Doppler Filtering and Detection Strategies for Multifunction Radar

by

GLEN DAVIDSON

A thesis submitted to the University of London for the Degree of Doctor of Philosophy
in Electronic Engineering

Department of Electronic and Electrical Engineering
University College London
January 2001

Abstract

This thesis concerns the analysis and processing of sea clutter from a Multiband Pulsed Radar - a land based research system operated by the British Defence Evaluation and Research Agency. This radar serves as a model for a class of Multi Function Radars (MFR) that offer extensive computer controlled adaptive operation.

A fast Sequential Edge Detector (SED) is formulated which, accounting for locally exponential speckle, allows the spatial inhomogeneity within a scene to be segmented. This simultaneously identifies high intensity areas and the noise dominated shadowed regions of the scene using an adaptively sized analysis window. The high resolution data is thus shown to contain discrete scatterers which exist in addition to the compound modulation from the wave surface.

The discrete component means the measured statistics cannot be considered homogenous or stationary. This is crucial for high resolution MFR as a priori information can no longer be relied upon when viewing a scene for the first time in order to make a detection decision.

Considering the returns to be discrete in nature leads to a potential Doppler detection scheme operable at low velocities within the clutter spectrum. A physically motivated test statistic, termed persistence, is demonstrated based upon the lifetime of scattering events determined via the Continuous Wavelet Transform.

When operated in coastal regions at low resolution, strong returns from the land-sea interface (edges) are expected which will seriously degrade the performance of radar detection models tuned to homogenous scenes. Explicit operational bounds are determined for the strength of these edges which show that simultaneous operation of an edge detector is required when assessing compound statistics such as the K-distribution using typical texture estimators.

Additionally a method for accurately determining the N-sum PDF of K-distributed statistics within noise is constructed using a numerical inverse Laplace transform. The SED is also applied to Synthetic Aperture Sonar data to detect the large shadows cast by targets rather than their point intensity.

To my Mother and Father; I owe them everything

Acknowledgement

Firstly, my thanks must go to Professor H.D. Griffiths. His help and supervision throughout this research have been unending, and he has made my stay at UCL both pleasant and memorable. Thanks too must go to Mr J Branson from DERA Portsdown, who has aided and encouraged. I would also like to express my gratitude to Professor K. Milne and Professor R. Benjamin for their expert advice and stimulating discussions during their visits to UCL. I appreciate the help from my friends and colleagues who have contributed to this thesis: Joe, Jonathan, Lluís, Moon, Richard, Rob, Simon, Trevor, Wookie and Yves from the Antennas and Radar Group at UCL and thanks to Nick Clayton who kept me sane on many occasions. I would like to acknowledge the financial assistance from the Engineering and Physical Sciences Research Council and the Defence Evaluation and Research Agency. Last, and far from least, I wish to thank Norie - I would never have finished this without you.

Contents

List of Figures

List of Tables

Nomenclature

α	Threshold Multiplier
β	Wavelet Translation Parameter
Δf	Bandwidth
Δf_s	Frequency Step Size
Δf_{dop}	Doppler Spectral Width
η	Subset of set N
γ	Euler-Gamma Number
λ	Variable Threshold
λ	Wavelength
μ	Arithmetic Mean
μ_A	Arithmetic Mean
μ_G	Geometric Mean
ν	K / Gamma Distribution Shape Parameter
ω_s	Sampling Frequency
Φ	Doppler Spectrum
ϕ	Signal Phase
Ψ	Wavelet Basis Function
$\psi^{(n)}$	Polygamma Function Of Order n
ρ	Significance Level
σ	Radar Cross Section
σ	Weibull Scale Parameter
σ^2	Variance
τ	Correlation Length of Target Motion

τ	Delay Time
τ	Transmit Pulse Duration
θ^m	m Parameter Vector
$\xi(\tau)$	Autocovariance Function
ζ	Doppler Frequency
${}_2F_1\{\cdot\}$	Hypergeometric Function
A	Gamma Distribution Scale Parameter
a	Amplitude
a	Weibull Shape Parameter
b	K Distribution Scale Parameter
c	Speed of Light
$C(x)$	Cumulative Probability of x
d	Target Dimension
$E[x]$	Expectation of x
$f(x)$	Statistical Distribution of x
f_c	Cutoff Frequency of Target Motion
f_p	Pulse Repetition Frequency
f_t	Transmit Frequency
G	Gain
H	Hazard Rate
H_i	Hypothesis i
I	Intensity
$I^{(m)}$	mth Moment of I
I_n	Modified Bessel Function of Order n
j	Complex Component
K	Kuiper Statistic
k	Chaotic Parameter
k	Vector Index
K_n	Modified Bessel Function of Order n
N	Set Population
p	Proportion of Step Over Sample Length

$P(x)$	Probability of x
P_d	Probability of Detection
P_j	Thermal Noise Power
P_r	Received Power
P_t	Transmit Power
P_{fa}	Probability of False Alarm
R	Range
R	True Intensity Ratio
r	Sample Intensity Ratio
s	Dominant:Small Scatterer Ratio
s	Magnitude of Intensity Step
s	Wavelet Scale Parameter
T	Duration
T	Fixed Threshold
T	Temperature
t	Sampled Time
U	Normalised Log Shape Estimate
U_0	True Underlying U
U_T	Distribution Change of U
$U_{\Delta I}$	Change in U due to Intensity Step
$U_{\Delta U}$	Change in U due to Distribution Step
U_ρ	Change in U due to Modulating Distribution
U_ρ	Estimated Value of U Subject to Finite Sample Number and/or Inhomogeneity
V	Contrast Shape Estimate
v	Velocity
V_a	Amplitude Contrast Shape Estimate
W	Variance of Log Shape Estimate
W	Window Size
$W_s(n)$	Wavelet Power Coefficient of Scale s at Time n
x	Signal Envelope
$x_{m:n}$	Vector of Signal m to n

Z	Resultant Signal
Z	Test Statistic
A/D	Analogue to Digital (Converter)
ACF	Auto-Correlation Function
CA	Cell Averaging
CAGO	Cell Averaging Greatest Of
CDF	Cumulative Probability Distribution Function
CFAR	Constant False Alarm Rate (Processor)
CW	Continuous Waveform
CWT	Continuous Wavelet Transform
dB	Decibels
DERA	Defence Evaluation and Research Agency
DWT	Discrete Wavelet Transform
FFT	Fast Fourier Transform
flops	Floating Point Operations per Second
FWHM	Full Width at Half Maximum
FWSE	Fixed Window with Scanning Edge
HH	Transmit Horizontal, Receive Horizontal Polarisation
HV	Transmit Horizontal, Receive Vertical Polarisation
I	In Phase Component of Quadrature
IID	Independent Identically Distributed
K	K Distribution
K-S	Kolmogorov-Smirnov Test
KP	Kuiper Test
MBPR	Multi-Band Pulse Radar
MFR	Multi Function Radar
MLE	Maximum Likelihood Estimate
MTD	Moving Target Detection
MTI	Moving Target Indicator
NECAPS	Naval Environmental Clutter Attenuation and Propagation Specification
NP	Neyman-Pearson

O(n)	Order of Magnitude n
OS	Order Statistic
PDF	Probability Distribution Function
PRF	Pulse Repetition Frequency
Q	Out of Phase Component of Quadrature
RTI	Range Time Intensity
SED	Sequential Edge Detector
SNR	Signal to Noise Ratio
SWCE	Scanning Window with Central Edge
VH	Transmit Vertical, Receive Horizontal Polarisation
VV	Transmit Vertical, Receive Vertical Polarisation
WFT	Windowed Fourier Transform
WT	Wavelet Transform

Chapter 1

Introduction

1.1 History

Radar (*Ra*dio *D*etection *A*nd *R*anging) equipment transmits radio waves and receives the backscattered radiation to allow detection and analysis of distant targets. In 1886 Hertz demonstrated that radio waves are reflected from both metallic and dielectric objects. The field of naval radar could be said to have begun in 1939 when prototype equipment was fitted to the battleship *Rodney* and the cruiser *Sheffield;* however as early as 1904 a patent was filed by the German engineer Christian Hülsmeyer for detection of radio waves reflected from ships.

In 1922 Marconi made a speech urging the use of short radio waves for detection of objects at the Institute of Radio Engineers (now the Institute of Electrical and Electronics Engineers) whilst in the same year A.H. Taylor and L.C. Young experimentally demonstrated detection of wooden ships at the Naval Research Laboratory using Continuous Waveform (CW) radar. In 1930 L.A. Hyland at NRL discovered that a moving aircraft could be detected due to 'wave-interference' between the transmit and receive signals.

The limitation of CW radar was that the range of the reflecting target is difficult to determine, this was solved by transmitting radar signals as a series of short pulses. Successful demonstration of a pulse radar was achieved first in Britain by Sir Robert Watson-Watt in June 1935 which by September that year could detect bomber aircraft at ranges greater than 40 miles. The onset of the Second World War led to an urgency to develop methods of detecting low-flying aircraft and small surface targets. Accuracy demands that the wavelength employed is as short as possible to enable

a concentrated beam to remain close to the sea surface, when the wavelength reached the physical dimensions of the conventional valves then used, progress in radar was temporarily halted.
The work of Randall and Boot at Birmingham University in 1940 led to the development of the cavity magnetron which operated with a power output 100 times that previously achievable at a wavelength of 10cm. By 1943 the detection of U-boat periscopes was a reality which in combat could prove effective at a range of 6 miles (Wylie 1978). In the radar field counter-measures were established that adopted chaff, jamming and stealth so as to mask a target behind noise and because of this signal processing has evolved into an esoteric mixture of tools for the radar designer to employ.

1.2 Previous Work

Detection of 'small' targets close to the sea surface is complicated by background environmental noise or 'sea clutter'. Gaussian statistics naturally arise from the central limit theorem to give an expected Rayleigh amplitude and exponential intensity clutter distribution - observed as 'speckle' in coherent imaging systems. The statistics of radar backscatter from the ocean surface is known to deviate significantly from Gaussian, prompting strong debate in the literature upon the form that deviation takes. An obvious problem from an experimental point of view is the range of environmental parameters that may have an effect, countered by the theoretical desire to form physical theories that lead to optimum target detection schemes. This Section gives a brief history of this field, and is by no means exhaustive.

Intensity Distribution of Sea Clutter

The early history of sea clutter observation can be found in Long (1983). By the late 1960s a number of studies had established general dependencies such as an increase in expected signal strength with increasing grazing angle, increased wind strength, vertical polarisation and in upwind directions (Skolnik 1981). Ship-borne maritime surveillance radar generally transmits and receives from the ship's mast and is necessarily operated at a low grazing angle; this in particular has shown wide variation in reported statistics. A particular problem in target detection is predicting the probability of high intensity returns from a limited number of samples; many researchers have reported on the nature of the high intensity tails from sea clutter returns with varying conclusions.

Trunk and George (1970) showed that a lognormal distribution could approximate the clutter at a grazing angle of 4.7° in sea state 2-3 when a 200ns pulse was used. Whilst lognormal provided better fits to some data, it does not reduce to Gaussian statistics and consequently overestimates the low probability returns in some cases (Skolnik 1981).

Fay et al. (1977) operated an X-band radar at grazing incidence which could transmit or receive in horizontal or vertical polarisation (HH and VV), using a 70-270 ns pulse, in a wind of 10-15 knots and found the Weibull distribution to be an excellent fit. Both lognormal and Rayleigh were dismissed. The Weibull distribution is a model that can handle the long low probability tails of sea clutter and also includes the exponential distribution as a special case. The model is versatile enough to fit both long tailed data and noise dominated data which consequently led to wide application within detection theory; covered by Sekine and Mao (1990)

The Weibull model has no strong physical justification and as early as 1957 high magnitude returns were associated with steep and breaking waves (Katzin 1957); these intermittent returns are generally termed 'sea spikes'. Using simultaneous video Lewis and Olin (1980) confirmed the link to breaking waves. They also showed a polarisation dependence that sea spikes occurred less frequently within horizontally polarised observations - giving greater variance to the distribution when compared with vertical polarisation.

A composite surface model (Wright 1968) was one of the first attempts to relate the observations of sea spikes with theory. It suggested that a Bragg backscattering resonance condition was responsible for the spikes; caused by enhanced backscatter from surface waves with a wavelength equal to half that of the radar transmission. If Bragg resonance is present, a physically large sea swell may only give a large backscatter if wind-formed capillary waves are present upon its surface (Skolnik 1981). With the capillary surface waves suitably modulated by the tilting surface, a complex interdependency upon environment is present which could account for some of the trends seen in data (Wright 1966).

In the light of the composite surface model, Hansen and Cavaleri (1982) analysed low grazing angle X band data and found that HH polarised data showed an abrupt deviation from an overall Weibull distribution for large intensities, no such effect was present in VV polarisation. Olin (1982) found that whilst the VV polarised data could be fitted by a single Weibull distribution, the HH polarisation conformed to a different Weibull shape at low intensities than at high. The suggestion that a different scattering mechanism was responsible for the extremely high magnitude returns was further explored by Trizna (1991). Quantifying the dependence upon grazing angle and wind speed, Trizna suggested the Weibull fit to high backscatter regions was associated with discrete scatters linked to whitecap coverage; while that fitting low regions was caused by distributed roughness accounted by the composite scatter model.

Whilst associating individual spikes with surface features can provide some explanation of the observed backscatter, a typical detection scenario may be based upon observation of such a large range extent or resolution cell that many spiking events are included. A large enough number of discrete scatterers is ultimately generalised to a continuous overall backscatter distribution; however the coherent imaging effects resulting from interference between scatterers (speckle) must still be considered. This 'compound formulation' has gained significant popularity recently, due in part to its physical justification.

Jakeman, Pusey and Tough (Jakeman and Pusey 1976, Jakeman and Tough 1981) considered optical scattering from a random media when the population of scatterers fluctuates according to a negative binomial distribution and showed that this resulted in a K-distribution. The K-distribution has, in the limit of infinite scatterer population, the exponential distribution as a special case. The negative binomial distribution arises from consideration of birth-death-(im)migration population statistics (Jakeman 1980) and is a discretised form of the gamma distribution.

Ward (1981) gave the K-distribution firm experimental justification by operating an airborne X-band radar, using a 30ns pulselength with 1.2° beamwidth at a grazing angle of 0.75° observed in sea state 3. Summing 36 consecutive range samples and utilising frequency agility, the speckle was averaged and after output to a TV recording was suitably characterised by a gamma distribution. This was said to be acceptable over sea states 1-5 and varying swell directions. The physical compound model proposed by Ward is that the underlying mean backscatter is slowly varying gamma distributed - this backscatter is sampled subject to speckle with a correlation time of the order of 10ms. The underlying correlation is particularly useful in modelling clutter; Watts (Watts 1996) in particular has quantified the performance *gain* available through knowledge of the correlation structure.

The gamma distribution also arises from general statistical consideration of a family of infinitely divisible distributions (Ward et al. 1990) as it is well known to be closed under convolution such that the sum of N similar gamma variates will itself be gamma distributed. If the mean backscatter from a (large) low resolution imaged area arises from a summation of several (smaller) high resolution imaged patches then observation of both resolutions, subject to speckle, will be K distributed with a shape parameter related to resolution.

Ward, Baker and Watts (1990) using X-band equipment operating at 4m range resolution with 1.2° azimuth at various range and grazing angles, gave examples where 250ms time sequences are locally Rayleigh but the mean value was gamma distributed confirming an overall K-distribution. Analysis of coherent returns demonstrated a fast speckle component at 5-10ms with a slower, presumably physical, periodic modulation of the order of seconds. It was reported that the distribution of Doppler returns was dependent upon velocity, but that their correlation properties were similar. The form of range-time images was shown to depend on resolution and look angle and an empirical K-distribution fit was shown to relate the radar parameters to the distribution shape. A large number of analysed data sets showed horizontal polarisation was significantly spikier than vertical. From a moment based analysis the vertical polarisation showed a good fit to the K distribution but the horizontal polarisation also showed some deviation. No significant statistical trend was established for variations with sea state, wind speed or aspect angle relative to wind although complex interdependencies were not dismissed.

Any 'diversity' between polarisation channels offers further methods for target discrimination. Farina et al. (1997) operating an X band radar with 30m range resolution and 0.9° beam width using data from sea state 3 at 0.645° grazing angle operating in coastal waters of significant wave height 1.42m found that only the vertical polarisation (VV) was fit by a K-distribution and that even accounting for thermal noise, horizontal (HH) and cross polarised channels (HV or VH where the received pulse is at a different polarisation to transmit) fell somewhere between K and lognormal.

Whilst two parameter distributions are naturally easier to analyse, thus preferable for establishing environmental relations, a number of higher parameterisations have been suggested. Azzarelli (1995) suggests a general class of models that have a fluctuating number of non-Gaussian scatterers in the presence of Gaussian noise; resulting in a four parameter distribution that is shown to fit a wide range of observed data distributions, with the K distribution as a limiting case. Further suggestions have been the generalised gamma (Anastassopoulos et al. 1999) and the generalised K, applied to sonar (di Bisceglie et al. 1999).

Velocity Distribution of Sea Clutter

The surface backscatter often has a spatially wave-like form, with a period linked to that of the ocean waves causing variation in range and time. This is not always the case, HH polarisation in particular can show far less periodicity than VV when simultaneously observed (Ward et al. 1990); however it is reasonable to assume that the observed scatterer motion has a strong dependence on the physical ocean wave motion. Determination of velocity requires the phase relationship between transmit and received waveforms and is commonplace for detecting high velocity airborne targets. The added complexity means that high resolution coherent radars are not common despite the potential for improved detection of slow moving targets.

A range of velocities is expected from the surface due to the finite (observed) lifetime of scattering events (Lee et al. 1995), commonly assessed by a Doppler spectrum. Differences in the spectra of HH and VV were analysed over 30 years ago (Pidgeon 1968, Valenzuela and Laing 1970), reporting that HH has a higher spectral peak than VV. Whilst literature is available from large grazing angle observation (Duncan et al. 1974), Lee and his colleagues (Lee et al. 1995, Lee et al. 1996, Lee et al. 1998) have dominated the research for typical marine surveillance grazing angles. Lee et al. report that the peak separation of the Doppler spectrum is a result of 'fast scatterers'. The existence of non-Bragg scattering is postulated as 'super events' where HH backscatter dominates VV, the number of super events reportedly increases with wind speed (Lee et al. 1996). Lee et al. (1998) verifies with laboratory wave tank experiments that the fast scatterers are due to breaking-wave backscatter.

The intensity distribution of individual velocity components of sea clutter is more complex than that of incoherent intensity measurements. Ward et al. (1990) reported that the low intensity tails of the Doppler spectrum, an obvious area for target detection, consisted of the spikiest intensity distribution which is difficult to characterise. The peak of the Doppler spectrum is also reported to be a function of wind direction with HH consistently having a larger Doppler offset than VV (Ward et al. 1990).

Posner (1998) utilised a high resolution 0.3m polarimetric radar with a pulse repetition frequency of 2000Hz at a grazing angle of 0.2°. Individually observed scattering events showed large variation in magnitude but had a constant velocity over their lifetime, suggesting that relative oscillations of a few strong scatterers or whitecap formation was responsible. Posner (1998) also reported large scale wave phenomena travelled at the predicted phase speed of the ocean waves with smaller

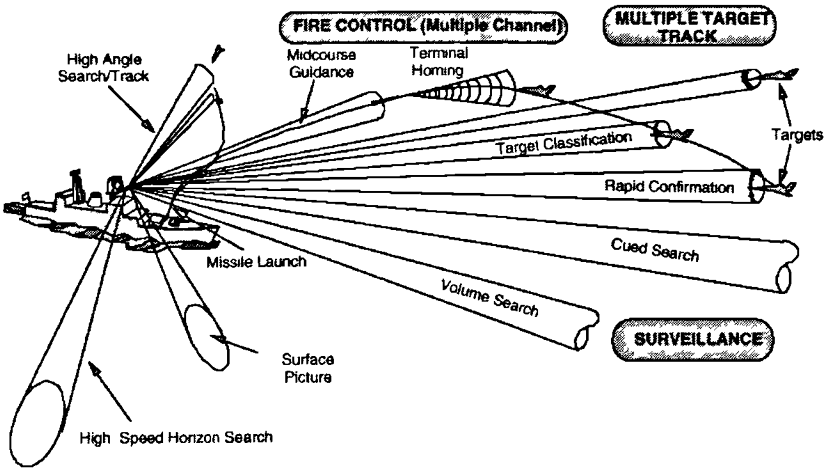

Figure 1.1: In a Multi-Function Radar, surface target detection is a small proportion of the required tasks, (after Butler, 1998)

scale wave phenomena at the group speed, observed in upwind HH and VV.

1.3 Motivation

The Multiband Pulsed Radar (MBPR) is a land based radar operated by the British Defence Evaluation and Research Agency (DERA). Whilst still a research tool, it serves as a model for a class of Multi Function Radars (MFR) which offer extensive computer controlled adaptive operation; including:

- A phased array antenna can form a beam at a variable resolution and dwell time.
- A number of pulse frequencies and polarisations can be chosen.
- Priority tasks can reduce the available number or frequency of pulses available to a particular process.
- Processing is achieved using high performance microprocessors.

Figure 1.1 shows the wide range of tasks required of an MFR. It is evident that surface target detection covers only a small proportion of the radar workload.

To detect targets, normally a deviation is identified in the intensity or velocity statistics due to the observed target radar returns. Usually a threshold is set from information about the target-free environment which requires prior knowledge from analysis of large amounts of data. Returns from a scene contain clutter - defined as any unwanted backscatter. For sea clutter in particular, a large amount of effort has been made by researchers to characterise the cause and effect of 'sea spike' features which cause the clutter returns to deviate from the basic scattering theory with a resultant loss in efficiency.

This thesis concerns the analysis and processing of sea clutter from the MBPR. It covers the physical aspects of the radar operation and discusses the mathematical tools necessary to analyse the clutter statistics. After analysis of real data, various methods of detection are discussed followed by simulation and use of actual targets in order to test these methods. When operated in coastal regions, strong returns from the land-sea interface are expected which will seriously degrade the performance of radar detection models tuned to homogenous scenes. Whilst this work only presents results from pure ocean scenes, the presence of discontinuities is considered throughout; although physical observation was not possible due to the land based radar employed.

1.4 Contributions Made by this Thesis

This work comes to the conclusion that at high resolution the wide variability of sea clutter is caused by inhomogeneous nonstationary clutter statistics due to the presence of discrete scatterers. This is crucial for MFR as we can no longer use *a priori* information to such an extent when viewing a scene for the first time in order to make a detection decision. Considering the returns to be discrete in nature leads to a potential Doppler detection scheme operable at low velocities within the clutter based upon the lifetime of the discrete scattering events.

The following aspects represent an addition to the body of knowledge about characterisation of radar clutter observed at low grazing angles and the detection problems that arise:

- A simple texture estimate U, already proposed in the literature as the approximate MLE solution to the K distribution shape estimate, is insufficient to characterise all aspects of backscatter fluctuation; however it is a useful tool to analyse events as it is Gaussian when operated at small sample population, within noise and for censored samples. Explicit determination of edge effects upon U allows correction of the determined statistics for inhomogeneous scenes.
- The potential contribution by Gaussian statistics due to system noise over a significant spa-

tial proportion of the high resolution images and time proportion of the Doppler spectra is not appreciated in the literature.

- A Sequential Edge Detector is proposed as a means of simultaneously detecting edges and targets in many models of clutter. This can additionally identify the shadowed regions in high resolution data and determine the noise component of time varying Doppler bins.
- Whilst modulation is present within sea clutter as in the compound formulation, discrete scatterers can be present whose magnitude is not related to the underlying compound distribution.
- The lifetime of the individual discrete scattering events contributing to the Doppler spectrum can be measured in the Wavelet domain yielding an expected exponential distribution.
- Target detection based on the lifetime of events is shown to be possible which leads to a scheme suited for relatively long observation times of slow moving targets at clutter velocities. Validation is performed in real clutter for both real and simulated targets which suggests the method is complementary to simple Radar Cross Section (RCS) based thresholding.
- Assuming that a low resolution is used so the RCS statistics *can* be described by a continuous distribution, explicit determination of the stability and sample size required for operational accuracy within arbitrary K-distributed clutter is performed. This shows that the region of weakly spiky data is the most susceptible to edges when determining the shape parameter. Thus concurrent edge detection is shown to be necessary before target detection is performed when operating in real environments.
- The distribution of an arbitrary sum of K-distributed clutter within noise is shown to be obtainable via a fast numerical inverse Laplace transform method which is accurate to false alarm probabilities down to 10^{-7}.

Additionally, in Appendix A, the Sequential Edge Detector is applied to Synthetic Aperture Sonar data. This leads to a proposed detection scheme where the shadows arising from targets are detected rather than the targets themselves. A single scene is tested which shows an apparent improvement in detection rate.

Chapter 2

Physical Aspects

2.1 The Generalised Radar Model

Before a detailed analysis of the data is presented, the physical aspects of the radar must be explored. This Section covers properties common to pulsed radar systems where a pulse at radar frequencies is emitted and its interaction with the environment is assessed from analysis of the echo.

2.1.1 The Rayleigh Scattering Model

When a radar pulse impinges upon a surface, it is assumed that the reflected pulse is contributed by N multiple scattering centres, each of which can be viewed as giving an amplitude a with a phase ϕ distributed uniformly $0 < \phi \leq 2\pi$. Due to the wave nature of the pulse these add vectorially, shown in Figure 2.1, such that the resultant signal Z is

$$Z = \sum_{n=1}^{N} a_n \exp\left(j\phi_n\right) \tag{2.1}$$

which can be viewed as a random walk upon the complex plane. The amplitude distribution $P\left(a\right)$ is unknown, however assuming N is large and invoking the Central Limit Theorem means both

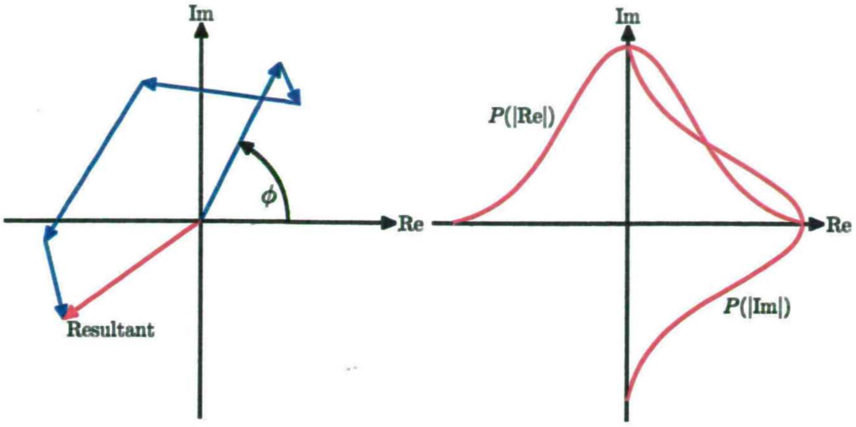

Figure 2.1: Received signal formed by a random walk in the complex plane

the real and imaginary parts of Z will conform to a Gaussian distribution

$$
\begin{aligned}
P(\mathrm{Re}(Z)) &= \frac{1}{\sqrt{\pi\langle a^2\rangle}}\exp\left(\frac{-|\mathrm{Re}(Z)|^2}{\langle a^2\rangle}\right), \; -\infty < \mathrm{Re}(Z) < \infty && (2.2)\\
P(\mathrm{Im}(Z)) &= \frac{1}{\sqrt{\pi\langle a^2\rangle}}\exp\left(\frac{-|\mathrm{Im}(Z)|^2}{\langle a^2\rangle}\right), \; -\infty < \mathrm{Im}(Z) < \infty && (2.3)
\end{aligned}
$$

where $\langle a^2\rangle$ is the expected mean square amplitude.
Assuming independent components

$$
\begin{aligned}
P(\mathrm{Re}(Z), \mathrm{Im}(Z)) &= P(\mathrm{Re}(Z)) \times P(\mathrm{Im}(Z)) && (2.4)\\
\delta\mathrm{Re}\,\delta\mathrm{Im} &\equiv x\,\delta x\,\delta\phi
\end{aligned}
$$

shows the envelope x of the received radar echo $|Z|$ is Rayleigh distributed

$$
P_{Rayleigh}(x) = \frac{2x}{\langle x^2\rangle}\exp\left(-\frac{x^2}{\langle x^2\rangle}\right), \; 0 \leq x < \infty \qquad (2.5)
$$

this is seen as multiplicative noise in coherent imaging systems such as radar, and is usually termed speckle in reference to the grainy appearance of images.
Several assumptions are made in the above argument which may be violated in high resolution systems:

- The number of scattering centres N is large.
- Scattering centres are such that the phase ϕ is spread uniformly.

- a is constant and statistically independent from ϕ over the imaging time.
- The radar pulse is narrow bandwidth such that the physical interaction with respect to changing frequency is not discontinuous.

2.1.2 The Radar Equation

A simple equation can be derived to give the relationship between the received signal strength and the radar characteristics. Due to the military origins of radar it is common to talk of a 'target' irrespective of the type of imaged scene. Making the assumptions (Skolnik 1981):

1. An antenna transmission of beam power P_t reaching range R_t is achieved.
2. Directionally dependent gain G from a directive transmitting antenna is employed.
3. The target interacts with the beam, re-radiating an amount equivalent to that which would be seen from a perfectly reflecting cross sectional area σ.
4. A proportion of the reflected signal is intercepted dependent upon the receiving antenna's area A_r at a distance R_r from the target.

Thus the received power P_r can be determined as

$$P_r = \frac{P_t}{4\pi R_t^2} \times G \times \frac{\sigma}{4\pi R_r^2} \times A_r \tag{2.6}$$

which in the case of a monostatic radar utilising a single antenna for transmit and receive reduces to

$$P_r = \frac{P_t G \sigma A}{(4\pi)^2 R^4} \tag{2.7}$$

Since the gain G and the effective area A_e of an antenna are related by

$$G = \frac{4\pi A_e}{\lambda^2} \tag{2.8}$$

then the monostatic radar equation can be expressed as

$$P_r = \frac{P_t G^2 \sigma \lambda^2}{(4\pi)^3 R^4} \tag{2.9}$$

Additionally the surface radar footprint (the area illuminated in range and azimuth) is linearly proportional to R. When viewing a distributed scene, rather than a point target, this causes an R^{-3} dependence upon range. The radar equation only gives a first approximation to the performance of the radar due to variation in σ, the concern of the next Section.

2.1.3 Cross Section Fluctuation

The cross section σ is accurately termed the Radar Cross Section (RCS) as it is the target cross section as viewed by the radar. A target detection strategy is usually based upon the determined statistics of σ; measurement difficulties arise from factors including:

- Material properties of the target such as absorption and reflection characteristics varying with frequency. There are three scattering regions dependent upon wavelength λ and target dimension d (Rees 1990):
    - $2\pi d/\lambda \ll 1$ is the Rayleigh region where $\sigma \propto \lambda^{-4}$
    - $2\pi d/\lambda \sim 1$ is the Mie region where the σ is oscillatory or erratic
    - $2\pi d/\lambda \gg 1$ is the optical region where σ approaches the optical area of the object

  Analysis of complex targets is difficult but analytic results can be determined for simple shapes such as a sphere (Rheinstein 1968).

- Propagation effects due to the atmosphere consist of (Skolnik 1981):
    - Attenuation (Rees 1990):
        - $\lambda \lesssim 0.4\mu$m dominated by Rayleigh scattering by molecules
        - 0.8μm $\lesssim \lambda \lesssim 15$mm dominated by molecular absorption
        - $\lambda \gtrsim 10$m dominated by ionosphere
    - Refraction due to atmospheric density gradients

  Refraction is largely responsible for anomalous propagation due to a phenomenon known as ducting. This can cause significant interference as antenna sidelobes in elevation are channeled, potentially giving a second propagation route to the imaged target, however it can be advantageous in determining the elevation of the imaged scene as explored by Money et al. (1997a). Intentional Over The Horizon (OTH) effects can occur from ionospheric reflection.

- Aspect ratio has a significant effect (Dunn and Howard 1970) since most targets consist of multiple scattering centres, all of which interact to modulate the RCS between total reinforcement or total cancellation dependent upon their relative positions.

Image removed due to third party copyright

Figure 2.2: RCS fluctuation models by Swerling with Ricean for comparison

Whilst not affecting the statistics of RCS, the following can affect the recorded signal:

- The antenna beam shape has a lobed structure which causes external targets to be visible within the central beam.
- Receiver noise due to Johnson noise within the circuits (Johnson 1928). This gives an available thermal noise power P_J in terms of Boltzmann's constant k, ohmic component temperature T and bandwidth Δf

$$P_J = kT\Delta f \tag{2.10}$$

As received power is proportional to RCS (Equation 2.7), the instantaneous recorded intensity is used to record variation in RCS over time; discussed next.

2.1.4 Fading Models

If the imaged scene is stationary with respect to the radar then the received signal will be constant. In a typical scene the effects of the previous Section will have an unknown overall effect and so several case models have been proposed by Swerling (1960) to account for the RCS variation seen over time. Taken literally from Skolnik (1981) the four types of fading are:

1. *The echo pulses received from a target on any one scan are of constant amplitude throughout the entire scan but are independent (uncorrelated) from scan to scan. This assumption ignores the effect of the antenna beam shape on the echo amplitude. An echo fluctuation of*

this type will be referred to as scan-to-scan fluctuation. The probability density function for the cross section σ is given by the density function

$$P(\sigma) = \frac{1}{\sigma_{av}} \exp\left(-\frac{\sigma}{\sigma_{av}}\right) \ , \ \sigma \geq 0 \tag{2.11}$$

where σ_{av} is the average cross section over all target fluctuations.

2. *The probability density function for the target cross section is also given by Equation 2.11 but the fluctuations are more rapid than in case 1 and are taken to be independent from pulse to pulse instead of from scan to scan.*

3. *In this case, the fluctuation is assumed to be independent from scan to scan as in case 1, but the probability density function is given by*

$$P(\sigma) = \frac{4\sigma}{\sigma_{av}^2} \exp\left(-\frac{2\sigma}{\sigma_{av}}\right) \ , \ \sigma \geq 0 \tag{2.12}$$

4. *The fluctuation is pulse to pulse according to Equation 2.12.*

In addition a notional 'Case 0' can be defined to represent a non-fluctuating signal of constant strength. Cases 1 and 2 consider the target as a collection of many independent scatterers of approximately equal echoing areas and is just the power equivalent of Equation 2.5. Case 3 and 4 are said to represent one large target with other small reflectors (Skolnik 1981); however this assumption leads to a Ricean distribution (Jao and Elbaum 1978)

$$P(\sigma) = \frac{1+s}{\sigma_{av}} \exp\left[-s - \frac{\sigma}{\sigma_{av}}(1+s)\right] I_0\left(2\sqrt{\frac{\sigma}{\sigma_{av}} s(1+s)}\right) \ , \ \sigma > 0 \tag{2.13}$$

where s is the ratio of dominant scatterer RCS to the total of the small scatterers and I_0 is the modified Bessel function of zero order. Equations 2.11 to 2.13 are shown in Figure 2.2.

A fixed length scan time is not applicable to MFR - only in rotating systems are the durations clearly defined. The observed cross section fluctuations from clutter rarely correspond to one of the above physical models and so a number of empirical solutions have been put forward, dealt with in Section 3.1. The motivation for using the above is that it provides standard target models with which to compare detection strategies.

The fading models make no attempt to characterise the second order statistics - the variation in the instantaneous statistics with time; but allow the two extremes of pulse and scan correlation. If a single target is being imaged, neglecting the effects mentioned in Section 2.1.3, it is reasonable

Table 2.1: Specifications of the DERA MPR radar

Transmit Frequency - f_t	2.9-3.1 GHz (S), 8-18 GHz (X-K$_u$), 33-37 GHz (K$_a$)
Respective Beamwidth	1.5°, 2°, 3°,
Transmit Pulse Width - τ	1μs
Polarisation	Horizontal, Vertical, Co-polar. Alternation in S & X-K$_u$
Pulse Repetition Frequency - f_p	5, 10, 20, 25, 30, 40 & 50 KHz
Stepped Frequency Bandwidth - Δf	0, 100, 200, 400 & 800 MHz
Frequency Step Size - Δf_s	0.39, 0.78, 1.56 & 3.13 MHz
Calibration	0.3m aluminium sphere flown in free space by kite or balloon

to assume that the decorrelation time is primarily related to the relative position of the scatterers contributing to a target return. Whilst they are stationary the resultant backscatter will remain constant; simplistically a $\lambda/2$ resultant movement will cause fading due to cancellation between scatterers introducing a dependence upon the surface velocity of the target. Sea clutter is often quoted to have a decorrelation length of order 10ms at X band (Croney 1970) which can be approximately derived from the width of the Doppler spectrum. This makes an important assumption that the Doppler spectrum is stationary; if not then there will be a spectrum of correlation lengths in addition to the velocity dependence.

2.2 Specifics of the Multi-Function Radar

Whilst the previous Section dealt with issues common to the majority of pulsed radars, details specific to the DERA Multi-Band Pulsed Radar (MBPR) are now covered. Designed as a research radar to investigate aspects of Multi Function Radar (MFR) processing, it can operate in many different configurations - detailed in Table 2.1. The transmit frequencies are labelled by the letter band system adopted since World War II; accepted by the IEEE as a standard although British band designations may differ slightly.

All imaging involves the sequential transmission of $\tau = 1\mu$s duration pulses forming a range gate of extent $c\tau/2$. These pulses are emitted at a rate f_p per second - the Pulse Repetition Frequency (PRF). The radar is coherent so the phase ϕ of both the transmit and receive pulse is known, this allows two distinct methods of imaging - Doppler and Hi-resolution (Hi-res). Doppler imaging uses a single transmit frequency f_t to record the wideband response of the sea surface, primarily to identify the velocity spectrum of the scene. Hi-res modes utilise a set of pulses swept over a range of frequencies centred at f_t with a bandwidth Δf via a sequential linear frequency step Δf_s. After frequency transformation this yields an interpretation of the range profile within the range gate; for this reason the mode is often called range resolution.

The pulse can be polarised Horizontal or Vertical on transmit or receive - usually labelled as co-polar HH or VV, cross-polar HV or VH. The co-polar returns can be simply referred to as H and V. Considerable differences can be seen between H and V polarised data however co-polar returns HV and VH should theoretically give the same returns since they must follow the same paths upon time reversal. The radar is capable of switching polarisations pulse to pulse but was operationally switched every 256 pulses.

2.2.1 Doppler Operation

This is the simplest mode of operation where the radar emits a set of single frequency pulses. By time gating the returns a specific range of observation can be selected. There is the possibility of 'multiple-time-around' echoes which arrive after the second pulse has been emitted. These will generally be of lower intensity due to their originating from beyond the imaged range - dropping off as R^{-4}. The maximum range that can be imaged is $c/2f_p$ which corresponds to 3000m for the highest operable PRF.

If the received signal is assumed to have been scattered from a moving target then, provided the distance moved in time $1/f_p$ is less than a wavelength, the phase change $\Delta\phi$ of the received signal can be determined unambiguously. Assuming a constant single target velocity (narrowband) the rate of change of phase $\partial\phi/\partial t$ is proportional to the target's radial velocity v_{dop}, with the convention that positive v_{dop} corresponds to an inbound target, such that

$$v_{dop} = \frac{\partial\phi}{\partial t}\lambda \equiv \frac{f_{dop}c}{f_t} \tag{2.14}$$

where c is the speed of propagation.

Using a single detector the tangential velocity is unknown. In general there will be more than one scatterer, with a range of velocities (wideband) within the range gate. Frequency analysis of the data, commonly achieved by an FFT upon a number of received pulses, will yield a Doppler spectrum of velocities with an associated frequency width Δf_{dop}.

When analysing the clutter intensity distribution, where the phase (and hence velocity information) is discarded, the data can show considerable correlation due to the component scatters moving little in time $1/f_p$ such that even the Rayleigh speckle is correlated. To obtain independent samples, only pulses separated by approximately the inverse of the Doppler frequency linewidth can be used. The residual correlation, usually determined by the Auto-Correlation Function (ACF), is the expected physical correlation within the scene.

As a moving target will produce an oscillation in the sequentially measured relative phase of the received signal it is common to talk of an associated (Doppler) frequency which simplifies discussion when treating the returns as a complex signal to be analysed. It is perhaps more correctly the *phase Doppler frequency* that arises from a moving scatterer; similarly the Power Spectrum or Power Spectral Density (PSD) of a signal is equivalent to the Doppler velocity spectrum.

2.2.2 High Resolution Operation

Emitting pulses similarly to Doppler operation but using N sequentially stepped transmit pulse frequencies, a linear 'chirp' is formed. Assuming that the range scene has not moved in the time taken to emit N pulses then the wideband frequency response of the range gate is recorded. The transmitted pulse is effectively 'compressed' to subdivide the range gate into N compressed range gates. A range profile, giving both the amplitude and phase of the compressed range gates, is formed by taking the FFT of the recorded chirp pulses. This can be viewed as the 'beating' of N frequencies to form as many nodes defining the compressed range gate spacing. Repeating the process at a rate f_p/N gives identical 'pulse trains', shown in Figure 2.3, allowing a Range-Time-Intensity (RTI) image to display the movement of the scene over time.

The variation in the range profile depends only on the relative phasors present within the received pulses and is circularly symmetric due to the Fourier transform method used. It is possible to obtain a reference phase with which to centre the image with respect to the transmitted pulse by measuring the delay within the system - this is termed a 'loopback frame'. The high intensity central region of the range compressed pulse is now always presented in the middle of the range profile.

Nathanson (1991) covers the mathematical proof and shows that if a rectangular pulse is utilised $\tau\Delta f_s = 1$ is desirable. Forming a range gate of extent $c\tau/2$; by emitting N pulses over bandwidth Δf a compressed resolution of $c/2N\Delta f_s$ is realised. If $\tau\Delta f_s = 1$ the range gate is resolved into N subdivisions giving no redundancy to the output of the FFT. In general the MBPR does not

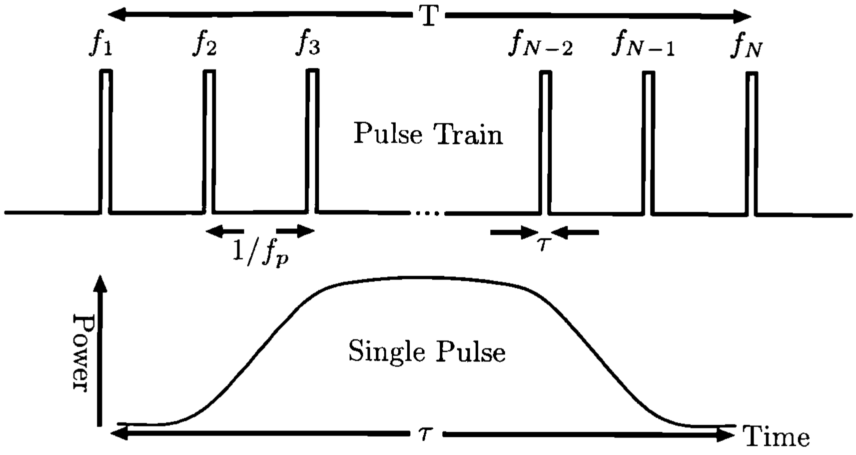

Figure 2.3: The stepped frequency pulse train used to form range profiles

maintain this constraint, usually $\tau \Delta f_s \leq 1$. This is primarily to allow a wide window to be applied to the data prior to the FFT to prevent aliasing in range, and also to account for the pulse shape not being rectangular.

In analysis of the high resolution intensity distribution, decorrelation of the data is not so critical. After the FFT the range gates are separated in time by N/f_p which is enough to decorrelate the speckle. If the operation has been carried out correctly, each point on the resultant range profile will have a physically meaningful phase. Using consecutive range profiles it is now possible to build up a Doppler spectrum for a particular compressed range gate.

2.2.3 Improving Accuracy and Reducing Noise

To improve the data quality obtained by the MBPR, processing is done before and after pulse recording. To reduce the effect of system noise, 6× oversampling is made on receive by a 12 bit A/D converter. After filtering at 10MHz and combining signals 15m apart this gives a simulated 16 bit output with reduced noise. The final pulseshape has a plateau ±37.5m and a curved fall off to the 3dB point at ±75m reaching zero at ±112.5m (Branson 1999).

In addition, the operation of an FFT to obtain either the Doppler spectrum or the range profile must be modified to account for windowing effects. This is a standard signal processing technique necessary with a finite data set of length N - corresponding to a rectangularly windowed infinite set. After an FFT the transform of the window is convolved over every FFT output, but the transform of a rectangular window w has the undesirable property of very large oscillations beyond

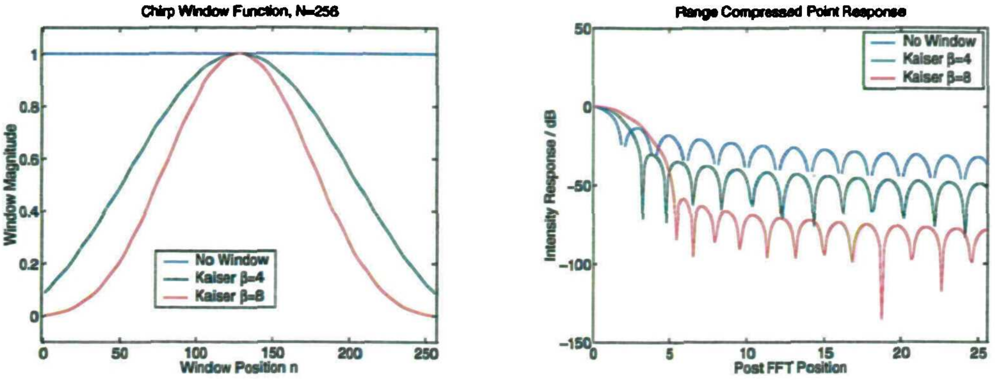

Figure 2.4: A Kaiser window lowers sidelobe levels at the expense of the mainlobe point response

the mainlobe called sidelobes.

$$w_{Rect}(n) = 1,\ 0 < n \leq N \tag{2.15}$$

Sidelobes mean that a high valued FFT output bin can raise the other output bins even though no frequency is present there. For example in Doppler mode a high amplitude low velocity target can still raise the value of the high velocity FFT output giving an inaccurate Doppler spectrum; in Hi-res mode a single high intensity compressed range gate would cause the sidelobe structure to be visible in neighbouring range gates.

The level and shape of the sidelobes can be controlled by choosing to further window the data. The total intensity of the resultant FFT must remain constant and so reducing the sidelobe levels must broaden the mainlobe. The Kaiser window was chosen, used in both DERA software (ARCANE) and in a report by TWR (TW Research Ltd. 1999). The Kaiser window is an approximation to the prolate spheroidal window for which the ratio of the mainlobe to sidelobe energy is maximised

$$w_{Kaiser}(n) = \frac{I_0\left[\beta\sqrt{1-\frac{4(n+1/2)^2}{(N-1)^2}}\right]}{I_0(\beta)},\ 0 < n \leq N \tag{2.16}$$

where β is a parameter determining sidelobes, N is the size of the weighting array and I_0 is a Bessel function. The windows and point response functions of the rectangular, Kaiser($\beta = 4$) and Kaiser($\beta = 8$) are shown in Figure 2.4 which give -13, -30 and -60dB sidelobe levels respectively.

2.3 Summary

Theory has been presented to give the expected returns from a general radar system showing the returns to depend considerably upon the environment. However there are a many sources of variation in propagation, target aspect ratio, relative motion and overall scene movement which cannot be easily accounted for. The mathematical tools necessary to analyse these variations are the concern of the next Chapter.

Chapter 3

Mathematical Aspects

3.1 Intensity Probability Distributions

The increased probability of high intensity returns obtained in comparison to the expected exponential distribution described in Section 2.1.1 has prompted various other RCS distributions to be suggested.

Parameter estimation of an expected probability distribution is usually made by a Maximum Likelihood Estimate (MLE) which gives the estimate of the m parameter vector θ^m that maximises the joint probability of the length n Independent Identically Distributed (IID) vector $x_{1:n}$ occurring. Thus

$$\widehat{\theta^m} = \max \left\{ \prod_{i=1}^{n} P\left(x_i; \theta^m\right) \right\} \tag{3.1}$$

The distributions considered here are two parameter distributions, defined by a *scale* and *shape* parameter. The scale estimate is largely a function of the first moment of the distribution and can be determined relatively accurately from the sample mean. The shape cannot be determined independently since it inherently measures a deviation from the mean. A moment based estimate where shape is some function of variance or second moment

$$Var\left[x\right] = E\left[x^2\right] - \left(E\left[x\right]\right)^2 \tag{3.2}$$

shows that the relative uncertainty must be at least as great as that in the mean - the problem in the use of higher order statistics for finite sample sizes.

3.1.1 The Weibull Distribution

The Weibull distribution, with intensity I is a two parameter distribution dependent upon a shape a and a scale parameter σ (defined independently of RCS).

$$P_{Weibull}(I) = \frac{a}{2\sigma^2} I^{a-1} \exp\left[-\frac{I^a}{2\sigma^2}\right], \; 0 \leq I < \infty \, ; a > 0 \tag{3.3}$$

First used to model breaking strengths of materials (Weibull 1939) by setting the hazard rate H to vary as a power of the probability variable; it was first empirically applied to sea clutter data by Fay et al. (1977), where

$$H(t) \equiv \frac{P(t)}{\int_t^\infty P(t)\, dt} \tag{3.4}$$

thus

$$H_I(I) = \frac{\frac{a}{2\sigma^2} I^{a-1} \exp\left[-\frac{I^a}{2\sigma^2}\right]}{\exp\left[-\frac{I^a}{2\sigma^2}\right]} = \frac{a}{2\sigma^2} I^{a-1} \tag{3.5}$$

Equation 3.3 shows the Weibull reduces to the Exponential distribution when $a = 1$, the Rayleigh distribution is obtained when $a = 2$. A Rayleigh distributed amplitude is exponentially distributed in power, a result of the Weibull distribution's invariance under an exponent transformation.

The Weibull MLE estimate is in closed form (Harter and Moore 1965, Cohen 1965), more recently a thorough analysis was made by Oliver (1993) who gave the iterative solutions as

$$\widehat{\sigma^2} = \frac{\widehat{I^a}}{2} \tag{3.6}$$

$$\frac{-2}{\hat{a}} + \frac{\log \widehat{I^{\hat{a}}}}{\widehat{I^{\hat{a}}}} = 2\widehat{\log I} \tag{3.7}$$

Sekine and Mao (1990) state from a review of several papers that a is expected to lie in the region $0.5 - 2$ for a variety of land or sea conditions. A non-iterative procedure is demonstrated by Oliver (1993) in a normalised log estimator which gives similar performance to the MLE solution for $a \geq 0.5$ with acceptable performance down to $a = 0.1$. This method is covered in the Section concerning the K-distribution.

3.1.2 Log Normal

The log normal distribution is a two parameter distribution dependent upon the mean μ and variance σ^2 of a log transformed Gaussian distribution to yield

$$P_{LogN}\left(I\right) = \frac{1}{I\sigma\sqrt{2\pi}}\exp\left[\frac{-\log\left(I-\mu\right)^2}{2\sigma^2}\right] \tag{3.8}$$

Applied by Trunk and George(1970) to model sea clutter it is still often used to provide a comparative fit to data (Farina et al. 1997). The parameter estimates are usually made by transformation to the normal distribution and using the sample values for μ and σ^2.

3.1.3 The K Distribution

The K Distribution, with amplitude x is a two parameter distribution dependent upon a shape ν and a scale b

$$P_K\left(|x|\right) = \frac{4b^{(\nu+1)/2}}{\Gamma\left(\nu\right)}x^{\nu}K_{\nu-1}\left(2x\sqrt{b}\right)\ ,\ 0\leq|x|<\infty\,; b,\nu>0 \tag{3.9}$$

where K_n is the modified Bessel function of order n. A compound form exists where the underlying cross section A varies according to a gamma distribution. A large number of scatterers are assumed to be present in each resolution cell so the gamma distribution modulates the Rayleigh speckle.

$$\begin{aligned} P_{Rayleigh}\left(x|A\right) &= \frac{2x}{A}\exp\left(-\frac{x^2}{A}\right)\ ,\ 0\leq x\leq\infty & (3.10)\\ P_{Gamma}\left(A\right) &= \frac{b^{\nu}}{\Gamma\left(\nu\right)}A^{\nu-1}\exp\left(-bA\right)\ ,\ 0\leq A\leq\infty; b,\nu>0 & (3.11)\end{aligned}$$

It was initially applied in the field of lasers (Jakeman and Pusey 1976) and applied to sea clutter by Ward (1981). If the gamma distributed cross section is assumed to vary slowly in comparison to the speckle, this gives some physical justification for the K distribution. Both IID and compound realisations of a K($\nu = 1$) variate are shown in Figure 3.1.

The K distribution does not have an analytic MLE, which has caused several estimators to be proposed. Using the intensity domain I, Blacknell (1994) compares the numerical solution of the MLE with three schemes:

1. $E\left(I^{(1)}; I^{(2)}\right)$ - first and second order moments, the *contrast estimate.*

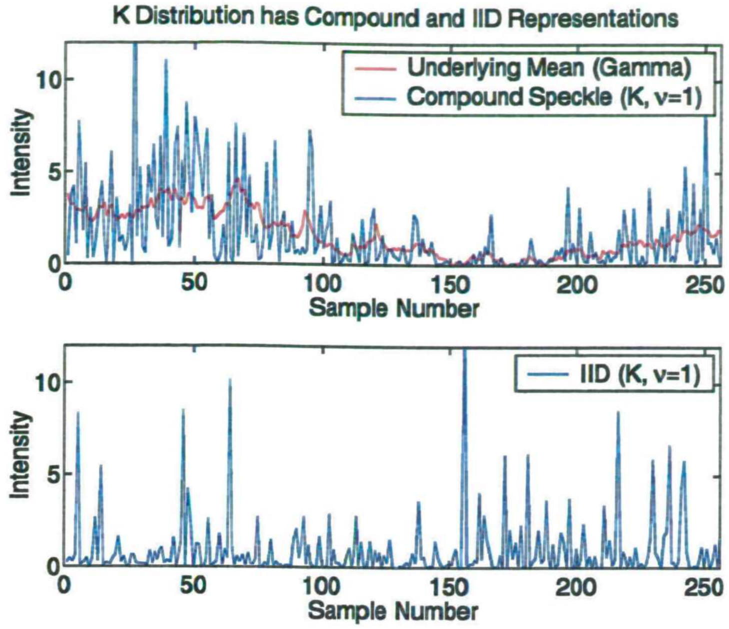

Figure 3.1: The K distribution has some physical justification as a modulating gamma variate subject to Rayleigh speckle

2. $E\left(\log(I)^{(1)}; \log(I)^{(2)}\right)$ - first and second order moments of the log data, the *normalised log estimate.*

3. $E\left(\bar{I}; \overline{\log(I)}\right)$ - sample mean and sample mean of the log data, the *variance of the log estimate.*

Parameter fitting for $E\left(I^{(1)}; I^{(2)}\right)$

Using the first two moments of the K distribution and solving for ν gives the simplest estimate available

$$V = \frac{\widehat{I^2}}{\widehat{I}^2} - 1 = 1 + \frac{2}{\widehat{\nu_V}} \tag{3.12}$$

The direct solution of the above has an associated bias and variance (Oliver 1993) to order $1/n$

$$\frac{\langle \Delta\nu \rangle_V}{\nu} \approx \frac{\nu}{n}\left(1+\frac{1}{\nu}\right)\left(1+\frac{6}{\nu}\right) \tag{3.13}$$

$$\frac{\langle \sigma_\nu^2 \rangle_V}{\nu^2} \approx \frac{\nu^2}{n}\left(1+\frac{1}{\nu}\right)\left(1+\frac{4}{\nu}\right)\left(1+\frac{5}{\nu}\right) \tag{3.14}$$

Equation 3.12 is commonly termed the contrast estimate.

Parameter fitting for $E\left(\bar{I};\overline{\log(I)}\right)$

The derivation of $E\left(\bar{I};\overline{\log(I)}\right)$ by Oliver (1993) uses the Mellin transform as an approximation in the limit of many looks to derive an approximation to the MLE,

$$U=\widehat{\log(I)}-\log\left(\hat{I}\right)=\psi^{(0)}\left(\widehat{\nu_U}\right)-\log\left(\widehat{\nu_U}\right)-\gamma \tag{3.15}$$

where $\gamma \approx 0.577$ is the Euler-Gamma number and $\psi^{(0)}(\bullet)$ is the polygamma function of order 0. The iterative solution of the above has an associated bias and variance (Oliver 1993) to order $1/n$

$$\frac{\langle\Delta\nu\rangle_U}{\nu} \approx \frac{\left[\frac{\langle I^2\rangle}{\langle I\rangle^2}-1\right]}{2n\left(1-\nu\psi^{(1)}(\nu)\right)} \tag{3.16}$$

$$\frac{\left\langle\sigma_\nu^2\right\rangle_U}{\nu^2} \approx \frac{\frac{1}{n-1}\left(\frac{\langle I^2\rangle}{\langle I\rangle^2}-\frac{2\langle I\log(I)\rangle}{\langle I\rangle}+\left\langle(\log(I))^2\right\rangle-\langle\log(I)\rangle^2+1\right)}{\left(1-\nu\psi^{(1)}(\nu)\right)^2} \tag{3.17}$$

Equation 3.15 is commonly termed the normalised log estimate.

Parameter fitting for $E\left(\log(I)^{(1)};\log(I)^{(2)}\right)$

The compound K distribution arises from a product of two components; this motivates taking the logarithm to form additive components to give an estimate

$$W\equiv\widehat{\log^2 I}-\widehat{\log I}^2=\psi^{(1)}\left(\widehat{\nu_W}\right)+\frac{\pi^2}{6} \tag{3.18}$$

The iterative solution of the above has an associated bias and variance (Oliver 1993) to order $1/n$

$$\frac{\langle\Delta\nu\rangle_W}{\nu} \approx \frac{\psi^{(1)}(\nu)+\frac{\pi^2}{6}}{n\nu\psi^{(2)}(\nu)} \tag{3.19}$$

$$\frac{\left\langle\sigma_\nu^2\right\rangle_W}{\nu^2} \approx \frac{\psi^{(3)}(\nu)+2\psi^{(1)2}(\nu)+\frac{11\pi^4}{90}+\frac{2\pi^2}{3}\psi^{(1)}(\nu)}{n\left(\nu\psi^{(2)}(\nu)\right)^2} \tag{3.20}$$

Equation 3.18 is commonly termed the variance of log estimate.

Inversion Problems

Whilst the variance of the estimators leads to the variance in ν, some values of the estimator cannot be inverted. As an example Figure 3.2 demonstrates hard limiting in the contrast estimate for $V < 1$ where ν_V takes on an unphysical negative value. Due to minimisation routines used

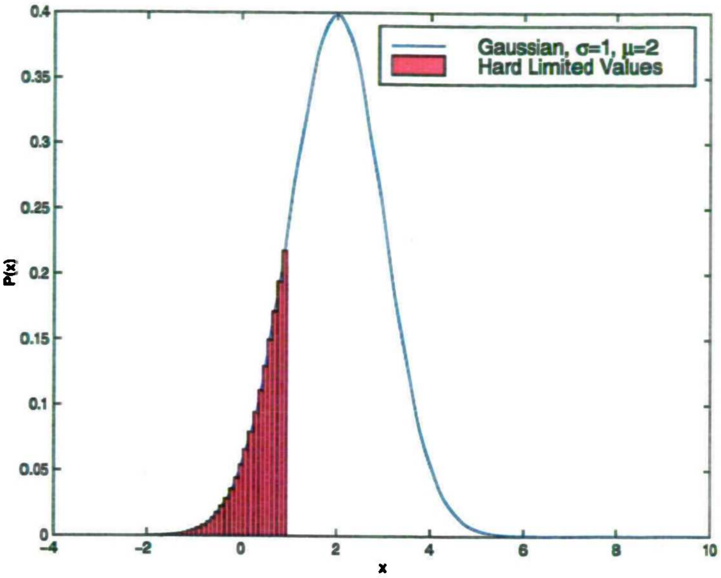

Figure 3.2: An inversion limit exists for each K distribution estimator

for U and W the exact hard limited ν values are determined by machine accuracy but a similar case arises. Discarding hard limited values will obviously affect the bias and variance of the estimators; possibly more importantly a discontinuity will be seen in the sample distribution of the shape parameter ν. It seems unnecessary to calculate this effect probabilistically as it is easier to record how many values were discarded and then directly handle the effect.

To gain an estimate of the probability of inversion failing, the estimator E is modelled (Oliver 1993) as a Gaussian distribution based upon the bias ΔE and variance σ_E^2 given above. Tests of this hypothesis were made and show that it is not completely accurate (see later Sections), but applicable for reasonably large sample populations.

$$P(E)\partial E \simeq \frac{1}{\sigma_E\sqrt{2\pi}}\exp\left[\frac{-\{E-(\langle E\rangle+\Delta E)\}^2}{2\sigma_E^2}\right]\partial E\,,\ -\infty < E < \infty \tag{3.21}$$

enables numeric calculation through a change of variable noting the conditions for inversion:

$$U+\Delta U \leq -\gamma \tag{3.22}$$

$$V+\Delta V \geq 1 \tag{3.23}$$

$$W+\Delta W \geq \pi^2/6 \tag{3.24}$$

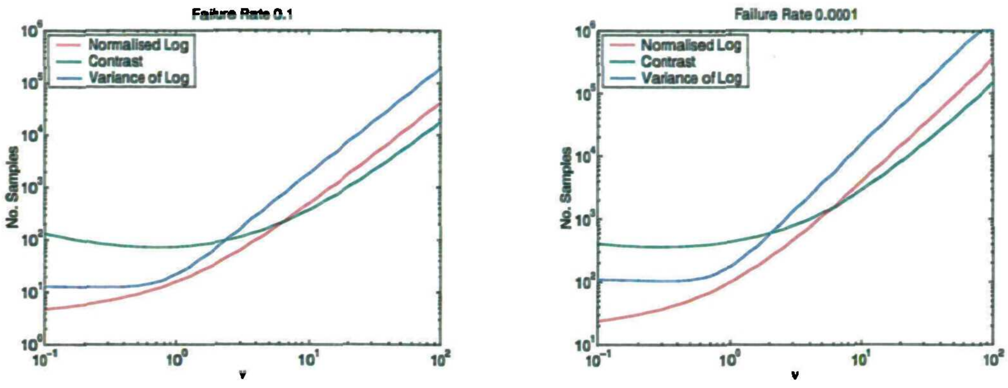

Figure 3.3: Expected minimum samples necessary to sucessfully invert contrast estimators

Figure 3.3 shows the expected minimum samples necessary to maintain a failure rate of 0.1 and 0.0001 (a similar graph is found in Lombardo and Oliver (1994)), the variance of log estimator W is inferior over the considered range of ν. For low values of ν the normalised log estimate U is superior but an upper limit exists where the contrast V achieves the failure rate with a lower sample population.

Best Performance Under Ideal Conditions

From analysis of a large amount of recorded data, Ward et al. (1990) published a graph suggesting that ν can vary between 0.1 and 10 with a modal value of about 1; under ideal conditions the choice of estimator must be based on its accuracy in determining the correct shape parameter over this range. Several papers have investigated this and confirm that for typical conditions the normalised log estimator U is the best overall choice as it approaches the numerical MLE solution for ν. Blacknell (1994) found that for $N = 256$ samples U performed best in the region $0.1 < \nu \lesssim 8$ but interestingly V performed better for higher order parameters, the relevant figure from this paper is reproduced in Figure 3.4 and it is obvious that the contrast estimator V is very poor overall. This range of performance was similarly observed in the conditions for inversion in the previous Section; this is expected since both criteria are largely determined by the sample variance in the estimator - U usually having the lowest sample variance.

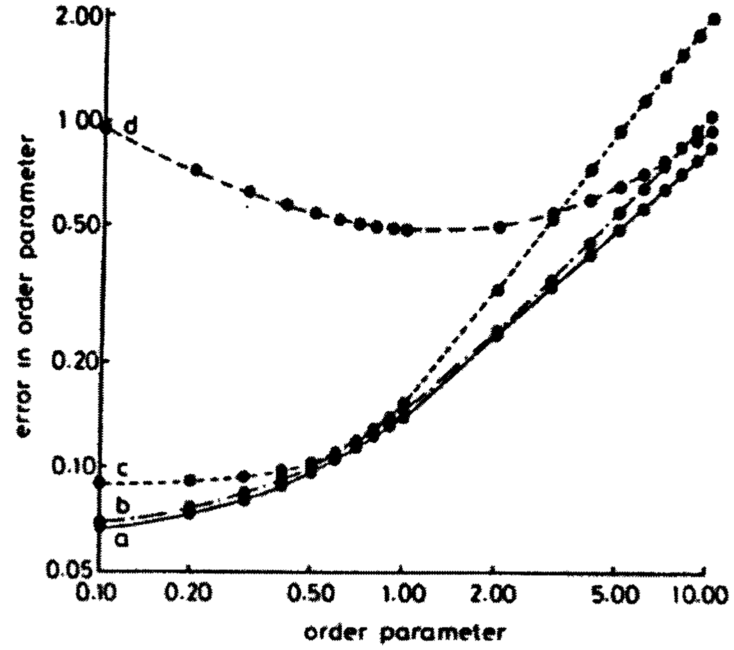

Figure 3.4: Error in ν order parameter estimate for 256 samples. a - Maximum Likelihood Estimate, b - U estimate, c - W estimate, d - V estimate (after Blacknell, 1994)

3.1.4 Deviation from the Ideal

In this subsection the effects of non-ideal observed data are considered. Problems may arise from (but not limited to) additive thermal noise, small sample sizes, censoring and inhomogeneity. No further consideration is made of V due to its poor performance and, although not shown, it is extremely sensitive to censoring with a long tailed distribution over some of the region of interest.

Censored Samples

In some target detection schemes it is necessary to censor the higher valued data, under the hypothesis that this is a target, and then form a decision based upon the remaining data assuming this to be background clutter. The effect of progressive censoring is tested upon the estimators U and W based upon an initial sample size of 256 with ten thousand repetitions. A histogram is calculated for U and W for the uncensored case and then the highest sample is removed and the tests re-applied. This was repeated until the largest 16 samples were censored from the population or equivalently 'censored at the 6% level'.

The plots shown in Figure 3.5 cover $\nu = 0.1$, $\nu = 1$ and $\nu = 10$ and are plotted with their natural abscissa (U or W - not converted to a shape parameter such as ν) this allows the effect of censorship to be compared with respect to the estimator's intrinsic variance. These plots have been overlaid with a Gaussian fit and show:

- The sample distribution of U is very close to Gaussian regardless of censorship.
- W deviates from Gaussian at low values of ν towards a longer tailed distribution.
- In comparison to their sample variance, especially at low ν, W is less sensitive to censorship.

This provides some justification for using the estimator W in a censored situation but the stable sample distribution of U is extremely useful for accurately calculating performance figures.

Limited Sample Sizes

The previous Section showed that U has a more stable sample distribution with respect to shape parameter alone. When comparing shape estimates from *different* sample sizes the resultant distribution change must be known. Additionally censoring is performed at the 3% and 6% level in each case and sample populations of $N = 32$ and 64 were chosen at a shape parameter of $\nu = 0.1$ shown in Figure 3.6. Again the plots have been overlaid with a Gaussian fit and show

- The distribution of W is far from Gaussian at low sample numbers and low ν - moving towards a long tailed distribution.
- U conforms to a Gaussian distribution even under censorship at low sample numbers and low ν.

Obtaining confidence levels for the estimator U is greatly simplified as the sample distribution is stable Gaussian; the instability towards a longer tail in the distribution of W presents difficulty.

Presence of Noise

The radar receiver will in general add a component of thermal noise to any measured signal. The exact amount of additive noise is system dependent and could in theory be calculated, but a significant component could also arise from the environment itself to contribute and overall SNR. A

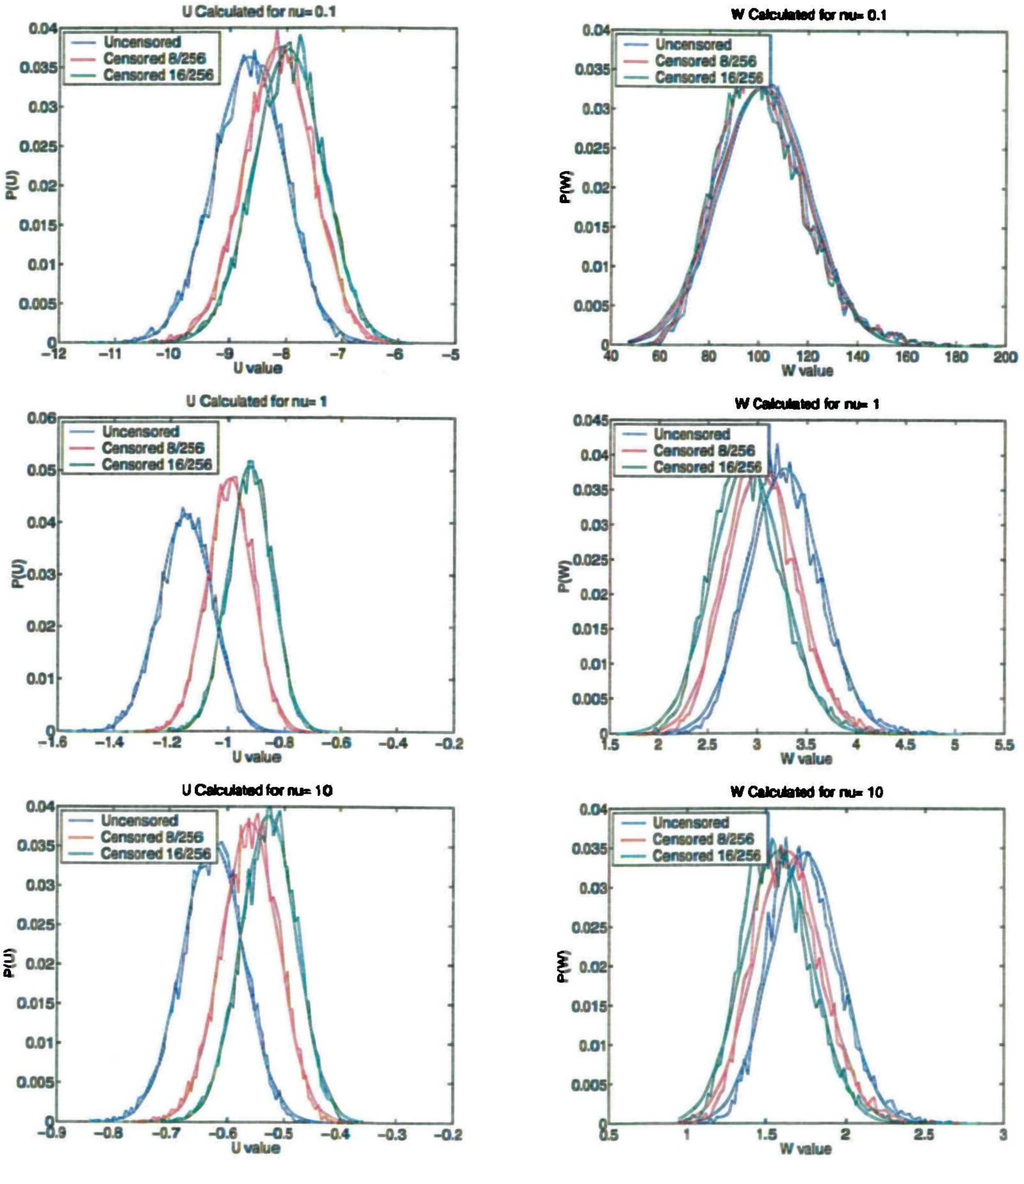

Figure 3.5: Sample distribution of the U and W estimates for varying shape parameter and censorship ($N = 256$)

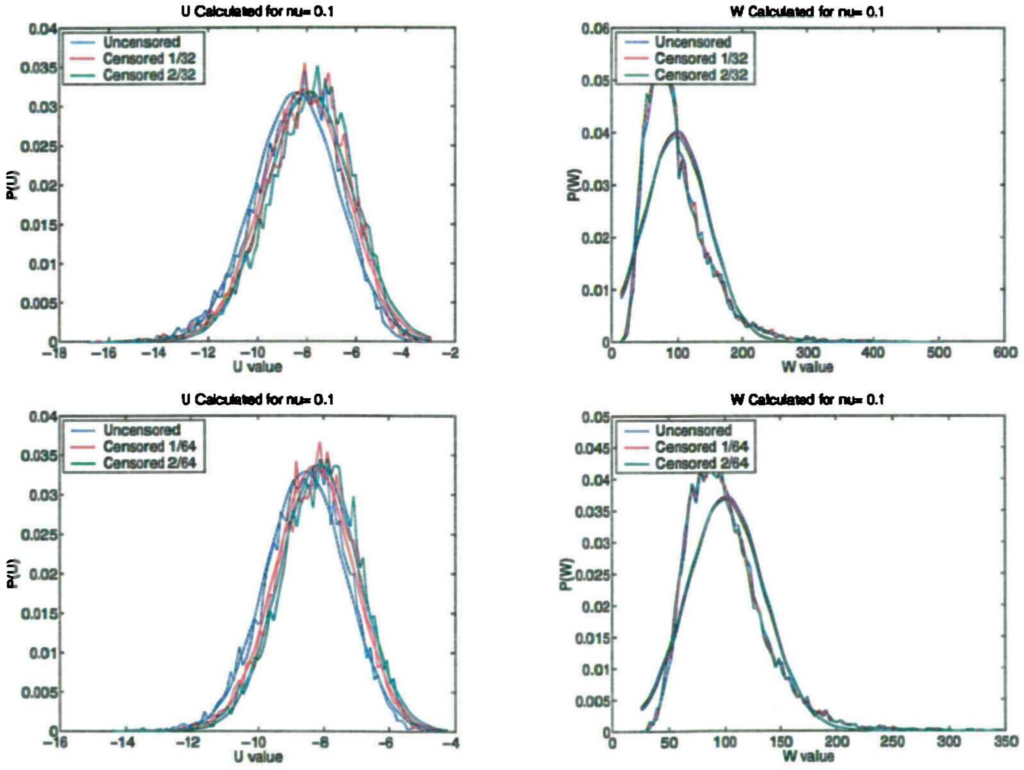

Figure 3.6: Sample distribution of the U and W estimates for varying sample number at $\nu = 0.1$

traditional view of scattering is that one mechanism is responsible for the observed statistics; however it may be that the emitted pulse samples a mixture of scattering sources - of which a component originates from multiplicative Rayleigh noise (Section 2.1.1). A Rayleigh scattered component is statistically indistinguishable from thermal noise unless discrimination is based upon:

- Forming the Doppler spectrum of the signal, thermal noise is white (it has a flat frequency spectrum). It can be isolated from Rayleigh scattering with a presumed physical origin confined within the observed Doppler spectrum.
- Observing the long term high resolution range profile, the mean noise level is determined by measuring the mean of the recorded signal outside the pulse shape. This relies on thermal noise being white.

Both the above assume well formed data without aliasing from the Fourier transform operation. In principal the magnitude of an assumed Rayleigh scattering component can be calculated even when combined with another non-Gaussian distribution. This is made through *cumulant analysis* which is dependent on the noise being completely described by its variance (Swami et al. 1999); the probability distribution of the non-Gaussian component distribution must be known to within a scale factor.

The distribution of U when contaminated by noise was investigated by Lombardo et al. (1995) who tested for SNR of 5, 10, 20 and 30dB. The effect is to add a significant bias to the determined U value with slight change to the variance; however the sample distribution remains suitably Gaussian (Appendix B.2).

Inhomogeneity

In this Section the effect upon the parameter estimator when an edge is present in otherwise homogenous statistics is derived, this has not been discussed in the literature.

An edge is defined to be a discontinuity from a single abrupt change in:

1. Intensity alone.
2. Distribution shape alone.
3. A simultaneous change of both intensity and distribution shape.

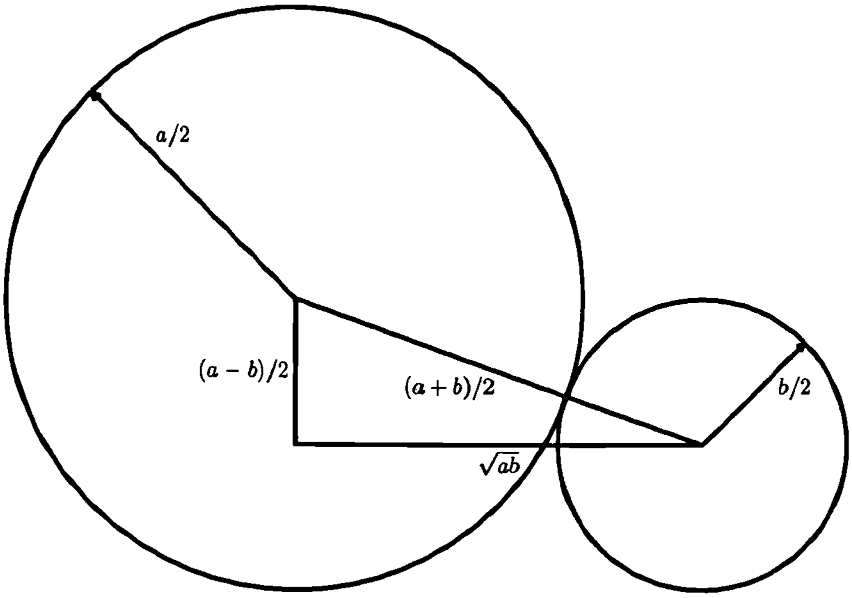

Figure 3.7: The texture estimator $U \leq 0$ from consideration of the Pythagorean triangle as $U = \log(\mu_G/\mu_A)$

The normalised log estimate U is used to determine the edge effects as it is the most accurate of the considered estimators. A data series of length N has a shape parameter U defined as

$$\begin{aligned} U &= \widehat{\log I} - \log \hat{I} && (3.25) \\ &\equiv \log\left(\prod_{i=1}^{N} I_i\right)^{1/N} - \log\left(\frac{1}{N}\sum_{j=1}^{N} I_j\right) && (3.26) \\ &= \log \mu_G - \log \mu_A && (3.27) \end{aligned}$$

where μ_G is the geometric mean and μ_A is the conventional arithmetic mean. The well known arithmetic-geometric means inequality, demonstrated in Figure 3.7, proves that $U \leq 0$ by using the two data values a and b with means μ_G and μ_A applied using Pythagoras' Theorem - the proof is extended to N data by induction.

The task is to determine $\langle U_\Delta \rangle$, the expected change in U arising from a step in otherwise homogenous statistics.

Intensity Change An intensity edge is present from point $k \leq N$ of relative scale s with no change in the distribution shape defined as

$$\begin{aligned}\langle I_{1:k-1}\rangle &= \frac{1}{s}\langle I_{k:N}\rangle && (3.28)\\ \langle U[I_{1:k-1}]\rangle &= \langle U[I_{k:N}]\rangle && (3.29)\end{aligned}$$

where I refers to the intensity of individual samples, and $\langle\cdot\rangle$ is the expectation over an ensemble average. Explicitly including this scale change, with the notation changed to reflect an assumed underlying IID variable X, the means are formed as

$$\begin{aligned}\mu_G(s,k) &= (X_1 \times X_2 \times \ldots \times sX_k \times sX_{k+1} \times \ldots \times sX_N)^{1/N} && (3.30)\\ \mu_A(s,k) &= \frac{(X_1 + X_2 + \ldots + sX_k + sX_{k+1} + \ldots + sX_N)}{N}\end{aligned}$$

causing the expectation of both μ_G and μ_A to change by

$$\begin{aligned}\Delta\mu_G &\equiv \mu_G(s,k) - \mu_G = \langle X\rangle s^{k/N} && (3.31)\\ \Delta\mu_A &\equiv \mu_A(s,k) - \mu_A = \langle X\rangle \frac{sk-k}{N}\end{aligned}$$

Forming the normalised log estimator U and redefining the quantities such that p refers to the proportion k/N over which a relative intensity scale change of s gives the change $U_{\Delta I}$,

$$\langle U_{\Delta I}\rangle = \log\frac{s^p}{sp-p+1} \qquad (3.32)$$

this is plotted in Figure 3.8 and indicates:

- The effect of an intensity scale change, *regardless of sign*, is to decrease measured U - this is intuitive but can be proven from an extension of Figure 3.7 where equality is expected iff $a = b$.
- A point-like target (low p) will have little effect on the assessment of U even at a relatively high dB - thus 'self masking' is not a strong effect in shape estimation.

Distribution Shape Change A distribution edge is present from point $k \leq N$ such that the distribution changes from an underlying U_0 to one of shape U_T with no change in the mean

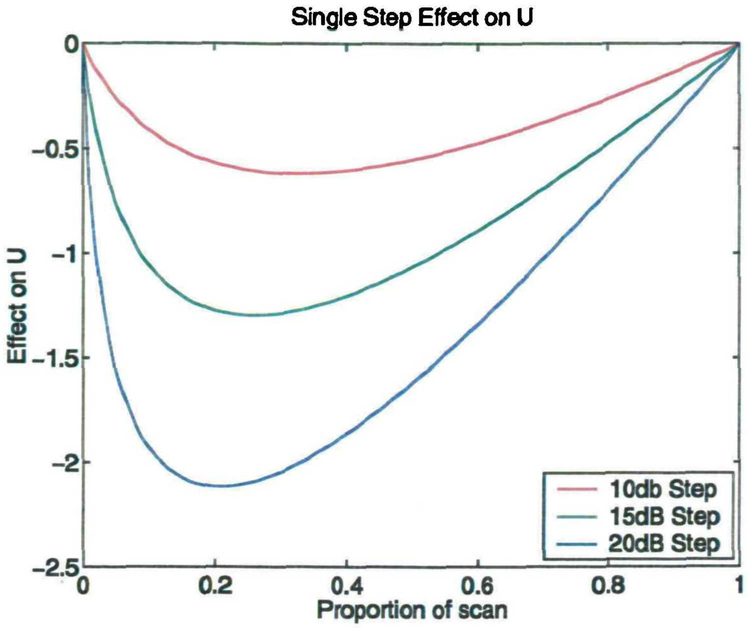

Figure 3.8: The effect an intensity step has upon determined normalised log estimator U

$$\begin{aligned}\langle U\left[I_{1:k-1}\right]\rangle &= U_0 && (3.33)\\ \langle U\left[I_{k:N}\right]\rangle &= U_T && (3.34)\\ \langle I_{1:k-1}\rangle &= \langle I_{k:N}\rangle && (3.35)\end{aligned}$$

With no change in the arithmetic mean, the change in U will be from the geometric mean change alone and the measured U will be

$$U = \frac{kU_T + (N-k)U_0}{N} \tag{3.36}$$

upon redefining the quantities as before, where p refers to the proportion k/N over which a distribution shape change from U_0 to U_T, gives the change $U_{\Delta U}$

$$\langle U_{\Delta U}\rangle = p\left(U_T - U_0\right) \tag{3.37}$$

- The effect of a distribution shape change is a simple linear dependency upon the contaminating edge distribution.

Simultaneous Change in Intensity and Distribution Shape In reality a change in conditions would probably cause a discontinuity in both mean and distribution shape, in summary:

- The magnitude of $U_{\Delta I}$ is independent of an assumed underlying distribution and the number of samples.
- The magnitude of $U_{\Delta U}$ is proportional to the change in U_0 at the step.
- Recasting in terms of ν is less tractable and will cause ν_Δ to be a function of ν.

The changes can be combined to give the overall effect on a measured U such that

$$\langle U_{\Delta I,\Delta U}\rangle = \log \frac{s^p}{sp - p + 1} + p\left(U_T - U_0\right) \tag{3.38}$$

gives the expected change in a measured U from an underlying homogenous distribution of shape U_0 subject to a simultaneous edge over a proportion p which consists of a distribution U_T of relative scale s.

- If $U_T > U_0$ the effect upon measured U will be lessened and for suitable values a simultaneous step change in mean and distribution will leave measured U unaffected.

Modulation of Intensity By Another Distribution The K-distribution has been physically justified by casting it in a compound form (Ward 1981) in Section 3.1.3. This links the exponential speckle expected from the imaging process to the physical modulation from ocean waves. To generalise this compound form to other distributions, define:

- A local intensity 'speckle' distribution $f_0\left(\mu\right)$ with zero correlation between samples defined by a local mean $\widetilde{\mu}$, with an individually measured U_0.
- A global modulating distribution f_ρ which has some correlation over a number of samples ρ, defining $\widetilde{\mu}$, with an individually measured U_ρ.
- The resultant global distribution measured over $N \gg \rho$ samples $f_g = f_0\left(\widetilde{\mu}\right)$, giving measure U_g.

The expected value of U_g is determined from the means

$$\begin{aligned} \langle \mu_G [f_g] \rangle &= \mu_G [f_0] \times \mu_G [f_\rho] & (3.39) \\ \langle \mu_A [f_g] \rangle &= \mu_A [f_0] \times \mu_A [f_\rho] & (3.40) \end{aligned}$$

such that

$$\begin{aligned} \langle U_g \rangle &= \left\langle \log \left(\frac{\mu_G [f_g]}{\mu_A [f_g]} \right) \right\rangle = \left\langle \log \left(\frac{\mu_G [f_0] \times \mu_G [f_\rho]}{\mu_A [f_0] \times \mu_A [f_\rho]} \right) \right\rangle & (3.41) \\ &= U_0 + U_\rho & (3.42) \end{aligned}$$

thus the effect of a modulation upon the local measure U_0 causes a bias of U_{Δ_ρ}

$$\langle U_{\Delta\rho} \rangle = U_\rho \tag{3.43}$$

where the expectation is taken over an ensemble of sample numbers $N \gg \rho$.

- A modulation can only serve to decrease measured U; it is the continuous analogue of a large number of abrupt intensity changes.

3.1.5 Measuring Distribution Parameters

In the literature there is no strong consensus on the distribution to which radar clutter conforms. Weibull and K are popular choices but each has its own texture parameter with associated estimators. As shown, the K distribution can be characterised well by the normalised log estimate in the expected region of interest; the previous tests were repeated using Weibull statistics and showed similar effects and so it is reasonable to assess U only. This can be justified by:

- The arithmetic and geometric means are simple to calculate and therefore very fast in operation.

- The normalised log estimate is near optimal for Weibull and K in the parameter regions and sample sizes of interest; if the true distribution is between the two then it will be applicable also.

- It is only for convenience that the one-to-one mapping is performed from any texture estimator to the shape parameter used in the distribution. A simple conversion graph can be used when this is necessary.

- The mapping may not be invertible as shown in Section 3.1.3. This causes a discontinuity in the sample distribution of the determined shape parameter.

- This mapping is non-linear and can be shown to cause significant errors in estimating the bias and error in the determined shape parameter. This is explored by Lombardo and Oliver (1994) and an example is given in Appendix B.1.

- Uniquely the distribution of U is suitably Gaussian for varying sample sizes, under censoring, at low sample numbers and in the presence of thermal noise. This allows accurate prediction of associated uncertainty in the estimate and is crucial in later sections.

- The expected effect upon U of any step or modulation in the data is explicitly calculable, does not depend on the sample size and is derived to be simple equations (3.32, 3.37 & 3.43). This gives a tool for extending the analysis to non-stationary data - covered in subsequent Chapters.

Appendix B.2 discusses a further estimator - the amplitude contrast V_a proposed by Lombardo, Oliver and Tough (1995) which is designed to offer resistance to thermal noise whilst maintaining the accuracy of determining the shape parameter. If the K-distribution represents the absolute underlying statistics, Lombardo et al. showed that for estimation of ν from clutter data contaminated by noise, V_a would be a good overall choice. Several points are suggested concerning this:

- It is arguable whether noise-free statistics from the surface are ever observed; if so, the commonly applied K or Weibull distribution may be an accurate approximation to the distribution of 'clutter plus noise'.

- As mentioned previously, some of the Rayleigh noise could be from the surface itself and indistinguishable from the thermal noise; sensitivity to this noise would be desirable for improved segmentation of scenes.

- The comparison between estimators is made assuming a linear dependence between the estimator and the reciprocal of the shape parameter ν. According to Lombardo and Oliver (1994) whilst this is true for all V, only in the limit of large ν can it be applied for U and W. As shown above, the distribution of the estimators is not identical with the result that discrimination on the basis of minor changes in the relative biases and variances is complicated.

- The sample distribution of V_a is not normal for low sample numbers and low ν as shown in Appendix B.2.

The decision is made to retain U as the preferred estimator; its consistent Gaussian distribution is a strong reason as the method of determining an associated confidence in the value will be more robust and this is relied upon in Chapter 8.

3.1.6 Additional Properties of U

It remains only to relate the chosen estimate U to the shape parameter used in the distributions. Equation 3.15 obtains the K shape parameter ν, equivalent Weibull shape parameter a is determined in a similar manner (Oliver 1993).

$$U(\nu) = \psi^{(0)}\left(\widehat{\nu_U}\right) - \log\left(\widehat{\nu_U}\right) - \gamma \tag{3.44}$$

$$U(a) = \frac{-\gamma}{a} - \log\left(\Gamma\left(1 + \frac{1}{a}\right)\right) \tag{3.45}$$

Equations 3.44 and 3.45 are plotted in Figure 3.9. It is possible to define a Weibull distribution that equates to $U > -\gamma$ in the case of the Weibull distribution; this is because the limiting distribution of $a \to \infty$ is not exponential as in the K-distribution. Realistically it is expected that the shape parameter $a \leq 1$.

For $\nu = a = 0.5$ the Weibull and K distributions coincide ($U \approx -1.85$) as the Bessel function can be simplified for half-integer orders (Armstrong 1992); thus any practical fit to real data near this parameter will be unable to distinguish between a Weibull and K distribution.

If locally stationary K distributed statistics are assumed, the primary cause of uncertainty in U will be finite sample size. From Section 3.1.4 the sample distribution of U is known to be suitably Gaussian

$$f(U) = N[\mu, \sigma^2] \equiv \frac{1}{\sigma\sqrt{2\pi}} \exp\left[\frac{-(U-\mu)^2}{2\sigma^2}\right]$$

and using the results of Lombardo and Oliver (1994) who derive the bias and variance in terms of ν as

$$\mu = \frac{\left[\frac{1}{2N}\left(1+\frac{1}{\nu}\right) + \frac{1}{2N^2}\left(1-\frac{1}{\nu}-\frac{1}{\nu^2}\right)\right]}{\left[\psi^{(0)}(\nu) - \ln(\nu) - \gamma\right]} + U \tag{3.46}$$

$$\sigma^2 = \frac{\left[\frac{1}{N}\left(\psi^{(1)}(\nu) + \frac{\pi^2}{6} - 1\right) + \frac{1}{2N^2}\left(1 - 2\left(1+\frac{1}{\nu}\right)\left(1+\frac{4}{\nu}-\frac{2}{\nu+1}\right)\right)\right]}{\left[\psi^{(0)}(\nu) - \ln(\nu) - \gamma\right]^2} \tag{3.47}$$

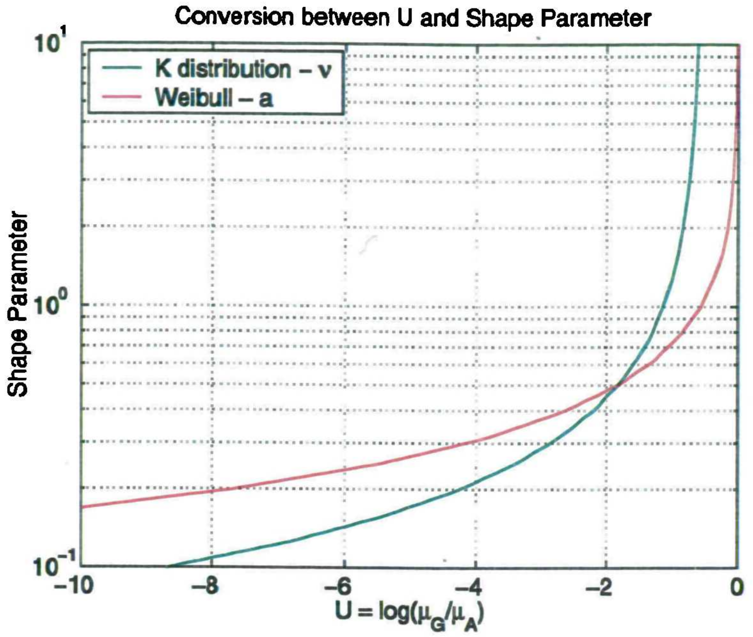

Figure 3.9: How the U estimator relates to Weibull and K distribution shape parameters

the sample distribution parameters are determined and verified through simulation - shown in Figure 3.10.

3.2 Physical Doppler Spectra

3.2.1 Origin

The Doppler spectrum from the sea surface, $\Phi(\zeta)$ where ζ is the Doppler frequency, can be rationalised as coming from three physical processes (Lee et al. 1995): Doppler broadening, damping dominated broadening and a mixture of the two.

Doppler broadening results from a fixed population of scatterers with a Gaussian distribution of speeds yielding

$$\Phi_G(\zeta) = \frac{1}{\zeta_e\sqrt{\pi}} \exp\left[-\frac{(\zeta-\zeta_G)^2}{\zeta_e^2}\right] \tag{3.48}$$

where ζ_G is the modal frequency and ζ_e is the Doppler width. Assuming a fast moving varying

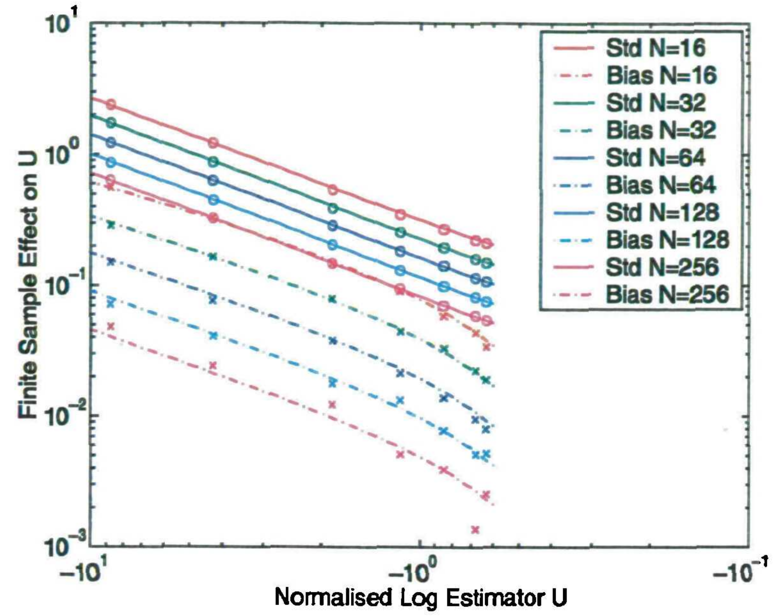

Figure 3.10: The bias and standard deviation in the sample distribution of U for varying sample number. Monte Carlo simulation shown as symbols upon a theoretical line

population with an exponentially distributed lifetime yields the Lorentzian distribution

$$\Phi_L(\zeta) = \frac{\Gamma/2\pi^2}{(\zeta - \zeta_L)^2 + (\Gamma/2\pi)^2} \tag{3.49}$$

where ζ_L is the modal frequency and Γ^{-1} is the characteristic scatterer lifetime giving a full width at half maximum (FWHM) of Γ/π. It is reasonable to assume that the true form of the spectrum will be a combination of the above effects and so the overall effect will be a convolution of the two, yielding the Voigtian spectrum

$$\Phi_V(a, u) = \frac{a}{\pi} \int_{-\infty}^{+\infty} \frac{e^{-y^2}}{(u - y)^2 + a^2} dy \tag{3.50}$$

where $u = (\zeta - \zeta_V)/\zeta_e$ is the normalised frequency, ζ_V is the modal value and $a = \Gamma/2\pi\zeta_e$ is the Voigt parameter which represents the ratio of the Lorentzian FWHM to the full-width at one e-folding of the Gaussian (Lee et al. 1995).

An extensive approximation suitable for curve fitting can be found in Lee et al. (1995) who further suggest using a Levenberg-Marquardt algorithm to fit spectral data via a linear mixture of the three components. One criticism which may be levelled is that a wide range of spectra could fit to a distribution modelled by 3 separate parameterised components without necessarily conforming to the physical models they represent.

Work by Walker (2000), again using wavetank data, has led to a 3 component Doppler model to include Bragg, whitecap and spiking events. Each component is assumed to contribute a Gaussian spectrum but the parameterisation is reduced by including a composite surface Bragg scattering model.

3.2.2 Analysis using the Windowed Fourier Transform

A coherent pulsed radar is capable of recording the phases of individual reflected pulses. At high transmit frequencies and natural environments the motion of the illuminated area has little effect on the received frequency. A Doppler spectrum of the scene is built up from analysing the phase evolution of the signal over time. Typically this is achieved using the Fast Fourier Transform (FFT) which in its standard form requires a complex data length of $N \in \{2^i; i \in \aleph\}$, resulting in the discrete frequency space representation. The resolution of the FFT is determined by the sampling frequency ω_s which, via the Nyquist sampling criterion, can only resolve N evenly spaced frequencies $\leq \omega_s/2$.

In a stationary environment with well defined periodic content the FFT gives an acceptable method of analysis; increasing the data length increases the number of distinct frequencies analysed. In a non-stationary environment the spectral content may change in time; windowing the data in order to localise the frequency content in time produces the Windowed Fourier Transform (WFT) in the left of Figure 3.11.

The window size is chosen to match the radar environment and is usually fixed for a given sampling frequency ω_s. It can be matched to the decorrelation time of the scene, the scan time or as a function of radar processing power since the FFT algorithm is order $N \log N$. This imposes a scale of analysis on the scene where all events are deemed to occur over the length of these windows. Any change in spectral content over the window is smeared into a single FFT; if the original periodic signal is not a pure sinusoid then incidental frequencies are introduced.

Windowing causes the measurement of frequencies which may not be of interest, accompanied by the masking of events with duration shorter or longer than the window size. In general signals are not composed of well defined frequencies and the resultant FFT needs smoothing in order to distinguish the shape of the Doppler spectrum. A common method is to simply use a boxcar

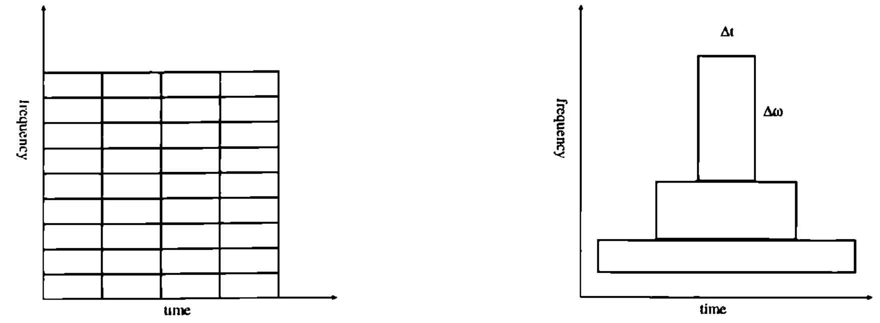

Figure 3.11: In the FFT sinusoidal basis functions are arranged at fixed intervals in frequency determined by the fixed window in time. The Continuous Wavelet Transform maintains a constant uncertainty in frequency and time

average, motivated by any frequency lying between two adjacent discrete FFT bins being spread into each. The averaging is often extended until the spectrum looks smooth, then consecutive windows are compared to view changes in the environment.

3.2.3 Analysis using the Wavelet Transform

The problem with the Windowed Fourier Transform is that fixed uncertainty in frequency is imposed by using a discrete algorithm and smoothing. This is accompanied by fixed uncertainty in time by windowing - yielding little scope for adaptivity to a changing environment. An alternate method of handling the uncertainty is from the classical form of the Heisenberg uncertainty principle which is not unique to quantum mechanics but a general property of functions (Kaiser 1994). If a function $g(t)$ is small outside a time localised to Δt and the Fourier transform $\widehat{g}(\omega)$ is small outside a frequency bandwidth $\Delta\omega$ then (Sarakar and Su 1998):

$$\Delta\omega\Delta t \geq \text{constant} \sim 1 \tag{3.51}$$

By normalising the function $g(t)$, hence $\widehat{g}(\omega)$, and identifying it with a probability function then the most probable values in time and frequency can be found from

$$t_0 \equiv \int_{-\infty}^{\infty} t|g(t)|^2 dt \ , \ \omega_0 \equiv \int_{-\infty}^{\infty} \omega|\hat{g}(\omega)|^2 d\omega \tag{3.52}$$

Looking at the variances from t_0 and ω_0

$$(\Delta t)^2 \equiv \int_{-\infty}^{\infty} (t - t_0)^2 |g(t)|^2 dt \ , \ (\Delta\omega)^2 \equiv \int_{-\infty}^{\infty} (\omega - \omega_0)^2 |\hat{g}(\omega)|^2 d\omega \tag{3.53}$$

it can be shown (Sarakar and Su 1998) that the limiting uncertainty product is

$$\Delta\omega\Delta t = \text{constant } \geq 0.5 \tag{3.54}$$

where the exact value can be calculated from the specific window shape.

As frequency cannot be measured instantaneously then Equation 3.54 gives the degree of certainty in a local estimate. The subtlety is that when looking at a range of frequencies it would be convenient to have the same overall uncertainty for each; Equation 3.54 gives the window size necessary as a function of frequency. This defines the Wavelet Transform (WT) of a signal, where a signal is analysed by scale to maintain constant uncertainty product shown in the right of Figure 3.11. Choosing a 'mother' wavelet and using scaled versions as basis functions not only allows a choice of frequencies to analyse, but also the shape of the periodic components to look for. To efficiently look for these frequencies at all window positions the convolution is implemented with the FFT. The Wavelet Transform thus expands the Doppler spectrum as an instantaneous evaluation of frequencies present within a sample length. As an aside the self similarity of the basis functions could reveal fractal nature in any signal (Wornell 1996)

The wavelet transform used is a discretised form of the Continuous Wavelet Transform (CWT) which introduces redundancy and is not orthogonal. A non-redundant transform exists in the form of the Discrete Wavelet Transform (DWT) which operates on a dyadic grid. The motivation for using a DWT is that the basis functions are truly localised, orthogonal and can be discontinuous. In addition frameworks are present to adapt the basis functions in response to the input signals for matched filtering. A major failing of the DWT is that it is not time-invariant; a single time shift in the input signal can produce a vastly different transform (see Daubechies (1993) for more detail on the DWT).

The standard FFT algorithm has computational complexity of

$$O[FFT] = N_{window} N_{freq} \log N_{freq} \tag{3.55}$$

where N_{window} refers to the number of windows used to analyse a data sequence. Consecutive windows can be overlapped or computed at every time step. N_{freq} refers to the size of that

window, determined by the desired time localisation but incidentally determining the frequency components calculated..
The WT has complexity of

$$O[WT] = \eta_{freq} N_{total} \log N_{total} \tag{3.56}$$

by analysing a chosen subset of frequencies η_{freq} the cost of increasing the transform length to the sample length N_{total} is offset; although in general the WT is more expensive.

Implementation

A possible input to a detector based on the wavelet transform is constructed in Chapter 7. The redundancy in the transform is accounted for, and is indeed a necessary result, from the uncertainty principle in Equation 3.54.
The choice of mother wavelet is a Morlet (Grossmann and Morlet 1985) because of the simple symmetric form. A further reason is that events of similar shape were observed in the data and have been demonstrated previously in the returns from a Loch (Werle 1995). The normalised form of the Morlet wavelet is shown in Figure 3.12; mathematically in time (t) and frequency (ω) (Torrence and Compo 1998) as

$$\Psi_0(t) = \pi^{-1/4} \exp(j\omega_0 t - \frac{t^2}{2}) \text{ , } \widehat{\Psi_0}(\omega) = \pi^{-1/4} \exp[\frac{-(\omega - \omega_0)^2}{2}] \tag{3.57}$$

Where $\widehat{\Psi}\,(\omega)$ indicates the Fourier transform of $\Psi\,(t)$. Strictly speaking the Morlet wavelet is not *admissible* since it extends to negative infinity and has non-zero mean (Farge 1992); however by choosing $\omega_0 \geq 5$ then the negative frequencies are small and the mean is negligible. The value chosen is that which made the first oscillation of the real component touch the envelope at half height ($\omega_0 = 5.336$) which also gives time-reversal symmetry.
Particular frequencies within a signal are analysed by dilating the mother wavelet by a scale s and shifting to a position β such that

$$\Psi_s = \frac{1}{\sqrt{s}} \Psi_0 \left(\frac{t - \beta}{s} \right) \tag{3.58}$$

If $x(t)$ is the signal then the wavelet coefficient of x at scale s and time t is defined from the Continuous Wavelet Transform shown in Figure 3.11 and defined as

$$W_{t,s} = \int x(\tau) \frac{1}{\sqrt{s}} \Psi_0(\frac{\tau - \beta}{s}) d\tau \tag{3.59}$$

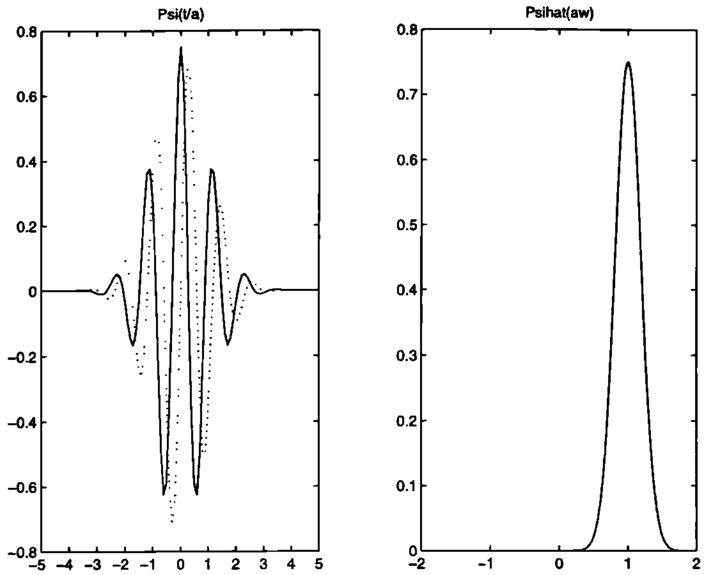

Figure 3.12: The Morlet mother wavelet in time and frequency (with respect to the Fourier frequency)

normalised to unit energy for each scale. Discretising this to N data with a signal $x(t) = x_{1:N}$ and using the convolution theory, the coefficients over the entire signal at a particular scale are calculated (Torrence and Compo 1998) via

$$\widehat{W}_s(n) = \sum_{k=0}^{N-1} \widehat{x}_k \widehat{\Psi}^*(s\omega_k)\exp(i\omega_k n) \qquad (3.60)$$

$$\omega_k = \begin{cases} \frac{2\pi k}{N} : k \leq N/2 \\ -\frac{2\pi k}{N} : k > N/2 \end{cases} \qquad (3.61)$$

where Ψ is now taken to be normalised.

Essentially a basis function defined by a particularly scaled mother wavelet is being applied at every time point n, this is done by convolution using Equation 3.60 where the FFT is used incidentally for speed. Upon inversion the complex value of the component W represents the certainty of a particular phase frequency being present at that time in the signal and so the modulus of this is taken as relative phase is unimportant (similarly, only the magnitude of the Doppler spectrum is useful). This is plotted as a 2 dimensional shaded image to reflect the magnitude of W along the dimensions of time and scale.

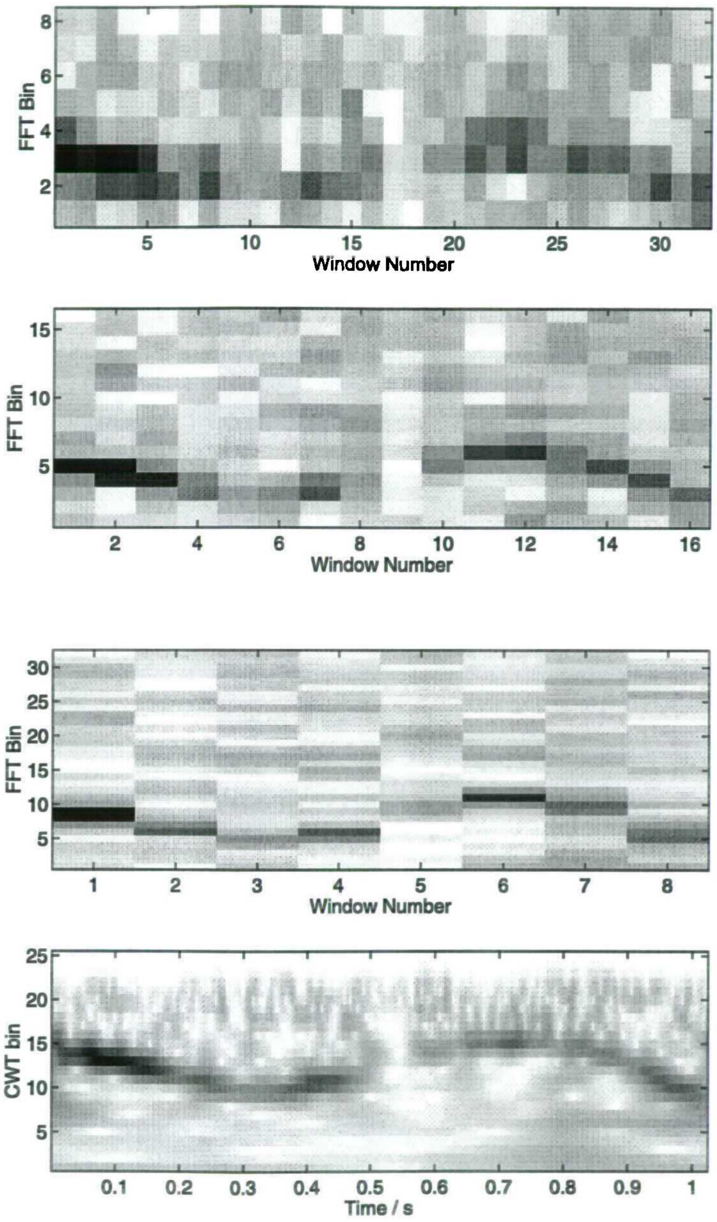

Figure 3.13: The Wavelet transform (bottom) identifies the best windowing in time and frequency which is difficult to obtain in the Windowed Fourier for nonstationary Doppler and power

Figure 3.13 demonstrates the Continuous Wavelet Transform in comparison to the Windowed Fourier Transform to illustrate the presence of a simulated target within system noise. Without knowing the velocity of the target, optimum windowing is unknown; the upper three plots show the result obtained from a WFT with time windows of 16, 32 and 64 samples from a simulated 512 samples. The lower plot shows the CWT which, whilst computationally more involved, seems to convey the information of varying velocity and intensity much better - this particular analysed signal is discussed in the next subsection.

3.2.4 Application

The choice of wavelet scale s is in some ways arbitrary but justification for using a logarithmic spacing is given in Chapter 7 as the width of the wavelet in frequency space is a linear function of its central frequency. This defines the natural scale of the CWT to be the logarithm of the velocity, unlike the linear relationship of the FFT.

A real target will have varying velocity and intensity, an accurate simulation (Chapter 7) of a typical target presented in Figure 3.14 is that analysed in the previous Section. The overall Doppler obtained using the FFT does not convey the situation well as there appear to be three targets at different velocities; taking the largest instantaneous CWT coefficient obtains an accurate representation.

There seems to be an advantage in performing the CWT for analysing the time varying Doppler of the sea surface as it removes the need to arbitrarily choose the size of analysis window. By minimising the uncertainty in declaring a velocity component to be present over time, the time-varying Doppler statistics of a signal are well represented - discussed further in Chapter 7.

3.3 Assessing Applicability of Statistical Models

Before any detection scheme can be developed based upon a statistical model for the data it is important that the degree of conformity to that model is assessed. In real world situations, noise (either statistical, system or environmental) usually causes some departure from the theoretical predictions which can mask the true model. Alternate models may be so similar that it is difficult to distinguish which is better, additionally the stability of the model should be assessed - any nonstationary departure from the assumed model may have considerable effects.

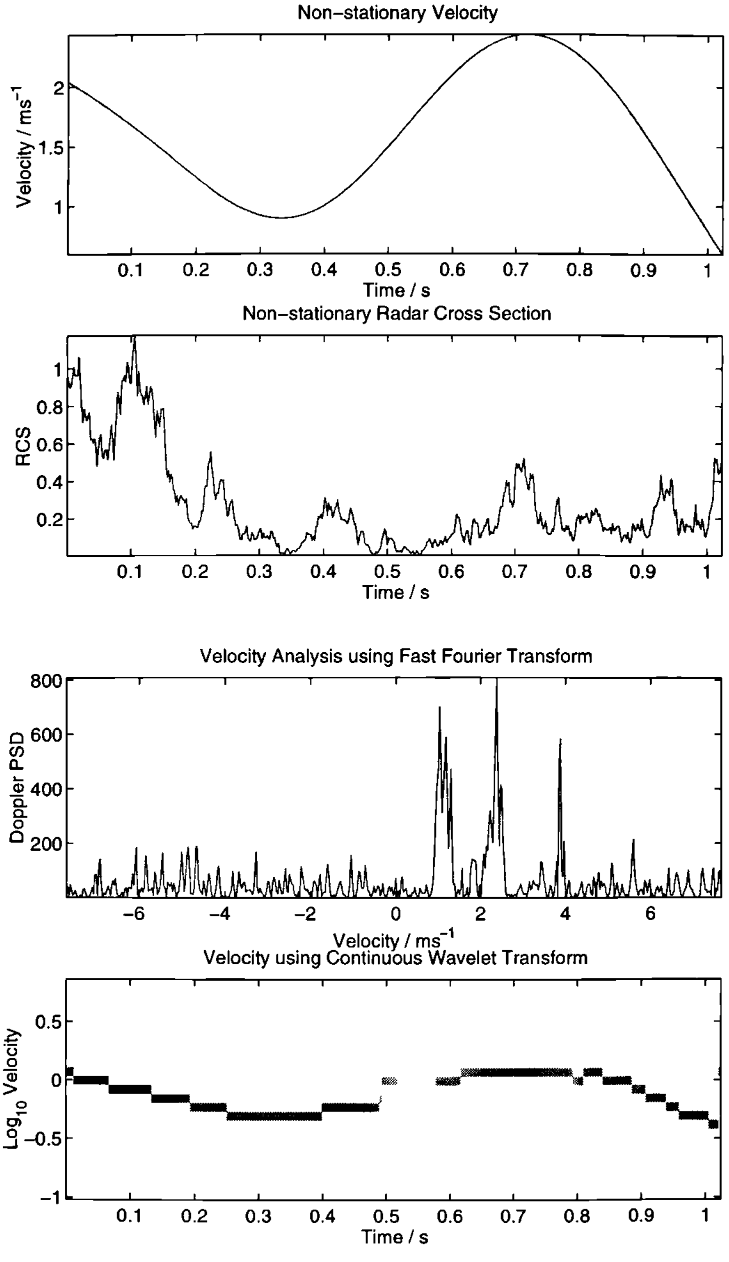

Figure 3.14: Wavelet analysis conveys information better than Fourier for nonstationary Doppler and power

3.3.1 Distribution Theory

Determining if two distributions are the same requires a search for some significance level ρ with which to discriminate between various parametrically known hypotheses H_i (PDFs) based upon the data x. The Maximum Likelihood Estimate (MLE) for the parameter vector θ_i of H_i is usually employed. Neglecting quantisation, a continuous distribution is tested to determine the significance of any departure of the data from that of the hypothesis.

The MLE is based upon maximising the probability of x by varying θ within a single H. As such this has no mathematical basis, being based on 'intuition'(Press et al. 1992). As the MLE is based upon H being true then any deviation from this will give an invalid estimate for θ.

The Kolmogorov-Smirnov (K-S) test (Press et al. 1992) measures the absolute difference D between the data CDF C_x and the hypothesis H CDF C_H

$$D = \max |C_H - C_x| \tag{3.62}$$

To determine a significance level ρ requires a property of the K-S test - that its distribution in the case of H being true is invariant under a transformation of x and can be approximated for $N > 4$ samples (Press et al. 1992) as

$$\begin{aligned} \rho_D &\approx Q_{KS}\left[\left(\sqrt{N} + 0.12 + \frac{0.11}{\sqrt{N}}\right) D\right] & (3.63) \\ Q_{KS}(\lambda) &= 2\sum_{j=1}^{\infty} (-1)^{j-1} e^{-2j^2\lambda^2} & (3.64) \end{aligned}$$

However, an important point is that Equation 3.63 is only valid if θ is not estimated from the data. Whilst this is embodied in the statement 'H being true', often the assumption is taken that it is true based upon an MLE estimate of the data which has an associated variance.

The test is insufficiently sensitive to departures from the tails (Cox and Hinkley 1974) which is the particular region of interest. The reason is that in the case of H being true then the probability distribution of D is not independent of x. The variance of D is proportional to $P(x)[1 - P(x)]$ and thus a deviation that may be significant at its own value of x is compared to the expected chance deviation at the maximum variance $P = 0.5$ (Press et al. 1992).

A way of overcoming the insensitivity to the tails of the distribution is to use a distance measure which has a constant variance over x. A simple choice is the Kuiper (KP) statistic (Press et

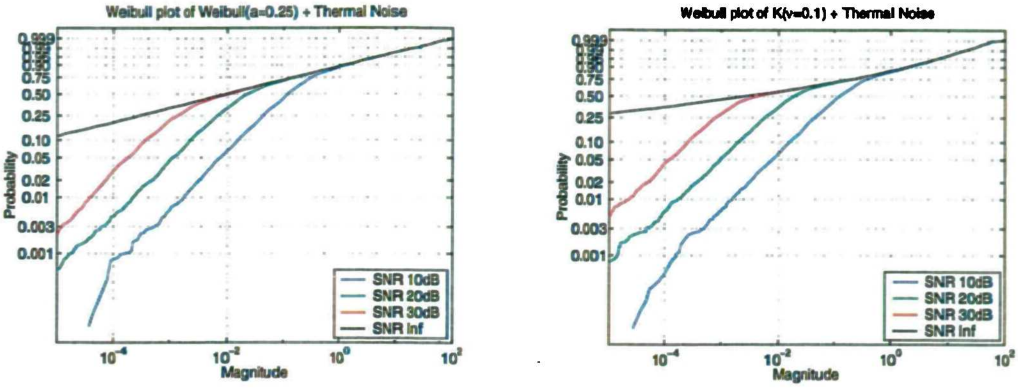

Figure 3.15: Weibull and K distributions show similar behaviour in the presence of thermal noise

al. 1992) defined as K

$$K = D_+ + D_- = \max\left[C_H - C_x\right] + \max\left[C_x - C_H\right] \tag{3.65}$$

with an associated significance level

$$\begin{aligned} \rho_K &\approx Q_{Kuiper}\left[\left(\sqrt{N} + 0.155 + \frac{0.24}{\sqrt{N}}\right)K\right] \qquad (3.66) \\ Q_{Kuiper}(\lambda) &= 2\sum_{j=1}^{\infty}\left(4j^2\lambda^2 - 1\right)e^{-2j^2\lambda^2} \end{aligned}$$

3.3.2 Alternate Methods

Weibull Analysis

So called 'Weibull paper' can be used to plot and analyse long tailed distributions. Weibull statistics are transformed to yield a straight line with a slope defined by the shape parameter, achieved via Equation 3.69 where $\mathcal{F}$ is the complementary CDF.

$$\begin{aligned} P_{Weibull}(I) &= \frac{a}{2\sigma^2} I^{a-1} \exp\left[-\frac{I^a}{2\sigma^2}\right] && (3.67) \\ \mathcal{F}_{Weibull}(I) &= 1 - \int_0^I P_{Weibull}(I)\, dI = \exp\left[-\frac{I^a}{2\sigma^2}\right] && (3.68) \\ \log\left[-\log\left[\mathcal{F}_{Weibull}(I)\right]\right] &= a\log(I) - \log\left(2\sigma^2\right) && (3.69) \end{aligned}$$

The appearances of Weibull and K distributed statistics plotted in this way are shown for shape parameters of $a = 0.25$ and $\nu = 0.1$ respectively in Figure 3.15, which include the expected

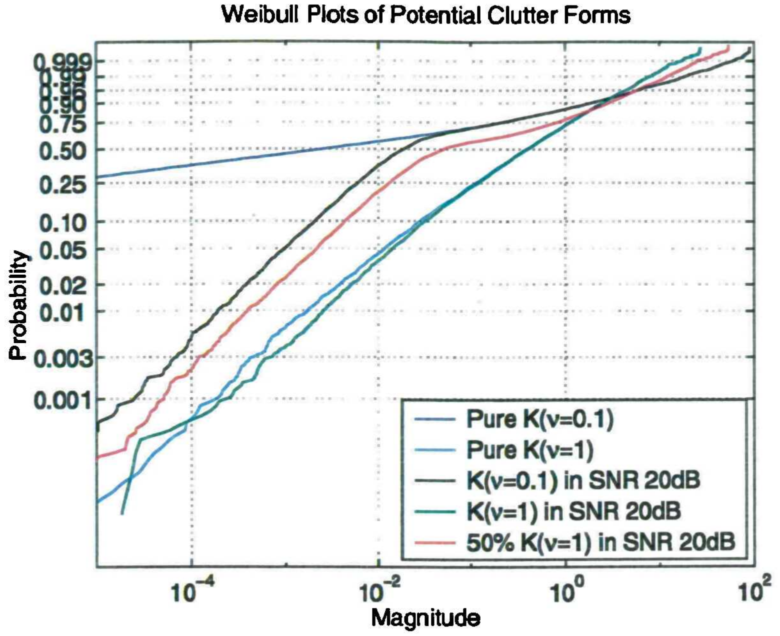

Figure 3.16: The data plotted on Weibull paper resembles extremely spiky data in noise or a mixture distribution of less spiky data

effects of thermal noise. This is included in a coherent manner and shows that even at a fairly good SNR of 30dB there is an obvious elbow in the plots where the lower magnitude has the expected slope of an exponential distribution but the higher values approach the noise free case. The plots show that in the presence of noise it is extremely difficult to distinguish K and Weibull statistics even at parameter values far from where they coincide ($a = \nu = 0.5$).

The Weibull plots of 3.15 shows shape parameters that are extremely spiky and on the limit of the expected values according to both the DERA clutter model ($a = 0.25$) and Ward et al. ($\nu = 0.1$) however these parameters are necessary to achieve an elbow in the Weibull plot that resembles the real data from Chapter 4; if a spatial mixture of distributions is formed then similar plots can be generated with much less spiky parameter values ($\nu = 1$).

3.3.3 Testing for Stationarity

The underlying statistics are assumed to be stationary in most detection schemes. The large body of literature on sea clutter shows that the statistics can vary considerably upon factors including grazing angle, frequency, polarisation, wind speed, sea state and radar type. Treating each data set as being a time series individually representative of a particular environment, temporal stationarity over the order of seconds should be tested as this is a possible detection scenario.

According to Harvey (1993), for a stochastic process X to be stationary several criteria must be satisfied at all time t by the discrete values x_t:

$$\langle x_t \rangle = \mu \tag{3.70}$$

$$\left\langle (x_t - \mu)^2 \right\rangle = \sigma_x^2 = \xi(0) \tag{3.71}$$

$$\langle (x_t - \mu)(x_{t-\tau} - \mu) \rangle = \xi(\tau)\ ,\ \tau = 1, 2, \ldots \tag{3.72}$$

where $\langle \cdot \rangle$ is the expectation. The above equations are recognised as mean, variance and autocovariance respectively. Harvey further notes that to test for stationarity in this way the time series must be 'ergodic' such that observations sufficiently apart should be almost uncorrelated; not the case in a cyclic time series for example.

Equations 3.70 - 3.72 define 'weak' stationarity. A further constraint defining 'strict' stationarity is that the joint probability distribution of the observations is constant, i.e. a multivariate condition. Using stationarity to test the applicability of a PDF only requires univariate observations and so this is too strong a constraint.

A problem in applying the stationary criteria to a single set of data is how to window the data in order to assess a change in statistics. This is normally determined by practical means, in a rotating radar the number of data points will be determined by the antenna angular velocity and PRF, but in MFR there is no restriction in stare mode and the PRF can be adaptively changed.

In real data even the weak stationary criteria are often too strict and so it is normal to apply a transformation to render the data stationary. Commonly a windowed processor defines a local region of the data and the data is normalised with respect to its local mean, this forces Equation 3.70. Similarly, dividing Equation 3.72 by the local variance yields the autocorrelation function which measures the form of the correlation decay rather than the magnitude.

Further transformations can be used to obtain stationarity, these include 'differencing' to remove linear trends by using the transformed variable $\Delta x_t = x_t - x_{t-1}$ and also a log transformation which is applicable when the standard deviation is proportional to the mean (Cromwell et al. 1994). Some mathematical tests for stationarity, discussed by Cromwell et al. using the Dickey-Fuller test (Dickey and Fuller 1979), rely on assuming an autoregressive model *a priori* and testing for the null hypothesis of a random walk. Priestly (1988) discusses others that compare two sections of a series' covariance or spectral properties, making the comment that these sections must be locally stationary for this to work. A test to assess the *overall* stationarity of the time series is required and is given by Priestly via an estimate of the 'evolutionary' spectrum over time. However, this evolutionary spectrum is determined using a window and is thus subject to the problems discussed in Section 3.2.3.

If a change or evolution in the statistics are possible then one must carefully choose detection schemes with respect to their *a priori* assumptions and robustness.

3.4 Further Possibilities

Two proposals concerning chaotic and fractal properties of sea clutter have been made in the literature. A brief introduction is necessary to make them accessible but no further consideration has been made in this thesis, the somewhat esoteric maths and inability to relate directly to the physical process does not justify it (although significant progress in applying fractal modelling to sea clutter has been made recently by Berizzi et al. (2000)).

3.4.1 Chaotic Model

A chaotic system can satisfy all tests for randomness and yet be entirely based on deterministic nonlinear dynamics. Haykin (1997) has advocated that sea clutter is a chaotic process.

The Logistic Equation

The logistic equation demonstrates some simple properties of chaos:

$$x_{t+1} = kx_t(1 - x_t)\,,\, 0 < k < 4 \tag{3.73}$$

Using the same initial x value of 0.5, graphs are shown for various parameter k in Figure 3.17. For $k < 1$ the graph tends to zero, i.e. zero is the *fixed point attractor* for the system. Increasing k towards 4 shows various trends, the fixed point attractor for $k \geq 1$ is non-zero. When $k > 3$ a bifurcation takes place whereby there are now 2 attractors, i.e. a two point attractor. Putting $k > 3.4494\ldots$ shows periodic doubling with a 4 point attractor. Increasing k beyond this shows $8, 16, 32, 64\ldots$ point attractors until arriving at a state which can show correlation and periodicity on many scales - generally referred to as chaotic.

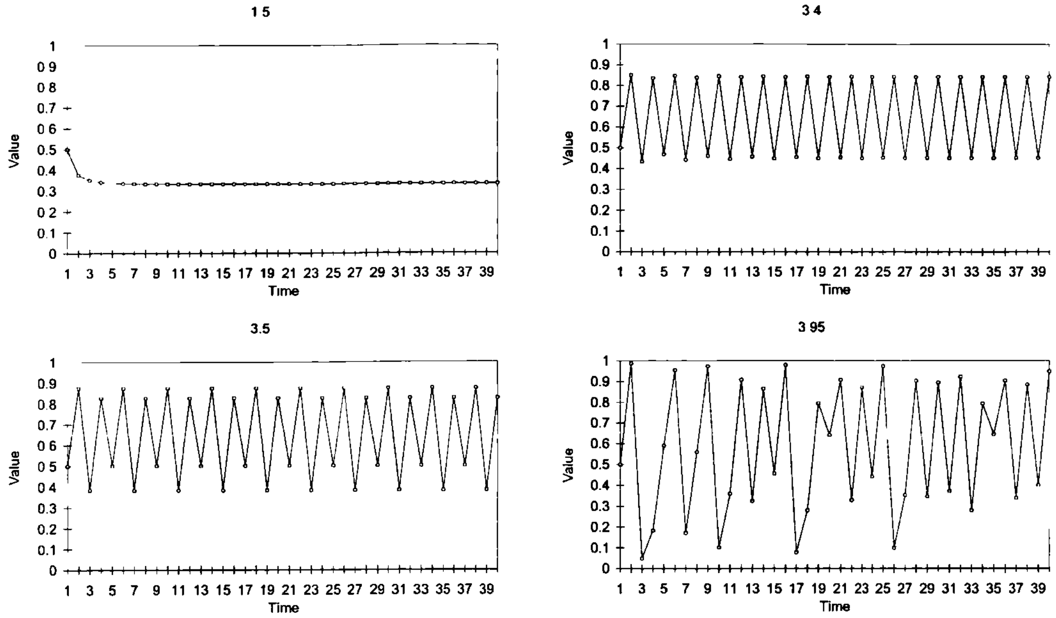

Figure 3.17: Evolution of the logistic equation dependent upon control parameter

The simple one dimensional equation shows a range of differing behaviour controlled by a single parameter. Periodic doubling is not immediately obvious from the equation, and is often an indication of the onset of chaos in simulated systems. Additional concepts in chaos are discussed in Appendix B.3.

3.4.2 Fractal Modelling

Fractals are systems whose invariance lies not with time but scale. A classic example is found in the bronchi of the human lung where the same branching structure is seen at every scale. Typically, analysis of structure is limited at large scales by the sample length and at small scales by the sampling rate. If the spatial statistics of the system are similar over this range this is termed 'statistical self similarity'. One such family is the $1/f$ process generally defined (Wornell 1996) as processes having a measured power spectrum obeying

$$\Phi(f) \sim \frac{\sigma^2}{|f|^\gamma} \tag{3.74}$$

which gives a correlation structure via Fourier transform (Champeney 1987)

$$\frac{\sigma^2}{|\omega|^\gamma} \longleftrightarrow \frac{\sigma^2 |\tau|^{\gamma-1}}{2\Gamma(\gamma)\cos(\gamma\pi/2)} \tag{3.75}$$

3.5 Summary

Probability distributions commonly applied to sea clutter have been covered, accompanied by their parameter estimators. The popularity of the K-distribution arises from its physically attractive compound representation which defines the statistics as coming from the speckle of an underlying power modulation. A commonly applied alternative is the Weibull distribution, despite it lacking a strong physical justification.

Both the Weibull and K shape parameters have iterative and non-analytic maximum likelihood estimators respectively and so an estimator suitable for real time processing is identified in the normalised log estimator U. Importantly this is known to be accurate for both distributions and has a simple interpretation related to the geometric and arithmetic mean.

Justification has been given for the subsequent use of a single texture measure U. A novel contribution was made investigating the sample distribution of the texture estimators subject to progressive censoring, thermal noise and low sample numbers. Uniquely U has been shown to be Gaussian distributed for all considered cases, this allows robust statistics to be used for determining the confidence in a particular measured value. The effect of edges and modulation upon the measurement of U was explicitly determined to allow further investigation of non-stationary data.

The physical origin of the Doppler spectrum was considered but it is suggested that except under the controlled conditions of a wave tank, a three component mixture could reasonably fit a wide range of measured spectra without necessarily confirming their physical nature. Explicit problems have been stated in the usual analysis of Doppler by fast Fourier transform methods. Analysis via the wavelet transform has been suggested as a way of overcoming the simultaneous uncertainty in frequency and time of a signal by achieving a lower bound to their uncertainty product. This gives an elegant way of analysing non-stationary signals without the arbitrary windowing and smoothing commonly associated with Fourier methods.

Whilst chaotic and fractal analysis offer appealing sources of novel research, they are still a new science and do not have the robustness of established mathematics, and so no further consideration will be made.

Simple methods of identifying an applicable statistical distribution have been shown with cautionary advice, and the definitions of stationarity have been stated.

Chapter 4

Assessment of Real Data

This Chapter covers the analysis of real data recorded in varying ocean conditions. It must be borne in mind that the results are an *interpretation* of the scene by the radar and whilst strong mathematical models are desirable, the experimental and environmental conditions are critical.
The environment is illustrated by two digitised stills from the continuous video recordings in Figure 4.1 which show typical surfaces from a 4.3m and 6.1m significant wave height (the mean height of the upper third vertical displacement as measured with a waverider buoy). It is difficult to spatially compare (register) these optical frequency images with those derived from the radar data without a visible distance scale in the images, but some general points can be made:

- The 4.3m scenes show the presence of numerous abrupt small-scale whitecaps. Although their effect at operable radar frequencies is presumably different it indicates a change in the underlying physical processes responsible for scattering.
- Modulation can be seen in both images due to the swell present, but further wave structure is apparent without breaking.
- Within the 6.1m scene significant areas are hidden from view by the much larger breaking structure, hence the modulation is partly in 'shadow' as viewed by the radar.
- Larger scale wave events can be seen towards the horizon which are explained in part by the popular belief that the seventh ocean wave is the largest. It is not clear how this arises but fully developed waves can be calculated to break when their height is approximately

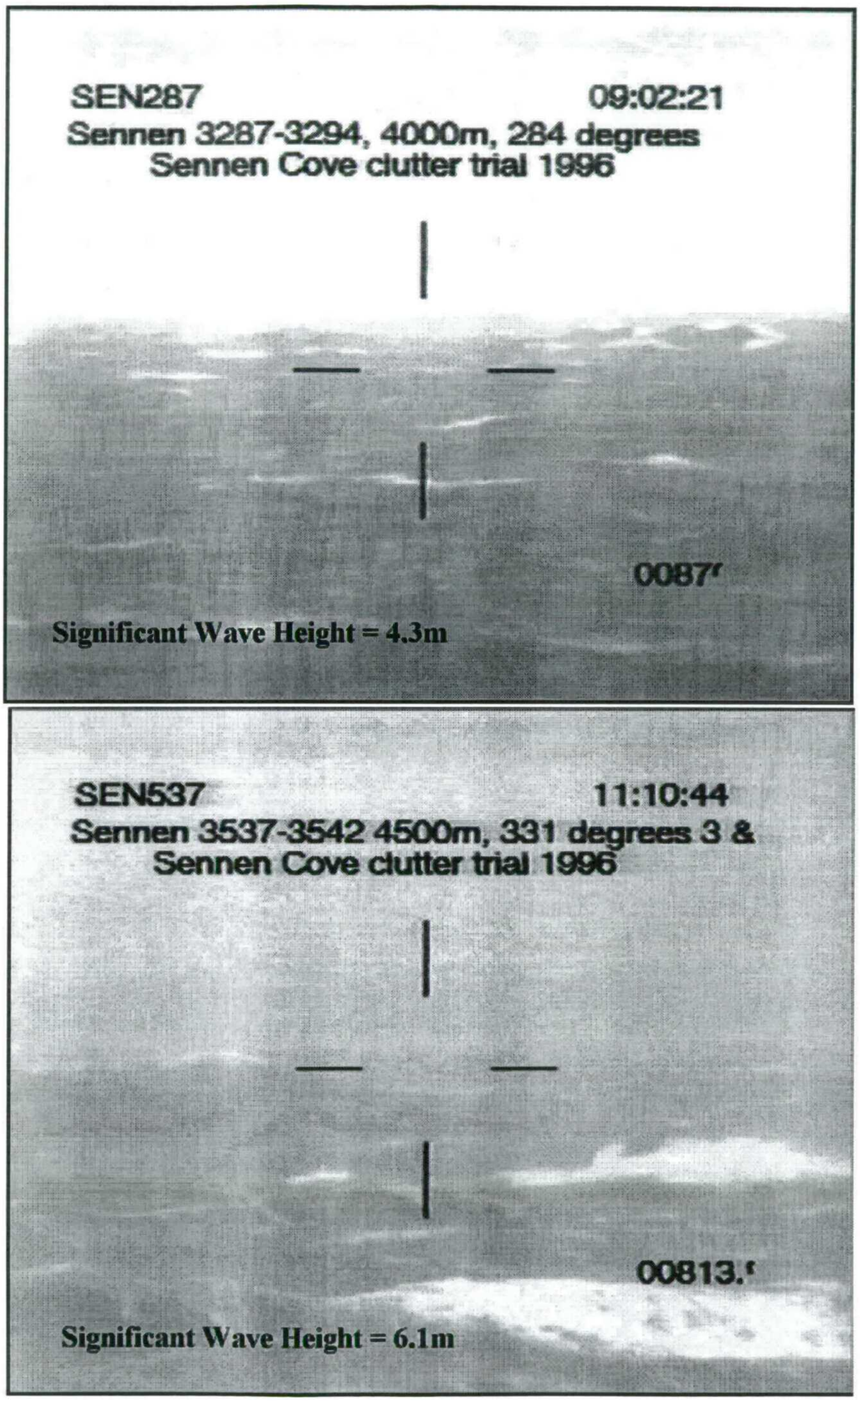

Figure 4.1: Simultaneous video still from a) sen3287 - sen3294 with significant wave height of 4.3m and b) sen3537 - sen3542 with significant wave height of 6.1m

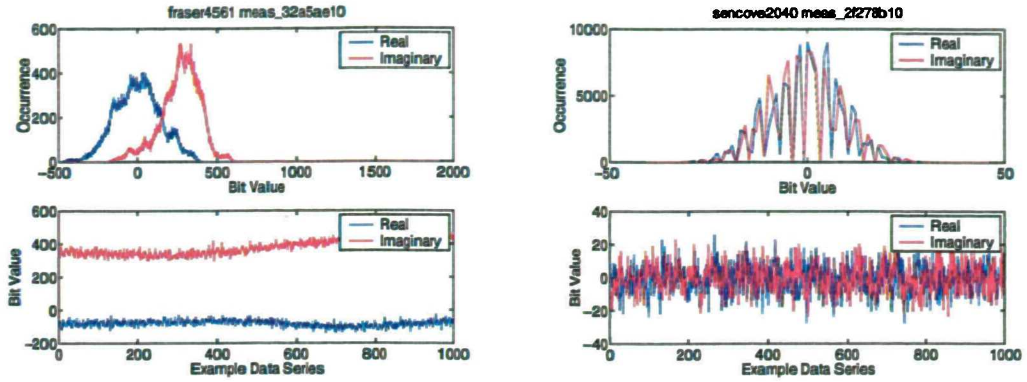

Figure 4.2: Raw data histograms showing A/D problems

one seventh of their wavelength (Komen et al. 1994) which, if any resonance is present, is a possible explanation.

- The majority of data is from high sea states which are at the extreme of usual operating parameters for practical radars.

4.1 Data Integrity

The data is output digitally by the radar in real time to Exabyte tapes. A large amount of data taken in varying operating conditions between 1994 and 1996 is available at the DERA Portsdown site of which several gigabytes were transferred to a PC for analysis using MATLAB (Mathworks 1999). User friendly software was provided by DERA in the form of Matlab programs to format, display and process the data. These were not suitable for fast development, being based upon a Graphical User Interface (GUI), so a parallel set of command line procedures were written.

The raw coherent data is stored as a complex integer defined in the complex plane of in-phase (I) and out-of-phase (Q) components. Each component is obtained from a separate Analogue to Digital (A/D) converter. By histogramming the I and Q values separately it is possible to check the operation of the A/D. In addition this checks to see whether the correct attenuation has been used on the radar to give the data a reasonable dynamic range without introducing spurious A/D noise. In a large number of cases the data was degraded by the absence of certain values, this is indicative of a 'sticky bit' where the binary representation of the integer consistently has one or more bits at the same value. In addition some integers were of lower than expected occurrence, indicative of an intermittent or partial sticky bit. The shape of the histogram was also variable, obviously some dependence on environment is expected but some were severely deformed with no obvious explanation.

Occasionally I and Q were not centred at zero, this is due to a bias voltage entering the system which can be corrected in pre-processing. Voltage spikes also contributed to data 'glitches' where a single data point would be far greater than expected. In all other reasonable cases the I and Q were balanced and had the same dynamic range. Two particularly bad examples are shown in Figure 4.2 but all files were screened by hand before use in this work. A further observation, presumably due to stray AC voltages, is frequency contamination of the data appearing as small spikes in the Doppler spectra. Another complication is the presence of birds in the data, this has been confirmed from simultaneous video capture for some runs.

A problem unique to the formation of a high resolution image is that the set of individual pulse frequencies necessary to form the chirp (Section 2.2.2) was often not sequentially transmitted quickly enough (cycled) to form an accurate range profile. Determining when this happened is made by examining the Doppler of a single chirp frequency or that of the imaged high resolution range cell. A suitably formed Doppler spectrum will be seen only if the scene has not moved significantly over the imaging time. Unfortunately, this fact, compounded with the other problems meant that no 'perfect' 8GHz high resolution data is available and very few perfect 3GHz files; however if the data is viewed as being representative of the output from an operational radar then it cannot be ignored. Indeed this situation may frequently arise if a fast moving ship is moving in the opposite direction to the waves and so it is still instructive to analyse, although a simple interpretation of the received statistics is lost.

4.2 High Resolution Range Profile Data

The high resolution scenes were formed by stepped frequency chirp imaging as discussed in Section 2.2.2. The 1.5m resolution scenes were achieved over a bandwidth of 100Hz with 256 frequency steps, the 19cm resolution data utilised 800Hz bandwidth with 512 frequency steps. Pulse repetition frequency, centre frequency and polarisation are detailed in Tables 4.1 and 4.2.

Table 4.1: High Res Range-Time files used at 8GHz. H_w is Significant Wave Height, V_w is Wind Velocity, Az refers to Azimuth Angle

Run	F/GHz	PRF	Res	Pol	H_w/m	V_w/ms^{-1}	Az_{wind}	Az_{wave}	Az_{look}	Rng/m
sen3363	8.4	20k	1.5m	HH/VV	4.08	12	250	273	284	4005
sen3464	8.4	20k	1.5m	HH/VV	4.17	21	230	257	331	5010
sen3502	8.4	20k	1.5m	HH/VV	4.46	22	230	260	284	5010
sen3470	8.4	20k	19cm	HH/VV	4.40	28	240	241	331	5010
sen3508	8.4	20k	19cm	HH/VV	4.46	22	230	260	284	5010

4.2.1 8 GHz Scenes

Five range time files in horizontal and vertical polarisations are used as an example. It must be stressed that the derived Doppler spectrum is nearly flat due to excessive aliasing of the spectrum, this implies the chirp frequencies were not cycled fast enough to form an entirely valid image. Figure 4.3 shows scenes at 1.5m resolution and Figure 4.4 shows scenes at 19cm resolution, each with 26 seconds of data. The data has been range calibrated and the RCS obtained from reference to a calibration sphere. Environmental and system parameters are given in Table 4.1.

Suitably comparable images were obtained by using calibrated RCS images and displaying them with a dB scale. Physically interpretable scenes are still produced indicating that the chirp cycling is not critical in obtaining an image. Only those low velocity scatterers whose phase can be measured unambiguously for the entire chirp will be imaged correctly. High velocity scatterers which move a distance greater than the transmit wavelength are 'smeared' into neighbouring range cells; this environmental dependence may have a complicated time varying effect upon the statistics. In particular the 19cm resolution data used a longer chirp length of 512 pulses which further increased the minimum chirp cycle time.

4.2.2 3 GHz Scenes

The larger wavelength of the 3GHz scene means that the frequency cycling is less of a problem. The stepped frequency chirp will image correctly if it can be formed whilst the scene is relatively stationary, this usually means movement of less than about $\lambda/2$. Large files whose Doppler spectra were not completely flat are detailed in Table 4.2. The conditions of sen3287 were shown in the video still of Figure 4.1.

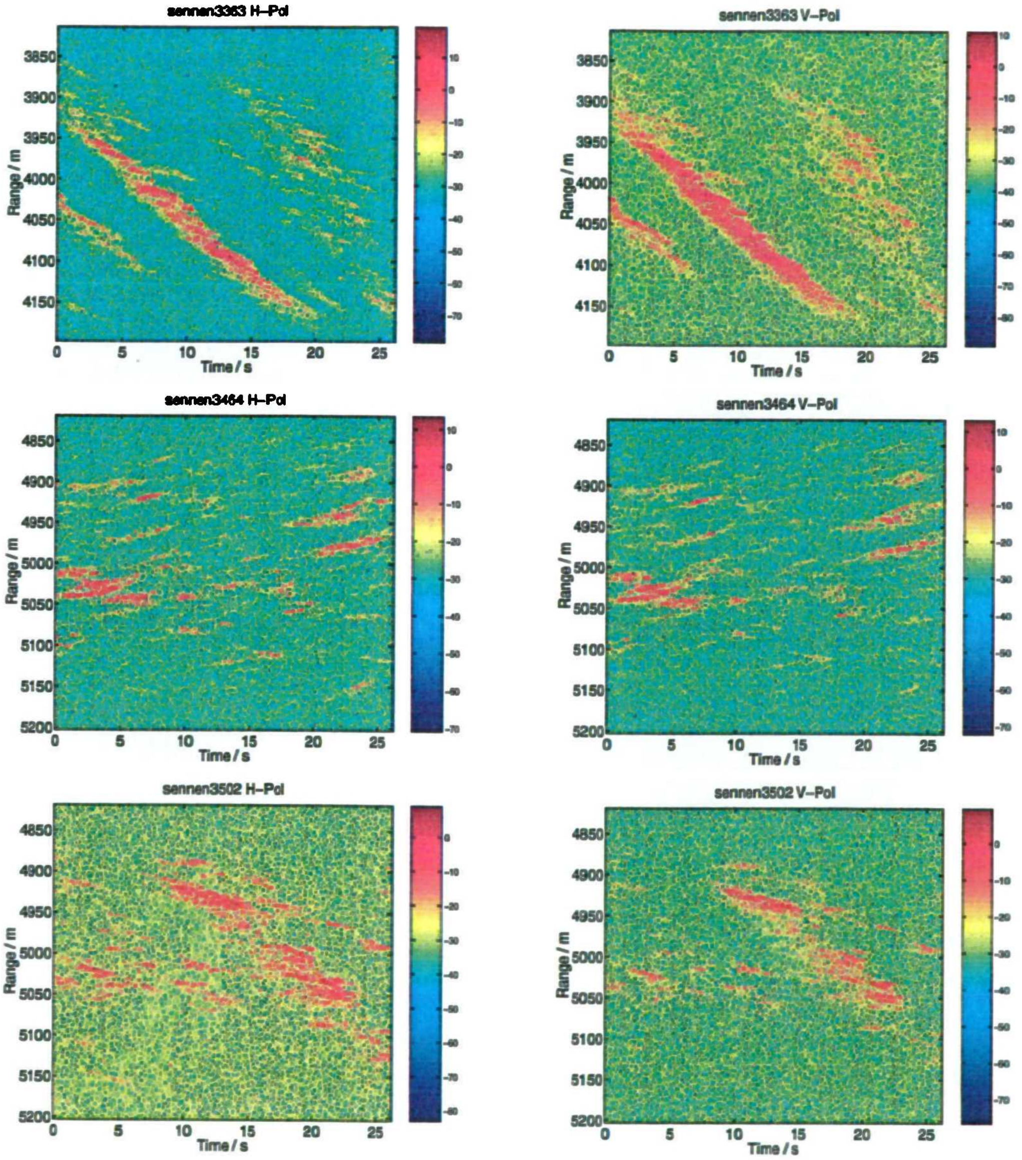

Figure 4.3: Calibrated dB RCS range-time plots for 8GHz 1.5m resolution (simultaneous H & V polarisation)

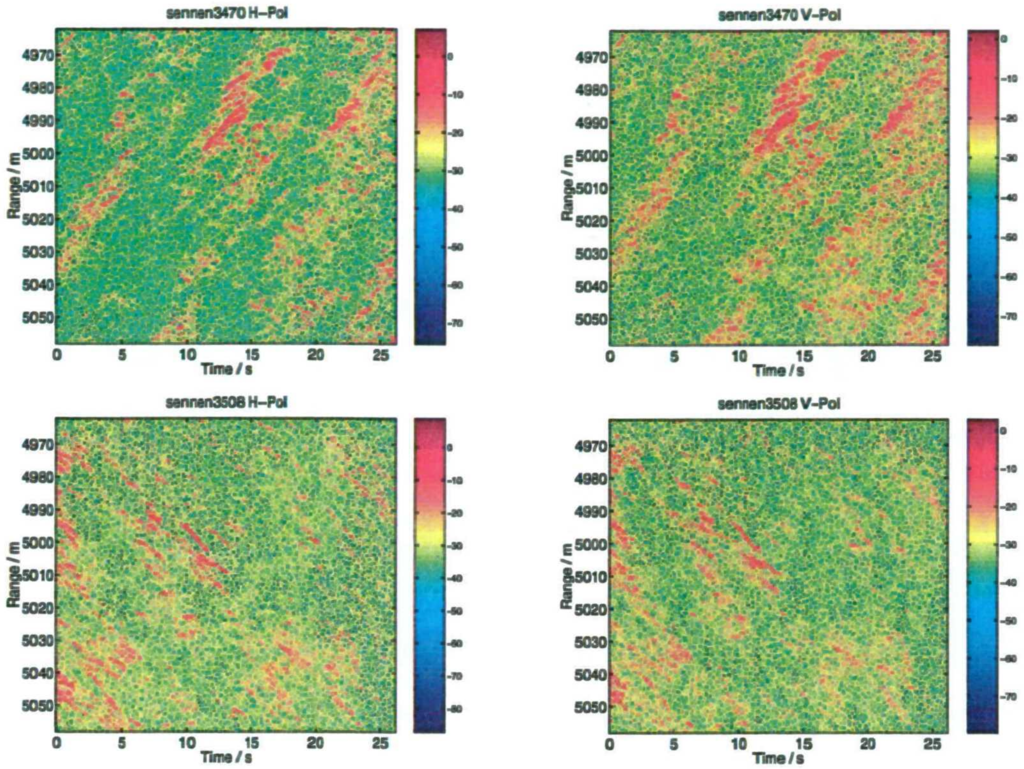

Figure 4.4: Calibrated db RCS range-time plots for 8GHz 19cm resolution (simultaneous H & V polarisation)

Table 4.2: High Res Range-Time files used at 3GHz. H_w is Significant Wave Height, V_w is Wind Velocity, Az refers to Azimuth Angle

Run	F/GHz	PRF	Res	Pol	H_w/m	V_w/ms^{-1}	Az_{wind}	Az_{wave}	Az_{look}	Rng/m
sen3277	3.0	20k	1.5m	HH	3.94	15	210	232	331	1995
sen3287	3.0	20k	1.5m	HH	4.25	30	330	276	284	4005
sen3352	3.0	40k	1.5m	HH	3.95	16	350	266	285	2010
sen3353	3.0	40k	1.5m	VV	3.95	16	340	266	285	2010
sen3613	3.0	20k	1.5m	VV	2.56	13	50	267	284	2505
sen3734	3.0	20k	1.5m	VV	1.37	9	320	280	284	2505

All files are singly polarised as no alternately polarised data was acceptable, however several consecutive files are available which should allow long term comparison. Each file is about 20 minutes long, coupled with their high PRF this gives a reasonably large measured data set which accounting for pulse shape gives over 10^7 samples of data per run, of which some is shown in Figure 4.5.

4.2.3 Discussion of Observed Scenes

Initial Impressions

Common to each frequency, three separable areas of homogeneity seem to be present in the scenes:

- Areas of near constant low RCS at the edges of the range gate. These are attributed to the received pulse being of constant high power over the central area of the range gate only.
- Wave fronts of high RCS with obvious correlation in range and time. The correlation varies between scenes and perhaps within a single scene.
- Areas with variable correlation and abrupt RCS changes.

The presence of separable regions of varying correlation suggest that it may not be possible to treat the scene as a single process. In particular the application of any estimator suited to a simple PDF with an associated measure of fit may be inappropriate if applied to the scene globally.

There is obviously a problem with sen3613 and sen3734 as they appear particularly noisy and their RCS scale is about 20dB below the other images. This suggests excessive attenuation was used however nothing in their notes supports this.

Figure 4.6 shows the RCS probability distribution from all scenes using Weibull paper as detailed in Section 3.3.2. The resultant plots have a definite elbow which separates two approximately straight lines. This is an important result as it indicates a change in distribution for large RCS returns; the straight lines show that a distribution similar to Weibull is applicable in two separate regions.

Crucially any estimator based upon a single global PDF is likely to result in a value midway between that of the two areas. To illustrate the relevance of this, the Weibull MLE estimator was applied to data from H-Pol Sennen3363. Figure 4.7 shows the MLE obtained PDF overlaid on that of the histogrammed data. Amplitude is used for effect, it shows a vast difference and the form

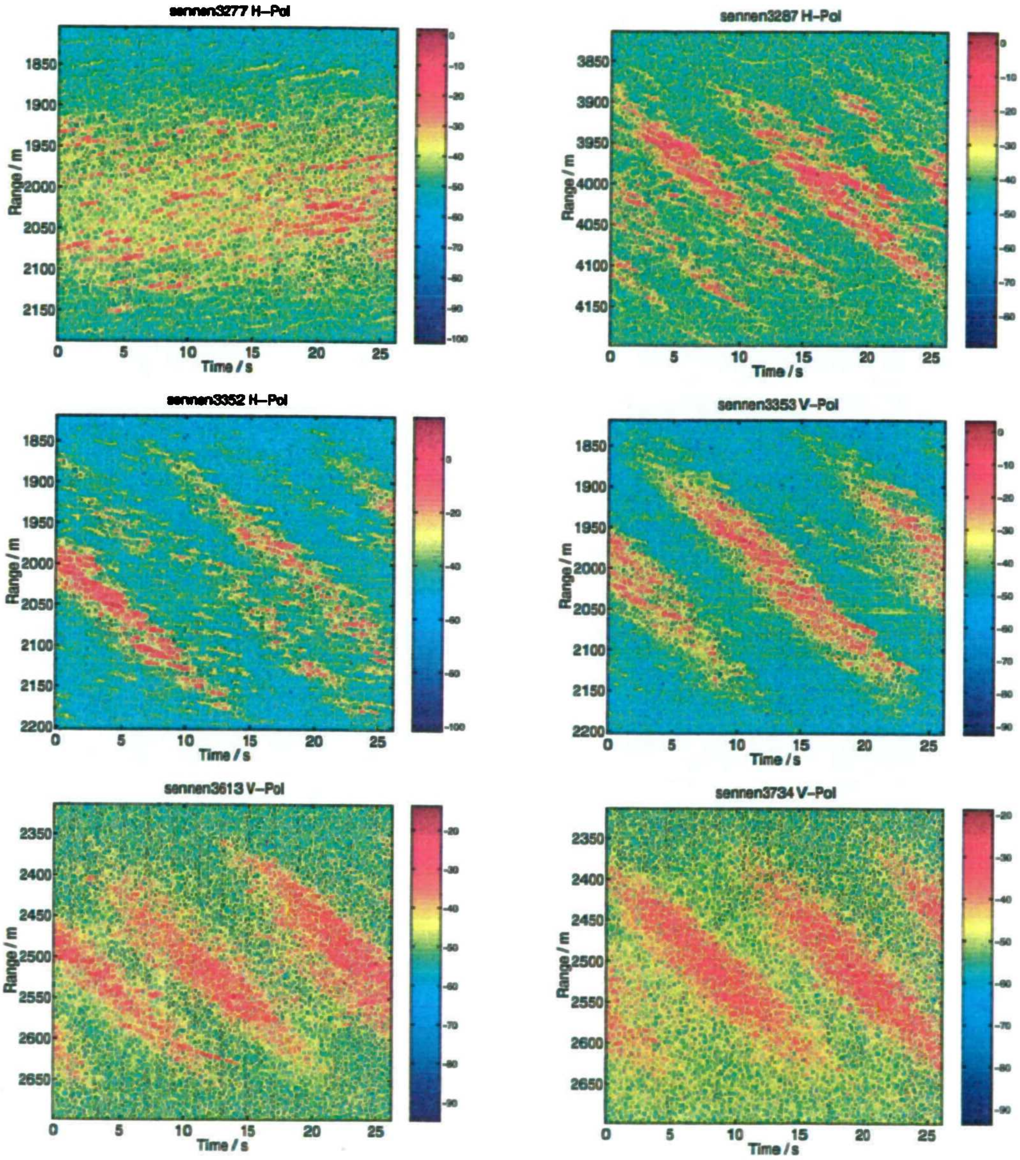

Figure 4.5: Calibrated dB RCS range-time plots for 3GHz 1.5m resolution (singly polarised H or V)

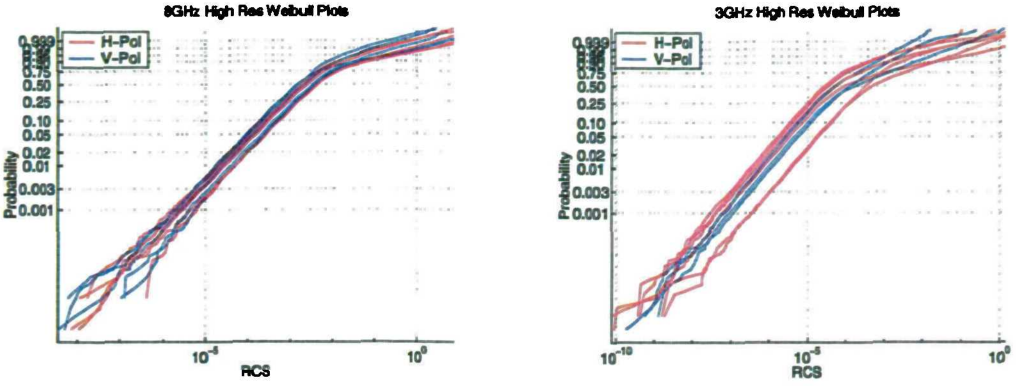

Figure 4.6: Weibull plots of RCS distributions from high resolution scenes

is not correct - the KP statistic is, unsurprisingly, zero. The problem is that the distribution is not Weibull and thus the Weibull MLE is not applicable.

Polarisation Differences

Using the 8GHz alternating polarisation scenes it is interesting to see if H and V have similar distributions. Although not shown explicitly, separating the H and V plots in Figure 4.6 shows the horizontally polarised data to have a much higher probability of large returns, as observed in many papers on sea clutter. The extent of this is variable, but without any perfect 3GHz alternate data further analysis is not performed.

Noise Component

The noise figure can be estimated by plotting the mean returns from the range swath over a long time. This has been performed in Figure 4.8 and the relative noise level can be determined from the edge regions. The axis shows the range bins averaged over time; the overall extent should be about 384 metres with a centred, approximately Gaussian, pulse shape corrected for range differences (discussed in Section 2.1.2) but:

- The edges are not equal; this could indicate range correction is not perfect due to poor mechanical calibration of the radar.
- The pulse is not always exactly centred; there could be a slight problem with the loopback (Section 2.2.2).

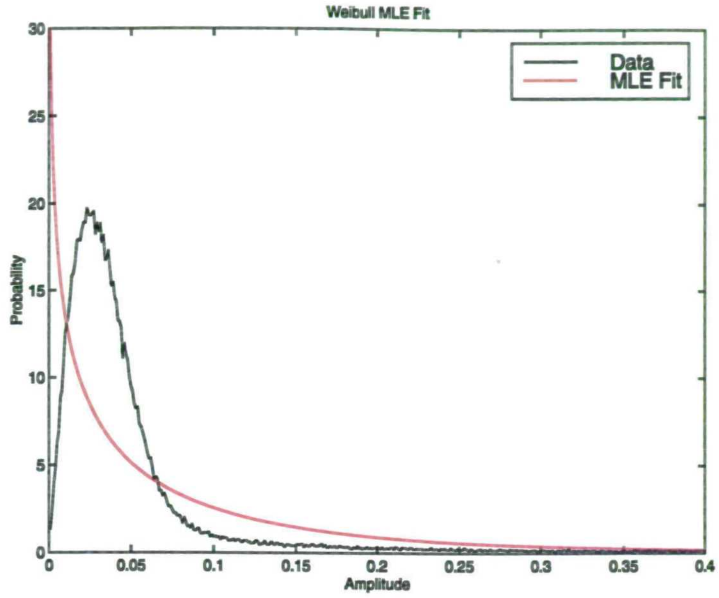

Figure 4.7: A Weibull MLE fit to H-Pol Sennen3363

- The pulse is suitably Gaussian over the central region and further analysis accounts for this.
- The SNR is variable but indeed sen3734 and sen3613 show the lowest SNR at about 10 and 15dB respectively, with other scenes better than 20dB.

As discussed in Section 3.1.4, in addition the noise component can be estimated from forming the Doppler spectrum. This is only possible if the chirp was cycled fast enough and Figure 4.9 shows that sen3277 and sen3287 are very poor. For the remaining files, the SNR calculated from the spectrum agrees with that calculated from the derived range profile above.

Comparing Figure 4.9 with the previous plot of Figure 4.5 shows that modulated wave-like images are obtained only when the chirp has been cycled fast enough and that the SNR has an obvious effect upon the image quality.

4.2.4 Further Corruption of the Data

Problems arising from analysis of the high resolution data are:

- Quantisation errors, presumably due to faulty A/D.

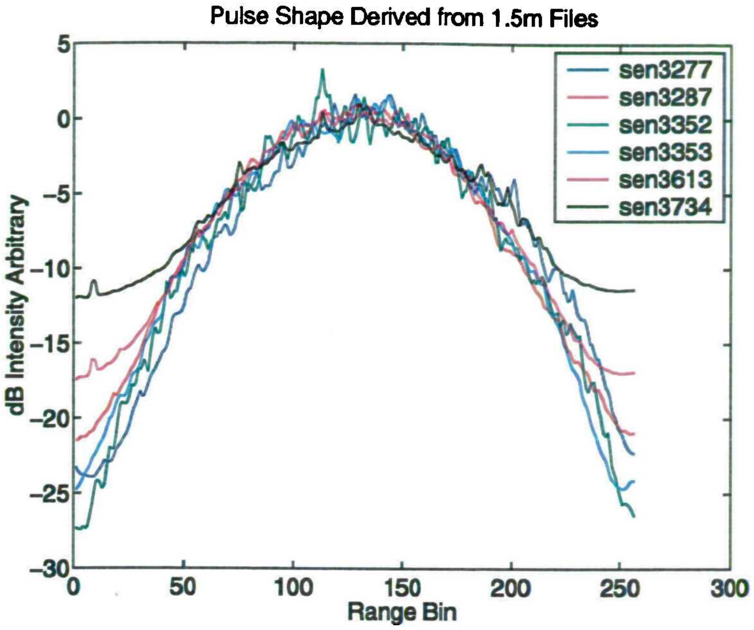

Figure 4.8: The mean intensity from high resolution range swath over 20 minutes, arbitrary dB scale

Table 4.3: The two selected High res files for overall analysis. H_w is Significant Wave Height, V_w is Wind Velocity, Az refers to Azimuth Angle

Run	F/GHz	PRF	Res	Pol	H_w/m	V_w/ms^{-1}	Az_{wind}	Az_{wave}	Az_{look}	Rng/m
sen3352	3.0	40k	1.5m	HH	3.95	16	350	266	285	2010
sen3353	3.0	40k	1.5m	VV	3.95	16	340	266	285	2010

- Mismatched I & Q channels.
- Excessive system noise, which the calibrated RCS suggests could be due to an attenuation of 20dB.
- Chirp cycling was not performed fast enough, especially in the 8GHz files, smearing high velocity components into neighbouring range cells.

Extensive checks leave only two acceptable files and two more with significant noise contamination. This is really not enough to discriminate between existing clutter models, and the decision is taken to concentrate on files sen3352 and sen3353. Their parameters are repeated in Table 4.3.

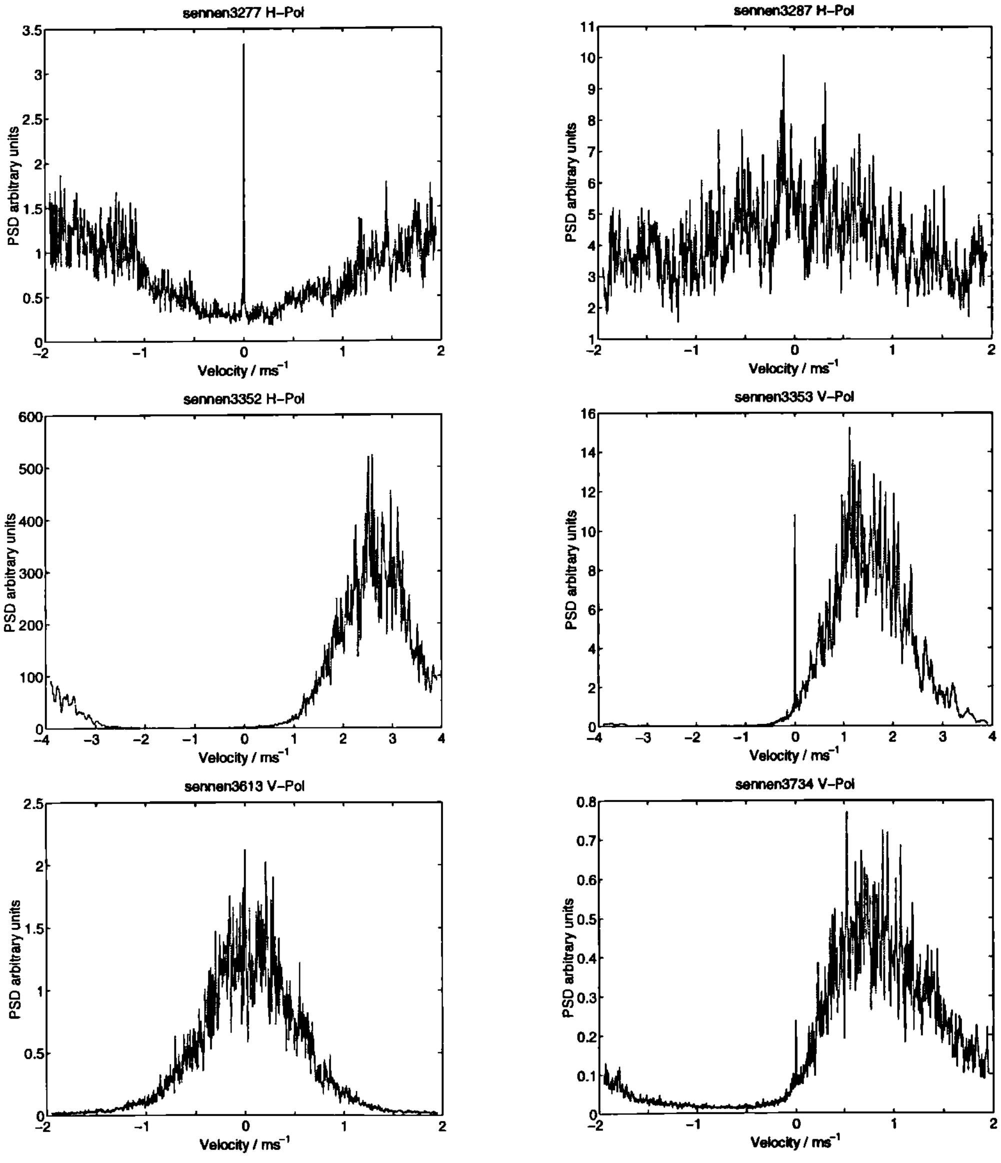

Figure 4.9: Doppler plots for 3GHz 1.5m resolution (singly polarised H or V)

Upon examination of these two files, intermittent high intensity range profiles were noticed termed ‘clicks’. These did not seem to be random with some periodicity present, examining the raw data showed these to arise from a single high intensity value in the received chirp. If the returns are accepted to be random on a pulse to pulse basis (since sweeping the frequency effectively decorrelates the scene) then this is perfectly possible in spiky clutter, but one further test was carried out:

1. Obtain the raw I & Q values of the 256 pulse chirp.
2. Locate the frequency number (1 - 256) of the maximum absolute value within the chirp.
3. Calculate the normalised maximum value as a ratio with respect to the mean amplitude of the chirp.
4. Plot the frequency number versus the normalised value as a scatter plot.

Note that step 4 is preferred rather than a histogram or an averaging method as the clicks are so intermittent they may not show up. Some modulation is expected, the chirp return does form a range profile FFT after all, and so the overall mean maximum frequency number value is calculated. The results from V polarisation are presented in Figure 4.10 where an obvious discontinuous area around frequency numbers 50-60 is present; single frequency bins show more than two orders of magnitude excess over the mean chirp level. The overall mean values are continuous and show that the chirp is usually well formed which could indicate that the A/D converters in frequency numbers 50-60 are intermittently contaminated with voltage glitches. It has been suggested (Branson 2000a) that these are due to intermittent interference from another 3GHz rotating system which could explain their slight time periodicity (not shown).

Without knowing the exact cause of the problem, the data has been processed by removing any chirp which has a single maximum frequency value with an RCS ten times that of the chirp average. This should not have a significant effect on the post-FFT range profile statistics as it removed only 2.5% of samples which were otherwise distributed evenly across the chirp (Figure 4.10 is clipped to show this).

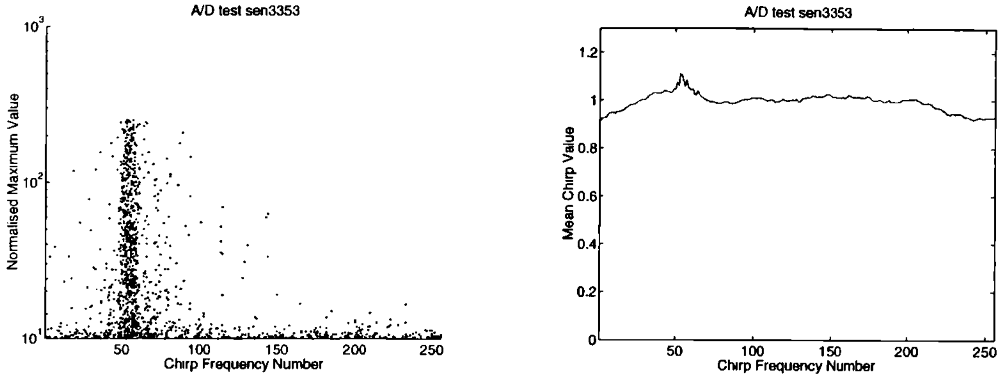

Figure 4.10: Chirp frequency bins 50-60 show erroneous intermittent large values

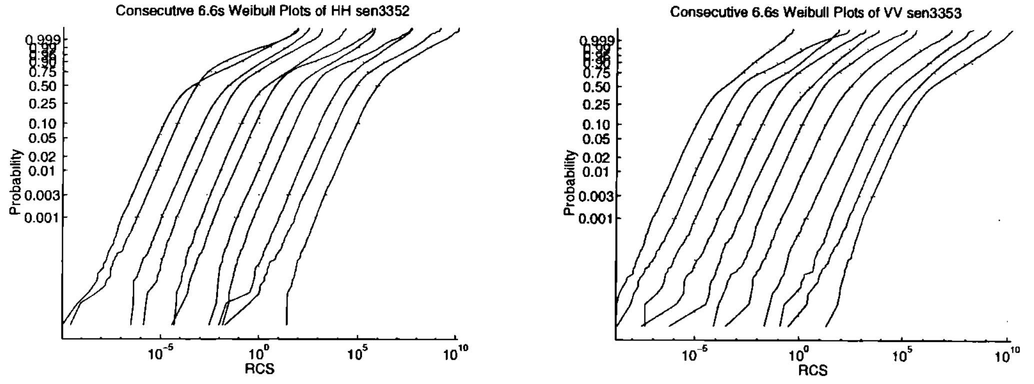

Figure 4.11: Consecutive 6.6s (displaced by 10dB) of Hi-res data shows the high intensity RCS distribution is variable

4.2.5 Determining the PDF of the RCS

Stability of the PDF for the Perfect Data

Large variations in the distribution are seen over time. One minute of consecutive 6.6s blocks of data are shown in Figure 4.11 which, when plotted on Weibull paper, demonstrates the high intensity RCS distribution fluctuates considerably whereas the low intensity is relatively stable. The low intensity returns are presumably dominated by thermal noise which explains the stability, but the high intensity backscatter is variable in its extent (from the position of the elbow on the plot) and its distribution (the gradient of the upper tail).

An exact fit to a simple distribution is unlikely based upon the Weibull plots of Figure 4.11. Distributions previously considered (Section 3.1) would appear as a continuous curve without the

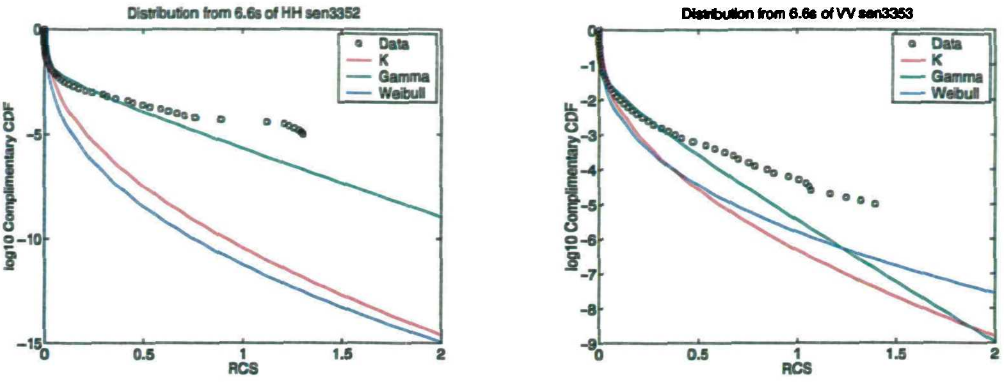

Figure 4.12: Various MLE fits for 6.6 seconds of 1.5m Hi-res data

distinct elbow, however their MLE solutions when applied to this data may be more robust and could fit the data reasonably well. By using the MLE for Weibull and Gamma with the normalised log estimate for K (Equation 3.15) the results of processing a randomly chosen 6.6 seconds of HH sen3352 and VV sen3508 are shown in Figures 4.12. The complementary CDF is plotted since it is the high RCS tail of the distribution that is crucial for setting a target detection threshold. None of the distributions have characterised the data over the RCS range with order of magnitude errors at typical threshold points, this was generally the case when processed over the entire data set.

It is clear that ocean wave events are present to such a degree as to render a 6.6 second duration distribution unstable. If the observation window is increased a stable distribution may eventually be reached however this was still not obvious over data blocks of 1 minute (not shown) which is unlikely to be obtainable in an operational radar. The intrinsic cause is that the high intensity wave events seen in Figure 4.5 are dominating the determined statistics.

The Extent of Thermal Noise

By determining the local mean (see Figure 4.14 in later sections) the events have a peak mean RCS of about 20dB above the low RCS regions. From the typical SNR of the files (10dB - 30dB) this would suggest that the low regions are dominated by thermal noise. This is confirmed by the slope of the low RCS part of the Weibull plots indicating an exponential distribution. Segmenting the images by hand showed low RCS areas to have a flat Doppler spectrum, this indicates that uncorrelated thermal noise is present over a considerable range extent of the image. This observation is discussed in depth in Chapter 6 where a statistical method of segmentation is suggested for individual range profiles.

Mixture Analysis

In the perfect 3GHz data, variation in the high RCS returns is a part of the backscatter statistics. A single Weibull PDF will be suboptimum at both low and high intensity; it is especially inaccurate if applied as part of a target detector based solely on magnitude statistics as the predictions for rare high intensity returns are then orders of magnitude in error.

The Weibull analysis in Figure 4.6 suggested a mixture model of two separate Weibull distributions could be applied for all the data, including the 8GHz slow cycled data. This is incompatible with the compound view of sea clutter, which views the underlying RCS as continuously modulating the speckle, but the unknown statistical effects from slow chirp cycling mean it is still instructive to see if single wave events can be classed in this way. Statistically this is realised by two distinct and unequal scattering populations each subject to different Weibull parameters, forming a parameterised 5 dimensional distribution, whose distribution is

$$\begin{aligned} P(I) &= p\frac{\beta_1}{\alpha_1} I^{\beta_1 - 1} \exp(-\frac{I^{\beta_1}}{\alpha_1}) + (1-p)\frac{\beta_2}{\alpha_2} I^{\beta_2 - 1} \exp(-\frac{I^{\beta_2}}{\alpha_2}) && (4.1) \\ C(I) &= 1 - p \exp\left[1 - \left(\frac{I}{\alpha_1}\right)^{\beta_1}\right] - (1-p)\exp\left[1 - \left(\frac{I}{\alpha_2}\right)^{\beta_2}\right] && (4.2) \end{aligned}$$

Graphical methods have been suggested to solve this (Kao 1959), and Falls (1970) give the necessary moments as

$$\mu_r = p\alpha_1^{r/\beta_1}\Gamma\left(\frac{r}{\beta_1} + 1\right) + (1-p)\,\alpha_2^{r/\beta_2}\Gamma\left(\frac{r}{\beta_2} + 1\right), r = 1, \dots, 5 \tag{4.3}$$

The method of moments requires a prohibitively large population for an accurate result. A numerical MLE search is suggested by Oliver (1995) when applied to a K-mixture but this is computationally expensive over 5 dimensions. Jiang and Murthy (1998) explore in detail the Hazard rate of this distribution (cf Equation 3.5) to show 8 separate cases exist, each of which can be parametrically described, this may lead to a less general estimator.

A mixed distribution is complicated by the additional number of parameters to be estimated. Numerical searches over the parameter space are required which may be expensive. A Nelder-Mead simplex search over five dimensions was carried out to fit the data to these mixtures of K and

Table 4.4: Mixture distribution fit - Kuiper statistics ρ with the mixing parameter p, hat indicates log averaging

File	$E\left(\widetilde{\rho_{WW}}\right)$	$\sigma\left(\widetilde{\rho_{WW}}\right)$	$E\left(p_{WW}\right)$	$\sigma\left(p_{WW}\right)$	$E\left(\widetilde{\rho_{KK}}\right)$	$\sigma\left(\widetilde{\rho_{KK}}\right)$	$E\left(p_{KK}\right)$	$\sigma\left(p_{KK}\right)$
sen3363	-0.763	0.973	0.511	0.177	-1.42	2.46	0.543	0.196
sen3464	-0.394	0.886	0.617	0.247	-0.720	1.16	0.574	0.301
sen3502	-0.619	1.09	0.584	0.195	-0.980	1.60	0.576	0.239
sen3470	-3.92	4.03	0.632	0.183	-4.22	3.58	0.548	0.242
sen3508	-6.58	6.22	0.556	0.151	-5.98	5.81	0.568	0.198

Weibull PDFs

$$
\begin{aligned}
P_{WW} &= pP^{\alpha}_{Weibull}\left(A\right) + \left(1-p\right)P^{\beta}_{Weibull}\left(A\right) && (4.4)\\
P_{KK} &= pP^{\alpha}_{K}\left(A\right) + \left(1-p\right)P^{\beta}_{K}\left(A\right) && (4.5)\\
P_{KW} &= pP^{\alpha}_{K}\left(A\right) + \left(1-p\right)P^{\beta}_{Weibull}\left(A\right) && (4.6)
\end{aligned}
$$

Various numerical methods exist to find a fit, least squares is a popular choice but does not give a realistic fit to the tail of the distribution. A numerical MLE failed to converge without careful choice of starting parameters, the converged result was very poor. As the Kuiper statistic is invariant over the distribution it is reasonable to minimise this and in doing so the significance level ρ_V is maximised (see Section 3.3.1).

Convergence was achieved by making an initial least squares estimate to the data using the 2 parameter distribution(s) and then initialising the 5 parameter estimate using these values and a mixing parameter p of 0.5. A least squares fit was made to the 5 parameter distribution before invoking a minimisation routine based upon the Kuiper statistic. Note that no constraints were used in the routines but an obvious one is that there should only be one root to the differential of the distribution. Results for V-polarisation Weibull/Weibull and K/K are shown in Table 4.4

Using data blocks of 1.6 seconds and using a sliding 1/4 window approximately 100 seconds of data was analysed. For every block the mixture parameters were determined and the Kuiper statistic ρ_V calculated. Defining $\widetilde{\rho}$ $=\log_{10}\left(\rho\right)$, the expectation of this $E\left(\widetilde{\rho}\right)$ and the standard deviation $\sigma\left(\widetilde{\rho}\right)$ are determined over the entire set. The mixing parameters are calculated normally as the expectation $E\left(p\right)$ and the standard deviation $\sigma\left(p\right)$. An example WW fit for one block of sen3363 is shown in Figure 4.13 which had an excellent fit to the higher amplitude data with overall ρ_V of 0.7.

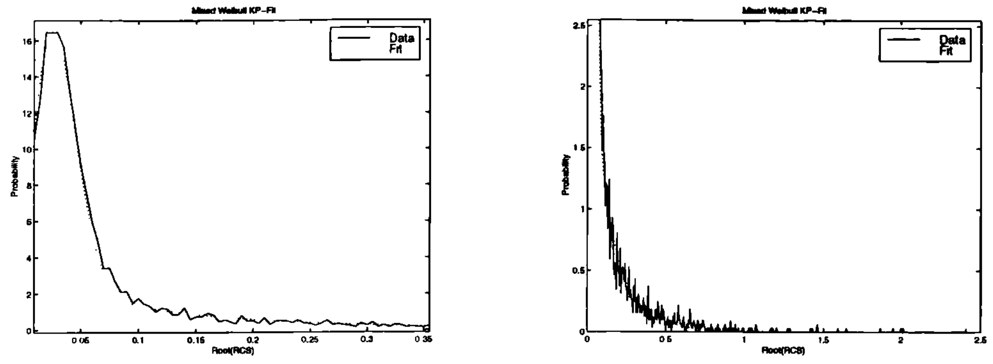

Figure 4.13: A mixed Weibull fit via minimisation of ρ_{KP} for 1 second sen3363

The fit is good for the 1.5m resolution but there is no overall fit to the 19cm data. This may seem in conflict with the straight lines and elbows seen in Figure 4.6 but it could show that the statistics vary in time, and over 100 seconds the form of the distribution deviates significantly from what can be achieved by a numerical search.

Whilst the fit to the data can be very good for the Weibull mixture, one must be cautious in assigning any significance to this due to the number of free parameters. Attempting to form the fit upon longer duration windows is less successful, indicating further heterogeneity or an incorrect probability model for large windows.

The extension from two to five dimensional parameter space is large but the reasonably constant mixture value over time ($\sigma(p)$ is low) indicates some stability. It would be interesting to see if greater stability could be attained by constraining p to vary slowly. Justification for this model could be attempted by partitioning the scenes in accordance with the mixing parameter p based upon contours set in intensity.

The statistical mixture analysis takes no account of the spatial extent of the mixture, the usual way of determining this is to use some form of windowed estimate. An arbitrary choice of some power-2 window is used for this type of analysis and in the subsequently developed detection schemes; Figure 4.14 shows a 16 by 16 sliding window in range and time (24 metres by 0.1 seconds) from the first 6 seconds of sen3353V, corrected for pulse shape. Using the expected errors in estimating U from Section 3.1.4 and assuming that the mean is accurately calculated for this window size, analysis of sen3352 and sen3353 shows:

- The mean RCS has large areas of homogeneity with definite wave event steps of the order

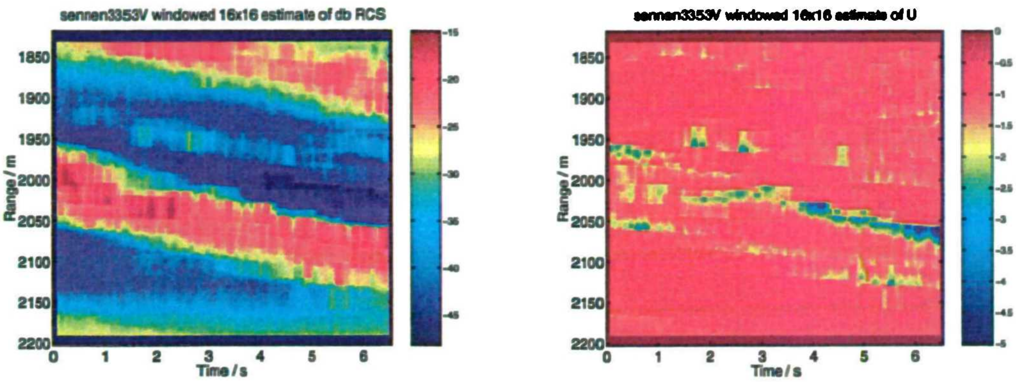

Figure 4.14: Arbitrarily chosen 16 by 16 sliding window estimate of the spatial form of sen3353 in terms of mean RCS and shape estimate U. The lowest value of U is caused by an abrupt edge in the RCS

of 20dB present.

- A particular wave event appears to have a stable mean RCS over the entire imaged time.
- The areas of low RCS were characterised as exponentially distributed well within error bounds.
- The areas of high RCS could often be characterised as exponential, but were usually spikier.
- The lowest U values were not from homogenous areas, but due to the wave event causing an edge effect upon the windowed U estimate (Section 3.1.4).

Overall, a particular range profile is often dominated by a single wave event and any further modulation is hidden by this. Reducing the window size so as to lessen the edge effects gave similar results but with increased uncertainty in the distribution of the high RCS areas.

4.2.6 Conclusions from High Resolution Analysis

Preliminary high resolution analysis has shown:

- Numerous problems with the data have been demonstrated, which without careful analysis and rejection could give spurious results:
  - The slow chirp cycling in the 8GHz data will smear any fast moving scatterer, but this could occur operationally for fast moving observer or target.

  - The intermittent voltage ‘glitches’ do not cause a major problem after careful filtering to remove them from the frequency domain before the image is formed, but they could give hugely erroneous results for the Doppler statistics in the next Section. If they are from another 3GHz radar then this may occur operationally and could even arise purposefully in a jamming situation.
  - General bit errors and poorly chosen attenuation are present in many data sets which would not occur operationally.

- Ocean wavefronts are present in those scenes with acceptable chirp cycling, producing separable areas of varying correlation in both range and time.
- As observed in many papers, horizontally polarised returns have a greater probability of high intensity returns. Commonly HH is found to be spikier than VV but without perfect alternately polarised data detailed analysis has not been performed.
- The commonly chosen Weibull and K distributions are unstable and offer a poor fit both locally and globally regardless of the observation time. Using observation windows of 1.6s, 6.6s and one minute upon the high quality sen3352 and sen3353 gave highly unstable fitted distribution parameters over the 20 minute files.
- A statistical search for a binary mixture of Weibull distributions fits the 1.5m resolution data extremely well on 1.6s observation windows; the 19cm data is not fitted well but the greater number of pulses in the chirp caused an especially slow chirp cycle time with expected image degradation.
- The proposed binary mixture is reinforced by the large homogenous areas seen in the windowed analysis. Although edge effects from the 20dB step are considerable, a large low RCS spatial area of the data is thermal noise whereas the higher RCS area is spikier.

Although results are based upon relatively few acceptable data sets, they indicate that the high resolution statistics cannot be characterised locally or globally by a simple distribution. Statistical analysis assuming a time varying binary mixture of heterogenous distributions, assessed on the order of 1 second was the only method that gave acceptable fits. Forming the binary mixture needed computationally expensive minimisation routines but segmentation ‘by eye’ suggested that

Table 4.5: Low res Doppler files used to assess RCS PDF. H_w is Significant Wave Height, Az refers to Azimuth Angle

Run	F/GHz	PRF	Pol	H_w/m	Az_{wave}	Az_{look}	Rng/m
sen2029	8.0	10k	HH/VV	3.30	310	284	1005
sen2030	8.0	10k	HH/VV	3.30	310	331	1500
sen2035	8.0	10k	HH/VV	3.69	319	284	1005
sen2036	8.0	10k	HH/VV	3.69	319	284	1500

a large extent of the scene consists of system noise. Later Chapters of this thesis will implement a fast method of determining the boundaries of the mixture that confirms this Section's results.

The conclusion is that increased observation time does not help in determining the local distribution when the presumed application will be to assess small areas of the RCS data for the presence of targets. Indeed, this is present in the compound theory of sea clutter where the description of the local distribution is exponential but the global distribution is K. A method of determining the local distribution is required that specifically includes the effects of the high intensity edges from the wave events.

4.3 Low Resolution Doppler Data

The raw intensity returns from the Doppler data are received from a larger range cell - 150m in range extent, two orders of magnitude larger than the high resolution data. From the scattering theory of Section 2.1.1 this will contain more scatterers and thus the PDF might be expected to be closer to exponential. The Doppler returns do not have the constraint of a minimum chirp cycling time as in Hi-res and consequently a greater number of acceptable files are available. Upon forming the Doppler spectrum it is possible not only to examine the form with respect to the theory of Section 3.2.1, but also to determine how the RCS distribution varies with respect to velocity.

4.3.1 8GHz Files

Four alternately polarised and two VV polarised 8GHz Doppler files are used as an initial example, environmental and system parameters are given in Table 4.5; wind data was not recorded.

Plots of both HH and VV intensity are given in Figure 4.15, both polarisations seem to be recording

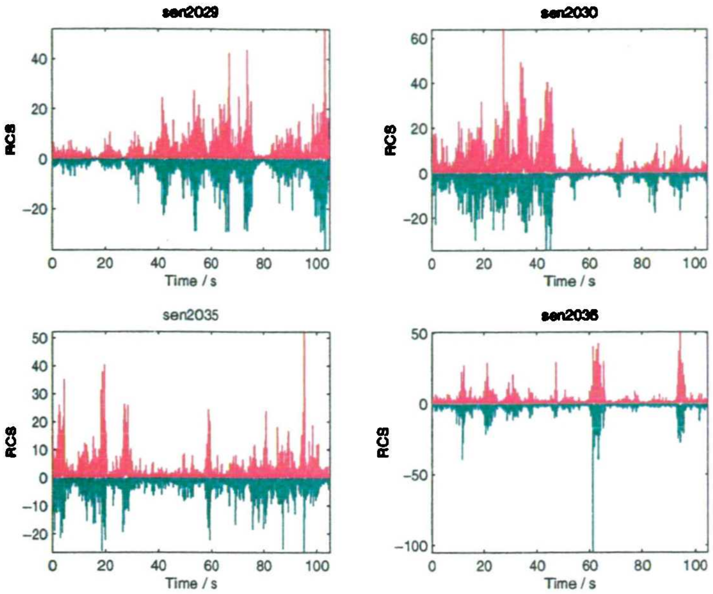

Figure 4.15: RCS magnitude plots (m^2/m^2 of 8GHz Doppler data (HH polarisation in red, VV polarisation in green - inverted for clarity so please ignore sign)

the same intensity structure. There are discernible areas of differing intensity over the order of seconds; from the Hi-res analysis these are probably due to the ocean wave events. Additionally there are shorter instances of high intensity returns commonly termed 'spikes' and file sen2030 shows a distinct step in the long term RCS at about 45 seconds. The large individual return in sen2036 VV polarisation is most probably due to faulty A/D conversion, analysis at such a high PRF means that the speckle is correlated to some extent and this return was too abrupt.

Similar to the Hi-Res analysis, Weibull plots of the data are used and consecutive 6.6s of alternate 8GHz data is shown on Figure 4.16. Compared to the high resolution plots of Figure 4.6 the elbow is immediately less obvious in the low resolution files and there is much less variation over time and between polarisations. Testing many data sets confirmed that, in general, the Weibull plots of the Doppler data gave less justification to attempt a mixed Weibull fit.

Some papers have found the Doppler spectra to show significant variation between polarisations. In particular, Ward et al. (1990) reported that HH mean Doppler was always larger than VV regardless of wind direction in high resolution. For a large number of low resolution data files this was found not to be the case; the upper plots of Figure 4.17 show not only the changing

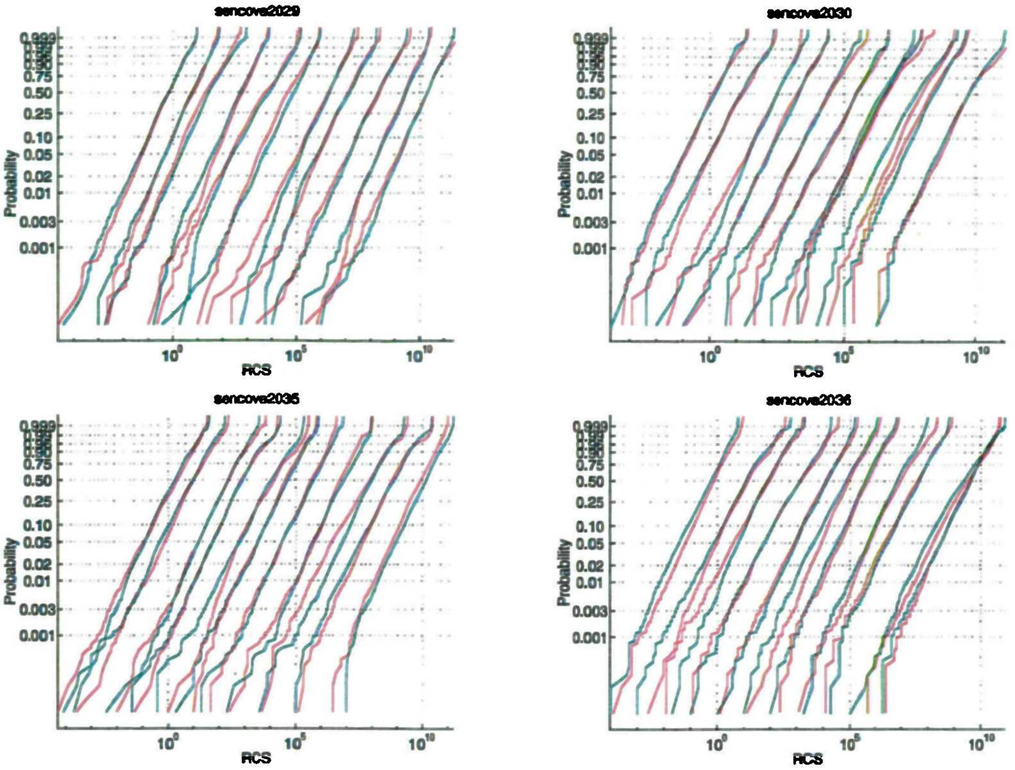

Figure 4.16: Weibull plots of consecutive 6.6 seconds of 8GHz low resolution data displaced by 10db (HH - Red, VV - Green)

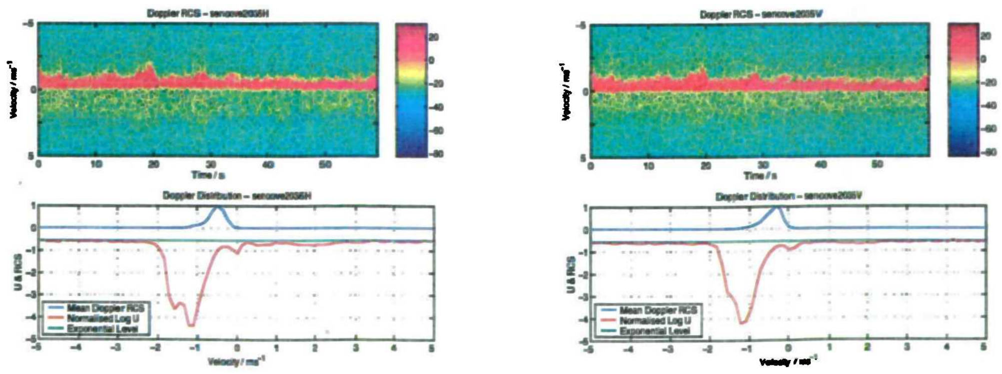

Figure 4.17: Low resolution 8GHz sen2035 file. Upper plots show HH and VV have similar Doppler spectra. Lower plots show U estimate vs velocity with normalised mean linear RCS

Table 4.6: Low res Doppler files used to test variability. H_w is Significant Wave Height, V_w is Wind Velocity, Az refers to Azimuth Angle

Run	F/GHz	PRF	Pol	H_w/m	V_w/ms^{-1}	Az_{wind}	Az_{wave}	Az_{look}	Rng/m	Notes
sen3537	3.0	10k	VV	6.09	32	290	284	331	4515	Events
sen3558	3.0	10k	VV	6.08	31	320	286	284	4515	Normal

sen2035 Doppler spectrum over time (1 minute) but also that the HH and VV polarisations give near identical spectra. The normalised log estimator was applied in each velocity bin and shows a strong dependence upon velocity, this is plotted simultaneously with normalised linear RCS in the lower plots of Figure 4.17 using a 1 second sliding window analysis. It would appear from this that the tails of the Doppler spectra have a spikier distribution; however from the upper Doppler image an approximately 40dB instantaneous SNR can be seen, the shape of the Doppler spectrum is not stable over time and so in effect a 40dB step is seen intermittently. Section 3.1.4 shows that in locally exponential noise this intensity of step only needs to be present for less than 1% of the scan time to cause a measured U as low as that seen here which explains some of the results.

- A simple stationary distribution shape estimate such as U may not be applicable upon individual Doppler velocity bins since the presence of the such velocities is intermittent and the samples are far from being Independent Identically Distributed (IID).

4.3.2 3GHz Files

Files sen3537 and 3558 were as similar as possible, taken on the same day 3 hours apart. File sen3537 has identifiable intermittent Doppler events attributed to birds (TW Research Ltd. 1999) or wind 'whipping off wave crests' (Money et al. 1997b) the sea state of these files is shown in the video still of Figure 4.1. Usually birds can be distinguished because they tend to change their direction and have a preference for flying against wind when feeding upon the surface.

TW Research reported that whilst sen3537 showed spiky character attributed to birds, sen3558 had a K-distribution parameter ν which was '*surprisingly low (i.e. the clutter is spiky) and needs further investigation*' calculated using shape estimator V or U. This is performed in Figures 4.18 and 4.19 where, as in the previous Section, the U value is calculated for each Doppler velocity bin. Additionally the proportion of the signal above system noise is plotted, calculated via a Sequential Edge Detector algorithm discussed in the next Chapter.

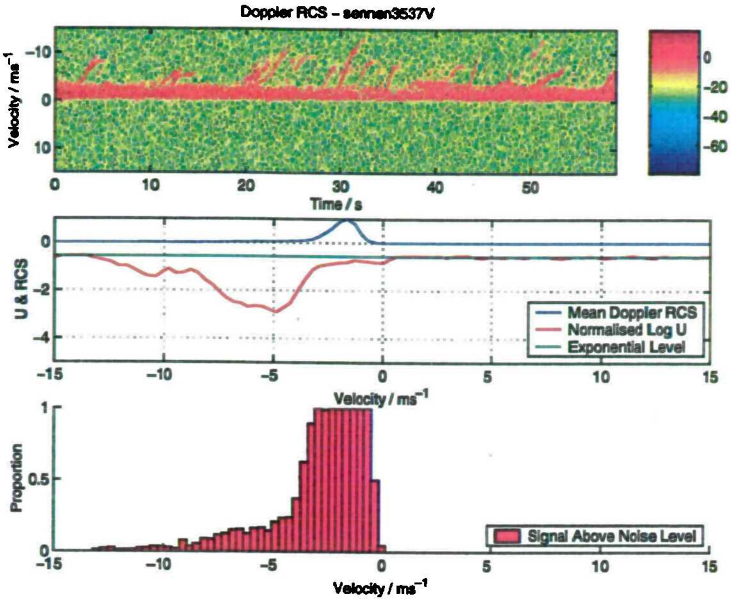

Figure 4.18: Wind caused events are probably responsible for the discrete fast scatterers seen in sen3537

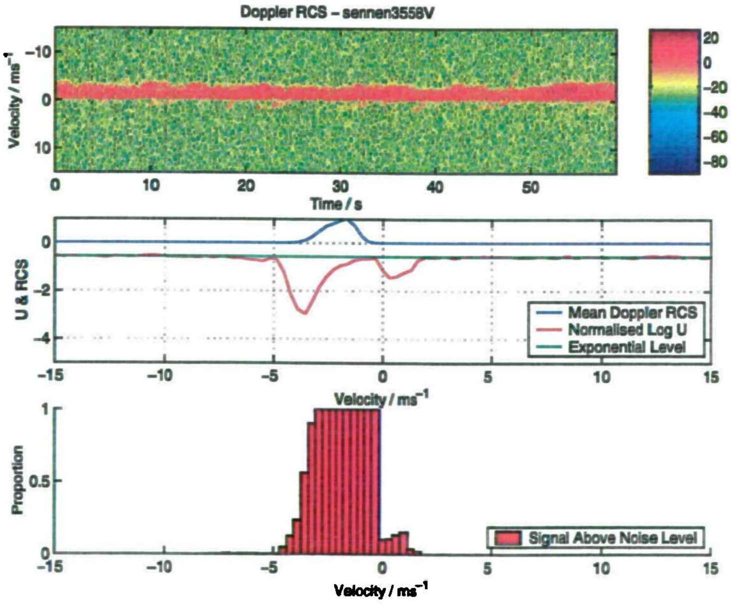

Figure 4.19: The non-stationary Doppler spectra is the primary cause of low U in sen3558 *not* homogenous spiky statistics

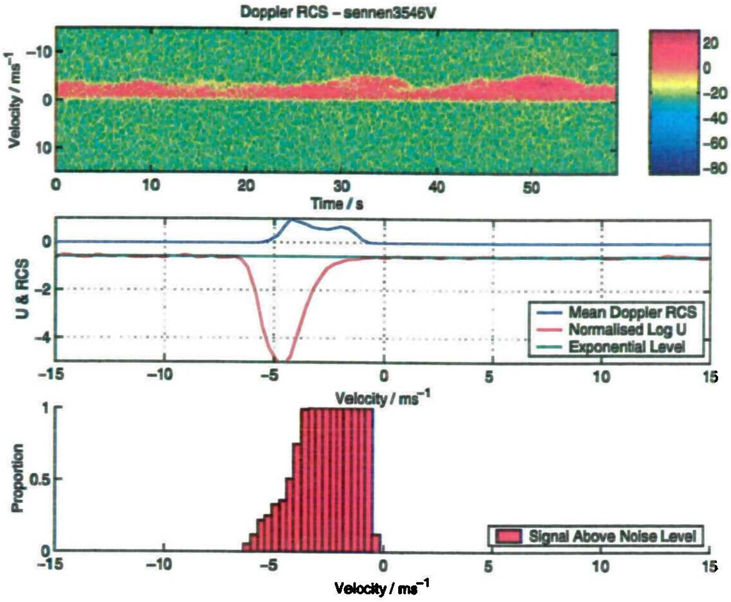

Figure 4.20: sen3546 shows extremely low U. This is due to the intermittent rolling of the waves rarely above system noise

From analysis of this data, it is suggested that the excessive spikiness in the tails of the Doppler spectra as observed by TWR and Baker (1991) are not due to homogenous statistics but by a modulation so slow that it appears as intermittent discrete intensity steps in a particular Doppler bin. This is known because the Doppler bin with the lowest U has a signal above the system noise for only about a fifth of the analysis time - from Section 3.1.4 this was the proportion required to cause maximum effect upon U when a single intensity edge was present in exponential noise.
To further emphasise this, Figure 4.20 shows the spikiest file encountered (i.e. lowest determined U) is just rolling waves of high RCS relative to system noise.

4.3.3 RCS PDF Analysis

Out of the few files analysed, Figures 4.16 & 4.15 shows that sen2036 has the least variation in form over time. Fits are performed to a Weibull, K, exponential and gamma using the normalised log estimator for K, and the MLE otherwise, upon 100 seconds of data. The complementary CDF for both H and V is shown in Figure 4.21. The K distribution shows a closer fit than any other but is still two orders of magnitude in error at typical threshold levels.

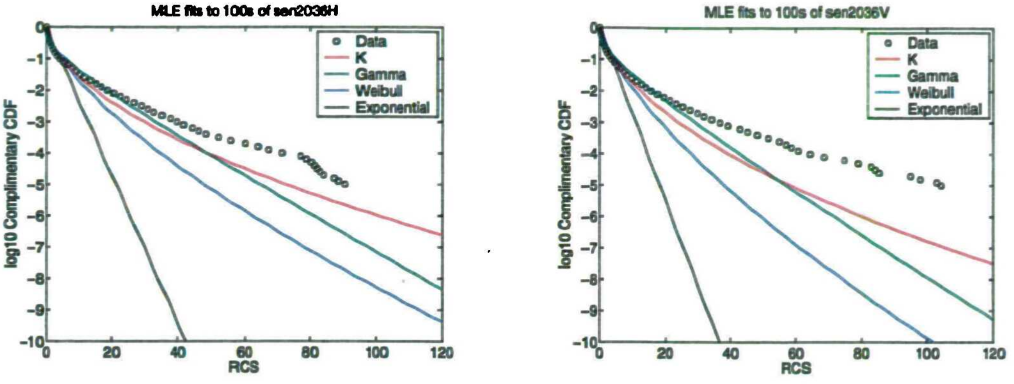

Figure 4.21: Various MLE fits to 8GHz sen2036 low resolution data

In a final effort, a statistical spatial mixture of exponential and K-distributed clutter subject to coherent noise was used to fit the data

$$P(I) = p\frac{1}{\alpha_1}\exp(-\frac{I}{\alpha_1}) + (1-p)\,P_{K+N}\,(I) \tag{4.7}$$

where P_{K+N} was calculated using a numerical inverse Laplace method covered in Section 5.4. The obtained fit could often fit the distribution well (expected due to the number of free parameters) but was not a consistently reliable estimate of the critical tail regions.

4.4 Summary

This Section has provided evidence that the statistics of backscatter returns are nonstationary in many ways. The Hi-res range profiles of Figure 4.5 immediately suggest at least a 2 component scene present, strengthened by the elbows seen in the Weibull plots of the data. The poor robustness of MLE methods has been shown in cases where the *a priori* model is not met. The discontinuity in the Weibull plot motivates mixture analysis which obtained an effectively exact fit to the distribution, however with 5 variable parameters this is not surprising and could be achieved by many such distributions.

System noise is present over a significant spatial portion of the range-time scenes, which is presumably due to large waves shadowing this region. The effect of system noise is also seen in the Doppler data where the varying proportion with respect to velocity affects the determined statistics considerably.

Doppler analysis provides further evidence of the nonstationary hypothesis. Whilst H and V are more similar than expected, the RCS statistics are shown to vary in magnitude and distribution not only between data sets but also over a timescale of the order of seconds.

The observation that commonly applied distributions were in error by up to 5 orders of magnitude in the tails of the RCS distribution casts serious doubt upon the use of a single *a priori* distribution to describe clutter. This is in stark contrast to the literature which seems obsessed in basing its detector performances upon known, fully specified stationary clutter statistics. Fast adaptation to the data is necessary so as to avoid excessive false alarms or missed detections.

Inhomogeneous nonstationary statistics require a different analysis method. After discussion of statistical detection schemes in the next Chapter, explicit nonstationary detection schemes are suggested for real data in Chapters 6 and 7.

As a final note to this Chapter, it is worth quoting Kazakos and Kazakos (1990) from their definition of the rules of statistical decision theory:

> *Careful experimental control to assure that the total observed outcome [from a realisation of an underlying stochastic process over some time T] represents a realisation from the same physical phenomenon or stochastic process, rather than a mixture of partial outcomes from different processes.*

Chapter 5

Statistical Detection Schemes

Analysis of data in order to detect a deviation from the assumed statistics is a common task within radar. By rigorously defining the probability distribution it is possible to give a detection scheme robustness and accurately determine its efficiency. This Chapter explores previously applied point target detectors and an attempt is made to generalise this to a Sequential Edge Detector which simultaneously flags point targets and edges within exponential speckle.

A numerical inverse Laplace transform is demonstrated to accurately determine the arbitrary N-fold sum of a K-distributed variate embedded in noise.

5.1 Assets Available

Assuming a number of samples beginning at time m and ending at time n, the following assets are available upon which to base a detection scheme:

1. A realisation $x_{m:n} = \{x_m, x_{m+1}, \dots , x_n\}$ from the radar output.
2. A null hypothesis H_0 indicative of the expected target free environment.
3. M parametrically known hypotheses H_i ; $i = 1, ..., M$ chosen from theory or observation.
4. A set p_i of a priori probabilities on the M hypotheses determined from previous observation.
5. A set c_{ki} ; $k, i = 1, ..., M$ of penalty coefficients to inhibit incorrect decisions.

6. A measure of performance based on probability of detection (P_d), probability of false alarm (P_{fa}), a response or efficiency measure, and time to false alarm τ_{fa}.

5.2 Constant False Alarm Rate Processors

Limiting the number of false alarms that occur gives robustness to any detection scheme. Commonly a Constant False Alarm Rate (CFAR) Processor is achieved by defining a fixed probability of miss-classification in *homogeneous statistics.* Any deviation from homogeneity over the observation time will severely decrease the effectiveness of these schemes. An edge, where both the power and type of statistics can vary abruptly, results in a finite response time during which the detector is lowered in performance.

A common performance measure is CFAR loss - the increase in target SNR required to maintain P_d and P_{fa}; the reduction in P_d in comparison to that achieved from complete knowledge of the statistics is another possible measure. This takes no account of inhomogeneity and the loss can be negative when the correlation is known. With reference to the previous Chapter more appealing measures are:

1. Sensitivity of the processor to errors in determining the clutter characteristics.
2. Response of the processor to abrupt variation in clutter characteristics (edges).
3. Occurrence of target masking due to edges or multiple targets.

Historically CFAR processors are of low complexity, their origins lying in acoustic delay lines. Emphasis is placed on keeping these processors simple despite computational power rising by orders of magnitude since their conception. The CFAR processor should be at least as complex as the method of determining clutter characteristics when this is being calculated concurrently. Commonly the CFAR is applied to the raw data but if the PDF from a filter output is known the CFAR can be applied to this.

5.2.1 Neyman-Pearson Tests

Utilising the following assets:

- A realisation $x_{1:n}$ from the radar output with fixed n.

- Hypothesis H_0 indicative of no target
  versus hypothesis H_1 indicative of target present.
- A measure of performance based on P_{fa} and P_d.

If the H_1 hypothesis is fully specified and selected in advance as the most important criterion, maximisation of P_d for fixed P_{fa} yields the Neyman-Pearson (NP) detection scheme (Kazakos and Kazakos 1990). When H_1 is not fully specified (unknown target strength or distribution) then the NP scheme cannot in general be realised (Kazakos and Kazakos 1990). Assuming a class of PDF for H_0 defines a *parametric* detector that given $x_{1:n}$ has to decide on parameters such as variance or scale. Estimating the PDF from $x_{1:n}$ or using a technique independent of PDF defines a *non-parametric* class of detector. Increased performance is achieved by the parametric detector as more prior information is included.

Usually a CFAR processor is parametric, formed by setting a threshold λ upon a test sample x so as to maintain a constant P_{fa}, the detection rule being

$$\begin{aligned} x &\geq \lambda \quad H_1 \text{ Declared} \\ x &< \lambda \quad H_0 \text{ Declared} \end{aligned} \tag{5.1}$$

Whilst one could set a fixed threshold from full knowledge of H_0, Finn and Johnson (1968) proposed an adaptive threshold that can follow a slowly varying change in the background mean which can dramatically improve P_d. Commonly a threshold is set in signal magnitude based upon a number of reference cells surrounding a test cell, this defines an analysis with window size W (the previously, similarly defined, shape estimator of Chapter 3 is no longer considered). P_d is improved with increasing W under homogenous conditions but consequently the response time of the detector to inhomogeneity is lengthened. Specifying W defines a correlation length for the statistics of which no *a priori* knowledge is usually available.

A target is determined when the test cell exceeds the threshold derived from the reference cells. The false alarm rate is found from

$$P_{fa} = \int_0^\infty \left(1 - C_X(\alpha x)\right) p_{Z|X}(x) dx \tag{5.2}$$

where $C_X(\alpha x)$ is the CDF of the probability distribution X at the point αx and $p_{Z|X}(x)$ is the PDF of the test statistic Z given X. The threshold is the value $\lambda = \alpha Z(x_{1:n})$. The type of CFAR

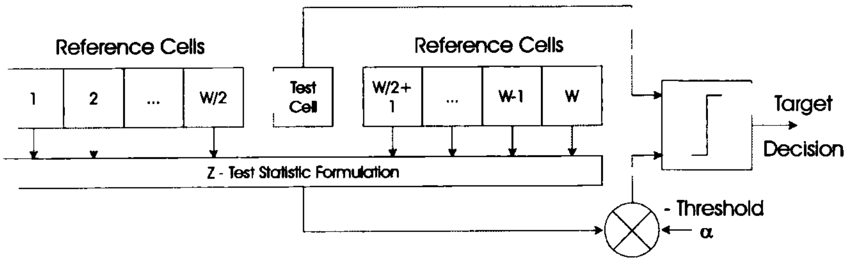

Figure 5.1: Generalised CFAR Process

is determined by the test statistic, a reference window of size W is assumed to lie either side of the test cell in these examples:

Cell Average (CA)

$$Z_{CA} = \frac{1}{W} \sum_{i=1}^{W} x_i \tag{5.3}$$

Cell Average Greatest Of (CAGO)

$$Z_{CAGO} = \frac{1}{W} \max \left\{ \sum_{i=-1}^{-W/2} x_i, \sum_{i=1}^{W/2} x_i \right\} \tag{5.4}$$

Order Statistic (OS) of order k

$$Z_{OS,k} = x_{(Wk)} \tag{5.5}$$

$$\left\{ x_{(W1)} \leq \cdots \leq x_{(WW)} \right\}$$

Thus the windowed CFAR uses a detection rule

$$\begin{aligned} x &\geq \alpha Z \quad \text{Target Declared} \\ x &< \alpha Z \quad \text{Target Absent} \end{aligned} \tag{5.6}$$

where the threshold multiplier α is a constant for a particular test and distribution, chosen so as to achieve a fixed P_{fa} in homogenous conditions.

Numerous modifications of these forms exist to counter mixtures of clutter edges, multiple targets and heterogenous environments some of which are discussed in Section 5.5. Calculations necessary to formulate CFAR upon various IID statistics are given in (Minkler and Minkler 1990) and

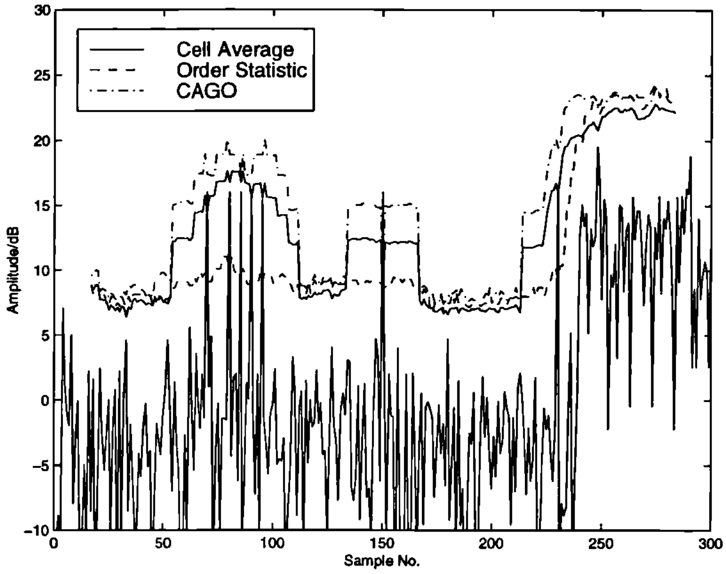

Figure 5.2: Example of NP windowed CFAR; after Armstrong(1992)

particular emphasis is placed upon the K distribution by Armstrong (1992). The CA detector is optimum in Rayleigh clutter for the detection of Swerling I targets (Gandhi and Kassam 1994). Whilst CAGO is often said to offer resistance to edges, a large edge ratio will cause a significant target loss over $W/2$ samples before the edge occurs as even one sample from the higher side will dominate the statistics. In the literature CFAR is usually taken to mean fixed window CFAR; however a constant false alarm rate can be achieved by several types of processors detailed in Appendix C.

5.3 Application of CFAR

5.3.1 Fixed Window CFAR

The commonly applied detection test is fixed window Neyman Pearson (FWNP) CFAR. A simulation of IID K-distributed data of shape $\nu = 0.5$ is used to test CA, CAGO and OS for a window size of 32, shape parameter is assumed known *a priori*. Using 300 samples, multiple 16dB point targets have been added at samples 70-95, isolated target at 150 and a potentially masked target at 230 by the 14dB edge at 240 onwards. WIth reference to Figure 5.2:

- CA has the lowest threshold, therefore greatest P_d in homogeneous clutter. Multiple targets cause masking with some going undetected. The target near the clutter edge is only just detected and abrupt changes in the threshold are seen.
- CAGO has the greatest threshold in homogeneous areas so lowest P_d. It is extremely sensitive to multiple targets - none of which are detected as the detector views them as an edge. The target near the clutter edge is completely masked and very abrupt changes are seen in the threshold.
- OS has a threshold midway between CA and CAGO, no target masking is seen - however this is due to judicious choice of settings and more multiple targets or an extended target would cause problems. Threshold changes are gradual.

Common to all algorithms is an initialisation time equal to that of the window size, edge effects will be apparent at the end of the data since we only have a one sided window to compare.
To specifically overcome some of the missed targets it is necessary to identify the causes, these are deemed to be:

- Self masking - a large persistent target is present for more than one sample.
- Edge masking - an edge is close to the cell under test and pollutes the window.
- Multiple target masking - targets artificially raise the mean value of the window.
- Edge - when a true edge is seen in the data a response time is seen, or the abrupt change is at the wrong point.

5.3.2 Choice of Window Size

The justification for choosing a particular window size W does not seem to be emphasised in the literature, obviously a power of 2 is preferable for fixed memory storage and one would like W to be short enough to be homogenous, whilst long enough to give a reliable estimation of the mean. In exponential noise one can easily compare the required threshold multiplier α normalised to the mean for varying window sizes at constant P_{fa} since the distribution of the sum of N exponential variates is easily calculable. In general this is not the case in the K and Weibull distributions but a numerical solution is proposed in the next Section.

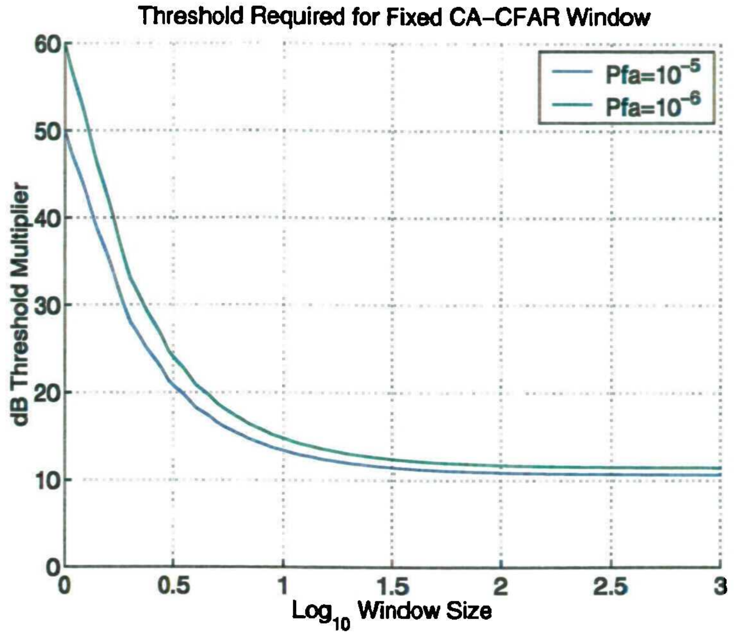

Figure 5.3: A factor determining the size of fixed window CFAR is the threshold multiplier required for a constant false alarm

Figure 5.3 shows why in exponential noise a window size of 32 is usually chosen - little benefit is seen above this due to the shallow curve. The greater variance of spiky clutter means that the window size required to reach this point (on a presumably similar shaped curve) will be much greater suggesting that the window size should be a function of the observed statistics for similar performance under variable clutter distribution.

Intuitively one would like to match the window size to the largest homogenous region local to the test cell - which is some function of the correlation length of the statistics. An exponential decay autocorrelation function (ACF) is often assumed for simplicity and has been observed in SAR clutter (Lombardo and Oliver 1995) who also show the effect of the distribution upon the ACF. A large discontinuity in the data and the presence of targets could easily pollute the ACF.

5.4 Clutter Plus Noise Distributions

The intensity PDF of K-distributed noise in the presence of zero-mean Gaussian noise has been shown (Watts 1987) to be given by

$$\begin{aligned} P_{C+N}(I) &= \int_0^\infty \frac{1}{\Gamma(\nu)}\frac{\nu}{\mu}\left(\frac{\nu\sigma}{\mu}\right)^{\nu-1} \\ &\times \exp\left[-\frac{\nu\sigma}{\mu}\right]\frac{1}{\sigma+a}\exp\left[-\frac{I}{\sigma+a}\right]d\sigma \end{aligned} \tag{5.7}$$

where ν and μ are the shape and mean of the K distribution, with the noise having a variance of a. It may be that the extent of thermal noise is known, or that a Swerling II target model is to be evaluated but the primary task is to evaluate the CDF at a particular value of intensity. In addition, for an ordinary CFAR the output is required from N fold convolution of this PDF for arbitrary parameters.

Numerical evaluation of Equation 5.7 will require very small integration steps if a Pfa threshold is to be set, it will become progressively more difficult to maintain accuracy against processing time considerations for N fold numerical convolution. Monte-Carlo simulation is a viable option when numerical accuracy is a problem but it is time consuming for reasonable values of N.

Convolution by FFT can be implemented, but as the mean of an N fold convolved PDF will be N times the original, one must start with an array much longer than required to satisfy the accuracy of the original PDF. For reasonable values of N and Pfa of the order of 10^{-5} there is little or no benefit over Monte-Carlo simulation due to the number of times Equation 5.7 must be evaluated in conjunction with an unreasonably large FFT.

Performing the *analytic* transform of Equation 5.7 to the frequency domain reduces some of the FFT processing constraints but accuracy may be a problem when raising to large powers of N in performing the convolution.

This Section investigates the possibilities offered by numerical Laplace transforms. This has not been considered in the radar literature and is rarely applied in the general literature due to its reputation for being unstable - largely undeserved as shown by d'Amore et al. (1999).

5.4.1 Frequency Domain Representation

Equation 5.7 for $P(I)$ requires numerical evaluation in σ, however performing the frequency transform

$$\widetilde{P}_{C+N}(s) = \int_0^\infty \exp(-sI)\, P(I)\, dI \tag{5.8}$$

and reversing the order of the integration it can be shown using Mathematica (Wolfram 1996)

$$\begin{aligned} \widetilde{P}_{C+N}(s) &= \int_0^\infty \frac{\exp\left(-\frac{\nu\sigma}{\mu}\right)\left(\frac{\nu\sigma}{\mu}\right)^\nu}{\sigma(1+as+s\sigma)\Gamma(\nu)} d\sigma \\ \operatorname{Re}\left[s+\frac{1}{a+\sigma}\right] &> 0, \operatorname{Re}\left[-\frac{1}{a+\sigma}-\frac{as}{a+\sigma}-\frac{s\sigma}{a+\sigma}\right] < 0 \end{aligned} \tag{5.9}$$

where, accepting the constraints, the integration with respect to σ can now be performed, and using Abramowitz & Stegun 6.1.17 (1972) is simplified to

$$\begin{aligned} \widetilde{P}_{C+N}(s) &= \frac{\exp(Y)\, Y^\nu \Gamma(1-\nu, Y)}{1+as} \\ Y &= \frac{(1+as)\nu}{s\mu} \\ \operatorname{Re}[\nu] &> 0, \arg\left[\frac{1+as}{s}\right] \neq \pi \end{aligned} \tag{5.10}$$

The transform in Equation 5.8 is recognised as the Laplace transform and the conditions are all met when the inverse is evaluated using the Bromwich integral

$$P(I) = \frac{1}{2\pi i} \int_{c-i\infty}^{c+i\infty} \widetilde{P}(s) \exp(sI)\, ds \tag{5.11}$$

Standard textbooks (Boas 1983) show how to evaluate this integral in terms of residues of the poles. In general this is only possible for simple transforms with finite partial fraction expansions and so a method due to de Hoog (1982) was implemented that takes a minimum amount of information about the form of $\widetilde{P}(s)$ and is relatively fast:

1. Equation 5.11 is discretised using a trapezoidal rule.
2. An analytic Fourier expansion is performed upon the approximation.
3. Successive Padé approximations are evaluated using a modified quotient-difference algorithm to improve convergence and speed.

Existing implemented algorithms (Hollenbeck 1998) attempt to invert the entire CDF at once for speed, however this degrades accuracy at low Pfa and so the intensity values are inverted individually. Further improvements with numerical tests on a variety of functions have been presented recently (d'Amore et al. 1999) which allow lower Pfa to be evaluated.

The frequency domain is particularly advantageous. Prior to the inversion, the PDF of the sum of N pulses is achieved by raising $\widetilde{P}(s)$ to the power N and the CDF of the distribution is obtained through division by s. Obviously some computational numerical limit is present when raising to large powers and in the necessary evaluation of limiting large and small numbers using standard arithmetic on a 32 bit system.

Typically evaluating the Laplace inversion for a single intensity at any value of N for a Pfa of 10^{-6} required $O(10^5)$ floating point operations (flops), a Pfa of 10^{-6} requires a minimum of about $N \times 10^8$ flops for the Monte Carlo samples to be generated (with substantial statistical noise present). This is a huge time saving.

Evaluation of the CDF for $\nu = 0.1, 0.2$ and 0.5 is performed as lower values represent the distributions that are difficult to approximate; additionally performed in the presence of 0dB, 10dB and 20dB noise for $N = 1, 10$ and 100 pulses in Figure 5.4 where the distribution is normalised to unit mean intensity. Each plot takes about a second to evaluate on a desktop PC.

This algorithm is stable up to a numerical Pfa limit of about 10^{-7}. The accuracy was confirmed through Monte Carlo simulation and some improvement may be available through an optimised algorithm or the use of variable precision arithmetic - however this will cause a substantial speed penalty. Alternatively an intrinsic 64 bit system such as a DEC Alpha or a modern Sun system could be used.

A particular filter output may be the result of a number of weighted inputs (incoherently summed) drawn from different K + noise distributions which can be easily incorporated into the above scheme by multiplying their respective transforms prior to inversion.

The benefit of using this method is that generation of arbitrary output PDFs is significantly quicker compared to Monte Carlo when N is large. In cases where the Laplace transform is not analytically available it can be evaluated as a series expansion; this is perfectly adequate since in practice *all* special functions are evaluated in a similar manner.

Although not shown, the sum of N Weibull distributed variates can be calculated although convergence problems are more difficult for certain shape parameters. The coherent inclusion of noise is

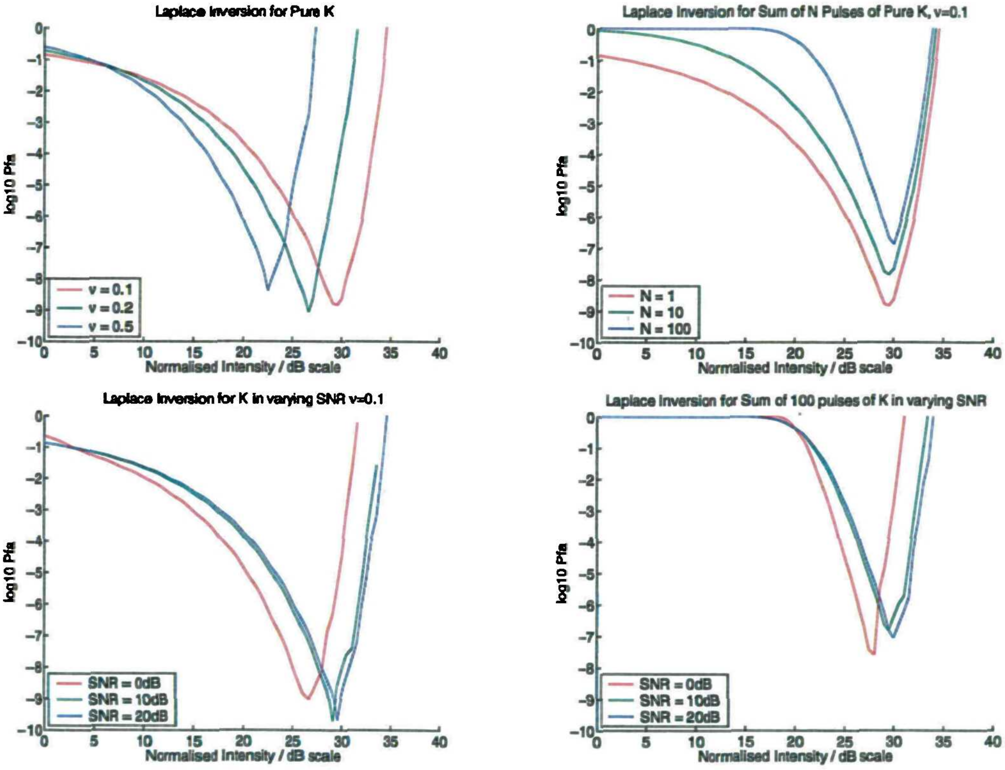

Figure 5.4: The numerical Laplace transform is accurate to a Pfa of 10^{-7} for $N < 100$ pulses of a K distribution in varying SNR. The abrupt discontinuity indicates the numerical accuracy limit. Each curve took the order of seconds to evaluate

not as simple as the K distribution and is the subject of an ongoing DERA contract. The numerical inverse Laplace transform is certainly applicable to many statistical problems encountered in radar.

5.5 The Problem of Simultaneous Target and Edge Detection

All fixed window algorithms use the *a priori* assumption that the window size is a function of the 'coherence length' of the data. For target detection the target test sample is almost always compared with respect to the samples either side of it. The expected environment of an MFR in littoral (coastal) regions will include the land-sea interface. A generic case is a single high intensity edge present within the scene similar to that shown previously in Figure 5.2. The magnitude of expected edges can be determined from experimental observations resulting in specification documents such as the Naval Environmental Clutter Attenuation and Propagation Specification (NECAPS) available at DERA Portsdown (Branson 2000b). Although mean RCS for land and sea can be determined from reports such as NECAPS, geographical effects from areas such as cliffs are not accounted for. These would give a large planar reflecting area over some range cells which could cause 20dB edges.

In an environment expected to contain edges then the CA architecture would not be used and some form of edge resistant CFAR would be necessary such as CAGO; as shown above they suffer heavily in the presence of interfering or extended targets. To compensate for this to varying degrees a large number of CFARs have been proposed whose operation include censoring, averaging and ranking of the samples. In addition to CAGO and OS are:

- CMLD - Censored Mean Level Detector: The samples are ranked and the upper k removed. The test statistic is the mean of the remaining samples, encompasses CA as special case. Analysis by Rickard and Dillard (1977).
- GCMLD - Generalised Censored Mean Level Detector: as CMLD but k is chosen adaptively. Analysis by Himonas and Barkat (1992).
- TM - Trimmed Mean: The samples are ranked and the upper N_1 and lower N_2 removed. The test statistic is the mean of the remaining samples, encompasses CA, OS and CMLD as special cases. Analysis by Mashade (1996).

- OSGO, OSSO - Order Statistics Greatest Of, Order Statistics Smallest Of: Only one side of the window is ordered and the test statistic is the kth ranked sample. Analysis by Elias-Fuste et al. (1990).
- GOSGO, GOSSO - Generalised Order Statistics Greatest Of, Generalised Order Statistics Smallest Of: As OSGO, OSSO but k is chosen adaptively.
- GOSCA - Generalised Order Statistics Cell Averaging: Both the leading half of the window and the lagging half of the window are ordered separately, the test statistic is the kth ranked sample of the leading plus the lth ranked sample of the lagging window. Analysis by He (1994).
- LCOS - Linearly Combined Order Statistics: r ranked samples are chosen, the test statistic is then a weighted sum of the r samples. Analysis by Nagle and Saniie(1995) where weighting parameters are chosen using Censored Maximum Likelihood and Best Linear Unbiased estimates.
- EXGO, EXCA - Excision Greatest Of, Excision Cell Average - All samples above a fixed threshold λ_X are discarded before GO and CA operations are applied. Analysis by Han and Kim (1996).

In addition most can be enhanced by multiple looks, storage of a clutter map and estimation of the correlation between samples. It is rare that these CFAR are tested in anything but simulated exponential clutter or when the shape parameter is unknown. The analysis of these processors is laborious and it is not clear in the literature how many have been implemented on real radars or real data although Farina and Studer (1986) comprehensively cover their general implementation. To simplify discussion a general case is given in the left plot of Figure 5.5 - a CFAR of window size $W = 2M$ centred at the target x_T with an edge from position x_e onwards which is coincidentally located at the censoring point of the CFAR x_k. Capital letters A, B, C, D and E are marked for reference to avoid confusing use of subscripts; however any fixed window CFAR will apply a test based upon $x_{T-M:T-1}$ and $x_{T+1:T+M}$ using some derived test statistic based on partitioning the data window (via censoring, weighted mean, ranking) related to some fraction k in magnitude or position with respect to the test sample.

The decision to be made will include some of these hypotheses:

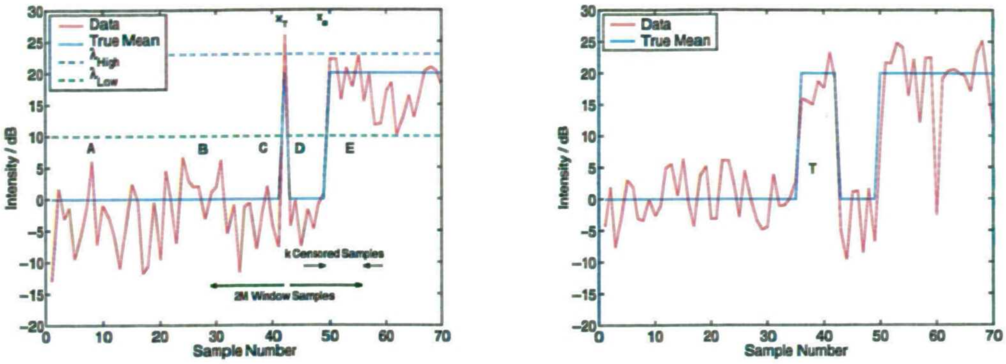

Figure 5.5: Generic cases for CFAR processing of littoral regions a) Point Target b) Extended Target

- x_T is within the low region $x_{T-M:T+M-k}$ (Region A-C) thus a threshold λ_{Low} is defined from these samples:
    - x_T will be flagged as a target if $x_T \geq \lambda_{Low}$.
    - x_T is a member of region A-C if $x_T < \lambda_{Low}$.
- x_T belongs to the high region $x_{T+k:T+M}$ (Region E) thus a threshold λ_{High} (where $\lambda_{High} \gg \lambda_{Low}$) is defined from these samples:
    - x_T will be flagged as a target if $x_T \geq \lambda_{High}$.
    - x_T is a member of region E if $x_T < \lambda_{High}$.

The censoring point k will usually be a fixed choice of the designer, chosen from the contending factors of:

- Pfa increase at the edge due to falsely declaring x_T as a target, when it is in fact part of the high edge (Region E).
- Pd decrease due to falsely declaring x_T as part of the high edge, when it is in fact a target within low clutter (Region C & D).

The various CFARs show interplay between their performance, where they will excel in one particular environment (tuned to their assumptions) but necessarily suffer some loss elsewhere. Note that in all algorithms the existence of an edge, whilst inferred from the statistics, is never flagged to the operator.

- Flagging a potential edge, regardless of form is desirable as it could enable a clutter map to be built up that informs the operator of the expected form of an edge for future trials.

It is interesting that even though an edge may be expected, CFAR algorithms are optimised to maintain the Pfa in homogenous clutter which in littoral regions will be rare.
The right plot of Figure 5.5 shows a similar scenario but where the 'target' T can no longer be considered as a single sample. This would arise from at least 4 distinct cases:

1. The target is spread over a number of samples.
2. Multiple close targets are present.
3. Multiple edges are present (in fact the 'target' is an island of high level clutter).
4. The target is a point target but the imaging chirp was of necessarily finite bandwidth giving a finite width point target response.

Point 4 can be handled by conventional CFAR via the use of guard cells surrounding the test sample, chosen to match the calculated response. This then gives reduced sample numbers and slower response to an edge but a range of possible situations will confuse any particular CFAR configuration for some scenes.

5.6 Generalising CFAR to Variable Window Size

It is clear that without exact knowledge of the expected size, magnitude and occurrence of edges then different scenarios could be envisaged where any of the discussed CFAR could either excel or perform poorly. Additionally, if an edge is present in the data then the estimated statistics whether it be mean, shape parameter or ACF will be some average of the distinct distributions and so it is interesting to investigate detectors that maximise the number of test samples used whilst simultaneously testing for homogenous conditions.

- *It is believed that a good indicator for deviations present within locally homogenous regions is the ratio between the means of all regions deemed suitably homogenous to have a well defined mean.*

This problem is probably ill posed. Even if a modulated exponential distribution is present, it is difficult to envisage simple tests that can maintain a fixed probability of false alarm for all edge configurations. The remainder of this Chapter details a novel attempt to apply this test.

5.6.1 Sequential Edge Detector

In formulating a detector for inhomogeneous environments:

- Fixed window CA-CFAR represents a simple way of detecting targets in homogenous clutter without the presence of edges.
- Simple floating point operations limited to $+, -, \times, \div$ and comparison operators will give the greatest operational speed.
- Multiple targets are not a primary concern of this thesis and are unlikely in a low observable detection scenario.
- Dead time must be handled correctly, present in most fixed window CFARs at the beginning and end of sample sequences.

From current literature and analysis of real data the clutter distribution is often considered as exponential, Weibull or slowly modulated exponential as in the K-distribution. As Weibull clutter can be transformed to an exponential distributed variable via a simple power transformation, detecting a change in the mean of a locally exponential variable offers a convenient starting point to characterise inhomogeneous environments.

An analysis by Oliver et al. (1996) covers optimum edge detection for a fixed window case with comparison of several edge detectors including the Maximum Likelihood Estimator (MLE) for a SAR application. Their concern is in accurately detecting the presence of single edges and their position; they consider two approaches - a Fixed Window with Scanning Edge (FWSE) and a Scanning Window with Central Edge (SWCE). In segmenting a scene, a step-up is of equal value to a step-down and so the MLE estimate is shown to perform best for each approach. They suggest a scheme to simultaneously maximise the probability of edge detection and edge position accuracy via a SWCE followed by a FWSE respectively. It is noted by Oliver et al. (1996) that '*In any application it is important that the window size should be sufficiently small that only a* *single edge would, in general, be expected*' and if this is not the case then an adaptive approach is required.

Greatest flexibility would be to allow for any number of changes in the mean either up or down with simultaneous target detection. In the preceding discussion the 'target' could easily be a point-like edge to be detected. In a multifunction radar there is no fixed resolution and thus no consensus upon the range extent of the target that poses a threat. If the resolution of the system is sufficient, even the aspect angle of the target will affect its range extent. Allowing for detection of extended targets (spanning multiple samples) means that the distinction between what is an 'edge' and what is a 'target' is dependent upon the operator. One can no longer define the expected number of samples within which a single edge is to be detected and so an adaptively windowed approach must be considered whereby both the tested target region window and homogenous region window are varied.

Adaptive windowing could be achieved through several methods:

- Run increasingly larger windowed CFAR in succession, starting at the smallest possible using a rule based system to progressively remove targets, this will be severely affected by false alarms at a strong clutter edge.
- Run decreasingly smaller windowed CFAR in succession, starting at the largest possible using a system to progressively remove edges, this will be severely affected by strong targets.
- Convolve a range profile with the expected edges to determine approximate positions of problem areas, probably affected by a combination of edges and targets and difficult to analyse.

A solution is to start at one end of the sample vector and use a post comparison window which progressively enlarges *but* every combination of samples that could contribute to a single edge are tested statistically. This will encompass any number of edges and any size of target; the only assumption being that speckle is decorrelated between samples.

5.6.2 SED Algorithm

The algorithm is shown graphically in Figure 5.6, based on the following available assets:

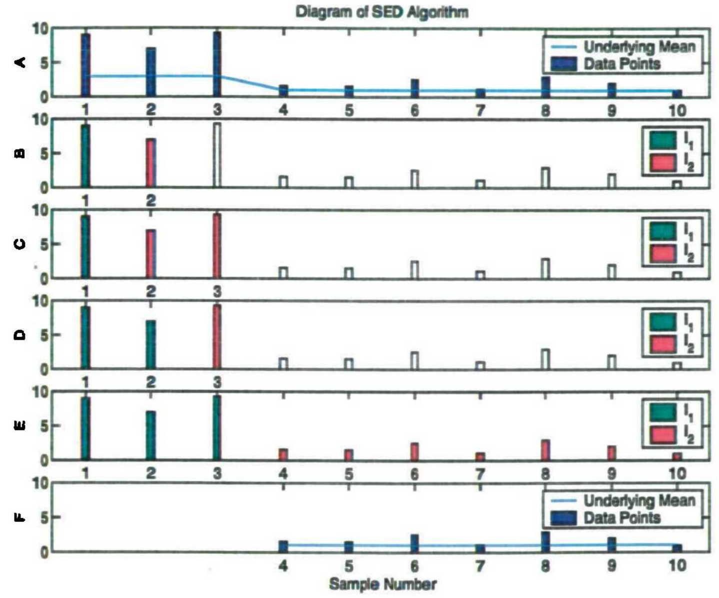

Figure 5.6: The Sequential Edge Detection algorithm: From an initial vector (A) the smallest subsample of $n = 2$ is chosen (B) to form the test between I_1 and I_2. As n increases the number of tests increases (C and D). When a significant ratio I_1/I_2 is declared (E) the test algorithm restarts (F)

- A realisation $x = \{x_{1:n}, n \rightarrow N\}$ from the radar output, where N is bounded due to a limited look time or a single range profile.
- Hypothesis H_0 indicative of no target, represented by PDF f_0 *versus* hypotheses H_i indicative of edge present represented by a single change point in the mean level after some point $1 \leq k < n$.
- A measure of performance based on P_{fa}.

The proposed algorithm is below; note that an edge is defined as an increase or decrease in the local mean level.

1. Obtain a realisation $x_{1:n}$, $n = 2$ initially (the minimum test region possible)
2. Test hypotheses $H_{i(n)}$

   $H_{0,(2)}$ indicative of homogenous statistics such that $f\left(x_{1:2}\right) = f_0$

   versus

   $H_{1,(2)}$ indicative of edge after sample x_1.
3. If any $H_{i,(n)}$, $0 < i < n$ are valid then flag edge and restart at Stage 1 with realisation $x_{i:i+1}$.
4. $H_{0,(n)}$ is valid, assert homogenous and increase test region via $n = \max\{n+1, N\}$ and obtain realisation $x_{1:n}$.
5. Test multiple hypotheses $H_{i(n)}$

   $H_{0,(n)}$ indicative of homogenous statistics such that $f\left(x_{1:n}\right) = f_0$

   versus

   $H_{k,(n)}$, $1 \leq k < n$ indicative of an edge after sample x_n.
6. Goto Stage 3 until $n = N$.

The algorithm has several drawbacks, the foremost is one of speed due to an order N^2 method; this has proved not to be a problem for typical sample lengths of $N = 256$ (in comparison to the fixed window NP-CFAR to be of order $W \times N$). The algorithm completion time is not constant as it can restart from a position $k < n$; although it is bounded for every possibility and fixed for homogeneous clutter. There is also a significant mathematical problem such that the non-null hypotheses are independent of the null hypothesis but not of each other.

The loss of independent sampling is unavoidable, at every point n a hypothesis of every change point is tested and thus the $n-1$ tests within $x_{1:n}$ are made using a varying proportion of identical samples. It is not clear how to handle this loss of independence mathematically, but it is actually present within fixed window CFAR when the windows overlap; as $W \ll N$ usually the effect is not significant. The only way to guarantee some stability is to ensure that every hypothesis test is made such that *if it were an independent test* then it would have the desired property of fixed P_{fa} (the implicit case of fixed window). Obviously the more tests made, the greater chance the false alarm rate, but due to loss of independence this increase may be slight, and easily warranted if one can detect edges or extended targets with greater accuracy.

5.6.3 Hypothesis Tests Used in SED

Whilst an MLE approach is optimum when all hypotheses are equally important (as in segmentation) for target detection the test is necessarily biased towards the non-null hypothesis of point target detection (as in the NP formulation). As the CA-CFAR is known to achieve optimum point target detection in locally homogenous clutter (Gandhi and Kassam 1994) then the generalised SED test must reduce to this in some form so as to maximise P_d. The ratio measure, closely related to a test proposed by Touzi et al. (1988) has such a form, although it has previously only been applied for fixed windows:

From a realisation $x_{1:n}$, define an edge position k such that region 1 is defined as samples $x_{1:k}$ with region 2 defined as $x_{k+1:n}$, from this, form the ratio r of the mean intensity within each region $\widehat{I_1}$ and $\widehat{I_2}$

$$r \equiv \frac{\widehat{I_1}}{\widehat{I_2}} \tag{5.12}$$

$$r \geq \lambda_1 \quad \text{Declare Edge Present with } I_1 > I_2 \tag{5.13}$$

$$r \leq \lambda_2 \quad \text{Declare Edge Present with } I_2 > I_1 \tag{5.14}$$

it can be seen that for $k = 1$, the test is equivalent to a one sided CA-CFAR test with a threshold multiplier $\alpha \equiv r$ and window size $W \equiv n-1$.

In homogenous exponential statistics, the PDF of the sample ratio r with knowledge of the true mean ratio R and the region sizes k and $N-k$ can be shown to be (Oliver et al. 1996)

$$P_r\left(r|R,k,N-k\right) = \frac{\Gamma\left(N\right)}{r\Gamma\left(k\right)\Gamma\left(N-k\right)}\left[\frac{\left\{\frac{rk}{R(N-k)}\right\}^k}{\left\{1+\frac{rk}{R(N-k)}\right\}^N}\right] \tag{5.15}$$

which has edge detection probability, when carried out at the correct position,

$$\begin{aligned} P_d &= 1-\int_{\lambda_2}^{\lambda_1} P\left(r\right)dr \qquad (5.16)\\ &= 1-\left|\frac{\Gamma\left(N\right)}{k\left(N-k\right)\Gamma\left(k\right)\Gamma\left(N-k\right)}\left[Y^k \times {}_2F_1\left\{k,k-N+1;k+1;Y\right\}\right]\right|_{\lambda=\lambda_2}^{\lambda=\lambda_1}\\ Y &\equiv \frac{k\lambda}{R\left(N-k\right)+k\lambda} \end{aligned}$$

where ${}_2F_1\left\{\cdot\right\}$ is the hypergeometric function, and the integration is to be evaluated between the limits λ_1 and λ_2.

Typical sample lengths are not expected to exceed 1024 and so Equation 5.16 is numerically inverted to obtain λ for $2 < N \leq 1024$ with $1 < k \leq N$ and the results stored in a triangular matrix. This operation was not trivial, problems with numerical accuracy arose which depended on the particular case of the hypergeometric function; standard code such as that in Numerical Recipes (Press et al. 1992) could not handle every case. Using the linear transformation formulae of the hypergeometric function in Abramowitz & Stegun 15.3 (1972) and observing their convergence criteria a suitable evaluation method can be applied for all cases but took the order of hours to evaluate the entire matrix. Figure 5.7 shows how the ratio threshold changes for geometry and Pfa.

When the ratio test is included in the SED algorithm, only the threshold matrix is necessary and so the algorithm is suitably efficient when implemented with a lookup table. An example of its operation is shown in Figure 5.8 for an exponential scene with multiple change points processed at individual test $P_{fa} = 10^{-3}$ and 10^{-6}. A gradual mean change is determined as a discrete step but this is unavoidable without specifying an expected sample correlation, the important point is that the variable correlation across the scene is determined.

To demonstrate the difficulties caused by the loss of independence, the simplest SED case of $N = 3$ is examined in detail. The six edge hypotheses applied upon these samples are below (The capital letter subscripts *do not* refer to the labelling in Figure 5.6 - although the second, third and

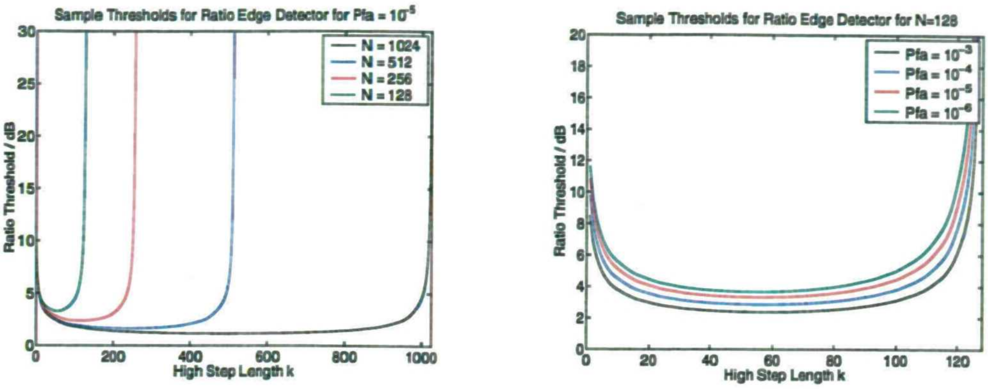

Figure 5.7: Sample ratio threshold as a function of geometry and Pfa

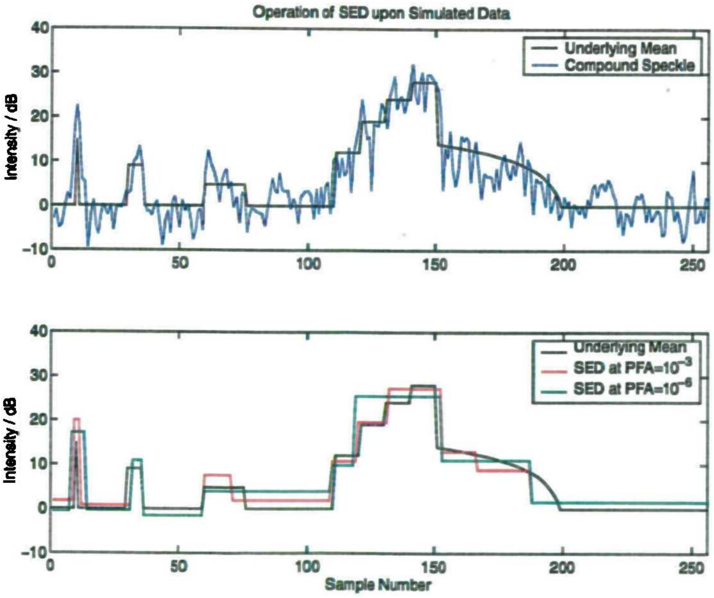

Figure 5.8: The SED test with Touzi ratio hypotheses segments a non-stationary scene reasonably well assuming locally exponential speckle

Table 5.1: Proportion of those sample failures common to two of the six edge hypotheses in a 3 sample SED, Pfa=10^{-3} in homogenous exponential noise

	H_A	H_B	H_C	H_D	H_E	H_F
H_A	1	0	<0.01	0.75	<0.01	0
H_B	0	1	<0.01	0	<0.01	0.06
H_C	<0.01	<0.01	1	0.03	0	0
H_D	0.75	0	0.03	1	<0.01	0
H_E	<0.01	<0.01	0	<0.01	1	0.03
H_F	0	0.06	0	0	0.03	1

fourth subfigures are relevant); note the first two hypotheses are tested without knowledge of the third sample.

$$
\begin{aligned}
\text{Step Up} &: H_A = \textit{low high} \\
\text{Step Down} &: H_B = \textit{high low} \\
\text{Step Up} &: H_C = \textit{low low high} \\
\text{Step Up} &: H_D = \textit{low high high} \\
\text{Step Down} &: H_E = \textit{high high low} \\
\text{Step Down} &: H_F = \textit{high low low}
\end{aligned}
$$

One would expect there to be strong positive correlation between H_A and H_D independent tests. The overall degree of correlation is presented by means of a comparison of those sample failures common to two hypotheses in Tables 5.1. Results were initially obtained based upon 10^8 runs at a single independent hypothesis $P_{fa} = 10^{-3}$, this produced the expected number of false alarms on the order of their statistical error ($100,000 \pm \sqrt{100,000}$). All the expected correlations are there - almost negligible except for the strong correlation present between H_A and H_D. A Pfa of 10^{-3} is not particularly low, and so Table 5.2 shows that for $P_{fa} = 10^{-3}$ the strong correlation effectively vanishes for all except H_A and H_D.

What is not shown is the existence of negative correlation, this will be most obvious between $H_A : H_B$, $H_C : H_D$, $H_E : H_F$ but also extends between tests carried out at different sample populations. The overall effect of the correlation between tests will require complicated analysis and so a simulation of 3.6x10^6 of a length 128 vector was performed to determine the probability of declaring any type of edge. A simplified SED algorithm was used where upon declaring an edge,

Table 5.2: As the Pfa decreases (here 10^{-5}), the dominant correlation between edge tests is fairly stable whilst others disappear

	H_A	H_B	H_C	H_D	H_E	H_F
H_A	1	0	0	0.75	0	0
H_B	0	1	0	0	0	<0.01
H_C	0	0	1	<0.01	0	0
H_D	0.75	0	<0.01	1	0	0
H_E	0	0	0	0	1	<0.01
H_F	0	<0.01	0	0	<0.01	1

the position of this edge was stored and a new vector was tested; in reality the SED algorithm would restart upon the remaining samples of the vector after the declared edge.

The results of this simulation are shown in Figure 5.9 for a Pfa of 10^{-3}. Neglecting edge effects, the overall form of the Pfa is linear with respect to the sample number. The probability of falsely declaring a step down is significantly greater than that of a step up, this is a direct result of the strong correlation between the tests for steps up shown in the previous tables. Note that the tests on the final few samples actually give a Pfa less than that of the independent hypothesis, this is because there is a reduced probability of even reaching the final sample numbers and less 'look back' where a sample is tested repeatedly once it has been passed.

The Pfa values are most importantly stable and of similar order of magnitude to the design based upon independent hypothesis tests. Several points with regard to this can be made:

- The performance in homogenous statistics does not need to be rigorously calculated since in practice one would never operate the SED in homogenous conditions.
- The robustness to edges of *any length* and configuration is most important which is hopefully what has been achieved.
- If an edge is declared it would always have been declared *had we decided to choose to test for that particular edge* at a fixed Pfa level, thus one can consider the SED as performing many tests in parallel.
- It is not clear from the literature whether the mathematics are available to handle the loss of independence, when Oliver (1996) encountered this using a fixed length sliding window, the required thresholds were calculated by Monte Carlo simulation.

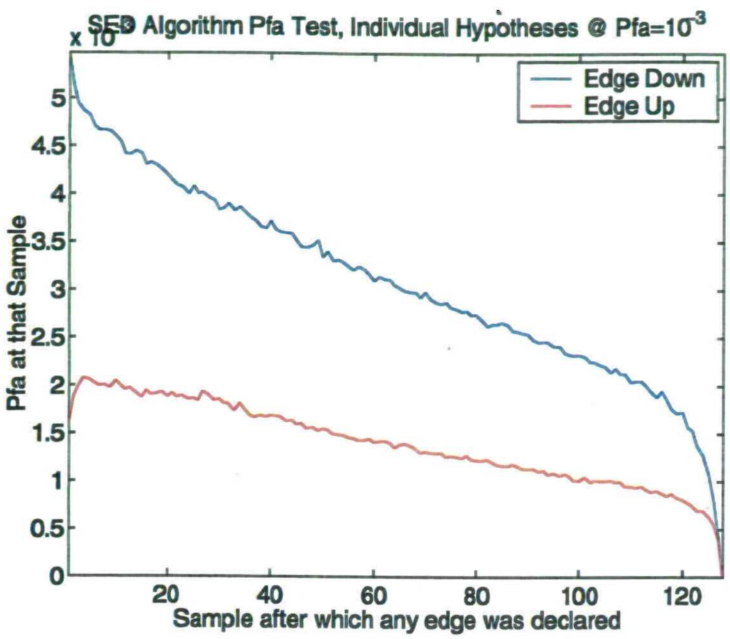

Figure 5.9: Edge detection using the SED algorithm with ratio hypotheses

In the next Chapter the SED is used to suggest that the high resolution images have discrete events which are not drawn from the underlying modulation responsible for the compound description of sea clutter, this causes the data to be nonstationary. It is very difficult to fully specify a stationary null hypothesis, due to the uncertainty of the actual sea clutter distribution and the presence of large edges from system noise (see next Chapter).

For some comparison with the tests performed in the next Chapter, exponential speckle is applied to a slowly modulated underlying exponential variate to give a simulated compound K distribution of shape parameter $\nu = 1$. The modulation was generated via two realisations of an Ornstein-Uhlenbeck process with an associated correlation length τ. Discussion of this process can be found in Section 7.1 where it is used to generate realistic time-varying target returns.

An equivalent of 12 seconds of high resolution scene samples were simulated to determine the length distribution of up and down steps when a continuous modulation is present, this is shown in Figure 5.10 for a correlation length of $\tau = 32$. The important point is that the SED produces a unimodal distribution of positive (step up) and negative (step down) lengths, the positive steps presumably being detected in greater numbers due to the form of individual hypotheses used.

- The histogram of step lengths gives an indication of the correlation present within a sample

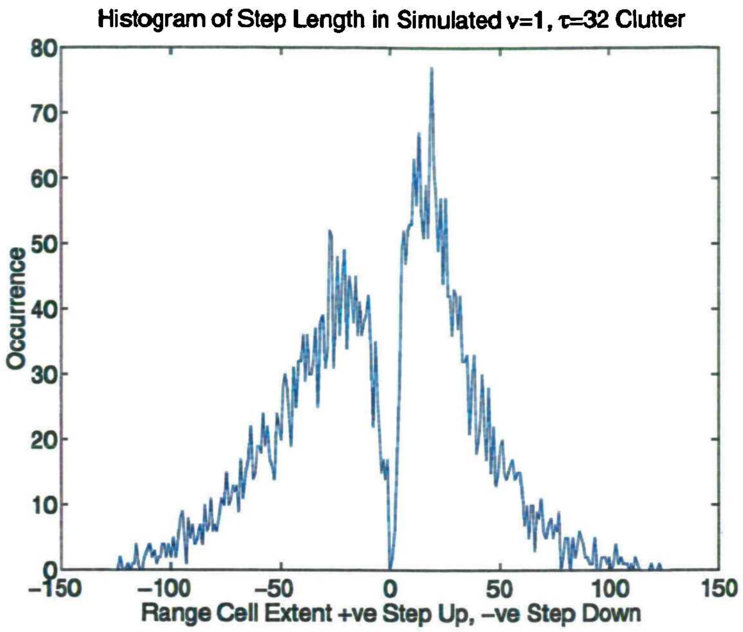

Figure 5.10: Whilst far from providing an ideal null hypothesis, the SED measured step length up and down is unimodal in simulated clutter of $\nu = 0.1$

vector which treats any modulation as a discrete number of *individual* steps without assuming stationary statistics over the vector.

This is far from an ideal null hypothesis with which to compare the real data results, one of the problems is that it is difficult to specify both the overall PDF (first order statistics) and the autocorrelation function (second order). This has been solved recently by Tough and Ward (1999) but has not been implemented here. The effectiveness of the SED is empirically justified in the next Chapter for further evaluation of the high resolution data - not as a target detector, but for evaluating the locally stable exponential regions.

5.7 Summary

As the data is usually subject to a point response function then a target is equivalent to consecutive finite length edges. The specification of a hypothesis to determine an edge is dependent upon the distribution of the underlying noise. As shown in Section 3.1.4, one cannot simply estimate this underlying noise when edges are present and so a fixed length CFAR is difficult to specify, especially when the correlation length is unknown. An attempt to overcome these problems is developed by a Sequential Edge Detection which assumes locally stationary exponential statistics and declares both positions and values of locally mean intensity. In addition no *a priori* decision is made upon the analysis length which, whilst ultimately suboptimum in perfectly specified clutter, should give a robust method for further analysis of the data.

The SED is difficult to analyse for the homogenous case as the repeated statistical tests carried out are not independent of each other, only the null hypothesis. Use of SED is justified by:

- One would not operate the SED in homogeneous clutter so its performance there is not a factor.
- Whilst the analysis is intractable, it is similar to testing for every possible edge through fixed window cell averaging type tests. An edge declaration would always arise *had we decided to choose to test for that particular edge* at a fixed Pfa level.
- The maximum likelihood hypotheses were not used in SED as a biased test towards point targets is preferable and the resulting cell averaging tests are far quicker in operation.
- Something which is never mentioned in the literature, but is applicable in this data, is that the probability of having a true edge within the data is far greater than typical false alarm rates and so the exact behaviour of the low false alarms will be swamped by this.

Classical CFAR statistical schemes have been discussed but robustness to their strong *a priori* assumptions is rarely tested. In particular a point target is always assumed. If an MFR can adaptively change its resolution then point targets are no longer always expected. If the SED can be viewed as a target (finite length edge) detector then it has the appealing property of being able to adapt to a potential target size even in the presence of strong edges. No claim to optimality is made (the hypothesis tests were chosen for processing speed) but the algorithm does have some interesting uses which are applied in the next Chapter.

Chapter 6

Further Analysis of the Data

The range profile was suggested to be a heterogeneous mixture of distributions in the preliminary analysis of the data. Large blocks of data were analysed which would be unavailable in an operational radar. Typical operating conditions for an MFR would be to sample an entire range profile from an area once every second (Webb 1999). As the high resolution scenes showed, movement of the wave events through the range extent caused the distribution to be unstable over this time period. Prior information on the observed distribution is of limited use and would not even be available if one is taking a first look at a particular area. If the ship is moving then it may not be possible to view the same area more than once and difficulty in aligning range cells causes problems when forming a clutter map.

6.1 Analysis using the SED

The previous Chapter described a Sequential Edge Detector which was implemented to avoid any dependence upon a chosen window size for analysis of the data and to simultaneously indicate areas that would flag as a point target and as edges. The 'perfect' high resolution data can now be re-evaluated to separate the regions. Note that although images are shown in this Section *all processing is performed on a single range profile basis* as this is how the radar would analyse a scene - probably at intervals of a second. All SED processing was made using mean ratio hypotheses at an independent Pfa of 10^{-6} assuming locally exponential noise.

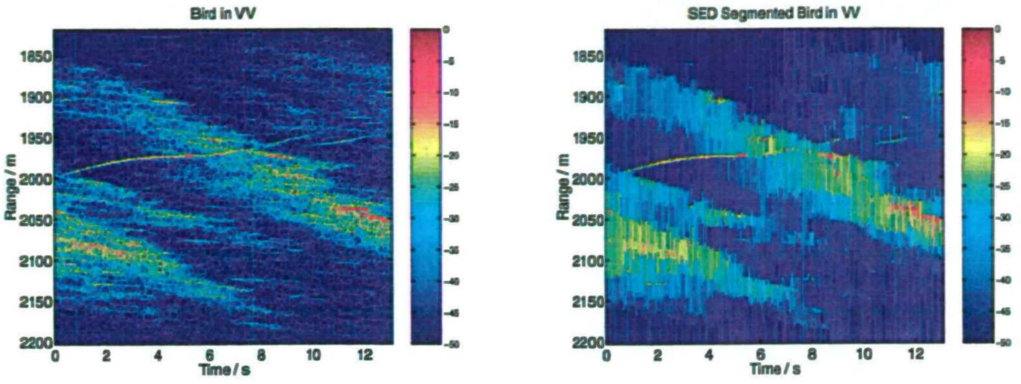

Figure 6.1: A bird in the VV data served as a point target (calibrated dB RCS)

6.1.1 Scene With a Bird Present

The radar operated in a littoral region, and so it is not surprising that a number of birds are sometimes present. These are obvious from inspection of the Hi-res scenes, as they are effectively point scatterers. Figure 6.1 shows part of sen3353 with a bird track present; both the raw RCS scene and the SED processed version are shown. The range profile has not been normalised with respect to the pulse shape, this is to illustrate the level of system noise at the edges of the range swath. Several points can be made from Figure 6.1:

- The SED has made a good attempt at identifying the regions with constant local mean. The edges are accurately determined over time with 256 samples per time slice.
- Although the wave event moves through the range cells, as the SED operates on a single time slice, the correlation within the underlying wave events is accurately shown.
- Variable correlation is seen across a single range cell however this is stable over the order of a second.
- As the underlying wave event is so stable, obvious edges between this and the system noise of the order of 20dB are inferred.
- From Section 4.2.2 the noise level was determined as an equivalent RCS of -40dB, it is now obvious that a large spatial extent of the scene is around this level and the SED highlights this.

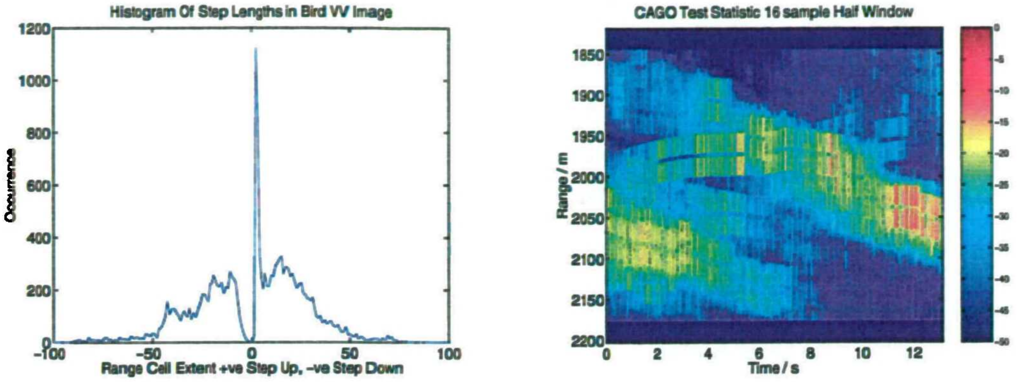

Figure 6.2: Histogram of the extent of locally stable mean regions suggests a useable fixed 16 sample half window CAGO-CFAR could be used. Resultant calibrated dB RCS image is shown on the right

- Striations are present upon the wave event from 8 seconds onwards between 1975m and 2050m in range. These seem to move at a lower speed than the wave event itself.
- The bird is easily observable in the low RCS areas which are dominated by system noise.
- Within the high RCS areas, the bird is indistinguishable from the striations over the order of a second.

From the SED analysis, one can form a histogram from the lengths of the locally stable mean regions which indicates a possible fixed window length to perform a CAGO comparison. The SED edge positions are stored and the measured lengths of the segmented regions regarded as either a step up or a step down relative to the locally determined mean.

Figure 6.2 illustrates this histogram which shows the effect of the target in the scene as numerous short high intensity steps. The modal step length is the width of the point target response determined by the Kaiser window applied prior to range compression. Other than the point scatterer peak, the histogram indicates that a reasonable size of CAGO half window would be 16 - with 2 adjacent guard cells to account for the point response of the test sample. Figure 6.2 shows the results of this CAGO:

- The local mean is severely in error when the bird is near the wave event - the classic problem with fixed window CFAR.

- The correlation effects are no longer visually apparent.
- A dead area is present at the start and end of the CAGO run.

So the standard method of identifying point targets does not perform well here, primarily because of the large edges determined by the SED. If one has a 20dB edge present over a significant part of the window then severe errors in identifying point scatterers will occur.

If there are spatially large noise dominated areas then the use of a single threshold multiplier applied with fixed window CFAR will not give a constant Pfa over the whole scene. If the high RCS regions are drawn from a different (spikier) distribution then the concept of a *variable* threshold multiplier α (Section 5.2.1) could be of use. Within the low RCS noise dominated areas one would use a significantly lower α than in the (presumably spikier) high RCS areas; exact α values to use could be calculated based on 'knowledge' of the local mean and distribution shapes (for CA and CAGO achieved relatively easily using numerical inverse Laplace methods of Section 5.4). This knowledge could never be achieved in practice, especially due to the large (20dB) edges present in the scene which would corrupt the local mean estimate of a fixed window.

If a target is present over the order of seconds, and the noise dominates a significant spatial area of the scene, a target detection method could be envisaged where *only* the system noise areas are tested since the statistics are fully known. Sampling a range profile at a frequency related to the swell movement would ensure that at some point a target would be revealed in locally noise dominated clutter.

6.1.2 Target Free Scene

The first target free region subsequent to the bird scene was chosen to illustrate the low velocity striations or 'slow events' as marked in Figure 6.3. This is typical of the VV scenes and again large areas of system noise are present in the central region of the scene - the striations can now be seen to extend into this area.

Figure 6.4 shows the scene segmented into point scatterers via SED (accounting for the point response) and standard CAGO (10dB above the local mean) both showing approximately 1% of the scene. This demonstrates that the majority of the false alarms are caused by the striations.

In particular the SED is now seen to show many point scatterers which are distributed across the ocean wave event - this highlights the effect that the striations have. Initially the SED was designed

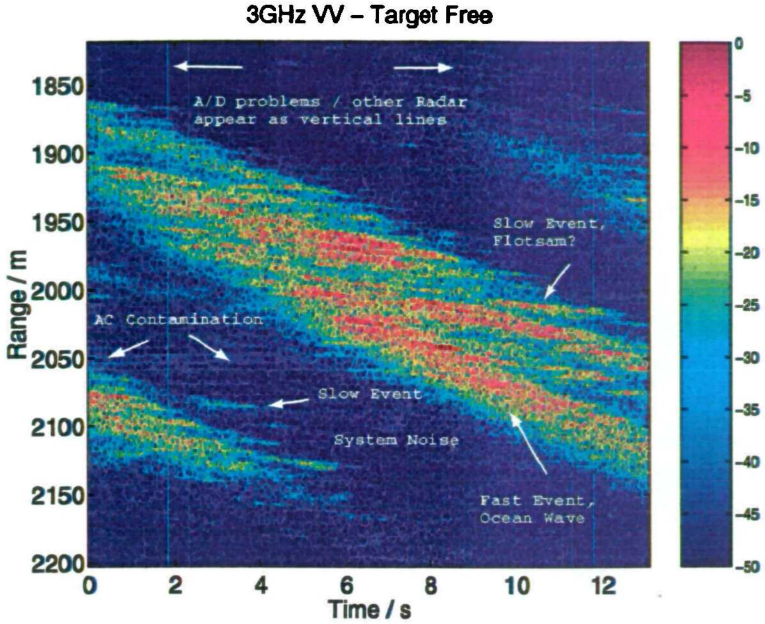

Figure 6.3: An illustration of a typical calibrated dB VV scene. The slow events can extend into the system noise

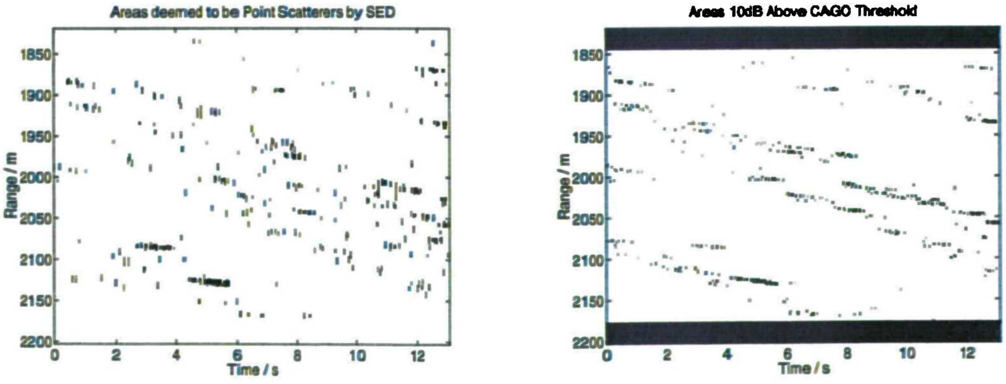

Figure 6.4: Segmenting the scene to observe point scatterers is possible by SED and standard CAGO

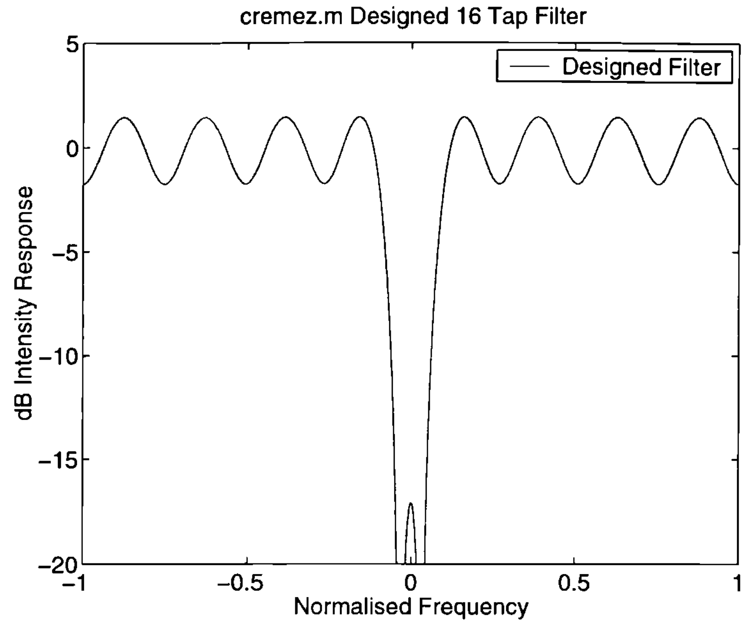

Figure 6.5: The 16 tap Matlab designed filter was operated on 10% of the Doppler spectrum in an attempt to remove the striations

to detect locally high intensity areas accounting for speckle without them necessarily being a point target; with the striations present it precludes the use of the SED as a target detector because at this resolution the returns cannot be completely described as locally exponential. There appears to be an additional component upon the overall wave modulation in the scene that is responsible for the point scattering - the striations are a major contributor to this.

6.1.3 The Striations

As this file was recorded so late in the day, the accompanying video recording was not taken due to poor light conditions. For this reason, the physical cause of the striations is unknown but as it is moving slower than the swell, it could be due to the foam left from a breaking wave or a wind generated wave.

To determine if any particular velocity component of the clutter caused the striations a suitable 16 tap filter was designed (`cremez.m` in Matlab, a complex equiripple FIR designer) which can operate on a 10% fraction of the Doppler spectrum with approximately 20dB suppression shown in Figure 6.5. This method, if proven useful, would remove any targets present in that fraction of

Doppler space - it did not remove the striation point scatterers. This suggests that the striations do not have a well defined constant Doppler velocity or that they scatter diffusely.

From observing the 20 minute VV run the striations are usually seen behind a wave front; in conjunction with their lower velocity this suggests they are flotsam. If this is the case then it casts doubt upon the ability to accurately predict performance figures for this type of data. Whilst the flotsam is probably associated with a breaking wave, its physical nature means that it may still be upon the surface a large distance from that wave, such that existing locally windowed detection methods will not be reliable.

6.1.4 Correlation Due to Striations

Figure 6.6 gives a histogram of the extent of the constant intensity regions from the high intensity central region of the scene, determined by SED, and it suggests two components of correlation even when the speckle is accounted for. The SED process assuming locally exponential clutter was ran at an individual Pfa of 10^{-6}, the segmented regions were then recorded as being a step up or down and their size (range cell extent) was histogrammed. Figure 6.6 demonstrates statistically significant evidence of there being discrete scatters upon a modulating correlation length because this individual Pfa would produce an expected maximum 2 false alarms per range cell extent in locally exponential speckle. The usual method of viewing correlation is via the autocorrelation function, also shown in Figure 6.6, which gives uncertain information since it is primarily describing the correlation of the speckle - of course the SED inherently accounts for this.

As highlighted in Figure 6.3 some AC frequency contamination is present; files showed this around 100Hz which suggests it is from some rectified 50Hz mains source. As it enters the signal prior to the FFT for range compression a particular range cell is corrupted, in this case around 2050m by a small amount. Although the effect of the apparent RCS is of the order of 5dB and does not affect the previous CAGO or SED analysis, in conjunction with the intermittent corrupted range profiles it could have serious effects upon any complex test (such as chaos or fractal) based upon data from this radar. This effectively precludes performance tests using this data as a background.

6.1.5 Scene Probability Distributions

If the scenes consist in general of a finite number of events then it is difficult to obtain a PDF that fits the overall statistics even for an observation time far greater than would be possible in

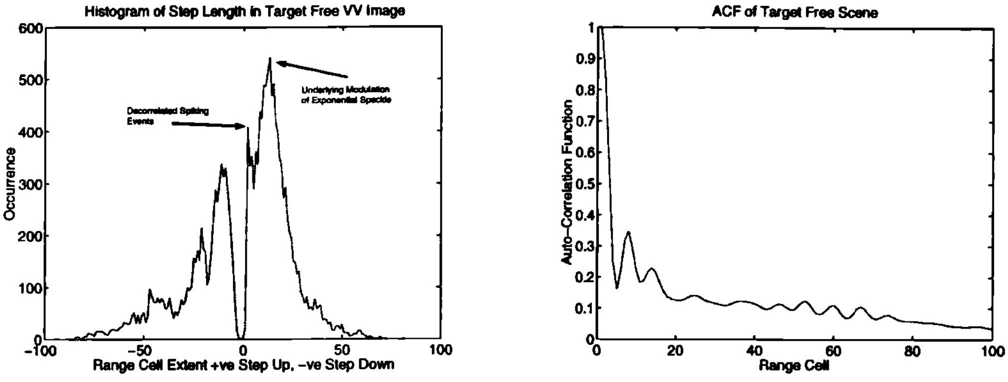

Figure 6.6: The existence of point scatterers in the nonstationary target free scenes and the spatial extent of the correlation is clearly shown by the SED. The ACF gives an uncertain picture as it views the scene as stationary

an operational scenario. The sample PDF is usually a smooth function especially if some form of windowing is used to average out the speckle, prompting attempts to fit the whole of its extent which are ultimately unreliable as shown in Chapter 4. If the SED is used to segment the scene then this averaging is performed in a more intuitive way that avoids the problems of large edges ‘smearing’ the resultant histogram and thus reveals the discrete nature of the scene. Figure 6.7 shows this graphically where the SED implies a multimodal distribution subject to speckle; such that the potential false alarms are from a discrete component of the scene that has no obvious relation to parameters such as mean and variance of the overall PDF determined on a pointwise basis.

- The high intensity tails are determined by a discrete component of the scene which is only obvious after initial segmenting based upon the mean RCS and assuming locally exponential speckle.

This SED analysis was extended over the entire 20 minute dataset to determine if a stable multimodal distribution was apparent. The peak positions and their form varied such that the swell events could eventually be considered as coming from a continuous distribution; this masked any obvious spiking events over this time.

As the large scale swell events appear to be relatively stable over the scene it is interesting to see if they can be analysed individually by spatially removing the areas of pure system noise. This is possible through segmenting by hand, but by thresholding the scene based upon the SED the

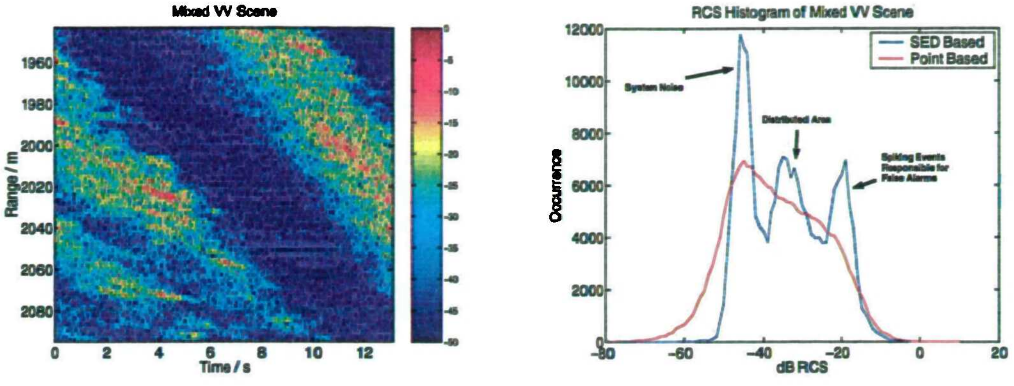

Figure 6.7: Obtaining the PDF of the SED segmented scene reveals the discrete nature of scenes. Without the SED, well defined areas are not apparent

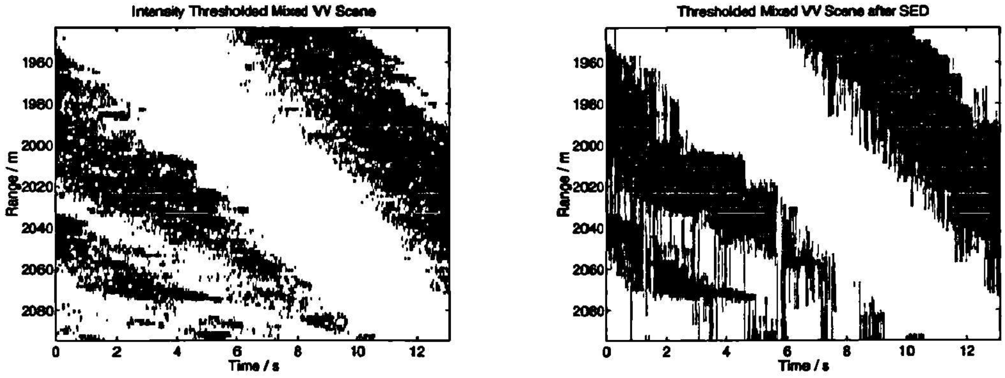

Figure 6.8: Thresholding the mixed VV scene at -30dB allows separation of the large scale RCS events from the system noise

spatial area to be sampled can be determined automatically. A simple intensity threshold does not provide well defined areas due to the effects of speckle shown in Figure 6.8.

This segmentation method was applied on a large number of scenes to look at individual swell areas and gave unstable results; while the K distribution could often be adjusted to fit the tail regions well it could not characterise the whole RCS range and the shape estimator U underestimated the tails even though the rest of the distribution was fitted well - this is a direct consequence of the presence of discrete events within each wave swell.

- The presence of discrete intermittent events means that statistical comparison of individual swell waves is difficult. Whilst they are correlated in underlying mean RCS on the order of several seconds their probability distributions are variable and their returns cannot be viewed as being from a single class of continuous unimodal distributions such as K.

Accepting that at a high enough resolution the backscatter is discrete in nature motivates a detection scheme based upon the lifetime of individual events, determined in the Doppler domain, presented in the next Chapter.

6.2 Summary

Chapter 4 showed that previously considered distributions could not fully account for the sea clutter RCS distribution. Every scene had large spatial areas of system noise presumably due to geometric shadowing effects from the low grazing angle and high sea state. By operating an adaptively windowed edge detector which assumed locally exponential noise these areas could be segmented and further revealed that the scenes consist of two forms of spatial correlation. The large correlation extent is the ocean wave modulation which is included in the compound model of clutter but upon this discrete point scatterers are present in the form of striations. The RCS distribution of the point events is not obviously related to the underlying swell, can appear away from breaking waves and does not have a fixed velocity which suggests that flotsam is a possible cause.

Chapter 7

Wavelet Detection Methods

The previous Chapter dealt with incoherent detection within a range profile, this Chapter considers a novel coherent detection technique based upon the lifetime of the events responsible for the Doppler spectrum. Classical techniques, such as acoustic delay Moving Target Indication (MTI) formed detection, are based upon identifying a target from the amplitude difference between successive pulses (Shrader and Gregers-Hansen 1990). In coherent digital systems this has evolved into Moving Target Detection (MTD) where successive Doppler spectra are processed by a filter bank, Figure 7.1 demonstrates this where filter number 4 is relatively clutter free.

If the desired target lies within the clutter and the Doppler spectrum remains constant it should still be possible to identify it, however the Doppler spectrum is not necessarily stable and so thresholds must be high to avoid excessive false alarms. The events responsible for the high resolution false alarms were close to the clutter peak and so it is interesting to see if a coherent target detection method can be operated within this constraint. To distinguish these from targets it is deemed necessary to observe the time evolution of these events with respect to their Doppler velocity. As discussed in Section 3.2.3 this can be achieved by the Continuous Wavelet Transform which adaptively changes the analysis window with which to form a Doppler spectrum.

The literature commonly refers to a *correlation length* as an average measure of coherence estimated from the autocorrelation function. A single correlation length is especially inappropriate for sea clutter when it is obvious that the sea has many components based on tide, swell, wave and ripples each varying individually in time.

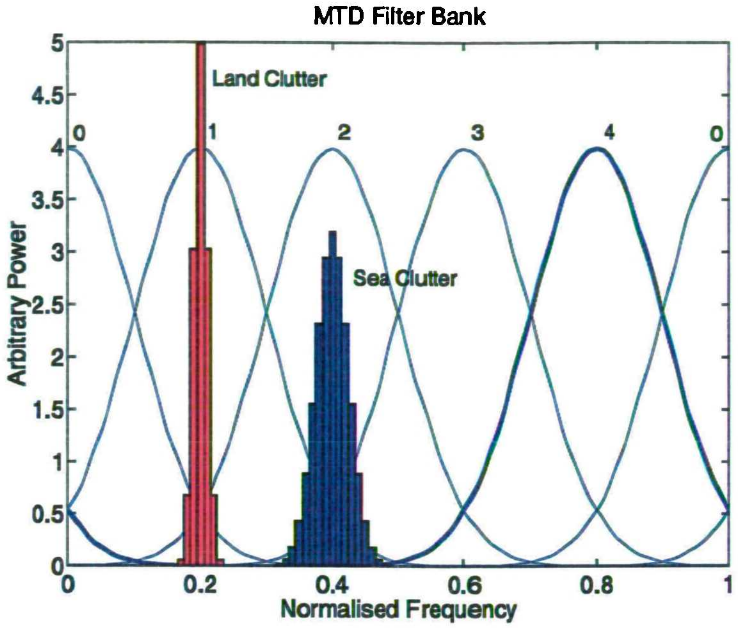

Figure 7.1: An MTD Filter bank can be used to detect fast moving targets outside the clutter. Filter 4 could be the input to a CFAR

7.1 Target Model Formulation

A method is required that can model a slow moving target effectively hidden within the Doppler spectrum of the sea surface to determine an effective detection technique. A paper by Tonkin and Dolman (1990) explored the RCS PDF of a periscope subject to:

- Shadowing by the sea surface.
- Specular multipath.
- Lobed backscatter structure from a cylinder.
- Corner reflector effects between vertical periscope and sea surface.

A combination of analytic and Monte Carlo methods determined the overall RCS PDF to be largely Swerling I with occasional periods of very high return. Whilst the returns were said to be 'pulse to pulse correlated' and an oscillating motion is to be inferred from the diagrams, the explicit form of the correlation over time is not explored. Incidental effects such as reflected waves or wake production were not included.

A generic time varying target is proposed:

1. The RCS is modelled as Swerling I generated via Ornstein-Uhlenbeck processes with a correlation time τ determined from

$$\begin{aligned} x_{t+1} &= \lambda x_t + \sqrt{1-\lambda^2} N(0,1) \ , \ \lambda^2 < 1 \\ \lambda &= \exp\left(-\frac{1}{\tau}\right) \text{ st } \langle x_t x_{t+1} \rangle = \exp\left(-\frac{t}{\tau}\right) \end{aligned} \tag{7.1}$$

   where $N(0,1)$ (and strictly x_1) is a Gaussian random number of zero mean, unit variance. Squaring and adding two independent realisations of Equation 7.1 gives an exponentially distributed variable whose correlation length is related to τ.

2. The Velocity is modelled by lowpass filtered Gaussian noise of cutoff frequency f_c which gives an expected Gaussian shaped Doppler spectrum of central velocity v_0 with associated width Δv, a 5th order Butterworth filter was suitable.

One can model different targets by varying τ and f_c but for slowly varying targets a time constant of order 1 second for both correlation parameters is physically reasonable, e.g. τ is set to 1 second and f_c is 1Hz. To emphasise the intended use of the Wavelet analysis v_0 and Δv are always chosen to be within the clutter spectrum, thus detection based on estimation of the background clutter shape would be of limited value since to a first approximation this is Gaussian itself.

It is this type of surface target that would be difficult to detect using a standard windowed CFAR operating on the FFT Doppler bins or in Range resolution since it is expected to emerge and fade slowly without causing a discernible high magnitude return. Additionally, associated surface effects such as a wake could be present which would corrupt any necessary local estimate of the statistics.

7.1.1 Targets Within the Doppler Spectra

Figure 7.2 shows the Doppler spectrum assessed from the centre of the Hi-res VV target free scene; the simulated target RCS and velocity are alongside. Based on the Doppler shape alone there is no obvious target present. The RCS variation means that the varying velocity is not apparent. Classical target injection of a single velocity with a well defined FFT peak will never be realised since both target and observer will have varying relative motion over the time necessary to discern a small target against the clutter background.

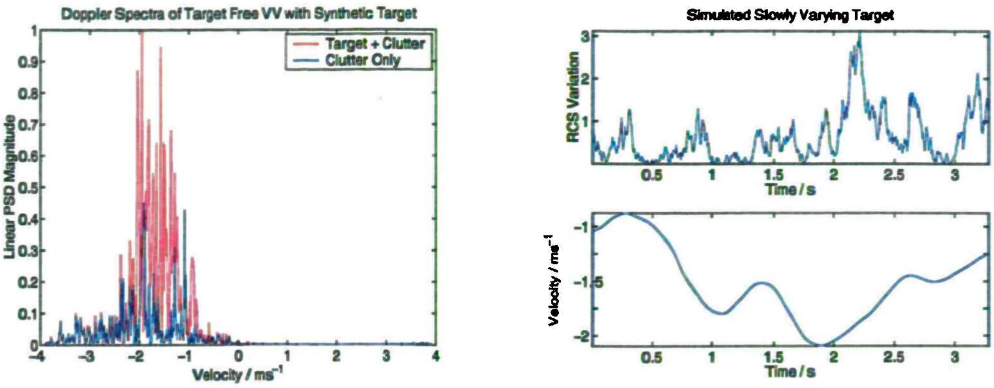

Figure 7.2: Slowly varying target model within high resolution clutter

The first assumption one makes is that the FFT window is simply too long and so shorter duration windows should be used to observe the varying velocity. This is the motivation behind using the Wavelet Transform described in Section 3.2.3 since *a priori* one cannot know the correlation time of the expected target; if detection is to be made upon differences in the correlation structure over time then a chosen fixed window will have a dramatic effect upon the determined spectra.

Low resolution sen3558 and sen3537 VV data is used for illustration. Taking a random 0.4 seconds of data, a plot is made such that Doppler velocity is along the abscissa with multiple ordinates showing the effect of window size from 1024, 512, 256 and 128 sample windows i.e. 'dyadic'. Subsequent Doppler plots can be seen in time on the lower plots, as 8 of the lowest resolution FFTs can be carried out in the time of the largest. A synthetic stationary 0dB target has been introduced into the central 0.2 seconds in the right hand plots of Figure 7.3. Additionally, Figures 7.4 show similar plots for a scene with birds present - a real target not unlike the model. Confusing spectra are observed where it is difficult to imagine a target detection scheme if this is a 'first look' at a potential hostile scene.

A potential solution is to use a clutter 'map' based upon *a priori* knowledge of the Doppler spectrum. This requires long look times and is unsuitable for a fast detection decision on a new scene. Sixty seconds of data was used to determine the 'complete' (but effectively unavailable) knowledge of the clutter Doppler spectrum. This was overplotted with its standard deviation to give an indication of a possible threshold in Figures 7.3 and 7.4.

From these figures it is obvious that conventional processing biased towards identifying high velocity targets away from the clutter is not suited to detection within clutter due to:

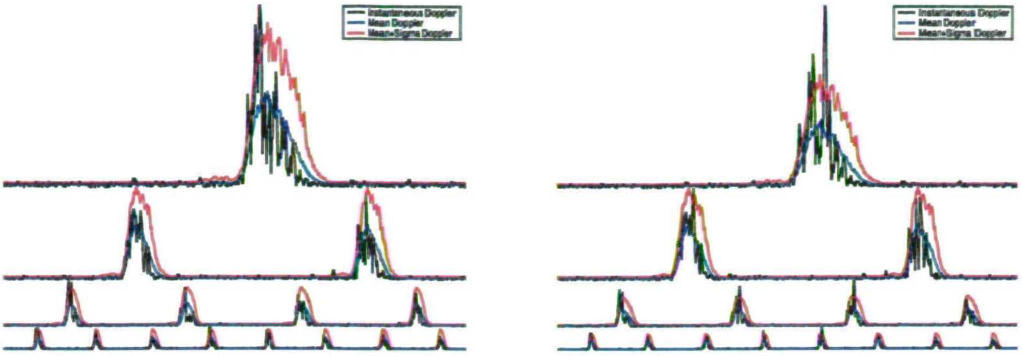

Figure 7.3: A dyadic plot of sen3558 Doppler. Right plot has an unlikely 0dB constant velocity target added

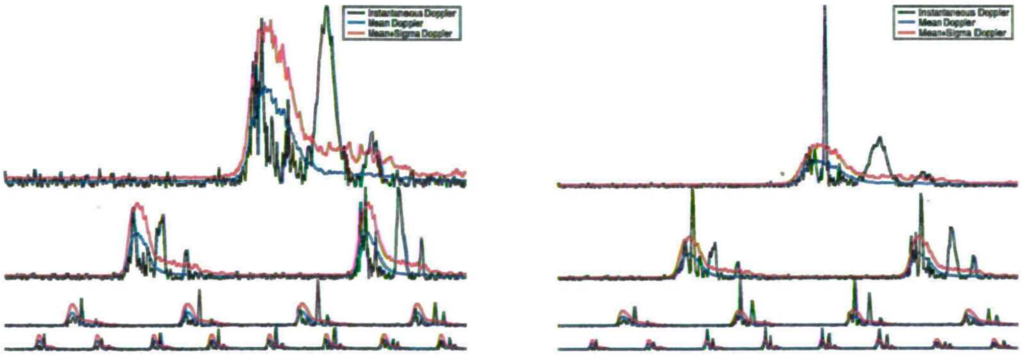

Figure 7.4: A dyadic plot of sen3537 Doppler with a single bird present. Right plot has an unlikely 0dB constant velocity target added

- Doppler variation over time will cause MTD to exceed the low threshold required to detect the target.
- Variation between window sizes means selection of the optimum is difficult.
- Objects such as birds could pollute a clutter map easily.
- An accelerating target such as the bird is seen as several discrete targets of highly variable amplitude due to fading.

Such that detection of targets, whether of unlikely constant velocity or those accelerating, is difficult within the clutter based on the instantaneous Doppler spectrum.

7.2 Application of the Wavelet Filters

Accepting that the return from the sea surface is as a result of several variably interacting components means that the traditional view of a Doppler ‘spectrum’ when analysed over short time frames is a misnomer when we are looking at a finite collection of individual scattering events. Isolating these events is naturally performed in the Wavelet domain since this considers localisation in time and frequency to be equally important in choosing window functions to minimise the combined uncertainty.

7.2.1 Frequency Response of the Wavelet Transform

Recall (Section 3.2.3) that the normalised form of the Morlet wavelet is defined (Torrence and Compo 1998) with a parameter ω_0 that determines the number of oscillations in the time domain (chosen as $\omega_0 = 5.336$)

$$\Psi_0(t) = \pi^{-1/4} \exp(j\omega_0 t - \frac{t^2}{2}) \tag{7.2}$$

which is then dilated in time by a scale factor s, represented in frequency where $\widehat{\Psi}\,(\cdot)$ is the Fourier transform of $\Psi\,(\cdot)$, as

$$\widehat{\Psi}(s\omega) = \pi^{-1\ 4} \exp[\frac{-(s\omega - \omega_0)^2}{2}]\ ,\ \omega > 0 \tag{7.3}$$

This defines a flexible time-frequency window which automatically scales to observe approximately ω_0 cycles of a particular frequency whilst still being localised in time. As the effective

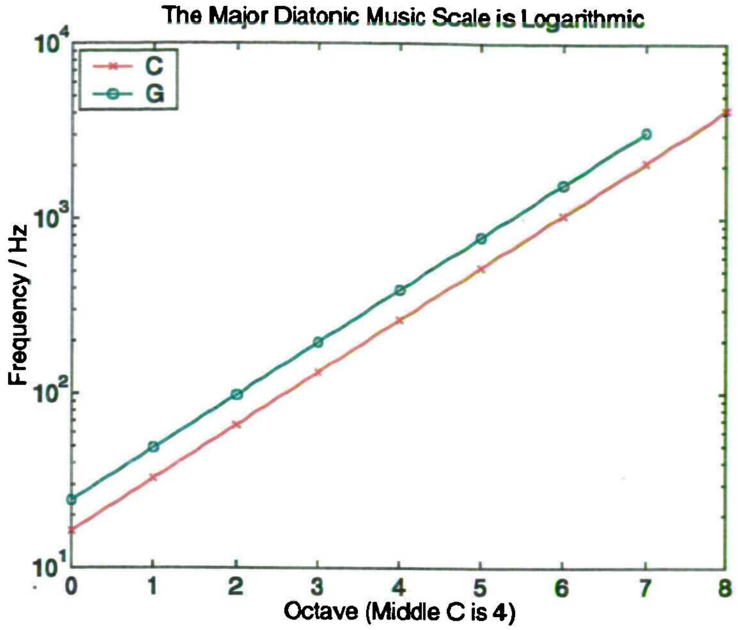

Figure 7.5: The Western diatonic music scale is spaced logarithmically, the wavelet filters mirror this to give a flat frequency response

bandwidth of the wavelet filter is a linear function of its central frequency, a logarithmic subset of frequencies is required to give a flat frequency response. This has similarities with the major diatonic pitch scale of Western music where octaves are spaced logarithmically, shown in Figure 7.5.

The scales are chosen so that the mth scale s_m is defined as

$$s_m = s_0 2^{m\Delta m} \ , \ 1 < m < M \tag{7.4}$$

where s_0 samples close to Nyquist (half the sampling frequency) and Δm determines the spacing between scales s_m and s_{m+1}. The parameter Δm is chosen to minimise the number of wavelets used, whilst being small enough to give a reasonably ripple free central frequency response, set as $\Delta m = 0.25$. Also note that in this form the transform is only defined for positive frequencies; negative frequencies are obtained by operation upon the conjugate of the original time series or equivalently the conjugate of the filter.

The algorithm to compute the Wavelet Transform can be summarised as:

1. Obtain time series $x_{1:N}$ and transform to frequency domain via FFT to yield $\widehat{x}$.

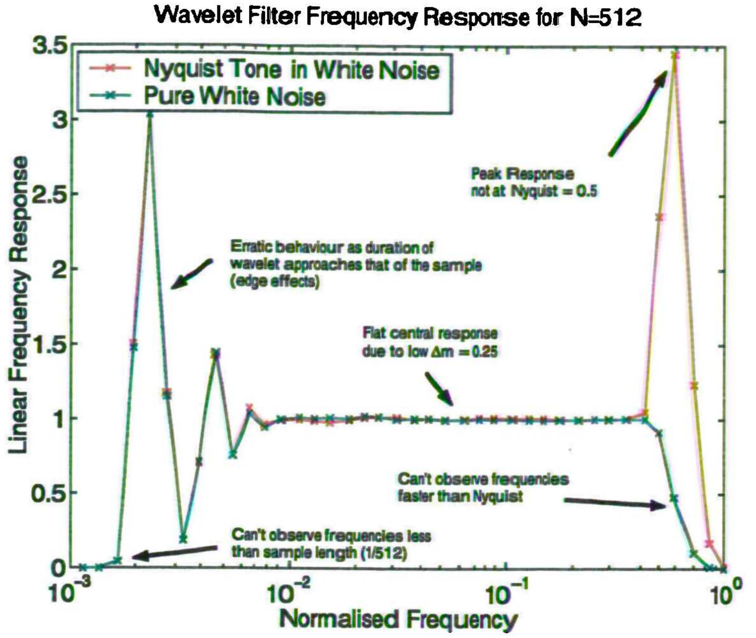

Figure 7.6: The frequency response of the Morlet CWT viewed as a filter bank is flat when operated correctly but several artefacts can arise

2. Calculate wavelet filter $\widehat{\Psi}(sw)$ in frequency domain for a particular scale s_m.

3. Convolve the filter over all time positions to operate the wavelet filter in the frequency domain

$$\widehat{W} = \widehat{x} \times \widehat{\Psi} \tag{7.5}$$

4. Calculate the inverse FFT of $\widehat{W}$ to obtain the length N, scale s component of the wavelet transform W_m.

5. Repeat from Stage 3 until all M scales are calculated.

The set of all M wavelets can now be viewed as a filter bank, from which a frequency response curve can be derived. As the filters are only implemented digitally via an FFT the response from white Gaussian noise is averaged for a typical sample length of $N = 512$ using the wavelet transform routines.

Care must be taken in choosing M and s_0:

- The wavelet has a variable duration in the time domain so edge effects are important as the scale s approaches that of the sample length N.

- One cannot measure frequencies faster than Nyquist or slower than the duration of the sample.

- Meyers et al. (1993) show that the peak response frequency f_r is not simply the inverse of the scale but

$$f_r = \frac{\omega_0 + \sqrt{2 + \omega_0^2}}{4\pi s} \approx \frac{0.864}{s} \quad (\omega_0 = 5.336) \tag{7.6}$$

  and so the derived Doppler spectrum must be scaled accordingly to correctly determine velocities.

These effects are summarised in the frequency response curve of Figure 7.6 which shows the normalised response to white noise and a tone at the Nyquist frequency embedded in noise.

7.2.2 Wavelet Transform Input to a Detector

From a time series $x_{1:N}$ the wavelet transform separately determines both the positive and negative Doppler velocities present. As the observed clutter spectra are all one sided this means only one transform need be operated but otherwise this would raise the question of how to handle events whose velocity changes sign (although a similar problem arises at the DC component of the standard FFT).

The resultant WT is now recognised as providing an instantaneous Doppler spectrum which has been ideally smoothed in the time-frequency plane.

- The instantaneous Doppler spectrum, as determined by the CWT *is not some arbitrarily smoothed average of multiple FFTs* but an explicit method of weighting the filters that balances the equally important criteria of localising an intermittent event in both frequency and time.

The global wavelet spectrum when averaged from the instantaneous wavelet spectrum is an unbiased estimation of the true power spectrum of the signal (Percival 1995). Figure 7.7 shows the mean WT compared to the mean FFT spectrum (calculated with a Kaiser-4 window) for 60

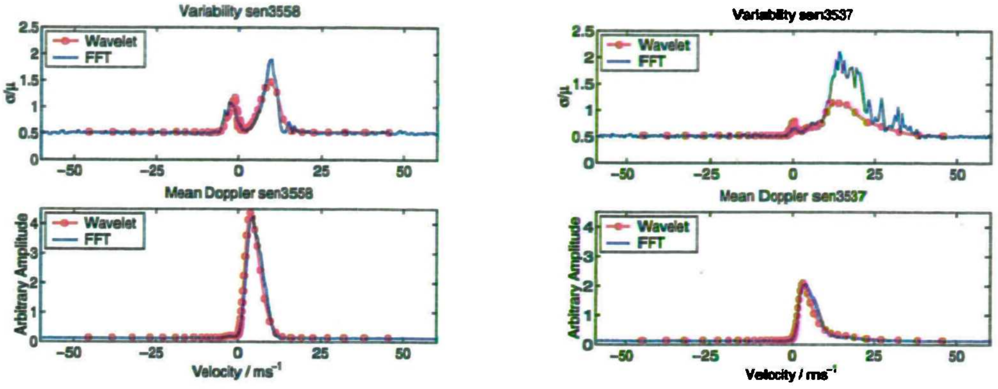

Figure 7.7: The mean wavelet spectrum represents the Doppler spectrum as well as the FFT but the filter bins show less variance

seconds of sen3558 and sen3537 and additionally shows the variability of the filter bins (standard deviation divided by mean). The variability is considerably smoother for the WT of sen3537 even though it contains intermittent events discussed in Section 4.3.2.

It is the smoothness of the WT that allows analysis of the instantaneous events causing the overall Doppler spectrum over time. Figure 7.8 gives a cartoon view of how these events may appear:

- The left of Figure 7.8 demonstrates the appearance of two separate scattering events at different velocities which are unlikely to be targets, many of these short events will contribute to form an overall Doppler spectrum.

- The right of Figure 7.8 shows a single event changing velocity but dominating the spectrum. This is a potential 'persistent' target which is present within the clutter spectrum.

By following the largest maxima within the instantaneous WT derived Doppler spectrum the statistical distribution of each WT bin is not required. Measuring the length of time these events are continuous (the length of the blue arrows in Figure 7.8) gives a physically motivated threshold with which to form a detection criterion that will hopefully be more stable than the RCS fluctuation within the clutter. It may be argued that selecting only the maxima is throwing away information but a detectable target will necessarily add to the magnitude of the Doppler bins within the clutter, thus identifying the largest maxima gives a simple initial threshold.

Obviously detection of small targets outside the clutter spectrum using this method is not viable but this area is dominated by system noise (and sidelobes from the clutter due to filtering) which

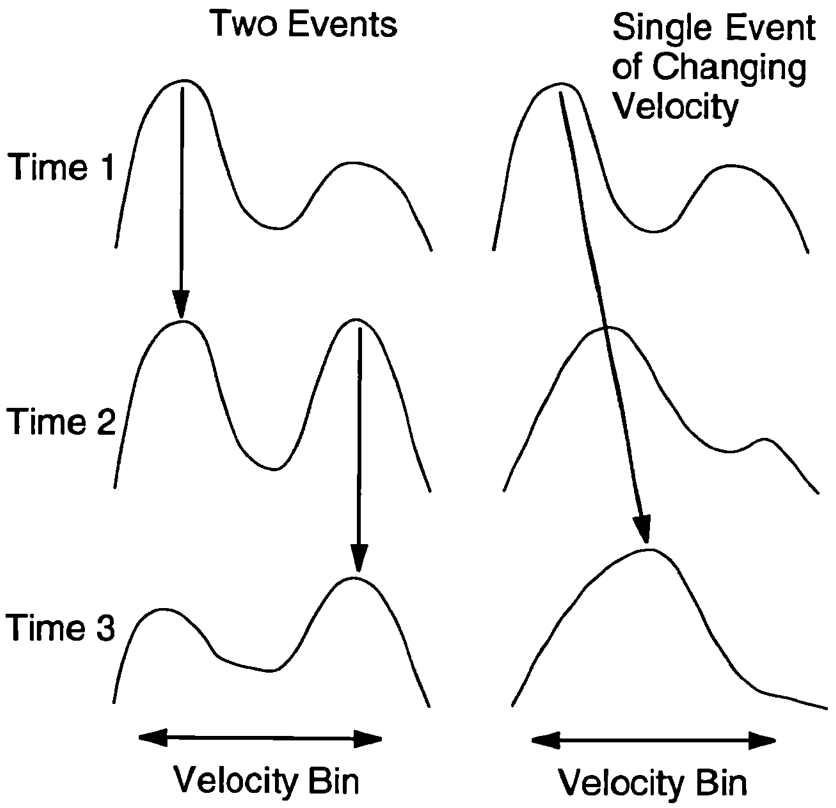

Figure 7.8: Cartoon showing the instantaneous WT velocity bins (Red) at 3 distinct times for two different scenario. Tracking the maxima (Blue) will measure the physical time that an event dominates the spectra

is not as variable as the clutter. Additionally this area represents the fast velocity targets whose 'lifetime' is determined not by any physical properties but the time they are within the sampled range cell. Over a certain velocity there is no use for the Wavelet Transform since an upper bound exists on the duration of a Doppler frequency.

- A detection statistic termed the '*persistence*' is proposed that is determined by measuring the length of time the largest WT Doppler component is continuous with respect to velocity. This is a physical parameter with units of time that is biased towards detection of slow moving targets at velocities within the long term clutter spectrum.

- The persistence *is not* related to any measure based on the autocorrelation length. The ACF can be determined from the FFT of the entire windowed sample and is thus subject to all the practical difficulties of the FFT for nonstationary signals discussed in Sections 3.2 and 7.1.1.

One caveat must be stated; as the time extent of each WT filter is a function of its matched velocity then the defined persistence measurement will not be completely independent of velocity. Torrence and Compo (1998) consider this in terms of the decorrelation time a particular wavelet has in response to a point discontinuity in the time series (such as random impulsive noise or edge effects). This '*cone of influence*', defined as the time taken for the filter intensity to drop by a factor e^{-2}, discriminates between random noise and a true frequency being present. This effectively puts a lower bound upon the measured persistence times of $\sqrt{2}s$ (Torrence and Compo 1998), as it turns out a usable target threshold is much larger than this value for typical well sampled time series and so the effect is minimal.

7.3 Observation of Real Data

To determine the statistical distribution of the Doppler event persistence, a large data file is necessary. Recall that the high resolution files are constructed from 256 individual frequency pulses swept through a chirp. By extracting a particular frequency from the chirp then an apparent 256 individual low resolution observations are made of the same scene. Long observations can not be made at a high PRF since this would dominate a multifunction radar processor, but this PRF of $40000/256 = 156.25$ represents an operational use of this detection method which additionally ensures the speckle is decorrelated. The 20 minute high resolution files in V(sen3353) & H(sen3352) now offer a large statistical population albeit covering slightly different transmission frequencies.

Figures 7.9 illustrates a short (3s) observation of sen3353 processed by the maximised wavelet method in low resolution with a 0dB time varying synthetic target introduced, several points can be made:

- For this particular data the clutter only Doppler spectrum resembles a 'traditional' point target *more* than the target plus clutter.

- The wavelet spectrum confirms the smoothness suggested in the cartoon Figure 7.8. The scale is defined in Equation 7.4, which for 512 samples bounds the Doppler spectrum well

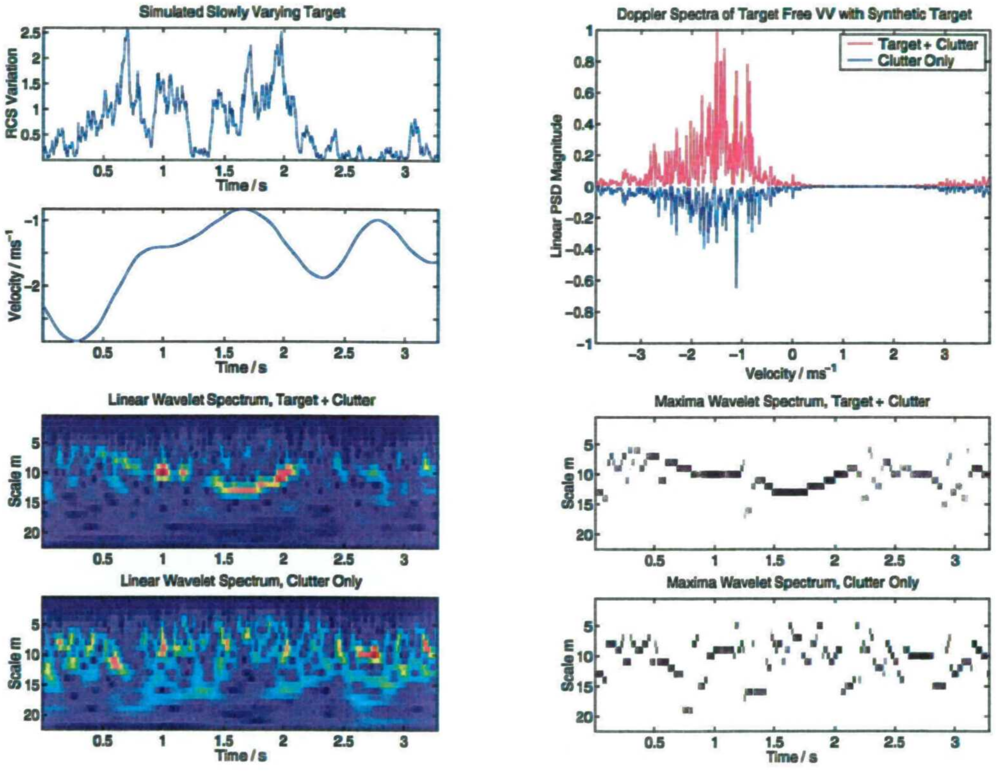

Figure 7.9: Wavelet processed low resolution VV data at PRF 156Hz demonstrates discrete nature of the clutter returns and detectability of a synthetic target based on the persistence of Doppler components

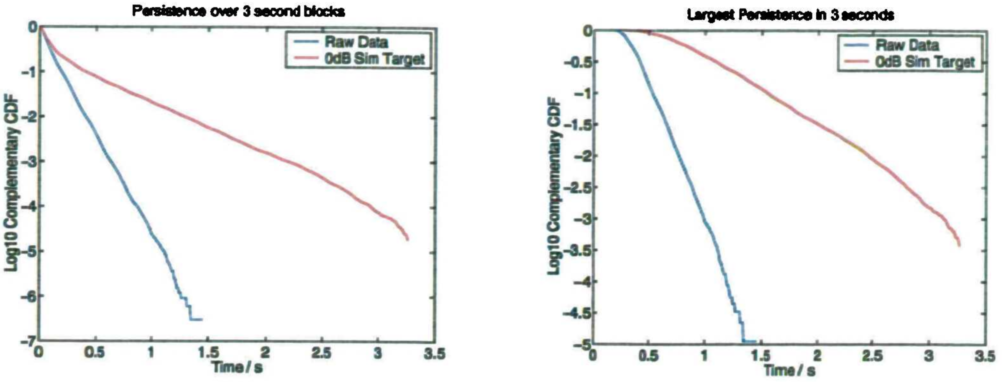

Figure 7.10: The persistence of events determined from raw VV sen3353 with addition of a 0dB synthetic target

without explicit edge effects.

- This smoothness allows the maxima to be extracted well and their continuity can be measured quickly with fast array operations (the entire processing takes the order of 100ms).
- The mean persistence of the clutter events is the order of 0.1s which is physically reasonable. The target plus clutter has far longer persistence events as expected.

The measured persistence lengths are drawn from a statistical distribution. As the persistence is presumably measuring some form of physical clutter event lifetime then one would expect a classical exponential distribution. This is exactly the assumption made by Lee et al. (1995) who assumed a fast moving varying scatterer population with an exponentially distributed lifetime to yield a Lorentzian distributed component within the long term Doppler spectrum.

Processing of the VV 20 minute data of sen3353 in 3 second blocks confirms that the distribution of the persistence is exponential and, whilst not shown, the form was stable over time. Figure 7.10 shows this graphically and also demonstrates the distribution obtained when the largest persistence value is taken as a detection statistic - additionally the effect of a simulated 0dB target is shown, detailed analysis is performed in the next Section.

Note that HH was not processed in this manner since the Doppler spectrum was not single sided due to the undersampling discussed previously.

Whilst it is tempting to conclude that the exponential observed distribution confirms the physical lifetime of events in the data, white (or bandlimited to the wavelet filter bank) noise would produce

Table 7.1: Specifications of the High Bandwidth Malvern Data

Run	F/GHz	PRF	Pol	H_{wave}/m	V_{wind}/ms^{-1}	Rng/m	Resolution
3092257	~9GHz	1000Hz	All	1.2 - 2.4	8	1578	0.3m x 1024

an exponential persistence distribution since in this case the probability of any particular filter being a maximum is the same. If the probability of a particular maxima's presence is constant then the classical exponential decay distribution arises.

7.3.1 Processing a Real Target

The MBPR high resolution data is not reliable enough to give accurate performance figures for low Pfa. Whilst 20 minute runs were taken in high resolution mode, the low resolution runs were usually less than a minute long. This causes the sample size to be insufficient; the lowest Pfa will be the reciprocal of the sample size but significant statistical noise will be present.

It is regrettably accepted that different data must be sourced to experimentally demonstrate the use of this method. This was kindly provided by DERA Malvern and has properties detailed in Table 7.1.

Exact details cannot be quoted but the radar is operated from a clifftop at a low grazing angle - about 1.5°, which is significantly larger than the DERA Portsdown data. A major difference between this and the DERA Portsdown MBPR is that the high resolution is achieved by a single high bandwidth chirp (500MHz) which avoids the synthetic chirp cycling problems of the MBPR. Crucially a real target is present within this data - an oil drum floating upon the sea surface at the edge of the range profile. This gives an immediately applicable test for the persistence algorithm upon a typical target.

The oil drum is floating on its side which would suggest VV target returns would be very different to the HH since the latter is exposing more cross section in the direction of the radar. In practice there was significant movement of the drum and over 30 seconds one cannot make firm conclusions.

The Doppler spectra from 1 second of the VV and HH are shown in Figure 7.11 which demonstrates that the mean Doppler velocity of the VV returns is about $2ms^{-1}$ whilst that of HH is $1ms^{-1}$. This is in contrast with the literature, but both raw and processed files from DERA were checked and confirmed the labelling. There has been a suggestion (Lycett 2000) that the radar

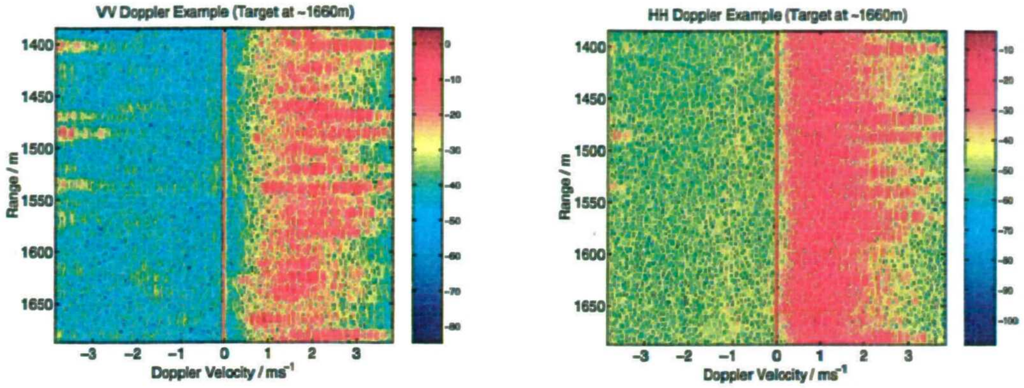

Figure 7.11: The Doppler spectra from 1 second of VV and HH Malvern data. Scale is arbitrary dB Intensity

channels were swapped in hardware but without firm evidence the file labels are retained.

A ship-borne radar may be unable to achieve range gate accuracy of 0.3m over 1 second due to platform motion as inertial sensors used for motion compensation are less sensitive to low frequency movements. At these Doppler velocities the persistence method cannot be expected to work since any potential target will have left the imaged rangecell before its lifetime can be assessed. Reprocessing the raw data to obtain 6m range cells (using DERA programs) means that any event lifetime can be measured reliably without excessively decreasing the sample size. The persistence method was then implemented at a PRF of 500Hz (the minimum necessary to resolve the Doppler spectrum) in 1 second blocks of data operating on a single range cell. The intensity was incoherently averaged on those same blocks to provide a comparison to ideal RCS thresholding (RCS calibration data was not available but the scenes were corrected for range dependence). Figure 7.12 and 7.13 presents processed images of the data showing arbitrary integrated intensity and the measured persistence values for the 1 second blocks. This Figure crucially demonstrates that for these conditions the persistence method is in some ways complementary to that of intensity thresholding. This scene is substantially more homogenous than previously analysed using the MBPR due to the lower wave height and higher grazing angle. A physical threshold has now been realised for target detection which will hopefully be more stable than that of RCS for low observable targets in sea clutter.

- The persistence threshold has been shown to work effectively in real data with a real target. The low intensity target areas are revealed which is the intended operation of this novel

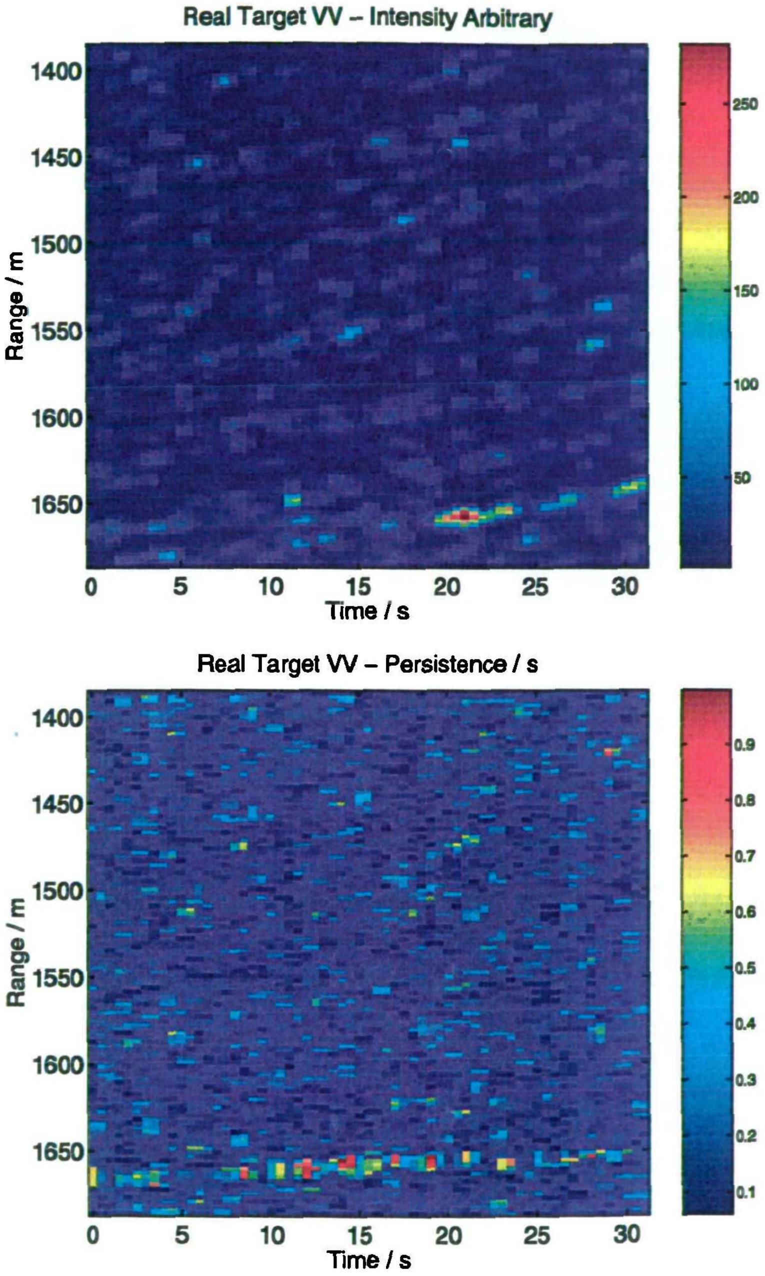

Figure 7.12: Malvern VV data. Using real data with a real target in the lower range cells the persistence method is complementary to Intensity and has the advantage of thresholding based upon a physical measure of time

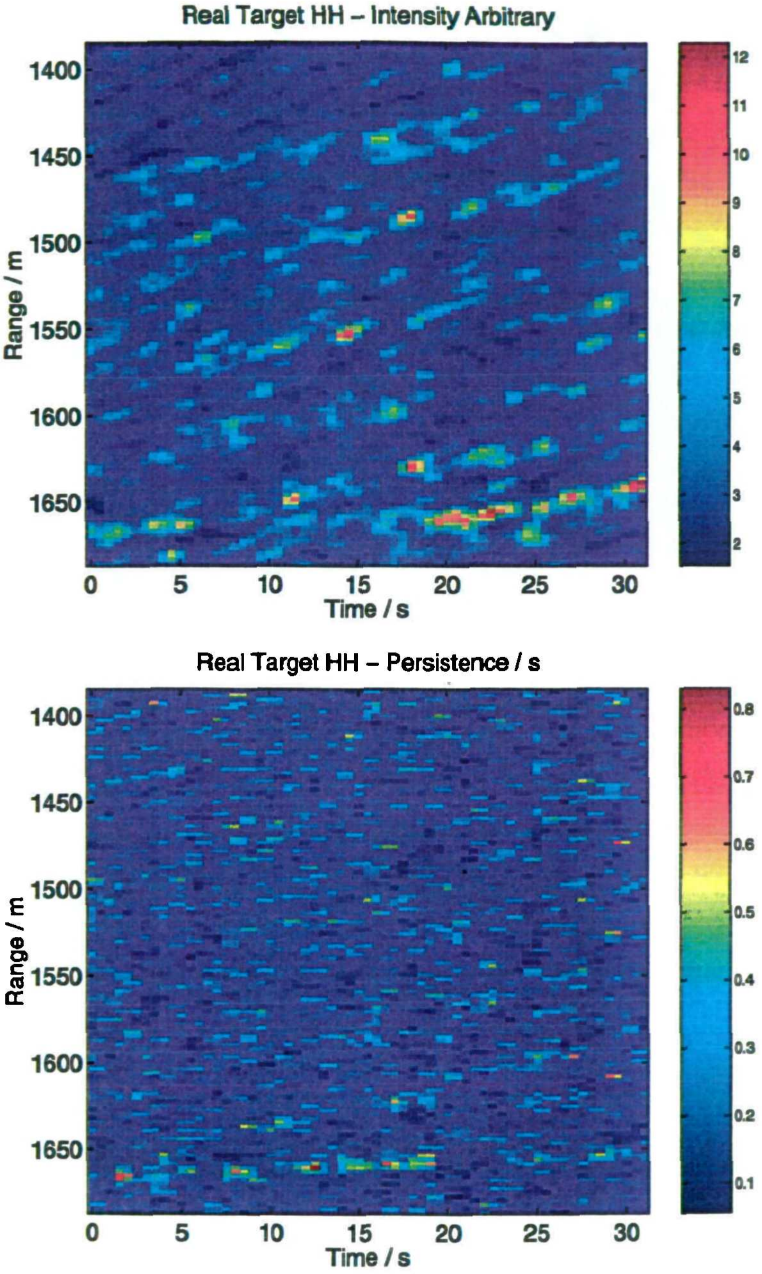
Real Target HH – Intensity Arbitrary

Real Target HH – Persistence / s

Figure 7.13: Malvern HH data. Using real data with a real target in the lower range cells the persistence method is complementary to Intensity and has the advantage of thresholding based upon a physical measure of time

method.

- Although only a single real target scene is available, the persistence method is complementary to intensity thresholding which gives strong justification for further testing.

In Figures 7.12 and 7.13 in the areas of low intensity from 5 - 20 seconds the persistence is large, however there is some indication that it is at a different angle to the intensity. This could indicate that the persistence algorithm is actually picking up a different velocity *disturbance* due to the target rather than the target itself, such as the wake.

- Identifying an effect due to a target presence may be just as effective as detecting the target itself but will be difficult to simulate.

7.3.2 Detection Performance of Simulated Target in Real Clutter

Radar Operating Characteristic (ROC) curves showing Pd versus Pfa are difficult to construct for real targets as the population of target samples is too low. Operating the detection method within the target free portion of the real data upon simulated targets injected into the clutter spectrum gives a suitable statistical population. The upper half of the scenes was chosen as the target free portion to avoid any residual range sidelobes from the target and for every second of each rangecell a random simulated target was injected at 0dB and 3dB relative power (determined individually and repeated 100 times). As described in Section 7.1, HH targets were injected at 1ms^{-1} with a width of 0.25ms^{-1} whilst VV targets were injected at 2ms^{-1} with a width of 0.5ms^{-1} each with a velocity time constant of 1 second. The RCS correlation time was matched to that of the observed target as 0.4s - a reasonable test since this is well below the 1 second observation time. The sample Pd and Pfa were then calculated by varying the ideal threshold for the entire scene to produce the ROC curves in Figure 7.14, note that the Pfa is calculated *per burst* (512 samples) to reflect the method.

- The persistence method is shown to significantly improve upon simple intensity thresholding for simulated targets in real VV data, an order of magnitude improvement is seen in Pfa over most of the ROC curve.

The performance curve for the HH is seemingly worse, however this is entirely in keeping with the existing theories - HH is spikier and so one could assume it has relatively more discrete events.

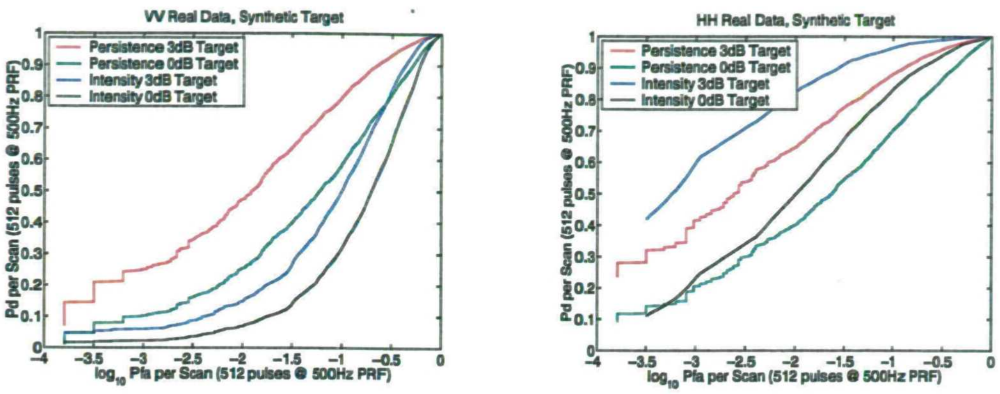

Figure 7.14: ROC curves for simulated target in real 9.75GHz data (1 second look time)

The straight addition of a target is perhaps not the correct way to analyse the performance. The striking point is the complementary performance shown by Figures 7.12 and 7.13 where the low intensity areas of the real target are revealed well by the persistence detector in both HH and VV polarisations.

7.4 Summary

This Chapter has implemented an entirely novel detection algorithm based upon the observed lifetime of discrete scatterer events, performance has been demonstrated on a real target in real clutter which is rarely shown in the literature. It is suggested that detection is not necessarily being performed on the target, but rather a surface disturbance relating to the target's presence. The complementary detection performance when compared with intensity alone is striking and the simulated ROC curves may not be applicable to performance calculations for this reason.

- It is a possibility that the radar files were incorrectly labelled and that VV and HH should be swapped. In any case conclusions cannot be drawn from a single file and the benefit of a persistence thresholding method, if any, requires further analysis of high quality data.

Chapter 8

CFAR Performance Limits in Littoral Clutter

Evidence has been provided that the observed high resolution sea clutter is not stationary with respect to its shape parameter and intensity. The operating conditions specify that a land edge may be present at an unknown position. It is necessary to determine when to adapt the processor to inhomogeneous statistics via the use of change point detection. This Section assumes that the resolution cell is large enough so that the discrete events are drawn from a continuous distribution such that the observed backscatter conforms to a K-distribution.

8.1 Operating Conditions

This Section outlines the conditions for which this analysis is applicable. A detection algorithm based on an uncertain distribution shape for is determined via:

1. Sweep for shape estimation for certain length of time, then global process.
2. Alternate sweep and process.
3. Concurrent estimation through censored statistics.

All at an operating throughput of order:

1. Desired Probability of False Alarm= 10^{-5}.

2. Desired Probability of Detection= 0.5.
3. 40 kHz PRF.

The significant difference of MFR is that it can operate in 'stare mode' so there is no physical limit on the number of samples available from a particular area such as in a rotating radar. Adaptive processing may include binary ('N from M' type detection), coherent (FFT bins) or incoherent (RCS summation) integration over time which could be applied to outputs such as:

- Single Hi-res Range-Intensity over a fixed 256 samples.
- Hi-res Range-Time-Intensity concerning time evolution of the above.
- Hi-res Range-Doppler possible from FFT of the above giving individual Doppler recordings from each range cell.
- Low res Intensity-Time.
- Low res instantaneous Doppler.
- Low res Doppler-Time such as FFT or WT.

In each case *a priori* knowledge of the distribution shape is unlikely and so this must be incorporated into the processing as a pre-detection stage. There is no reason not to simultaneously check for a target during this pre-detection stage but without accurate knowledge of the current statistics this will be sub-optimum.

The likely statistics from any of the environmental modes will be subject to:

- Abrupt edges due to land, birds, changing sea conditions, filter noise or extended targets
- Modulation from ocean waves, undulating land or slowly fluctuating targets
- A combination of the above from a discontinuous land-sea interface for example

The task is to determine the distribution shape reliably whilst these effects are occurring

8.2 Estimating the Distribution Shape

If clutter can be represented as locally stable Weibull or K-distributed, it must be known when adaptation to any deviation from these statistics is necessary. Chapter 3 concluded that the shape estimator U was of use in estimating the underlying statistics, and explicitly determined the effect of edges within the statistics.

$$\begin{aligned} U &= \widehat{\log I} - \log \hat{I} && (8.1)\\ &\equiv \log\left(\prod_{i=1}^{N} I_i\right)^{1/N} - \log\left(\frac{1}{N}\sum_{j=1}^{N} I_j\right) && (8.2)\\ &= \log \mu_G - \log \mu_A && (8.3)\end{aligned}$$

Whilst the fractional accuracy of the U parameter estimator can be calculated with respect to the underlying distribution, an accuracy relating to the performance of the detection algorithm is preferable - two bounds must be set for this. After discussion with DERA (Branson 1999), adaptation is defined to be necessary when:

- Probability of False Alarm is increased by an order of magnitude from the design as the spikiness has been underestimated. In an MFR this represents the level where an automated processor would be overloaded.
- Detection threshold is increased causing the required target SNR for detection to be 3dB greater than necessary as the spikiness has been overestimated. This would represent a significant stealth capability and is broadly equivalent to raising the CFAR threshold by 3dB i.e. a CFAR loss.

8.2.1 Performance Limits

As each of the bounds will have a dependence on the exact form of CFAR processor used, a simplification assumes that absolute knowledge of the local mean μ_0 is available. This can never be true but the uncertainty in the shape parameter will always be greater than the local mean if estimated concurrently. This Chapter indicates the loss from poor knowledge of the distribution shape and would apply in estimating U from a finite number of samples and then processing with a relatively large CFAR window in homogenous clutter.

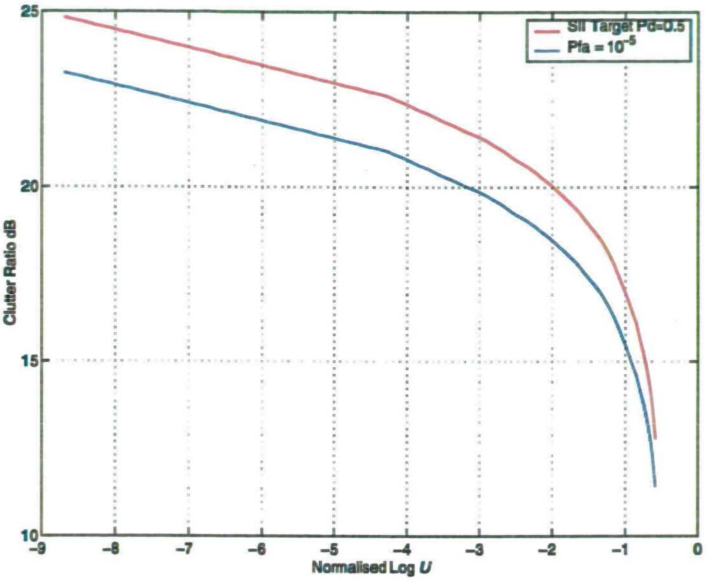

Figure 8.1: Threshold multiplier for Pfa= 10^{-5} and SNR required for Swerling II Pd= 0.5 with respect to U

Assume testing for a point-like target in K clutter, the clutter will have a CDF of

$$C_C\left(I;\nu,\mu\right)=\frac{2}{\Gamma\left(\nu\right)}\left(\frac{\nu I}{\mu}\right)^{\nu/2}K_\nu\left(2\sqrt{\frac{\nu I}{\mu}}\right)\tag{8.4}$$

from which inversion gives the threshold multiplier α that sets the critical threshold $I_c=\alpha\mu$ for a specific $\mathrm{P_{fa}}$.

Further assuming a Swerling II target model (Section 2.1.4) of cross section r, the target plus clutter distribution can be shown to be (Watts 1987)

$$P_{T+C}\left(I;\sigma,\nu,\mu\right)=\int_0^\infty\frac{1}{\Gamma\left(\nu\right)}\frac{\nu}{\mu}\left(\frac{\nu\sigma}{\mu}\right)^{\nu-1}\exp\left[\frac{-\nu\sigma}{\mu}\right]\frac{1}{\sigma+r}\exp\left[\frac{-I}{\sigma+r}\right]d\sigma\tag{8.5}$$

redefining to normalise target SNR $R=r/\mu$ and in the operating conditions of large $\mathrm{P_d}$ and low $\mathrm{P_{fa}}$ the above is adequately approximated by

$$P_{T+C}\left(I;\sigma,\nu,\mu\right)\approx\frac{1}{1+R}\exp\left[\frac{-I}{1+R}\right]\tag{8.6}$$

Figure 8.1 shows the required threshold I_c and target strength R with respect to U.

Before setting this threshold the degree of uncertainty must be known in the estimate of U and furthermore it is better to relate that degree of uncertainty to the chance that the processor is outside the operational bounds of U defined from the previous Section.
To calculate the positions of the bounds, notation is defined:

U' : The estimated value of the normalised log measure, subject to variation from finite sample numbers or inhomogeneity

U_0 : The actual normalised log measure of the underlying statistics

ν' : The estimated shape parameter determined from U', for this exercise K distributed statistics are assumed

ν_0 : The actual shape parameter of the K distributed statistics

α' : The threshold multiplier applied from the estimated statistics such that the threshold used is $I_c = \alpha' \mu_0$

α_0 : The threshold multiplier required such that a threshold at $I_c = \alpha_0 \mu_0$ would give a Pfa = 10^{-5}

The operational bounds of U are calculated from the high and low limits where U' causes either Pfa to increase by a factor of 10 ($U' = U_H$) or required target SNR is increased by 3dB ($U' = U_L$) when the true statistics are U_0.

1. $U' > U_0$ such that $P_{fa}\left(U = U'_H\right) = 10 \times P_{fa}\left(U = U_0\right)$. U_H is calculated from a numerical search by inverting Equation 8.4 to find which U' will give an α' that in statistics of ν_0 produces $P_{fa} = 10^{-4}$.

2. $U' < U_0$ such that $R\left(U = U'_L\right) = 3dB + R\left(U = U_0\right)$ is calculated from a numerical search by inverting Equation 8.6 to find which U' sets a threshold I_c that requires R to be 3dB greater than necessary to be detected when embedded in statistics of ν_0

Figure 8.2 shows the above graphically for the case of spiky and exponential clutter. Figure 8.3 plots the upper and lower bounds to U' over $-10 < U \leq -\gamma$ (including the exponential case). Remarkably these bounds are a near linear function of U for these values and presumably could be calculated by finding some linear approximation of the above procedures. Where U' is evaluated

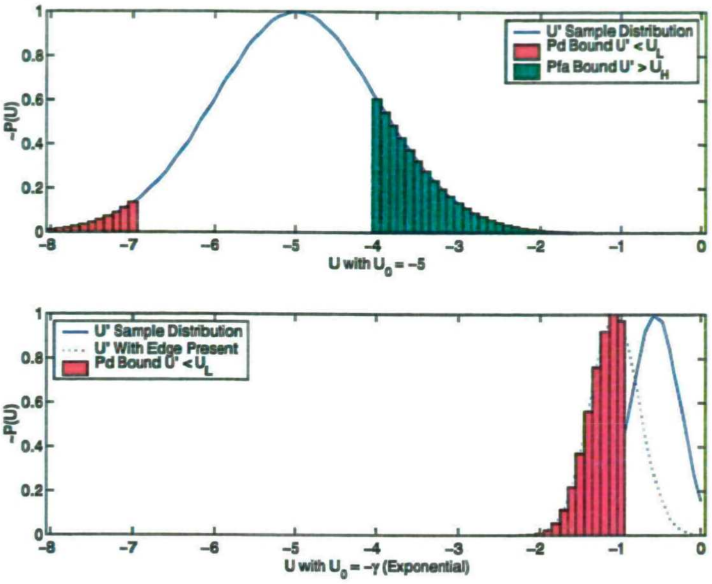

Figure 8.2: Upper plot: Maintaining Pfa by estimating the distribution from finite sample number is most difficult. Lower plot: In exponential clutter intensity edges cause major Pd loss

at the upper or lower bound this suggests

$$\frac{\partial U'}{\partial U} \equiv \frac{\partial U'}{\partial \nu'}\frac{\partial \nu'}{\partial \nu}\frac{\partial \nu}{\partial U} = \text{constant} \tag{8.7}$$

Equation 8.7 is difficult to prove but is not necessary since the calculation takes less than a second; it only needs to be performed once to determine just *how* accurate U' need be. Operating at $U' = -\gamma$ (i.e. estimate is of exponential clutter) the limit where excessive false alarms occur is $U_0 \simeq -0.62$ ($\nu_0 \simeq 12.5$)

From Figure 8.3, the Pfa bound U_H is much closer to U_0 than that of probability of detection U_L. Knowing that the sample estimate of U is Gaussian (Section 3.1.4), it is reliably concluded that the major problem of using finite sample number is maintaining the Pfa bound over all considered regions. In fact the necessary accuracy required to maintain the Pfa bound gives an expected target SNR loss of about 1dB, which is relatively small.

The major contribution to target loss is not from the finite sample number, but rather from the presence of intensity edges in a scene which cause $U' < U_0$ in an otherwise stable local distribution with a single change in underlying mean. The effect was explicitly determined in Section 3.1.4 to

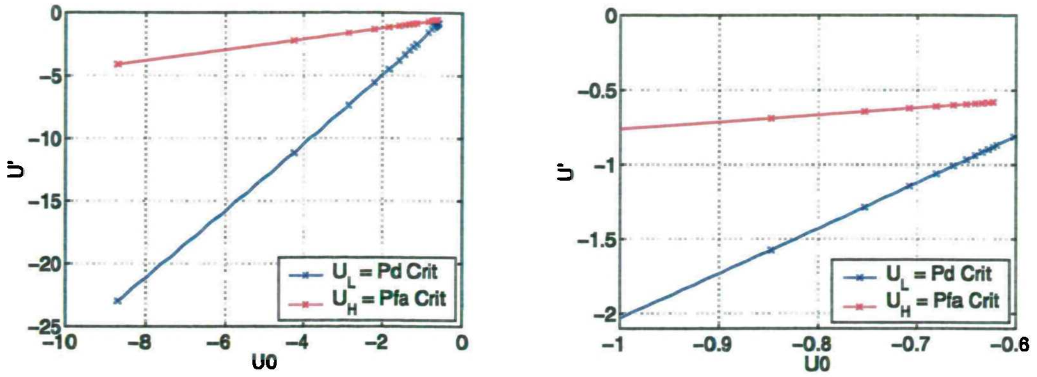

Figure 8.3: The critical values of estimated U for reduced Pfa and increased Target SNR are linear in U

cause a change in U whose expected value is

$$\langle U_{\Delta I}\rangle = \ln\frac{s^p}{sp-p+1} \tag{8.8}$$

where a step of magnitude s is present over a proportion p of the scene.

Knowing the required accuracy bounds for U derived in the previous Section, the above equation (previously shown in Figure 3.8) can be inverted for particular step sizes and proportions that cause the Pd bound to be reached . Figure 8.4 shows that a 5dB step over a significant fraction of exponential clutter causes a target loss of 3dB (due to erroneously determining $U' \simeq -0.72$, $\nu' \simeq 3.6$)

Whilst intensity edges can only cause U_0 to be overestimated producing a target loss from the Pd bound, edges arising from a change in distribution shape are dependent on the direction of change and so both the Pd and Pfa bound can be met. Section 3.1.4 showed the expected change in U to be

$$\langle U_{\Delta U}\rangle = p\,(U_T - U_0) \tag{8.9}$$

where a step of distribution shape U_T is present over a proportion p of the scene that is of underlying distribution U_0.

Again this can be inverted to determine when the Pd and Pfa bounds are reached and plotted in Figure 8.5 which shows that in estimating the distribution of very spiky clutter a corrupting exponential edge is required over more than half the scan to cause U to be underestimated enough to reach the Pfa bound. Operating a K estimator in exponential clutter shows that a corrupting

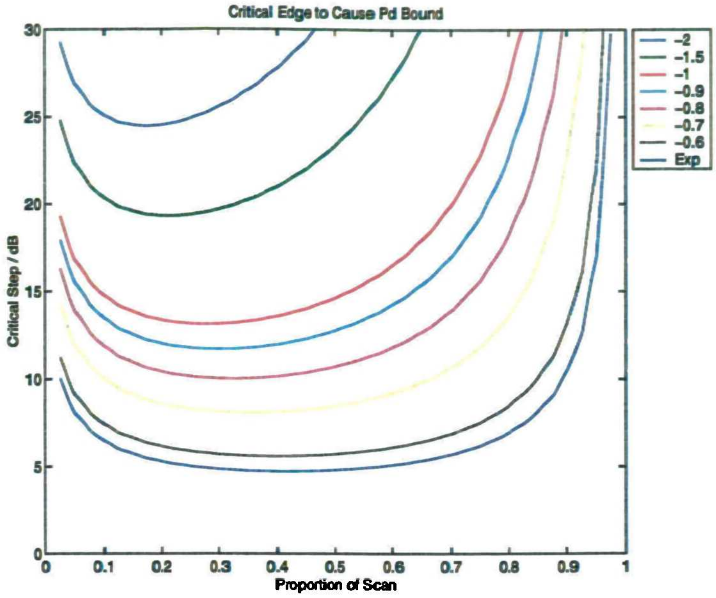

Figure 8.4: The required intensity and extent of a step to cause a 3dB loss in detection for varying U

spiky edge is required to be $U_T \lesssim -0.85$ ($\nu \lesssim 2$) over about half the scan to cause a critical Pd loss.

It can be concluded that estimating the distribution shape through U is relatively robust to typical distribution edges encountered but is fairly sensitive to intensity edges that can only cause the Pd bound to be reached. The Pfa bound is largely due to the finite number of samples available causing excessive variance in the estimate of U.

The effect of an edge showing a simultaneous change in both intensity and distribution can be calculated similarly but this has not been explicitly plotted since the parameter range is difficult to show. The necessary inversion is the combined effect of both $U_{\Delta I}$ and $U_{\Delta U}$ to give

$$\langle U_{\Delta I,\Delta U}\rangle = \log\frac{s^p}{sp-p+1} + p(U_T - U_0) \qquad (8.10)$$

which gives the expected change in the measured U' from an underlying homogenous distribution of shape U_0 subject to a simultaneous edge over a proportion p which consists of a distribution U_T of relative scale s. If $U_T < U_0$ the combined effect is greater than an intensity change alone; if $U_T > U_0$ the effect is lessened. The overall effect is dependent upon the nature of the changing environment.

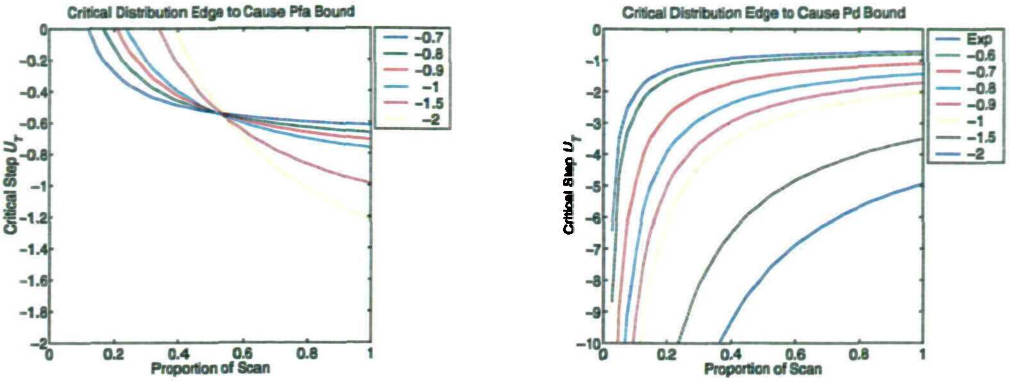

Figure 8.5: A contaminating edge of distribution U_T within the sample can cause a critical Pfa and Pd bound dependent upon U_0

The simple effect of a global modulation upon the local distribution was also determined in Section 3.1.4. An underlying homogenous distribution of shape U_0 modulated by a global distribution U_ρ (which may be correlated over a number of samples ρ) will give a bias to measured U' of

$$\langle U_{\Delta\rho} \rangle = U_\rho \tag{8.11}$$

if the expectation is carried out over a number of samples $N \gg \rho$. The operational bounds can be easily calculated from previously determined Figure 8.2.

8.2.2 Operating Implications from Shape Estimation

The derivations covered in the previous Section define rules for operating in potentially spiky clutter. To summarise:

- Required threshold, and thus absolute target detectability, is a function of the distribution shape.
- A desired Pfa of 10^{-5} is required, the final processor can withstand 10^{-4} before being overloaded.
- The shape estimator must not cause an additional loss of greater than 3dB in target detectability.

Which shows:

- Assuming exponential clutter ($U' = -\gamma$), an ideal exponential processor can withstand homogenous $U_0 \gtrsim -0.62$ ($\nu \gtrsim 12.5$) without excessive false alarms.
- Assuming K distributed clutter ($U' = U$), an ideal U estimator is susceptible in exponential clutter with intensity edges greater than about 5dB intensity over a significant proportion of the scan. This causes the estimator to overestimate spikiness as $U \lesssim -0.72$ ($\nu \lesssim 3.6$) causing a Pd loss in homogenous exponential clutter.
- Assuming K distributed clutter ($U' = U$), an ideal U estimator is susceptible in spiky clutter with a contaminating exponential edge over more than half the scan. This causes the estimator to underestimate the spikiness causing excessive false alarms in homogenous spiky clutter.
- Assuming K distributed clutter ($U' = U$), an ideal U estimator is susceptible in exponential clutter when, for example, the contaminating edge has distribution $U_T \lesssim -0.85$ ($\nu \lesssim 2$) over more than half the scan or $U_T \lesssim -2$ ($\nu \lesssim 0.5$) over 10% of the scan. This causes the estimator to overestimate spikiness causing excessive Pd loss in homogenous exponential clutter.

The numerical values are obviously a function of the goal Pfa and the bounds set for robustness, however due to the inherent uncertainty in estimating the statistics and allowing for the presence of edges the rule will always be encountered that describes *when* to operate detectors tailored for spiky clutter. A decision has to be made that balances the increased false alarm rate with lower target detection.

The type of edges present will be dependent upon the environment. Their major effect is Pd loss from intensity edges, these cannot contribute to Pfa loss. Pfa loss can only occur in the unlikely case of a distribution edge occupying more than half the scan.

The primary contribution to Pfa loss is from estimating U from finite sample number and the probability of this occurring is shown in Figure 8.6 which additionally shows that the probability of Pd loss is insignificant in comparison. If $U_0 \simeq -0.7$ and an estimate U' is made from 256 samples there is about a 10% probability of suffering Pfa loss on subsequent processing in homogenous U_0 based on this figure.

For comparison, the key result of Figures 8.4 have been repeated for goal Pfa of 10^{-2}, 10^{-3} and 10^{-4} in Figures 8.7 to 8.9. This indicates that if steps of order 10dB are expected in the data then a

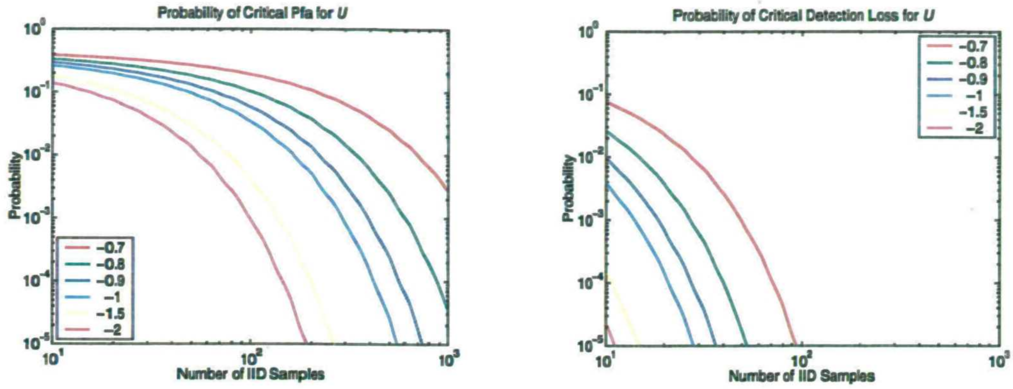

Figure 8.6: The probability of critical Pfa varies with U and number of samples N, probability of critical Pd is negligible in comparison

goal Pfa of 10^{-2} is more realistic without a form of simultaneous edge detection *before* any target detection algorithm based on spiky clutter occurs.

8.3 Summary

From this Chapter conclusions are:

- The region most sensitive to operating conditions is the transitional stage where clutter is expected to be spiky and one must operate an accurate estimator U', but where the true distribution U_0 is unknown due to finite sample numbers - the primary risk is Pfa loss.
- The number of samples observed must be maximised, but in the process the statistics could change. Small changes in the local mean intensity risk Pd loss.
- The optimum estimator with regard to minimising Pd and Pfa loss is one which repeatedly samples the statistics until a change occurs.

The results demonstrate that a step as small as 5dB can have a dramatic effect in mildly spiky data and so simultaneous edge detection must be performed in littoral regions *before* operating target detection based on a determined background shape parameter otherwise a critical Pd loss of 3dB will result.

Further decisions must be made based on the resultant local stationarity of the distribution, since a change in the intensity could be an inherent modulation of the statistics as in the compound formu-

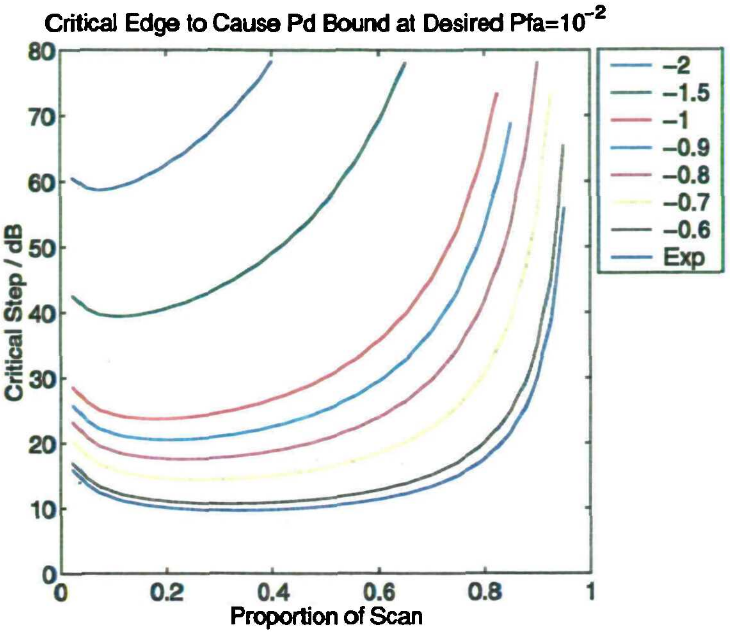

Figure 8.7: Required step to cause a 3dB loss in detection; goal Pfa= 10^{-2}

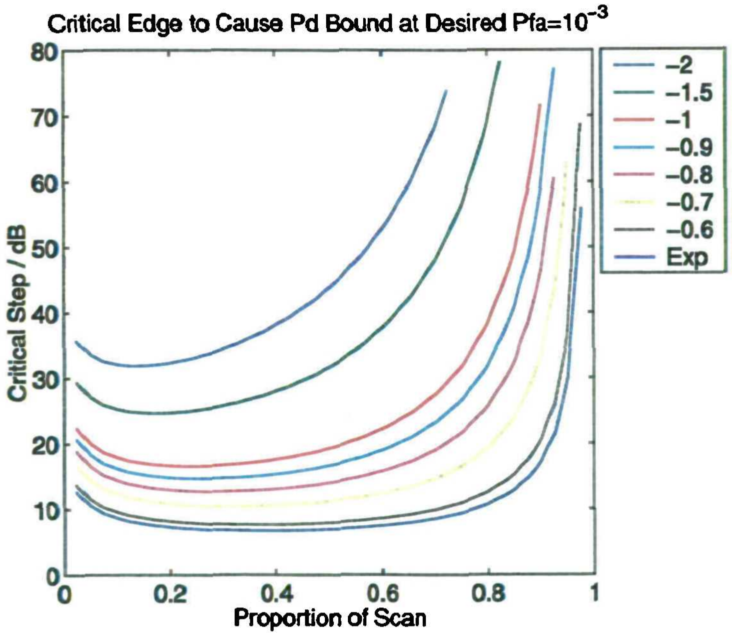

Figure 8.8: Required step to cause a 3dB loss in detection; goal Pfa= 10^{-3}

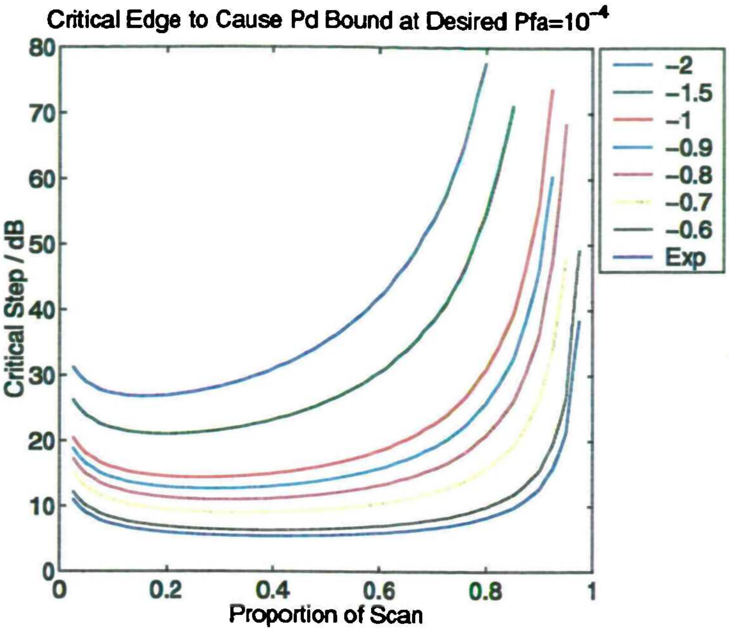

Figure 8.9: Required step to cause a 3dB loss in detection; goal Pfa= 10^{-4}

lation of the K distribution. There is still a requirement to know the extent of the local stationarity, as this would presumably define a suitable fixed CFAR window to use in target detection.

Chapter 9

Final Conclusions and Further Work

9.1 Summary

This thesis was begun with the aim of analysing a large amount of sea clutter taken in varying conditions to determine the best way of filtering the data for detection of slow moving targets. As with all experiments, some problems are expected with the data and this reduced the number of acceptable files considerably so that a large comparative study was not feasible.

It is hoped that the analysis has shown the data to vary considerably over typical detection times. Throughout this thesis emphasis has been placed on the need to account for this nonstationary behaviour. Commonly applied probability distributions may apply to radars that observe large range cells where so many scattering events are occurring that it can be viewed as a continuous distribution; this assumption breaks down at the high resolution of the MBPR as it is in a position to observe the individual scattering events. Whilst the compound formulation of the K distribution gives a way of modelling the underlying swell, any detection scheme relying on explicit knowledge of such a continuous distribution will not be applicable at this high resolution.

One contribution of this thesis is that the spatial extent of system noise has not been recognised before, often about a third of a high resolution scene was pure system noise and thus one could apply standard detection techniques if this area could be detected. A variable threshold multiplier is suggested if standard CFAR is used, but some form of sequential edge detector is required for this to be feasible since the large 20dB steps due to sea clutter easily pollute a fixed window scheme. A method of calculating the distribution of K distributed clutter plus system noise was proposed for calculating such thresholds.

Viewing the scene as a finite collection of events means that the traditional method of intensity thresholding based upon the local mean may not be feasible and additional criteria may be better at distinguishing a target from the background. Based upon a wavelet analysis, measurement of the scattering event lifetimes was performed to give a simple physical test statistic - the persistence. This is far from a final solution, it primarily gave a way of handling a coherent signal that can not be viewed as stationary over the order of the look time. Critically the wavelet based detector was complementary to that based upon intensity alone which, in conjunction with the processed scenes, suggested that a disturbance due to the target was detected.

If a large enough rangecell is used so that the returns are approximated by a K distribution, the littoral environment presents a particular problem in estimating the shape parameter where a land edge is present in the data. The performance limits of this were explicitly calculated and demonstrate the need for care when assessing a region such that edge detectors should always be employed in conjunction with shape estimators.

Two techniques were proposed to handle nonstationary behaviour in range resolution and Doppler mode, and they have necessarily been tested on a small amount of acceptable data. This thesis demonstrates the difficulty in analysing sea clutter when the environment has such a strong impact upon the observed statistics. For this reason it is sincerely hoped that a large amount of experimental data is released into the community so that hypothesised detection schemes can be assessed for robustness in achieving their claimed performance in fully specified statistics.

9.2 Further Work

It is suggested that the persistence algorithm is tested on many more real targets to determine if the performance is always complementary to intensity thresholding. A better time-varying target model could be adopted to describe the statistics and thus obtain applicable performance curves. Due to the expected presence of edges in littoral range swaths, an edge detector must be operated in parallel with any shape estimator and thus a scheme could be derived that balances the risk of an edge being present with the necessity of using more range samples to improve statistical accuracy. To maximise accurate sample population used, suggested further research is to formulate a fully sequential detection scheme which can adapt on a pointwise basis to the environment without the notion of a sample window. This would be ideally suited to the MFR since a decision can be declared at any moment with an associated *risk* balanced by a known *cost* of taking another sample to decide. This is a fruitful area because it links the notion of radar load and scheduling directly with the detection schemes. If multiple airborne targets are present then it is vital to assess potential naval targets faster than normal.

For the high resolution scenes, if 1% of the scene is dominated by 'events' which persist of the order of seconds but are difficult to characterise in RCS, then the Pfa will be discontinuous at the 1% level. If an event based analysis is used rather than a continuous spectrum then time varying properties of the events may provide better discrimination and the CWT method was presented as a first attempt. The test statistic (persistence) is a physical time and so it may turn out to be more stable than intensity based statistics. Further analysis of high quality (target free) data is required to determine this as the analysis has been based on only one real target scene.

9.3 Endnote

As a final note I believe that this work has achieved my personal objective of 'exploring' a subject. Far from my expectation that my view of science would narrow - I have been exposed to a much greater breadth of work than my degree studies and this thesis represents a small fraction of what I have learned since joining UCL.

References

Abramowitz, M. and Stegun, I. (eds) (1972). *Handbook of Mathematical Functions*, Vol. 10, Dover, New York. ISBN: 0-486-61272-4.

Anastassopoulos, V., Lampropoulos, G., Drosopoulos, A. and Rey, M. (1999). High resolution radar clutter statistics, *IEEE Trans. AES* **35**(1): 43–60.

Armstrong, B. (1992). *Processing Techniques for improved Radar Detection in Spiky Clutter*, PhD thesis, University College London.

Azzarelli, T. (1995). General class of non-Gaussian coherent clutter models, *IEE Proc. Radar Sonar Navig.* **142**(2): 61–70.

Baker, C. (1991). K-distributed coherent sea clutter, *IEE Proc. Pt. F* **138**(2): 89–92.

Berizzi, F., Greco, M. and Verrazzani, L. (2000). Fractal approach for sea clutter generation, *IEE Proc. Radar Sonar Navig.* **147**(4): 189–198.

Blacknell, D. (1994). Comparison of parameter estimators for K-distribution, *IEE Proc. Radar Sonar Navig.* **141**(1): 45–52.

Blake, A., Blacknell, D. and Oliver, C. (1995). High resolution SAR clutter textural analysis, *IEE Colloquium Recent Developments in Radar and Sonar Imaging Systems: What Next?*, IEE Colloquium Digest 1995/239, Savoy Place, London.

Boas, M. (1983). *Mathematical Methods in the Physical Sciences*, 2 edn, John Wiley & Sons, New York. ISBN 0-471-09960-0.

Branson, J. (1999). Personal communication from DERA Portsdown West.

Branson, J. (2000a). Personal communication from DERA Portsdown West.

Branson, J. (2000b). Naval environmental clutter attenuation and propagation specification - NECAPS 3, *Unclassified*, DERA Sea Systems.

Butler, J. (1998). *Tracking and Control in Multi-Function Radar*, PhD thesis, University of London.

Champeney, D. (1987). *A Handbook of Fourier Theorems*, Cambridge Univ. Press.

Cohen, A. (1965). Maximum likelihood estimation in the Weibull distribution based on complete and censored samples, *Technometrics* **7**: 579–588.

Cox, D. and Hinkley, D. (1974). *Theoretical Statistics*, Chapman & Hall, London, UK. ISBN: 0-412-16160-5.

Cromwell, J., Labys, W. and Terraza, M. (1994). *Univariate Tests for Time Series Models*, SAGE Publications Ltd., 6 Bonhill St, London. ISBN: 0-8039-4991-X.

Croney, J. (1970). Civil marine radar, *in* M. Skolnik (ed.), *Radar Handbook*, McGraw-Hill, New York, chapter 31.

d'Amore, L., Laccetti, G. and Murli, A. (1999). An implementation of a Fourier series method for numerical inversion of the Laplace transform, *ACM Transactions on Mathematical Software* **25**(3): 279–305.

Daubechies, I. (ed.) (1993). *Proceedings of Symposia in Applied Mathematics*, Vol. 47, American Mathematical Society, Providence, Rhode Island.

de Hoog, F., Knight, J. and Stokes, A. (1982). An improved method for numerical inversion of Laplace transforms, *SIAM J. Sci. and Stat. Comput.* **3**: 357–366.

di Bisceglie, M., Galdi, C. and Griffiths, H. (1999). Statistical scattering model for high-resolution sonar images: characterisation and parameter estimation, *IEE Proc. Radar Sonar Navig.* **146**(5): 264–272.

Dickey, D. and Fuller, W. (1979). Distribution of the estimators for autoregressive time series with a unit root, *Journal of the American Statistical Association* **74**: 427–431.

Duncan, J., Keller, W. and Wright, J. (1974). Fetch and wind speed dependence of Doppler spectra, *Radio Science* **9**: 809–819.

Dunn, J. and Howard, D. (1970). Target noise, *in* M. Skolnik (ed.), *Radar Handbook*, McGraw-Hill, New York, chapter 28.

El Mashade, M. (1996). Monopulse detection analysis of the trimmed mean CFAR processor in nonhomogeneous situations, *IEE Proc. Radar Sonar Navig.* **143**(2): 87–94.

Elias-Fuste, A. and de Mercado, M.G.and Davo, E. (1990). Analysis of some modified order statistic CFAR: OSGO and OSSO CFAR, *IEEE Trans. AES* **26**(1): 197–202.

Falls, L. (1970). Estimation of parameters in compound Weibull distributions, *Technometrics* **12**: 399–407.

Farge, M. (1992). Wavelet transforms and their application to turbulence, *Annu. Rev. Fluid Mech.* **24**: 395–457.

Farina, A. and Studer, F. (1986). A review of CFAR detection techniques in radar systems, *Microwave Journal* **29**(9): 179–185.

Farina, A., Gini, F., Greco, M. and Verrazzani, L. (1997). High resolution sea clutter data: statistical analysis of recorded live data, *IEE Proc. Radar Sonar Navig.* **144**(3): 121–139.

Fay, F., Clarke, J. and Peters, R. (1977). Weibull distribution applied to sea clutter, *Proc. IEE Conference Radar '77*, IEE Conf. Publ. 155, London, pp. 101–103.

Finn, H. and Johnson, R. (1968). Adaptive detection model with threshold control as a function of spatially sampled clutter-level estimates, *RCA Review* pp. 414–464.

Gandhi, P. and Kassam, S. (1994). Optimality of the cell averaging CFAR detector, *IEEE Trans. Inform. Theory* **40**(4): 1226–1228.

Gordon, L. and Pollak, M. (1997). Average run length to false alarm for surveillance schemes designed with partially specified pre-change distribution, *Ann. Statist.* **25**(3): 1284–1310.

Grossmann, A. and Morlet, J. (1985). *Decomposition of Functions Into Wavelets of Constant Shape, and Related Transforms*, World Scientific Publishing, Singapore.

Han, Y. and Kim, T. (1996). Performance of excision GO-CFAR detectors in nonhomogeneous environments, *IEE Proc. Radar Sonar Navig.* **143**(2): 105–112.

Hansen, J. and Cavaleri, V. (1982). High-resolution radar sea scatter, experimental observations and discriminants, *Technical Report 8557*, Naval Research Laboratory, Washington D.C.

Harter, H. and Moore, R. (1965). Maximum likelihood estimation of the parameters of Gamma and Weibull populations from complete and censored samples, *Technometrics* **7**: 639–643.

Harvey, A. (1993). *Time Series Models*, Harvester Wheatsheaf, London. ISBN: 0-7450-1200-0.

Haykin, S. and Puthusserypady, S. (1997). Chaotic dynamics of sea clutter, *Chaos* **7**(4): 777–802.

He, Y. (1994). Performance of some generalised modified order statistics CFAR detectors with automatic censoring technique in multiple target situations, *IEE Proc. Radar Sonar Navig.* **141**(4): 205–212.

Himonas, S. and Barkat, M. (1992). Automatic censored CFAR detection for multiple target situations, *IEEE Trans. AES* **28**(1): 286–304.

Hollenbeck, K. (1998). A Matlab function for numerical inversion of Laplace transforms by the de Hoog algorithm, http://www.isva.dtu.dk/staff/karl/invlap.htm.

Jakeman, E. (1980). On the statistics of K-dsitributed noise, *J. Phys. A.* **13**(2): 31–48.

Jakeman, E. and Pusey, P. (1976). A model for non Rayleigh sea echo, *IEEE Trans. Antennas Prop.* **24**(6): 806–814.

Jakeman, E. and Tough, R. (1981). Non-Gaussian models for the statistics of scattered waves, *Adv. Phys.* **37**: 471–529.

Jao, J. and Elbaum, M. (1978). First-order statistics of a non-Rayleigh fading target and its detection, *Proc. IEEE* **66**: 781–789.

Jiang, R. and Murthy, D. (1998). Mixture of Weibull distributions - parametric characterisation of failure rate function, *Applied Stochastic Models and Data Analysis* **14**: 47–65.

Johnson, J. (1928). Thermal agitation of electricity in conductors, *Phys. Rev.* **32**: 97–109.

Kaiser, G. (1994). *A Friendly Guide to Wavelets*, Birkhäuser, Boston.

Kao, J. (1959). A graphical estimation of mixed Weibull parameters in life-testing electron tubes, *Technometrics* **1**: 389–407.

Katzin, M. (1957). On the mechanism of radar sea clutter, *Proc. IRE* **45**: 44–54.

Kazakos, D. and Kazakos, P. (1990). *Detection and Estimation*, Computer Science Press, New York. ISBN: 0-7167-8181-6.

Komen, G., Cavaleri, L., Donelan, M., Hasselmann, K., Hasselmann, S. and Janssen, P. (1994). *Dynamics and Modelling of Ocean Waves*, Cambridge University Press, Cambridge.

Lee, P., Barter, J., Beach, K., Caponi, E., Hindman, C., Lake, B., Rungaldier, H. and Shelton, J. (1995). Power spectral lineshapes of microwave radiation backscattered from sea surfaces at small grazing angles, *IEE Proc. Radar Sonar Navig.* **142**(5): 252–258.

Lee, P., Barter, J., Caponi, E., Caponi, M., Hindman, C., Lake, B. and Rungaldier, H. (1996). Wind-speed dependence of small grazing-angle microwave backscatter from sea surfaces, *IEEE Trans. Antennas Prop.* **44**(3): 333–340.

Lee, P., Barter, J., Lake, B. and Thompson, H. (1998). Lineshape analysis of breaking-wave Doppler spectra, *IEE Proc. Radar Sonar Navig.* **145**(2): 135–139.

Lewis, B. and Olin, I. (1980). Experimental study and theoretical model of high-resolution radar backscatter from the sea, *Radio Science* **15**: 815–826.

Lombardo, P. and Oliver, C. (1994). Estimation of texture parameters in K-distributed clutter, *IEE Proc. Radar Sonar Navig.* **141**(4): 196–204.

Lombardo, P. and Oliver, C. (1995). Estimating the correlation properties of K-distributed SAR clutter, *IEE Proc. Radar Sonar Navig.* **142**(4): 167–178.

Lombardo, P., Oliver, C. and Tough, R. (1995). Effect of noise on order parameter estimation for K-distributed clutter, *IEE Proc. Radar Sonar Navig.* **142**(1): 33–40.

Long, M. (1983). *Radar Reflectivity of Land and Sea*, Artech House, Norwood, Mass.

Lycett, S. (2000). Personal communication from DERA Malvern.

Mathworks (1999). *MATLAB, The language of technical computing, PC Version 5.3*, The MathWorks, 24 Prime Park Way, Natick, MA 01760-1500.

Meyers, S., Kelly, B. and O'Brien, J. (1993). An introduction to wavelet analysis in oceanography and meteorology: with application to the dispersion of Yanai waves, *Mon. Wea. Rev.* **121**: 2858–2866.

Minkler, G. and Minkler, J. (1990). *CFAR*, 1st edn, Magellan Book Company, Baltimore, MD. ISBN: 0-9621618-1-0.

Money, D., Branson, J., Hooker, M. and Mobogunje, A. (1997a). Radar littoral environment predictions and measurements, *Proc. IEE Conference Radar '97*, IEE Conf. Publ. 449, Edinburgh, pp. 164–168.

Money, D., Mabogunje, A., Webb, D. and Hooker, M. (1997b). Sea clutter power spectral lineshape measurements, *Proc. IEE Conference Radar '97*, IEE Conf. Publ. 449, Edinburgh, pp. 85–89.

Nagle, D. and Saniie, J. (1995). Performance analysis of combined order statistic CFAR detectors, *IEEE Trans. AES* **31**(2): 522–533.

Nathanson, F. (1991). *Radar Design Principles*, 2nd edn, McGraw Hill. ISBN: 0-07-046052-3.

Olin, I. (1982). Amplitude and temporal statistics of sea spike clutter, *IEE Conf. Publ. 216*, pp. 198–202.

Oliver, C. (1993). Optimum texture estimators for SAR clutter, *J. Phys. D* **26**: 1824–1835.

Oliver, C., Blacknell, D. and White, R. (1996). Optimum edge detection in SAR, *IEE Proc. Radar Sonar Navig.* **143**(1): 31–40.

Page, E. (1955). A test for a change in a parameter occuring at an unknown point, *Biometrika* **42**(4): 523–527.

Percival, D. (1995). On estimation of the wavelet variance, *Biometrika* **82**: 619–631.

Pidgeon, V. (1968). Doppler dependence of radar sea return, *J. Geophys. Res.* **73**: 1333–1341.

Posner, F. (1998). Experimental observations at very low grazing angles of high range resolution microwave backscatter from the sea, *Technical Report NRL/MR/5310–98–8326*, NRL, Washington, D.C.

Press, W., Vetterling, W., Teukolsky, S. and Flannery, B. (1992). *Numerical Recipes in C*, 2 edn, Cambridge University Press. ISBN: 0-521-43108-5.

Priestly, M. (1988). *Non-Linear and Non-Stationary Time Series Analysis*, Academic Press, Oval Road, London. ISBN: 0.12.564910.X.

Rees, W. (1990). *Physical Principles of Remote Sensing*, Cambridge University Press. ISBN: 0521359945.

Rheinstein, J. (1968). Backscatter from spheres: A short pulse view, *IEEE Trans. Antennas Prop.* **16**: 89–97.

Rickard, J. and Dillard, G. (1977). Adaptive detection algorithms for multiple target situations, *IEEE Trans. AES* **13**(14): 338–343.

Roberts, S. (1966). A comparison of some control chart procedures, *Technometrics* **8**: 411–430.

Sarakar, T. and Su, C. (1998). A tutorial on wavelets from an electric engineering perspective, part 2: The continuous case, *IEEE Antennas and Propagation Magazine* **40**(6): 36–49.

Sekine, M. and Mao, Y. (1990). *Weibull Radar Clutter*, IEE Radar, Sonar, Navigation and Avionics Publ. RA003, IEE, London. ISBN 0 86341 191 6.

Shiryayev, A. (1963). On optimum methods in quickest detection problems, *Theory Probab. Appl.* **8**: 22–46.

Shrader, W. and Gregers-Hansen, V. (1990). MTI radar, *in* M. Skolnik (ed.), *Radar Handbook*, 2 edn, McGraw-Hill, chapter 15, pp. 15.1–15.7.

Skolnik, M. (1981). *Introduction to Radar Systems*, International Edition, McGraw-Hill, Singapore. ISBN: 0-07-066572-9.

Swami, A., Mendel, J. and Nikias, C. (1999). *MATLAB, Higher-Order Spectral Analysis Toolbox, Version 2*, The Mathworks, 24 Prime Park Way, Natick, MA 01760-1500.

Swerling, P. (1960). Probability of detection for fluctuating targets, *IRE Trans. IT* **6**: 269–308.

Tartakovsky, A., Kligys, S. and Petrov, A. (1999). Adaptive sequential algorithms for detecting targets in heavy IR clutter, *SPIE Proceedings: Signal and Data Processing of Small Targets*, Vol. 3809, Denver, CO.

Tonkin, S. and Dolman, D. (1990). Sea surface effects on the radar return from a periscope, *IEE Proc. Pt. F* **37**(2): 149–156.

Torrence, C. and Compo, G. (1998). A practical guide to wavelet analysis, *Bull. AMS ii* **79**(1): 61–78.

Tough, R. and Ward, K. (1999). The correlation properties of gamma and other non-gaussian processes generated by memoryless non-linear transformation, *J. Phys. D* **32**(23): 3075–3084.

Touzi, R., Lopes, A. and Bousquet, P. (1988). A statistical and geometrical edge detector for SAR images, *IEEE Trans. Geosci. Remote Sens.* **26**: 764–773.

Trizna, D. (1991). Statistics of low grazing angle radar sea scatter for moderate and fully developed ocean waves, *IEEE Trans. Antennas Prop.* **39**(12): 1681–1690.

Trunk, G. and George, S. (1970). Detection of targets in non Gaussian sea clutter, *IEEE Trans. AES* **6**: 620–628.

TW Research Ltd. (1999). Clutter classification to improve radar clutter rejection capabilities, *Technical Report R039/01*, DERA Contract SSDW1/793, Malvern.

Valenzuela, G. and Laing, M. (1970). Study of Doppler spectra of radar sea echo, *J. Geophys. Res.*

Wald, A. (1947). *Sequential Analysis*, Wiley, New York.

Walker, D. (2000). Experimentally motivated model for the low grazing angle radar doppler spectra of the sea surface, *IEE Proc. Radar Sonar Navig.* **147**(3): 114–120.

Ward, K. (1981). Compound representation of high resolution sea clutter, *Elect. Lett.* **17**(16): 561–563.

Ward, K., Baker, C. and Watts, S. (1990). Maritime surveillance radar; part 1: Radar scattering from the ocean surface, *IEE Proc. Pt. F* **137**(2): 51–62.

Watts, S. (1987). Radar detection prediction in K-distributed sea clutter and thermal noise, *IEEE Trans. AES* **23**: 40–45.

Watts, S. (1996). Cell-averaging CFAR gain in spatially correlated K-distribution clutter, *IEE Proc. Radar Sonar Navig.* **143**(5): 321–327.

Webb, D. (1999). Personal communication from DERA Portsdown West.

Weibull, W. (1939). *A Statistical Theory of the Strength of Materials*, number 153 in *Ingeniors Vetenskaps Akademiens Handlingar*, Royal Swedish IER.

Werle, B. (1995). Sea backscatter, spikes and wave group observations at low grazing angles, *IEEE International Radar Conference*, IEEE Publ. 95CH3571-0, Virginia, pp. 187–195.

Williams, G. (1997). *Chaos Theory Tamed*, Taylor & Francis, London. ISBN: 0-7484-0749-9.

Wolfram, S. (1996). Mathematica, a system for doing mathematics by computer.

Wornell, G. (1996). *Signal Processing with Fractals: A Wavelet-based Approach*, Prentice Hall, New Jersey.

Wright, J. (1966). Backscattering from capillary waves with application to sea clutter, *IEEE Trans. Antennas Prop.* **14**: 749–754.

Wright, J. (1968). A new model for sea clutter, *IEEE Trans. Antennas Prop.* **16**: 217–223.

Wylie, F. (ed.) (1978). *The Use of Radar at Sea*, 5 edn, Royal Institute of Navigation, Hollis & Carter, London, pp. 3–4. ISBN: 0 370 30091 2.

Appendix A

An Appropriate Use for the SED

The original need of the SED was to evaluate whether the returns could be considered locally exponential but without specifying any *a priori* knowledge about the correlation or an analysis window, this was an attempt to process the compound formulation of the statistics. Originally it was hoped that targets could be identified simply by being short extent steps in the local mean. It was concluded that the flotsam present upon the waves causes highly variable returns within the wave region that are locally uncorrelated and spikier than exponential which gave too many false alarms.

The system noise present in the scenes was however, accurately segmented which leads to a proposed Synthetic Aperture Sonar (SAS) target detection method.

A.1 Brief Outline of SAS

Obviously the primary difference is that acoustic waves are now used to image the scene. Synthetic Aperture processing, commonly applied in radar, gives improved azimuth resolution by utilising the motion of the sound emitter (hydrophone) as it passes a stationary target and coherently reconstructing the effect of a much larger aperture.

The particular SAS configuration was rail mounted to give accurate motion reconstruction and an array of 32 hydrophones were used with a range-Doppler imaging algorithm to obtain the image in Figure A.1 of 2 targets present upon the sea bed. The exact specifications of the SAS cannot be quoted, but the resolution is of the order of centimetres such that the range and azimuth extent of the image is about 10m in extent. The image is courtesy of DERA Bincleaves.

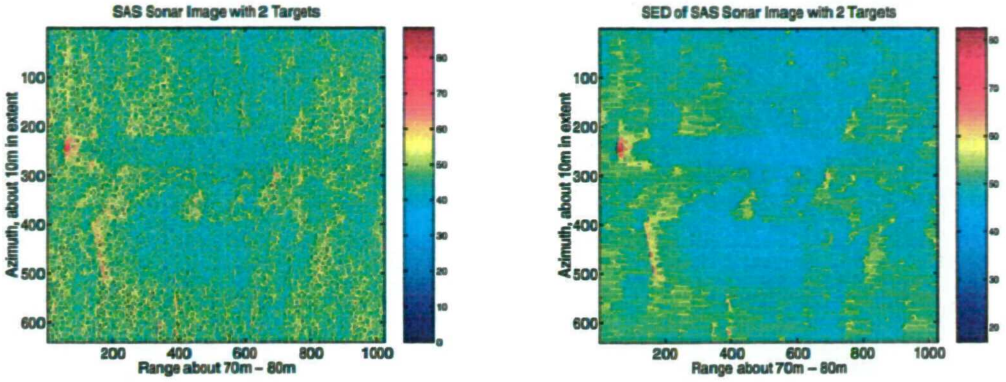

Figure A.1: Synthetic Aperture Sonar scene with 2 targets present; arbitrary dB intensity

A.2 The SED Applied to Detect Shadows

Figure A.1 shows that apart from the slowly modulating seabed texture the targets show as bright points, but crucially *they cast large shadows behind them.* The shadows are due to the nature of the targets, such that sound does not propagate behind them. The SED should be able to detect the regions if operated horizontally.

The initial detection procedure would be to use some sort of threshold in intensity to detect the bright targets, the SED however is applied to detect *long regions at the system noise level.* A suitable comparison is to threshold so that no false alarms are produced in the azimuth lines where no target is present. This is illustrated in Figure A.2 which clearly shows that processing the scene in both ways can detect the targets. The SED method produces a much better detection rate, especially for the lower target which shows up so clearly that it is complementary to simple high intensity thresholding.

This is believed to be a novel implementation of SAS target detection and the length of the shadows may give information about the height of the object.

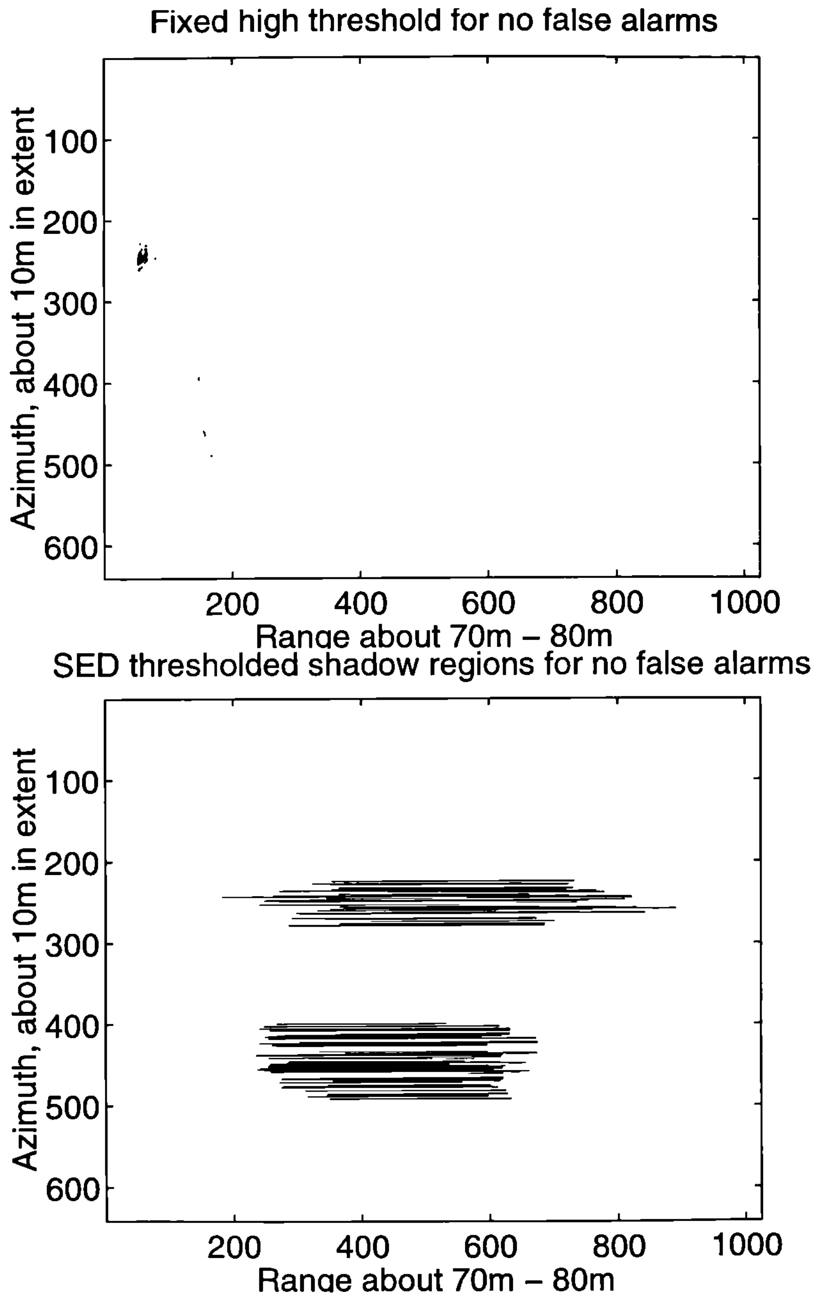

Figure A.2: The shadows of a target are detected more efficiently than the target itself using the SED

Appendix B

Mathematical Annex

B.1 Nonlinearity in the Normalised Log Estimate of ν

Lombardo and Oliver (1994) show that the first order expansion is insufficient to predict the U estimator's bias and variance and that a second order correction is only acceptable for $\nu < 1$. Importantly they identify the nonlinearity in the inversion of Equation 3.15 as causing the predicted error to deviate from simulation and suggest an improved texture measure of $t \equiv 1/\nu$ which tends to a linear dependence for large ν. Whilst individual assessments of texture are unaffected, the average bias error is reduced significantly and follows the theoretical prediction closely. The improved texture measure with associated bias and error is achieved by substitution of t and the expectation moments into Equation 3.15 - 3.17 (Lombardo and Oliver 1994).

$$U = \widehat{\log(I)} - \log \hat{I} = \psi^{(0)}\left(t^{-1}\right) + \log(t) - \gamma \tag{B.1}$$

$$\frac{\langle \Delta t \rangle}{t} \approx \frac{1 + 2t}{2n\left(1 - \frac{\psi^{(1)}(t^{-1})}{t}\right)} \tag{B.2}$$

$$\frac{\langle \partial t^2 \rangle}{t^2} \approx \frac{\psi^{(1)}\left(t^{-1}\right) + \frac{\pi^2}{6} - 1}{n\left(1 - \frac{\psi^{(1)}(t^{-1})}{t}\right)^2} \tag{B.3}$$

Simulation results using 10^4 trials based on $0.1 \leq \nu \leq 10$ are presented for various sample sizes. Figure B.1 shows the rms error in the estimator with the predicted value overlaid using Equation 3.17. Figure B.2 shows the differing behaviour of the bias for $\nu < 1$ and $\nu > 1$, with the predicted value overlaid using Equation 3.16.

An inability to invert some values of the estimator causes spurious results. For large values of ν

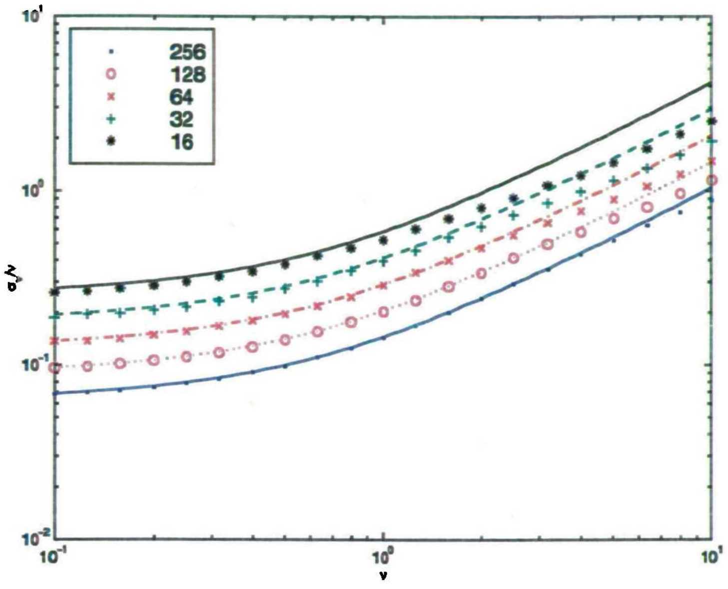

Figure B.1: RMS error in estimation of ν from U (Normalised Log Estimator)

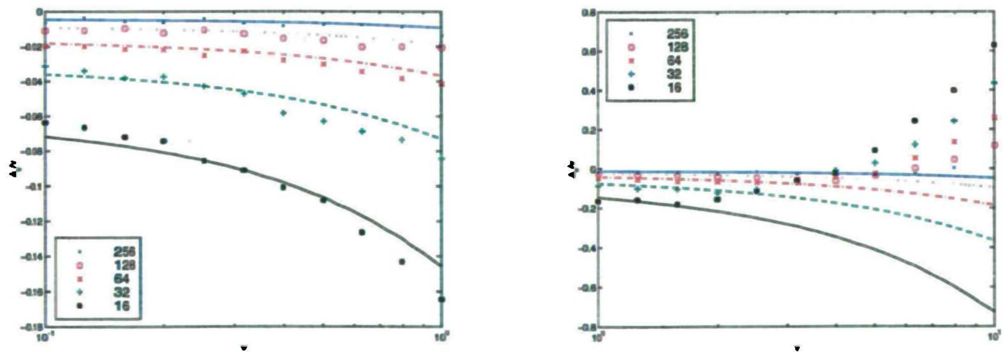

Figure B.2: Bias error in estimation of ν from U (Normalised Log Estimator)

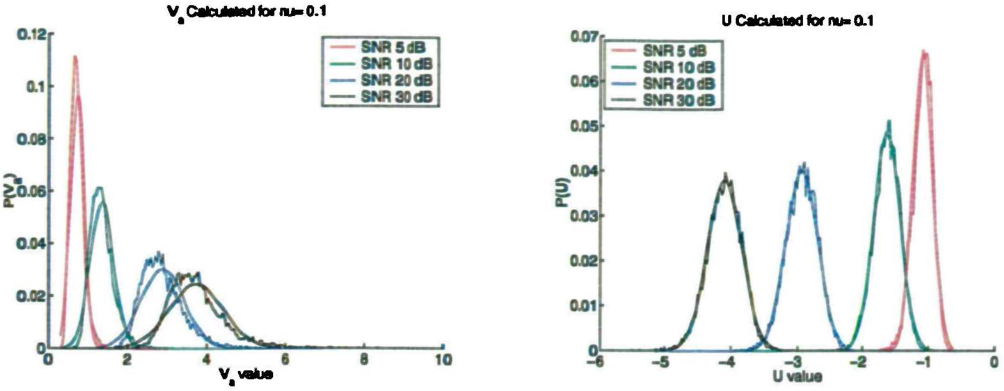

Figure B.3: Whilst V_a may be resistant to thermal noise, the sample distribution is not Gaussian - unlike U

the rms error is overestimated, reducing the sample population has the same effect. A gradient sign-change is seen in the bias indicating where the effect occurs.

B.2 Texture Estimate with Resistance to Additive Thermal Noise

Lombardo, Oliver and Tough (Lombardo et al. 1995) suggest that whilst the normalised log estimate U is an accurate estimator in pure K distributed clutter, it is particularly sensitive to additive thermal noise. To completely characterise the distribution of K plus noise they introduce the concept of a 'set of sufficient statistics' - mean, normalised log and intensity contrast. They propose a modified estimator, the amplitude contrast V_a, and show it has slightly less accuracy than U but has improved resistance to noise.

$$V_a \equiv \frac{\langle I \rangle}{\left\langle \sqrt{I} \right\rangle^2} - 1 \tag{B.4}$$

Figure B.3 shows the distribution for 256 samples at $\nu = 0.1$ has a significant deviation from normal in the sample distribution of V_a in contrast to U. The text of Section 3.1.5 justifies that U is still the preferred measure.

B.3 Chaos

Section 3.4.1 briefly introduced chaos. This annex discusses some of the concepts applied in the chaotic analysis of sea clutter. Many of the definitions are taken from Williams (1997).

B.3.1 Terminology

Extending to two or more dimensions requires the use of phase space to plot each independent feature recorded. Pseudo phase space is for representing a single feature under different conditions, for example lag space plots x_t vs x_{t+i} where i represents the lag time. Commonly, embedding dimension refers to the number of lags compared - for time prediction this is normally the number of points required before making a prediction.

An attractor is the locus of stability in a system's phase space. As shown in Figure 3.17 though any small deviation off this attractor may cause the system to flip onto another part of the attractor. The basin of attraction refers to the volume of phase space in which the attractor has an effect on the system, this leads to trajectories along which the system moves (asymptotically) onto the attractor. Multiple attractors can exist with overlapping basins causing complex fluid-like flow of the system within phase space.

Topology plays a major rôle in defining attractor shapes, including:

- point attractor: zero dimensions.
- periodic attractor or limit cycle: two or more values that recur in order occupying 2 dimensions in phase space.
- torus: combination of limit cycles. This can be further subdivided into periodic whereby the trajectories upon the torus exactly repeat themselves, or slightly mismatch as quasi-periodic.

The definitions of chaos are numerous, including:

- Chaos results from a deterministic process.
- It happens only in non linear systems.
- It can usually pass all tests for randomness.
- The range of variables have finite bounds restricting the attractor volume.
- Chaotic behaviour is hyper-sensitive to (but has no memory of) the initial conditions.
- Short term prediction is possible.
- Fourier spectrum is broadband but can have significant peaks.

B.3.2 Trajectories and Simulation

Trajectories in non-chaotic and chaotic systems for non-linear dissipative systems are fixed for a particular control parameter (k in Equation 3.73). In the non-chaotic regime the trajectories converge or remain equidistant. In the chaotic domain these trajectories diverge. The particular trajectory a system takes can be determined by any vanishingly small finite difference in phase space position, this causes problems in the computer simulation of chaos due to the inherent finite precision of floating point calculations and cannot be avoided by increasing this computational precision.

B.3.3 Order Within Chaos

Chaos consists of various regimes of order masked by random-like behaviour. Typical forms of order within chaos are:

- Windows: Regions of periodicity.
- Routes: Period doubling at the onset of chaos.
- Chaotic or Strange Attractor: complex phase space surface to which the trajectory is asymptotic in time.
- Zones of popularity: Regions of high probability density upon the chaotic attractor.
- Fractal structure: the chaotic attractor is a fractal.
- Self organisation: regular patterns emerging in space, time or function.

B.3.4 Reconstruction of Phase Space

This is generally only graphically possible in systems embedded in 3 or less dimensions. When systems of higher embedding dimension are compacted into lower dimensionality pseudo-phase space they produce false nearest neighbours. Upon increasing the embedding dimension they will abruptly move apart when the correct embedding dimension is used. This can be seen by an analogy of a cubic crystal viewed slightly offset in 2D, sites appear closer to each other than one would intuitively expect from the uniform planar positions. Viewing the crystal in 3D immediately reveals the orderly structure. A mathematical way of determining the correct embedding dimension is to maximise the mean square distance between points in phase space. This is the essence of the 'singular system analysis', 'singular value decomposition', 'Karhunen-Loeve decomposition' etc. methods. If an attractor exists then it should appear from the probability density distribution in the phase space reconstruction, however at erroneously low embedding dimensions spurious attractor shapes can be revealed. With any reconstruction it is important to provide a control dataset formed from 'surrogate data', this is non-deterministic artificially generated data that mimics certain features of the measured time series. Computer generated randomness is deterministic and extreme care must be exercised when generating surrogate data.

B.3.5 Dimensions

Dimensions are central to chaos theory. Authors seem to take various terms to mean different things; based on definitions in Williams(1997):

- Euclidean: the number of co-ordinates of a body needed to describe its shape.
- Topological: 1+euclidean dimension of simplest shape that can subdivide the body (usually equal to the euclidean dimension).
- Variable: the degrees of freedom of a system.
- N-d vector: above as applied to pseudo-phase space (i.e. the number of lags used).
- Embedding: the number of lagged values used for the purpose of reconstruction.
- Scaling exponent: not necessarily an integer, used to describe many properties of a system that conform to a 'simple' power law e.g. the Fractal dimension, Information, Correlation and Lyapunov.

Appendix C

Additional Statistical Detectors

In Chapter 5 Neyman Pearson detection schemes were used as Constant False Alarm Rate (CFAR) processors but several other methods are available.

C.1 Wald Tests

By extension of the Neyman-Pearson formulation the observation interval can be allowed to vary given the following assets:

- A realisation $x = \{x_{1:n}, n \rightarrow \infty\}$ from the radar output.
- Hypothesis H_0 indicative of no target, represented by PDF f_0
  versus hypothesis H_1 indicative of target present, represented by PDF f_1.
- A measure of performance based on P_d and τ_{fa}.

In conjunction with a *stopping rule* $\Delta_n(x)$ and a *decision rule* $\delta_n(x_{1:n})$

$$\Delta_n(x) = \Pr\{\text{decide to stop} \mid x_{1:n} \text{ observed}\} \tag{C.1}$$

$$\delta_n(x^n) = \Pr\{H_1 \text{decided} \mid \text{stopped}\} \tag{C.2}$$

a *sequential detection* scheme results. The explicit inclusion of P_d implies a lower bound on the number of samples required for decision, an upper bound is also set dependent on radar load. The basic Wald CFAR can be extended to detect a change in distribution as well as deciding upon the class of distribution, and is based a scheme due to Wald (Wald 1947):

Select two thresholds λ_0 and λ_1 such that $\lambda_0 \leq \lambda_1$. Start observing data $x_1, x_2, \ldots, x_n$ sequentially and calculate a measure T based upon the likelihood, stop at the *first* n such that

$$\begin{aligned} \text{Either } T\left(x_{1:n}\right) &= \log \frac{f_1\left(x_{1:n}\right)}{f_0\left(x_{1:n}\right)} \leq \lambda_0 && \text{(C.3)} \\ \text{Or } T\left(x_{1:n}\right) &= \log \frac{f_1\left(x_{1:n}\right)}{f_0\left(x_{1:n}\right)} \geq \lambda_1 && \text{(C.4)} \end{aligned}$$

where the subscripts on λ refer to the chosen hypothesis H.

The logarithm in Equations C.3 and C.4 allow a sequential updating step

$$T\left(x_{1:n}\right) = T\left(x_{1:n-1}\right) + \log \frac{f_1\left(x_n | x_{1:n-1}\right)}{f_0\left(x_n | x_{1:n-1}\right)} \qquad \text{(C.5)}$$

As detailed in Kazakos and Kazakos (1990) the λ threshold can be calculated as

$$\begin{aligned} \lambda_0 &\geq \log\left(\frac{1-\beta}{1-P_{fa}}\right) && \text{(C.6)} \\ \lambda_1 &\leq \log\left(\frac{\beta}{P_{fa}}\right) && \text{(C.7)} \end{aligned}$$

where β is a *power level* related to P_d. This provides an interesting comparison to NP where only P_{fa} is specified. Obviously increasing P_d will necessitate longer observation before deciding upon H. Additionally in sequential detection theory the probability of false alarm defines the *average run length* (ARL or τ_{FA}), giving the expected number of observations before a false alarm is observed.

The difficulty lies in determining T easily since it is not independent of $x_{1:n-1}$. Complications arise from:

- the class of the background PDF f_0 may be complicated.
- f_0 may be known only up to a *nuisance parameter* such as scale or shape.
- the test may be sub-optimal when H is estimated from the first observations.
- T may not be invariant to the parameters defining H.
- The target may only be present for a limited or variable number of samples.

C.2 Change Point Tests

As in Wald, a variable observation interval is available. The hypotheses are modified to allow a *change point* in the statistics, given the following assets:

- A realisation $x = \{x_{1:n}, n \to \infty\}$ from the radar output.
- Hypothesis H_0 indicative of no target, represented by PDF f_0 *versus* hypothesis H_1 indicative of target present from time k onwards, represented by PDF $f_0(x_{1:k-1})$ and $f_1(x_{k:n})$ $(1 < k \leq n)$.
- A measure of performance based on P_d, τ_{fa} and response time for detection of change.

Observations are assumed to be taken sequentially $x_1, x_2, \ldots, x_n$ with the requirement of detecting a change in distribution at a point k. Page (1955) initially proposed a form based on maximum likelihood termed the CUSUM procedure, an alternate form based on a Bayesian argument was presented by Shiryayev (1963) and Roberts (1966). Both suggest computing a sum of likelihood ratios based on the measure $P_\infty \equiv f_0(x_{1:n})$ and $P_k \equiv f_0(x_{1:k-1}) f_1(x_{k:n})$ with completely specified f_0 and f_1 such that

$$\text{Max Likelihood} \quad : \quad L_n = \max_{k<n} \left\{ \frac{P_k}{P_\infty}(x_{1:n}) \right\} \tag{C.8}$$

$$\text{Bayesian} \quad : \quad R_n = \sum_{k=1}^{n} \frac{dP_k}{dP_\infty}(x_{1:n}) \tag{C.9}$$

and asserting a change in distribution when L or R first exceed a threshold level A.

Difficulties highlighted in the previous Section can be partially overcome by the use of *scale invariant statistics* (SIS) that can identify a change point in the mean of a known class of PDFs without having complete knowledge of the initial mean. Assuming a known class of PDF f_0 one can form a statistic based upon the likelihood of there being a transition from the pre-change distribution $f_0(\eta x)$ to post-change distribution $f_1 = f_0(\eta \alpha x)$ with unknown nuisance parameter $\eta > 0$ and specified $\alpha \neq 1$. The invariance refers to the statistic being independent of the unknown parameter η.

The Gamma distribution is of particular interest, recall

$$f_{Gamma}(x) = \Gamma^{-1}(\beta)\, x^{\beta-1} e^{-x} \tag{C.10}$$

with shape parameter β and scale 1. Gordon and Pollak (1997) give a method of determining R_n with an associated ARL using a scale invariant Shiryayev-Roberts procedure. If f_0 and f_1 are known to be Gamma distributed with known β, differing in scale with both pre and post-change values unknown and relative scale change α specified, then

$$\begin{aligned} R_n &= \sum\nolimits_{k=1}^{n} \Lambda_k^n && \text{(C.11)} \\ \Lambda_k^n &= \alpha^{\beta(n-k+1)} \left[\alpha + (1-\alpha) \prod_{j=k}^{n} T_j \right]^{-n\beta} && \text{(C.12)} \\ \text{where } T_1 &= 0 \text{ and } T_n = \frac{\sum_{j=1}^{n-1} x_j}{\sum_{j=1}^{n} x_j} \text{ , } n > 1 && \text{(C.13)} \end{aligned}$$

Change point tests could be useful if the relative magnitude of change to be detected is known, but this is often not the case.

C.3 Target Presence Tests

The above can be further modified to allow the detection of not only a target's appearance but its anticipated subsequent disappearance. A variable observation interval is available. Introducing an additional hypothesis to allow two *change points* in the statistics, given the following assets:

- A realisation $x = \{x_{1:n}, n \to \infty\}$ from the radar output.
- Hypothesis H_0 indicative of no target, represented by PDF f_0

  versus hypothesis H_1 indicative of target present from time k onwards, represented by PDF $f_0(x_{1:k-1})$ and $f_1(x_{k:n})$ $(1 < k \leq n)$

  versus H_2 indicative of target present from time k until l, represented by $f_0(x_{1:k-1})$, $f_1(x_{k:l})$ and $f_0(x_{l+1:n})$ $(1 < k < l \leq n)$.
- A measure of performance based on P_d , τ_{fa} and response time for detection of change.

This test has been analysed by Tartakovsky et al. (1999) within a CUSUM framework applied to Gaussian noise.

The added complication seems unnecessary as it should be possible to achieve equivalent results by simultaneously running two change point tests to detect both increase and decrease in a parameter.

This would have the added benefit of allowing simultaneous multiple targets to be recognised such as

- Hypothesis H_3 indicative of target present from time k onwards, further target present from l onwards represented by $f_0(x_{1:k-1})$, $f_1\left(x_{k:l}\right)$ and $f_2\left(x_{l+1:n}\right)$ $(1 < k < l \leq n)$.

A further problem with this type of test is the event of not detecting the targets disappearance. This can cause the test to 'lock' in the target present state unless explicit knowledge of the expected target lifetime is available.