Université d'Ottawa • University of Ottawa

Digital Noise Cancellation for xDSL

David Fenton

A Thesis Submitted to the School of Graduate Studies and Research
in Partial Fulfillment of the Requirements for the Degree of

Master of Applied Science
in Electrical Engineering

Ottawa-Carleton Institute for Electrical and Computer Engineering

School of Information Technology and Engineering
(Electrical and Computer Engineering)

August, 1999

Your file Votre référence
ISBN: 0-612-90069-X
Our file Notre référence
ISBN: 0-612-90069-X

Abstract

Digital subscriber line (DSL) modems operate in frequency bands which coincide with many significant radio-frequency interference sources, particularly commercial AM radio. In these bands, the balance of most twisted-pair cables is low enough to allow substantial interference to transfer to differential mode, disrupting the transmitted information signal.

To compensate for this handicap, xDSL receivers will require a front-end circuit which uses the common-mode signal as a reference to cancel out as much differential interference as possible. Unfortunately, design of such a canceller is complicated by the statistical properties of the interference, as well as the complex interactions between the differential and common-mode signals on the loop. If a digital canceller is desired, compensation for finite-precision effects poses an additional challenge.

This thesis examines the feasibility of wideband digital noise cancellation for xDSL. Emphasis is placed on the stability of the adaptive algorithm, the number of bits of precision required, the adaptive filter parameters which lead to the best noise reduction, and the number of filter taps required. Simulation results indicate that digital cancellation is possible, giving a maximum noise reduction of 30 - 40 dB. Some hardware implementation problems are also identified in the course of the simulations.

Table of Contents

List of Abbreviations

A/D	analog-to-digital
ADC	analog-to-digital converter
ADSL	asymmetric digital subscriber line
AGC	automatic gain control
AM	amplitude modulation
ASIC	application-specific integrated circuit
AWG	American wire gauge
AWGN	additive white Gaussian noise
BER	bit error rate
CAP	carrierless amplitude modulation / phase modulation
CO	central office
DFE	decision feedback equalizer
DMT	discrete multi-tone
DNE	digital noise estimator
DSL	digital subscriber line
DSP	digital signal processing
FDD	frequency-division duplexing
FEXT	far-end crosstalk
FFE	feedforward equalizer
FFT	fast Fourier transform
FIR	finite impulse response
FPGA	field-programmable gate array
f_s	sampling frequency
IFFT	inverse fast Fourier transform
ISI	inter-symbol interference
LMS	least-mean-square
Mbps	mega-bits-per-second
MMSE	minimum mean-squared error
MSE	mean-squared error
MSPS	mega-samples-per-second
NEXT	near-end crosstalk
ONU	optical network unit
PAR	peak-to-average ratio
PSD	power spectral density
QAM	quadrature amplitude modulation
RFI	radio-frequency interference
RLS	recursive least squares
rms	root-mean-square
SDMT	synchronized discrete multi-tone
SIR	signal-to-interference ratio
SNR	signal-to-noise ratio
SQNR	signal-to-quantization-noise ratio
TP	twisted pair
UTP	unshielded twisted pair
VDSL	very-high-speed digital subscriber line
xDSL	generic designation for ADSL/HDSL/VDSL

1. Introduction

1.1 Thesis Motivation and Organization

Faced with increasing competition from alternate service providers, telephone companies are beginning to offer broadband telecommunications services over existing twisted-pair cables. Asymmetric Digital Subscriber Line (ADSL) deployment has commenced, and standardization of Very-High Speed Digital Subscriber Line (VDSL) equipment is underway. On short loops, VDSL will be able to sustain data rates as high as 52 Mbps. On longer loops, DSL signals are so badly attenuated that externally generated radio-frequency noise on the loop is often more powerful. This interference commonly forces modems to lower their data rates in order to maintain adequate margins against other noise sources.

While it is possible for a receiver to cancel out much of the interference using its equalizer, this task is greatly simplified if the noise level is reduced at the receiver's front end. Fortunately, most interference couples to the loop in common mode before partially converting to differential mode and corrupting the information signal. The common-mode signal can be used as a noise reference for a front-end adaptive noise canceller which removes the interference component of the differential signal.

The front-end canceller can be implemented using either analog or digital circuitry. Both options have advantages and disadvantages. An analog adaptive noise canceller conserves analog-to-digital converter resolution, but is more difficult to design than a digital device (this is particularly true in systems with low power supply voltages). It also remains susceptible to additive noise. A digital canceller is easier to design and miniaturize for production, but may require the signals to be represented using a very high number of bits, and could be susceptible to stability problems if improperly implemented. Since the noise canceller design will not need to be included in the ADSL or VDSL standards documents, either implementation could be selected by a modem manufacturer.

This thesis will examine the feasibility and potential performance of a digital noise canceller. Severe VDSL loop conditions are assumed for all tests, but extension of the same ideas to ADSL loop conditions is simple. The document is laid out as follows:

a. Chapter 2 describes the basic differences between differential and common-mode signal propagation, and gives the balance characteristics of a typical loop. It also describes some of the characteristics of the radio-frequency interference (RFI), impulse noise and crosstalk noise which might be encountered on the line. The channel and noise models used for the project's simulations are also described;

b. Chapter 3 briefly surveys some issues that are currently under discussion at the T1E1.4 subcommittee. In particular, it describes the different ways in which the two competing VDSL line codes handle noise cancellation. Both techniques will require a wideband front-end canceller, which is the central concern of this project. However, they have very different methods of cancelling residual noise at the equalizer, and it is important to describe these differences in order to understand some of the design choices made for the front-end canceller. This chapter also explains the subcommittee's latest thinking on analog-to-digital converters (ADCs) for VDSL, and analyzes the circumstances under which a digital front-end canceller will be feasible;

c. Chapter 4 is an overview of previous work conducted in the field of noise cancellation, and focuses mostly on least-mean-square (LMS) adaptive filter theory. Conditions for convergence, numerical stability, and performance in non-stationary environments are discussed. Other noise cancellers designed for a variety of telecommunications applications are also briefly surveyed;

d. Chapter 5 outlines the proposed noise canceller architecture, and identifies the goals of the experiments to be carried out in the following chapter;

e. Chapter 6 describes the program and performance measurements used to test the noise canceller in simulation, and analyzes the results; and

f. Chapter 7 summarizes the conclusions derived from the chapter 6 experiments, and gives some recommendations for future work in the field.

Throughout the document, all power measurements assume that the loop is terminated in a 100 Ω resistor.

1.2 Thesis Contributions

Original research work is distributed throughout this document to preserve the flow of ideas and allow later chapters to build on earlier results in a cumulative manner. The main contributions of the thesis are:

a. the proposal of a simple model for the twisted-pair common-mode channel;
b. some observations on the importance of capacitive leakage through the hybrid transformer;
c. an analysis of the conditions under which the use of 12-bit analog-to-digital converters will allow digital noise cancellation for VDSL;
d. the proposal of a simple method to simulate the transfer of signals between the differential and common modes of transmission;
e. the design of a digital noise canceller for VDSL; and
f. through simulation testing, an evaluation of the noise canceller's capabilities, required design parameters and potential hardware implementation problems.

2. Some Characteristics of Twisted-Pair Loops

In order to develop a noise canceller for digital subscriber lines, it is first necessary to present some background information on typical twisted-pair (TP) loop characteristics. This chapter describes a general lumped-parameter model for transmission lines, and then defines the parameters necessary to simulate the differential mode of transmission in TP loops. The differences between common-mode and differential transmission are then discussed, along with the concept of loop imbalance. Next, a method of generating a filter to model the loop's attenuation and phase distortion is presented. The technique uses two-port matrices, and can simulate the effects of wire gauge changes, bridged taps, and termination mismatches. The parameters it uses are well-defined only for differential transmission, and so a similar model is developed for the common-mode case. Last, the imperfections of the transformers used in modem front ends are briefly discussed.

Also included in this chapter is an overview of the different types of interference which are likely to be present on the loop. For noise cancellation, the most important of these is induced radio-frequency interference. Subsequent sections describe impulse noise and crosstalk noise to the level of detail required for this project's simulation work.

2.1 The Physical Plant

Twisted-pair loops are used to connect customer equipment to the telephone company's central office (CO), either directly or using an intermediate optical network unit (ONU) connected to the CO by fiber-optic cable. The average North American loop is approximately 2.44 km (8,000 ft) long, with 25% of loops longer than 3.66 km (12,000 ft) [67]; VDSL will only be offered over loops under 1.5 km long, but ADSL will be able to operate over many of the longer cables. By convention, the two wires of the loop are referred to as Tip and Ring.

Almost all of the cables in the deployed plant are copper, and can be classified as one of four gauges. American wire gauge (AWG) 26 cable is the narrowest of the four, with a diameter of only 0.404 mm, and is the medium with the worst attenuation characteristics [57]. Accordingly, it is the only cable type considered in this document.

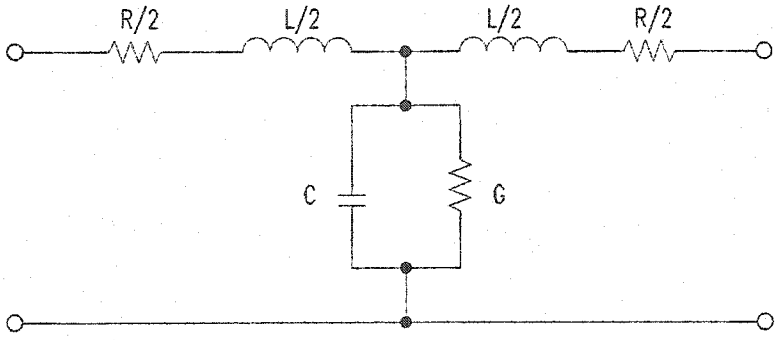

Figure 2-1: T-equivalent circuit in terms of primary constants

A loop can be quickly characterized by a set of parameters known as its primary and secondary constants (which actually aren't constant, but vary with frequency). The primary constants are the series inductance *L*, the shunt capacitance *C*, the shunt conductance *G*, and the series resistance *R*, all expressed in standard electrical units per unit length. An incremental section of a transmission line can be modelled using the equivalent circuit shown in Figure 2-1 [56].

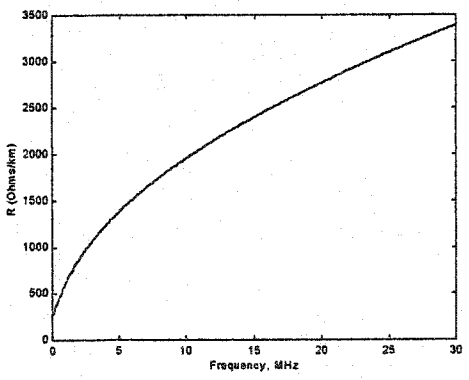

Figure 2-2: Primary constant R vs. frequency (differential mode)

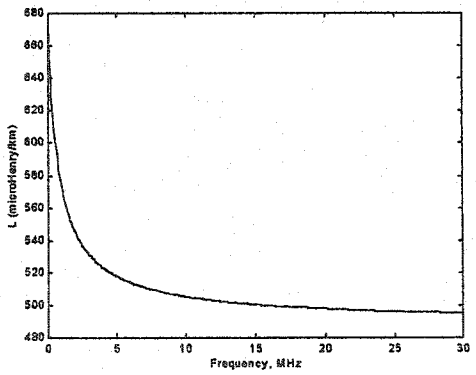

Figure 2-3: Primary constant L vs. frequency (differential mode)

The series resistance is the most important of these parameters in defining the attenuation which occurs over the loop. As the frequency of the signal propagating through the cable increases, the skin effect causes the signal currents to flow closer to the outside surfaces of the wires [14]. This means that the resistance of the loop rises with frequency. Resistance also varies significantly with temperature, rising by about 4% for each additional 10°C [57]. Above 300 kHz,

R is approximately proportional to $\sqrt{f}$ [67]. Below that frequency, it climbs at a slower rate. Altogether, for a AWG 26 cable it can be modelled by [9]:

$$R(f) = \sqrt[4]{(286.17578)^4 + 0.1476962 \cdot f^2} \quad \Omega/\text{km}, \tag{2.1}$$

which gives the curve plotted in Figure 2-2.

L and C are the most important parameters in defining the speed at which the signal propagates through the cable. The skin effect causes the inductance L to decrease with rising frequency, and it can be modelled by [9]:

$$L(f) = \frac{675.36888 + 488.95186 \cdot \left(f / 806.33863\,\text{Hz}\right)^{0.92930728}}{1 + \left(f / 806.33863\,\text{Hz}\right)^{0.92930728}} \quad \mu\text{H/km}. \tag{2.2}$$

A plot of this function is shown in Figure 2-3.

The capacitance C depends on the distance between, and twist of, the wires, as well as the dielectric medium between the wires [14]. It is approximately 49 nF/km at all frequencies for 26-gauge cable. Last, the conductance G is also responsible for attenuation effects, but is usually negligible compared to the much larger series resistance R. It depends primarily on the dielectric medium, and can be modelled as [9]:

$$G(f) = 43 \cdot f^{0.7} \quad \text{nS/km}, \tag{2.3}$$

which is shown in Figure 2-4.

There are two secondary parameters which can be expressed in terms of the primary constants: the propagation constant and the characteristic impedance. These are developed in [56] by applying the Kirchoff voltage and current laws to an incremental section of the loop. The propagation constant is given by:

$$\rho = \alpha + j\beta = \sqrt{(R + j\omega L)\cdot(G + j\omega C)} \tag{2.4}$$

where ω is the radian frequency. The real part of this quantity, α, is called the attenuation constant and expressed in Nepers per unit length (1 Neper = 8.686 dB). The attenuation constant for AWG 26 cable is plotted, in dB/km, in Figure 2-5. The imaginary part of ρ , β, is called the phase constant, and is expressed in radians per unit length. The other secondary constant is the characteristic impedance Z_0, which is independent of cable length and expressed in Ohms:

$$Z_0 = \sqrt{\frac{(R + j\omega L)}{(G + j\omega C)}} \tag{2.5}$$

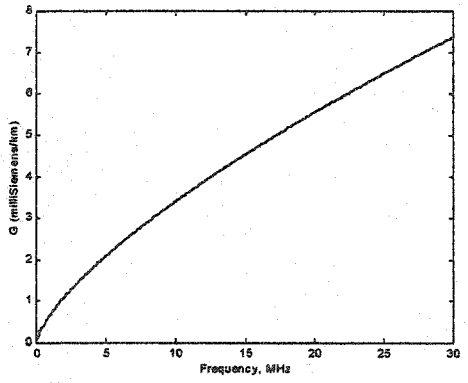

Figure 2-4: Primary constant G vs. frequency (differential mode)

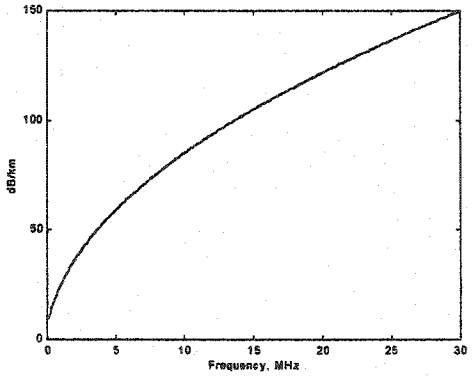

Figure 2-5: Attenuation constant vs. frequency (differential mode)

The role of these parameters is most easily understood by examining the following equation, which describes signal propagation in a lossy transmission line [55]:

$$I(d) = I_0 \cdot e^{-\rho d} = I_0 \cdot e^{-\alpha d} \cdot e^{-j\beta d} \tag{2.6}$$

Here, *I(d)* is a phasor quantity describing the differential current in the loop, and *d* is the distance from the transmitter. I_0 is the current at the transmitter end of the loop, and can be estimated from the input power P_{in} using:

$$I_0 = \sqrt{\frac{2 \cdot P_{in}}{Z_0}} \tag{2.7}$$

This equation assumes that I_0 is not given in rms units and that the input power is not applied through any device which causes an insertion loss. The model described by (2.6) and (2.7) will be expanded into a full two-port matrix representation in Section 2.4.

2.2 *The Basics of Signal Transmission and Reception*

A signal is transmitted over a twisted-pair cable in differential mode (also called metallic, or transverse, mode). It is injected on Tip and Ring with equal but inverted amplitude. The voltage difference between the two wires causes a differential current, which travels down one wire of the loop, through the hybrid device at the receiver end, and back along the other wire; see Figure 2-6. As the current passes through the primary side of the hybrid device, which usually includes a transformer to protect the receiver from excessive currents, the differential signal is coupled to the load.

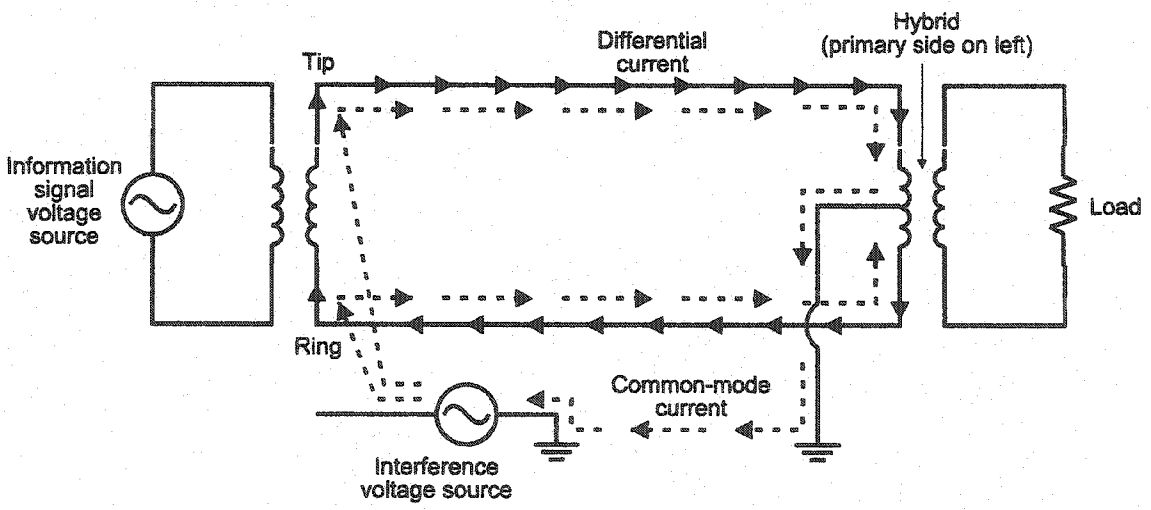

Figure 2-6: Loop currents

If the twisted-pair wire is placed in an electric field, then an additional voltage will be induced on each wire of the pair. Coupling of the electric field may occur over the entire length of the cable, or perhaps over only a small length of it. If the cable is tightly twisted, it is highly likely that the field will illuminate both wires with equal strength, and thus the voltage induced between Tip and ground will be equal to the voltage induced between Ring and ground. This situation is referred to as common-mode (or longitudinal) voltage.

As an example, consider the case where a tightly twisted cable is illuminated by RF energy with strength E V/m. Under worst-case coupling conditions, the common-mode voltage induced will (to a first-order approximation) be equal to E volts [23]. The resulting common-mode currents travel down both wires simultaneously, with a return path through ground. The currents enter both ends of the hybrid transformer's primary coil at the same time, and meet each other at the centre tap. As a result, two equal and opposite voltages appear at the secondary coil and cancel each other out. In this manner, the common-mode signal is perfectly rejected while the differential (desired) signal is passed to the receiver.

2.3 Loop Imbalance

Two conditions are primarily responsible for loop imbalance. First, as the frequency of the interfering electric field increases, or the number of twists per unit length of the cable decreases, the voltages induced between ground and the two wires are less likely to be equal. The difference in the two potential levels represents a differential voltage, whose associated current will induce a net voltage at the secondary side of the hybrid transformer. Second, if the interfering signals couple to the wires with equal strength but encounter unbalanced impedances to ground as they travel toward the receiver, some of the common-mode current will convert to differential mode [14]. This also causes some of the induced energy to pass to the load. In either case, the end result is the addition of an unwanted noise signal to the desired differential signal that was originally transmitted over the loop.

Formally, cable balance is defined by the formula:

$$B = 20 \cdot \log_{10} \left| \frac{V_{cm}}{V_d} \right| \tag{2.8}$$

where V_{cm} is the common-mode voltage between the cable and ground, and V_d is the differential voltage between the two wires of the cable [5]. To meet specifications for telephone use, it must be at least 60 dB at 1,004 Hz, but no required levels have been set for other frequencies [57]. As a result, it is difficult to find a general model for cable balance [61]. Most loops are likely to have a balance of 50 to 60 dB at frequencies lower than 100 kHz. Above that, B degrades at about 12

dB/decade, and may be as little as 20 - 30 dB at VDSL frequencies [5, 12, 14]. In the worst case, a loop may have only 30 dB balance at low frequencies, and as little as 10 or 15 dB at the upper end of the VDSL spectrum [32, 61]. The balance of a typical twisted-pair cable is shown in Figure 2-7.

Most external interference signals couple to the loop in common mode. However, impedance imbalances along the loop can cause some portion of these common-mode signals to be converted into differential signals, and vice-versa. Obviously, the more noise that converts to differential mode, the more difficult it will be to receive the transmitted signal correctly. To better understand the interaction between the two modes of transmission, it is necessary to analyze the characteristic impedance of the loop.

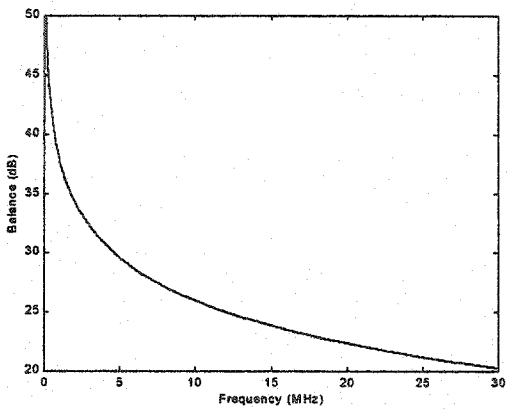

Figure 2-7: Balance of a typical TP cable

At high frequencies, the characteristic impedance given by equation (2.5) is dominated by the reactive elements, and can be approximated as:

$$Z_{0d} = \sqrt{\frac{L}{C}} \tag{2.9}$$

where the subscript *'d'* has been used to indicate that this parameter refers to the differential mode. Using equation (2.2) or reading an approximate value of L from Figure 2-3, it can be verified that at 1 MHz this estimate gives $Z_{0d} \cong 100\ \Omega$. At higher frequencies, Z_{0d} drops only very slightly from this value [57].

In contrast, a common-mode signal propagating down a shielded cable sees the pair as a single wire only half as long as the differential loop [14]. The current return path is through the cable shield, which is a thick conductor with much less resistance than the 26-gauge wire. It is therefore expected that the common-mode resistance per unit length will be much less than its differential counterpart. Similarly, the inductive coupling between the pair and the shield is much smaller than the coupling between the two wires of the pair, so the common-mode inductance per unit length of each wire is higher than for the differential loop. Last, the common-mode capacitance per unit length is only slightly less than the capacitance of the differential line, and the conductance is usually negligible [14].

The higher value of L and unchanged value of C mean that the common-mode characteristic impedance of a single wire should be larger than Z_{0d}. [14] gives 430 Ω as a good approximate value for AWG 26 cable at 1 MHz. Since the two wires are in parallel, the overall common-mode impedance is approximately $Z_{0c} \cong 215\ \Omega$. This is more than twice the magnitude of Z_{0d}. Electromagnetic waves travel down the transmission line by continually exchanging energy between the electrical and magnetic elements (ie. the capacitor and inductor in the lumped-parameter model). Since $Z_{0c} > Z_{0d}$, the common mode presents a greater impedance to this energy transfer [20].

This fact has consequences for the speed at which the two waves travel along the transmission line. At high frequencies, the envelope delay of the wave can be approximated by $\sqrt{L \cdot C}$ [67]. The group velocity is simply the inverse of the envelope delay [20]:

$$v = \frac{1}{\sqrt{L \cdot C}} = \frac{1}{C \cdot Z_0} \tag{2.10}$$

Given that the primary constant C is approximately equal for both the common and differential modes, it is clear that the common-mode wave will travel at approximately half the speed of the differential wave. At high frequencies, the speed of the differential wave is found to be:

$$v_d = \frac{1}{\left(49 \times 10^{-9}\ \text{F/km}\right) \cdot \left(1\ \text{km}/1000\ \text{m}\right) \cdot \left(100\ \Omega\right)} = 2.04 \times 10^8\ \text{m/s} = 0.68c \tag{2.11}$$

where c is the speed of light. Assuming an identical value of C for the common mode gives $v_c = 0.32c$. Discounting the inaccuracies caused by using high-frequency approximations, these results are consistent with the measured values given in [14] for an AWG 26 cable at 1 MHz, where v_d and v_c were found to be about $0.65c$ and $0.34c$ respectively.

A mathematical model for signal attenuation will be given in the next section. However, at this point it is appropriate to highlight the difference in the attenuation experienced by the two modes. Simulation results given in [14] found that at 1 MHz the differential attenuation over a 1-km AWG 26 loop is 26.1 dB, while the common-mode value is 7.4 dB. If the cable does not have proper longitudinal termination, the common-mode attenuation figure is even smaller.

Knowing the speed and attenuation characteristics of both the differential and common-mode waves, it is easy to see how impedance imbalances in the line will affect the quality of the received signal. Consider a cable in which there is a region of impedance imbalance far away from the receiver. A differential signal travelling down the wire will partly convert to common mode at the imbalance (which is more likely to be a diffuse "region" of imbalance than a nicely defined point). Over the length of the cable, the common-mode wave will lag behind the differential one. It will also attenuate more slowly, making it appear stronger relative to the differential signal.

The same logic can also be applied in reverse. Interfering signals which couple to the cable in common mode will be partly converted to differential mode by an impedance imbalance. The differential portion of this noise will arrive at the receiver before the corresponding common-mode signal is present at the centre tap of the hybrid transformer. If multiple interferers couple to the cable at different distances from the receiver, and there are multiple impedance imbalances along the cable, it is clear that the differential and common-mode signals at the end of the loop can be related by a highly complicated correlation function.

When the cable is unshielded, it is possible that there will be no speed difference between the differential and common-mode signals because there is no easy current return path for the common mode. This situation simplifies the timing problem for noise cancellation, but also implies that there will probably be more noise to cancel because of the lack of shielding.

Fortunately, much of the loop plant is buried underground, where it is isolated from outside electromagnetic influences (with the exception of other cables). Even most overhead cable is well-shielded up until the final service drop to the customer premises [57]. Therefore, in most cases it can be expected that much of the external interference will couple to the unshielded cable close to the receiver, and that the two modes will remain closely synchronized in time. However, it is prudent to prepare for the worst case, when the cable imbalance occurs at a fault in the cable shielding at some distance from the receiver.

2.4 Attenuation and Phase Distortion

2.4.1 Differential Mode

The attenuation and phase distortion of differential signals is easily modelled using the equivalent circuit for lossy transmission lines which was introduced previously. In this approach, all of the components of the loop (sections with different gauge wire, bridged taps, etc) are represented as two-port circuits. By multiplying all of the corresponding two-port matrices together, a single matrix modelling the entire loop can be found. In turn, the transfer function of the loop can be found from this matrix.

A good development of two-port modelling for twisted-pair wire is given in [9], and will not be fully reproduced here. In general, the loop is represented as:

$$\begin{bmatrix} V_1 \\ I_1 \end{bmatrix} = \begin{bmatrix} A & B \\ C & D \end{bmatrix} \cdot \begin{bmatrix} V_L \\ I_L \end{bmatrix} \tag{2.12}$$

where V_1 and I_1 are the voltage and current on the loop at the transmitter end, and V_L and I_L are the voltage and current at the load (on the secondary side of the receiver's hybrid device). The two-port matrix for a homogeneous section of loop with length d (in km) may be expressed as [9]:

$$\begin{bmatrix} A & B \\ C & D \end{bmatrix} = \begin{bmatrix} \cosh(\rho d) & Z_0 \cdot \sinh(\rho d) \\ \frac{1}{Z_0} \cdot \sinh(\rho d) & \cosh(\rho d) \end{bmatrix} \tag{2.13}$$

at any given frequency, where ρ is the propagation constant.

Using (2.12) and $Z_L = V_L / I_L$, the following transfer function can be found:

$$T(f) = \frac{V_L}{V_1} = \frac{Z_L}{A \cdot Z_L + B} \tag{2.14}$$

The magnitude of $T(f)$ is called the insertion loss. Note that this transfer function does not relate the voltage at the load to the voltage at the transmitter, but to the voltage at the transmit end of the loop. A transfer function which factors in the source voltage V_S and the source impedance Z_S is given by:

$$H(f) = \frac{V_L(f)}{V_S(f)} = \frac{V_L(f)}{V_1(f)} \cdot \frac{V_1(f)}{V_S(f)} = \frac{Z_1}{Z_1 + Z_S} \cdot \frac{Z_L}{AZ_L + B} \tag{2.15}$$

where $Z_1 = V_1 / I_1$. In practice, equation (2.14) is usually more important than (2.15) because restrictions on the transmit power apply to V_1 and not to V_S.

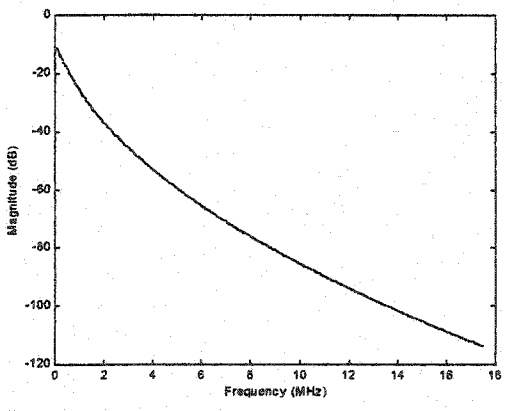

Figure 2-8: Differential insertion loss vs. frequency, 1-km AWG 26 loop

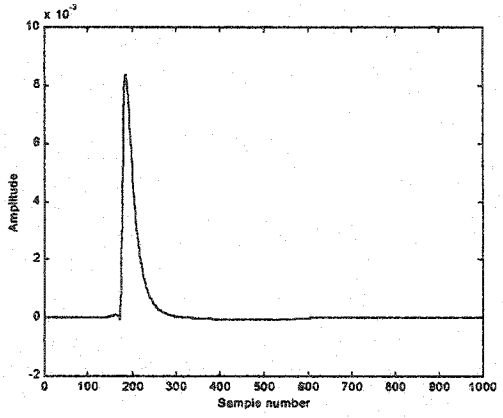

Figure 2-9: Impulse response of 1-km AWG 26 channel (differential mode, 35 MHz sampling frequency)

For a 1-km 26-gauge cable terminated in 100 Ω, $T(f)$ appears as shown in Figure 2-8. Appending the conjugate of $T(f)$ to itself symmetrically and taking an inverse Fourier transform yields a real impulse response for the loop. For a sampling frequency of 35 MHz and the same

cable as was used for Figure 2-8, this gives the impulse response shown in Figure 2-9. This plot clearly shows the propagation delay caused by the loop (about 180 samples, or 5 microseconds, in this case). Also, because the impulse response is not a single spike, some phase distortion of the transmitted signal will occur. This dispersion is caused primarily by the slight dependence of L on frequency, and is the reason why inter-symbol interference (ISI) occurs [9].

2.4.2 Common Mode

Common-mode attenuation and phase distortion can also be modelled using two-port theory, defining the voltages from each wire to ground. Unfortunately, however, there is very little published information on the common-mode characteristics of twisted-pair cable. Furthermore, in the lab it was not possible to reproduce the conditions which cause slower signal propagation in common mode. There were probably two reasons for this. First, the cable was unshielded, denying the common mode current a low-resistance return path. Second, it was wrapped around a spool instead of being extended to its full length; even though its direction of winding was reversed every few turns, this left the loop susceptible to inductive effects, and the wrapped cable definitely did not resemble the electromagnetic model of a long conductor suspended over a ground plane.

It is important to test potential noise canceller designs with misaligned differential and common-mode signals, as phase alignment is a canceller's most difficult function. To compensate for the lack of available information on twisted-pair cable's common-mode characteristics, a simple model has been derived from the differential two-port characterization.

Knowing that the common-mode signal sees the two wires of the cable as a single conductor, it can be expected that the loop itself will have half the resistance as it does in the differential case [14]. The return path through the cable shield is known to have much less resistance than the loop. It is therefore reasonable to estimate that the total common-mode resistance R_c will be 0.5 - 0.55 times the differential resistance R_d. As discussed previously, it is also known that the common-mode capacitance per unit length is very slightly less than its differential counterpart, and that the characteristic impedances of the two modes at 1 MHz are Z_{0c} $\cong$ 215 Ω and $Z_{0d} \cong 100\ \Omega$. Since $Z_{0c} = \sqrt{L_c/C_c}$ at high frequencies, if C_c is estimated to be $0.95 \cdot C_d$, then it is reasonable to expect $L_c = 4.4 \cdot L_d$. [14] neglects the conductance for both

modes, but it is likely to be slightly higher in common mode because the shield is further away from the pair than the two wires are from each other.

To test this model, it must be compared with the attenuation and velocity characteristics of a common-mode signal. Once again from [14], it is known that the cable's attenuation coefficient is about 4.4 dB/km at 1 MHz, without accounting for termination. Most VDSL documents assume a common-mode termination of 50 Ω because there is no need to receive the common-mode signal. However, this resistance is badly mismatched with the characteristic impedance of the channel, and will cause reflections and standing waves on the loop. A noise canceller will therefore need a special front-end circuit to provide correct termination in both modes (for an example of such a circuit, see [15]). According to the model developed above, the limiting value of Z_{0c} at very high frequencies is about 210 Ω. Using this value as the termination resistance will give a total insertion loss of 7.4 dB on a 1-km loop (in actual hardware, both wires would be terminated in 420 Ω and experience a 6 dB loss; considering both wires together, there is only a 3 dB loss). This model gives the common-mode insertion loss depicted in Figure 2-10, and the impulse response shown in Figure 2-11.

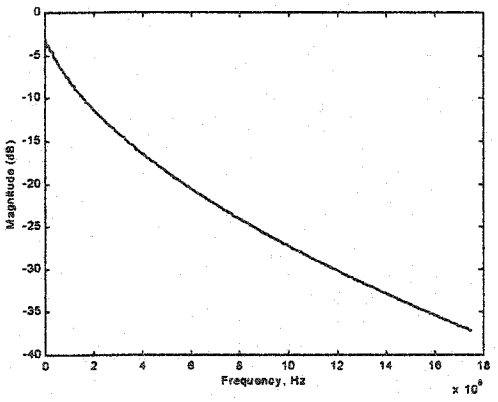

Figure 2-10: Common-mode insertion loss over 1-km AWG 26 cable

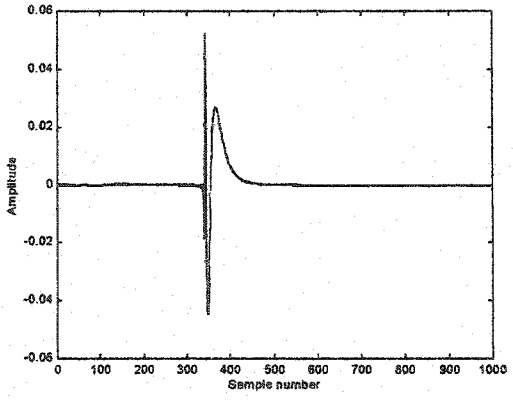

Figure 2-11: Common-mode impulse response for 1-km AWG 26 cable (35 MHz sampling frequency)

This impulse response is an intuitively satisfying result. It stands to reason that a mechanical relay closing at the central office will result in a single energy spike which couples to the loop in common mode. If this impulse travels to the receiver mostly in common mode, it is clear that the common-mode impulse response should resemble a burst of impulse noise. Also

reassuring is the fact that the delay is approximately 340 samples for a 35 MHz sampling frequency. This corresponds to a velocity of approximately $0.34c$, which agrees with the number given in [14].

2.5 Bridged taps

Bridged taps represent another impairment to high-speed transmission. These are additional open-circuited twisted pairs which branch off from operational wire pairs that were originally installed to allow the loop plant to be easily reconfigured. One major North American carrier has estimated that about 80% of its non-loaded loops contain bridged taps, with an average length of 400 m [67]. Similarly, a loop survey in Taiwan found that 65% of loops had bridged taps, of which 69% were under 350 metres long [8]. Public carriers are reluctant to remove all of these spare cables, partly for economic reasons and partly because the taps do give them greater flexibility in loop maintenance. VDSL researchers must therefore assume that the bridged taps will remain in place.

The presence of a bridged tap has two principal effects. First, a signal transmitted over the loop will reflect from the end of the tap and echo back to the transmitter. Second, this reflected signal will also appear at the receiver as an attenuated and out-of-phase copy of the main signal (which propagates directly from the transmitter to the receiver). Altogether, the main signal suffers a net loss in power, both from the attenuation experienced in the bridged tap and from the destructive interference which occurs on the loop [67].

If the reflected signal appears at the receiver 180° out of phase with the main signal, the maximum possible amount of destructive interference will occur. There will therefore be certain frequencies at which the power spectral density curve dips significantly. These are referred to as "nulls" even though the magnitude drop is typically only 3 - 5 dB. If a loop has a single bridged tap less than 600 metres long, the nulls will occur at the frequencies for which the bridged tap length d_{bt} is equal to an odd multiple of a quarter wavelength; this approximation is valid because the phase delay is approximately constant at high frequencies (about 5.4 μs/km). More succinctly, the nulls are located approximately at odd multiples of:

$$f_0 = \frac{45{,}000}{d_{bt}} \text{ kHz} \tag{2.16}$$

where d_{bt} is expressed in metres [67].

Bridged taps are easily incorporated into the two-port model of a twisted-pair wire. An unterminated wire pair with propagation constant ρ_{bt}, characteristic impedance Z_{0bt}, and length d_{bt} can be represented by the matrix:

$$\Phi_{bt} = \begin{bmatrix} 1 & 0 \\ \frac{1}{Z_{0bt}} \cdot \tanh(\rho_{bt} d_{bt}) & 1 \end{bmatrix} \tag{2.17}$$

When multiplied in sequence with the matrices representing the remainder of the loop, this matrix will create appropriate spectral nulls in the transfer function.

A recent study tested a field-deployable ADSL modem over loops with a variety of different bridged tap lengths [25]. The modem was forced to drastically reduce its data rate for tap lengths of 300 - 450 feet (91.5 - 137 metres). Unsurprisingly, these lengths correspond to the primary nulls for frequencies between 330 - 500 kHz. These frequencies are near the lower end of the downstream spectrum allocation for ADSL, where a discrete multi-tone modem would normally be able to support large constellations at each subcarrier (these modems will be further discussed in chapter 3).

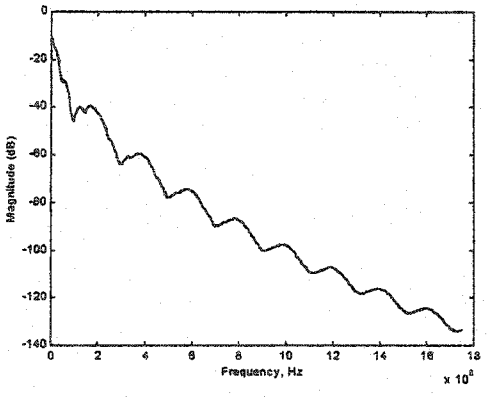

Figure 2-12: Differential insertion loss for 1-km AWG 26 loop with 100-m and 50-m bridged taps

In contrast, the VDSL spectrum will only overlap a primary null if the bridged tap is 25 metres or less in length. Obviously, these short taps will cause very deep notches in the spectrum because the attenuation over a total path length of less than 50 metres is small. For longer bridged-tap lengths, the nulls at harmonics of f_0 will fall in the VDSL frequency band but will be less deep. Figure 2-12 shows the insertion loss for a 1.1-km AWG 26 loop with 100 and 50-metre bridged taps.

2.6 Leakage Through the Hybrid Transformer

In most ADSL or VDSL systems, the primary source of imbalance will be the loop. However, it is also possible for the front end of the modem to be poorly balanced. In practice, the physical imperfections of the hybrid transformer often let some portion of the common-mode signal leak to the secondary coil.

In general, the high-frequency imperfections of a transformer can be modelled by a "leakage inductance" and one or more capacitances [26, 47]. The leakage inductance accounts for imperfect magnetic coupling between the primary and secondary windings, and is not a major concern when considering the transformer's common-mode characteristics. The capacitances account for a variety of locations where voltage gradients can exist in the transformer, specifically [26]:

a. between turns;
b. between layers of individual windings;
c. between windings;
d. between terminals;
e. between the core and the end turn of a layer; and
f. between the core and each terminal.

Of these, the capacitance between windings, called the "bridge capacitance", is usually the most significant [26]. If the information signal is much weaker than the common-mode interference, the additional noise capacitively coupled to the secondary side of the transformer can be significant even when the bridge capacitance is very low.

A simple laboratory experiment was conducted to confirm that transformer leakage is predominantly a capacitive effect. A sinusoidal waveform with amplitude 1 V was injected into a two-metre cable in common mode. In turn, the cable was connected to a prototype hardware board equipped with an inexpensive Mini-Circuits TT1-6 transformer and a 100-Ω differential termination. The amplitude of the voltage which leaked to the secondary side of the transformer for 1.5 - 10 MHz input signals is plotted in Figure 2-13; all measurements were made using a 150-MHz analog oscilloscope. The mismatched termination caused some resonance on the cable, and clock noise and other interference on the board made accurate readings difficult. However, the leaked voltage clearly rises with the signal's frequency. If the loop is unterminated, a much smoother curve showing an approximately linear voltage increase versus frequency is obtained.

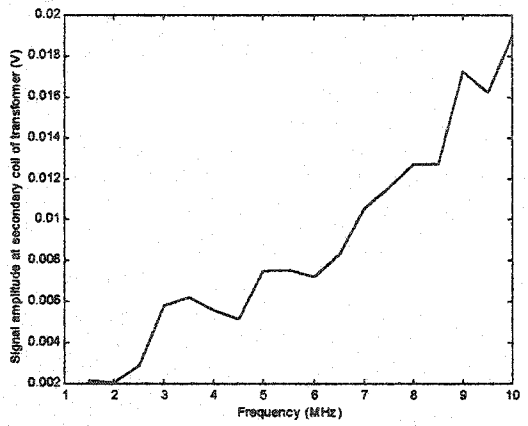

Figure 2-13: Leakage of 1V Common-Mode Signal through Mini-Circuits TT1-6 Transformer (100-W differential termination)

A review of current vendor information shows that good telecommunications transformers typically have a minimum balance in the order of 60 - 80 dB at their design frequencies (upper frequency limits of 4000 Hz - 2.0 MHz were observed, depending on the application for which the transformer was intended). Most of the time, the loop balance will be much lower, meaning that the cable will transfer far more interference to differential mode. However, in cases where the interference transfers to differential mode at a great distance from the receiver, the common-mode noise which leaks through the transformer could have a magnitude comparable to the interference contributed by the loop.

The simplest way to fix this capacitive coupling problem is to find a better transformer. If installing a better transformer is not cost-effective, some preliminary hardware work has indicated that capacitive noise can be mitigated using fixed analog circuitry: for example, the entire front end of the board can be built using well-balanced differential circuits (this also helps to suppress clock noise from elsewhere on the board), or a simple cancellation circuit which mimics the transformer's common-mode characteristics can be installed. However, for the purposes of the simulation work carried out in this document, the transformer is considered to have a minimum balance of 60 dB, and the capacitively-coupled noise is not suppressed before it arrives at the digital canceller.

2.7 *Radio-Frequency Interference (RFI)*

The radio-frequency interference which couples to a twisted-pair cable may originate from many sources. These include transmissions from broadcast AM stations, amateur AM radio, shortwave radio, maritime and aeronautical HF communications, and navigational aids.

ADSL has a maximum transmitted power of about 14 dBm, and a downstream spectrum which extends to 1.104 MHz before cutoff [9]. It will be most susceptible to RFI from broadcast AM radio, although it may also experience interference from radio-navigation aids below the commercial AM bands. VDSL's spectrum extends from 1.104 to 20 MHz, and its maximum transmitted power is 11.5 dBm [9]. This band also overlaps with commercial AM frequencies. In addition, almost the entire spectrum up to 20 MHz is licensed for amateur AM, shortwave, and mobile HF radio transmissions [34, 49].

Broadcast AM is by far the most ubiquitous and persistent form of interference. In North America, commercial AM stations occupy 10 kHz bands at centre frequencies from 540 - 1600 kHz, and typically have 50 kW or 5 kW transmitters. Worst-case conditions can be severe: loops in New York City are exposed to at least 22 stations with transmitter powers of 1 kW or more [31]. Since multiple stations are often transmitted from one site, it is also to be expected that there will be multiple interferers with equal or similar received power.

A demographic analysis of AM RFI in North America is given in [18]. Some of the key results have been reproduced in Figure 2-14. Here, the author determined the differential power

that would likely couple to drop wires with 40 and 50 dB balance. With a 40 dB balance, he found that almost –20 dBm of differential power can be induced on a UTP line within 1 km of a 50 kW AM transmitter; this corresponds to an rms voltage of approximately 31.6 mV. Up to 30 km away, –50 dBm of differential noise power could appear on the loop. By making some assumptions concerning the number of homes served by overhead drop wires, and the balance of those drop wires, it was estimated that up to 30% of homes could be at threat from an interference of – 50 dBm or higher. At the average house, the received powers of the six strongest interferers were within 10 dB of each other.

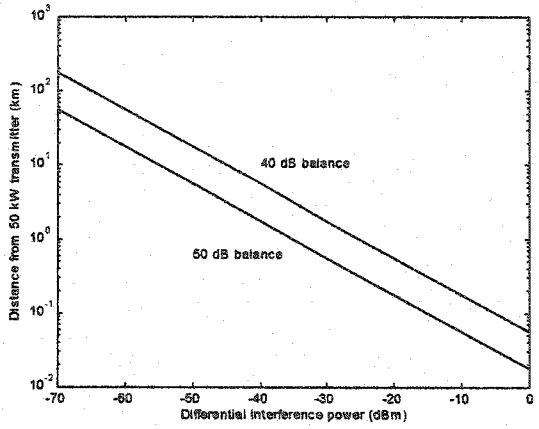

Figure 2-14: Interference power vs. proximity to AM transmitter

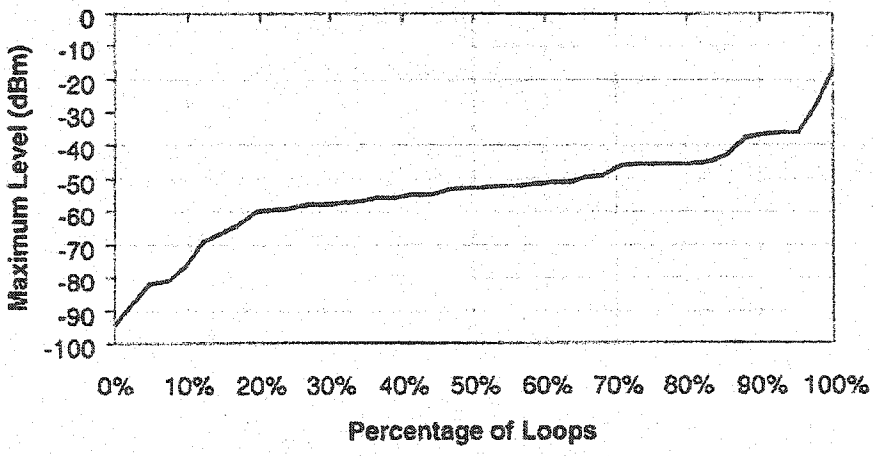

Figure 2-15: Cumulative Percentage of Loops with Differential AM Ingress Levels Below a Maximum Level

Measurements show that this demographic analysis is realistic. After performing tests on 42 loops in two Canadian cities, it was found that 50% of loops carried at least one interferer with a differential power of –53 dBm or higher [21]. The 90th and 99th percentile figures were –37 and –20 dBm respectively; these results are reproduced in Figure 2-15. Similar to the predictions given by [18], it was found that on average, the second most powerful interferer was only 1.3 dB weaker than the primary one. In all measurements, this ratio never exceeded 6 dB.

Amateur radio signals use 2.5 kHz slots in the nine international amateur bands which fall between 1 and 30 MHz [3, 9]:

a. 1.81-2.0 MHz;
b. 3.5-4.0 MHz;
c. 7.0-7.3 MHz;
d. 10.1-10.15 MHz;
e. 14-14.35 MHz;
f. 18.068-18.168 MHz;
g. 21-21.45 MHz;
h. 24.89-24.99 MHz; and
i. 28-29.7 MHz.

Interference from amateur radio is rarely encountered, but can be devastating. Since the transmitter may be located arbitrarily close to the twisted-pair cable, the induced differential RFI may be very high: 40 metres from a 400 W transmitter, up to 100 mV_{rms} of differential voltage may be observed on the loop (giving an average power of –10 dBm) [23]. Cancellation of these signals is complicated by their bursty nature and variable carrier frequencies. Furthermore, in order to prevent VDSL signals from interfering with amateur radio reception, the transmitted VDSL power spectral density will probably be limited to –80 dBm/Hz in the amateur bands. This PSD limit will render VDSL particularly susceptible to amateur RFI. In [3], it is estimated that over 80 dB of RFI suppression may be required to allow reliable VDSL operation in the presence of worst-case amateur radio signals.

While high-powered amateur radio signals are encountered only very rarely, small-scale noise from maritime and aeronautical HF radio, as well as shortwave radio, is almost always

present [34]. To date, there are no comprehensive field measurements or simulation models available for such noise sources [9]. However, it has been suggested that the VDSL simulation model should include several shortwave tones between 6 and 9 MHz, and at powers of -50 to -60 dBm [33]. Similarly, the T1E1.4 subcommittee is currently considering mobile HF radio signals, even though in most cases they arrive at the receiver with powers in the order of only -120 dBm [38]. For ADSL, it is also important to consider the effects of the LORAN-C navigation system, which employs 2 MW transmitters at 100 kHz.

The VDSL system requirements document defines three RFI threats for testing purposes [9]. Each one consists of 13 signals: 10 AM broadcast signals, one amateur radio signal, and two other noise sources which have yet to be characterized. All of the broadcast sources are 30% AM modulated, and use a Gaussian noise source bandlimited to 5 kHz to generate the modulating signal. The amateur radio signal is single sideband, and changes carrier frequency by at least 50 kHz every two minutes. Its modulating signal is a weighted noise designed to simulate the major attributes of speech, and its received differential power is -10 dBm.

The models range from a severe noise environment, with two 50 kW broadcast signals transmitted from a close distance, to a relatively benign suburban environment. The characteristics of the AM interferers are summarized in Table 2-1 [9]. A 60 dB cable balance is assumed at all broadcast frequencies, and the cable is assumed to be terminated with 100 Ω in differential mode (DM), and 50 Ω in common mode (CM).

Table 2-1: AM Radio Noise Threats for VDSL

Frequency (kHz)	Model 1		Model 2		Model 3	
	CM power (dBm)	DM power (dBm)	CM power (dBm)	DM power (dBm)	CM power (dBm)	DM power (dBm)
660	0	-60	+5	-55	0	-60
710	+30	-30	+10	-50	+10	-50
770	-10	-70	0	-60	-10	-70
1050	+5	-55	+10	-50	+5	-55
1130	+30	-30	+30	-30	+10	-50
1190	0	-60	0	-60	0	-60
1280	+5	-55	+5	-55	+5	-55
1330	0	-60	+5	-55	+5	-55
1480	-10	-70	-10	-70	-10	-70
1600	0	-60	0	-60	0	-60

The impact of RFI on VDSL can be quickly appreciated by analyzing the signal-to-interference ratios (SIRs) likely to exist on the cable. On a loop of length between 500 m and 1 km, the received power of a VDSL signal will fall between –50 and –20 dBm [50, 58]. In the presence of a powerful amateur interferer, the SIR on the loop could approach –40 dB; in the more likely case of severe broadcast AM ingress, SIRs as bad as – 30 dB may be expected.

2.8 Impulse Noise

Impulse noise can originate both inside and outside the communications system. Internal sources include the opening and closing of old mechanical relays at the central office, or the operation of ringing/busy signalling circuits over adjacent loops [15, 57, 67]. These unwanted signals usually couple to the loop of interest through familiar crosstalk mechanisms. NEXT-coupled impulse noise at the CO has been the focus of a number of previous studies [eg. 53, 68].

In VDSL, where the interface between optical and electrical signalling will in most cases occur outside the CO, the noise caused by the operation of old relays should not be as severe. External noise sources may prove to be more significant. These include such diverse sources as lightning, power system switching activity, fluorescent lights, arc welders and vehicle starter motors [16, 67].

A review of published information on impulse noise shows that it remains poorly understood [22, 57, 63]. This is primarily because the statistics of the noise vary widely with both the configuration and the location of the loop plant. The results of a loop survey can also be seriously skewed by the experimental setup; the equipment must be triggered to initiate storage of impulse events, and setting an appropriate trigger level is not trivial. The results of earlier surveys are also limited by the sampling rates and resolution of the analog-to-digital converters employed. However, since the advent of ISDN at least three major carriers have undertaken serious efforts to better understand impulse noise [15, 29, 30, 62]. Unfortunately, none of these studies has analyzed the noise above 5 MHz. A preliminary result given in [63] indicates that significant power density levels may be reached at frequencies as high as 10 MHz, so further study will be required.

In order to properly characterize impulse noise, at least four different attributes must be recorded and analyzed:

a. voltage magnitude;
b. burst shape (which determines the frequency spectrum);
c. burst duration; and
d. burst inter-arrival time.

A complete simulation model for impulse noise should include a separate probability distribution for each of these characteristics. Developing such simulations is very complicated, and is often avoided: at present, the European Telecommunications Standards Institute is using pulsed white noise for its testing, and other studies have simply used recordings of real impulse events [9]. It is reasonable to simplify the model such that it still exhibits the key characteristics of impulse noise, but with lower complexity. This approach is most easily implemented by selecting a "representative burst" and varying its amplitude according to a probability distribution for the peak received differential voltage. This method has been adopted for the simple impulse noise tests conducted in chapter 6, and is described in the following pages. All impulses are assumed to couple to the loop entirely in common mode, and then partially convert to differential mode as they travel towards the receiver.

2.8.1 Burst Inter-Arrival Time

The time between bursts varies widely with time of day, as well as with the day of the week. In general, peak activity tends to occur during workdays, but few other conclusions can be drawn. [15] shows a peak of about four pulses per minute at the noon-hour on weekdays. [57] and [68] indicate peak activity of six or seven bursts every ten minutes at 10 a.m., 4 p.m. and during the evening on weekdays. In contrast, a NYNEX loop survey found that up to 10 bursts/minute occurred at one central office [62]. A modified exponential probability density function for burst inter-arrival time is given in [29], but it too varies widely from location to location.

Since a simulation with only a few bursts every minute would have a very long run time, no probability distribution has been used for the inter-arrival time in this research. Instead, impulse events are spaced at even time intervals, with enough time between bursts to allow the noise canceller to recover from the previous impulse.

2.8.2 Burst Duration

The comprehensive survey carried out in [29] found that impulse length was the statistic which varied most from location to location. In some cities, it was almost always less than 50 μs, while in others it regularly exceeded 200 or 300 μs. Other studies have found that a significant number of impulses extend to 1 ms, and are sometimes as long as 3 ms [19, 15]. For simplicity, however, the simulations in Chapter 6 use a fixed pulse length of 100 μs.

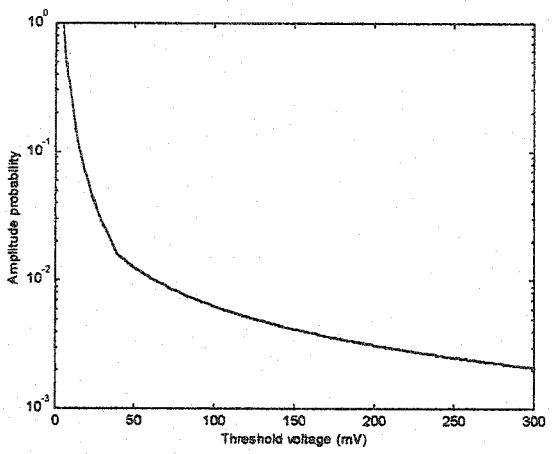

Figure 2-16: Amplitude probability distribution of impulse noise events

2.8.3 Peak Voltage

For impulses with peak differential voltages between 5 and 40 mV, a hyberbolic probability distribution is frequently used [68, 57]. The probability that the peak magnitude $|v|$ will exceed some threshold value v_t can be written as:

$$P(|v| > v_t) = \left(\frac{5}{v_t}\right)^2 \qquad 5 \le v_t \le 40 \ \text{(mV)} \tag{2.18}$$

Above 40 mV, the distribution function becomes:

$$P(|v| > v_t) = \frac{0.625}{v_t} \qquad v_t > 40 \ \text{(mV)} \tag{2.19}$$

The combined distribution is illustrated in Figure 2-16. In simulation, the transformation method is used to generate a sequence of peak voltages in accordance with this distribution.

2.8.4 Burst Shape (Frequency Spectrum)

Many different frequency spectra for impulse events have been reported. In most cases, impulse bursts have over half of their energy concentrated below 40 kHz [67]. At higher frequencies, different references report different characteristics, but it is fair to estimate that the power spectral density of a typical impulse event is highest at low frequencies, and falls by approximately 5 to 15 dB/decade as the frequency climbs [15, 29].

A simple symbolic pulse which exhibits a 5 dB/decade decline was proposed in [15]. Called the Cook pulse after its inventor, it is given by the equation:

$$p(t) = \begin{cases} V_p |t|^{-3/4} & t > 0 \\ 0 & t = 0 \\ -V_p |t|^{-3/4} & t < 0 \end{cases} \tag{2.20}$$

where V_p is a scaling factor to get the desired peak voltage and t is time. In fact, t is not a continuous-time sequence. The author had digital simulation in mind when the pulse was proposed, and his scaling factor V_p was designed to work properly when t was a discrete-time sequence spaced according to:

$$\ldots, -\frac{3}{2f_s}, -\frac{1}{2f_s}, \frac{1}{2f_s}, \frac{3}{2f_s}, \ldots \tag{2.21}$$

Here, f_s is the sampling frequency of the receiver, and is assumed to be twice the signalling rate of the system. This report does not make use of the formula for V_p, but it is comforting to know that the discrete-time nature of t puts some bounds on the size of the discontinuity in the centre of the burst. Figure 2-17 shows a plot of a Cook pulse.

This pulse has been criticized because it does not resemble a real impulse, both because of its discontinuity and its short duration [29, 30]. These are legitimate concerns, but the pulse does disrupt the receiver circuitry significantly, and is a reasonable test if the bit error rate (BER) is not being evaluated.

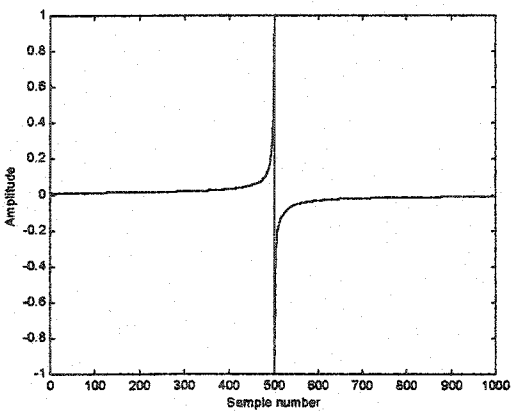

Figure 2-17: Shape of Cook pulse

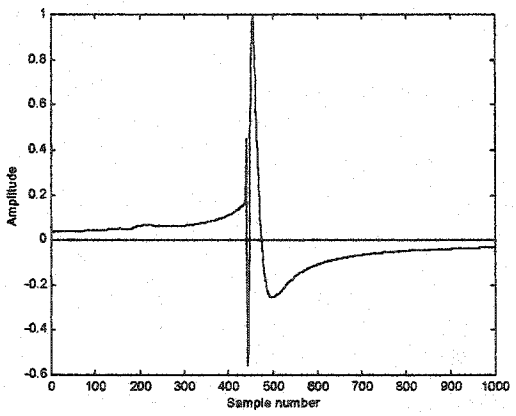

Figure 2-18: Shape of cook pulse after convolution with common-mode impulse response

2.8.5 Other Notes on Impulse Noise

The Cook pulse was originally intended to be applied directly to the differential input of the receiver. However, in testing noise cancellation, it is far more important to examine the worst-case scenario: impulses in both modes of transmission which are misaligned in time. Accordingly, the simulation model for this research merely uses the Cook pulse as a method to excite the common-mode channel.

The simulations in Chapter 6 are designed for short tests under worst-case conditions. Every 0.3 ms, a Cook pulse is coupled to the loop in common mode at the VDSL transmitter. The

burst propagates down the common-mode channel, and experiences some distortion from its original shape; Figure 2-18 shows a close-up view of a Cook pulse which has passed through a 1-km channel. A portion of the impulse couples to differential mode at the same location that the RFI does.

At the transmitter, the impulses are scaled such that they will appear at the receiver in differential mode in accordance with the amplitude distribution given by Figure 2-16. The scaling factor was determined empirically using both the common-mode and differential impulse responses, and the resulting pulses were invariably within 1 mV of the intended level.

This model is very simple, but still overstates the severity of the impulse noise threat. The simulations conducted in chapter 6 are not comprehensive. Their chief purpose is to determine whether or not the canceller worsens the impulse noise problem: there is an obvious danger that a cancellation device could introduce duplicate impulses into the differential signal. For this purpose, the impulse model described here is sufficient.

2.9 Crosstalk Noise

RFI and impulse noise can be classified as performance-limiting noise, whereas crosstalk noise is capacity-limiting [16]. RFI and impulse noise are intermittent in nature and geographically variable. Theoretically, their effects can be mitigated simply by increasing the transmitted power. On the other hand, crosstalk is almost always present, and at predictable levels. Increasing the transmitted power on all loops will not reduce the power of the crosstalk relative to the received information signal.

Crosstalk can be divided into two categories: near-end crosstalk (NEXT) and far-end crosstalk (FEXT). NEXT appears on a loop at the same end of the cable as the source of the interference. FEXT appears at the end of the cable opposite from the source of the interference. NEXT is usually the more severe of the two types, particularly at a central office, where there are many co-located transmitters. However, in both ADSL and VDSL, frequency-division duplexing and spectral management will be used in order to drastically reduce the level of NEXT [10]. For this reason, NEXT is not considered in this report.

Far-end crosstalk can be modelled by [9]:

$$PSD_{FEXT} = PSD_{disturber} \cdot |H(f)|^2 \cdot K_{FEXT} \cdot \left(\frac{N}{49}\right)^{0.6} \cdot d \cdot f^2 \qquad (2.22)$$

where $H(f)$ is the channel's differential transfer function, K_{FEXT} is the FEXT coupling constant, N is the number of disturbers, d is the length of the loop in feet, and f is the frequency in Hz. K_{FEXT} is a function of the twisted pair cable quality; a typical VDSL test value is 7.99×10^{-20} for AWG 24 cable [9]. A recent study has found that this model remains accurate at frequencies up to 40 MHz [64]. However, the K_{FEXT} values used by ANSI are too conservative, leading to errors of up to 4 dB. Instead, a value of 6.36×10^{-19} has been proposed for AWG 24.

If the information signal is assumed to have the same power spectral density as the disturbers, then the PSD of the information signal at the receiver is:

$$PSD_{sig} = PSD_{disturber} \cdot |H(f)|^2 . \qquad (2.23)$$

Therefore the ratio of the FEXT power to the received information signal power can be written as:

$$10\log_{10}\left(\frac{PSD_{FEXT}}{PSD_{sig}}\right) = 10\log_{10}\left(K_{FEXT} \cdot \left(\frac{N}{49}\right)^{0.6} \cdot d \cdot f^2\right) \qquad (2.24)$$

This quantity is plotted in Figure 2-19 for a 1-km AWG 26 loop. It can be observed that the FEXT power is at least 10 dB smaller than the received signal power at all frequencies below 20 MHz.

FEXT will not be further considered in this report. There are several reasons for this:

a. Figure 2-19 shows that the crosstalk power will be much smaller than the information signal power unless the loop is extremely badly balanced. An adaptive filter will not be able to achieve any cancellation in this case (see chapter 4). It is a far greater priority to

ensure that very strong RFI interferers can be cancelled, and that impulse noise will not be worsened by the canceller;

b. crosstalk levels will depend on the line code selected. This is an added complication for simulation; and

c. it is difficult for an adaptive filter to cancel out a large number of interferers which have approximately the same power spectra. This fact will also be demonstrated in chapter 4.

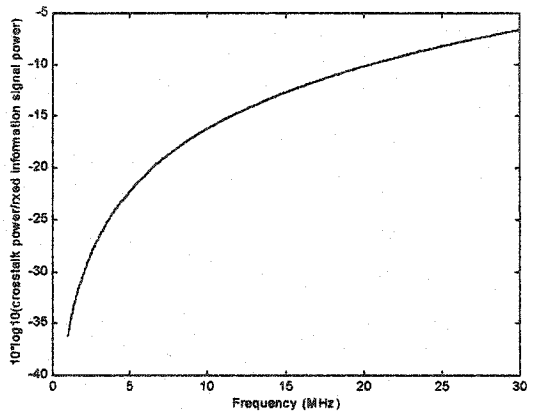

Figure 2-19: Ratio of crosstalk power to received information signal power

2.10 Summary

The ADSL and VDSL channels pose a number of significant impediments to the successful transmission and reception of high-bandwidth signals. The information signal is transmitted in differential mode, and experiences very bad attenuation. Meanwhile, a variety of strong noise sources, particularly AM radio signals and impulse noise, often couple to the DSL in common mode and experience very little attenuation. The gradual balance degradation of the loop with rising frequency causes the interference signals to transfer some of their energy to differential mode. Even more interference transfers to differential mode if the receiver's front end is poorly balanced. This interference can dwarf the original signal and make correct reception of high-data-rate signals impossible unless a noise canceller is used.

Noise cancellation is rendered difficult by possible differences in the speeds at which the differential and common-mode signals travel, and by different coupling paths to the loop. The

methods currently used for such cancellation depend on the line code used, and are briefly surveyed in the next chapter.

3. Some Notes on the Emerging VDSL Standard

This chapter provides a quick overview of some pertinent subjects under study at the T1E1.4 standards committee. In particular, the two contending line codes are described, focusing on their different approaches to the noise cancellation problem. The required analog-to-digital converter resolution is also discussed.

3.1 Discrete Multi-Tone vs. Quadrature Amplitude Modulation

The battle over the VDSL standard has broken the T1E1.4 subcommittee into two blocs. On one side are the proponents of quadrature amplitude modulation (QAM) or carrierless AM/PM (CAP), and on the other are the proponents of discrete multi-tone (DMT). At present, these two implementations are the only two line codes under serious consideration. A good assessment of their relative strengths and weaknesses is given in [59]; less objective analyses can be found in a variety of T1E1.4 documents, for example [37] and [60].

The QAM proposal can be viewed as an extension of the technology used in current voiceband modems (CAP is similar to a digital implementation of QAM, but uses FIR filters to eliminate the need for explicit modulation). The fact that QAM is already well understood is one of its main advantages, along with its relatively low power consumption and simple control circuitry [59]. However, a QAM VDSL system has a very high symbol rate and a wide frequency spectrum. The dispersive effects of the UTP channel are therefore severe, and may cause ISI spanning hundreds of symbols. As a result, a very long equalizer is required, and it must conduct many multiplication operations at high processing speeds. In an attempt to limit the attenuation and dispersion across the signal spectrum to a lower range, recent proposals have suggested that two frequency bands be used for two downstream QAM signals, each operating at half the signal rate [39]. This change allows the use of equalization-by-subchannel, which in general does not reduce the overall number of mathematical computations to be performed, but does slow down the equalizers' speed of operation. An extension of this idea to higher numbers of bands is presented in [6].

A QAM VDSL system will combat ingress noise with two devices: a wideband front-end noise canceller, and the equalizer. A decision feedback equalizer (DFE) can very nicely compensate for any narrowband interference which remains after the wideband canceller. The feedforward portion of the equalizer is used to create a deep notch at the frequency of the RF interferer, and the signal energy lost in the "notched-out" band of frequencies is then restored by the decision-feedback portion of the equalizer [28]. However, if the feedforward equalizer is unable to remove all of the RFI, many decision errors may result and cause the DFE to introduce additional errors into the detected symbol stream.

A recent field trial of a VDSL QAM modem without a front-end noise canceller found that it could sustain 25.92 Mbps at signal-to-interference (SIR) ratios as low as –4 to –6 dB, provided that the interference consisted solely of narrowband AM radio signals [35]; 4 dB of extra margin was still available for other noise sources. As the interference power increased, driving the SIR below –6 dB, the number of bits available to describe the information signal also fell. This introduced additional quantization noise, reducing the SNR margin available to combat other noise and eventually forcing the modem to communicate at a lower bit rate.

It was previously established that the equalizer length per subchannel can be reduced by increasing the number of QAM subchannels and decreasing the symbol rate in each of those subchannels. Taking this argument to the limit results in a system with a very large number of subchannels, each with a one-tap equalizer. This configuration is the basic idea behind the other competing line code, DMT.

DMT systems with thousands of subchannels have been proposed. In practice, modulation to this large number of carrier frequencies is carried out digitally using an inverse fast Fourier transform (IFFT); most current VDSL designs use a 4096-point transform. Because the input data is manipulated such that the outputs of the IFFT bank are real, very efficient implementations of the IFFT are possible. Furthermore, because the system symbol rate is much slower than for a comparable QAM system, the number of computations required per unit time is fewer, and the system has greater immunity to impulse noise [60]. Another major advantage of DMT is that the number of information bits assigned to each subcarrier (ie. the size of the QAM constellation at each subcarrier) can be independently varied according to the channel quality. Its disadvantages include longer delay, a high peak-to-average ratio (PAR) which increases the system's

susceptibility to clipping, greater sensitivity to timing jitter, more complex control circuitry and higher power consumption in the modem's analog sections [59].

A DMT VDSL system will deal with ingress noise using a front-end wideband noise canceller, and by adjusting the bit allocation of any subchannel which contains a narrowband interferer to zero. Originally, this ability to notch out certain bands was seen as a major advantage for DMT.

Unfortunately, complications have arisen in two of the major promises of DMT: simple equalization and easy cancellation of narrowband interference. First, since many symbols are transmitted at different carrier frequencies, dispersion is still a problem and a time-domain equalizer at the front end of the modem is required to simplify synchronization. Second, the spectra of adjacent DMT bands cross at only –3 dB, and a band's first sidelobe overlaps the main lobe 2 bands away only 13 dB down from the peak power density. As a result, one narrowband interferer can corrupt many subchannels. Even worse, if the narrowband interferer does not have a stable carrier frequency, the bit allocation algorithm may have difficulty adapting in a timely manner as the interferer's centre frequency shifts between bands. Three different solutions have been proposed to solve this problem [13]:

a. use the time-domain equalizer to insert a notch at the interferer frequency;
b. reduce the subchannel sidelobes by applying a windowing filter to the DFT filter bank at the receiver; or
c. leave a number of subchannels unused in the frequency regions most susceptible to RFI. When an interferer is present in a neighbouring band, the modem will be able to calculate its frequency, amplitude and phase by cross-referencing the interference signals received in two of the nulled bands. With this information, it is possible to generate a set of signals which mimic the effect of the interferer in neighbouring bands, and then subtract those signals from the adjacent bands (see [61] for further details).

All of these improvements come at the cost of additional receiver complexity, although the first two methods are easy to design and implement.

No published information on the ability of deployed DMT modems to operate in heavy interference environments could be found. DMT modem designers still hotly contest that the

method is more susceptible to RFI [2, 60]. Unfortunately, much of their argument is based on the relative performance of QAM and DMT without any noise cancellation: in this case, DMT performs better in strong RFI because it limits the errors to a number of tones, while the entire QAM bandwidth would be affected. This argument neglects the ability of the QAM equalizer to suppress RFI without any extra special circuitry. Altogether, DMT is unlikely to perform as well as QAM in the presence of RFI [59]. It therefore has more to gain from good front-end interference cancellation.

3.2 Analog-to-Digital Converters for VDSL

In current QAM and DMT modems, the achievable bit rate is limited by the number of A/D bits available to characterize the desired (information) signal after front-end cancellation has been carried out. There are two possible solutions to this problem:

a. implement a front-end noise canceller at least partly in the analog domain [12, 13, 35]; or
b. increase the A/D resolution.

The second of these two options is easier, but is limited by the maximum sampling speed of current analog-to-digital converters (ADCs). In practice, both solutions must be considered. The following pages quantify the conditions under which digital noise cancellation can be conducted without adversely affecting the rest of the modem receiver, and identify the noise sources which must be cancelled using analog circuits.

Given the very wide dynamic range of a received VDSL signal, the selection of an appropriate analog-to-digital converter is key to the success of the technology. At present, the state-of-the-art high-speed, high-resolution ADC has 14-bit resolution and can operate at 50 - 100 MSPS [17]. Unfortunately, this device consumes far too much power to meet the design objectives of VDSL. Meeting this performance level with lower power consumption is not likely in the near term. ADC development has progressed much slower than the development of other digital devices: in the past eight years, only 1.5 output bits have been added at any given sampling frequency [65, 66]. By the time of commercial VDSL deployment, it is estimated that low-power 12-bit, 35 MSPS devices will be available [17].

3.2.1 Acceptable Quantization Noise Floor

In an ideal uniform ADC, the only source of noise is quantization error. If the ADC input signal has a wide frequency spectrum, this error can be modelled as a uniform distribution between $-\delta/2$ and $\delta/2$, where δ is the size of the elemental quantization step. In the frequency domain, noise with this characteristic appears as a white noise with variance [11]:

$$\sigma_q^2 = \frac{\delta^2}{12} \tag{3.1}$$

The one-sided PSD of the quantization noise is given by:

$$S_{\sigma_q^2}(f) = \begin{cases} \frac{\delta^2}{12} \times \frac{2}{f_s} & 0 \le f \le \frac{f_s}{2} \\ 0 & \text{otherwise} \end{cases} \tag{3.2}$$

where f_s is the sampling frequency. Using equation (3.1) and assuming that the ADC input signal is a full-scale sinusoid, it is possible to show that the signal-to-quantization noise ratio (SQNR) is [40]:

$$SQNR = 6.02N + 1.76 + 10\log_{10}\left(\frac{f_s}{2 \cdot BW}\right) \quad \text{(dB)} \tag{3.3}$$

where N is the number of output bits and BW is the bandwidth of the input signal. In a practical device, other noise sources will exist, including circuit noise, aperture jitter and comparator ambiguity. These additional degradations typically reduce the effective number of output bits by about 0.6 - 1.5 bits from the nominal value N [66].

The analysis of A/D requirements in VDSL is best started by looking at the thermal noise floor, which is presumed to be at −140 dBm/Hz. Given that the transmit PSD will be constrained to be below −60 dBm/Hz, on longer loops the received signal will only achieve a positive noise margin over the AWGN floor across a limited frequency range. See Table 3-1 for examples of actual figures, assuming underground cables with no bridged taps and no crosstalk [9, 50].

Table 3-1: Positive Noise Margin on Long VDSL Loops

Loop	Positive Noise Margin achieved in frequency range:
1 km AWG 26	1.1 - 8 MHz
1.5 km AWG 26	1.1 - 4 MHz
1 km AWG 24	1.1 - 12 MHz
1.5 km AWG 24	1.1 - 7 MHz

Since service over long loops is desired, the −140 dBm/Hz noise floor should not be raised significantly in the process of quantization. [36] suggests that a rise of 0.5 dB is tolerable. Given that all measurements assume a 100 Ω termination, the acceptable quantization noise power density is therefore:

$$10\log_{10}\left(S_{\sigma_q^2}(f)\right) = 10\log_{10}\left[10^{\frac{-139.5+20-30}{10}} - 10^{\frac{-140+20-30}{10}}\right] - 20 + 30 \qquad (3.4)$$
$$= -149.14\ \text{dBm/Hz}$$

Substitution of this number into equation (3.2) yields a step size of δ = 160 μV for an ideal 12-bit, 35-MHz ADC, and gives a maximum range of [−328, 328] mV (or [−232, 232] mV_{rms}).

3.2.2 ADC Considerations for Digital Noise Cancellation

Dynamic range calculations for a variety of loops are shown in Table 3-2. The ADC numbers in the table assume an ideal 12-bit device operating at 35 MHz; the 17.5 MHz upper frequency limit for the 500-metre loops was chosen to meet the Nyquist criterion for such a device.

The required upper limit of the ADC's dynamic range may depend on several factors:

a. the transmit power leakage (if digital filters are to be used to provide spectral isolation between the transmit and receive bands in a FDD system);
b. the power of the ingress RFI;
c. the power of the received information signal; and

d. the probability of clipping deemed acceptable by the system designer.

Of these, the transmit power leakage would be by far the most significant. Trans-hybrid attenuation is usually less than 12 dB [51], and the maximum transmit power allowed on the loop is 11.5 dBm. Accounting for the hybrid's insertion loss, the transmitter will operate at about 17.5 dBm, giving a leakage power in the vicinity of 5.5 dBm. From equation (3.3), the SQNR of an ideal 12-bit, 35 MHz ADC whose input signal has bandwidth $f_s/2$ is 74 dB. Referring to the desired A/D quantization noise powers given in Table 3-2, it is clear that a VDSL system will not be able to incorporate digital FDD filtering without increasing the quantization noise considerably, thus reducing the bandwidth available for transmission. To preserve the thermal noise floor, this report will assume that analog duplexing filters are used; such filters are favoured by the proponents of single-carrier modulation for VDSL [35, 36, 50, 51].

Table 3-2: Sample ADC Calculations for VDSL (ideal 12-bit, 35 MHz ADC)

Loop	Frequency band	Thermal noise power	Lower limit of ideal ADC range	Upper limit of ideal ADC range
500 m AWG 26	1 - 17.5 MHz	-67.8 dBm	-77.0 dBm	-2.7 dBm
1 km AWG 26	1 - 8 MHz	-71.5 dBm	-80.7 dBm	-2.7 dBm
1.5 km AWG 26	1 - 4 MHz	-75.2 dBm	-84.4 dBm	-2.7 dBm
500 m AWG 24	1 - 17.5 MHz	-67.8 dBm	-77.0 dBm	-2.7 dBm
1 km AWG 24	1 - 12 MHz	-69.6 dBm	-78.7 dBm	-2.7 dBm
1.5 km AWG 24	1 - 7 MHz	-72.2 dBm	-81.4 dBm	-2.7 dBm

The next most powerful signals which must be handled are amateur radio transmissions, for which a worst-case power level of −10 dBm has been assigned in the VDSL requirements document [9]. Table 3-2 shows that such a signal lies at the outside limit of an ideal 12-bit ADC's capabilities: the voltage of the interferer is 100 mV_{rms} before PAR is considered, while the maximum level of the ideal ADC is 231 mV_{rms}. Unfortunately, a real 12-bit ADC will likely have only 10.5 to 11 effective bits, and so the effective top of its dynamic range will likely rest somewhere between −12 and −9 dBm. This is probably not crippling: amateur radio transmissions are rarely detected in field measurements, and are extremely rare at −10 dBm levels [35].

The largest broadcast AM signals required for VDSL testing have −30 dBm differential power, although a −20 dBm interferer represents a better 99th percentile threat. Even if there are

several –30 dBm interferers, the total received interference power is unlikely to exceed –24 dBm. On a 500-metre AWG 26 loop, the information signal will be comparable, at about –23 dBm. This upper limit gives the ADC just enough extra capacity to account for small impulsive events and the peak-to-average ratios of the information signal and interferers. For a 10^{-7} probability of clipping, the PAR for DMT is 14.6 dB, assuming that the transmit and receive signals have the same PAR [51]; for single-carrier modulation, the PAR will be at least 6 dB lower than the DMT figure [36]. Considering these numbers, a digital canceller in a QAM system operating over a long loop should be able to handle even the worst-case –20 dBm interferer with only a small rise in the noise floor.

When the loop is shorter than 500 metres, the signal power will be so high that the noise floor will not be a major concern. In these situations, even the largest RFI interferers can probably be cancelled digitally without affecting the capacity of the communications link.

The main drawback of shifting the ADC range as high as possible is that the ratio of the received information signal's power to the quantization noise power will drop, particularly for long loops. However, the thermal noise floor is unavoidable, and extending the dynamic range more than 18 - 20 dB below it is of little benefit. In general, a good rule of thumb would be to position the bottom of the ADC's range 20 dB below the thermal noise floor whenever the interference environment permitted it. This is a sufficient buffer against excessive quantization noise arising from A/D conversion, and against finite-precision errors introduced by the noise canceller.

Altogether, in differential mode a 12-bit ADC will be sufficient to represent all signals except transmit power leakage, extremely strong RFI and large impulse events without raising the thermal noise floor appreciably. On the other hand, the common-mode signal could have a total power 40 to 50 dB higher than its differential counterpart, so there is no chance that the thermal noise can be represented digitally. However, this is not a serious problem, as the differential and common-mode SQNRs will be similar. Simulation results in Chapter 6 will show that a 12-bit ADC can be used for the common-mode input with no adverse effects.

3.3 Summary

The two line codes in contention for VDSL both have some inherent noise cancellation capabilities. However, at extremely low SIRs, a QAM modem's equalizer may have difficulty

achieving reliable steady-state operation. Similarly, a DMT modem will probably not be able to fully compensate for strong interferers which corrupt many tones. If the interferer frequency varies often (as it might for an application such as FSK over HF radio), DMT also may not be able to adapt the constellations on its tones in a timely manner.

Considering that VDSL modems will likely use ADCs of at least 12-bit resolution, almost all interference signals can be represented digitally without raising the thermal noise floor with quantization noise. Since the common-mode signal is an excellent reference for the RFI which has corrupted the differential signal, a digital front-end canceller can be implemented. The front-end device will be especially valuable for DMT, as it will drastically reduce the number of tones that are badly corrupted by each narrowband interferer. In a QAM system, the canceller should simplify equalizer startup in high-noise environments, and will remove the requirement for the FFE to notch out significant bands of signal energy.

From this discussion, two disadvantages of a digital canceller are immediately obvious. The first is the very fast speed at which the canceller must operate, and the second is its inability to handle powerful amateur AM interference without reducing the bandwidth over which the system can achieve positive noise margin. These issues will be further addressed when the canceller architecture is discussed in Chapter 5. First, it is necessary to review some of the noise cancellation techniques that have been developed for communications applications.

4. Previous Work in the Field of Noise Cancellation

While noise cancellation for twisted-pair loops has only attracted interest with the advent of the digital subscriber line, the field of research is over forty years old and cancellers have been designed for many other applications. This section will explore and comment on some of the ideas which have been advanced in recent years. First, however, it is necessary to review the basic principles behind adaptive filter theory.

4.1 Adaptive Noise Cancellation

The noise cancellation problem for TP cable is complicated by two principal factors:

a. the receiver equipment cannot determine the differential and common-mode channel characteristics, or the points at which external noise signals couple to the loop; and
b. the statistics of the noise signal will likely be time-varying. This is particularly true when amateur radio is involved, as its transmissions are bursty and its carrier frequencies are variable [12].

It is therefore not possible to design a fixed filter which can process the common-mode input into an accurate estimate of the differential noise component. An adaptive noise canceller is required.

The general structure of an adaptive noise canceller for communication over twisted-pair wire is shown in Figure 4-1. The canceller is assumed to be digital, where j is a discrete-time index. The common-mode signal $\mathbf{X}_j$ is the reference noise for the system, and is assumed to be uncorrelated with the information signal. Because cable imbalance causes some conversion of the differential information signal to common mode, some correlation will in fact exist; this issue will be addressed later. The differential input $\mathbf{D}_j$ is called the "desired signal", an unfortunate piece of terminology which relates to other uses of adaptive filters. In this case $\mathbf{D}_j$ is actually a combination of the distorted information signal $\mathbf{S}_j$ and the induced noise $\mathbf{N}_j$. $\mathbf{X}_j$ and $\mathbf{N}_j$ are correlated in some unknown manner, so the tap weights $\mathbf{W}_j$ of the adaptive filter (assumed to be

FIR) may be adjusted such that the filter's output $\mathbf{Y}_j$ closely approximates $\mathbf{N}_j$. If this operation is performed correctly, the output signal $\mathbf{E}_j$ will contain much less noise than the originally received signal $\mathbf{D}_j$.

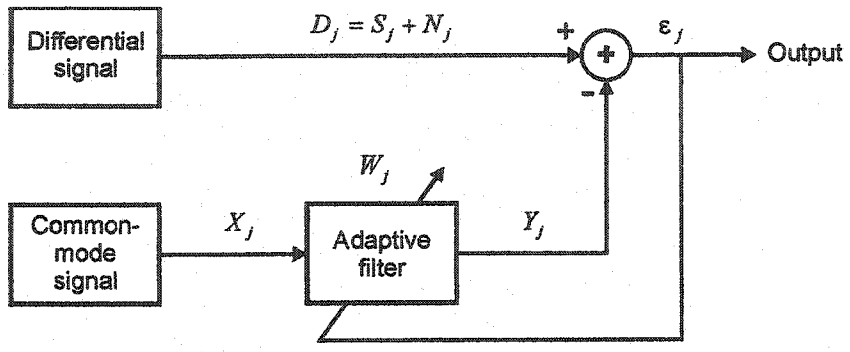

Figure 4-1: Adaptive noise canceller for communications over UTP cable

Adaptive noise cancellers must be designed very carefully to ensure that they do not actually degrade the overall performance of the communications system. The following pages describe a suitable algorithm, its capabilities and a number of implementation issues.

4.1.1 The Least-Mean-Square (LMS) Algorithm

In order to adjust the adaptive filter's tap weights, it is necessary to develop a cost function which can be used to gauge the progress of the adjustments. Most frequently, the cost function is taken to be the mean-squared error between the desired response $\mathbf{D}_j$ and the filter output $\mathbf{Y}_j$. It is a function of the tap weights. If the system inputs $\mathbf{X}_j$ and $\mathbf{D}_j$ are wide-sense stationary and the filter has M taps, the cost function may be viewed as an M-dimensional paraboloid with a uniquely defined minimum point [27]. The tap weights which correspond to this minimum point on this "error-performance surface" give the optimum estimate of $\mathbf{D}_j$ which can be generated using $\mathbf{X}_j$ (the Wiener solution).

In theory, the Wiener solution can be found analytically by solving a system of M equations in M unknowns (the Wiener-Hopf equations). In practice, M is often a very large number, and this approach can pose an intolerable computational burden. Furthermore, the

statistical properties of $\mathbf{X}_j$ and $\mathbf{D}_j$ usually vary slowly with time. In this non-stationary environment, the minimum point of the error-performance surface shifts its location as a function of time. Accordingly, it is necessary to find a method which can track the movement of the minimum point, and also reduce the mathematical complexity required to solve for the filter coefficients [46].

Assuming wide-sense stationary statistics, the necessary mathematical complexity can be dramatically reduced by using the method of steepest descent to solve the Wiener-Hopf equations recursively rather than analytically. This method implements a deterministic feedback system which uses iterative computations of a gradient vector to converge to the minimum point [27]. However, it assumes knowledge of the autocorrelation matrix of $\mathbf{X}_j$ ($\mathbf{R}_X$), and of the cross-correlation vector between $\mathbf{X}_j$ and $\mathbf{D}_j$ ($\mathbf{P}$). These calculations are usually computationally intensive.

The least-mean-square (LMS) algorithm instead estimates an instantaneous error gradient by assuming that ε_j^2, the square of a single error sample, is equal to the mean-squared error [70]. This allows the tap weights to be calculated using simple estimates for $\mathbf{R}_X$ and $\mathbf{P}$, greatly reducing the system complexity. However, the algorithm no longer converges continually towards the minimum point. It follows a noisy path, occasionally steering in the wrong direction. For this reason, it is called a *stochastic gradient* algorithm, and it has a slower rate of convergence than the method of steepest descent [27].

Altogether, the LMS algorithm has been selected for this project because:

a. it is the simplest adaptive system to implement for simulation purposes. Its complexity increases linearly with the number of tap weights, M [27]; and
b. in non-stationary environments, it can reliably track the bottom of the error-performance surface, provided that the input data statistics vary slowly compared to the learning rate of the system [27]. By contrast, the other prominent *O(M)* algorithm, fast recursive-least-squares, uses a state model to describe the system. Great care must be taken to ensure that the model chosen allows reasonable tracking performance, and this additional complication is undesirable at this early stage of investigation.

The LMS algorithm will now be developed mathematically, following the analysis given in [70]. First, the vectors for the common-mode input and the adaptive filter tap weights are defined as:

$$\mathbf{X}_j \equiv \begin{Bmatrix} x_j \\ x_{j-1} \\ \vdots \\ x_{j-M+1} \end{Bmatrix}, \text{ and } \quad \mathbf{W} \equiv \begin{Bmatrix} w_1 \\ w_2 \\ \vdots \\ w_M \end{Bmatrix}. \tag{4.1}$$

respectively. At this stage of the analysis, $\mathbf{X}_j$ and $\mathbf{D}_j$ are assumed to be stationary. It is also assumed that the tap weight vector is to be solved for analytically, and therefore $\mathbf{W}$ is not yet time-dependent. Following this convention, the filter output at time j is simply the inner product of these two matrices, and the system output at time j is:

$$\varepsilon_j = d_j - y_j = d_j - \mathbf{W}^T\mathbf{X}_j \tag{4.2}$$

where T denotes matrix transposition (all matrix entries are assumed to be real). Knowing that $\mathbf{W}^T\mathbf{X}_j = \mathbf{X}_j^T\mathbf{W}$, the mean-squared error between the filter output and the desired response can be written as:

$$MSE_j = E\{\varepsilon_j^2\} = E\{d_j^2\} - 2\cdot\mathbf{P}^T\mathbf{W} + \mathbf{W}^T\mathbf{R}_X\mathbf{W} \tag{4.3}$$

where:

$$\mathbf{P} \equiv E\{d_j\mathbf{X}_j\} = E\begin{Bmatrix} d_j x_j \\ d_j x_{j-1} \\ \vdots \\ d_j x_{j-M+1} \end{Bmatrix}, \text{ and} \tag{4.4}$$

$$\mathbf{R}_X \equiv E\{\mathbf{X}_j\mathbf{X}_j^T\} = E\begin{bmatrix} x_j x_j & x_j x_{j-1} & x_j x_{j-2} & \cdots & x_j x_{j-M+1} \\ x_{j-1}x_j & x_{j-1}x_{j-1} & x_{j-1}x_{j-2} & \cdots & x_{j-1}x_{j-M+1} \\ x_{j-2}x_j & x_{j-2}x_{j-1} & x_{j-2}x_{j-2} & \cdots & x_{j-2}x_{j-M+1} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ x_{j-M+1}x_j & x_{j-M+1}x_{j-1} & x_{j-M+1}x_{j-2} & \cdots & x_{j-M+1}x_{j-M+1} \end{bmatrix}. \tag{4.5}$$

To find the minimum mean square error (MMSE), the gradient of the error function is required. From (4.3):

$$\nabla \equiv \begin{Bmatrix} \frac{\partial E\{\varepsilon_j^2\}}{\partial w_1} \\ \vdots \\ \frac{\partial E\{\varepsilon_j^2\}}{\partial w_M} \end{Bmatrix} = -2 \cdot \mathbf{P} + 2 \cdot \mathbf{R}_X \mathbf{W} \tag{4.6}$$

The tap weights resulting in the MMSE are found by setting $\nabla = 0$. For almost all engineering applications, $\mathbf{R}_X$ is positive definite [27]. This implies that it is almost always non-singular and therefore invertible. Accordingly, for $\mathbf{X}_j$ and $\mathbf{D}_j$ real and stationary, the MMSE will occur for the tap weights:

$$\mathbf{W}_{Wiener} = \mathbf{R}_X^{-1} \mathbf{P}. \tag{4.7}$$

This equation is the matrix form for the *M*-dimensional system of Wiener-Hopf equations. As mentioned previously, these equations are difficult to solve because they require the evaluation of correlation functions and the inversion of a very large matrix. To eliminate the matrix inversion, the tap weight vector can be solved for iteratively using the method of steepest descent:

$$\mathbf{W}_{j+1} = \mathbf{W}_j - \mu \nabla_j \tag{4.8}$$

where μ, frequently called the *step size*, is a parameter that controls the algorithm's stability and rate of convergence. Referring to equation (4.6), the gradient calculated at iteration *j* is given by:

$$\nabla_j = -2 \cdot \mathbf{P} + 2 \cdot \mathbf{R}_X \mathbf{W}_j \tag{4.9}$$

The LMS algorithm removes the expectation operations from the method of steepest descent by assuming that ε_j^2 is a good instantaneous estimate of $E\{\varepsilon_j^2\}$. This gives a gradient estimate:

$$\hat{\nabla}_j = \begin{Bmatrix} \frac{\partial \varepsilon_j^2}{\partial w_{1j}} \\ \vdots \\ \frac{\partial \varepsilon_j^2}{\partial w_{Mj}} \end{Bmatrix} = 2\varepsilon_j \cdot \begin{Bmatrix} \frac{\partial \varepsilon_j}{\partial w_{1j}} \\ \vdots \\ \frac{\partial \varepsilon_j}{\partial w_{Mj}} \end{Bmatrix} = 2\varepsilon_j \cdot \begin{Bmatrix} \frac{\partial \left(d_j - \mathbf{W}^T \mathbf{X}_j \right)}{\partial w_{1j}} \\ \vdots \\ \frac{\partial \left(d_j - \mathbf{W}^T \mathbf{X}_j \right)}{\partial w_{Mj}} \end{Bmatrix} = -2\varepsilon_j \mathbf{X}_j \qquad (4.10)$$

Combining (4.10) with (4.8) gives the equation for filter tap updates:

$$\mathbf{W}_{j+1} = \mathbf{W}_j + 2\mu\varepsilon_j \mathbf{X}_j \qquad (4.11)$$

This is the LMS algorithm, also called the Widrow-Hoff algorithm after its inventors.

The step-size parameter μ obviously must be chosen appropriately. A very small μ can lead to an unacceptably slow rate of convergence, and a high μ can cause the entire algorithm to become unstable. Convergence analysis of the LMS algorithm is still an active area of research [27]. However, it is possible to quickly develop a conservative estimate of the allowable upper bound on μ. The analysis in the following paragraphs follows the development given in [71].

Rather than converging to the exact Wiener solution, the LMS algorithm approaches and then fluctuates around it. However, the conditions for initial convergence of the LMS algorithm are the same as for the initial convergence of the method of steepest descent [70]. The convergence of the method of steepest descent can be examined by analyzing the MSE of the method at each iteration. Combining equations (4.3) and (4.7), the minimum mean-square error at the bottom of the error-performance surface simplifies to:

$$MSE_{\min} = E\left\{d_j^2\right\} - \mathbf{P}^T \mathbf{W}_{Wiener} \qquad (4.12)$$

It is necessary to measure the amount of additional MSE that is incurred over the minimum. Once again considering the method of steepest descent, (4.3) becomes:

$$MSE_j = MSE_{\min} + \left(\mathbf{W}_j - \mathbf{W}_{Wiener}\right)^T \mathbf{R}_X \left(\mathbf{W}_j - \mathbf{W}_{Wiener}\right) \qquad (4.13)$$

Now, since $\mathbf{R}_X$ is positive definite, a unitary similarity transformation may be used:

$$\mathbf{R}_X = \mathbf{Q}_X \mathbf{L}_X \mathbf{Q}_X^{-1} = \mathbf{Q}_X \mathbf{L}_X \mathbf{Q}_X^T \tag{4.14}$$

where $\mathbf{L}_X$ is the diagonal eigenvalue matrix of $\mathbf{R}_X$:

$$\mathbf{L}_X = diag[\lambda_1, \lambda_2, ..., \lambda_M] \tag{4.15}$$

and $\mathbf{Q}_X$ is the $M \times M$ matrix which contains all the associated eigenvectors as column vectors, each normalized to a squared length of 1. All the eigenvalues of $\mathbf{R}_X$ are real and non-negative [27].

Defining a new matrix:

$$\mathbf{V}_j \equiv \mathbf{Q}_X^T \left(\mathbf{W}_j - \mathbf{W}_{Wiener} \right), \tag{4.16}$$

(4.13) can be rewritten as:

$$MSE_j = MSE_{\min} + \mathbf{V}_j^T \mathbf{L}_X \mathbf{V}_j \tag{4.17}$$

and the error gradient with respect to $\mathbf{V}_j$ can be written as $\nabla_j = 2\mathbf{L}_X \mathbf{V}_j$ (4.18). Using (4.8), (4.16) and (4.18), the method of steepest descent can be reformulated in terms of $\mathbf{V}_j$:

$$\mathbf{V}_{j+1} = \mathbf{V}_j - 2\mu \mathbf{L}_X \mathbf{V}_j = (\mathbf{I} - 2\mu \mathbf{L}_X)\mathbf{V}_j \tag{4.19}$$

With repeated applications of (4.19) it is clear that convergence will only be achieved if the maximum magnitude of the elements in the diagonal matrix $(\mathbf{I} - 2\mu \mathbf{L}_X)$ is less than 1. This is achieved using a μ which satisfies:

$$0 < \mu < \frac{1}{\lambda_{\max}} \tag{4.20}$$

where $\lambda_{\max}$ is the largest eigenvalue of $\mathbf{R}_X$. The gradient estimates accumulated over time by the LMS algorithm are close enough to the true gradient that (4.20) is a sufficient condition for

convergence of the LMS algorithm, provided that $\mathbf{R}_X$ is well-conditioned [71]. A more detailed analysis presented in [27] gives identical results.

In practice, the eigenvalues of $\mathbf{R}_X$ are not likely to be known, and it is necessary to find an approximation to (4.20). Given that the trace of a matrix is the sum of the elements on its main diagonal:

$$\mathrm{tr}[\mathbf{R}_X] = \sum_{i=1}^{M} r_X(i,i) \tag{4.21}$$

a result that holds for all square matrices is [71]:

$$\mathrm{tr}[\mathbf{R}_X] = \sum_{i=1}^{M} \lambda_i \,. \tag{4.22}$$

If $\mathbf{X}_j$ is wide-sense stationary, $\mathbf{R}_X$ is Toeplitz with all elements on the main diagonal equal to $R_X(0)$, the autocorrelation of $\mathbf{X}_j$ with no time shift. $R_X(0)$, in turn, is equal to the average power of the reference noise stored in the M delays of the adaptive filter:

$$\mathrm{tr}[\mathbf{R}_X] = M \cdot R_X(0) = \sum_{i=0}^{M-1} E\left\{\left|x_{j-i}\right|^2\right\} \tag{4.23}$$

This quantity is referred to as the "tap-input power", and is easily calculated by the noise canceller. Referring to (4.22) and (4.23), a conservative estimate of the allowable bounds on μ can be given by:

$$0 < \mu < \frac{1}{(\text{tap - input power})} < \frac{1}{\lambda_{\max}} \tag{4.24}$$

Another performance measure of interest is a unitless quantity called "misadjustment". This is simply the ratio of the excess MSE caused by stochastic gradient noise to the optimum MSE of the Wiener solution:

$$\Pi = \frac{\text{average excess MSE}}{MSE_{\min}} \tag{4.25}$$

Without proof, the misadjustment of the LMS algorithm can be stated as [71]:

$$\Pi = \mu \cdot \mathrm{tr}[\mathbf{R}_X] = \mu \cdot (\text{tap - input power}) \tag{4.26}$$

This result agrees very well with experimental results for values of Π under 0.25 [27, 71]. Equation (4.26) implies that when the inputs are wide-sense stationary and infinite-precision arithmetic is used, a low value of μ will yield a lower steady-state MSE.

4.1.2 Required Number of Adaptive Filter Taps

The total time delay of the filter should be selected as the reciprocal of the desired filter frequency resolution [70]. A suitable value can therefore be determined using:

$$\text{Number of taps} = \frac{\text{A/D sampling frequency}}{\text{Desired filter frequency resolution}} \tag{4.27}$$

Another factor which must be considered for the case of twisted-pair noise cancellation is the possible phase separation between the differential and common-mode signals. If they diverge at $0.31c$ for the entire length of a 1.5-km loop, the differential wavefront will arrive at the receiver 7 μs before the common-mode reference signal. At a sampling frequency of 35 MHz, this implies that the differential signal should be delayed by at least 245 samples, and that the adaptive filter should have 245 taps. Fortunately, RFI will usually transfer to differential mode close to the customer premises, and even if it does transfer at great distance from the receiver the differential attenuation of the loop will greatly mitigate its severity. These considerations should allow a considerable reduction in the number of filter taps required for the majority of loops.

4.1.3 Finite Precision Effects

The adaptive filter of Figure 4-1 suffers quantization errors from two sources [11]:

a. the A/D quantization errors for the reference noise $\mathbf{X}_j$ and the desired signal $\mathbf{D}_j$; and

b. the quantization error arising from the fact that the filter coefficients and output are only represented by a finite number of bits.

The first of these was discussed in Chapter 3, and is uniformly distributed with zero mean and variance:

$$\sigma_q^2 = \frac{\delta^2}{12} \tag{4.28}$$

where δ is the voltage difference across one least significant bit of the A/D converter. The most significant cause of the second type of error is the repeated multiplications carried out for the LMS algorithm. Each multiplication results in a double-precision number which must then be either truncated or rounded back to some lesser precision. With successive iterations, these errors can compound quickly, and since they often do not have zero mean they can accumulate until overflow occurs; after overflow, the filter must reset itself and re-converge [11]. Overflow may take hundreds of millions of iterations to occur, but such conditions are possible in equipment which must run continuously. In contrast, a stable system will bound the maximum excess MSE to some finite limit. In practice, forcing the LMS algorithm to be stable is simple, but the modifications incur some extra computational cost and increase the excess MSE over the infinite-precision case [27].

The interrelationship among the reference noise, the step-size parameter, and the quantization error can be quickly illustrated with an example from [11]. First, it is noted that the tap-weight update equation for the LMS algorithm can be written as:

$$\mathbf{W}_{j+1} = \mathbf{W}_j + \mathbf{G}_j \tag{4.29}$$

where

$$\mathbf{G}_j = 2\mu\varepsilon_j\mathbf{X}_j . \tag{4.30}$$

In hardware, the evaluation of $\mathbf{G}_j$ will start with the scalar multiplication of $2\mu\varepsilon_j$, and then proceed with the multiplication of the vector $\mathbf{X}_j$ (note that in an efficient implementation, 2μ is

combined as a single variable rather than requiring an extra multiplication). Explicitly identifying the quantization errors, the limited-precision gradient estimate is:

$$\hat{\mathbf{G}}_j = \left(2\mu\left(\varepsilon_j + \Delta\varepsilon_j\right) + \Delta_{1j}\right)\cdot\left(\mathbf{X}_j + \Delta\mathbf{X}_j\right) + \Delta_{2j} \tag{4.31}$$

where $\Delta\varepsilon_j$ is related to the number of storage bits allotted to ε_j, $\Delta\mathbf{X}_j$ relates to the A/D quantization error of the reference noise, and Δ_{1j} and Δ_{2j} are the errors due to truncation or rounding during the multiplication operations. These cumulative effects can be isolated into one compact term:

$$\mathbf{b}_j = \hat{\mathbf{G}}_j - \mathbf{G}_j \tag{4.32}$$

and (4.29) can be modified to include the effects of the errors:

$$\mathbf{W}_{j+1} = \mathbf{W}_j + \hat{\mathbf{G}}_j = \mathbf{W}_j + \mathbf{G}_j + \mathbf{b}_j \tag{4.33}$$

Assume that the system inputs are stationary, and that $\mathbf{X}_j$ is independent of $\mathbf{W}_j$ (this is not always true, but is frequently done to make the analysis more tractable [27]). Substituting (4.2) and (4.30) into (4.33) and taking the mean results in:

$$E\{\mathbf{W}_{j+1}\} = [\mathbf{I} - 2\mu\mathbf{R}_X]E\{\mathbf{W}_j\} + 2\mu\mathbf{P} + E\{\mathbf{b}_j\} \tag{4.34}$$

If the vector mean of the finite-precision errors $E\{\mathbf{b}_j\}$ is not zero, it is clear from (4.34) that the system will not be stationary, and the expected value of the filter tap weights will grow steadily with time until overflow. However, if $E\{\mathbf{b}_j\}$ is assumed to be zero or very close to zero for all time intervals j, then the algorithm will operate in a stable manner at steady state. Solving (4.34) with the condition $E\{\mathbf{W}_{j+1}\} = E\{\mathbf{W}_j\}$ gives:

$$E\{\mathbf{W}_j\} = \mathbf{R}_X^{-1}\mathbf{P} + \frac{1}{2\mu}\mathbf{R}_X^{-1}E\{\mathbf{b}_j\} \tag{4.35}$$

This is a combination of the Wiener solution and an additional corruption term [Ci3]. Unfortunately, this example is not flawless. First, $E\{\mathbf{W}_j\}$ will obviously not tend to infinity when μ is set to zero, as suggested by (4.35); instead, the adaptive filter will simply freeze its coefficients and update them no further. Second, the equation is only strictly valid when $E\{\mathbf{b}_j\}=0$. These two flaws will be temporarily overlooked to provide some insight into the rate at which finite-precision errors accumulate.

Before analyzing equation (4.35), it is first necessary to note that the eigenvalue spread of the correlation matrix $\mathbf{R}_X$ is bounded by the ratio of the maximum and minimum values of the power spectral density of $\mathbf{X}_j$ [27]:

$$\text{eigenvalue spread} = \frac{\lambda_{\max}}{\lambda_{\min}} \leq \frac{S_{X,\max}(f)}{S_{X,\min}(f)} \tag{4.36}$$

A high eigenvalue spread implies that $\mathbf{R}_X$ is ill-conditioned (nearly singular). In this case, if $\mathbf{R}_X$ is normalized such that its largest elements (on the main diagonal) are equal to 1, $\mathbf{R}_X^{-1}$ will contain some much larger elements and have a maximum eigenvalue $1/\lambda_{\min}$. Any errors which occur in the direction of the eigenvector corresponding to this maximum eigenvalue will be greatly magnified when multiplied by $\mathbf{R}_X^{-1}$.

Knowing this, it is clear that the second term on the right-hand side of (4.35) will be very sensitive to the statistical properties of $\mathbf{X}_j$. If the common-mode input is highly correlated, $\mathbf{R}_X$ will be very ill-conditioned. This is the case for audio-controlled AM-modulated signals, whose envelopes vary extremely slowly in comparison with the xDSL system's sampling frequency. This intuitive understanding can be supported by the bound of (4.36): a narrowband interferer's PSD consists of one very high spike surrounded by very low power density levels, so a large eigenvalue spread is likely even if the bound is loose. Since the LMS algorithm implicitly inverts $\mathbf{R}_X$, $\mathbf{R}_X^{-1}$ will likely contain large values, and any deviation of $E\{\mathbf{b}_j\}$ from zero could quickly lead to numerical instability.

Both [11] and [27] emphasize that the factor $1/2\mu$ in (4.35) is an additional destabilizing influence. However, in the simulations conducted for Chapter 6, this factor was found to be of

little consequence. A quick inspection of equations (4.30) to (4.32) confirms that $\mathbf{b}_j$ will be heavily dependent on 2μ, greatly reducing the importance of the division by 2μ. Furthermore, increasing the value of μ to reduce the magnification of finite-precision effects only serves to increase the misadjustment (4.26). Since the misadjustment affects the performance of the algorithm even when there has been no accumulation of quantization error, it is usually viewed as more significant than concerns over finite-precision effects [11].

4.1.4 The Leaky LMS Algorithm

The leaky LMS algorithm can be used to prevent overflow by changing the cost function such that it minimizes a combination of the MSE and the energy of the filter's impulse response. The new cost function can be written as:

$$J_j = \varepsilon_j^2 + \alpha \left\| \mathbf{W}_j \right\|^2 \tag{4.37}$$

where α is a control parameter. Effectively, this function augments the autocorrelation matrix to $\mathbf{R}_X + \alpha \mathbf{I}$, ensuring that the lowest eigenvalue is greater than α. Minimization of J_j with respect to $\mathbf{W}_j$ leads to the leaky LMS tap-weight update formula:

$$\mathbf{W}_{j+1} = (1 - 2\mu\alpha)\mathbf{W}_j + 2\mu\varepsilon_j \mathbf{X}_j \tag{4.38}$$

where, for convergence:

$$0 \le \alpha < \frac{1}{2\mu} \tag{4.39}$$

The extra multiplication operations add complexity to the LMS algorithm, although it is possible to economize by leaking only one tap at each iteration [11]. The excess MSE will also rise slightly; see [11] for further details. Conceptually, the leakage factor $-2\mu\alpha\mathbf{W}_j$ in (4.38) represents a zero-mean white noise process with variance α added to the reference noise input [27]. In other words, it is equivalent to the use of a dither signal, a familiar technique in many other applications.

4.1.5 Tracking Performance in Non-Stationary Environments

Analysis of the LMS filter's ability to track the time-varying Wiener solution for non-stationary inputs is complex [71, 27], and will not be reproduced here. Generally speaking, the time-dependent environment causes a new term to be added to the misadjustment Π. In (4.26), Π was related only to the gradient noise which caused the algorithm to waver about the optimum solution. Assuming that this gradient noise is Gaussian with zero mean and variance σ_g^2, that the autocorrelation matrix of the tap weights is denoted by $\mathbf{R}_W$, and that a large number of iterations have occurred, the new misadjustment is [27]:

$$\Pi = \mu \cdot \mathrm{tr}[\mathbf{R}_X] + \frac{1}{4\mu\sigma_g^2} \cdot \mathrm{tr}[\mathbf{R}_W] \tag{4.40}$$

The first term on the right-hand side is simply equation (4.26), the misadjustment for the stationary case. The second term relates to the time that the system needs to "catch up" with the changing input statistics [71]. It is proportional to the number of filter coefficients, but inversely proportional to the rate of adaptation. More significant, however, is the fact that the first term varies with μ while the second varies inversely with μ. [71] shows that Π is minimized when the first and second terms are forced equal to each other. This gives an optimum μ of:

$$\mu_{opt} = \frac{1}{2\sigma_g} \sqrt{\frac{\mathrm{tr}[\mathbf{R}_W]}{\mathrm{tr}[\mathbf{R}_X]}} \tag{4.41}$$

Since this expression would be difficult to evaluate analytically in a modem receiver, another gradient descent technique, the LMS algorithm with adaptive gain, has been developed.

4.1.6 The LMS Algorithm with Adaptive Gain

The idea for the LMS algorithm with adaptive gain started as an exercise in [1] and has since been treated in greater depth, particularly in [41]. The cost function to be minimized is:

$$J_j = \frac{1}{2} E\left\{ \left| \varepsilon_j \right|^2 \right\} \tag{4.42}$$

which means that the equations for the error calculations and the tap weight updates remain the same as the original LMS algorithm (4.2 and 4.11). To update μ, an iterative calculation of the form:

$$\mu_{j+1} = \mu_j - \beta \hat{\nabla}_{\mu,j} \tag{4.43}$$

is desired, where β is another step-size parameter. Assuming that the system inputs and filter coefficients are all real numbers, the true gradient of the cost function with respect to μ can be expressed as:

$$\nabla_{\mu,j} = \frac{\partial J_j}{\partial \mu} = \frac{1}{2} E\left\{ 2\varepsilon_j \frac{\partial \varepsilon_j}{\partial \mu} \right\} = E\left\{ \varepsilon_j \cdot \frac{\partial \left(d_j - \mathbf{X}_j^T \mathbf{W}_j \right)}{\partial \mathbf{W}_j} \cdot \frac{\partial \mathbf{W}_j}{\partial \mu} \right\} \tag{4.44}$$

Defining

$$\mathbf{T}_j = \frac{\partial \mathbf{W}_j}{\partial \mu} \tag{4.45}$$

gives:

$$\nabla_{\mu,j} = E\left\{ -\varepsilon_j \mathbf{X}_j^T \mathbf{T}_j \right\} \tag{4.46}$$

Assuming that taking instantaneous values of $\nabla_{\mu,j}$ and averaging them over time will give a sufficient estimate of the true gradient:

$$\hat{\nabla}_{\mu,j} = -\varepsilon_j \mathbf{X}_j^T \mathbf{T}_j \tag{4.47}$$

(4.43) can be re-written as:

$$\mu_{j+1} = \mu_j + \beta \varepsilon_j \mathbf{X}_j^T \mathbf{T}_j . \tag{4.48}$$

The updated value for μ_j must be treated carefully. Specifically, the optimum value of μ_j is often near the point at which the algorithm cannot converge [41]. It is therefore important to ensure that the system is implemented such that μ_j has a maximum permissible value, μ_+.

Similarly, it is prudent to select a minimum permissible value, μ_-. In practice, μ_+ is more difficult to set, but a suitable value can usually be found by running simulations during the design process [41].

Updates of $\mathbf{T}_j$ can be found by differentiating the tap-weight update equation:

$$\begin{aligned}\frac{\partial \mathbf{W}_{j+1}}{\partial \mu} &= \frac{\partial}{\partial \mu}\left[\mathbf{W}_j + 2\mu_j \varepsilon_j \mathbf{X}_j\right] \\ \mathbf{T}_{j+1} &= \mathbf{T}_j + 2\varepsilon_j \mathbf{X}_j + 2\mu_j \mathbf{X}_j \frac{\partial \varepsilon_j}{\partial \mu} \\ &= \mathbf{T}_j + 2\varepsilon_j \mathbf{X}_j - 2\mu_j \mathbf{X}_j \mathbf{X}_j^T \mathbf{T}_j\end{aligned} \tag{4.49}$$

Packaging this equation with (4.2), (4.11) and (4.48), the algorithm can be summarized as:

$$\begin{aligned}\varepsilon_j &= d_j - \mathbf{W}_j^T \mathbf{X}_j \\ \mathbf{W}_{j+1} &= \mathbf{W}_j + 2\mu_j \varepsilon_j \mathbf{X}_j \\ \mu_{j+1} &= \mu_j + \beta \varepsilon_j \mathbf{X}_j^T \mathbf{T}_j \\ \mathbf{T}_{j+1} &= \left[\mathbf{I} - 2\mu_j \mathbf{X}_j \mathbf{X}_j^T\right]\mathbf{T}_j + 2\varepsilon_j \mathbf{X}_j\end{aligned} \tag{4.50}$$

4.1.7 Adaptive Noise Cancellation Performance Analysis

This section will present an analysis of the adaptive noise canceller's performance under two conditions: with a single interferer, and with a large number of similar interferers. These two case studies are shortened and slightly modified versions of those given in [70] and [54]. For convenience, all signals are assumed to be wide-sense stationary, and the LMS algorithm is assumed to be capable of achieving performance arbitrarily close to that of the method of steepest descent.

Consider Figure 4-2. A message signal $m[n]$ and N interfering signals $i_1[n], \ldots, i_N[n]$, all pass through a twisted-pair channel. $g_m[n]$ and $h_m[n]$ denote the differential and common-mode channel impulse responses respectively for the message signal, while $g_k[n]$ and $h_k[n]$ correspond

to interferer k. Separate white Gaussian noise signals $\eta_{diff}[n]$ and $\eta_{cm}[n]$ are added to the two inputs, and the remainder of the system consists of the same adaptive filter as was presented in Figure 4-1. Note that the channel impulse responses may be non-linear because of the physical properties of the hybrid transformer. In this case, perfect cancellation of the interfering signals will not be achieved using a linear adaptive filter. However, in order to allow mathematical analysis, the entire system model is assumed to be linear in the following pages.

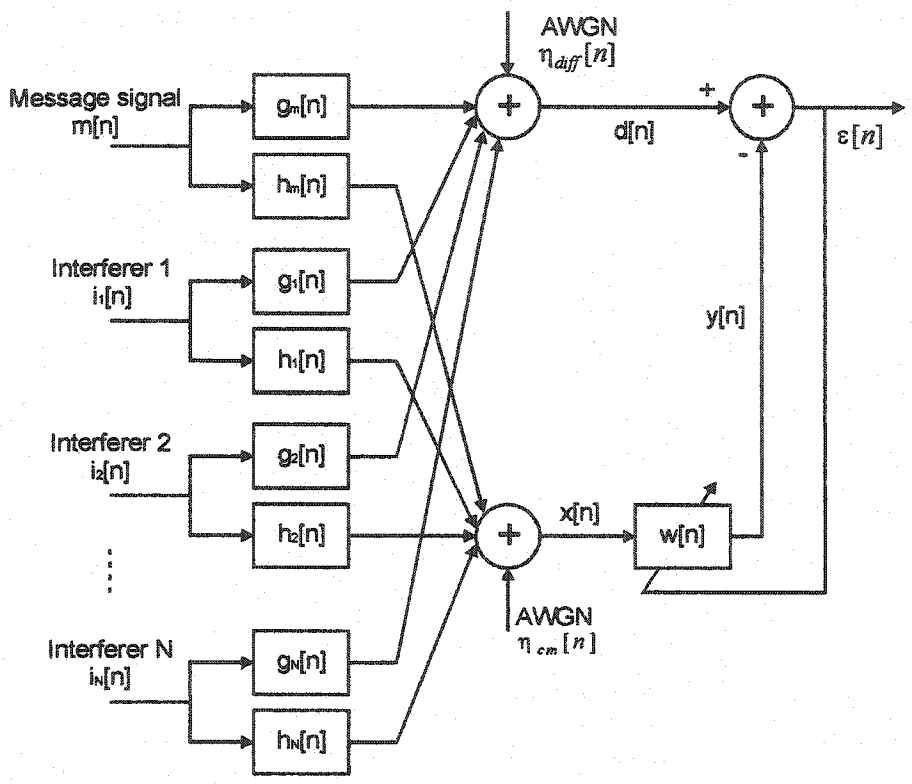

Figure 4-2: Mathematical Model of Noise Cancellation System

Looking at the model, it is clear that the desired signal and reference noise can be expressed as:

$$d[n] = m[n] * g_m[n] + \left(\sum_{k=1}^{N} i_k[n] * g_k[n] \right) + \eta_{diff}[n] \tag{4.51}$$

and

$$x[n] = m[n] * h_m[n] + \left(\sum_{k=1}^{N} i_k[n] * h_k[n] \right) + \eta_{cm}[n] \tag{4.52}$$

where the * operator represents a convolution. The optimum filter coefficients $w[n]$ can be found using the Wiener-Hopf equations:

$$R_{dx}[p] = \sum_{k=-\infty}^{\infty} w[k] \cdot R_x[p-k] \qquad \forall p \in (-\infty, \infty) \tag{4.53}$$

where $R_{dx}[p]$ is the cross-correlation between $d[n]$ and $x[n]$, and $R_x[p]$ is the autocorrelation function of $x[n]$. In this form, the equations are not constrained to be causal, but they are more mathematically tractable because they can be z-transformed, giving:

$$S_{dx}[z] = W[z] \cdot S_x[z] \tag{4.54}$$

where $S_{dx}[z]$ is the cross-power spectral density between $d[n]$ and $x[n]$, and $S_x[z]$ is the power spectral density of $x[n]$. This gives the optimum filter coefficients as:

$$W_{Wiener}[z] = \frac{S_{dx}[z]}{S_x[z]} \tag{4.55}$$

To evaluate the two PSDs, assume that the interference signals are uncorrelated with the message signal. The autocorrelation function of $x[n]$ is:

$$\begin{aligned} R_x[p] &= E\{x[n] \cdot x[n+p]\} \\ &= R_m[p] * h_m[p] * h_m[-p] + \left(\sum_{k=1}^{N} R_{i_k}[p] * h_k[p] * h_k[-p] \right) + R_{\eta_{cm}}[p] \end{aligned} \tag{4.56}$$

where the three R terms in the second line are the autocorrelation functions of the message signal, the interference signal, and the common-mode AWGN respectively. Taking the z-transform yields:

$$S_x[z] = S_m[z] \cdot |H_m[z]|^2 + \left(\sum_{k=1}^{N} S_{i_k}[z] \cdot |H_k[z]|^2 \right) + S_{\eta_{cm}}[z] \tag{4.57}$$

Using similar notation, the cross-correlation of $d[n]$ and $x[n]$ can be expressed as:

$$R_{dx}[p] = E\{d[n] \cdot x[n+p]\}$$
$$= R_m[p] * g_m[p] * h_m[-p] + \left(\sum_{k=1}^{N} R_{i_k}[p] * g_k[p] * h_k[-p] \right) \tag{4.58}$$

giving:

$$S_{dx}[z] = S_m[z] \cdot G_m[z] \cdot H_m[z^{-1}] + \left(\sum_{k=1}^{N} S_{i_k}[z] \cdot G_k[z] \cdot H_k[z^{-1}] \right) \tag{4.59}$$

Substitution of (4.57) and (4.59) into (4.55) gives an optimum solution:

$$W_{Wiener}[z] = \frac{S_m[z] \cdot G_m[z] \cdot H_m[z^{-1}] + \left(\sum_{k=1}^{N} S_{i_k}[z] \cdot G_k[z] \cdot H_k[z^{-1}] \right)}{S_m[z] \cdot |H_m[z]|^2 + \left(\sum_{k=1}^{N} S_{i_k}[z] \cdot |H_k[z]|^2 \right) + S_{\eta_{cm}}[z]} \tag{4.60}$$

Case study 1: a single interferer

This example, taken from [70], illustrates the best improvement in signal-to-noise ratio that can be achieved using an ideal adaptive noise canceller. Consider the case where there is only a single interferer, $i_1[n]$. For convenience of notation, both of the differential transfer functions ($g_m[n]$ and $g_1[n]$) are set equal to 1, while the common-mode transfer functions $h_m[n]$ and $h_1[n]$ are arbitrary. Both AWGN signals are assumed to be negligible.

Referring to (4.60), it can be seen that the optimum filter tap weights for this case are given by:

$$W_{Wiener}[z] = \frac{S_m[z] \cdot H_m[z^{-1}] + S_{i_1}[z] \cdot H_1[z^{-1}]}{S_m[z] \cdot |H_m[z]|^2 + S_{i_1}[z] \cdot |H_1[z]|^2} \tag{4.61}$$

Note that when $H_m[z] \to 0$, this reduces to the intuitive result:

$$W_{Wiener}[z] = \frac{1}{H_1[z]} \tag{4.62}$$

From Figure 4-2, it can be seen that the PSD of the signal component at the system output is given by:

$$S_{m_{out}}[z] = S_m[z] \cdot \left| 1 - H_m[z] \cdot W_{Wiener}[z] \right|^2 \tag{4.63}$$

Similarly, the spectrum of the interference at the system output is:

$$S_{i1_{out}}[z] = S_{i_1}[z] \cdot \left| 1 - H_1[z] \cdot W_{Wiener}[z] \right|^2 \tag{4.64}$$

Combining (4.61), (4.63) and (4.64), the signal-to-noise density ratio at the output can, after some simple algebra, be expressed as:

$$\rho_\varepsilon[z] = \frac{S_{m_{out}}[z]}{S_{i1_{out}}[z]} = \frac{S_{i_1}[z] \cdot |H_1[z]|^2}{S_m[z] \cdot |H_m[z]|^2} \tag{4.65}$$

Again referring to Figure 4-2, it is easily verified that the signal-to-noise density ratio at the reference noise input to the adaptive filter is:

$$\rho_x[z] = \frac{S_m[z] \cdot |H_m[z]|^2}{S_{i_1}[z] \cdot |H_1[z]|^2} \tag{4.66}$$

Comparing (4.65) and (4.66), a significant result is obtained:

$$\rho_\varepsilon[z] = \frac{1}{\rho_x[z]} \tag{4.67}$$

The signal-to-noise density ratios at the canceller's output is the inverse of the ratio at the common-mode input. In other words, if the signal-to-noise density ratio at the reference input is -20 dB for some frequency *f*, the ratio at the output is +20 dB at that frequency. In actual operation, performance will not reach this ideal limit because of stochastic gradient noise, finite precision effects and receiver thermal noise. However, on a cable with a 40 dB balance at some AM carrier frequency, ρ_x will typically be lower than –40 dB at that frequency, especially if the RFI is strong or if the VDSL signal is badly attenuated before transferring between modes. While a ρ_ε of 40

dB would not be possible in this case, it would not be unrealistic to hope for values in the vicinity of 25 to 30 dB. For more than one interferer, the ρ_ε value will drop further [7]; this effect will be discussed briefly in the next example.

Case study 2: multiple similar interferers

Now consider the case where there is a large number N of disturbers, and for simplicity the message signal is set to zero. Assuming that the AWGN is negligible, (4.60) reduces to:

$$W_{Wiener}[z] = \frac{\sum_{k=1}^{N} S_{i_k}[z] \cdot G_k[z] \cdot H_k[z^{-1}]}{\sum_{k=1}^{N} S_{i_k}[z] \cdot |H_k[z]|^2} \tag{4.68}$$

If the spectra of all the interfering signals are the same, as is possible in the case of VDSL crosstalk, (4.68) becomes:

$$W_{Wiener}[z] = \frac{\sum_{k=1}^{N} G_k[z] \cdot H_k[z^{-1}]}{\sum_{k=1}^{N} |H_k[z]|^2} \tag{4.69}$$

To gauge the performance of the adaptive noise canceller, [54] proposes a processing gain, PG:

$$\begin{aligned} PG &= \frac{\text{Expected noise power with no noise cancellation}}{\text{Expected noise power with noise cancellation}} \\ &= \frac{\sum_{k=1}^{N} S_{i_k}[z] \cdot |G_k[z]|^2}{\sum_{k=1}^{N} S_{i_k}[z] \cdot \left| G_k[z] - W_{Wiener}[z] H_k[z^{-1}] \right|^2} \end{aligned} \tag{4.70}$$

Once again assuming that the spectra of all interferers are the same, and using (4.69), this may be simplified to:

$$PG = \frac{1}{1 - \frac{\left(\sum_{k=1}^{N} G_k[z] \cdot H_k[z^{-1}]\right)^2}{\left(\sum_{k=1}^{N} |G_k[z]|^2\right) \cdot \left(\sum_{k=1}^{N} |H_k[z]|^2\right)}} \tag{4.71}$$

The second term in the denominator of (4.71) is the squared sample correlation coefficient between the mode-to-mode transfer functions [54]. If there is a high degree of correlation, this term will tend to 1 and the processing gain will be large. However, in a large cable bundle containing tens or hundreds of wires with different twists, it is not reasonable to expect a high correlation coefficient. In this case, the second term of (4.71) will tend towards 0, and the processing gain will be negligible. Note that the processing gain cannot fall under 1; in other words, the filter can do no worse than no cancellation using infinite precision arithmetic [12]. [54] does point out that the processing gain will climb if a few crosstalkers are much more powerful than the others. Unfortunately, this scenario does not seem likely for VDSL, whose transmitters will all transmit at the same power level. Furthermore, the received crosstalk levels will be far lower than the received AM levels, making it a low priority for cancellation. For all these reasons, mitigation of crosstalk will not be further considered in this document.

To a certain extent, the results discussed here also apply when the interfering signals have equal power, but do not have the same spectrum. In this case, the adaptive filter will be able to provide near-perfect cancellation of the interferer with the highest power [11]. By setting its filter coefficients to perform this task, however, it biases itself against full cancellation of other interferers. Accordingly, the processing gain as defined in (4.70) drops in value as the number of interferers increases.

4.2 *Wideband RFI Noise Cancellation for SDMT*

Synchronized DMT (SDMT) is a time-division duplexed version of DMT, and will not become the VDSL line code because the T1E1.4 subcommittee has already selected frequency-division duplexing [4, 12]. However, Amati Communications Corporation conducted some basic research into front-end noise cancellation for SDMT before this decision was made. The canceller was envisioned as an analog adaptive filter, although a digital implementation was not ruled out.

The analog implementation was preferred in order to prevent A/D converter saturation in the event that a strong amateur radio interferer was present. It is interesting to note, however, that one of the designers has since recommended the use of digital FDD for VDSL [10]. Since the trans-hybrid loss is usually assumed to be 12 dB, the modem's own transmit signal would arrive at the receiver with an average power of about 5 dBm, 15 dB greater than the worst-case amateur AM interferer in VDSL tests [9]. The digital FDD recommendation would seem to remove the most advantageous trait of an analog canceller.

Since the reference common-mode noise extracted from the loop will normally contain trace levels of the message signal, the SDMT canceller carries out filter adaptation only during the time-domain duplexing guard intervals, when it is guaranteed that there will be no information signal present. Unfortunately, this technique has been overtaken by the selection of FDD for VDSL. However, it also overstates the importance of the tiny quantities of message signal that transfer to common mode on most loops [58]; as some of the simulations in Chapter 6 will show, the corruption of the reference signal is negligible when the loop has at least 30 dB balance at low frequencies.

4.3 *Interference Rejection Using Programmable Notch Filters*

In Chapter 3, it was found that analog cancellation of very strong narrowband interferers would be required. One approach to this task is to use the analog adaptive filter described in Section 4.2. However, satisfactory cancellation may also be achieved using simpler devices. This section describes an analog blanker which can suppress multiple narrowband interferers, and the next section outlines the operation of a narrowband analog canceller.

The design given in reference [48] is intended for interference rejection in systems employing high-speed HF radio links to send serial digital data. HF radio is usually used for long-distance mobile communications, and so interference will often come from unknown remote sources. Furthermore, HF antennas are almost always omnidirectional, so there will be no "clean" noise reference to use for cancellation. As a consequence, strong interferers can be removed only by nulling the received signal at the interferer frequencies.

The design works as follows. The bandwidth of the information signal is divided into a number of bins, and a spectrum analyzer is used to determine the received signal's power level in each bin. Since the spectral shape of the information signal will be known, the received signal can be compared to an appropriately scaled version of this spectral template. If the received signal exceeds the template in a particular bin, then that frequency slot is identified as containing excessive interference. A programmable notch filter is then set to nullify the received signal in the bin.

This technique is conceptually simple, but uses complex hardware to carry out spectral analysis. The fact that it cancels out all of the received signal's energy in the interference bins is also undesirable, as the information signal loses energy along with the noise. For the twisted-pair channel, which supplies a relatively clean noise reference in common mode, a better alternative will now be considered.

4.4 Narrowband Analog Suppression of RFI

A narrowband canceller to remove individual AM disturbers from a DSL signal was proposed by Lefebvre and Yeap in [45]. The fundamental ideas behind the device can be explained with the aid of Figure 4-3. The differential and common-mode signals are extracted from the loop using a hybrid device and a summer respectively. The common-mode signal is passed to a noise detection unit which uses a sliding bandpass window and some logic circuitry to locate the most powerful source of AM interference. Once the centre frequency of this noise source is found, the noise control circuitry uses both the received differential and common-mode signals to determine the attenuation and phase shift that must be applied to the noise to cancel its effect on the differential signal. With these values calculated, an inverted noise estimate is generated using a tunable bandpass filter, an adjustable phase inverter, and an adjustable gain block. Actual cancellation is carried out by a summer, which adds the differential signal with the inverted noise estimate. The output differential signal is fed back to the noise control circuitry to allow the controller to adjust the phase and gain of the noise estimate as appropriate. Obviously, more than one narrowband signal can be cancelled if multiple noise estimator circuits are used and the noise detection and control circuitry is expanded.

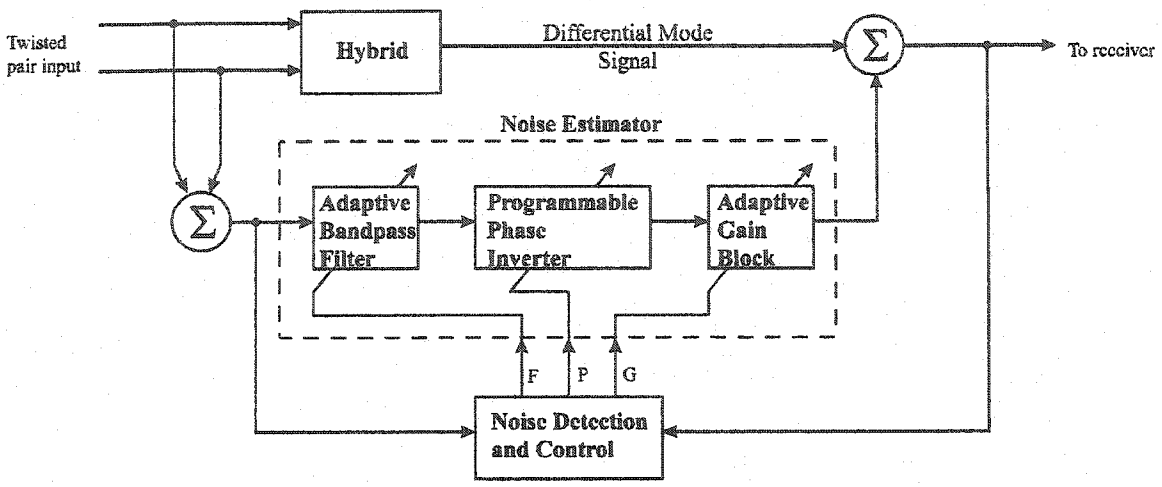

Figure 4-3: Narrowband analog canceller

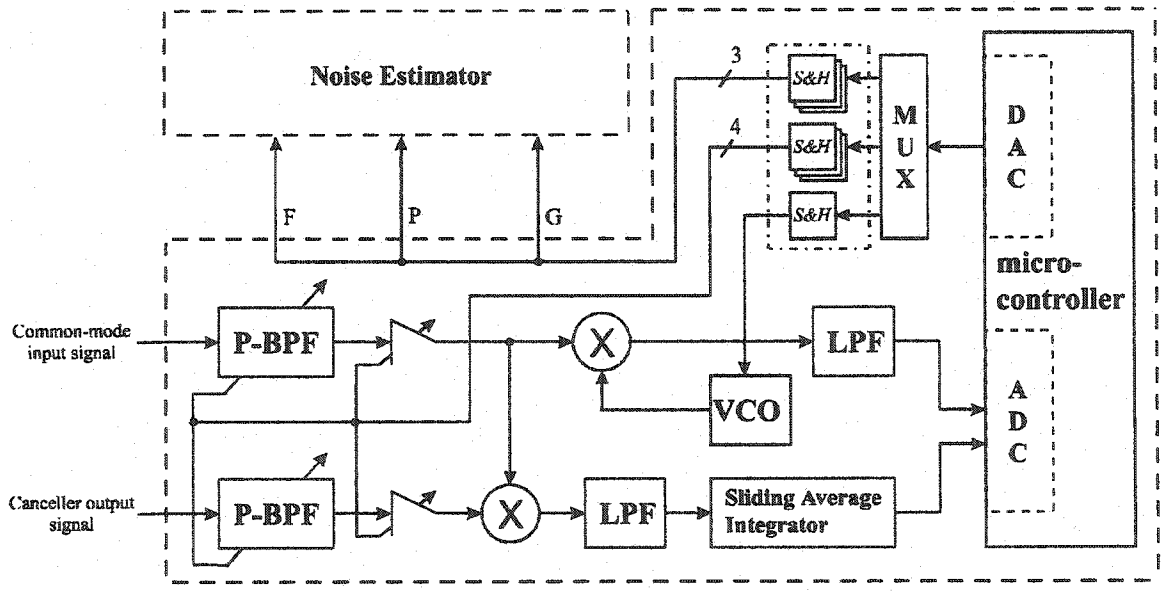

Figure 4-4: Noise Detection and Control Circuitry

The controller circuit used to make feedback adjustments to the phase and gain is shown in Figure 4-4. The common-mode input signal and canceller output signal are both filtered by bandpass filters set to the same centre frequency as the tunable bandpass filter in Figure 4-3. Both filter outputs are subsequently amplified by a common gain. The common-mode signal is then down-converted and lowpass filtered to obtain the envelope of the AM interferer. It is also multiplied with the output signal to obtain a continuous-time cross-correlation. The residual signal at the output of the multiplier is then lowpass filtered to remove spurious noise and applied to a sliding average integrator, which extracts the DC component from it. The envelope of the

common-mode signal and the residue signal are both A/D converted, and a microcontroller determines whether or not the residue tracks the noise, and whether its amplitude is positive or negative. Based on these measurements, the logic circuitry can make changes to the noise estimate's phase and gain, and then communicate these changes back to the noise estimator through a DAC and a serial-to-parallel converter.

The main advantages of this analog canceller are:

a. it conserves ADC resolution for the information signal;
b. if several narrowband cancellers are used, it is possible to obtain very precise phase alignment and gain control for several interferers, allowing for very good cancellation. By contrast, a wideband canceller can match the performance of a narrowband one when there is only a single interferer, but will fare more poorly if additional noise sources and coupling paths are present; and
c. control functions are relatively simple and can be carried out by a low-speed microcontroller.

The main disadvantages of this technique are:

a. in the event that two AM interferers arrive at the receiver with approximately equal power, it will be difficult to keep the canceller from switching back and forth between them;
b. the analog circuits used must run with relatively high currents to keep additive noise from becoming a significant factor; and
c. the analog canceller will not be as easily miniaturized as an equivalent digital canceller.

This device has been developed by Nortel Networks and is known to work satisfactorily [72]; it is capable of suppressing a narrowband interferer by as much as 40 dB.

4.5 Final Notes

A number of other noise cancellation devices have been invented to suppress AM RFI or impulse noise. A system described in [42] cancels impulse noise in LAN modems using

differential analog circuitry. Obviously, very good balance can be achieved on a small ASIC, and impulses caused by clock or switch noise will be greatly reduced. This canceller is merely taking advantage of the use of differential circuitry in the same way that emitter-coupled logic or twisted-pair cable does, and does not solve the problem of declining balance at higher frequencies. [7] uses adaptive filtering to carry out RFI cancellation for coaxial cables. This is very similar to the twisted-pair techniques already discussed in this chapter, except the reference noise is generated by setting up an antenna somewhere near the receiver site. Last, [43] proposes an impulse noise canceller with a unique implementation. Its basic premise is that the impulsive bursts on most twisted-pair cables take on only a few different shapes. The author proposed that adaptive pattern recognition be used to identify incoming impulses and cancel them out. This method requires many complex circuits to train the impulse detector and determine the shape, magnitude and arrival time of received noise bursts, and will not be considered further.

This chapter has addressed some of the basic theories of adaptive noise cancellation, concentrating on variants of the least-mean-square algorithm. It has also briefly outlined some canceller designs that have been proposed for a variety of applications. The design trade-offs introduced in this chapter will now be discussed in the development of this paper's proposed solution.

5. Proposed Digital Noise Canceller

5.1 Analysis of Limitations and Trade-offs

From the previous chapters, it is clear that the following key limitations will be imposed on any digital noise canceller design:

a. the ADC speed and resolution;
b. the permissible power consumption and board space;
c. the theoretical limitations on adaptive filter performance;
d. the number of bits of precision used in fixed-point arithmetic; and
e. the permissible complexity of the adaptive filter's control algorithm.

The ADC capabilities are obviously crucial to the successful performance of the VDSL system as a whole. If large signals, such as the portion of the transmit signal that leaks through the hybrid or high-power amateur AM interferers, are to be treated digitally, then the ADC's quantization noise will limit the bandwidth over which the system can achieve positive noise margin. For this reason, it is desirable to do FDD filtering and maybe some noise cancellation in the analog domain.

Unfortunately, the design and construction of a wideband adaptive analog canceller are not trivial. Its signals remain susceptible to additive noise unless greater currents are allowed to flow, or very well balanced differential circuitry is used. Also, its tap weights must be updated while the reference noise is being processed through the filter, a procedure likely to introduce additional transient noise.

It cannot be stated definitively that a digital adaptive filter will have lower power consumption than its analog counterpart; similar to the high-power equalizer of a single-carrier system, it must operate at extremely high speeds. However, it is easily updated every clock cycle without disrupting its input, is more easily miniaturized, and has no susceptibility to external noise. On the other hand, it does generate its own internal noise through accumulated numerical inaccuracy, and may require a very large number of bits to adequately represent its inputs and outputs. The control algorithm used to update the filter coefficients may also be prohibitively

complex. For example, it is clear that the LMS algorithm with adaptive gain will not be useful for high-speed processing if there are a large number of filter coefficients; the update equation for $\mathbf{T}_{j+1}$ given by (4.49) will quickly mushroom to an unsustainable number of computations unless a very fast processor is used.

The most appealing characteristic of an adaptive filter is that it essentially designs itself, and no explicit measurements on the incoming reference noise or desired signal are required (except to set the step size μ). Assuming numerical stability, it can also do no worse than no cancellation; ie. aside from finite precision effects, it will not add noise to its output. However, as outlined in Chapter 4, it has some serious performance limitations. In particular, it can provide near-perfect cancellation of a single narrowband interferer, but will cancel other interferers to a lesser degree. In contrast, a parallel configuration of narrowband cancellers can provide excellent performance against all interferers; its major drawback is that it requires "intelligent" control circuitry to align the noise reference in both phase and amplitude with the desired signal. Phase mismatches of only a few degrees mean that little cancellation will be achieved, and large mismatches will actually increase the noise component of the differential signal.

Altogether, it is impossible to foresee any noise canceller design which is clearly superior to all others. The advantages and disadvantages of analog versus digital, and narrowband versus wideband, are largely complementary. This project concentrates on a digital structure mostly because it is easiest to design and simulate, and on a wideband implementation because it is anticipated that multiple narrowband cancellers would incur an excessive hardware cost.

5.2 Proposed Noise Canceller Design

In the course of this project, a large number of different canceller designs were quickly tested in the lab. In general, it was found that very simple designs work extremely well when the differential and common-mode signals arrive at the hybrid synchronized in time. When this synchrony does not exist but the interference is a single narrowband signal, it is also relatively easy to design an analog circuit which aligns the reference noise with the differential noise in amplitude and phase (similar to [45]). However, when many interferers are present, the only methods which

show promise are a parallel configuration of multiple narrowband cancellers, or the use of a wideband adaptive filter.

The following pages outline two designs which conduct wideband noise cancellation. The first is a basic digital canceller which can handle moderate RFI levels. The second incorporates front-end narrowband analog cancellation to reduce very strong RFI and conserve ADC resolution. Based on these designs, the main objectives of this project's simulation testing are then described.

5.2.1 Basic digital cancellation

Figure 5-1 shows a generic digital noise canceller, capable of handling noise which converts to differential mode as a result of both loop and hybrid device imbalances. The differential signal is extracted from the loop using a hybrid device (pictured here as a transformer), amplified by gain G1, and digitized using an ADC. The A/D block is assumed to include a lowpass filter to prevent aliasing. Similarly, the common-mode signal is extracted from the loop by a summer, amplified by gain G2 and digitized. Gains G1 and G2 are set by automatic gain control circuitry in the system's control unit, and should be maintained at a fixed ratio to prevent adverse effects on the adaptive filter.

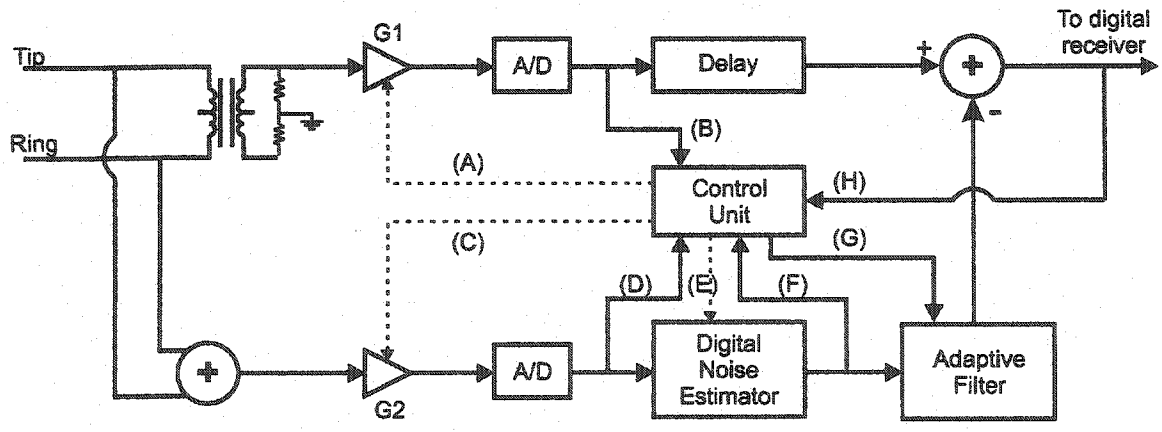

Figure 5-1: Basic digital noise canceller

The common-mode signal passes into a digital noise estimator (DNE), whose function is to pass the narrowband interference signals and block as much of the leaked QAM information signal as possible; in theory, this should make it less likely that the adaptive filter will remove information

signal energy from the differential signal. The DNE consists of a linear-phase filter bank which splits the common-mode signal into a number of frequency bands. The system control unit determines which bands are to be summed and passed to the adaptive filter, and which bands are to be zeroed. In practice, this is accomplished by estimating the power spectral density of the received common-mode signal and comparing it with a spectral mask giving the expected shape of the received QAM signal to find undesired narrowband peaks; see Figure 5-2 for a functional flow diagram of the system control unit. If the PSD exceeds the mask for a given percentage of measurements in a given frequency band, then that band is output from the DNE; bands which consistently fall under the mask are zeroed.

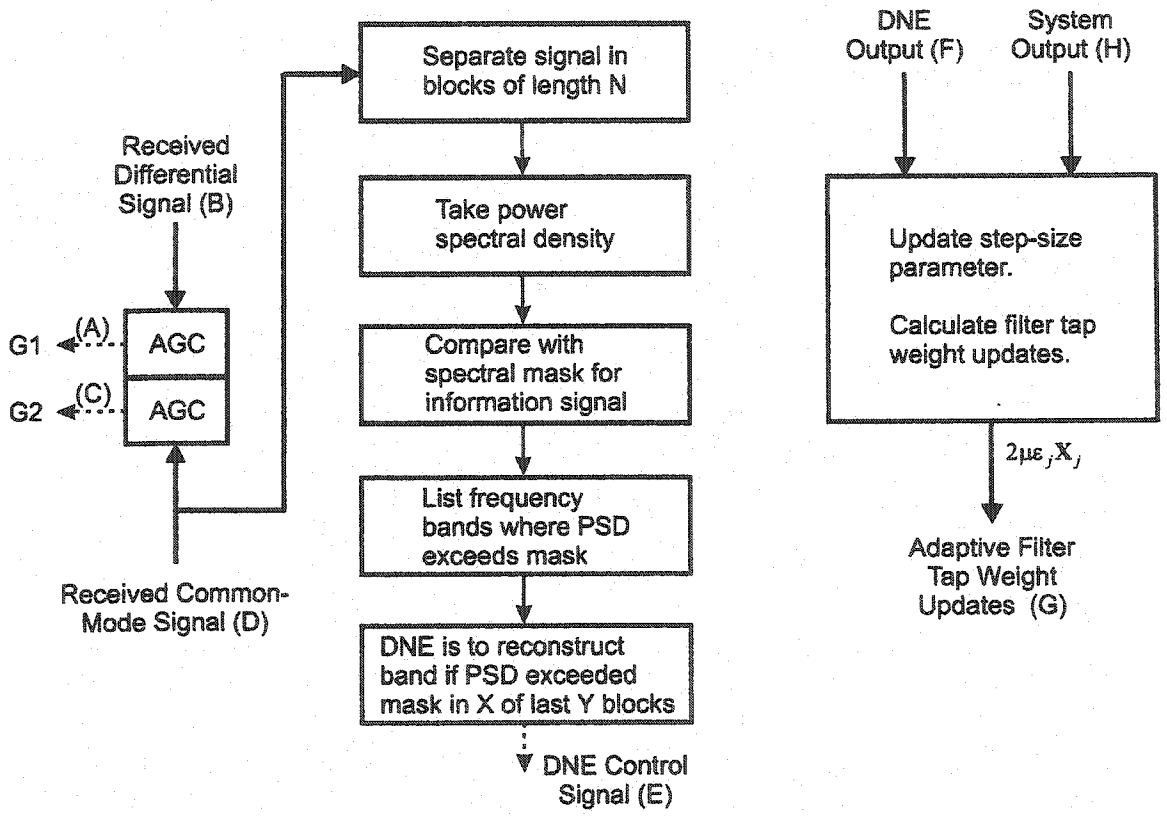

Figure 5-2: Functional diagram of control unit

There is a very important tradeoff to be considered in the DNE control algorithm. The fact that the adaptive filter's convergence is affected by the DNE output favours only very occasional changes in the DNE configuration. On the other hand, bursty signals such as amateur AM radio can appear suddenly, and the DNE should ensure that they are passed to the adaptive filter as fast as possible. Altogether, it is clear that the DNE and its requisite control circuitry add substantial

complexity to the system. The simulations in this research project do not address these complicated control issues, and instead concentrate on how great a problem the transfer of energy from differential mode to common mode really is. If the DNE does not provide any significant performance advantage, it should be omitted entirely.

The remainder of the system from Figure 5-1 is a conventional adaptive filter, with an additional fixed delay in the path of the differential signal to compensate for any delay incurred by signal divergence on the loop, or by DNE processing. The control unit is used to calculate the tap-input power using the DNE output, and the value of the step-size parameter μ is varied in inverse proportion with this number.

This system is clearly limited by the dynamic range of its analog-to-digital converters. A large interference component in the differential signal will force the AGC circuitry to adjust the gain G1 such that the ADC's quantization noise surpasses the thermal noise floor. This problem can be mitigated by allowing for some front-end analog cancellation.

5.2.2 Digital cancellation with a narrowband front-end analog canceller

Figure 5-3 shows a canceller with some additional front-end analog circuitry: bandpass filters to carry out spectral isolation for FDD, and a narrowband RFI canceller of the type disclosed in [45].

Referring to the figure, both the differential and common-mode signals are bandpass filtered to restrict them to the frequencies at which positive noise margin can be achieved and to remove any DC bias which the common-mode signal might have; this task would obviously be performed in any receiver, but the VDSL standard has yet to decide whether this filtering will be conducted digitally. These analog filters would introduce phase distortion that the equalizer would have to compensate for. However, the operation of the noise canceller would not be seriously affected by this additional distortion because both the reference and desired signals would experience it equally.

The reception of an amateur AM interferer with differential power of -10 dBm will be a very rare event; the reception of two such signals will be extremely rare. Accordingly, a single

narrowband canceller will almost always be sufficient to guard against occasions when an amateur radio signal might cause a –40 dB SIR. If there are concerns that an analog canceller would carry too high a hardware cost, or adapt too slowly to amplitude bursts or slight carrier frequency shifts, then bandstop filters could simply be used to blank the amateur signal at both the differential and common-mode inputs; this is what the feedforward equalizer in a QAM system would do in any case.

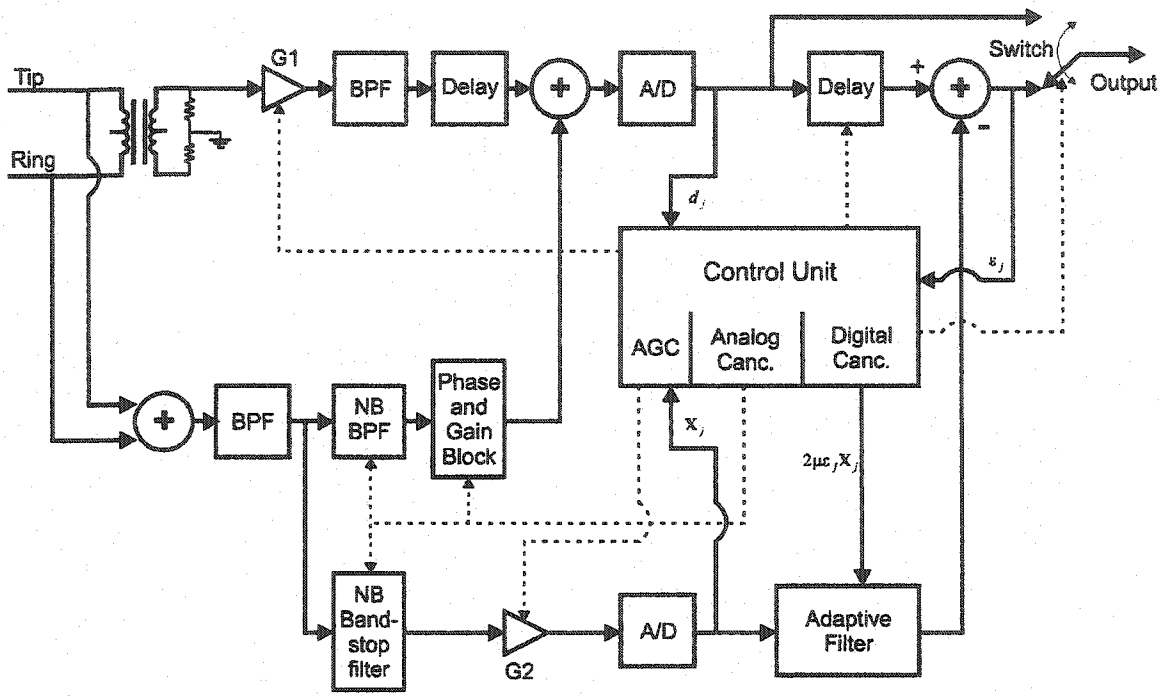

Figure 5-3: Digital noise canceller with front-end narrowband analog canceller

The system's operation can be described as follows. The differential signal is extracted from the loop and immediately amplified. It is then bandpass filtered to retain only those frequencies coinciding with the main lobe of the received signal. A fixed analog delay follows to compensate for the small delay incurred by the analog canceller. The common-mode signal is received and bandpass filtered to the same bandwidth as the differential signal. Note that it may first be necessary to attenuate the common-mode signal in order to ensure that it fits within the modem's power rails.

At startup, the analog noise canceller is disabled, and the automatic gain control circuits are allowed to adjust gains G1 and G2 such that ADC saturation is avoided. If possible, the common-mode and differential gains should be held at a fixed ratio to avoid excessive adjustment of the

adaptive filter (in fact, the AGC circuits, noise canceller and equalizer might have to adapt one at a time to prevent undesired interactions between the devices, and possible system instability). Assuming that the top and bottom reference voltages of the ADCs are fixed, it is possible to estimate whether or not quantization noise is raising the differential noise floor by monitoring the gain G1. If G1 is high, then there will be no need for analog cancellation. If G1 is low, then the control circuitry must determine whether this because there is a powerful narrowband interferer, or because the loop is very short and the received signal power is high. This function is most conveniently performed by the swept bandpass filtering arrangement described in [45], and included as part of the analog canceller's control circuitry (it is highly desirable to avoid performing digital PSD calculations from a hardware complexity point of view). If a narrowband interferer is present, the analog canceller may be activated according to the specifications programmed into the control algorithm.

The use of the analog canceller should be determined by the ambient noise environment and power consumption restrictions. If there is one overwhelming narrowband interferer, then it is clear that it should be activated. If there are two narrowband interferers of roughly equal power, then it should also be activated. This situation would allow the analog canceller to cancel one interferer, and leave the digital filter to carry out maximum cancellation of the other. However, care would have to be taken to ensure that the analog canceller remained locked to one of the interference signals, and did not "hunt" back and forth between the two. Last, if the noise environment is mild and fits cleanly into the ADC's dynamic range without raising the noise floor, then the analog canceller should be deactivated to alleviate power consumption and heat dissipation concerns.

If the analog canceller is enabled, it will operate in the same manner as described in Chapter 4. The controller will set the common-mode narrowband BPF parameters, as well as the phase and gain adjustments necessary to carry out cancellation. As mentioned before, communication between the controllers for the digital and analog controllers will be necessary to ensure that the analog device stays fixed on one interferer, thus preventing undue stress on the AGC circuitry or the adaptive filter. Communication between the AGC circuits and analog canceller is also necessary to ensure that the latter unit can make gain adjustments simultaneously with changing values of G1. The analog canceller has one other feature that was not necessary in [45]. Since the common-mode noise signal must still be used to conduct wideband cancellation, a narrowband bandstop filter is used to blank the frequencies at which the analog canceller operated.

The digital canceller which encompasses the remainder of Figure 5-3 is the same as described in Section 5.2.1, and may include a dynamic noise estimator (the DNE has been omitted here for simplicity). Also pictured in the figure is an output switch to select either the raw differential output or the adaptive filter output; when the induced noise levels are low, the numerical errors introduced by the adaptive filter are detrimental to the performance of the receiver, and the output switch allows the digital canceller to be bypassed.

5.3 Research Objectives

The primary goal of this research is to evaluate the performance of a wideband digital noise canceller for VDSL. Given that interference levels above –20 dBm are very rare, this project will not attempt to evaluate the performance of a front-end analog canceller. Simulations of the system depicted in Figure 5-1 are to be carried out, with the addition of two front-end analog BPFs to conduct FDD filtering. The simulations will attempt to answer the following questions:

a. Is the reference noise likely to be so ill-conditioned that digital noise cancellation will not be possible? How many bits of precision are needed for the adaptive filter if fixed-point arithmetic is to be used?

b. What values of the step-size parameter μ (or if LMS with adaptive gain is used, β) give the best performance?

c. What is the best reduction in noise power that can be achieved for a severe AM ingress threat? At what SIR level will the canceller serve only to introduce extra error into the received data stream?

d. Is the inclusion of a digital noise estimator (DNE) worth the significant hardware cost which it will incur?

e. How many taps should the adaptive filter have for best performance on a given length of loop?

f. Does LMS with adaptive gain significantly outperform the much simpler normal (leaky) LMS algorithm?

g. Will analog front-end filtering be able to bring the magnitude of most impulse events into the ADC range? If so, will the canceller likely be able to achieve any cancellation of these impulses?

6. Simulated Performance of the Digital Noise Canceller

6.1 Description of the Simulation

The simulation program is written in Matlab, and tests a digital noise canceller in a VDSL environment. Extension of the idea to ADSL is simple: on long loops with very high insertion loss, the ADSL scenario will be identical to VDSL, focusing on preserving the external noise floor at −140 dBm/Hz. On shorter loops, the ADSL signal will likely be significantly more powerful than the AM noise, and noise cancellation will be less critical.

The program carries out the following tasks in succession:

a. Generation of a QAM-16 signal. Two random streams of 4-PAM symbols are generated and raised-cosine filtered with roll-off 0.4. The two symbol sequences are then modulated on the I- and Q-channels of a given carrier frequency. A sampling frequency of 35 MHz is used for all simulations, and the carrier frequency and total data rate are selected according to the loop length:

Table 6-1: QAM signal parameters

Loop length	Carrier frequency	Data rate
500 m	9 MHz	46.67 Mbps
1 km	4.5 MHz	20 Mbps
1.5 km	2.5 MHz	10 Mbps

The signal is scaled to fit under the −60 dBm/Hz transmit power mask, and its average power is calculated. If the user desires it, the program can retain the same symbol stream as was used in a previous simulation run.

b. Generation of AM interferers. Ten broadcast AM interferers are generated in accordance with the VDSL requirements document [9], and as described in chapter 2. All of the noise is produced at common-mode power levels, and is later coupled over to differential mode. It is initially generated under the assumption that the loop has a constant balance of 60 dB up to 1.6 MHz, as detailed in [9]. In cases where lower

balance is desired, the AM noise is appropriately scaled to ensure that it couples to the loop with the differential strength given in Table 2-3.

When broadcast AM threat 1 is used, three amateur radio interferers are generated at frequencies of 3.9, 7.05 and 10.1 MHz, and with common-mode average powers of 25, 30, and 25 dBm. For threat 2, only the 3.9 and 7.05 MHz interferers are used (with average powers 5 and 10 dBm respectively), and for threat 3 there is no amateur radio interference. The amateur signals are interrupted periodically to simulate speech, as described in [9]. No mobile or shortwave interference sources are included in the noise model. Altogether, the PSD of the threat 1 RFI in common mode takes the shape:

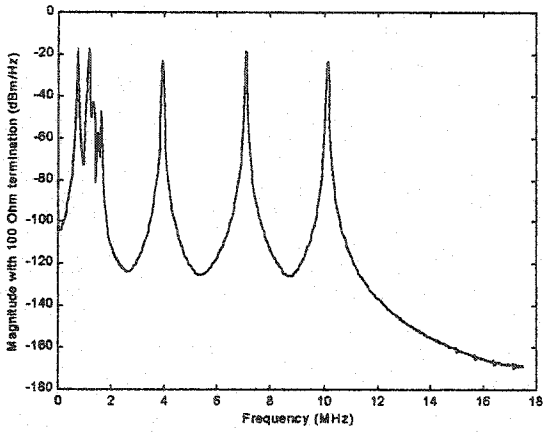

Figure 6-1: PSD of RFI in common-mode

c. Generation of the channel impulse responses. The impulse responses of the differential and common-mode channels are derived using two-port theory and the parameters described in chapter 2. The simulation is capable of constructing the ten different loops described in Table 6-2, where all cable is AWG 26.

In the cases where the RFI must be injected to the loop somewhere other than the receiver, the channel is broken into two impulse responses, with the first one using the loop's characteristic impedance Z_{0d} or Z_{0c} as its load. When the two impulse responses found in this manner are convolved, the result is extremely close to the end-to-end impulse response that would have been found without separating the loop into two parts.

Table 6-2: Loop models used for the simulations

Loop designator	Description	Distance from receiver at which RFI couples to loop (metres)
11	500-m cable	0
12	500-m cable	250
13	500-m cable, 100-m bridged tap, 50-m cable, 50-m bridged tap, 50-m cable	100
21	1-km cable	0
22	1-km cable	250
23	1-km cable, 100-m bridged tap, 50-m cable, 50-m bridged tap, 50-m cable	100
24	1-km cable	175
25	1-km cable	500
31	1.5-km cable	0
32	1.5-km cable	250

d. Generation of impulse noise. Impulse noise is generated as described in chapter 2. Each impulse is 3,500 samples long, and the spacing between impulses is 10,000 samples.

e. Cable attenuation, imbalance and injection of RFI and impulse noise. Attenuation and phase distortion is carried out using the impulse responses generated in part (c). The RFI and impulse noise are coupled to differential mode with a user-specified constant balance from DC to 1.6 MHz, and declining at 12 dB/decade after that.

The cable imbalance must be simulated carefully to ensure that the principle of conservation of energy is not violated. Cable imbalance usually occurs because the two wires of the pair see different impedances to ground. In other words, one signal is either attenuated and/or phase-shifted differently than the other. It is possible to work out a general solution relating both the attenuation and phase shift to the balance, but it is easiest to simply assume that the phase shift is zero. For example, consider two signals propagating down the two wires of a cable:

$$\begin{aligned} s_1(t) &= \cos(\omega t), \text{ and} \\ s_2(t) &= A\cos(\omega t) \end{aligned} \tag{6.1}$$

where A is obviously a scaling factor describing the amplitude of signal 2 relative to signal 1, and is assumed to be between 0 and 1. The common-mode signal at the end of the loop is $s_1(t)+s_2(t)$, and the differential signal is $s_1(t)-s_2(t)$. Referring back to equation (2.8), it can be seen that the balance in this situation is:

$$B = 20 \cdot \log_{10}\left(\frac{1+A}{1-A}\right) \quad \text{(dB)}. \tag{6.2}$$

Using this formula, it is possible to generate a graph giving A as a function of the balance:

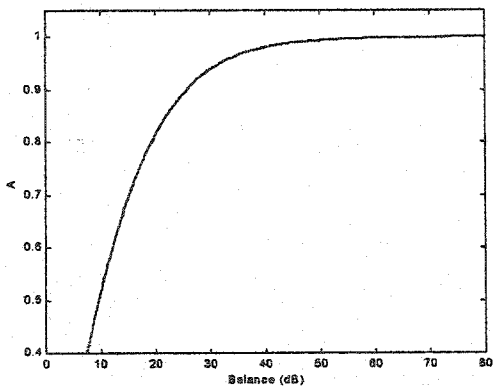

Figure 6-2: Scaling factor A vs. cable balance

Given any desired characteristic of cable balance versus frequency, it is possible to interpolate appropriate values of A from Figure 6-2. The vector of A versus frequency can then be used to design a linear-phase filter. The filter is used to scale a second copy of the common-mode signal, and then the signal and its copy are subtracted to generate the component that leaks to differential mode, and added to find the component which remains in common mode. The same approach is used with the differential signal, and the total power on the two wires remains constant. In practice, this method has been found to achieve balances very close to the desired characteristic, provided that the filter has a sufficient number of coefficients.

Aside from the common-mode and differential signals, a "benchmark" signal is also generated. This is an attenuated but otherwise uncorrupted version of the information

signal, and it is used as the correct value of the signal when MSE calculations are performed.

f. Capacitive leakage across the hybrid transformer. The bridge capacitance of the transformer in this simulation is set at 10 pF, which gives it about 83 dB balance at 1.1 MHz and 60 dB balance at 17.5 MHz. The 90° phase shift is carried out by multiplying the FFT by $\sqrt{-1}$ between $\omega = 0$ and $\omega = \pi$. After performing symmetric extension of the FFT, it is inverted to get a real output.

g. Front-end analog filtering. Front-end analog filters bandlimit both the common-mode and differential signals to the region between 1.1 MHz and the upper frequency detailed in Table 3-2, depending on the loop length. The benchmark signal is also passed through the same filter. The filters are designed with slightly over 60 dB of stopband attenuation and a moderate roll-off, both realistic for actual analog filters (an eighth-order elliptic lowpass filter can achieve 80 dB rolloff at about 1.25 times the cutoff frequency [61]). However, they are implemented as finite impulse response filters in the simulation for convenience. The filters therefore will not distort the amplitude and phase of their input signals as a real analog filter would. This is acceptable because it is the receiver equalizer's job to compensate for the filter distortion, and the equalizer is not simulated in this research. Furthermore, the same filter distortion will apply to both the differential and common-mode signals, so it will not affect the operation of the noise canceller.

h. Additive white Gaussian noise. Separate sequences of AWGN with PSDs at −140 dBm/Hz are generated and added to the differential and common-mode signals.

i. Analog-to-digital conversion. In a hardware implementation, amplification of the two inputs would be performed in the receiver front end. However, it is skipped in the simulation to allow calculations pertaining to the ADC step size to be performed using the received power levels; AGC functions for the ADCs have not been simulated. To compensate for the various precision errors introduced by non-ideal 12-bit ADCs, only 11-bit ideal ADCs are used for both inputs. It is assumed that both slowly adjust their step sizes δ to use 85% of their ranges, not including impulse activity. On the loops with

very low SIRs, this occasionally forced the differential device to adjust its δ up to 20% higher than the maximum desired level set in chapter 3 (160 μV).

j. Digital noise estimator. A variety of different DNE designs were tried in the course of the simulations. Unfortunately, this application favours filters which have sharp transitions and the linear-phase property. The first property allows the DNE to pass as little information signal energy as possible, and the second property allows the DNE output to be reconstructed through simple addition of the frequency bands containing narrowband interferers. These two properties imply that FIR filters with many coefficients must be used.

The DNE filter is implemented as a 32 or 64-band discrete Fourier transform filter bank with a windowing filter capable of 60 dB stopband attenuation. Complementary bands from the filter are then added together to generate either 17 or 33 bands of real numbers. Those bands which exceed a user-specified PSD threshold are then added together and passed to the adaptive filter.

For reasons that will become apparent shortly, not much effort was put into advanced development of the DNE. A full control algorithm was never completed, so in simulation the unit reconstructs its output based on the PSD of the entire noise sequence. In addition, the DNE was never implemented in fixed-point mathematics.

k. Adaptive filter. The adaptive filter implements either the leaky LMS algorithm or LMS with adaptive gain. It can be implemented with floating-point ("infinite-precision") or fixed-point mathematics. The following filter characteristics are user-selectable:

i. the number of filter taps, and whether or not the differential signal is to be delayed by a number of samples equal to the number of taps;

ii. β. For LMS with adaptive gain, this is the step-size parameter used in the calculation of new values for μ. For normal LMS, β is used to set μ as follows:

$$\mu = \frac{\beta}{\text{tap - input power}} \tag{6.3}$$

iii. the number of extra bits of precision that the filter will use internally. Even though both inputs use 11 bits, it is beneficial to carry some additional precision inside the canceller before outputting a 11-bit number;
iv. whether rounding or truncation is to be used when fixed-point mathematical operations are carried out; and
v. the leakage factor $2\mu\alpha$ used for leaky LMS.

Fixed-point math is implemented using the "round" or "floor" functions in Matlab. Both input streams are divided through by 2^{11} to give sequences of numbers between -0.5 and 0.5. Filter overflow occurs if any number exceeds +1 or falls below −1 (except the tap-input power, which carries several extra bits of precision). For the multiplication of numbers X and Y using A bits of precision and rounding, the following command was carried out:

$$\text{mult}(X, Y) = \frac{\text{round}\left(X \times Y \times 2^{A}\right)}{2^{A}} \tag{6.4}$$

This implementation was grossly inefficient, and future simulation work should be carried out using proper binary math and a compiled programming language such as C/C++.

The tap-input power is calculated, assuming that the reference noise is ergodic, with the following recursive relationship in infinite precision:

$$\text{tap - input power} = 0.999 \times \text{tap - input power} + 0.001 \times \mathbf{X}^{T}\mathbf{X} \tag{6.5}$$

The formula is recalculated with every new input to the adaptive filter, although in an actual hardware implementation this rate of computation could be relaxed. In fixed-point math, this equation is honoured as closely as possible. Aside from the imprecision incurred by the individual multiplications in the inner product $\mathbf{X}^{T}\mathbf{X}$, the scaling factors 0.999 and 0.001 will vary with the number of bits of precision. The lowest number of

bits considered will be 10, so the scaling factors are never allowed to drop all the way to 1 and 0 respectively.

l. Digital-to-analog conversion. Although an actual hardware implementation of the canceller would pass its digital output directly to the rest of the receiver, here the output signal is processed through an ideal 11-bit digital-to-analog converter. The D/A converter uses the same step size as the differential ADC, facilitating the comparison of the output signal with the benchmark signal.

m. Calculation of mean-squared error and noise floor. The mean-squared error between the DAC output and the benchmark signal is computed. Calculations begin after 1000 samples have been processed by the adaptive filter. A second calculation is performed to measure the MSE between the received differential signal and the benchmark signal. Since the information signal energy is approximately the same at the input and the output of the canceller, the ratio of the two MSEs can be viewed as the reduction in noise power achieved by the canceller. In the simulation results, this ratio is referred to as the "noise reduction", and is expressed in dB:

$$\text{Noise reduction} = 10 \cdot \log_{10}\left(\frac{\text{MSE between rxed differential signal and benchmark signal}}{\text{MSE between noise canceller output and benchmark signal}}\right) \tag{6.6}$$

Note that this parameter is directly analogous to the processing gain given by equation (4.70).

The ratio of the two MSEs is not a sufficient statistic by itself. As the size of μ increases, so does the level of gradient noise, which usually appears as a white noise floor. If μ gets too large, the gradient noise can completely eclipse the differential signal. In less extreme circumstances, the rising noise floor will raise the noise floor significantly. For this report, approximate measurements of the output noise floor were obtained by taking a 1000-point PSD of the output signal and measuring the top and bottom values. The maximum of these two values is retained, and the other is discarded.

6.2 A Note on Statistical Accuracy

Since the LMS algorithm is a stochastic gradient method, its performance should be evaluated using ensemble-averaged values resulting from a large number of independent trials. Given the very long run times of some of the simulations used in this study, extensive ensemble averages were seldom taken; in most cases, the figures presented are the averaged results of only two simulation runs. It is therefore important to develop some concept of the experimental error introduced by this shortfall.

Twenty-five independent simulations were run on two different loops to test the noise reduction capabilities of leaky LMS in the presence of RFI and AWGN. For each run, new data and RFI sequences were generated. Table 6-3 presents the parameters used; the reasons for their selection will become clear in the following pages. The results achieved are summarized in Table 6-4. For these two cases, the mean value of the noise reduction was biased slightly towards the maximum value, but the 25 values were too few to allow recognition of any familiar probability distribution. The minimum noise reduction values on the loops were 1.86 dB and 1.78 dB away from the mean values. Adding a safety margin of an extra decibel, an individual simulation run has a very good probability of being within 3 dB of the mean. Until more comprehensive follow-up testing of the noise canceller is completed, large error margins of ±3 dB should be assumed.

Table 6-3. Summary of statistical accuracy example

# of QAM symbols	50,000 (loop 11), 25,000 (loop 21)	Digital noise estimator	no
Sampling frequency	35 MHz	DNE frequency bands	N/A
Loops	11 (500 m), 21 (1 km)	DNE stopband atten.	N/A
Low-frequency balance	40 dB	Differential delay	no
RFI	Threat 1	# of adaptive filter taps	50
Impulse noise	none	Filter leakage factor $2\mu\alpha$	0.0002
Leakage at hybrid	yes	β	β = 0.007 (loop 11), β = 0.04 (loop 21)
AWGN	−140 dBm/Hz	Total bits precision	14
LMS algorithm	leaky	Flooring or rounding	rounding

Table 6-4: Results of statistical accuracy experiment

Loop	Mean noise reduction	Maximum noise reduction	Minimum noise reduction	Difference between min and max
11	21.86 dB	23.34 dB	20.00 dB	3.34 dB
21	31.29 dB	32.78 dB	29.51 dB	3.27 dB

6.3 Case 1: Best values for β and number of bits of precision

The first set of simulations was performed using relatively simple loop conditions, and was meant to assess:

a. the best noise reduction that could be achieved by the canceller;
b. the values of β that led to peak performance; and
c. the number of bits of precision that would be required for the internal calculations of the noise canceller.

The test conditions for the 500-m loops are summarized Table 6-5. Two runs of each simulation were conducted. The noise consisted solely of severe RFI and standard AWGN, and there was capacitive leakage at the hybrid. The digital noise estimator was not used. The adaptive filter was implemented with 150 taps because 150 was a round number that gave a filter resolution of approximately 100 kHz. No differential delay was used because it was presumed that the worst SIRs would occur when the noise transferred from common mode to differential mode right at the hybrid.

Table 6-5: Case 1A Summary

# of QAM symbols	50,000	Digital noise estimator	no
Sampling frequency	35 MHz	DNE frequency bands	N/A
Loops	11, 12, 13 (500 metres)	DNE stopband atten.	N/A
Low-frequency balance	40 dB	Differential delay	no
RFI	Threat 1	# of adaptive filter taps	150
Impulse noise	no	Filter leakage factor $2\mu\alpha$	0
Leakage at hybrid	yes	β	$0.0004 < \beta < 0.4$
AWGN	–140 dBm/Hz	Bits of precision	∞, 10 to 18
LMS algorithm	normal	Flooring or rounding	rounding

The low-frequency balance of the line was dropped to 40 dB simply to keep the level of the interference realistic. With a 60 dB balance, RFI threat 1 would have implied peak common-mode voltage levels in the vicinity of 70 V. This amplitude would be possible in very close proximity to a powerful radio transmitter, but such a situation is very rarely encountered.

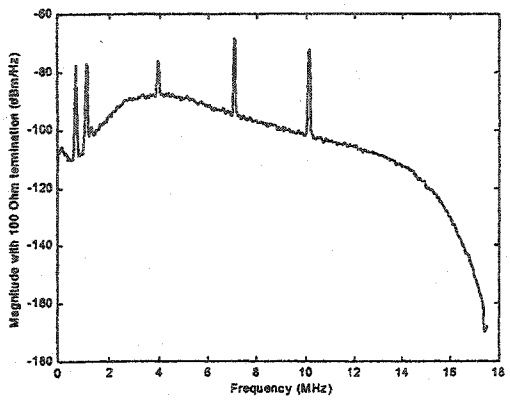

Figure 6-3: PSD of differential signal at input to hybrid device

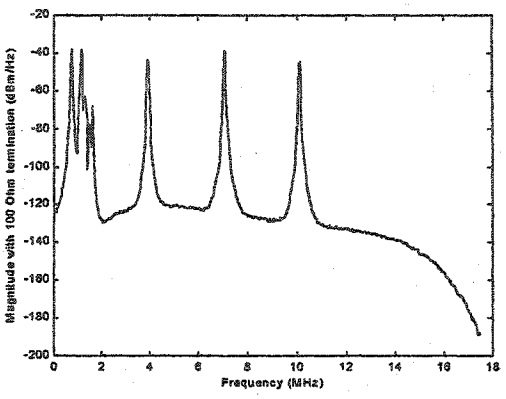

Figure 6-4: PSD of common-mode signal at input to hybrid device

Before analyzing the results, it appropriate to provide a qualitative example of a typical simulation run. The power spectral densities in Figures 6-3 to 6-7 are taken from the test of loop 11, which experiences very powerful RFI but not much attenuation. At the hybrid, the differential signal has a SIR of –4.3 dB and the PSD shown in Figure 6-3. The narrowband interference is clearly visible. Similarly, since the loop's balance is only 40 dB or lower, traces of the QAM signal are visible in common mode (Figure 6-4).

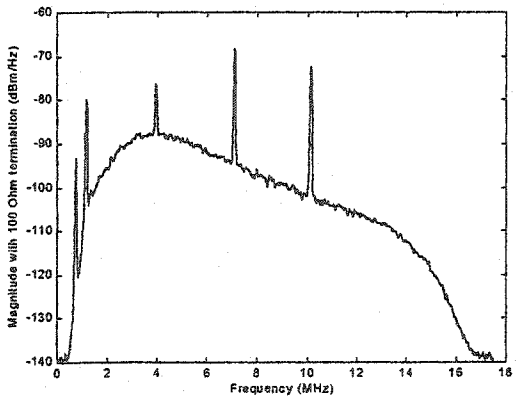

Figure 6-5: PSD of differential signal after analog BPF and ADC

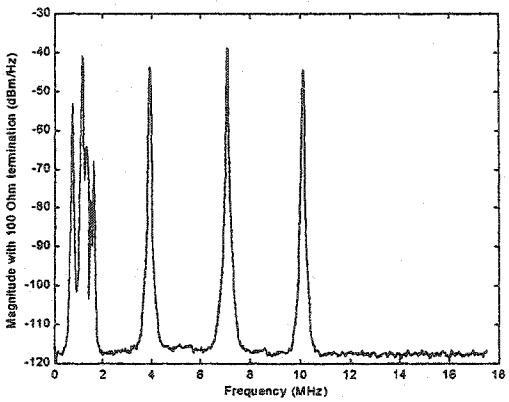

Figure 6-6: PSD of common-mode signal after analog BPF and ADC

After the analog BPF, the SIR of the differential signal improves slightly, to –4.0 dB, because the low-frequency broadcast AM stations are attenuated. When AWGN is added and the signal is converted to digital form, the PSD appears as in Figure 6-5. The lower-frequency AM

stations are also partially removed from the common-mode signal. After the ADC, the common-mode quantization noise floor falls at approximately −117 dBm/Hz (Figure 6-6). The leaked components of the QAM signal are almost completely eclipsed by thermal and quantization noise.

At the noise canceller output (Figure 6-7), the powerful broadcast AM station at 1.13 MHz and the amateur signals at 3.9 and 7.05 MHz are almost perfectly cancelled. Residual levels of the other two primary interferers remain.

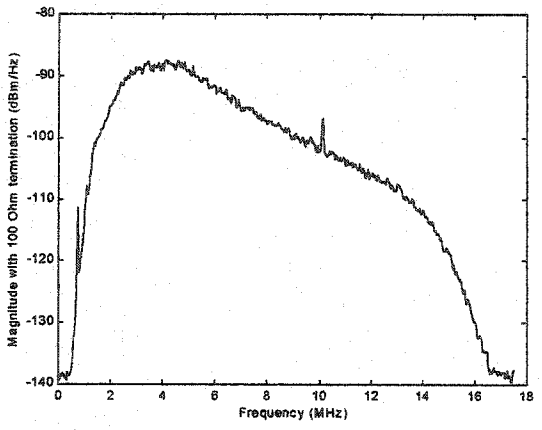

Figure 6-7: PSD of noise canceller output

In this case, the cancellation is very good, and the noise floor only rises by one or two decibels. On longer loops, which experience a more severe asymmetric loss (ie. a more pronounced slope from low frequency to high frequency), one signal will still be cancelled very well. However, the residual levels of other signals at the canceller output relative to their strengths at the input tend to be higher than shown for loop 11. In such cases, the back-end noise cancellation of the QAM or DMT receiver would have to compensate for this performance shortcoming.

The first observation arising from the Case 1 simulations was that the value of β had to be kept quite small. Almost all values above 0.2 caused unacceptable gradient noise which raised the noise floor by at least 5 dB. Above $\beta = 0.4$, the signal was obscured by gradient noise so badly that it was completely unrecoverable. Once β reached levels above 0.6, the system was very likely to become unstable. Note that the values of μ which guaranteed convergence in equation (4.24) were obtained by assuming that $\mathbf{R}_X$ was well-conditioned; in the Case 2 simulations, this presumption will be shown to be false, thus explaining the early onset of instability.

The decision concerning the required number of bits of precision must be made while considering both the noise reduction and the noise floor. Aside from ADC quantization, the noise floor can be raised by gradient noise (primarily dependent on β) or finite-precision effects in the adaptive filter (dependent on the number of bits of precision). For these simulations, it was assumed that a noise floor up to 3 dB higher than the thermal noise floor was permissible; in practice, this will probably be far lower than the crosstalk levels on most loops. In the plots of noise reduction versus β given throughout this chapter, the points are encoded as follows:

a. circle: noise floor under –137 dBm/Hz (good);
b. square: noise floor between –137 dBm/Hz and –135 dBm/Hz (marginal); and
c. asterisk: noise floor above –135 dBm/Hz (unacceptable).

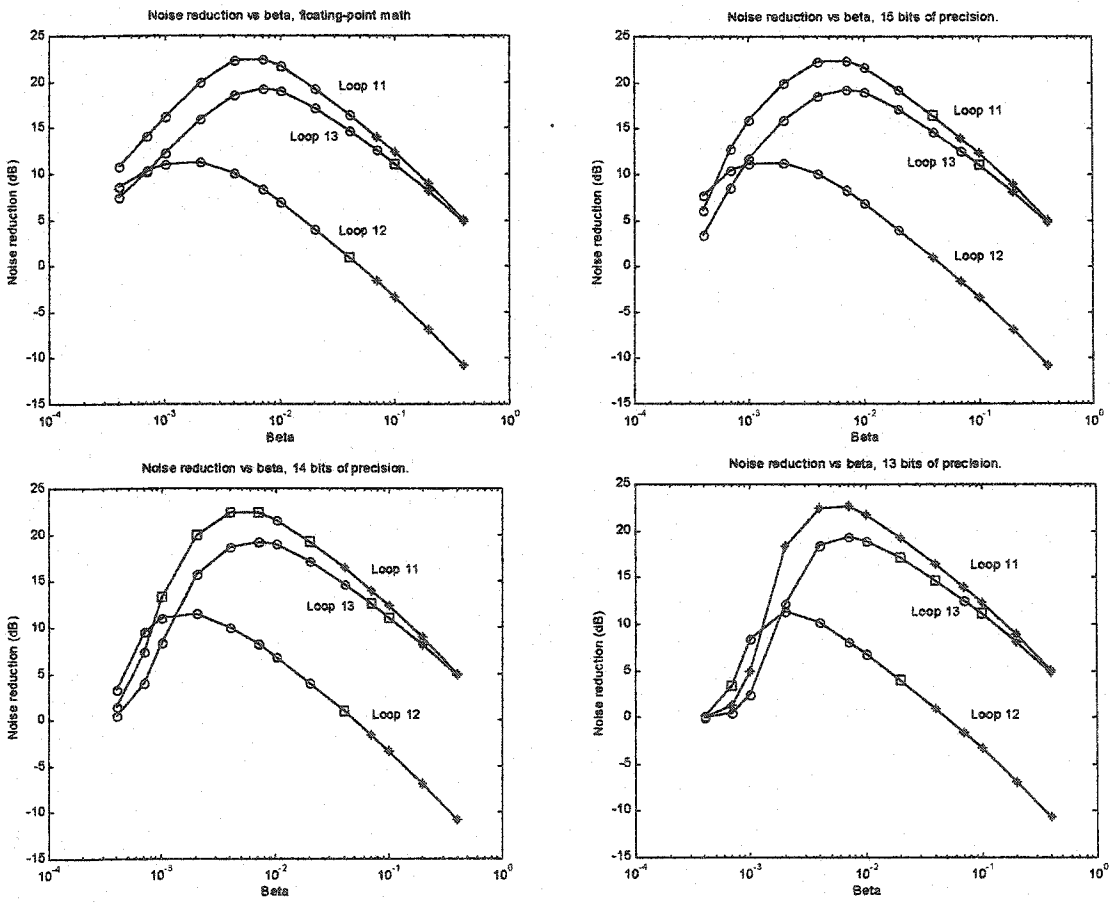

Figure 6-8: Case 1A, Noise reduction vs. β

Figure 6-8 shows the noise reduction achieved on the 500-metre loops. Using IEEE floating-point arithmetic, the noise reduction on the loop with the most interference reached a maximum of 22.5 dB; this value is within the ensemble bounds given previously in Table 6-4. For β values of 0.07 and higher, the noise floor rose unacceptably. In fact, at these β values the system added noise to loop 12, which had a relatively high SIR of 11.82 dB at the canceller's differential input.

Evaluations were also conducted using 10 to 18 bits of precision. In general, the results for 12 bits and under were completely unacceptable, and those for 16 bits and over were very good. The results for 13 to 15 bits are also shown in Figure 6-8. It can be seen that the noise floor for loop 11 rises unacceptably when 13 bits are used. For 14 bits, the floor rises by 3 - 5 dB, but the values are quite close to the 3 dB limit. A quick glance at Figure 6-7 confirms that this short loop has enough signal energy to provide a good margin against noise even if the floor rises to −135 dBm/Hz, and so 14-bit math is acceptable. 15-bit math gives good performance on all three loops.

Table 6-6: Case 1B Summary

# of QAM symbols	25,000	Digital noise estimator	no
Sampling frequency	35 MHz	DNE frequency bands	N/A
Loops	21 - 25 (1 km)	DNE stopband atten.	N/A
Low-frequency balance	40 dB	Differential delay	no
RFI	Threat 1	# of adaptive filter taps	150
Impulse Noise	no	Filter leakage factor $2\mu\alpha$	0
Leakage at hybrid	yes	β	$0.0004 < \beta < 0.4$
AWGN	−140 dBm/Hz	Bits of precision	∞, 10 to 18
LMS algorithm	normal	Flooring or rounding	rounding

Case 1B ran the same test on the five 1-km loops. Since the loss experienced by the differential signal was more extreme than for the 500-m loops, better noise reduction was achieved. However, as mentioned earlier, higher levels of residual interference were usually apparent at the upper end of the output signal's frequency spectrum. Figure 6-9 shows the noise reduction factors obtained. For differential input SIRs ranging from −20.4 to 5.4 dB, maximum noise reductions between 31.7 and 15.8 dB were observed. Also, 13-bit mathematics once again led to an undesirable increase in the noise floor, while 14 bits provided adequate accuracy.

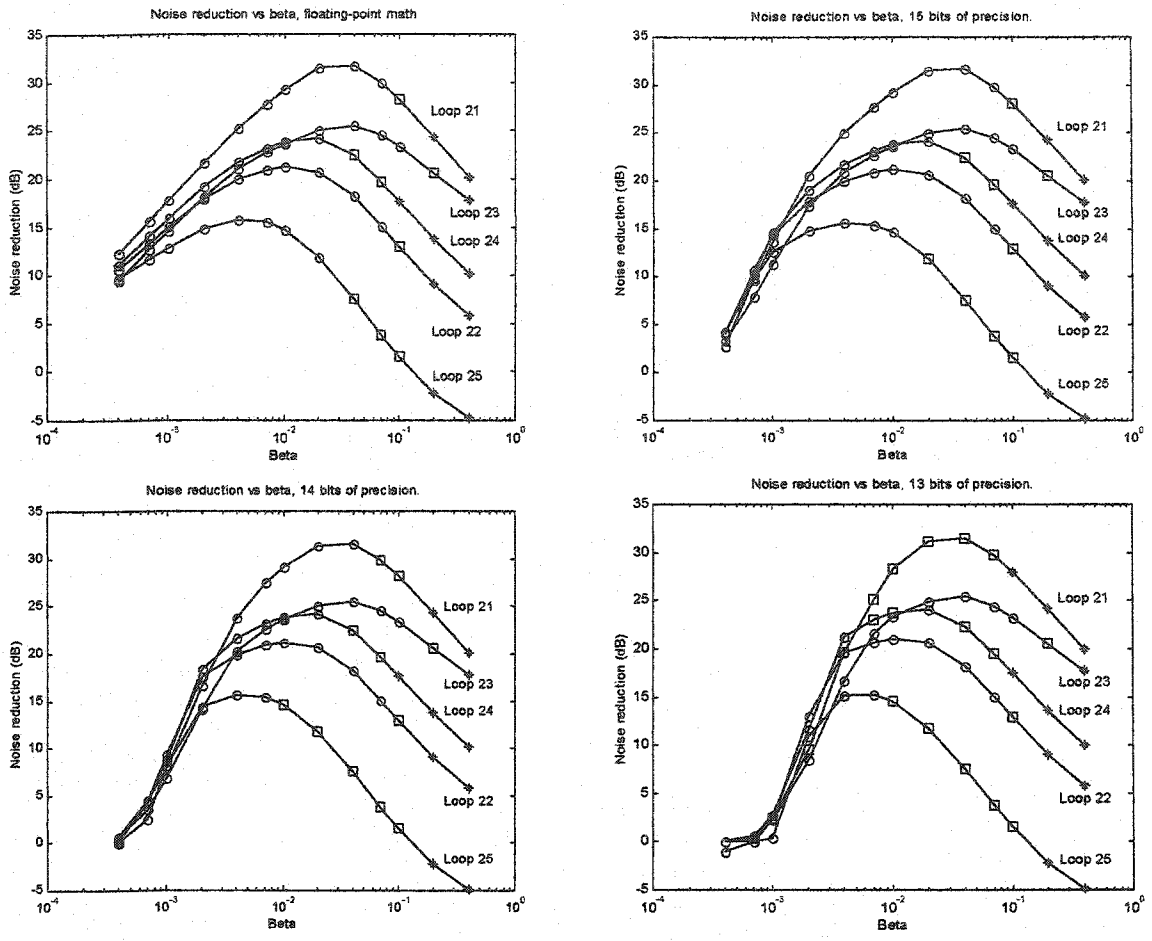

Figure 6-9: Case 1B, Noise reduction vs. β

Table 6-7: Case 1C Summary			
# of QAM symbols	12,500	Digital noise estimator	no
Sampling frequency	35 MHz	DNE frequency bands	N/A
Loops	31, 32 (1.5 km)	DNE stopband atten.	N/A
Low-frequency balance	40 dB	Differential delay	no
RFI	Threat 1	# of adaptive filter taps	150
Impulse Noise	no	Filter leakage factor $2\mu\alpha$	0
Leakage at hybrid	yes	β	$0.0004 < \beta < 0.4$
AWGN	−140 dBm/Hz	Bits of precision	∞, 10 to 18
LMS algorithm	normal	Flooring or rounding	rounding

Case 1C tested the two 1.5-km loops, which had very low signal energy (approximately −50 dBm after the analog BPF). This gave them differential input SIRs of −23.5 dB and −12.4 dB, as compared to −20.4 and −5.0 dB for the similar 1-km loops (21 and 22). However, since the signal power was much closer to the thermal noise power, the noise reduction achieved was several dB poorer than for the 1-km loops (see Figure 6-10). On the other hand, the low signal power

allowed the ADC to place the bottom end of its dynamic range so low that the noise floor did not become a concern until only 12 bits of precision were used.

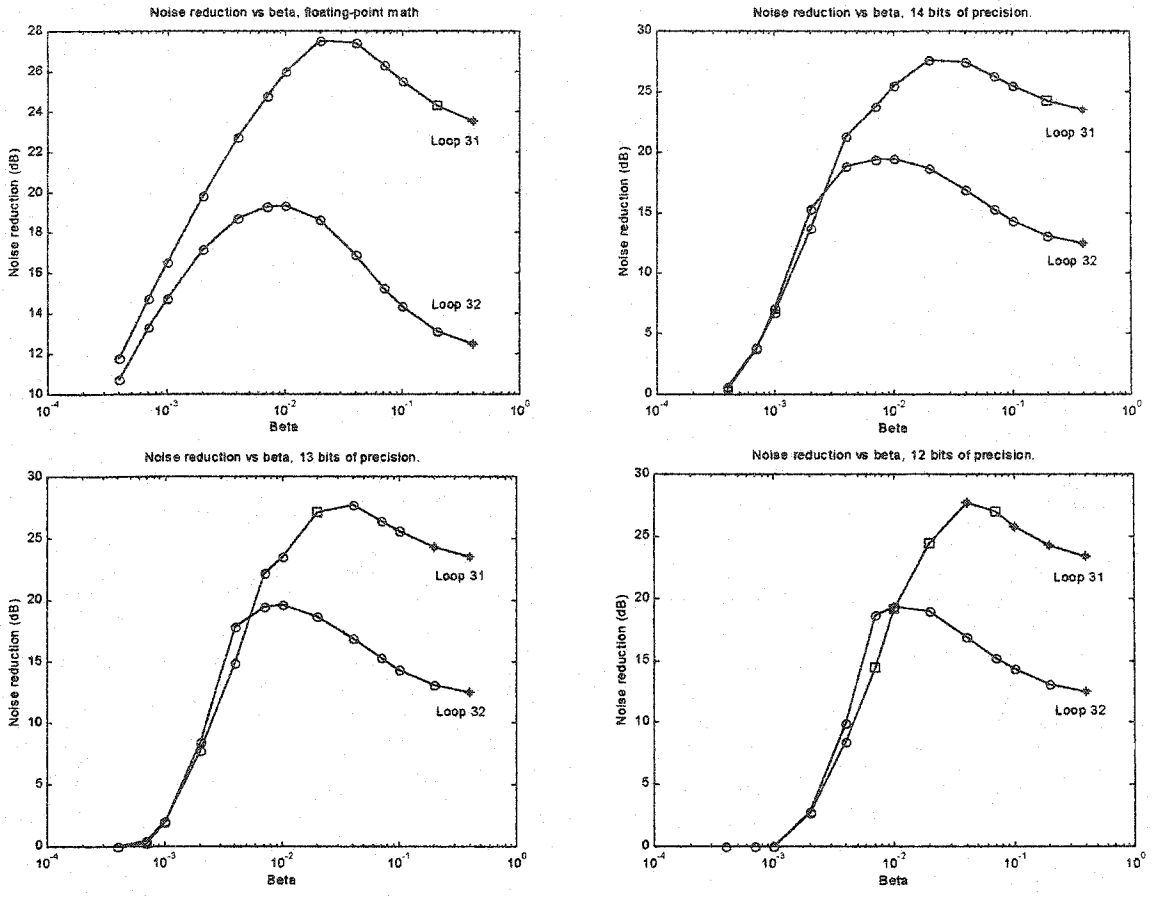

Figure 6-10: Case 1C, Noise reduction vs. β

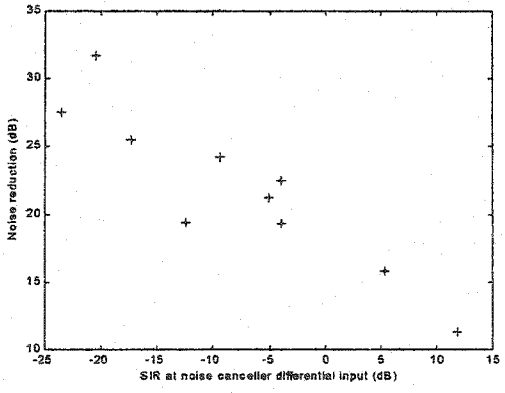

Figure 6-11: Best noise reduction vs. differential SIR, all 10 loops

The best noise reduction achieved for each of the ten loops is plotted against the differential SIR in Figure 6-11. The points on the graph are too widely spread to allow curve-fitting, but, as expected, they do show that the canceller achieved better cancellation as the differential SIR decreased.

Combining the results from all loops less than 1 km long (and including some data not otherwise presented in this section), Figure 6-12 shows a relationship between the differential SIR and the value of β which gave the best noise reduction for RFI threat 1. The line shown in the figure was fitted to the points on the log-log plot in a least-squares sense, and is described by:

$$\beta = 10^{(-0.0365 \cdot SIR - 2.1027)}, \qquad -27 \le SIR\,(\text{dB}) \le 35 \tag{6.7}$$

Since the 1.5-km loops were more sensitive to gradient noise, they required lower β values for best performance. These could be approximately found by halving the numbers calculated from equation (6.7).

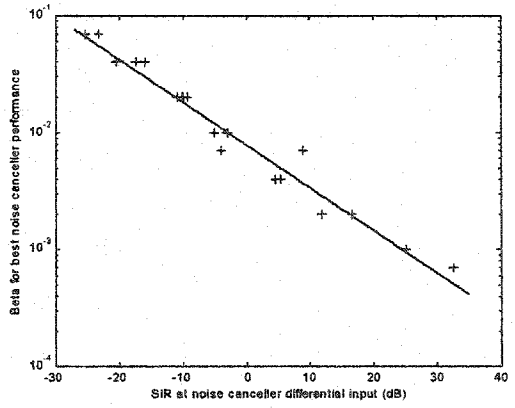

Figure 6-12: Best value of β vs. SIR of differential input to canceller

The model described by equation (6.7) was found to work very reliably when the noise threat included many narrowband interferers with similar average powers (within 20 dB). Under these circumstances, the probability density of the common-mode input samples was approximately Gaussian, a consequence of the central limit theorem. The broadcast AM threat can be expected to have a Gaussian distribution in urban areas, where both [18] and [21] found that there were likely to be several signals with similar average powers.

When a single interferer dominates all others, equation (6.7) usually yields a β for adequate, but not optimal, noise reduction. This is because an individual interferer does not have a Gaussian distribution. Figure 6-13A shows a histogram of 100,000 30%-modulated broadcast AM voltage samples. Its large carrier component gives it two peaks at its outer limits. In contrast, the histogram of 300,000 SSB amateur AM voltage samples is shown in Figure 6-13B. This signal has a density function which appears Laplacian; this is sensible, given that the signal's envelope is shaped by speech, which is usually modelled by a Laplacian distribution [46]. The problem of reconciling the single-interferer scenario with equation (6.7) has not yet been extensively studied, as the appearance of one dominant interferer is a rare event. However, this issue will need to be addressed in future research before a robust control algorithm for the canceller can be designed.

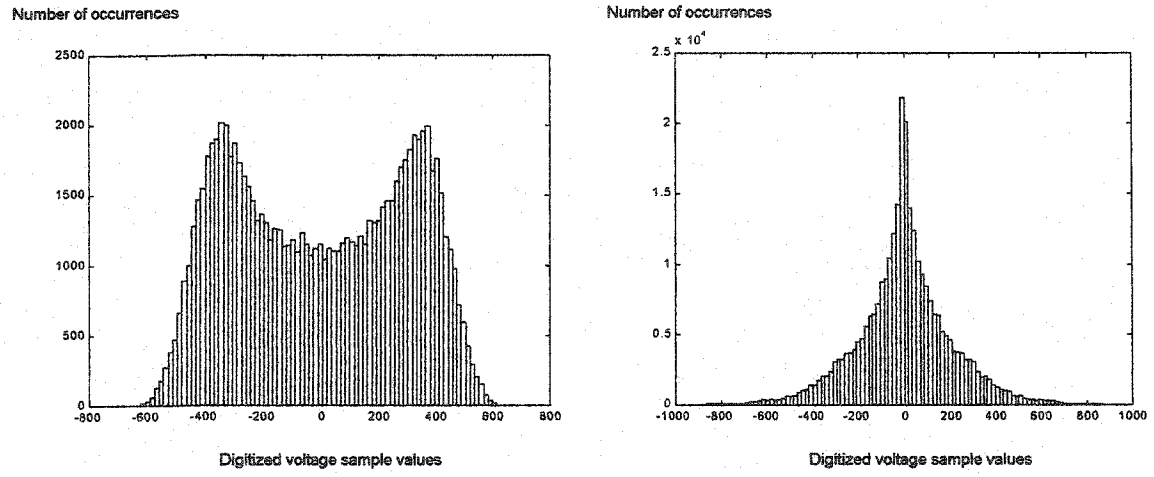

A. 30%-modulated broadcast AM *B. SSB amateur AM*

Figure 6-13: Histograms of individual narrowband interferers

6.4 Case 2: Truncation vs rounding in fixed-point math

This test was similar to the test carried out on loop 21 in case 1, except that truncation was used with fixed-point math, and only 15,000 QAM symbols were transmitted. Where possible, the use of truncation is much more desirable than rounding; in hardware, the use of rounding requires greater circuit complexity, and incurs extra delay as carry bits ripple through the adder circuit. However, the benefit of rounding is that the mean value of the finite-precision errors ($E\{b_j\}$ from equations 4.32 - 4.35) is much more likely to be close to zero.

Table 6-8: Case 2 Summary			
# of QAM symbols	15,000	Digital noise estimator	no
Sampling frequency	35 MHz	DNE frequency bands	N/A
Loop	21	DNE stopband atten.	N/A
Low-frequency balance	40 dB	Differential delay	no
RFI	Threats 1 & 3	# of adaptive filter taps	150
Impulse Noise	no	Filter leakage factor $2\mu\alpha$	0
Leakage at hybrid	yes	β	$0.002 < \beta < 0.2$
AWGN	−140 dBm/Hz	Bits of precision	14, 15
LMS algorithm	normal	Flooring or rounding	flooring

This test was aborted early because the finite-precision errors were found to compound very quickly, causing overflow of the registers containing the filter coefficients. The number of samples processed until the first overflow is shown in Figure 6-14 for RFI threats 1 and 3, and using 15- and 14-bit arithmetic. With 14-bit math, numerical instability occurred after only 25,000 - 35,000 samples. As could be expected, 15-bit arithmetic almost doubled the operational lifetime of the filter, but the overflow times were still unacceptably low.

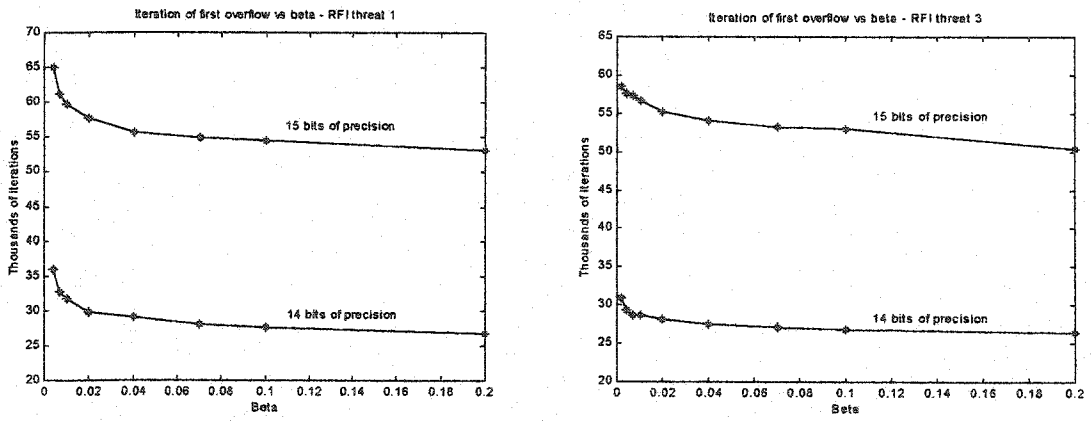

Figure 6-14: Number of iterations until storage overflow - loop 21 with truncation

From equation (4.35), it is known that the numerical stability of the system depends primarily on the inverse of the common-mode input's autocorrelation matrix $\mathbf{R}_X^{-1}$, and the expected value of the finite-precision errors $E\{\mathbf{b}_j\}$. With truncation, the last of these factors was clearly not zero. The situation was further exacerbated by the condition of $\mathbf{R}_X$. This matrix was estimated, assuming that the common-mode input sequence $\mathbf{X}_j$ was ergodic, by averaging a large number of instantaneous values of $\mathbf{R}_X$ (ie., equation (4.5) without the expectation operator). The

maximum eigenvalue of $\mathbf{R}_X^{-1}$ typically fell between 50 and 60 million for RFI threat 1. The cumulative effect of the non-zero $E\{\mathbf{b}_j\}$ and the ill-conditioned $\mathbf{R}_X$ was an inherently unstable system, and thus truncation will not be further discussed until the capabilities of leaky LMS are tested in Case 5.

6.5 Case 3: Use of the digital noise estimator (DNE)

Table 6-9: Case 3 Summary			
# of QAM symbols	50,000	Digital noise estimator	yes
Sampling frequency	35 MHz	DNE frequency bands	32, 64
Loop	12	DNE stopband atten.	60 dB
Low-frequency balance	30 dB	Differential delay	no
RFI	Threats 1 & 3	# of adaptive filter taps	150
Impulse Noise	no	Filter leakage factor $2\mu\alpha$	0
Leakage at hybrid	yes	β	$\beta = 0.002$
AWGN	−140 dBm/Hz	Bits of precision	∞
LMS algorithm	normal	Flooring or rounding	N/A

One of the major concerns expressed in [12] was that a portion of the information signal would corrupt the noise reference, causing a reduction in signal energy at the system's output. The purpose of the digital noise estimator is to remove as much of the wideband information signal as possible from the reference noise. To provide conditions under which the DNE was most likely to operate successfully, the balance of the loop below 1.6 MHz was reduced to only 30 dB, falling at 12 dB/decade beyond that. Loop 12 was selected for the tests because the transmitted signal would only travel 250 metres in differential mode before the point of imbalance. Since the common-mode attenuation over the remaining 250 metres of the loop was low, a large component of the information signal was clearly visible in the reference noise (Figure 6-15). After processing by the DNE, much of the underlying QAM spectrum was notched out (also shown in Figure 6-15).

In a relatively benign noise environment (AM threat 1 coupling to the loop 250 metres from the receiver), the SIR of the noise canceller's desired signal was 11.8 dB, and the peak power density levels of the QAM interference were 50 dB below the narrowband AM noise peaks. Clearly, if the loop is better balanced than 30 dB, or if stronger RFI signals couple to it, the QAM interference will become negligible. (In most of the tests conducted in Case 1, the signal

component in the reference noise was so small that it was completely eclipsed by the quantization noise of the ADC.)

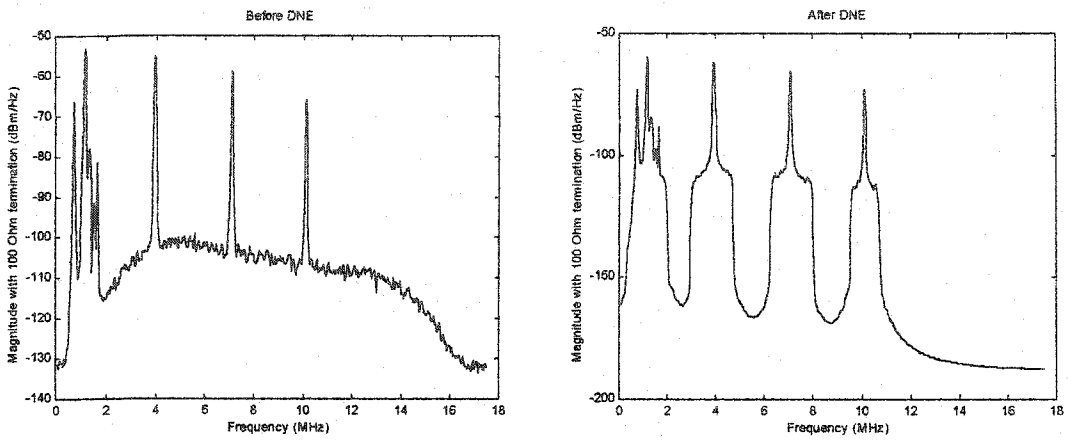

Figure 6-15: PSD of reference noise before and after 64-band DNE

A second test was carried out using the same loop with broadcast AM threat 3 and no amateur radio interference. In this case, the information signal constituted a larger component of the reference noise. However, the impact of the AM interference on the differential signal was very small; the SIR at the canceller input was 20.9 dB. Under these circumstances, the VDSL receiver could easily have operated without any noise canceller at all.

Using floating-point math, it was found that adding the DNE to the noise canceller was of no benefit; see Table 6-10, which gives the averaged results of three simulation runs for each of the two noise threats. The noise reduction figures with and without the DNE were very similar, and the noise floor remained within 1 dB of the thermal noise floor in all cases. These numbers do not justify the extra complexity and delay incurred by the use of the DNE. Furthermore, if fixed-point math had been used, the numerical inaccuracies introduced by the DNE would certainly have led to a performance degradation instead of the slight improvement described here. In view of these results, it is possible to conclude that there is no need to "clean" the reference noise.

Table 6-10: System Performance with Digital Noise Estimator

RFI threat	Noise Reduction with:		
	no DNE	64-band DNE	32-band DNE
threat 1, 3 amateur interferers	8.455 dB	8.456 dB	8.455 dB
threat 3, no amateur interferers	1.114 dB	1.116 dB	1.121 dB

6.6 Case 4: Performance in low or moderate noise environments

Table 6-11: Case 4 Summary			
# of QAM symbols	50,000	Digital noise estimator	no
Sampling frequency	35 MHz	DNE frequency bands	N/A
Loops	11, 12 (500 m)	DNE stopband atten.	N/A
Low-frequency balance	60 dB	Differential delay	no
RFI	Threats 2 and 3	# of adaptive filter taps	150
Impulse Noise	no	Filter leakage factor $2\mu\alpha$	0
Leakage at hybrid	yes	β	$0.0001 < \beta < 0.2$
AWGN	−140 dBm/Hz	Bits of precision	∞, 13 - 15
LMS algorithm	normal	Flooring or rounding	rounding

As the SIR of the differential input climbs, there will be little noise to cancel, and at some point the canceller's only contribution to the system will be the introduction of additional finite-precision and gradient noise. Accordingly, the canceller's control algorithm should include the ability to pass the received differential signal directly to the receiver. The exact conditions under which this bypass should occur will depend heavily on the performance of the VDSL receiver. In a QAM system, the canceller should be bypassed if the differential SIR is high enough to ensure that the FFE/DFE combination can operate without feeding back many decision errors. In a DMT system, the canceller bypass will depend on the windowing filter used, and on the performance of the adjacent-band noise cancellation method employed.

This quick experiment was run to check the canceller's performance in low or moderate interference environments. Two 500-metre loops were used to ensure that the signal energy at the receiver was relatively high (around −20 dBm), and RFI threats 2 and 3 were selected to keep the interference power low. The SIRs achieved for the four different test cases are shown in Table 6-12; two simulation runs were conducted for each scenario.

Table 6-12: Differential input SIRs for case 4

	RFI threat 2	RFI threat 3
Loop 11	8.9 dB	25.2 dB
Loop 12	16.6 dB	32.6 dB

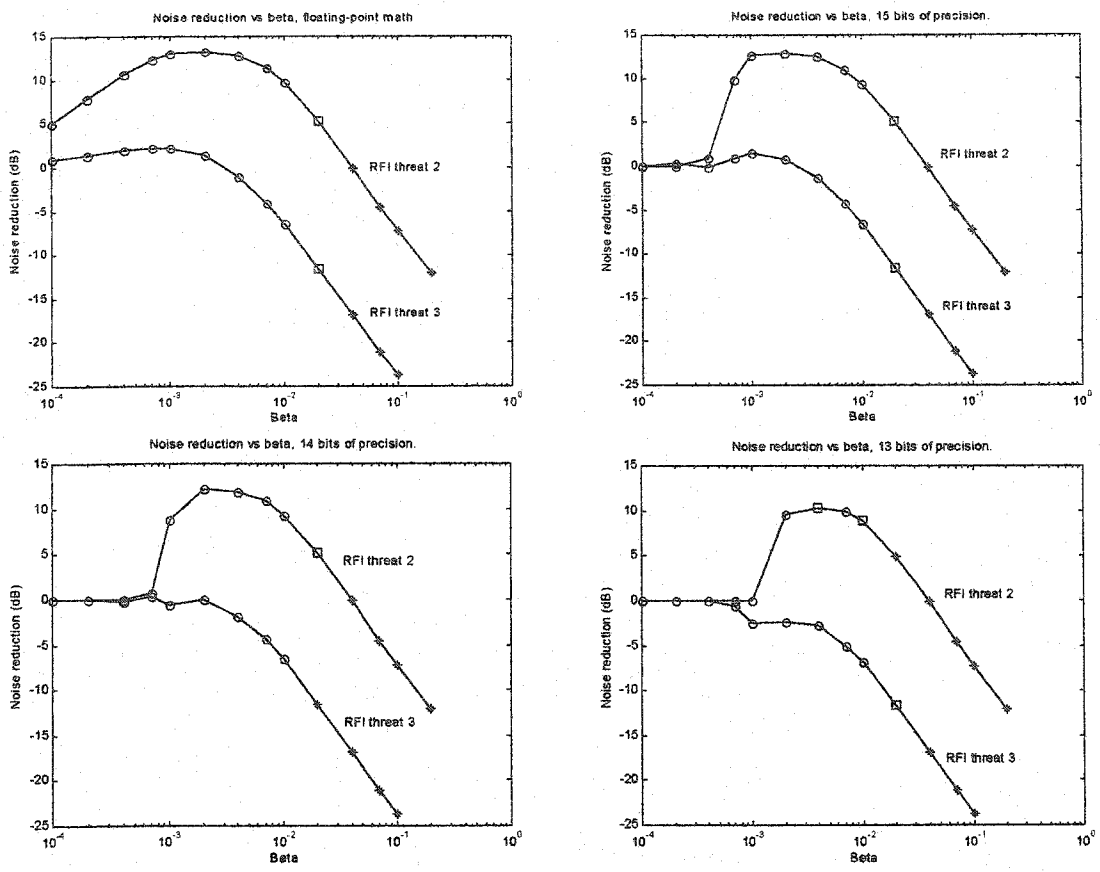

Figure 6-16: Noise reduction vs. β, loop 12, case 4

Figure 6-16 shows a plot of the noise reduction achieved for the two loop 12 cases. Even with IEEE floating-point arithmetic, it is clear that the canceller is of minimal benefit for the threat 3 scenario (SIR 32.6 dB). Using 14- or 13-bit fixed-point math, the unit consistently degrades the quality of the communications link. For the threat 2 case (SIR 16.6 dB), the floating-point performance is good: β values up to 0.01 may be used before the gradient noise starts to rise seriously. However, in finite-precision math, the noise reduction curve is gradually squeezed by the inability of the filter to adapt itself at low β values. Faced with the narrow window of β that achieves positive cancellation, and with the fact that the DFE will perform very well when its input has an SIR of 16.6 dB, it makes sense to bypass the canceller in this situation as well.

The loop 11 results with floating-point and 14-bit fixed-point math are shown in Figure 6-17. It can be seen that both curves exhibit the same general shape as the upper curve in Figure 6-16. Once again, the same level of reliable noise suppression could be expected from a DFE.

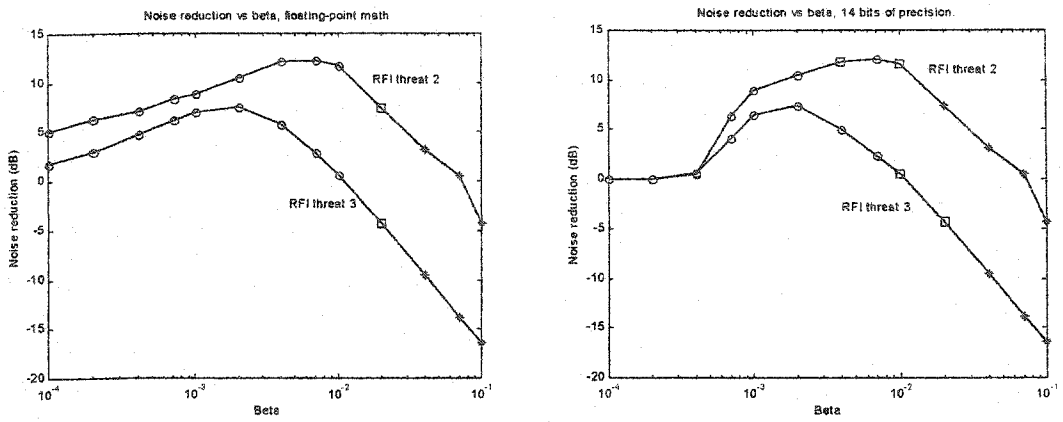

Figure 6-17: Noise reduction vs. β, loop 11, case 4

In the course of this research, it was found whenever the loop had a differential SIR greater than 0 dB, the canceller degraded the communications link for some range of β under 0.2; loops 12 and 25 from case 1 are other examples of such performance. This suggests that a good rule of thumb might be to bypass the canceller for SIRs higher than 0 dB, and to add a safety margin by allowing a maximum β of 0.1. Note that this dependence on β assumes that the AGC rules will be the same as those used in the simulations; if not, an additional correction factor will be required. While this rule of thumb has not yet been adequately tested, it would be a good starting point for someone attempting to design a complete VDSL receiver.

6.7 Case 5: Endurance tests with leaky LMS

Table 6-13: Case 5 Summary			
# of QAM symbols	200,000 / indefinite	Digital noise estimator	no
Sampling frequency	35 MHz	DNE frequency bands	N/A
Loops	21, 22	DNE stopband atten.	N/A
Low-frequency balance	40 dB	Differential delay	no
RFI	Threat 1	# of adaptive filter taps	150
Impulse Noise	no	Filter leakage factor $2\mu\alpha$	$0 \le 2\mu\alpha \le 0.002$
Leakage at hybrid	no	β	$\beta = 0.04$ (loop 21), $\beta = 0.01$ (loop 22)
AWGN	−140 dBm/Hz	Bits of precision	14, 15
LMS algorithm	normal / leaky	Flooring or rounding	both

A wide variety of tests were performed in an attempt to force a canceller which used 14-bit math and rounding into numerical instability. Interference signals ranging from RFI threat 1 to an

unmodulated −20 dBm sinusoid were used, ensuring that $\mathbf{R}_x$ was very poorly conditioned. However, none of these tests was able to provoke numerical instability within 30 million samples. In all cases, the maximum tap weight settled at a value far from 1. Some significant adjustments of this tap weight did occur, probably in response to the bursty activity of the amateur AM interferers, but it never grew above a certain maximum value. One reason for this prolonged stability was the assumption of an ideal ADC. A real ADC may cause quantization errors with non-zero mean.

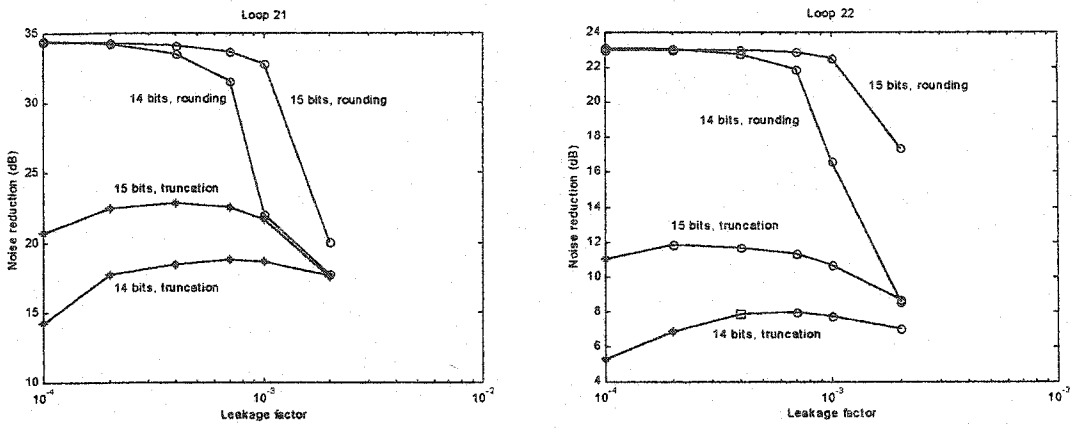

Figure 6-18: Noise reduction vs. leakage factor $2\mu\alpha$, case 5

To provide some insurance against overflow, whatever its eventual cause might be, it is sensible to use the leaky LMS algorithm. An experiment was therefore run to examine whether this algorithm could stabilize a canceller which used truncation, and to see if tap leakage would badly degrade the performance of a canceller which used rounding. Loops 21 (differential SIR −20.4 dB) and 22 (SIR −5.1 dB) were selected for testing with RFI threat 1. The noise reductions obtained using 15- and 14-bit arithmetic are plotted in Figure 6-18.

The first point to note about these results is that the noise reduction achieved for loop 21 was about 3 dB higher than the ensemble average calculated earlier. This difference occurred because a much greater number of symbols was used in this test, reducing the importance of the initial convergence errors. It is also obvious that leaky LMS makes the use of truncation feasible, but not desirable. In all tests, the noise reduction achieved with truncation was over 10 dB worse than that accomplished with rounding. Furthermore, when truncation was used, a small spike often appeared in the canceller output's PSD at $\omega = 0$; this condition was particularly pronounced on loop 21, where the spike rose 20 dB above the noise floor for low leakage factors. The same symptom was witnessed in the outputs of the case 2 tests, suggesting that the filter taps were again

developing an undesired DC bias and tending towards overflow. The only solution to this problem is more leakage, which implies poorer noise reduction and more noise at the system's output.

In contrast, the use of the leaky LMS algorithm performed well with rounding, so long as the leakage factor was 0.0004 or lower. Above that, the limiting factor in the system became the leakage noise, and the system's performance fell to the levels achieved with truncation. Considering that a leakage factor of 0.0002 was sufficient to stabilize the algorithm with truncation to at least 1.4 million data samples, and that it caused no appreciable decline in the achievable noise reduction with rounding, it should be sufficient to prevent the accumulation of rounding errors over many millions of iterations.

Three endurance tests were conducted on loop 21 to check the long-term performance of the leaky LMS algorithm with rounding. The first used RFI threat 1 as its interference signal, the second used a continuous –20 dBm 6 MHz sinusoid, and the last used a single bursty –30 dBm amateur interferer which cycled on and off every 10 milliseconds. For all three cases, the canceller showed no signs of instability. The tests were terminated at 30, 60 and 25 million iterations respectively, having achieved noise reductions of 35.2, 44.0 and 29.6 dB.

6.8 Case 6: Number of adaptive filter taps required

Table 6-14: Case 6 Summary			
# of QAM symbols	25,000	Digital noise estimator	no
Sampling frequency	35 MHz	DNE frequency bands	N/A
Loops	21, 22, 25	DNE stopband atten.	N/A
Low-frequency balance	40 dB	differential delay	yes
RFI	Threat 1 / 1 amateur sig	# of adaptive filter taps	5 - 500
Impulse Noise	no	Filter leakage factor $2\mu\alpha$	0.0002
Leakage at hybrid	yes	β	best β found in case 1
AWGN	–140 dBm/Hz	Bits of precision	14
LMS algorithm	leaky	Flooring or rounding	rounding

Because the VDSL bandwidth is so large, this system has been assumed to operate at 35 MHz. Processing 14-bit data through a large filter at this speed will require a specially-designed circuit which consumes considerable power. Accordingly, it is desirable to use as few filter taps as possible. This experiment tested the noise canceller with 5 to 500 taps, and on three different

loops. The loops were all 1 km long, and the RFI coupled to them at distances of 0, 250 and 500 metres from the receiver. Two different interference signals were used. The first was RFI threat 1, and the second was a single –30 dBm interferer which burst on and off every 0.5 ms and switched frequency from 3.9 to 7.9 MHz halfway through the simulation. This quick bursty behaviour is not representative of a real voice signal, but the frequent on/off activity in concert with the divergence of the differential and common-mode signals was intended as an additional challenge for the filters.

The filters were also implemented in two configurations. Consider loop 25, which has RFI couple to it 500 metres from the receiver. The common-mode and differential wavefronts will diverge at about $0.31c$ over the time it takes the differential signal to travel these 500 metres. When the differential signal arrives at the receiver, the common-mode signal will be about 2.33 μs, or 82 samples, out of alignment. Since the differential signal arrives first, a conventional adaptive filter will have to estimate the differential RFI component using a common-mode reference that is at least 82 samples "stale". Given the high correlation of the common-mode input, this is usually possible. However, the problem suggests a second filter configuration, in which the differential input is delayed long enough to allow the common-mode signal to reach the adaptive filter. For simplicity, in these simulations the differential signal was buffered by a number of samples equal to the number of adaptive filter tap weights.

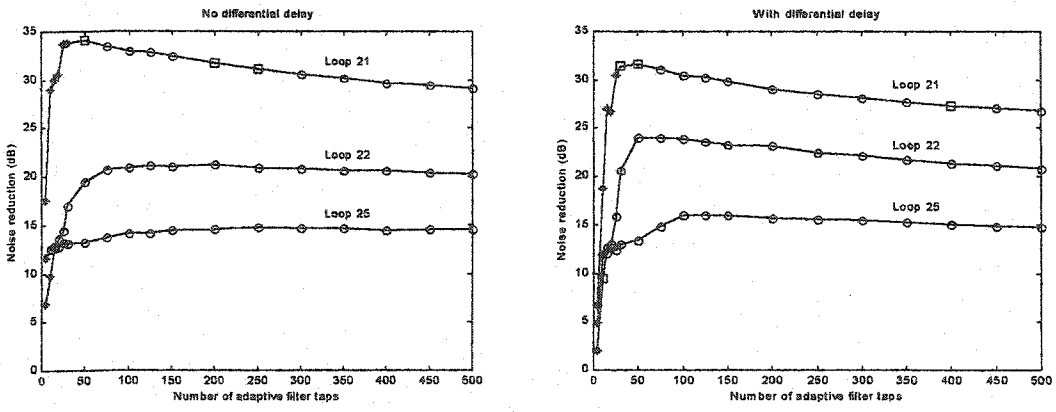

Figure 6-19: Noise reduction vs. number of filter taps, RFI threat 1

The test results for RFI threat 1 are shown in Figure 6-19. The noise cancellation achieved on loops 22 and 25 (with coupling distances of 250 and 500 metres) was between 1 and 3 dB better when the differential signal was delayed. Loop 21, on the other hand, suffered a performance

degradation of about 2 dB when the delay was employed; since the two signals arrived at the receiver in synchrony on this line, the second configuration effectively forced the filter to conduct cancellation using *future* samples of the reference noise. It will likely be possible to find a happy medium between the two filter configurations tested. Since the RFI will couple to the line close to the home in most cases, it is probable that only a short differential buffer will be required. The precise length of the buffer will depend on the RFI coupling characteristics of the average loop, and will have to be selected based on hardware tests using typical customer DSLs.

As expected, the best filter performance with differential delay was achieved when the number of filter taps slightly exceeded the divergence between modes (41 samples on loop 22, and 82 samples on loop 25). Beyond this peak, the performance curves gradually fell as a function of the number of filter coefficients. Finite-precision effects appear to be the reason for this behaviour: as the number of taps increased, more multiplications were required, generating larger errors at the filter output. For the wideband noise threat used in this test, the best combination of noise reduction and noise floor preservation was achieved with 75 to 100 taps. This number of taps ensures that the filter can handle signals which have coupled to the line anywhere in the 500 metres closest to the receiver, and gives it enough frequency resolution to avoid noisy performance at low SIRs (such as that displayed by loop 21 for low numbers of taps in Figure 6-19).

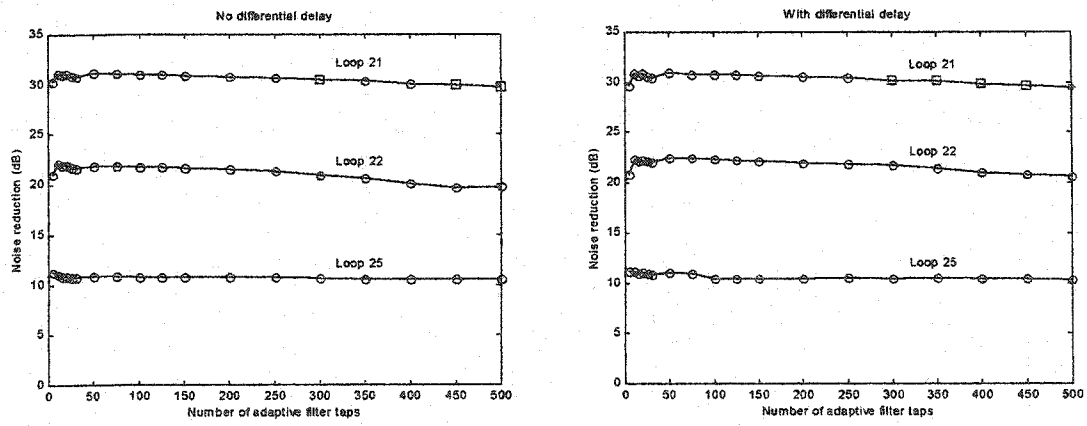

Figure 6-20: Noise reduction vs. number of filter taps, single amateur interferer

The suppression achieved against the single amateur interferer is shown in Figure 6-20. It was expected that the differential delay would be especially beneficial to loops 22 and 25; without it, a differential RFI component would occasionally appear while the only reference noise was

AWGN. In fact, the differential delay gave only a very slight performance improvement for loops 22 and 25, and a very slight degradation on loop 21 (three simulations of each loop were conducted, with three different information and interference signals). While the divergence between signals does cause some extra error, it is not significant compared to the error allowed by the filter in the longer periods when it does have a workable reference signal. It is also worth noting that the gradual negative slope of the performance curve is not as pronounced in the narrowband case. This is most likely due to the long periods during which there was no RFI, and for which the adaptive filter would have adjusted its output to almost zero.

6.9 Case 7: LMS with adaptive gain

Table 6-15: Case 7 Summary			
# of QAM symbols	50,000 / 25,000	Digital noise estimator	no
Sampling frequency	35 MHz	DNE frequency bands	N/A
Loops	11 - 13, 21 - 25	DNE stopband atten.	N/A
Low-frequency balance	40 dB	Differential delay	no
RFI	Threat 1	# of adaptive filter taps	50
Impulse Noise	no	Filter leakage factor $2\mu\alpha$	0
Leakage at hybrid	yes	β	$1 \times 10^{-6} < \beta < 0.31$
AWGN	−140 dBm/Hz	Bits of precision	∞, 13 - 15
LMS algorithm	LMS with adaptive gain	Flooring or rounding	rounding

One of the major challenges in noise cancellation will be to design a controller which enables the device to achieve optimum performance. So far, the performance of the canceller has been related to the SIR at the differential input. Unfortunately, approximating this SIR would require complex circuitry. Other potential methods to control μ, such as calculating the cross-correlation between the canceller output and the common-mode input and attempting to force it to zero, are also costly.

The LMS algorithm with adaptive gain offers some hope that a simple method exists to steer μ towards its optimum value, even in non-stationary environments. Unfortunately, it is only suitable for use with a low number of tap weights M, as equation (4.49) contains an $M \times M$ matrix. This test makes the optimistic assumption that up to 50 filter taps can be used, and evaluates the algorithm's performance against RFI threat 1. A maximum μ of $\mu_+ = 0.5/$(tap-input power) was allowed, along with a minimum of $\mu_- = 0.0001/$(tap-input power).

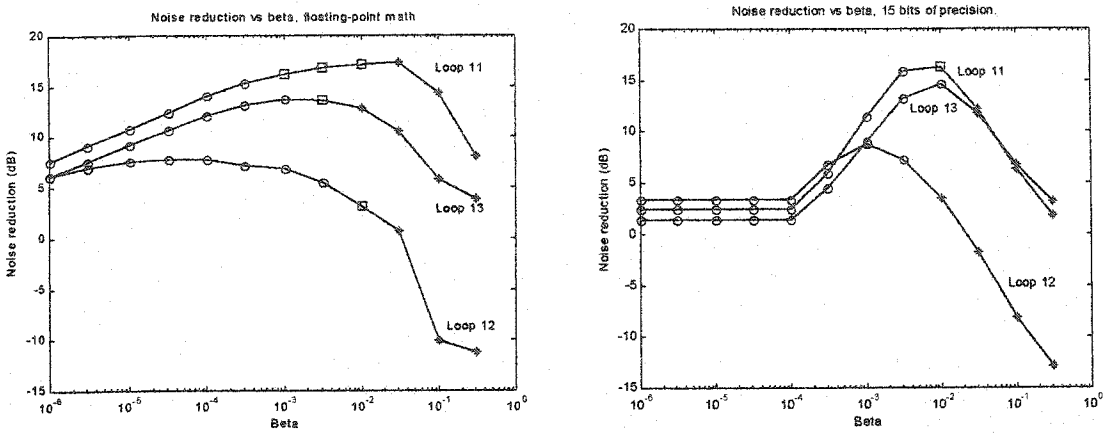

Figure 6-21: Noise reduction vs. β, 500-metre loops, case 7

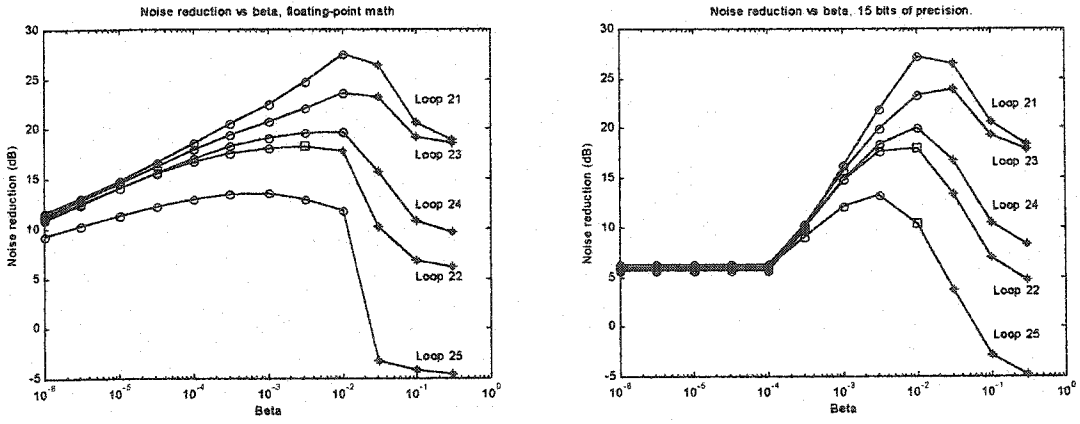

Figure 6-22: Noise reduction vs. β, 1-km loops, case 7

The noise reduction achieved on the 500-metre test loops is plotted versus β in Figure 6-21, and similar results for the 1-km loops are shown in Figure 6-22. Even with floating-point math, this algorithm was unable to match the noise reduction levels attained with the standard LMS algorithm. Also, as mentioned in [41], the transition between the method's best performance and the onset of unacceptably high gradient noise is very sudden; to ensure satisfactory operation of the system, β would have to be backed away from its optimum value. Last, the algorithm is ill-suited for implementation in fixed-point math, as it requires very small β values. Using 15-bit arithmetic, the lowest β which can be represented is about 3×10^{-5}. The plots given in Figures 6-21 and 6-22 show that the system operates almost exclusively at μ_- until β climbs past 0.0001. For all of these reasons, the LMS algorithm with adaptive gain is unsuitable for digital noise cancellation in VDSL.

6.10 Case 8: A note on canceller performance in impulse noise

Table 6-16: Case 8 Summary			
# of QAM symbols	200,000 / run	Digital noise estimator	no
Sampling frequency	35 MHz	DNE frequency bands	N/A
Loops	22, 25	DNE stopband atten.	N/A
Low-frequency balance	40 dB	Differential delay	yes
RFI	Threat 1	# of adaptive filter taps	25 - 200
Impulse Noise	yes	Filter leakage factor $2\mu\alpha$	0.0002
Leakage at hybrid trans	no	β	$\beta = 0.01$ (loop 22), $\beta = 0.004$ (loop 25)
AWGN	−140 dBm/Hz	Bits of precision	14
LMS algorithm	leaky	Flooring or rounding	rounding

Since impulse noise tends to have most of its energy at low frequencies, analog front-end FDD filters can reduce most impulses to levels that will not cause ADC saturation. If these small impulses arrive at the two canceller inputs at the same time, the canceller may be able to partially suppress them. If they are not time-aligned and the adaptive filter does not buffer the differential signal, or does not have a sufficient number of taps, it is theoretically possible that the noise canceller might create a second copy of the impulse at its output. This introduction of extra errors - potentially very large ones - is obviously undesirable.

This experiment was run to test how the canceller would react to impulse noise. Specifically, it is important to know if the canceller will propagate extra impulsive errors into the received differential signal, and if not, whether partial suppression of the impulses can be achieved. Cook pulses were transmitted in common mode every 10,000 samples, and loops 22 and 25 were used (with transfer points to differential mode at 250 and 500 metres from the receiver respectively). Five sequences of 200,000 symbols (7×10^6 samples total) were simulated.

The noise reduction achieved with and without delay of the differential signal is shown in Figure 6-23. When there is no differential delay, the results are very similar to those presented in Figure 6-19. With a differential delay equal to the delay of the adaptive filter, the results are about 2 - 3 dB better than those in Figure 6-19. Time-domain plots showed that very good cancellation of small impulses (under 10 mV) was achieved when the canceller included a differential delay chain and had enough filter taps to compensate for the lag in the common-mode signal. Partial

cancellation of larger impulses was possible, although in some cases up to 90% of the pulse's peak voltage was still present at the canceller output.

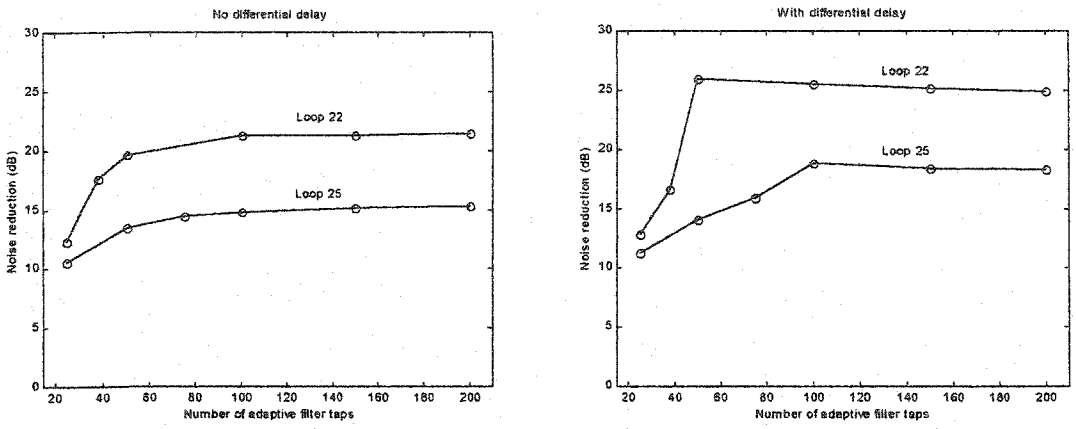

Figure 6-23: Noise reduction vs. number of adaptive filter taps, case 8

Figure 6-24 shows an example of such partial cancellation. The top left diagram shows 200 samples of the benchmark signal, which represented the best output that could be achieved by the canceller. The top right plot shows the received differential signal, which was badly corrupted by RFI and a high-frequency spike with 30 mV magnitude; the experiment was run on loop 22, and had a differential SIR of –4.5 dB. Note that the impulse had lost much of its expected shape because the front-end analog filters had already removed its low-frequency energy. After processing this input through an 80-tap noise canceller, the spike was reduced by about one-third. Furthermore, after the impulse had occurred the output very quickly returned to a shape similar to the benchmark signal.

When the filter had an insufficient differential delay, the impulse survived the canceller intact, and was sometimes augmented by a small amount. However, no duplicate impulses of similar amplitude propagated through the filter to the output. When the pulses were small and the filter was short, the canceller returned to peak performance very quickly (often in only 50 - 100 samples). Large impulses could cause the canceller output to deviate from the benchmark signal for 500 or more samples, and often caused high-frequency distortion in the output signal's PSD. The noise reduction figures shown in Figure 6-23 do not accurately reflect these disruptions, which would probably cause a significant number of symbol errors in an uncoded system.

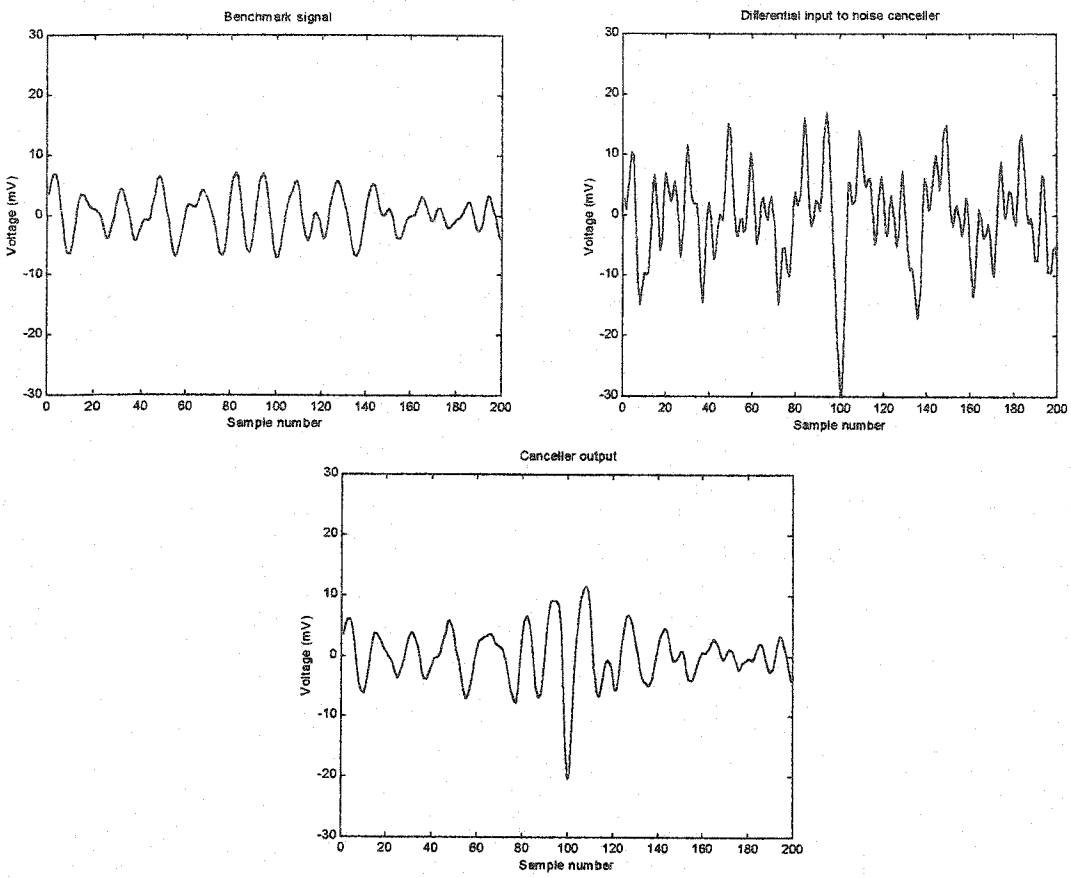

Figure 6-24: Time-domain plots of impulse cancellation

If the impulse events transfer to differential mode at some location different than the RFI coupling point, the adaptive filter will not be able to suppress them. In the short time it takes an impulse to reach its maximum amplitude, the filter will be unable to train itself to compensate for another mode-to-mode transfer function. Fortunately, it is likely that at least some of the impulse energy will switch modes at the same place as the RFI. Furthermore, impulses which make the switch further down the cable will be seriously attenuated by the differential channel, reducing their impact at the receiver.

While this simulation was a bit simplistic, it is reasonable to conclude that a canceller with a differential delay chain has a realistic chance of suppressing impulse noise which transfers between modes at the same point as the RFI, and close to the receiver.

6.11 Final notes

While most of the results in this study are qualitative rather than quantitative because of the lack of ensemble averaging, digital cancellation for VDSL should be feasible. The two critical problems which remain to be investigated are the very fast speed at which the canceller must operate, and possible long-term stability issues caused by physical imperfections in hardware components. Also to be determined is a good low-complexity control unit for the canceller.

The final chapter of this document is dedicated to summarizing the findings of the experiments conducted in this chapter, discussing potential problems, and outlining some plans for future research work.

7. Conclusions and Comments on Follow-up Research

7.1 Summary of Results

Wideband digital noise cancellation shows considerable promise as a solution to the interference problems experienced by xDSL. Simulation results show that the interference noise power can be reduced by 30 - 40 dB in bad noise environments. This document has also identified some other capabilities and computational costs associated with digital cancellation, in particular:

a. the ill-conditioned nature of the common-mode input's autocorrelation matrix makes the canceller sensitive to finite-precision effects. This implies that fixed-point math circuits which simply truncate the double-precision numbers created by multiplication operations cannot be used;

b. reliable canceller operation is possible when the double-precision number is rounded rather than truncated after each multiplication. The finite-precision errors generated by rounding are much more likely to have zero mean, and can therefore counteract the high correlation between the common-mode input's samples;

c. ADCs with at least 12-bit resolution are required for the two canceller inputs if digital noise cancellation is to prove feasible for VDSL;

d. 14- or 15-bit arithmetic will be necessary for the internal operations of the adaptive filter. If fewer bits are used, finite-precision effects are likely to reach intolerable proportions;

e. low values for the adaptive filter's step size μ (< 0.1/tap-input power) are required to keep gradient noise levels low and reduce the risk of instability;

f. the values of μ leading to the best interference reduction can be approximately related to the SIR at the canceller's differential input by a simple equation when the probability distribution of the common-mode input amplitude is Gaussian;

g. a small amount of filter tap leakage will be necessary to ensure the stable long-term performance of the canceller. Further stability testing under less ideal conditions is required;

h. preliminary studies show that the canceller should be bypassed completely when the differential SIR is 0 dB or higher;

i. to provide good cancellation of wideband interference, the adaptive filter should have at least 50 taps, and preferably 75 to 100 taps;

j. there is no requirement for a digital noise estimator or any other device which removes information signal power from the noise reference. The DNE introduces a tremendous cost in hardware complexity for a negligible improvement in cancellation performance;

k. the LMS algorithm with adaptive gain should not be considered for this application because of poor performance and excessive hardware cost; and

l. if impulse noise transfers from common mode to differential mode at the same impedance imbalance as the RFI, the canceller will probably be able to perform partial suppression of it.

7.2 *Issues to Resolve*

There are a number of outstanding issues to be resolved in follow-up research:

a. speed of operation. A digital canceller will have to complete about 100 multiplications every sample. At a VDSL sampling rate, this will require a special high-speed circuit. At an ADSL sampling rate, prototype testing can be carried out using modern DSP boards, and therefore the design of a ADSL digital canceller would be a prudent starting point. Two other benefits can be reaped from the use of an ADSL sampling frequency: first, ADCs with greater resolution can be used, guaranteeing that all interference will fall in the ADC's dynamic range without raising the noise floor, and second, no amateur radio interference falls in the ADSL band.

b. hardware testing. Hardware tests will allow experiments to be run at higher speeds, and with more realistic RFI coupling conditions. They will also allow long-term stability testing with actual hardware deficiencies. In particular, the chapter 6 simulations assumed an ADC whose quantization error had zero mean; real-world devices might not come close enough to this ideal property to ensure numerical stability. Also, the simulation of complex noise signals and loop coupling characteristics, particularly for impulse noise, involves many assumptions. It will be much easier to test the canceller's capabilities against impulse noise with an actual AWG 26 cable.

c. control. A suitable control algorithm for the canceller must be developed. In principle, the control functions should be as simple as possible, and should avoid PSD estimates or lengthy cross-correlations. An approximate relationship between the optimum value of μ and the differential SIR has been found, but it applies only when the common-mode signal has a Gaussian amplitude distribution, and the SIR will not be easy to estimate. It is likely that μ will have to be set by conducting occasional cross-correlations between the canceller output and the common-mode input. Further experiments must be conducted to find out how well this technique will work, and to simplify it so that it can be implemented in hardware without excessive cost.

d. interface with the narrowband analog canceller. Once an analog canceller design has been selected, it will be necessary to investigate a common control algorithm with the digital device. The two primary issues to consider are the development of decision criteria which prevent the analog canceller from frequently switching from one interferer to another, and a communication link from the board's differential AGC circuits to ensure that the canceller can adjust its noise estimator gain in synchronization with the AGC.

e. multiple or diffuse RFI coupling points. One very likely source of inaccuracy in the chapter 6 simulations is the assumption that all of the cable imbalance occurs at a single impedance fault. In actual fact, there are likely to be multiple impedance imbalances, or diffuse regions of imbalance along the line. The ability of the canceller to compensate for more complicated loop transfer functions needs to be tested.

f. length of the differential delay. Chapter 6 showed that the canceller provided up to 3 dB better noise reduction when the differential signal was delayed long enough to compensate for the lag of the common-mode signal. Finding the optimum value of this delay for an operational noise canceller will require hardware testing on a representative sample of customer loops. Alternately, it may be desirable to make the delay another adaptive factor, adjusted by the modem on a loop-by-loop basis.

7.3 Future Research Plans

Hardware and software experiments will continue through the next year. Simulation tests to be carried out include:

a. more runs on the chapter 6 tests, to develop better ensemble averages;
b. tests with multiple impedance imbalances on the loop; and
c. further tests to search for a simple quantity which can be evaluated to select the optimum value of μ.

For hardware testing, the construction of a prototype noise canceller is nearing completion; some corrective work on the analog front end and some programming remain to be done. Figure 7-1 shows a photo of the prototype. The board's design is centred on an Altera field-programmable gate array (FPGA) with 100,000 usable gates; this device will be used to implement the adaptive filter and control functions. All of the board's design parameters are reduced from the levels required for DSLs: the ADCs are only 10-bit, the sampling frequency that the FPGA can sustain will probably be under 2 MHz, and the filter will not have many taps because only one or two multipliers are likely to fit on the FPGA. These simplifications save costs, and reflect the realization that stellar performance cannot be expected from a single-layer board. However, the board should be sufficient to achieve some basic results under real RFI coupling conditions.

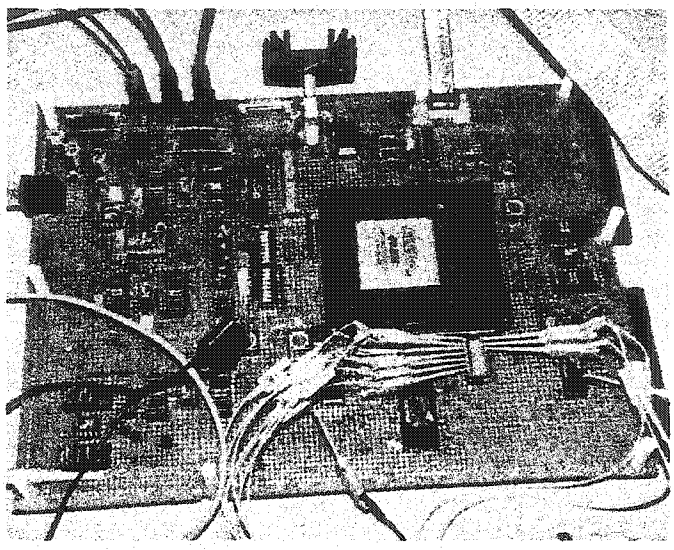

Figure 7-1: Prototype canceller

The main purpose of the prototype will be endurance testing for stability. Wideband RFI signals similar to those used in the simulations will be coupled to the cable using a wideband amplifier and antenna, and the wire pair will pass through a special board designed to introduce extra loop imbalance. Under different loop conditions, the RFI will be pulsed on and off in attempts to provoke numerical stability. Basic tests with impulse noise may also be conducted. Altogether, the lessons learned with the board should give some insight into the hardware problems which would have to be avoided in a follow-on ASIC design.

8. References

[1] Albert Benveniste, Michel Métivier, and Pierre Priouret, *Adaptive Algorithms and Stochastic Approximations*, Springer-Verlag, 1990.

[2] John Bingham, "A Comparison Between the Effects of RFI on Multicarrier and on Single Carrier", *ANSI Contribution T1E1.4/96-322 (Amati Communications)*, 11 November 1996.

[3] John Bingham, John Cioffi, and Mark Mallory, "Quantifying the Problem of RF Ingress in VDSL", *ANSI Contribution T1E1.4/96-082 (Amati Communications Corp.)*, 22 April 1996.

[4] John Bingham, John Cioffi, and Mark Mallory, "Digital RFI Cancellation with SDMT", *ANSI Contribution T1E1.4/96-083 (Amati Communications Corp.)*, 22 April 1996.

[5] British Telecom presentation at "Very-high-speed Digital Subscriber Lines (VDSL) Workshop" (leader: John Cioffi), *IEEE Global Communications Conference*, 18 November 1996.

[6] G. Cherubini, E. Eleftheriou and S. Oelcer, "CAP/QAM, DMT, and Filter-Bank Modulation Techniques for VDSL", *ANSI Contribution T1E1.4/99-329 (IBM)*, Ottawa, Ontario, 7 June 1999.

[7] Walter Chen and Alan Gatherer, "HFC Upstream Ingress Noise Cancellation Techniques", *IEEE 802.14/96-027 (Texas Instruments)*, Dallas, Texas, 18 January 1996.

[8] Wu-Jhy Chiu, Ming-Jung Wu, Wen-King Hwang, and Shyue-Win Wei, "Loop Survey in the Taiwan Area and Feasibility Study for HDSL", *IEEE Journal on Selected Areas in Communications*, vol. 9, no. 6, August 1991, pp. 801-809.

[9] John Cioffi, ed. "Very-high-speed Digital Subscriber Lines - System Requirements", *ANSI Contribution T1E1.4/98-043R8*, Plano, Texas, 30 November 1998.

[10] John Cioffi, "Recommendation of Digital Frequency-Division Duplexing for VDSL", *ANSI Contribution T1E1.4/99-274R2 (Stanford University)*, Ottawa, Ontario, 7 June 1999.

[11] John Cioffi, "Limited Precision Effects in Adaptive Filtering", *IEEE Transactions on Circuits and Systems*, vol. CAS-34, no. 7, July 1987.

[12] John Cioffi, Mark Mallory, and John Bingham, "Analog RF Cancellation with SMDT", *ANSI Contribution T1E1.4/96-084 (Amati Communications Corp.)*, 22 April 1996.

[13] John Cioffi, Vladimir Oksman, Jean-Jacques Werner, Thierry Pollet, Paul Spruyt, Jacky Chow, and Krista Jacobsen, "Very-High-Speed Digital Subscriber Lines", *IEEE Communications Magazine*, vol. 37, no. 4, April 1999, pp. 72 - 79.

[14] Richard A. Combellack, "Improving Range and Bandwidth of Telco Loop Plant", *Proceedings of the 47th International Wire and Cable Symposium*, Philadelphia, Pennsylvania, 16 - 19 November 1998.

[15] John W. Cook, "Wideband Impulsive Noise Survey of the Access Network", *BT Technology Journal*, vol. 11, no. 3, July 1993, pp. 155 - 162.

[16] John Cook, Rob Kirkby, Martin Booth, Kevin Foster, Don Clarke, and Gavin Young, "The Noise and Crosstalk Environment for ADSL and VDSL Systems", *IEEE Communications Magazine*, vol. 37, no. 5, May 1999.

[17] C. Conroy, B. Kim and J. Cioffi, "Some Implementation Issues Relevant to VDSL A/D Converters for Digital Duplexing", *ANSI Contribution T1E1.4/99-272R2 (VDSL Alliance)*, Ottawa, Ontario, 7 June 1999.

[18] I.K. Czajkowski, "Demographic Analysis of AM Broadcast RFI for a North American Scenario", *ANSI Contribution T1E1.4/97-083 (Nortel)*, Harlow, England, 3 February 1997.

[19] Stan Davies et al., "Australian Wideband Impulsive Noise Survey Data", *ANSI Contribution T1E1.4/95-125 (Telstra)*, Orlando, Florida, 13 November 1995.

[20] Richard Dorf, editor-in-chief, *The Electrical Engineering Handbook*, CRC Press, 1993.

[21] N. Drew and K. Harris, "AM Ingress on xDSL Loops", *ANSI Contribution T1E1.4/99-132 (Nortel Networks)*, Costa Mesa, California, 8 March 1999.

[22] Robert Fano, "A Theory of Impulse Noise in Telephone Networks", *IEEE Transactions on Communications*, vol. COM-25, no. 6, June 1977, pp. 577 - 588.

[23] Kevin Foster, "Practical Measurements of the Levels of Induced RFI from Amateur Radio Transmissions on Various Lengths of Dropwiring Commonly Found in BT's Access Network", *ANSI Contribution T1E1.4/95-097 (British Telecom)*, Silver Creek, Colorado, 21 August 1995.

[24] Kevin Foster and John Cook, "A Symbolic Pulse for Impulsive Noise Testing", *ANSI Contribution T1E1.4/92-143 (British Telecommunications Laboratories)*, 19 August 1992.

[25] Greg Gough, "Proposed Test Loops to be included in ADSL Forum WT-030v3", *ANSI Contribution T1E1.4/99-344 (Ameritech)*, Ottawa, Ontario, 7 June 1999.

[26] Nathan Grossner, *Transformers for Electronic Circuits*, McGraw-Hill Book Company, 1983.

[27] Simon Haykin, *Adaptive Filter Theory*, Prentice-Hall Inc., 1996.

[28] Dale Harman, Gi-Hong Im and Jean-Jacques Werner, "Steady-State Performance and Blind Training of a DFE in the Presence of RF Interference", *ANSI Contribution T1E1.4/97-104 (Lucent Technologies)*, Austin, Texas, 3 February 1997.

[29] Werner Henkel and Thomas Keßler, "A Wideband Impulsive Noise Survey in the German Telephone Network: Statistical Description and Modelling", *AEÜ*, vol. 48, no. 6, November/December 1994, pp 277-288.

[30] Werner Henkel, Thomas Keßler, Hong Chung, "Coded 64-CAP ADSL in an Impulse-Noise Environment – Modeling of Impulse Noise and First Simulation Results", *IEEE Journal on Selected Areas in Communications*, vol. 13, no. 9, December 1995, pp. 1611 – 1621.

[31] Mahbub Hoque, "Radio Frequency Aspects of VDSL", *ANSI Contribution T1E1.4/95-132 (Bellcore)*, Orlando, Florida, 13 November 1995.

[32] L. D. Humphrey, "Balance Measurements on BT Drop Wire 10", *ANSI Contribution T1E1.4/95-142 (Nortel)*, Orlando, Florida, 13 November 1995.

[33] L. D. Humphrey, "Capacity Modelling for HSAS Model A Noise", *ANSI Contribution T1E1.4/95-144 (Nortel)*, Orlando, Florida, 13 November 1995.

[34] David Jones, "Proposal for Specification of Broadcast and Mobile Ingress Sources in the ANSI VDSL Systems Requirements Document", *ANSI Contribution T1E1.4/99-312 (Broadcom Corporation)*, Ottawa, Ontario, 7 June 1999.

[35] David Jones, "Laboratory Measurements of a QAM-based VDSL Transceiver Being Used in Commercial Mass Market Service Deployments", *ANSI Contribution T1E1.4/99-313 (Broadcom Corporation)*, Ottawa, Ontario, 7 June 1999.

[36] David Jones, "Analog Front End Dynamic Range Requirements for Programmable Spectrum VDSL Systems", *ANSI Contribution T1E1.4/99-315 (Broadcom Corporation)*, Ottawa, Ontario, 7 June 1999.

[37] David Jones, "Summary Comparison of QAM and DMT for VDSL with Recommendation", *ANSI Contribution T1E1.4/99-343 (Broadcom Corporation)*, Ottawa, Ontario, 7 June 1999.

[38] Question and answer session with David Jones, of Broadcom Corporation, T1E1.4 Committee meeting, Ottawa, Ontario, 7 June 1999.

[39] David Jones, "Service Time-to-Market Implications of ADSL/VDSL Interoperability", *ANSI Contribution T1E1.4/99-316 (Broadcom Corporation)*, Ottawa, Ontario, 7 June 1999.

[40] Walt Kester, "High-Speed Sampling and High-Speed ADCs", *Analog Devices Technical Reference Book on High-Speed Design Techniques*, Prentice-Hall, 1996.

[41] Harold Kushner and Jichuan Yang, "Analysis of Adaptive Step-Size Stochastic Approximation Algorithms for Parameter Tracking", *IEEE Transactions on Automatic Control*, vol. 40, no. 8, August 1995, pp. 1403 - 1410.

[42] Charles Lawrence, Floyd McDougall, and Jay Lundschen, "Technique for Cancelling Common-Mode Switching Noise To Achieve Reduced Error Rates in a Local Area Network", *US Patent 5,555,277*, 10 September 1996.

[43] Joseph W. Lechleider, "Impulse Noise Cancellation for ADSL", *ANSI Contribution T1E1.4/92-154 (Bellcore)*, Portland, Oregon, 19 August 1992.

[45] Pierre Lefebvre and Tet Yeap, "Adaptive Multiple Sub-band Common-Mode RFI Suppression", *Canadian and US Patent (agent's docket number AP 549)*, 13 May 1998.

[46] Alberto Leon-Garcia, *Probability and Random Processes For Electrical Engineering*, Addison-Wesley Publishing Company, 1994.

[47] Leander Matsch, *Capacitors, Magnetic Circuits and Transformers*, Prentice-Hall Inc., 1964.

[48] Len Mays and Edmund P. Darbyshire, *European Patent Application EP 0 453 213 A2 (GEC Marconi)*, 15 April 1991.

[49] Office of Spectrum Management, National Telecommunications and Information Administration, "Tables of Frequency Allocations", www.ntia.doc.gov/osmhome/alloctbl/alloctbl.html.

[50] Vladimir Oksman, "Reducing ADC Resolution by Using Analog Band-pass Filters in FDD-based VDSL", *ANSI Contribution T1E1.4/99-334 (Lucent Technologies)*, Ottawa, Ontario, 7 June 1999.

[51] Vladimir Oksman, "On ADC and DAC Resolution Requirements for DMT/FDD-based VDSL", *ANSI Contribution T1E1.4/99-290 (Lucent Technologies)*, Ottawa, Ontario, 7 June 1999.

[52] Martin Pollakowski and Stephan Heuser, "On the Validity of Loop and Crosstalk Models for VDSL", *Globecom 96 Workshop on Very High Speed Digital Subscriber Lines*, London, England, 18 November 1996.

[53] Philip Potter and Bernard Smith, "Statistics of Impulsive Noise Crosstalk in Digital Line Systems on Multipair Cable", *IEEE Transactions on Communications*, vol. COM-33, no. 3, March 1985, pp.259-270.

[54] Dennis J. Rauschmayer, "Cross-talk Reduction using Common-Mode Signal Correlation", *ANSI Contribution T1E1.4/96-069 (Pulse Communications Inc)*, Herndon, Virginia, 22 April 1996.

[55] Jeffrey Rakos and Mahbub Hoque, "Ingress and Egress Model Analyses for Broadband Deployment", *ANSI Contribution T1E1.4/97-279 (Bellcore)*, 22 September 1997.

[56] Whitham Reeve, *Subscriber Loop Signalling and Transmission Handbook - Analog*, IEEE Press, 1992.

[57] Whitham Reeve, *Subscriber Loop Signalling and Transmission Handbook - Digital*, IEEE Press, 1995.

[58] Ilan Reuven and Ehud Rokach, "On Meeting RF Ingress Requirements in a QAM-based VDSL Transmission System", *ANSI Contribution T1E1.4/97-061 (Orckit Communications)*, 3 February 1997.

[59] Burton Saltzberg, "Comparison of Single-Carrier and Multitone Digital Modulation for DSL Applications, *IEEE Communications Magazine*, vol. 36, no. 11, November 1998.

[60] Nicholas Sands, "A Comparison of Multi-carrier and Single-carrier modulation on the basis of Performance and Complexity", *ANSI Contribution T1E1.4/99-297 (Texas Instruments)*, Ottawa, Ontario, 7 June 1999.

[61] Thomas Starr, John Cioffi and Peter Silverman, *Understanding Digital Subscriber Line Technology*, Prentice-Hall, 1999.

[62] Loren Swingle, "Impulse Noise Measurements in New York City and Their Possible Ramifications on Selection of DSL Line Code Standard", *ANSI Contribution T1D1.3/86-144 (NYNEX)*, 25 August 1986.

[63] Craig Valenti, "Survey Plan for Characterization of Impulse Noise and Background Noise at VDSL Frequencies", *ANSI Contribution T1E1.4/97-295 (Bellcore)*, Morristown, New Jersey, 22 September 1997.

[64] Craig Valenti, "Cable Cross-talk Parameters and Models", *ANSI Contribution T1E1.4/97-302 (Bellcore)*, Morristown, New Jersey, 22 September 1997.

[65] Robert Walden, "Performance Trends for Analog-to-Digital Converters", *IEEE Communications Magazine*, vol. 37, no. 2, February 1999, pp. 96 - 101.

[66] Robert Walden, "Analog-to-Digital Converter Survey and Analysis", *IEEE Journal on Selected Areas in Communications*, vol. 17, no. 4, April 1999, pp. 539 - 549.

[67] Jean-Jacques Werner, "The HDSL Environment", *IEEE Journal on Selected Areas in Communication*, vol. 9, no. 6, August 1991, pp. 785-800.

[68] Jean-Jacques Werner, "Impulse Noise in the Loop Plant", *Proceedings of the International Communications Conference*, 1990, pp 348.1.1 - 348.1.4.

[69] Jeffery Wepman, "Analog-to-Digital Converters and Their Applications in Radio Receivers", *IEEE Communications Magazine*, vol. 33, no. 5, May 1995, pp. 39 - 45.

[70] Bernard Widrow et al., "Adaptive Noise Cancellation: Principles and Applications", *Proceedings of the IEEE*, vol. 63, no 12, December 1975, pp. 1692-1716.

[71] Bernard Widrow et al., "Stationary and Nonstationary Learning Characteristics of the LMS Adaptive Filter", *Proceedings of the IEEE*, vol. 64, no. 8, August 1976, pp 1151 - 1162.

[72] Conversation with Dr. Tet Yeap, August 1999.