations specified in Table 6-2 above.

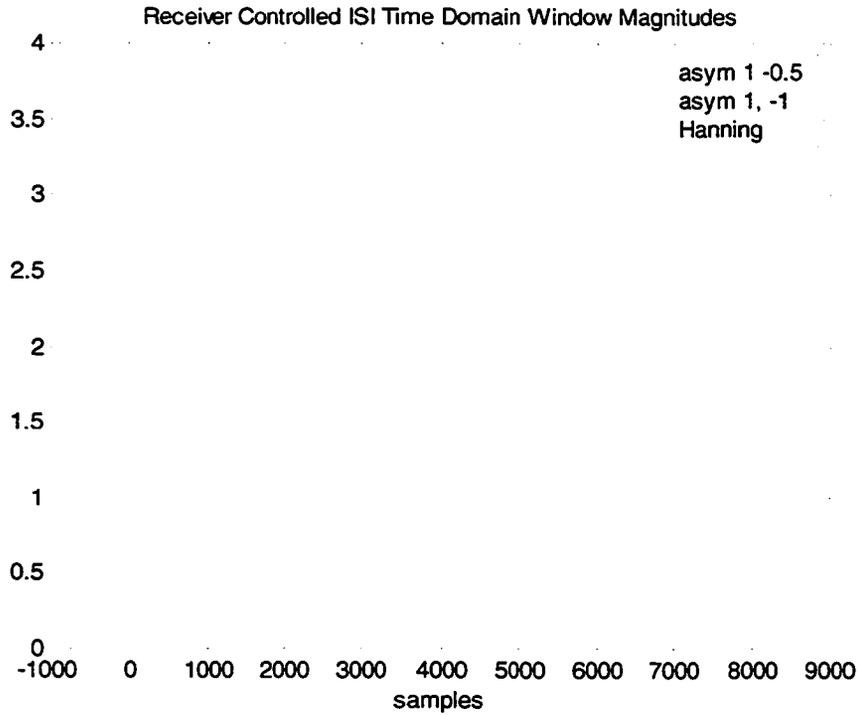


Figure 6-9: Magnitude of Constrained Time-Domain Windows

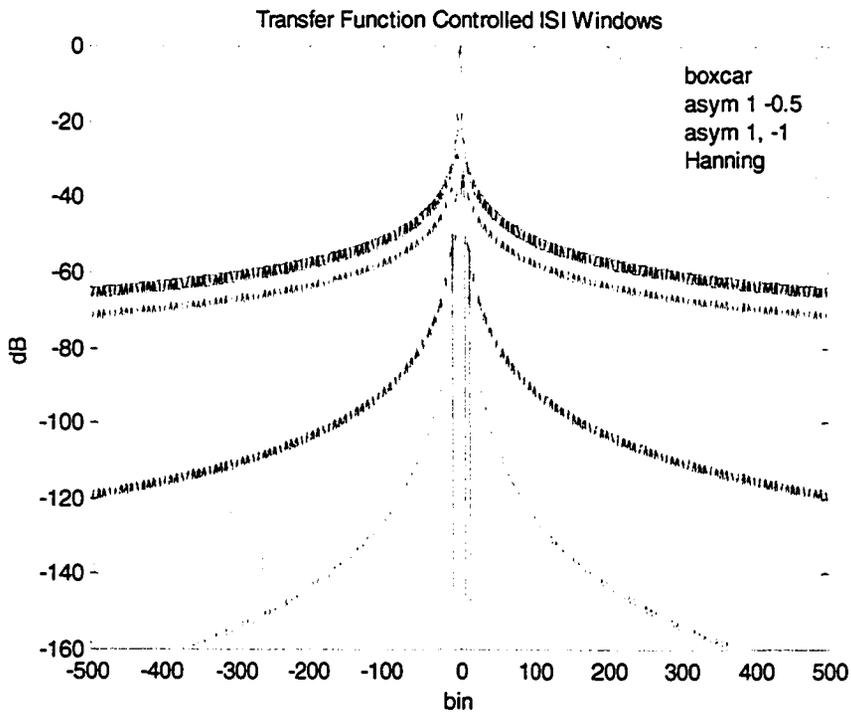


Figure 6-10: Transfer Function Magnitude of Constrained Windows

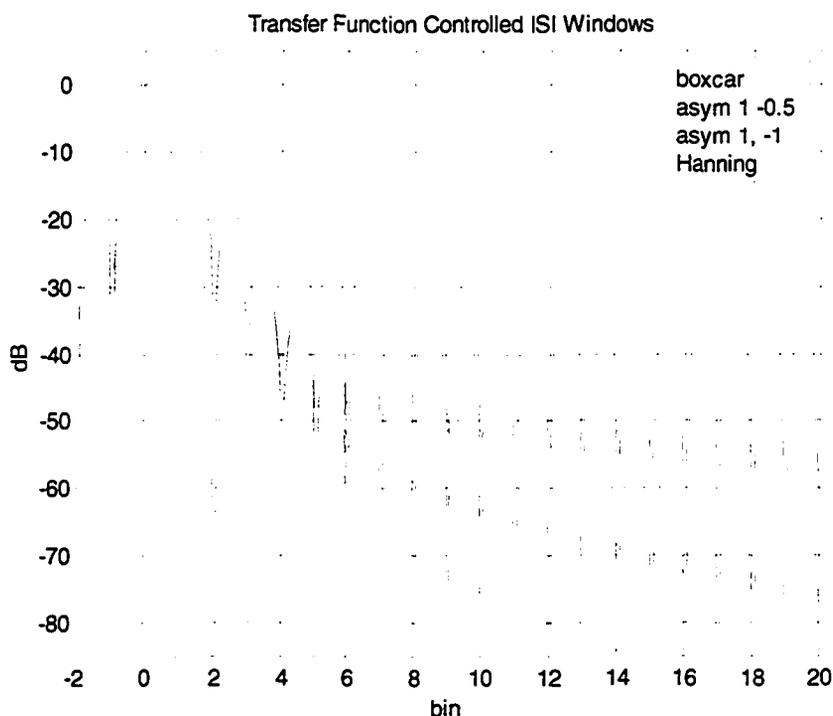


Figure 6-11: Bins 0 through 20 of Constrained Window Transfer Functions

The enhanced sidelobe attenuation brought on by constrained windowing, specifically near the main lobe, is significantly more than that provided by the raised cosine windows. For example, at a distance of 20 bins, the asymmetric{ 1, -1 } window provides 25.4 dB of enhanced attenuation compared to 8 dB for the 8576 sample raised cosine. At a distance of 100 bins however, the extra attenuations provided by the asymmetric{ 1, -1 } and the 8576 sample raised cosine windows are approximately the same (40dB). Thus, although constrained windowing might be better at protecting bins from nearby RFI than raised cosine windowing, the effectiveness of both techniques should be about the same for bins located at a further distance from the RF interferer.

The following characteristic of the constrained windowing algorithm is worth mentioning. While the receiver windowing operation affects the complete signal spectrum, the DFE can only be applied to active bins as it requires input from the slicer, which does not operate on inactive bins. Although this has no impact on performance, a number of non-active bins (one or two depending on the type of windowing) adjacent to

one side of each active band contains substantial interference power after the windowing processing. The reason for this is that the window induced ICI originating from the adjacent active bins is not cancelled-out on these one or two non-active bins.

Figure 6-12 shows components of the receive signal PSD prior and after constrained asymmetric $\{1, -1\}$ windowing for the case of a 2000 foot loop and one RFI located at bin 870.5. For clarity, the near-end echo component at the receiver is not shown. Note the RFI level, prior and after windowing, relative to the upstream signal. Receive SNRs are plotted in Figure 6-13. Four situations are considered: 1) the SNR at the AFE in the absence of RFI, 2) the SNR at the AFE in the presence of one -10 dBm RFI interferer located at bin 870.5, 3) the SNR after constrained windowing in the absence of RFI, and 4) the SNR after constrained windowing in the presence of the -10 dBm RFI at bin 870.5.

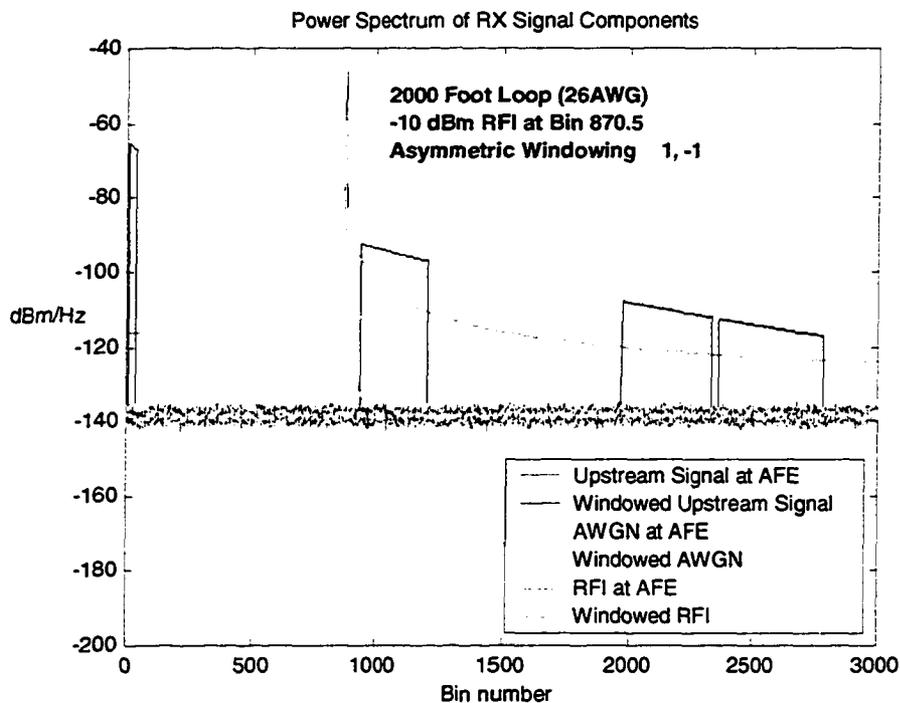


Figure 6-12: Receive Signal PSD with Asymmetric $\{1, -1\}$ Windowing

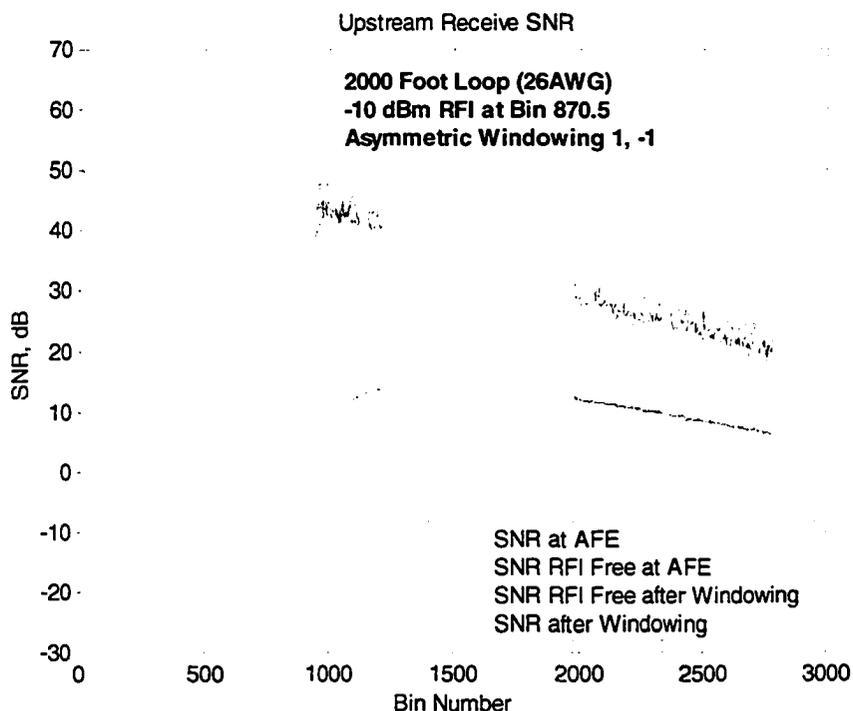


Figure 6-13: Receive SNR with and without Asymmetric {1,-1} Windowing

6.1.3.1 Constrained Windowing Computational Requirements

The computations requirements for executing the constrained windowing algorithm arise from the windowing and DFE processes as well as from the possible added FFT complexity required to handle the time-domain signal which becomes complex-valued instead of real-valued due to the windowing operation. For the windowing operation, one complex multiplication is required for each sample undergoing shaping, which in this case is equivalent to 2 real multiplications per sample (while the window coefficients may be complex, the received sample is real). As all the samples of an elementary symbol require shaping, the multiplications need in practice to be executed at the sampling rate, which translates in a rate requirement of 70.6×10^6 real multiplications per second or 17664 real multiplications per symbol.

Next, the computational requirement for performing the decision feedback equalization of a tone as per equation (6-6) depends on the number of non-zero coefficients of the window transfer function. For the case of the Hanning window that generates ICI in two bins, 2 complex multiplications are required to generate the data vector, which translates into 8 real multiplications and 4 real additions. Another 2 complex multiplications plus 2 complex additions are required to perform the equalization. As the window transfer function is real, the 2 complex multiplications translate into 4 real multiplications while the 2 complex additions translate into 4 real additions.

The DFE for the asymmetric windows requires half of the above computations because they only generate ICI on one bin. Assuming 1174 upstream active bins as per the ANSI frequency plan (see Ref. [4]), the computational requirements for the Hanning window DFE translates into 14088 real multiplications and 9392 real additions per symbol, while again half of this computational amount is required for the asymmetric window DFE.

6.2 RFI Suppression using an adaptive notch filter

As mentioned in section 3.3, the industry has generally avoided the use of digital notch filters to suppress ingress RFI in DMT-based VDSL systems. One important reason for this is that the use of digital filters increases the overall channel dispersion time, which has been shown to hinder DMT performance (see section 7.1.2). In fact, De Clercq et al. [5] suggest RFI notch filtering causes “unacceptable high intersymbol and intercarrier interference levels”. To what extent this is true is not immediately clear and needs to be investigated. For this purpose, the RFI detection and frequency tracking problems will be ignored for the time being. Detection and tracking can always be implemented based on the algorithms presented in [13], [25] and [29].

With the RFI source frequency known, simulation will be performed to evaluate the impact of notch filtering on system performance.

The single second order RFI notch filter presented in [15] and [22] has the following form:

$$H(z^{-1}) = \frac{1 - 2 \cos \omega_k z^{-1} + z^{-2}}{1 - 2r \cos \omega_k z^{-1} + r^2 z^{-2}} \quad (6-7)$$

where ω_k is the notch frequency and r is the pole contraction factor which determines the 3 dB bandwidth of the notch approximated as:

$$\text{bandwidth} \approx \pi(1 - r) \quad (6-8)$$

A large notch filter bandwidth is desirable for providing good attenuation over an 8 kHz band, which is necessary to effectively suppress a -10 dBm RF interferer and to allow for a frequency tracking error equivalent to 1 bin. Furthermore, in order to mitigate TD-ISI, a notch filter with the shortest possible impulse response should be employed, which is equivalent to selecting a filter with a bandwidth as large as possible. On the other hand, attenuation of the transmitted tones is to be avoided as much as possible. This translates into a requirement to minimise the attenuation beyond 10 bins with respect to the filter's centre frequency, which corresponds to the distance covered by the 10 RFI guard bins on each side of a HAM band (see section 2.2.7).

The transfer function of a notch filter with a pole contraction factor of 0.96 is shown in Figure 6-14. The 3 dB bandwidth of this filter is 700 kHz, corresponding to 164 bins. This may at first seem large for the purpose of filtering a 4 kHz bandpass RFI signal, but it only provides -35 dB of attenuation over an 8 kHz band. An 8 kHz band is considered here to allow for a centre frequency error offset of ± 2 kHz. This offset corresponds to the maximum error introduced by the trivial RFI frequency-tracking algorithm consisting of simple bin energy peak detection. Substantial tracking improvement can of course be achieved by using a more elaborate RFI frequency estimation algorithm. Assuming the worst case scenario of an RFI located on the edge of a HAM band, such a filter located at this position will attenuate the first active bin adjacent to it by 15 dB (the first 10 bins with respect to the filter's centre frequency are used as RFI guard bins). This attenuation

is relatively little when considering that without the notch filter, the RFI impairment at this position can potentially be as high as -40 dBm (-10 dBm RFI picked up by the -30 dB sidelobe of the rectangular window's transfer function -- see section 7.1.2.1). The impulse response of the notch filter is shown in Figure 6-15. Its dispersion time is approximately 75 samples which is similar to the dispersion time of a 1000 foot twisted pair (see Table 5-2). Figure 6-16 depicts the combined impulse response of two consecutive notch filters. As can be seen, the dispersion time of the convolved impulse responses is not significantly greater than the dispersion times of the individual impulse responses. All in all, it would appear that a notch filter with a pole contraction factor of 0.96 provides a good bandwidth compromise for the purpose of suppressing RFI in DMT-based VDSL. For severe cases of RFI, a more aggressive notch filter may be preferable.

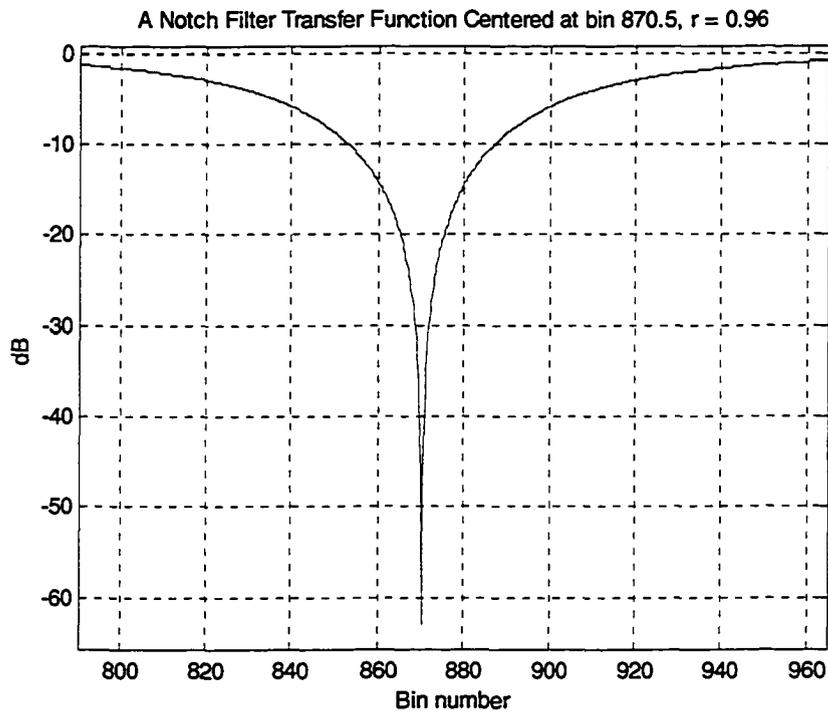


Figure 6-14: Transfer Function of Notch Filter ($r = 0.96$)

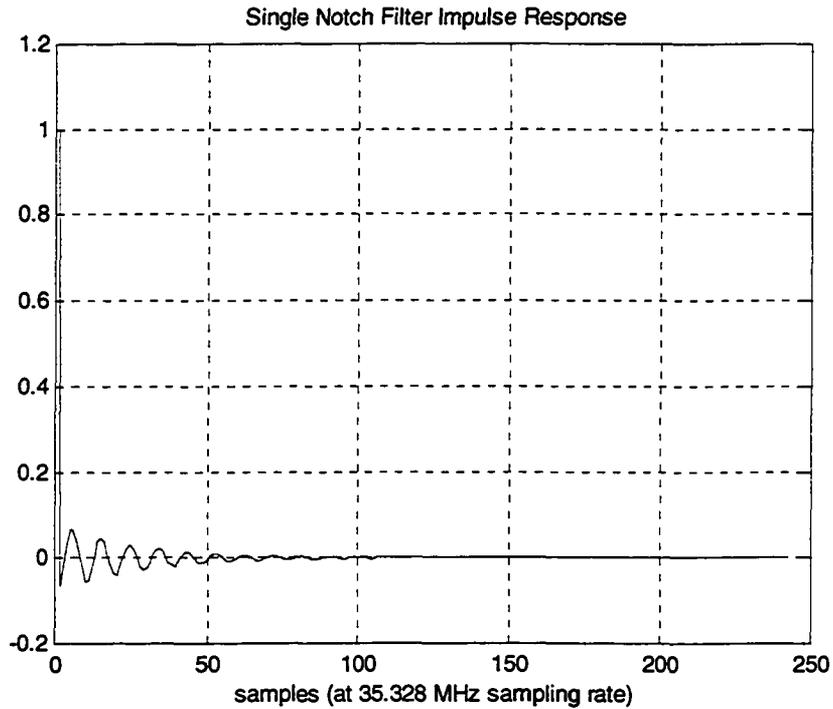


Figure 6-15: Impulse Response of Single Notch Filter ($r = 0.96$)

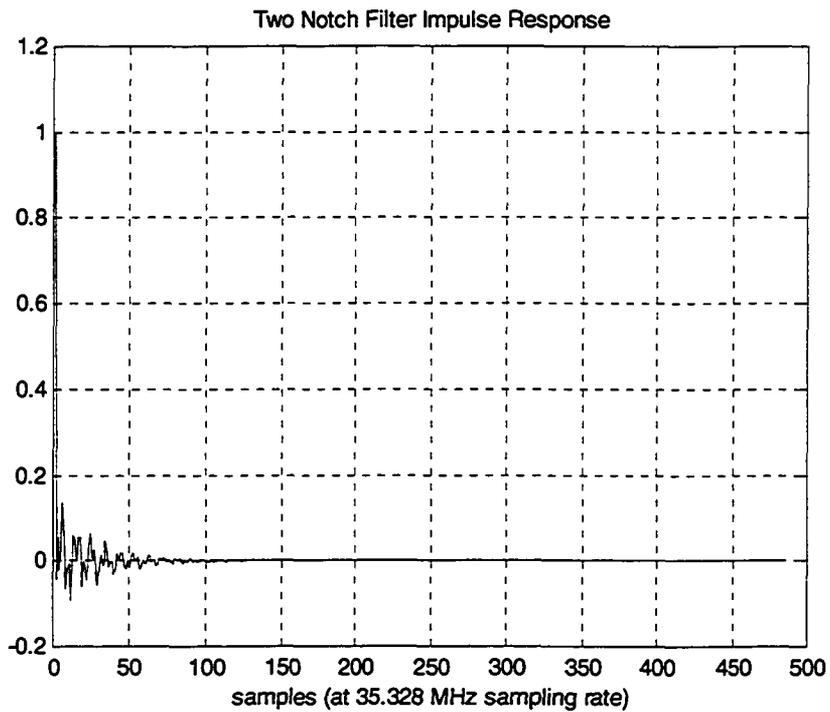


Figure 6-16: Impulse Response of Two Consecutive Notch Filters ($r = 0.96$)

Figure 6-17 shows components of the receive signal PSD before and after the application of notch filtering with a pole contraction factor of 0.91, for the case of a 2000 foot loop and one -10 dBm RFI located at bin 870.5. A pole contraction factor of 0.91 instead of 0.96 was selected in view of the severe RFI scenario. For clarity, the near-end echo component at the receiver is not shown. Note the RFI level before and after filtering relative to the upstream signal. Receive SNRs are plotted in Figure 6-18. Four situations are considered: the SNR at the AFE in the absence of RFI, the SNR at the AFE in the presence of one -10 dBm RFI interferer located at bin 870.5, the SNR after notch filtering in the absence of RFI, and the SNR notch filtering in the presence of the -10 dBm RFI at bin 870.5. Although notch filtering attenuates active bins in the vicinity of the filter's centre frequency, the noise floor is also reduced there thus preserving the SNR.

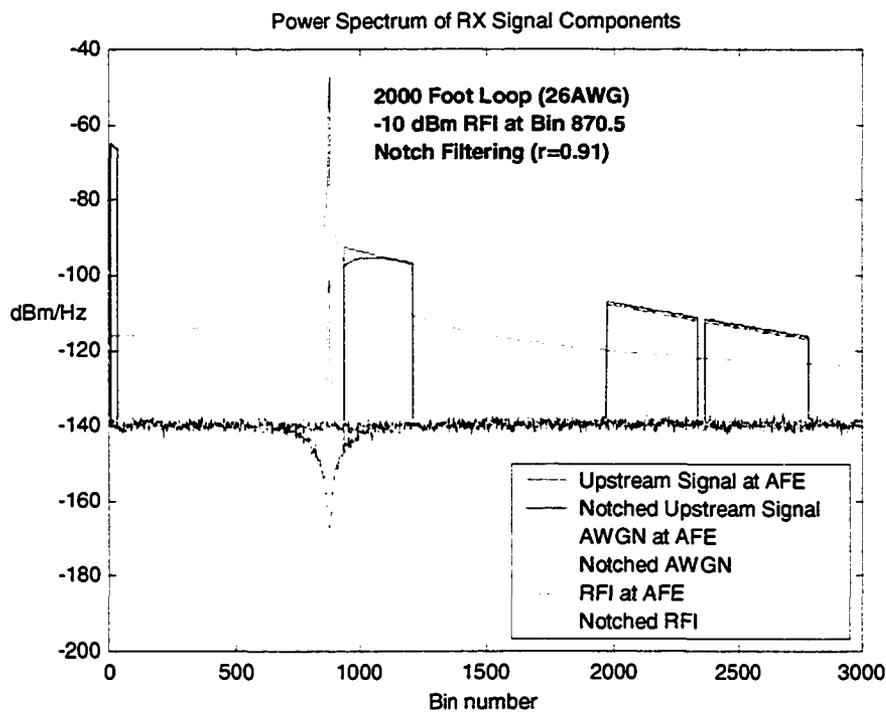


Figure 6-17: Receive Signal PSD with Notch Filtering

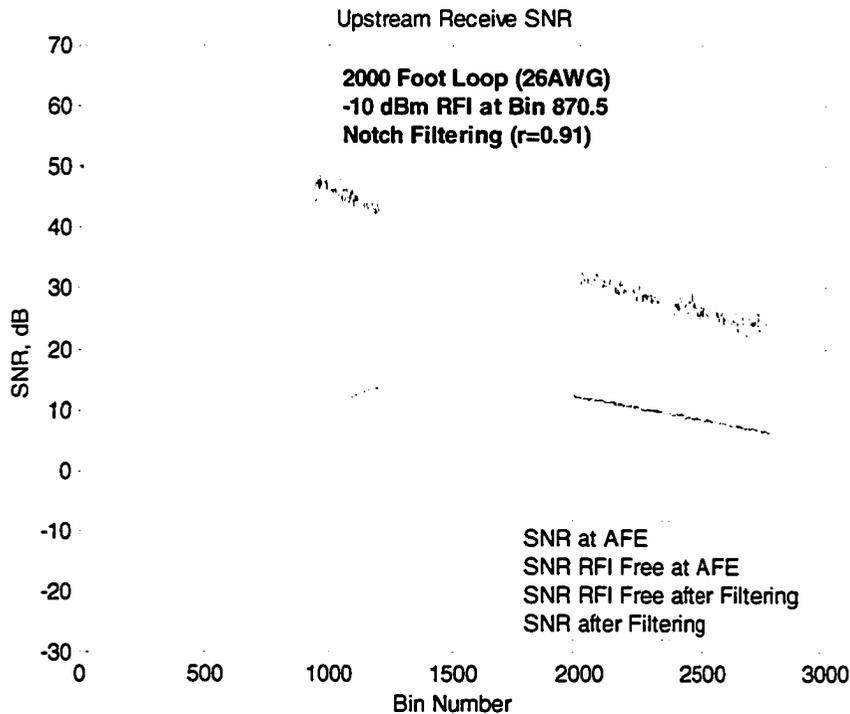


Figure 6-18: Receive SNR with and without Notch Filtering

6.2.1.1 Notch Filtering Computational Requirements

The computations requirements for performing second order notch filtering in the time domain is determined by observation of equation (6-7). For each filtered sample, the numerator contributes one real multiplication and two real additions, while the denominator contributes two real multiplications and two real additions for a total of three real multiplications and four real additions. With 8832 received samples per symbol period ($2*N_{sc} + L_{CE}$), this translates into 26496 real multiplications and 35328 real additions per symbol. The above development does not take into account the computations required to perform frequency tracking of the RFI.

6.3 RFI Cancellation using Time Domain Modelling

As mentioned in section 3.4, the RFI cancellation algorithms described in the referenced documents ([13], [17], [18] and [29]) rely on zero or first order time domain modelling of the RFI envelope to cancel its effect on a symbol by symbol basis. The first order model is given by (3-1), where the two parameters a and b are to be solved in the frequency domain with the use of RFI contribution measurements made on two bins in the vicinity of the RFI centre frequency. It is assumed that these bins are located in the HAM bands where VDSL bins are inactive.

For RFI cancellation based on first order modelling to work, the time-domain envelop of the RFI signal should have a straight line behaviour over the extended symbol period. This is equivalent to saying that the RFI bandwidth should be much less than the VDSL symbol transmission frequency. This condition turns out to be false, as both the RFI bandwidth and VDSL symbol transmission frequency are 4 kHz. It is thus suspected that RFI cancellation using a first order model will not meet its intent. This will be confirmed through the simulation analysis presented in section 7.4.

To get an idea of the model order required to achieve a given maximum error, consider the case of an RFI envelop equal to a pure 4 kHz sine wave. An estimation of the model error based on Taylor's Formula (see ref [27]) is given as:

$$Error_k(n) = \left| \frac{\sin^k(2 * \pi * 4000 * z / f_s)}{(k+1)!} * (n-c)^{k+1} \right| \quad (6-9)$$

where k is the model order, $\sin^k()$ is the k^{th} derivative of the $\sin()$ function, n is the index of the sample to be modelled, c is the index about which the model parameters are derived, z is a number between n and c , and f_s is sampling frequency equal to 35.328 MHz. Furthermore, both n and c lie in the 0 to 8831 ($2*N_{SC}+L_{CE} - 1$) interval. The error estimate attains its maximum when n and c are set at both extreme of the interval

resulting in $(n-c) = 8831$. Also, as the interval covers a full cycle period of the sinusoidal, the numerator will always attain its maximum value within the interval. Obviously this error estimate is very loosely upper-bounded. Nonetheless, it serves to emphasise the impact of the RFI's relatively large bandwidth. For example, an RFI with a much smaller bandwidth (e.g. 400 Hz) would result in a numerator upper bound orders of magnitude smaller with k even.

It is also worth mentioning that RFI cancellation as proposed is not well suited for combating more than one source at a time. Aside from requiring a complete second processing block to cancel the second RFI, the measurements used for modelling the individual RFIs in a two RFI scenario would contain contributions from both RFI sources thus corrupting the models.

Chapter 7

Simulation Results

This chapter presents simulation results for different scenarios. Performance benchmarks are first established, followed by performance analysis of the various RFI suppression algorithms under consideration.

A description of the simulation model is given in Chapter 5. Data rate results presented in the following sections are based on channel SNR profiles calculated over 20 symbols (SNR averaging), which has been shown to provide very good estimates.

7.1 Performance Benchmarks

The RFI suppression algorithms evaluated in the following sections may hinder or improve the VDSL receiver's ability to deal with TD-ISI and/or non-orthogonal echo. For example, the raised cosine windowing technique presented in section 6.1.1 demodulates part of the cyclic extensions to combat RFI, thus reducing the guard time and possibly increasing the receiver's vulnerability to TD-ISI and echo symbol transitions. On the other hand, receiver windowing, with its lower spectral sidelobes, may reduce the negative impact of non-orthogonal echo on receiver performance much in the same way as transmitter windowing [10] reduces out-of-band egress. To qualify potential benefits, it is necessary to first establish through simulation the system's susceptibility to near-end echo and receiver symbol misalignment in an RFI-free environment. Such a benchmark is helpful for making the appropriate compromises and for evaluating the performance of the various RFI suppression techniques in an RFI free environment.

Nominal receiver alignment is performed as outlined in section 5.4.2. In this case, the cyclic prefix is entirely used to contain TD-ISI while the cyclic suffix is used to contain echo symbol transitions. This arbitrary partitioning of the cyclic extension is selected to simplify the analysis and does not affect the validity of the observation due to the cyclic nature of the extended VDSL symbol. Ideally, the partitioning of the extended symbol into portions to contain TD-ISI and echo symbol transitions is made based on a data rate maximisation criterion. This would entail however complex analysis in the frequency domain. In practice, one or more simple time-domain criteria, such as interference energy, are used to determine receiver symbol alignment.

7.1.1 Impact of Near-end Echo on VDSL Performance

Upstream simulations were performed with and without downstream signal present to quantify the effect of near-end echo on upstream rates. Furthermore, simulations with downstream signal present were performed twice, once with transmit shaping applied and once without. Rate degradation results are given in Table 7-1 for the selected loop models presented in section 5.3.1.

The simulation results show that rate degradation due to echo occurs on loops of length greater than 3000 ft. It can be surmised that at the 3000-ft threshold, the cyclic suffix length becomes insufficient to completely contain echo symbol transitions. Rate degradation is especially significant on loops of length equal or greater than 4000 ft, which is consistent with the duration of the echo path impulse responses presented in section 5.5.6. With these loops, rates are degraded by approximately 50% when non-shaped downstream signal is present. The application of transmit shaping to the downstream signal allows for a substantial amount of the rate degradation to be recovered, except for the cases of very long loops (i.e. > 5000 ft), although even these cases show significant improvement when using transmit shaping.

Table 7-1: Impact of Near-end Echo with and without Transmit Shaping Applied

VDSL Loop Model	Loop Length (feet)	Rates with no Echo (kbps)	Rate Degradation Due to Echo				Rate Improvement Due to Tx Windowing	
			No Tx Windowing		Tx Windowing		Absolute (kbps)	Relative (%) To Rates without Tx Windowing
			Absolute (kbps)	Relative (%)	Absolute (kbps)	Relative (%)		
1 (26 AWG)	500	64500	-52	-0.08%	-52	-0.08%	0	0.00%
1 (26 AWG)	1000	60648	-52	-0.09%	-52	-0.09%	0	0.00%
1 (26 AWG)	1500	45872	-52	-0.11%	-52	-0.11%	0	0.00%
1 (26 AWG)	2000	29256	-48	-0.16%	-48	-0.16%	0	0.00%
1 (26 AWG)	2500	13472	-80	-0.59%	-72	-0.53%	8	0.06%
1 (26 AWG)	3000	7272	-80	-1.10%	-64	-0.88%	16	0.22%
1 (26 AWG)	3500	4408	-348	-7.89%	-148	-3.36%	200	4.93%
1 (26 AWG)	4000	1524	-588	-38.58%	-148	-9.71%	440	47.01%
1 (26 AWG)	4500	1220	-528	-43.28%	-120	-9.84%	408	58.96%
1 (26 AWG)	5000	1096	-528	-48.18%	-104	-9.49%	424	74.65%
1 (26 AWG)	5500	1004	-524	-52.19%	-172	-17.13%	352	73.33%
1 (26 AWG)	6000	920	-484	-52.61%	-220	-23.91%	264	60.55%
4	3300	3612	-364	-10.08%	-168	-4.65%	196	6.03%
5	950	60884	-48	-0.08%	-48	-0.08%	0	0.00%
6	3250	6704	-72	-1.07%	-68	-1.01%	4	0.06%
7	4900	1224	-496	-40.52%	-104	-8.50%	392	53.85%

7.1.2 VDSL Performance Sensitivity to Receiver Alignment

To investigate the effect of symbol misalignment on TD-ISI, rate simulations were performed with receiver alignment offsets of plus and minus 112 samples with respect to nominal alignment, with no echo present. A plus 112 sample alignment offset translates into a cyclic prefix increase of 25% and a corresponding cyclic suffix decrease of 25%, while the opposite is true for a negative 112 sample alignment offset. It is worth keeping in mind that 256 samples of each cyclic extension (nominally 57%) are used for transmit shaping and should not be fed to the demodulator. Results and relative deltas with respect to nominal rates are presented in Table 7-2 for selected loops. Rates versus VDSL1 loop lengths are plotted in Figure 7-1 for various receiver alignments.

With no near-end echo present, a bias partitioning of the cyclic extensions to the advantage of the cyclic prefix benefits containment of TD-ISI at no cost, since the cyclic suffix is not required. This leads to a performance gain for loop lengths greater than 3500 feet, which is reflected in the simulation results for the case of the positive alignment offset. Conversely, a negative alignment offset decreases the receiver's ability to contain TD-ISI, resulting in performance degradation and rate deterioration with these same loops. It is apparent from the simulation results that the impact of receiver misalignment on rates is more pronounced on longer loops. This is not surprising, since these loops exhibit a loop dispersion time (see section 5.5.4) approximately equal to or greater than the nominal cyclic prefix length.

Table 7-2: Loop rates with various receiver symbol alignments and no echo present

VDSL Loop Model	Loop Length (feet)	Rates with Nominal Alignment (kbps)	Rate Change with -112 Samples Alignment Offset		Rate Change with +112 Samples Alignment Offset	
			Absolute (kbps)	Relative (%)	Absolute (kbps)	Relative (%)
1 (26 AWG)	500	64500	0	0.00%	0	0.00%
1 (26 AWG)	1000	60648	-24	-0.04%	28	0.05%
1 (26 AWG)	1500	45872	-16	-0.03%	16	0.03%
1 (26 AWG)	2000	29256	80	0.27%	40	0.14%
1 (26 AWG)	2500	13472	-20	-0.15%	-40	-0.30%
1 (26 AWG)	3000	7272	-204	-2.81%	16	0.22%
1 (26 AWG)	3500	4408	-436	-9.89%	40	0.91%
1 (26 AWG)	4000	1524	-304	-19.95%	276	18.11%
1 (26 AWG)	4500	1220	-204	-16.72%	216	17.70%
1 (26 AWG)	5000	1096	-164	-14.96%	176	16.06%
1 (26 AWG)	5500	1004	-136	-13.55%	144	14.34%
1 (26 AWG)	6000	920	-120	-13.04%	124	13.48%
4	3300	3612	-400	-11.07%	260	7.20%
5	950	60884	28	0.05%	4	0.01%
6	3250	6704	-228	-3.40%	8	0.12%
7	4900	1224	-168	-13.73%	212	17.32%

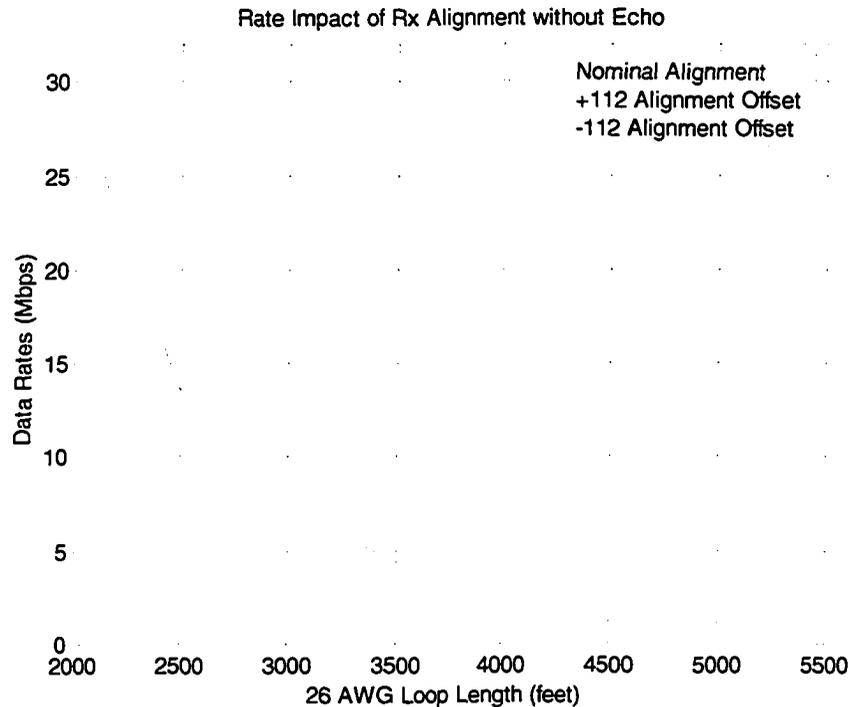


Figure 7-1: Rate Reach with no Echo for Various Receiver Alignments

Rate simulations were also performed to investigate the impact that symbol misalignment has on the receiver's ability to contend with near-end echo, with and without transmit shaping applied. For the case of no transmit shaping, results and relative deltas with respect to nominal rates are presented in Table 7-3 for selected loops. Rates versus VDSL1-x loop lengths are plotted in Figure 7-2 for various receiver alignments. Results show that a positive alignment offset degrades rates for loop lengths between 2000 ft to 4500 ft, while a negative alignment offset improves rates for loop lengths between 4500 ft to 6000 ft. It thus appears that the gains achieved in an echo free environment by biasing the cyclic extensions for mitigating TD-ISI (i.e. positive alignment offset) are more than offset by the shortfalls caused by non-orthogonal echo when transmit shaping is not applied. In this case, system performance, with loops greater than 4000 feet specifically, actually benefit from an opposite bias of the cyclic extension (i.e. negative alignment offset), which increases containment of near-end echo symbol transitions.

Table 7-3: Loop rates with no transmit shaping and various symbol alignments

VDSL Loop Model	Loop Length (feet)	Rates with Nominal Alignment (kbps)	Rate Change with -112 Samples Alignment Offset		Rate Change with +112 Samples Alignment Offset	
			Absolute (kbps)	Relative (%)	Absolute (kbps)	Relative (%)
1 (26 AWG)	500	64448	0	0.00%	0	0.00%
1 (26 AWG)	1000	60596	-44	-0.07%	20	0.03%
1 (26 AWG)	1500	45820	-208	-0.45%	16	0.03%
1 (26 AWG)	2000	29208	-464	-1.59%	-9552	-32.70%
1 (26 AWG)	2500	13392	-620	-4.63%	-7628	-56.96%
1 (26 AWG)	3000	7192	-736	-10.23%	-4404	-61.23%
1 (26 AWG)	3500	4060	-596	-14.68%	-3376	-83.15%
1 (26 AWG)	4000	936	-44	-4.70%	-292	-31.20%
1 (26 AWG)	4500	692	128	18.50%	-120	-17.34%
1 (26 AWG)	5000	568	188	33.10%	-56	-9.86%
1 (26 AWG)	5500	480	236	49.17%	0	0.00%
1 (26 AWG)	6000	436	168	38.53%	-40	-9.17%
4	3300	3248	-540	-16.63%	-2540	-78.20%
5	950	60836	24	0.04%	16	0.03%
6	3250	6632	-520	-7.84%	-4624	-69.72%
7	4900	728	144	19.78%	-92	-12.64%

Finally, Table 7-4 presents the simulation results for the case of downstream signal present with transmit shaping applied, which is the most likely scenario. Rates versus VDSL1-x loop lengths are plotted in Figure 7-3 for various receiver alignments. Biasing the cyclic extensions with a positive alignment offset at the receiver has similar consequences to those of the previous scenario where transmit shaping was not applied, except that rate degradations in this case are less severe. But while in the previous scenario a negative alignment offset improves rates on some loops, in this case such a bias toward mitigating near-end echo leads to rate degradations. These results support the choice of the nominal alignment based on the channel's propagation delay (see sections 5.4.2) as an appropriate compromise between allocating the cyclic extensions for TD-ISI containment versus near-end echo transient containment. They also concur with the loop propagation delay and dispersion time measurements presented in sections 5.5.4 and 5.5.5.

It is also worth mentioning that the performance in the presence of near-end echo and the performance in the presence of synchronised NEXT should interact with receiver misalignments in a similar fashion. It could therefore be beneficial to bias the alignment towards increasing echo/NEXT containment (i.e. negative alignment offset) when both of these impairments are present. Investigation of this aspect is beyond the defined scope of this thesis.

Table 7-4: Loop rates with transmit shaping and various symbol alignments

VDSL Loop Model	Loop Length (feet)	Rates with Nominal Alignment (kbps)	Rate Change with -112 Samples Alignment Offset		Rate Change with +112 Samples Alignment Offset	
			Absolute (kbps)	Relative (%)	Absolute (kbps)	Relative (%)
1 (26 AWG)	500	64448	0	0.00%	0	0.00%
1 (26 AWG)	1000	60596	-24	-0.04%	20	0.03%
1 (26 AWG)	1500	45820	-16	-0.03%	16	0.03%
1 (26 AWG)	2000	29208	44	0.15%	8	0.03%
1 (26 AWG)	2500	13400	-40	-0.30%	-300	-2.24%
1 (26 AWG)	3000	7208	-376	-5.22%	-1168	-16.20%
1 (26 AWG)	3500	4260	-468	-10.99%	-2036	-47.79%
1 (26 AWG)	4000	1376	-316	-22.97%	-412	-29.94%
1 (26 AWG)	4500	1100	-180	-16.36%	-240	-21.82%
1 (26 AWG)	5000	992	-136	-13.71%	-316	-31.85%
1 (26 AWG)	5500	832	-48	-5.77%	-284	-34.13%
1 (26 AWG)	6000	700	40	5.71%	-220	-31.43%
4	3300	3444	-476	-13.82%	-936	-27.18%
5	950	60836	24	0.04%	16	0.03%
6	3250	6636	-352	-5.30%	-1452	-21.88%
7	4900	1120	-168	-15.00%	-224	-20.00%

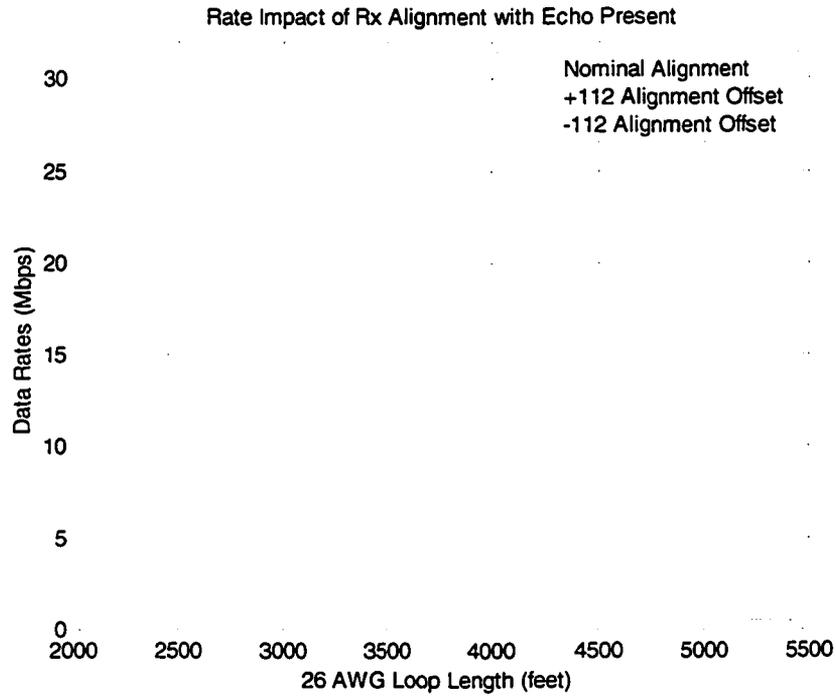


Figure 7-2: Rate Reach without Transmit Shaping for Various Receiver Alignments

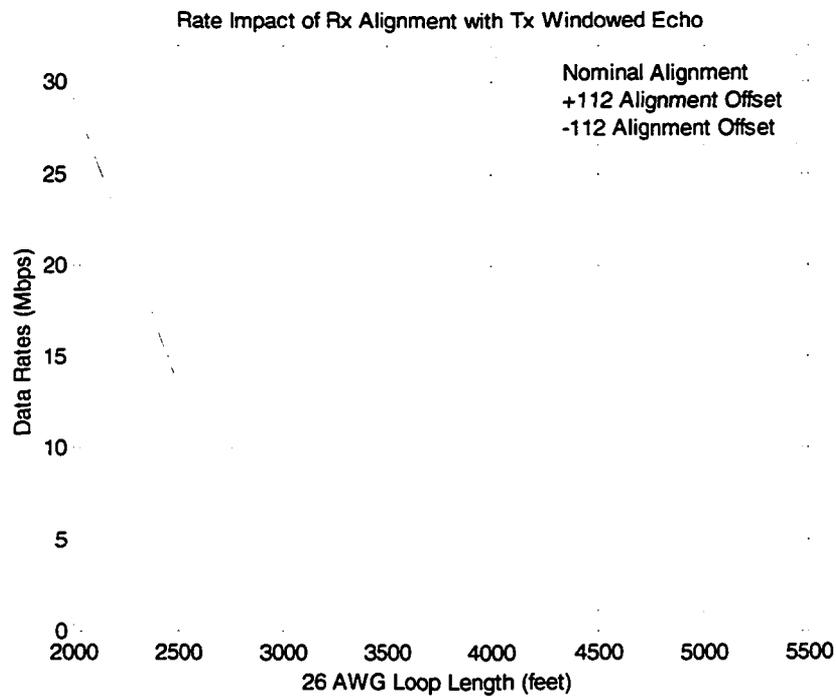


Figure 7-3: Rate Reach with Transmit Shaping for Various Receiver Alignments

7.1.2.1 Nominal Benchmark

The rates with nominal alignment presented in Table 7-4 are used as the nominal benchmark from hereon in. These rates were achieved in simulation using the following settings:

- downstream echo with transmit windowing applied
- nominal receiver alignment
- AWGN at -140 dBm/Hz

7.1.3 RFI impact on VDSL Performance

Figure 7-4 depicts the transfer function of the demodulator's nominal (rectangular) window. As can be seen, the sidelobes are relatively high, even at a distance of 1000 bins from the main lobe. To illustrate the significance of the sidelobes at this distance, a -10 dBm RFI signal would only be attenuated down to -80 dBm. Considering that a VDSL signal is typically transmitted at a power of -60 dBm/Hz and that it undergoes loop attenuation, such an impairment could significantly impact the ability for that bin, as well as the ability of any other bin closer to the RFI source, to carry data.

Figure 7-5 shows the spectrums at the receiver of the upstream signal and of a -10 dBm RFI centred on bin 870.5. The receive SNRs, with and without the RFI present, are depicted in Figure 7-6. From these results, it is clear that VDSL performance suffers a significant deterioration in the presence of a single -10 dBm RFI.

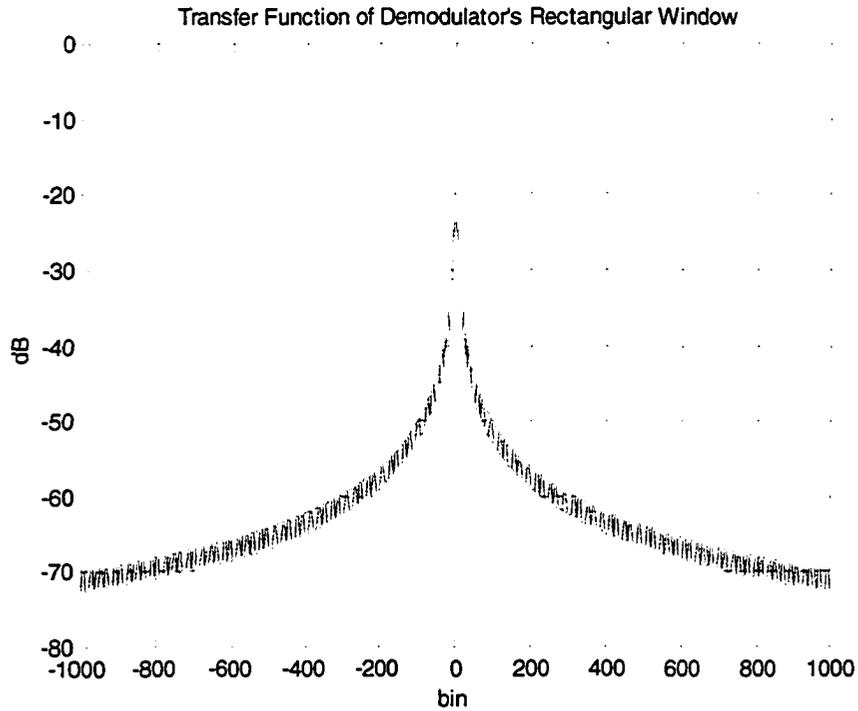


Figure 7-4: Transfer Function of the Demodulator's Nominal Window

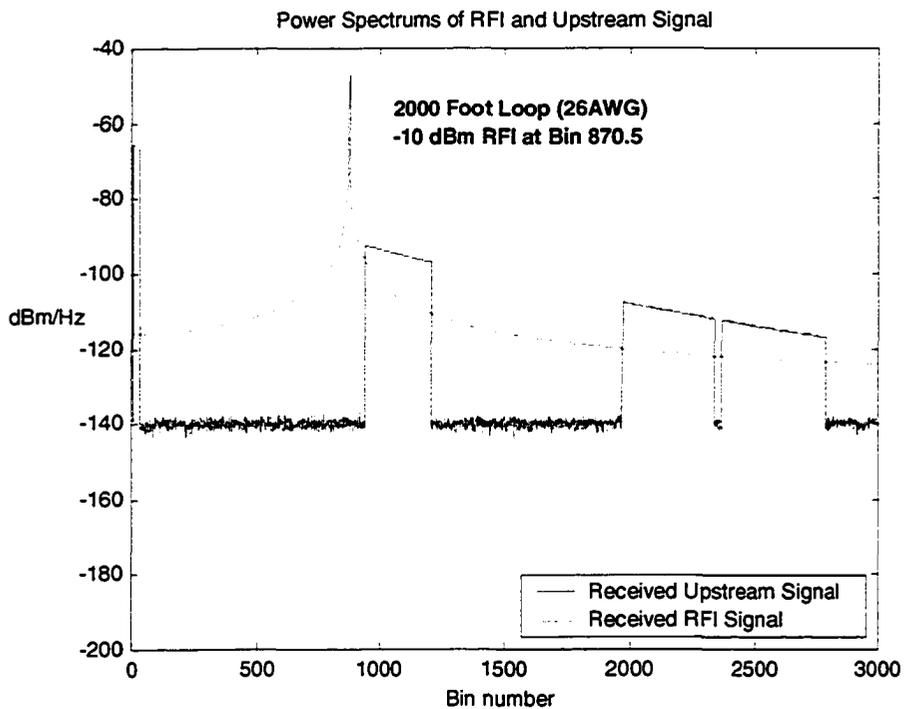


Figure 7-5: RFI and Upstream PSD

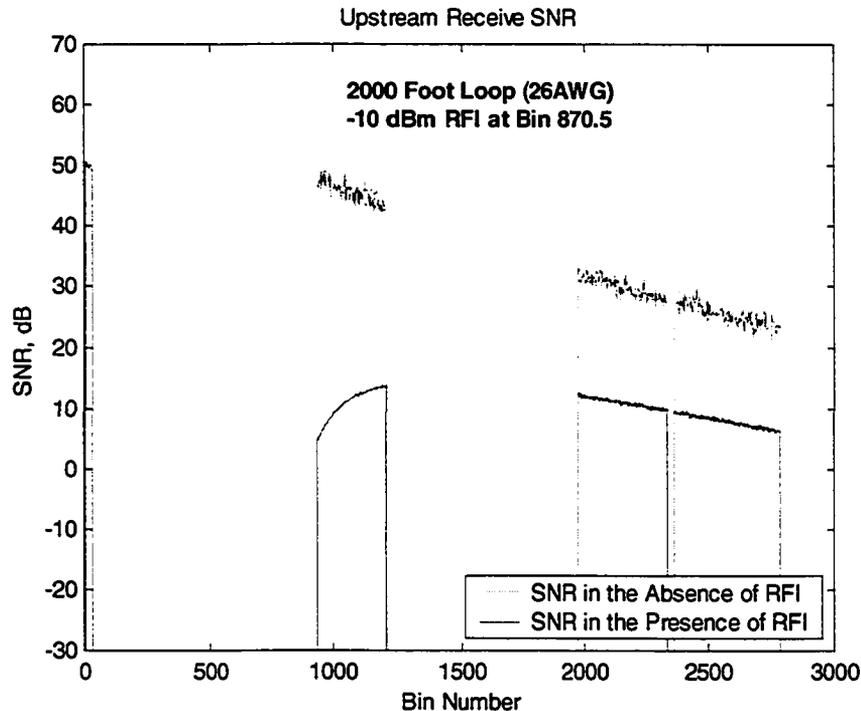


Figure 7-6: Receive SNR With and Without RFI Present

Figure 7-7 shows rate degradation versus loop length for three cases of RFI. One curve refers to a single -10 dBm RFI source centred at bin 870.5, another refers again to a single -10 dBm RFI source but this time centred at bin 2343.5, and the final curve refers to the case of two distinct -10 dBm RFI sources located at bins 870.5 and 2343.5 respectively. On short loops, all active bins carry close to the maximum number of bits permissible, but the lower bins tend to have more SNR margins due to the low channel attenuation at these frequencies. These bins are therefore less susceptible to RFI than higher bins. This is why performance degradation with these short loops is more pronounced for the case of a high frequency RFI than for the case of a low frequency RFI. The opposite is true however for longer loops. In these cases, the lower bins account for the majority of the channel's capacity since the higher bins are heavily attenuated. The lower bins therefore tend to carry significantly more bits than the higher bins and the overall data rate will thus be more impacted by an interferer at lower frequency.

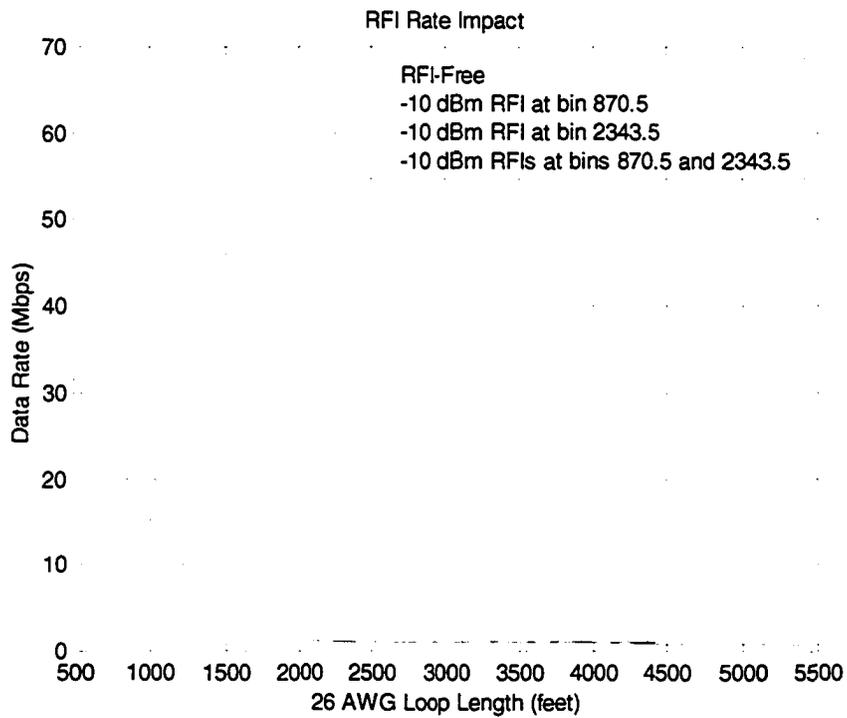


Figure 7-7: RFI Impact on VDSL Rates

Table 7-5 lists the corresponding rate values and degradations relative to nominal rates.

Table 7-5: RFI Impact on Performance

VDSL Loop Model	Loop Length (feet)	Nominal Rates (kbps)	Rate Change with -10 dBm RFI at bin 870.5		Rate Change with -10 dBm RFI at bin 2343.5		Rate Change with two -10 dBm RFI at bins 870.5 and 2343.5	
			Absolute (kbps)	Relative (%)	Absolute (kbps)	Relative (%)	Absolute (kbps)	Relative (%)
1 (26 AWG)	500	64448	-16396	-25.4%	-28024	-43.5%	-32816	-50.9%
1 (26 AWG)	1000	60596	-28460	-47.0%	-39964	-66.0%	-45340	-74.8%
1 (26 AWG)	1500	45820	-29760	-64.9%	-36056	-78.7%	-42624	-93.0%
1 (26 AWG)	2000	29208	-27908	-95.5%	-22464	-76.9%	-28072	-96.1%
1 (26 AWG)	2500	13400	-12124	-90.5%	-9648	-72.0%	-12332	-92.0%
1 (26 AWG)	3000	7208	-5996	-83.2%	-5752	-79.8%	-6200	-86.0%
1 (26 AWG)	3500	4260	-3116	-73.1%	-3008	-70.6%	-3288	-77.2%
1 (26 AWG)	4000	1376	-336	-24.4%	-196	-14.2%	-392	-28.5%
1 (26 AWG)	4500	1100	-164	-14.9%	-24	-2.2%	-184	-16.7%
1 (26 AWG)	5000	992	-56	-5.6%	-4	-0.4%	-176	-17.7%
1 (26 AWG)	5500	832	-36	-4.3%	0	0.0%	-144	-17.3%
1 (26 AWG)	6000	700	-76	-10.9%	-4	-0.6%	-60	-8.6%
4	3300	3444	-2404	-69.8%	-2340	-67.9%	-2480	-72.0%
5	950	60836	-28124	-46.2%	-38176	-62.8%	-43964	-72.3%
6	3250	6636	-5444	-82.0%	-5180	-78.1%	-5468	-82.4%
7	4900	1120	-184	-16.4%	-36	-3.2%	-184	-16.4%

7.2 Windowing Performance

A set of scenarios was simulated to investigate the performance of both raised cosine and constrained windowing under various conditions. Aside from evaluating RFI-suppression performance, it is of interest to determine the impact of windowing in the following RFI-free settings:

- nominal setting (as defined in section 7.1.2.1)
- negative receiver misalignment to emphasise the impact of TD-ISI
- positive receiver misalignment to emphasise the impact of near-end echo
- various levels of AWGN

7.2.1 Nominal Windowing Performance

Table 7-6 lists the relative performance of the various windowing algorithms in a nominal environment (i.e. -140 dBm/Hz AWGN, no symbol misalignment, no RFI).

The raised cosine windowing performances corresponding to the three window sizes under consideration are illustrated in Figure 7-8. As can be observed, raised cosine windowing provides some marginal rate improvements, especially on longer loops. This is interesting because the cyclic extensions, which are not sufficiently long in normal circumstances to properly contain TD-ISI and echo transients with these longer loops, are effectively reduced in length due to the raised cosine windows extending into them. The gain in performance is most likely due to sidelobe attenuation of the TD-ISI and echo interference. This conjecture is investigated in following sections. It is also worth noting that raised cosine windowing performance is maximised with the mid-size raised cosine window (8384 samples). The 8384 sample window thus appears to provide the best compromise between sidelobe attenuation and cyclic extension preservation in the nominal scenario.

The constrained windowing performance is depicted in Figure 7-9. Rate degradation is apparent for loops of length less than 4000 feet, which can be attributed to the noise enhancement characteristic of the constrained windowing technique. As expected, rate degradation is most significant with the Hanning window, which picks off noise from two adjacent bins as opposed to one for the two other windows. On the other hand, constrained windowing improves rates on longer loops, which again can be credited to sidelobe attenuation of the TD-ISI and echo interference. The gain is least apparent for the Asymmetric{1, -0.5} window, which also provides the least sidelobe attenuation.

Table 7-6: Windowing Impact on Rates in RFI-Free Environment

VDSL Loop Model	Loop Length (feet)	Nominal Rates with Rectangular Window (kbps)	Relative Rate Change with respect to Nominal Rates					
			R-Cos 8288	R-Cos 8384	R-Cos 8576	Asym 1, -0.5	Asym 1, -1	Hanning
1-x	500	64448	0.0%	0.0%	0.0%	0.0%	-0.1%	-0.1%
1-x	1000	60596	0.0%	0.0%	0.1%	-1.6%	-5.1%	-14.8%
1-x	1500	45820	0.1%	0.1%	0.2%	-3.0%	-9.0%	-24.1%
1-x	2000	29208	0.1%	0.1%	0.2%	-4.5%	-14.1%	-35.8%
1-x	2500	13400	0.3%	0.5%	0.6%	-7.9%	-22.6%	-43.9%
1-x	3000	7208	0.1%	0.3%	0.5%	-4.8%	-14.9%	-36.1%
1-x	3500	4260	0.8%	0.9%	0.9%	-6.7%	-24.0%	-57.4%
1-x	4000	1376	3.5%	6.1%	5.8%	8.4%	13.4%	13.4%
1-x	4500	1100	1.5%	2.9%	2.5%	19.3%	41.8%	41.1%
1-x	5000	992	0.8%	4.0%	2.0%	21.8%	57.3%	53.2%
1-x	5500	832	7.7%	8.7%	6.7%	26.0%	84.6%	76.4%
1-x	6000	700	4.0%	10.9%	10.3%	29.1%	110.9%	105.1%
4	3300	3444	4.2%	5.0%	8.4%	0.9%	-7.5%	-44.9%
5	950	60836	0.0%	0.0%	0.1%	-1.6%	-4.9%	-13.7%
6	3250	6636	0.1%	0.2%	0.3%	-4.8%	-15.0%	-38.1%
7	4900	1120	1.1%	2.5%	5.0%	19.3%	39.3%	39.3%

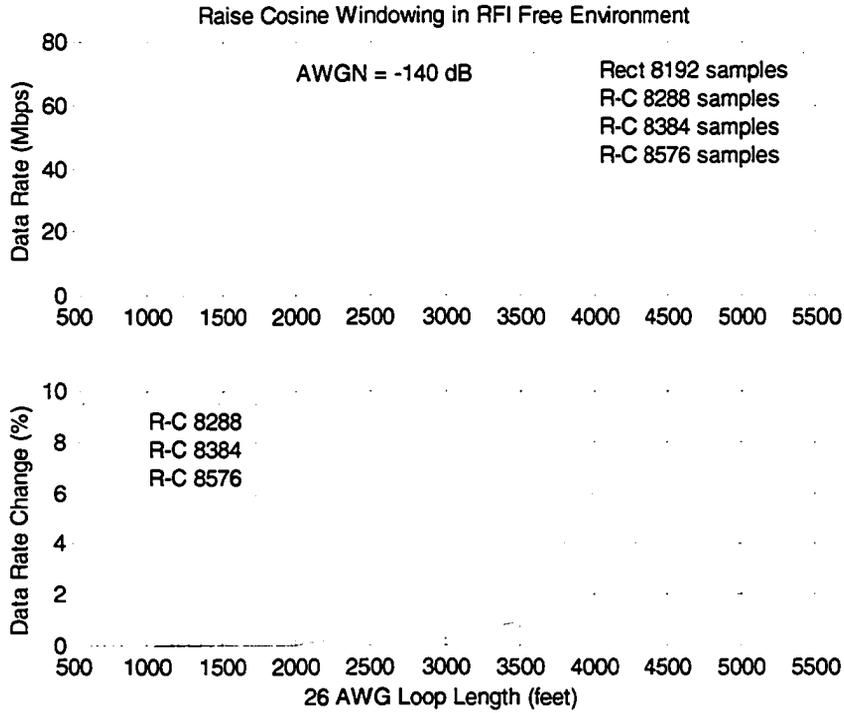


Figure 7-8: Raised Cosine Windowing Impact in RFI-Free Environment

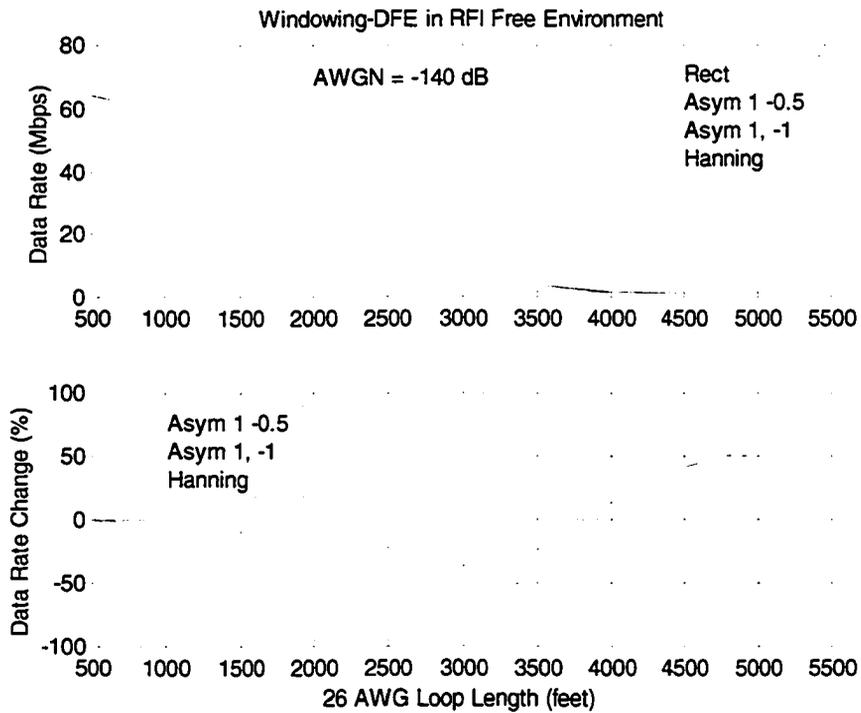


Figure 7-9: Constrained Windowing Impact in RFI-Free Environment

7.2.2 Windowing Performance with Symbol Misalignments

This section looks at the impact of symbol misalignment at the receiver on windowing performance, with no RFI present. A negative misalignment serves to increase the effect of TD-ISI, while a positive misalignment increases the effect of near-end echo transients. The performances of raised cosine windowing relative to rectangular windowing for -112, 0, and +112 alignment offsets are shown in Figure 7-10. Similar information for constrained windowing is depicted in Figure 7-11.

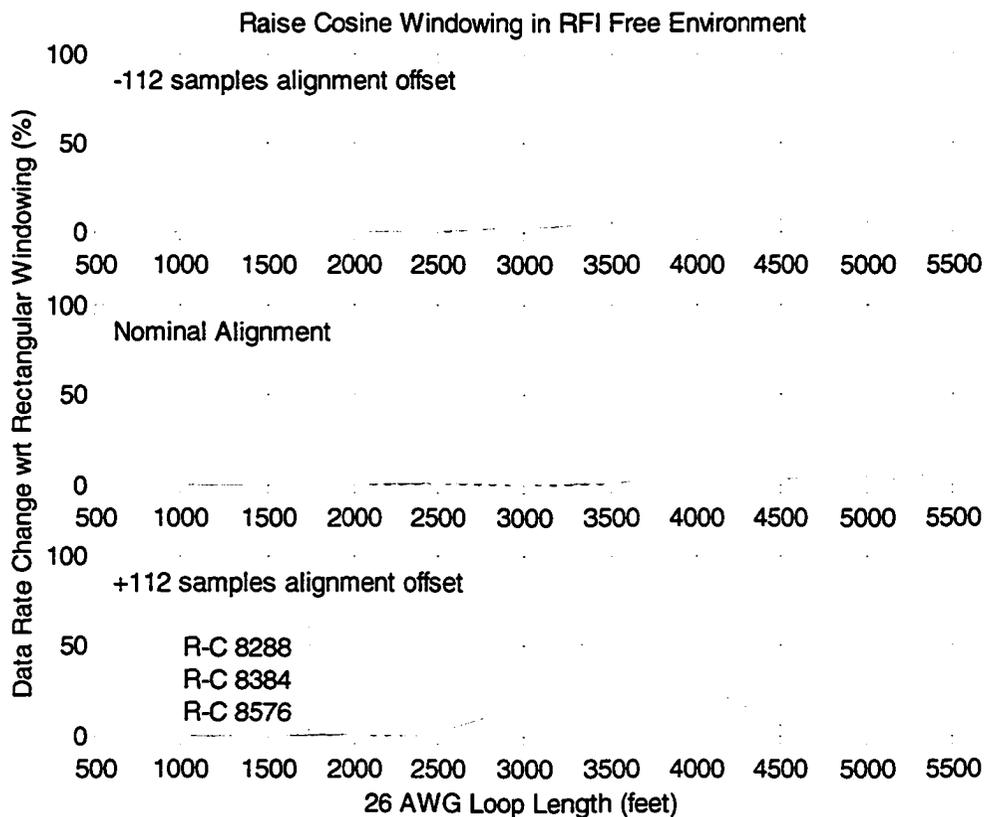


Figure 7-10: Raised Cosine Windowing Impact with Symbol Misalignments

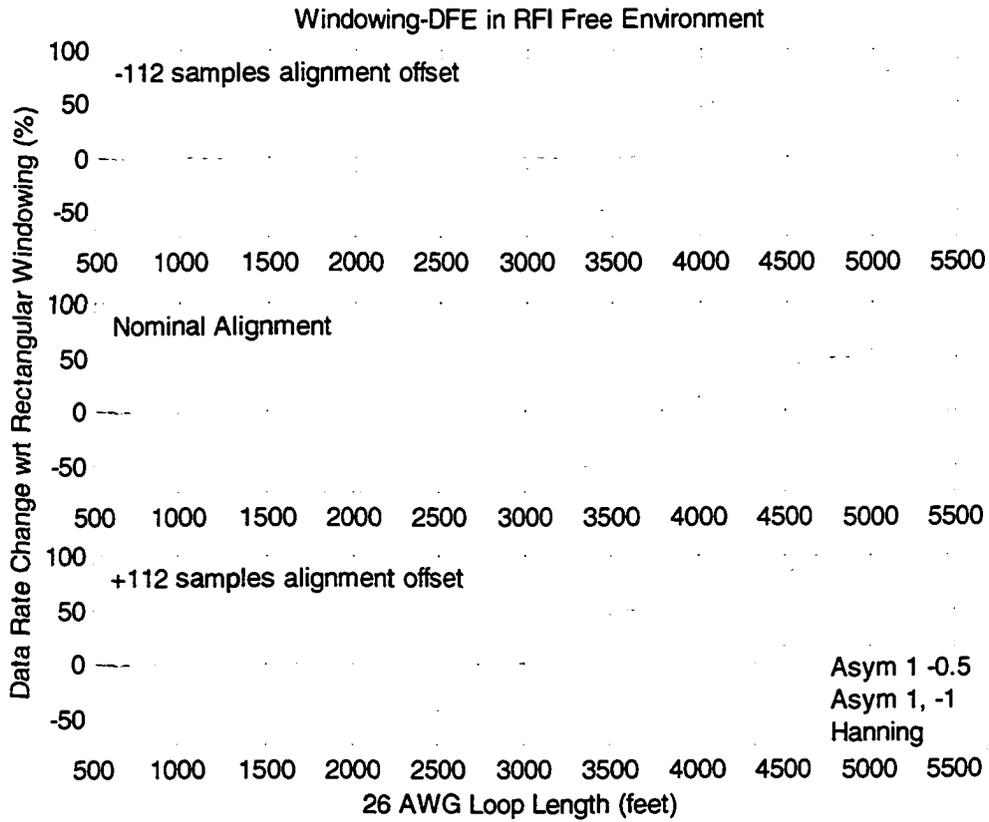


Figure 7-11: Constrained Windowing Impact with Symbol Misalignments

Table 7-7 lists the performance of windowing with a -112 alignment offset (TD-ISI enhancement) relative to nominal performance.

Table 7-7: Windowing Impact on TD-ISI (alignment offset of -112 samples)

VDSL Loop Model	Loop Length (feet)	Nominal Rates with Rectangular Window (kbps)	Relative Rate Change with respect to Nominal Rates						
			Rectangular	R-Cos 8288	R-Cos 8384	R-Cos 8576	Asym 1, -0.5	Asym 1, -1	Hanning
1-x	500	64448	0.0%	0.0%	0.0%	-4.0%	0.0%	-0.1%	-0.1%
1-x	1000	60596	0.0%	-0.1%	0.0%	-0.8%	-1.6%	-5.2%	-14.8%
1-x	1500	45820	0.0%	0.1%	0.1%	-0.2%	-3.0%	-9.1%	-24.2%
1-x	2000	29208	0.2%	0.3%	0.3%	0.1%	-4.6%	-14.1%	-35.8%
1-x	2500	13400	-0.3%	0.1%	0.1%	-0.6%	-8.1%	-22.4%	-43.9%
1-x	3000	7208	-5.2%	-3.3%	-3.2%	-2.9%	-5.9%	-14.8%	-36.1%
1-x	3500	4260	-11.0%	-7.4%	-6.6%	-3.9%	-13.0%	-23.7%	-57.6%
1-x	4000	1376	-23.0%	-13.7%	-9.9%	-0.3%	-10.2%	13.4%	13.4%
1-x	4500	1100	-16.4%	-15.3%	-11.3%	3.3%	4.4%	41.1%	41.5%
1-x	5000	992	-13.7%	-13.3%	-8.9%	8.1%	8.5%	53.6%	53.2%
1-x	5500	832	-5.8%	-3.4%	2.9%	21.6%	21.2%	78.8%	76.4%
1-x	6000	700	5.7%	7.4%	13.1%	33.7%	38.3%	104.6%	103.4%
4	3300	3444	-13.8%	-1.5%	-0.3%	3.4%	-5.8%	-7.5%	-45.2%
5	950	60836	0.0%	0.1%	0.1%	-1.8%	-1.6%	-4.9%	-13.6%
6	3250	6636	-5.3%	-3.7%	-3.2%	-2.5%	-6.2%	-14.9%	-38.1%
7	4900	1120	-15.0%	-13.2%	-9.3%	1.1%	4.6%	38.6%	39.3%

For loops under 2000 feet, rates with rectangular windowing are not affected by a -112 alignment offset and remain optimum, most likely because the TD-ISI is still contained within the shortened cyclic prefix. For loops over 2500 feet, the TD-ISI environment becomes more severe and the raised cosine windowing generally provides an improvement over rectangular windowing. The improvement on long loops is most significant with the 8576 sample raised cosine window due to its lower sidelobes, but this window's considerable utilisation of the cyclic extensions also results in the worst performance on short loops. With the use of this window, not only is nominal performance mostly preserved, but rates actually improve for loops 4500 feet and longer, even when compared to windowed rates with nominal alignment (Table 7-6). This

suggests that in this situation, the window's sidelobe attenuation is effective enough at dealing with TD-ISI to make it worthwhile to bias alignment in order to combat near-end echo, which is not the case for rectangular windowing. The fact that performance is better with the 8576 sample window as opposed to the mid-size 8384 window, which is optimal with nominal alignment, further suggests that it is more important to attenuate sidelobes than to preserve the cyclic extension. This notion is intuitive since the misalignment already reduces the ability for the cyclic prefix to deal with TD-ISI, making the impact of a further cyclic prefix length reduction of lesser consequence, especially for longer loops. What is revealing is the extent to which sidelobe attenuation provided by windowing can suppress interference caused by TD-ISI. This tends to imply that frequency ICI is more of a factor than frequency ISI (see section 2.2.3), which is aligned with the main lobe of the window's transfer function.

A similar reasoning to the above can be applied to understand the constrained windowing results. On longer loops, absolute rates achieved with constrained windowing are practically the same whether a -112 misalignment is present or not, and they tend to be superior to the rates achieved with rectangular or raised cosine windowing. This indicates that the sidelobe attenuation provided by the constrained windows is very effective at handling TD-ISI. Nevertheless, the use of constrained windowing with loops less than 4000 feet still results in degraded rates with respect to performance with rectangular windowing. This is to be expected since TD-ISI on these loops is less serious and more easily contained, making additive noise a more significant impairment with constrained windowing due to the algorithm's susceptibility to it.

Table 7-8 lists the relative performance of windowing with a +112 alignment offset (echo transient enhancement).

Table 7-8: Windowing Impact on Echo (alignment offset of +112 samples)

VDSL Loop Model	Loop Length (feet)	Rates with Nominal Rectangular Window (kbps)	Relative Rate Change with respect to Nominal Rates						
			Rectangular	R-Cos 8288	R-Cos 8384	R-Cos 8576	Asym 1, -0.5	Asym 1, -1	Hanning
1-x	500	64448	0.0%	0.0%	0.0%	0.0%	0.0%	-0.1%	-0.1%
1-x	1000	60596	0.0%	0.1%	0.1%	-0.3%	-1.6%	-5.1%	-14.8%
1-x	1500	45820	0.0%	0.1%	0.1%	-0.1%	-2.9%	-9.0%	-24.1%
1-x	2000	29208	0.0%	0.1%	0.1%	0.0%	-4.4%	-14.1%	-35.9%
1-x	2500	13400	-2.2%	-0.6%	-0.5%	-1.3%	-8.5%	-22.7%	-43.9%
1-x	3000	7208	-16.2%	-1.2%	-2.2%	-4.4%	-9.0%	-15.1%	-35.9%
1-x	3500	4260	-47.8%	-12.3%	-10.6%	-10.9%	-23.1%	-24.1%	-57.2%
1-x	4000	1376	-29.9%	-10.5%	-8.7%	-10.8%	-14.8%	13.4%	13.4%
1-x	4500	1100	-21.8%	-20.0%	-17.8%	-16.0%	-3.3%	41.8%	41.1%
1-x	5000	992	-31.9%	-27.8%	-23.4%	-18.1%	-11.3%	56.9%	53.2%
1-x	5500	832	-34.1%	-29.8%	-22.6%	-15.4%	-9.6%	74.5%	76.4%
1-x	6000	700	-31.4%	-27.4%	-22.3%	-13.7%	-3.4%	89.1%	104.6%
4	3300	3444	-27.2%	5.7%	6.7%	5.1%	-4.9%	-7.5%	-44.7%
5	950	60836	0.0%	0.0%	0.1%	-0.2%	-1.5%	-4.9%	-13.7%
6	3250	6636	-21.9%	-3.7%	-3.7%	-5.5%	-11.1%	-14.9%	-38.2%
7	4900	1120	-20.0%	-17.9%	-16.1%	-13.2%	-1.8%	39.3%	39.3%

The use of windowing with of a +112 alignment offset offers similar benefits to those presented for the case of a -112 alignment offset. Raised cosine windowing provides an improvement over rectangular windowing, except for the case of the 8576 sample window with loops shorter than 2000 feet. While this particular window offers the greatest benefit with loops of 4500 feet or longer, it is the mid-sized 8384 sample window which provides the best performance for loops under 4500 feet. The performance difference between rectangular and raised cosine windowing is most apparent with loops of 3500 feet, which coincides with the greatest performance degradation for rectangular windowing with respect to nominal alignment. Recall that the impulse response of the near-end echo path typically has an initial concentrated pulse (Figure 5-13). It is thus

normal for enhanced echo to suddenly degrade performance at a loop length threshold where the concentrated pulse is no longer contained within the cyclic suffix. At this point, significant echo transient energy spills into the samples to be demodulated, exposing the merit of sidelobe attenuation provided by windowing. This degradation behaviour as a function of loop length is different from the degradation behaviour due to enhanced TD-ISI. In that case, it is the relatively flat tail end of TD-ISI that is no longer contained in the cyclic prefix as loop length is increased, leading to rates that tend to decrease more gradually.

The performance of constrained windowing in the enhanced echo environment is also very similar to the performance behaviour exhibited with a -112 alignment offset. The difference is that the loop length threshold at which constrained windowing offers an advantage over rectangular windowing occurs sooner, and more specifically, as early as 3000 feet for the asymmetric windows.

7.2.3 Windowing Performance in Additive White Gaussian Noise

This section looks at the windowing performance sensitivity to additive wideband noise (i.e. AWGN). Table 7-9 lists the relative performance of windowing in a noise environment lower than nominal, while lists the relative performance of windowing in a noise environment higher than nominal.

Table 7-9: Windowing Performance with -150 dBm/Hz AWGN

VDSL Loop Model	Loop Length (feet)	Rates with Nominal Rectangular Window (kbps)	Relative Rate Change with respect to Nominal Rates						
			Rectangular	R-Cos 8288	R-Cos 8384	R-Cos 8576	Asym 1, -0.5	Asym 1, -1	Hanning
1-x	500	64448	0.0%	0.0%	0.0%	0.0%	0.0%	-0.1%	-0.1%
1-x	1000	60596	6.4%	6.4%	6.4%	6.4%	6.3%	6.3%	3.6%
1-x	1500	45820	25.4%	25.4%	25.4%	25.4%	22.9%	18.3%	5.7%
1-x	2000	29208	47.5%	47.7%	47.7%	47.8%	42.1%	32.7%	8.4%
1-x	2500	13400	99.3%	99.7%	99.8%	100.0%	88.5%	68.8%	17.0%
1-x	3000	7208	69.1%	70.5%	70.5%	70.0%	57.5%	36.7%	10.5%
1-x	3500	4260	71.8%	77.7%	77.9%	74.9%	70.3%	55.0%	17.9%
1-x	4000	1376	190.1%	222.7%	220.9%	206.7%	210.5%	169.8%	55.2%
1-x	4500	1100	0.0%	39.6%	41.8%	41.8%	25.5%	44.7%	41.8%
1-x	5000	992	0.0%	1.2%	4.0%	2.0%	21.8%	57.3%	57.3%
1-x	5500	832	0.0%	7.2%	8.7%	6.7%	25.0%	86.1%	87.5%
1-x	6000	700	0.0%	4.0%	11.4%	10.3%	29.7%	117.7%	122.9%
4	3300	3444	77.2%	103.9%	105.3%	106.5%	92.9%	86.8%	37.2%
5	950	60836	5.8%	5.8%	5.8%	5.8%	5.6%	5.1%	2.3%
6	3250	6636	62.0%	63.7%	64.2%	63.1%	54.4%	37.6%	11.6%
7	4900	1120	0.0%	48.6%	52.9%	55.0%	21.1%	40.7%	39.3%

Table 7-10: Windowing Performance with -130 dBm/Hz AWGN

VDSL Loop Model	Loop Length (feet)	Rates with Nominal Rectangular Window (kbps)	Relative Rate Change with respect to Nominal Rates						
			Rectangular	R-Cos 8288	R-Cos 8384	R-Cos 8576	Asym 1, -0.5	Asym 1, -1	Hanning
1-x	500	64448	-0.9%	-0.9%	-0.9%	-0.8%	-2.4%	-5.6%	-15.5%
1-x	1000	60596	-19.7%	-19.7%	-19.6%	-19.6%	-21.9%	-26.8%	-37.7%
1-x	1500	45820	-30.1%	-30.0%	-30.0%	-29.9%	-33.2%	-39.3%	-53.5%
1-x	2000	29208	-45.8%	-45.7%	-45.7%	-45.6%	-51.1%	-60.0%	-75.7%
1-x	2500	13400	-49.8%	-49.8%	-49.8%	-49.8%	-52.0%	-57.0%	-69.0%
1-x	3000	7208	-45.4%	-45.6%	-45.4%	-45.2%	-51.5%	-63.9%	-80.4%
1-x	3500	4260	-66.9%	-66.8%	-66.5%	-66.9%	-64.4%	-64.6%	-68.4%
1-x	4000	1376	-10.8%	-10.2%	-8.7%	-8.1%	-0.6%	6.1%	-5.2%
1-x	4500	1100	-0.4%	0.0%	2.9%	1.8%	14.9%	27.6%	13.1%
1-x	5000	992	-0.4%	0.8%	2.4%	1.2%	18.5%	35.1%	21.4%
1-x	5500	832	0.0%	6.7%	8.7%	6.7%	25.0%	55.3%	36.5%
1-x	6000	700	-1.1%	4.0%	9.1%	9.1%	27.4%	74.9%	57.7%
4	3300	3444	-66.8%	-66.4%	-65.5%	-62.8%	-61.2%	-56.0%	-60.6%
5	950	60836	-18.7%	-18.6%	-18.6%	-18.6%	-20.9%	-25.6%	-36.6%
6	3250	6636	-48.6%	-48.5%	-48.5%	-48.2%	-55.6%	-72.4%	-78.5%
7	4900	1120	-1.1%	0.7%	2.1%	3.9%	16.4%	30.0%	16.4%

Figure 7-12 depicts the performance of raised cosine windowing relative to rectangular windowing with different levels of AWGN. At first glance, it appears that the benefit of raised cosine windowing is greatest in the lowest additive noise environment. However, the performance gain is in reality due to the mitigation of TD-ISI and echo transients. By lowering the level of AWGN, the performance limitation due to noise is made relatively less important than the limitation imposed by TD-ISI and echo transients. This in turn puts emphasis on the ability of windowing to reduce the impact of TD-ISI and echo transients on long loops, and also compounds the effect of SNR estimate errors. Without this mechanism in play, the performance enhancement provided by raised cosine windowing should not depend on the level of AWGN.

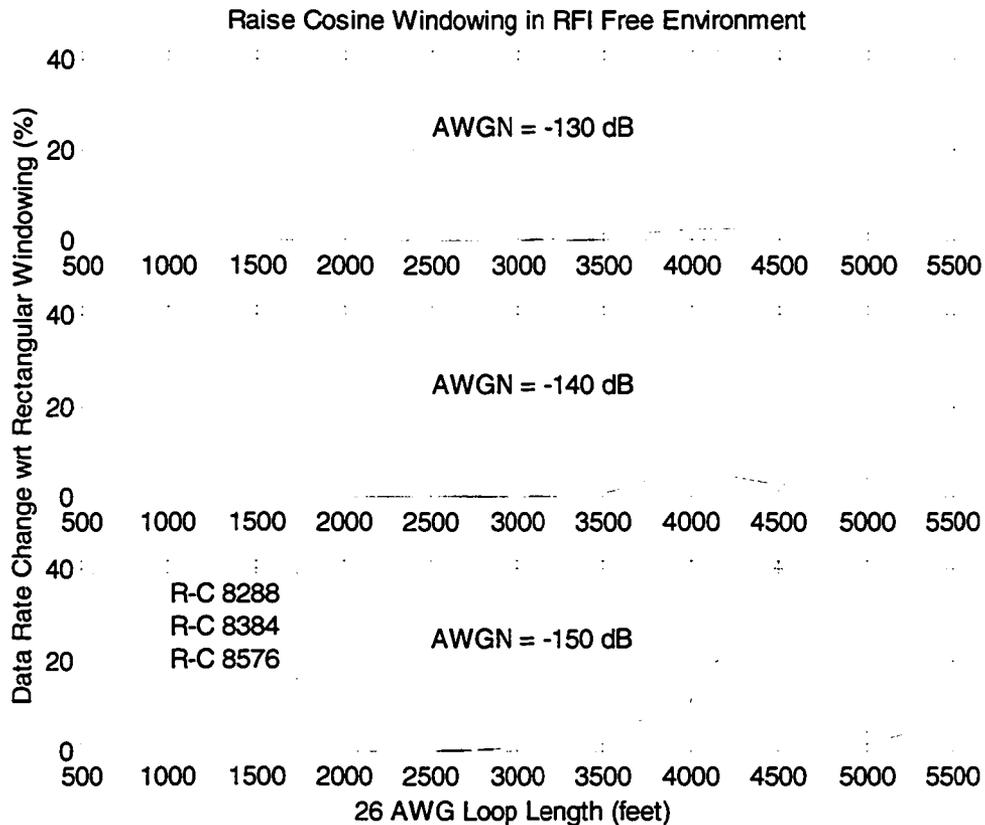


Figure 7-12: Raised Cosine Windowing Impact in Different Noise Environment

Figure 7-13 depicts the performance of constrained windowing relative to rectangular windowing with different levels of AWGN. As expected, relative performance with constrained windowing is generally inversely proportional to the level of AWGN. What is interesting is that the loop length threshold at which constrained windowing provides a performance gain instead of a loss tends to decrease with higher levels of AWGN. This suggests that for bins still capable of carrying data in higher noise environments, the benefit of TD-ISI and echo transient mitigation outweighs the drawback related to the algorithm's susceptibility to AWGN, for midsize loops specifically.

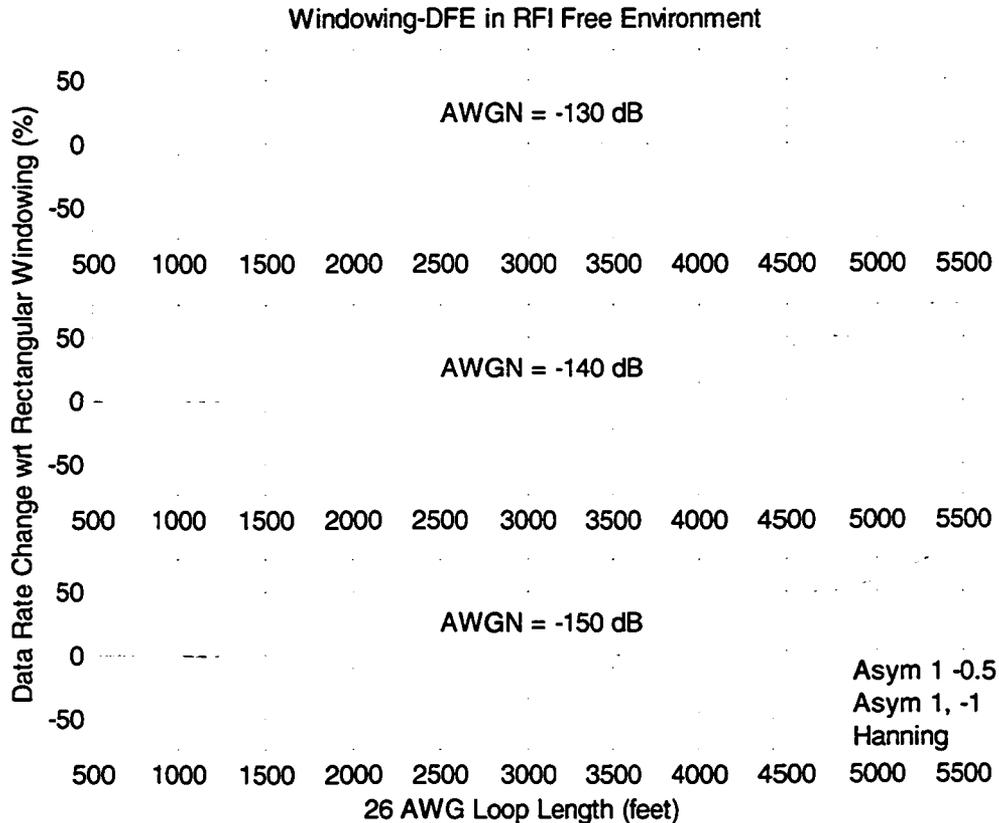


Figure 7-13: Constrained Windowing Impact in Different Noise Environments

7.2.4 Ability of Windowing to Combat RFI

The ability of windowing to combat ingress RFI is investigated in this section. Table 7-11 lists the various windowing performances in the presence of a -10 dBm RFI source centred on bin 870.5 relative to nominal rates in an RFI free environment. Figure 7-14 shows the rates achieved with raised cosine windowing and the RFI present, along with the corresponding windowing performances relative to rectangular windowing. Figure 7-15 depicts the same information for constrained windowing with an added reference to raised cosine windowing with the 8576 sample version.

The raised cosine windowing method, with the 8576 sample version specifically, is the most effective at achieving near-nominal performance in the presence of an RFI at bin 870.5. Using the 8384 sample raised cosine window provides the next best result. In

terms of purely combating the RFI, the constrained windowing method using the Hanning window, and the Asymmetric{1,-1} to a lesser extent, offer the best performance. In fact, performance with the Hanning window is the same whether the RFI at bin 870.5 is present or not. Unfortunately, this method suffers so much from its susceptibility to AWGN, that its performance with loops less than 4000 feet is still inferior to that of raised cosine windowing, even with the RFI present. As for performance with the 8288 raised cosine and the asymmetric {1, -0.5} windows, they are really not up to the task of mitigating the effect of RFI due to their relatively high sidelobes.

Table 7-11: Rates after Windowing with -10 dBm RFI at Bin 870.5

VDSL Loop Model	Loop Length (feet)	Rates with Nominal Rectangular Window (kbps)	Relative Rate Change due to RFI Located at Bin 870.5						
			Rectangular	R-Cos 8288	R-Cos 8384	R-Cos 8576	Asym 1, -0.5	Asym 1, -1	Hanning
1-x	500	64448	-25.4%	-1.3%	-0.1%	0.0%	-12.4%	-0.1%	-0.1%
1-x	1000	60596	-47.0%	-3.0%	-0.7%	0.0%	-32.9%	-5.2%	-14.8%
1-x	1500	45820	-64.9%	-6.6%	-2.3%	-0.2%	-47.7%	-9.5%	-24.1%
1-x	2000	29208	-95.5%	-10.4%	-3.5%	-0.5%	-77.1%	-14.9%	-35.8%
1-x	2500	13400	-90.5%	-21.0%	-7.7%	-1.2%	-88.9%	-24.3%	-43.9%
1-x	3000	7208	-83.2%	-33.0%	-15.0%	-2.7%	-80.4%	-17.3%	-36.1%
1-x	3500	4260	-73.1%	-43.0%	-19.5%	-3.7%	-68.4%	-28.2%	-57.4%
1-x	4000	1376	-24.4%	-9.3%	-6.1%	-5.8%	-9.6%	13.4%	13.4%
1-x	4500	1100	-14.9%	1.5%	2.9%	2.5%	4.0%	41.8%	41.1%
1-x	5000	992	-5.6%	0.8%	4.0%	2.0%	15.3%	57.3%	53.2%
1-x	5500	832	-4.3%	7.7%	8.7%	6.7%	20.7%	84.6%	76.4%
1-x	6000	700	-10.9%	4.0%	10.9%	10.3%	20.6%	110.9%	105.1%
4	3300	3444	-69.8%	-52.1%	-18.6%	2.4%	-63.8%	-14.9%	-44.9%
5	950	60836	-46.2%	-2.9%	-0.8%	0.1%	-32.7%	-4.9%	-13.7%
6	3250	6636	-82.0%	-33.6%	-14.5%	-2.8%	-79.0%	-17.5%	-38.1%
7	4900	1120	-16.4%	1.1%	2.5%	5.0%	2.1%	39.3%	39.3%

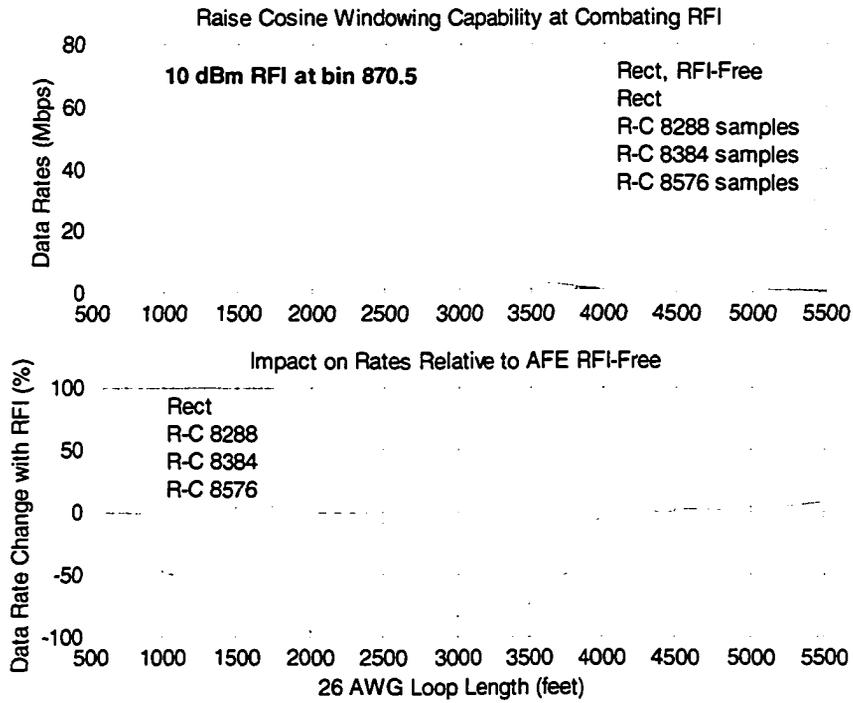


Figure 7-14: Raised Cosine Windowing Impact with RFI at bin 870.5

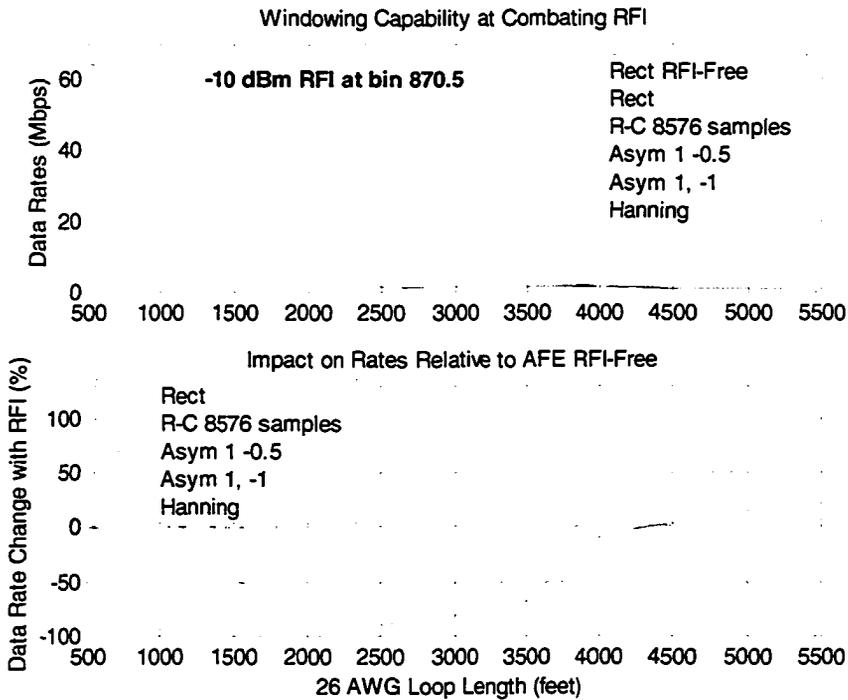


Figure 7-15: Windowing Impact with RFI at bin 870.5

Simulations were repeated for the case of two -10 dBm RFI sources, centred at bin 870.5 and bin 2343.5. Table 7-12 lists the corresponding windowing performances relative to nominal. Achieved rates and performances relative to rectangular windowing are shown in Figure 7-16 and Figure 7-17 for raised cosine windowing and constrained windowing respectively.

In this severe RFI environment, constrained asymmetric $\{1, -1\}$ windowing provides the best performance, although rates are still degraded by up to the 33.5% with the worst case loop. The next best performance is achieved with the 8576 sample raised cosine windowing, followed closely by 8384 sample raised cosine windowing.

Table 7-12: Rates after Windowing with -10 dBm RFIs at Bins 870.5 and 2343.5

VDSL Loop Model	Loop Length (feet)	Rates with Nominal Rectangular Window (kbps)	Relative Rate Change due to two RFIs Located at Bins 870.5 and 2343.5						
			Rectangular	R-Cos 8288	R-Cos 8384	R-Cos 8576	Asym 1, -0.5	Asym 1, -1	Hanning
1-x	500	64448	-50.9%	-12.1%	-5.9%	-2.4%	-37.7%	-1.2%	-0.2%
1-x	1000	60596	-74.8%	-21.5%	-10.9%	-6.5%	-61.8%	-8.4%	-14.9%
1-x	1500	45820	-93.0%	-27.9%	-14.1%	-14.8%	-80.8%	-14.0%	-24.3%
1-x	2000	29208	-96.1%	-32.9%	-17.6%	-30.5%	-91.5%	-20.6%	-36.1%
1-x	2500	13400	-92.0%	-37.3%	-25.6%	-31.0%	-90.5%	-26.1%	-43.9%
1-x	3000	7208	-86.0%	-42.2%	-25.9%	-19.7%	-83.1%	-20.2%	-36.1%
1-x	3500	4260	-77.2%	-63.9%	-57.8%	-46.9%	-72.2%	-33.5%	-57.6%
1-x	4000	1376	-28.5%	-9.6%	-8.4%	-7.8%	-13.1%	13.4%	13.4%
1-x	4500	1100	-16.7%	0.7%	2.9%	1.5%	1.8%	41.8%	41.1%
1-x	5000	992	-17.7%	1.2%	4.0%	1.6%	2.8%	57.3%	53.2%
1-x	5500	832	-17.3%	7.7%	9.1%	6.7%	7.7%	84.6%	76.4%
1-x	6000	700	-8.6%	4.0%	11.4%	9.7%	20.6%	110.3%	105.1%
4	3300	3444	-72.0%	-59.6%	-43.7%	-31.4%	-66.0%	-20.3%	-44.9%
5	950	60836	-72.3%	-20.2%	-10.2%	-5.4%	-60.2%	-8.2%	-13.8%
6	3250	6636	-82.4%	-43.6%	-28.9%	-23.4%	-79.3%	-20.3%	-38.1%
7	4900	1120	-16.4%	1.1%	2.1%	4.3%	2.1%	39.3%	39.3%

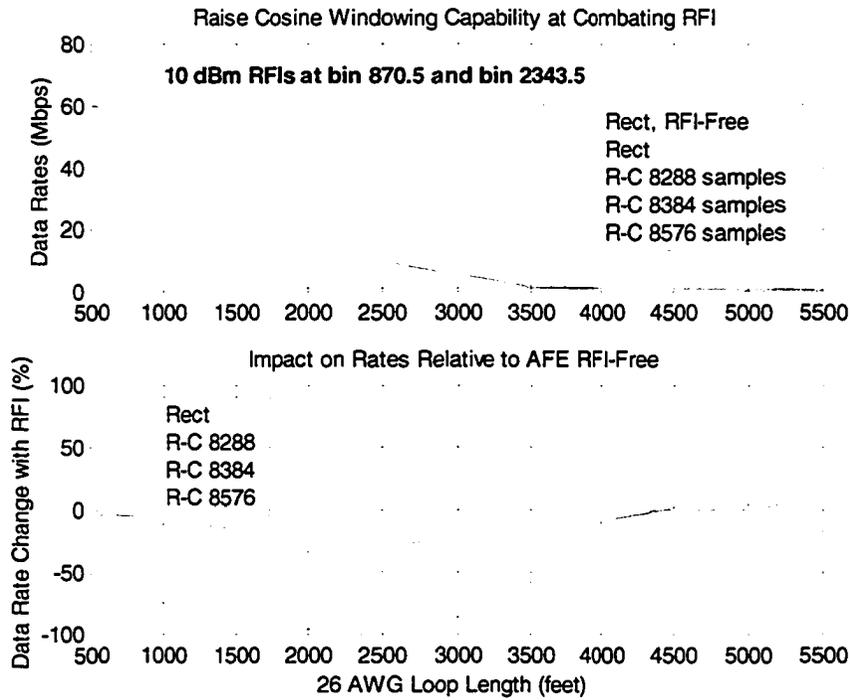


Figure 7-16: Raised Cosine Windowing Impact with RFIs at bins 870.5 and 2343.5

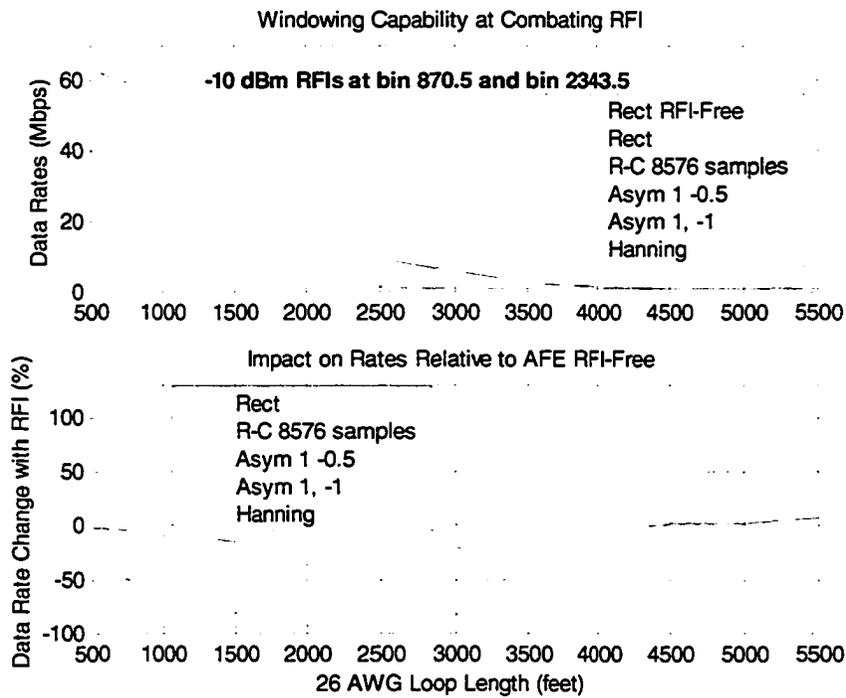


Figure 7-17: Windowing Impact with RFIs at bins 870.5 and 2343.5

Figure 7-18 and Figure 7-19 depict the achieved rates and relative performance with raised cosine and constrained windowing respectively in the presence of a single -10 dBm RFI centred at bin 2343.5 instead of 870.5. As for the case of a single RFI centred at bin 870.5, the 8576 sample raised cosine windowing algorithm is the most effective at achieving near-nominal performance in the presence of an RFI at bin 2343.5.

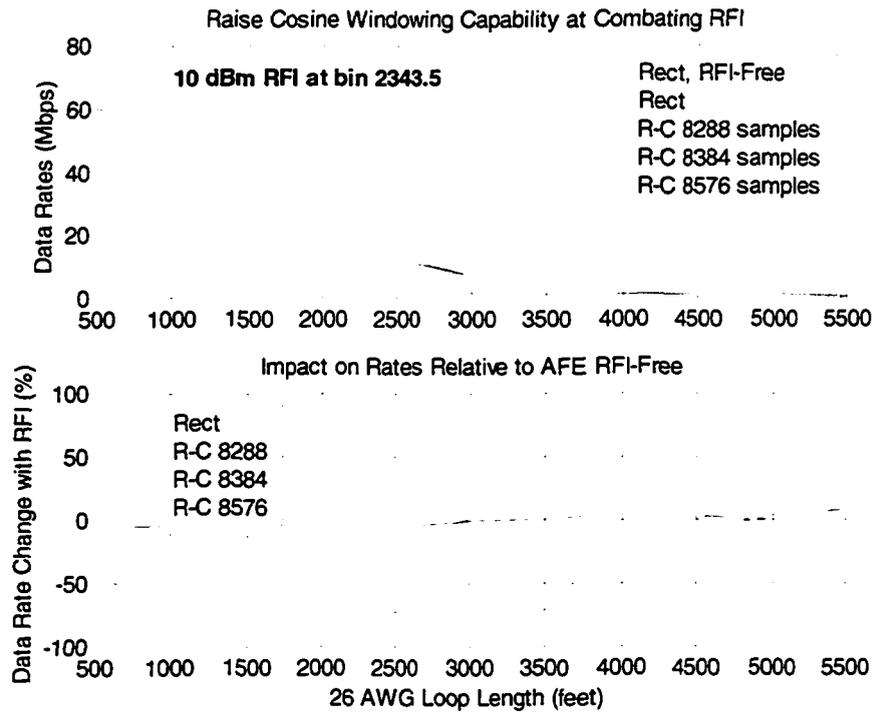


Figure 7-18: Raised Cosine Windowing Impact with RFI at bin 2343.5

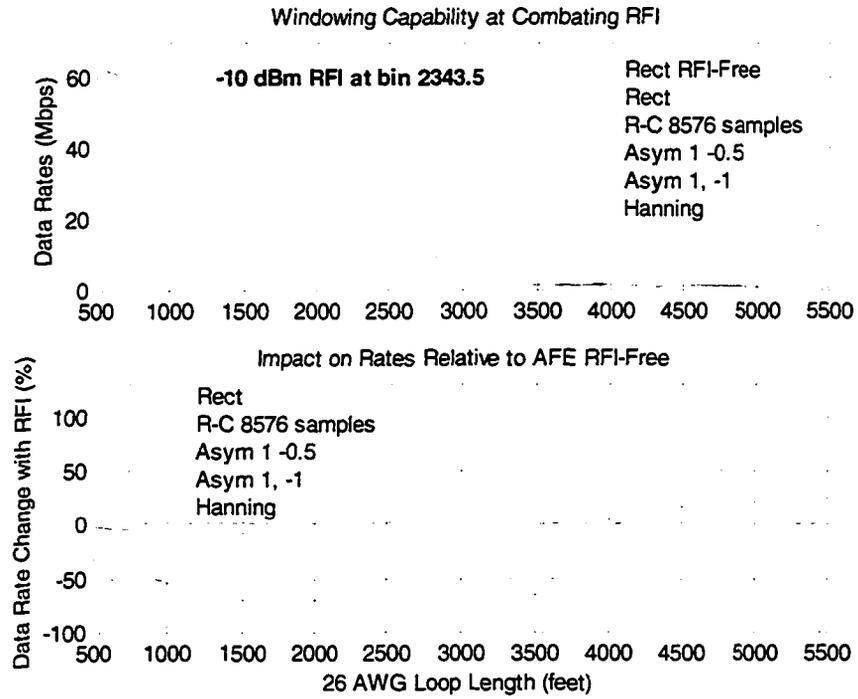


Figure 7-19: Windowing Impact with RFI at bin 2343.5

7.3 Notch Filtering Performance Analysis

Simulations were performed to explore the possibility of using notch filters to combat ingress RFI in VDSL DMT systems. First, the effect of a single notch filter on system performance in an RFI-free environment was evaluated. Table 7-13 lists the relative performance degradation for notch filtering with various pole contraction factors (i.e. r). The filter's negative impact on RFI-free rates increases as the pole contraction factor is decreased, although the effect is insignificant for loops shorter than 3000 feet.

Table 7-14 shows the resulting rate impact when two notch filters are implemented, representing the case of two RFI sources present simultaneously. Relative performance plots with one and two notch filters are depicted in Figure 7-20.

Table 7-13: Single Notch Filter Impact on Rates in RFI-Free Environment

VDSL Loop Model	Loop Length (feet)	Nominal Rates (kbps)	Relative Deltas in RFI-Free Environment with Notch Filter at Bin 870.5 (%)								
			$r = 0.99$	$r = 0.97$	$r = 0.95$	$r = 0.93$	$r = 0.91$	$r = 0.89$	$r = 0.87$	$r = 0.85$	
1-x	500	64448	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
1-x	1000	60596	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
1-x	1500	45820	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
1-x	2000	29208	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
1-x	2500	13400	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
1-x	3000	7208	0.0%	-0.1%	-0.2%	-0.2%	-0.2%	-0.2%	-0.2%	-0.2%	-0.3%
1-x	3500	4260	-0.1%	0.0%	-0.4%	-0.5%	-0.6%	-0.7%	-0.9%	-1.0%	-1.0%
1-x	4000	1376	-0.6%	-0.6%	-1.7%	-2.9%	-3.5%	-4.1%	-5.8%	-7.0%	-7.0%
1-x	4500	1100	-0.4%	-0.4%	-0.4%	-0.4%	-0.4%	-0.4%	-0.4%	-0.4%	-0.4%
1-x	5000	992	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
1-x	5500	832	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
1-x	6000	700	0.0%	0.0%	0.0%	0.0%	0.0%	0.6%	0.6%	0.6%	0.6%
4	3300	3444	-0.3%	-0.5%	-1.0%	-2.3%	-2.8%	-4.2%	-4.5%	-5.3%	-5.3%
5	950	60836	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
6	3250	6636	0.0%	0.0%	0.0%	0.0%	0.1%	0.0%	-0.1%	-0.1%	-0.1%
7	4900	1120	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%

Table 7-14: Dual Notch Filter Impact on Rates in RFI-Free Environment

VDSL Loop Model	Loop Length (feet)	Nominal Rates (kbps)	Relative Deltas in RFI-Free Environment With Notch Filters at Bins 870.5 and 2646.5 (%)								
			r = 0.99	r = 0.97	r = 0.95	r = 0.93	r = 0.91	r = 0.89	r = 0.87	r = 0.85	
1-x	500	64448	-0.1%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
1-x	1000	60596	-0.2%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	-0.1%
1-x	1500	45820	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	-0.1%	0.0%
1-x	2000	29208	0.0%	0.0%	0.0%	0.0%	0.0%	-0.1%	-0.1%	-0.1%	-0.1%
1-x	2500	13400	0.0%	0.1%	0.0%	-0.1%	-0.3%	-0.4%	-0.4%	-0.3%	-0.4%
1-x	3000	7208	0.1%	-0.1%	-0.2%	-0.2%	-0.2%	-0.3%	-0.3%	-0.3%	-0.2%
1-x	3500	4260	0.0%	-0.1%	-0.1%	-0.6%	-0.7%	-0.8%	-0.8%	-0.8%	-1.2%
1-x	4000	1376	-0.6%	-1.2%	-1.7%	-2.3%	-4.1%	-5.2%	-6.4%	-7.6%	-7.6%
1-x	4500	1100	-0.4%	-0.4%	-0.4%	-0.4%	-0.4%	-0.4%	-0.4%	-0.4%	-0.4%
1-x	5000	992	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
1-x	5500	832	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
1-x	6000	700	0.0%	0.0%	0.0%	0.0%	0.0%	0.6%	0.6%	0.6%	0.6%
4	3300	3444	0.1%	-0.5%	-1.0%	-2.4%	-2.8%	-4.3%	-4.6%	-5.1%	-5.1%
5	950	60836	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%
6	3250	6636	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	-0.1%	-0.1%	-0.1%
7	4900	1120	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%

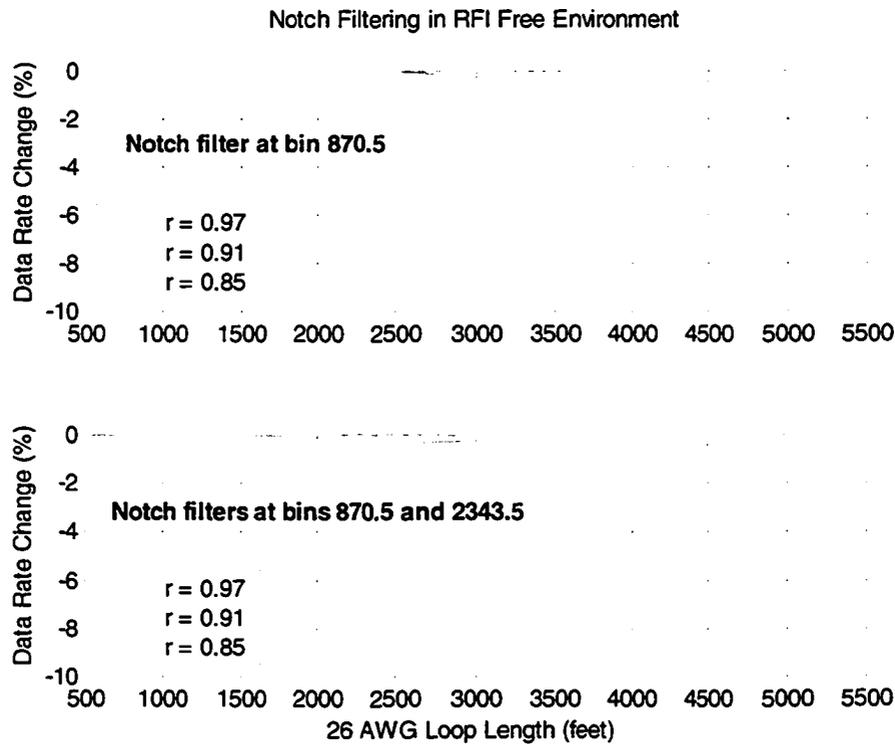


Figure 7-20: Notch Filtering Impact in RFI-Free Environment

Table 7-15 lists the notch filtering performance achieved in the presence of a -10 dBm RFI source centred on bin 870.5. The stated performances are relative to nominal rates in an RFI free environment and correspond to filters with various pole contraction factors and either of two centre (notch) frequencies. The first centre frequency is matched to the RFI source while the second is offset by 2 kHz to approximate a worst case RFI frequency tracking error. In general, decreasing a filter's pole contraction factor will allow it to better handle frequency tracking errors, but decreasing it too much will result in overall performance degradation. The simulation results show that notch filtering offers RFI suppression performance similar or superior to the windowing performances in the same conditions (section 7.2.4), especially with shorter loops (< 4000 feet). Figure 7-21 compares the relative performance of 8576 sample raised cosine windowing and notch filtering with the pole contraction factor set to 0.85, in the presence of one RFI.

Table 7-15: Rates after Notch Filtering -10 dBm RFI at Bin 870.5

VDSL Loop Model	Loop Length (feet)	Nominal Rates (kbps)	Rate Deltas with one RFI at Bin 870.5 Relative to Nominal RFI-Free Rates (%)									
			No filter	r = 0.91	2 kHz offset	r = 0.89	2 kHz offset	r = 0.87	2 kHz offset	r = 0.85	2 kHz offset	
1-x	500	64448	-25%	0%	0%	0%	0%	0%	0%	0%	0%	0%
1-x	1000	60596	-47%	0%	0%	0%	0%	0%	0%	0%	0%	0%
1-x	1500	45820	-65%	0%	0%	0%	0%	0%	0%	0%	0%	0%
1-x	2000	29208	-96%	0%	-1%	0%	-1%	0%	-1%	0%	0%	0%
1-x	2500	13400	-90%	-1%	-1%	-1%	-1%	-1%	-1%	-1%	-1%	-1%
1-x	3000	7208	-83%	-2%	-2%	-1%	-2%	-1%	-2%	-1%	-2%	-2%
1-x	3500	4260	-73%	-3%	-5%	-3%	-4%	-3%	-4%	-3%	-4%	-4%
1-x	4000	1376	-24%	-8%	-8%	-8%	-8%	-8%	-8%	-9%	-9%	-9%
1-x	4500	1100	-15%	0%	0%	0%	0%	0%	0%	0%	0%	0%
1-x	5000	992	-6%	0%	0%	0%	0%	0%	0%	0%	0%	0%
1-x	5500	832	4%	0%	0%	0%	0%	0%	0%	0%	0%	0%
1-x	6000	700	-11%	0%	0%	1%	1%	1%	1%	1%	1%	1%
4	3300	3444	-70%	-4%	-5%	-5%	-5%	-5%	-6%	-6%	-7%	-7%
5	950	60836	-46%	0%	0%	0%	0%	0%	0%	0%	0%	0%
6	3250	6636	-82%	-2%	-3%	-2%	-2%	-2%	-2%	-2%	-2%	-2%
7	4900	1120	-16%	0%	0%	0%	0%	0%	0%	0%	0%	0%

Table 7-16 lists the notch filtering performance achieved in the presence of two -10 dBm RFI sources centred on bins 870.5 and 2343.5. In this severe RFI environment, notch filters with a smaller pole contraction factor tend to provide the best overall performance. Notch filtering also seems to outperform windowing in this scenario. This can be observed in Figure 7-22, which presents the relative performance plots of the optimum windowing and notch filtering algorithms in this scenario.

Table 7-16: Rates after Notch Filtering -10 dBm RFIs at Bins 870.5 and 2343.5

VDSL Loop Model	Loop Length (feet)	Nominal Rates (kbps)	Rates Deltas with RFIs at Bins 870.5 and 2343.5								
			Relative to Nominal RFI-Free Rates (%)								
			No filter	r = 0.91	2 kHz offset	r = 0.89	2 kHz offset	r = 0.87	2 kHz offset	r = 0.85	2 kHz offset
1-x	500	64448	-51%	-1%	-1%	-1%	-1%	-1%	-1%	-1%	-1%
1-x	1000	60596	-75%	-4%	-5%	-4%	-5%	-4%	-5%	-4%	-5%
1-x	1500	45820	-93%	-6%	-8%	-6%	-8%	-6%	-7%	-5%	-7%
1-x	2000	29208	-96%	-8%	-11%	-8%	-10%	-8%	-10%	-7%	-10%
1-x	2500	13400	-92%	-9%	-12%	-8%	-11%	-8%	-10%	-8%	-10%
1-x	3000	7208	-86%	-6%	-10%	-5%	-8%	-5%	-8%	-5%	-7%
1-x	3500	4260	-77%	-10%	-15%	-8%	-12%	-8%	-11%	-7%	-11%
1-x	4000	1376	-28%	-11%	-11%	-11%	-11%	-11%	-11%	-11%	-11%
1-x	4500	1100	-17%	0%	0%	0%	0%	0%	0%	0%	0%
1-x	5000	992	-18%	0%	0%	0%	0%	0%	0%	0%	0%
1-x	5500	832	-17%	0%	0%	0%	0%	0%	0%	0%	0%
1-x	6000	700	-9%	-1%	-1%	-1%	-1%	0%	0%	0%	0%
4	3300	3444	-72%	-12%	-14%	-11%	-13%	-11%	-12%	-12%	-12%
5	950	60836	-72%	-4%	-5%	-4%	-5%	-4%	-5%	-4%	-5%
6	3250	6636	-82%	-7%	-9%	-6%	-8%	-6%	-7%	-5%	-7%
7	4900	1120	-16%	0%	0%	0%	0%	0%	0%	0%	0%

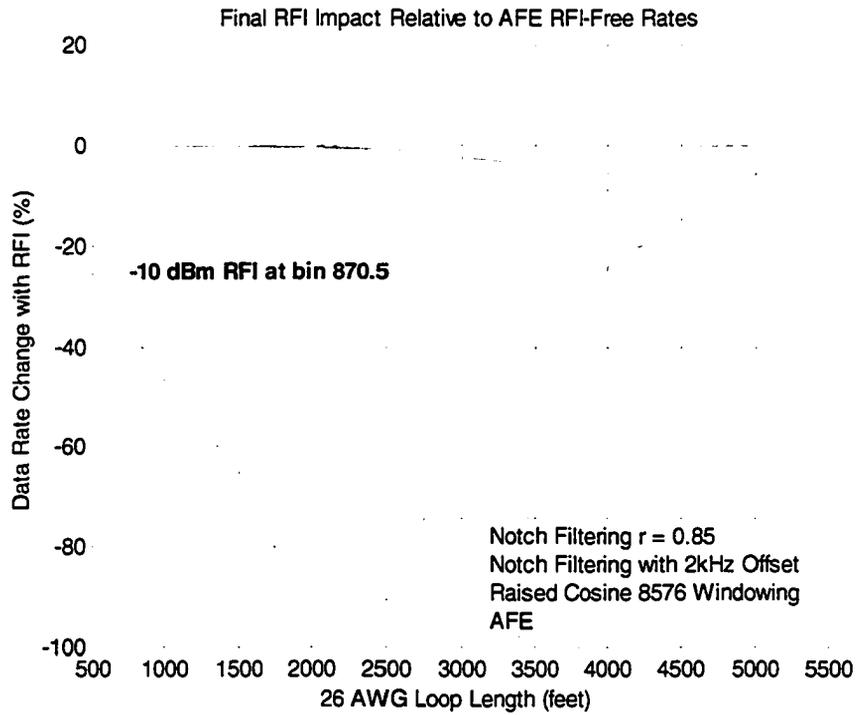


Figure 7-21: Notch Filter and R-C Windowing Impact with Single RFI

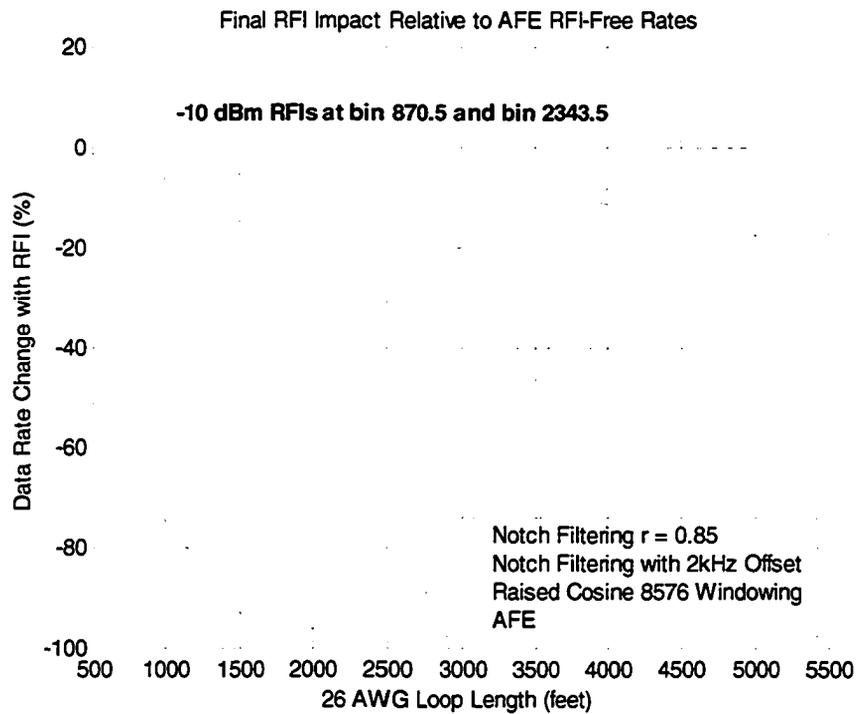


Figure 7-22: Notch Filter and R-C Windowing Impact with Dual RFI

7.4 RFI Cancellation Performance Analysis

Simulations were performed to investigate the RFI Cancellation algorithm presented in [29]. To simplify analysis, the true RFI location was fed to the cancellation algorithm instead of using an estimate as a real system would. Results are shown in Table 7-17 for the case of the standard RFI located at bin 870.5. In this instance, RFI cancellation actually deteriorates rates compared to nominal rates achieved with the RFI present.

Figure 7-23 depicts the spectrums of the RFI signal, RFI model, and RFI residue, along with the upstream signal spectrum received over a 2000-foot loop.

Table 7-17: RFI Cancellation Rates with 4 kHz RFI at bin 870.5

VDSL Loop Model	Loop Length (feet)	Nominal Rates with Rectangular Window (kbps)	Rates with 4 kHz wide RFI at bin 870.5	
			With Nominal Rectangular Window (kbps)	With RFI Cancellation (kbps)
1-x	500	64448	48052	44496
1-x	1000	60596	32136	27624
1-x	1500	45820	16060	12124
1-x	2000	29208	1300	1216
1-x	2500	13400	1276	1184
1-x	3000	7208	1212	1112
1-x	3500	4260	1144	1068
1-x	4000	1376	1040	956
1-x	4500	1100	936	876
1-x	5000	992	936	832
1-x	5500	832	796	728
1-x	6000	700	624	632
4	3300	3444	1040	936
5	950	60836	32712	28096
6	3250	6636	1192	1152
7	4900	1120	936	936

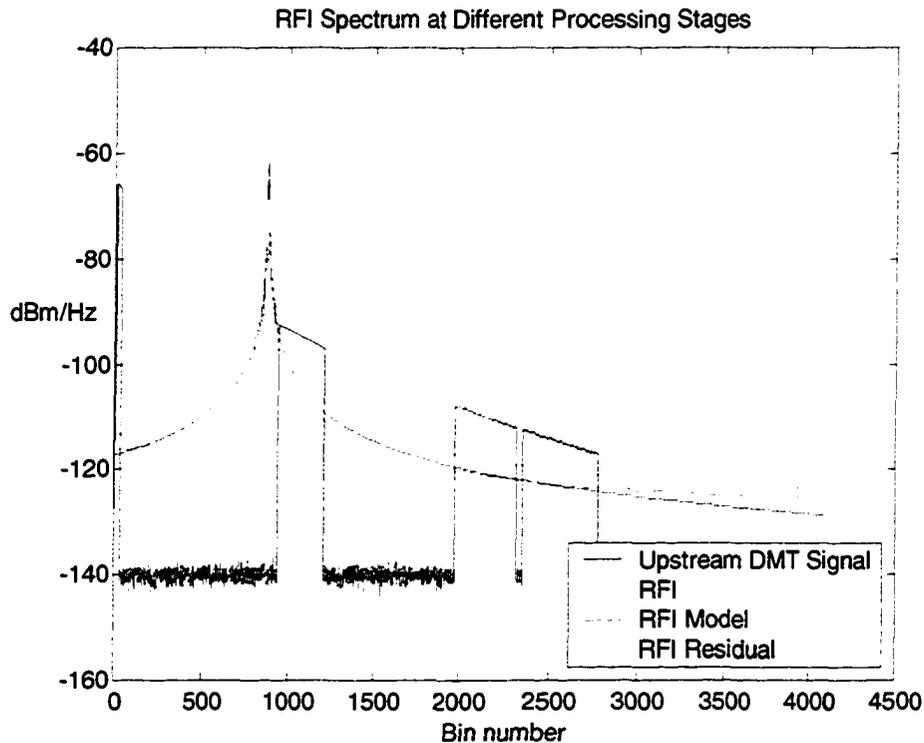


Figure 7-23: RFI Cancellation of 4 kHz RFI at bin 870.5 with 2000 foot Loop

Next, the nature and bandwidth of the RFI were changed in an attempt to prove the algorithms validity under the condition that the RFI bandwidth is much less than the 4 kHz symbol transmission frequency. Specifically, the RFI carrier was modulated with a combination of two sine waves of frequencies 300 Hz and 400 Hz. This RFI scenario was simulated for two receiver cases: nominal receiver and receiver with RFI cancellation. Results are presented in Table 7-18. Figure 7-24 depicts the spectrums of the RFI signal, RFI model, and RFI residue along with the upstream signal spectrum received over a 2000-foot loop. Simulation results show that RFI cancellation does indeed suppress a narrowband RFI, but not to the degree of previous methods such as raised cosine windowing. This is not surprising as [29] states that RFI cancellation in conjunction with receiver windowing is required to achieve satisfactory performance.

Table 7-18: RFI Cancellation Rates with 400 Hz RFI at bin 870.5

VDSL Loop Model	Loop Length (feet)	Nominal Rates with Rectangular Window (kbps)	Rates with 400 Hz wide RFI at bin 870.5	
			With Nominal Rectangular Window (kbps)	With RFI Cancellation (kbps)
1-x	500	64448	20892	30604
1-x	1000	60596	6976	14212
1-x	1500	45820	632	2880
1-x	2000	29208	632	736
1-x	2500	13400	564	648
1-x	3000	7208	520	612
1-x	3500	4260	456	552
1-x	4000	1376	348	416
1-x	4500	1100	328	416
1-x	5000	992	320	340
1-x	5500	832	292	264
1-x	6000	700	204	156
4	3300	3444	472	564
5	950	60836	6408	14024
6	3250	6636	532	624
7	4900	1120	576	452

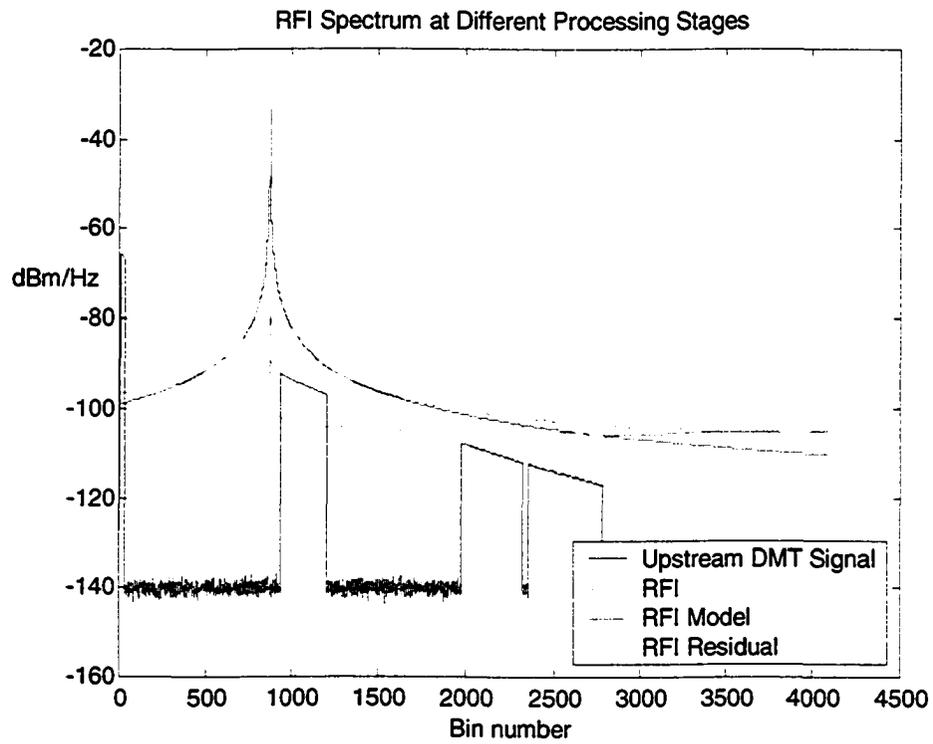


Figure 7-24: RFI Cancellation of 400 Hz RFI at bin 870.5 with 1000-foot loop.

7.5 Summary

In terms of combating two severe RFI sources, notch filtering generally outperforms windowing as is observed in Figure 7-22. Specifically, notch filters with smaller pole contraction factors tend to provide the best overall performance, with a worst case rate degradation of 12% for the filter with $r = 0.85$. Of the windowing approaches, constrained asymmetric{1,-1} windowing provide the best performance in this severe RFI environment, with a worst case rate degradation of 33.5%, followed by raised cosine windowing (8576 sample version) with a worst case rate degradation of 47%.

With only one RFI source present, the worst case rate degradation for each of, notch filtering ($r = 0.85$), raised cosine windowing (8576 sample), and constrained windowing (Asymmetric{1,-1}), are respectively 9%, 6%, and 28%. Thus, with RFI reduced, it is the raised cosine windowing algorithm that provides the best performance, followed closely by notch filtering. In fact, simulation results show that notch filtering offers RFI suppression performance similar or superior to raised cosine windowing on shorter loops (< 4000 feet). As for using the constrained windowing algorithm in this less severe RFI environment, the additive noise enhancement associated with it outweighs the superior RFI suppression achieved with the asymmetric window, except for the cases of very long loops.

It can be argued that even the single RFI environment simulated is rather harsh since the impairment was generated with the maximum expected power (-10 dBm). Given a less severe environment, the advantage of raised cosine windowing over constrained windowing should increase. The difference in performance between raised cosine windowing and notch filtering based on simulation results is however harder to predict for this scenario. One reason for this is that as the RFI power is reduced, so is the required level of RFI suppression, meaning that the notch filter's pole contraction factor could be increased. This would in turn reduce the negative aspects of notch filtering (i.e. attenuation of active bins and contribution to ISI). Further investigation on this matter

seems warranted to establish the optimum pole contraction factors for various RFI scenarios. It should also be mentioned that the performance results pertaining to the constrained windowing technique are based on ideal decision feedback equalization with no error present at the output of the slicer. In reality, some error at the slicer appears from time to time and leads to error propagation in the DFE. This topic is further discussed in [14].

Windowing also offers benefits in an RFI-free environment. In particular, raised cosine windowing with nominal alignment either sustains or enhances rates on all loops investigated; due to its ability to mitigate the effect of TD-ISI and echo transients. The benefit of raised cosine windowing is even more apparent with misalignments of ± 112 samples, although the use of the 8576 sample window version in these cases causes a slight rate decline on very short loops when compared to rates achieved with rectangular windowing. Otherwise, it is this window version which enhances rates the most, especially on longer loops.

The use of constrained windowing in an RFI-free environment is another matter. With nominal conditions, all versions of constrained windowing suffer from performance degradation with loops of 3500 feet or less, again due to the noise enhancement they introduce. Nevertheless, constrained windowing with the asymmetric $\{1, -1\}$ window in particular, increases performance substantially on longer loops, even when compared to raised cosine windowing results. This performance improvement over rectangular windowing is even greater when a misalignment of ± 112 samples is present. Thus, the asymmetric $\{1, -1\}$ windowing technique can be considered for improving performance, but only with loops of 4000 feet or more. Many VDSL applications (e.g. fibre to the curb) however tend to target shorter loops.

The number of computations required to implement the various RFI suppression techniques are summarised in Table 7-19. Raised cosine windowing with folding incorporated requires by far the least amount of computation.

Table 7-19: Additional Computations Estimates for RFI Suppression Algorithms

Algorithm	Real Multiplications per Symbol	Real Additions per Symbol
8576 Sample Raised Cosine Windowing with Folding	384*	768
Constrained Hanning Windowing	31752**	9392**
Constrained Asymmetric Windowing	24708**	4696**
Single Notch Filtering***	26496	35328
Dual Notch Filtering***	52992	70656

* Assumes one symbol delay

** Does not take into account increase in FFT complexity

*** Does not take into account the computations required for RFI detection and frequency tracking.

Finally it was shown in section 7.4 that RFI cancellation algorithms proposed in the literature are not suitable in the context of the current VDSL standard and expected RFI impairments.

Chapter 8

Conclusions and Recommendations

This work has identified and investigated a variety of approaches to combating ingress RFI in DMT VDSL systems. In particular, the performance and complexity of the following four RFI suppression techniques were analysed:

- receiver windowing extending beyond the symbol boundaries
- receiver windowing constrained within the symbol boundaries followed by decision feedback equalization
- digital RFI notch filtering
- RFI time modelling and digital cancellation

8.1 Thesis Conclusions

A system model was elaborated in Chapter 5 in order to provide a relevant context in which to evaluate the various RFI suppression techniques. The features and characteristics of the model were determined based on the available literature, the ANSI VDSL standard, and the thesis author's insights in the matter, which were obtained through industry experience.

Based on the system model and the RFI suppression techniques studied in Chapter 6, a complete end to end time-domain simulation tool was developed to investigate the effect of the following aspects on system performance with the various RFI suppression techniques in place:

- loop characteristics
- time-domain intersymbol interference
- near-end echo
- receiver symbol misalignments
- additive white Gaussian noise
- presence of one and two RFI sources located in different bands

Based on the theoretical analysis presented in Chapter 6 and the simulation results of Chapter 7, the following is concluded:

- The notch filtering techniques developed for combating RFI in single carrier VDSL are also applicable in DMT VDSL. The increase time-domain intersymbol interference and the attenuation of active bins expected with this filtering is of little significance compared to the provided RFI suppression gains.
- The proposed RFI cancellation techniques using modelling do not work. They were developed based on assumptions that are not valid in the context of the VDSL standard.

-
- Selection of the optimum RFI suppression approach depends on the number and severity of the interference as well as the length of the loop. With a single severe interferer present, raised cosine windowing (8576 sample and 8384 sample versions) provides the best performance overall, followed closely by notch filtering. In this scenario, constrained windowing suffers too much from its noise floor enhancement to be a viable alternative. However, with two severe interferers present instead of one, the use of the constrained asymmetric $\{1, -1\}$ window provides performance that is superior compared to raised cosine windowing performance, although notch filtering is clearly the best approach to take in this scenario.
 - Raised cosine windowing increases the system's ability to deal with intersymbol interference, near-end echo, and receiver symbol misalignment. In particular, the achievable rates with the 8384 sample raised cosine window in an RFI-free environment are either maintained or enhanced with respect to nominal performance, for all loops and for various receiver misalignments. The use of the 8576 sample window version causes a slight rate decline on very short loops
 - Constrained windowing in an RFI-free environment tends to deteriorate performance with loops less than 4000 feet. They can however improve rates on longer loops, particularly with the asymmetric $\{1, -1\}$ window, if the noise floor is sufficiently low.
 - Raised cosine windowing followed by folding is by far the least computational intensive of all proposed RFI suppression techniques.
 - The performance of constrained windowing and notch filtering is dependent on the validity of estimates. For the case of constrained windowing, the operation of the DFE relies on the channel estimate and the decoded samples (slicer output) being correct. Notch filtering requires a fairly good centre frequency estimate of the RFI, and perhaps other RFI metrics if the notch bandwidth is to be adaptive.

8.2 Recommendations for further study

Considering the low computational complexity and the overall performance advantages of the raised cosine windowing algorithm, it would be worthwhile to further investigate its performance in the presence of other impairments such as synchronous and asynchronous NEXT. It is expected that windowing will mitigate the effect of NEXT much in the same way that it mitigated the effect of TD-ISI and near-end echo. It may even help confine to a few bins the interference due to an in-band narrowband noise source. The algorithm's performance should also be evaluated with varying levels of RFI present. As pointed out in section 7.5, the simulation results presented in this thesis correspond to severe RFI scenarios, with RFI levels at their maximum expected value. In practice, the occurrence of RFI in a controlled environment such as at the central office should be limited to much lower power levels.

For more harsh RFI conditions, the application of notch filtering could be considered as a complimentary RFI suppression technique to raised cosine windowing. Further analysis of the various algorithms for estimating the RFI centre frequency and for performing notch bandwidth adaptation would be needed however to optimise this approach.

References

- [1] Al-Dhahir, N., and Cioffi, J., "A Low-Complexity Pole-Zero MMSE Equalizer for ML Receivers," Proc. Allerton Conf. Commun., Control, Comput., pp. 623-632, September 1993.
- [2] Al-Dhahir, N., and Cioffi, J., "Optimum Finite-Length Equalization for Multicarrier Transceivers," *IEEE Transactions on Communications*, vol. 44, no. 1, pp. 56-64, January 1996.
- [3] ANSI Standard T1.413, *Network and customer installation interfaces – Asymmetrical Digital Subscriber Line (ADSL) Metallic Interface*, ANSI, New York, March 1995.
- [4] ANSI T1E1.4 Working Group on DSL Access, *T1.424 Trial Use Standard: Interface Between Networks and Customer Installations -Very-high-bit-rate Digital Subscriber Line (VDSL) Metallic Interface*, Alliance for Telecommunications Industry Solutions (ATIS), Washington, DC, 2002.
- [5] De Clercq, L., et al., "Mitigation of Radio Interference in xDSL Transmission," *IEEE Communication Magazine*, pp. 168-173, March 2000.
- [6] De Clercq, L., and Spruyt, P., "Ingress measurements on twisted pairs," *ANSI Contribution T1E1.4/97-315*, September 1997, Minneapolis, MN.
- [7] Cuypers, G., et al., "Combining Per Tone Equalization and Windowing in DMT Receivers," *IEEE International Conference on Acoustics, Speech, and Signal Processing Proceedings*, vol. 3, pp. 2341-2344, 2002.
- [8] Foster, K., and Cook, J., "The Radio Frequency Interference (RFI) environment for very high-rate transmission over metallic access wire-pairs," *ANSI Contribution T1E1.4/95-020*, February 1995, Dallas, TX.

-
- [9] Im, G.-H., and Won, H.-C., "RF Interference Suppression for VDSL System," *IEEE Transactions on Consumer Electronics*, vol. 47, no. 4, pp. 715-722, November 2001.
- [10] Isaksson, M., et al., "Pulse Shaping with Zipper – Spectral Compatibility and Asynchrony," *ANSI Contribution T1E1.4/98-041*, (Telia), March 2, 1998, Austin, TX.
- [11] ITU, "Draft Revision of Recommendation G.993.1 – Very high speed digital subscriber line", *ITU-T Study Group 15 – Contribution 107R1*, February 2004.
- [12] ITU-T Standard G.992.5, *Asymmetric Digital Subscriber Line (ADSL) transceivers – Extended bandwidth ADSL2 (ADSL2plus)*, ITU, Geneva, 2005.
- [13] Jeong, B. J., and Yoo, K. H., "An Efficient Digital RFI canceller for DMT based VDSL," *ANSI Contribution T1E1.4/97-422*, December 1997, Sacramento, CA.
- [14] Kapoor, S., and Nedic, S., "Interference Suppression in DMT Receivers Using Windowing," *IEEE Proc. ICC*, pp. 778-782, 2000.
- [15] Liu, Y., Laakso, T., and Diniz, P., "Adaptive RFI Cancellation in VDSL Systems", *ECCTD'01 - European Conference on Circuit Theory and Design*, vol. 3, pp. 217-220, August 2001, Espoo, Finland.
- [16] Liu, Y., Diniz, P., and Laakso, T., "Adaptive Steiglitz-McBride Notch Filter Design for Radio Interference Suppression in VDSL Systems," *IEEE Proceedings Globecom'01*, pp. 359-363, December 2001, San Antonio, TX.
- [17] Locatelli, R., et al., "An Optimized Digital RFI Cancellation Scheme for DMT-Based VDSL Systems," *IEEE International Conference on Communication Systems (ICCS 2002)*, vol.1, pp. 229-233, November 2002, Singapore.
- [18] Locatelli, R., et al., "RFI Cancellation in DMT VDSL: A Digital Frequency Domain Scheme," *IEICE Transactions on Fundamentals*, vol. E86-A, no.8, pp. 1993-2000, August 2003.

-
- [19] Ödling, P., et al., "An Approach to Analog Mitigation of RFI", *IEEE Journal On Selected Areas In Communications*, vol. 20, no. 5, pp. 974-986, June 2002.
- [20] Proakis, G., *Digital Communications*, Third Edition, McGraw-Hill, Boston, Massachusetts, 1995.
- [21] Proakis, G., and Manolakis, D., "Digital Signal Processing, Principles, Algorithms, and Applications", Third Edition, Prentice Hall, Upper Saddle River, New Jersey, 1996.
- [22] Pazaitis, D.I., et al., "Equalization and Radio Frequency Interference Cancellation in Broadband Twisted Pairs Receivers," *IEEE Proceedings Globecom'98*, vol. 6, pp. 3503-3508, Nov. 1998, Sydney.
- [23] Sehlstedt, M., "RFI Cancellation in VDSL," Master's Thesis, Luleå University of Technology, 2000.
- [24] Sjöberg, F., et al., "Zipper: A Duplex Method for VDSL Based on DMT," *IEEE Transactions On Communications*, vol. 47, no. 8, pp. 1245-1252, August 1999.
- [25] Sjöberg, F., et al., "Digital RFI suppression in DMT-based VDSL systems," *Proc. Int. Conf. Telecommunications (ICT)*, vol. 2, pp. 189-193, June 1998, Chalkidiki, Greece.
- [26] Starr T., Cioffi, J. M., and Silverman, P. J., *Understanding Digital Subscriber Line Technology*, Prentice Hall, Toronto, 1999.
- [27] Stewart, J., *Second Edition Calculus Early Transcendentals*, Brooks/Cole Publishing Company, Pacific Grove, 1991.
- [28] Van Acker, K., "Equalization and Echo Cancellation for DMT-Based DSL Modems," Doctoral Thesis, Katholieke Universiteit Leuven, January 2001.
- [29] Wiese, B., and Bingham, J., "Digital Radio Frequency Cancellation for DMT VDSL," *ANSI Contribution T1E1.4/97-460*, December 1997, Sacramento, CA.