Frequency Shift Filtering for Cyclostationary Signals

This thesis is submitted in partial fulfilment of the requirements for the degree of Doctor of Philosophy (D.Phil.)

Jonathan F. Adlard
Communications Research Group
Department of Electronics
University of York

1st September 2000

Abstract

The frequency-shift (FRESH) filter is a structure which exploits the spectral correlation of cyclostationary signals for removing interference and noise from a wanted signal. As most digital communication signals are cyclostationary, FRESH filtering offers certain advantages for interference rejection in a communications receiver.

This thesis explores the operation and application of FRESH filters in practical interference scenarios. The theoretical background to cyclostationarity is clarified with graphical interpretations of what cyclostationarity is, and how a FRESH filter exploits it to remove interference. The effects of implementation in a sampled system are investigated, in filters which use baud rate related cyclostationarity, leading to efficiency improvements. The effects of varying the wanted signal pulse shape to enhance the cyclostationarity available to the FRESH filter are also investigated.

A consistent approach to the interpretation of the FRESH filter's operation is used throughout, while evaluating the performance in a wide range of realistic channel conditions.

VLF radio communication is proposed as one area where interference conditions are particularly suitable for the use of FRESH filtering. In cases of severe adjacent channel interference it is found that a FRESH filter can almost completely remove the interferer. The effects of its use with an impulse rejection technique are also investigated.

Finally, blind adaptation of FRESH filters through exploitation of carrier related cyclostationarity is investigated. It is found that one existing method loses the advantage of FRESH filtering over time invariant linear filtering. An improvement is proposed to the latter which restores its performance to that of a trained FRESH filter, and also reveals that carrier related cyclostationarity can be exploited, in some cases, by a simpler method.

Contents

List of Figures

Acknowledgements

I would like to thank the Defence Evaluation and Research Agency for their funding of parts of this work. In particular, I thank Stuart Bennet of DERA for the useful technical discussions. I also thank my supervisors Tim Tozer and Alister Burr for their help throughout this project, and especially to Tim for keeping my salary coming in for all these years, and for making the Communications Research Group happen.

I am also indebted to many other people who have helped on the way. In particular, Dan Verdin, Steve Baines, Alek Kolcz, Colin Brown and Yuriy Zakharov for the technical discussions, David Grace for the coffee and the computers, Brian for the computers, George for the football, Neil for the enthusiasm, Zelda for the dancing, Noelle for the English and the AUT for job security.

I would also like to thank Peter Brand of the University of Edinburgh for accidentally starting my interest in signal processing many years ago.

Declaration

Some of the research presented in this thesis has resulted in publications in conference proceedings [1, 2, 3, 4].

All work presented in this thesis as original is so, to the best knowledge of the author. References and acknowledgements to other researchers have been given as appropriate.

Glossary

ACI	Adjacent Channel Interference
AM	Amplitude Modulation
AWGN	Additive White Gaussian Noise
BER	Bit Error Rate
BPSK	Bi-polar (or Binary) Phase Shift Keying
BT	Bandwidth-Time product (for Gaussian filter)
CCI	Co-channel Interference
CDMA	Code Division Multiple Access
CMA	Constant Modulus Algorithm
CW	Continuous Wave
DFE	Decision Feedback Equaliser
EBW	Excess bandwidth
FIR	Finite Impulse Response
FSE	Fractionally Spaced Equaliser
FRESH	Frequency Shift (filter)
GSM	Global System for Mobile Communications
GMSK	Gaussian Minimum Phase Shift Keying
HOS	Higher Order Statistics
IID	Independent Identically Distributed
ISI	Intersymbol Interference
LCL	Linear Conjugate-Linear (filtering)
LMS	Least Mean Square
MF	Matched Filter
MSE	Mean Squared Error
MSK	Minimum Phase Shift Keying
NWMF	Noise Whitening Matched Filter
QAM	Quadrature Amplitude Modulation
QPSK	Quadrature Phase Shift Keying
RLS	Recursive Least Squares
SCD	Spectral correlation density (section 2.2.4)
SNR	Signal to Noise Ratio
SOI	Signal of Interest
SPW	Signal Processing Worksystem
SS	Spread Spectrum
SSE	Symbol Spaced Equaliser

VLF	Very Low Frequency radio
R_x^α	Cyclic autocorrelation function of x with frequency shift α (section 2.1.4)
$R_{xx^*}^\alpha$	Conjugate cyclic autocorrelation function of x with frequency shift α (section 2.1.4)

Chapter 1

Introduction

Contents

1.1 Background

This thesis examines a particular range of techniques which can be used in digital communication systems for interference rejection. In the last three decades there has been a dramatic growth in the use of digital communications. Initially this growth was centred in wire based systems, as public telephone networks began to carry digital data. Although there is still rapid growth in wired communications, the majority of technical challenges lies in developing wireless, or radio, systems, be they mobile or fixed, which can carry the ever-increasing load of data traffic.

Interference is "the intrusion of electrical disturbances which interfere with reception" [5] and "interference rejection" is the process of removing, or reducing, the effect of such disturbances to allow reliable reception of the wanted signal, or *signal of interest* (SOI). In this work, the source of the "electrical disturbances" is assumed to be man-made, and is usually another communications signal.

As radio bandwidths become more heavily used, interference becomes a more significant problem. As use of the radio spectrum is now very expensive for network operators (3rd generation mobile spectrum licenses have been sold recently in the UK for a total of £22 billion), so they understandably attempt to maximise the capacity of any network. By doing so, they may reduce the margins

of immunity to interference. In fact, the 3rd generation UMTS systems will use code-division multiple access (CDMA), which has a capacity limited by the interference by users of the systems to each other.

Similarly, as more communications satellites are put into service, the separation between them apparent to an earth station is reduced; thus narrower beamwidth antennas have to be used, or the interference of one satellite signal to its neighbours will increase.

Wired systems are not immune to this problem. Crosstalk is the term given to interference between signals carried on cables which are bundled together [6]. Crosstalk becomes more of a problem as higher data rates are sent through electrical cables. It is a significant factor limiting capacity in modern Digital Subscribed Line (DSL) systems where high speed data are sent over cables initially only intended to carry a 4 kHz bandwidth [7].

Jamming, or deliberate interference transmitted to prevent the reception of the jammed signal by its intended recipient, is a closely related problem in military radio communications. Some of the interference scenarios described in this thesis could also be interpreted as jamming scenarios, and anti-jamming techniques are usually very similar to interference rejection techniques.

Clearly, interference is a common and significant problem in modern communications systems. One can also say that the great majority of modern communications signals carry digitally modulated data; this means that they have statistics which are *cyclostationary*. This description, explained in greater detail later on in this thesis, is equivalent to saying that the signals have *spectral correlation*, which means that they possess an inherent form of frequency diversity. In other words, the information carried by the signals is repeated in the spectral domain in some way.

This property is what is exploited by the techniques investigated in this thesis. There are many different approaches to interference rejection, or mitigation. Methods are available of exploiting any differences between the signals, be they temporal, spectral, spatial or statistical in nature. The property exploited in the current work could be described as spectral or statistical. A review of the range of techniques available follows in section 1.3.

1.2 Scope of the Thesis

This work follows on from a series of papers, published mainly by one group at the University of California (Davis), on exploiting cyclostationarity (for example, [8, 9, 10, 11, 12, 13]). In these papers, and others cited later, the concept and application of cyclostationarity is presented, and in particular, the structure known as the *frequency-shift filter* (FRESH filter) is proposed for interference rejection.

The purpose of this thesis is to examine in more detail, and from a more practical point of view, the use of the FRESH filter for interference rejection in digital communications. It should be made clear at this point, that this is a single antenna technique - it does not use multiple sensors to exploit the spatial properties of signals. Here it is considered in the context of narrow band signals, although it can also be applied to spread spectrum (SS) signals.

In the literature, FRESH filtering appears to be a highly promising technique, offering significant improvement in performance over the simpler Wiener filtering. It is particularly attractive because

it uses the diversity often already present in the signal for other reasons. This means that it can often be used as an "add-on" to an existing system: it could improve performance to an existing communications system without requiring a redesign of the whole system. However, the existing work has often used unrealistic SOI and interference scenarios, and has focussed on the technique as a generalisation of Wiener filtering, and so measured performance in terms of mean squared error (MSE), rather than the ultimately more important bit error rate. Here, the aim is to use more likely signal scenarios, and to explore possible practical applications of the technique.

Also, the existing descriptions of the concepts and techniques are often complex and unnecessarily abstract. One aim of this thesis is to explain what are rather simple concepts and techniques in a way which is straightforward and more intuitive.

It will be seen later that explicit use of spectral correlation achieved by the FRESH filter is also done implicitly in some cases by a standard matched filter. This is highlighted to help explain how FRESH filtering relates to, and compares with, standard filtering techniques.

1.3 Interference Rejection Context

There are many different approaches to the general problem of interference rejection in digital communications. These are described in an excellent review of the subject by Laster and Reed [14]. The main points from this review are summarised in this section, and some additional work particularly related to FRESH filtering and other uses of cyclostationarity, is also described.

Broadly speaking, one can divide interference rejection techniques into three areas:

- methods which use antenna arrays to exploit different angles of arrival of the signals;
- spread spectrum (SS) techniques;
- narrow band techniques.

Of course, techniques in the second and third categories can also be applied to the outputs of antenna arrays, and there are many similarities between wide and narrow band single antenna techniques.

1.3.1 Spatial filtering

The main area of interest in using antenna arrays is in finding methods of training the array to have high gain towards the SOI and low gain in the direction of any interferers. This may be done "blindly" [15, 16] or with a training signal [17]. There are many different algorithms for this (reviewed in [18] and [19]) but those of relevance to the current work are the algorithms based on the cyclostationarity of the signals [12, 20], where the array is adapted to maximise the correlation of the filtered signal with a frequency shifted version of the received signal, and [21] where a non-linearity generates spectral lines representing the cyclostationarity of the input, and the strength of these are maximised. The spatial filtering is combined with additional temporal filtering using cyclostationary properties in [22, 23] for beamforming, and in [24] the combination of the spatial and temporal approach is used for direction finding of cyclostationary signals.

No more is said about the spatial approach to interference rejection in this thesis, other than to point out that the approach used in [12] was used for blind adaptation of FRESH filters (for a single antennal receiver) in [25]. There, frequency shifted versions of the input signal are filtered in such a way as to maximise their correlation. However it is shown in chapter 7 of this thesis that this method is flawed.

1.3.2 Interference rejection with spread spectrum signals

There has been a large amount of work done on interference rejection in SS systems. This work is surveyed in [14]. Here, again, some major points are noted, and work which exploits cyclostationarity is highlighted as particularly relevant to this thesis. Interference to a SS signal can be categorised into two areas: narrow band interference (relative to the SS SOI) assumed to be external to the system, and wideband, or multiple access, interference, assumed to be from other users of the system.

SS signals are inherently resistant to narrow band interference, essentially reducing the spectral density of the interferer by the spreading factor after de-spreading. However, there is a range of additional techniques for use when this de-spreading is not sufficient on its own. The only one of particular interest here is [26], in which the spectral redundancy of a Binary Phase Shift Keyed signal (BPSK) is used to reduce the corruption to the SOI caused by notch filtering against an interferer. This process, although it is not described as such, is equivalent to linear-conjugate-linear (LCL) filtering described by [27], which is the baseband equivalent of Gardner's FRESH filtering for carrier related cyclostationarity as described in chapter 3 and further discussed in chapter 7.

Wideband interference in direct sequence SS usually means interference from other users of the system, so it will be occupying the same bandwidth as the SOI, and have the same baud rate and chip rate. There are again many different approaches to mitigating its effect; these are either single user methods, where it is assumed only one user's spreading sequence is known at the receiver, or multi-user methods, where all users' sequences are known. Multi-user methods (see for example, [28]) attempt to detect each user and remove its effect from the remaining users, thus reducing the interference. Single user methods use more general properties of the signals, and so are closer to the narrow band scenarios considered in this thesis. An important factor in limiting multiple access interference is coping with the near-far problem, where nearby users may cause a higher power than others at the receiver, and so cause worse interference. Agee [29] has apparently used spectral correlation to aid in the management of this problem. [30] and [31] also exploit spectral correlation in direct sequence SS. In [32], Reed uses the spectral correlation of the interference to improve reception of frequency hopping spread spectrum.

1.3.3 Interference rejection with narrow band signals

Again, [14] reviews this area thoroughly. This is the category into which the current work falls. The use of spectral correlation for interference rejection can be thought of as an extension to the linear equalisation immediately below. Other techniques are then summarised before returning to specific papers on the use of spectral correlation.

Linear equalisation

A basic way of removing interference is to use *linear equalisation*, which can be either *symbol-spaced* or *fractionally-spaced* [33] depending on whether the delay between filter coefficients (in a transversal implementation) is equal to, or less than, one symbol period. The term "equalisation" reveals that the main motivation for such a process is to equalise a channel which introduces intersymbol interference (ISI). However, it also has a beneficial effect against additive interference. Qureshi gives a thorough description of adaptive equalisation in general in [34], and at a simpler level in [35]. The performance of linear equalisation against interference is examined in [36] and in [37] where, although cyclostationarity is not mentioned, the interferer has the same baud rate as the SOI, so baud rate sampling in the system exploits spectral correlation of both the interferer and SOI (see section 3.6.4). Linear equalisation of adjacent channel interference (ACI) is the subject of [38], where both the signal of interest (SOI) and the interferer are Gaussian filtered BPSK. The approach then taken is to look at the equaliser parameters (such as the number of taps) required to receive various signal and interference scenarios most efficiently. In particular, [38] is interested in the effect of varying the time/bandwidth product of the transmit filters for the SOI and interferer. Again, the cyclostationarity of the SOI and interferer is implicitly exploited by the baud rate sampling in the receiver structure. In the work that follows, the *fractionally-spaced equaliser* (FSE) performance is often used as a benchmark against which to measure the FRESH filter.

Decision feedback equalisation

The *decision feedback equaliser* (DFE) is a non-linear structure which can offer significant improvements over linear equalisers in certain channel conditions [39]. The reception of Quadrature Phase Shift Keying (QPSK) with a DFE in continuous wave (CW) (i.e. sinusoidal) interference is addressed in [40]. This paper includes analytical calculation of the ideal tap weights for a known interferer, an analysis of the effects of using the Least Mean Squares (LMS) adaptation algorithm, and an investigation of the effects of error propagation in the decision-feedback loop of the equaliser. The feed-forward part of the equaliser uses symbol spaced taps.

In [41] the DFE is compared with the FSE and the symbol-spaced equaliser (SSE) in channels with fading and interference (either in-band CW, or adjacent channel modulated interference). The DFE is shown to perform significantly better than the FSE, which in turn is better than the SSE. Other work examining the performance of DFE structures in the presence of interference is reported in [42], where ACI and co-channel interference (CCI) are added to a signal in a static frequency selective channel. The authors of [42] point out that using shorter pulses to widen bandwidth (and enhance cyclostationarity) reduces the span of the equaliser required for a given channel, while increasing the adjacent channel interference problem in the system overall.

Similar problems are addressed in [43], but with an exploration of the effect of exploiting different signal bandwidths. The authors of [44] look at crosstalk interference in a cable system, and examine the situation where there is either one dominant interferer (or equivalently several symbol synchronous interferers) or several asynchronous interferers. Effectively this is a comparison of cyclostationary and stationary interference. They demonstrate that cyclostationary interference can be more effectively removed than stationary interference when using a DFE. Again, the interference here has the same baud rate as the SOI. As the DFE contains symbol rate sampling, this

is showing how the baud rate cyclostationarity is exploited by the filter. Also of interest is [45], where the DFE is used to remove interference which, with other parameters constant, has varying degrees of cyclostationarity.

Non-linear methods

Other than the DFE, which is a non-linear method, there is a wide range of techniques which fall into this category, most of which can be categorised as using neural networks. These are reviewed in [46]. The main advantages over other classes of techniques include better performance in non-Gaussian interference (which can apply to impulsive noise, as in chapter 6, or to digital interference), better performance against non-linear distortion (which can be introduced by system hardware) and the potential use of new, perhaps superior, types of adaptive algorithm.

Blind adaptation against interference

A number of papers describing the use of the constant modulus algorithm (CMA) for interference rejection is described in [14], but these papers describe the use of the CMA in (blindly) adapting a filter (usually a linear equaliser) in an interference environment; the CMA is not an interference rejection algorithm as such. In this work, the focus is on the static performance of different filter structures, not the rates or methods of adaptations, with the exception of chapter 7. The only published works specifically examining the blind adaptation of FRESH filters are [11, 25, 47, 48, 49]. [25] proposes a blind algorithm for FRESH filters; this is combined with a blind equalisation algorithm in [47] to operate in a fading and interference environment. Chapter 7 demonstrates, however, that the algorithm of [25] does not work successfully. [48] is a modification of [25] which works, but with severe limitations. Again, this is discussed in chapter 7. [48] is essentially the same approach described as the *spectral correlation discriminator* in [11]. A completely different approach is shown in [49], where the CMA is applied to FRESH filters, and similarly in [50], where the stop-and-go Bussgang algorithm is used for the same purpose.

Cyclostationarity based methods

The majority of the work on cyclostationary, or spectral correlation, based techniques for interference rejection has been carried out by W. A. Gardner and his co-workers.

Gardner has generally concentrated more on the theoretical background to cyclostationarity, and its representation and exploitation, starting with his PhD thesis, [51] and [8]. The work on representation continued with [52] and [53]. The spectral correlation interpretation of cyclostationarity was presented in [9], although most of this material also appears in Gardner's book [54], along with many of the results concentrating on analogue and digitally modulated communications signals from [55] and [56]. For the QAM signals, rectangular pulse shaping is used. Here, in contrast, the properties of raised cosine filtered signals are shown.

Other publications of Gardner and others take a more applied approach to filter structures which exploit cyclostationarity for various applications. For example, in [57] and [58], Reed looks at interference rejection in a range of scenarios with additive interference only, but using signals

with either rectangular pulses or pulses bandlimited with a Butterworth filter, neither of which are widely used in practice in narrowband communications. In [57], different modulation types of SOI and interference are considered, for signals with the same carrier frequency and different baud rates. [58] combines the demodulation into the (frequency domain) FRESH filter by adding the SOI carrier frequency to the filters' frequency shifts. The FRESH filter is compared with the matched filter in a simple BPSK scenario, and is found to be better, except when the interferer is very weak. Notice that Reed uses the term *time-dependent adaptive filter* instead of frequency-shift filter (or, more strictly, adaptive frequency shift filter). In [59, 60], more theoretical development of FRESH filtering is given with some more examples of simulated performance, but in the former, the important equivalence is shown between the FRESH filter and the FSE when both are followed by baud rate sampling. The spectral shapes considered there include 20% and 100% excess bandwidth (EBW) raised cosine, and the performance is measured for filters containing different numbers of frequency shifts.

Gardner summarises the previous work on FRESH filtering and its theoretical background in [13] and does the same at a simpler level in the magazine article [61]. There is little published work to date which examines performance variation with interferer carrier frequency, power or symbol timing, all of which are considered here.

In later work, there are various attempts to combine FRESH filtering with other techniques, such as array processing [58] and the DFE [62].

1.3.4 Other uses of cyclostationarity

In the subject of communications there are four areas, other than interference, where cyclostationarity can be a useful property. These are:

- blind channel equalisation or identification
- blind antenna array beam/null steering (see section 1.3.1)
- signal identification and detection
- synchronisation

Blind channel equalisation is a large subject which cannot be described here, other than to say that cyclostationary based techniques [63, 64, 65, 66, 67, 68] offer an alternative to higher-order statistics (HOS) based methods. HOS methods tend to require very long training periods compared to cyclostationary techniques, but the latter suffer from not being able to identify some channels [69]. The original suggestion that second-order cyclostationary statistics could be used for blind identification of dispersive channels came from Gardner [70].

Signal identification or detection (i.e. determining whether or not a signal is present, as opposed to receiving its data) can use the spectral correlation properties of cyclostationary signals by treating their spectral correlation density functions or cyclic autocorrelation functions (see chapter 2) as signatures. For example, the appropriate function can be calculated for the received, noisy, signal, and a search for matching features from a set of signatures of expected signals can be made [71, 72, 73].

Various synchronisation processes could be improved by reducing the noise in a signal with a FRESH filter before attempting to synchronise, but there have been some efforts to use the cyclostationarity of the received signal more directly [10, 74, 75].

1.3.5 Current Work

The use of spectral correlation or cyclostationarity looks attractive because it can exploit structure in the spectra of both the SOI and interference. Many other techniques use less specific properties, and so are likely to be less effective. The structures required for exploiting cyclostationarity are relatively simple (amounting to little more than a bank of FSEs) and make use of excess bandwidth which is likely to be present in the signals anyway.

The existing work is lacking in its analysis of FRESH filtering in realistic communications scenarios. This thesis uses either raised cosine filtered QPSK or BPSK, or Gaussian Minimum Shift Keying (GMSK), and, specifically in chapter 6, uses a scenario which closely matches the conditions encountered in a real communications system. This thesis also attempts to improve the current state of the literature on the subject, by highlighting the simplicity of the concepts and techniques involved. In particular, it is made clear that matched filtering exploits SOI baud rate spectral correlation, that BPSK carrier related correlation is equivalent to LCL filtering at baseband, and that this can be achieved by discarding the imaginary part of the baseband noise.

1.4 Structure of the Thesis

In chapter 2 the fundamental ideas of stationarity and cyclostationarity, and the frequency domain manifestation of cyclostationarity, spectral correlation, are defined and explained. The *cyclic autocorrelation function* and *spectral correlation density function* (SCD) are introduced, and several examples of SCD plots of digital communication signals are given. This chapter forms a background to the research and is based largely on pre-existing material (in particular [54, 13]) with the exceptions of the graphical representation of the two parameter autocorrelation function for a rectangular pulse sequence (section 2.1.5), and the interpretation of baud rate related spectral correlation as a result of the sampling theorem (section 2.2.2).

Chapter 3, "Frequency Shift Filtering", introduces the frequency shift filter (FRESH filter) structure and describes how it may be used in a receiver to replace, for example, a FSE or matched filter. Its adaptive form is shown, as is the equivalent structure, the periodically time-varying filter. The effect of sampling before or after the filter is also investigated both for the purpose of improving the efficiency of implementation, and of highlighting the important equivalence of the FSE and FRESH filters when they are followed by baud rate sampling. How, then, the FSE and FRESH filter exploit spectral correlation, is also shown. The exploitation of the properties of cyclostationary interference is shown, and the chapter finishes by describing the links between FRESH filtering and linear-conjugate-linear (LCL) filtering, and the property of circularity. The original work in this chapter comprises the graphical representations of how the FRESH filter and FSE reject interference (sections 3.5.1 and 3.6.7), and the analysis of the effect of sampling in FRESH filters (sections 3.6.3 to 3.6.6). The explanation of how the LCL filter is the baseband representation of the FRESH filter with a frequency shift of twice the carrier, is also novel, as is the link between

circularity and carrier related spectral correlation.

The implementation of FRESH filtering for a QPSK SOI and interferer and where both the SOI and interferer are raised cosine filtered, but where the interferer has a much higher power than the SOI is shown in chapter 4. This difference in performance is shown between a filter which exploits only SOI properties, and one which also uses interferer cyclostationarity. To verify the correct operation of the simulations, the cyclic Wiener filter equations [13] are solved for some example scenarios and the filter responses compared with those from the simulations, which were obtained by using an adaptive algorithm to drive the filters to the minimum MSE solution. The resulting values of the MSE are also compared with the theory and found to give good agreement. This material is an implementation of previously published filter structures, and a specific solution of general known filter equations, but as such is all original.

The next chapter (chapter 5) explores the possibility of improving FRESH filter performance by choosing a SOI pulse shape to enhance the signal's cyclostationarity. This also provides an opportunity to test FRESH filter performance in some other interference scenarios, including an interferer of equal power to the SOI, an interferer with equal or different baud rate to the SOI, and an interferer with different carrier frequencies. It is found that there is no obvious improvement to be gained from manipulating the pulse shape of the SOI.

Very Low Frequency (VLF) radio communications provide a potential application for FRESH filters and this is described in chapter 6. Interference scenarios where there is ACI up to 40 dB more powerful than the SOI provide an opportunity to test the most effective frequency shifts to use. The modulation scheme used is GMSK. As VLF communications routinely use techniques to mitigate impulsive noise, the possibility of operating FRESH filters in a system with one of these techniques (hole-punching) is investigated and some problems are identified.

Finally the issue of blind adaptation of the FRESH filter is considered in chapter 7. The aim here is to look at methods for adaptation which arrive at the optimum filter solution, without being concerned about adaptation rate. Two published adaptation algorithms are implemented: one is found to be flawed in that it does not adapt to the correct solution; the other is shown to have severe limitations, and to result in a structure which does not have the interference rejection ability of the FRESH filter. A change to the filter is proposed which, although in practical terms is very simple, is sufficient to give performance as good as a trained FRESH filter.

The conclusions of the work are then summarised in chapter 8.

1.5 Simulation Tools and Methods

The majority of the results shown in this thesis were obtained through Monte-Carlo simulation using commercial software called Signal Processing Worksystem (SPW).

This a simulation package which allows a Monte-Carlo simulation program to be built up graphically from a large collection of standard library blocks. These blocks may represent functions such as random signal sources, arithmetic processes, general filtering, error counting and so on. It also allows extra functionality to be included by writing new blocks in C. SPW converts the graphical representation of the system to C code, which is then compiled and run. SPW can be

described as a "waveform" or "dataflow" simulator [76] meaning that the entire waveform is simulated through discrete time samples, and the simulation runs processes on each sample as required by the program.

It is important to be aware of the limitations of the simulation methods used. When representing a continuous quantity, a discrete time equivalent signal must be used, subject to the limitations of the sampling theorem (i.e. that the sampling frequency must be at least twice the bandwidth of the signal). The dataflow method of simulation has implicit synchronisation built in, in that every sample "transmitted" through a channel can be "received" as if with perfect time, phase and frequency synchronisation, unless specific efforts are made to simulate synchronisation errors. This implicit synchronisation can be useful when one is interested in other processes, or sources of error, but it must be recognised that such simulations do not represent the real world conditions a receiver will operate under.

In this thesis the simulations are validated through comparison with theoretical methods where possible (as in chapter 4). In this chapter it was possible to compare values from analytic expressions for the mean squared error (MSE) of the output of the systems, which were then compared with the simulated value. However most of the rest of the present work measures performance in terms of bit error rate (BER) this being a more useful performance indicator for communications systems. Analytic expressions for BER are only available for simple situations such as QPSK received in the presence of AWGN only. Wherever possible and appropriate, simulation systems were tested with such simple scenarios so that a comparison with accepted theory could be made.

In general, other than the synchronisation implicit in the method, no information was used in the simulated receivers which would not be available to a real receiver, other than for the purposes of quantifying performance. For example, adaptive filters were used, trained with a standard algorithm, rather than using filter solutions based on ensemble averages.

The channels modelled for this work were all fairly simple and well understood. All were purely additive: the corruption introduced was either white Gaussian noise, or impulsive noise (in chapter 6) with an interfering signal also present.

Counting bit errors in a system containing several random processes, such as signal and noise sources, leads to results which are also fundamentally random. However, one can be confident (in a quantifiable way) that the resulting BER is a good representation of the "true" result which one is aiming for. The theory underlying the reliability of BER measurements is well described in [76]. For many practical measurements, the "rule of thumb" that counting 10 errors gives a BER within a factor of 2 with 95% confidence was used. Good computing resources were available for the simulations presented in this thesis, so all BER measurements were taken with at least 100 bit errors. This gives a BER within about 25% of the true value with a confidence of 95%.

The confidence limits above assume independent error events, so they do not apply, for example, to situations where a single fade in a channel may cause multiple errors. In this thesis the only possibility for such slowly changing conditions leading to correlated errors is in the effect of beat frequencies, when, for example an interferer has a baud rate very close to that of the SOI. In such situations, longer simulations were run, if required to ensure that BER vs. E_b/N_0 curves were smooth. This comparison of independent estimates of BER with different SNR is a recognised way of improving confidence in BER estimates, albeit one whose effect is hard to quantify.

MATLAB was also used in a limited way for producing the spectral correlation density plots. Here the desired random signal was produced and filtered, and spectral analysis was performed, using standard functions, on the result. The number of symbols in each signal was varied, until a reasonably smooth graph was obtained. These graphs are intended to illustrate the properties of the signals in a qualitative way only.

Chapter 2

Stationarity, Cyclostationarity and Spectral Correlation

Contents

This section describes the property of cyclostationarity of signals, in comparison to stationarity, in both the time and frequency domains. This forms a background to the following chapters where filter structures which exploit cyclostationarity are investigated. Many digital communications signals are cyclostationary; the property is a result of the implicit periodicity of these signals, related to the baud rate, carrier frequency or any other periodic component. This modelling of signals as cyclostationary is a departure from the traditional stationary model, on which much of classical signal processing theory is built.

Cyclostationary signals have periodically time-varying second order statistics (in particular, they have a periodic autocorrelation); stationary signals have second order statistics which are constant with time. As Wiener filtering theory is based on the second order statistics of the signal being processed, the more specific cyclostationary model allows an improvement in filtering performance over classical Wiener filtering. This is described in chapter 3, but it is first necessary to describe the property of cyclostationarity and how it compares with the stationary model. A time domain approach to this, which concentrates on the autocorrelation function, is presented in section 2.1,

while in section 2.2 the frequency domain manifestation of cyclostationarity, spectral correlation, is described. The frequency domain approach allows a more intuitive understanding of how the filters work which exploit cyclostationary.

The majority of the initial work on cyclostationarity related to communications was done by W. A. Gardner and his co-workers [8, 9, 53, 52, 54, 59, 60, 61, 77, 61, 13]. However to clarify what can be a difficult subject, a number of new interpretations and representations of cyclostationarity are presented in this chapter, in particular the 2D surface representing the two-parameter autocorrelation function in section 2.1.5 and the link between spectral correlation and the sampling theorem described in section 2.2.2. Cyclostationarity based techniques are well known in the areas of antenna array beam steering and blind channel estimation and equalisation, but the interest here is in the less well researched area of interference rejection in a single antenna receiver.

2.1 Time Domain Representation of Cyclostationarity

This section examines some details of the time domain manifestation of cyclostationarity, and in particular compares the two parameter cyclic autocorrelation with the more conventional autocorrelation function.

This examination of the two parameter autocorrelation function is illustrated by deriving the two-parameter autocorrelation function surface for a random bipolar rectangular pulse sequence. The standard autocorrelation function of this waveform is well known and the fact that the signal only takes the values of -1 and 1 means it is relatively easy to calculate its time-varying autocorrelation values.

The two parameter autocorrelation function is the starting point in the analysis of cyclostationary signals. Periodicity of this function with time results in spectral redundancy which can be exploited by advanced signal processing methods to give an enhanced immunity of a signal to interference.

2.1.1 Autocorrelation and mean definitions

Autocorrelation is a measure of how closely a signal, or sequence, is correlated with time shifted versions of itself. Two relatively time shifted versions of the signal are multiplied together and then some form of ăveraging is used to give a result which is a function of time or time shift.

We have three types of autocorrelation to consider:

- the autocorrelation of a specific finite signal or sequence, where the average is a time average over the whole signal;
- the autocorrelation of a *random signal* (or *stochastic process*) which is a probabilistic concept involving an average over an ensemble of different realisations of the random process;
- the autocorrelation of a random signal, where the signal has infinite duration and a time average in the limit of infinite time replaces the ensemble average used for random signals.

The first of these gives the real properties of a specific signal which will affect the performance of, for example, a filter used to process the signal. However it is a unique function for each signal.

The other two types of autocorrelation give the "average" properties of a group of signals - the group being defined by what signals make up the ensemble, or in what way the time averaging is done.

Finite sequence definition

For a signal $x(t)$ of duration $2P$ there are three equivalent definitions of autocorrelation. The first is:

$$R_x(t,\tau) = \frac{1}{2P}\int_{t-P}^{t+P} x(t')x^*(t'-\tau)dt' \tag{2.1}$$

or, shifting the origin of t:

$$R_x(t,\tau) = \frac{1}{2P}\int_{t-P}^{t+P} x(t'+\frac{\tau}{2})x^*(t'-\frac{\tau}{2})dt' \tag{2.2}$$

Notice that unlike in standard treatments of this subject, the autocorrelation is written explicitly as a function of the time origin of the calculation, t. This is because peridocity of the autocorrelation with respect to t will be considered later. τ represents the lag between the two instances of the signal.

This can also be written in terms of t_1 and t_2, where $t_1 = t - \tau/2$ and $t_2 = t + \tau/2$:

$$R_x(t_1,t_2) = \frac{1}{2P}\int_{\frac{t_1+t_2}{2}-P}^{\frac{t_1+t_2}{2}+P} x(t')x^*(t'-t_1+t_2)dt' \tag{2.3}$$

so now the difference between t_1 and t_2 is the relative time shift or lag, τ, between the two versions of the signal.

This is a time-average of the signal, multiplied by a time shifted version of itself, averaged over the duration of the signal. The exact value of this function will depend on the particular sequence used, but as this sequence becomes longer the finite autocorrelation will tend towards the function given by the probabilistic and infinite time-averaging definitions below.

Probabilistic definition

The autocorrelation $R_x(t,\tau)$ of a random process $x(t)$ is the expected value of the product $x(t)x^*(\tau)$, that is,

$$R_x(t,\tau) = E[x(t)x^*(t-\tau)] \tag{2.4}$$

or

$$R_x(t,\tau) = E[x(t+\frac{\tau}{2})x^*(t-\frac{\tau}{2})] \tag{2.5}$$

As before we can express this in terms of t_1 and t_2:

$$R_x(t_1,t_2) = E[x(t_1)x^*(t_2)] \tag{2.6}$$

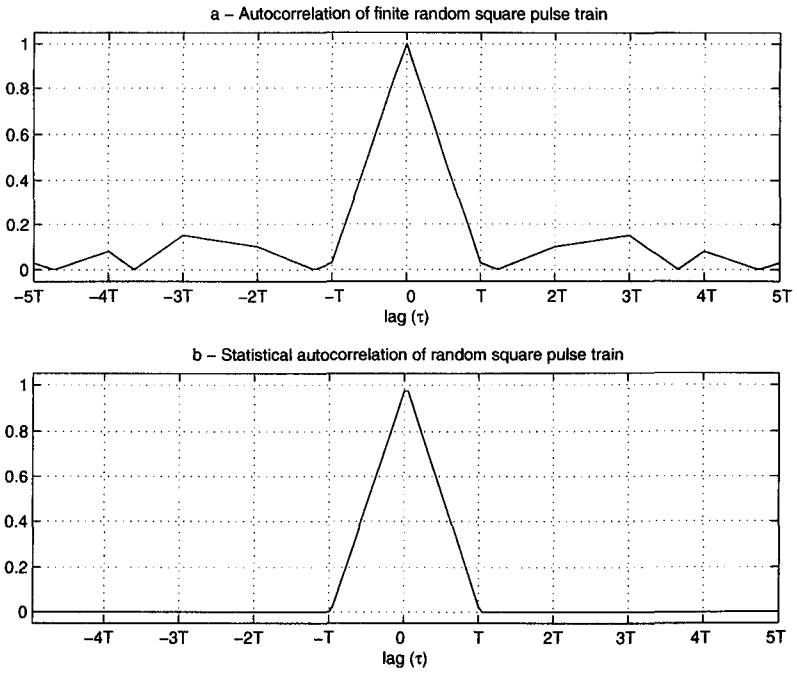

Figure 2.1: Probabilistic and finite autocorrelation for rectangular pulse sequence

where $E[]$ indicates the *expected value* or *expectation* [78]. This means an average over an imaginary infinite set of occurrences of $x(t)$. This infinite set is the *ensemble*.

Infinite time-averaging definition

The time averaging definition of the autocorrelation of a signal $x(t)$ is

$$R_x(t,\tau) = \lim_{2P\to\infty} \frac{1}{2P} \int_{t-P}^{t+P} x(t')x^*(t'-\tau)dt' \tag{2.7}$$

or

$$R_x(t,\tau) = \lim_{2P\to\infty} \frac{1}{2P} \int_{t-P}^{t+P} x(t'+\frac{\tau}{2})x^*(t'-\frac{\tau}{2}dt' \tag{2.8}$$

or

$$R_x(t_1,t_2) = \lim_{2P\to\infty} \frac{1}{2P} \int_{\frac{t_1+t_2}{2}-P}^{\frac{t_1+t_2}{2}+P} x(t')x^*(t'-t_1+t_2)dt' \tag{2.9}$$

Mean definition

For completeness, the three definitions of the mean of a signal, which correspond to the three autocorrelation definitions above, are given here:

$$\mu_x(t) = \frac{1}{2P} \int_{t-P}^{t+P} x(t')dt' \tag{2.10}$$

$$\mu_x(t) = E[x(t)] \tag{2.11}$$

$$\mu_x(t) = \lim_{2P\to\infty} \frac{1}{2P} \int_{t-P}^{t+P} x(t')dt' \tag{2.12}$$

2.1.2 Autocorrelation of a random rectangular pulse sequence (stationary)

To illustrate the difference between the various definitions of autocorrelation, and the difference between stationarity and cyclostationarity, the autocorrelation is shown of a signal composed of rectangular pulses carrying random data. In this section the standard (time-invariant) autocorrelation of such a signal is presented. In section 2.1.5 the surfaces representing the two parameter autocorrelation of the same signal are shown, using both the stationary and cyclostationary models.

The first example, figure 2.1*a* shows the finite sequence autocorrelation of a random rectangular pulse sequence of 100 symbols which can take the values of -1 or $+1$. The symbol period is T. Figure 2.1*b* shows the graph of the probabilistic or infinite time-averaging definition. Both show the similar triangular shape, but whereas b is identically zero outside the range $-T$ to T, the finite sequence definition is "noisy". This is a result of the sequence being of finite length. Averaging over an infinite time would give the "ideal" graph of 2.1*b*.

For the probabilistic definition it is the averaging over an *infinite* set of possible realisations of the process that removes the noise shown in 2.1*a* to give the noise free function of figure 2.1*b*. The sequence could be finite or infinite. Conversely, the structure of the noise in figure 2.1*a* is determined by the particular data values in the finite data sequence analysed.

However for practical signals, it is only possible to measure a finite sequence, so this function is what is used as an estimate of the probabilistic function. If a sufficiently long sequence can be used in the averaging, then good results can be obtained.

The fact that the autocorrelation is plotted as a function of τ only, means that stationarity has been assumed and any variation with t will not be seen.

2.1.3 Definitions of stationarity and cyclostationarity

Stationarity

A wide sense stationary signal or process is one in which the mean value and autocorrelation are invariant in time. This is expressed mathematically as:

$$E[x(t)] = constant \tag{2.13}$$

$$R_x(t,\tau) = constant \tag{2.14}$$

for all values of t.

A common description of this property is that the second order statistics of the signal are constant with time (or more correctly the first order (mean) and second order (autocorrelation) are constant). There is a more restrictive stationarity definition, narrow-sense stationarity, which requires the signal to have a probability density function which is invariant with time. Effectively this means

that as well as constant second order statistics, the signal should have all higher statistical moments constant with time. In the context of Wiener filtering however, the higher order statistics are not used, so the wide-sense definition is sufficient, and is, of course, simpler to work with. This also applies later to the narrow and wide sense definitions of cyclostationarity. For the work presented in this thesis, the wide sense definitions are sufficient, so all further uses of the terms stationary and cyclostationary imply wide-sense.

Figure 2.1 uses the conventional way of displaying an autocorrelation function; it is assumed that it is a function of τ only, that is it is independent of t. This is equivalent to assuming that the signal is stationary. Alternatively we may say that if we are not interested in any variation with t, then the autocorrelation can be averaged over all time t to give a function of lag, τ only.

This style of display is therefore unsuitable for examining the autocorrelation of non-stationary signals. In the following sections, cyclostationary signals are defined and time-varying autocorrelation is plotted.

Cyclostationarity

A wide sense cyclostationary signal or process is one in which the mean value and autocorrelation are periodic. This applies to all the definitions of these quantities given in section 2.1.1. That is (using the probabilistic definitions):

$$E[x(t)] = E[x(t + mT_0)] \tag{2.15}$$

$$R_x(t,\tau) = E[x(t)x^*(t-\tau)] = E[x(t+mT_0)x^*(t+mT_0-\tau)] \tag{2.16}$$

where m is an integer. The subscript x on R_x indicates that R is the autocorrelation of the signal $x(t)$. $R_x(t,\tau)$ is the probabilistic autocorrelation function expressed as a function of two variables, τ is "parametric time" - the lag between the two signals, and t is "real time" - the time origin of the autocorrelation calculation. Equation 2.16 describes an autocorrelation function periodic in t with period T_0. That is, the fundamental period is T_0, and R_x is also periodic at harmonics of T_0. In a communications signal T_0 will be related, for example, to a carrier, symbol or chip frequency. If R_x has more than one fundamental period and they are not harmonically related (as in a digitally keyed carrier where the keying and carrier frequencies are not harmonically related) then the signal is said to be *polycyclostationary* [54].

There is a corresponding narrow sense definition of cyclostationarity which requires that the probability density function of the signal is periodic in time. However for filtering applications the wide sense definition is sufficiently restrictive.

Ensemble averaging and cyclostationarity

The properties of the autocorrelation (and mean) based on the probabilistic definition are of course dependent on how the ensemble is defined. An ensemble is an imaginary set of an infinite number of instances of the process or signal under consideration, but it is important to be aware of the similarities and differences between these different instances.

If we take the example of the autocorrelation of a random rectangular pulse sequence with symbol period T, which can take the values -1 or 1, all members of the ensemble used in this case will also be random rectangular pulse sequences with period T, and amplitude -1 or 1, but with different random data sequences. Whether the sequences are synchronous or not (that is, whether symbol transitions occur at the same time in each ensemble member) determines if the process is stationary or cyclostationary.

Consider figure 2.2: graph *a* illustrates a member of the ensemble representing a random rectangular pulse sequence; graph *c* shows another member. Both have the same pulse shape and period, but carry different data. There is also an arbitrary shift in their symbol timing relative to each other. Over the whole ensemble the symbol timing shifts would be uniformly distributed from 0 to T.

Graph *b* shows the same waveform as in *a* but delayed by τ. In the calculation of the autocorrelation, *a* is multiplied by *b* (for every value of τ), *c* is multiplied by *d*, and so on for every member of the infinite ensemble. This describes the conventional calculation of probabilistic autocorrelation

In contrast, figure 2.3 shows how the autocorrelation is calculated to reveal the cyclostationarity of the rectangular pulse sequence. In this case, graph *c* shows a waveform with different data from that in graph *a*, but with synchronised transitions between the symbols. All the members of the ensemble are constrained to have their symbol transitions occurring together. It is this constraint which reveals the waveform to be cyclostationary with a fundamental period of T, in contrast with the uniform distribution of symbol timings in the stationary case.

Notice that applying a similar constraint to the averaging of a signal with no inherent periodicity (such as a white noise process) would not result in a periodic autocorrelation function. The periodicity of the autocorrelation is not a result of the constraint, but it is necessary to apply the correct constraint to reveal a particular signal's cyclostationarity.

Time averaging and cyclostationarity

In the literature it is normal to use the probabilistic approach to autocorrelation and to assume that it applies to real random signals of long enough duration. As the duration of a finite signal increases, the finite deterministic autocorrelation function obviously tends towards the infinite deterministic autocorrelation. As time tends to infinity, the deterministic autocorrelation tends towards the probabilistic autocorrelation as long as the signal or process is ergodic [1]. In practice ergodicity requires only uncorrelated data ([17] page 99), which is a common and reasonable assumption for communications signals. For signals where this is not the case, the data can be encoded to ensure it is uncorrelated (see, for example, [79] chapter 3). However, much of the theoretical work on cyclostationarity has been done by W.A. Gardner, who prefers to use the non-probabilistic, time averaging theory [54], so this approach is also described.

We have seen that an ensemble with symbol synchronous waveforms gives a periodic autocorrelation which shows the waveform is cyclostationary; and that an ensemble with asynchronous symbol transitions gives a stationary waveform. The analogy using time-averaging is that to reveal the cyclostationarity of a waveform, the time-averaging must by synchronised, that is, the

[1] This is in fact the definition of ergodicity. See [78] page 427.

waveform is averaged with versions of itself, shifted in time by its cyclic period. To make such a waveform stationary, there must be random changes in the symbol clock phase. As the time averaging is taken over a period tending to infinity, there will be enough random phase jumps to remove the innate periodicity of the whole signal (see also section 2.1.5).

Gardner uses the following definition (the *limit periodic autocorrelation*) for analysing cyclostationary signals:

$$R_x(t_1, t_2) = \lim_{N\to\infty} \frac{1}{2N+1} \sum_{n=-N}^{N} x(t_1 + nT_0)x^*(t_2 + nT_0) \tag{2.17}$$

where n is an integer, T_0 is the cyclic period, which in the current example is T, the symbol period.

This is *synchronised averaging* - instead of a symbol synchronous ensemble average, the averaging is done over the same signal shifted by an integer number of symbols. This shifting is done an infinite number of times. Effectively we are then averaging over all symbols in the signal. As long as the data are uncorrelated, this is equivalent to an infinite number of different realisations of the signal; that is, a conventional ensemble average. This definition is used by Gardner in [54], because he believes that it avoids conceptual difficulties associated with the abstract idea of a statistical ensemble. However the statistical ensemble is a well understood tool in signal processing and the current author believes that the time-averaging approach is cumbersome and obscures the otherwise interesting content of [54]. The statistical approach is used henceforth in this work. It is interesting to note that a review [80] of [54] led to strongly worded correspondence attacking and defending Gardner's approach [81, 82, 83, 84, 85, 86]. This correspondence developed into an argument of the usefulness of the concept of *circularity* [85, 87, 88, 89, 90] which is addressed in section 3.8.

The equivalent process of synchronised averaging applied to the time averaging definition of cyclostationarity is described in [54] chapter 10, which is based on old techniques, apparently published as early as 1847 [91, 92].

The equivalent non-probabilistic situation is a signal carrying data in the form of rectangular pulses, where there is no change in the symbol clock phase. In that case a time translation of an integer number of symbols will always result in a signal which is symbol synchronised with the original.

2.1.4 Cyclic autocorrelation

Cyclic autocorrelation function

Definitions and descriptions now follow of a function which is useful in quantifying and illustrating the cyclostationary properties of a signal: the *cyclic autocorrelation function* [13, 54]. The related *spectral correlation density function* is described in section 2.2.4.

As $R_x(t, \tau)$ (equation 2.16) is periodic in t, we can write it in the form of a Fourier series expansion:

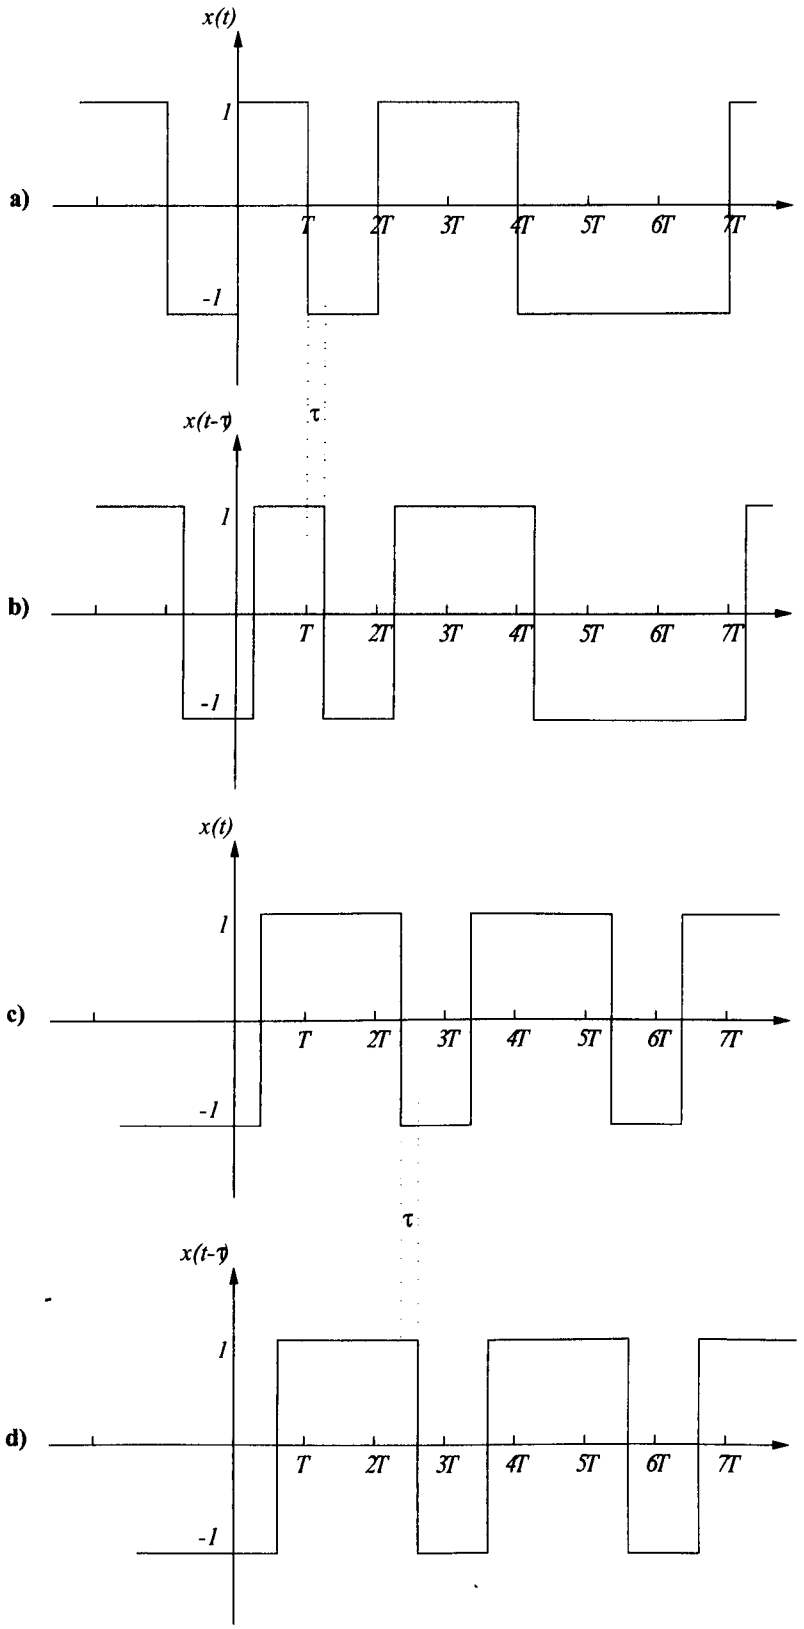

Figure 2.2: Two members (a and c) of an ensemble for the calculation of the stationary autocorrelation of a rectangular pulse sequence, and their delayed versions (b and d)

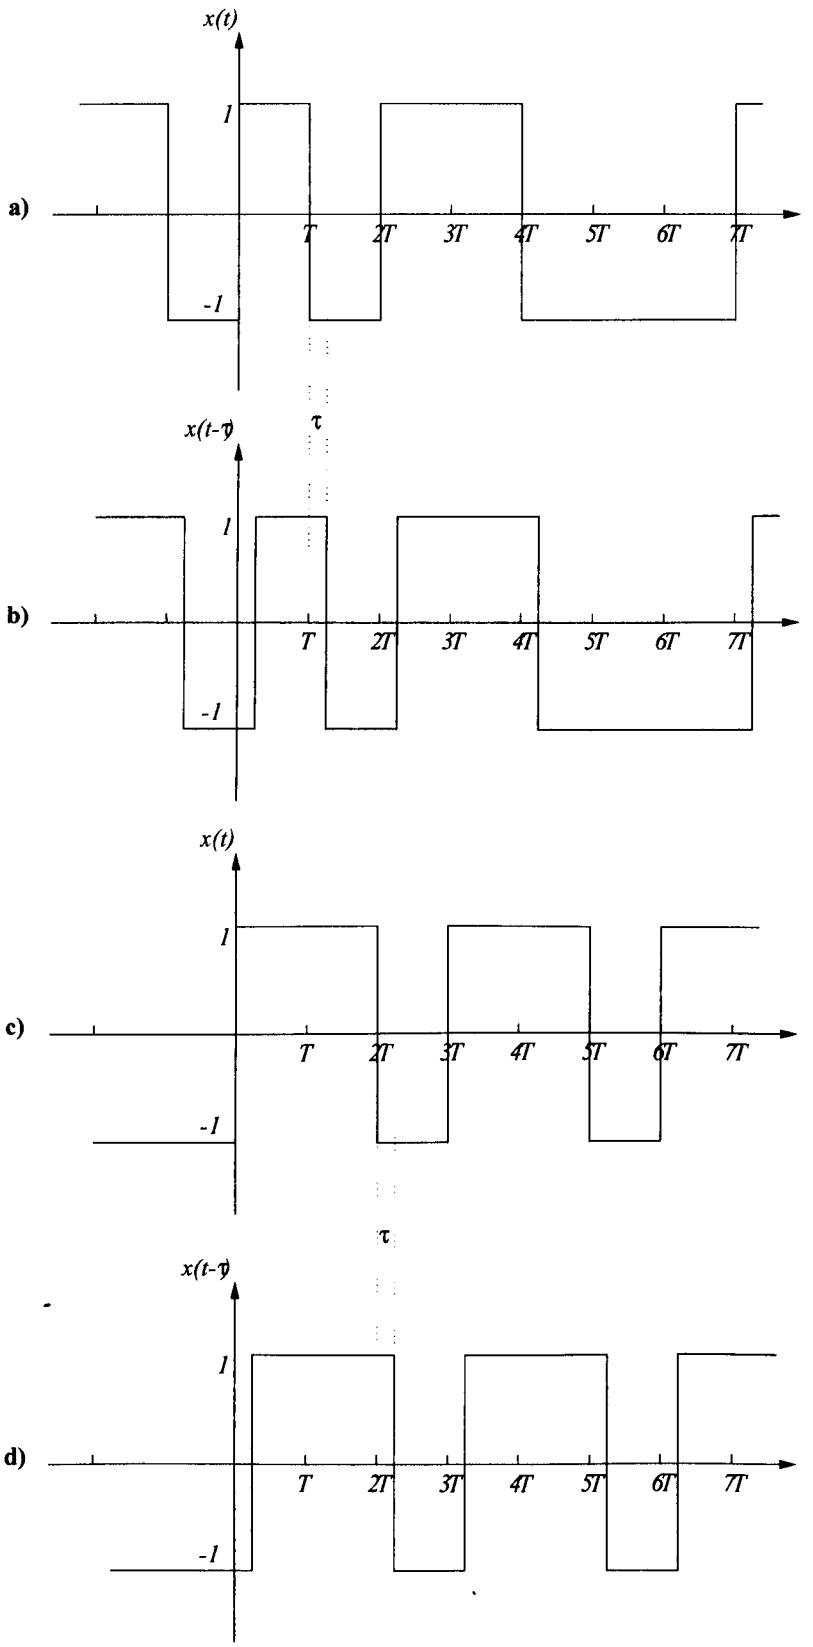

Figure 2.3: Two members (a and c) of an ensemble for the calculation of the cyclostationary autocorrelation of a rectangular pulse sequence, and their delayed versions (b and d)

$$R_x(t,\tau) = \sum_\alpha R_x^\alpha(\tau)e^{j2\pi\alpha t} \tag{2.18}$$

Here the summation can be over all values of α although the coefficient R_x^α will be zero unless α is equal to a period of the autocorrelation function. For example if the baud rate $1/T$ is a cyclic frequency, then α will take values $0, 1/T, 2/T, ...$ because the baud rate harmonics are also cyclic frequencies.

If $R_x(t,\tau)$ is not periodic, then all coefficients of the summation will be zero except for R_x^0. If there is periodicity present then at least one other R_x^α will be non-zero. The quantities R_x^α are therefore useful as a fundamental parameter of cyclostationarity. This parameter is called the *cyclic autocorrelation function* [54], and is defined as any other Fourier coefficient:

$$\begin{aligned} R_x^\alpha(\tau) &= E[R_x(t,\tau)e^{-j2\pi\alpha t}] &(2.19)\\ &= E[x(t+\frac{\tau}{2})x^*(t-\frac{\tau}{2})e^{-j2\pi\alpha t}] &(2.20)\end{aligned}$$

However, we can interpret this equation in a slightly different way, by recognising that $x(t-\tau/2)e^{j2\pi\alpha t}$ is simply the signal $x(t-\tau/2)$ shifted in frequency by α. This means that if a signal x is cyclostationary with a cyclic frequency α, then there is non-zero correlation between x and x shifted in frequency by α. This phenomenon is known as *spectral correlation* and is simply the frequency domain manifestation of cyclostationarity. This is discussed further in section 2.2.

In digitally modulated signals, the cyclic frequencies are usually related to the baud rate and the carrier frequency. Spread spectrum signals may have additional cyclic frequencies present such as a chip rate in direct sequence spread spectrum, or the hopping frequency in frequency hopping spread spectrum. As an example, a rectangular pulse BPSK signal with baud rate $1/T$ and carrier frequency f_c has cyclic frequencies $\pm 2f_c, \pm 1/T, \pm 2/T, \pm 3/T, ..., \pm 2f_c \pm 1/T, \pm 2f_c \pm 2/T, \pm 2f_c \pm 3/T, ...$ [54, 13].

Conjugate cyclic autocorrelation function

Another function which can describe the cyclostationarity of some signals is the *conjugate cyclic autocorrelation function* [54] which is defined as:

$$R_{xx^*}^\alpha(\tau) = E[x(t+\frac{\tau}{2})x(t-\frac{\tau}{2})e^{-j2\pi\alpha t}] \tag{2.21}$$

This function represents the correlation between the signal and its complex conjugate frequency shifted by α. A simple example of the conjugate cyclic ACF is for purely real signals, such as baseband BPSK or any real modulated carrier. As the signal is purely real, then $x = x^*$ and clearly $R_x^\alpha = R_{xx^*}^\alpha$ for all α. For many complex signals, such as baseband QPSK, the conjugate cyclic ACF will be zero for all values of α while for some, baseband MSK, for example, the function is more interesting. However it is easier to understand the differences between cyclic correlation and conjugate cyclic correlation by examining signals' properties in the frequency domain. This is done in section 2.2.

2.1.5 Autocorrelation surfaces for stationary and cyclostationary signals

The sections above contain definitions of autocorrelation, and stationary and cyclostationary processes with reference to a random rectangular pulse sequence. In this section, three dimensional graphical representations are shown of the two parameter autocorrelation function for these two types of rectangular pulse sequence.

The autocorrelation of a signal is normally shown as a function of one variable: the lag, τ, in equation 2.4. In a stationary signal there is no dependency on t, the value of the time origin, so it is not plotted. We display the autocorrelation here as a function of t and τ, or equivalently t_1 and t_2, to illustrate the difference between stationary and cyclostationary signals. For cyclostationary signals the autocorrelation function varies with both t and τ.

Stationary signal

The autocorrelation of a random rectangular pulse sequence has the well known triangular shape seen in figure 2.1*b* when the stationary model of the process is used. The horizontal axis in this graph is τ where τ is as in equation 2.4. That is, the autocorrelation depends only on the lag between the signal and its time shifted representation, not on the value t, the absolute time variable. It is instructive, however, to plot it as a function of t_1 and t_2: figure 2.4 represents equation 2.6 for a random rectangular pulse sequence with a limited set of values of t_1 and t_2. In figures 2.4 to 2.8 the symbol period is T_0.

The values of the function are represented by four lines on the graph, which are shown as the edges of two planes to make the diagram clearer. In the limited range shown, the autocorrelation values are obvious by extension from the traditional autocorrelation plot in figure 2.1*b*. The triangular shape of that figure is now seen twice, on the edges of the two planes. The symmetry of equation 2.6 results in the symmetrical graph.

Figure 2.6 only shows values for t_1 or t_2 equal to 0. We can extend this further to the full surface representing the two variable autocorrelation: this is shown in figure 2.5. The four lines shown in figure 2.4 lie on the surface of a "tent" shape. It can be seen that this is the correct shape by remembering that the autocorrelation must be 1 for all $t_1 = t_2$, and that the value of the autocorrelation depends only on the difference between t_1 and t_2, not on their absolute values. The function extends to infinity in both t_1 and t_2. We can also plot this as a function of t and τ as in figure 2.6.

Cyclostationary signal

The cyclostationary case is different as a direct result of the symbol alignment of members of the ensemble. The value of the autocorrelation is either 0 or 1; there is no gradual change from 0 to 1 as in the stationary case.

Consider just one pulse of the pulse sequence (see figure 2.7), for convenience centred around $t = 0$ (using the equation 2.5 definition of autocorrelation). When $t = 0$ and $\frac{\tau}{2} \leq \frac{T_0}{2}$, $x(t + \tau/2)$ and $x(t - \tau/2)$ always have the same value, as they lie within the same pulse. Therefore for these values of t, τ and T_0 we have $R_x(t, \tau) = 1$. When $t = 0$ and $\frac{\tau}{2} > \frac{T_0}{2}$, $x(t+\tau/2)$ and $x(t-\tau/2)$ no

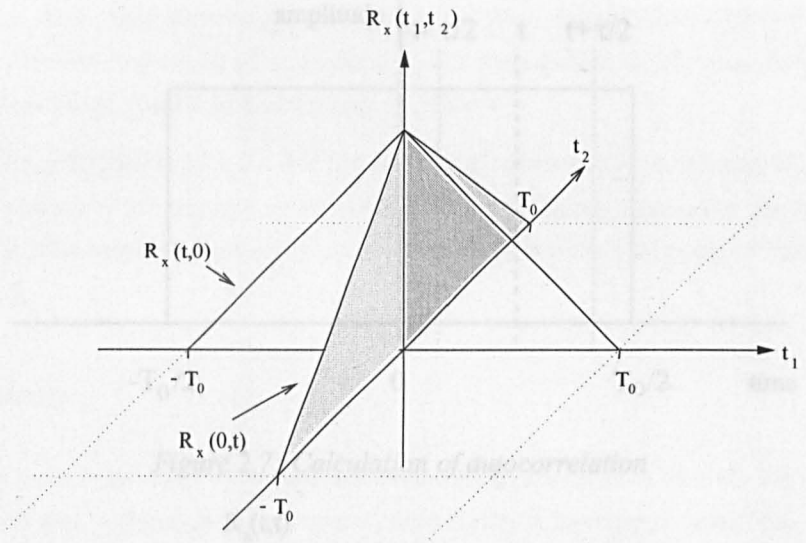

Figure 2.4: Part of autocorrelation of stationary rectangular pulse sequence

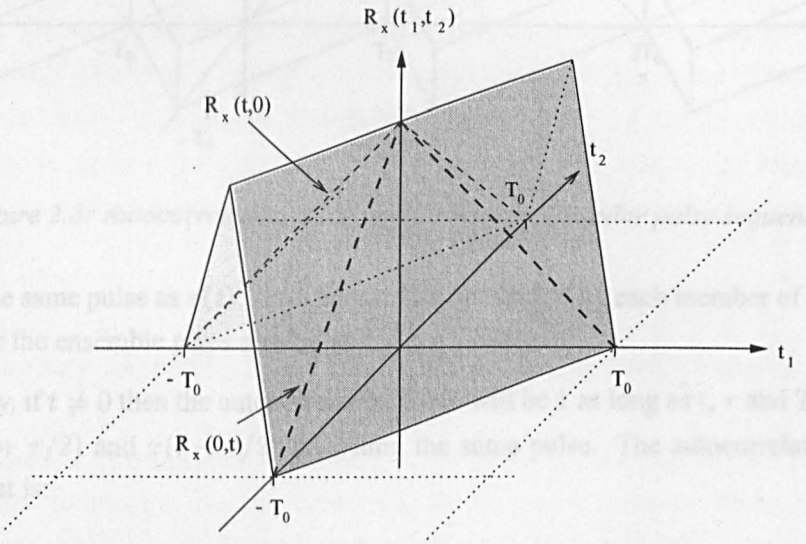

Figure 2.5: Autocorrelation of stationary rectangular pulse sequence

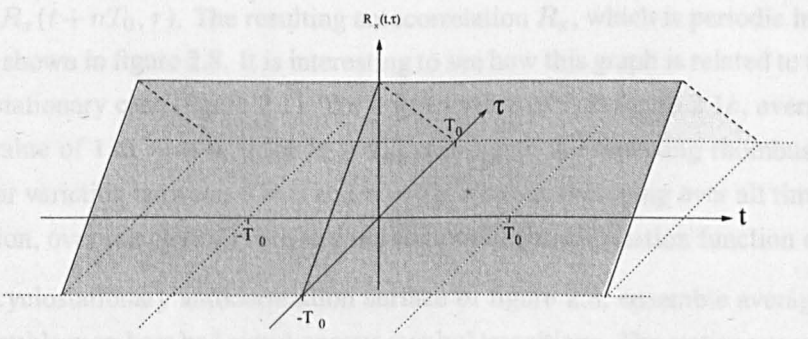

Figure 2.6: Autocorrelation of stationary rectangular pulse sequence

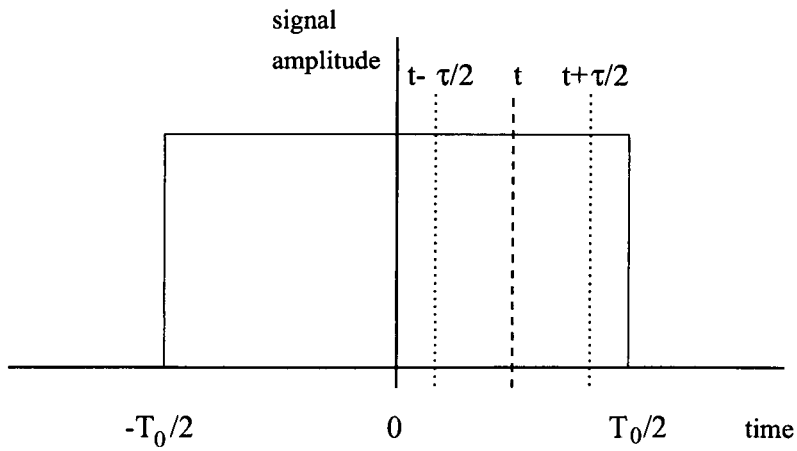

Figure 2.7: Calculation of autocorrelation

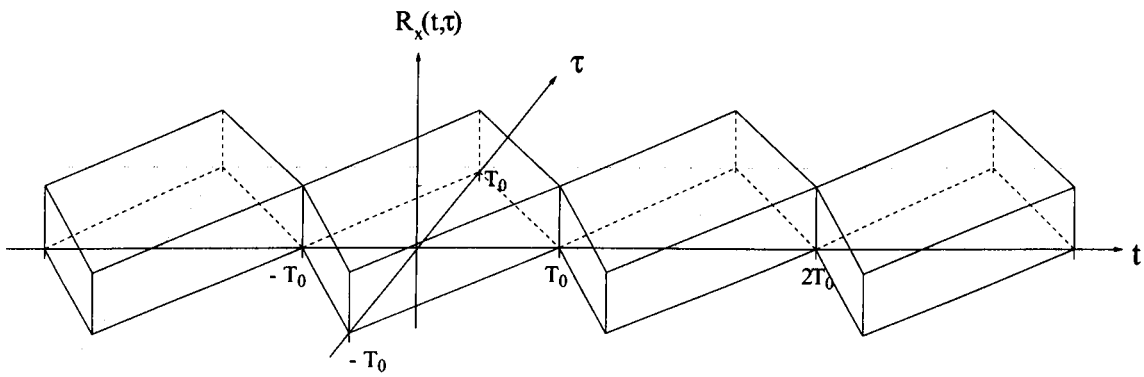

Figure 2.8: Autocorrelation of cyclostationary rectangular pulse sequence

longer lie in the same pulse as $x(t)$. As different data are carried by each member of the ensemble, averaging over the ensemble gives a value of $R_x(t,\tau) = 0$.

More generally, if $t \neq 0$ then the autocorrelation value will be 1 as long as t, τ and T_0 have values such that $x(t+\tau/2)$ and $x(t-\tau/2)$ lie within the same pulse. The autocorrelation will be 0 otherwise. That is:

$$R_x(t,\tau) = \begin{cases} 1 & \tau < T_0 - 2|t| \\ 0 & \text{otherwise} \end{cases} \tag{2.22}$$

The same argument can be applied to other pulses, that is for values of $t = t + nT_0$, which results in $R_x(t,\tau) = R_x(t+nT_0,\tau)$. The resulting autocorrelation R_x, which is periodic in t for fixed τ, $0 < \tau < T_0$ is shown in figure 2.8. It is interesting to see how this graph is related to the equivalent graph for the stationary case (figure 2.1). For a fixed value of τ in figure 2.1*b*, averaging over all t will give a value of 1 at $\tau = 0$, 0 for $\tau \geq T_0$, and due to the repeating rhombus shape of the surface, a linear variation between $\tau = 0$ and $\tau = T_0$. That is, averaging over all time (or as it is a periodic function, over one period) recovers the stationary autocorrelation function of 2.1*b*.

To create the cyclostationary autocorrelation surface of figure 2.8, ensemble averaging was used where the ensemble members had simultaneous symbol transitions. The stationary autocorrelation function results from an ensemble where each member's symbol transition times are not linked. This has a physical significance: if we add together a large number of cyclostationary waveforms,

which have the same cyclic frequencies, and in addition have simultaneous symbol transitions, then the sum is also cyclostationary, with the same cyclic frequencies; however adding many cyclostationary waveforms which have the same cyclic frequencies, but have uniformly distributed symbol transition times, results in a stationary waveform.

This is the effect described in [45, 93, 44] for the case of cross-talk in telephony, where linking the symbol clock phases of the transmitter in the telephone exchange causes the the cross-talk to be cyclostationary. The same phenomenon is described from a statistical point of view in [78] page 374, Theorem 2.

2.1.6 Summary

The time domain description of cyclostationarity and its representation through the cyclic autocorrelation function and two-parameter autocorrelation function have been described, with reference to the example of a signal carrying random data with a rectangular pulse shape.

It has been shown that the two parameter autocorrelation surface for a random rectangular pulse sequence is as shown in figure 2.8. The important feature of this graph is that $R_x(t, \tau)$ is periodic in t for any fixed τ. This periodicity means that the signal exhibits spectral correlation, as described in section 2.2. The exploitation of this spectral correlation, which is demonstrated in chapters 4 and 6, can improve interference immunity considerably. The two parameter correlation function was demonstrated with a sequence of random rectangular pulses as a simple example. Obviously such plots could be generated for any cyclostationary signal, but more useful representations of cyclostationarity are given by the cyclic autocorrelation and spectral correlation density functions described below.

2.2 Frequency Domain Representation of Cyclostationarity

To fully appreciate the significance of the periodicity of the autocorrelation of a cyclostationary signal, it is useful to interpret this phenomenon in the frequency domain. The frequency domain manifestation of the periodically time varying statistics of a cyclostationary signal is *spectral correlation*, that is, the different frequency components of the signal are correlated with each other, or, put another way, the signal is correlated with itself shifted in frequency by certain particular values.

2.2.1 Spectral correlation

There are two fundamentally different types of spectral correlation which are commonly encountered, which can be loosely described as *carrier frequency related* and *baud rate related.* In this context, chip rate related cyclostationarity in direct sequence spread spectrum signals is similar to baud rate related cyclostationarity. These types of correlation are described in this section. Their existence is explained as a result of the properties of the Fourier transform (carrier related correlation) and the spectral properties of sampled signals (baud rate related correlation).

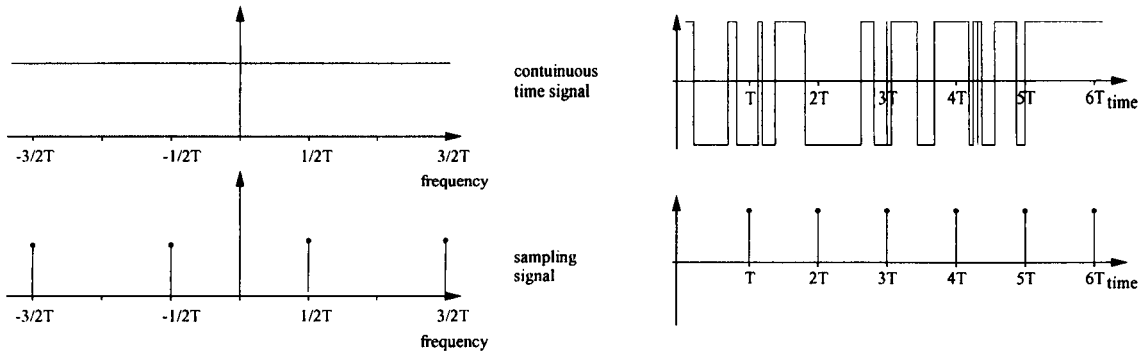

Figure 2.9: Continuous time and sampling signals; frequency and time domain

2.2.2 Baud rate related correlation and the sampling theorem

From [54] we know that QAM signals exhibit cyclostationarity related to the baud rate. Cyclic frequencies associated with the carrier frequency are ignored for the moment. A description follows of how the baud rate related correlation can be seen to be a direct result of the properties of sampled signals [94].

Let $x_c(t)$ be a continuous time signal which changes instantaneously, at random times, between the values $(+1, -1)$. It has equal probability of taking each these two values. The signal changes in such that at every instant, it is uncorrelated with every other instant. This, purely real, signal is chosen because it is simple to represent graphically. The argument which follows is equally valid for any complex or real constellation.

With no minimum time between changes in value, this signal has a (standard, single parameter) autocorrelation which is an impulse, and therefore a white spectrum, as shown in figure 2.9. Note that the time domain representation of this signal is not actually possible to draw, but the signal sketched has similar properties.

A sequence of impulses $x_s(t)$ carrying data values of $(+1, -1)$ can be constructed by multiplying $x_c(t)$ by a sampling function $s(t)$ which is a sequence of impulses:

$$s(t) = \sum_{n=-\infty}^{\infty} \delta(t - nT) \tag{2.23}$$

so

$$x_s(t) = \sum_{n=-\infty}^{\infty} x_c(t)\delta(t - nT) \tag{2.24}$$

$s(t)$ is shown in figure 2.9, and $x_s(t)$ in figure 2.10. Let $X_s(f)$, $X_c(f)$ and $S(f)$ be the Fourier transforms of $x_s(t)$, $x_c(t)$ and $s(t)$ respectively. $X_s(f)$ is therefore given by the convolution of $X_c(f)$ and $S(f)$. It is well known that the Fourier transform of an infinite sequence of time domain impulses is an infinite sequence of frequency domain impulses, that is:

$$S(f) = \frac{2\pi}{T} \sum_{k=-\infty}^{\infty} \delta(f - \frac{k}{T}) \tag{2.25}$$

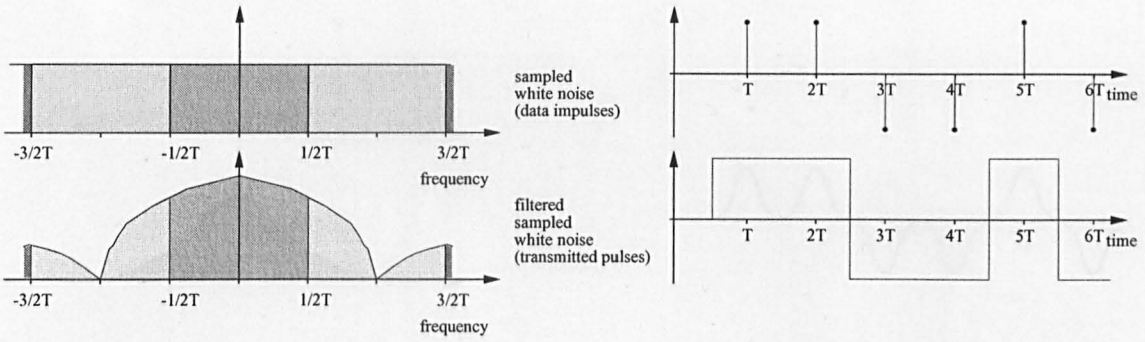

Figure 2.10: Impulses and rectangular pulses; frequency and time domain

so

$$X_s(f) = \frac{1}{T} \sum_{k=-\infty}^{\infty} X_c(f - \frac{k}{T}) \tag{2.26}$$

where k is an integer and T is the symbol period.

Equation 2.26 shows that the spectrum $X_s(f)$ is invariant under frequency shifts of $\frac{k}{T}$ and made up of aliased copies of the spectrum $X_c(f)$. This is represented in figure 2.10 by the different shadings in the top spectrum. Each rectangle contains the same spectrum - it repeats at intervals of $\frac{1}{T}$. These "images" are created by the sampling of the original continuous signal.

So the spectrum $X_s(f)$ exhibits spectral correlation as a direct result of its regular symbol period.

An impulse sequence has a spectrum of infinite bandwidth. Figures 2.10 and 2.11 illustrate the effects on the spectral correlation of filtering an impulsive signal to provide a more practical waveform. Figure 2.10 illustrates filtering which gives a rectangular pulse shape and a $sinc^2$ power spectrum which is also of infinite bandwidth. Figure 2.11 illustrates raised cosine filtering, which is commonly used in communications. The spectrum has a raised cosine shape, with the total bandwidth being controlled by the roll-off parameter. The diagram illustrates a roll-of of 1 which gives 100% excess bandwidth relative to the Nyquist bandwidth. It is clear that for the information within the Nyquist bandwidth ($-\frac{1}{2T}$ to $\frac{1}{2T}$) to be repeated completely, a similar extra bandwidth is required. So 100% excess bandwidth is the minimum required to duplicate all the information in the signal.

All three signals described above (the impulse sequence, rectangular pulse sequence and raised cosine filtered pulses) have spectral correlation and are therefore cyclostationary. If a signal's bandwidth is no greater than the Nyquist limit, then there is no spectral correlation, and therefore no cyclostationarity; such a signal is stationary.

From this also follows the important point that if a signal is cyclostationary, to maintain its spectral correlation, a sufficiently high sampling frequency must be used. A signal sampled at one sample per symbol cannot exhibit baud rate related spectral correlation so any processing which exploits cyclostationarity must be performed at a higher sampling rate.

Another example of the spectral correlation of a QAM signal with rectangular pulses is shown in figure 2.12. A wider bandwidth is shown than in figure 2.10. The signal components in the region labelled A are correlated with those in B' after a frequency shift of 1 Hz, and with those in C' after a frequency shift of 2 Hz, and so on. This is exploited by the FRESH filtering described in

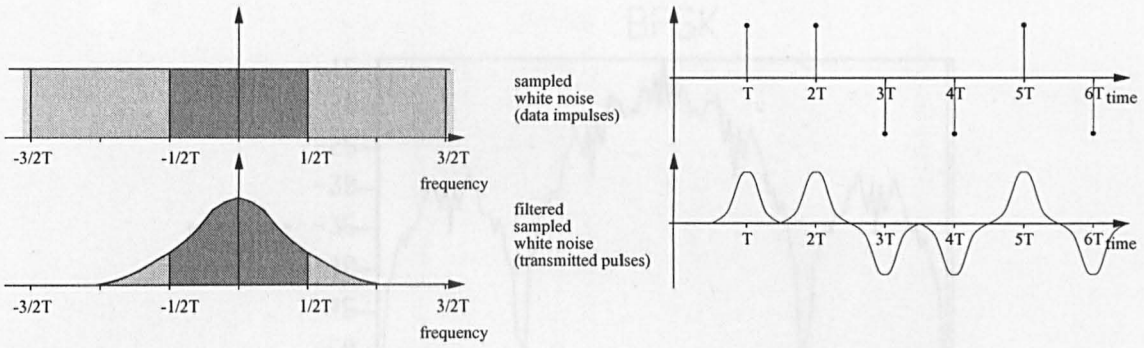

Figure 2.11: Impulses and 100% excess bandwidth raised cosine pulses; frequency and time domain

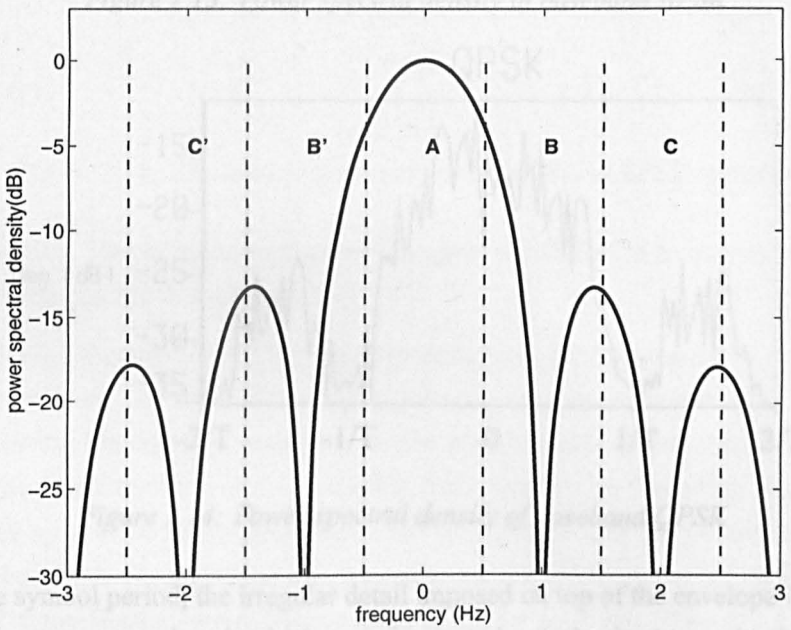

Figure 2.12: Spectral correlation of a rectangular pulse

chapter 3. Any other signals present (including thermal noise) which are not cyclostationary with the same cyclic frequencies, will not be correlated under these frequency shifts.

2.2.3 Carrier frequency related correlation

While a digitally modulated carrier is cyclostationary with cyclic frequencies related to twice the carrier frequency, we can, in general, model all the properties of a modulated signal at complex baseband, that is, at zero carrier frequency. It is shown here that the carrier related correlation is equivalent to reflective symmetry about the carrier frequency (or zero frequency at complex baseband).

The Fourier transform of a purely real signal is symmetric about zero frequency while a complex signal shows no such symmetry. In figure 2.13 and figure 2.14 we see the power spectral density of BPSK (which can be expressed as a purely real signal) and QPSK (which is complex) pseudo random baseband data sequences, 250 symbols long. The pulse shape of both these signals is rectangular, so the shape of the spectrum envelope is $sinc^2(x)$ which has a relatively high level of side lobes making the spectral correlation easier to see. Here we show the spectrum up to $\pm 2/T$

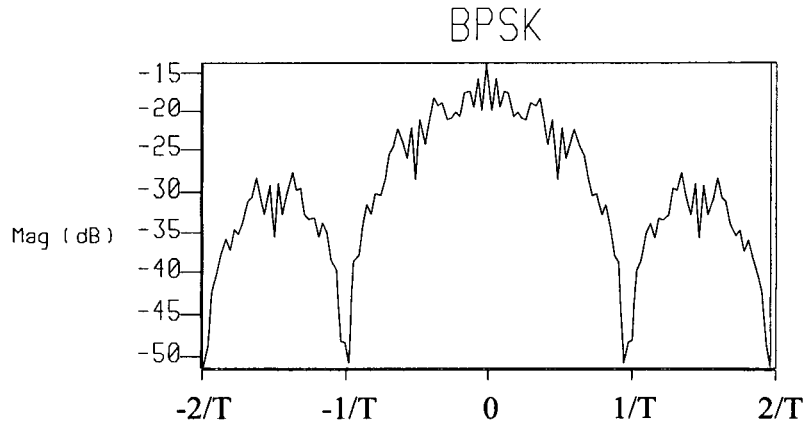

Figure 2.13: Power spectral density of baseband BPSK

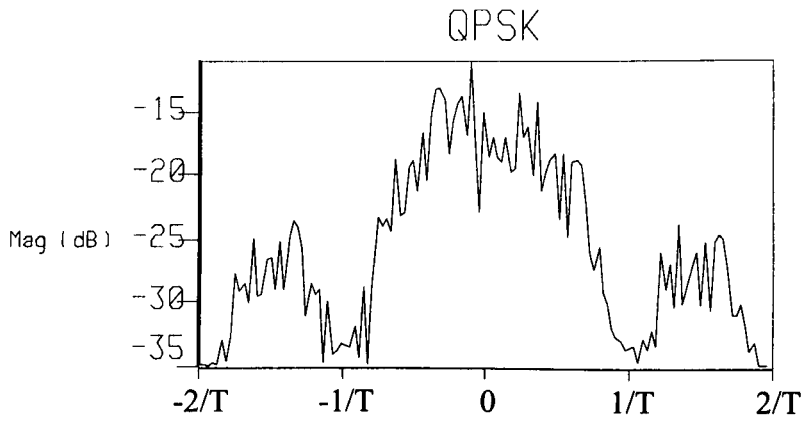

Figure 2.14: Power spectral density of baseband QPSK

where T is the symbol period; the irregular detail imposed on top of the envelope is defined by the particular data sequence analysed, which was chosen to be fairly short to make this fine structure more obvious. We can see that in the QPSK signal the irregular pattern repeats every $1/T$, that is, a translation of the spectrum of an integer multiple of $1/T$ would result in correlated signal components being overlaid. Effectively in the frequency range shown we have 4 copies of the data present.

However there is even greater redundancy in the BPSK spectrum, as there is also reflective symmetry present about the zero frequency axis. This coupled with the baud rate spectral correlation means that there are 8 copies of the data present in the frequency range shown (for example: in frequencies 0 to 0.5, 0.5 to 1.0, 1.0 to 1.5, 1.5 to 2.0, and the negative images of these 4 ranges).

In practical situations a more bandwidth efficient spectrum is normally used, so figure 2.15 shows a schematic representation of the correlation in a 100% excess bandwidth raised cosine filtered BPSK signal. Here we have labelled 4 regions A, B, C and D. A is correlated with C under a frequency shift of $1/T$, as is B with D. B and C have reflective symmetry about the zero frequency axis, as do A and D. So we have the same information appearing 4 times in this restricted frequency range.

If this were a QPSK spectrum the symmetry in the fine structure would disappear, leaving correlation only between A and C, and between B and D.

In chapter 3 a filter structure is described which includes frequency shifts to exploit the baud rate

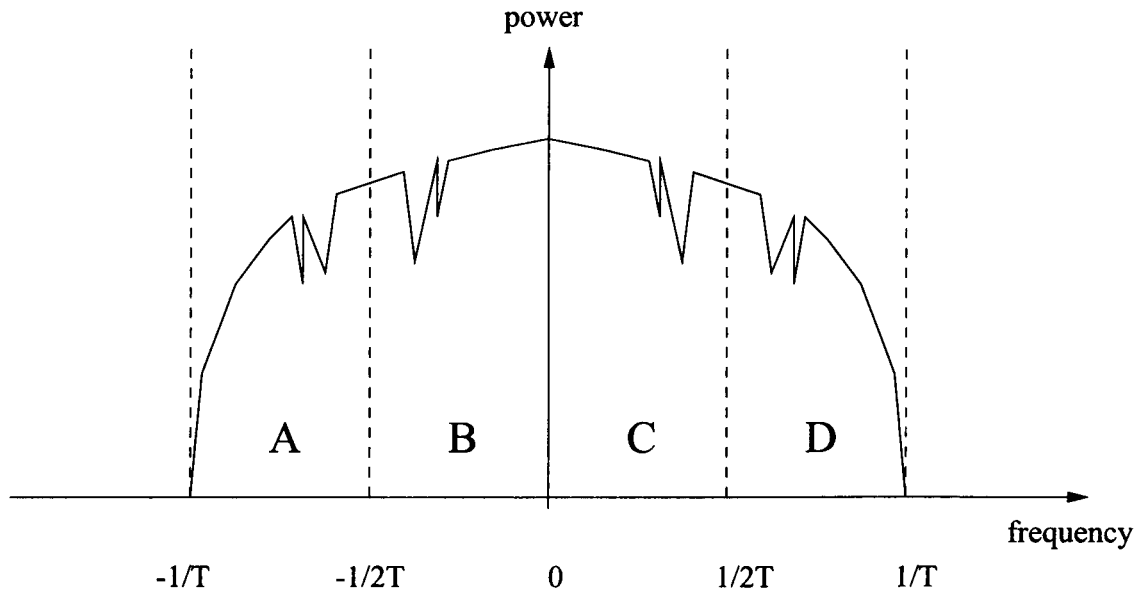

Figure 2.15: Schematic representation of baseband BPSK spectrum

spectral correlation. Two or more versions of the signal are shifted in frequency so that correlated frequency components are overlaid.

However exploiting the *symmetry* of the spectrum of *baseband* BPSK cannot be done by frequency shifting, as there is no one frequency spacing between correlated components. It can be done, however, by taking the complex conjugate of the signal, which has the effect of reflecting its spectrum in the zero frequency axis. To exploit the symmetry by frequency shifting, the signal must be moved up to a carrier frequency and then a frequency shift can be applied to overlay the correlated components. Figure 2.16 shows the spectrum of a 100% excess bandwidth raised cosine filtered BPSK signal with a carrier frequency of twice the baud rate. The low carrier frequency was chosen to allow display of both positive and negative frequencies of the signal together. It is clear now that a frequency shift of twice the carrier frequency will move the correlated frequency components together, by moving the negative and positive frequency bands together.

As before, a schematic version of a BPSK bandpass spectrum is given in figure 2.17. The symmetry of the spectrum about the carrier frequency means that AB is correlated with CD under reflection about f_c. The symmetry of the whole spectrum (which follows because the bandpass signal is purely real) means that all of $ABCD$ is correlated under reflection with $A'B'C'D'$. Therefore shifting by $2f_c$ overlays $D'C'$ and AB; and $A'B'$ and CD.

This has overlaid the correlated components arising due to the symmetry of the equivalent baseband spectrum. This could not be done with a non-symmetric spectrum such as QPSK. There is also correlation under additional shifts of $\pm 1/T$.

If there is stationary noise added to the signal then it will be symmetric about the y-axis, but will not be correlated under the $2f_c$ frequency shift, except at the carrier frequency itself.

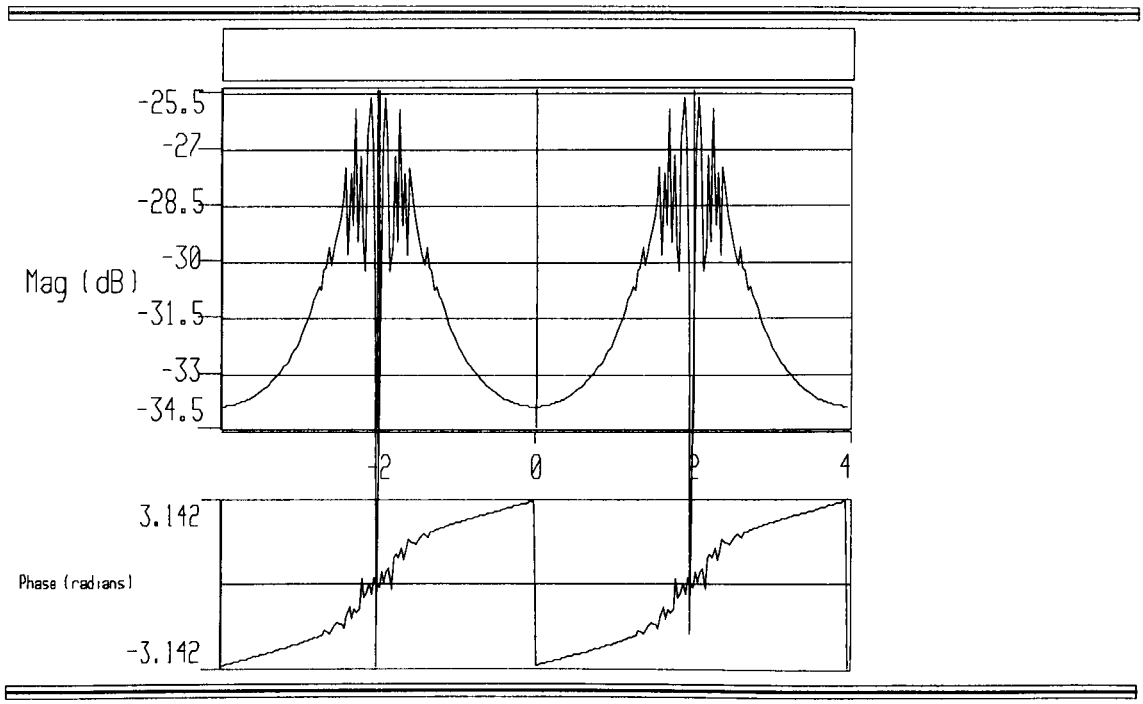

Figure 2.16: Power spectral density of bandpass BPSK signal

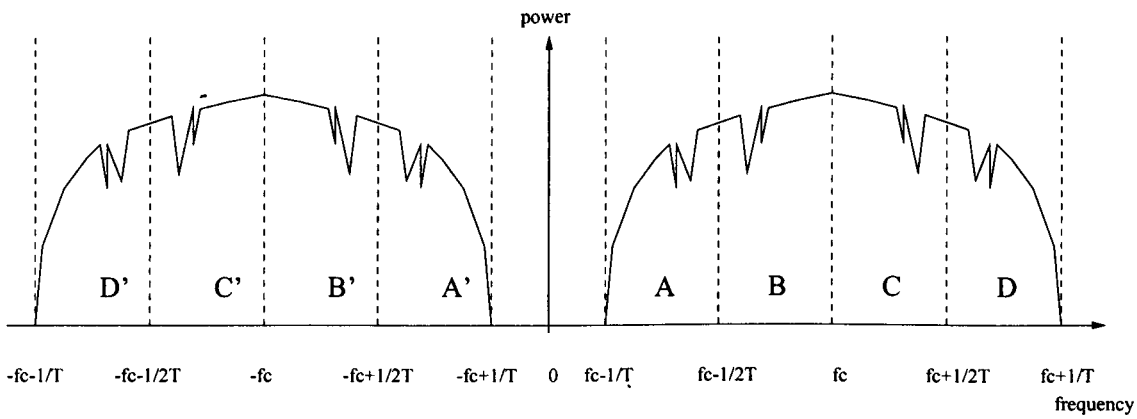

Figure 2.17: Schematic spectrum of BPSK modulated carrier at f_c

2.2.4 Spectral correlation density function

A useful way of displaying the cyclostationary properties of a signal is to plot the *spectral correlation density* (SCD). This function is defined as the Fourier transform of the cyclic autocorrelation function, as follows:

$$S_x^\alpha(f) = \int_{-\infty}^{\infty} R_x^\alpha(\tau) e^{j2\pi f\tau} d\tau \tag{2.27}$$

The conjugate SCD is similarly defined:

$$S_{xx^*}^\alpha(f) = \int_{-\infty}^{\infty} R_{xx^*}^\alpha(\tau) e^{j2\pi f\tau} d\tau \tag{2.28}$$

The cyclic autocorrelation function and spectral correlation density Fourier transform pair for cyclostationary signals are analogous to the autocorrelation function and power spectral density pair for stationary signals.

When $\alpha = 0$ the SCD is equal to the power spectral density of the signal. At other values of α the SCD is the cross-spectral density of the signal, and the signal shifted in frequency by α.

An important feature of second order cyclostationary statistics is that they contain phase information about the signal. This is particularly useful for methods of blind channel identification or equalisation [70, 64, 63, 65]. In contrast, second order stationary statistics contain no phase information. The phase information is contained in the mathematical representations of the second order cyclostationary statistics, the cyclic autocorrelation function, conjugate cyclic autocorrelation function, the SCD and the conjugate SCD. However for simplicity in what follows we plot only the magnitude of the complex SCD and conjugate SCD functions, as this is sufficient to illustrate the form of the correlation exploited in the filters of the next chapter.

Cyclostationarity, cyclic frequencies and the subsequent spectral correlation are likely to arise from any periodic process contributing to the signal generation or distortion. The most common periodicities are related to the carrier frequency, the baud rate, the chip or hopping rate (for spread spectrum signals). Here we are concerned with carrier frequency and baud rate related effects.

The graphs in this section represent the correlation of signals calculated by computer simulation. This was done in MATLAB, by generating a signal of the appropriate modulation type and pulse shape, with random data, and performing a cross-correlation of this with frequency-shifted versions of itself. Using finite simulated signals in this way, as opposed to an analytical ensemble average approach, means that noise is apparent in the plots as well as the main cyclostationary features. The length of the signal used was chosen in each case to be sufficient to reduce this noise to a level where it is visible, but does not obscure the features of interest. Mathematical derivation of these functions is given in [54].

Sine wave

As a simple example, the spectral correlation density of a sine wave with frequency $f = 3.5$ is shown in figure 2.18. The spectral correlation density is displayed for a range of different values

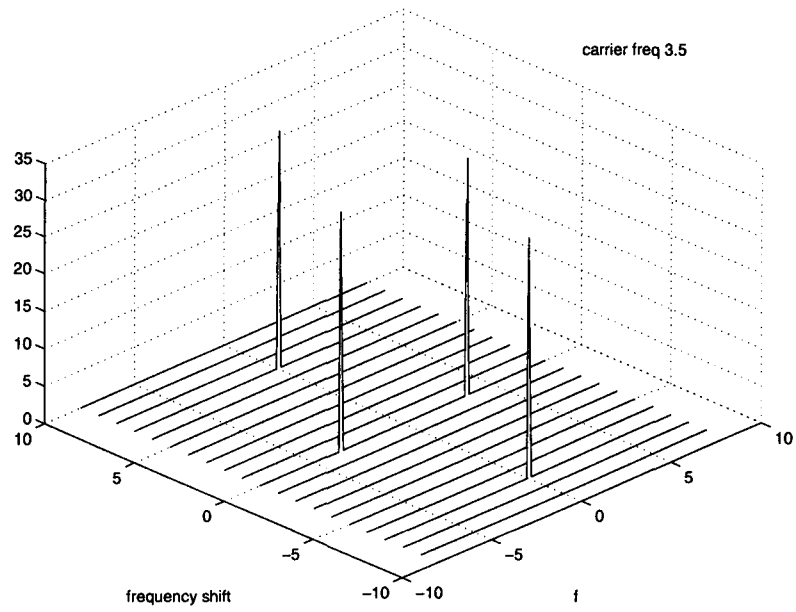

Figure 2.18: SCD of a sine wave

of the frequency shift α. At a frequency shift of 0 we have the normal power spectral density of a sine wave which consists of two impulses, one at f and the other at $-f$.

The other two peaks correspond to cross correlations after frequency shifting the $sin(2\pi ft)$ signal so that impulses originally at $\pm f$ are at the same frequency. To generate one, the original signal is shifted by $-f$ and f (so the relative frequency shift is $2f$), and cross-correlated. The other peak's existence is obvious from symmetry arguments. This correlation is a trivial example of the carrier frequency related correlation which also applies to some modulated carrier signals.

BPSK

Another example of a spectral correlation density plot is shown in figure 2.19. The signal used for this was rectangular pulse BPSK, carrying independent, identically distributed (IID) data, with a baud rate of 1 and a carrier frequency of 3.5.

As before, with zero frequency shift, this spectral correlation density is equivalent to the standard power spectral density function. In figure 2.19 the $\alpha = 0$ line shows the $sinc^2(x)$ shape that is expected from data with a rectangular pulse shape. For $\alpha \neq 0$ the graph function represents the cross spectral density of the signal with itself shifted in frequency by α. It is clear that correlation exists at frequency shifts of $\pm 1, \pm 2, \pm 3...$ (i.e. integer multiples of the baud rate) and that a greater correlation exists at ± 7 (twice the carrier frequency), and that similar correlation to that at $\pm 1, \pm 2, \pm 3...$ is present at $\pm 7 \pm 1, \pm 7 \pm 2, \pm 7 \pm 3...$. For each value of the frequency shift parameter α the line in figure 2.19 is the cross correlation between the original BPSK signal, and that signal shifted in frequency by α.

The conjugate SCD for this, and all other purely real signals, is identical to the SCD, as a real signal is equal to its complex conjugate.

The SCD plot for the equivalent baseband BPSK signal is as shown in figure 2.20. It is clear that the structure of 2.19 is a combination of that in figures 2.18 (representing the carrier related

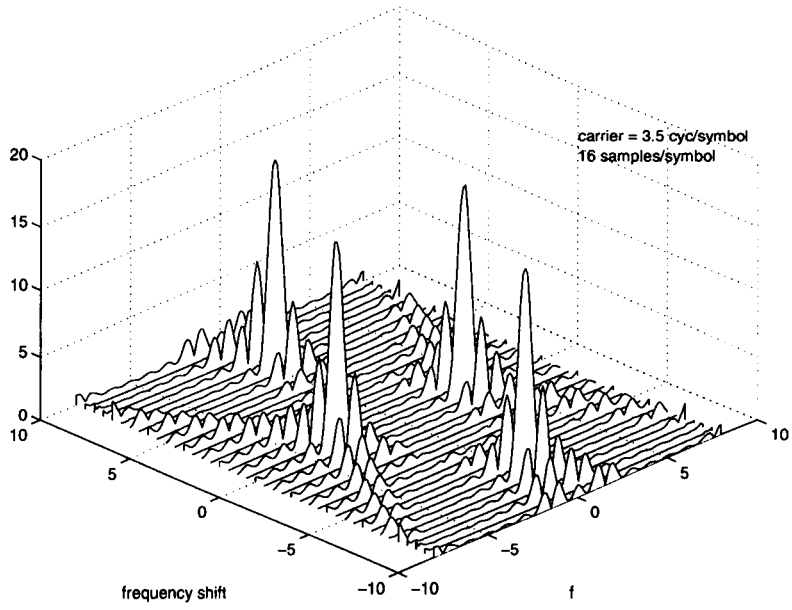

Figure 2.19: SCD of a BPSK modulated carrier

correlation) and 2.20 (representing baud rate related correlation). Again the conjugate SCD of baseband BPSK is identical to the SCD as the signal is purely real.

Notice that correlation occurs at frequency shifts of specific, discrete values; the SCD is not a continuous function.

AM

The SCD of an amplitude modulated (AM) carrier (figure 2.21) is similar to that of the sine wave but the impulses are replaced by the spectral shape of the modulating signal. Again, the carrier frequency is $f = 3.5$ and correlation exists under frequency shifts of $2f$, or ± 7. Obviously the only correlation present in the AM signal or the sine wave is carrier frequency related; there is no keying frequency in AM as in the digital signals which follow.

QPSK

It is interesting to compare the SCD of BPSK (figure 2.19) with that of QPSK (figure 2.22). This graph represents the SCD of QPSK with a rectangular pulse shape, IID data, a baud rate of 1 and a carrier frequency of 3.5. It is clear that there is no correlation associated with the carrier frequency in QPSK, whereas the baud rate related correlation is similar to that of BPSK.

Baud rate related correlation is dependent on the excess bandwidth of the signal relative to the Nyquist limit. As an example of this, figure 2.23 shows the SCD of QPSK with a raised cosine spectrum. Note that this graph has a different scale on the "frequency shift" axis from figure 2.22. The excess bandwidth of the signal was 100%, which limits any spectral correlation to a frequency shift of 1 times the baud rate. Any larger frequency shifts result in the two correlated signals having no overlap in the frequency domain. Whereas a rectangular pulse shape (as used in figure 2.22) results in a $sinc^2(x)$, infinite bandwidth spectrum.

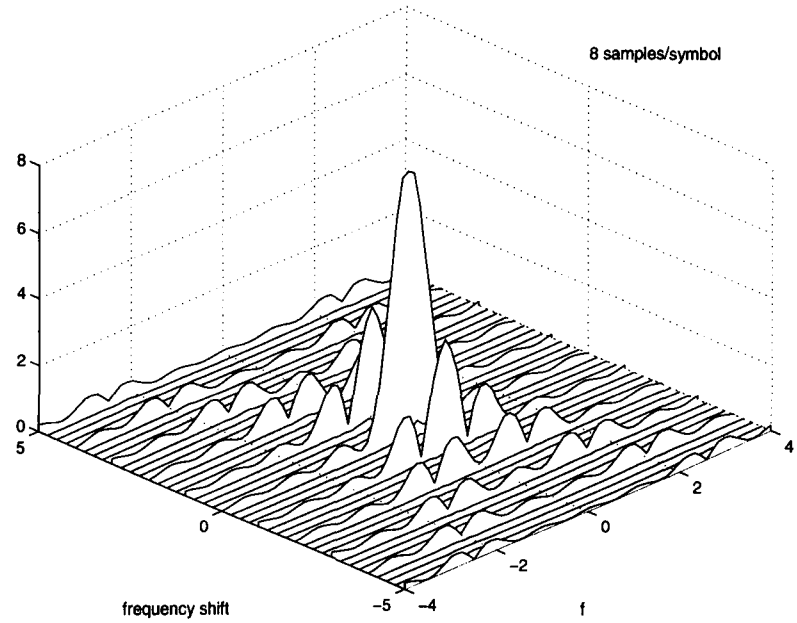

Figure 2.20: SCD of a baseband BPSK signal

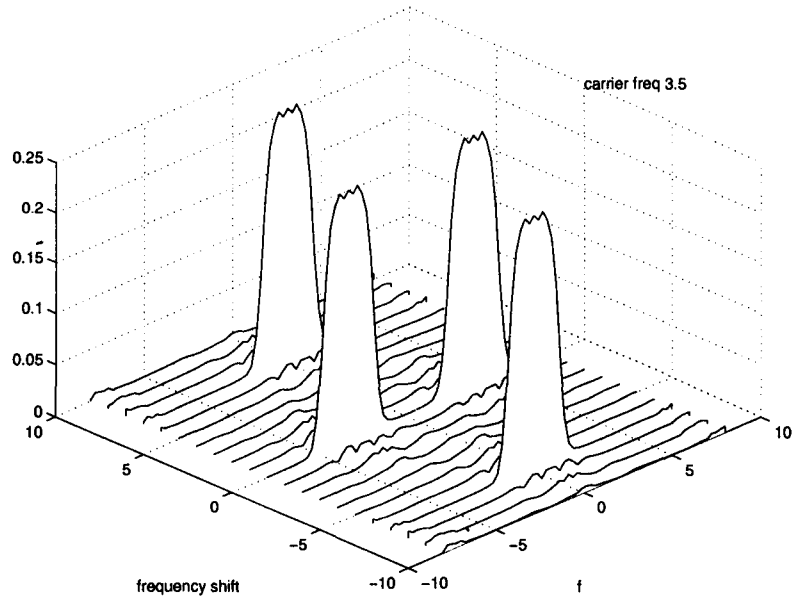

Figure 2.21: SCD of an amplitude modulated carrier

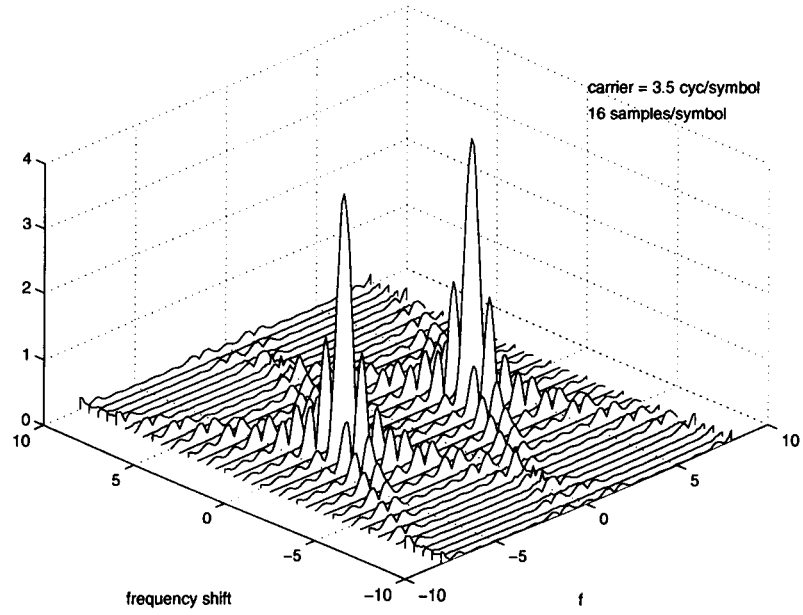

Figure 2.22: SCD of QPSK modulated carrier

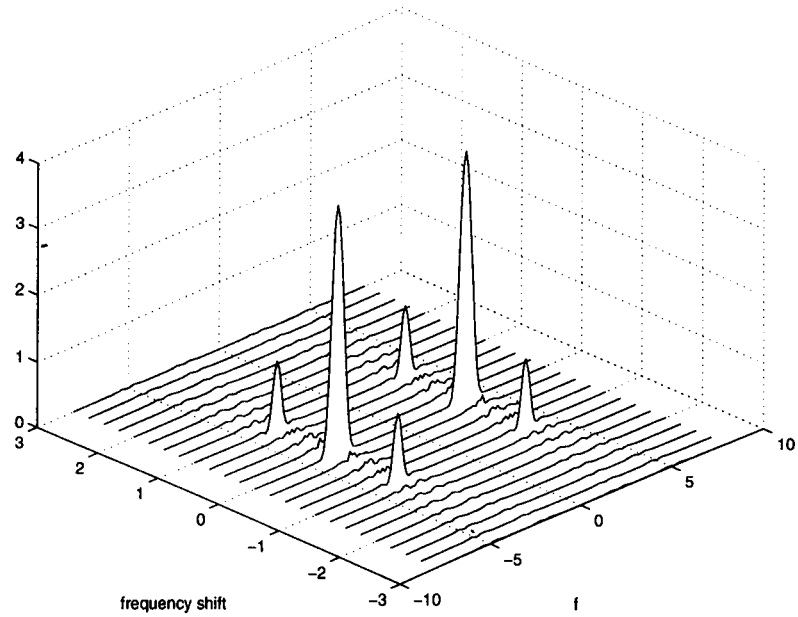

Figure 2.23: SCD of raised cosine filtered 100% excess bandwidth QPSK modulated carrier

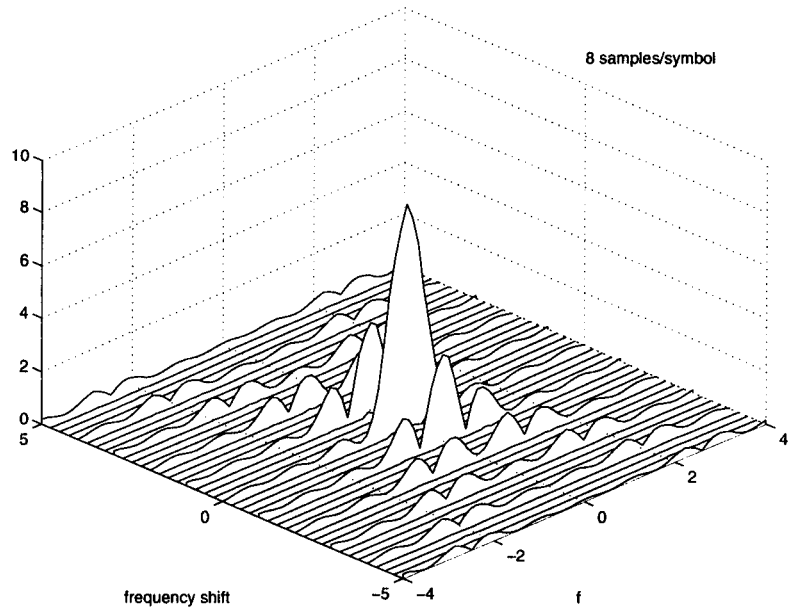

Figure 2.24: SCD of baseband QPSK signal

The baseband QPSK signal (figure 2.24) has similar structure to the BPSK example, as this is dependent on the pulse shape, which was again chosen to be rectangular. However in this case the conjugate SCD is different, and is identically zero for all values of f and α. This is plotted in figure 2.25. The small non-zero values in the centre of the graph can be attributed to the finite signal length used in the simulation. This graph shows that a QPSK signal is not correlated with its complex conjugate under any frequency shift, or equivalently that a passband QPSK signal will have no carrier related correlation.

MSK

Minimum Shift Keying (MSK) results in a signal with different spectral correlation properties again (see figure 2.26). Here there is correlation related to the baud rate, but it occurs at $\pm 2/T, \pm 4/T...$, whereas the carrier and baud rate related correlation is at $2f_c \pm 1/T, 2f_c \pm 3/T, ...$ There is no correlation at $2f_c$.

The baseband MSK SCD and conjugate SCD are shown in figures 2.27 and 2.28. In this case we cannot identify the SCD of the baseband signal with baud rate related correlation and the conjugate SCD with carrier related correlation, as before, because the conjugate correlation depends on both the carrier frequency and the baud rate.

Other signals

The SCD functions of a number of other signals are given in [54]. In general one can say that other QAM or PSK signals (for example, 16-QAM or 8-PSK) have a similar SCD structure to that of QPSK with the same pulse shape.

These different patterns of spectral correlation are important when designing filters which can exploit such correlation for interference removal. They indicate at what frequency shifts correlation

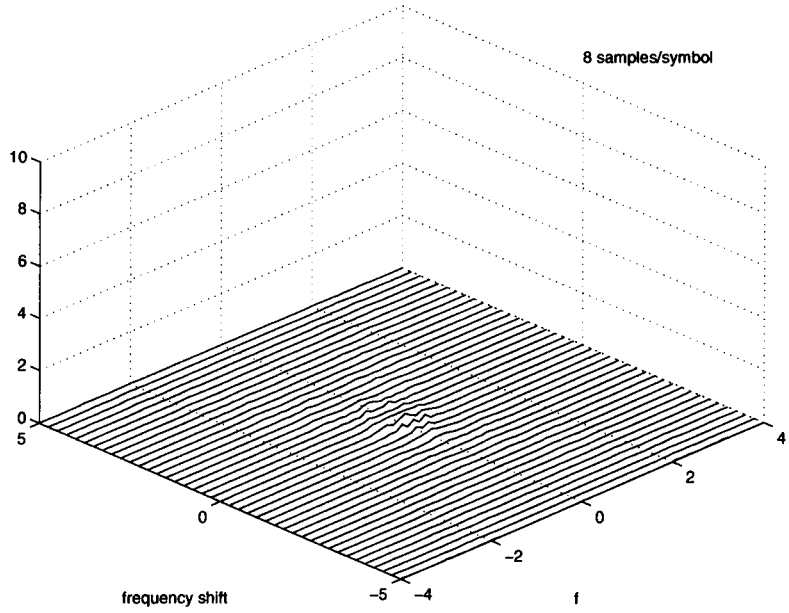

Figure 2.25: Conjugate SCD of baseband QPSK signal

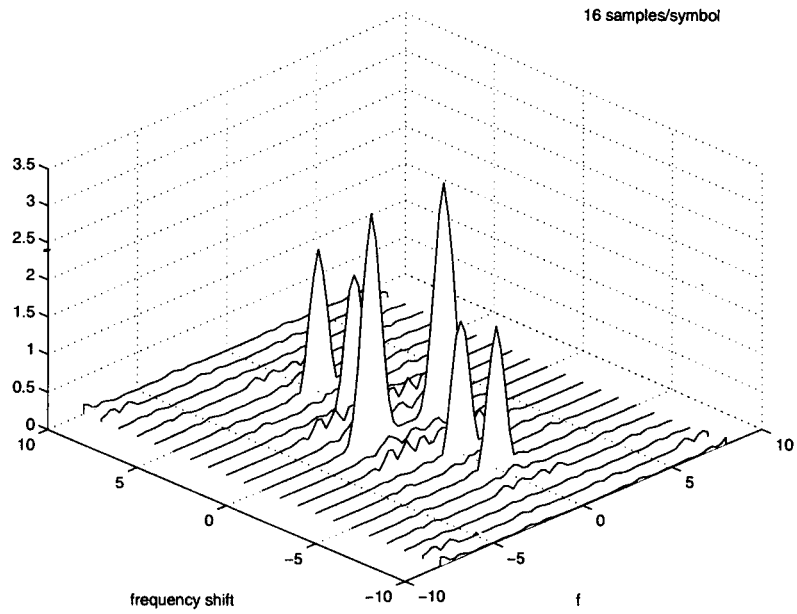

Figure 2.26: SCD of MSK modulated carrier

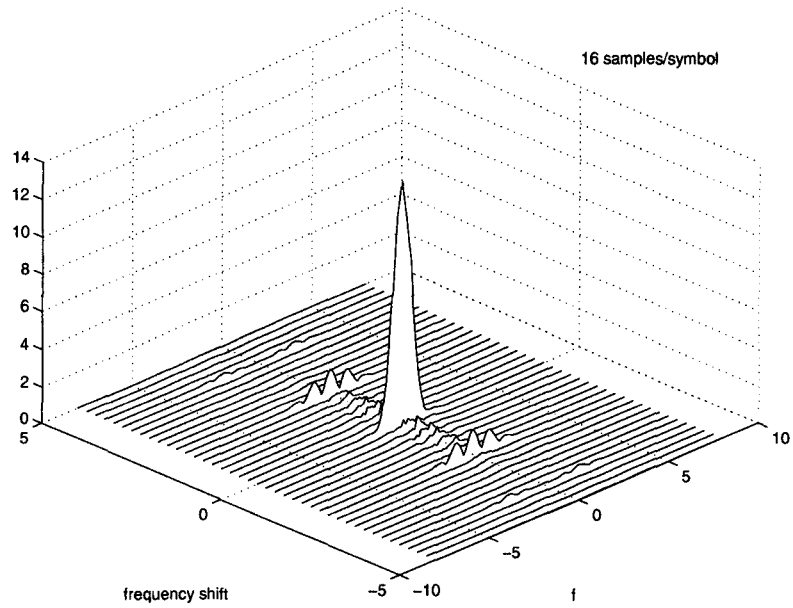

Figure 2.27: SCD of baseband MSK signal

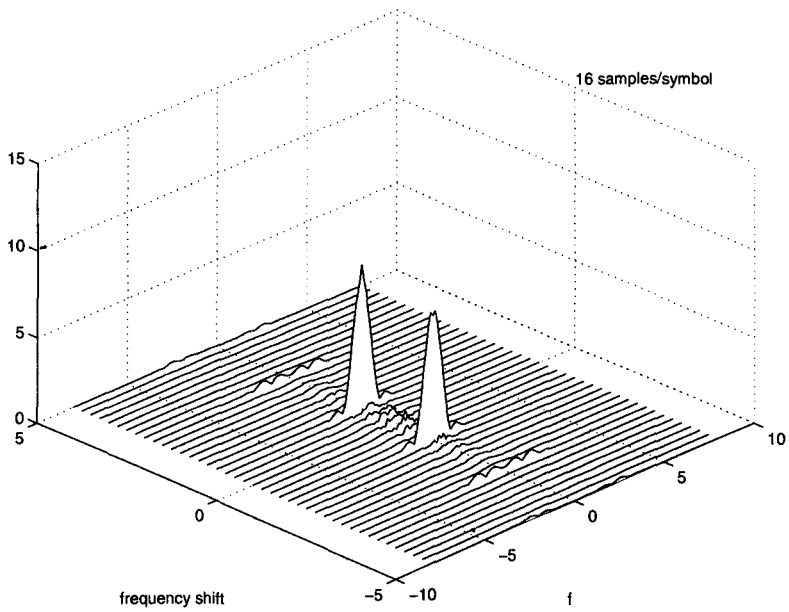

Figure 2.28: Conjugate SCD of baseband MSK signal

occurs, and the frequency ranges of this correlation. This determines the frequency shifts which are useful in the FRESH filter, and the bandwidth required in the filter following each frequency shift.

In particular, chapter 3 is concerned with removing QPSK interference from a QPSK wanted signal by exploiting the spectral correlation shown above. The effect of shaping the spectrum in enhancing the correlation is examined in chapter 5 and in chapter 6, a realistic VLF interference scenario is improved by exploiting the correlation inherent in MSK modulation. In chapter 7, there is an examination of the correlation in BPSK signals in the context of blindly adapting interference rejection filters.

It is also clear from the SCD plots shown above that they can form a useful "signature" of modulation types and pulse shapes which can be used for signal identification [71, 72, 73].

Chapter 3

Frequency Shift Filtering

Contents

Frequency shift filtering, or *FRESH* filtering is a technique which uses the spectral correlation of cyclostationary signals to reduce the effect of noise and interference on the signal of interest. The filter structure which achieves this, and its role in a communications receiver, is described in this chapter. Initially only the exploitation of the signal of interest (SOI) correlation is considered (section 3.5) but in section 3.7 we extend this to making use of the correlation of a cyclostationary interferer.

However, first some background material is presented to put FRESH filtering into a communications context. Section 3.3 describes the optimum filter for extracting a stationary signal from stationary noise: the Wiener filter as a background to the filters used in communication receivers. Some basic techniques in receiver filtering are first described to show what role a FRESH filter can take in a communications receiver.

Generally, as the FRESH filter exploits the cyclic frequencies of the SOI and/or interference, the more different cyclic frequencies the signals have, the more effective the filter will be. This suggests good performance is possible when the signals have different baud rates, different carrier frequencies, and different modulation schemes.

3.1 The Wiener Filter and Matched Filtering

In a digital communications receiver, the aim is to decode the data sequence as accurately as possible, which requires detecting, or making a decision on each transmitted symbol. Usually some form of *matched filter* [79] is used to help this process - this is a filter (described below) which maximises the signal to noise ratio at the ideal symbol sampling points, with the expectation that symbol rate sampling will occur afterwards.

This is a different process from *Wiener filtering*, in which the aim is to extract the entire waveform from noise. In a sampled system, this means minimising the MSE of all the samples, not just symbol spaced samples.

In what follows, it is shown how a Wiener filter, or the cyclic Wiener filter can be incorporated into a communications system, where the aim is to determine the symbol values.

For simplicity in analysis and implementation, particularly with adaptive systems, a finite impulse response (FIR) filter structure will often be used ([94], section 6.5).

3.2 Processing in a Receiver

The reception of a digitally modulated signal requires suitable processing to allow, ultimately, sampling of the signal at the symbol rate. These symbol rate samples can then be decoded to recover the transmitted data. There are a number of problems to be overcome to allow this process to work effectively. The three most significant of these can be summarised as the effects of noise and interference, intersymbol interference (ISI) and synchronisation errors. In this thesis the primary interest is in interference rejection, but the subject of equalisation (removing ISI) is closely related.

Synchronisation

"Synchronisation" describes a range of techniques for determining the exact frequency, and perhaps phase, of the carrier frequency oscillator in the receiver, and of the correct sample times for the symbol rate sampling. It can also cover code synchronisation, or determining which data symbols correspond to separate code words, in a system which includes some form of error correction

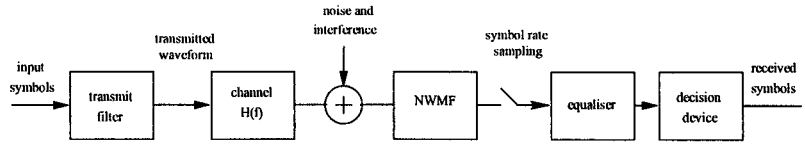

Figure 3.1: Communications system

or detection coding. In this thesis the issue of synchronisation is not considered further. In all simulations of communication systems presented here, perfect carrier frequency and phase, and perfect symbol timing recovery is assumed. It is however worth noting that all techniques for synchronisation are degraded by the effects of noise and ISI, and so filtering which can mitigate these effects can be useful from a synchronisation point of view. Some work has also been done on the more direct application of cyclostationarity to synchronisation [10, 74].

The signal values at these sampling points are then used in a decision device to identify the received symbols.

Additive noise and the matched filter

The term "noise" usually includes the effects of all corrupting signals which combine additively with the wanted signal. The principal sources are atmospheric noise from the channel, interference from other man made signals and noise from the circuitry in the receiver. The optimum filter for detecting data symbols in white noise is the MF. There can be confusion as to whether this term applies to the optimum filter for detecting symbols in non-white noise. In this thesis, the MF is the filter matched in white noise. The effects of non-white additive noise are be minimised by the *noise whitening matched filter* (NWMF) whose response depends on the transmitted waveform and the spectrum of the noise present. The NWMF is defined to be the filter which minimises the SNR at the output of the filter at the ideal symbol spaced sampling instants; it is assumed that such a filter is followed by symbol rate sampling.

The position of the NWMF in a typical communication system structure is shown in figure 3.1 (this figure does not show synchronisation operations). The NWMF could be implemented as a continuous time or a sampled data filter. Should a sampled data implementation be used it would be preceded by a (continuous time) anti-aliasing filter, and sampling at a rate sufficiently high to support the bandwidth of the wanted signal. In figure 3.1 it is implicit that a continuous time filter is used.

The term "noise whitening" describes the best that can be done by a linear time invariant filter against additive interference. The filter is a cascade of two processes: a noise whitening filter which has an output made up of a a distorted version of the SOI added to white noise, and then a MF for the waveform created by the action of the noise whitening filter on the transmitted waveform ([95] section 7.2).

In most communications systems the noise and interference will not be precisely known, or it will be variable in time. This means that the best receive filter to use may be adaptive. The adaptive filter would adapt towards a NWMF.

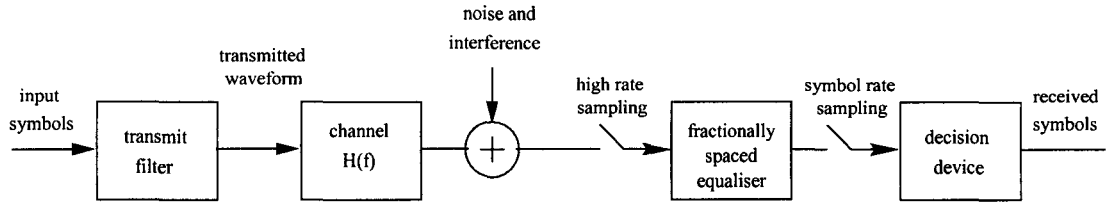

Figure 3.2: Communications system using a FSE

Intersymbol interference and the symbol spaced equaliser

ISI occurs when the channel is bandlimited or when there are multiple propagation paths through the channel with different delays.

A common assumption made in modelling communications channels is that they introduce ISI and AWGN, but no non-white noise from other communications signals. In this case the optimum receiver is a MF followed by baud rate sampling and then a symbol spaced equaliser (SSE). The SSE is a transversal linear filter operating at a sampling rate equal to the symbol, or baud, rate with a response which is approximately the inverse of the channel response. Theoretically this will often require an infinite impulse response but for most channels a good approximation can be achieved with a reasonable number of filter taps.

This filter combination could be implemented as an analogue MF with the sampling combined with analogue to digital conversion, and then a sampled, digital SSE, or the MF could also be a sampled digital filter having been preceded by an anti-aliasing filter and sampling above the Nyquist rate for the SOI bandwidth. Should the channel add noise which is not white (e.g. interference) then one simply replaces the MF with a NWMF in the structure above.

Referring to figure 3.1, the equaliser removes ISI from the symbol rate samples after the received signal has been filtered with a NWMF. Usually an adaptive equaliser is used, if the channel is time-varying, as this allows the receiver to automatically compensate for changes in the channel multipath characteristics.

This receiver structure is useful in a system where the matched filtering would be carried out on the analogue received signal, or where it must be, for some other reason, a separate entity from the adaptive digital filter which is the adaptive equaliser. Otherwise a FSE (section 3.2.1) is usually preferred, due to its ability to correct symbol timing errors [34].

The equaliser is constrained in that it can only produce symbol rate samples which are weighted sums of other symbol rate samples. Equivalently the bandwidth of the equaliser is $\pm\frac{1}{2T}$. It is not hard to believe that in a system receiving a signal with bandwidth $\pm\frac{1}{T}$ an adaptive filter with bandwidth $\pm\frac{1}{T}$ would perform better than a fixed filter with bandwidth $\pm\frac{1}{T}$ followed by an adaptive filter of bandwidth $\pm\frac{1}{2T}$. This improvement over the matched filter/equaliser combination can be realised as an FSE, described in section 3.2.1.

3.2.1 Adaptive fractionally spaced equaliser

If all the receive filtering is being done in the sampled domain, then it is possible to combine the adaptive equaliser and MF or NWMF into one digital adaptive filter.

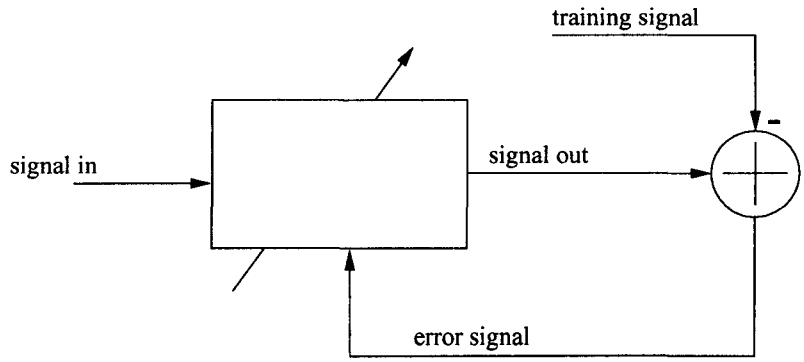

Figure 3.3: Adaptive filter structure

If the channel conditions are known and fixed then this will not improve performance over the separate MF and equaliser combination. If, however, the channel is varying then it is important to be able to vary the filtering appropriately and this is much easier with digital than analogue filters. If the ISI varies but the additive noise is constant, then a fixed MF or NWMF followed by an adaptive SSE is sufficient. If however the noise also varies then optimum performance would require the NWMF to also be adaptive. In this case it is more efficient to combine the NWMF and SSE into one adaptive filter. The NWMF must have a bandwidth at least as wide as the SOI so it must be operating at higher sampling frequency than the SSE, or equivalently, it has a tap spacing which is a fraction of a symbol period, hence the term fractionally spaced equaliser.

The adaptive symbol spaced equaliser cannot act as an adaptive matched filter on its own because it has a bandwidth of $\pm\frac{1}{2T}$, where T is the symbol duration. This bandwidth is too narrow to accurately manipulate all realisable signals keyed at the $\frac{1}{T}$ rate. What is required is an adaptive filter with a bandwidth at least as wide as the received signal.

This filter would adapt towards the NWMF for the channel (with symbol rate updates). In performing this filtering (assuming perfect adaption) the filter also removes any ISI. Such a filter can be implemented as a *fractionally spaced equaliser* (FSE).

The term FSE can be used to describe any digital filter which has more than one tap per symbol. The use of the word "equaliser" however suggests its best known use which is to correct, or "equalise" the effects of some transmission channel. The term FSE distinguishes it from a symbol spaced equaliser which has only one tap per symbol. The fractional spacing of the taps causes the fractionally spaced equaliser to have a wider bandwidth.

Fixed and adaptive equalisers in general, including those with fractional spacing and their advantages over symbol spaced equalisers, are described in detail in [34] or at a simpler level in [35].

3.2.2 Adaptation of filters

Although it may be possible to calculate a statistical representation of a channel, the effects of time-varying channels can be better mitigated by using a time-varying filter. This is especially useful when the channel properties vary rapidly in time (as they do for example in the land mobile, or HF radio channels). Often this will be achieved by sending a known "training signal" through the channel. The received (and corrupted) version of this signal is compared in the receiver to the receiver's copy and any differences are used to adapt the filter weights to undo the channel

corruption. The general structure of a trained adaptive filter is shown in (figure 3.3). The adaptive filter will, in suitable conditions, converge to the Wiener filter. In time-varying conditions, if the variation is slow enough to be tracked by the adaption algorithm, the filter will stay close to the (time-varying) optimum filter [17]. In terms of implementation, this is likely to be a FIR transversal digital filter, with tap weights adjusted using an algorithm such as LMS or recursive least squares (RLS).

Normally adaptive filters are digital, operating on sampled data. The structure described above, of an analogue MF followed by a SSE was common because often the ISI in a channel is time varying but the interference is not. When digital technology was slow and expensive it was sensible to minimise the complexity of any adaptive digital part of the receiver. Now however the FSE (section 3.2.1) is often preferred.

Any digital filter can be made adaptive and the theory and practice of linear transversal (FIR) adaptive filters is a well developed one, particularly for systems where a training sequence is available. In this thesis, apart from in chapter 7, the interest is in trained filters. With a fixed filter one can calculate the response from knowledge of the noise spectrum, and the transmitted pulse shape (for a MF or NWMF) or from knowledge of the transmitted, desired and received signal statistics (for a Wiener filter, see section 3.3). In a trained adaptive filter, the final filter response is determined by the statistics of the received and training signals, and by the timing of the filter updates. For example, if the filter taps are updated every sample, then the MSE will be minimised for every sample, and the filter will adapt to be the Wiener filter for the signals concerned. If the taps are only updated once every symbol, then the filter will minimise the MSE once per symbol, and the filter will adapt to the MF.

The structure shown in figure 3.1 is optimum if the interference conditions do not vary, but the channel ISI does, and is tracked effectively by the equaliser.

3.2.3 Decision feedback equaliser

The symbol spaced equaliser and FSE are both linear structures; their outputs are linear combinations of their input data. A significant improvement in the equalisation of ISI channels can be achieved by the used of a non-linear structure known as the *decision feedback equaliser* (DFE). This exploits the knowledge of data symbols which have already been detected. If post-cursor ISI is created by the channel then more accurate removal of this ISI requires manipulating the data from the decision device (which is noise free) instead of manipulating the unprocessed received signal, as the linear filters do. An adaptive structure which does this is shown in figure 3.4.

This filter can take the place in a receiver of a symbol spaced equaliser or a FSE depending on whether the feedforward filter is symbol spaced or fractionally spaced.

It can also be enhanced by using a FRESH filter as the feedforward filter [62]. This structure then has good performance in channels with interference and ISI.

There are many other techniques for equalisation which are not described here. The FSE is however of interest as a baseline for comparison with the FRESH filter.

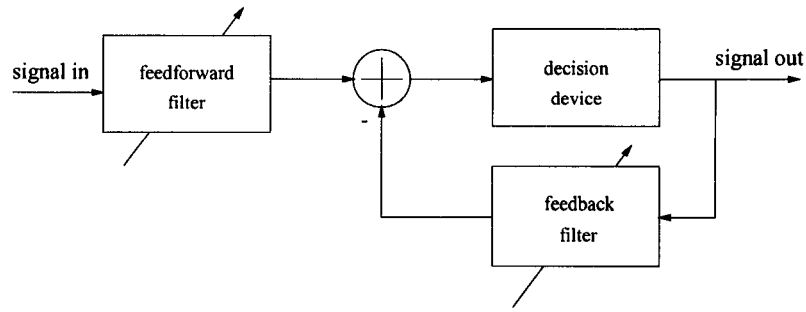

Figure 3.4: Decision feedback equaliser

3.2.4 The role of the filters

It is important to distinguish between the roles of the different filters that have been introduced so far. In this discussion the term Wiener filter will be used loosely to mean either the true ideal filter, or an implementable filter which is the best practical approximation to the true Wiener filter.

The Wiener filter theory is based on the extraction of a signal from noise where the statistics of the signal and noise are known. So the output of a Wiener filter is the closest (in the mean square sense) estimate possible of the corrupted signal, using a linear time-invariant filter. It attempts to extract the entire waveform - there is no mention in Wiener filtering theory of digital data or a symbol rate.

The Wiener filter performs a different job from the matched filter. The matched filter maximises the SNR at symbol spaced sampling instants. The filter output between these instants is not considered because it is assumed that symbol rate sampling is performed on the output. Notice that the definition of the matched filter does not require stationarity, or any other statistical assumptions, about the signals.

Another way of distinguishing between the two filters is to say that the Wiener filter is used for estimation - it provides an estimate of the noise corrupted signal. A matched filter is for detection - it produces an output on which a decision can be based as to whether a particular waveform (or any one of a set of waveforms) is present or not.

In terms of a practical implementation using a FSE, the distinction between the two filters comes down to timing the updates of the adaptive algorithm. If new tap values are calculated for each output sample, then the algorithm will attempt to minimise the cost function (for example, the mean squared error) at every sample, and will so minimise the MSE over the whole waveform. This is an approximation to a Wiener filter. If however new tap values are calculated only once per symbol, then the cost function is minimised at one sample in every symbol. If the cost function is MSE then this is the NWMF.

The FRESH filter, described below, is primarily intended as an interference rejection filter, and (for a filter exploiting baud rate related cyclostationarity) exploits excess bandwidth over the Nyquist minimum. This allows two different ways of implementing it in a receiver: it may perform an interference rejection role (as an improved version of a noise whitening filter) followed by an appropriate MF, or it may replace the NWMF, and incorporate the MF role. In terms of implementation, in the first case a FSE or FRESH filter would precede a MF, with an equaliser after that; in the second a single FSE or FRESH filter carries out all the filtering in the receiver.

For an adaptive implementation, the first case would require the filter (FSE or FRESH) to be trained with a training signal identical to the transmitted signal. That is, if the transmitted signal were square root raised cosine filtered, the training signal would also be square root raised cosine filtered, and would be followed by a fixed square root raised cosine filter (the MF).

The second case would require a transmitted signal to make the filter adapt to the NWMF. In the example of a square root raised cosine transmitted signal, the training signal would be raised cosine, and the filter could then be followed immediately by baud rate sampling and a decision device.

In this thesis, the first example is used (with GMSK signals) in chapter 6. The second implementation, with square root raised cosine QPSK, is used everywhere else.

The frequency shift filter described below is based on a generalisation of the Wiener filter and as such is used for estimation, not detection.

3.3 Wiener-Hopf Filter Equations

The FRESH filter is a structure which can implement the cyclic Wiener filter. This filter is a generalisation of the standard Wiener filter, which is described very briefly for comparison.

A fractionally spaced equaliser operating adaptively using the LMS algorithm, will adapt towards the minimum mean square error solution for a time-invariant filter. This corresponds to a Wiener filter solution, except for the "adaptation noise" caused by the taps jumping around close to the ideal solution. The Wiener filter is defined by the Wiener-Hopf equations, the normal presentation of which is in a time domain form:

$$\mathbf{R_{xx}w_0} = \mathbf{p} \tag{3.1}$$

where $\mathbf{R_{xx}}$ is the correlation matrix of the input signal $\mathbf{x}$, $\mathbf{w_0}$, is the optimum tap weight vector and $\mathbf{p}$ is the cross correlation vector between the desired response d, and the input signal, $\mathbf{x}$, ([96] chapter 5). So the optimum tap weight vector is given by

$$\mathbf{w_0} = \mathbf{R_{xx}}^{-1}\mathbf{p} \tag{3.2}$$

This vector describes the optimum filter for detection of a stationary signal in stationary noise. An equivalent representation can be given in the frequency domain:

$$S_x(f)W(f) = S_{dx}(f) \tag{3.3}$$

where $S_x(f)$ is the power spectral density of the input signal (the Fourier transform of one row of $\mathbf{R_{xx}}$), $W(f)$ is the frequency response of the optimum filter (the Fourier transform of $\mathbf{w_0}$) and $S_{dx}(f)$ is the cross-spectral density of the desired response and input signal. This representation and notation is compatible with Gardner's development of the cyclic Wiener filter equations (or FRESH filter equations) [13].

3.4 Cyclic Wiener Filter Equations

The cyclic Wiener filter, as described by Gardner [13] is a generalisation of the Wiener filter so that, as well as $\mathbf{R_{xx}}$ (the autocorrelation of the input), the cross-correlations of the input with the frequency shifted inputs are also used. Similarly, as well as $\mathbf{p}$ (the cross-correlation of the desired and input signals) the cross-correlations of the desired and frequency shifted input signals are used.

First define the vector concatenations

$$\mathbf{h}(t) = [\mathbf{a_1}(t), \ldots, \mathbf{a_M}(t), \mathbf{b_1}(t), \ldots, \mathbf{b_N}(t)]' \tag{3.4}$$

$$\mathbf{z}(t) = [x_{\alpha_1}(t), \ldots, x_{\alpha_M}(t), x^*_{-\beta_1}(t), \ldots, x^*_{-\beta_N}(t)]' \tag{3.5}$$

where x_{α_i} is the input signal shifted in frequency by α_i and $x^*_{-\beta_i}$ is the complex conjugate of the input shifted in frequency by $-\beta_i$. The set of frequency shifts $\alpha_1, \ldots, \alpha_M$ and $\beta_1, \ldots, \beta_N$ are the cyclic frequencies and conjugate cyclic frequencies of the input. $a_1(t), \ldots, a_M(t)$ and $b_1(t), \ldots, b_N(t)$ are the impulses responses of filters used to filter each frequency shifted version of the input, such that the estimate $\hat{d}(t)$ of the desired signal $d(t)$ is

$$\hat{d}(t) = \sum_{m=1}^{M} a_m(t) * x_{\alpha_m}(t) + \sum_{n=1}^{N} b_n(t) * x^*_{-\beta_n}(t) \tag{3.6}$$

or

$$\hat{d}(t) = \mathbf{h}'(t) * \mathbf{z}(t) \tag{3.7}$$

A filter structure which produces this estimate is shown in figure 3.6. If we define the cyclic cross-correlation matrices:

$$\mathbf{R}_{zz}(\tau) = E[\mathbf{z}(t+\tau/2)\mathbf{z}^H(t-\tau/2)] \tag{3.8}$$

$$\mathbf{R}_{dz}(\tau) = E[d(t+\tau/2)\mathbf{z}^*(t-\tau/2)] \tag{3.9}$$

then the filter solution can be written as

$$\mathbf{S}_{zz}(f)\mathbf{H}(f) = \mathbf{S}_{dz}(f) \tag{3.10}$$

where $\mathbf{S}_{zz}$, $\mathbf{S}_{dz}$ and $\mathbf{H}(f)$ are the Fourier transforms of $\mathbf{R}_{zz}(\tau)$, $\mathbf{R}_{dz}(\tau)$ and $\mathbf{h}(t)$. Writing equation 3.10 in terms of $x(t)$ using equations 3.8 and 3.9 the result is

$$\begin{aligned}
&\sum_{m=1}^{M} S_x^{\alpha_k-\alpha_m}\left(f-\frac{\alpha_k+\alpha_m}{2}\right)A_m(f) + \sum_{n=1}^{N} S_{xx^*}^{\beta_n-\alpha_k}\left(f-\frac{\alpha_k+\beta_n}{2}\right)^* B_n(f) = S_{dx}^{\alpha_k}, \quad k=1,2,...M \\
&\sum_{m=1}^{M} S_{xx^*}^{\beta_k-\alpha_m}\left(f-\frac{\beta_k+\alpha_m}{2}\right)A_m(f) + \sum_{n=1}^{N} S_x^{\beta_k-\beta_n}\left(-f+\frac{\beta_k+\beta_n}{2}\right)B_n(f) = S_{dx^*}^{\alpha_k}, \quad k=1,2,...N
\end{aligned} \tag{3.11}$$

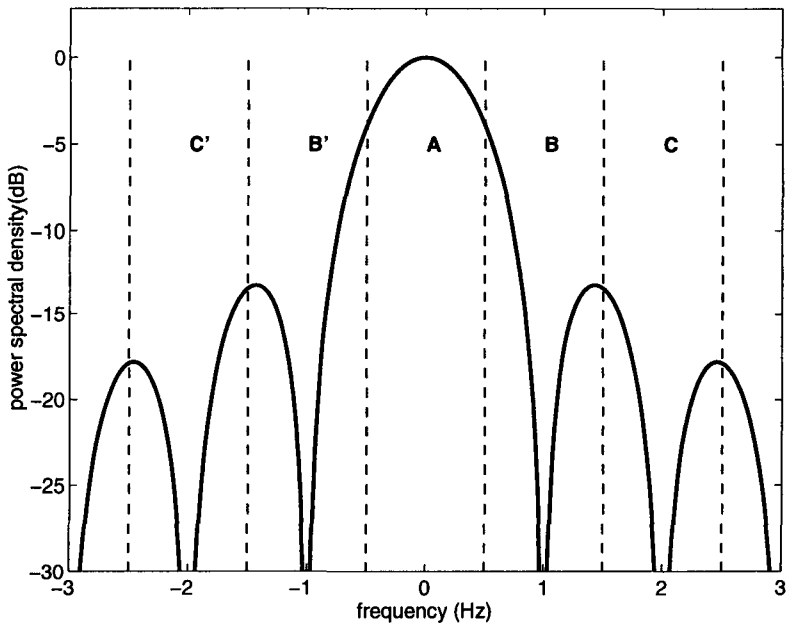

Figure 3.5: Power spectral density of a rectangular pulse

which specifies the optimum filter for cyclostationary signals. Examples of how these equations can be applied to specific signals are given in chapter 4.

These filter equations describe the minimum MSE solution for a given input signal, so an LMS adaptive FRESH filter will adapt towards this solution.

3.5 Frequency Shift Filtering

It has been shown in section 2.2 that cyclostationary signals exhibit spectral correlation in regularly spaced frequency bands. This is illustrated in the case of a rectangular pulse shaped signal with a symbol rate $\frac{1}{T}$ of 1 Hz (figure 3.5). The spectral components in the frequency ranges B and B' are correlated with those in A, (as are those in C, C' etc.). The region A which is known as the Nyquist bandwidth, is all that is required to recover all the samples with zero mean square error (if there is no noise or interference) ([79], section 6.2.1), but if some excess bandwidth is received (i.e. outside the Nyquist zone) then there is redundancy in the received signal which can be used for mitigating the effects of interference.

Region A contains all the information required to decode the data in the signal. The other bands are scaled, frequency shifted, replicas of the signal in band A (see section 2.2.2), so each individual band contains all the information needed to decode the signal.

Note that the boundaries between the regions in figure 3.5 have been drawn with an arbitrary origin. The signal could be detected ISI free from any $\frac{1}{T}$ bandwidth in the spectrum.

To exploit this correlation a filter must shift the correlated frequencies to the same frequency and add the resulting signals together. Provided that any noise present is not also correlated under the same frequency shift, this will increase the signal to noise ratio. The greater the number of correlated components that can be combined in this way, the greater the SNR increase will be.

The frequency shifts between correlated components are related to properties of the signal, such as

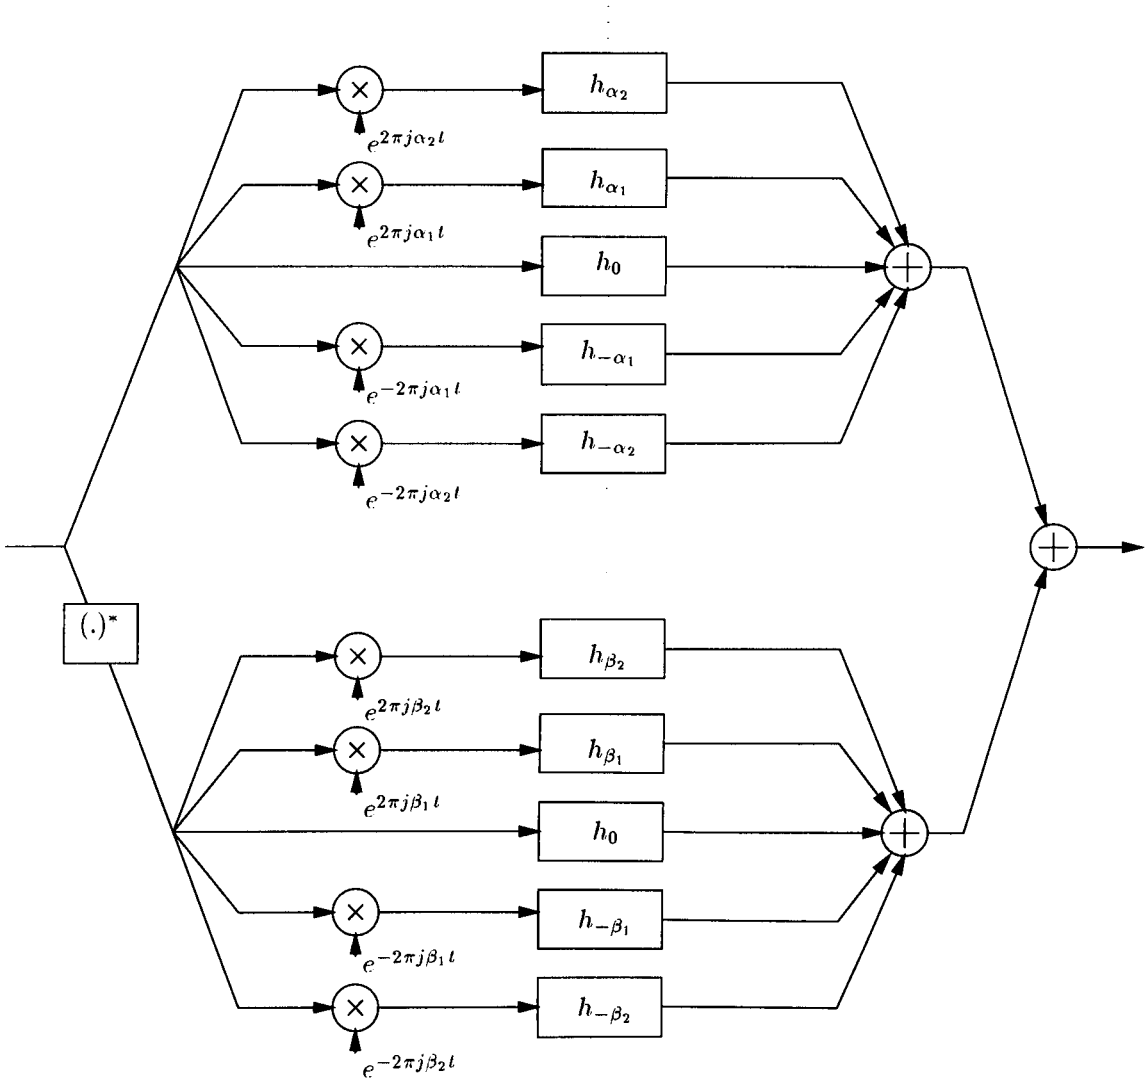

Figure 3.6: Frequency shift filter

the baud rate and carrier frequency. Therefore in general, other signals present, including thermal noise, will not be correlated under the same frequency shifts. A filter structure which combines correlated components after frequency shifting is shown in figure 3.6. This is known as a *frequency shift filter (FRESH filter)* [59, 60], or a *periodically time varying filter* [97, 98] or a *time-dependent filter* [57, 58]. It is also an approximation to, and can be equivalent to, the *cyclic Wiener filter* [13]. This name suggests a relationship with the better known Wiener filter, which is described below. This filter is also known as the *Optimal Time Dependent Filter* [57].

3.5.1 Exploiting the correlation of the SOI

Figure 3.6 shows a filter which can be used for extracting from noise a SOI with cyclic frequencies of $0, \pm\alpha_1, \pm\alpha_2$ (note that 0 is a trivial cyclic frequency of all signals), and conjugate cyclic frequencies of $0, \pm\beta_1, \pm\beta_2$. No assumptions are made about the noise other than than it does not have cyclic frequencies of $\pm\alpha_1, \pm\alpha_2, \pm\beta_1, \pm\beta_2$. The h_α and h_β are time-invariant filters which have bandwidth at least as great as the input signal. They are therefore implemented as individual FSEs. They are referred to here as *subfilters*.

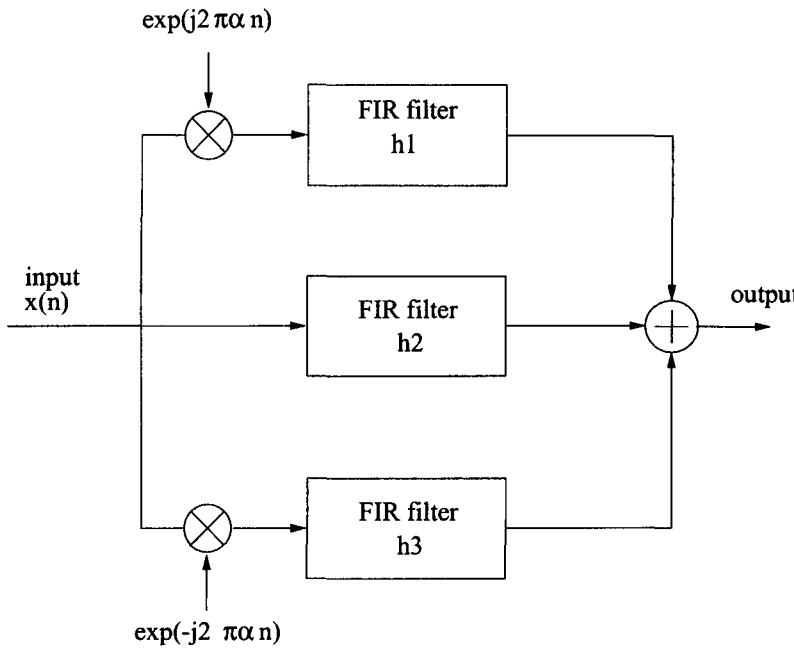

Figure 3.7: FRESH filter for 100% excess bandwidth QPSK

When filtering a signal such as QPSK, which has no conjugate spectral correlation, the bottom half of the filter can be omitted.

The signal input is shifted in frequency in each of the filter branches and then passed through an optimum time-invariant filter. This is equivalent to shifting every band in figure 3.5 to the zero frequency position (the band A position), filtering it, then summing all the branches together.

For best performance the number of branches is infinite, if there are an infinite number of cyclic frequencies. This would require an infinite bandwidth signal such as baseband BPSK with rectangular pulse shaping, which has cyclic, and conjugate cyclic, frequencies of $\pm\frac{n}{T}$ for any integer n (see section 2.2.4, and figure 2.20). We would expect that the greater the number of frequency shifting branches, the greater the increase in SNR relative to a FSE, but clearly even for a spectrally inefficient spectrum such as the $sinc^2(\omega)$ of the rectangular pulse shape, there is little signal energy shifted on top of the main lobe with the higher cyclic frequency shifts.

Exploiting non-conjugate spectral correlation

As another example, figure 3.7 shows a FRESH filter suitable for exploiting all the spectral correlation in a 100% excess bandwidth QPSK signal. Correlation in this signal exists under frequency shifts which are multiples of the symbol rate, but for shifts of more than $\frac{1}{T}$ or $-\frac{1}{T}$ there is no overlap in the spectrum, so a three shift filter overlays all the correlated components present (see figure 2.23).

Figure 3.8 illustrates schematically how this filter reduces the effect of a narrow band interferer. The first graph shows the spectra of the SOI and the noise; the second shows the response of the central subfilter in figure 3.7. This filter removes the interference but corrupts the SOI. The third graph shows the response of the subfilter in the $-\frac{1}{T}$ shifted branch. Here the SOI energy which is correlated with the notched region is amplified so it can replace the SOI energy removed by the

central branch.

Exploiting conjugate spectral correlation

As a further example, figure 3.9 shows a filter suitable for exploiting all the spectral correlation present in a 100% excess bandwidth baseband BPSK signal. The additional three frequency shifts and subfilters are required to exploit the conjugate correlation resulting from the symmetry of the BPSK spectrum.

The idea of a frequency shift filter which exploits the carrier frequency related cyclostationarity of a signal such as BPSK is described in [13], however an earlier work describes essentially the same process for the equivalent baseband signals [27] using *linear-conjugate-linear* filtering. This is also discussed further in sections 3.8 and 7.5.

It is well known that the Fourier transform of a baseband signal which can be expressed as purely real, is symmetric about the zero frequency axis, which is a manifestation of the fact that without introducing complex signals, the concept of negative frequency is meaningless. In fact the spectrum of a purely real signal is often plotted from zero frequency upwards only.

If such a signal is corrupted by additive noise which is also purely real, then the noise spectrum will also be symmetrical about zero frequency. We then have two identical copies of the SOI and noise at positive and negative frequencies. As the correlated SOI components are combined with correlated noise components, there is no way of exploiting the SOI symmetry.

If however we have a real baseband signal shifted up to some carrier frequency f_c as in figure 3.11, corrupted by real noise, there is now a difference in the structure of the symmetry in the SOI and the noise. The spectra of the SOI and the noise are both symmetrical about zero frequency, but there is also symmetry within each side band of the SOI about the carrier frequency. The noise will not be symmetric about the carrier frequency. If we shift one side band on top of the other, and add the signals, then we are adding correlated SOI components and uncorrelated noise components.

This does not however result in an improvement over processing the signal only at baseband. We have added complex white noise to a fundamentally real signal, and then used the symmetry of the SOI to remove the imaginary noise components. Moving the passband signal to baseband and setting the imaginary part to zero, which is possible with coherent reception, has the same effect.

However when the noise is not white, there is an advantage to be gained from exploiting the symmetry of the SOI spectrum. If we filter the two side bands of the SOI before shifting and adding them, (see figure 3.12) then the power of the interferer can be reduced, and the power of the SOI at frequencies correlated with the corrupted frequencies can be increased.

It is of course possible to model this process at baseband as well. In this case (as shown in figure 3.13) the SOI spectrum would be filtered, attenuating the interference affected frequencies and adding gain to the correlated SOI components, and then the signal would be added to its conjugate to combine the correlated components.

So we can see that there is a benefit in exploiting the spectral symmetry of a purely real signal such as BPSK. If the noise is complex and white, then this is equivalent to discarding the imaginary part of the noise, but if the noise is complex and non-white, then the filtering can remove interference

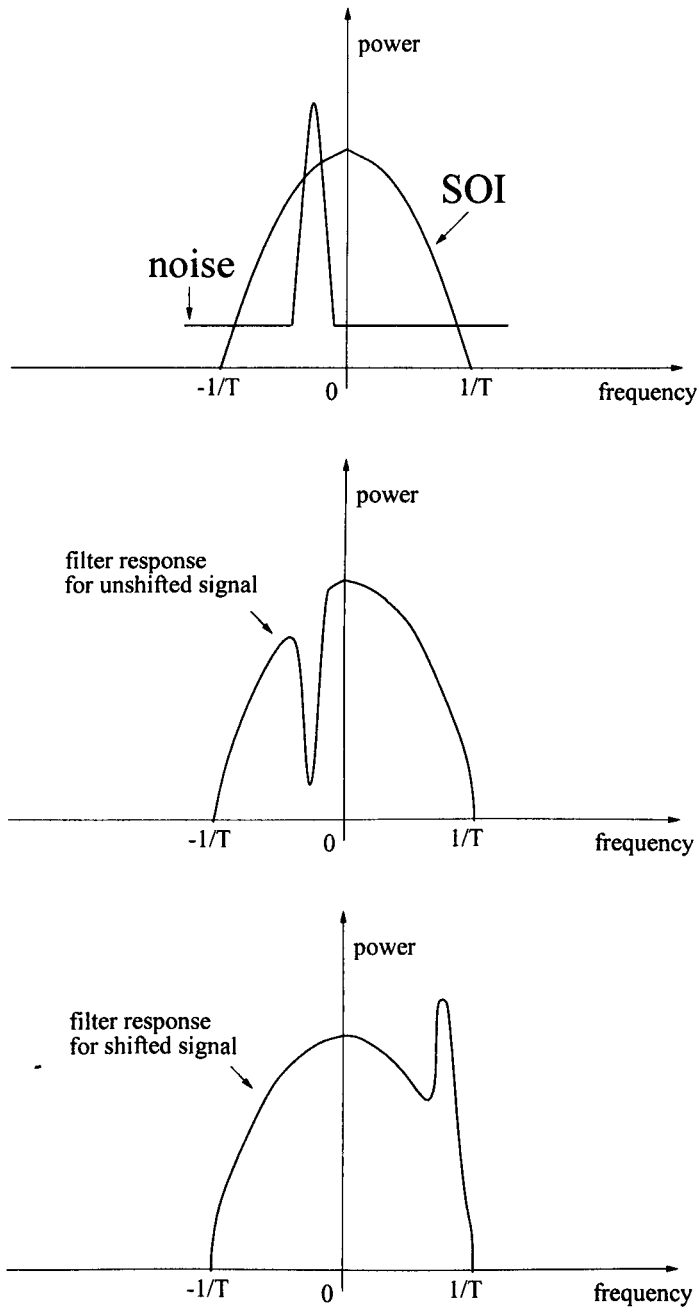

Figure 3.8: Removal of interference by a FRESH filter

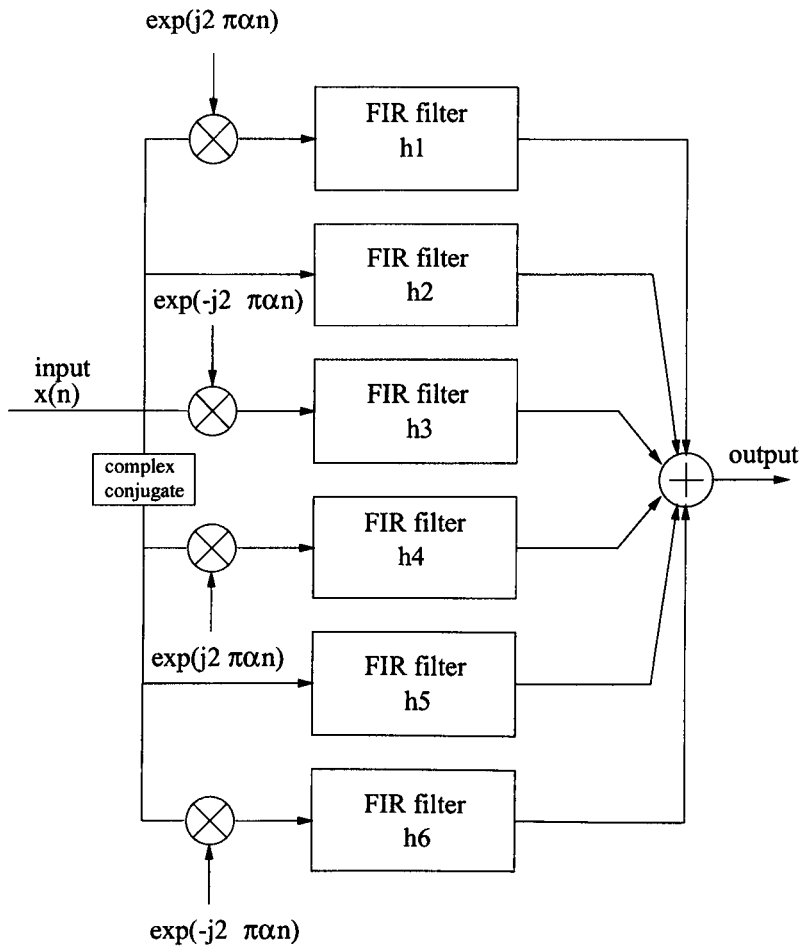

Figure 3.9: FRESH filter for 100% excess bandwidth BPSK

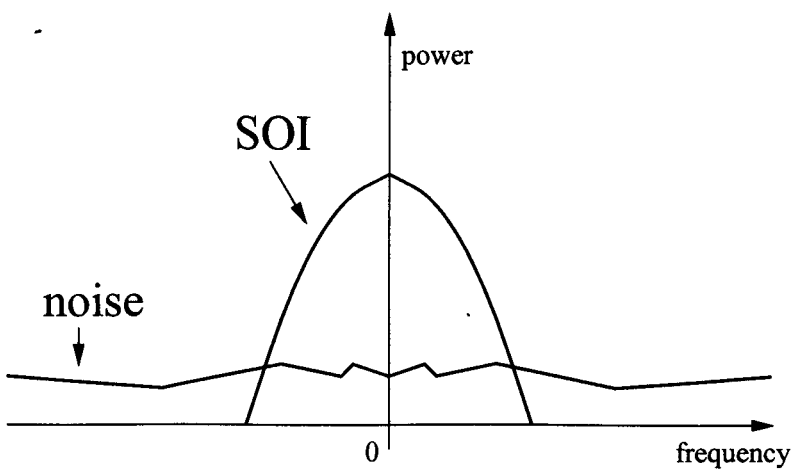

Figure 3.10: Spectrum of real baseband SOI in real white noise

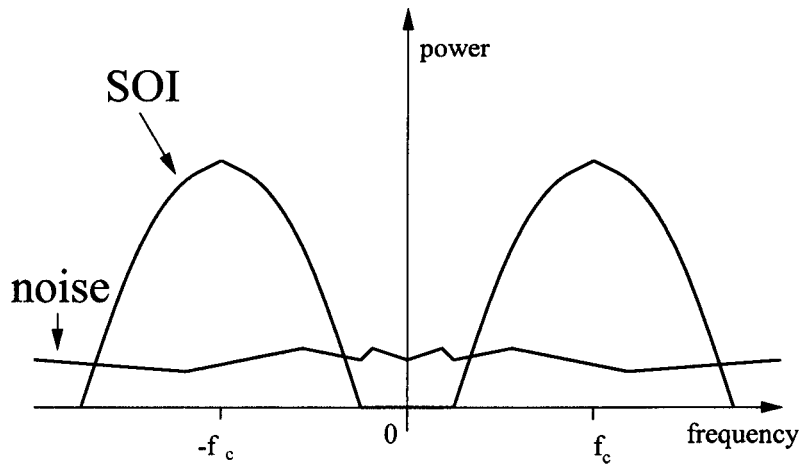

Figure 3.11: Spectrum of passband SOI in real white noise

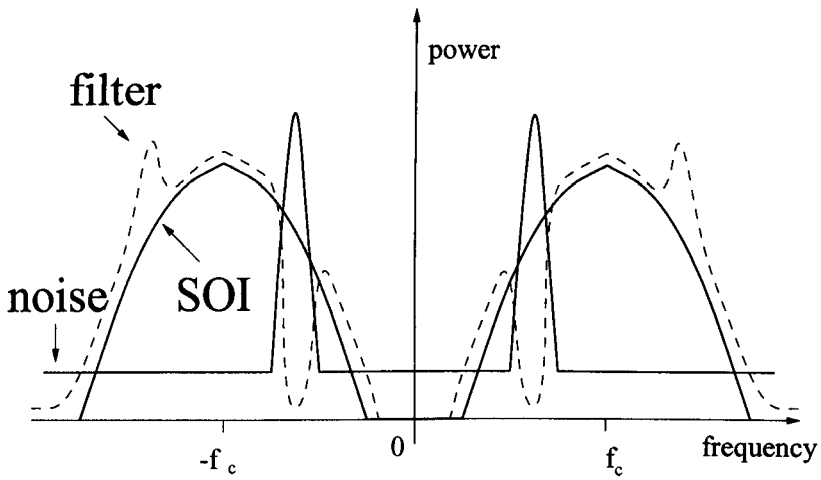

Figure 3.12: Spectrum of passband SOI in real coloured noise

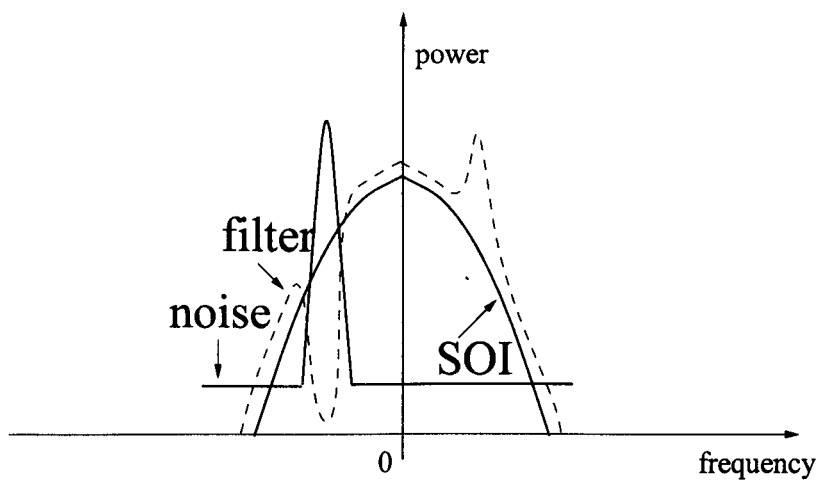

Figure 3.13: Spectrum of baseband SOI in coloured noise

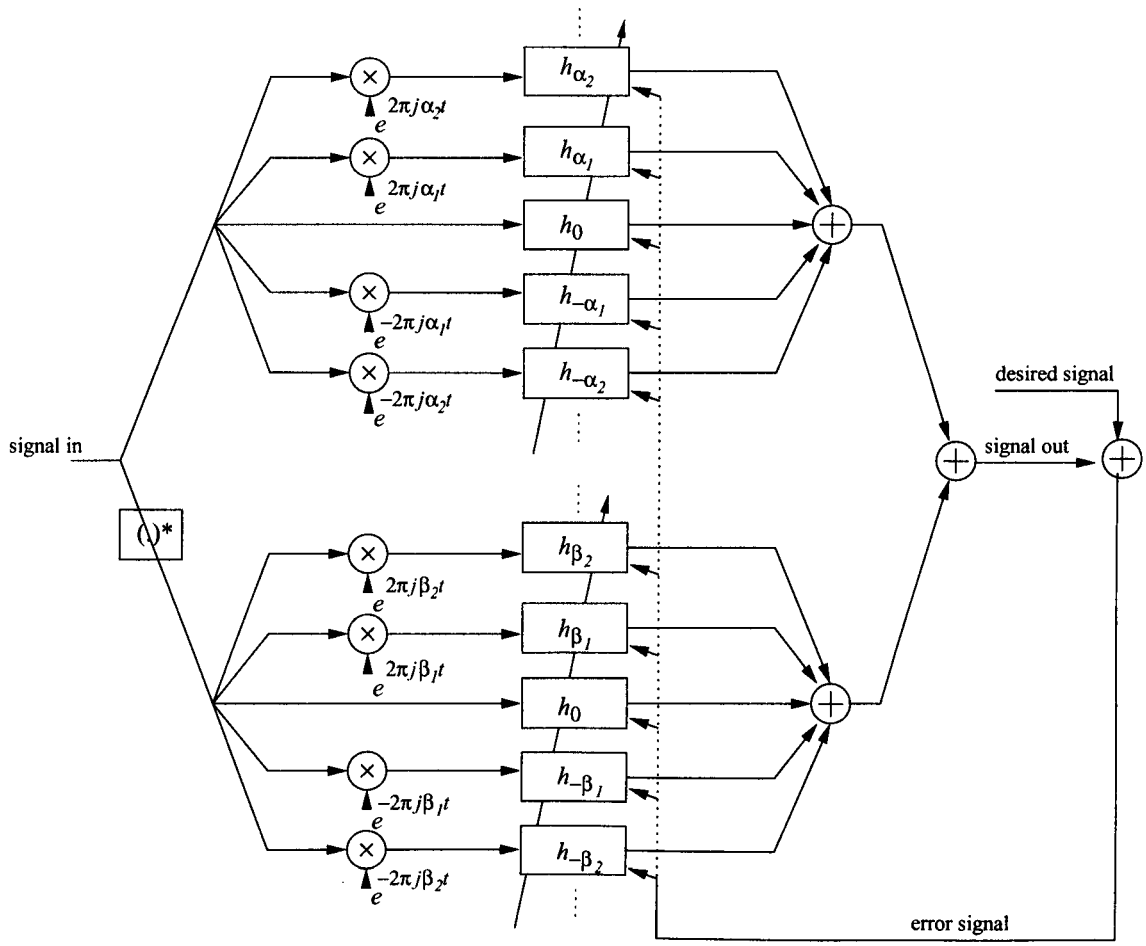

Figure 3.14: Adaptive FRESH filter

and replace the corrupted SOI components with correlated components from the other side band. This effect can be exploited by using passband processing and $2f_c$ frequency shifts, or by treating the signal as complex and using conjugate filtering.

It is also clear that the symmetry exploited at baseband by reflection (or by the LCL filter of [27]) is the same symmetry exploited at passband by a $2f_c$ frequency shift. Therefore one can describe the LCL filter as the baseband equivalent of the carrier frequency exploiting FRESH filter. In section 7.5 it is shown that such a filter can also use the SOI spectral symmetry for blind adaptation.

3.5.2 Adaptive FRESH filter

This (as described in [57]) is the adaptive filter implementation of the cyclic Wiener filter (see figure 3.14). It allows exploitation of spectral correlation in the signal or noise, while adjusting to varying channel conditions. The adaptive filters in each shifted branch, use standard adaptation algorithms, and all use the same error signal. Such an adaptive filter would converge towards the cyclic Wiener filter solution of equation 3.11.

3.5.3 Role of the FRESH filter in a receiver

It has been stated that the FRESH filter is an implementation of the cyclic Wiener filter, which is designed for extracting the whole SOI waveform from noise.

This means that a FRESH filter should incorporate a MF or be followed by a MF when symbol rate samples are required. In fact should the FRESH filter, with SOI symbol rate frequency shifts only, be followed by symbol rate sampling then the combined effect is identical to that of a FSE followed by symbol rate sampling. This is discussed in section 3.6.4.

A FRESH filter using only carrier frequency related shifts (or equivalently using a complex conjugate at baseband) gives no better performance than a system which identifies the SOI as purely real and discards the imaginary component of the received signal, as long as the filtering is first done with the complex signals. Simulation results to support this statement are shown in chapter 7. The simplest structure to do this is given in section 7.5

So it may appear that a FRESH filter has little use in a communications receiver but this is not the case. Firstly, up to now exploitation of SOI correlation only has been considered. It is also possible to exploit the spectral correlation properties of cyclostationary interference (see chapters 4 and 6). Secondly, the ability of a FRESH filter to remove the imaginary part of noise from a real signal may be preferred to other methods (for example, coherent detection). Thirdly, it may be useful to extract the whole waveform from the received signal with minimum distortion for use in operations such as synchronisation, although for such applications it may be preferable to exploit cyclostationarity more directly in the estimation algorithm [10, 74, 75].

A FRESH filter can be viewed as a development of the FSE which gives it a greatly improved interference rejection ability. The filter's role in a receiver is then the same as that of the FSE - it replaces a white noise matched filter and symbol spaced equaliser combination. It can also take the place of the feedforward filter in a DFE.

3.6 Variations and Simplifications of FRESH Filters

3.6.1 Periodieally time-varying filters

There are two intuitively different (but mathematically equivalent) ways of considering the cyclostationary property of excess bandwidth digitally modulated signals. They are the time domain approach and the frequency domain, or spectral correlation, approach [97] (see chapter 2). There are two distinct but equivalent filter structures which exploit the cyclostationarity of the digital signal, corresponding to these two different manifestations of cyclostationarity.

It will be assumed here, for the sake of simplicity that the filters considered here are only attempting to exploit the correlation due to the symbol rate. This removes the need for any conjugate filtering [13].

The first filter has an explicitly time-varying structure as shown in figure 3.15. It is assumed that the input is a sampled signal, so the number of subfilters is equal to the number of samples in one symbol interval. The commutator at the output means that the output is taken from each subfilter in turn, so making the overall filter time varying. Each filter would have the optimum transfer

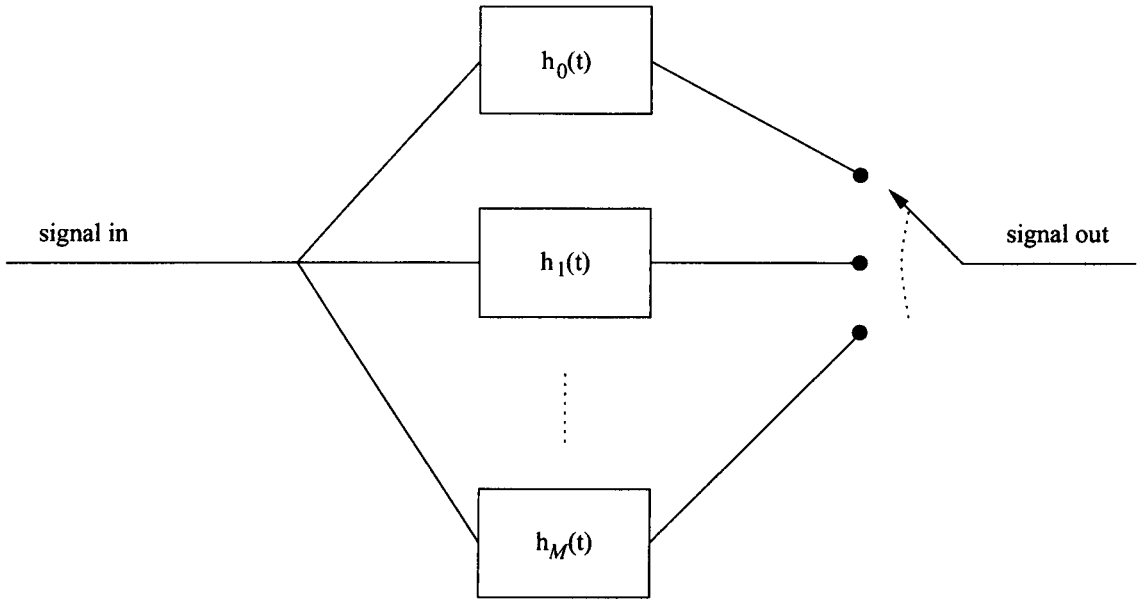

Figure 3.15: Explicitly time-varying filter

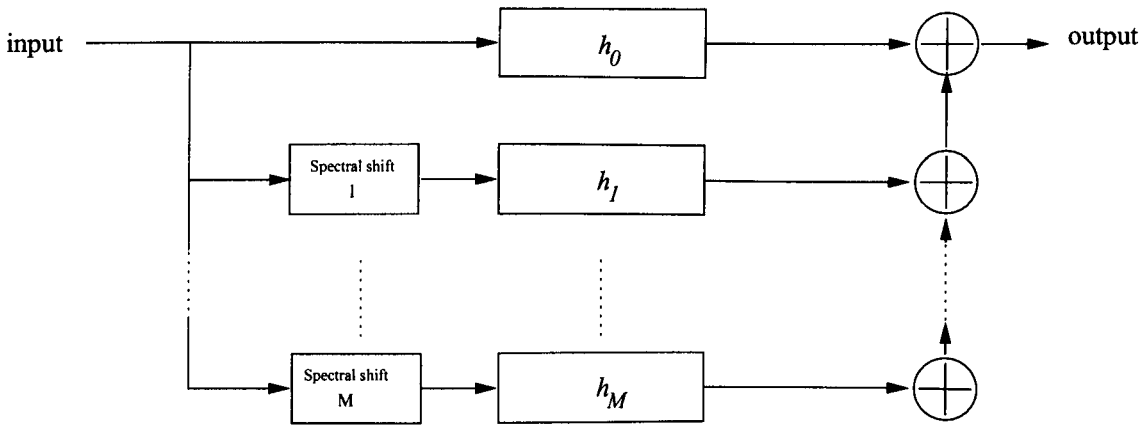

Figure 3.16: FRESH filter

function for the signal statistics at that point in the cycle. The frequency shift filter in figure 3.16 is more explicitly exploiting spectral correlation, by shifting the correlated components to the same frequency, filtering them independently and then adding them. There now follows a proof, taken from [97], that the two structures are equivalent.

The commutator which switches position with each sample, effectively decimates the signal in each branch with a staggered delay in the relative downsampling times of each branch. This is equivalent mathematically to multiplying the signal in each branch by $\delta_m(n)$ where:

$$\delta_m(n) = \begin{cases} 1 & \text{if } n = m, m \pm M, m \pm 2M... \\ 0 & \text{otherwise} \end{cases}$$

where there are M samples per symbol and m is an arbitrary integer.

This can also be written in a Fourier series representation:

$$\delta_m(n) = \frac{1}{M} \sum_{i=0}^{M-1} e^{\frac{j2\pi i(n-m)}{M}} \tag{3.12}$$

To demonstrate that the filters in figures 3.15 and 3.16 are equivalent, we use equation 3.12 to write

$$y(n) = \sum_{m=0}^{M-1} \{x(n) * h_m(n)\} \frac{1}{M} \sum_{i=0}^{M-1} e^{\frac{j2\pi i(n-m)}{M}} \tag{3.13}$$

where $y(n)$ is the output sequence from the filter, $x(n)$ is the input sequence and the $h_m(n)$ are the subfilter impulse responses as indicated in figure 3.15. $*$ indicates discrete convolution.

If we write $Y(z)$ for the z-transform of $y(n)$ (and likewise $X(z)$ and $H_m(z)$), then, taking the z-transform of both sides of equation 3.13 gives

$$Y(z) = \frac{1}{M} \sum_{i=0}^{M-1} \sum_{m=0}^{M-1} e^{\frac{j2\pi im}{M}} \sum_{n=-\infty}^{\infty} \{x(n) * h_m(n)\} e^{\frac{j2\pi in}{M}} z^{-n}$$

$$= \frac{1}{M} \sum_{i=0}^{M-1} \sum_{m=0}^{M-1} e^{\frac{j2\pi im}{M}} X(ze^{\frac{j2\pi i}{M}}) H_m(ze^{\frac{j2\pi i}{M}})$$

$$= \sum_{i=0}^{M-1} X(ze^{\frac{j2\pi i}{M}}) \frac{1}{M} \sum_{m=0}^{M-1} H_m(ze^{\frac{j2\pi i}{M}}) e^{\frac{j2\pi im}{M}}$$

Taking the inverse z-transform of both sides of the equation above gives

$$y(n) = \sum_{i=0}^{M-1} \{x(n)e^{\frac{j2\pi in}{M}}\} * \left\{ \frac{1}{M} \sum_{m=0}^{M-1} h_m(n) e^{\frac{j2\pi i(n-m)}{M}} \right\} \tag{3.14}$$

We can now define a new set of filter impulse responses by

$$w_i(n) = \frac{1}{M} \sum_{m=0}^{M-1} h_m(n) e^{\frac{j2\pi i(n-m)}{M}} \tag{3.15}$$

so equation 3.14 becomes

$$y(n) = \sum_{i=0}^{M-1} \{x(n)e^{\frac{j2\pi in}{M}}\} * w_i(n). \tag{3.16}$$

The $w_i(n)$ are the filter coefficients shown in figure 3.16. Equation 3.16 therefore shows that the output of the explicitly time varying filter is equivalent to a sum of frequency shifted versions of the input $(x(n)e^{j2\pi in/M})$ filtered by the $w_i(n)$ subfilters. The two filter structures in figures 3.15 and 3.16 are therefore equivalent.

Note that the number of subfilters in each filter representation is the same. If a critically sampled 100% excess bandwidth signal was being processed then $M = 2$. That is, the sampling rate is twice the symbol rate, and each filter contains two subfilter branches. This shows that only one frequency shift is required in the frequency shift implementation. This shift could be either $+\frac{1}{T}$ or $-\frac{1}{T}$.

If the signal is oversampled, it might appear that the corresponding increase in the number of branches in the time-varying implementation would have to be matched with extra frequency shifts that are not cyclic frequencies. However the extra samples which result from oversampling would

be correlated, and the extra filters in the time varying implementation would therefore be linearly dependent on those which appear in a critically sampled implementation.

3.6.2 The FRESH filter with sampling

There are two ways in which the sampling of a system can affect the best implementation of the FRESH filter. Both of these rely on the frequency-shifting property of sampling.

If the input to the FRESH filter is sampled at twice the highest magnitude cyclic frequency present then the filter requires only one shift for each magnitude of cyclic frequency present in the signal. For example, in a system with excess bandwidth of greater than 0% but less than or equal to 100%, the highest cyclic frequency present is $\frac{1}{T}$. A system sampled at $\frac{2}{T}$ (which is a practically sensible rate) requires only the $+\frac{1}{T}$ or the $-\frac{1}{T}$ frequency shift in the FRESH filter.

If the signal is sampled at the symbol rate immediately after FRESH filtering and the frequency shifts used are only integer multiples of the SOI baud rate, then the FRESH filter is exactly equivalent to a fractionally spaced equaliser (see section 3.6.4). This is the same as removing all the frequency shifted branches from the filter.

3.6.3 Effect of Nyquist-rate sampling

We consider the FRESH filter suitable for filtering a signal with 100% excess bandwidth, shown in figure 3.7. The entire input signal is processed by each branch so the bandwidth of each subfilter must be $\pm\frac{1}{T}$. This means that each subfilter is equivalent to a fractionally-spaced equaliser.

If the filter is in an analogue implementation then the two explicit frequency shifts are necessary to cause the correlated spectral components to overlap. However, if the filter is in a Nyquist-rate sampled system, then only one of the frequency shifts is necessary to cause the same overlap, because the spectrum of the shifted signal will wrap around half the sampling frequency - no benefit is gained by using all the theoretical cyclic frequencies.

This is apparent from consideration of the spectrum of the shifted signals. As we are dealing with a sampled system, frequency shifting will wrap the spectrum from $2f_s$ to $-2f_s$, where $2f_s$ is the sampling frequency. A positive frequency shift of $2f_s - m$ will give the same result as a negative shift of m. A good description of this wrapping of the spectrum in a sampled systems, with reference to filtering, is given in [34].

For the sake of simplicity this is shown only for the case of a 100% excess bandwidth signal. The Nyquist sampling rate for this signal is $\frac{2}{T_s}$ where T_s is the symbol period. In this case the effect of the sampling is to make the $\frac{1}{T_s}$ and $-\frac{1}{T_s}$ frequency shifts equivalent.

This can be shown mathematically as follows:

Let $x(t)$ be the received signal, and $x_1(t)$ and $x_{-1}(t)$ the versions of $x(t)$ shifted in frequency by $\frac{1}{T_s}$ and $-\frac{1}{T_s}$ respectively. $\frac{1}{T_s}$ is the symbol frequency. Sampling these signals at the rate $2f_s = \frac{2}{T_s}$ is equivalent to multiplying them by

$$\delta_s = \sum_{n=-\infty}^{\infty} \delta(t - \frac{nT_s)}{2} \tag{3.17}$$

Let the sampled versions of $x_1(t)$ and $x_{-1}(t)$ be $x_{1s}(t)$ and $x_{-1s}(t)$ respectively. Then:

$$x_{1s} = x_1(t).\delta_s(t)$$

$$x_{-1s} = x_{-1}(t).\delta_s(t)$$

If $X_{1s}(f)$ and $X_{-1s}(f)$ are the Fourier transforms of $x_{1s}(t)$ and $x_{-1s}(t)$ respectively, then, as multiplication in time domain is equivalent to convolution in the frequency domain:

$$X_{1s} = X_1 * \mathcal{F}[\delta_s(t)] \tag{3.18}$$

where $*$ indicates convolution and $\mathcal{F}[y]$ indicates the Fourier transform of y.

So:

$$X_{1s} = X_1 * \mathcal{F}[\sum_{n=-\infty}^{\infty} \delta(t - \frac{nT_s}{2})]$$

$$X_{1s} = X_1 * (f_s \sum_{n=-\infty}^{\infty} \delta(f - 2nf_s))$$

$$X_{1s} = f_s \sum_{n=-\infty}^{\infty} X_1(f - 2nf_s) \tag{3.19}$$

Similarly

$$X_{-1s} = f_s \sum_{m=-\infty}^{\infty} X_{-1}(f - 2mf_s) \tag{3.20}$$

But $X_1(f) = X_{-1}(f+2f_s)$ because they are frequency shifted versions of each other, so replacing n in equation 3.19 with $m - 1$:

$$X_{1s} = f_s \sum_{m=-\infty}^{\infty} X_1(f - 2mf_s + 2f_s)$$

$$X_{1s} = f_s \sum_{m=-\infty}^{\infty} X_{-1}(f - mf_s) \tag{3.21}$$

Therefore:

$$X_{1s} = X_{-1s} \tag{3.22}$$

As their Fourier transforms are equal, so x_{1s} and x_{-1s} are equal, which shows that for a 100% excess bandwidth signal, sampled at $\frac{2}{T_s}$, frequency shifts of $\frac{1}{T_s}$ and $-\frac{1}{T_s}$ give the same signal.

It is helpful also to consider the effect of Nyquist-rate sampling when the filter is implemented in the explicitly time-varying form shown in figure 3.15. As we have 100% excess bandwidth the sampling rate is twice the symbol rate, but the commutator operates at the symbol rate. There can therefore only be two filters in this implementation, as in the FRESH implementation.

3.6.4 Effect of symbol-rate sampling

If the FRESH filter is followed by a symbol rate sampler, then the folding effect of the sampling on the spectrum means that no frequency shifts are necessary [59], and the filter simplifies to a fractionally-spaced equaliser. This applies whether the FRESH filter is in a sampled system or an analogue implementation.

Using the FRESH filter as a receive filter in a communications system will normally involve sampling its output at the symbol rate, so this simplification is highly significant. It shows that the same exploitation of spectral correlation can be achieved with a much simpler structure. In addition it shows that in one filter we have the benefits of exploiting cyclostationarity, and the other known advantages of fractionally spaced equalisers. In a communications system the most significant of these is that the fractionally spaced equaliser can very successfully correct for any error in the sampling timing of the following symbol rate sampler [34].

Again for simplicity we prove the equivalence of the FRESH filter and the fractionally spaced equaliser for a 100% excess bandwidth signal. Effectively this means proving that all three branches of the filter in figure 3.7 are the same, which we do by showing that a frequency shift of $\frac{1}{T_s}$ or $-\frac{1}{T_s}$ followed by sampling at $\frac{1}{T_s}$ is the same as simply sampling the original signal at $\frac{1}{T_s}$.

We start by rewriting equations 3.19 and 3.20 for the lower sampling rate of $f_s = \frac{1}{T}$:

$$X_{1s} = f_s \sum_{n=-\infty}^{\infty} X_1(f - nf_s) \tag{3.23}$$

$$X_{-1s} = f_s \sum_{m=-\infty}^{\infty} X_{-1}(f - mf_s) \tag{3.24}$$

and we have similarly:

$$X_s = f_s \sum_{p=-\infty}^{\infty} X_(f - pf_s) \tag{3.25}$$

for the original signal.

As above,

$$X_{1s}(f - f_s) = X(f) = X_{-1}(f + f_s)$$

So we replace n with $p + 1$ in 3.23 and m with $p - 1$ in 3.24, then

$$X_{1s} = f_s \sum_{p=-\infty}^{\infty} X_1(f - pf_s - f_s)$$

$$X_{1s} = f_s \sum_{p=-\infty}^{\infty} X(f - pf_s)$$

and

$$X_{-1s} = f_s \sum_{p=-\infty}^{\infty} X_1(f - pf_s + f_s)$$

$$X_{-1s} = f_s \sum_{p=-\infty}^{\infty} X(f - pf_s)$$

That is

$$X_{1s} = X_s = X_{-1s}$$

This demonstrates that the two shifted, sampled signals are equivalent to the sampled version of the original signal, which in turn shows that if the filter is followed by symbol rate sampling,

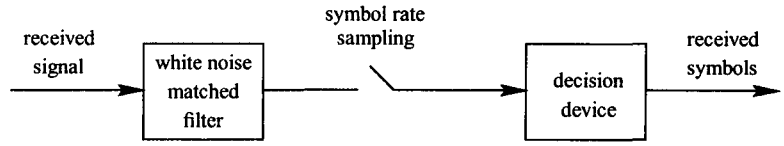

Figure 3.17: Simple receiver

the frequency shifts are redundant and the filter simplifies to the centre branch only, which is a fractionally spaced equaliser.

This is illustrated in figure 3.18 which shows a 100% excess bandwidth raised cosine spectrum before and after baud rate sampling. We know (see, for example, [99] section 2.5) that the spectrum of a sampled signal contains an infinite number of images of the spectrum of the original signal. These images are spaced $\frac{1}{T}$ apart where $\frac{1}{T}$ is the sampling frequency.

We can ignore all of the spectrum outside $\pm\frac{1}{2T}$ as it contains repeated versions that within $\pm\frac{1}{2T}$, so it gives us no additional information. The sampled signal spectrum has been *aliased.* This term normally implies that information has been lost but in this case the overlapping signal components are perfectly correlated, as described in [59].

This effect also means that a NWMF or a MF followed by symbol rate sampling also exploits spectral correlation related to the SOI symbol rate. To see why this is so, we consider the simple receiver shown in figure 3.17, receiving a signal with 100% excess bandwidth.

The sampling which follows the matched filter creates symbol rate spaced images in the spectrum exactly as described above. The explicit frequency shifts used in the FRESH filter are carried out here by the sampling. With the symbol rate updates, we have a FSE which adapts to be the NWMF for the particular interference scenario. The sampling overlays the correlated components, and it is the job of the filter to adjust the phases so that these components add constructively, and to weight them according to each component's signal to noise ratio [59]. A matched filter is the optimum filter for detecting a particular waveform in noise - there are no assumptions about stationarity as in Wiener filter theory; we cannot improve on matched filter theory by refining the model, as we can refine the Wiener filter to the cyclic Wiener filter by removing the assumption of stationarity.

3.6.5 Sampling after a periodically time-varying filter

It is also useful to consider symbol rate sampling following the explicitly time-varying implementation of the filter (figure 3.15). The commutator operates at the symbol rate, so if the filter is followed by symbol rate sampling, then only one of the subfilters contributes to the output symbol-rate samples. The other subfilters are therefore redundant and can be removed, leaving one filter operating at the sampling rate. This remaining filter is the fractionally spaced equaliser.

This simple argument combined with the equivalence of the filter structures in figures 3.15 and 3.16 of section 3.6.1 amounts to a proof of the equivalence of a FRESH filter and a FSE when both are followed by symbol rate sampling. This was first shown in [59].

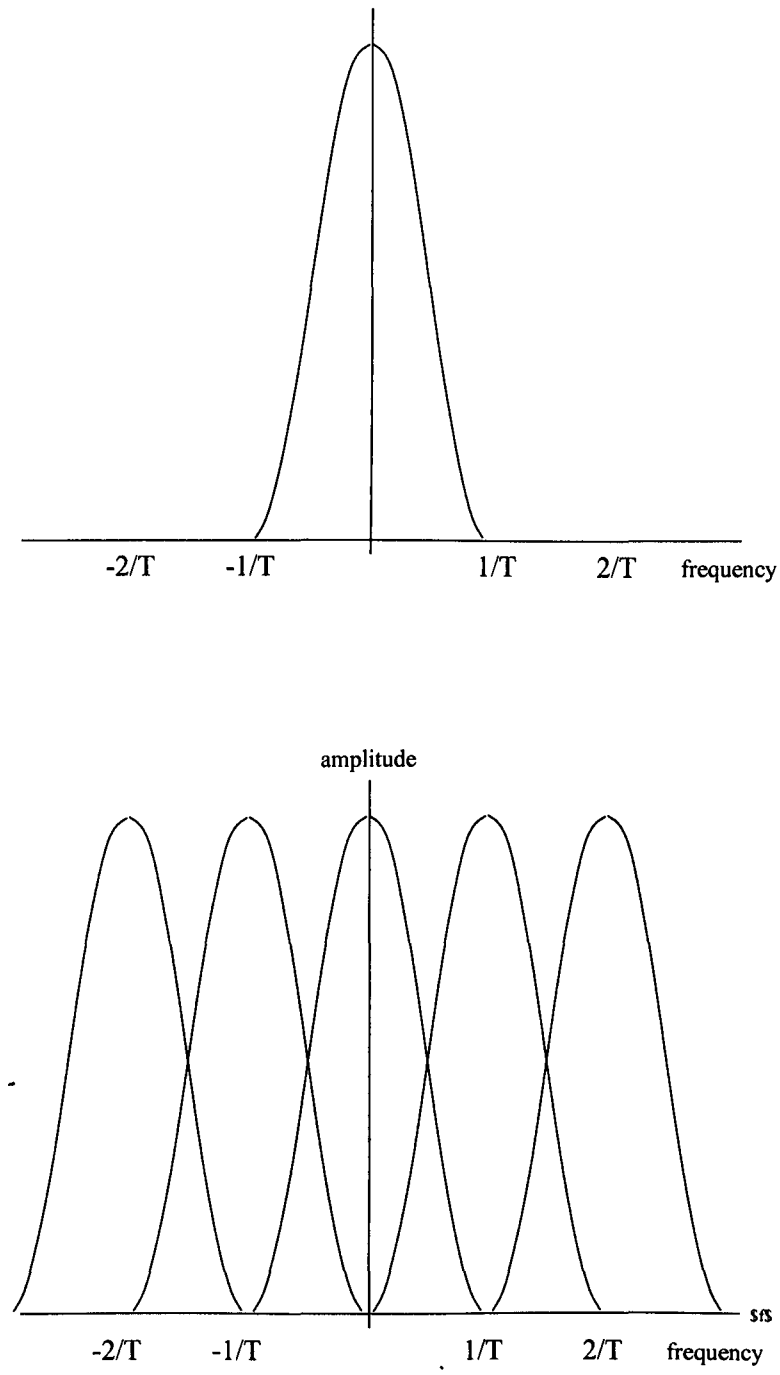

Figure 3.18: Sampled raised cosine spectrum components

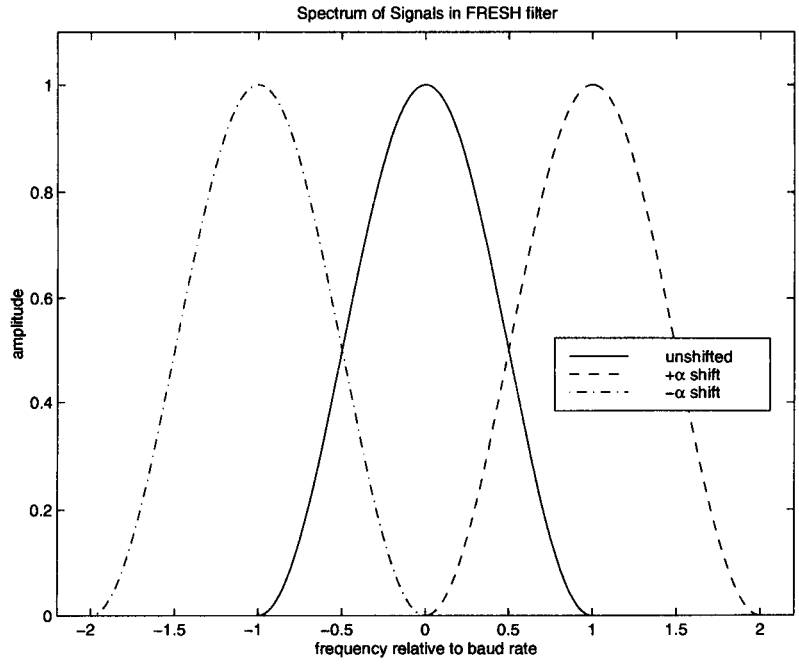

Figure 3.19: Internal spectra for 3 branch FRESH filter

3.6.6 Removing unnecessary frequency shifts

This section describes a modification to the FRESH filter structure which results in a more efficient filter in terms of processing load, especially when using signals with up to 100% excess bandwidth. This is relevant to filters which exploit baud rate related spectral correlation. The reason for the modification is apparent from considering the spectra of the signals in each branch of the filter. Figure 3.19 shows the spectra of the signals entering the three subfilters of the FRESH filter shown in figure 3.7. It is assumed that the sampling rate for the system is high enough to support the total bandwidth of all the frequency shifted signals, that is at least 4 times the baud rate. The input signal in this case is raised cosine filtered with 100% excess bandwidth, although the discussion which follows is equally valid for any excess bandwidth up to 100%.

To exploit the baud rate related spectral correlation, each subfilter is required to manipulate the components of its input which are correlated with the signals in other branches. This means that (using the labels from figure 3.7) h_2 must operate on frequencies from -1 to 1 (using the frequency scale from figure 3.19), while h_1 and h_3 need only operate on frequencies from -1 to 0 and from 0 to 1 respectively. That is, two of the filters only need half the bandwidth of the h_2 filter.

It is apparent then that having three filters operating at the same $(4/T)$ bandwidth is inefficient. At first sight it may seem that a better approach is to downsample each filter input to $2/T$. This means the h_2 branch is critically (that is, efficiently) sampled while the two other input signals are aliased, but in such a way that both inputs are identical. In that case only one of the branches is necessary, and the inputs to the two remaining branches are as shown in figure 3.20.

However it is not normally advisable to use critical sampling, as in many functions of the receiver it removes any margin for error, and can make filtering more difficult. But the FRESH filter complexity can still be reduced to containing only 2 subfilters by noticing that the spectra of the two shifted signals do not overlap. This means that they can be added together and filtered by one

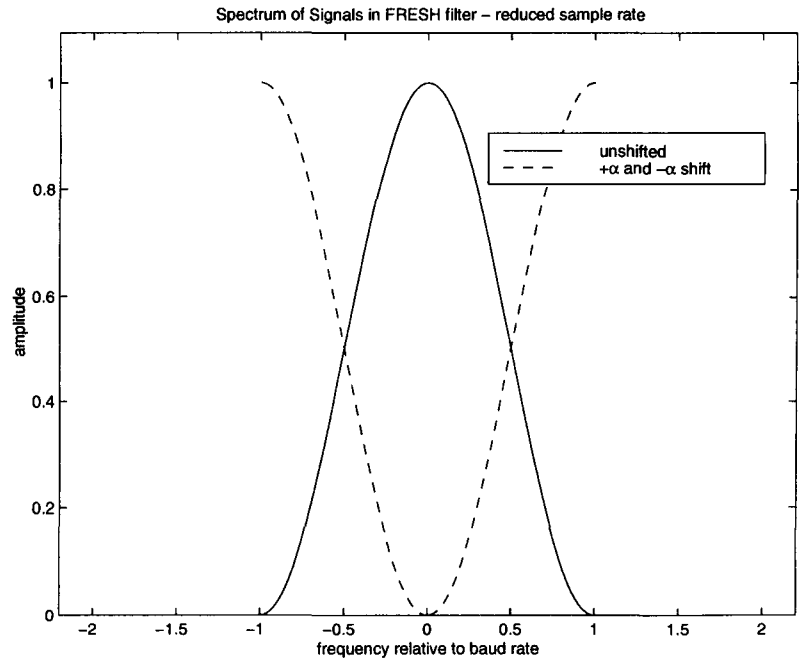

Figure 3.20: Internal spectra for critically sampled 3 branch FRESH filter

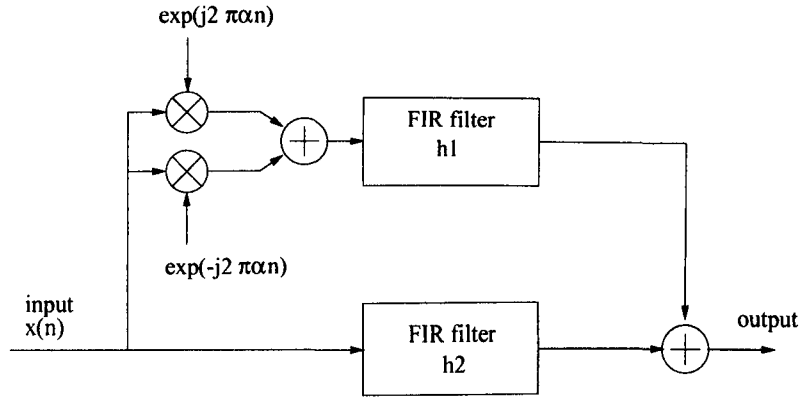

Figure 3.21: Efficient FRESH filter for 100% excess bandwidth

filter with bandwidth from $-2/T$ to $2/T$. We can then use a sampling rate of $4/T$ and a structure as shown in figure 3.21. The result is a filter which contains two subfilters instead of three, but exploits the spectral correlation at $1/T$ and $-1/T$. The filter proposed by Gardner (figure 3.7) has two subfilters (h_1 and h_3) with double the necessary bandwidth, and hence double the necessary processing load. This can be extended to wider signal bandwidths: for example a signal with excess bandwidth between 100 and 200% exhibits spectral correlation at $-2/T, -1/T, 1/T$ and $2/T$. This correlation could be fully exploited by a filter containing four subfilters, one operating on the unshifted signal, one on the $1/T$ shifted signal, one on the -1/T signal (these two cannot be added as their spectra overlap) and one on the $2/T$ and $-2/T$ signals added together. For these examples Gardner states incorrectly that three and four subfilters are required respectively [13].

3.6.7 Exploitation of spectral correlation by an FSE

We know that a FRESH filter with SOI baud rate frequency shifts exploits the spectral correlation at those frequencies by overlaying the correlated components and adding the signals. But we also

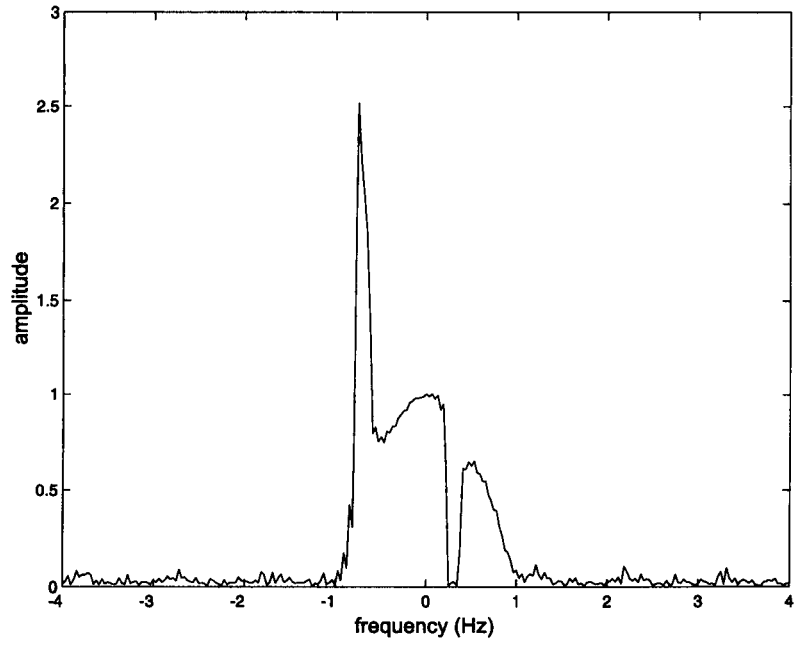

Figure 3.22: Frequency response of FSE with narrow band interference (symbol rate updates)

know that a symbol rate updated FSE is equivalent to such a FRESH filter, if both are followed by baud rate sampling (section 3.6 and [59]. This means that the FSE must also be exploiting the spectral correlation associated with the SOI baud rate.

For this equivalence to hold, the FSE response must be such as to minimise the mean squared error at symbol rate samples, not at every sample. In terms of implementation of an adaptive FSE, this means that the tap update process should be carried out once per symbol, and not every sample. The FSE which minimises MSE every symbol adapts to the matched filter for the SOI; but if it minimises the MSE for every sample it is trying to recover the entire waveform, and will adapt towards the Wiener filter for the SOI.

This is illustrated in figure 3.22. Here a baseband SOI (100% excess bandwidth square root raised cosine filtered QPSK) with baud rate 1 Hz is corrupted with a low level of white noise and a QPSK interferer with similar properties except it has a baud rate rate of 0.1 Hz and a carrier frequency of 0.3 Hz. The figure shows the magnitude frequency response of a FSE which was adapted using the LMS algorithm updating every symbol. We see that a notch in the response at 0.3 Hz removes the interferer, but there is also a peak in the response at -0.7 Hz, where the correlated components are enhanced to replace the corrupted SOI energy at 0.3 Hz. The baud rate sampling causes the correlated components to be overlaid.

Compare this with figure 3.23 which shows the response of the same FSE which has been allowed to update its taps every sample in exactly the same conditions as above. Here the filter cannot exploit spectral correlation, so it simply puts a notch around the interference.

3.7 Exploiting the Correlation of the Interferer

We have concentrated so far on the exploitation of the cyclostationarity of the SOI. A FRESH filter can shift correlated components to the same frequency, weight them appropriately according to the

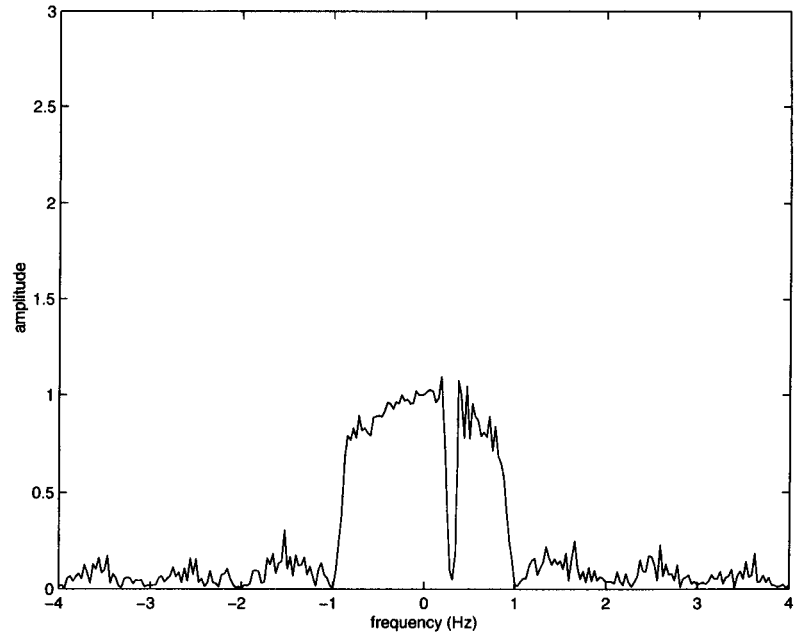

Figure 3.23: Frequency response of FSE with narrow band interference (sample rate updates)

SNR and add them together. If a cyclostationary interferer is present then knowledge of the cyclic frequencies of such an interferer allows a FRESH filter another degree of freedom in removing the interferer. In this case, frequency shifts are included which shift correlated interferer components to the same frequency. They can then be weighted and subtracted, so that the interferer effectively removes itself.

The effectiveness with which this can be done depends mainly on the interferer to noise to ratio. This means that the application of FRESH filters is especially appropriate in situations where the interferer power is very high. High powered interference overlapping with the SOI will cause major corruption of the SOI, which cannot be easily undone by other methods. FRESH filtering, on the other hand, benefits from the high interferer to noise ratio - the theoretical performance of the FRESH filter does not decrease with increasing interferer power.

The exploitation of interferer properties is already implicitly included in the cyclic Wiener equations 3.11. The equations include terms for every combination of cyclic frequencies present in the input signal, including those cyclic frequencies which are belong to the interferer.

A practical implementation of an interferer exploiting FRESH filter is likely to be in an adaptive structure such as in figure 3.14. It is necessary to know the cyclic frequencies of the interferer, but this issue has not been addressed here.

Results showing the performance of interference exploiting FRESH filters in QPSK and GMSK communications are given in chapters 4 and 6.

3.8 Circularity and Conjugate Filtering

It is interesting to consider the work published by Picinbono on the subject of circularity in light of the discussion on spectral symmetry in section 2.2.3. In [87] [100] Picinbono shows that in general the estimation of complex data is *widely linear* rather than linear. That is, when estimating

a random variable y in terms of an observation vector x, where both are complex, he says the optimum solution is not in general of the form

$$y = h^H \mathbf{x} \tag{3.26}$$

as in the real case, but

$$y = h^H \mathbf{x} + g^H \mathbf{x}^* \tag{3.27}$$

That is, the optimum solution is a linear combination of $\mathbf{x}$ and $\mathbf{x}^*$, not just a linear combination of $\mathbf{x}$ as given by Wiener filtering theory. This part of Picinbono's work was published earlier by Brown [27]. However, equation 3.27 reduces to equation 3.26 when $\mathbf{x}$ and y are jointly circular. This condition is expressed mathematically as

$$\begin{aligned} E[\mathbf{x}\mathbf{x}^T] &= 0 \\ E[y\mathbf{x}] &= 0 \end{aligned} \tag{3.28}$$

Note that these equations are not saying that the autocorrelation of x and the cross-correlation of x and y are both zero, because the normal definition of correlation for complex variables involves the complex conjugate of one of the variables. So in equation 3.28 Picinbono is in fact saying that the cross-correlation of $\mathbf{x}$ and $\mathbf{x}^*$ is zero, and that the cross-correlation of y and $\mathbf{x}^*$ is zero.

A random complex signal such as QPSK is not correlated with its complex conjugate and is therefore circular. Taking the scalar version of the first equation of 3.28:

$$\begin{aligned} E[x.x] &= E[a^2 - b^2 + 2jab] \text{ where } x = a + jb \\ &= E[a^2] - E[b^2] + 2jE[ab] \\ &= 0 \\ \text{because } E[a^2] &= E[b^2] \\ \text{and } E[ab] &= 0 \end{aligned}$$

The penultimate line of the above is a result of the symmetry of the QPSK spectrum, whereas the last line expresses the independence of the real and imaginary parts of the signal, which follows from the assumption of transmitting IID data. This establishes that QPSK is circular, but BPSK (or any purely real signal) is clearly not, as $\mathbf{x} = \mathbf{x}^*$ so the condition $E[x^2] = 0$ requires that $\mathbf{x}$ be 0 for circularity.

This is a general statement of something which for most practical situations is almost trivially simple: if we are estimating a signal such as QPSK in complex noise, a linear filter is sufficient. If however we are estimating a purely real signal like BPSK from complex noise, then a linear-conjugate-linear is significantly better because it allows the imaginary part of the noise to be discarded. When estimating BPSK which is not purely real (for example, using the constellation (1+j,-1-j)) the noise orthogonal to the SOI can be removed completely by the LCL structure,

although equally it could be removed by phase rotating the signal until the SOI is purely real, and then discarding the imaginary noise. This is perhaps why Johnson [88] describes the term "circularity" as "unnecessary and overblown".

From a practical point of view, one can say that most complex communication signals, such as QPSK, are circular, so the traditional equation 3.26 applies. Non-circular signals are those which have linear dependence between their real and imaginary part. For symmetric constellations this restricts us to signals which can, after a phase shift, be expressed as purely real. In this case the effect of equation 3.27 is to remove the imaginary part of the remaining noise, which of course could also be done by coherent detection. This topic is addressed again in section 7.5. As far as FRESH filters are concerned, one can say that non-circular signals have exploitable carrier related spectral correlation.

Chapter 4

Interference Rejection Performance of FRESH Filters

Contents

In this section we demonstrate the use of the FRESH filter in interference rejection, where the properties of the SOI only, or of both the SOI and the interferer can be used. QPSK signals are used as SOI and interference, so there is no carrier related correlation present. It is shown that given the knowledge that a cyclostationary interferer is present, and given the symbol rate of that interferer, a filter can be constructed which can almost completely remove the interferer in low noise conditions. If the noise level is higher, the filter will give a significant improvement over a FSE (with symbol rate tap updates) which does not exploit the cyclostationary properties of the input. In the literature, most examples of FRESH filtering use signals with rectangular pulses. Here, raised cosine filtered signals are used as they are much more common in practical situations.

4.1 FRESH Filter Using SOI Properties Only

Considering a 100% excess bandwidth (EBW) QPSK signal, which contains enough correlation to be equivalent to a frequency diversity system of order 2, it is clear that a 3 branch FRESH filter is sufficient to exploit all the correlation present in that signal. Such a filter is shown in figure 4.1(b).

This filter consists of three subfilters, each of which has bandwidth equal to that of the entire received signal. One subfilter (following "α shift" in figure 4.1(b)) operates on a version of the signal which has been frequency shifted by $+\alpha$ where α is the symbol rate of the signal, one operates on the unshifted signal, and the third operates on a version which has been shifted in frequency by $-\alpha$. This structure therefore causes each subfilter to act on versions of the signal of interest (SOI) which are correlated, but corrupted by noise and interference. The noise and interference are uncorrelated between each filter branch.

It is known ([59] and section 3.6.4) that the structure shown in figure 4.1(b) is equivalent to a fractionally spaced equaliser (figure 4.1(a)) when the filtering operation is followed immediately by symbol rate sampling and the filter is designed to minimise the MSE of those symbol rate samples. The effect of sampling is to create frequency shifted images of the original signal which are spaced in frequency by the sampling frequency. The symbol rate frequency shifts of the FRESH filter are made redundant by the imaging introduced by the symbol rate sampling downstream.

This equivalence exists only when the filter is producing symbol rate samples. In some cases a higher sampling rate may be required after the filter, so the 3-branch FRESH filter and the fractionally spaced equaliser are no longer equivalent. A higher rate of sampling could be used, for example, to enable frequency tracking of the carrier. If a higher sampling rate is used, then the filter is no longer recreating only the symbol rate samples, but is producing some estimate of the transmitted waveform. As stated above, the filter in figure 4.1 can exploit all correlation in a 100% EBW signal. More frequency shifts would be required if the signal had a wider bandwidth. For example a 200% EBW signal would need a filter with frequency shifts of $+f_s$, $+2f_s$, $-f_s$, and $-2f_s$, to exploit its spectral correlation fully.

4.2 FRESH Filter for Exploiting the Interferer

The filter described in section 4.1 exploits the correlation of the SOI by shifting the signal by the frequencies at which the correlation occurs. If a cyclostationary interferer is present with the same symbol rate as the signal of interest, then this filter can use the spectral correlation of the interferer to more effectively remove the effect of the interferer on the signal of interest. In this case no redesign of the filter is required, and it is still equivalent to a FSE. An example of a situation where a MSK SOI and interferer share cyclic frequencies is given in chapter 6.

If, however, the cyclostationary interferer, has a symbol rate which is different from that of the signal of interest, then a FRESH filter with additional frequency shifts and associated subfilters will improve performance significantly. The frequency shifts must match the frequencies of spectral correlation of the interfering signal, so for QPSK they will be harmonics of the symbol rate (see chapter 2).

The filter structure used in the simulations is shown in figure 4.1(c). This is the optimum filter for receiving a 100% EBW signal with symbol rate α in the presence of a 100% EBW interferer with symbol rate γ, assuming both signals show no conjugate correlation. The two symbol rates need not be harmonically related, and the filter requires no knowledge of the carrier frequency of the interference.

The number of additional frequency shifts required for optimum performance depends on the

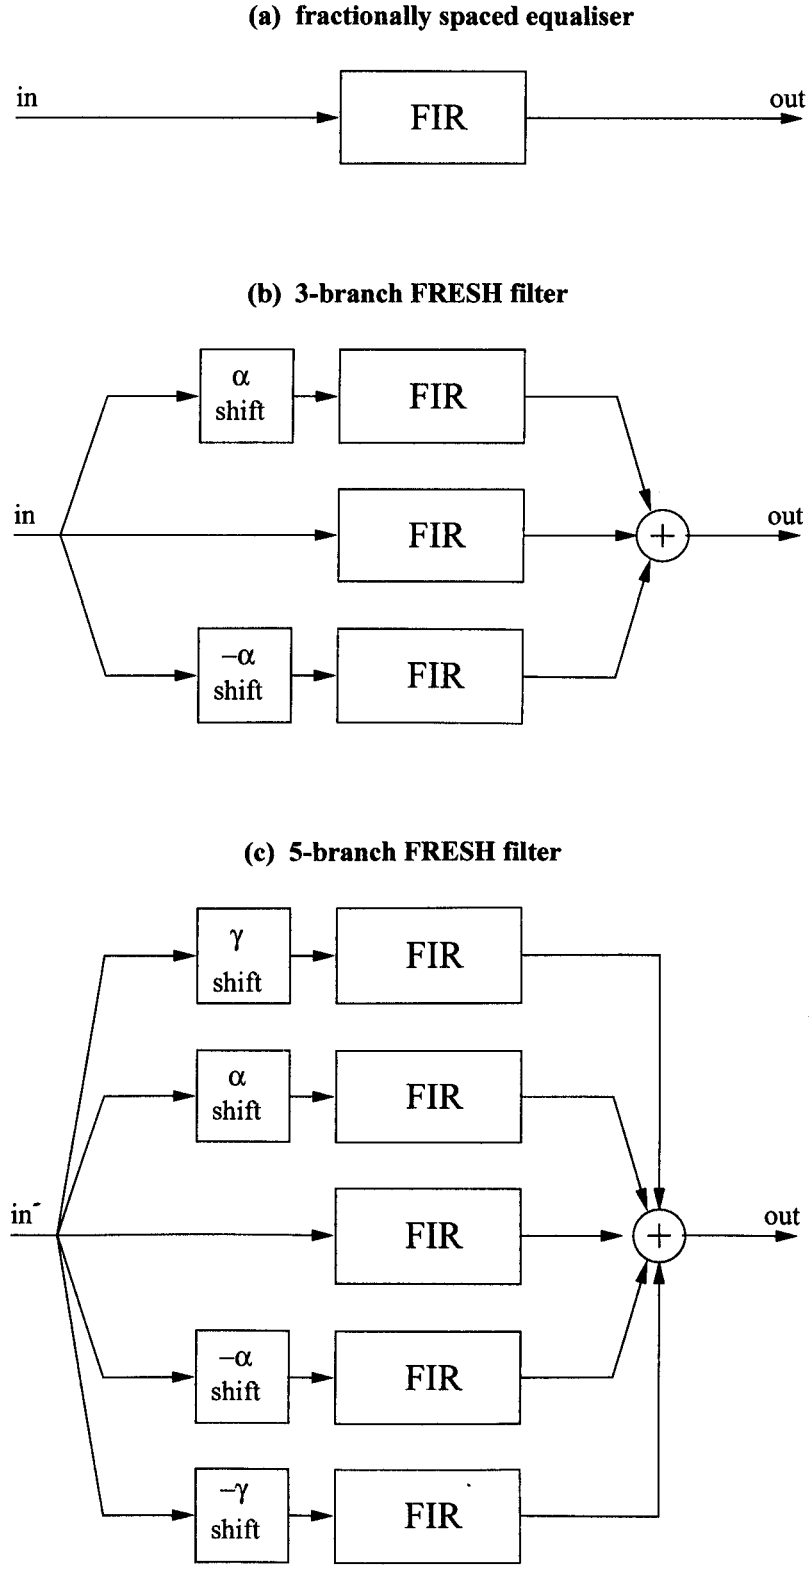

Figure 4.1: Structure of (a) FSE, (b) 3- and (c)5-branch FRESH filters

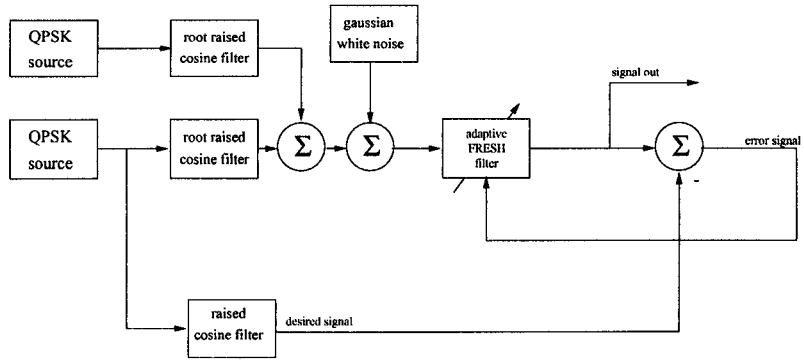

Figure 4.2: System used for simulation

bandwidth of the interferer. In the simulation results presented below the interferer has 100% EBW so frequency shifts of $+\gamma$ and $-\gamma$ are used. There is no advantage in using additional shifts unless the interferer has a wider bandwidth.

4.3 Description of Simulation and Results

Simulations were run to compare the performance of a FSE, a FRESH filter exploiting SOI properties, and a FRESH filter exploiting SOI and interferer properties. A block diagram of the system which was simulated is shown in figure 4.2. A QPSK signal is the SOI, and a separate QPSK signal forms the interference. AWGN is also added. The block labelled "adaptive FRESH filter" was implemented as either a fractionally spaced equaliser, a 3-branch FRESH filter or a 5-branch FRESH filter as shown in figure 4.1. The mean squared error of the output was measured for these three filters in the presence of different levels of white noise, and with different carrier offsets of the interferer. Noise was added to give an E_b/N_0 (for the SOI) of 0 dB to 15 dB. To make the advantage of using the 5-branch FRESH filter clearer, the interferer power was set to 10 times higher than the SOI power.

The FSE here is updating every sample, not every symbol, so will not exploit symbol rate spectral correlation, and will not therefore give the optimum BER performance. However the 3-branch FRESH filter BER results are equivalent to those for a symbol rate adapted FSE (see section 3.6.4).

The interferer carrier offset was used to model a difference in carrier frequency between the signal of interest and the interferer. As the simulation used the complex baseband, the carrier of the SOI was 0. The symbol rates of the SOI and interferer were 1 and 0.94 respectively. Carrier offsets of 0, 0.3, 0.7, 1.1 and 2.0 were used (that is 0, 0.3 etc. times the SOI symbol rate). These figures can be scaled to values likely to be seen in practice. For example a SOI symbol rate of 1 kHz could be assumed, so the interferer symbol rate is 940 Hz. If the SOI carrier were 1 MHz then the interferer carrier frequencies used were 1, 1.0003, 1.0007, 1.0011 and 1.0020 MHz. These values were chosen to show the extremes of performance: a 0 offset means the two signals are almost completely overlapping ("almost" because the different symbol rates mean that the extreme edges of the signal of interest will be uncorrupted by the interferer); an offset of 2 puts the two signals far enough apart so they do not overlap at all. The relative spectra of the signal of interest and interferer are illustrated in figure 4.3.

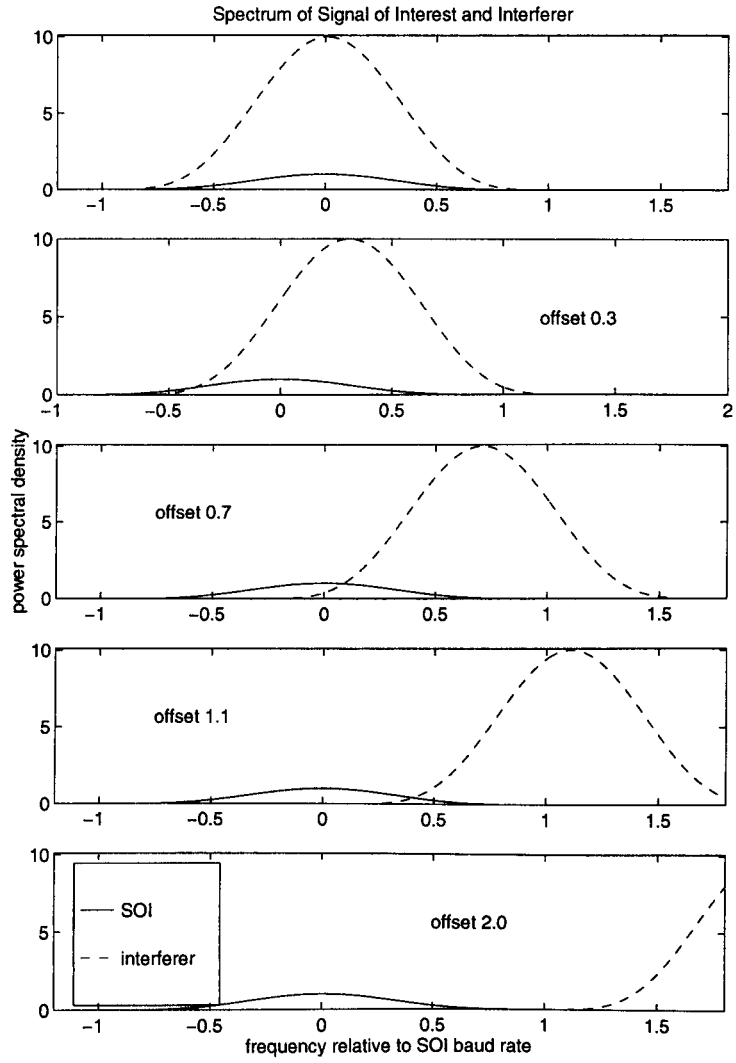

Figure 4.3: Relative spectra of SOI and interferer used in simulations

Figures 4.4 to 4.9 show the comparative performance of the three filter structures described above. Each graph represents one of the interferer scenarios illustrated in figure 4.3. With the interferer and SOI power and carrier frequency fixed, the noise level was varied to give E_b/N_0 (for the SOI) of 0 to 15 dB. In each case, the theoretical lowest BER for receiving QPSK in AWGN only is also plotted for comparison.

These results show that the 5-branch FRESH filter gives a dramatic improvement in performance over the 3-branch FRESH filter and single branch filter (fractionally spaced equaliser). This is particularly apparent for partially overlapping interference (figures 4.6, 4.7 and 4.8), where the improvement over the fractionally spaced equaliser is up to 13 dB, and over the 3-branch FRESH filter up to 10 dB at high SNR.

This means that the 5-branch filter is effectively removing nearly all the interference effects in low noise. It performs better than the other structures for all cases except high noise and completely overlapping signal.

As expected when the signals do not overlap at all (figure 4.9) there is little advantage in using

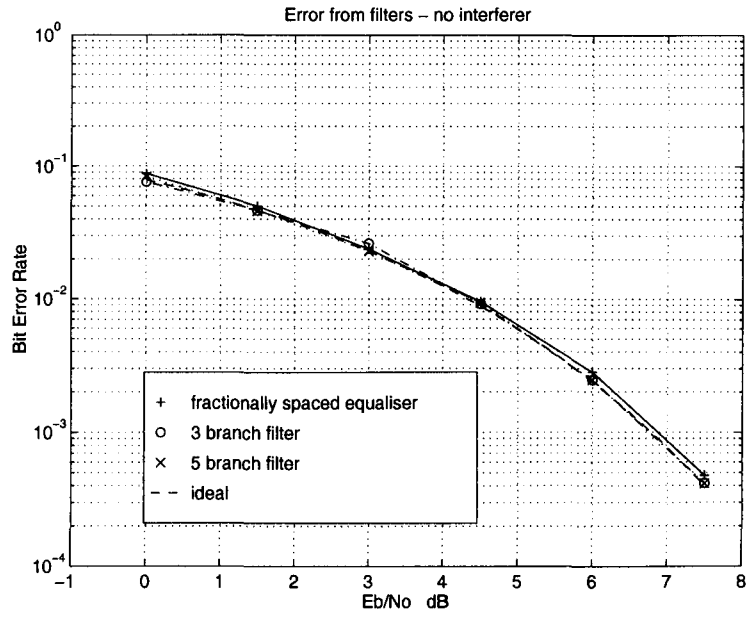

Figure 4.4: BER of QPSK in white noise only

the 5-branch filter, as a simple fractionally spaced equaliser can eliminate the interferer without corrupting the signal of interest.

It is interesting that at high noise the fractionally spaced equaliser sometimes performs better than the 3-branch FRESH filter. It is thought that this is due to the adaptive nature of the filters. As they are continually adapting, even in the steady state there is a small noise component generated by the changes in tap positions around the optimum at each iteration. This effect will be worse when there are 3 adaptive systems instead of just one, and so 3 times the number of filter taps in total. This effect will therefore depend on the value of the feedback coefficient used in the adaptation. The LMS algorithm was used here, with a feedback coefficient of 10^{-5} and a signal power of 1.

4.4 Effect of Number of Taps

All the results in the previous section used a filter of length equivalent to 4 symbols of the SOI. That is, using 16 samples per symbol and 1 tap per sample, the FSE and each subfilter of the FRESH filters had 65 taps.

For one of the above filter scenarios (interference carrier frequency 0.7 Hz, 5-branch FRESH filter) the effect of varying the number of taps in the filter is shown in figure 4.10. The BER was calculated for a filter with 17, 33, 65 and 257 taps in each of the subfilters. Given that the SOI symbol rate was 1 Hz and the sampling frequency 16, it is clear that the filter "lengths" correspond to approximately 1, 2, 4 and 16 symbols. With 17 taps the filter does not have sufficient frequency resolution to manipulate the signal as required so the performance is poor. 33 and 65 taps give similar results to each other, and the optimum tap number is in this region. With 257 taps the effects of "adaptation noise" start to be seen. The filter was adapted using the LMS algorithm so the taps will be close to, but not exactly at the optimum values.

Although the filter was fixed while the results were being taken, adaptation noise is present due to

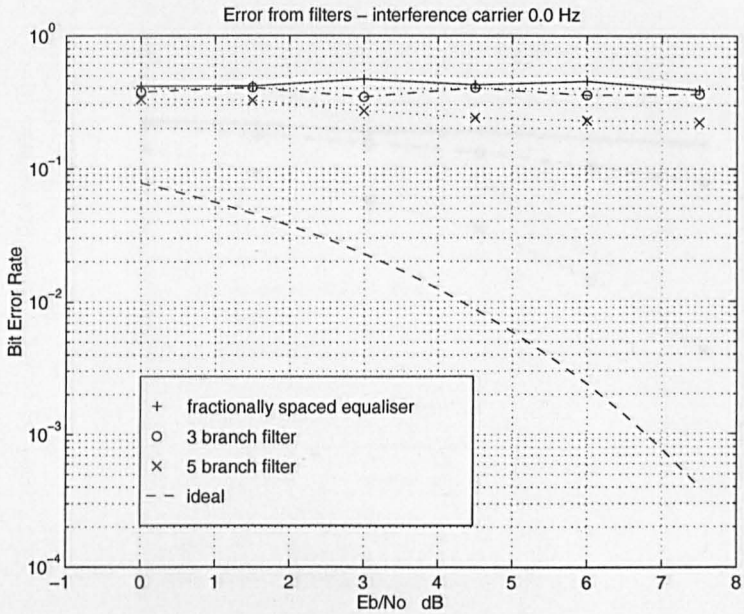

Figure 4.5: QPSK with QPSK interferer on same carrier

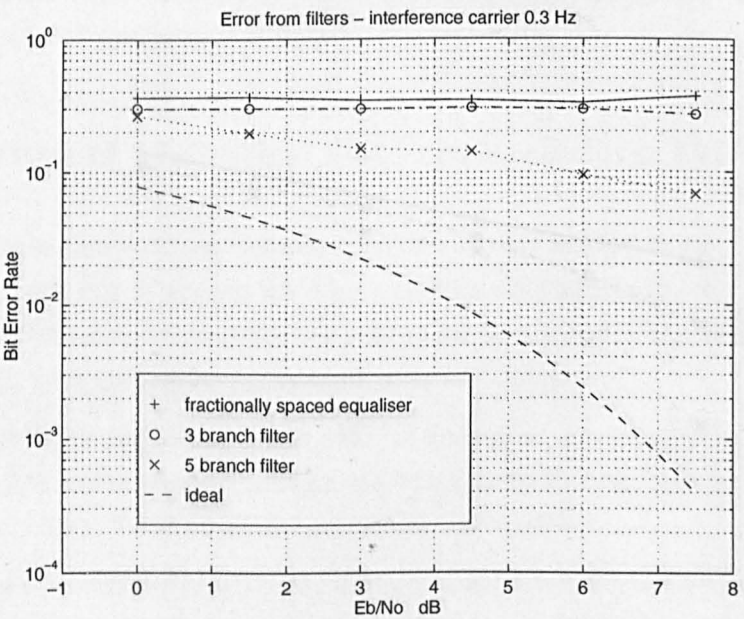

Figure 4.6: Interferer carrier offset 0.3

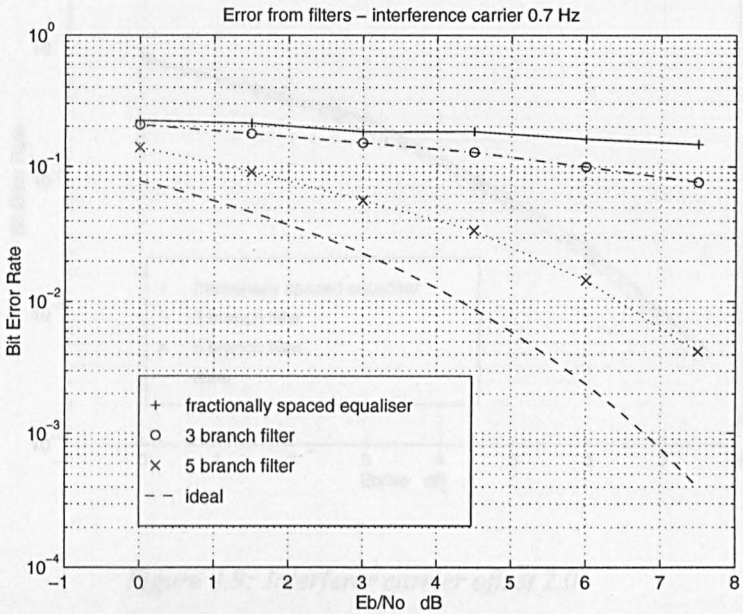

Figure 4.7: Interferer carrier offset 0.7

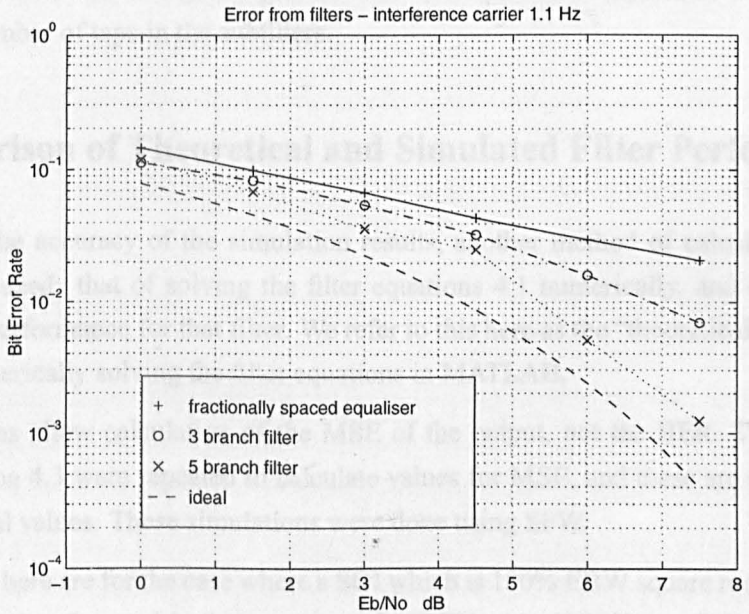

Figure 4.8: Interferer carrier offset 1.1

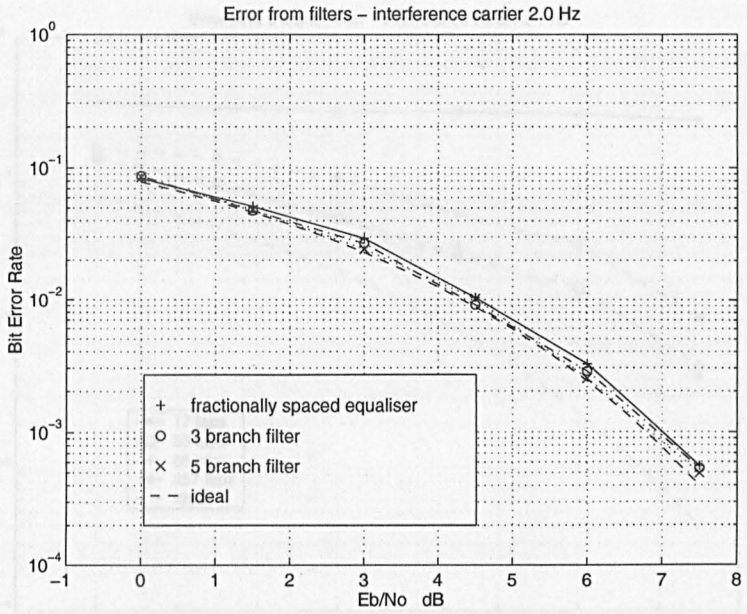

Figure 4.9: Interferer carrier offset 2.0

the particular tap configuration which existed when the adaptation was frozen. This configuration will be slightly different from the optimum and the effect of the error increases with the number of taps, as the taps at the end of the filter tend to have (ideally) values close to zero, and the missing taps in a short filter may be closer to the exact solution than the adaptation noise value left on the taps of a long filter. This explains why the 257×5 tap filter performs worse than the 65×5 tap filter. Note that the amount of adaptation noise is dependent on the value of the feedback coefficient in the LMS algorithm. Here a value of 10^{-5} was used. A larger coefficient, which gives faster adaptation, will also give a greater level of adaptation noise and so would optimally need a smaller number of taps in the subfilters.

4.5 Comparison of Theoretical and Simulated Filter Performance

To help confirm the accuracy of the simulation results, another method of calculating the filter performance was used: that of solving the filter equations 4.1 numerically, and calculating the theoretical MSE performance for that filter. We refer to this here as the "theoretical" result. It was calculated by numerically solving the filter equations in MATLAB.

The filter equations allow calculation of the MSE of the output, not the BER. The simulations described in section 4.3 were repeated to calculate values for MSE, and these are compared here with the theoretical values. These simulations were done using SPW.

The results shown here are for the case where a SOI which is 100% EBW square root raised cosine filtered QPSK is being detected in the presence of AWGN and a QPSK interferer, which is also 100% EBW square root raised cosine filtered. The SOI has a baud rate of 1 and a carrier frequency of 0; the interferer has a baud rate of $0.9412... = \frac{16}{17}$, and a carrier frequency of 0.7. The power of the interferer was 10 times that of the SOI, and the AWGN was set to give a SOI E_b/N_0 of 7.5 dB. The baud rate of the interferer was chosen to be 0.9412 because this is not harmonically related

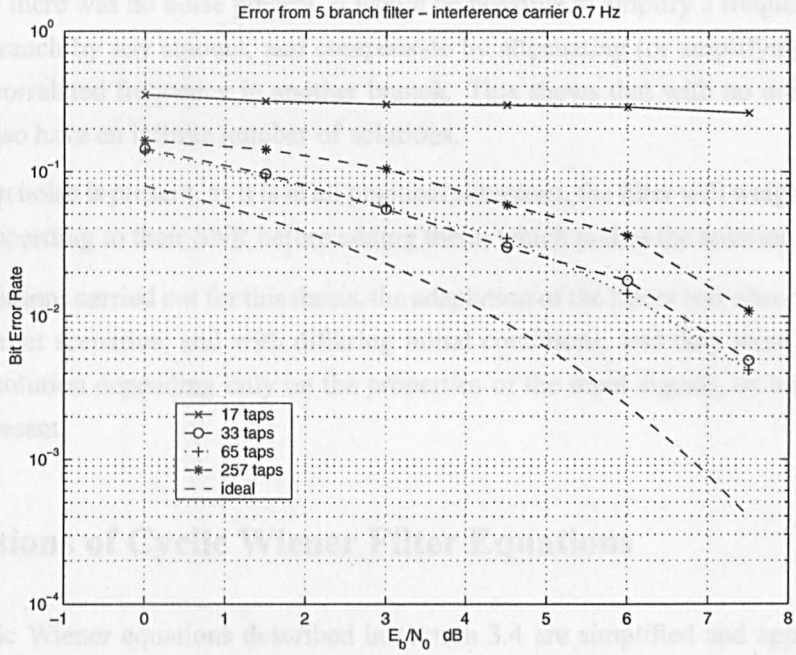

Figure 4.10: Effect of number of taps on FRESH filter performance

to the SOI baud rate but can be represented by an integer number of samples per symbol. SPW cannot represent symbols with non-integer sample numbers. The sampling rates used here were 16 samples per symbol for the SOI and 17 samples per symbol for the interferer.

The reference (or desired) signal in each case was 100% EBW raised cosine filtered QPSK, so the FRESH filter is expected to perform the role of matched filtering as well as interference rejection.

Some assumptions were made in solving the filter equations: the data carried by the SOI and by the interferer (if present) are both assumed to be perfectly white sequences. There is also assumed to be no correlation between the two data sequences.

The solution of the filter equations for the fractionally spaced equaliser, the 3-branch FRESH filter (with frequency shifts matching the SOI baud rate) and the 5-branch FRESH filter (with frequency shifts matching the SOI and interferer baud rates) are described in sections 3.3, 4.7.1 and 4.7.2.

4.6 Uniqueness of Filter Solutions

An issue of relevance here, which has not yet been thoroughly investigated, is that there is not necessarily a unique solution to the fractionally spaced equaliser, or FRESH filter equations. This can cause problems in both the numeric solution and in the adaptive implementation of the filter.

It is known that a FSE does not have a unique solution if the tap spacing is such that the filter bandwidth is greater than that of the signal being filtered (as it usually is) and if there is no noise present in the out of band section of the filter bandwidth [34]. The second condition means that in most practical applications there is a unique solution. An infinite set of solutions exists when these two conditions are met because the filter could arbitrarily amplify any of the out-of-band portion of its own bandwidth without affecting the output signal.

With FRESH filters the situation is more complicated because several versions of the signal exist

in the filter. If there was no noise present, it would be possible to amplify a frequency component in one filter branch by any amount, and compensate by attenuating (or amplifying and reversing the phase) a correlated frequency in another branch. This shows that with no noise, the FRESH filter would also have an infinite number of solutions.

However when noise is present, as it is in all practical situations, the filter will weight the correlated frequencies according to their SNR before adding them, which makes the solution unique.

In the investigations carried out for this thesis, the adaptation of the filters was observed graphically in many different scenarios, and with differing initial conditions, and they were seen to always adapt to one solution depending only on the properties of the input signals, as long as there was some noise present.

4.7 Solutions of Cyclic Wiener Filter Equations

Here the cyclic Wiener equations described in section 3.4 are simplified and applied to specific FRESH filters to allow comparison with simulation results.

4.7.1 3-branch FRESH filter

Gardner's cyclic Wiener filter equations are:

$$\begin{aligned}
&\sum_{m=1}^{M} S_x^{\alpha_k-\alpha_m}\left(f-\frac{\alpha_k+\alpha_m}{2}\right)A_m(f)+\sum_{n=1}^{N} S_{xx^*}^{\beta_n-\alpha_k}\left(f-\frac{\alpha_k+\beta_n}{2}\right)^* B_n(f)=S_{dx}^{\alpha_k}, \quad k=1,2,...M\\
&\sum_{m=1}^{M} S_{xx^*}^{\beta_k-\alpha_m}\left(f-\frac{\beta_k+\alpha_m}{2}\right)A_m(f)+\sum_{n=1}^{N} S_x^{\beta_k-\beta_n}\left(-f+\frac{\beta_k+\beta_n}{2}\right)B_n(f)=S_{dx^*}^{\alpha_k}, \quad k=1,2,...N
\end{aligned} \tag{4.1}$$

where S_x^α is the SCD of the input signal x with a frequency shift of α, $S_{xx^*}^\alpha$ is the conjugate SCD of x with a frequency shift of β, A_m and B_n are the subfilter frequency responses corresponding the the filters labelled h_α and h_β in figure 3.6. S_{dx}^α is the cross-spectral density of the reference signal d and the input x under a frequency shift of α.

These can be simplified for the specific example under consideration here by noting that the signal of interest is QPSK which exhibits no conjugate spectral correlation. This means that $N = 0$ and the second equation of 4.1 disappears. It also means that $S_{xx^*}^\gamma = 0$ for any frequency shift γ, so the second term of the first equation disappears. We have a filter containing three frequency shifts (including the zero shift of the central branch) so $M = 3$ and the equation represents a system of three equations (due to the $k = 1, 2, \ldots M$ condition), each with three terms on the left hand side (due to the summation $\sum_{m=1}^{M}$). This filter is the middle one shown in figure 4.1.

Equations 4.1 therefore simplify to

$$\begin{aligned}
&S_x^{\alpha_1-\alpha_1}\left(f-\frac{\alpha_1+\alpha_1}{2}\right)A_1(f)+S_x^{\alpha_1-\alpha_2}\left(f-\frac{\alpha_1+\alpha_2}{2}\right)A_2(f)+S_x^{\alpha_1-\alpha_3}\left(f-\frac{\alpha_1+\alpha_3}{2}\right)A_3(f)=S_{dx}^{\alpha_1}\\
&S_x^{\alpha_2-\alpha_1}\left(f-\frac{\alpha_2+\alpha_1}{2}\right)A_1(f)+S_x^{\alpha_2-\alpha_2}\left(f-\frac{\alpha_2+\alpha_2}{2}\right)A_2(f)+S_x^{\alpha_2-\alpha_3}\left(f-\frac{\alpha_2+\alpha_3}{2}\right)A_3(f)=S_{dx}^{\alpha_2}\\
&S_x^{\alpha_3-\alpha_1}\left(f-\frac{\alpha_3+\alpha_1}{2}\right)A_1(f)+S_x^{\alpha_3-\alpha_2}\left(f-\frac{\alpha_3+\alpha_2}{2}\right)A_2(f)+S_x^{\alpha_3-\alpha_3}\left(f-\frac{\alpha_3+\alpha_3}{2}\right)A_3(f)=S_{dx}^{\alpha_3}
\end{aligned}$$

(4.2)

The $\alpha_1, \alpha_2, \alpha_3$ represent the frequency shifts of $\alpha, 0, -\alpha$ where α is the baud rate of the signal of interest. Substituting these values into equation 4.2 gives

$$\begin{aligned} S_x^0(f+\alpha)A_1(f)+ \quad & S_x^{-\alpha}(f+\tfrac{\alpha}{2})A_2(f)+ \quad & S_x^{-2\alpha}(f)A_3(f) = \quad & S_{dx}^{-\alpha}\left(f+\tfrac{\alpha}{2}\right) \\ S_x^{\alpha}(f+\tfrac{\alpha}{2})A_1(f)+ \quad & S_x^{0}(f)A_2(f)+ \quad & S_x^{-\alpha}(f-\tfrac{\alpha}{2})A_3(f) = \quad & S_{dx}^{0}(f) \\ S_x^{2\alpha}(f)A_1(f)+ \quad & S_x^{\alpha}\left(f-\tfrac{\alpha}{2}\right)A_2(f)+ \quad & S_x^{0}(f-\alpha)A_3(f) = \quad & S_{dx}^{\alpha}\left(f-\tfrac{\alpha}{2}\right) \end{aligned}$$

(4.3)

As the signal of interest is 100% EBW we can say $S_x^{2\alpha} \equiv S_x^{-2\alpha} \equiv 0$, because x shifted by $+2\alpha$ or -2α has none of its spectrum overlapping x (unshifted). From symmetry we know that $S_x^{\alpha} \equiv S_x^{-\alpha}$, so finally we have

$$\begin{aligned} S_x^0(f+\alpha)A_1(f)+ \quad & S_x^{\alpha}(f+\tfrac{\alpha}{2})A_2(f)+ \quad & 0 = \quad & S_{dx}^{-\alpha}\left(f+\tfrac{\alpha}{2}\right) \\ S_x^{\alpha}(f+\tfrac{\alpha}{2})A_1(f)+ \quad & S_x^{0}(f)A_2(f)+ \quad & S_x^{\alpha}(f-\tfrac{\alpha}{2})A_3(f) = \quad & S_{dx}^{0}(f) \\ 0+ \quad & S_x^{\alpha}\left(f-\tfrac{\alpha}{2}\right)A_2(f)+ \quad & S_x^{0}(f-\alpha)A_3(f) = \quad & S_{dx}^{\alpha}\left(f-\tfrac{\alpha}{2}\right) \end{aligned}$$

(4.4)

The subfilter frequency responses are therefore defined in 4.4 in terms of the cyclic spectral density of the input signal and the cross spectral density of the desired signal and the input signal, with various relative frequency shifts. These spectral densities are defined by Gardner as follows:

$$S_x^{\alpha}(f) = \lim_{T\to\infty} T\langle X_T(t, f+\tfrac{\alpha}{2})X_T^*(t, f-\tfrac{\alpha}{2})\rangle$$

and

$$S_{xy}^{\alpha}(f) = \lim_{T\to\infty} T\langle X_T(t, f+\tfrac{\alpha}{2})Y_T^*(t, f-\tfrac{\alpha}{2})\rangle \tag{4.5}$$

for any two signals x, y, where $\langle . \rangle$ indicates averaging over all time T. The more general form, which does not require ergodic signals, is:

$$S_x^{\alpha}(f) = E[X_T(t, f+\tfrac{\alpha}{2})X_T^*(t, f-\tfrac{\alpha}{2})]$$

and

$$S_{xy}^{\alpha}(f) = E[X_T(t, f+\tfrac{\alpha}{2})Y_T^*(t, f-\tfrac{\alpha}{2})] \tag{4.6}$$

The signals (e.g. X_T) are the product of the pulse shape and the data spectrum. We are assuming that the data are a white noise sequence, so its Fourier transform has equal magnitude over the whole spectrum. The cyclic spectral density functions 4.6 then simplify to:

$$S_x^\alpha(f) = g(f + \frac{\alpha}{2})g^*(f - \frac{\alpha}{2})$$

and

$$S_{xy}^\alpha(f) = g(f + \frac{\alpha}{2})h^*(f - \frac{\alpha}{2}) \tag{4.7}$$

for values of α which are cyclic frequencies of x or joint cyclic frequencies of x and y, and 0 for all other α, where g and h are the pulse shaping functions applied to x and y.

The calculation of all the cyclic spectral density functions in the system of equations 4.4 involves multiplying frequency shifted versions of the Fourier transforms of the pulse shapes of the signals involved. It is assumed that all the data are a white noise sequence (on the SOI and the interferer) and that the SOI data have no correlation with the interferer data. This means that in many of the terms, the cross correlations are simplified because there is no correlation between most of the signals. For example, the desired signal is correlated with the signal of interest but not with the interferer or the white noise. This means that S_{dx}^0 is calculated by multiplying the spectra of the SOI and the desired signal. Under a relative frequency shift of α this correlation still stands, so the shifted spectra are multiplied. However under a relative frequency shift of γ there is no correlation between these signals so S_{dx}^γ is zero.

It can be seen from figures 4.11 and 4.12 that the MSE predicted from the two methods is very close. The simulation value is always slightly higher which is to be expected due to the slight tap misadjustment which is inevitable in an adaptive system. This error can be reduced by using a smaller value of the LMS feedback coefficient, although the filter then adapts more slowly. Figure 4.11 shows the predicted MSE in receiving the SOI waveform in the presence of white noise only. Figure 4.12 shows the same in the presence of an interferer with a carrier frequency of 0.7. Figures 4.13 and 4.14 show the filter magnitude frequency responses for each subfilter of the FRESH filter corresponding to the 2 scenarios described. The SNR used for these filter solutions was $E_b/N_0 = 7.5$ dB.

In each of figure 4.13 and 4.14 the top graph shows the responses of the top subfilter of the middle FRESH filter in figure 4.1. This subfilter is preceded by a frequency shift of α, the SOI baud rate). The next two graphs correspond to the other filter responses in the obvious way.

From figure 4.13 one can see how the FRESH filter improves the MSE in white noise only. The energy from the outer half of the spectrum is shifted to correlated frequencies, and the 3 subfilters weight the components according to their SNR. Adding the outputs of these filters results in correlated SOI energy and uncorrelated noise being added, so the SNR is increased.

There is a small difference visible in the two solutions, presumably due to the fact that the analytic solution was calculated purely in the frequency domain where zero gain (e.g. in the third subfilter for $f < 0$) is easy to achieve. The simulated solution is the Fourier transform of a finite length time domain filter so it is impossible to achieve such a sharp drop to zero gain. Using a longer filter (257 taps were used in this case) would result in a closer match between the two solutions.

In figure 4.14 there is a high powered interferer at a carrier frequency of +0.7 Hz present. The third subfilter is amplifying energy from the lower frequency part of the SOI spectrum, which is unaffected by the interferer. The correlated SOI energy in the upper half of the spectrum is

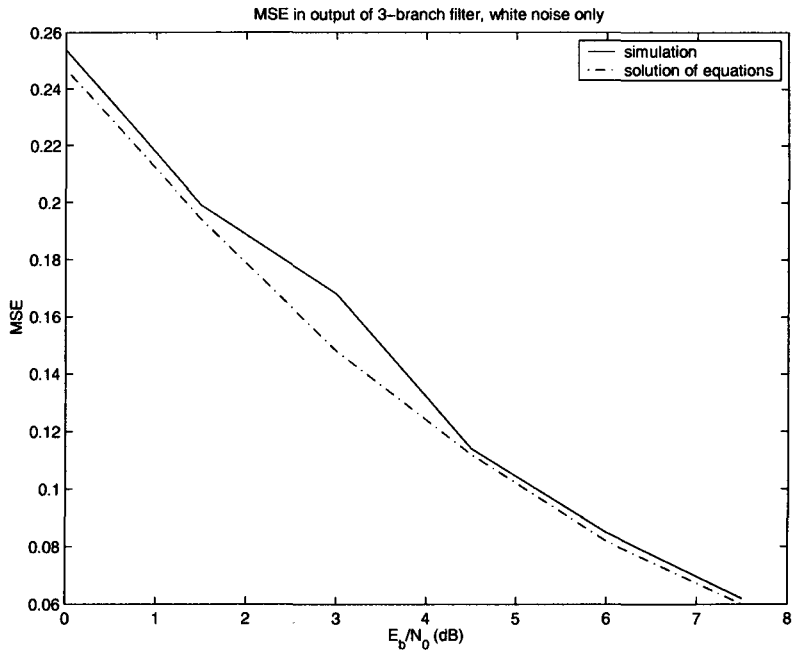

Figure 4.11: Comparison of two methods of calculating 3-branch filter MSE in white noise only

attenuated by the subfilter, along with the interferer. The top subfilter passes very little as it can only manipulate SOI energy which is severely corrupted by the interferer.

Notice that the numerically calculated responses for the subfilters are all limited in bandwidth to 1. This is due to the fact that the total signal bandwidth is 2 and the frequency shifts are 1. This is linked with the issues of efficiency of implementation raised in section 3.6.

4.7.2 5-branch FRESH filter

The FRESH filter in the previous section was able to exploit only the spectral correlation of the SOI. In this section a 5-branch filter is considered which allows exploitation of both SOI and interferer spectral correlation

Again we start from equation 4.1. We are assuming in this case that an interferer is present and that the filter contains branches with frequency shifts equal to γ (the baud rate of the interferer) and $-\gamma$. The SOI and interferer are both QPSK which means that again the filter equations simplify as there is no conjugate correlation in the input signal. Here, $M = 5$ so the filter equations are

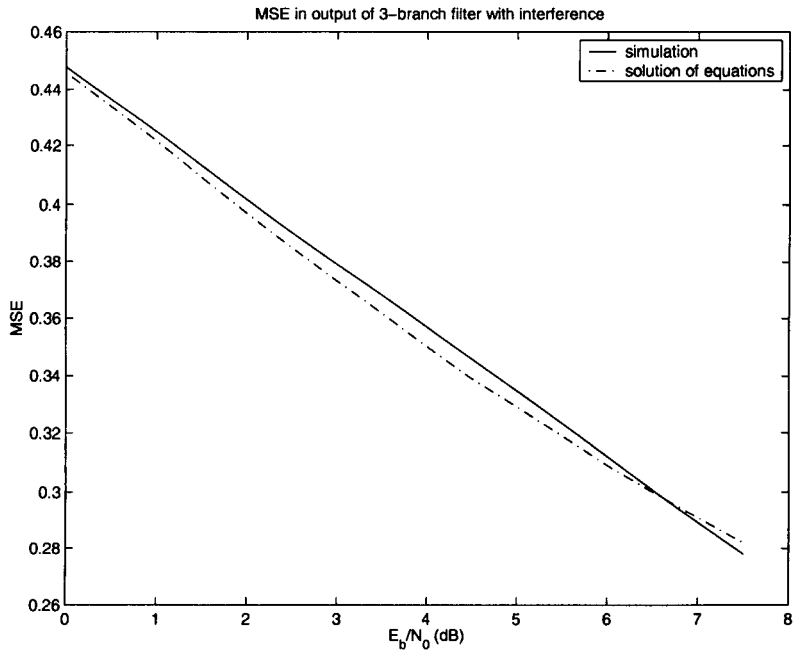

Figure 4.12: Comparison of two methods of calculating 3-branch filter MSE in interference and white noise

$$
\begin{aligned}
&S_x^0(f+\gamma)A_1(f) \quad +S_x^{-\gamma+\alpha}(f+\tfrac{\gamma+\alpha}{2})A_2(f) \quad +S_x^{\gamma}(f+\tfrac{\gamma}{2})A_3(f) \quad +S_x^{-\gamma-\alpha}(f+\tfrac{\gamma-\alpha}{2})A_4(f)\\
&\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad +S_x^{-2\gamma}(f)A_5(f) \quad = S_{dx}^{-\gamma}(f+\tfrac{\gamma}{2})\\
&S_x^{-\alpha+\gamma}(f+\tfrac{\gamma+\alpha}{2})A_1(f) \quad +S_x^0(f+\alpha)A_2(f) \quad +S_x^{-\alpha}(f+\tfrac{\alpha}{2})A_3(f) \quad +S_x^{-2\alpha}(f)A_4(f)\\
&\qquad\qquad\qquad\qquad\qquad\qquad +S_x^{-\alpha-\gamma}(f+\tfrac{\alpha-\gamma}{2})A_5(f) \quad = S_{dx}^{-\alpha}(f+\tfrac{\alpha}{2})\\
&S_x^{\gamma}(f+\tfrac{\gamma}{2})A_1(f) \quad +S_x^{\alpha}(f+\tfrac{\alpha}{2})A_2(f) \quad +S_x^0(f)A_3(f) \quad +S_x^{-\alpha}(f-\tfrac{\alpha}{2})A_4(f)\\
&\qquad\qquad\qquad\qquad\qquad\qquad +S_x^{-\gamma}(f-\tfrac{\gamma}{2})A_5(f) \quad = S_{dx}^0(f)\\
&S_x^{\alpha+\gamma}(f+\tfrac{\gamma-\alpha}{2})A_1(f) \quad +S_x^{2\alpha}(f)A_2(f) \quad +S_x^{\alpha}(f-\tfrac{\alpha}{2})A_3(f) \quad +S_x^0(f-\alpha)A_4(f)\\
&\qquad\qquad\qquad\qquad\qquad\qquad +S_x^{\alpha-\gamma}(f-\tfrac{\alpha+\gamma}{2})A_5(f) \quad = S_{dx}^{\alpha}(f-\tfrac{\alpha}{2})\\
&S_x^{2\gamma}(f)A_1(f) \quad +S_x^{\gamma+\alpha}(f+\tfrac{\alpha-\gamma}{2})A_2(f) \quad +S_x^{\gamma}(f-\tfrac{\gamma}{2})A_3(f) \quad +S_x^{\gamma-\alpha}(f-\tfrac{\gamma+\alpha}{2})A_4(f)\\
&\qquad\qquad\qquad\qquad\qquad\qquad +S_x^0(f-\gamma)A_5(f) \quad = S_{dx}^{\gamma}(f-\tfrac{\gamma}{2})
\end{aligned}
\tag{4.8}
$$

Again we can make further simplifications by noticing that many of the cross correlation terms on the left hand side are zero. $S_x^{2\alpha}$ and $S_x^{-2\alpha}$ are zero because under a frequency shift of 2α there is no correlation in the input signal due to the limited bandwidth of the SOI. Similarly $S_x^{2\gamma}$ and $S_x^{-2\gamma}$ are zero because of the limited bandwidth of the interferer. There is some overlap between the SOI and the interferer under these frequency shifts but these signals are not correlated. There is also no correlation in the SOI except at $\pm\alpha$ or in the interferer at $\pm\gamma$ so the cross correlations with frequency shifts of $\alpha+\gamma$, $\alpha-\gamma$, $-\alpha+\gamma$ and $-\alpha-\gamma$ are all zero. Also the cross correlation of the desired signal and the SOI is zero except under frequency shifts of 0 or $\pm\alpha$, so the terms S_{dx}^{γ} and $S_{dx}^{-\gamma}$ are also zero. Finally we can also say from symmetry that $S_x^{\alpha} = S_x^{-\alpha}$ and $S_x^{\gamma} = S_x^{-\gamma}$.

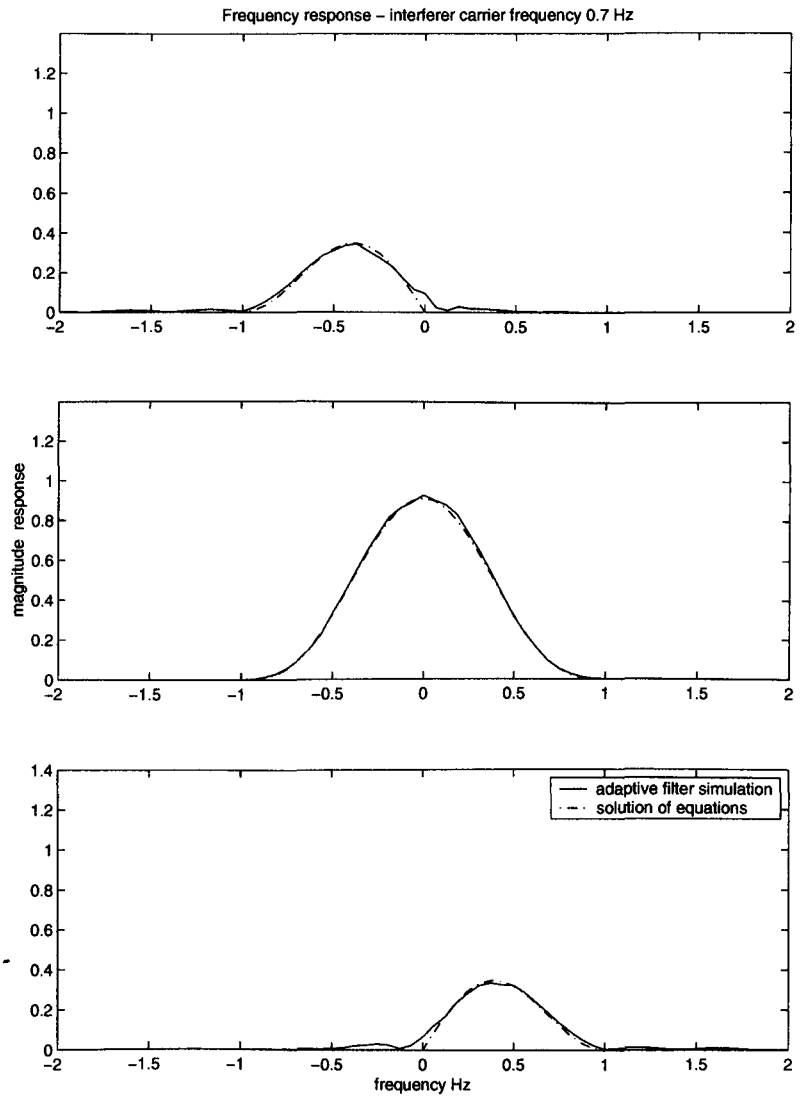

Figure 4.13: Magnitude frequency response of 3-branch FRESH filter with no interference

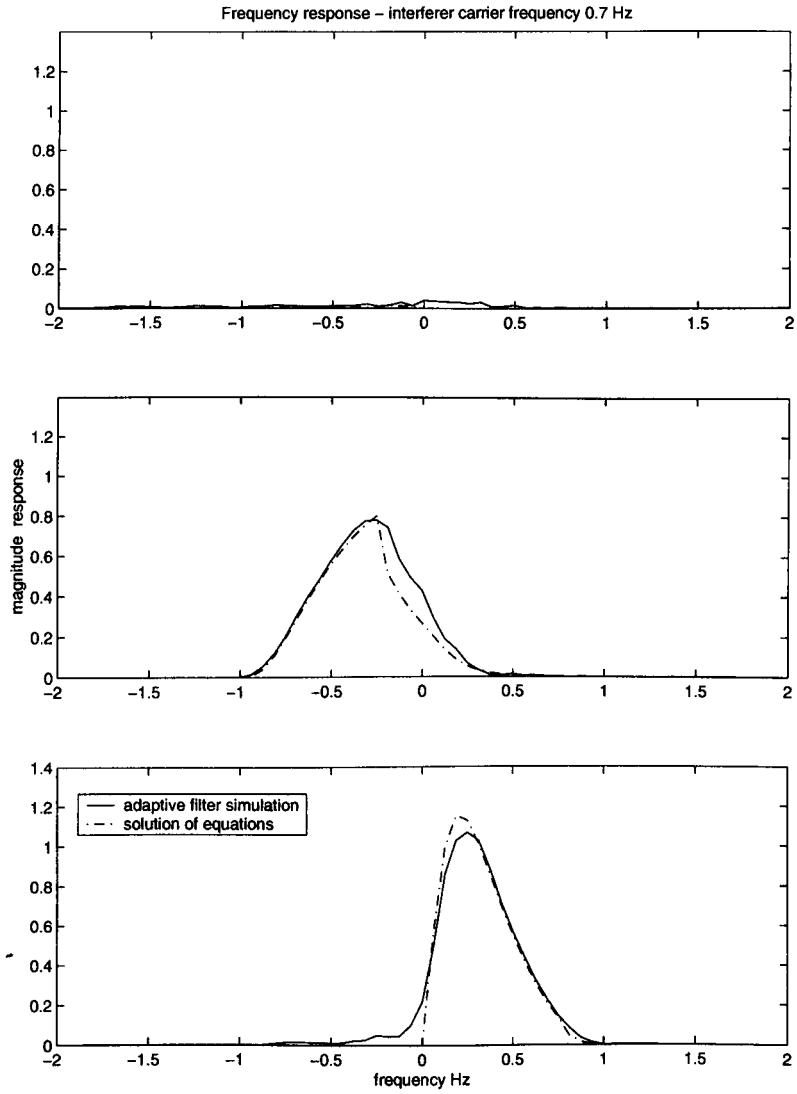

Figure 4.14: Magnitude frequency response of 3-branch FRESH filters with interferer carrier 0.7

Equations 4.8 then simplify to

$$\begin{array}{ccccc}
S_x^0(f+\gamma)A_1(f) & +0 & +S_x^\gamma(f+\frac{\gamma}{2})A_3(f) & +0 & +0=0 \\
0 & +S_x^0(f+\alpha)A_2(f) & +S_x^\alpha(f+\frac{\alpha}{2})A_3(f) & +0 & +0=S_{dx}^{-\alpha}(f+\frac{\alpha}{2}) \\
S_x^\gamma(f+\frac{\gamma}{2})A_1(f) & +S_x^\alpha(f+\frac{\alpha}{2})A_2(f) & +S_x^0(f)A_3(f) & +S_x^\alpha(f-\frac{\alpha}{2})A_4(f) & +S_x^\gamma(f-\frac{\gamma}{2})A_5(f) \\
 & & & & =S_{dx}^0(f) \\
0 & +0 & +S_x^\alpha(f-\frac{\alpha}{2})A_3(f) & +S_x^0(f-\alpha)A_4(f) & +0=S_{dx}^\alpha(f-\frac{\alpha}{2}) \\
0 & +0 & +S_x^\gamma(f-\frac{\gamma}{2})A_3(f) & +0 & +S_x^0(f-\gamma)A_5(f)=0
\end{array} \tag{4.9}$$

These equations determine the five filter responses in terms of the cross spectral densities, with frequency shifts of the SOI with itself and SOI with the desired.

The mean squared error which results from filtering the input x with such a filter can be calculated from

$$MSE = \int_{-\infty}^{\infty} S_e(f)df \tag{4.10}$$

where $S_e(f)$ is the power spectral density of the error signal. This is given by

$$S_e(f) = S_d(f) - \sum_{m=1}^{5} S_{dx}^{\alpha_m}\left(f - \frac{\alpha_m}{2}\right) A_m(f) \tag{4.11}$$

As with the 3-branch filter, there is good agreement between theoretical and simulated MSE of the filter output, as shown in figure 4.15. The filter frequency responses in figure 4.16 show that the subfilters following the SOI related frequency shifts behave in a similar way to those in the 3-branch filter (figure 4.14) but that there is an additional contribution from the filter following the interferer related shift (the top graph of figure 4.16). This shows that the interferer is correlated under this frequency shift with the unshifted signal in the central branch, and the two are being weighted so their magnitudes match as closely as possible. There will also be a π radian phase shift between the two filter responses so that the interferer components are subtracted instead of added.

4.8 Summary and Further Work

It has been shown that by taking advantage of the spectral correlation in an interferer, a significant improvement is made in the performance of the FRESH filter. This filter can then effectively remove the effects of the interference. The only knowledge assumed in the construction of the filter is the symbol rates of the SOI and interferer. If this knowledge is available then this filter can offer extremely effective interference rejection performance.

A number of issues arise from this work. In the simulations, exact knowledge of the interferer symbol rate was assumed. Before practical use of this technique is considered, it is necessary to examine how any errors in this information affect the filter performance. It is expected that this is related to the feedback coefficient of the filter adaptation algorithm which is in effect the rate

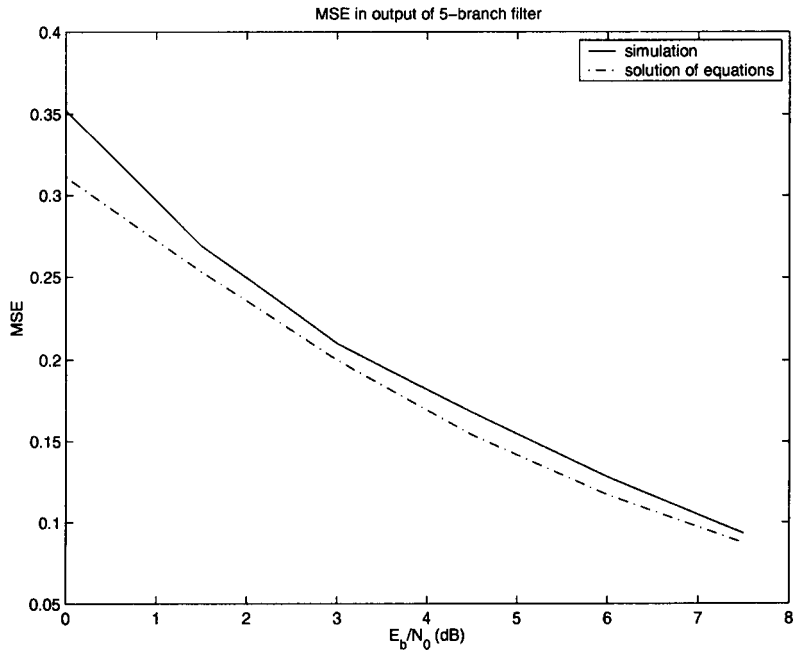

Figure 4.15: Comparison of two methods of calculating 5-branch filter MSE

of adaptation of the filter. The assumption was also made that the interference lies completely within the bandwidth of the receiver. The rejection of the interference will be affected if all of its spectrum is not available to the filter. It would be useful to look into how partial reception of the interferer affects performance. Related to this is the problem of a high powered interferer, limiting the effective dynamic range of a receiver available to the signal of interest. An important area of study is how to extract the useful information from the interferer without saturating the receiver

Numerical solutions to the cyclic Wiener filter equations have been demonstrated, and compared with solutions obtained by training the filter in simulations using the LMS algorithm. The resulting MSE values have also been compared. There is generally good agreement between these two methods of calculating the filter responses which confirms that the filters are operating in simulation as intended by the theory. The subfilter responses shown are useful for visualising how the filters manipulate the signals to remove the interference.

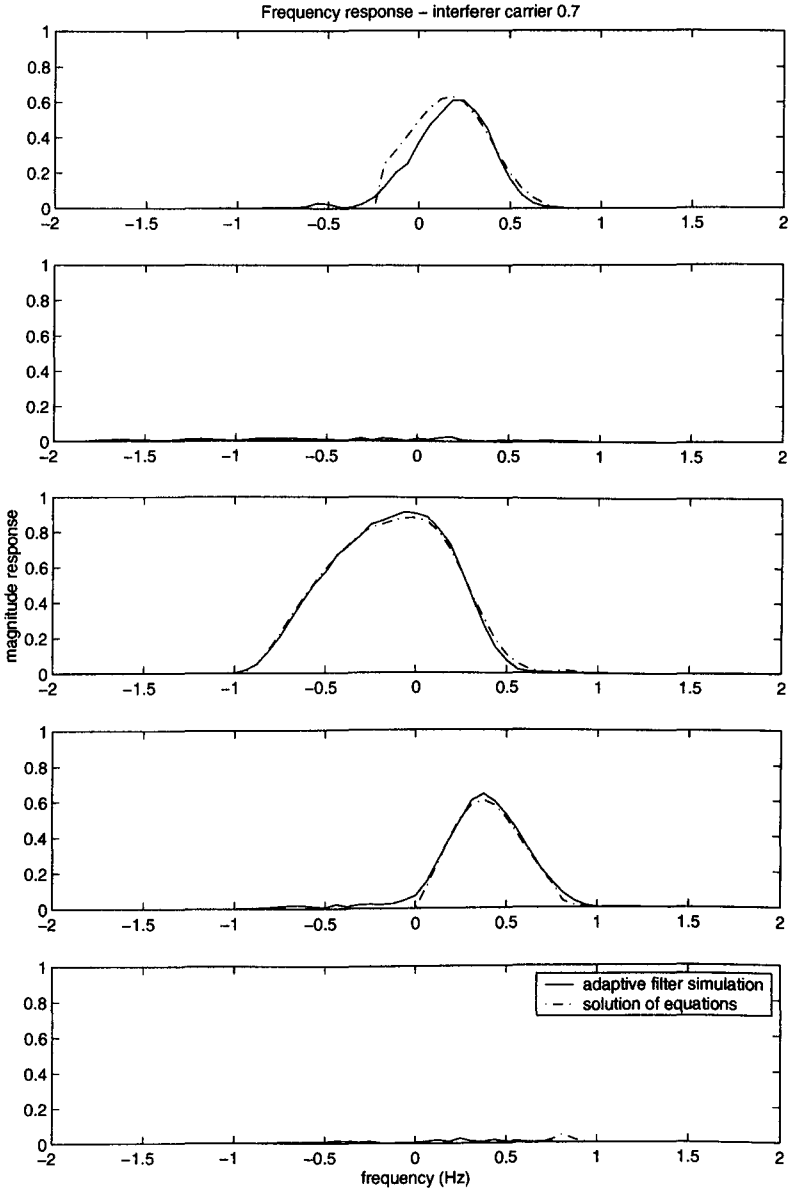

Figure 4.16: Magnitude response of 5-branch FRESH filter with interferer carrier offset 0.7 Hz

Chapter 5

Spectral Shaping to Enhance Cyclostationarity

Contents

5.1 Introduction

It is clear from chapter 2 that the spectral correlation properties of a signal are strongly affected by the pulse shape, or equivalently, the shape of the spectrum. As the FRESH filter exploits spectral correlation, it is reasonable to assume that in a given interference scenario, the effectiveness of the FRESH filter, and hence the immunity of the system to interference, is dependent on the spectrum of the SOI. It will also, of course, be dependent on the spectrum of the interferer, but it is assumed here that the interference is from another communication system, and so its spectrum cannot be changed.

So an investigation was carried out on the possibility of improving the interference immunity of the SOI by enhancing its cyclostationary properties. It was hoped that, for example, by taking a signal with 100% excess bandwidth (EBW) and shaping the spectrum so that the energy was equally distributed over the whole bandwidth, a FRESH filter would be able to make better use of the spectral correlation, and so would be more effective against an interferer. However the results reported here suggest that this is not a useful technique in general.

5.2 Simulation Description

Figure 5.1 shows three spectra which are idealised representations of the spectra of the SOI used in the simulations reported here. One is a signal with no excess bandwidth, which has a perfectly rectangular spectrum, another has a 100% EBW raised cosine spectrum and the third has

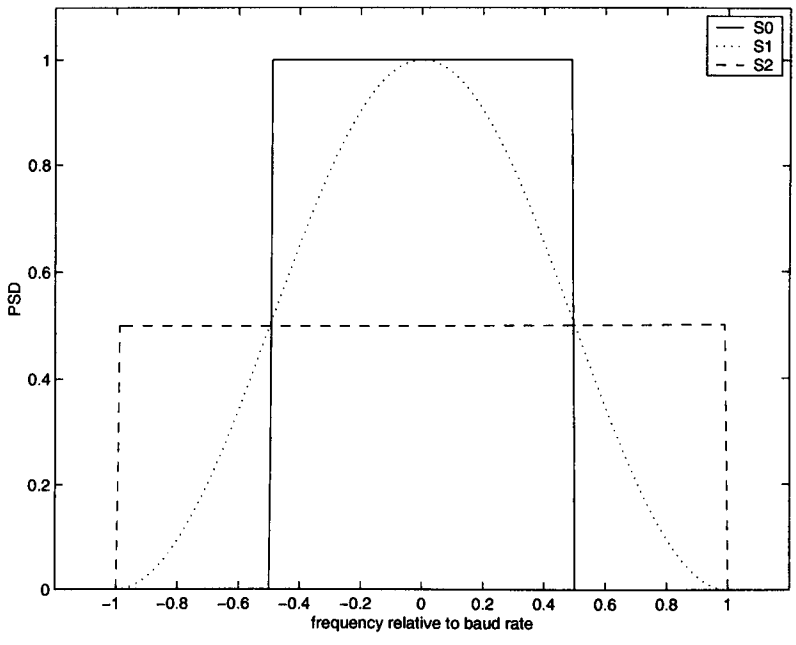

Figure 5.1: PSD of the three SOI spectra tested

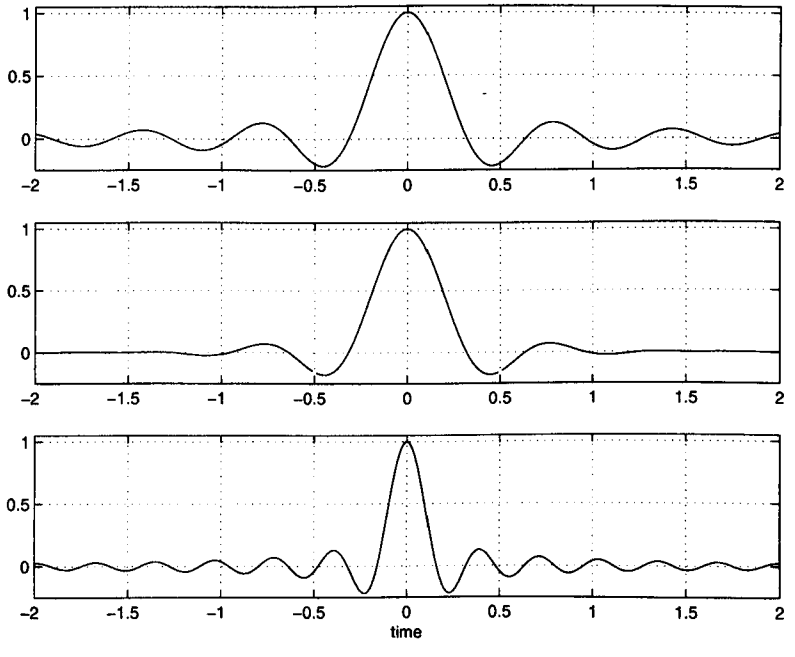

Figure 5.2: Pulse shapes of the three SOI tested (from top: S0, S1, S2)

100% EBW with a perfectly rectangular spectrum. For convenience, these three spectral shapes are referred to as S0, S1 and S2 respectively. With the powers shown in figure 5.1 these signals have the same energy per bit transmitted.

The corresponding pulse shapes for the three spectra are shown in figure 5.2. It can be seen that the shapes of S0 and S2 are similar, but with S2 (100% EBW) oscillating at twice the frequency of the S0 pulse. All the pulses have zeros at, at least, symbol spaced intervals, and so are free of intersymbol interference. It was hoped that using the rectangular signal spectrum S2 would allow the equaliser to reject the interference more effectively. Because the energy in the signal is spread evenly across the band, there is more available in the edges of the spectrum to replace any corrupted signal energy in the centre of the band.

To produce a perfectly rectangular spectrum as in S0 and S2 is impossible as it requires an infinitely long pulse in the time domain. Strictly speaking the raised cosine pulse is also not exactly obtainable with a finite pulse, but a practical realisation of a very close approximation is possible, as the magnitude of the pulse falls off rapidly with time.

It is clear from figure 5.2 that truncated realisations of the three pulse shapes will be less accurate for S0 and S2 than S1 because the impulse response of S2 tends to be closer to 0 for times away from the centre of the pulse. The actual pulse shapes used were such truncations; no smoothing window was used. The time span shown in figure 5.2 corresponds to a 65 tap filter at 16 samples per symbol. In all the simulations described this was the sampling frequency used, and the number of taps in the transmit, or pulse shaping, filters of the SOI and interferer. The number of taps in each branch of the FRESH filter was varied, but was one of 33, 65 and 257.

To illustrate that S0 does not contain spectral correlation but that S1 and S2 both do, the SCDs and conjugate SCDs of the signals are shown here. In each of these graphs (figures 5.3 to 5.8) there is some low level noise visible. This is a result of producing these graphs by simulation with a finite amount of data. The noise can be reduced to an arbitrarily small level by increasing the length of the simulation. Figure 5.3 shows that there is no correlation at any frequency shift from 0 to 2 times the baud rate with S0. This is as expected as the Nyquist bandwidth is the minimum required to transmit the data so there can be no redundancy in the signal. S1 contains a small amount of correlation at a frequency shift of 1 (figure 5.5). The magnitude of the graphs are related to the product of the two frequency shifted components used for each line on the frequency shift axis, so it gives an indication of proportion of the signal power used in providing diversity. The S2 graph shows the highest correlation level of the three signals. As there is no correlation between the real and imaginary parts of a QPSK signal, the conjugate SCDs are all zero for all frequency shifts.

The interference considered here is square root raised cosine filtered QPSK with 100% EBW. It has power equal to that of the SOI. A range of scenarios were tested when the interferer baud rate is either the same or slightly different to that of the SOI. The interferer carrier frequency was varied from 0 to 2 times the SOI baud rate although in most cases 0.7 was used. This means that the signals' spectra are largely overlapping as shown in figures 5.9, 5.10 and 5.11. The powers of the SOI and the interferer were equal.

The filter used to remove the interference was a 3-branch FRESH filter (as in figure 3.7) with frequency shifts equal to ± 1 (the symbol rate of the SOI was 1). This filter is therefore only capable of exploiting SOI spectral correlation unless the interferer symbol rate is also 1.

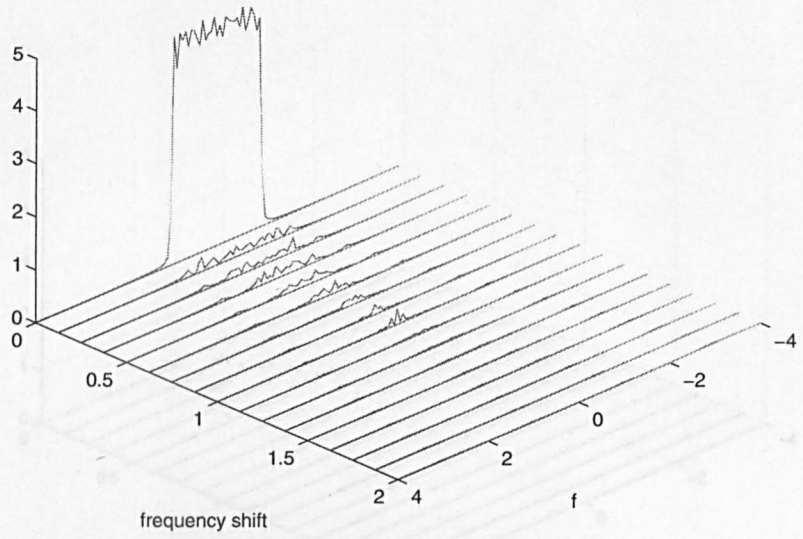

Figure 5.3: SCD of 0% excess bandwidth baseband QPSK

A simplified block diagram of the simulation system is shown in figure 5.12. The transmit filters for both the SOI and interferer had 65 taps. The sampling frequency was 16 samples per symbol and the FRESH filter was trained using the LMS algorithm.

The training signal used in the FRESH filter carried the same data as the SOI, but was 100% EBW raised cosine filtered (as opposed to square root raised cosine). This means that the FRESH filter should adapt towards the cyclic Wiener filter for the transmitted signal in the presence of the interference and also perform the role of the matched filter. In a real system the adaptation is likely to use a training signal and thereafter possibly a decision-directed process.

The simulation was used to calculate the bit error rate when receiving S0, S1 and S2 in the presence of various levels of Gaussian white noise and with and without an interferer present. The signal to noise ratio was specified in terms of E_b/N_0, and took values from 0 to 8 dB. The BER results were the average of 10 runs of each scenario, where 300 errors were counted in each run.

5.3 Results and Discussion

Before considering the effect of interference, the simulation was tested with the SOI corrupted by AWGN only. From matched filter theory it is expected that in this case, the BER depends only on E_b/N_0, and that the pulse shape has no effect [79].

This is shown in figure 5.13: each line represents the BER performance for one of the chosen SOI pulse shapes. It is clear that the three spectra do not give identical performance as required by the theory. It was assumed that this was due to the truncation of the ideal impulse response in the transmit and FRESH filters. The filter had 65 taps per branch, the transmit filters had 65 taps.

This assumption was confirmed by repeating the simulation with 257 taps in each FRESH filter branch (figure 5.14). Here the difference between the three spectra is reduced. Note however that

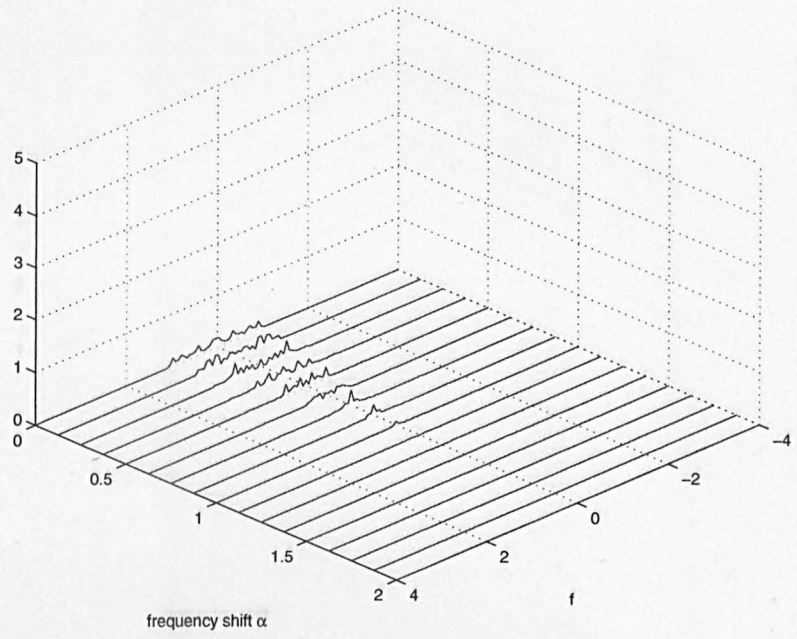

Figure 5.4: Conjugate SCD of 0% excess bandwidth baseband QPSK

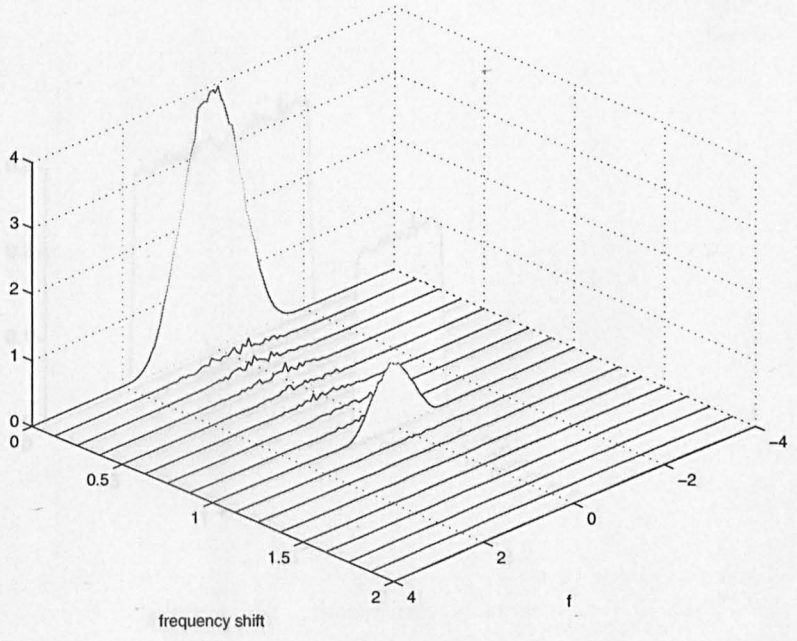

Figure 5.5: SCD of 100% excess bandwidth raised cosine filtered baseband QPSK

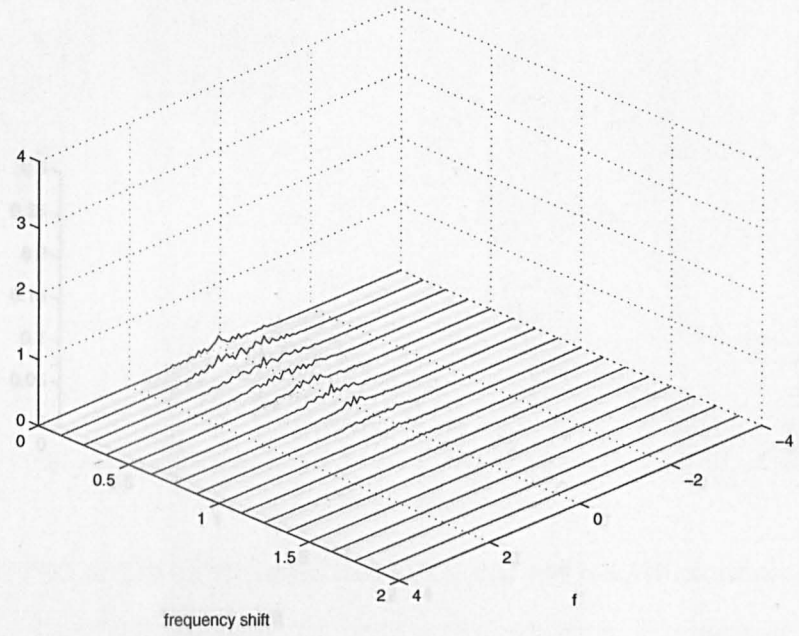

Figure 5.6: Conjugate SCD of 100% excess bandwidth raised cosine filtered baseband QPSK

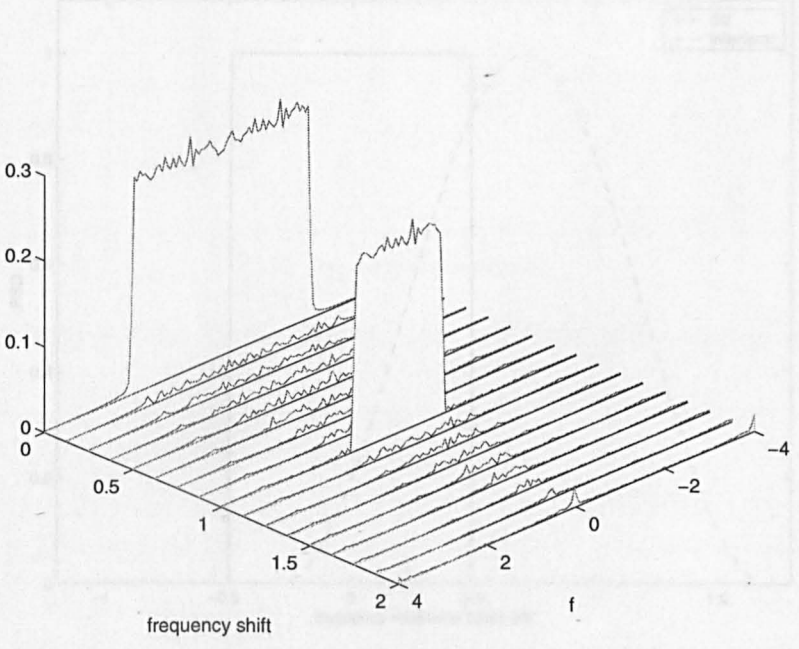

Figure 5.7: SCD of 100% excess bandwidth rectangular filtered baseband QPSK

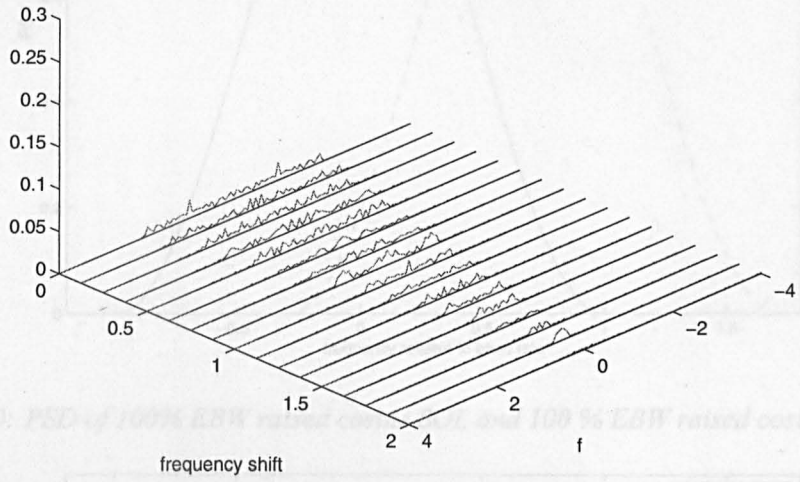

Figure 5.8: Conjugate SCD of 100% excess bandwidth rectangular filtered baseband QPSK

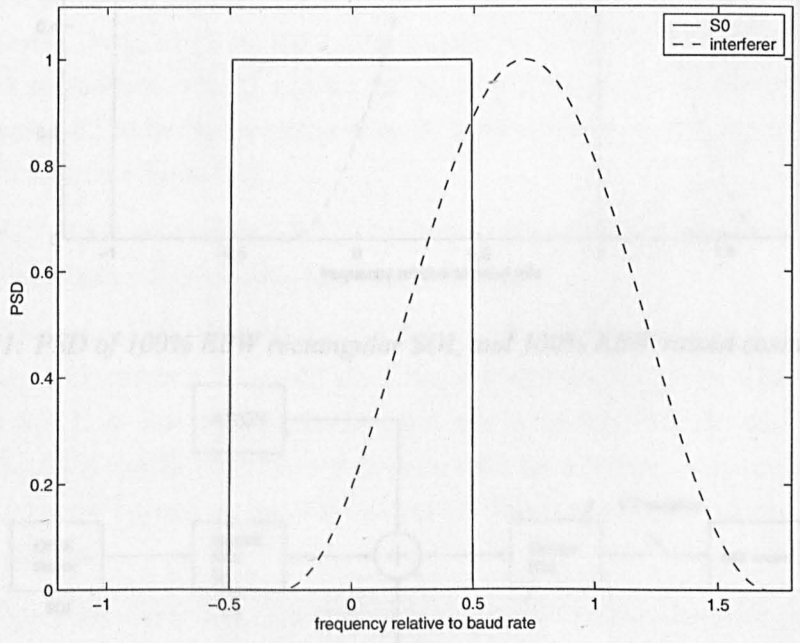

Figure 5.9: PSD of SOI with no EBW, and 100% EBW raised cosine interferer

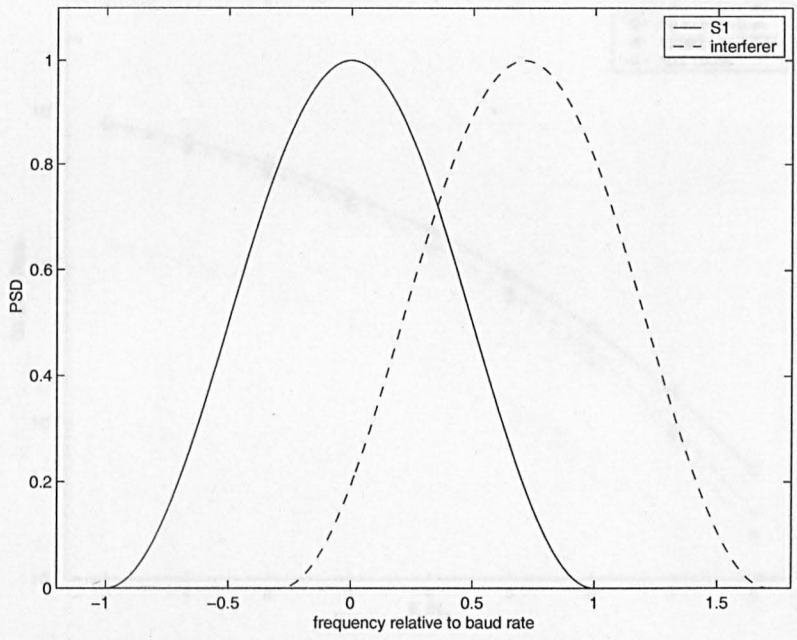

Figure 5.10: PSD of 100% EBW raised cosine SOI, and 100 % EBW raised cosine interferer

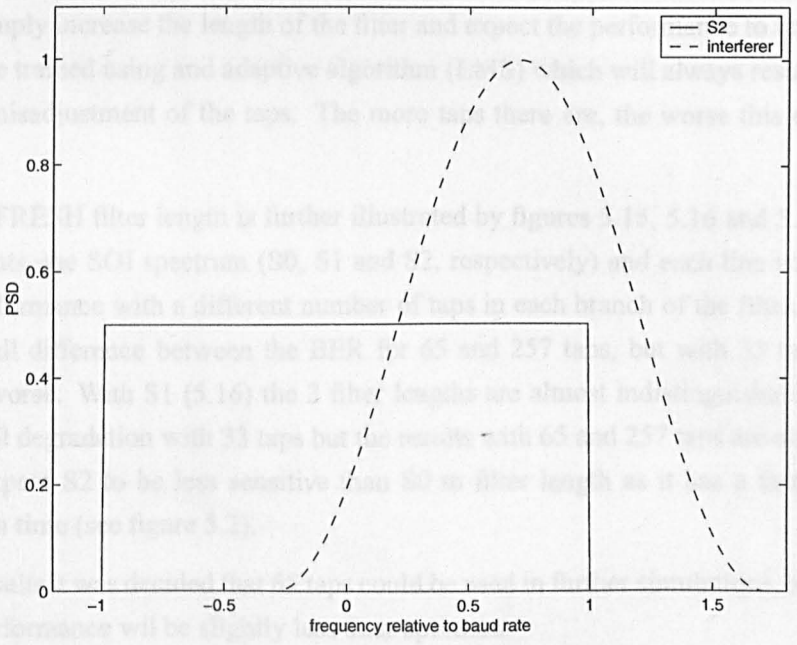

Figure 5.11: PSD of 100% EBW rectangular SOI, and 100% EBW raised cosine interferer

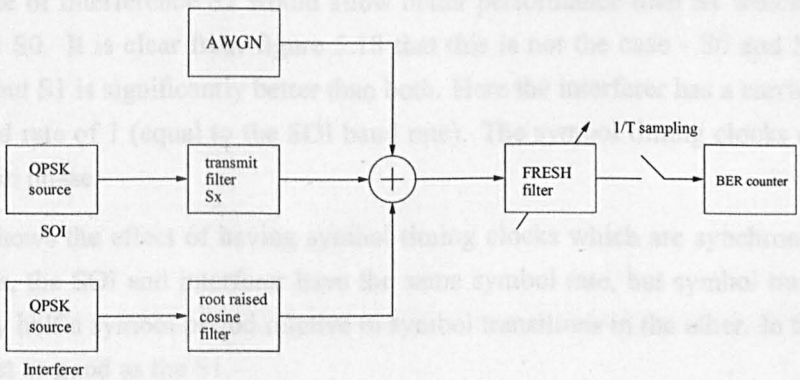

Figure 5.12: System used for simulations

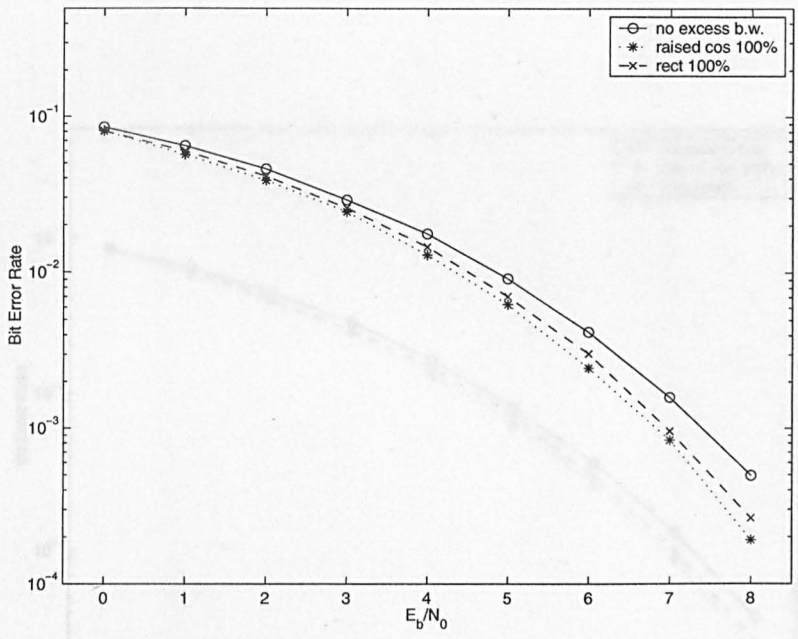

Figure 5.13: Effect of SOI spectrum, AWGN only, 65 taps

one cannot simply increase the length of the filter and expect the performance to improve because the filters were trained using and adaptive algorithm (LMS) which will always result in errors due to the slight misadjustment of the taps. The more taps there are, the worse this effect becomes [17].

The effect of FRESH filter length is further illustrated by figures 5.15, 5.16 and 5.17 where each graph represents one SOI spectrum (S0, S1 and S2, respectively) and each line within the graph shows the performance with a different number of taps in each branch of the filter. For S0 (5.15) there is a small difference between the BER for 65 and 257 taps, but with 33 taps the BER is significantly worse. With S1 (5.16) the 3 filter lengths are almost indistinguishable and with S2 there is a small degradation with 33 taps but the results with 65 and 257 taps are almost identical. One would expect S2 to be less sensitive than S0 to filter length as it has a faster decrease in amplitude with time (see figure 5.2).

From these results it was decided that 65 taps could be used in further simulations, bearing in mind that the S0 performance wil be slightly less than optimum.

The graphs above represent extracting the SOI from AWGN only. However it was expected that in the presence of interference S2 would allow better performance than S1 which in turn would be better than S0. It is clear from figure 5.18 that this is not the case - S0 and S2 give similar performance but S1 is significantly better than both. Here the interferer has a carrier frequency of 0.7 and a baud rate of 1 (equal to the SOI baud rate). The symbol timing clocks of the SOI and interferer are in phase.

Figure 5.19 shows the effect of having symbol timing clocks which are synchronised but not in phase. That is, the SOI and interferer have the same symbol rate, but symbol transitions in one are delayed by half a symbol period relative to symbol transitions in the other. In this case the S2 result is almost as good as the S1.

In figure 5.20 the symbol rate of the interferer is different from that of the SOI (16/17 instead of

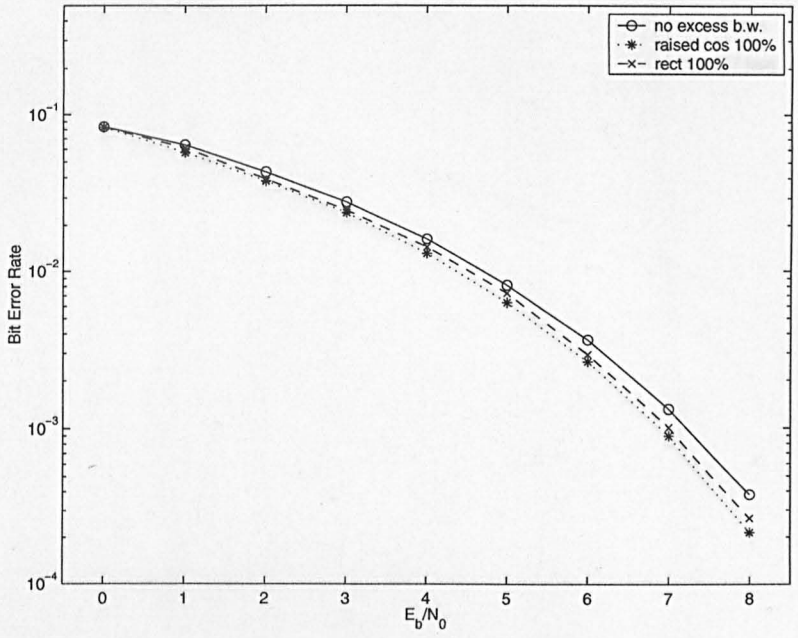

Figure 5.14: Effect of SOI spectrum, AWGN only, 257 taps

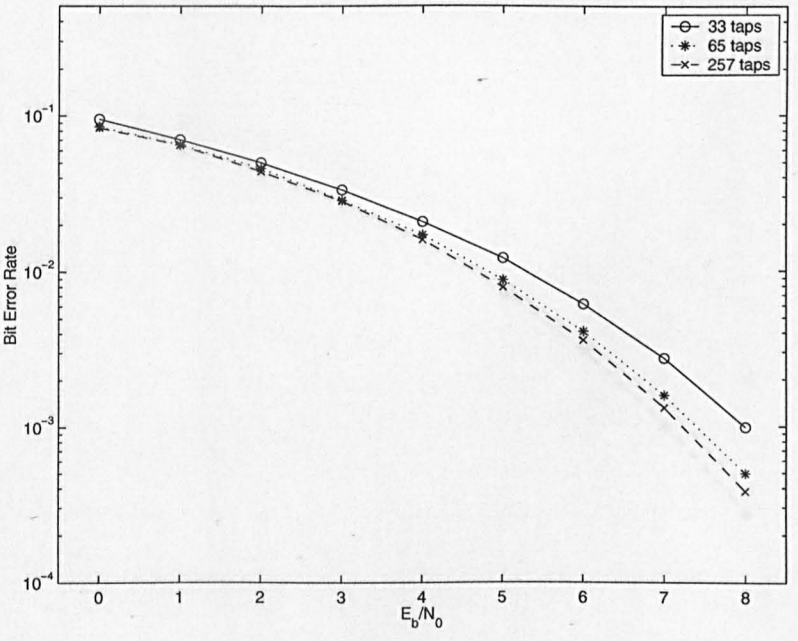

Figure 5.15: Effect of number of filter taps, 0% EBW, AWGN only

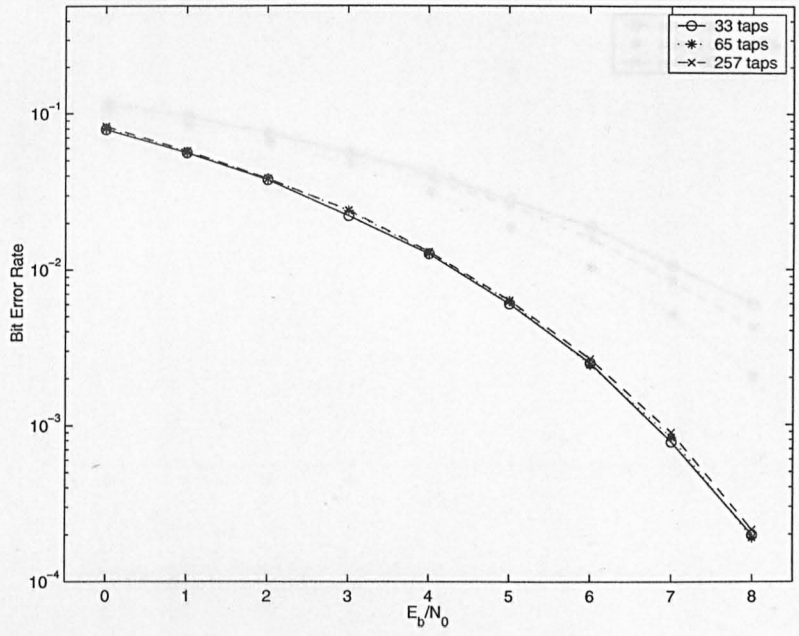

Figure 5.16: Effect of number of filter taps, 100% raised cos, AWGN only

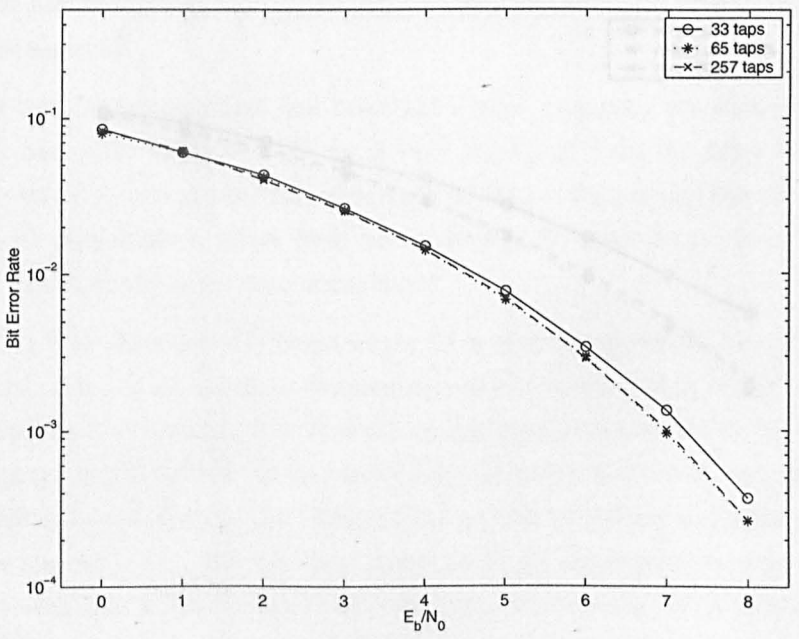

Figure 5.17: Effect of number of filter taps, 100% rectangular, AWGN only

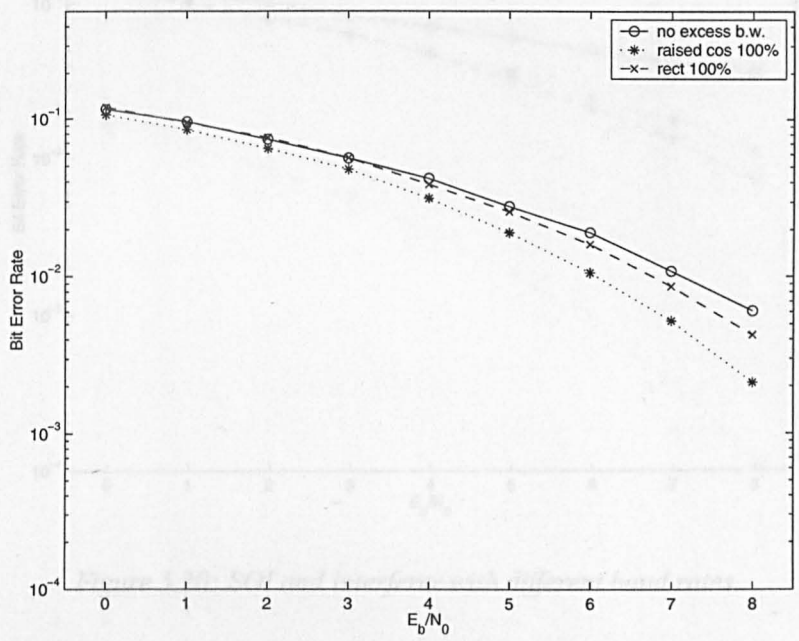

Figure 5.18: Effect of SOI spectrum with interference

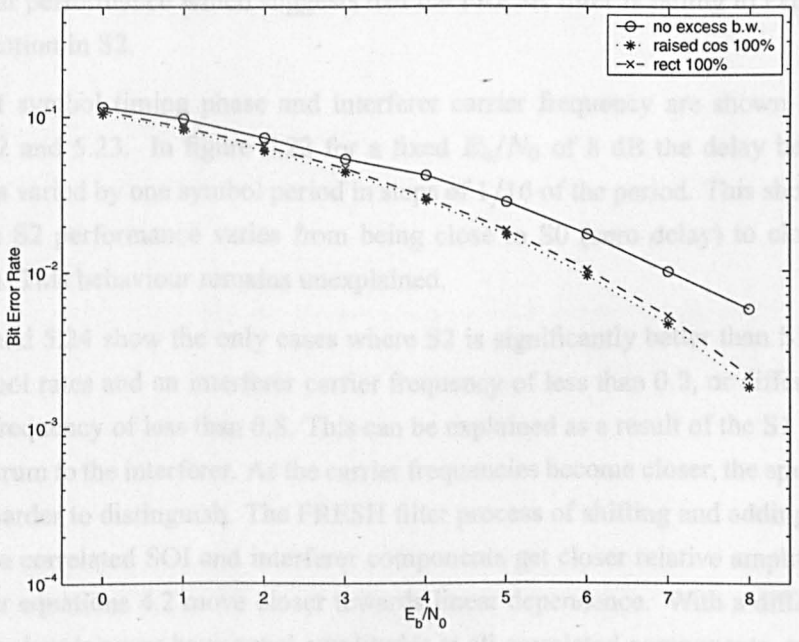

Figure 5.19: SOI and interferer symbol clocks out of phase

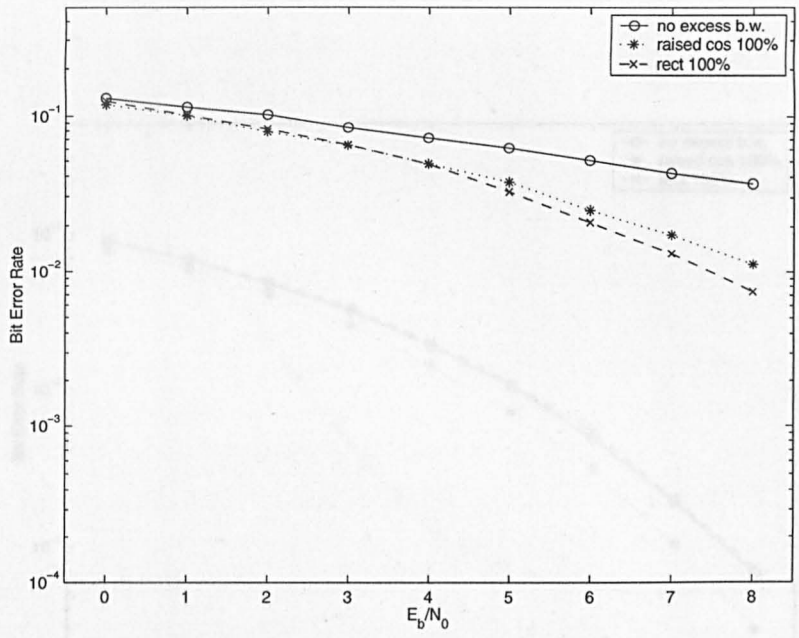

Figure 5.20: SOI and interferer with different baud rates

1). In this case only the spectral correlation of the SOI is used, as the filter has frequency shifts of ± 1. S0 is significantly worse here than in figures 5.18 and 5.19 because there is no spectral correlation in the S0 spectrum. S1 and S2 are also worse as the filter can only exploit the SOI correlation instead of both SOI and interferer in the earlier graphs.

For a scenario with identical symbol rates, with the clocks in phase (as in figure 5.18) but a different interferer carrier frequency of 1.3, the results are shown in 5.21. The interferer overlaps the SOI less, so all BERs are lower than in figure 5.18, but again S1 is better than S2. In fact S2 and S0 have similar performance which suggests that the FRESH filter is failing to exploit any of the spectral correlation in S2.

The effects of symbol timing phase and interferer carrier frequency are shown in more detail in figures 5.22 and 5.23. In figure 5.22 for a fixed E_b/N_0 of 8 dB the delay between symbol transitions was varied by one symbol period in steps of $1/16$ of the period. This shows the curious effect that the S2 performance varies from being close to S0 (zero delay) to close to S1 (half symbol delay). This behaviour remains unexplained.

Figures 5.23 and 5.24 show the only cases where S2 is significantly better than S1, which is for identical symbol rates and an interferer carrier frequency of less than 0.2, or different baud rates and a carrier frequency of less than 0.8. This can be explained as a result of the S1 SOI having an identical spectrum to the interferer. As the carrier frequencies become closer, the spectra of the signals become harder to distinguish. The FRESH filter process of shifting and adding becomes less effective as the correlated SOI and interferer components get closer relative amplitudes. Equivalently the filter equations 4.2 move closer towards linear dependence. With a different spectrum such as S2 the signals never have equal amplitudes at all correlated components, and the FRESH filter can weight each component in a way which maximises the SOI without also maximising the interferer.

But it is clear that using S2 does not, in general, give an improvement in performance over S1.

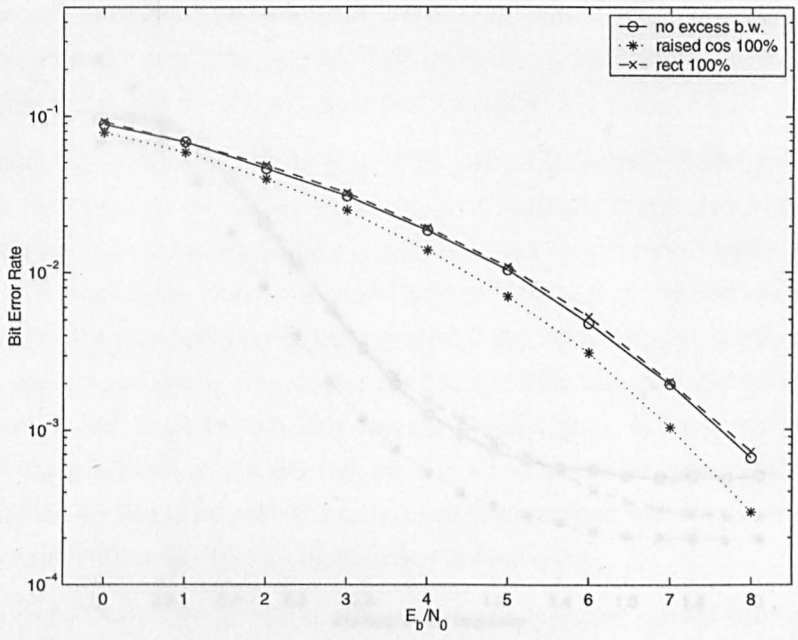

Figure 5.21: Interferer with carrier frequency 1.3

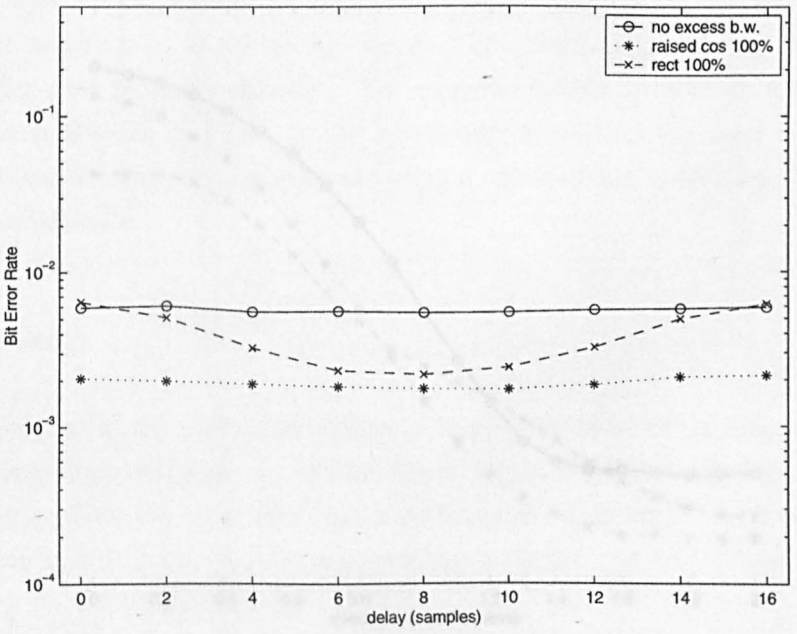

Figure 5.22: Effect of varying SOI and interferer symbol relative delay

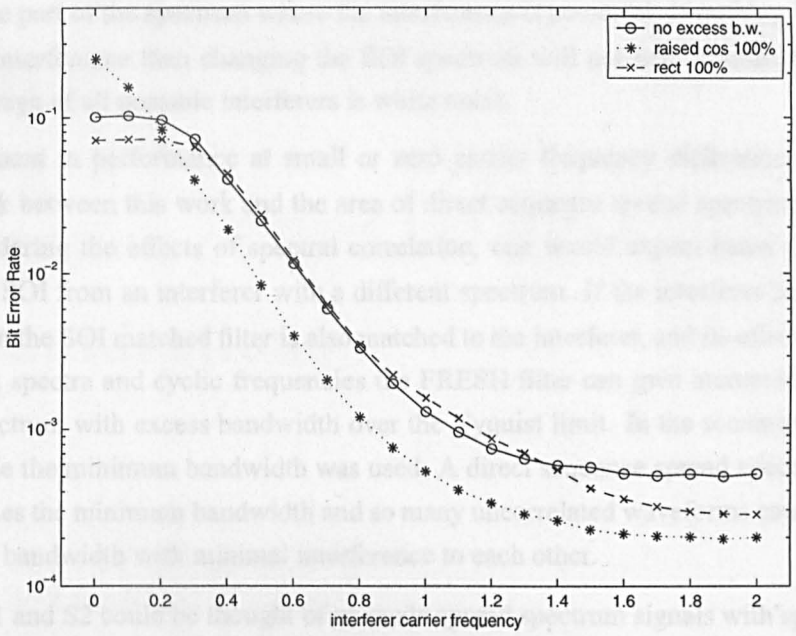

Figure 5.23: Effect of varying interferer carrier frequency, same baud rates

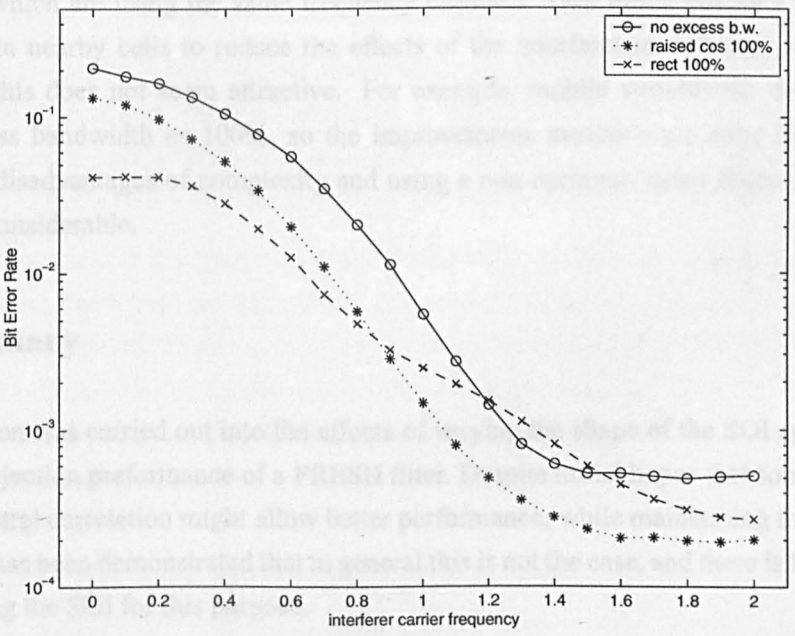

Figure 5.24: Effect of varying interferer carrier frequency, different baud rates

There are particular interference scenarios where each is preferable, but the improvements in BER are not dramatic. If a receiver is operating in a predictable interference scenario it may be worth considering designing the SOI spectrum accordingly. A simple example of this is, of course, to avoid using the part of the spectrum where the interference is powerful. If nothing is known about the potential interference then changing the SOI spectrum will not help immunity; after all, the ensemble average of all possible interferers is white noise.

The improvement in performance at small or zero carrier frequency differences introduces an interesting link between this work and the area of direct sequence spread spectrum signals. Even without considering the effects of spectral correlation, one would expect better performance in receiving and SOI from an interferer with a different spectrum. If the interferer has the spectrum of the SOI then the SOI matched filter is also matched to the interferer, and its effect is maximised. With different spectra and cyclic frequencies the FRESH filter can give increased benefit, but it requires a spectrum with excess bandwidth over the Nyquist limit. In the scenarios tested in this chapter, double the minimum bandwidth was used. A direct sequence spread spectrum waveform uses many times the minimum bandwidth and so many uncorrelated waveforms can be transmitted over the same bandwidth with minimal interference to each other.

The signals S1 and S2 could be thought of as crude spread spectrum signals with spreading factor of 2 (as they are double the Nyquist bandwidth). The use of FRESH filtering with spread spectrum signals is described in [30]. From this perspective, linear multi-user detection of direct sequence spread spectrum signals is a related process to interference rejection using spectral differences in the signals, although in direct sequence spread spectrum it is differences in the code domain which are used, instead of the frequency domain in frequency shift filtering.

One situation where interference is at the same carrier frequency and baud rate, and with the same spectrum as the SOI, is in cellular mobile communications. Such interference comes from the nearest cells which are using the same frequency channels. One might consider using different pulse shapes in nearby cells to reduce the effects of the interference, although there are many reasons why this does not seem attractive. For example, mobile systems do not normally use as much excess bandwidth as 100%, so the improvements available are even less than shown here, and the disadvantages of complexity and using a non-optimum pulse shape for a particular situation are considerable.

5.4 Summary

An investigation was carried out into the effects of varying the shape of the SOI spectrum on the interference rejection performance of a FRESH filter. Despite initial hopes that some spectra with enhanced spectral correlation might allow better performance, while maintaining the same overall bandwidth, it has been demonstrated that in general this is not the case, and there is little advantage in manipulating the SOI for this purpose.

Chapter 6

Application of FRESH filters to VLF Communications

Contents

It was shown in chapter 3 that FRESH filters are particularly effective in removing interference when they can exploit interferer cyclostationarity. The advantage over a FSE becomes more obvious as the signal to interference ratio decreases. A practical application to which FRESH filters may be suited is that of Very Low Frequency (VLF) radio communication.

VLF radio, which covers the frequency range 3 to 30 kHz, provides a useful means of communication with submarines because it has stable and predictable propagation and can penetrate some distance through seawater. However there are a number of factors which can seriously degrade its effectiveness, such as impulsive noise from lightning strikes, and interference from other users of the frequency band. VLF communication often uses Gaussian Minimum Shift Keying (GMSK) to modulate the carrier, as transmit antenna bandwidths are limited, and because it is a constant amplitude waveform.

As VLF is used for long range communications, land based transmitters may use very high powers. The adjacent channel interference (ACI) problem arises from the fact that reception may take place on a vessel which is operating far from the SOI transmitter, but close to an interfering transmitter.

Despite the use of GMSK modulation which limits the energy in the signal's side lobes, if another signal is on an adjacent channel and at a much higher power, then these side lobes can still present a major problem.

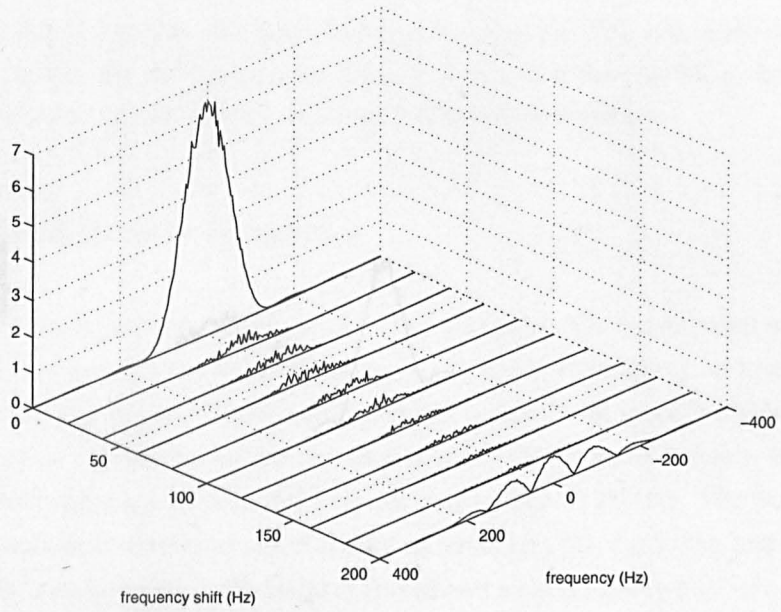

Figure 6.1: Spectral correlation density of GMSK

It is not expected that FRESH filtering can provide any protection from high powered impulsive noise, but we are interested in whether the presence of impulses, or the presence of the effects of imperfectly removed impulses causes significant degradation to the FRESH filter's interference rejection function.

The work presented in this chapter has been published in three conference papers [2, 3, 4].

The spectral correlation of Minimum Shift Keying (MSK) and GMSK is first described for the purpose of choosing the appropriate FRESH filter structure for this application. The number of filter branches required is tested and it is found that for high powered interferers one is sufficient with a frequency shift related to the interferer carrier and baud rate. The BER performance is first tested in AWGN, and then in impulsive noise, with and without the use of an impulse rejection technique.

6.1 Spectral Correlation of GMSK

The spectral correlation exhibited by MSK is of a more complicated form than that of QPSK or BPSK. GMSK signals are generated from Gaussian pre-filtering of the data in a MSK modulator, so one cannot assume that the structure of the SCD is identical for the two signals in general. However, in this case examined here the bandwidth of the Gaussian filter is relatively wide, resulting in a signal which differs little from MSK. It is shown in [54] that for MSK conjugate correlation exists under frequency shifts of $\alpha = \pm 2f_c \pm \frac{k}{T_0}$ for odd values of k and non-conjugate correlation for $\alpha = \pm\frac{k}{T_0}$ for even k, where k are integers, f_c is the carrier frequency, and T_0 is the baud period. When the signal is modelled at complex baseband, the useful correlation occurs at $\alpha = \frac{k}{T_0}$ with k odd for conjugate correlation and k even for non-conjugate correlation. In figures 6.1 and 6.2 the spectral correlation of GMSK is shown - the corresponding graph for MSK is similar but with slightly different shapes of the peaks.

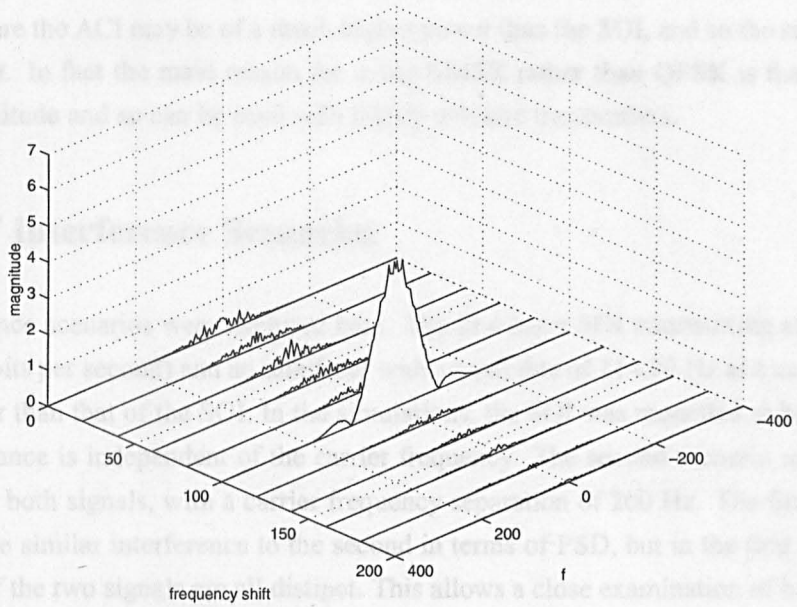

Figure 6.2: Conjugate spectral correlation density of GMSK

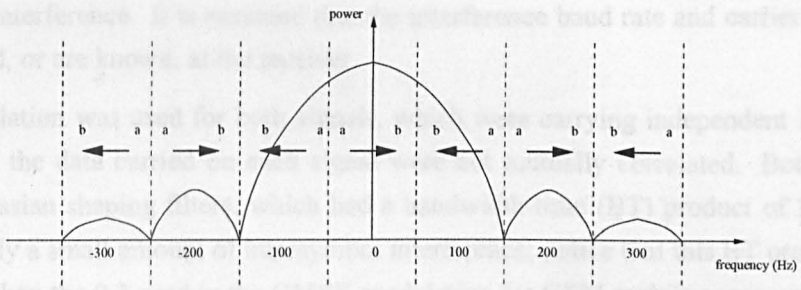

Figure 6.3: Spectral correlation (conjugate and non-conjugate) of GMSK

Figure 6.3 illustrates the structure of the spectral correlation in a different way. Here the spectrum of GMSK is shown from -350 Hz to 350 Hz and the baud rate is 100 Hz. the power spectral density of GMSK is shown with the correlated frequency bands indicated by the vertical dashed lines. Conjugate correlation is indicated by arrows in opposite directions; non-conjugate correlation exists in the bands with arrows in the same direction. For example the central region (-50 to 50 Hz) is correlated with the regions at frequency shifts of -200 Hz and at 200 Hz (that is, frequency shifts of -2 and 2 times the baud rate). However correlation of the centre with the region from 50 to 150 Hz occurs only after reflection about the y-axis and then shifting by 100 Hz, in other words, this is conjugate correlation under a baud rate frequency shift.

This can be clarified with the example that the signal component at 49 Hz exhibits correlation with the components at -151 and 249 Hz, and conjugate correlation with the components at -349, -149, 51 and 251 Hz (ignoring larger shifts for simplicity). If the carrier frequency is not zero then an extra $2f_c$ must be added to the shifts exploiting conjugate correlation.

Practically, this means that the most useful frequency shifts appear to be at $\pm 2f_c \pm \frac{1}{T_0}$ and $\pm 2/T_0$. At larger shifts than this, the energy in the relevant side lobes of the signal is very low. As GMSK is often used in situations where spectral efficiency is important it may seem strange that there is significant redundancy. But in fact the width of the main lobe of GMSK is greater than that of

100% EBW QPSK. Although the proportion in the side bands is very low, we are considering here scenarios where the ACI may be of a much higher power than the SOI, and so the side lobe powers are significant. In fact the main reason for using GMSK rather than QPSK is that GMSK has a constant amplitude and so can be used with highly efficient transmitters.

6.2 VLF Interference Scenarios

Two interference scenarios were explored here. The first has a SOI transmitting at a baud rate of 100 Hz (200 bits per second) and an interferer with a baud rate of 114.29 Hz at a carrier frequency 210 Hz higher than that of the SOI. In the simulations, the SOI was modelled at baseband, as the filter performance is independent of the carrier frequency. The second scenario uses a baud rate of 100 Hz for both signals, with a carrier frequency separation of 200 Hz. The first scenario was chosen to have similar interference to the second in terms of PSD, but in the first case the cyclic frequencies of the two signals are all distinct. This allows a close examination of how the FRESH filter exploits frequency shifts related to the properties of the SOI and interferer. In the second case, which is more likely to occur in practice, some cyclic frequencies are common to both the SOI and the interference. It is assumed that the interference baud rate and carrier frequency can be determined, or are known, at the receiver.

GMSK modulation was used for both signals, which were carrying independent identically distributed data; the data carried on each signal were not mutually correlated. Both signals used identical Gaussian shaping filters, which had a bandwidth-time (BT) product of 2.5. This filter introduces only a small amount of intersymbol interference; notice that this BT product is significantly larger than the 0.3 used in the GMSK modulation for GSM mobile communications [101]. Normally an equaliser would be used to remove this, however the ISI introduced by the Gaussian filter is insignificant compared to the ACI so no equalisation was performed here. In fact the spectrum of GMSK with such a large Gaussian filter bandwidth is close to that of MSK.

Interferer powers from 0 to 40 dB relative to the SOI were used. The spectra of the signals with the interferer at 30 dB are shown in figure 6.4. The general structure of the simulation is shown in figure 6.5. Note that here, unlike in the QPSK systems presented earlier, the FRESH filter does not incorporate the matched filter; the signal used to train the FRESH filter was the same as the transmitted SOI, and the FRESH filter was followed by a matched filter.

Two different white noise scenarios were tested. First the simple AWGN case (section 6.3) and then an impulsive noise model was used, to give a more realistic representation of atmospheric noise in the VLF band (section 6.4).

The filters were implemented as trained adaptive filters. The LMS algorithm was used to adapt the filters to a steady state then adaptation was stopped while bit errors were counted. The SOI power was 1 and the LMS feedback coefficient, μ, was 0.001. It was found that lower values of μ than this made practically no difference to the final BERs calculated.

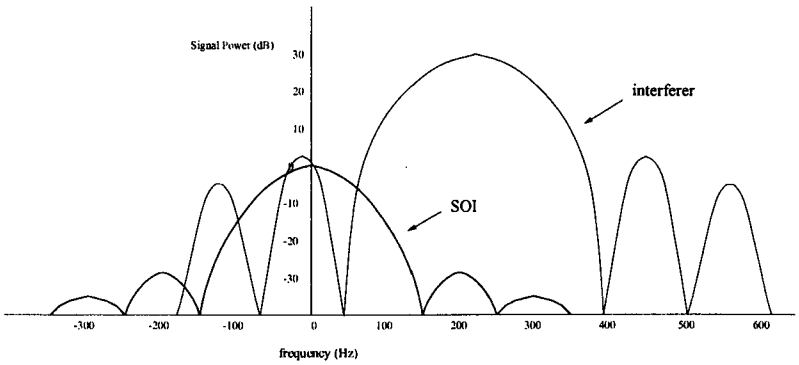

Figure 6.4: Spectrum of interference scenario (example)

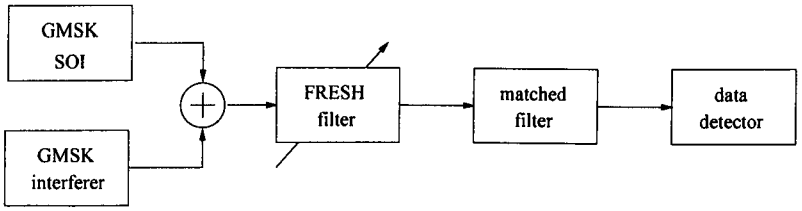

Figure 6.5: Structure of simulation

Related to which signal	Frequency shift (Hz)	Relation to signal properties
Interferer	305.71	$2f_i - b_i$
	-228.57	$-2b_i$
	534.28	$2f_i + b_i$
	228.57	$2b_i$
	77.14	$2f_i - 3b_i$
SOI	200.00	$2b_s$
	-200.00	$-2b_s$
	100.00	b_s
	-100.00	$-b_s$
	-300.00	$-3b_s$

Table 6.1: Frequency shifts tested with different baud rates (b_i=baud rate of interferer, b_s=baud rate of SOI, f_i=interferer carrier)

Related to which signal	Frequency shift (Hz)	Relation to signal properties
Interferer	300	$2f_i - b$
	-200	$-2b$
	500	$2f_i + b$
	200	$2b$
	100	$2f_i - 3b$
SOI	200	$2b$
	-200	$-2b$
	100	b
	-100	$-b$
	-300	$-3b$

Table 6.2: Frequency shifts tested with equal baud rates (b=baud rate of interferer and SOI, f_i=interferer carrier)

6.3 Interferer and Gaussian Noise

6.3.1 Selection of frequency shifts

First scenario - different baud rates

An important practical decision to be made in implementing a FRESH filter is the number of branches and which frequency shifts to use. The relative effectiveness of particular frequency shifts depends on many factors, for example, the cyclic frequencies of the SOI and interferer, the relative powers of the signals and the separation of the SOI and interferer carrier frequencies. The best choice of filter design is therefore highly scenario dependent.

For the chosen scenario detailed in section 6.2 a study was carried out on which frequency shift is the most effective in a single shift FRESH filter, and of what benefit there is in having a greater number of frequency shifts.

Table 6.1 shows the frequency shifts used in the first scenario (SOI and interferer with different baud rates), and how they are derived from the signals' properties. These 10 shifts (5 SOI related and 5 interferer related) were chosen as the most likely to be effective by inspection of figure 6.4. The equivalent frequency shifts for the second scenario (SOI and interferer with the same baud rate, and channel spacing equal to twice the baud rate) are shown in table 6.2.

Simulations of the VLF system were designed to identify the most effective number and value of frequency shifts. Initially, a system using a single shift was constructed, and the BER was measured for different values of that single shift. This was repeated for interferer powers from 0 dB to 40 dB relative to the SOI. Figure 6.6 shows the ranking of the best three frequency shifts for each interferer power; the highest column indicates the most effective shift in reducing BER. This graph shows that with higher powered interferers more of the best frequency shifts are related to the interferer properties. At lower interferer powers the best frequency shifts are SOI related.

Figure 6.7 shows the actual improved error rates that result from the use of four of these frequency shifts. The best two SOI related and interferer related shifts were used, and the resulting BERs

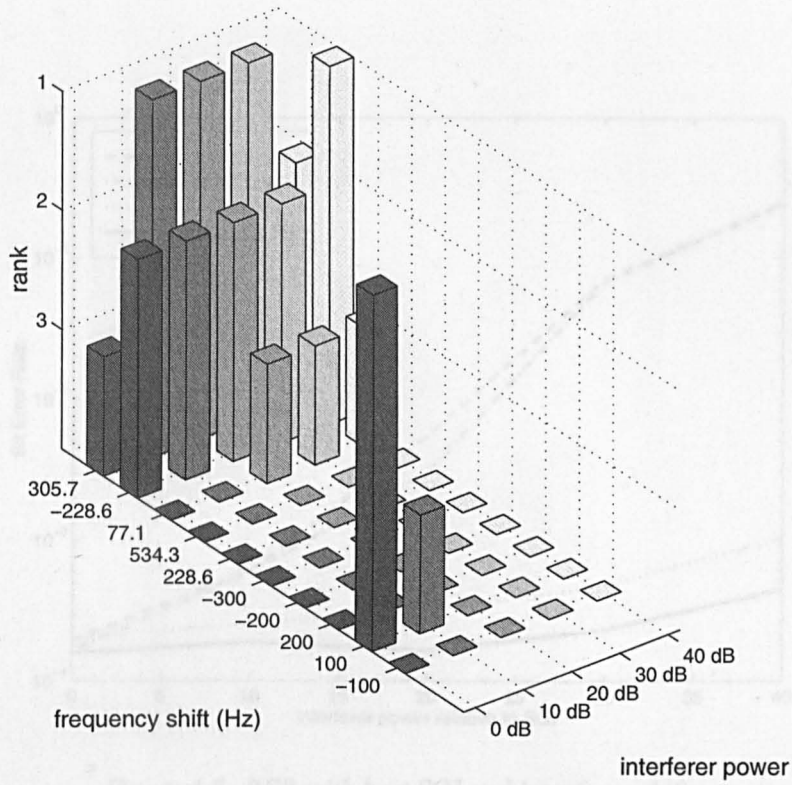

Figure 6.6: Ranking of frequency shifts for effect on BER (different baud rates)

plotted against white noise power for each interferer power. The interferer related shifts are more useful for all interferer powers greater than that of the SOI. Of course the actual error rate is higher when the interferer power is higher, but the benefit in using a FRESH filter also increases, as is shown below.

The previous two graphs were based on a FRESH filter with two branches and one frequency shift. The effect of combining several different frequency shifts in a multi-branch filter were also measured. Figure 6.8 shows the performance of filters with between zero and four frequency shifts. Zero shifts is equivalent to a fractionally spaced equaliser; four shifts require a structure with 5 branches and subfilters. The frequency shifts were added in order of their effectiveness in a single shift filter, that is 305.7, -228.6, 77.1, 534.3 Hz. It is clear from figure 6.8 that there is almost no advantage in using more than one frequency shift.

This shows that only a moderate increase in complexity is required: a filter with one frequency shift has two branches, and two subfilters, so the filtering computational load is double that of the corresponding FSE.

Second scenario - equal baud rates

The comparison of the effectiveness of individual frequency shifts was repeated for the second interference scenario where the SOI and interference have the same baud rate and have carrier frequencies separated by twice the baud rate. In this case the cyclic frequencies are as shown in table 6.2.

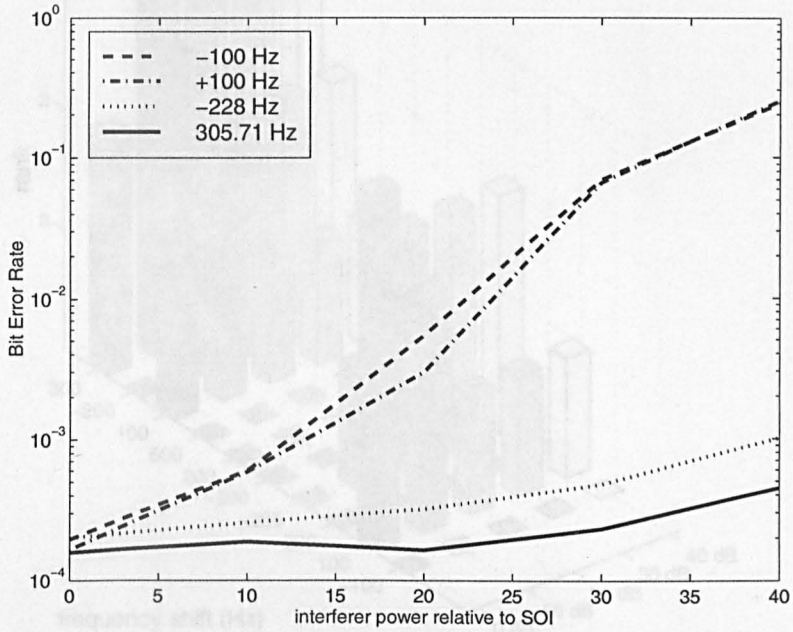

Figure 6.7: BER with best SOI and interferer shifts

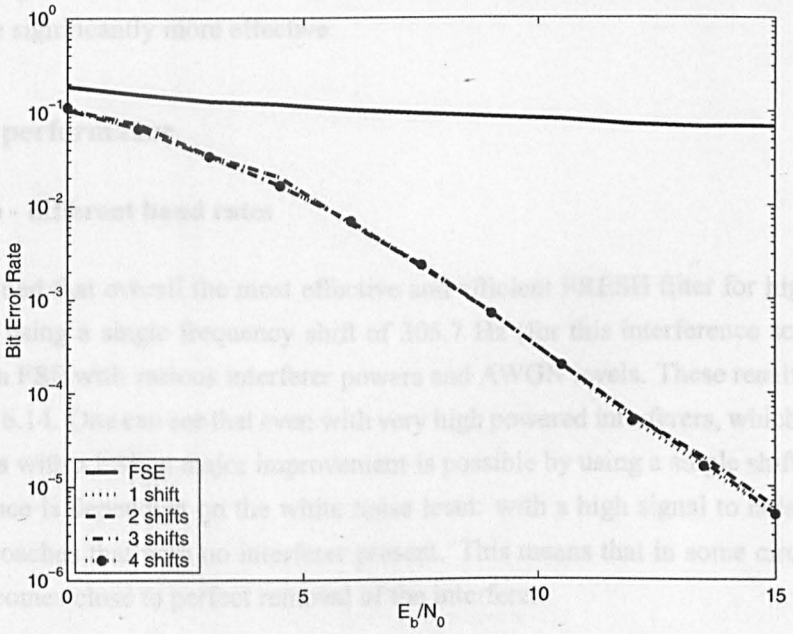

Figure 6.8: BER with different numbers of shifts in filter

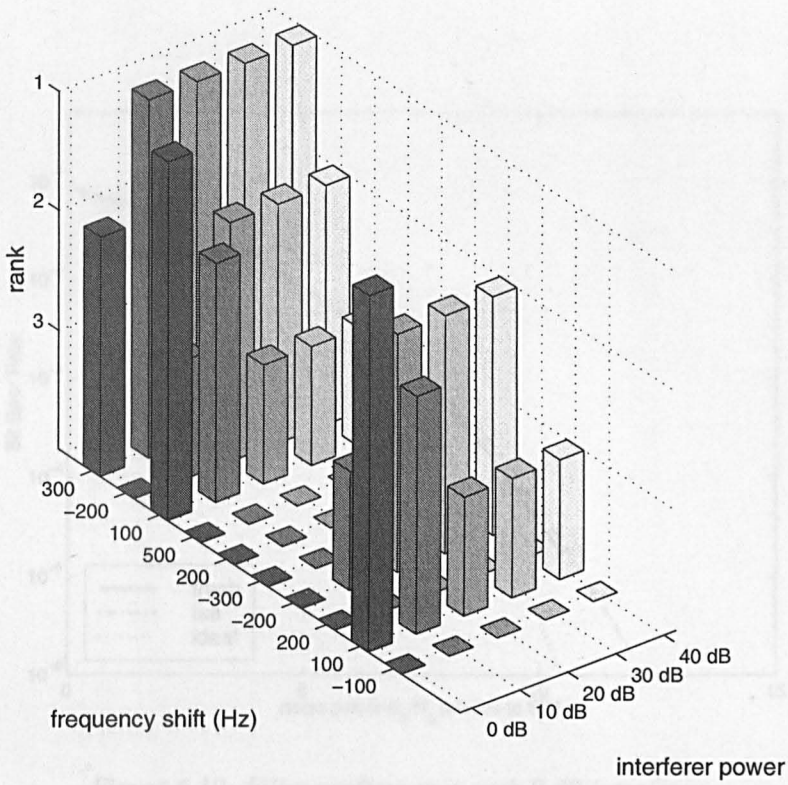

Figure 6.9: Ranking of frequency shifts for effect on BER (same baud rates)

Clearly, some of these cyclic frequencies can be interpreted as either SOI or interferer related. The ranking of the effect of these frequency shifts is shown in figure 6.9. For comparison with figure 6.6, the frequency shifts with two interpretations are plotted twice. We see that there is no significant change in the rankings from figure 6.6: shifts related to both the SOI and the interferer do not become significantly more effective.

6.3.2 BER performance

First scenario - different baud rates

Having confirmed that overall the most effective and efficient FRESH filter for high interference powers is one using a single frequency shift of 305.7 Hz (for this interference scenario), it was tested against a FSE with various interferer powers and AWGN levels. These results are shown in figures 6.10 to 6.14. One can see that even with very high powered interferers, which give unusably high error rates with a FSE, a major improvement is possible by using a single shift FRESH filter. The performance is dependent on the white noise level: with a high signal to noise level, the bit error rate approaches that with no interferer present. This means that in some circumstances the FRESH filter comes close to perfect removal of the interferer.

Second scenario - equal baud rates

The actual BERs achieved with this scenario, using a single shift of 300 Hz and an interference power of 30 dB, are shown in figure 6.15 along with the data from figure 6.13. This shows the

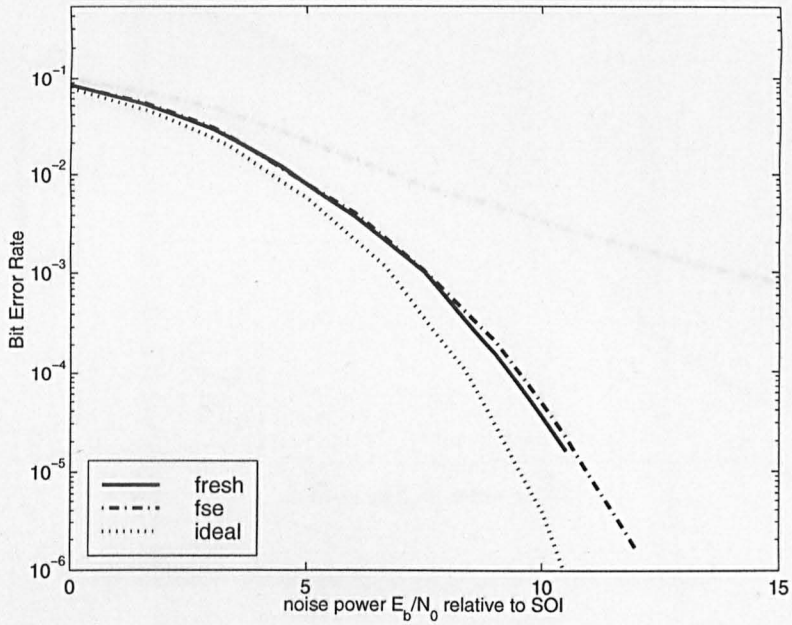

Figure 6.10: Filter performance with 0 dB interferer

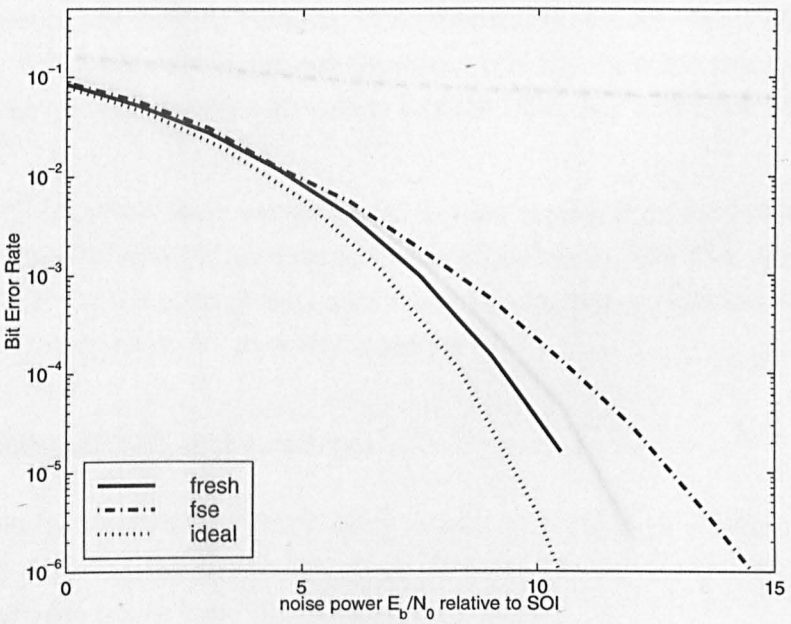

Figure 6.11: Filter performance with 10 dB interferer

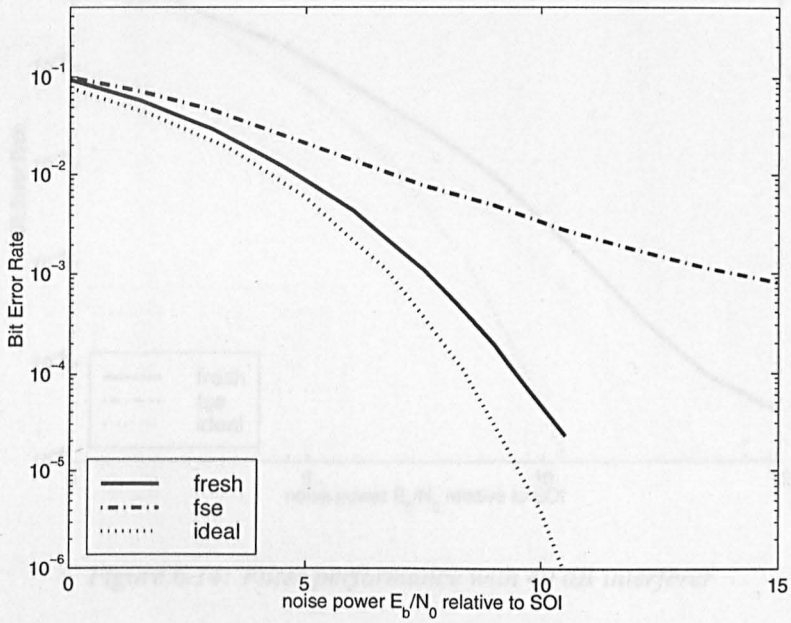

Figure 6.12: Filter performance with 20 dB interferer

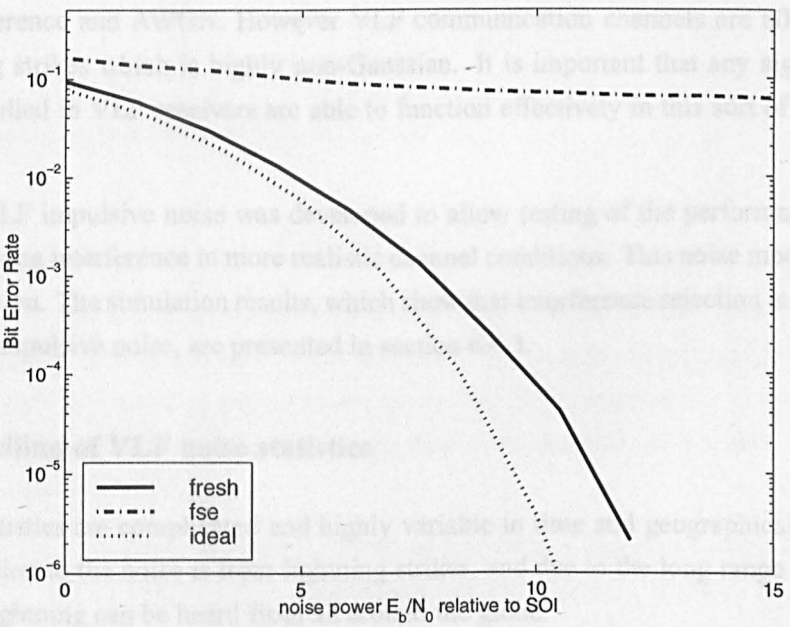

Figure 6.13: Filter performance with 30 dB interferer

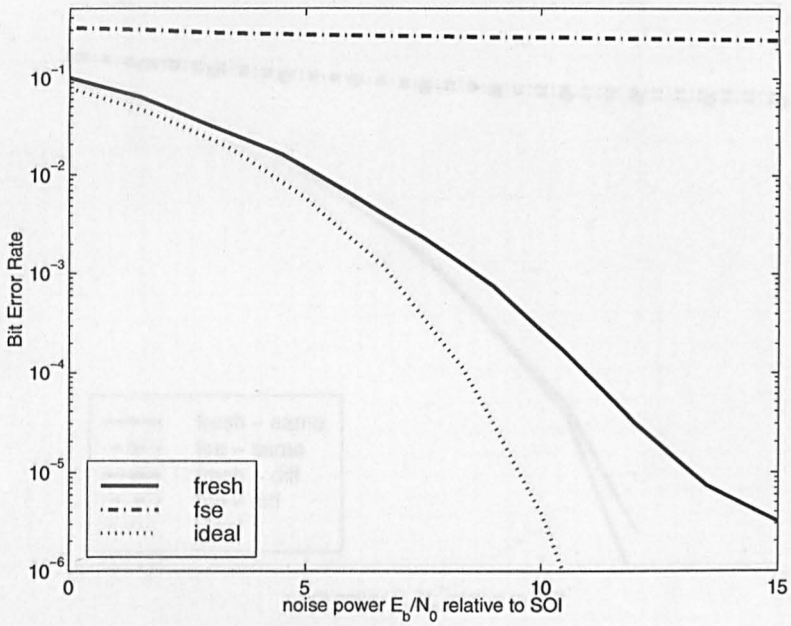

Figure 6.14: Filter performance with 40 dB interferer

change in interference scenario has little effect on the BER. This was also the case for other interferer powers, so figures 6.10 to 6.14 also illustrate performance under the second interference scenario.

6.4 Interferer and Impulsive noise

The results presented previously in this chapter were based on the reception of GMSK in the presence of interference and AWGN. However VLF communication channels are affected by noise from lightning strikes which is highly non-Gaussian. It is important that any signal processing techniques applied in VLF receivers are able to function effectively in this sort of noise environment.

A model of VLF impulsive noise was developed to allow testing of the performance of FRESH filters in removing interference in more realistic channel conditions. This noise model is described in the next section. The simulation results, which show that interference rejection is still performed effectively in impulsive noise, are presented in section 6.4.3.

6.4.1 Modelling of VLF noise statistics

VLF noise statistics are complicated and highly variable in time and geographical location. The main contribution to the noise is from lightning strikes, and due to the long range propagation of VLF signals lightning can be heard from all around the globe.

The purpose of this study was to investigate general the effects of impulsive noise, and its rejection, on FRESH filtering, so a detailed time-varying noise model is not required. A model was developed which was based on a graph supplied by DERA Portsdown (figure 6.16). This graph displays the cumulative distribution function (CDF) of the amplitude of VLF signals received un-

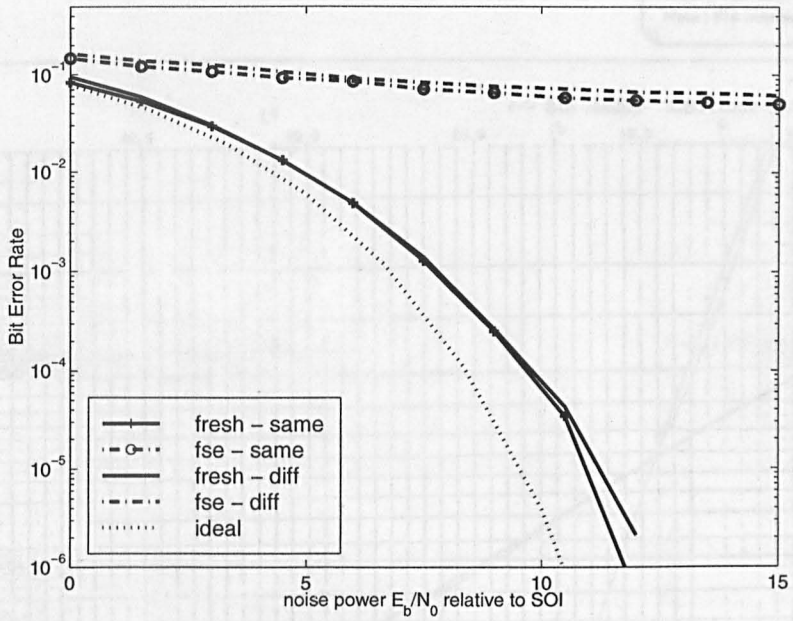

Figure 6.15: Comparison of filter performance with interferer at same and different baud rates

der some "typical" conditions. In fact the graph represents the output of a VLF atmospheric noise simulator, but the mathematical model used by this simulator was not available.

The points on the graph are reasonably described by two straight lines. The steeper of these two represents the Gaussian noise floor; the flatter line is the impulsive, or "super-Gaussian" component. Complex Gaussian noise (that is, noise which has two independent real and imaginary Gaussian components) has an amplitude distribution which is described by the Rayleigh distribution function [76] as shown:

$$f(x) = \frac{x}{\sigma^2} e^{-x^2/2\sigma^2} \tag{6.1}$$

The corresponding cumulative distribution function is:

$$F(x) = 1 - e^{-x^2/2\sigma^2} \tag{6.2}$$

We wish to relate this to the Gaussian line in figure 6.16, using the left hand y-axis of figure 6.16, "Percentage of time amplitudes exceed abscissa". The x-axis of figure 6.16 is signal level in terms of "bin number" which is a logarithmic scale. Bin 32 corresponds to a signal level of 2 V peak-to-peak, with each bin corresponding to an amplitude threshold of 1.5 dBV. Bin 0 therefore represents an amplitude of $30.5\mu V$.

The standard CDF is defined in terms of time, or proportion, below a certain value, so figure 6.16 is in fact a plot of 100(1-CDF). This means that the Gaussian line can be represented by $100e^{-x^2/2\sigma^2}$. It was found that a good agreement was achieved by setting σ to be 5.81×10^{-4}, which is equivalent to bin 8.5 on the x-axis of figure 6.16.

The equations of the two straight lines were calculated, and using equation 6.2, it was found that the super-Gaussian component could be represented by the CDF shown in equation 6.3:

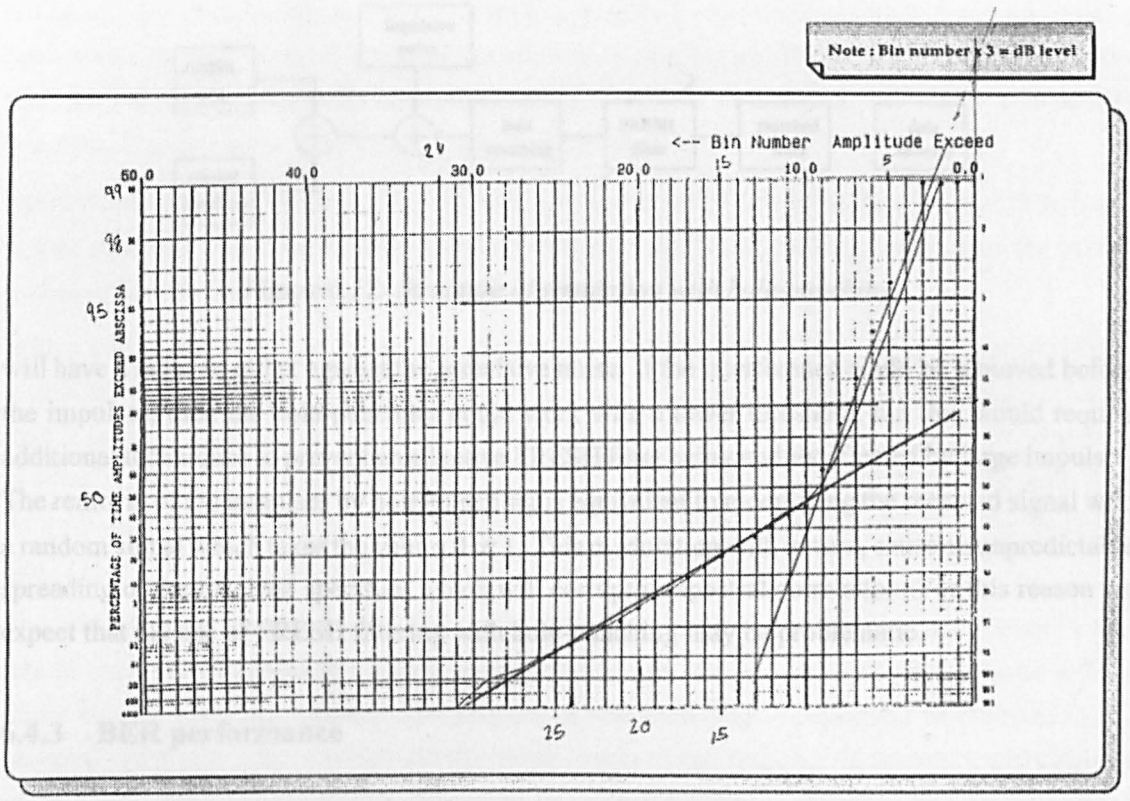

Figure 6.16: VLF atmospheric noise amplitude cumulative distribution

$$F(x) = 100e^{-19.64x^{0.428}} \tag{6.3}$$

This gives reasonable agreement up the 99% level. A block was written for SPW which produces noise with a distribution which is a combination of the Gaussian and super-Gaussian components as required, and this was used in the simulations described below. It was found that the total noise power was about 20 dB higher than the Gaussian component only. The noise signal given by this model is spectrally white, and so its power can still be described in terms of noise power spectral density, which is labelled N_i to distinguish it from the AWGN noise level N_0 used previously.

6.4.2 Impulse removal by hole punching

It was not expected that FRESH filters would be particularly effective in removing the impulses in the simulated VLF noise, but it is important that they are able to perform their interference removal function alongside other techniques which would be applied to moderate the effects of the impulsive noise.

A useful and simple technique for impulse removal is "hole-punching" where it is assumed that whenever the signal magnitude goes over a certain threshold, it is as a result of a noise impulse, and this impulse is removed by setting the signal magnitude to zero. A block was written for SPW which performs this function. It was incorporated into the simulation as shown in figure 6.17.

The use of hole-punching before FRESH filtering raises some interesting questions. With a high powered interferer, the hole-punching threshold must be set high, so that the normal level of the interferer does not cause a complete loss of the received signal. This means that the hole-punching

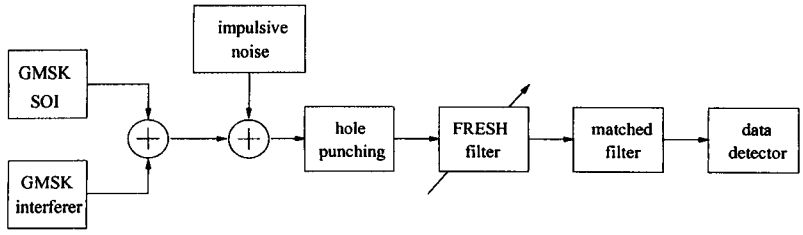

Figure 6.17: Structure of simulation with hole-punching

will have a reduced effect against the impulsive noise. If the interference could be removed before the impulses, then the hole-punching could work with a lower threshold, but this would require additional techniques to prevent an adaptive FRESH filter being unduly affected by large impulses. The removal of the impulses by hole-punching is equivalent to modulating the received signal with a random signal which takes the values 0 or 1. This modulation will, in turn, cause an unpredictable spreading of the received spectrum which will corrupt its spectral correlation. For this reason we expect that the use of FRESH filtering with hole-punching may be problematic.

6.4.3 BER performance

There is no reason to expect the impulsive noise statistics to affect the spectral correlation of the signal, so the same FRESH filters were used here as with AWGN, that is, a single shift (two branch) FRESH filter with a frequency shift of 305.71 Hz for the first interference scenario and 300 Hz for the second.

As in section 6.3 it was found that there was very little difference in the results obtained from the two interference scenarios. For the sake of brevity, only the second, equal baud rate, results are included here. The corresponding graphs for the first scenario are in [3].

In the simulations the interferer power was varied between 10 dB and 30 dB above the SOI level. This causes the first side lobe of the interferer to vary from -10 dB to +10 dB relative to the SOI main lobe, which it overlays. The impulsive noise level was varied to give E_b/N_i values for the SOI from 0 to 20 dB. The hole-punching threshold used was dependent on the interferer power used.

In one case (figure 6.21) the FRESH filter performance was also compared to that of a FSE. When the FSE was used it replaced the FRESH filter in the structure shown in figure 6.5.

Before examining interference rejection performance, we first show how the choice of hole-punching threshold affects the BER in impulsive noise only with no interference present. Figure 6.18 shows the BER for thresholds of 5, 10 and 30 and for no hole-punching. These thresholds were used with the interference, and are chosen by inspection to be greater than, but close to the magnitudes of the SOI (power=1), interferer and non-impulsive noise component. The threshold of 5 was used with the 10 dB interferer, 10 with the 20 dB interferer and 30 with the 30 dB interferer. It is clear that the lower hole-punching thresholds cause a major reduction in BER, with no interference present. A threshold of 30, however, which is necessary for an interference of 30 dB, has almost no effect on the impulsive noise.

Figure 6.19 shows the performance of the FRESH filter in removing interference, with and without

hole-punching being performed too. The BER is roughly 2 orders of magnitude lower when using hole-punching than when all impulses are allowed to pass to the FRESH filter. The performance with no interference, but with hole-punching, is plotted as an indicator of "ideal" interference rejection. The results do not extend to BERs of less than 10^{-6}.

However, figure 6.20, which is for an interference power of 20 dB above the SOI, shows that with higher powered interference, hole-punching increases the BER. That is, it degrades the performance of the FRESH filter in removing the interference.

With a still higher interferer power (30 dB) and a hole-punching threshold of 30 (figure 6.21) the results are more extreme - here the hole-punching can degrade the FRESH performance by more than 2 orders of magnitude (in terms of BER). Notice that (from figure 6.18 this hole-punching is actually having very little beneficial effect on the impulsive noise. This graph also shows the performance of a FSE, which is very poor, as it can do little to mitigate the effects of the interferer. See [2] for a more detailed comparison of FSE and FRESH performance.

The FRESH filter tested here is using the spectral correlation of the interference to remove that interference from the received signal. This explains why the hole-punching has a more serious effect at high interferer powers. The process of hole-punching is equivalent to modulating the received signal with signal which has the value 1 most of the time, but at random times switches to zero for a short interval. The correlation of the GMSK signal shown in figures 6.2 and 6.3 is seriously affected by this non-linear process. With higher interference to SOI ratios, to maintain accurate reception of the SOI, the filter requires a more accurate "match" of the correlated interferer components. The random modulation corrupts this correlation.

In general we can say that at higher interferer powers the FRESH filter is seriously degraded by the hole-punching process, but at lower powers both can be used together successfully.

The hole-punching was placed before the FRESH filter, as this would allow easier integration with existing technology. If the order were reversed, the degradation of the FRESH performance would not occur, but it would be necessary to have a receiver and FRESH filter with a very high dynamic range to allow the accurate processing of the impulsive signal.

There are many other approaches to impulsive noise rejection which may have a less serious effect on the interference spectral correlation. For example instead of setting the signal to zero when it exceeds a threshold, it could be replaced with interpolated values. This is of course a more complicated process, but feasible, especially if the interpolation is mainly to replace the interferer and the interference to noise ratio is high. A more sophisticated order statistics approach might also improve matters [102, 103] as could a wavelet based impulse rejection technique [104]. The main disadvantage of such techniques is that they would require a major redesign of existing VLF receivers. Currently hole-punching can be done on the analogue signal, before sampling.

6.5 Conclusions and Further Work

In section 6.3 it has been shown that in the high interference scenarios tested, FRESH filtering is highly effective in removing adjacent channel interference. This could be done effectively with only one frequency shift in the filter, this shift being matched to the cyclic frequencies of the

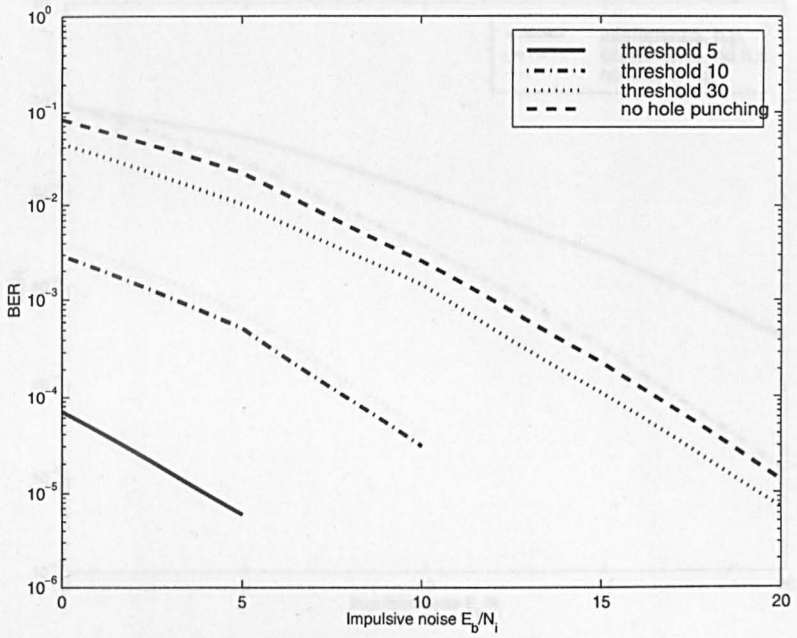

Figure 6.18: Hole-punching performance with no interference

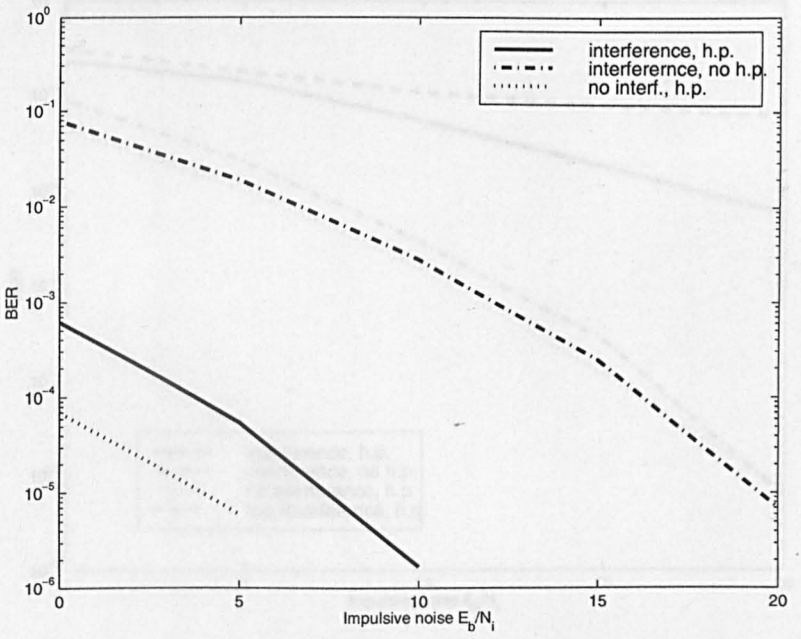

Figure 6.19: FRESH interference removal performance with interference at 10 dB above SOI

interferer.

FRESH filters cannot mitigate the effects of impulsive noise, but they are highly effective at interference removal. We have shown that impulsive noise in itself does not prevent the FRESH filters from removing interference, but the simple hole-punching technique does not work well in the presence of interference, as a threshold must be set above the interference magnitude, which can allow the majority of impulsive noise to pass. Also, the use of hole-punching degrades the FRESH filter performance by corrupting the spectrum of the received signal. It would be interesting and useful to investigate if other impulse removal techniques would allow the FRESH filter to operate more effectively. A technique which identifies an impulse in a more sophisticated way than

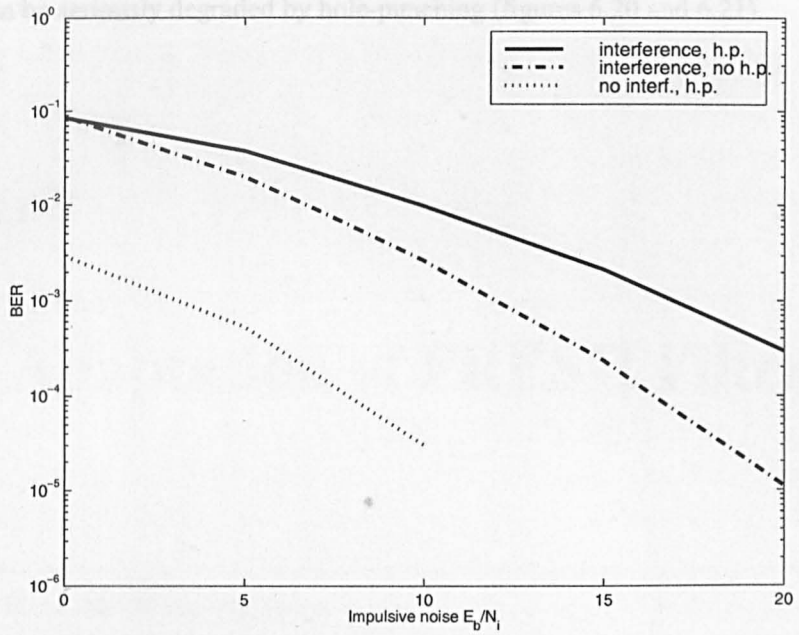

Figure 6.20: FRESH interference removal performance with interference at 20 dB above SOI

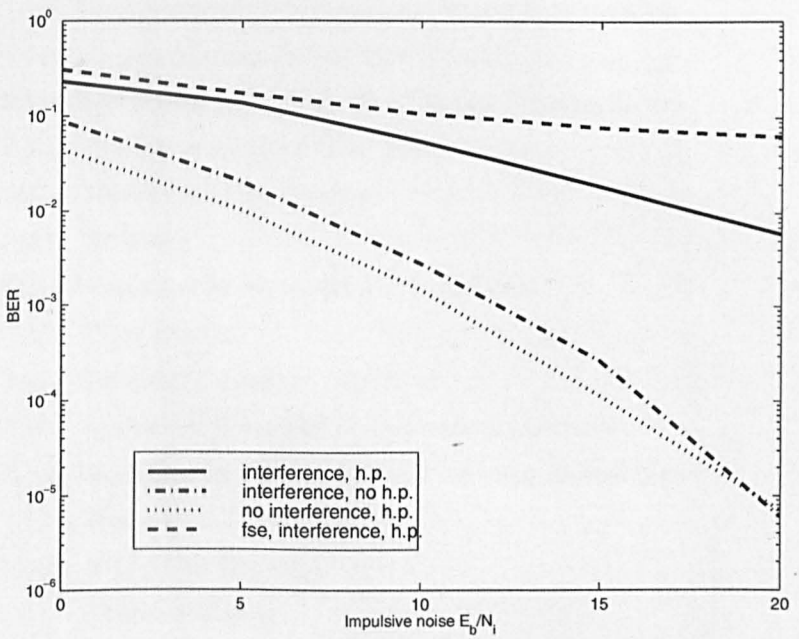

Figure 6.21: FRESH interference removal performance with interference at 30 dB above SOI

a simple magnitude threshold may help, as may a method of repairing the effects of the impulse excision, such as linear interpolation of the remaining waveform.

The conclusions drawn from the results in section 6.4 can be summarised as follows:

- in the absence of interference, hole-punching is very effective against impulsive noise (figure 6.18);
- when the interferer power is low, impulsive noise effects dominate and FRESH filtering can be used successfully with hole-punching (figure 6.19);
- when the interferer power is high, interference dominates; the performance of the FRESH

filter can be seriously degraded by hole-punching (figures 6.20 and 6.21).

Chapter 7

Blind Adaptation of FRESH Filters

Contents

In many communications scenarios, the channel conditions (that is the dispersion and interference) change in an unpredictable way. Traditional techniques to overcome the corruption introduced by the channel usually rely on some form of adaptive filtering which is "trained" by initially or

periodically sending a signal which is known *a priori* at the receiver. Comparison of the received signal with the known signal allows adaptation of the filter to remove the channel distortion and interference.

Blind adaptation techniques do not require any training sequence; they instead use other properties of the signal which are known in advance, for example, its finite alphabet nature or its statistical properties. This can mean using third or higher order statistics or using cyclostationary statistics.

Digital communication requires channel magnitude and phase distortion to be equalised. The second order cyclostationary statistics of a received signal contain magnitude and phase information for the distorting channel [54, 70] while stationary second order statistics contain only magnitude information. For this reason, cyclostationary statistics are used in a number of blind adaptation techniques [64, 63, 65, 105]. The interest in equalisation is largely because dispersive channels are a severe problem in mobile communications. However, the strength of FRESH filtering is its ability to mitigate the effect of additive interference. As the structure contains a number of linear transversal filters, it can also cope with the equalisation problem, at least as well as a simple linear equaliser, but we do not address this problem here, except in a limited way in section 7.2.

It is shown in section 7.1 that in a channel with a flat frequency response, but unknown interference, the cyclic Wiener filter equations can be solved by estimation and inversion of the input signal statistics. No training signal is required. However, this direct approach was not successful in practice due to the numerical problems of inverting a poorly conditioned matrix. So in section 7.2 a new technique is described for blind interference rejection (this does not attempt to correct multipath distortion) which can be applied to FSEs or FRESH filters, for use in situations where the interference can be detected separately from the SOI. In the current work, the concern is not blind adaptation of filters in general, but the blind adaptation of FRESH filters. In sections 7.3 and 7.4 two published algorithms for blindly adapting FRESH filters are found, in the case of the former [25], to be faulty, and in the case of the latter [48], to be severely limited in the possible areas of application, and in fact only equivalent in performance to blind adaptation of a FSE. An improved version of the algorithm described in section 7.4 is proposed in section 7.5.

7.1 Blind Interference Rejection

A particular strength of FRESH filters is their ability to exploit the spectral correlation of a cyclostationary interfering signal. Up to this point when implementing a filter which includes interferer related frequency shifts, we have always adapted all the sub-filters by attempting to drive the output of the whole filter to be, by some criterion, close to a desired signal.

If the situation is such that the channel frequency response is flat, or known, (or equivalently, if there is no, or known, intersymbol interference) then a simpler approach to blindly adapting the receiver filter can be employed. This is based on the replacement of one of the terms in the Wiener filter equation with an equivalent term calculated a priori for the particular intersymbol interference environment expected.

In this section we describe this algorithm and its relation to the equivalent idea applied to a fractionally spaced equaliser. This involves developing a time-domain form of the cyclic Wiener filter equations and a consideration of practical methods of solving these equations. This is shown to be

directly related to standard linear least squares theory, so although we have encountered numerical stability problems, it is expected that these could be overcome by modifying standard techniques.

A simple LMS based algorithm has also been developed which allows blind removal of interference in a system where the channel and interference can be monitored before SOI transmission has started. This is described in section 7.2.

7.1.1 Fractionally spaced equaliser in a flat channel

We first describe a blind adaptation method for a fractionally spaced equaliser in a flat channel. It is well known (see, for example [11]) that the optimum filter for estimating a stationary signal in stationary noise is given by the Wiener filter equation. This has already been stated as equation 3.1 but is given here using a different notation:

$$\mathbf{R_{xx}w} = \mathbf{R}_{dx} \tag{7.1}$$

where $\mathbf{w}$ is the impulse response of the filter, $\mathbf{R_{xx}}$ is the autocorrelation matrix of the input signal vector $\mathbf{x}$, and $\mathbf{R}_{dx}$ is the cross-correlation vector of $\mathbf{x}$ and the desired signal $\mathbf{d}$. $\mathbf{x}$(n) here is the vector of signal samples contributing to the filter output at time n, and is given by

$$\mathbf{x}(n) = [x(n), x(n-1), ..., x(n-N+1)]^T \tag{7.2}$$

where N is the number of filter taps. $d(n)$ is the (scalar) desired signal value at time n. The other terms are defined as follows:

$$\mathbf{R_{xx}} = E[\mathbf{x}(n)\mathbf{x}^H(n)] \tag{7.3}$$

$$\mathbf{R}_{dx} = E[\mathbf{x}(n)d^*(n)] \tag{7.4}$$

A training algorithm such as LMS or RLS will drive a linear filter towards this solution, by directly comparing the filter output with the desired signal. We are interested in the situation where there is no desired signal available in the receiver, so a training algorithm cannot operate in the normal way, and the terms in equation 7.1 cannot be directly calculated.

The approach is to notice that $\mathbf{R}_{dx}$ is not dependent on any additive noise or interference. In equation 7.4 we can write $\mathbf{x}(n) = \mathbf{t}(n) + \mathbf{a}(n)$ where $\mathbf{t}$ is the component of the signal which is correlated with transmitted signal, (which is of course correlated with the desired signal) and $\mathbf{a}$ is the additive noise (which is not). If there is no intersymbol interference (ISI) then $\mathbf{t}$ is equal to the transmitted signal; if there is known ISI then $\mathbf{t}$ is the convolution of the transmitted pulse shape and the channel impulse response.

$$\mathbf{R}_{dx} = E[d(n)(\mathbf{t}(n) + \mathbf{a}(n))] \tag{7.5}$$

$$\mathbf{R}_{dx} = E[d(n)\mathbf{t}(n)] + E[d(n)\mathbf{a}(n)] \tag{7.6}$$

$$\mathbf{R}_{dx} = \mathbf{R}_{dt} \tag{7.7}$$

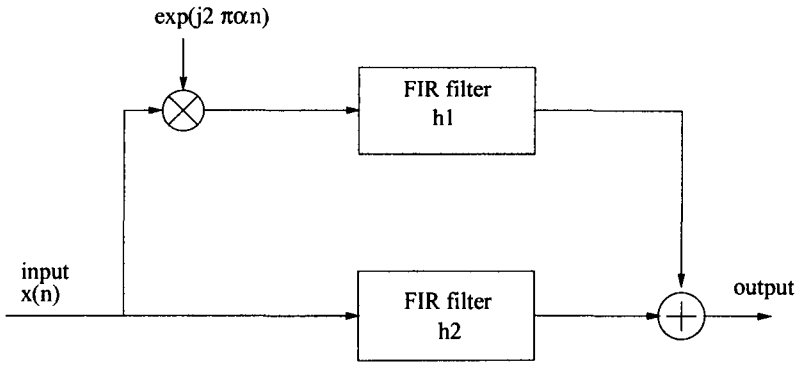

Figure 7.1: Simple FRESH filter

If one assumes that the transmitted data are IID (that is, a white noise sequence) then $\mathbf{R}_{dt}$ is given by the convolution of the pulses for d and $\mathbf{t}$. Equation 7.1 then becomes

$$\mathbf{w} = \mathbf{R}_{\mathbf{xx}}^{-1}\mathbf{R}_{dt} \tag{7.8}$$

and only $\mathbf{R}_{\mathbf{xx}}$ is not known a priori. An estimate of this term is easily calculable by time averaging the outer product of the input vector $\mathbf{x}$. This allows an iterative version of equation 7.8 to be used to update the filter taps $\mathbf{w}$.

This filter is then a blindly adaptable version of the Wiener filter. The "adaptable" part is the noise whitening term $\mathbf{R}_{\mathbf{xx}}$, so any changes to the additive noise can be coped with during reception. The filtering performance of such a filter will be no better than a standard Wiener filter; the advantage is the fact that no training sequence is required to solve the equations. Section 7.1.4 describes the application of this idea to cyclic Wiener filters, or FRESH filters, which can deal more effectively with interference than the fractionally spaced equaliser.

7.1.2 Time domain cyclic Wiener filter equations

A time domain expression of the filtering equations is appropriate for an adaptive version of the cyclic Wiener filter, although an adaptive frequency domain implementation would be possible, if inefficient. Gardner [13] uses a frequency domain form; we present here the time domain equivalent, although for the sake of clarity in presentation, this is done for a two branch FRESH filter only. It is possible to write down the equations for a general M branch filter, but they are cumbersome and do not make the techniques involved any clearer.

The simplest FRESH filter is one consisting of two branches and sub-filters, and one frequency shift (figure 7.1). This frequency shift which is labelled α here, could be at twice the carrier frequency for a purely real signal, or an integer multiple of the baud rate for any bauded signal. Simplifying equations 4.1 gives the frequency domain cyclic Wiener equations for such a filter:

$$S_x(f+\alpha)h_1 + S_x^{\alpha}(f)h_2 = S_{dx}^{\alpha}(f+\frac{\alpha}{2}) \tag{7.9}$$

$$S_x^{\alpha}(f+\frac{\alpha}{2})h_1 + S_x(f)h_2 = S_{dx}^{\alpha}(f) \tag{7.10}$$

which in the time domain are

$$\mathbf{R_{uu}h_1} + \mathbf{R_{xu}h_2} = \mathbf{R}_{du} \tag{7.11}$$

$$\mathbf{R_{xu}h_1} + \mathbf{R_{xx}h_2} = \mathbf{R}_{dx} \tag{7.12}$$

where $\mathbf{u}(n)$ is $\mathbf{x}(n)$ shifted in frequency by α,

$$u(n) = x(n)e^{j2\pi\alpha n} \tag{7.13}$$

and

$$\mathbf{u}(n) = [u(n), u(n-1), ..., u(n-N+1)]^T \tag{7.14}$$

The new correlation matrices are defined in the usual way:

$$\mathbf{R_{uu}} = E[\mathbf{u}(n)\mathbf{u}^H(n)] \tag{7.15}$$

$$\mathbf{R_{xu}} = E[\mathbf{x}(n)\mathbf{u}^H(n)] \tag{7.16}$$

and

$$\mathbf{R_{xx}} = E[\mathbf{x}(n)\mathbf{x}^H(n)] \tag{7.17}$$

as before.This system of equations can be written as a single matrix equation

$$\begin{pmatrix} \mathbf{R_{uu}} & \mathbf{R_{xu}} \\ \mathbf{R_{xu}} & \mathbf{R_{xx}} \end{pmatrix} \begin{pmatrix} \mathbf{h_1} \\ \mathbf{h_2} \end{pmatrix} = \begin{pmatrix} \mathbf{R}_{du} \\ \mathbf{R}_{dx} \end{pmatrix} \tag{7.18}$$

7.1.3 Practical issues in solving cyclic Wiener filter equations

The terms $\mathbf{h_1}$ and $\mathbf{h_2}$ represent the impulse responses of the two sub-filters of the two branch FRESH filter. They are therefore both $N \times 1$ vectors, where N is the number of taps in each filter. For the matrix multiplications to be defined, the $\mathbf{R_{xx}}$ matrices must therefore be of dimension $N \times N$. The $\mathbf{R}_{dx}$ terms are similarly constrained to be $N \times 1$ vectors. The large $\mathbf{R}$ matrix on the left hand side of equation 7.18 is therefore of dimension $2N \times 2N$ and other two terms in this equation are $2N \times 1$ vectors.

Writing out the solution of equation 7.18 explicitly

$$\begin{pmatrix} \mathbf{h_1} \\ \mathbf{h_2} \end{pmatrix} = \begin{pmatrix} \mathbf{R_{uu}} & \mathbf{R_{xu}} \\ \mathbf{R_{xu}} & \mathbf{R_{xx}} \end{pmatrix}^{-1} \begin{pmatrix} \mathbf{R}_{du} \\ \mathbf{R}_{dx} \end{pmatrix} \tag{7.19}$$

it is clear that the inverse of this $2N \times 2N$ cyclic correlation matrix is required. In the adaptive filter application it is useful to be able to solve this equation iteratively, by updating and improving the filter solution and the terms contributing to it with the arrival of each new signal sample.

It is necessary to use a time-averaged approximation to the statistically defined correlation matrices and vectors. For example

$$\mathbf{R}_{xx} = \sum_{i=N}^{M} (\mathbf{x}(i)\mathbf{x}^H(i)) \tag{7.20}$$

However to allow the filter to adapt to changing channel conditions, a definition is used which allows the older data's influence on the result to decrease with time. This uses an exponential weighting factor λ^{M-i} so that the autocorrelation matrix is defined as

$$\mathbf{R}_{xx} = \sum_{i=N}^{M} \lambda^{M-i}(\mathbf{x}(i)\mathbf{x}^H(i) \tag{7.21}$$

which can be written in the following iterative form

$$\mathbf{R}_{xx}(n) = \lambda \mathbf{R}_{xx}(n-1) + (\mathbf{x}(i)\mathbf{x}^H(i)) \tag{7.22}$$

Equation 7.22 gives a straightforward method for calculating all the correlation terms in equation 7.19. The factor (λ is a parameter to be adjusted to best suit the rate of change of channel conditions. An obvious way of approaching the solution of equation 7.19 would be as each sample is received, calculate the auto- and cross-correlation matrices, and then invert the cyclic correlation matrix. Note that although this matrix is a 2×2 matrix of $N \times N$ matrices, it can be inverted by treating it as a normal $2N \times 2N$ matrix of scalar elements.

Matrix inversion lemma

The inversion of large matrices requires considerable processing time, so for a real time application recalculating a correlation matrix and inverting it with every new sample would not normally be suitable. The inversion of a normal autocorrelation matrix is part of the recursive least squares algorithm (RLS) and there the approach is to update the inverse of the matrix directly, using the matrix inversion lemma, which reduces the amount of computation required considerably. This lemma is described below, firstly as applied to a standard autocorrelation matrix, and then to a cyclic correlation matrix. This description is taken from [17].

Let $\mathbf{A}$ and $\mathbf{B}$ be two positive-definite $M \times M$ matrices related by

$$\mathbf{A} = \mathbf{B}^{-1} + \mathbf{C}\mathbf{D}^{-1}\mathbf{C}^{\mathbf{H}} \tag{7.23}$$

where $\mathbf{D}$ is a positive-definite $N \times N$ matrix, and $\mathbf{C}$ is an $M \times N$ matrix. The matrix lemma states that the inverse of the matrix $\mathbf{A}$ may be expressed as

$$\mathbf{A}^{-1} = \mathbf{B} - \mathbf{B}\mathbf{C}(\mathbf{D} + \mathbf{C}^{\mathbf{H}}\mathbf{B}\mathbf{C})^{-1}\mathbf{C}^{\mathbf{H}}\mathbf{B} \tag{7.24}$$

The lemma is applied to a correlation matrix by defining the four matrices as follows

$$\begin{aligned} \mathbf{A} &= \mathbf{\Phi}(n) \\ \mathbf{B}^{-1} &= \lambda\mathbf{\Phi}(n-1) \\ \mathbf{C} &= \mathbf{x}(n) \\ \mathbf{D} &= 1 \end{aligned}$$

to give this recursive equation for the inverse of the correlation matrix

$$\mathbf{\Phi}^{-1}(n) = \lambda^{-1}\mathbf{\Phi}^{-1}(n-1) - \frac{\lambda^{-2}\mathbf{\Phi}^{-1}(n-1)\mathbf{x}(n)\mathbf{x}^H(n)\mathbf{\Phi}^{-1}(n-1)}{1+\lambda^{-1}\mathbf{x}(n)\mathbf{\Phi}^{-1}(n-1)\mathbf{x}(n)} \tag{7.25}$$

This equation gives the current value of the correlation matrix from the previous value and the current value of the input tap vector.

This lemma can be applied to the cyclic autocorrelation matrix as well. First of all to simplify the notation $\mathbf{\Phi}$ is assigned to be the cyclic autocorrelation matrix. That is

$$\mathbf{\Phi} = \begin{pmatrix} \mathbf{R}_{\mathbf{uu}} & \mathbf{R}_{\mathbf{xu}} \\ \mathbf{R}_{\mathbf{xu}} & \mathbf{R}_{\mathbf{xx}} \end{pmatrix} \tag{7.26}$$

Notice that by defining $\mathbf{z}(n) = \begin{bmatrix} \mathbf{u(n)} \\ \mathbf{x}(n) \end{bmatrix}$, $\mathbf{\Phi}$ can be written as

$$\begin{pmatrix} \mathbf{R}_{\mathbf{uu}} & \mathbf{R}_{\mathbf{xu}} \\ \mathbf{R}_{\mathbf{xu}} & \mathbf{R}_{\mathbf{xx}} \end{pmatrix} = E[\mathbf{z}(n)\mathbf{z}^H(n)] = E[\mathbf{u}(n)\mathbf{x}(n)\mathbf{x}^H(n)\mathbf{u}^H(n)] \tag{7.27}$$

and to apply the matrix inversion lemma the following identifications are also made:

$$\begin{aligned} \mathbf{A} &= \mathbf{\Phi}(n) \\ \mathbf{B}^{-1} &= \lambda\mathbf{\Phi}(n-1) \\ \mathbf{C} &= \begin{bmatrix} \mathbf{u(n)} \\ \mathbf{x}(n) \end{bmatrix} = \mathbf{z}(n) \\ \mathbf{D} &= 1 \end{aligned}$$

The recursive relationship for calculation of the inverse of the correlation matrix is therefore

$$\mathbf{\Phi}^{-1}(n) = \lambda^{-1}\mathbf{\Phi}^{-1}(n-1) - \frac{\lambda^{-2}\mathbf{\Phi}^{-1}(n-1)\mathbf{z}(n)\mathbf{z}^H(n)\mathbf{\Phi}^{-1}(n-1)}{1+\lambda^{-1}\mathbf{z}(n)\mathbf{\Phi}^{-1}(n-1)\mathbf{z}(n)} \tag{7.28}$$

What we have shown here is effectively the first stage of applying the recursive least squares (RLS) algorithm to a FRESH filter. This adaptation method is much faster than LMS but also requires a training signal. This could be a future line of enquiry: is this approach more efficient than running a separate version of the algorithm for each filter branch? However the purpose of presenting this method of calculating the inverse of the cyclic autocorrelation matrix is to indicate how it may be used in a blind adaptation scheme.

7.1.4 Noise whitening FRESH filter

A solution of the FRESH filter equations can be found without a training signal when the channel is flat, or of other known frequency response. The arguments behind this are the same as those applied in section 7.1.1 to solve the Wiener filter equation without the use of a training signal.

It is clear that the cross-correlation terms $\mathbf{R}_{du}$ and $\mathbf{R}_{dx}$ in the cyclic Wiener filter equations 7.19 are unaffected by any additive noise corrupting the received signal x. We define a signal t to be the result of transmitting the data over the channel without any additive noise. s is defined to be t shifted in frequency by the cyclic frequency α. If there is no ISI introduced by the channel, then t is the transmitted signal; otherwise t is the convolution of the transmitted signal and the channel impulse response. $\mathbf{R}_{du}$ and $\mathbf{R}_{dx}$ can then be replaced with $\mathbf{R}_{ds}$ and $\mathbf{R}_{dt}$ where $\mathbf{R}_{dt}$ is the cross-correlation matrix of the desired signal with the signal t, and $\mathbf{R}_{ds}$ is the cross-correlation matrix of the desired signal with s.

These matrices can be worked out a priori, and, providing the data transmitted are independent and identically distributed, they are given by the cross-correlation of the pulse shapes of d and t, and of d and s.

So to blindly adapt the FRESH filter it is necessary to solve the following filter equation:

$$\begin{pmatrix} \mathbf{h_1} \\ \mathbf{h_2} \end{pmatrix} = \begin{pmatrix} \mathbf{R_{uu}} & \mathbf{R_{xu}} \\ \mathbf{R_{xu}} & \mathbf{R_{xx}} \end{pmatrix}^{-1} \begin{pmatrix} \mathbf{R}_{ds} \\ \mathbf{R}_{dt} \end{pmatrix} \tag{7.29}$$

that is,

$$\begin{pmatrix} \mathbf{h_1} \\ \mathbf{h_2} \end{pmatrix} = \mathbf{\Phi}^{-1} \begin{pmatrix} \mathbf{R}_{ds} \\ \mathbf{R}_{dt} \end{pmatrix} \tag{7.30}$$

The inverse autocorrelation matrix $\mathbf{\Phi}^{-1}$ can be calculated efficiently as described in section 7.1.3.

MATLAB and SPW implementations

Unexpected results were obtained when implementing this algorithm in MATLAB and SPW. It appears that these are due to numerical instabilities inherent in inverting the cyclic autocorrelation matrix.

This matrix inverse was calculated in two different ways, to explore the numerical errors. One was a straightforward inverse using the MATLAB inv() function. The other was to treat the $2N \times 2N$ matrix as a 2×2 matrix of matrices. That is the inverse was calculated as:

$$(\mathbf{R_{uu}R_{xx}} - \mathbf{R_{xu}R_{xu}})^{-1} \begin{pmatrix} \mathbf{R_{xx}} & -\mathbf{R_{xu}} \\ -\mathbf{R_{xu}} & \mathbf{R_{uu}} \end{pmatrix} \tag{7.31}$$

Mathematically this is equivalent to the normal inverse but in practice it gives much poorer results. It shows however that the inverse cyclic correlation matrix should exhibit the same sort of symmetry as the cyclic correlation matrix; for example the upper right quadrant should be equal to the lower left one. This is not the case for either form of inverse. At the edges of each quadrant, anomalous values appear, presumably as a result of the discontinuity in values between the three different correlation functions contributing to the matrix. The situation is further complicated by

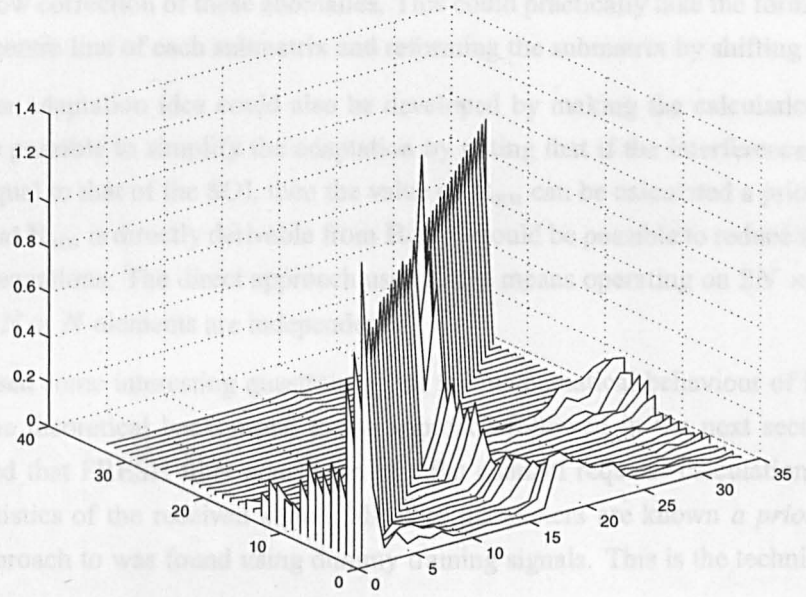

Figure 7.2: Cyclic correlation matrix multiplied by its inverse

the fact that at low white noise levels, the condition number of the whole matrix is high - indicating that any solution to equation 7.30 will be highly susceptible to inaccuracies in the cyclic autocorrelation. As this quantity is a time-averaged approximation to a statistical mean, such inaccuracies will be invariably present. As an example, figure 7.2 shows the result of multiplying the cyclic autocorrelation matrix by its inverse. The result should be the identity matrix. In this plot the x, y values indicate the position within the matrix, and the z-axis indicates the value at that position. Ideally this would be 1 on the main diagonal and zero everywhere else.

For these reasons a successful implementation of direct solution of the time domain equations could not be achieved. There are however a number of ways in which this work could be continued. The poor conditioning of the standard autocorrelation matrix is a well recognised one in the field of adaptive signal processing [17]. One approach in minimising the problems is to operate on the data matrix, $\mathbf{A} = [\mathbf{x}(n), \mathbf{x}(n-1)...\mathbf{x}(n-m)]$ where $\mathbf{x}(n)$ is a vector of input values to the filter at time n. The autocorrelation matrix can be expressed as $\mathbf{A}^H\mathbf{A}$ which allows a reformulation of the Wiener filter equations in terms of $\mathbf{A}$ rather than $\mathbf{\Phi}$. The condition number of A is the square root of that of $\mathbf{\Phi}$, so the numerical problems can be reduced by this approach.

Other ideas

It can be shown that the autocorrelation matrix is Toeplitz if the signal concerned is stationary. If we can assume we are dealing with a signal which is stationary except for the cyclostationarity we are interested in, then we can see that the cyclic autocorrelation matrix has a block Toeplitz structure. Each quadrant of the matrix is individually Toeplitz. It is in fact inefficient to represent the cyclic autocorrelation in this matrix form in this case, because each line of each sub-matrix is a time shifted version of the other lines in the submatrix. The full matrix representation allows non-cyclic time variations to be accounted for. It appeared to be the case that the errors described in section 7.1.4 were due to anomalies towards the edge of each submatrix of the inverse cyclic autocorrelation matrix. Knowing however that these submatrices should exhibit certain symme-

tries could allow correction of these anomalies. This could practically take the form, for example, of taking the centre line of each submatrix and reforming the submatrix by shifting this line.

The blind filter adaptation idea could also be developed by making the calculations simpler. It would also be possible to simplify the adaptation by noting that if the interference has no cyclic frequencies equal to that of the SOI, then the value of $\mathbf{R_{xu}}$ can be calculated a priori. Given this, and the fact that $\mathbf{R_{uu}}$ is directly derivable from $\mathbf{R_{xx}}$ it should be possible to reduce the complexity of the matrix equations. The direct approach used above means operating on $2N \times 2N$ matrices, whereas only $N \times N$ elements are independent.

This work raised some interesting questions about the mathematical behaviour of FRESH filters and gave some theoretical background to the approach followed in the next section, in that it has highlighted that FRESH filter adaptation in a flat channel requires calculation of the cyclostationary statistics of the received signal; all other parameters are known *a priori*. A simpler alternative approach to was found using dummy training signals. This is the technique described in section 7.2.

7.2 Blind Interference Rejection with a Dummy Training Signal

This novel technique is suitable for use in cases where a channel can be observed for some time before SOI transmission starts. If interference is present during that period, and remains present with the same second order statistical parameters during SOI transmission, then a FRESH filter can be adapted to minimise the effect of the interferer before SOI transmission starts. The constraints referred to above are such that no extra distortion is introduced into the SOI when it is received. This is achieved by adding a dummy training signal at the receiver, which has the same second order statistics as the SOI, and generating a dummy desired signal carrying the same data as the dummy training signal. The LMS algorithm can then be used to adapt the filter output to the dummy desired signal.

The result is that if the SOI is received in the presence of white noise and interference only (that is, with no channel dispersion) then the FRESH filter can be adapted to its optimum solution without any training signal being transmitted. If the channel introduces dispersion, then further trained adaptation will be required, but a shorter training time is required, and the initial MSE in reception is much lower than without pre-adaptation. The latter property means that initial detection of the SOI (for synchronisation) is much more reliable. This technique has been published as [1].

Starting again from the FSE/Wiener filter case, we first note that the Wiener filter equation 3.1 simplifies to equation 7.8 when the only corruption to the transmitted signal is from additive noise or interference (which is equivalent to saying that the training signal gives no information about the interference). The solution $\mathbf{w}$ to equation 7.8 is unchanged if $\mathbf{t}$ is replaced by $\mathbf{t'}$, a dummy training signal, and d with d', a dummy desired signal, as long as $\mathbf{R_{t't'}} = \mathbf{R_{tt}}$ and $\mathbf{R_{d't'}} = \mathbf{R_{dt}}$ (because $\mathbf{R_{xx}} = \mathbf{R_{tt}} + \mathbf{R_{aa}}$).

To make these two conditions hold, we make the common assumption that the data carried by $\mathbf{t}$, $\mathbf{t'}$, d and d' are independent and identically distributed. It is then only necessary to make the pulse shapes of $\mathbf{t}$ and $\mathbf{t'}$, and of d and d' the same. If there is no channel dispersion then choosing the pulse shape of $\mathbf{t'}$ to be the same as the transmitted shape will give the optimum filter response.

If there is dispersion, then further trained adaptation is required for optimum response. Note that symbol transitions need not be synchronised in $\mathbf{t}$ and $\mathbf{t}'$ as $\mathbf{R}_{\mathbf{tt}'}$ is time invariant.

7.2.1 Pre-adaptation of a FRESH filter

The technique described above can be applied to a receiver where a FRESH filter replaces the FSE above. In this case equation 3.1 becomes

$$\mathbf{W} = \mathbf{R}_{\mathbf{zz}}{}^{-1}\mathbf{R}_{\mathbf{z}d} \tag{7.32}$$

where $\mathbf{W}$ is now a concatenation of the impulse response vectors of all the sub-filters of the FRESH filter. $\mathbf{z} = [\mathbf{x}_{\alpha 1}, \mathbf{x}_{\alpha 2}, ..., \mathbf{x}_{\alpha M}, \mathbf{x}_{-\beta 1}, \mathbf{x}_{-\beta 2}, ..., \mathbf{x}_{-\beta N}]^T$ is a concatenation of the input vectors to each sub-filter which are frequency shifted by the values $\alpha 1, \alpha 2, ...\alpha M, -\beta 1, -\beta 2, ..., -\beta N$, which are some or all of the cyclic frequencies present in the input signal. d is defined as before. This equation gives the optimum filter for extracting the desired signal when the input is cyclostationary [2].

Again it is possible in the absence of the SOI to add a dummy training signal to the filter input. We write $\mathbf{z} = \mathbf{u} + \mathbf{v}$, where $\mathbf{u}$ represents the SOI component of the received signal concatenated with frequency shifted versions of itself and $\mathbf{v}$ represents a similar structure for the noise and interference components. So equation 7.32 becomes

$$\mathbf{W} = (\mathbf{R}_{\mathbf{uu}} + \mathbf{R}_{\mathbf{vv}})^{-1}\,\mathbf{R}_{\mathbf{u}d} \tag{7.33}$$

$\mathbf{W}$ is optimum if a dummy training signal $\mathbf{u}'$ is added such that $\mathbf{R}_{\mathbf{u}'\mathbf{u}'} = \mathbf{R}_{\mathbf{uu}}$ and $\mathbf{R}_{\mathbf{u}'\mathbf{d}} = \mathbf{R}_{\mathbf{u}d}$. This requires that the dummy training signal has the same cyclic autocorrelation as the real one. However, unlike its stationary equivalent, this quantity is time variant so ideally the dummy training signal should be synchronised with the real one. However a frequency shift filter works both by adding correlated components of the SOI, and by subtracting correlated components of the cyclostationary interferer from each other. Its performance based on the latter technique is not affected by the dummy signal timing. So, for example, when the interference power is high, this technique is particularly useful.

7.2.2 Simulation of the algorithm

The ideas described above were verified by a simulation structured as in figure 7.3.

The SOI and interferer were both QPSK, with transmit filters which gave a root raised cosine spectrum with 100% excess bandwidth. IID data was used for both and for the dummy training signal and dummy desired signal. The desired signal and dummy desired signal had 100% excess bandwidth raised cosine spectra.

Figures 7.4 to 7.6 show the mean squared error of the filter output varying with time as the filter adapts. The curves showing the performance of pre-adaptation against the interference start at $t = 0$. The noise, interference, dummy training signal and dummy desired signal all start at $t = 0$. From $t = 0$ to $t = 3$ the MSE represents the error in extracting the dummy desired signal from the

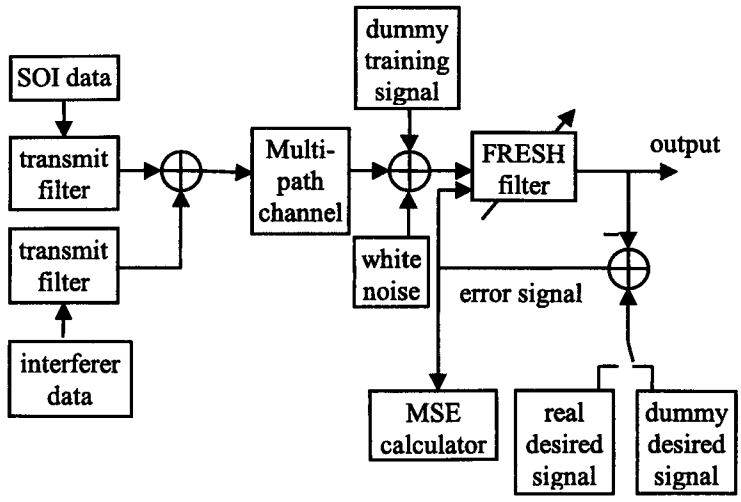

Figure 7.3: Simulation block diagram

filter input. At $t = 3$ the real SOI starts, as does the real desired signal. The two dummy signals are turned off at this point. From $t = 3$ the MSE represents the error in extracting the real desired signal from the filter input. There are two pre- adaptation curves: one where the dummy signals are symbol synchronous with the real signals they replace (solid line), and one when they are half a symbol delayed (dot-dash line). These two situations are the best and worse cases respectively for pre-adaptation performance.

Each graph also has a curve representing the normal method of adaptation, relying on the real training signal for all adaptation. This starts at $t = 3$ to allow comparison of the length of real training sequence required in each case.

The symbol rate of the SOI and the interferer was 1200 Hz. The interferer carrier frequency was 840 Hz, while the SOI was at baseband; the signals therefore overlap more than half of each other's spectrum. The white noise level was set to give E_b/N_0 of 20 dB relative to the SOI.

The filter consisted of two sub-filters. The input to one was the received signal with no frequency shift applied. The input to the second was the received signal shifted by 1200 Hz added to the shifted signal shifted by -1200 Hz. This filter exploits all the cyclostationarity of the SOI and interferer used here. Each sub-filter had 33 taps, spaced at 8 taps per symbol. The LMS algorithm was used to adapt the filter [17]. The sub-filters' initial impulse responses were set to the solutions of the cyclic Wiener filter equations for the case of no interference or multipath, and E_b/N_0 of 20 dB.

The multipath channel applied in figure 7.4 and 7.5, consisted of 2 paths, one delayed by 2 ms and with a gain of 0.5 times that of the main path.

Figure 7.4 shows the performance with multipath and an interferer of power twice that of the SOI. Figure 7.5 shows the result of the same scenario except the interferer power is 4 times that of the SOI. Figure 7.6 shows the effect of an interferer of 4 times the power of the SOI but with no multipath propagation present.

The horizontal line in each graph is at a value of 15% above the final MSE. This can be used as an indicator of when adaptation is effectively complete.

7.2.3 Summary

The technique described here offers two benefits: one is the increased likelihood of the real training sequence being detected when its transmission starts; the second is a reduced real training sequence length required to adapt to the multipath channel. The first of these is indicated by the value of each curve at $t = 3$. This shows the error at the start of the real training sequence. It is clear that pre- adaptation significantly reduces this error in all cases, but especially with the higher interferer power (figures 7.5 and 7.6).

The reduced conventional training time is shown by the times from $t = 3$ to where each curve crosses the horizontal line, which indicates completed adaptation. When the interferer power is high there is an advantage whether the dummy signal is in phase with the real signal or not. If the interferer power is lower then any advantage depends on correct dummy signal timing.

In the latter case, it would be possible have more than one filter adapted by dummy signals delayed by varying amounts. When the real transmission starts, the filter with the best output can be selected for further training against multipath.

In the case where there is no multipath (figure 7.6) then there is an advantage in pre-adaptation with the dummy signal out of phase with the real one, and when they are in phase, then perfect blind adaptation occurs.

It should be noted that the conventional training curves represent a best case scenario, where the real signal is synchronised with the initial filter response. The optimum response of a FRESH filter is dependent on the relative phase of the data symbols and the frequency shifting phasors in the filter. In the simulations presented here, the real training signal and phasors had the same relative timing as was used for calculating the initial filter tap settings. We have therefore shown the quickest conventional training curve of the FRESH filter.

The principle being used here, that *a priori* information about the interference can be used to initialise a filter is also used, in a different way, in [106] to prevent the CMA locking on to an interferer instead of the SOI.

7.3 Blind Adaptation by Maximising of Correlation

This section describes the implementation of an algorithm proposed by Zhang [25] and also used in [47] to adapt a FRESH filter without the use of a training sequence. However, it is shown that there are flaws in this algorithm. We explain here what the problems are and suggest a reason for the authors' error. A second closely related algorithm is described in section 7.4. There are limitations in the usefulness of this algorithm but a novel modification to it, using the same fundamental idea of maximising correlation, is proposed in section 7.5.

The algorithm in [25] uses the fact that if a signal $x(n)$ contains a cyclostationary signal with cyclic frequency α, then $x(n)$ and $x(n)e^{j2\pi\alpha n}$ are correlated. Any noise components of $x(n)$, provided they have different or no cyclic frequencies will not be correlated with $x(n)e^{j2\pi\alpha n}$. A structure such as that shown in figure 7.7 is proposed for blind adaptation. If the two FRESH filters are adapted so as to maximise the correlation of their outputs then the idea is that this will maximise the output of the SOI and minimise the output of any other signals, so maximising the

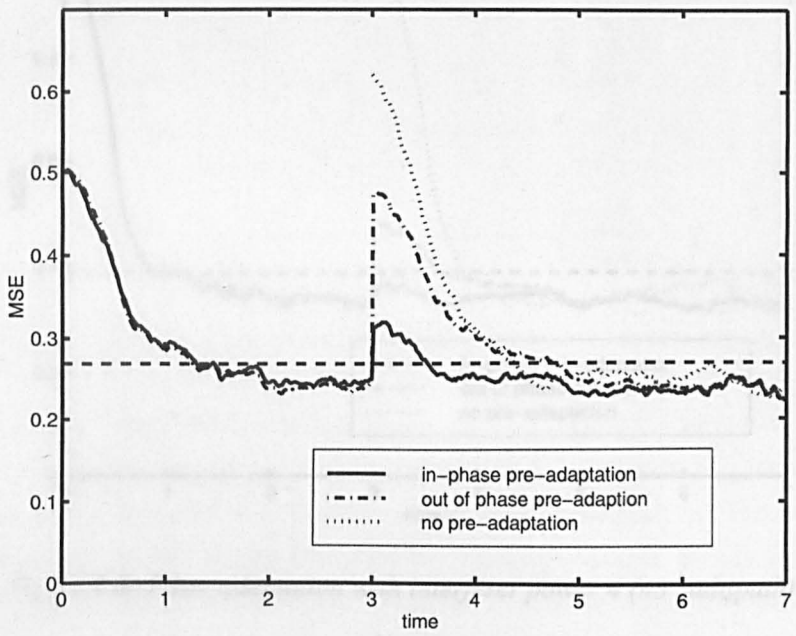

Figure 7.4: Filter adaptation with multipath and interferer power 2

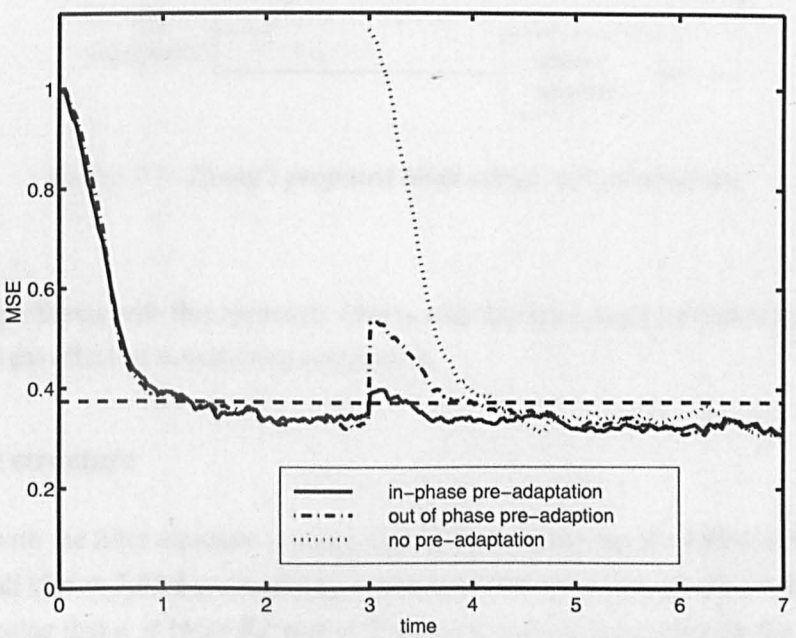

Figure 7.5: Filter adaptation with multipath and interferer power 4

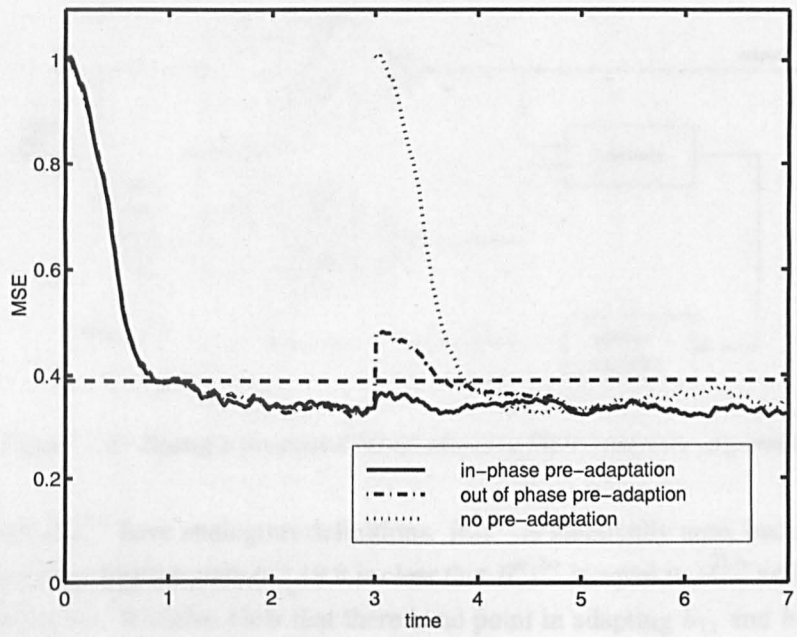

Figure 7.6: Filter adaptation with interferer power 4 (no multipath)

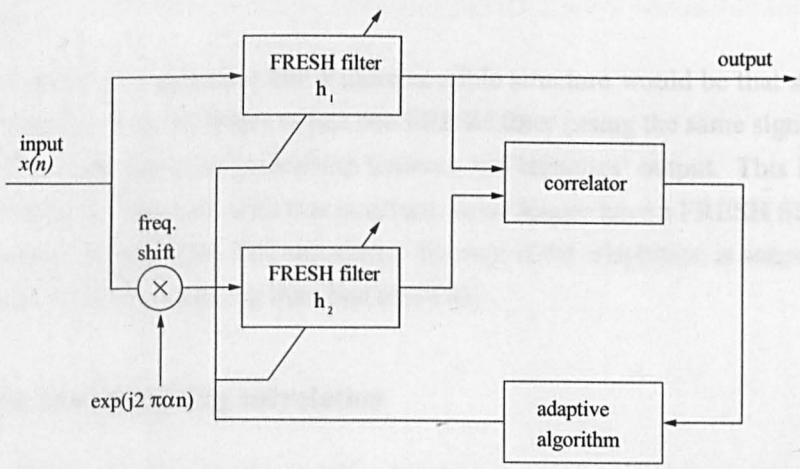

Figure 7.7: Zhang's proposed blind adaptive filter structure

SNR.

There are two problems with this approach. One is with the filter structure shown in figure 7.7, the other concerns the effect of maximising correlation.

7.3.1 Filter structure

The problem with the filter structure is made clearer by drawing out the FRESH filters shown in figure 7.7 in full (figure 7.8). For simplicity, we assume that a two branch filter is being used (for example, assuming that α is twice the carrier frequency, and we are exploiting the carrier related correlation of a BPSK signal). We see now that the sub-filters h_{11} and h_{22} are both operating on the same signal, $x(n)e^{j2\pi\alpha n}$, while h_{12} operates on $x(n)$ and h_{21} operates on $x(n)e^{j4\pi\alpha n}$.

We use the notation $R_{xx}^{0,\alpha}$ for the correlation between $x(n)$ and $x(n)e^{j2\pi\alpha n}$ as defined in section

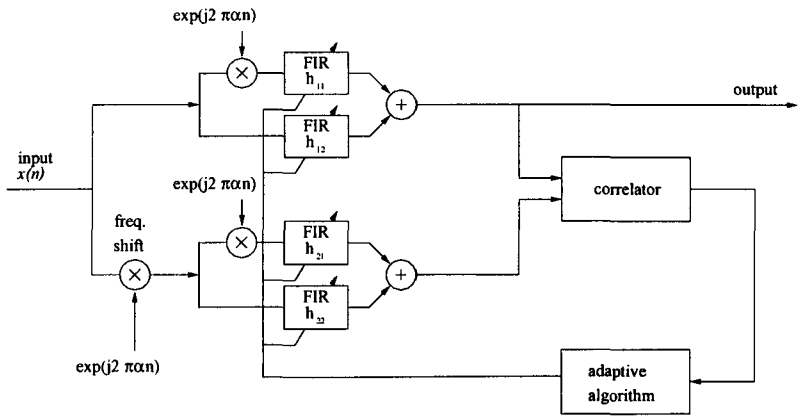

Figure 7.8: Zhang's proposed blind adaptive filter structure expanded

2.1.4. $R_{xx}^{0,2\alpha}$ and $R_{xx}^{\alpha,2\alpha}$ have analogous definitions. $R_{xx}^{0,2\alpha}$ is identically zero, because 2α is not a cyclic frequency of x. From equation 2.19 it is clear that $R_{xx}^{\alpha,2\alpha}$ is equal to $R_{xx}^{0,\alpha}$ and so provides no additional information. It is also clear that there is no point in adapting h_{11} and h_{22} to maximise the correlation of their outputs as they both have the same input signal. So any non-zero value of $h_{12} = h_{22}$ will give a maximum output correlation of 1, without necessarily passing any of the signal of interest.

From these arguments it is apparent that a more sensible structure would be that shown in figure 7.9. Here we adapt the two sub-filters of just one FRESH filter (using the same signal assumptions as in figure 7.8) to maximise the correlation between the branches' output. This is the structure adopted by Wong [47]. However, with this structure we no longer have a FRESH filter - the output is simply the output of one of the FIR sub-filters. So even if the adaptation is successful, the final output performance will be no better than that of a FSE.

7.3.2 Failure of maximising correlation

The second problem with the technique of maximising correlation is more serious. Although the optimum filter solutions given in [13] result in a high level of correlation between the two filter branches' outputs, this is not the maximum correlation achievable with the filter shown in figure 7.9. For this reason, maximising the correlation of the sub-filters' outputs is fundamentally flawed as a method of blindly adapting the FRESH filter.

The algorithm relies maximising a cost function J where:

$$\max_{\mathbf{h_1},\mathbf{h_2}} J(\mathbf{h}) = \max_{\mathbf{h_1},\mathbf{h_2}} \frac{|R_{yr}|^2}{|R_{yy}||R_{rr}|} \tag{7.34}$$

Here $\mathbf{h}$ is the vector which is the concatenation of the subfilter impulse responses, y is the output of subfilter h_1 and r is the output of subfilter h_2.

This filter is able to adapt both branches to pass just a narrow frequency range. We have seen that in the absence of noise there is perfect correlation between the frequency shifted signals where their spectra overlap, although there is, of course, a variation in the absolute value of the correlation function according to the amplitudes of the signals. In the presence of noise the correlation is no

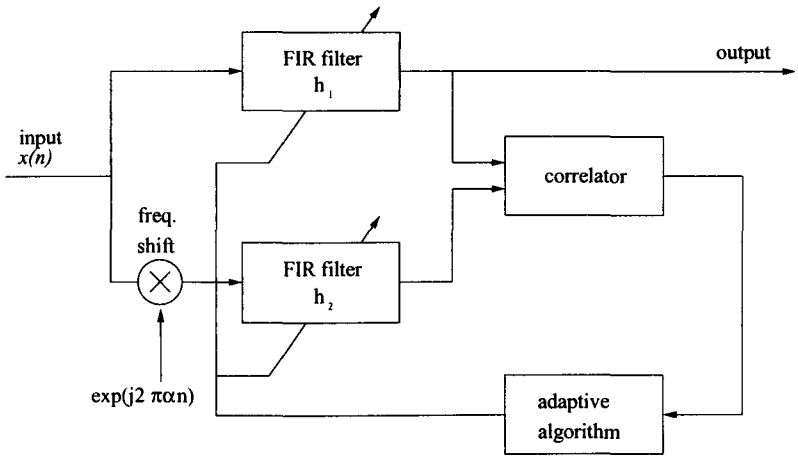

Figure 7.9: Wong's proposed blind adaptive filter structure

longer perfect, and its value will depend on the SNR of the two signals. If the noise is white then it is clear that the maximum correlation occurs at the maximum value of the cross-spectral density of the two signals. The filter can then achieve the maximum correlation of the two branches' outputs by passing only the frequency at which the cross-spectral density is maximum. Passing any other frequencies in addition to this will involve adding signals with lower correlation and will so reduce the overall correlation of the sub-filters' outputs.

For this reason the algorithm in [25] results in the filters passing a very narrow frequency band at the frequency with maximum cross-spectral density. This is illustrated by the simulation results in the next section.

7.3.3 Simulation of correlation maximisation algorithm

The unsuitability of maximising correlation as a method for blind adaptation was demonstrated by implementing the algorithm from [25] in MATLAB, but using the filter structure of [47] and figure 7.9. A BPSK SOI was corrupted with AWGN only - no modulated interferer was added. The output was as shown in figure 7.10 with E_b/N_0 of 12 dB, a baud rate of 0.05 Hz, a carrier frequency f_c of 0.1 Hz, a frequency shift of $2f_c$ with 200 taps in each sub-filter, and a sampling frequency of 1 Hz. The graph shows the spectrum of the unshifted input $x(n)$ (rectangular pulse shaped BPSK) and the resulting magnitude frequency response of the filter h_1. It is clear that the filter is passing only a narrow frequency range around the carrier frequency, where the cross-correlation with the frequency shifted input is maximum.

The experiment was repeated with a QPSK signal of interest and the frequency shift α equal to the SOI baud rate. All other parameters were as in the BPSK example above. The results in figure 7.11 show that the filters select the frequency at which the two frequency responses cross. This is because the maximum cross-spectral density of $x(n)$ and $x(n)e^{j2\pi\alpha n}$ is at this frequency.

This algorithm was taken from the field of antenna array processing where it is known as SCORE (spectral self-coherence restoral) [12]. This link provides another interpretation of why the filter algorithm is not suitable for adapting FRESH filters. SCORE is used to blindly steer an antenna array towards a cyclostationary signal.

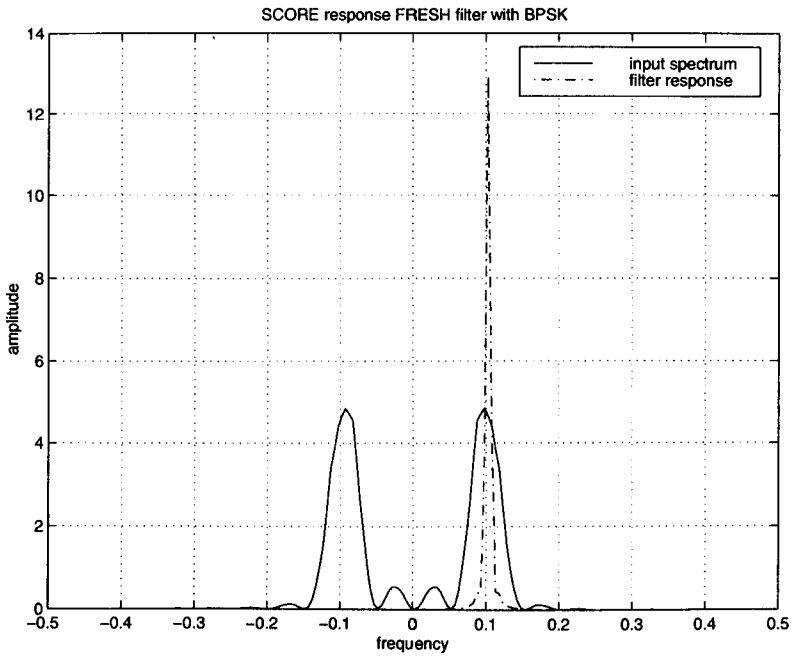

Figure 7.10: Zhang's filter: input spectrum and frequency response for BPSK (200 taps)

A plane wave arriving at a linear array at an angle θ will have a path difference to neighbouring array elements of $dsin\theta$. Weighting the signal from each antenna element steers the beam to the desired direction. So the antenna steering operation is weighting and summing several delayed versions of the signal. This is a spatial version of a standard (temporal) linear transversal filter. The delay elements of the temporal filter correspond to the physical separation of the elements in the spatial filter. Applying SCORE to two spatial filters (one with cyclic frequency shifted input) will cause the filters to steer the beam to the direction which maximises correlation. The "single frequency" output problem in temporal filtering is equivalent to a "single direction" in antenna array processing, which is exactly what is desired for beam steering.

The current author has been in correspondence with the authors of [47] and it seems that they only achieved successful adaptation when the interferer and SOI were well separated in frequency. It may be that they were using a filter with an insufficient number of taps to give a true picture of what was happening. Fewer filter taps means poorer frequency resolution, and a narrow frequency response could be spread out to one which looks like the response required to extract the SOI, if there were few enough taps. Figure 7.12 shows the result of repeating the simulation of figure 7.10 but with only 20 taps in each sub-filter.

Related algorithms are described in sections 7.4 and 7.5 which overcome this notch filtering problem, effectively by maximising correlation across the whole bandwidth of the SOI.

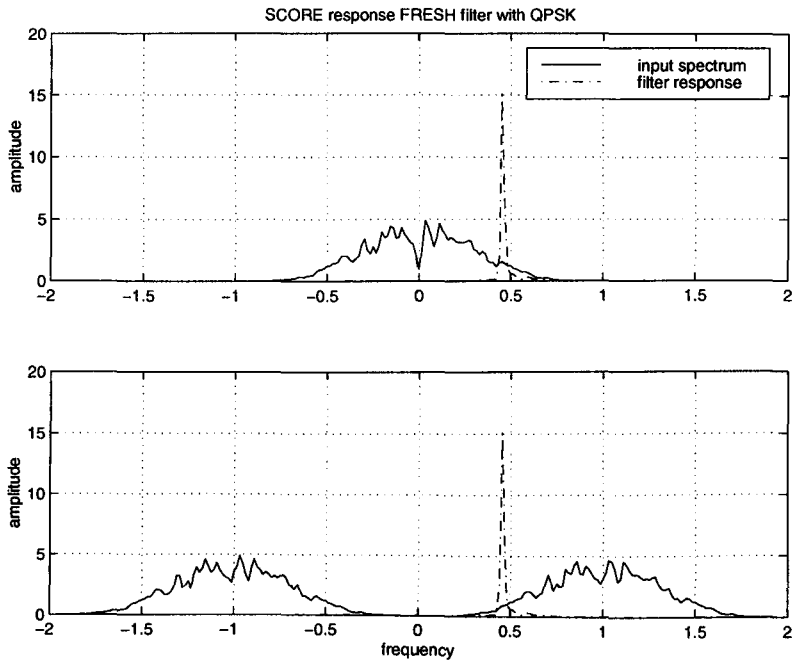

Figure 7.11: Zhang's filter: input spectrum and frequency response for QPSK (200 taps)

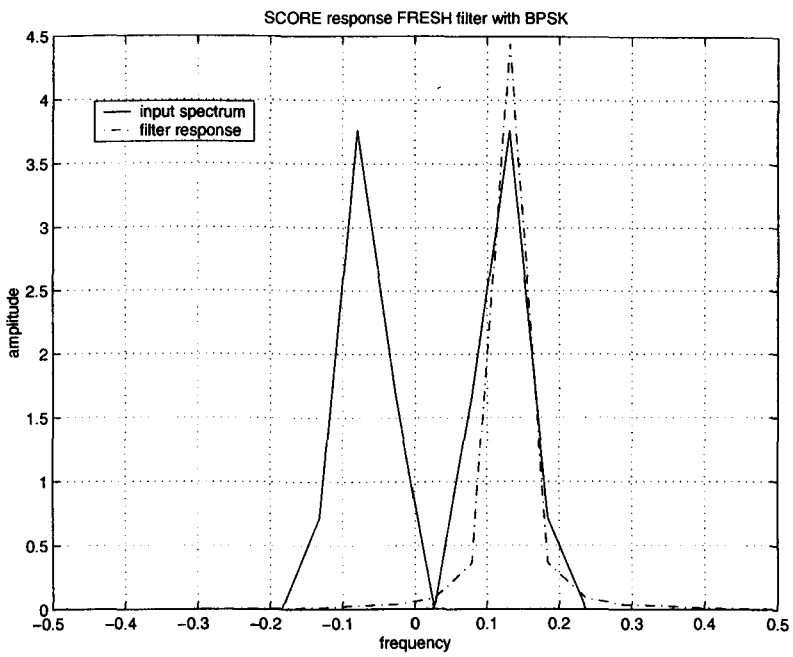

Figure 7.12: Zhang's filter: input spectrum and frequency response for BPSK (20 taps)

7.4 Blind Adaptation by Training with a Frequency Shifted Input

7.4.1 Principles

Ideally an adaptive filter will be trained using a locally generated reference signal which is an exact replica of the actual desired signal (the desired signal, in a communications context is normally of course the same as the transmitted signal). However if such a signal is not available we might assume that a noisy version of the desired signals could be used instead. The correlation between frequency shifted versions of a cyclostationary signal suggest that we might use a frequency shifted version of the received, noisy signal as a training signal

Consider a purely real BPSK signal; its spectrum is symmetric about f_c, so the positive and negative halves of the spectrum contain the same signal information, whereas any noise present will not, in general, possess such symmetry. A more general complex BPSK signal contains the same diversity, which can be transformed to symmetry by a phase shift. Adding the shifted and unshifted signals therefore gives an increase in SNR.

The blind adaptation of a FSE is first considered, rather than a FRESH filter. A FSE which is not adapting is a time invariant structure; it does not possess the implicit periodic time variation of a FRESH filter. It can therefore only exploit the stationary statistics of the input signal and so its optimum solution can be described by the Wiener filter equations.

Notice that it is not sufficient to have a reference signal which is merely "correlated" with the actual desired signal (the transmitted signal). Nor is it essential for the reference signal to be equal to the transmitted signal. We see below that the optimum performance is achieved if the cross-spectral density of the reference signal and the transmitted signal is the same as the power spectral density of the transmitted signal.

We can see this by considering the frequency domain form of the Wiener filter equations (see, for example, [107]). Let $W(f)$ be the optimum ("Wiener") filter frequency response for extracting the reference signal $d(t)$ from the filter input $x(t)$. $x(t)$ is assumed to contain the transmitted SOI, $s(t)$ with additive noise, so $x(t) = s(t) + n(t)$. We define $u(t) = x(t)e^{j2\pi\alpha t}$ to be a frequency shifted version of $x(t)$, where the frequency shift α, is a cyclic frequency of $s(t)$. The filter output is $y(t)$. We use corresponding letters (e.g. $Y(f)$) to represent the Fourier transforms of these quantities. The Wiener solution [107] is:

$$W(f) = \frac{S_{xd}(f)}{S_{xx}(f)} \tag{7.35}$$

where the cross spectral density of x and d is

$$S_{xd}(f) = E[X(f)D^*(f)] \tag{7.36}$$

and the power spectral density of x is

$$S_{xx}(f) = E[X(f)X^*(f)] \tag{7.37}$$

But if we are using a reference signal which is not exactly the wanted signal then we need to consider the filter output in more detail. The output is given by

$$Y(f) = W(f)X(f) \tag{7.38}$$

so the PSD of the output is

$$E[|Y(f)|^2] = |W(f)|^2 S_{xx}(f) \tag{7.39}$$

$$S_{yy} = \left|\frac{S_{xd}(f)}{S_{xx}(f)}\right|^2 S_{xx}(f) \tag{7.40}$$

$$= \frac{S_{xd}^2(f)}{S_{xx}(f)} \tag{7.41}$$

If $d(t) = s(t)$ then this is the best estimate possible of s from the signal x, with a linear, time-invariant filter. If $d \neq s$ then this is the best estimate using d as a reference signal.

How close this estimate is to s depends of course on the relationship between d and s. The condition on the reference signal is $S_{xd} = S_{xs}$ for the performance to be optimum (that is, equivalent to using s as the reference signal) even if $d \neq s$. It is normal to assume that the noise n is not correlated with either s or d, so then $S_{xs} = S_{ss}$ and $S_{xd} = S_{sd}$.

So the requirement for optimum performance is

$$S_{sd} = S_{ss} \tag{7.42}$$

that is, the cross-spectral density of transmitted and reference signals, is the same as the power spectral density of the transmitted signal. In this case, u is used as the reference signal, so the optimum performance is achieved if $S_{su} = S_{ss}$. But u is x shifted in frequency by α, so $S_{su} = S_{ss'}$ where $s'(t) = s(t)e^{j2\pi\alpha t}$, that is s shifted by the same frequency. This last step relies on the noise being stationary, or, more accurately, not having α as a cyclic frequency.

It is clear from consideration of the spectrum of a purely real signal such as BPSK, that if u is taken as the sum of two shifted versions of x, using shifts of $+2f_c$ and $-2f_c$, then the requirement of equation 7.42 is met. In general, with other cyclic frequencies, such as baud rate related frequencies, this is not the case.

With a BPSK signal we have two parts to the spectrum, in the negative and positive frequency halves of the spectrum, whose cross-spectral density is the same as the original signal PSD if frequency shifts of $\pm 2f_c$ are used. So if the SOI is BPSK a training signal of the correct spectral shape can be generated by adding signals shifted by $+2f_c$ and $-2f_c$. This is the approach described in section 7.4.2.

The discussion above concerns the Wiener filter, but an LMS adaptive filter will adapt towards the Wiener filter solution, so we can maximise the correlation between two signals, with constraints, by using the LMS algorithm.

This is achieved by the structures shown in figures 7.14 and 7.15. The first of these is a filter using the received signal, frequency shifted, as the input to the filter, and the unshifted signal as the desired signal for the filter adaptation. The second swaps the roles of these two signals: the unshifted signal is the filter input, and the shifted signal is the desired signal. The labels x and u refer to the unshifted and shifted signals respectively.

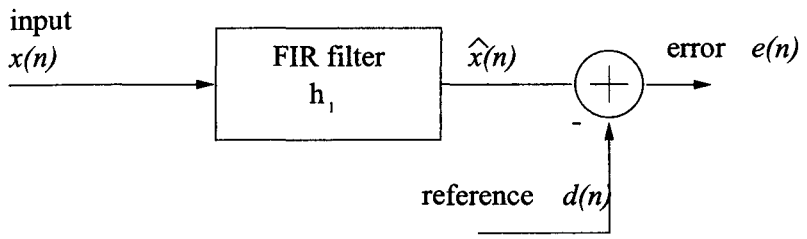

Figure 7.13: Signal identification for Wiener filter

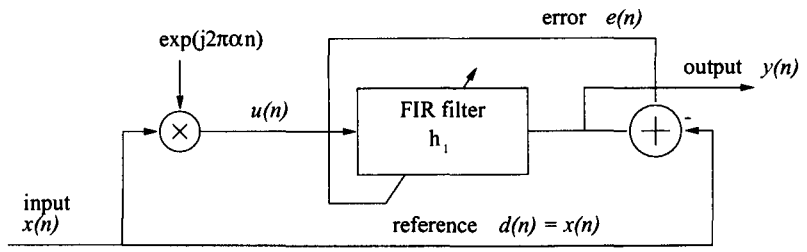

Figure 7.14: Maximising correlation of u relative to x

The output of each of these filters will represent a maximisation of correlation between $x(n)$ and $u(n)$ at all frequencies where the spectra of $x(n)$ and $u(n)$ overlap. The filter in figure 7.14 will give an output which is the components of $u(n)$ which are correlated with $x(n)$. The filter in figure 7.15 will give an output which is the components of $x(n)$ which are correlated with $u(n)$. These two signals are the same, as far as SOI components are concerned, but they will contain different noise signals. The constraints which were missing in [25] are present here because the desired signal is not adaptively filtered, so the filter output spectrum will be constrained by the reference signal spectrum.

7.4.2 BA-FRESH filtering

The technique described above is used in [48] for blind adaptation of a FRESH filter (the BA-FRESH filter). The example used is of two BPSK signals with overlapping spectra; one is treated as a wanted signal, the other as interference. The letter concentrates on the rate of adaptation towards the "optimum" filter solution. Unfortunately this solution is not optimum in the way that the standard FRESH filter would be. In fact the use of the cyclostationarity to train the filter leads, in this case, to a filter equivalent to a fractionally spaced equaliser, or a simple FIR with more than

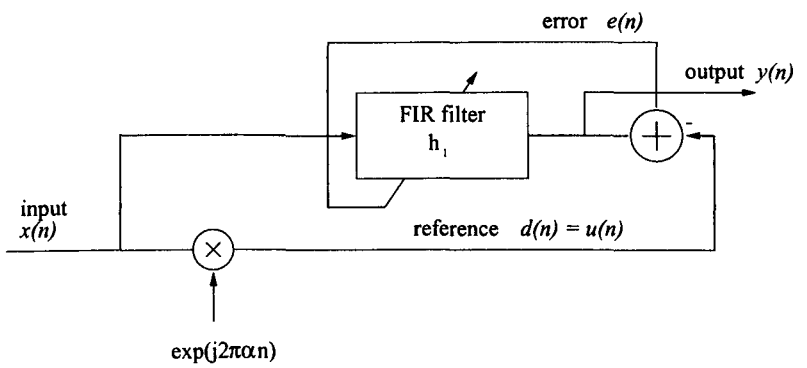

Figure 7.15: Maximising correlation of x relative to u

one sample per symbol.

There are also fairly severe limitations on the applicability of this algorithm - it can be used only for signals in which the frequency shifted components have a cross-spectral density the same as the transmitted signal's power spectral density. Much of [48] is concerned with adaptation rates of various algorithms; this aspect of the letter is not considered here, but the assumptions about the final filter solutions are.

In this section the limitations of the approach proposed in [48] are demonstrated, and in section 7.5 a modification is proposed which results in a method of blind adaptation which still gives some of the interference rejection benefit of FRESH filtering.

A BPSK signal has baud rate related cyclic frequencies of $\frac{\pm n}{T_b}$, where T_b is the symbol duration and n is an integer sufficiently small such that the original and shifted signals still have some part of their spectra overlapping. It also has cyclic frequencies of $\pm 2f_c$ where f_c is the carrier frequency. These carrier dependent cyclic frequencies are a direct consequence of the perfect symmetry of a BPSK signal spectrum about the carrier frequency. In this section, as in [48], we only consider this carrier related cyclostationarity.

7.4.3 Passband filtering

A description follows of the specific examples of the trained and blind FRESH filters Zhang [48] used in simulations. More general comments are made later. In [48] the filters shown in figures 7.16 and 7.17 are compared. The first is a trained FRESH filter as proposed by Gardner [13]; the second is Zhang's proposed blind adaptive FRESH filter (BA-FRESH) structure

There are a number of comments to be made on these systems The top two branches of both filters (shifted by $+2f_c$ and $-2f_c$) are shifted so they contain SOI components overlapping with the unshifted branch, but they do not overlap with each other.

The advantage of the FRESH filter comes from adding together SOI components at different frequencies, while the noise at the same frequencies is uncorrelated.

The $\pm 2f_c$ frequency shifts cause correlated SOI components to be at the same frequencies because of the symmetry of the BPSK spectrum: referring to figure 7.18.1, there is the symmetry of the positive frequency half of the spectrum under reflection about f_c (and of the negative frequency half about $-f_c$). This symmetry is due to the baseband equivalent BPSK signal being purely real. There is also symmetry of the whole spectrum under reflection about zero frequency, due to the fact that the modulated signal is purely real (see, for example, [108], chapter 3. A $+2f_c$ shift causes the negative frequency part labelled A to coincide with the positive frequency part labelled B in the unshifted signal. Similarly the negative frequency part B overlays the unshifted positive part A.

A $-2f_c$ shift causes A and B of the positive part to overlay B and A of the negative frequency part. Under both these shifts the noise is uncorrelated because it will not have, in general, symmetry around f_c (or $-f_c$). But although there is correlation between x and $xe^{-j4\pi f_c t}$, and there is correlation between x and $xe^{+j4\pi f_c t}$, there is no correlation between $xe^{-j4\pi f_c t}$ and $xe^{+j4\pi f_c t}$ because these signals have no parts of their spectra overlapping.

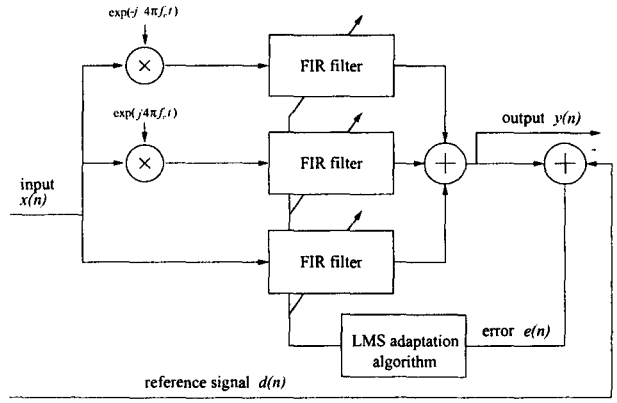

Figure 7.16: FRESH filter exploiting carrier related cyclostationarity (passband)

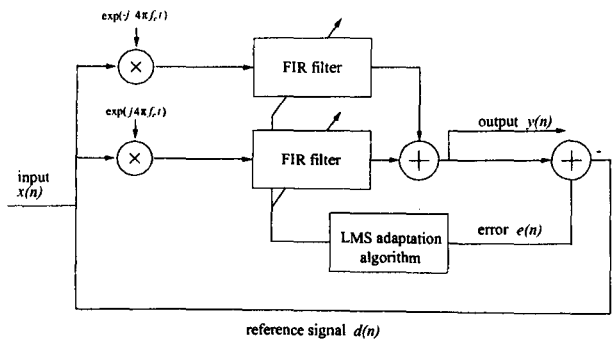

Figure 7.17: BA-FRESH filter exploiting carrier related cyclostationarity (passband)

However, due to the symmetry about zero frequency, the two frequency shifts produce correlation functions which are complex conjugates of each other. Let $u^{+2f_c} = xe^{+j4\pi f_c t}$ and $u^{-2f_c} = xe^{-j4\pi f_c t}$ and let R^α be the cyclic cross correlation of x under a frequency shift of α. The cross correlation of x and u^{+2f_c} is

$$\begin{aligned} R^\alpha &= E[x.u^{+2f_c}] \\ &= E[(x^*.u^{-2f_c})^*] \quad \text{because} \quad u^{-2f_c} = \left(u^{2f_c}\right)^* \\ &= E[(x.u^{-2f_c})^*] \quad \text{because} \quad x = x^* \\ &= R^{-\alpha} \end{aligned} \tag{7.43}$$

It is clear from section 7.4.1 that using both the frequency shifts is necessary to make $S_{sd} = S_{ss}$, but as the signals in the two branches carry the same information (one is the complex conjugate of the other) then the structure is clearly inefficient.

In fact figure 7.17, can be thought of as a blind adaptive FSE, or a trained FSE which is using a noisy training signal, but in either case, the positive and negative frequency halves of the filtered spectrum are swapped by the frequency shifting. The symmetry of each part around $2f_c$ and $-2f_c$ means that this swap does not impair performance.

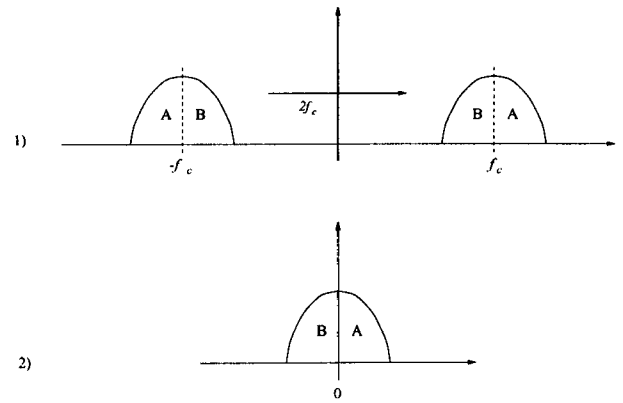

Figure 7.18: Spectrum of BPSK signal 1) passband 2) complex baseband

7.4.4 Complex baseband filtering

This discussion becomes simpler if it is based on the equivalent complex baseband signals instead of the original passband ones. The equivalent situation at the complex baseband is as shown in figure 7.18.2. The SOI spectrum is symmetric about zero frequency, but the noise is not, in general. To overlay the correlated parts of the signal one takes the complex conjugate, which is equivalent to reflection about the zero frequency axis.

The trained FRESH filter at baseband, analogous to the passband filter in figure 7.16, is therefore the structure shown in figure 7.19, while the baseband version of Zhang's filter (figure 7.17) is as shown in figure 7.20. In this structure we are using the positive frequency half of the spectrum of the signal to train the filter to pass the negative frequency half, and simultaneously using the negative half to train the filter to pass the positive half. As the SOI is symmetric about zero frequency, this is equivalent to a trained FSE, except that the training signal contains noise.

It is interesting to note that the filter structure in figure 7.19 was proposed as early as 1969 by Brown [27] and the underlying theory further discussed in [100]. [27] contains a thorough theoretical justification for the use of "conjugate linear filtering", but an intuitive explanation for its use is as follows: A signal such as BPSK can be represented at baseband (following, in general, a phase shift) as purely real. A realistic baseband representation of noise will be complex, so in filtering to recover the signal from the noise an obvious (non-linear) improvement is to discard the imaginary part of the noise. The sum of a linearly filtered and a conjugate linearly filtered version of the signal allows the removal of the imaginary part of the noise. At passband, this corresponds to perfect coherent reception, where the in phase component is selected and the quadrature component discarded.

It is shown in [27] that conjugate linear filtering is advantageous when $E[x(t+\tau)x(t)] \neq 0$ (the left hand side differs from autocorrelation in a missing conjugate operation on the second $x(t)$). Such signals were described in [87] as "circular" (see section 3.8). The maximum benefit is when the real and imaginary parts of $x(t)$ are linearly dependent, which, for most practical signal constellations, is another way of saying that the signal can be transformed to a purely real one by a phase shift.

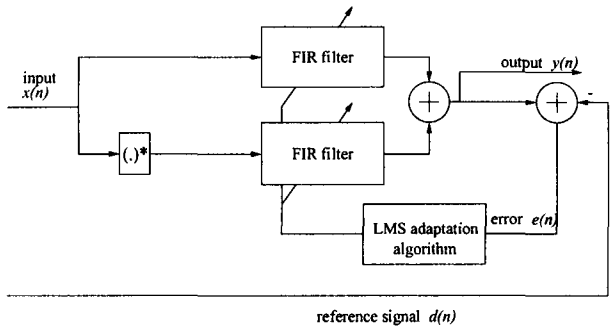

Figure 7.19: FRESH filter exploiting carrier related cyclostationarity (complex baseband)

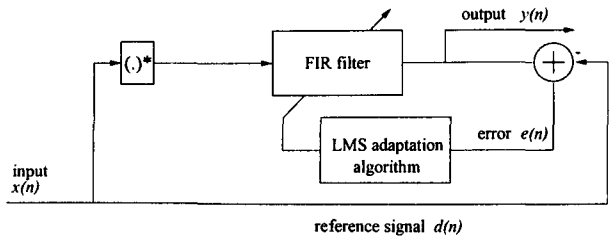

Figure 7.20: BA-FRESH filter exploiting carrier related cyclostationarity (complex baseband)

7.4.5 Simulation results

The trained FRESH filter (as shown in figure 7.19), the BA-FRESH filter (figure 7.20), and trained FSE were compared by simulation using scenarios similar to those used in [48]. That is, the SOI and interferer were of equal power and symbol rate. AWGN was added with power to give an SOI E_b/N_0 of 9 dB and 18 dB. In each case, the filters were trained using the LMS algorithm; adaptation times and step sizes were chosen to minimise the effect of tap misadjustment on the results. In realistic scenarios we can expect that the interferer and SOI have incoherent carrier phases; this was simulated by averaging MSE results over a range of relative phases, although in practice this only makes a difference at carrier frequency differences of 0 and the baud rate.

We expect that, in view of the discussion above, that the BA-FRESH filter output MSE should be exactly the same as that of the trained FSE, and that the trained FRESH filter should be significantly better. This is confirmed by the simulation results shown in figures 7.21 and 7.22: the FSE and BA-FRESH lines are almost indistinguishable, while the trained FRESH clearly has an improved MSE.

The departures from the smooth curves at carrier frequency differences of 0 and 1 Hz are addressed in section 7.5.1.

7.4.6 Conclusion

It is clear from the above that Zhang's filter (figures 7.17 and 7.20) cannot give an improvement in interference rejection over a FSE. At first it may appear from figure 7.17 that only one of three possible filter branches is being used for training so the remaining two branches can provide some FRESH filtering advantage. This is not the case, however, because the two remaining branches have no overlapping SOI components. It is immediately obvious from the equivalent baseband

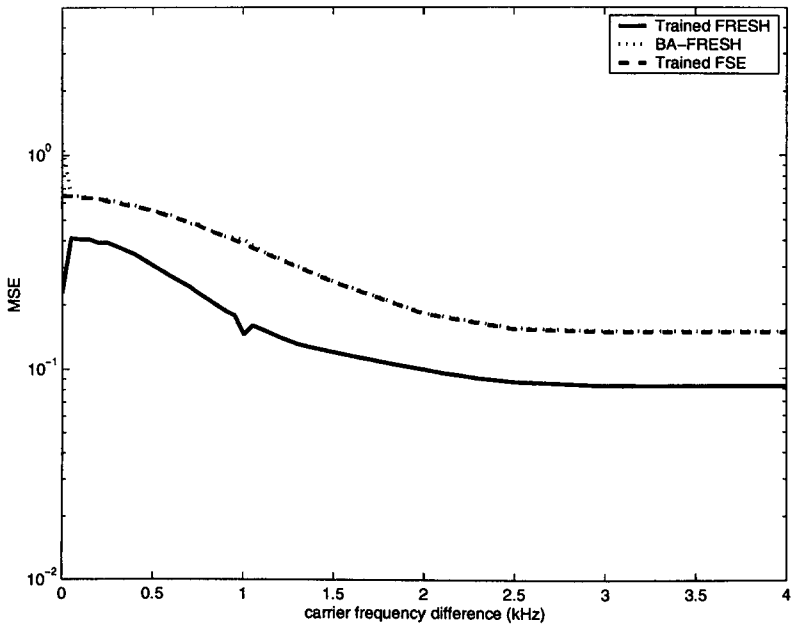

Figure 7.21: BA-FRESH; variation of MSE with carrier frequency

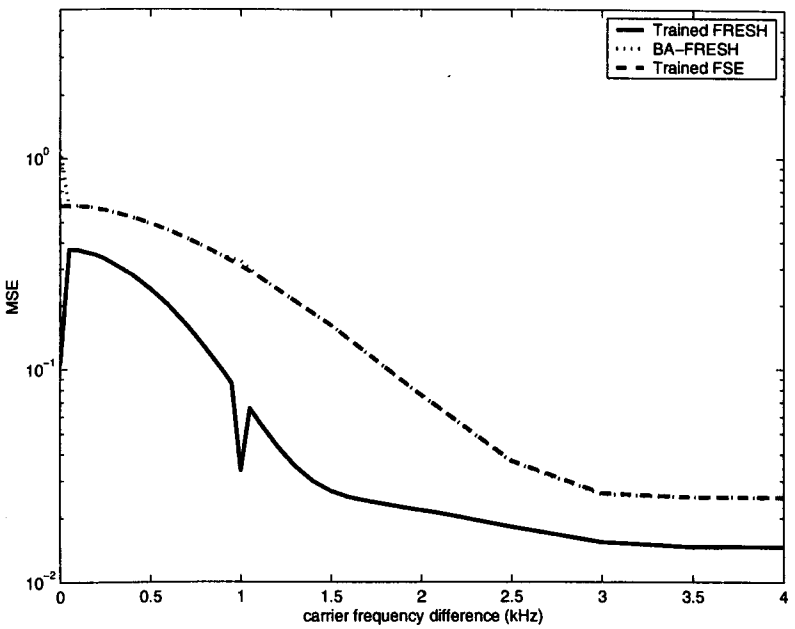

Figure 7.22: BA-FRESH: variation of MSE with carrier frequency

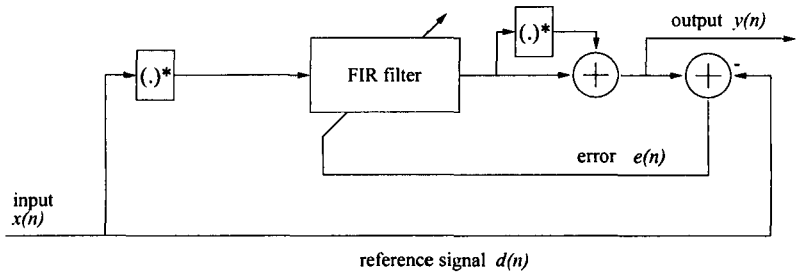

Figure 7.23: Improved BA-FRESH filter exploiting carrier related cyclostationarity (complex baseband)

version (figure 7.20) that no addition of correlated components is possible with such a structure. The filter is therefore equivalent to a blind adaptive FSE.

7.5 Improved Blind Adaptive FRESH Filtering

From the previous section it is apparent that a passband FRESH filter with carrier related frequency shifts is equivalent to a conjugate linear filter at complex baseband. The advantage of such a filter in extracting a purely real signal from complex noise is equivalent to that from discarding the imaginary part of the noise. So it is simple to construct a blind adaptive filter which has the same performance as a conjugate linear filter: such a filter is shown in figure 7.23. The positive and negative halves of the symmetric SOI spectrum are used to train each other as before, but the addition of the output to its complex conjugate removes the imaginary part of the noise without affecting the SOI.

Again simulation was used to test this filter structure in comparison to Zhang's BA-FRESH filter and a trained FRESH filter. the results are shown in figures 7.24 and 7.25. It is clear that the new filter structure has output MSE performance almost exactly equivalent to the trained FRESH filter. The obvious differences at carrier frequency separations of zero and 1 Hz (twice the baud rate) are discussed in section 7.5.1.

7.5.1 Effect of relative phase of interferer and SOI carriers

Figures 7.21, 7.22, 7.24 and 7.25 show that as well as the general trend that MSE decreases as carrier separation increases (as expected) and the differences between the four filter types discussed above, there are some interesting features apparent at carrier separations of 0 and 1 kHz. Some suggestions are made here for the causes of these curious results.

The fundamental issue involved here is the relative phase of the SOI carrier and the interferer carrier. In a realistic communications scenario one would expect that an interferer would be coming from a source independent from the SOI source and would therefore be incoherent with the SOI. That is, there would be no fixed phase difference between the two carriers. In simulation, it is necessary to explicitly force such incoherence onto the signals, and this was done for the results presented in figures 7.21, 7.22, 7.24 and 7.25 by repeating each simulation for a different initial phase difference, and averaging the resulting MSE values. Initial phases of 0 to 90° in steps of 10° were used.

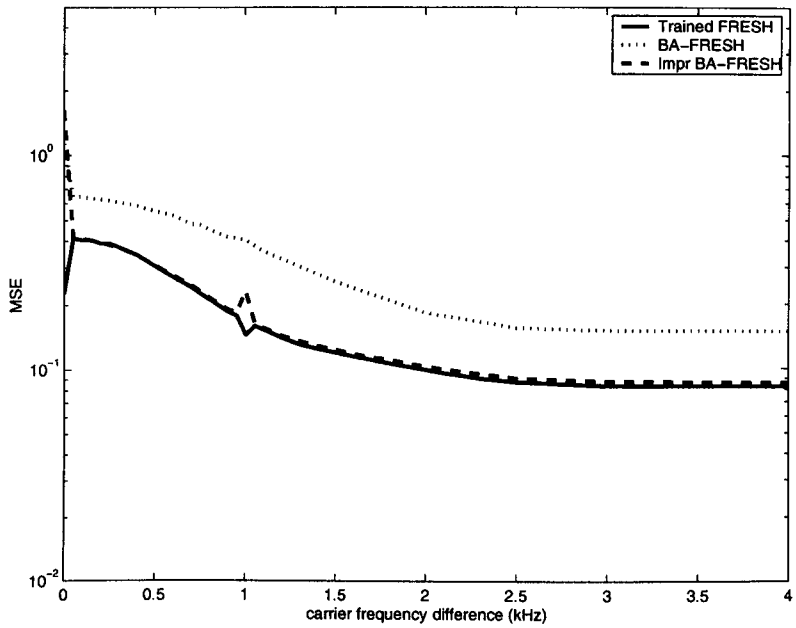

Figure 7.24: Improved BA-FRESH: variation of MSE with carrier frequency, E_b/N_0 = 9 dB

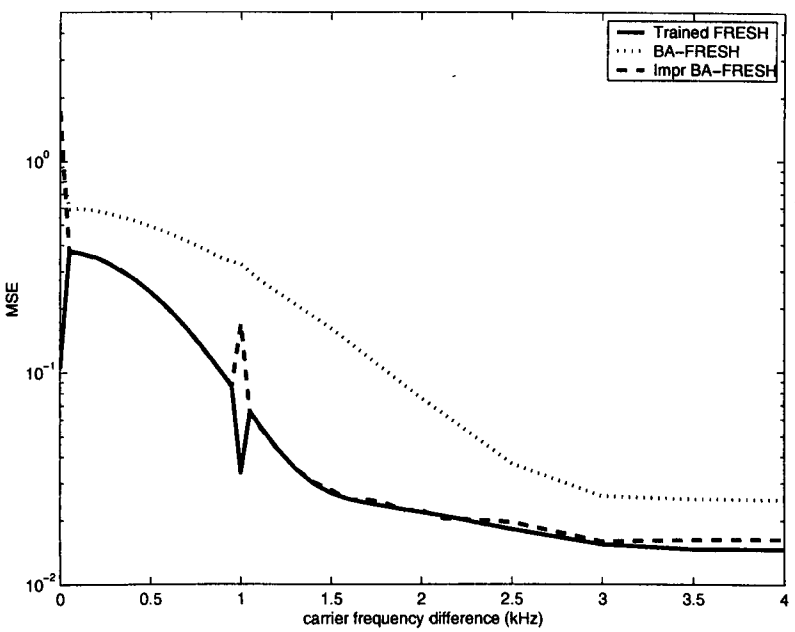

Figure 7.25: Improved BA-FRESH: variation of MSE with carrier frequency, E_b/N_0 = 18 dB

When the carriers have the same frequency the initial phase difference remains the phase difference throughout the simulation. In this case there is a major difference between an interferer with identical frequency and phase and one which is 90° out of phase with the SOI. The former is difficult to separate from the SOI (especially in the scenarios tested here where the two signals have identical power and both carry IID data). The latter, however, can be perfectly separated from the SOI by taking the real or imaginary part of the signal as required.

When the carriers are separated by half the baud rate (as is the case at 1 kHz separation in the simulations) then in one symbol period, the phase difference caused by the frequency separation is 180°. This maps one BPSK constellation point onto the other available point, so the interfering signal appears similar to a BPSK signal with a fixed relative phase difference to the SOI, and similar effects in the MSE are observed.

The FSE is a linear filter and cannot therefore operate any differently on the real or imaginary parts of the signal. The effects described above are therefore irrelevant and the filter performance has no dependence on the relative phase of the two signals. This is shown in figure 7.26 for two sample initial phase differences of 0° and 90°. There is no difference in the output MSE for these two cases.

The trained FRESH filter (figure 7.27) has much better performance when the phase difference remains constant (zero, and half baud rate carrier frequency difference) and that phase difference is 90°, as the signals are perfectly separable, and the FRESH filter is capable of operating independently on the real and imaginary inputs. In both figures 7.26 and 7.27 the E_b/N_0 is 18 dB.

The BA-FRESH structure is equivalent to a blindly adapted FSE, so it cannot perform any better than the FSE. However the reference signal used to train the filter contains the interfering signal. Normally the filter training relies on the lack of correlation between any interference components in the reference signal and the filter input. However with equal carrier frequencies and 90° phase difference, the interference in the input and reference signals are perfectly correlated (except for the effects of the AWGN) so the performance is worse than that of the FSE (as shown in figure 7.28).

The improved BA-FRESH (figure 7.29) filter has similar performance to the trained FRESH except that at equal carrier frequencies, and a 90° phase difference, the performance is significantly worse. This can again be attributed to the fact that the interferer in the reference signal is correlated with the interferer in the input to the filter, so the filter will be adapted to enhance the interferer as much as possible.

Figures 7.30, 7.31 and 7.32 show the variation of MSE with initial carrier phase difference for frequencies of 0, 0.5 and 1 kHz and E_B/N_0 of 18 dB. These results justify the technique of averaging MSE over 10° steps in phase value, as the variation is smooth. One exception to that is the jump from 0 to 10° in the improved BA-FRESH curve in figure 7.30. This suggests that the 0 carrier difference values for this filter in figures 7.24 and 7.25 are too low, but this will be a small effect.

In general we can see that the FSE is independent of the initial phase while the trained FRESH filter has a fairly strong phase dependence. The BA-FRESH filter has a weaker phase dependence, and the improved BA-FRESH phase dependence is complicated and unexplained. Figure 7.31 illustrates the lack of phase dependence when the carrier difference is any value other than zero or

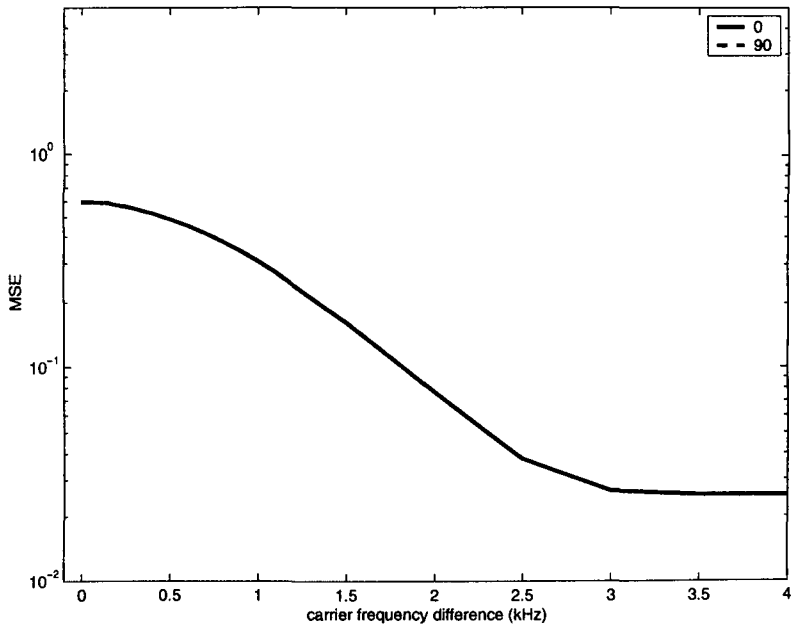

Figure 7.26: FSE - effect of carrier phase and frequency

half the baud rate.

7.5.2 Other frequency shifts

In [48] the theory for the BA-FRESH algorithm is presented in a general way, and the BPSK example considered above is used as an example. We have seen that in this is example we have no advantage in using BA-FRESH over a FSE, but the question remains as to whether the general technique is still valid. The answer is that if additional cyclic frequencies are exploited then it should be possible to achieve blind adaptation with simultaneous FRESH filtering advantage over interference, but further filtering will be required to correct the spectrum of the output signal.

Blind adaptation using the positive frequency half of the spectrum to train the negative frequency half is successful (in that it gives the correct output signal) because correlated halves of the spectrum have the same shape under the conjugate transformation. If, however, baud rate related frequency shifts are used, the situation is not so simple. Consider the simple case of figure 7.20 where the complex conjugate operation is replaced with a frequency shift equal to the baud rate. The output of the filter will be the filtered version of the input signal that has maximum correlation with the training signal. This will have a power spectral density equal to the cross spectral density of the unshifted training signal and the input shifted signal, which in general is not equal to the SOI power spectral density.

Note that frequency shifts equal to the cyclic frequencies of the interferer (if they are known) can also be used successfully [13, 2].

7.5.3 Conclusion

The BA-FRESH technique as presented in [48] is of limited use as it has none of the interference rejection advantages of FRESH filters in the examples given in [48]. It is in those cases equivalent

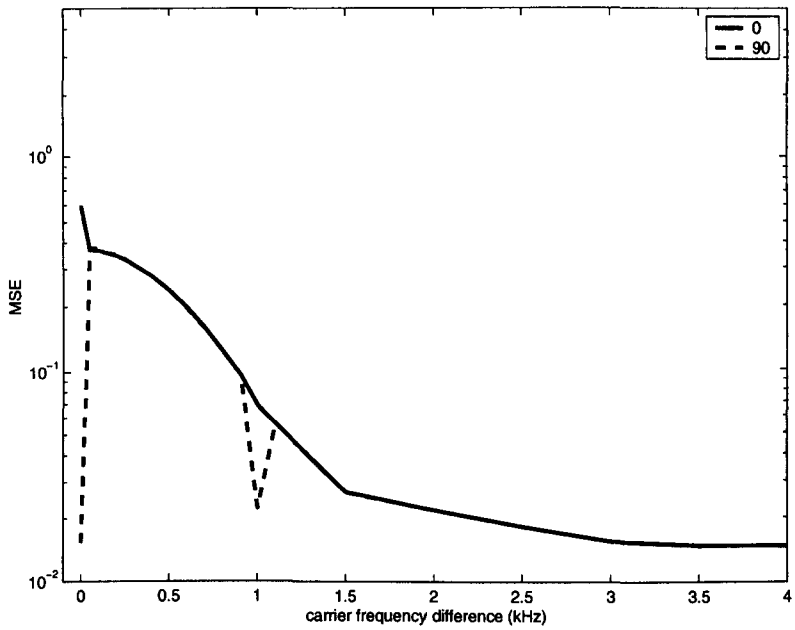

Figure 7.27: Trained FRESH - effect of carrier phase and frequency

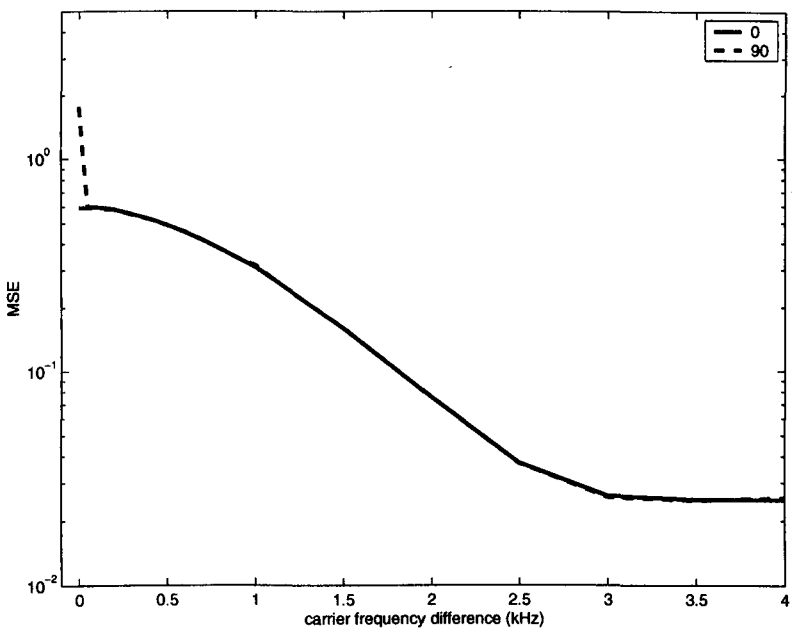

Figure 7.28: BA-FRESH - effect of carrier phase and frequency

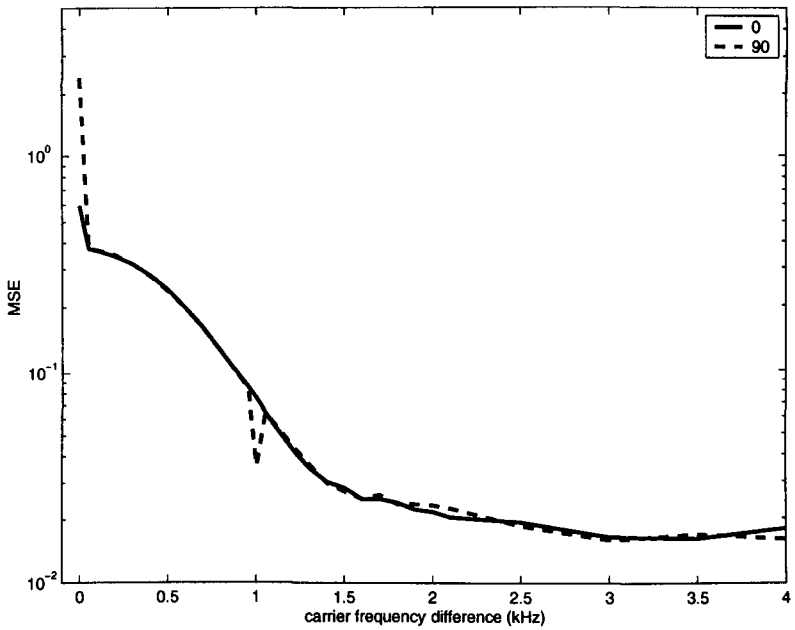

Figure 7.29: Improved BA-FRESH - effect of carrier phase and frequency

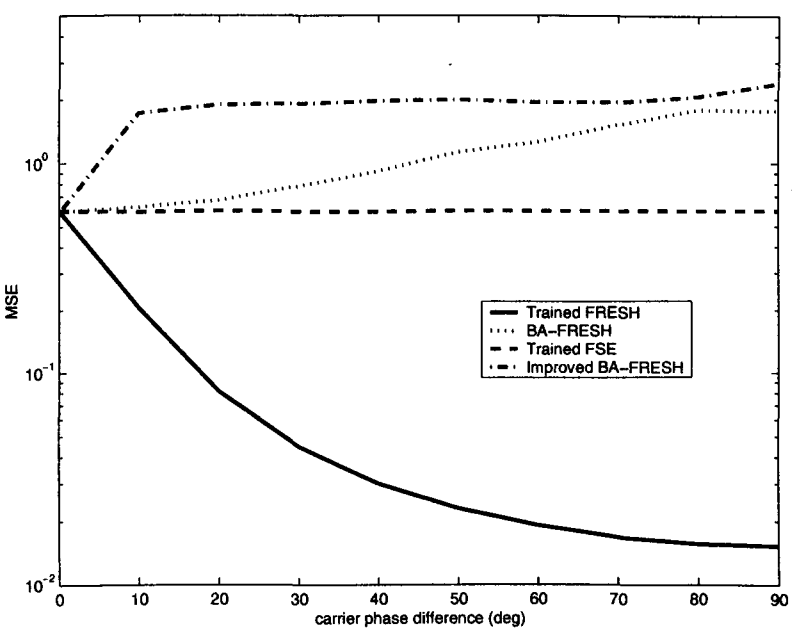

Figure 7.30: Variation of MSE with carrier phase, no carrier separation

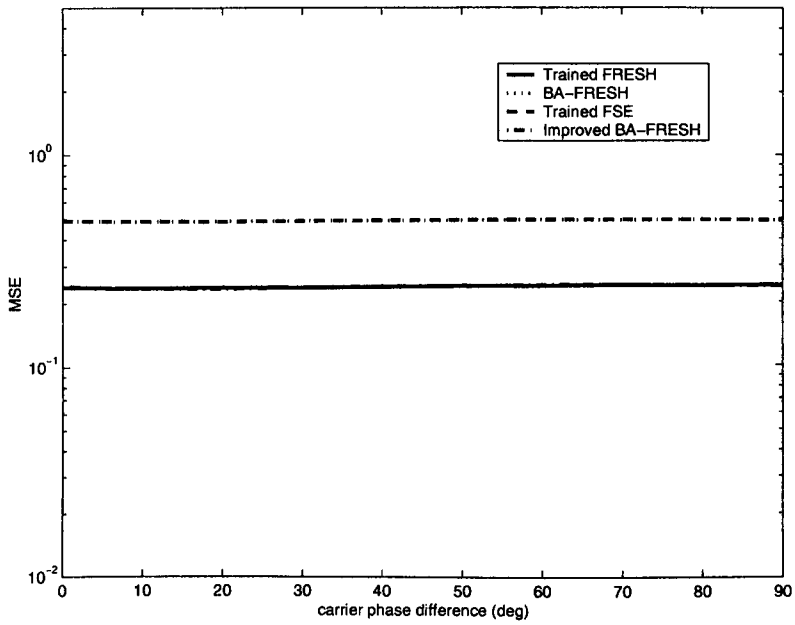

Figure 7.31: Variation of MSE with carrier phase, carrier separation 0.5 kHz

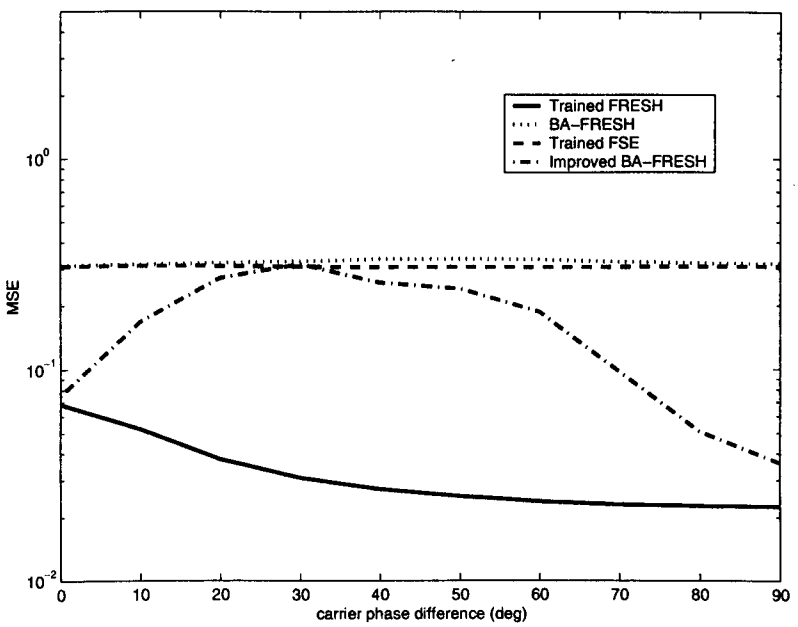

Figure 7.32: Variation of MSE with carrier phase, carrier separation 1 kHz

to a method for blindly adapting a FSE.

The same performance as a trained FRESH filter can be achieved however by taking the real part of the filter output. This is an illustration of the fact that exploiting the carrier related correlation of BPSK, which is the same as exploiting the symmetry of the BPSK spectrum, is equivalent to discarding the imaginary part of the complex noise.

7.6 Summary

The chapter has examined the problem of blind adaptation of FRESH filters, particularly in a channel with a white frequency response. For general cyclostationary signals, the mathematics describing the problem was presented in section 7.1 but did not give a practical algorithm for blind adaptation. A simple method for doing this using a dummy training signal was presented in section 7.2. This algorithm is proposed for use in systems where the channel can be monitored without the SOI present, such as those using some form of time division multiple access (TDMA). Simulation results confirm the effectiveness of this algorithm.

Another author has proposed two algorithms for blind adaptation of FRESH filters. One minimises the correlation between a shifted and unshifted version of the SOI. Section 7.3 shows that this method does not work. The second algorithm uses a frequency shifted version of the received signal as a training signal for the filter. In section 7.4 it was shown that this only has performance equivalent to a FSE and that it is limited to signals such as BPSK with symmetric spectra.

A simple change to this algorithm (which is to discard the imaginary part of the filter's output) improves the performance to the level of a trained FRESH filter. This is described in section 7.5. Put another way, sections 7.4 and 7.5 have demonstrated that FRESH filtering of BPSK using $\pm 2f_c$ frequency shifts only is equivalent to linear, conjugate-linear (LCL) filtering at baseband, which is equivalent to discarding the imaginary part of the output of a FSE.

It has been shown that two algorithms proposed in the literature for blindly adapting FRESH filters against interference do not perform successfully as claimed by the original authors.

The reasons for the failure of these algorithm have been explained. For the second algorithm the performance has been shown to be equivalent to that of a FSE, but a change to the filter structure has been proposed which gives performance as good as a trained FRESH filter. The fact that these filter structures are limited to exploiting the spectral symmetry of purely real signals has also been highlighted.

Chapter 8

Conclusion

In this thesis the use of FRESH filters for the excision of interference has been considered. A variety of interference scenarios have been tested, each with one digitally modulated signal interfering with a digitally modulated SOI. The emphasis throughout the thesis has been on using realistic scenarios, so band limited signals, such as raised cosine filtered QPSK or GMSK have been used in simulations to test the FRESH filters. The performance of FRESH filters has been shown to vary according to the carrier separation between SOI and interferer, their relative baud rates, their relative powers and their pulse shapes.

Specific points have been made, or results shown, in individual chapters as listed below.

Chapter 1 contains a literature review which covers the major categories of techniques for interference rejection to put FRESH filtering into context. It is shown that although there has been fairly thorough coverage of the theoretical background to FRESH filtering, there has been little work published on applying the techniques to practical situations. The impression gained from reading the existing publications on the subject is that FRESH filtering appears to be a very promising technique, but that it is not clear when or if it can be used successfully in practice. The work presented in the rest of the thesis has gone some way to suggest that in selected applications (particularly where there is reasonable predictable high powered interference) it can be a useful technique, but that the exploitation of spectral correlation is often accomplished by standard techniques such as matched filtering.

Chapter 2 reviews the theory of the cyclostationary model of digitally modulated signals, and the resulting spectral correlation. Also a comparison is made between Gardner's non-probabilistic approach to spectral analysis and the more accepted methods based on ensemble averaging. The previously published work in this subject tends to be presented in rather an abstract way, so this chapter concentrates on presenting the ideas as simply as possible, with graphical interpretations of the two-parameter autocorrelation function and of baud rate and carrier frequency related spectral correlation.

The exploitation of the properties described in chapter 2 is described in chapter 3. The FRESH filter is introduced, and its role in a communications receiver described with reference to the matched filter, and linear equalisers. The FRESH filter can be used as a replacement for the matched filter/equaliser combination, which may have improved performance against interference, or may be used as a stand-alone interference rejection filter placed before the matched filter. Following

on from the discussion of the relationship between baud rate related correlation and the sampling theorem, it is shown that there is a close relationship between the frequency shifting and sampling within the FRESH filter, and this can lead to improvements in the efficiency of implementing the filter. The important equivalence of the FSE and baud rate shift FRESH filter if both are followed by baud rate sampling is highlighted here (it is first shown in [59]), and this leads on to a demonstration of how the FSE exploits spectral correlation in such a situation. At this point it is clear that systems which include baud rate sampling (such as the matched filter followed by sampling, or the DFE) are implicitly using the baud rate related spectral correlation of the SOI. A distinction is drawn here between Wiener filtering which extracts the entire waveform from noise, and matched filtering in which baud rate samples are detected in the optimum way. The Wiener filter assumes that all signals present are stationary, so a more accurate cyclostationary model leads to the improved Wiener filter. However, the matched filter assumes nothing about the statistics of the SOI; it is defined according to the noise spectrum and the SOI pulse shape, so the SOI cyclostationary model does not lead to any improvement in the MF performance.

But using the spectral correlation of the interferer does lead to an improvement over the MF which only operates in the context of baud rate sampling. This is demonstrated through the simulation results in chapter 4. Here the implementation of the FRESH filter is described for removing a QPSK SOI from QPSK interference. The filter responses resulting from training the filters with the LMS adaptive algorithm are compared to the results of solving Gardner's cyclic Wiener filtering equations and found to be in good agreement. It is also shown how the choice of the number of filter taps can affect performance: a large number increases the degrading effect of tap misadjustment in an adaptive implementation, whereas too small a number does not give the filter sufficient frequency resolution to accurately manipulate the signals. The point is made here that in the presence of high powered interference it is important to exploit the interferer correlation, not just that of the SOI.

At first sight it might appear that as cyclostationarity in the SOI is beneficial for rejecting interference, then enhancing this cyclostationarity in some way should lead to an improvement in interference immunity if FRESH filtering is used. It is shown in chapter 5 that this is not the case in general. Given a fixed transmitter power, one can distribute the SOI energy over different bandwidths with different spectra, but spreading the energy further means overall a lower spectral density, which then means lower immunity at specific frequencies. So without having any knowledge of what a likely interference scenario is, it is not possible to design a particular signal which is robust - there is no signal which is robust in general. This chapter also examines more QPSK interference scenarios, with equal SOI and interferer power.

The application of FRESH filters to VLF communication is explored in chapter 6. This is a real application with the parameters chosen to represent real world problems as closely as possible. The problem addressed here is that of extremely high powered adjacent channel interference, which can be so powerful that the normally insignificant side lobes of the (G)MSK adjacent channel signal are of similar power to the main lobe of the SOI. By using a FRESH filter with a single frequency shift matched to the spectral correlation of the interferer, a large improvement in BER performance is possible. This was first modelled in AWGN, but then a more realistic impulsive atmospheric noise model was developed and used. Impulsive noise is the most serious problem in VLF communications, so any interference rejection technique must be able to operate alongside

an impulse rejection system. A simple but effective method of removing impulses known as hole-punching is used before the FRESH filter, and it is seen that depending on the relative powers of the SOI, interferer and impulsive noise, it can be better to use either the FRESH filter, or the hole-punching, or both, but that the hole-punching does degrade the performance of the FRESH filter.

The last chapter 7 looks at some methods of adapting the FRESH filter without a training signal (blind adaptation). There are many blind adaptation methods available, usually concerned with the more difficult problem of adapting a filter to equalise a dispersive channel response. Here the focus is on two methods which directly use the spectral correlation of the input signals to drive the adaptation. It is shown that one is faulty, in that the FRESH filter does not adapt to the desired solution, and the other is limited in its application to purely real signals, and results in a filter structure which is equivalent to a FSE. A modification to this filter is proposed which makes its performance equivalent to the trained FRESH filter. It is clear from this work that the FRESH filter which exploits only the carrier related correlation of a purely real signal such as BPSK, is equivalent at complex baseband to a linear-conjugate-linear filter, which in turn is equivalent to linear filtering followed by discarding of the imaginary part of the noise.

The last comment shows that one of the problems with the existing literature on FRESH filtering is that it appears to be a complicated topic, so the simple interpretations of what the filters are doing are often overlooked. It is hoped that this work has shown that the idea of spectral correlation or cyclostationarity, and its exploitation through the FRESH filter, are in fact rather simple.

It is believed that the current work has shown that, with some reservations, FRESH filtering can be used for solving real interference problems. As standard techniques such as the MF or the DFE implicitly exploit the baud rate correlation of a SOI, then there is no advantage in using a FRESH filter for this. Similarly, the carrier related correlation of BPSK could equally be exploited by coherent reception and taking the real part of the output. This can either be seen as an advantage of BPSK, or as an example of its spectral inefficiency which creates the redundancy. The more promising areas for FRESH filter application are those where, as in chapter 6, interferers may have high power, and there are cyclic frequencies of the interferer which do not correspond to baud rate multiples of the SOI.

As well as making these points, the fundamental concepts underlying FRESH filtering have been clarified, and various issues related to their implementation have been explored, including a more comprehensive examination of their performance under differing interference scenarios, with the selected modulation schemes, than is available in the literature to date.

8.1 Original Work

A number of original pieces of work have been presented in this thesis, and are indicated in chapter 1. They are summarised here for convenience:

- the graphical representation of the two parameter autocorrelation function for a rectangular pulse sequence (section 2.1.5) and the interpretation of baud rate related spectral correlation as a result of the sampling theorem (section 2.2.2);

- the graphical representations of how the FRESH filter and FSE reject interference (sections 3.5.1 and 3.6.7);
- the analysis of the effect of sampling in FRESH filters (sections 3.6.3 to 3.6.6);
- the explanation of how the LCL filter is the baseband representation of the FRESH filter with a frequency shift of twice the carrier;
- the link between circularity and carrier related spectral correlation;
- the comparison between the direct solution and the LMS adaptive solution of the FRESH filter equations (chapter 4);
- the use of FRESH filters for the specific interference scenarios tested in chapters 4, 5 and 6;
- the attempt to improve FRESH filter performance by manipulating the SOI pulse shape, while keeping the SOI bandwidth the same;
- the investigation of FRESH filter performance with impulsive noise and hole-punching (chapter 6);
- the analysis and criticism of the two algorithms for blind adaptation of FRESH filters [25, 48] (sections 7.3 and 7.4);
- the interpretation of the gain due to carrier related correlation with real signals as equivalent to selection of the real part of the filter output, and the corresponding improvement to the blind adaptive FRESH filter (section 7.5).

8.2 Further Work

Before using FRESH filters in a communications system there is obviously a great deal of investigation which needs to be done. Much of this would be scenario dependent, but there are some general issues which have not been considered here.

Synchronisation, be it symbol timing recovery, or carrier phase and frequency is a crucial process in any communications receiver. Here, perfect synchronisation has been assumed in all analysis and simulations. FRESH filters are bound to have the advantages of FSEs in correcting symbol timing errors [34], but a drift in symbol timing could cause problems, as it can be seen from chapter 7 that the filter performance (and of course ideal response) is dependent on the timing phase of the signals.

Accurate carrier phase recovery is also important in coherent receivers, although one could turn round the comment made above about discarding the imaginary part of the noise when receiving BPSK, and say that LCL filtering may provide an alternative way of (implicitly) recovering the carrier phase with a purely real signal.

Generally it would be expected that any jitter in the parameters which affect the cyclic frequencies of a signal (in the current examples, carrier frequency and baud rate, which could be caused by hardware imperfection or channel conditions such as Doppler shift) would affect the performance

of a FRESH filter. The sensitivity of the filter to errors in the frequency shifts relative to the cyclic frequencies of the input signal has not been analysed.

Also, from the most general point of view, the work on the co-existence of FRESH filtering with hole-punching needs to be extended to the co-existence with all other processes operating in the receiver. For example, VLF communications also routinely use a 2 element antenna array to remove interference by spatial discrimination [109], but this can only cope with one interfering signal. If a second interferer were present, then questions arise as to how to determine in which order to do the FRESH filtering and antenna null steering, and how to select which signal is most appropriate to be tackled by each method. Similarly, there is a range of alternative techniques to hole-punching which may be less disruptive to FRESH filtering, and work is required to determine which they are, and what, if any, effect they have on FRESH filtering, or what effect FRESH filtering has on them.

Bibliography

[1] J. F. Adlard, T. C. Tozer, and A. G. Burr, "Co-channel interference rejection in the absence of a signal of interest", in *Proceedings of the IASTED International Conference on Signal Processing and Communications*, Canary Islands, Spain, February 1998, pp. 195–198.

[2] J. F. Adlard, T. C. Tozer, and A. G. Burr, "Application of frequency-shift filtering to the removal of adjacent channel interference in VLF communications", in *Proceedings of the IEEE Global Telecommunications Conference (GLOBECOM)*, Sydney, November 1998, pp. 3515–3520.

[3] J. F. Adlard, T. C. Tozer, and A. G. Burr, "Frequency-shift filtering in VLF communications for interference rejection in impulsive noise", in *Proceedings of the 6th IEEE International Workshop on Intelligent Signal Processing and Communication Systems*, Melbourne, November 1998, vol. 2, pp. 722–726.

[4] J. F. Adlard, T. C. Tozer, and A. G. Burr, "Interference rejection in impulsive noise for VLF communications", in *Proceedings of the IEEE Military Communications Conference (MILCOM)*, Atlantic City, November 1999.

[5] "The Shorter Oxford English Dictionary", 1983.

[6] H. Cravis and T. V. Crater, "Engineering of T1 carrier system repeatered lines", *Bell System Technical Journal*, vol. 42, no. 2, pp. 431–486, March 1963.

[7] J. W. Cook, R. H. Kirkby, M. G. Booth, K. T. Foster, D. E. A. Clarke, and G. Young, "The noise and crosstalk environment for ADSL and VDSL systems", *IEEE Communications Magazine*, vol. 37, no. 5, pp. 73–78, May 1999.

[8] William A. Gardner and L. E. Franks, "Characterization of cyclostationary random signal processes", *IEEE Transactions on Information Theory*, vol. IT-21, no. 1, pp. 4–14, January 1975.

[9] William A. Gardner, "The spectral correlation theory of cyclostationary time-series", *Signal Processing*, vol. 11, pp. 13–36, 1986.

[10] William A. Gardner, "The role of spectral correlation in design and performance analysis of synchronizers", *IEEE Transactions on Communications*, vol. COM-34, no. 11, pp. 1089–1095, November 1986.

[11] J. H. Reed and T. C. Hsia, "A technique for sorting and detecting signals and interference", in *Proceedings of the IEEE Military Communications Conference (MILCOM)*, 1988, pp. 425–430.

[12] B. G. Agee, S. V. Schell, and William. A. Gardner, "Spectral self-coherence restoral: A new approach to blind adaptive signal extraction using antenna arrays", *Proceedings of the IEEE*, vol. 78, no. 4, pp. 753–767, April 1990.

[13] William A. Gardner, "Cyclic Wiener filtering: Theory and method", *IEEE Transactions on Communications*, vol. 41, no. 1, pp. 151–163, January 1993.

[14] J. D. Laster and J. H. Reed, "Interference rejection in digital wireless communications", *IEEE Signal Processing Magazine*, vol. 14, no. 3, pp. 37–62, May 1997.

[15] J. F. Cardoso, "Blind signal separation: statistical principles", *Proceedings of the IEEE*, vol. 86, no. 10, pp. 2009–2025, October 1998.

[16] K. Abed-Meraim, W. Z. Qui, and Y. B. Hua, "Blind system identification", *Proceedings of the IEEE*, vol. 85, no. 8, pp. 1310–1322, August 1997.

[17] Simon Haykin, *Adaptive Filter Theory*, Prentice Hall International, third edition, 1996.

[18] H. Krim and M. Viberg, "Two decades of array signal processing research - the parametric approach", *IEEE Signal Processing Magazine*, pp. 67–94, July 1996.

[19] T. H. Chen, A. Hero, P. M. Djuric, H. Messer, J. Goldberg, D. J. Thomson, M. G. Amin, H. Krim, J. C. Pesquet, G. Giannakis, A. Swami, J. K. Tugnait, J. F. Cardoso, L. Tong, and J. Krolik, "Highlights of signal processing for communications", *IEEE Signal Processing Magazine*, vol. 16, no. 2, pp. 14–50, March 1999.

[20] Q. Wu and K. M. Wong, "Blind adaptive beamforming for cyclostationary signals", *IEEE Transactions on Signal Processing*, vol. 44, no. 11, pp. 2757–2767, November 1996.

[21] L. Castedo and A. R. Figueiras-Vidal, "An adaptive beamforming technique based on cyclostationary signal properties", *IEEE Transactions on Signal Processing*, vol. 43, no. 7, pp. 1637–1650, July 1995.

[22] Paul Petrus and Jeffrey H. Reed, "Time dependent adaptive arrays", *IEEE Signal Processing Letters*, vol. 2, no. 12, pp. 219–222, December 1995.

[23] S. V. Schell and W. A. Gardner, "Blind adaptive spatio-temporal filtering for wide-band cyclostationary signals", *IEEE Transactions on Signal Processing*, vol. 41, pp. 1961–1964, May 1993.

[24] G. Xu and T. Kailath, "Direction-of-arrival estimation via exploitation of cyclostationarity - a combination of temporal and spatial processing", *IEEE Transactions on Signal Processing*, vol. 40, no. 7, pp. 1775–1786, July 1992.

[25] J. Zhang, K. M. Wong, Q. Jin, and Q. Wu, "A new kind of adaptive frequency shift filter", in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Detroit, USA, May 1995, vol. 2, pp. 913–916.

[26] F. Dominique and P. Petrus, "Spectral redundancy exploitation in narrow band interference rejection for a PN-BPSK system", in *Proceedings of IEEE Military Communications Conference (MILCOM)*, 1994, pp. 405–409.

[27] W. M. Brown and R. B. Crane, "Conjugate linear filtering", *IEEE Transactions on Information Theory*, vol. IT-15, no. 4, pp. 462–465, July 1969.

[28] S. J. Baines, *Linear multi-user detection in DS-CDMA cellular systems*, PhD thesis, University of York, 1998.

[29] B. G. Agee, "Solving the near-far problem: exploitation of spatial and spectral diversity in wireless personal communication networks", in *Virginia Tech's Third Symposium on Wireless Personal Communications*, June 1993.

[30] V. Aue and J. H. Reed, "An interference robust CDMA demodulator that uses spectral correlation properties", in *Proceedings of the 43rd IEEE Vehicular Technology Conference*, 1994, pp. 563–567.

[31] R. D. Holley and J. H. Reed, "Time dependent adaptive filters for interference cancellation in CDMA systems", in *Workshop on Cyclostationarity Signal Processing*, Monterey, August 1994.

[32] J. H. Reed and B. G. Agee, "A technique for instantaneous tracking of frequency agile signals in the presence of spectrally correlated interference", in *Proceedings of the 23rd Annual Asilomar Conference on Signals, Systems and Computers*, 1992.

[33] R. D. Gitlin and S. B. Weinstein, "Fractionally-spaced equalization: An improved digital transversal equalizer", *Bell System Technical Journal*, vol. 60, no. 2, pp. 275–296, February 1981.

[34] Shahid U. H. Qureshi, "Adaptive equalisation", *Proceedings of the IEEE*, vol. 73, no. 9, pp. 1349–1387, September 1985.

[35] Shahid Qureshi, "Adaptive equalization", *IEEE Communications Magazine*, March 1982.

[36] R. C. North, R. A. Axford, and J. R. Zeidler, "The performance of adaptive equalization for digital communication systems corrupted by interference", in *Proceedings of the 27th Annual Asilomar Conference on Signals, Systems and Computers*, 1993, vol. 2, pp. 1548–1554.

[37] E. Biglieri, M. Elia, and L. Lopresti, "The optimal linear receiving filter for digital transmission over nonlinear channels", *IEEE Transactions on Information Theory*, vol. 35, no. 3, pp. 620–625, May 1989.

[38] Norihito Kinoshita and Seiichi Sampei, "Method of rejecting adjacent channel interference using an adaptive equalizer", *Electronics and Communications in Japan*, vol. 72, no. 11, pp. 1119–1126, 1989, Translated from Denshi Tshusin Gakkai Ronbunshi, vol. 71-B, No. 10, Oct. 1988, pp. 1119-1126.

[39] Carlos A. Belfiore and John H. Park, "Decision feedback equalization", *Proceedings of the IEEE*, vol. 67, no. 8, pp. 1143–1157, August 1979.

[40] L. Li and L. B. Milstein, "Rejection of CW interference in QPSK systems using decision-feedback filters", *IEEE Transactions on Communications*, vol. COM-31, no. 4, pp. 473–483, April 1983.

[41] P. Niger and Patrick Vandamme, "Performance of equalization techniques in a radio interference environment", *IEEE Transactions on Communications*, vol. 39, no. 3, pp. 452–457, March 1991.

[42] Norm W. K. Lo, David D. Falconer, and Asrar U. H. Sheikh, "Adaptive equalization for co-channel interference in a multipath fading environment", *IEEE Transactions on Communications*, vol. 43, no. 2/3/4, pp. 1441–1453, February/March/April 1995.

[43] Brent R. Petersen and David D. Falconer, "Suppression of adjacent-channel, cochannel, and intersymbol interference by equalizers and linear combiners", *IEEE Transactions on Communications*, vol. 42, no. 12, pp. 3109–3118, December 1994.

[44] Majeed Abdulrahman and David D. Falconer, "Cyclostationary crosstalk suppression by decision feedback equalization on digital subscriber loops", *IEEE Journal on Selected Areas in Communications*, vol. 10, no. 3, pp. 640–649, April 1992.

[45] Brent R. Petersen and David D. Falconer, "Minimum mean square equalization in cyclostationary and stationary interference - analysis and subscriber line calculations", *IEEE Journal on Selected Areas in Communications*, vol. 9, no. 6, pp. 931–940, August 1991.

[46] I. Howitt, J. H. Reed, V. Vemuri, and T. C. Hsia, "Recent developments in applying neural nets to equalization and interference rejection", in *Proceedings of Virginia Tech's Third Symposium on Wireless Personal Communications*, June 1993.

[47] H. E. Wong and J. A. Chambers, "Two-stage interference immune blind equaliser which exploits cyclostationary statistics", *Electronics Letters*, vol. 32, no. 19, pp. 1763–1764, September 1996.

[48] J. Zhang, K. M. Wong, Z. Q. Luo, and P. C. Ching, "Blind adaptive FRESH filtering for signal extraction", *IEEE Transactions on Signal Processing*, vol. 47, no. 5, pp. 1397–1402, May 1999.

[49] R. Mendoza, J. H. Reed, T. C. Hsia, and B. G. Agee, "Interference rejection using the time-dependent constant modulus algorithm (CMA) and the hybrid CMA/spectral correlation discriminator", *IEEE Transactions on Signal Processing*, vol. 39, no. 9, pp. 2108–2111, September 1991.

[50] S. Pei and M. Shih, "Fractionally spaced blind equalization using polyperiodic linear filtering", *IEEE Transactions on Communications*, vol. 46, no. 1, pp. 16–19, January 1998.

[51] William A. Gardner, *Representation and estimation of cyclostationary processes*, PhD thesis, University of Massachusetts, 1972.

[52] William A. Gardner, "Rice's representation for cyclostationary processes", *IEEE Transactions on Communications*, vol. COM-35, no. 1, pp. 74–78, January 1987.

[53] William A. Gardner, "Common pitfalls in the application of stationary process theory to time-sampled and modulated signals", *IEEE Transactions on Communications*, vol. COM-35, no. 5, pp. 529–534, May 1987.

[54] William A. Gardner, *Statistical Spectral Analysis*, Prentice Hall, 1988.

[55] William A. Gardner, "Spectral correlation of modulated signals: Part I - Analog modulation", *IEEE Transactions on Communications*, vol. COM-35, no. 6, pp. 584–594, June 1987.

[56] W. A. Gardner, W. A. Brown, and C.-K. Chen, "Spectral correlation of modulated signals: Part II - Digital modulation", *IEEE Transactions on Communications*, vol. COM-35, no. 6, pp. 595–601, June 1987.

[57] Jeffrey H. Reed and Tien C. Hsia, "The performance of time dependent adaptive filters for interference rejection", *IEEE Transactions on Acoustics, Speech and Signal Processing*, vol. 38, no. 8, pp. 1373–1385, August 1990.

[58] J. H. Reed, N. M. Yuen, and T. C. Hsia, "An optimal receiver using a time-dependent adaptive filter", *IEEE Transactions on Communications*, vol. 43, no. 2/3/4, pp. 187–190, February/March/April 1995.

[59] William A. Gardner and William A. Brown, "Frequency-shift filtering theory for adaptive co-channel interference removal", in *Proceedings of the 23rd Annual Asilomar Conference on Signals, Systems and Computers*, 1989, pp. 562–567.

[60] William A. Gardner and S. Venkataraman, "Performance of optimum and adaptive frequency-shift filters for cochannel interference and fading", in *Proceedings of the 24th Annual Asilomar Conference on Signals, Systems and Computers*, 1990, pp. 242–247.

[61] William A. Gardner, "Exploitation of spectral redundancy in cyclostationary signals", *IEEE Signal Processing Magazine*, pp. 14–36, April 1991.

[62] F. Hendessi, H.M. Hafez, and A.U.H. Sheikh, "The structure and performance of FRESH-decision feedback equalizer in the presence of adjacent channel interference", in *Proceedings of the IEEE Vehicular Technology Conference*, 1993, pp. 641–644.

[63] Lang Tong, Guanghan Xu, and Thomas Kailath, "Blind identification and equalization based on second-order statistics: A time domain approach", *IEEE Transactions on Information Theory*, vol. 40, no. 2, pp. 341–349, March 1994.

[64] Eric Moulines, Pierre Duhamel, Jean-Francois Cardoso, and Sylvie Mayrargue, "Subspace methods for the blind identification of multichannel FIR filters", *IEEE Transactions on Signal Processing*, vol. 43, no. 2, pp. 516–525, February 1995.

[65] Lang Tong, Guanghan Xu, B. Hassibi, and T. Kailath, "Blind channel identification based on second-order statistics: A frequency-domain approach", *IEEE Transactions on Information Theory*, vol. 41, no. 1, pp. 329–334, January 1995.

[66] D. Gesbert, P. Duhamel, and S. Mayrargue, "On-line blind multichannel equalization based on mutually referenced filters", *IEEE Transactions on Signal Processing*, vol. 45, no. 9, pp. 2307–2317, September 1997.

[67] H. Pozidis and A. P. Petropulu, "Cross-spectrum based blind channel identification", *IEEE Transactions on Signal Processing*, vol. 45, no. 12, pp. 2977–2992, December 1997.

[68] C. B. Papadias and D. T. M. Slock, "Fractionally spaced equalization of linear polyphase channels and related blind techniques based on multichannel linear prediction", *IEEE Transactions on Signal Processing*, vol. 47, no. 3, pp. 641–654, March 1999.

[69] Zhi Ding, "Characteristics of band-limited channels unidentifiable from second-order cyclostationary statistics", *IEEE Signal Processing Letters*, vol. 3, pp. 150–152, May 1996.

[70] William A. Gardner, "A new method of channel identification", *IEEE Transactions on Communications*, vol. 39, no. 6, pp. 813–817, June 1991.

[71] William A. Gardner, "Signal interception - a unifying theoretical framework for feature detection", *IEEE Transactions on Communications*, vol. 36, no. 8, pp. 897–906, August 1988.

[72] W. A. Gardner and C. M. Spooner, "Signal interception - performance advantages of cyclic-feature detectors", *IEEE Transactions on Communications*, vol. 40, no. 1, pp. 149–159, January 1992.

[73] W. A. Gardner G. K. Yeung, "Search-efficient methods of detection of cyclostationary signals", *IEEE Transactions on Signal Processing*, vol. 44, no. 5, pp. 1214–1223, May 1996.

[74] Nelson M. Blachman, "Beneficial effects of spectral correlation on synchronization", in *Cyclostationarity in Communications and Signal Processing*, William A. Gardner, Ed., chapter II, Article 2. IEEE, 1994.

[75] Fulvio Gini and Georgios B. Giannakis, "Frequency offset and symbol timing recovery in flat-fading channels: A cyclostationary approach", *IEEE Transactions on Communications*, vol. 46, no. 3, pp. 400–411, March 1998.

[76] Michel C. Jeruchim, Philip Balaban, and K. Sam Shanmugan, *Simulation of Communication Systems*, Applications of Communications Theory. Plenum Press, 1992.

[77] William. A. Gardner, "On the spectral coherence of nonstationary processes", *IEEE Transactions on Signal Processing*, vol. 39, no. 2, pp. 424–430, February 1991.

[78] Athanasios Papoulis, *Probability, Random Variables and Stochastic Processes*, McGraw-Hill, 3rd edition, 1991.

[79] John G. Proakis, *Digital Communications*, McGraw-Hill, third edition, 1995.

[80] Melvin J. Hinich, "Review of 'Statistical spectral analysis: A non-probabilistic theory' by W. A. Gardner", *IEEE Signal Processing Magazine*, pp. 14–16, April 1994.

[81] William A. Gardner, "Author's comment [on book review]", *IEEE Signal Processing Magazine*, pp. 16–18, April 1994.

[82] William A. Gardner, "Ensembles in wonderland", *IEEE Signal Processing Magazine*, pp. 18–23, April 1994.

[83] Neil L. Gerr, "To the editor", *IEEE Signal Processing Magazine*, p. 12, October 1994.

[84] William A. Gardner, "To the editor", *IEEE Signal Processing Magazine*, pp. 12–14, January 1995.

[85] Neil L. Gerr, "To the editor", *IEEE Signal Processing Magazine*, p. 16, March 1995.

[86] William A. Gardner, "Gardner rests his case", *IEEE Signal Processing Magazine*, November 1995.

[87] B. Picinbono, "On circularity", *IEEE Transactions on Signal Processing*, vol. 42, no. 12, pp. 3473–3482, December 1994.

[88] G. Johnson, "Questioning 'distinction without a difference' debate", *IEEE Signal Processing Magazine*, p. 19, November 1995.

[89] B. Picinbono, "Too complex to be real", *IEEE Signal Processing Magazine*, pp. 18–19, July 1996.

[90] G. Johnson, "The complex (and circular) argument continues", *IEEE Signal Processing Magazine*, pp. 42–44, September 1996.

[91] S. Chapman and J. Bartels, *Geomagnetism, Vol II, Analysis of the Data and Physical Theories*, Oxford University Press, 1940.

[92] C. H. D. Buys-Ballot, *Les changements périodiques de température dépendants de la nature du soleil et de la lune mis en rapport avec le pronostic du temps, déduits d'observations néerlandaises de 1729 à 1846*, Kemink & Son, Utrecht, 1847.

[93] Michael L. Honig, Pedro Crespo, and Kenneth Steiglitz, "Suppression of near- and far-end crosstalk by linear pre- and post-filtering", *IEEE Journal on Selected Areas in Communications*, vol. 10, no. 3, pp. 614–629, April 1992.

[94] Alan V. Oppenheim and Ronald W. Schafer, *Discrete-Time Signal Processing*, Prentice-Hall, second edition, 1999.

[95] R. E. Ziemer and W. H. Tranter, *Principles of Communications*, Houghton Mifflin, fourth edition, 1995.

[96] Simon Haykin, *Introduction to Adaptive Filters*, Macmillan Publishing Company, 1984.

[97] Earl R. Ferrara, "Frequency-domain implementations of periodically time-varying filters", *IEEE Transactions on Acoustics, Speech and Signal Processing*, vol. ASSP-33, no. 4, pp. 883–892, August 1985.

[98] R. A. Meyer and C. S. Burrus, "A unified analysis of multirate and periodically time-varying digital filters", *IEEE Transactions on Circuits and Systems*, vol. CAS-22, no. 3, pp. 162–168, 1975.

[99] Rodger E. Ziemer and Roger L. Peterson, *Introduction to Digital Communications*, Macmillan Publishing Company, 1992.

[100] B. Picinbono and P. Chevalier, "Widely linear estimation with complex data", *IEEE Transactions on Signal Processing*, vol. 43, no. 8, pp. 2030–2033, August 1995.

[101] Lajos Hanzo, "The pan-European cellular system", in *The Mobile Communications Handbook*, Jerry D. Gibson, Ed. CRC Press, 1996.

[102] Eugene Dubossarsky, Sam Reisenfeld, and Thomas R. Osborn, "Adaptive channel equalisation with the multirate optimising order statistics equaliser (MOOSE)", in *Proceedings of the IEEE Global Telecommunications Conference (GLOBECOM)*, Sydney, November 1998, pp. 3509–3514.

[103] F. Palmieri and C. G. Boncelet Jr., "L1-filters - a new class of order statistic filters", *IEEE Transactions on Acoustics, Speech and Signal Processing*, vol. 37, no. 5, pp. 691–701, May 1989.

[104] S. Burley and M. Darnell, "Robust impulse noise suppression using adaptive wavelet denoising", in *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Munich, 1997, IEEE, vol. V, pp. 3417 – 3420.

[105] Zhi Ding, "Blind channel identification and equalization using spectral correlation measurements", in *Cyclostationarity in Communications and Signal Processing*, William A. Gardner, Ed., chapter II, Article 4. IEEE, 1994.

[106] R. P. Gooch and B. Daellenbach, "Prevention of interference capture in a blind (CMA-based) adaptive receive filter", in *Proceedings of the 23rd Annual Asilomar Conference on Signals, Systems and Computers*, 1989, vol. 2, pp. 898–902.

[107] Saeed V. Vaseghi, *Advanced Signal Processing and Digital Noise Reduction*, Wiley, 1996.

[108] Ferrel G. Stremler, *Introduction to Communication Systems*, Addison-Wesley, second edition, 1982.

[109] J. F. Adlard, D. Grace, and T. C. Tozer, "An improved null steering algorithm for polarised VLF", in *Proceedings of the First International Workshop on Image and Signal Processing and Analysis*, Pula, Croatia, 2000, pp. 217–222.