

**DETECTION OF GROWLERS IN SEA CLUTTER USING AN X-BAND
PULSE-DOPPLER RADAR**

By

TIMOTHY JOSEPH NOHARA, B. Eng., M. Eng. McMaster University

A Thesis

Submitted to the School of Graduate Studies

in Partial Fulfilment of the Requirements

for the Degree

Doctor of Philosophy

McMaster University

© Copyright by Timothy Joseph Nohara, February 1991

DETECTION OF GROWLERS IN SEA CLUTTER

DOCTOR OF PHILOSOPHY (1991)
(Ph.D)

MCMASTER UNIVERSITY
Hamilton, Ontario

TITLE: **Detection of Growlers in Sea Clutter Using an X-Band
Pulse-Doppler Radar**

AUTHOR: **Timothy Joseph Nohara**
B. Eng., M. Eng. McMaster University

SUPERVISOR: **Dr. Simon Haykin**

NUMBER OF PAGES: **xix, 273**

Abstract

Safe navigation of vessels through ice-infested waters in the Canadian Arctic and along the East and West Coasts of Canada is of vital importance to the exploration and development of natural resources. Current marine radar technology performs poorly in distinguishing first year ice from multiyear ice and in detecting small ice fragments (called growlers) in open water.

Previous studies by the Communications Research Laboratory, McMaster University have been successful in developing improved classification techniques for first year and multiyear ice using an X-band radar. This thesis aims to provide improved detection techniques of small ice targets in open water. With the recent go ahead of the Hibernia Project (a large oil-drilling project off of the coast of Newfoundland), the problem of reliable detection of growlers with a ship-based radar has become even more critical.

In this thesis, improved growler detection capabilities are presented. These results are based upon a detailed study of sea clutter and growler data obtained with the use of a dual-polarized, pulse-Doppler, X-band radar known as IPIX. Several contributions have resulted from the collection and analysis of real radar returns from growlers and the sea. We show that the amplitude statistics of sea clutter are K-distributed, and we quantify the poor performance that results in this non-Rayleigh clutter. Furthermore, we show that significant performance improvements can be obtained by exploiting the time-varying Doppler signatures of growlers in sea clutter. Several *medium dwell-time* coherent detectors are designed and their performances evaluated and presented in the form of receiver operating

characteristics (ROC). Finally, the comparative advantages and tradeoffs of these detectors are considered in detail.

Acknowledgement

This thesis could not have been completed without the help, guidance, patience, and labour of many others. Although there is not the space here to give my thanks and appreciation to all who deserve it, I would like to explicitly thank a few.

To my supervisor and friend, Dr. Simon Haykin, I thank you for your guidance and encouragement. The challenges that you brought forth to me each time we spoke were truly inspirational. I am indebted to you.

To the members of my Supervisory Committee, Dr. Des Taylor, Dr. Pat Yip, and Dr. Chuc Carter, thank you for your advice and your guidance. I have enjoyed many long conversations and fruitful discussions with you.

To the IPIX team, Brian Currie, Carl Krasnor, Oliver Slupecki and Dave Hamburger, I thank you for the weeks that were spent with the data collection effort in Newfoundland, and for the laughter and fun times that we shared. I also thank you for the help and guidance that you provided, and for the opportunities that were given to me for financial support.

To fellow CRL researchers Vytas Kezys, Peter Weber, and Andy Ukrainec, I thank you for the many hours of discussion and help that you have given me throughout the course of my research. I further thank Peter and Andy for the use of their image generation software that generated several of the images in this thesis. To Vytas, I especially thank you for the technical advice that you so freely offered me, time and time again, from which has resulted many of the ideas that I have developed in this thesis.

To the staff at the CRL, thank you so much for always being so helpful and cheerful. You brought a friendly personal touch to graduate student life.

To the Natural Sciences and Engineering Research Council of Canada, to the Department of Electrical Engineering at McMaster University, and to the Department of Fisheries and Oceans, I thank you for the financial support that you have provided me by way of scholarship, teaching, and contract work. Without such help, my financial burdens as a student would have been much more difficult.

To my wonderful family whom I love very much, thank you for everything. Your love, patience, support, and words of encouragement have helped me get through the last ten years of my life in ways you will never know. Mom and Dad, Theresa, Annette and Darryl, you have always been there for me, especially when I needed you most. I will never forget that.

To the love of my life, my darling Janet, where do I begin to thank you? You have been my best friend, my wife, and my crutch. Whenever I was down, you picked me up. Whenever I was tired, you helped me rest. You always encouraged me, and you were always patient with me. Thank you for all the joy and support you have brought to my life, and the impetus that you have brought to my work, especially recently, with the news of our first child.

And finally, to the source of all my strength, perseverance and good fortune, I thank you God.

Glossary

1G	1 Gaussian component
2G	2 Gaussian components
A/D	Analog to Digital
AES	Atmospheric Environment Service
AR	AutoRegressive
ARLPM	AutoRegressive Largest Pole Magnitude
AWGN	Additive White Gaussian Noise
BICE	Built In Calibration Equipment
CA-CFAR	Cell Averaging CFAR
CAGO-CFAR	Cell Averaging Greatest Of CFAR
CDF	Cumulative Distribution Function
CFAR	Constant False Alarm Rate
CNR	Clutter to Noise Ratio
CO	Complex Operation
CRL	Communications Research Laboratory
CRT	Cathode Ray Tube
DBI-GSW	Dual Binary Integration - GSW
DC	Direct Current
DMA	Direct Memory Access

FFT	Fast Fourier Transform
FM	Frequency Modulation
GCR	Growler to Clutter Ratio
GNR	Growler to Noise Ratio
GLRT	Generalized Likelihood Ratio Test
GSW	Gaussian Spectral Width
GSW/A	Gaussian Spectral Width / Amplitude
IA	Integrated Amplitude
IF	Intermediate Frequency
IPIX	Intelligent PIXel-processing (radar)
LLRT	Log Likelihood Ratio Test
LRT	Likelihood Ratio Test
ML	Maximum Likelihood
MTI	Moving Target Indicator
OS-CFAR	Ordered Statistic CFAR
PD	Probability of Detection
PDF	Probability Density Function
PFA	Probability of False Alarm
PM	Pierson-Moskowitz
PPI	Plan Position Indicator
PRF	Pulse Repetition Frequency
PRT	Pulse Repetition Time
RCS	Radar Cross Section
RF	Radio Frequency
ROC	Receiver Operating Characteristic
RSRE	Royal Signals and Radar Establishment
SAW	Surface Acoustic Wave

SNR	Signal to Noise Ratio
STC	Sensitivity Time Control
SVD	Singular Value Decomposition
TFE	Transcendental Function Evaluation
TTL	Transistor Transistor Logic
TWT	Travelling Wave Tube
UMP	Uniformly Most Powerful

Contents

Abstract	iii
Acknowledgement	v
Glossary	vii
1 The Growler in Sea Clutter Detection Problem	1
1.1 Introduction	1
1.2 Problem Description	3
1.3 Research Goals and Approach Taken	5
1.4 Overview of Results	7
1.5 Summary of Research Contributions	13
1.6 Survey of Chapters	15
2 The Detection of Targets in Clutter	18
2.1 Introduction	18
2.2 Matched Filtering in Clutter	20
2.2.1 Matched Filter Theory	21
2.2.2 Applicability of Signal-Matched Filters	24
2.3 Classical Detection and Estimation Theory	26
2.3.1 Bayesian Detectors	28

2.3.2	Neyman-Pearson Detectors	31
2.3.3	Maximum Likelihood Estimation	32
2.3.4	Composite Hypothesis Testing	35
2.3.5	Sufficient Statistic	36
2.4	Constant False Alarm Rate Detectors	38
2.4.1	Understanding CFAR Processors	38
2.4.2	CFAR Processors for Sea Clutter	43
2.5	Summary	44
3	Field Experiments and the IPIX Radar	46
3.1	Introduction	46
3.2	The IPIX Radar	46
3.2.1	System Features and Capabilities	47
3.3	Field Experiments	51
3.3.1	The Data Base	53
3.3.2	Surface Truthing	58
3.4	Calibration	60
3.5	Summary	61
4	Noncoherent Analysis of Sea Clutter and Growler Data	62
4.1	Introduction	62
4.2	The Radar Equation	63
4.2.1	Power Law	63
4.2.2	Radar Cross-Section	65
4.3	Amplitude Statistics	67
4.3.1	The Compound K-Distribution Model	68
4.3.2	Sea Clutter Empirical Amplitude Distributions	75
4.3.3	Growler in Clutter Empirical Amplitude Distributions	90

4.4	Temporal and Spatial Correlations	99
4.5	Summary	103
5	Coherent Analysis of Sea Clutter and Growler Data	105
5.1	Introduction	105
5.2	Phase Statistics of Sea Clutter	106
5.3	Time-Varying Spectra of Growlers in Sea Clutter	107
5.3.1	Dynamic Modeling of Growler Motion	115
5.4	Doppler Spectrum Estimates of Sea Clutter Returns	120
5.4.1	Nonparametric Spectrum Estimators	121
5.4.2	Pulse Pair Estimators	123
5.4.3	Autoregressive Spectrum Estimators	124
5.4.4	Gaussian Spectrum Modeling	132
5.5	Doppler Spectrum Estimates of Growler Returns	133
5.5.1	Pulse Pair Estimators	136
5.5.2	Autoregressive Spectrum Estimators	136
5.5.3	Gaussian Spectrum Modeling	139
5.6	Spectral Moment Data for Sea Clutter and Growler	141
5.7	Growler and Sea Clutter Power Ratios	149
5.8	Summary	154
6	Detector Philosophy and Design	158
6.1	Introduction	156
6.2	Noncoherent Detectors	158
6.2.1	Amplitude Detectors	159
6.2.2	CFAR Detectors	170
6.3	Coherent Detectors	175
6.3.1	Autoregressive Largest Pole Magnitude Detector	176

6.3.2	Gaussian Model Spectral Width Detector	178
6.4	Practical Considerations	180
6.4.1	IA Detector Complexity	182
6.4.2	ARLPM Detector Complexity	182
6.4.3	GSW Detector Complexity	183
6.5	Summary	184
7	Detector Performances	185
7.1	Introduction	185
7.2	Data Testbed	186
7.3	Detector Performances	188
7.3.1	Performance on a File by File Basis	188
7.3.2	Detector Performance Summary	207
7.4	Detector Performance Comparisons	216
7.5	Meeting The Goals	221
7.6	Summary	227
8	Summary and Recommendations	229
8.1	Summary of Results	229
8.2	Contributions of this Research	233
8.3	Conclusions	234
A	Weather and Wave Height Records	237
A.1	Weather Records	237
A.2	Wave Height Records	239
B	IPIX Calibrations	248
B.1	I and Q Calibration	248
B.2	Input/Output Calibration	253

C Confidence on Empirical Amplitude Statistics	259
C.1 Confidence Limits	259
D Nonlinear Least Squares Spectrum Estimation	262
D.1 Nonlinear Least Squares Solution	262
Bibliography	266

List of Tables

3.1	Radar parameters of experiments for noncoherent analysis	54
3.2	Weather and wave height data for noncoherent experiments	55
3.3	Radar parameters of experiments for coherent analysis	56
3.4	Weather and wave height data for coherent experiments	57
4.1	K-distributed Rayleigh speckle component	77
4.2	Shape parameters	88
5.1	Sea clutter spectral moments	150
5.2	Sea clutter and growler rcs data	153
7.1	Figure number and data file mapping	189

List of Figures

2.1	General detector receiver	19
2.2	Binary detection model	27
2.3	Parameter estimation model	33
2.4	A CFAR window spanning two clutter fields.	41
3.1	Radar site in Cape Bonavista	52
3.2	Waverider wave height data	59
4.1	Power versus range law for sea clutter	66
4.2	K-distributions versus scale parameter	70
4.3	K-distributions versus shape parameter	71
4.4	Normalized K-distributions versus scale parameter	72
4.5	Normalized K-distributions versus shape parameter	73
4.6	Normalized family of K-distributions	74
4.7	Underlying mean level of developed sea clutter data	78
4.8	Underlying mean level of choppy sea clutter data	80
4.9	Amplitude statistics of front-end noise	81
4.10	Sea clutter amplitude statistics (B198)	83
4.11	K normalized moments, like-polarized	84
4.12	K normalized moments, cross-polarized	86
4.13	K plus noise sixth moments	87
4.14	Amplitude images of Growler-1	91

4.15	Sea clutter amplitude statistics near Growler-1	93
4.16	Growler-1 amplitude statistics	94
4.17	Sea clutter amplitude statistics near Growler-2	95
4.18	Growler-2 amplitude statistics	96
4.19	Growler-1 and clutter cumulative densities	97
4.20	Growler-2 and clutter cumulative densities	98
4.21	Temporal/spatial correlations of sea clutter	101
4.22	Temporal/spatial correlations of sea clutter(2)	102
5.1	Phase statistics of sea clutter	108
5.2	Time/Frequency image of Growler HH	111
5.3	Time/Frequency image of Growler VV	112
5.4	Time/Frequency image of Sea Clutter HH	113
5.5	Time/Frequency image of Sea Clutter VV	114
5.6	Typical short-time clutter spectra	122
5.7	Pulse pair estimates for sea clutter	125
5.8	AR spectral estimates for sea clutter	128
5.9	AR pole plots for sea clutter	130
5.10	AR spectral width estimates versus pole magnitude	131
5.11	Gaussian spectrum model for sea clutter	134
5.12	Typical growler spectra	135
5.13	Pulse pair estimates for growler in sea clutter	137
5.14	AR spectral estimates for growler in sea clutter	138
5.15	AR pole plots for growler in sea clutter	140
5.16	Gaussian spectrum model for growler in sea clutter	142
5.17	Growler AR spectral moments versus time	143
5.18	Sea clutter spectral moments (HH)	145
5.19	Sea clutter spectral moments (VV)	146

5.20	Growler spectral moments (HH)	147
5.21	Growler spectral moments (VV)	148
5.22	Signal to clutter ratios	152
6.1	Clutter amplitude versus time	160
6.2	Growler amplitude versus time	161
6.3	Receiver operating characteristic from B98 data	163
6.4	Clutter integrated amplitude versus time	164
6.5	Growler integrated amplitude versus time	165
6.6	Integrated amplitude ROC	167
6.7	Growler and clutter temporal correlation	169
6.8	Prediction of CFAR performance	172
6.9	Prediction of CFAR performance	173
6.10	AR pole scatter plot	177
7.1	Performance curves for file B97 (HH)	190
7.2	Performance curves for file B98 (HH)	191
7.3	Performance curves for file B98 (VV)	192
7.4	Performance curves for file B99 (VV)	193
7.5	Performance curves for file B110 (HH)	194
7.6	Performance curves for file B111 (VV)	195
7.7	Performance curves for file B112 (HH)	196
7.8	Performance curves for file B112 (VV)	197
7.9	Performance curves for file B113 (HH)	198
7.10	Performance curves for file B114 (VV)	199
7.11	Performance curves for file B115 (HH)	201
7.12	Performance curves for file B115 (VV)	202
7.13	Performance curves for file B123 (HH)	203
7.14	Performance curves for file B124 (VV)	204

7.15 Performance curves for file B125 (HH)	205
7.16 Performance curves for file B125 (VV)	206
7.17 Composite IA performance curves (HH)	208
7.18 Composite IA performance curves (VV)	209
7.19 Composite ARLPM performance curves (HH)	210
7.20 Composite ARLPM performance curves (VV)	211
7.21 Composite GSW performance curves (HH)	212
7.22 Composite GSW performance curves (VV)	213
7.23 IA/ARLPM/GSW detector performances (HH)	214
7.24 IA/ARLPM/GSW detector performances (VV)	215
7.25 Amplitude and spectral width scatter plot (B98-HH)	219
7.26 GSW and GSW/A detector performance (B98-HH)	220
7.27 The GSW and DBI-GSW detector performances (HH)	225
7.28 The GSW and DBI-GSW detector performances (VV)	226
A.1 Experimental ID versus time	238
A.2 Wind speed versus time	240
A.3 Wind direction versus time	241
A.4 Visibility versus time	242
A.5 Temperature versus time	243
A.6 Pressure versus time	244
A.7 Wave height versus time	245
A.8 Wave period versus time	246
A.9 Wind speed and wave height versus time	247
B.1 Uncalibrated I/Q test data	254
B.2 Calibrated I/Q test data	255
B.3 Input/Output Characteristic	258

Chapter 1

The Growler in Sea Clutter Detection Problem

1.1 Introduction

Oil and other natural resources are of vital importance to the economic and social well being of Canadians. Particularly in times of world tension among oil-exporting nations, and in times of shortages in the world supply of these commodities, domestic production of these vital resources is essential.

Exploration and development of resources in the Canadian Arctic has provided Canadians with a valuable supply of energy sources, and is in part responsible for economic growth in Canada. Transportation of these resources from the Arctic fields to the large metropolitan centres is provided in a large part by the shipping industry. Large cargo-carrying ships must make their way safely and reliably up through the Arctic and back to the centres where their cargo is needed.

Solid ice floes and pack ice, and a large variety of icebergs of different shapes and sizes are a trademark of the Canadian Arctic. Although they are considered beautiful and majestic to tourists, they are viewed quite differently from ships which have to navigate

through them. Two distinct problems challenge vessel navigation in ice-infested waters. The first problem occurs when ice-breaking vessels navigating in full ice cover need to distinguish between first-year sea ice, which is comparatively easy to break, and multiyear sea ice and icebergs, which are significant obstacles that are hazardous to shipping. The second problem concerns the reliable detection of icebergs and iceberg fragments in open water. In heavy seas, these fragments are very difficult to detect and pose a real danger to navigating vessels.

Most vessels rely on their marine radar systems to detect these dangerous obstacles in a localized fashion, so that they can establish a safe path for navigation through ice infested waters. However, current marine radars perform poorly in both situations described above. Although airborne synthetic aperture surveillance radars can provide much of the required ice classification capability, they are not economically feasible in many applications, and furthermore, they cannot be flown continuously or in severe weather. Ship captains are therefore forced to rely on their marine radar as their primary sensor. Hence, an improved ship-based radar is of vital importance to safe navigation in ice-infested waters!

One reason for the poor performance of conventional marine radars is that they are capable of measuring only the amplitude of the return echoes, in order to make a decision. In the early 1980's, The Communications Research Laboratory (CRL), McMaster University, embarked on a program to investigate potential improvements in marine radar that could overcome this shortcoming. Experimental programs involving the CRL in 1983 [1] and 1984 [2] using a land-based site in the Canadian Arctic clearly showed that the use of dual linear polarization significantly improved the ability to detect and classify multiyear ice and icebergs. Using an X-band radar, a 6 to 7 dB improvement in target-to-clutter ratio was obtained using the cross-polarized returns. The icebergs and multiyear ice were the desired targets and first year ice was considered to be the undesired target. This improvement occurs as a result of the different dielectric properties of first year and multiyear ice. First year ice has a high salt content and thus permits little penetration of the radar signal and

therefore little depolarization of the reflected signal. Multiyear ice and icebergs, on the other hand, have a low salt content and thus are low-loss dielectrics permitting the radar wave to penetrate the ice. Discontinuities within the ice cause reflection and depolarization of the radar wave, scattering some of the energy into the cross-polarization orientation.

In a 1986 experiment [3], a dual-polarized radar was used to investigate the value of polarization in the detection of icebergs in open water. Although it was found that the sea depolarized the radar returns to a similar degree as did the iceberg targets, indicating that no significant improvement in signal-to-clutter ratio could be expected using the cross-polarized returns, it was noted that the like- and cross-polarized returns scintillated independently, suggesting possible improvements in the joint detection probability using both polarization channels.

The experiments described above provided important results that are useful for improving marine radar capabilities in full ice cover and open water conditions. Although for the open water problem, improvements obtainable with dual-polarization are not as promising as in the case of full ice cover, both experiments were useful as they provided invaluable experience in designing and carrying out radar experiments for research purposes. In the next section, we outline specifically the problem addressed by this thesis.

1.2 Problem Description

The previous section outlined the two problems that challenge ship navigation in ice-infested waters. The problem of navigating in full ice cover was dealt with thoroughly by the 1984 experiments conducted by the CRL, and the results are described in depth in [4]. In this thesis, we are concerned with the latter problem; that is, the detection of ice targets in open water. More specifically, we are interested in detecting small ice fragments, as opposed to larger fragments and icebergs. These small ice fragments called *growlers* can weigh up to about 100 metric tonnes and if struck by a vessel, can cause considerable damage. They are

particularly difficult to detect with ship-based radars because they protrude only about a metre or so above the waterline. In calm seas, the radar echo from a growler has to compete with the echoes from the waves in the same resolution cell as the growler. These wave echoes are called *sea clutter*. Since the growler can be considered a point target for typical marine radar resolution cell sizes, the distributed sea clutter masks the growler return, making the growler difficult to detect. In rougher seas, the problem is even more severe, as the waves visually cover or shadow the growlers a large percentage of the time.

During the spring thaw, large icebergs and ice floes migrate from the Canadian Arctic southwards along the east coast of Canada. They reach as far as Newfoundland in late Spring, and are very hazardous to shipping in the region. During the melt, the large icebergs break into fragments, producing growlers and bergy bits. (Bergy bits are ice fragments that are larger than growlers and thus are more easily detected with radar). Thus, in the Atlantic region off Canada's East Coast, there is particular interest in improving the detectability of these growlers with radar.

Recently, the *Hibernia* project was given the go ahead by the Canadian federal government. This project involves developing new oil reserves off of Canada's East Coast by a consortium of Canadian oil companies. Oil drilling platforms will be built about 200 km offshore for the purpose of bringing oil reserves from the sea bottom to the surface. This oil will then be shipped to shore by several large tankers. Due to the damage that could result from a collision with an oil tanker and growlers, the need to improve growler detection with radar is more pressing than ever. In fact, the oil consortium responsible for the *Hibernia* project is presently seeking the development of new marine radars capable of detecting these growlers with high probability! In the next section, the goals of this research will be described, along with the approach that was taken in order to realize improved growler detection using a ship-based radar.

1.3 Research Goals and Approach Taken

The ultimate goal of this research is to develop techniques which are capable of providing improved detection performance of growlers in the sea, using a ship-based radar. More specifically, we wish to aim for the following goals:

1. To provide real measurements on both sea clutter and growler backscatter in order to enhance data reported in the literature.
2. To provide statistical descriptions capable of describing the data and useful for understanding the underlying processes.
3. To investigate a variety of potential solutions each with its own merit, performance, and associated costs, in order to provide a realistic comparison.
4. To provide a recommended solution that is modular in that future enhancements are easily accommodated and can be built on top of the solution.
5. To indicate where future research should be focussed.

In order to achieve the above goals, it was imperative to develop a well defined research plan. Clearly, an integral part of that plan would include field trials in order to obtain real measurements of sea clutter and growlers. The experiences obtained from the full ice cover experiments provided a lot of guidance in designing this research program. The experiments described above were performed with noncoherent radars, that is, radars capable of measuring only the amplitude of the return signal. Clearly, this is not sufficient for developing improved detection capabilities.

Along with polarization diversity, coherence is essential. A coherent radar is capable of measuring the phase changes in the return signal, as well as the amplitude. Visual observations of the ice/water interaction indicate that motion characteristics such as splashing, bobbing and translation may distinguish floating ice from its surroundings. This suggests

that Doppler processing, which computes target velocity by operating on the phase of the returned signal, could enhance the probability of detecting such targets. Although the addition of coherence seems obvious for this problem, no quantitative measurements of detection improvement using a coherent radar are known to have been published to date! There are however reports of the use of coherent radar in iceberg detection applications [5].

The ice research described earlier was performed using commercially available radar systems which had been adapted for research use. These modified radars were not ideal for conducting research experiments, as they rarely offered a full set of radar features that could easily be varied at will. For example, changing the polarization of the transmitted signal often required mechanical alterations, making the system unsuitable for making polarization scattering matrix measurements. Furthermore, most of these systems had CRT's as the output device, making data collection and storage very difficult. Issues such as Sensitivity Time Control (STC) were often manually adjusted by the operator, making accurate measurements very difficult to obtain. Commercial radars are designed with a particular end in mind, and it was determined that they are unsuitable for research use as they limit the range of experiments that can be performed, as well as the quality of the measurements.

The previous experiments conducted by the CRL indicated the need for an instrumentation-quality radar system that is designed specifically for research use. The requirements for such a system were formulated by considering the radar parameters to be investigated, as well as the ease with which those parameters needed to be varied. Furthermore, particular emphasis was given to the data collection effort, designed to facilitate further analysis of the data. With funding from the Natural Sciences and Engineering Research Council of Canada, development of the system began in 1984 and a prototype was tested in 1986. The system has matured since that time and is known as the IPIX (for Intelligent PIXEL-processing) radar. Details of IPIX are found in Chapter 3.

With the development of the IPIX radar, the research plan unfolded. It was decided that a land-based site was most suitable for carrying out research experiments. The

requirements for the site were quite stringent. The site would have to provide a large sector of sea coverage, and would have to be on a cliff-top that simulated the height the radar would be mounted on a large vessel. Furthermore, the sea would have to drop in depth very quickly to simulate deep sea conditions. A wide variety of weather conditions were required, as well as a wide variety of sea states. Finally, and most importantly, the site would have to be along ice migration paths, so that the radar could view a large variety of ice targets. Such a site was located at Cape Bonavista, Newfoundland.

The research plan consisted of making several field trips in order to collect data from various ice targets, and from the sea itself. These measurements would then be examined to provide the basis for improved detector designs. Between May 1988 and June 1989, four such field trips were carried out providing us with a large library of data. Following the experimental effort, the data was analyzed at the CRL in order to determine which radar features are most useful for detecting growlers in sea clutter. It is the results of this study that are reported here. An overview of these results is given in the next section, followed by a summary of the research contributions. The final section of this chapter describes the organization of the remainder of this thesis, and indicates where details of the various results can be found.

1.4 Overview of Results

The analysis of sea clutter and growler data can be divided into two areas, noncoherent analysis and coherent analysis. In order to give some perspective on the difficulty of the problem and on the kind of performance that can be expected using conventional marine radars, the noncoherent results will be reviewed here first.

It was discovered that the compound K-distribution developed by Ward [6] modelled the sea clutter amplitude statistics very well. Sea clutter data was collected for several pulse widths, look-directions, and sea conditions. Both linear horizontal and vertical polarizations

were used on transmit and receive, allowing for measurements of the full scattering matrix. Thus, like-polarized (HH and VV) and cross-polarized (HV and VH) measurements were made. In all cases, the K-distribution provided an accurate description of the amplitude statistics of like-polarized data. The cross-polarized amplitude statistics did not fit the K-distribution very well though. In the cross-polarized channels, the system receiver noise is much more predominant in the data since the cross-polarized echoes are typically 10 dB below the like-polarized echoes. Thus, by using a K-plus-noise distribution for the data, a much better fit was achieved!

The Rayleigh distribution is a special case of the K-distribution, and has often been used as the basis for detector design in clutter environments. The performance of these detectors in sea clutter has been much worse than that predicted by the Rayleigh model. That is, actual false alarm rates are much higher than predicted. This is due to the fact that a larger number of large clutter returns (called spikes) are received than predicted by the Rayleigh model. The term *spiky* is often used to describe clutter returns that behave in this manner. The K-distribution is capable of describing the *spikiness* of clutter with its shape parameter. The shape parameter and scale parameter are the two parameters that define the family of probability density functions known as K-distributions.

Fitting empirical amplitude statistics with a particular family of probability distributions is rather *ad hoc* in that often more than one family of distributions will describe the data. The particular significance of the K-distribution to our problem is due to the compound nature of the distribution. It is shown in Chapter 4 that the K-distribution can be broken down into two components. The first component represented by a root-Gamma distribution describes the *underlying mean level* of the data and has with it its own correlation time. The second component described by a Rayleigh distribution models the local speckle, and again has its own correlation time associated with it. Our measurements of sea clutter data support this compound model very closely. From a given patch on the sea surface, an overall *tilt* can be used to describe the mean sea height distribution and

this tilt is responsible for the mean level of backscatter. The tilt is partially attributed to the underlying sea swell and thus varies quite slowly (on the order of a second). On the otherhand, small capillary waves are continuously influenced by local winds and thus induce a Rayleigh component that has a much shorter correlation time (about 10 ms).

Using the compound K-distribution to describe the amplitude statistics of sea clutter is ideal because it models the nature of the sea clutter returns accurately. This compound modeling allows realistic performance prediction to be made when considering the improvements of integration and CFAR type processing. These aspects will be dealt with in more detail later.

Amplitude statistics for growlers in sea clutter were also studied and compared with the statistics from nearby sea clutter resolution cells. For the challenging growlers examined, the mean amplitude was very similar to that of the nearby clutter, making detection based solely upon the size of the reflected echo impractical. Also, the cross-polarized channels were examined to see whether there was more separation between growlers and sea clutter than in the like-polarized channels. It was determined, however, that when using only long-term amplitude statistics, the cross-polarized channels offered no advantage. With the poor performance resulting from the use of amplitude by itself, coherent data were examined in order to see what improvements could be gained.

During our field trials, we discovered a time-varying Doppler signature characteristic of returns from growlers in the sea. This signature was first observed while watching the inphase or quadrature signal with the radar fixed along a particular azimuth containing a growler. The inphase or quadrature channel was displayed on an oscilloscope and at the range position where the growler was located, the signal appeared to *wink* at various rates. The *winking* would occur slowly at first, and then speed up gradually, only to reach a maximum rate and slow down again. This pattern would repeat itself at intervals on the order of several seconds. We have called this discovery the *winking phenomenon* in the literature [7, 8]. The explanation of this discovery lies with particle and fluid dynamics. For

growlers that are small compared to the wavelength of the underlying sea swell (typically 100–150 m), the growler behaves like water particles in motion. As a wave crest passes through the growler, the growler rises and is accelerated forward. As the crest falls, the growler decelerates. The wave troughs cause the reverse effect. Therefore the growler actually has a cyclical motion whose period is controlled by the underlying wave period (5–12 seconds typically).

As alluded to earlier, these visual observations of growlers in the sea indicated to us that Doppler analysis could provide significant improvements in detection performance, because of its ability to characterize motion. Indeed, this hunch was correct! Analysis of this data using Doppler techniques confirmed the presence of the *winking phenomenon* that we had observed in the field.

Traditionally, marine radars operate in a scanning mode to provide 360 degree coverage, with a scanning rate of about 30 revolutions per minute. With typical pulse repetition frequencies (PRF) (1–2 kHz) and azimuthal beamwidths (1–2 degrees), a dwell time of 5–10 ms on a particular radar resolution cell might be typical. This translates to about 10–20 pulses that may be integrated in order to make a detection. We call scanning approaches like these *short dwell-time* approaches. Recalling that the *winking phenomenon* is physically related to the underlying wave mechanics, detection techniques which are designed to detect this phenomenon will require several seconds of data from each resolution cell, just to observe a complete cycle of the time-varying Doppler signature. These methods we call *long dwell-time* approaches. Clearly, data storage, computational complexity, and the minimum update time are going to be critical issues for any *long dwell-time* detector.

As a result of the discovery of the *winking phenomenon*, a large research effort has just begun at the CRL in order to develop *long dwell-time* techniques to exploit these time-varying Doppler signatures.

The research approach that is taken in this thesis differs from both the *short dwell-time* and *long dwell-time* approaches described above. Although the study described herein

is also concerned with exploiting the *winking phenomenon*, it will take a middle ground and look at *medium dwell-time* techniques. By *medium dwell-times*, we mean techniques that will base a decision on the coherent integration of data over a period of a fraction of a second. It is felt that an approach like this will serve many purposes:

1. It will provide faster detections which translate to faster updates.
2. It will offer improved growler detection performance over *short dwell-time* approaches.
3. It will be able to provide meaningful results based on the current data-base.
4. It will require less computation and less memory than *long dwell-time* approaches which translates to cost savings.
5. It will indicate the kind of performance improvements that can be expected from *long dwell-time* methods.
6. It will be modular in design in that it will provide the necessary features needed by a *long dwell-time* system.

In order for *medium dwell-time* techniques to be successful, coherent integration of sea clutter data and growler data over a fraction of a second must be noticeably different. Indeed, spectrum analysis has shown that there are differences between clutter and growler spectra. The clutter spectra can be described as being Gaussian-like in shape, having a mean velocity shift that can range between 0 and 5 knots. At X-band, this translates to a frequency shift between 0 and 150 Hz. Furthermore, the velocity spread for sea clutter, or frequency spectral width is typically about 20–30 Hz. The growler too has a spectrum that is Gaussian-like in shape, and although it occupies roughly the same range of mean frequency shifts as the clutter, it has a predominantly smaller spectral width, typically about 6 Hz.

In order to quantify the usefulness of spectral width in classifying sea clutter and growler returns, spectral width statistics were calculated. For sea clutter, a Gaussian model

was fitted to the estimated sea clutter spectra, in order to estimate the spectral width. A two-component Gaussian model (one Gaussian component for the clutter and another Gaussian component for the growler) was fitted to the estimated growler-in-clutter spectra, in order to determine growler spectral width statistics. As expected, a considerable separation resulted, indicating the usefulness of spectral width as a discriminant to distinguish between clutter resolution cells, and growler-in-clutter cells.

The GSW (Gaussian Spectral Width) detector was designed and its performance evaluated. The GSW detector performs the following operations. First, a spectrum estimate is computed using 0.5 seconds of data for each estimate. Then, a single Gaussian model and a double Gaussian model are fitted to the data. Each Gaussian has three free parameters that describe it, namely the center frequency, the amplitude, and the spectral width. These parameters are optimized with respect to the data in a minimum Least Squares Error sense. The model that provides the *best fit* to the data is chosen. If a single Gaussian model is chosen, then its spectral width is used as the sufficient statistic. Otherwise, for the double Gaussian model, the smaller spectral width is used. The sufficient statistic is then compared to a threshold. If the threshold is exceeded, the data set is said to be from sea clutter; otherwise the data set is said to be from a growler.

This detector is evaluated against a large testbed of real sea clutter and growler data for several different thresholds. Thus, performance curves are generated showing the probability of detection versus the probability of false alarm for this detector.

A second detector, called the ARLPM (AutoRegressive Largest Pole Magnitude) detector was also developed for comparison, and evaluated with the same database. This detector also makes use of the spectral differences between clutter and growler returns and works as follows. For each data sequence of 0.5 seconds in length, a sixth order Autoregressive model is evaluated and the six poles of the model are determined as described in [9]. Each pole is contained within the unit circle. If a growler is present in the data, one of the poles tends to cluster very close to the unit circle since the growler spectra is narrowband.

If no growler is present, then all the poles tend to lie further inside the unit circle. Therefore the largest pole magnitude is calculated and compared to a threshold. If this magnitude exceeds the threshold, a growler is declared; otherwise clutter is declared.

Finally, in order to quantify the performance improvements achieved by the two coherent detectors described above, a simple noncoherent detector, called the IA (Integrated Amplitude) detector, was also evaluated. This detector operates on the same 0.5 second data sequences by noncoherently integrating the samples and comparing the sum to a threshold. Again, if the threshold is exceeded, a growler is declared; otherwise, clutter is declared.

The performances of the three detectors are now described. The ARLPM and GSW detectors are shown to be far superior to the IA detector. With some careful analysis, the GSW detector is shown to be far more robust to changes in the growler-to-clutter ratio (GCR) than the ARLPM detector. Therefore, it is the preferred detector for growler detection. Since the IA detector is representative of the kind of performance one can expect from current marine radars, the performance improvement indicated by the GSW detector is very significant. For example, with a probability of false alarm (PFA) of 5 %, the IA detector offers a probability of detection (PD) of only 60 %, while the GSW detector offers a PD of 96 %. Furthermore, it is shown that if five consecutive GSW detection decisions are made, and a three out of five rule is used to establish the validity of growler detections, then the performance improves even more. For this case, a PD of 98 % is realized for a PFA of only 0.1 %, which is a remarkable improvement in performance!

1.5 Summary of Research Contributions

In the previous section, a brief description of the results of our research was provided. The contributions of this research are summarized as follows:

1. Real, coherent, dual-polarized, X-band sea clutter and growler-in-sea clutter radar backscatter data have been collected and organized into a database. This database

has been described in great detail in the thesis. The associated radar parameters, along with extensive surface truthing data have been provided. Furthermore, descriptive data such as GCR, growler radar cross section (RCS), and sea clutter σ_o data are also given. Such a complete description of data is rarely found in the literature, and will provided other researchers with valuable measurements for further analysis. Furthermore, this database represents very low grazing angle X-band sea clutter measurements along with growler measurements, that are very scarce in the literature.

2. The applicability of the K-distribution for modeling the amplitude statistics of very low grazing angle sea clutter has been verified. Furthermore, its use in performance prediction has been shown.
3. The discovery of the *winking phenomenon* associated with growler radar returns has been described. This discovery is essential to the understanding of growler dynamics and has been exploited in the development of improved growler detectors.
4. The time-varying Doppler spectra of sea clutter and growler data have been studied and spectral width is identified as a sufficient detection statistic.
5. Two detection algorithms have been developed that offer substantial improvements over conventional noncoherent techniques. These are the AutoRegressive Largest Pole Magnitude detector and the Gaussian Spectral Width detector. In particular, the GSW detector offers performance, robustness, and modularity giving the system designer maximum freedom. Furthermore, the timeliness and scope of this study suit well the needs of the recent Canadian effort in the development of the Hibernia oil fields.
6. Calibration, data validation, data handling, and analytical software has been developed for the radar system computer. The software is designed to operate on the data, in place, as they are collected. This library will be very helpful to graduate researchers

who are working on the IPIX project.

7. Several conference papers and journal papers have been published that report on the IPIX radar, the data collection effort, and results obtained [8, 10, 11, 12, 7, 13, 14, 15].

These results should have a significant impact on the commercial development of improved marine radars, capable of detecting growlers in sea clutter. Furthermore, they provide a sound basis for future studies and experimental efforts.

1.6 Survey of Chapters

In this chapter, the growler detection problem was described along with the experimental plan that was carried out in order to increase our understanding of the problem. A summary of the results of this research was also presented including the detection performance improvements achieved. Details of the results described in this chapter can be found elsewhere in the thesis. Although each chapter is an important part of this thesis as a whole, some effort has been made to keep each chapter self-contained. Hopefully, this will help the reader who is interested in just a particular section. We will now proceed to describe the contents of the remaining chapters.

Chapter 2 looks at the theory behind the detection of targets in clutter. We begin by proposing a model of a general detector receiver. Then, the chapter focusses on several detection and estimation theories to provide the component stages of the general receiver model. Matched filtering is examined and its applicability in cases where the clutter is non-additive is discussed. Then, classical Bayesian detection theory is reviewed. In radar problems, it is often difficult to design suitable cost functions as are required by Bayesian detectors. Therefore, Neyman-Pearson detectors are discussed. The chapter concludes by looking at several *constant false alarm rate* (CFAR) detection techniques which are suitable in nonstationary clutter environments.

In Chapter 3, the IPIX radar is discussed in depth. The system features and

capabilities are described in detail. Then, the field trials that were performed and the ground truthing that was available are described. Finally, we look at issues such as data calibration and translation of digitized data to real received power.

Details of the noncoherent analysis of sea clutter and growler data are given in Chapter 4. First, we look at the forms of the radar equation that are suitable to this problem. Then we present some measurements indicating the range power law of distributed sea clutter. Radar cross-section calculations are then described. Following this, the amplitude statistics of sea clutter and growlers are given, and the K-distribution and its description of these statistics are described. The chapter finishes by examining the temporal and spatial correlations of sea clutter.

In Chapter 5, the coherent analysis of sea clutter and growler data is given. The chapter begins by discussing the Doppler effect and spectrum estimation. Then the phase statistics of sea clutter are examined and shown to be uniformly distributed. Next, the time-varying nature of sea clutter and growler spectra are studied and parameterized using several different spectral models. Estimators are derived for the spectral moments of clutter and growler signals, and the differences in their spectral width statistics are examined in detail. Finally, growler to clutter ratios are evaluated along with radar cross sections for clutter and growlers.

Detector philosophies and designs are described in Chapter 6. Three different detectors are described. These include an integrated-amplitude detector, an autoregressive detector which is based on the largest pole magnitude, and a Gaussian spectral width detector. Other detector designs such as scan-to-scan and traditional CFAR detectors are also discussed.

In Chapter 7, the performance evaluations of the three detectors described in Chapter 6 are given, the testbed of data on which these evaluations are based is described, and finally practical considerations for each of the detectors are presented.

Finally, Chapter 8 concludes by summarizing the results and contributions of this

research, indicating future directions that can be taken.

Chapter 2

The Detection of Targets in Clutter

2.1 Introduction

As indicated in the first chapter, we are interested in improving detectors capable of detecting growlers in a sea clutter environment. Before proceeding then, a review of basic detection and estimation theory is in order. We begin by discussing *Matched Filtering*, which is basic to most detection systems today. Then, we will take a more general look at classical detection and estimation theory by considering *Bayesian* detection, *Neyman-Pearson* detection, and *Maximum Likelihood* detection theories. Following these, we consider *CFAR* detectors which are used when the received signal processes are nonstationary in space or time. These theories will be examined in the context of detection of targets in clutter, for use in later chapters.

The goal of every detection system is to make a decision out of a countable set of possible decisions, in some optimal fashion. In order to satisfy the optimality criterion for a given system, the minimization of a cost function usually results, the solution of which provides the basis for the detector design. The cost functions that describe the




Figure 2.1: General detector receiver. The receiving antenna feeds the front-end filter. The feature extractor operates on the output of the front-end filter, and supplies the input to the threshold detector.

decision criteria vary considerably from system to system. The function may be as simple as describing the output signal to noise ratio of the signal processor, where only signal energies and noise powers are involved, or it may be probabilistic in nature. For example, it may describe the probability of making an error in the decision making process, based on some knowledge of the statistics of the signals and noise received. It may also include a priori information about the signal and noise processes.

In any case, binary detection, which involves only two decision classes (H_0 : noise present, and H_1 : signal and noise present), results with a single test-statistic (the sufficient statistic) that is compared to a threshold in order to decide between the two classes. This is the type of detection system we are interested in.

Detection systems can be described as consisting of several stages in series, each performing a particular function on the output signal of the previous stage, and the final stage producing the decision. Consider the general receiver described in Figure 2.1. The first stage is a *front-end filter* that is tailored to the signals received by the antenna. This

stage is followed by a *feature-extractor* stage which extracts features of the received data upon which a decision can be based. Finally, that last stage is a *threshold device* that selects between H_0 and H_1 .

In the next section, we begin by examining the applicability of *Matched-Filter* theory to our detection problem. In particular, we consider the applicability of *matched filtering* in the design of the *front-end filter*.

2.2 Matched Filtering in Clutter

Often, the stage at the front-end of the receiver is common to a variety of detector designs. We are interested in designing a front-end filter that will operate on the data, and whose output will be suitable for a variety of detectors that will follow.

An X-band radar is employed that transmits pulses that are separated by the pulse repetition period, PRT. The receiver is coherent in that it measures both the amplitude and phase of the return signal. Consider the case where a pulse with duration τ seconds is used at some pulse repetition frequency (PRF) (Hz). The pulse is transmitted along a particular azimuth and the returned pulse is received. The sampling frequency is f_s (Hz). The pulse has a range ambiguity of $c\tau/2$ metres where c is the speed of light (3×10^8 m/s) and it is assumed that in general, the sampling frequency is such that there are N samples per pulse length τ .

If a growler is present at a particular range (considered a point target), it will reflect some of the pulse's energy back to the receiver, and the reflected pulse length will also be τ . For a distributed target such as sea clutter, the reflected pulse length will in general be larger than τ , due to the extended nature of clutter scatterers. We consider this in more depth in the section on homomorphic signal models.

In all cases, additive, zero-mean, white Gaussian noise (AWGN) is assumed to be present at the receiver front-end, and to have a bandwidth far greater than that of either

the target or clutter processes. We are interested in designing a filter that will combine the N samples per pulse in such a way that the following conditions will be satisfied:

- a) the filter can be used as a pre-processor for many different detector schemes.
- b) the filter offers a growler/clutter integration gain greater than unity.
- c) the filter still applies to cases where more than a single pulse is transmitted, allowing for Doppler processing.

A *matched filter*, matched to the transmitted signal is the classical solution that is applicable for pulsed-carrier target signals in additive noise. We discuss its applicability to our problem.

2.2.1 Matched Filter Theory

Classical filters that provide optimal signal visibility in a noise background date back to the pioneering work of North [16] in 1960, although it was originally published as a military-classified report in 1943. This work still forms the basis of matched-filter theory today, although it has been generalized by Turin [17] and Middleton [18] among others.

North was interested in determining what transfer characteristic between the antenna and the detector would provide a maximum output, when a pulsed-carrier input signal was received in additive white noise. Excluding a delay term, North showed that this filter had a transfer characteristic that was the conjugate of the spectrum of the echo at the antenna; hence, the *conjugate* or *matched* filter. Since ranging was also required along with detection, the receiver required range-gating hardware in order to isolate the return from a particular range bin. The matched filter was also shown to be equivalent to correlating the input signal with an image of itself. For a rectangular pulse of length τ , the output signal is triangular with length 2τ and peak output occurring at the center. Finally, this filter has the property of maximizing the output signal energy to rms noise power (SNR); and this maximum value is independent of the signal shape!



The case of nonwhite additive noise leads to a very similar structure [17]. For an input signal with a spectrum $S(j2\pi f)$ and a noise power spectral density $N(j2\pi f)$, the resulting optimum filter in the maximum output SNR sense is given by

$$\frac{kS^*(j2\pi f)e^{-j2\pi f\tau}}{|N(j2\pi f)|^2} \quad (2.1)$$

where \star denotes complex conjugation. This can easily be seen to be the cascade of a noise whitening filter followed by a filter matched to the *whitened* signal. Ideally, one would like to get a large impulse out of the optimal filter when a target signal is present. We have seen instead that the matched filter leads to an output which is not an impulse. Clearly, one could obtain an impulse simply by using an inverse filter, rather than the conjugate filter! Closer examination reveals the apparent paradox. Any physical signal has a spectrum that approaches zero for frequencies that approach infinity. Since the noise spectrum is assumed to be much wider in frequency than the signal spectrum, an inverse filter would have the effect of applying very large weights to regions of noise outside of the signal bandwidth. Therefore the matched filter is a compromise that weighs the signal more strongly in bands where it is large compared to the noise. Strictly speaking, the filter described in Equation 2.1 is matched to the input signal. This filter is actually made up of a bank of filters, each one tuned (matched) to a different expected Doppler frequency. In practice, a single matched filter is employed that is matched to the transmitted signal, rather than the received signal. This filter is then followed by a Doppler processor. This combination has approximately the same effect as the bank of frequency-tuned matched filters.

The matched filter descriptions given above were the result of considering mean-square or maximum SNR criteria. These criteria are useful since they only require knowledge of the second order statistics of the data. Since performance ratings of detection systems are usually characterized by parameters such as probability of detection, PD, and probability of false alarm, PFA, it is preferable to use criteria directly related to these performance parameters. (These criteria will be considered in more depth in Section 2.3.1 and Section 2.3.2).

We are now forced to call upon statistical detection and estimation theory. These theories lead to the formulation of a likelihood ratio test that determines which of two hypotheses is true, H_0 or H_1 :

$$\left. \begin{array}{l} \text{accept } H_1 \text{ if } \frac{p_1(x)}{p_0(x)} \geq \lambda \\ \text{accept } H_0 \text{ if } \frac{p_1(x)}{p_0(x)} < \lambda \end{array} \right\}. \quad (2.2)$$

The probability density functions under H_0 and H_1 are $p_0(x)$ and $p_1(x)$ respectively, and λ is the detection threshold. It is shown [17] that when the additive noise is white and has a probability distribution function that is Gaussian, the same North matched filter arises. When the receiver is coherent, the matched filter is followed by a sampler and a threshold comparator. If only the envelope is available, then under the assumption that the carrier phase is uniformly distributed, the North filter still arises, only in this case it is followed by an envelope detector and then the sampler!

The fact that the same matched filter arises when several different detection criteria are used lends justification to its broad use in most detection systems today [19, 20]. Therefore, we will now consider how to digitally implement this matched filter at baseband, before we proceed further.

In the IPLX radar, a pulsed-carrier transmit signal is employed, but on receive, the signal is demodulated down to baseband with a coherent carrier, resulting with inphase and quadrature baseband signals, I and Q . These waveforms are then sampled at the sampling frequency, f_s and stored as 8-bit words. Therefore any filtering operation must be performed on these digital signals. As described above, the basic filter operation is equivalent to correlating the received signal with a replica of itself. The correlation operation is equivalent to a multiplication of the received signal with its image, and then integrating the result over the duration of the signal. The transmit signal is given by

$$s(t) = p(t)\sin(2\pi f_c t) \quad (2.3)$$

and the received signal given by

$$r(t) = p(t)\sin(2\pi(f_d + f_c)t + \theta), \quad (2.4)$$

where the phase θ is assumed to be zero without loss of generality, and $p(t)$ is the pulse envelope function. The demodulation process is usually performed in quadrature in order to preserve the sign of the doppler frequency f_d . That is, $r(t)$ is multiplied by both $\cos 2\pi f_c t$ and $\sin 2\pi f_c t$ and the products are each lowpass filtered, resulting in the I and Q components, respectively. Clearly, all that is required to complete the matched filtering operation (in range) is to integrate the I and Q components over the length of the pulse $p(t)$. As mentioned earlier, I and Q are both sampled such that N samples are obtained per pulse length τ . The integration operation thus becomes a summation. The matched filtering is then described as

$$I_m = \sum_{i=1}^N I_i \tag{2.5}$$

$$Q_m = \sum_{i=1}^N Q_i$$

assuming that the pulse envelope function $p(t)$ is rectangular. It should be emphasized that this matched filter operation simply takes care of the necessary range matching. A bank of Doppler filters for example would then follow, matched to the set of expected Doppler frequencies in the received signal!

In this section, we have examined classical matched filter theory as it applies to signals in additive noise. Our problem concerns a signal in additive noise, but there is also an additive clutter component. Therefore, in order to design a proper matched filter, the clutter spectrum needs to be accurately modelled. In our problem, this is quite difficult because clutter characteristics vary as a function of time. Instead, we assess the benefits of simply using a filter matched to the transmitted waveform rather than the received waveform.

2.2.2 Applicability of Signal-Matched Filters

Recall that our main objective is to design a front-end filter that will satisfy the requirements described earlier. It is not necessary to provide optimal separation between the clutter and

target processes at this early stage.

Although from a theoretical viewpoint, filters matched to the transmit signal may not provide optimal reception of a desired signal in the presence of additive clutter, they are still used in most communication and detection systems today. The reason is that they seem to perform well, even in a cluttered environment. Turin [17] states the following:

No matter what formalism is used to view a given communication or detection situation, Gaussian noise statistics lead usually to some form of correlation or matched filtering as a part of the set of operations that will perform the desired function most efficiently. This appears to be true even when in addition to the noise there are other perturbing factors present, such as randomly varying multipath, uncertainties in signal delay or Doppler shift, Doppler or delay smearing, or unwanted clutter.

Intuitively, applying a filter matched to the transmitted signal should provide some improvement in SNR since the clutter will only be integrated in the region where there is signal energy. Furthermore, if the clutter has considerable spatial variance, integration over the pulse length may help reduce the average contribution from the clutter.

Westerfield *et al.* take a more quantitative look at the processing gains provided by matched filters against clutter [21]. They consider clutter to be composed of several overlapping echoes arising from the extended scatterers. The scatterers are considered to be randomly distributed in range and velocity and the distribution of the velocities is considered to be Gaussian. The signal power spectrum is also assumed to be Gaussian distributed. By making use of Woodward's ambiguity function, which is simply the squared magnitude of the cross-correlation between $s(t)$ and a delayed and Doppler shifted version of itself [22], the processing gain of the simple matched filter against clutter is evaluated. It is shown that large processing gains can be achieved when the ratio of the transmitted signal bandwidth to the spread of Doppler shifts in the clutter is either large or small. These results are intuitively satisfying. For large signal bandwidths, the improved range resolution

is responsible for the processing gain. On the other hand, for very long transmitted pulses, the resolution in the frequency domain is improved, narrowing the ambiguity function along the frequency axis.

In conclusion, applying the traditional matched filter to the transmitted pulse is a practical way of achieving the goals described earlier. Even though this matched filter is derived for additive signal and noise processes, it appears to offer some advantage even in cases with additive clutter. With the *front-end filter* designed, the next step is to design the *feature-extractor* and *threshold device* stages. For these stages, we call on classical decision and estimation theories.

2.3 Classical Detection and Estimation Theory

As described in the introduction, the detection of growlers in sea clutter is a binary hypothesis testing problem. The hypothesis H_0 is the case when no target (growler) is present, and the hypothesis H_1 is the case when a target is present. The problem we face is simply to design a ‘rule’ which can operate on the received measurement data in order to decide between H_0 and H_1 . The treatment given here will follow very closely that given by Van Trees in [23].

Any binary detection problem can be described very simply by the model given in Figure 2.2. The ‘Source’ represents the statistical model that generates the H_0 and H_1 events. In our case, the ‘Source’ represents the *a priori* probability that a growler will be located in the radar resolution cell under examination. The ‘Probabilistic transition mechanism’ describes the manner in which the two hypotheses are mapped into the ‘Observation space’. The ‘Observation space’ represents the measurements or observables that are available to the detector. In our problem, the ‘Observation space’ represents the space spanned by the radar features (amplitude, phase, and polarization) and the ‘Probabilistic transition mechanism’ describes the statistical distributions which govern these features.




Figure 2.2: Binary detection model

The ‘Decision rule’ refers to the receiver/detector that is used to operate on the measured data and the ‘Decision’ is the output of the receiver and indicates whether we have decided in favour of H_0 or H_1 .

Clearly, any ‘Decision rule’ is dependant on the stages that precede it. A good description of the ‘Source’ and ‘Probabilistic transition mechanism’ will allow for the optimal receiver design. As we shall see though, a complete description of all of these stages is not always available; this fact will clearly influence the kinds of decision rules that will work best. Since our measurements consist of a set of discrete measurements or samples, the observation space is finite-dimensional. This class of problems is referred to as *classical* because it has been studied for many decades by statisticians.

Let \mathbf{r} denote a vector in the N -dimensional observation space:

$$\mathbf{r} = \begin{bmatrix} r_1 \\ r_2 \\ \vdots \\ r_N \end{bmatrix} \quad (2.6)$$

and let \mathbf{R} be an observed realization of this random vector. Furthermore, let the two conditional probability densities $p_{\mathbf{r}|H_0}(\mathbf{R}|H_0)$ and $p_{\mathbf{r}|H_1}(\mathbf{R}|H_1)$ define the ‘Probabilistic transition

mechanism'. If these two probability density functions are known completely, then the 'Decision rule' is said to be a *Simple Binary Hypothesis Test*. We shall see in this case that the 'Decision rule' is in the form of a *likelihood ratio test* (LRT). If on the other hand, only the functional form of the two densities is known, and there is a vector of unknown parameters, then these parameters will need to be estimated first, before detection can proceed. This type of problem is called a *Composite Hypothesis Test* and leads to the formulation of a *generalized likelihood ratio test* (GLRT).

In any event, the detector design will be influenced by the detection criteria chosen. For example, the detector may be optimized to minimize the probability of an error, or it might be optimized to minimize the probability of a false alarm. We will examine two criteria that are important to our problem. They are the Bayes criterion and the Neyman-Pearson criterion. These are presented in Section 2.3.1 and Section 2.3.2 respectively, where only the *Simple Binary Hypothesis Test* is considered. In Section 2.3.3, the ML method is considered for the estimation of unknown parameters. These ideas are then used in Section 2.3.4 to develop the GLRT.

2.3.1 Bayesian Detectors

The function of the detector is simply to operate on the received data vector \mathbf{R} and to choose either H_0 or H_1 as the guess to which \mathbf{R} belongs. Since \mathbf{R} does in fact belong to H_0 or H_1 , there are four possible input/output combinations each time the receiver operates on a data vector. Detection criteria place a relative importance or weight on each of the four courses of action.

The Bayes criterion is based on two assumptions:

1. It is assumed that the 'Source' outputs are governed by *a priori* probabilities. For our problem, the chance of having a growler and clutter in a radar resolution cell is denoted by the probability P_1 , and the chance of having only clutter in a radar resolution cell is denoted by the probability P_0 . Clearly, $P_0 + P_1 = 1$.

2. Costs are assigned to each of the four possible courses of action. These costs are denoted by C_{00} , C_{01} , C_{10} , and C_{11} . The first subscript indicates the hypothesis chosen by the detector, and the second subscript indicates the hypothesis that was actually true.

Every time a decision is made by the detector, a certain cost will be incurred. The Bayes criterion is designed to minimize the *average cost* incurred by the detector. The *average cost* or *risk* is denoted by \mathcal{R} and is expressed as:

$$\begin{aligned}\mathcal{R} = & C_{00}P_0\text{Prob}(\text{say } H_0|H_0 \text{ is true}) + \\ & C_{10}P_0\text{Prob}(\text{say } H_1|H_0 \text{ is true}) + \\ & C_{01}P_1\text{Prob}(\text{say } H_0|H_1 \text{ is true}) + \\ & C_{11}P_1\text{Prob}(\text{say } H_1|H_1 \text{ is true}).\end{aligned}\tag{2.7}$$

Notice that we are averaging over two different probabilities in the evaluation of the average cost \mathcal{R} .

It is convenient to view detector design as a division of the observation space defined by the received vector \mathbf{r} . If we denote the total observation space by Z , then we must divide the observation space into two regions, Z_0 and Z_1 . When an observation \mathbf{R} falls into Z_0 we choose H_0 , and when \mathbf{R} falls into Z_1 we choose H_1 . The risk function can now be rewritten in terms of the transition probabilities:

$$\begin{aligned}\mathcal{R} = & C_{00}P_0 \int_{Z_0} p_{\mathbf{r}|H_0}(\mathbf{R}|H_0)d\mathbf{R} + \\ & C_{10}P_0 \int_{Z_1} p_{\mathbf{r}|H_0}(\mathbf{R}|H_0)d\mathbf{R} + \\ & C_{01}P_1 \int_{Z_0} p_{\mathbf{r}|H_1}(\mathbf{R}|H_1)d\mathbf{R} + \\ & C_{11}P_1 \int_{Z_1} p_{\mathbf{r}|H_1}(\mathbf{R}|H_1)d\mathbf{R}.\end{aligned}\tag{2.8}$$

Note that these integrals are N -dimensional. The Bayes test reduces to solving for the decision regions Z_0 and Z_1 . Recalling that $Z = Z_0 + Z_1$ and assuming that the cost of

making an error is greater than the cost of making a correct decision, minimization of Equation 2.8 reduces to

$$\frac{P_{\mathbf{r}|H_1}(\mathbf{R}|H_1)}{P_{\mathbf{r}|H_0}(\mathbf{R}|H_0)} \underset{H_0}{\overset{H_1}{\gtrless}} \frac{P_0(C_{10} - C_{00})}{P_1(C_{01} - C_{11})}. \quad (2.9)$$

The quantity on the left hand side is called the *likelihood ratio* and is denoted by $\Lambda(\mathbf{R})$ and the quantity on the right hand side is the threshold of the test and is denoted by η . Thus the Bayes criterion results with a *likelihood ratio test* given by

$$\Lambda(\mathbf{R}) \underset{H_0}{\overset{H_1}{\gtrless}} \eta. \quad (2.10)$$

Often, because of the mathematical form of the transition probability distributions, it is convenient to use the log likelihood ratio test (LLRT) given by

$$\ln \Lambda(\mathbf{R}) \underset{H_0}{\overset{H_1}{\gtrless}} \ln \eta. \quad (2.11)$$

This is justified because the natural logarithm is a monotonic function.

Let's examine the Bayes LRT more closely. Only the threshold η is dependent on the *a priori* probabilities and the cost assignments. Therefore the processor (feature-extractor) which operates on the data to form the likelihood ratio $\Lambda(\mathbf{R})$ can be designed independently. Recall in Figure 2.1 that only the final stage of the general detector receiver requires η . This property is significant from a practical standpoint since often the costs and *a priori* probabilities are not known or are mere guesses. The form of the Bayes LRT is intuitively satisfying. For example, if the *a priori* probabilities are equal, the cost of both errors are equal, and the cost of both correct decisions are the same, then $\eta = 1$. This implies that the test simply chooses the hypothesis that most likely produced the observed data vector \mathbf{R} ! Other Bayes tests arise simply by using different cost assignments or *a priori*

probabilities. For example, let the costs associated with making an error have a value of 1, and let the costs associated with making a correct decision have a value of 0. In this way, minimizing the risk is equivalent to minimizing the probability of error. It can be shown that this test in fact maximizes the ratio of the *a posteriori* probabilities. We will see in the next section what happens when cost assignments and *a priori* probabilities are unknown.

2.3.2 Neyman-Pearson Detectors

In many real problems, it is difficult if not impossible to assign realistic costs or *a priori* probabilities. Furthermore, even when the costs and *a priori* probabilities are known, the GCR and growler to noise ratio (GNR) are usually not known. In situations like these, an alternative method for setting the threshold in the LRT defined in Equation 2.9 must be sought. The Neyman-Pearson criterion is a simple way of dealing with cases like this, and as a result, it has found wide use in radar problems. The performance of radar systems are often described by the *probability of false alarm rate* and the *probability of detection*. These can be defined in terms of the transition probabilities and are given by

$$\begin{aligned} P_F &= \int_{Z_1} p_{\mathbf{r}|H_0}(\mathbf{R}|H_0)d\mathbf{R}, \\ P_D &= \int_{Z_1} p_{\mathbf{r}|H_1}(\mathbf{R}|H_1)d\mathbf{R}. \end{aligned} \quad (2.12)$$

Ideally, we would like to make P_F as small as possible and P_D as large as possible, but these are usually conflicting requirements. Alternatively, we could constrain P_F to some tolerable value and maximize P_D subject to the constraint. This type of requirement is solvable using Lagrange multipliers.

The Neyman-Pearson criterion is stated as follows. Constrain $P_F = \alpha' \leq \alpha$ and design a processor that will maximize P_D subject to this constraint. Solving this problem

with Lagrange multipliers gives the following LRT:

$$\Lambda(\mathbf{R}) \underset{H_0}{\overset{H_1}{\gtrless}} \lambda. \quad (2.13)$$

Satisfying the constraint so that $P_F = \alpha'$ is equivalent to requiring that

$$P_F = \int_{\lambda}^{\infty} p_{\Lambda|H_0}(\Lambda|H_0)d\Lambda = \alpha'. \quad (2.14)$$

In this equation, $p_{\Lambda|H_0}(\Lambda|H_0)$ is the density of Λ when H_0 is true. Solving Equation 2.14 for λ gives the threshold required for the Neyman-Pearson test. Notice that as the threshold λ decreases, P_D increases. Therefore, λ is decreased in Equation 2.14 until we get the largest possible $\alpha' \leq \alpha$.

For both the Bayes and Neyman-Pearson criteria, the same likelihood ratio results; only the threshold is different. This should come as no surprise since the Neyman-Pearson criterion is really a special case of the Bayes formulation. Many other LRT detectors are possible by defining new criteria and constraints.

Up to this point, we have only been concerned with detection. In the next section, we will review some classical estimation theory which will provide the basis for estimating unknown parameters in the received data vector \mathbf{R} . For example, in our problem, we might be interested in estimating the velocity of the growler.

2.3.3 Maximum Likelihood Estimation

In the previous two sections, we were concerned with the detection problem. In this section, we are interested in a closely related problem. The 'Source' in Figure 2.2 now generates a parameter in the 'Parameter space' as is shown in Figure 2.3. This parameter gets mapped onto the 'Observation space' by the 'Probabilistic mapping'. Then an 'Estimation rule' is used to obtain an estimate of the parameter value. Parameters of interest fall into two classes: random, or unknown and not random. In the random parameter case, a Bayes




Figure 2.3: Parameter estimation model

formulation similar to the detection problem can be derived where the cost function is minimized over both the parameter density functions as well as the ‘Probabilistic mapping’ functions. Since in our work, we are more interested in the unknown parameter problem, we do not review the random parameter case here. It is described in detail in [23].

Assume that based on an observation vector, \mathbf{R} , we wish to estimate the value of an unknown parameter \mathbf{a} whose actual value happens to be A , and let’s denote this estimate by $\hat{\mathbf{a}}(\mathbf{R})$. In order to come up with an estimation rule, some measure of quality must be defined. The Bayes criterion is not applicable in this case. One quality measure of value is that the estimator be unbiased. That is,

$$E\{\hat{\mathbf{a}}(\mathbf{R})\} \triangleq \int_{-\infty}^{+\infty} \hat{\mathbf{a}}(\mathbf{R}) p_{\mathbf{R}|a}(\mathbf{R}|A) d\mathbf{R} = A \quad (2.15)$$

where E is the statistical expectation operator. Even an unbiased estimate can perform poorly on a given realization if the variance of the estimate is large. The variance of an unbiased estimate is given by:

$$\text{Var}\{\hat{\mathbf{a}}(\mathbf{R})\} = E\{[\hat{\mathbf{a}}(\mathbf{R}) - A]^2\}. \quad (2.16)$$

Therefore, we would like to find an unbiased estimator with a small variance. Unfortunately, a straight-forward procedure for finding a minimum-variance, unbiased estimator does not

exist. We will have to be content with trying various procedures and then seeing how well they perform.

The *Maximum Likelihood Estimator* is one of the most famous and best-performing estimation procedures known. The criterion is to choose as the estimate the value that *most likely* caused the observation \mathbf{R} to occur. We begin by defining the *likelihood function* $p_{\mathbf{R}|a}(\mathbf{R}|A)$, viewed as a function of A since the observation vector \mathbf{R} is not variable once it has been observed. As with the LRT, it is often convenient to work with the *log likelihood function*. The maximum likelihood estimate of a is the value of A that causes the likelihood function to be maximum. This estimate is denoted by $\hat{a}_{ml}(\mathbf{R})$. Assuming that the maximum is interior to the range of A , and the log likelihood has a continuous first derivative, a necessary condition on $\hat{a}_{ml}(\mathbf{R})$ is given by:

$$\left. \frac{\partial \ln p_{\mathbf{R}|a}(\mathbf{R}|A)}{\partial A} \right|_{A=\hat{a}_{ml}(\mathbf{R})} = 0. \quad (2.17)$$

This equation is called the *likelihood equation* and its solution is the maximum likelihood estimate of a .

The effectiveness of the ML procedure can be appreciated by looking at its mean and variance. In many cases, it is difficult to compute the mean and variance of an estimator directly. It is often easier to look at the asymptotic properties. Before doing so, we will describe a lower bound on the variance of any unbiased estimator. This will give us a measure against which we can compare the performance of the ML estimator.

Theorem 2.3.1 *If $\hat{a}(\mathbf{R})$ is any unbiased estimate of A , then*

$$\text{Var}[\hat{a}(\mathbf{R})] \geq \left(E \left\{ \left[\frac{\partial \ln p_{\mathbf{R}|a}(\mathbf{R}|A)}{\partial A} \right]^2 \right\} \right)^{-1} \quad (2.18)$$

provided that $\partial p_{\mathbf{R}|a}(\mathbf{R}|A)/\partial A$ exists and is absolutely integrable. This result is known as the Cramér-Rao (CR) bound [24, 25, 26, 27].

We call any estimate that satisfies the bound with an equality an *efficient* estimate. It can be shown [23] that if an efficient estimate exists, it is the maximum likelihood estimate

$\hat{a}_{ml}(\mathbf{R})$. This is the reason why the ML estimator is so popular. It is important though to note the following:

1. If the ML estimate is biased, then the CR bound does not apply, and the asymptotic performance of the estimate is unknown.
2. If an efficient estimate does not exist, again the performance of the ML estimate or any other is not known. For an efficient estimate $\hat{a}(\mathbf{R})$ to exist, the following equation must be satisfied:

$$\frac{\partial \ln p_{\mathbf{r}|a}(\mathbf{R}|A)}{\partial A} = [\hat{a}(\mathbf{R}) - A]k(A) \quad (2.19)$$

for all \mathbf{R} and A .

In the next section, we combine maximum likelihood estimation and the likelihood ratio test in order to derive a generalized likelihood ratio test for binary hypothesis testing problems with unknown parameters.

2.3.4 Composite Hypothesis Testing

In this section, we are interested in performing *composite* hypothesis testing. A composite hypothesis is one where the form of the transition probability density function is known, but there are one or more parameters in the pdf that are not known. As in the last section, we are interested in the case where the unknown parameters are nonrandom. Again the Bayesian formulation does not apply since there are no *a priori* pdfs to average over, so we consider only Neyman-Pearson tests.

As in the last section, we must have a definition of the quality of the test. Consider first what is known as the *perfect measurement* bound. No test that we could design would be better than a receiver that could first measure the unknown parameters perfectly, and then design the optimum likelihood ratio test. Thus, this fictitious test can serve to bound the performance of any other test. Such bounds are usually formulated in the Neyman-Pearson sense. That is, for a given P_F , a curve is calculated which gives the P_D of the

perfect measurement bound as a function of the possible range of parameter values.

A test that meets the bound for the whole range of parameter values, having fixed the P_F is said to be a *uniformly most powerful* (UMP) test. Stated differently, for a given P_F a UMP test has a P_D greater than or equal to that of any other test, for the whole range of possible parameter values. It can be shown that a UMP test exists if and only if the likelihood ratio test including the threshold can be completely defined for all possible parameter values, without knowledge of the parameter values. Clearly, whenever a UMP test exists, it should be used as it performs just as well as if the unknown parameters were known perfectly. However, for many problems of practical interest, a UMP test does not exist. The perfect measurement bound suggests that we might first estimate the unknown parameters assuming each of the hypotheses is true, and then forming a likelihood ratio test as though the estimated values were the true values. Because of the excellent properties of the ML estimator, we use ML estimates in this procedure. The *generalized likelihood ratio test* can be written as:

$$\Lambda_g(\mathbf{R}) = \frac{\max_{\Theta_1} p_{\mathbf{R}|\Theta_1}(\mathbf{R}|\Theta_1)}{\max_{\Theta_0} p_{\mathbf{R}|\Theta_0}(\mathbf{R}|\Theta_0)} \underset{H_0}{\overset{H_1}{\gtrless}} \gamma \quad (2.20)$$

where the parameter values Θ_1 vary over the permissible range of the parameter vector Θ in hypothesis H_1 , and the parameter values Θ_0 vary over the permissible range of the parameter vector Θ in hypothesis H_0 . The numerator in Equation 2.20 is equivalent to making a ML estimate of Θ_1 assuming that H_1 is true and then calculating the likelihood using this estimate. In a similar manner, the denominator is evaluated. The threshold γ is the result of solving the Neyman-Pearson constraint equation as described in Section 2.3.2.

2.3.5 Sufficient Statistic

By incorporating detection and estimation theory, we have derived the generalized likelihood ratio test for a composite binary hypothesis problem. As we have seen, the transition

probabilities of the received data vector are central to the formation of the likelihood ratio. In many cases, one finds that the likelihood ratio $\Lambda_g(\mathbf{R})$ depends explicitly on a particular function of the received data vector $l(\mathbf{R})$. The function $l(\mathbf{R})$ is called a *sufficient statistic* because knowing it is equivalent to knowing \mathbf{R} in so far as the detector design is concerned. The point of a sufficient statistic is perhaps missed if it is merely stumbled across in the formation of a likelihood ratio. If $l(\mathbf{R})$ is recognized directly though, the problem can be simplified greatly.

The vector \mathbf{r} is a point in the N -dimensional space described by the coordinates r_1, r_2, \dots, r_N . The likelihood ratio in the form that we have described makes use of this coordinate system. There may indeed be a better coordinate system to work in for the purpose of evaluating $\Lambda_g(\mathbf{R})$. A sufficient statistic simply describes a single coordinate in this new coordinate system, and in fact, by its definition, it is the only coordinate that has any effect on the decision. Let l denote the sufficient statistic, and let the remaining $N - 1$ coordinates which do not affect the decision be denoted by the vector \mathbf{y} . Therefore

$$\Lambda_g(\mathbf{R}) = \Lambda(L, \mathbf{Y}) = \frac{p_{l, \mathbf{y} | H_1}(L, \mathbf{Y} | H_1)}{p_{l, \mathbf{y} | H_0}(L, \mathbf{Y} | H_0)}. \quad (2.21)$$

Since l is a sufficient statistic, this reduces to

$$\Lambda_g(\mathbf{R}) = \Lambda(L) = \frac{p_{l | H_1}(L | H_1)}{p_{l | H_0}(L | H_0)}. \quad (2.22)$$

The simplification resulting from the recognition of a sufficient statistic is now apparent. Only the probability density functions of l under both hypotheses are necessary to form the optimum processor. This simplification is significant when closed form expressions are not available for the transition probability densities and an empirical evaluation of the performance of a detector is required. We make use of this concept in Chapter 6.

In the next section, we look at alternative detector designs when the transition probability density functions vary spatially and temporally.

2.4 Constant False Alarm Rate Detectors

The theory presented in Section 2.3 are generally applicable to all forms of targets and clutter. However, in cases like ours where the clutter is spatially nonstationary, the detector performance becomes spatially (and hence temporally) dependent. This spatial variation can be handled in general by the composite hypothesis test; but an accurate description of the spatial variation of clutter is required. Since this description is usually not available in practice, other techniques must be sought which more easily account for the spatial variation of the clutter.

In this section, we examine techniques that are useful in maintaining constant false alarm rates (CFAR) in radar detection. Our primary interest concerns the detection of small targets in a marine environment, where the nonhomogeneous and nonstationary properties of sea clutter offer a particular challenge. After discussing the CFAR techniques that are in use today, we determine the most useful ones for our problem. Both noncoherent and coherent CFAR processors are considered, and use is made of our findings that sea clutter amplitude statistics are K-distributed [13, 14, 28].

2.4.1 Understanding CFAR Processors

There is no shortage of CFAR processors described in the literature. A wide range of CFAR detectors have been proposed for many different applications. Rather than giving an extensive review of all CFAR designs currently in use, we develop basic CFAR principles that are appropriate to our problem. In this way, we quickly obtain the understanding needed to apply CFAR principles to our problem. For a detailed look at individual CFAR designs, the interested reader is referred to [5, 29, 30, 31, 32, 33, 34, 35, 36, 37, 38, 39, 40, 41, 42, 43, 44].

Due to the large number of CFAR designs that have been proposed, developing a general understanding of the basic principles is essential. Therefore, our discussion begins with a common framework by which all CFAR processors can be described. This discussion

closely follows that given by Farina and Studer [29].

CFAR detectors are either noncoherent or coherent processors. The noncoherent CFAR schemes all involve adaptively setting the decision threshold in some manner. These processors can be subdivided into two groups, those that use spatial processing, and those that use temporal processing. The spatial CFAR's involve using the returns from neighbouring cells to determine the required detection threshold for a given probability of false alarm (PFA). On the other hand, the temporal methods use scan to scan processing at each resolution cell in order to determine the threshold. Clutter maps, two dimensional maps (azimuth vs range) where the mean amplitude in each cell is determined by successive scans, are used for comparison to each test cell. This method is particularly susceptible to self-masking by targets in the radar cell, and therefore the clutter map cell size is generally larger than the radar resolution cell size. Drawbacks result from making the clutter map cells too large as well. Other noncoherent methods include nonlinear receivers, an example of which is the binary integrator.

Coherent CFAR techniques are also divided into two main categories: coherent clutter maps and spectral methods; however, only the spectral methods are applicable for sea clutter (due to the uniform phase distribution). Therefore, coherent clutter maps are not considered any further. The spectral methods operate as follows. For each of several range cells, the spectrum is computed using either Doppler filter banks (FFT) or other spectrum estimators. Then each Doppler bin for the test range cell is compared with the Doppler bins at neighbouring range cells in order to determine whether or not a target is present. In this way, CFAR is achieved by setting the adaptive threshold for the test cell based upon the corresponding Doppler bins of neighbouring cells. Not only is a decision made, but also the Doppler frequency of the target is determined.

CFAR processors attempt to account for varying clutter fields by forming local estimates of the clutter statistics in the vicinity of the resolution cell under test. This is done by sliding a window through range (range CFAR's) or time (clutter map CFAR's).

We consider only range CFAR processors in the sequel, because the clutter map approach is more susceptible to target self-masking. Usually, the center cell in the CFAR window is the cell under test, and the neighbouring cells are used to estimate the local clutter power. The estimate of the local clutter power is then used to select a threshold for the test cell. The cells in the CFAR window that are used to estimate the clutter power in the test cell can be combined in several different ways. For example, they might be averaged, or the largest clutter value might be used in setting the local threshold. The critical assumption common to these approaches is that the clutter statistics in the test cell are the same as the statistics in the other cells in the CFAR window. Therefore, the CFAR window size must be picked with care.

An excellent paper by Rohling discusses several ways in which these local clutter power estimates can be made from the cells in a CFAR window [34]. He introduces the familiar CA-CFAR (Cell Averaging CFAR) in which a window in range of length N around the test cell is used to predict the average clutter power (Rayleigh envelope assumed). It is also assumed that there is a guard cell on either side of the test cell, in order to limit target leakage. He then introduces CAGO-CFAR (Cell Averaging Greatest of CFAR) where the clutter power is estimated independently on both sides of the test cell, and the greatest clutter power is used to determine the threshold. These methods are shown to be disappointing when there is a clutter field with edges within the CFAR window as shown in Figure 2.4. Both correlated and uncorrelated clutter fields are considered. Also, multiple target cases are considered and these techniques are shown to mask at least one of the targets. These shortcomings motivated the development of the order-statistic CFAR (OS-CFAR), which is a true CFAR in that it is independent of the clutter power, but which also works well under the situations described above, as well as being superior under the homogeneous clutter case. This scheme amounts to setting the threshold S according to the equation

$$S = T * Z \quad (2.23)$$

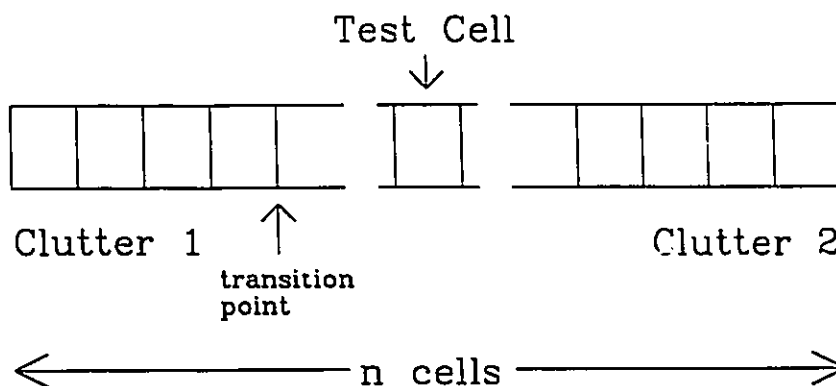


Figure 2.4: A CFAR window spanning two clutter fields.

where

$$Z = X_{(k)} \quad (2.24)$$

is the k^{th} order statistic (ie. the k^{th} largest sample in the CFAR window) and T is a constant that controls (along with Z) the PFA. The Rayleigh envelope or Negative Exponential power distribution is considered for uniform clutter, and it is shown that indeed the PFA expression is independent of the clutter power parameter μ in

$$p(\sigma) = \frac{1}{\mu} \exp(-\sigma/\mu) \quad (2.25)$$

where $p(\sigma)$ is the negative exponential density function which describes the clutter power. This processor is shown to be superior to both the CA-CFAR and CAGO-CFAR processors in multiple target situations and with jamming or clutter edges. Furthermore, no guard cells are required for this processor. It should be emphasized that the OS-CFAR processor makes full use of the underlying distribution of the clutter process. This is not a rank-order

method that only makes use of the order, and neglects the sample value at a given order! It is shown that for $N = 24$ and $N = 32$, a suitable value for k is in the range

$$\frac{N}{2} \leq k \leq \frac{3N}{4}. \quad (2.26)$$

A different approach to CFAR design is taken by Bucciarelli in [38]. Rather than designing a detector that works reasonably well for a number of probability density functions (pdfs), he uses the moments of the sample data to distinguish between a Weibull or lognormal pdf (the principle can work for any pdf) and the threshold is then selected accordingly. He uses the first three moments to govern his threshold selection. A window in range or azimuth is used rather than a clutter map. He does not, however, consider the problems of target self-masking and nonhomogeneous clutter fields. His concern is primarily a CFAR selection of the detection threshold, based on the identified distribution.

The detectors described by Rohling and Buciarelli are noncoherent CFAR detectors. The most promising coherent CFAR detectors for detection of growlers in sea clutter are of the spectral type, since growler motion is considerably different than the motion of the sea. Larsson *et al.* describe such a CFAR detector that they use with the ICERAD Radar [5]. A range/Doppler map forms the basis of their CFAR detector. Their processor collects 512 temporal hits for 32 consecutive range cells, and performs a 512-point FFT for each range bin. Then each FFT bin is compared with its neighbours. They make use of an order statistic detector; the test cell must exceed the M th order statistic for a target to be declared. However, rather than ordering their data in each CFAR window, they make use of a clever comparison scheme which requires only N comparisons rather than $N * \log(N)$ for sorting.

The CFAR detectors just described are representative of the kinds of CFAR detectors that have been proposed in the literature. Each of these methods fall into the general framework of CFAR processors referred to by Farina *et al* [29]. In the next section, the applicability of these CFAR designs to the growler detection problem is considered.

2.4.2 CFAR Processors for Sea Clutter

Our investigations (see Chapter 4) indicate that sea clutter amplitude statistics are described well by the K-distribution. The K-distribution is given by

$$p(x) = \frac{4c}{\Gamma(v)} (cx)^v K_{v-1}(2cx) \quad (2.27)$$

where $v \geq 0$ and $0 \leq x \leq \infty$. We can also write this distribution as a conditional probability density function (pdf) modulated by a mean level term y . That is,

$$p(x) = \int_0^{\infty} p(y)p(x|y)dy \quad (2.28)$$

where

$$p(y) = \frac{2b^{2v}}{\Gamma(v)} y^{2v-1} \exp(-b^2 y^2), \quad 0 < y < \infty \quad (2.29)$$

$$p(x|y) = \frac{\pi x}{2y^2} \exp\left(-\frac{\pi x^2}{4y^2}\right), \quad 0 < x < \infty \quad (2.30)$$

and $c = \sqrt{\pi/4b}$. The pdf $p(y)$ represents the mean or local level of the sea clutter amplitude and $p(x|y)$ represents the speckle component, which is Rayleigh distributed with a mean value of y . This decomposition indicates that if CFAR processing is done over a small enough region so that the mean level component y is relatively unchanged, then the local statistics will be Rayleigh, allowing the use of all of the CFAR processors considered in the previous section.

For noncoherent processing of the return amplitudes, the OS-CFAR processor described by Rohling [34] is the most suitable. It is the most robust against target masking problems and clutter edges. Furthermore, it makes full use of the underlying statistics of the data. It can also be implemented very easily. If range samples are used in the CFAR block, the size of the block must be small enough to keep a relatively constant mean level, so that the resulting statistics will be almost Rayleigh. Azimuthal samples may also be used to estimate the clutter power in the CFAR block. Time samples are not recommended since targets in the radar resolution cells will cause problems. At any rate, the OS-CFAR

should be more robust than any of the other methods described. One alternative would be to consider an approach such as that used by Bucciarelli et al. [38] where the scale and shape parameters of the K-distributed clutter are estimated from the CFAR block, and then an appropriate threshold is selected from table lookup, for a prescribed PFA. Although this type of approach would track slow changes in the overall clutter statistics, it is not sensitive to tracking the local changes in the underlying mean level component of the sea clutter returns.

For coherent CFAR processors, the method described by Larsson *et al.* [5] looks promising. A window of range bins would be analyzed simultaneously, surrounding the test cell. A DFT operation would be done at each range bin (implying multiple hits) and then the corresponding Doppler bins from each range cell would make up the CFAR block. An OS-CFAR would then be applied to the amplitudes of the corresponding Doppler bins in order to decide whether a target is present in the test cell, as indicated by a large component in any test cell Doppler bin, relative to the corresponding surrounding bins. In this way, the processor determines whether or not a target is present, as well as providing the target's velocity. An alternative procedure to the OS operation would be to empirically or theoretically derive the statistics of the Doppler bin amplitudes, and then apply a CFAR operation custom designed to those statistics.

2.5 Summary

In this chapter, we reviewed basic detection and estimation theories that are applicable to our problem. We began by considering the general receiver structure shown in Figure 2.1. Then, we called on Matched Filter theory in order to design a front-end filter for our detector. The benefits that could be achieved in nonadditive clutter environments were discussed. With the front-end filter designed, theories suitable for the design of the feature-extractor were considered. We began by reviewing classical detection and estimation theory. In

particular, the Bayes and Neyman-Pearson criteria were derived for the simple binary hypothesis test. In order to derive the generalized binary hypothesis test, maximum likelihood estimation procedures were examined. We concluded by describing the benefits afforded by sufficient statistics. In the last section, we surveyed CFAR methods that are in use today when the clutter or target processes are nonstationary. We concluded by recommending CFAR designs that seem to be most suitable to our problem of detection of growlers in sea clutter. The theory described in this chapter will serve as the basis for the methods that are described in the remainder of this thesis.

Chapter 3

Field Experiments and the IPIX Radar

3.1 Introduction

As described in Chapter 1, a key element in our research plan is the development of the IPIX radar. This radar has enabled us to make real measurements of growler and sea backscatter at X-band. Since the results of this research are based on the data collected with the IPIX radar, and since the contribution of the data library itself is significant, a detailed description of IPIX is in order. Although details of the IPIX radar are available in the literature [10, 11, 12], a summary will be given here for completeness. This chapter is devoted to describing the IPIX radar along with the experiments that were performed with this instrument.

3.2 The IPIX Radar

As a result of other studies [1, 2, 3], in 1983 the CRL identified the need for an instrumentation-quality radar system designed specifically for research use. At that time, several

system requirements were identified by considering the problem of research as well as the interaction between the end user and the radar. It was decided that the radar should have:

1. Dual-Polarization
2. Coherent Transmission/Reception
3. Pulse Compression
4. X-band Transmission
5. Digital Data Acquisition
6. Built-In Calibration
7. Flexible Operation and Modifications
8. Computer Control

System development began in 1984 with funding from the Natural Sciences and Engineering Research Council of Canada. In the summer of 1986, a prototype was tested and the present system known as the IPIX (for Intelligent PIXel-processing) radar emerged.

3.2.1 System Features and Capabilities

All of the design criteria were met with the IPIX radar. Each system requirement enumerated above will be dealt with here in detail.

Dual Polarization: The IPIX radar is dual-polarized. It has two identical receivers; one is connected to the vertical polarized antenna feed and the other to the horizontal polarized antenna feed. Therefore, both polarizations are always received. A high speed ferrite waveguide switch having 50 dB of isolation between channels is used to route the transmitted signal through either the horizontal or vertical channel. To achieve this kind of isolation, a combination of three ferrite devices are actually used in the switch. The switch can change state at a continuous rate of up to 2 kHz. Therefore the radar is capable

of near simultaneous transmission of both orthogonal polarizations, which is necessary for measuring the full polarization scattering matrix of a target.

Due to the high channel-to-channel isolation of the ferrite switch, the antenna is the primary determining factor in the horizontal to vertical channel isolation figure. Currently, a 2.2 degree parabolic dish antenna is used with about 25 dB isolation between channels. This antenna is mounted on an antenna positioner unit that is operated under computer control. The two-axis pedestal is capable of repeating absolute fixed positions in azimuth and elevation to better than 0.1 degree accuracy. It is capable of spin rates of up to 30 RPM.

Coherent Transmission/Reception: The IPIX radar is a coherent radar allowing for accurate measurements of the phase of the returned radar echo. A high-stability 5 MHz crystal oscillator is used as the master reference clock for the entire system. Both IF (150 MHz) and RF (9.24 GHz) sources are phase-locked to the master clock and are used to generate the transmitted 9.39 GHz signal as well as to downconvert and demodulate the received signals. When pulse compression is in operation, the SAW pulse expander is synchronized with the IF clock. Quadrature demodulators in each receiver channel provide the inphase (I) and quadrature (Q) video outputs for subsequent processing.

Pulse Compression: The IPIX radar has a pulse compression subsystem consisting of a SAW pulse expander and matched SAW compressors in each of the two receivers. The expander generates a nonlinear FM coded pulse 5 microseconds wide with a bandwidth of about 50 MHz. The compressed pulse has an effective width of 32 nanoseconds. The importance of pulse compression in IPIX is twofold. First, it makes IPIX a high resolution radar with a range resolution of 4.8 metres, allowing for improved visibility of small targets. Second, higher average power is delivered to the targets increasing the maximum range at which a given target can be detected. Since we have no control of the location of targets of opportunity with respect to the radar, this increase in usable range is significant.

X-band Transmit Signal: The IPIX radar is an X-band radar that transmits a peak power of 8 kW at 9.39 GHz. The wavelength of the transmitted signal is approximately

3 cm. This wavelength couples well with the small capillary waves on the ocean surface, allowing for useful sea clutter measurements. Furthermore, at this wavelength, a target velocity of 1 knot causes a Doppler shift of about 34 Hz. At our maximum PRF of 2 kHz, this allows for unaliased velocity measurements of up to 30 knots, which is a reasonable limit for the targets that we will study.

Digital Data Acquisition: The data acquisition system is probably the most important aspect governing the utility of a radar for research use. Much effort was put into the design of this system in IPIX. Digital data acquisition in radar is challenging because of the high data sampling rates and volume of data involved. Although high sampling rates are feasible with current technology, data storage is still of vital concern, especially with regard to offline processing and data retrieval. High-speed analog or digital tape recorders are capable of handling the continuous streams of data; however, they are not very suitable for offline recovery of a particular window of data that may be of interest. The IPIX radar produces data at a peak rate of 120 Megabytes per second and can sustain a rate of 16 Megabytes per second. The IPIX sampling system is programmed to only store data in a user-specified window defined in azimuth and range. Thus, the storage facilities only need to handle the data that is of interest to the user.

The IPIX data acquisition system operates as follows: After each pulse is transmitted, the selected range interval is digitized at a rate of 30 MHz into the sweep buffer. Each range sample consists of the four receiver channels: horizontal I (HI) and Q (HQ) and vertical I (VI) and Q (VQ). Before the next pulse is transmitted, the sweep buffer is written via a high-speed VMEbus DMA channel directly into the memory space (16 Megabytes currently) of a 68020-based computer system (now upgraded to 68040 system). Each sweep is stored sequentially in memory. The most significant benefit of this approach is that the data are available for processing immediately, and are addressable by a simple pointer to memory. A high-capacity digital tape drive is used to store large numbers of data sets inexpensively, either in raw form, or after processing.

Built-in Calibration: The IPIX radar has built-in calibration equipment (BICE) consisting of a computer controlled frequency-agile IF source that can be routed to the transmitter's upconverter to produce an RF signal. This source, along with a digitally controlled attenuator, can then be used to inject known signals into the receiver, in order to produce an input/output characteristic. Sensitivity time control (STC) is also available to ensure that the received signal falls within the linear range of the receiver.

Flexible Operation and Modification: The IPIX radar has been designed so that it is flexible and modifiable. For flexible operation, remote-controlled coaxial switches are used throughout the system to allow computer control of the RF signal path. These switches make it a simple matter to configure the system for normal operation, self-calibration and testing. Furthermore, hardware circuits monitor the state of these switches in order to ensure that they are in a safe configuration. The system is also easily modified. This is due to its modular design. For example, multi-frequency enhancements are now being made to the system, as well as sampling system upgrades without affecting any other part of the system. All modules are also rack-mounted and easily accessible to allow for in-the-field modifications and repairs.

Computer Control: Virtually all aspects of the IPIX radar are changeable by computer control. These include:

- The mode of operation which is affected by coaxial switches.
- The transmit signal characteristics that include pulse width, PRF, and polarization.
- The sampling system that specifies the azimuth and range windows required and the number of sweeps to store.
- The antenna positioner unit that controls the position and scan rate of the antenna.
- The data storage facility, and the generation of log files for each experiment.

The computer interface for most of the control operations is a simple parallel bus

that uses TTL signals and can be implemented with four 8-bit parallel ports. The control software is written in the C programming language for portability. Any computer with 32 bits of parallel I/O and a C compiler can serve as the control computer.

3.3 Field Experiments

With the IPIX radar developed, the experimental program began. A suitable site for conducting experiments with the IPIX radar was identified at Cape Bonavista, Newfoundland and is shown in Figure 3.1. A coast guard building there was used to house the radar and support equipment. A concrete pad immediately in front of the building was poured as the platform for the antenna positioner unit and the antenna. The platform was on a cliff edge overlooking the Atlantic ocean at a height of 22 metres above sea level. This height simulates the height of a ship-based antenna. The antenna had a viewing area that spans from a bearing of 20 degrees to 150 degrees, with north at 0 degrees. The ocean floor fell off quickly to a depth of about 300 feet so that deep seas were experienced.

The first field trials were conducted in May 1988 at the Cape Bonavista site. Several very large icebergs and a variety of bergy bits and growlers were sampled at this time, and the system was tested thoroughly. Many improvements were identified and the IPIX team returned to the CRL to implement them. Over the next year two other field trials were conducted at the site as the system matured. Finally, in June of 1989, our fourth and most successful field trip was conducted. A large database of sea clutter and growler measurements were made, along with measurements of fishing boats and life rafts. Furthermore, a 2 metre RCS (radar cross-section) spar reflector was moored on a nonreflective buoy for some man-made target measurements. The data files that are analyzed in this thesis were all collected during the June 1989 field trip. The radar parameters associated with each of these data files will now be described.




Figure 3.1: The map shows the island of Newfoundland on the East Coast of Canada. The IPIX radar test site is located at Cape Bonavista on the eastern shore of the island.

3.3.1 The Data Base

A variety of data files were collected during the last field trip conducted with IPIX between June 02, 1989 and June 14, 1989. Each data file consisted of 16 Megabytes of data, representing 4 Megasamples where each sample consists of an 8-bit HI, HQ, VI, and VQ channel measurement. Although scan data was collected, all of the data described here were collected with the antenna staring along a particular azimuth direction. Data were collected between a range start gate and a range stop gate that were specified in the sampling system. The sampling clock ran at 30 MHz so that a sample in range was made every 5 metres. For a given data file, consecutive samples were stored until the 16 Megabytes of memory were filled. The data were then transferred to tape, as another experiment was prepared. By keeping the antenna fixed, we were able to collect a large volume of data from a particular target or from a small patch of sea clutter. This mode of operation is useful for gathering statistics on target and clutter processes.

Each experiment is associated with an experimental identification number. Experiment ID's for the June 1989 trials started with B0 and ran consecutively to B400. That is, there were 401 individual experiments performed, each resulting with a 16 Megabyte data file. In this thesis, two sets of experiments are described. The first set of experiments was selected to gather amplitude statistics about sea clutter and growler targets. The second set of experiments was selected to make measurements of growler and neighboring sea clutter for the purpose of Doppler analysis. Thus, for each set, certain radar parameters were appropriate.

The parameters for the first set of experiments are given in Table 3.1 and Table 3.2. Table 3.1 gives the experiment ID associated with each file, the date the experiment was conducted, the target type, transmitted polarization, PRF, pulse width, the starting range position and range swath, and the azimuth position of the antenna. Table 3.2 gives the wind speed, wind bearing and the significant wave height at the time of collection.

The low PRF in these experiments allowed for the continuous collection of data for

ID	DATE	TYPE	POL	PRF	PULSE	RSTART	RSWATH	AZ
	m/d/y	c/g	h/v	(Hz)	(ns)	(m)	(m)	(deg)
B100	6/7/89	g	h	200	200	6030	200	116.0
B122	6/7/89	g	v	200	200	4440	160	67.0
B195	6/9/89	c	h	200	200	6000	160	30.0
B196	6/9/89	c	h	200	200	6000	160	75.0
B197	6/9/89	c	h	200	200	6000	160	120.0
B198	6/9/89	c	v	200	200	6000	160	120.0
B199	6/9/89	c	v	200	200	6000	160	30.0
B200	6/9/89	c	v	200	200	6000	160	75.0
B202	6/9/89	c	h	200	1000	8400	160	30.0
B203	6/9/89	c	h	200	1000	8400	160	75.0
B204	6/9/89	c	h	200	1000	8400	160	120.0
B205	6/9/89	c	v	200	1000	7050	160	120.0
B206	6/9/89	c	v	200	1000	7050	160	30.0
B207	6/9/89	c	v	200	1000	7050	160	75.0
B208	6/9/89	c	h	200	30	6000	80	30.0
B209	6/9/89	c	h	200	30	6000	80	75.0
B210	6/9/89	c	h	200	30	6000	80	120.0
B211	6/9/89	c	v	200	30	6000	80	120.0
B212	6/9/89	c	v	200	30	6000	80	30.0
B213	6/9/89	c	v	200	30	6000	80	75.0

Table 3.1: Radar parameters associated with experiments performed for noncoherent analysis of sea clutter and growler returns. Under TYPE, c = clutter and g = growler. PULSE refers to the pulse width, RSTART refers to the distance from the radar to the first sample collected, and RSWATH refers to the length of the range window that was sampled. POL refers to the transmit polarization, h or v, and AZ is the azimuth angle in degrees.

ID	WIND SPEED (knots)	WIND DIRECTION (degrees)	WAVE HEIGHT (metres)
B100	8	240	2.04
B122	21	170	1.50
B195	23	10	2.70
B196			3.20
B197	22	350	3.61
B198			3.10
B199			2.67
B200			2.89
B202	26	350	2.93
B203			2.70
B204			2.60
B205			3.08
B206			
B207	26	350	
B208			2.40
B209	25	350	
B210			3.09
B211			2.87
B212			2.73
B213	19	340	2.35

Table 3.2: Weather and significant waveheight data associated with experiments performed for noncoherent analysis of sea clutter and growler returns. The weather data was provided by the Atmospheric Environment Service and the waveheight data was provided by a nondirectional waverider.

ID	DATE	TYPE	POL	PRF	PULSE	RSTART	RSWATH	AZ
	m/d/y	c/g	h/v	(Hz)	(ns)	(m)	(m)	(deg)
B97	6/7/89	g	h	2000	200	6570	200	118.5
B98	6/7/89	g	h/v	2000	200	6420	200	119.0
B99	6/7/89	g	v	2000	200	6300	200	118.0
B110	6/7/89	g	h	2000	200	4170	200	86.0
B111	6/7/89	g	v	2000	200	4140	200	85.0
B112	6/7/89	g	h/v	2000	200	4080	200	81.0
B113	6/7/89	g	h	2000	200	4020	200	71.0
B114	6/7/89	g	v	2000	200	4050	200	69.5
B115	6/7/89	g	h/v	2000	200	4080	200	67.0
B123	6/7/89	g	h	1000	200	4500	150	63.5
B124	6/7/89	g	v	2000	200	4530	150	61.5
B125	6/7/89	g	h/v	2000	200	4530	150	59.5

Table 3.3: Radar parameters associated with experiments performed for coherent analysis of sea clutter and growler returns. Under TYPE, g = growler. PULSE refers to the pulse width, RSTART refers to the distance from the radar to the first sample collected, and RSWATH refers to the length of the range window that was sampled. POL refers to the transmit polarization, h or v, while AZ is the azimuth angle in degrees.

periods of about seven minutes at a time. Since amplitude statistics of sea clutter were sought from this data, the low PRF provided a maximum number of independent samples to be collected. The analysis of this data will be described in detail in Chapter 4.

Tables 3.3 and 3.4 provide similar data for the experiments that were performed with coherent analysis in mind. For these experiments, the pulse width was kept at 200 ns which was the smallest real pulse we had available. Furthermore, a PRF of at least 1 kHz in each of the polarization channels was used to allow for accurate estimation of sea clutter and growler spectra. The results of this analysis are given in Chapter 5.

ID	WIND SPEED (knots)	WIND DIRECTION (degrees)	WAVE HEIGHT (metres)
B97	8	240	
B98			1.40
B99			1.80
B110	20	160	1.57
B111			1.57
B112	18	170	1.46
B113			1.38
B114			
B115			1.50
B123			
B124			
B125	20	170	1.50

Table 3.4: Weather and significant waveheight data associated with experiments performed for coherent analysis of sea clutter and growler returns. The weather data was provided by the AES and the waveheight data was provided by a nondirectional waverider. Note: the wind speed and direction recorded for B110 and B125 were actually measured shortly before and after their respective experiments.

3.3.2 Surface Truthing

Surface truthing for each of the experiments was provided in a variety of forms. Weather reports were provided from the Atmospheric Environment Service's weather office in Cape Bonavista. These reports include wind speed and bearing measurements, gusts, temperature, pressure, visibility charts etc. Profiles of these measurements throughout the June 1989 field trip can be found in Appendix A. The entries given in Table 3.2 and Table 3.4 were taken from these AES weather charts. AES makes measurements periodically throughout the day, and for this reason, some fields in the tables are empty. In addition to AES's weather records, a local weather station installed at the radar site was used to provide instantaneous measurements. These were particularly useful when experiments were being conducted through the night.

Wave height information is important for correlating changes in sea clutter statistics with wave development. A nondirectional waverider was moored 6.75 km off the coastline at a bearing of 72 degrees in order to provide wave height and wave period measurements. The waverider sampled the sea height at a rate of 2.56 Hz and relayed the information back to a local receiving station at the site. Instantaneous wave height, wave period and fluctuation spectra were available. An example of the reports generated from the waverider system is shown in Figure 3.2. The measurements given in Tables 3.2 and 3.4 were extracted from these reports.

Along with weather and wave reports, photographic recordings were also made of the experimental setup and the targets that were sampled. Live video was taken of the ocean several times each day, as well as still photographs. Photographs were also taken of ice targets so their approximate sizes could be determined. The I and Q channels were also displayed on an oscilloscope and recorded. Finally, PPI images of the local scene were also videotaped.

Finally, with each experiment performed, a computer-generated log file was stored along with the data. The log file has all of the radar parameter and switch settings stored,

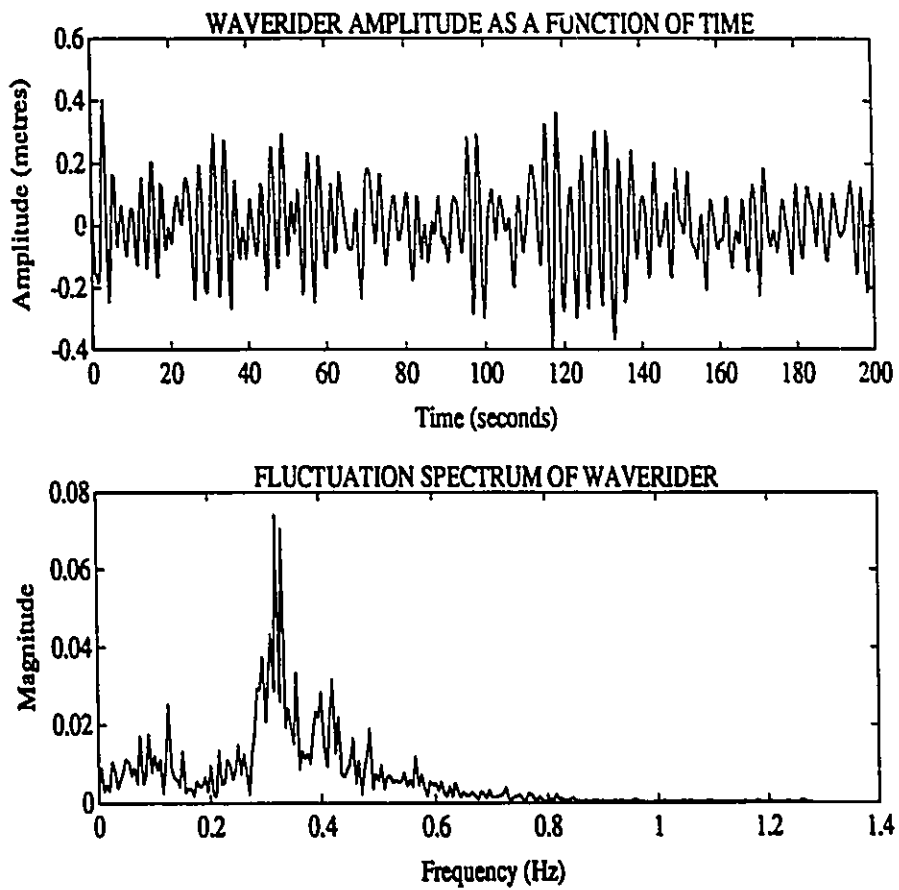


Figure 3.2: Waverider wave height measurements and wave fluctuation spectrum. Height is given in metres and time is given in seconds. Winds were 20-40 knots.

so that the complete state of the radar at the time of data collection can be determined. Also, the sampling system writes along with each sweep of data, a header that indicates the experimental ID, the antenna position and polarization, the range swath, and an absolute reference clock that can be used to derive timing information. Since the header is part of the data, this information can never be scrambled or lost.

3.4 Calibration

The purpose of a research radar is to make quantitative measurements on targets and the surrounding environment. It is important therefore that the measurements describe the characteristics of the targets being examined, free from contaminations from the receiver itself. The I and Q output signals that are collected by the IPIX radar must be mapped back to the input signals that were actually received from the target. This process we refer to as *calibration*. Clearly, any conclusions that can be drawn about the illuminated targets must take into account the input/output characteristics of the receiver performing the measurements.

We have divided the calibration process for the IPIX radar into two stages. The first stage describes the corrections that must be made to the I and Q signals themselves, so that they are *balanced* and *orthogonal* as they should be. These corrections are particularly important for spectral processing. Procedures for performing these I and Q corrections are given in Appendix B. The second stage of calibration then involves mapping the output power described by the corrected I and Q signal components to the proper input power level actually received. This procedure requires an input/output power characteristic for the receiver and is also described in Appendix B.

3.5 Summary

In this chapter, we have taken a detailed look at the IPIX radar that was used to collect sea clutter and growler data for our study. We described the features of IPIX that provide multiparameter measurements of the targets of interest, in a way that is appropriate for research. We then summarized the radar parameters associated with the experiments described in this thesis. The experiments were divided into two sets: one set for noncoherent analysis and the other for coherent analysis. The results of these experiments are described in Chapter 4 and Chapter 5, respectively. The various forms of surface truthing were also discussed. These include weather and waverider data, photography and video, along with computer-based logs to retain the state of the radar at the time an experiment is conducted. Finally we discussed the data calibration problem in the context of the IPIX radar. These calibrations are essential to remove radar system related effects from the recorded data.

Chapter 4

Noncoherent Analysis of Sea Clutter and Growler Data

4.1 Introduction

In this chapter, we begin our analysis of real IPIX data by looking at the properties of the amplitude of the return radar echoes for sea clutter and growlers. The amplitude is formed by taking the square root of the sum of the squares of I and Q. We begin by examining the radar equation as it applies to point targets such as growlers, and to distributed targets such as clutter, in order to develop power law relationships. The power law describes how the power returned from the target varies as a function of range to the target. We then use these laws to estimate the radar cross-section of some of the sea clutter and growlers that we studied. Then, we look at the amplitude statistics of sea clutter. In particular, we examine the usefulness of the compound K-distribution in describing these amplitude statistics. Next, we consider amplitude statistics of growlers in sea clutter. Finally, we look at temporal and spatial correlations of sea clutter data. This will provide the insight necessary to predict the performance of various noncoherent detection algorithms.

4.2 The Radar Equation

The radar equation is fundamental to our understanding of radar. It relates the received signal power to the transmitter, antenna, target and environmental characteristics. Its derivation is given in [20]. The simple form of the radar equation (this form doesn't explicitly model system losses, target cross section variability and meteorological conditions) is given by

$$P_g = \frac{P_t G A_e \sigma}{(4\pi)^2 R^4}, \quad (4.1)$$

where P_g is the received echo signal power in Watts, G is the antenna gain factor, A_e is the effective area of the receiving antenna in square meters, σ is the radar cross section of the target in square meters, and R is the range to the target in meters. If the same antenna is used for transmission and reception, then G and A_e are related by

$$G = \frac{4\pi A_e}{\lambda^2}. \quad (4.2)$$

where λ is the operational wavelength.

For a given set of radar parameters, Equation 4.1 can be used to determine the maximum range of the radar, the minimum detectable signal strength, the power law relating received signal power to range, and target cross section models. In the next two sections, we investigate the latter two.

4.2.1 Power Law

As described earlier, the two classes of targets that we are interested in are growlers and sea clutter. Growlers can be considered as point targets, and sea clutter as distributed targets. In order to make our measurements independent of the range of the target from the radar, the power law as a function of range needs to be determined. This allows for correction of the received power, and in many radar systems, this correction is implemented by means of Sensitivity Time Control (STC). Although the IPIX radar has STC capabilities, STC was not used in the data described in this thesis.

The radar equation given in Equation 4.1 was derived for a point target with cross section σ , a distance R from the radar. Therefore the received power from the point target, P_r is inversely related to the fourth power of range. This power law relationship is suitable for describing the received power from a growler sampled with the IPIX radar.

For distributed targets such as sea clutter however, the power law is a little more complicated. For a radar pulse of duration τ seconds, the range resolution of the pulse is $c\tau/2$ radial meters (ie. in the direction of the radar pulse) where c is the speed of light ($3e8$ m/s). Therefore, the radar is not capable of resolving two point targets that are separated by less than $c\tau/2$ radial meters when sampling the matched filter output every τ seconds. Stated in another way, a receiver output at a particular time t_o and hence a particular range R_o is actually the result of the superposition of echoes from all scatterers within a radial distance of $c\tau/2$ of R_o . Since the target is distributed in azimuth as well as range, a similar effect results in angle. The return echo at a given time is the result of the superposition of all the scatterers within the antenna beamwidth θ_B . Clearly, the target cross section σ in Equation 4.1 can now be seen to be a function of both the pulse width τ and the antenna beamwidth θ_B .

Consider a cliff-top radar such as the IPIX radar, that illuminates the sea surface with a depression angle of ϕ radians. The depression angle is the angle between the horizontal and the line of sight of the radar. Then, for a pulse width of τ seconds, the radial range resolution is $c\tau/2$ meters. Thus, the range resolution along the sea surface is $c\tau/2 \sec \phi$ meters. In azimuth, the surface extent of the scatterers that contribute to the received signal at a given time is $R\theta_B$. Therefore, the clutter patch area or resolution cell size is given approximately by

$$A_c = R\theta_B(c\tau/2) \sec \phi. \quad (4.3)$$

Since for distributed surface targets such as sea clutter, the received signal power is a function of the resolution cell size given by A_c , a target cross section density would be more descriptive than σ . Therefore, we define $\sigma^o = \sigma/A_c$ as the radar target cross section density

for surface clutter targets. Substitution of σ into Equation 4.1 gives the surface clutter radar equation:

$$P_c = \frac{P_t G A_e \sigma^\circ \theta_B (c\tau/2) \sec \phi}{(4\pi)^2 R^3} \quad (4.4)$$

where P_c is the received clutter power. We see from Equation 4.4 that the clutter received power now varies inversely as the third power of range as predicted by the radar equation.

In order to test the power versus range law described above for sea clutter, we evaluated this law empirically. A 2.5 km patch of the sea surface was illuminated using a pulse width of 200 ns. The seas were choppy with winds gusting from 20–35 knots. The transmit polarization was horizontal so HH and HV were collected. The received power was calculated at each range cell and averaged over a period of 0.5 s. Then, the best fit power law relationship was determined from the data using a least square regression. The results are given in Figure 4.1 for both the HH and HV channels. The curves indicate a power versus range law with an exponent of about -3.5, somewhere between the point target law predicted in Equation 4.1 and the surface clutter law predicted in Equation 4.4. Also note the separation of about 10 dB between the like-polarized and cross-polarized channels. The range law exponent of -3.5 suggests the presence of scatterer bunching in this instance. The grazing angle varies only by about 0.1 degrees over the 2.5 km range and hence should not affect the measured σ° or the power versus range law.

4.2.2 Radar Cross-Section

In the previous section, we derived the radar equations that are appropriate for point targets such as growlers, and for distributed targets such as sea clutter. Now we isolate the radar cross section σ and the radar cross section density σ° for point targets and distributed targets respectively. This will be useful for describing the target size.

Isolating σ from Equation 4.1 and making use of Equation 4.2, we get

$$\sigma = \frac{P_g (4\pi)^3 R^4}{P_t G^2 \lambda^2} \quad (4.5)$$

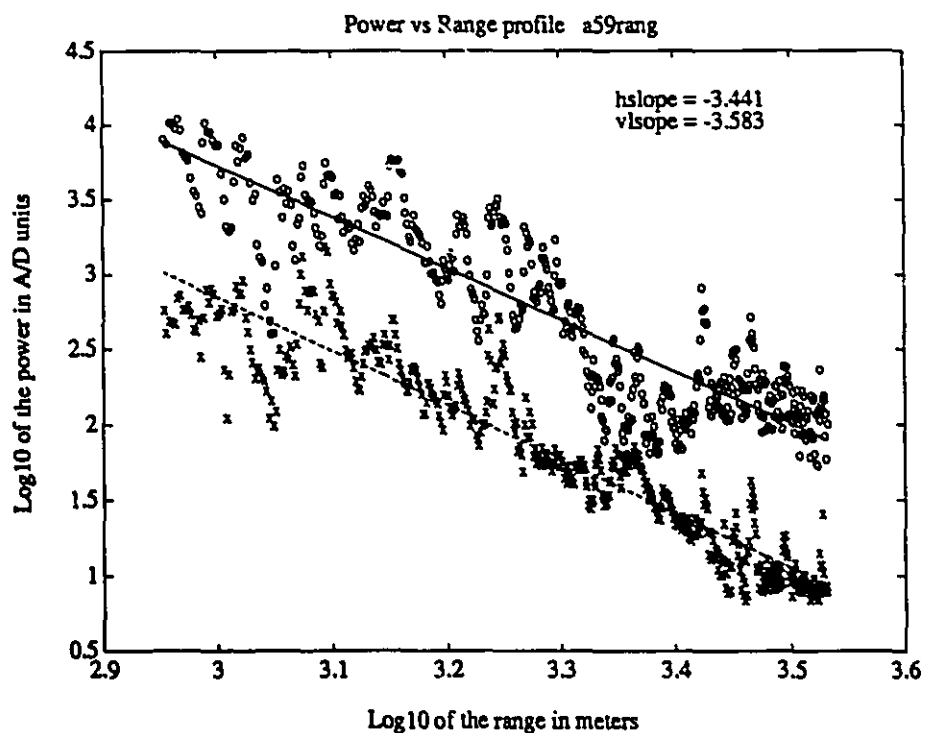


Figure 4.1: This figure describes the power versus range variation of sea clutter data. The data marked with 'o' is from the HH channel and the data marked with 'x' is from the HV channel. The range scale represents 2.5 km. The pulse width was 200 ns and PRF was 2 kHz. Each point represents 1000 averages.

for a point target's radar cross section. Similarly, using Equation 4.4, we get

$$\sigma^o = \frac{P_c(4\pi)^3 R^3 \cos \phi}{P_t G^2 \lambda^2 \theta_B \tau / 2} \quad (4.6)$$

for a distributed target's radar cross section density.

The IPIX radar has the following fixed parameter values:

- Peak Power $P_t = 8$ kW
- Antenna Gain $G = 35$ dB
- Operational Wavelength $\lambda = 3.19$ cm
- Azimuthal Beamwidth $\theta_B = 2.2$ degrees

These parameters will be used in Equation 4.5 and Equation 4.6 along with the pulse width τ in order to estimate radar cross sections for growlers and sea clutter, respectively. It is necessary to estimate the growler received power P_g and the clutter received power P_c in order to use these equations. These estimates must be determined from the received data and will be dealt with in detail in Chapter 5.

4.3 Amplitude Statistics

In this section, we investigate the use of the K-distribution to describe the amplitude statistics of sea clutter data collected with the IPIX radar. We show that the K-distribution is suitable for modeling the amplitude statistics of both like-polarized and cross-polarized configurations. The amplitude statistics of small ice targets (growlers) in spiky clutter are also examined, and our results indicate that using only long-term amplitude statistics is not sufficient for the reliable detection of such targets in the presence of sea clutter. Furthermore, we show that the cross-polarized channel does not offer any added benefit to small target detection, when considering only long-term amplitude statistics.

4.3.1 The Compound K-Distribution Model

Many ad hoc distributions have been proposed in the literature to model the amplitude statistics of spiky sea clutter. Although some of these models can be adequately fitted to the empirical amplitude statistics, they do not describe the temporal or spatial correlation in the data. As a result, it is difficult to accurately predict the detection probabilities when integration or CFAR processing is used. This problem is overcome with the K-distribution [45], defined as

$$p(\sigma) = \frac{b}{\sqrt{\sigma}\Gamma(v)}(b\sqrt{\sigma}/2)^v K_{v-1}(b\sqrt{\sigma}) \quad (4.7)$$

where Γ is the Gamma function, and K_{v-1} is the second order modified Bessel function. Equation 4.7 describes the statistics of the radar cross section of clutter, σ . In this thesis, we deal with amplitude statistics, and therefore adopt the form of the K-distribution given in [46, 47, 48]. This form,

$$p(x) = \frac{4c}{\Gamma(v)}(cx)^v K_{v-1}(2cx) \quad (4.8)$$

is related to Equation 4.7 by the transformation $x = \sqrt{\sigma}/2$; it describes the statistics of the clutter amplitude, x . In both cases, v is the so called shape parameter, and b and c represent an identical scale parameter (ie. $b = c$). The K-distribution is therefore a two-parameter distribution whose n^{th} amplitude moments are given by

$$m_n = E[x^n] = \frac{\Gamma(v + n/2)\Gamma(n/2 + 1)}{c^n \Gamma(v)} \quad (4.9)$$

where E is the expectation operator. The n^{th} normalized amplitude moments are given by

$$M_n = \frac{E[x^n]}{(E[x])^n}. \quad (4.10)$$

The K-distribution is based on an underlying physical model that treats the received signal as a superposition of returns from a number of independent patches or scatterers illuminated by the radar beam. The effective number of scatterers along with their relative bunching are critical in determining the overall statistics of the received data. By the Central Limit Theorem, as the number of scatterers becomes large, the complex clutter

process approaches a Gaussian form, and amplitude statistics therefore become increasingly Rayleigh. On the other hand, as the effective number of scatterers decreases and bunching among scatterers occur, the statistics become less Rayleigh-like. The amplitude statistics of sea clutter behave in a similar manner. As the radar resolution increases and the illuminated area becomes smaller (fewer scatterers), the sea clutter amplitude statistics become increasingly non-Rayleigh.

Figure 4.2 shows several K-distributions with the same shape parameter, $v = 1.2$, and scale parameter c that takes on values 0.1 to 1.1 in 0.2 increments. As the name implies, the parameter c scales the distribution. Figure 4.3 shows another set of K-distributions where the scale parameter c is the same for all of the curves, $c = 1.0$, but the shape parameter v varies from 0.6 to 5.6 in unit increments. If each K-distribution is normalized to its mean as is done for example in [45], then Figure 4.4 and Figure 4.5 result from Figure 4.2 and Figure 4.3, respectively. This normalization procedure allows various distributions to be compared, based only on their shape parameter v . In Figure 4.4, the curves redrawn from Figure 4.2 appear identical since these curves all have the same shape parameter v . In Figure 4.5, the effect of reducing the shape parameter is evident: the distributions have longer tails. It is easy to show that when the shape parameter v is equal to infinity, the K-distribution reduces to a Rayleigh distribution, and when v approaches zero, the tails of the K-distribution grow increasingly longer. Traditionally, Rayleigh distributions have been used to model speckle-like clutter and lognormal distributions have been used for modeling spiky (long tails) clutter. High resolution sea clutter is generally observed to have a shape parameter that lies in the region $0.1 \leq v \leq \infty$, where very spiky clutter has $v \simeq 0.1$ [49]. Amplitude histogram data presented in this thesis make use of the normalization procedure described above; only they will be plotted on a semilogarithmic scale as shown in Figure 4.6. Figure 4.6 shows a set of normalized K-distributions with a different shape parameters plotted in this fashion.

Perhaps the most important characteristic of the K-distribution is its ability to

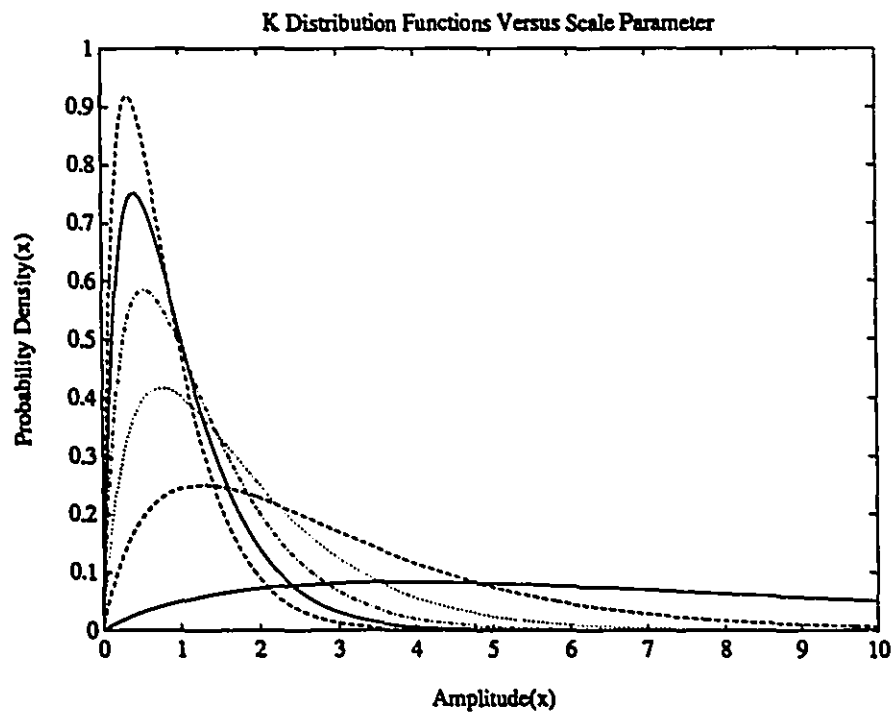


Figure 4.2: Various K-distributions for shape parameter $\nu = 1.2$ and scale parameter $c = 0.1, 0.3, 0.5, 0.7, 0.9$ and 1.1 . The curve with the largest peak has $c = 1.1$.

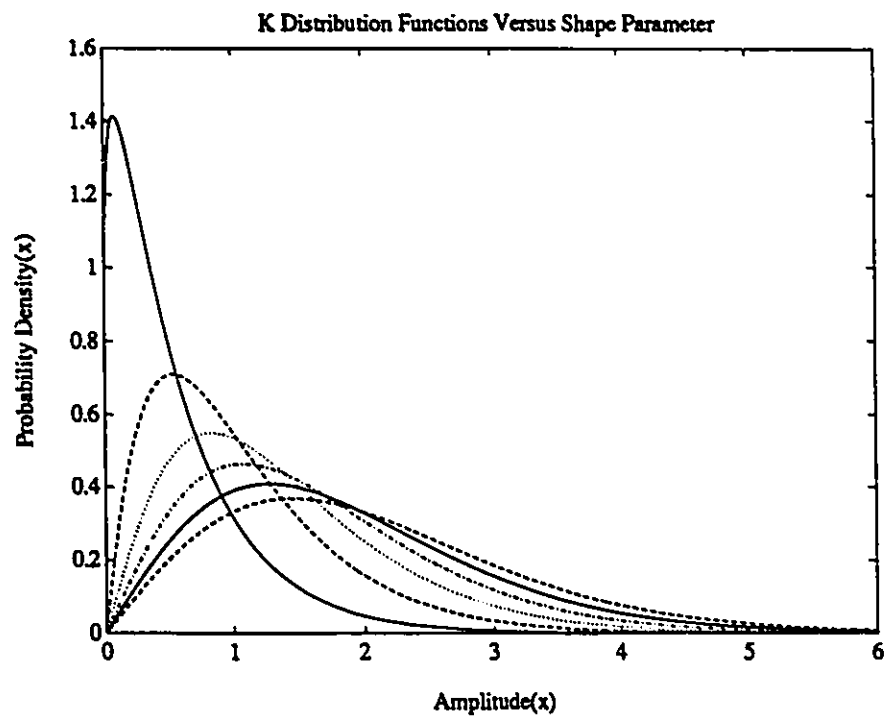


Figure 4.3: Various K-distributions for scale parameter $c = 1.0$ and shape parameter $v = 0.6, 1.6, 2.6, 3.6, 4.6$ and 5.6 . The curve with the largest peak has $v = 0.6$.

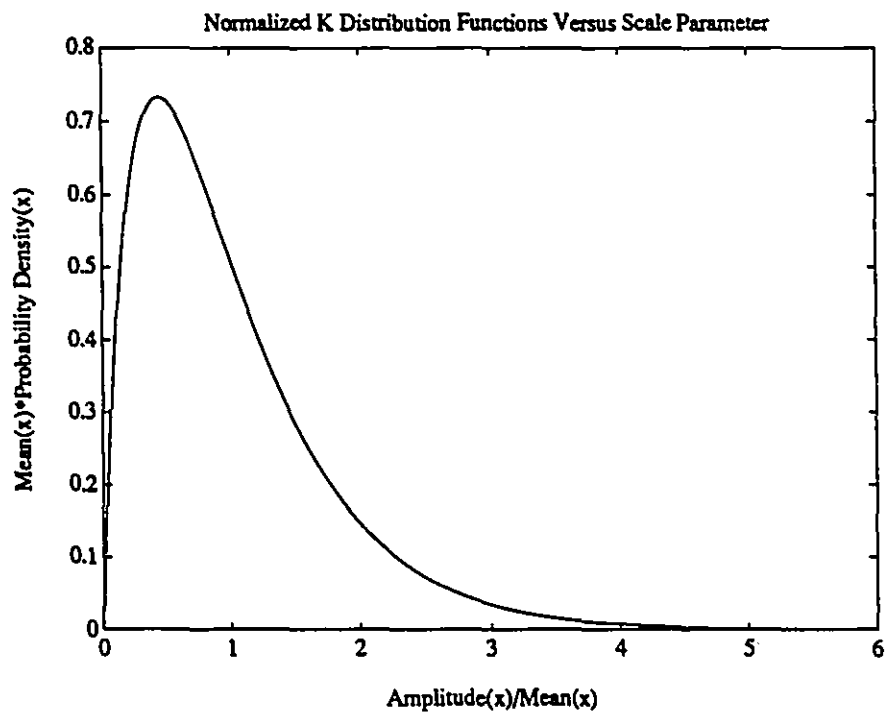


Figure 4.4: Various K-distributions for shape parameter $\nu = 1.2$ and scale parameter $c = 0.1, 0.3, 0.5, 0.7, 0.9$ and 1.1 . Each curve is normalized to its mean. All curves are identical. This is why c is called a scale parameter.

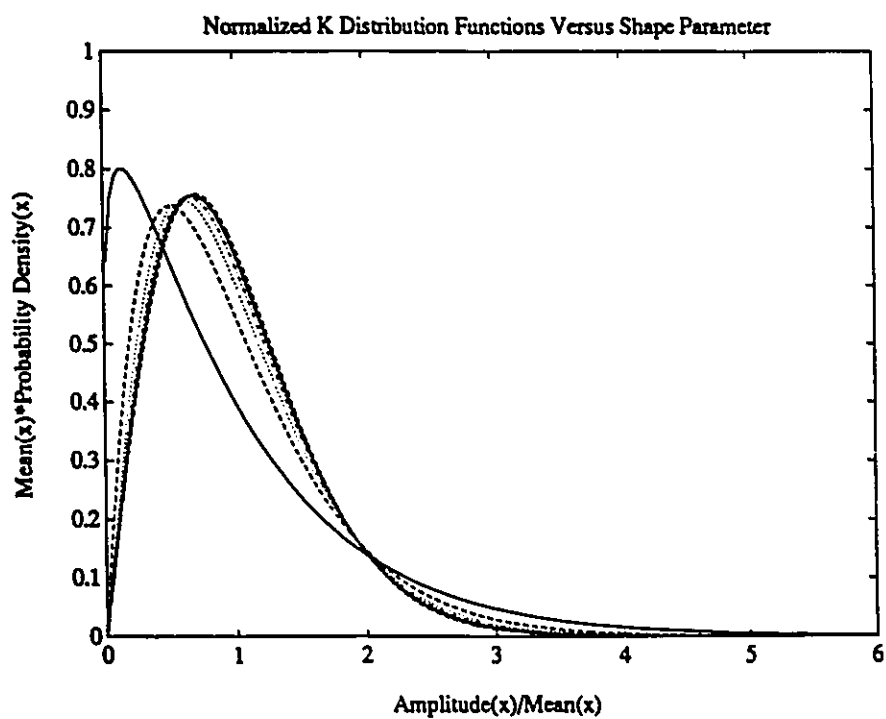


Figure 4.5: Various K-distributions for scale parameter $c = 1.0$ and shape parameter $v = 0.6, 1.6, 2.6, 3.6, 4.6$ and 5.6 . Each curve is normalized to its mean. The curve with the longest tail has $v = 0.6$.

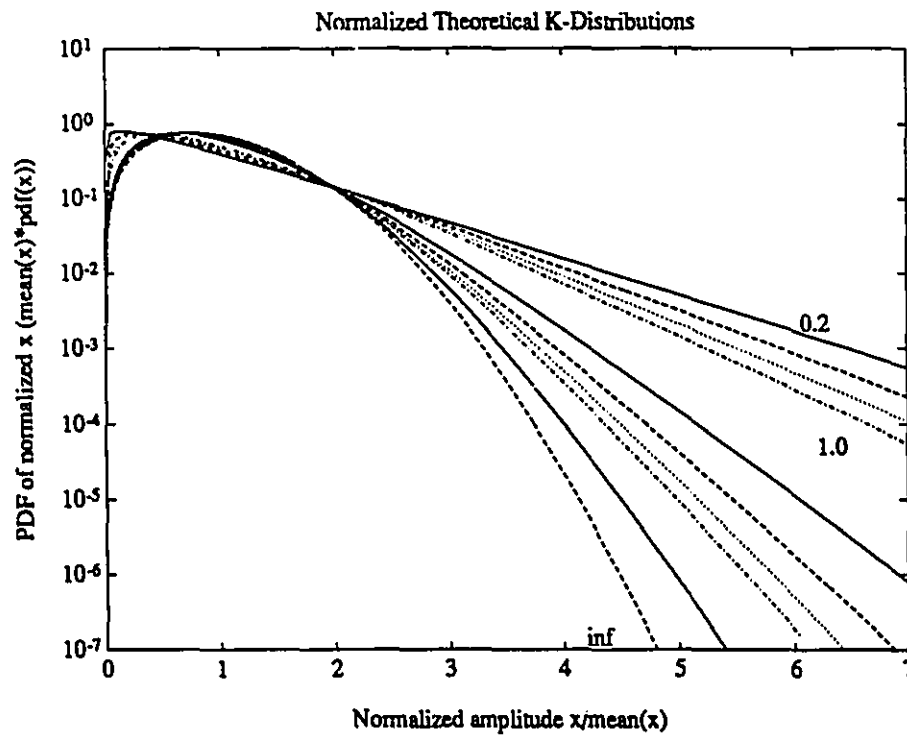


Figure 4.6: A family of K-distributions with shape parameters varying from $v = 0.20, \dots, \infty$. Each curve is normalized to its mean. The probability axis is logarithmic to accommodate the large dynamic range of probabilities encountered. The curves with shape parameters 0.2, 1.0, and ∞ are labelled. Note: there is a singularity at the origin for $v < 0.5$.

take into account the correlation properties of sea echos. This ability is a result of the fact that the K-distribution is a compound distribution, made up of a Rayleigh, speckle-like component $p(x|y)$, whose mean value y is the result of a modulating underlying (swell) component that has a root-Gamma distribution. That is,

$$p(x|y) = \frac{\pi x}{2y^2} \exp\left(-\frac{\pi x^2}{4y^2}\right) \quad (4.11)$$

and

$$p(y) = \frac{2b^{2v}y^{2v-1}}{\Gamma(v)} \exp(-b^2y^2), \quad (4.12)$$

where both x and y are bounded by $0 \leq x, y < \infty$. Now the amplitude distribution, $p(x)$ as described in Equation 4.8 is derived from Equation 4.11 and Equation 4.12 by the relation,

$$p(x) = \int_0^\infty p(x|y)p(y) dy. \quad (4.13)$$

The correlation properties are accounted for because each component can have a different decorrelation time. For example, the Rayleigh component of X-band sea clutter typically decorrelates in about 10ms, whereas the underlying modulating component that is attributed to the underlying sea swell can have decorrelation times on the order of several seconds. Watts [46, 47] evaluates the predicted radar detection performance under several target models, and compares these to the performance obtained when the correlation properties of sea clutter are ignored. The effect of ignoring the clutter correlation properties is significant. In the next section, we consider real sea clutter measurements made with the IPIX radar.

4.3.2 Sea Clutter Empirical Amplitude Distributions

It is well known that the K-distribution introduced by Jakeman and Pusey [45, 50] and the compound form of the K-distribution developed by Ward [6] provide a good model for the amplitude statistics of sea clutter. This distribution has been successfully applied to real, high resolution radar data for linear polarization [6, 49, 46, 51, 52]. However,

no experimental results appear to have been reported for the cross-polarized channel. In this section, amplitude statistics of sea clutter for both the like-polarized (HH and VV) and cross-polarized (HV and VH) channels are presented. Furthermore, we show that the K-distribution provides an accurate mathematical description of our measurements.

The data presented in this section were obtained during field trials conducted with the IPIX radar at a site located in Cape Bonavista, Newfoundland in June 1989. The parameters for most of the data sets are described in Table 3.1 and Table 3.2. Some data that we present in this section are not described in Table 3.1 and Table 3.2 and so their parameters will be given where appropriate. We begin by first examining the compound nature of our clutter data; then we analyze the amplitude statistics under a variety of conditions.

The data sets that we have analyzed indicate a compound clutter model such as the model the K-distribution provides, and we now present some typical results. The data set that is presented was collected with a PRF of 2kHz and a pulse width of 200 ns. The transmit polarization was horizontal, and a range window of 160 m (33 samples) was used. There were gale conditions with winds of 30 knots from a bearing of 120 degrees and the swell was from the same direction. The significant wave height was about 3 m and the wave period was about 7 s, as reported from the waverider. The antenna was fixed at 120 degrees so the aspect was upswell. About 6dB of RF attenuation was applied to the horizontal receive channel to avoid significant saturation, while maintaining sensitivity for the cross-polarized HV channel. Analysis has indicated the underlying swell component to be strongly correlated over a period of 0.25s. Therefore amplitude histograms were performed at several range positions over the 33-sample range window. For each range position 500 sweeps (0.25 seconds at 2kHz PRF) were taken to calculate the amplitude histogram. The mean, and normalized moments M_2 , M_3 , and M_4 were calculated for each range position; the results are presented in Table 4.1. Both like-polarized HH and cross-polarized HV data are presented. The theoretical normalized amplitude moments for a

File	Mean		M2 = 1.27		M3 = 1.91		M4 = 3.24	
	HH	HV	HH	HV	HH	HV	HH	HV
B347r18.dat	12.10	11.57	1.28	1.25	1.93	1.83	3.32	3.02
B347r21.dat	8.42	10.99	1.27	1.26	1.92	1.84	3.25	2.99
B347r24.dat	26.06	19.43	1.33	1.31	2.15	2.03	4.00	3.57
B347r27.dat	23.91	18.23	1.33	1.28	2.18	1.94	4.12	3.27
B347r30.dat	13.34	11.72	1.30	1.28	2.04	1.95	3.67	3.36
B347r33.dat	6.28	9.05	1.35	1.28	2.24	1.94	4.36	3.27
B347r37.dat	10.58	11.99	1.28	1.31	1.97	2.10	3.43	4.02
B347r40.dat	38.66	30.24	1.27	1.33	1.89	2.17	3.18	4.09

Table 4.1: Table of mean amplitude values and normalized amplitude moments calculated from sea clutter data by including only 0.25 seconds worth of data. These moments are calculated for 8 different range positions spanning a distance of about 110 m. This data set serves to expose the Rayleigh speckle component of sea clutter that is present at every range position, but the mean of the distribution varies as a function of range. Theoretical values for Rayleigh moments are indicated in the Table for comparison.

Rayleigh distribution are also shown for reference. The results indicate that the speckle component is indeed Rayleigh distributed, and that the underlying component, indicated by the mean, varies with range.

For the same data set, the underlying component was calculated following the procedure given in [52]. A window of 500 sweeps were averaged in amplitude to give an estimate of the underlying component. This procedure was continued throughout the data set, resulting in 160 averaged sweeps of the underlying component. This represents a duration of 40 s. The resulting like- and cross-polarized images are presented in Figure 4.7. The strongly correlated temporal/spatial swell structure of the ocean surface is evident. The data are scaled to span the grey-scale code (0-255). The image indicates a wave period of about 7.5 seconds which is in agreement with that reported from the nondirectional wavereider that was deployed. The white areas of the image indicate regions of high radar reflectivity, while the black areas indicate regions of little or no reflectivity.

It is worth while mentioning that the sea surface is not always as structured as




Figure 4.7: The underlying, mean, sea clutter backscatter of file B347. The figures show the well developed swell structure of the sea. (a) The HH polarization returns. (b) The HV polarization returns.

indicated in Figure 4.7. The structured sea surface shown in Figure 4.7 is the result of strong winds persisting for many hours over a long fetch of the ocean. The sea surface is often confused, and the well developed crests and troughs are not apparent. For example, consider the data shown in Figure 4.8. These images were created in an identical fashion to those in Figure 4.7, from another sea clutter data set. The point to be aware of is that although there is no well developed structure visible here, the clutter is still correlated in time for periods on the order of seconds.

These images emphasize the point that temporal and spatial correlations must be taken into account in order to properly determine the performance of a given signal processor. Now that the compound nature of sea clutter backscatter has been verified with this data, a detailed look at the overall statistical behaviour of the clutter amplitudes is in order. These results are now presented.

We consider amplitude statistics of sea clutter as a function of polarization, pulse width and look-direction. As a reference for the measurements taken, data sets were collected with the travelling-wave tube (TWT) amplifier disabled. In this way, the statistics of front-end receiver noise can be characterized separately from sea clutter and target returns. These noise-only files allow for an accurate measurement of receiver noise levels, which can be used to correct for noise contaminating sea clutter data sets. The amplitude distribution of a noise-only data set is shown in Figure 4.9. The statistics are distinctively Rayleigh, indicating the Gaussian nature of receiver noise in both the inphase and quadrature channels.

To characterize sea clutter, measurements pertaining to three different pulse widths, namely 32ns, 200ns and 1000ns, at bearings of 30, 75, and 120 degrees were taken and the PRF was fixed at 200Hz. For each case, vertical and horizontal polarizations were transmitted separately, while both vertical and horizontal polarizations were received simultaneously. In this way, data were collected for both like (HH and VV) and cross (HV and VH) polarizations. The sea clutter data described herein was collected at a range of 6 km, in 2.5 to




Figure 4.8: The underlying, mean, sea clutter backscatter of file B319. The figures show a choppy structure of the sea surface. (a) The HH polarization returns. (b) The HV polarization returns.




Figure 4.9: The amplitude statistics of front-end receiver noise are shown. The noise data is Rayleigh distributed due to the Gaussian noise in the I and Q channels. The solid line represents a theoretical Rayleigh distribution ($\nu = \infty$) and the dashed line is a normalized K-distribution with $\nu = 0.25$.

3 metre seas. The wind was typically 24 to 26 knots from the north (0 degrees) and the swell direction was at approximately 110–120 degrees. This set of experiments results in 24 like-polarized and 24 cross-polarized combinations; the data sets are described in detail in Tables 3.1 and 3.2.

Figure 4.10 shows the empirical distributions of one of these data sets. Rather than showing the empirical distributions associated with the other combinations described above, we take a more rigorous approach here by doing a moment analysis of the data in each data set. This is the same procedure done in [14] and similar to that done in [51]. If desired, the empirical distributions for each data set can be found in [13]. A treatment of the confidence on the empirical distributions that we derive is given in Appendix C.

The data sets are subdivided into two groups, those that are like-polarized and those that are cross-polarized. Receiver noise contaminates all data. The like-polarized data have a typical CNR (clutter to noise ratio) of about 12 dB. The cross-polarized data have a much lower CNR. A typical value is 6dB. Although the like-polarized data have reasonably high CNR, the cross-polarized data are contaminated substantially by noise due to the reduced clutter level. For each like-polarized data set, we compute the second, third, and fourth normalized moments, M_2, M_3 and M_4 , as given in Equation 4.10. The third and fourth normalized moments are shown in Figure 4.11, plotted against the second normalized moment. Theoretical third and fourth normalized K moments are also given for comparison, along with the theoretical lognormal moments. The leftmost data moment pair with a second normalized moment of 1.27 corresponds to the data from the noise file in Figure 4.9. Therefore, this data moment pair corresponds to the normalized Rayleigh moments ($v = \infty$). The spikier the clutter, the further along the second moment axis the data moment pairs (third and fourth normalized moments) will fall. The lognormal curves are given for reference since lognormal probability density functions have been used in the literature to describe spiky sea clutter. The like-polarized data sets show good agreement with the K-distribution.

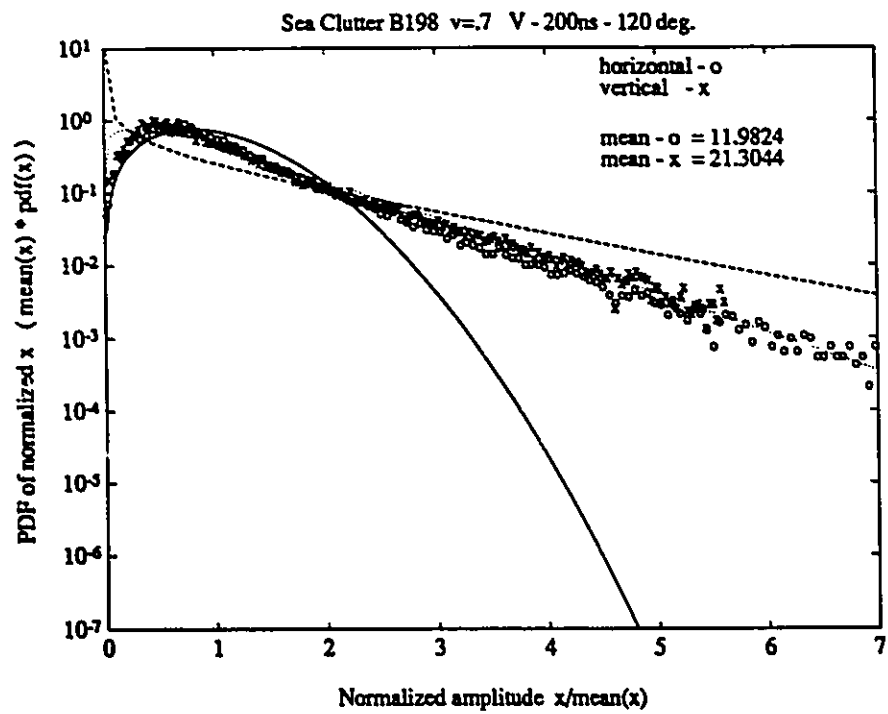


Figure 4.10: A typical empirical amplitude distribution of sea clutter. Polarizations VV and VH are shown. The VV data are indicated with an 'x' and the VH data are indicated with an 'o'.

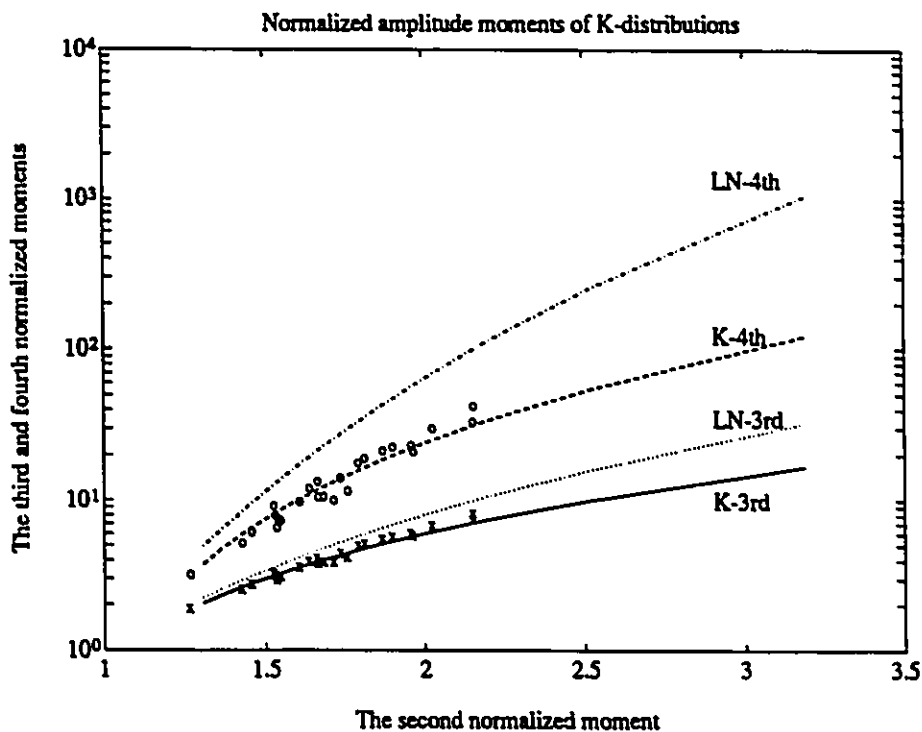


Figure 4.11: Third and fourth normalized amplitude moments for sea clutter data plotted against the second normalized moment. All data on this graph are for HH and VV polarizations. The theoretical noiseless K-moments and lognormal moments are also shown for reference. Third data moments are indicated with an 'x' and fourth moments are indicated with an 'o'.

The normalized cross-polarized moments are shown in Figure 4.12. For this case, the sixth normalized amplitude moments are also given. The effect of receiver noise contamination is evident due to the low CNR. As described in reference [51], low CNR measurements result with the moments given by Equation 4.10 falling between the K and lognormal curves. In [46], Watts derives the K plus noise distribution, and also provides the following unnormalized moments for this case:

$$m_2 = 2\sigma^2 + \frac{4v}{\pi b^2} \quad (4.14)$$

$$m_4 = 8\sigma^4 + \frac{32\sigma^2 v}{\pi b^2} + \frac{32v(v+1)}{\pi^2 b^4} \quad (4.15)$$

$$m_6 = 48\sigma^6 + \frac{288\sigma^4 v}{\pi b^2} + \frac{576\sigma^2 v(v+1)}{\pi^2 b^4} + \frac{384v(v+1)(v+2)}{\pi^3 b^6} \quad (4.16)$$

These equations can be manipulated to give the following:

$$v = 18(m_4 - 2m_2^2)^3 \quad (4.17)$$

and

$$b^2 = \frac{4v}{\pi(m_2 - 2\sigma^2)} \quad (4.18)$$

The cross-polarized data moments m_2 , m_4 , and m_6 are calculated. Then the K plus noise shape parameter v is calculated according to Equation 4.17, along with the scale parameter b from Equation 4.18. The noise-only file described earlier was used to estimate the noise power level $2\sigma^2$, where σ^2 is the noise power level in each of the inphase and quadrature channels. Finally, the theoretical moment m_6 is calculated according to Equation 4.16 and compared to the moment m_6 determined from the data. These values are plotted in Figure 4.13; they indicate that the cross-polarized measurements agree reasonably well with a K plus noise clutter model.

In order to determine relationships with respect to pulse width, look direction and polarization, the shape parameter is calculated for both the like and cross-polarized data sets according to Equation 4.17. These shape parameters are listed in Table 4.2.

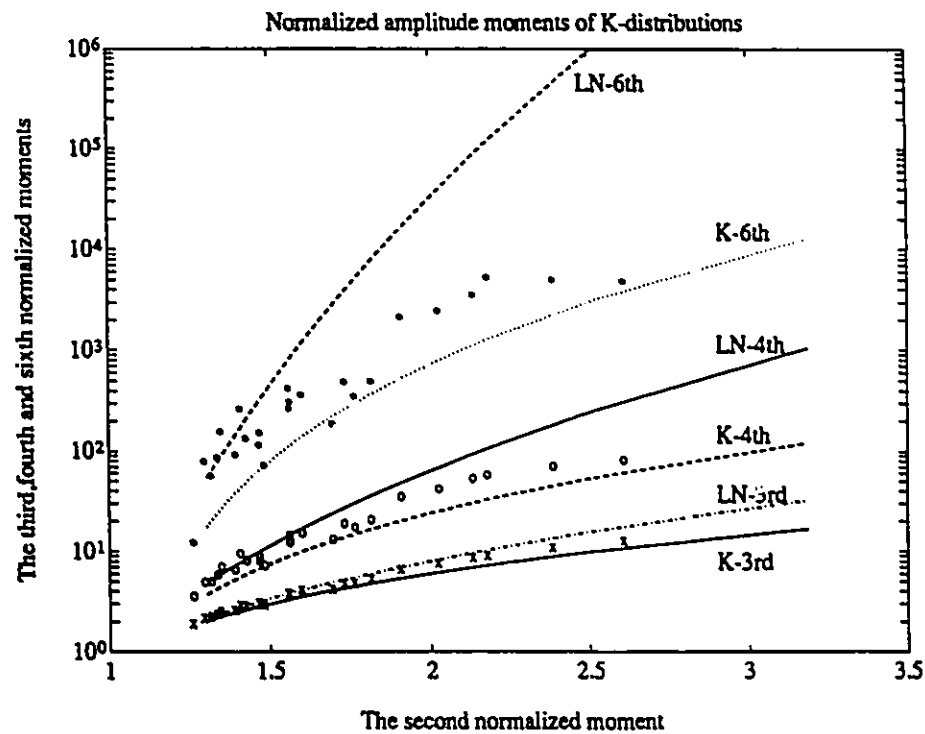


Figure 4.12: Third and fourth and sixth normalized amplitude moments for sea clutter data plotted against the second normalized moment. All data on this graph are for HV and VH polarizations. The theoretical noiseless K-moments and lognormal moments are also shown for reference. Third data moments are indicated with an 'x', the fourth moments are indicated with an 'o', and the sixth moments are indicated with a '*'.

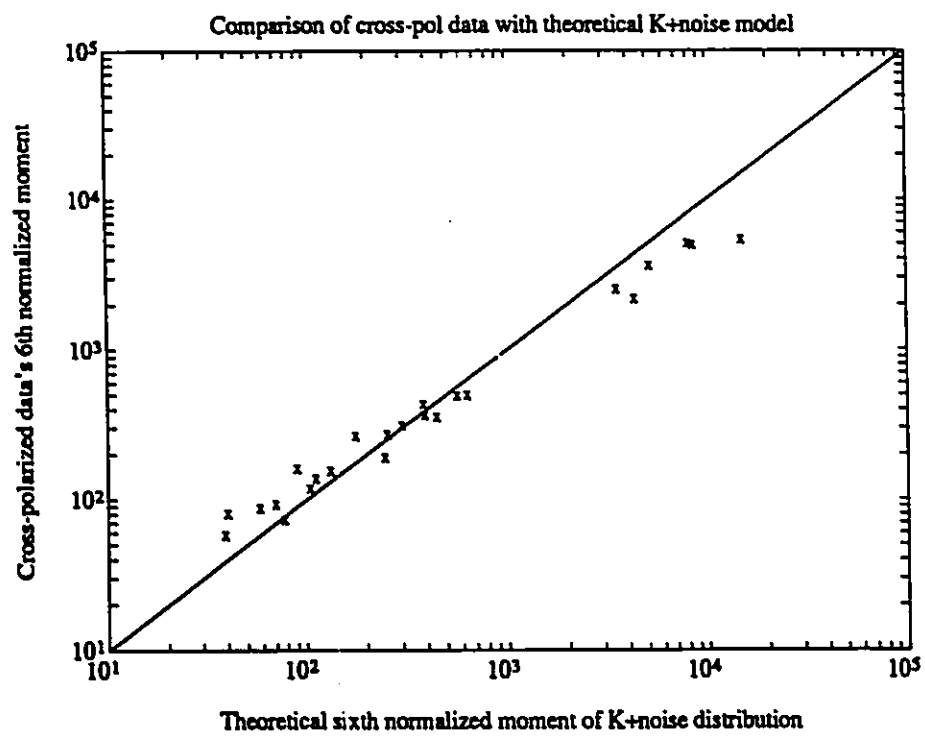


Figure 4.13: Cross-polarized (HV and VH) sea clutter sixth data moments plotted against the corresponding theoretical K plus noise distribution sixth moments.

File H_{TX}, V_{TX}	Pulse Width ns	Azimuth degrees	Shape Parameters				Mean Amplitude	
			HH	VV	HV	VH	HH	VV
B210 , B211	32	30	0.31	0.40	0.12	0.09	17.16	18.57
B208 , B212	32	75	0.29	0.45	0.09	0.09	15.84	18.95
B209 , B213	32	120	0.18	0.37	0.03	0.01	12.23	16.63
B195 , B199	200	30	1.06	1.00	0.37	0.36	29.15	28.21
B196 , B200	200	75	1.05	0.98	0.67	0.33	28.12	26.78
B197 , B198	200	120	0.42	0.67	0.78	0.30	18.61	21.30
B202 , B206	1000	30	1.74	2.34	0.61	0.66	32.62	36.81
B203 , B207	1000	75	1.53	3.39	0.88	1.20	29.75	36.90
B204 , B205	1000	120	0.59	2.01	0.38	0.70	20.58	30.63

Table 4.2: This Table presents a summary of the sea clutter shape parameters that were determined, and the radar parameters corresponding to each. The pulse width, look direction and polarization configurations for each data set are indicated in the table, along with the mean amplitude of the radar return clutter.

Relationships between the shape parameter v , and the various radar parameters such as polarization, pulse width and look direction, are easily seen by studying the results in Table 4.2. First, our results indicate that the HH channel is spikier than the VV channel as seen by the smaller values for v in the HH channel. This result is in agreement with previous results reported in [20]. Second, our measurements indicate that as the pulse width is reduced, the backscatter becomes spikier in nature. Again, this agrees with previously published data. Third, the cross-polarized channels HV and VH have similar values of shape parameter v as expected by antenna reciprocity, and they tend in general to be spikier than the like-polarized data, when corrected for noise contamination. We point out that noise correction was not taken into account for the data sets in [13]. Finally, our sea clutter measurements indicate that the backscatter statistics vary as a function of the look-direction. Both the local wind direction (0 degrees) and the swell direction (110–120 degrees) play an important role here. Our strongest returns occur at 30 degrees which correspond most closely to an upwind direction. Our spikiest returns occur at 120 degrees which correspond most closely to the upswell direction. Again, these results are in agreement

with that reported in references [51, 20].

It is instructive to compare the database described here (CRL database) with the database reported on by Ward, Watts, and Baker (Royal Signals and Radar Establishment, RSRE database) in references [46, 49, 6, 51]. The data published by these researchers, as well as the data reported here, support the compound K-distribution model as an accurate representation of sea clutter statistics. Understanding the similarities and differences between these two databases allows for a better appreciation of the general applicability of the K-distribution for sea clutter data. Perhaps the most notable difference between the RSRE and the CRL databases is that they represent different sea clutter measurements. The databases complement each other because they were collected at different times and in different locations. Therefore they represent two independent databases of clutter measurements. Furthermore, the manner in which the data were collected is unique. The RSRE database were collected using both an airborne and cliff-top, noncoherent radar, whereas the CRL database were collected using a cliff-top coherent radar. On the other hand, both databases used a similar operating frequency (9.3 to 10 GHz).

The radar parameters describing the RSRE and CRL databases are also diverse. For example, the RSRE database employed a 30 ns pulse width, whereas the CRL measurements were obtained using 32 ns, 200 ns and 1000 ns pulse widths. Furthermore, HH, HV, VH and VV polarization measurements were included in both databases, although the RSRE cross-polarized measurements do not appear to have been reported. The antenna beamwidths also differ between the two databases. RSRE used a 1.2 degree beamwidth, while CRL used a 2.2 degree beamwidth antenna.

Finally, the look-directions, grazing angles and sea states represented in the RSRE and CRL databases represent a wide range of conditions. The RSRE database contains a full set of look-directions (360 degree coverage) with respect to wind and swell conditions. Also, a large range of grazing angles (0.14 degrees to 10 degrees) are represented. The CRL database on the other hand is limited to a 130 degree viewing area, and has grazings angles

ranging from about 0.1 degrees to 3 degrees. Both databases are representative of a large range of sea states (sea state 0 to sea state 6).

It is evident that the two databases described represent a large number of independent sea clutter measurements, under a wide range of operating conditions. The K-distribution has been shown to provide an accurate description of the amplitude statistics of both the RSRE and CRL databases, supporting the general applicability of the K-distribution to sea clutter modeling.

4.3.3 Growler in Clutter Empirical Amplitude Distributions

In the previous section, we examined sea clutter amplitude statistics and found that the K-distribution provides an accurate description of these statistics. In this section, we examine growler-in-clutter amplitude statistics. In particular, we are interested in determining whether the growler-in-clutter statistics are sufficiently different from the sea clutter statistics that reliable detection of growlers in sea clutter would be possible by thresholding the returned signal amplitude.

We begin by examining an amplitude image that contains both a growler and sea clutter alone. Consider the images shown in Figure 4.14. The HH, VV, HV, and VH amplitudes are shown as a function of time. The growler is located at a range of about 6.52 km in each image. Most of the time, it is difficult to distinguish the growler amplitude from the neighbouring clutter. Figure 4.14 was generated from file B98 and the associated radar parameters are given in Table 3.3 and Table 3.4. For identification purposes, we call this growler 'Growler-1'. Growler-1 was determined to be about 1-2 m high (above the water) and about 10 m across.

The amplitude statistics for Growler-1 and neighbouring sea clutter were determined empirically in the same manner as was done for the sea clutter data in the last section. For these amplitude statistic calculations though, the Growler-1 was resampled with a different set of radar parameters given by file B100 in Table 3.1 and Table 3.2. The Growler-1




Figure 4.14: Amplitude images (HH, VV, HV, VH) of Growler-1 and sea clutter (B98). Growler-1 is located at a range of about 6.52 km. Notice that the strength of the clutter returns is comparably strong making the growler difficult to recognize most of the time.

amplitude statistics are presented in Figure 4.15 and Figure 4.16. Figure 4.15 gives the amplitude statistics for sea clutter near Growler-1, and Figure 4.16 shows the statistics for Growler-1 plus sea clutter.

A second growler was also studied, and we call this growler Growler-2. Growler-2 was sampled using a vertical polarized transmit wave (VV and VH collected) at a range of 4.4km and a bearing of 67 degrees. The winds for this case were from the south at 20 knots. Growler-2 was determined to be about 2 m high and about 20 m across. The amplitude statistics for this growler and neighbouring sea clutter are given in Figure 4.17 and Figure 4.18. The radar parameters associated with Growler-2 are found in Table 3.1 and Table 3.2 under file B122. Figure 4.17 corresponds to sea clutter data near Growler-2, while Figure 4.18 corresponds to the data for Growler-2 plus sea clutter. In both cases, the sea had an average height of 1.5 to 2 metres.

The Growler-1 and Growler-2 data sets serve to illustrate the nature of amplitude statistics from growlers. Some discussion is in order. It is important to determine how significantly the amplitude statistics change in a sea clutter cell when a growler drifts in. It is the difference between the growler statistics and the nearby clutter statistics that differentiate the two. Although Growler-1 and Growler-2 have mean amplitudes large enough to distinguish them from their sea clutter surroundings, the tails of the amplitude distributions of the growlers and those of the sea clutter surroundings are very similar, making high false alarm rates probable. To illustrate the point, the probability density curves for Growler-1 and Growler-2 are redrawn on linear axes in Figure 4.19 and Figure 4.20 respectively. The cumulative distribution functions are also given in order to show the type of performance that can be expected from single-pulse amplitude thresholding in K-distributed sea clutter.

For example, the cumulative distribution functions for Growler-1 and its neighbouring clutter shown in Figure 4.19 indicate that while suffering a false alarm rate of 20 %, a detection probability of only 60 % can be achieved in the HH channel. The HV channel

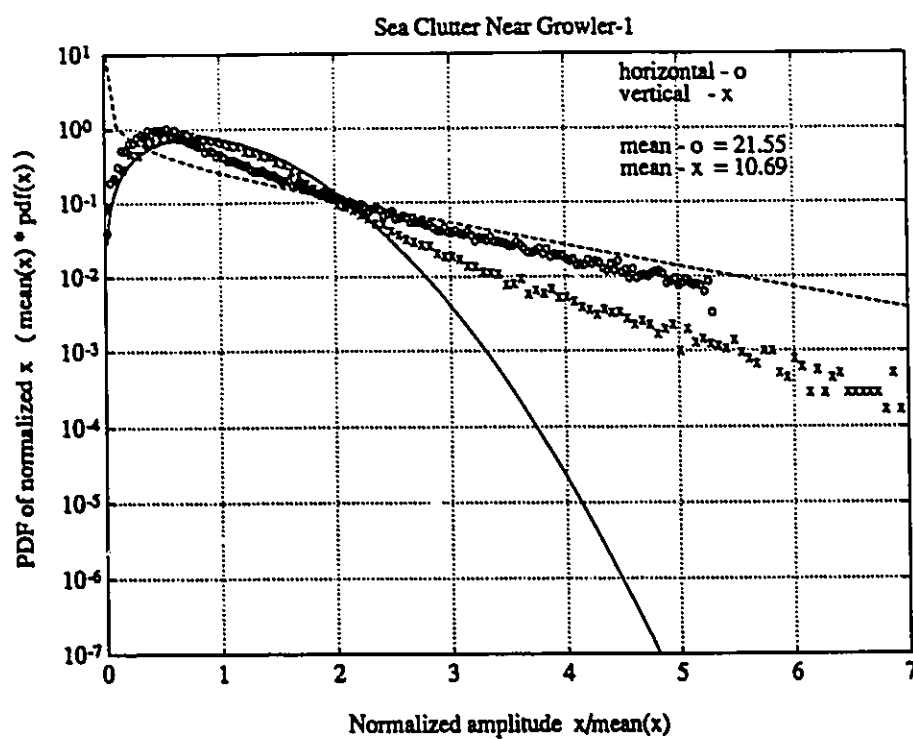


Figure 4.15: Amplitude statistics (HH and HV) of sea clutter near Growler-1 (B100). The HH channel is indicated by 'o' and the HV channel is indicated by 'x'. The Rayleigh curve (solid) and a K-curve with $\nu = 0.25$ (dashed) are also shown for reference.

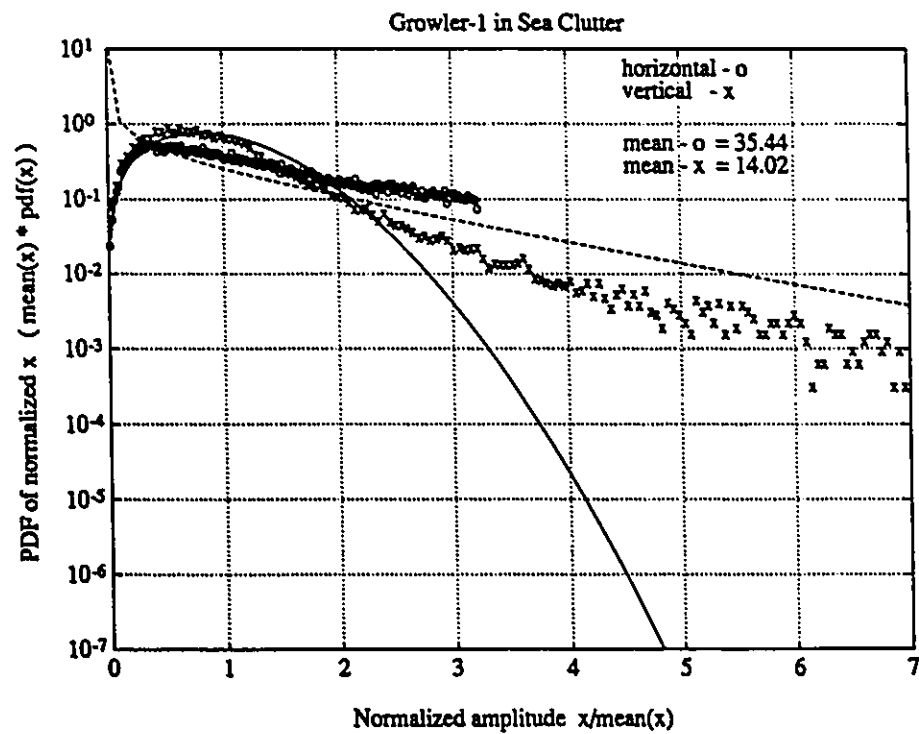


Figure 4.16: Amplitude statistics (HH and HV) of Growler-1 in sea clutter (B100). The HH channel is indicate by 'o' and the HV channel is indicated by 'x'. The Rayleigh curve (solid) and a K-curve with $\nu = 0.25$ (dashed) are also shown for reference.




Figure 4.17: Amplitude statistics (VV and VH) of sea clutter near Growler-2 (B122). The VH channel is indicated by 'o' and the VV channel is indicated by 'x'. The Rayleigh curve (solid) and a K-curve with $\nu = 0.25$ (dashed) are also shown for reference.




Figure 4.18: Amplitude statistics (VV and VH) of Growler-2 in sea clutter (B122). The VH channel is indicated by 'o' and the VV channel is indicated by 'x'. The Rayleigh curve (solid) and a K-curve with $\nu = 0.25$ (dashed) are also shown for reference.

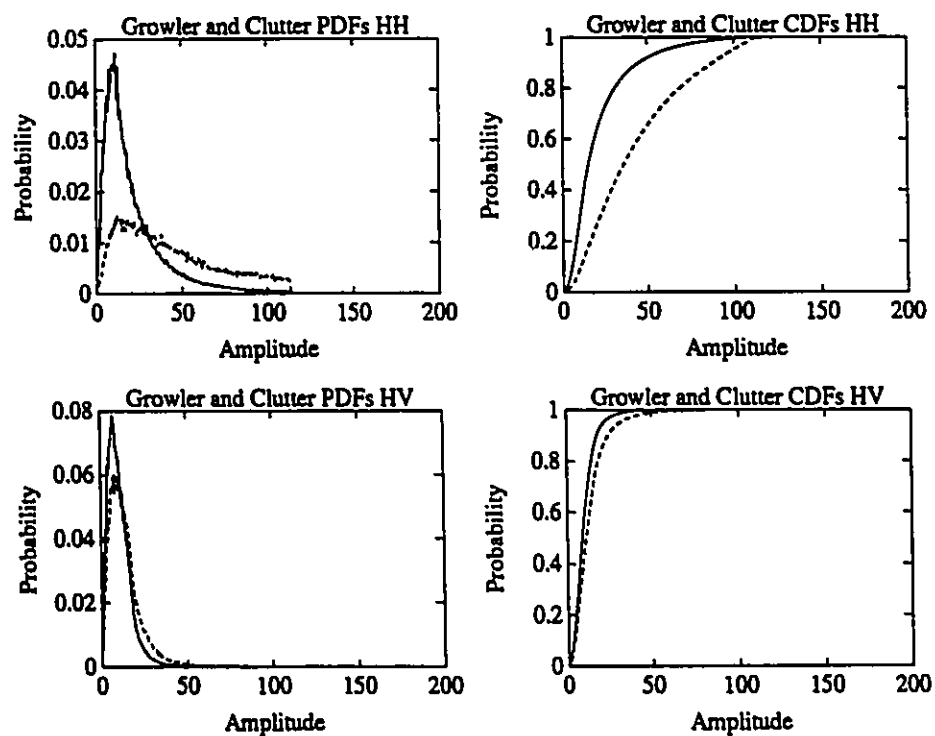


Figure 4.19: Amplitude statistics (HH and HV) of Growler-1 and neighbouring sea clutter. The probability density functions (PDF) as well as the corresponding cumulative density functions (CDF) are shown for each channel. The clutter distributions are indicated by solid curves and the growler distributions are indicated by dashed curves. The HH channel offers greater separation between the growler and sea clutter statistics. (B100).

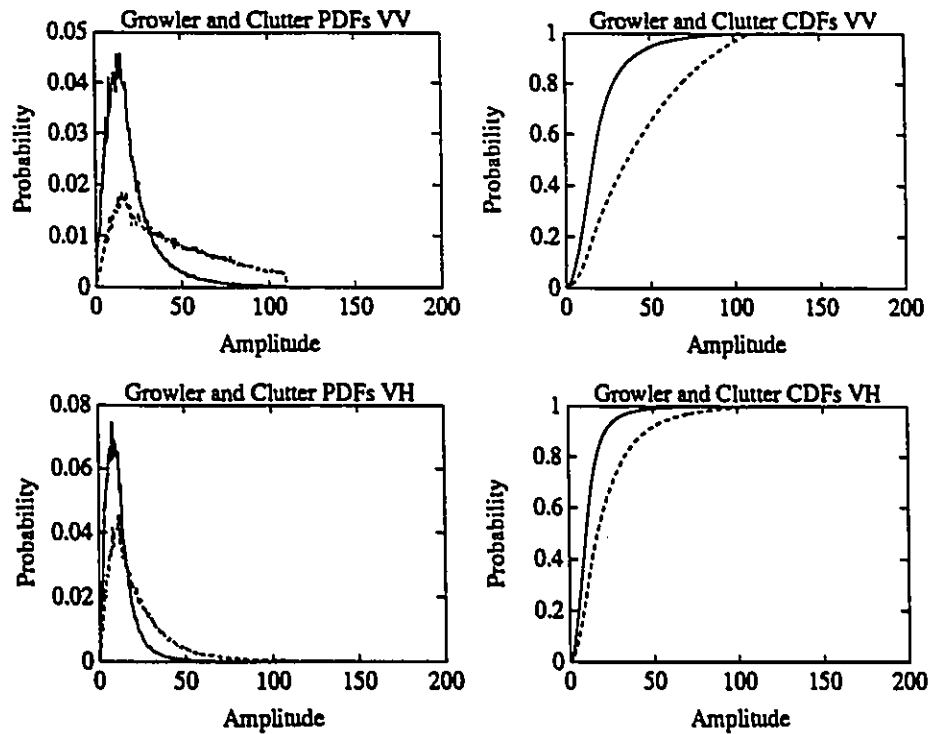


Figure 4.20: Amplitude statistics (VV and VH) of Growler-2 and neighbouring sea clutter. The probability density functions (PDF) as well as the corresponding cumulative density functions (CDF) are shown for each channel. The clutter distributions are indicated by solid curves and the growler distributions are indicated by dashed curves. The VV channel offers greater separation between the growler and sea clutter statistics. (B122).

performance is far worse. Similar performances are seen for Growler-2 in Figure 4.20. These results demonstrate the fact that in spiky clutter such as this with the presence of low lying targets like Growler-1 and Growler-2, using a single-pulse amplitude threshold detector result in high false alarm rates, or alternatively, a desensitized detector. Furthermore, these results also indicate that there is no advantage in using the cross-polarized channels for small ice target detection. In fact, the growlers stand out much more in the like-polarized channels than they do in the cross-polarized channels, at least when using only long-term amplitude statistics for each channel.

4.4 Temporal and Spatial Correlations

Knowledge of the amplitude statistics under H_0 and H_1 is very important in the evaluation of the performance of noncoherent detection strategies. Equally important, however, are the spatial and temporal correlations of the data. Knowledge of these second order statistics improves our understanding of detector performance, and often provides the basis for new detector designs. Correlation properties of sea clutter, for example, allow us to determine the frequency of occurrence of threshold crossings. If sea clutter was completely uncorrelated from sample to sample, then we would not expect a false alarm at a given instant of time to be followed by other false alarms in the next few sample times. However, correlated sea clutter result in several threshold crossings occurring together. Clearly, correlation in space results in false alarm rates that are not independent of position. In situations like this, CFAR processors excel as they try to estimate the *local* characteristics of the clutter in order to adaptively set a threshold.

Recall in Section 4.3.1, the compound nature of the K-distribution was described. In particular, we saw that the two-component densities given in Equation 4.11 and Equation 4.12 were said to have correlation times that were different. We said that the speckle component given by Equation 4.11 typically decorrelates in about 10 ms at X-band. The

mean level component described by Equation 4.12 could take several seconds to decorrelate. We look at sea clutter data collected with the IPIX radar in order to see how well our measurements agree with those reported in the literature. We also examine spatial correlations of sea clutter, since both spatial and temporal correlations are important for CFAR processors.


Consider the plots given in Figure 4.21. In this figure, we show calculations of the temporal and spatial correlations of sea clutter data given in file B195. The radar parameters associated with this file can be found in Table 3.1 and Table 3.2. The correlation estimates shown in Figure 4.21 were determined by averaging over the first 10 s of data. The entire range swath of 160 m was used for range correlation estimates.

Figure 4.21 supports the compound clutter model for sea clutter. Two distinct correlation times are evident. The fast speckle component described in Equation 4.11 is determined to have a correlation time of about 10 ms. This agrees with the results reported in [45]. Furthermore, the much longer correlation time of the second component is also observed. For a lag of 1 s, this component is still very correlated. This is due to the fact that the mean level of the sea is strongly correlated over a period of 1 s. The range correlations shown in Figure 4.21 are also interesting. Over a distance of only 30 m (the pulse-length), the spatial correlation has dropped to less than 60 %. This indicates that the mean level component varies much more quickly in range, than it does in time.

In Figure 4.22, 60 s of data were averaged in order to estimate temporal lags out to 10 s. The temporal lag spacing was chosen to be 50 ms since we are only interested in determining the time constant associated with the mean level component from this figure. From this figure, we see that the correlation of the mean level component is very strong and persists for several seconds.

The temporal and spatial correlation results of sea clutter described here will have considerable influence on our detection strategies. These will be discussed in detail in Chapter 6.

SPATIAL/TEMPORAL CORRELATION



(a)






Figure 4.21: This figure gives estimates of the temporal and spatial correlation properties of sea clutter (B195). In (a), an orthographic view of the two-dimensional time/range correlation of the sea clutter data is given. In (b), just the range correlation is shown and in (c), just the temporal correlation is shown.

SPATIAL/TEMPORAL CORRELATION



(a)





Figure 4.22: This figure gives estimates of the temporal and spatial correlation properties of sea clutter (B195). In (a), an orthographic view of the two-dimensional time/range correlation of the sea clutter data is given. In (b), just the range correlation is shown and in (c), just the temporal correlation is shown. This figure is intended to emphasize the mean level component. Therefore time lags which are large compared to the time constant of the speckle component are used.

4.5 Summary

In this chapter, a noncoherent analysis of sea clutter data collected with the IPIX radar was described. We began by considering the radar equation for point targets like growlers and distributed targets like sea clutter. This led to the determination of power versus range laws and radar cross section calculations, for both target models. Next, we introduced the K-distribution as a model suitable for describing the amplitude statistics of sea clutter and growler returns. The results on amplitude statistics of sea clutter and growlers may be summarized as follows:

1. The K-distribution is a good model for the amplitude statistics of sea clutter, for both the like-polarized and cross-polarized configurations.
2. The measurements on small ice targets indicate that amplitude alone does not provide an adequate discriminant for detection in spiky sea clutter, when considering only single pulse, fixed threshold processing. Furthermore, the long-term amplitude statistics in the cross-polarized channels do not offer any advantage for the detection of small ice targets. The like-polarized channels perform much better.

Finally, the temporal and spatial properties of sea clutter were examined. Again, these results support the compound model of sea clutter described by the K-distribution, and influence detector design.

The noncoherent measurements described in this chapter emphasize the challenge that faces a radar designer when considering the detection of targets whose radar cross section is small relative to that of the surrounding clutter, making their presence in a radar resolution cell difficult to detect. Although single pulse, fixed threshold amplitude detectors are not suitable for the reliable detection of challenging growlers such as Growler-1 and Growler-2, other multiple pulse, scan-to-scan type processors may perform better, since the sea variation is usually greater than the target variations. These and other processing structures such as CFAR processors deserve attention as they may provide a benchmark for

comparison to higher dimensional radar signal processors. Higher dimensional radar signal processors are considered in the next chapter.

Chapter 5

Coherent Analysis of Sea Clutter and Growler Data

5.1 Introduction

In the last chapter, we focussed on the noncoherent properties of sea clutter and growler data in both the like-polarized and cross-polarized channels. We saw that the cross-polarized channels did not offer any increased separation between growler and sea clutter amplitude statistics. Studies presently being done at the CRL indicate that when it comes to polarization processing, the like-polarized channels offer significantly more discrimination between growlers and sea clutter than do the cross-polarized channels. In this chapter, we will focus on the coherent properties of sea clutter and growler data. Because the like-polarized channels have shown to be better suited for growler detection in sea clutter, we will limit our coherent processing to the like-polarized channels HH and VV. We are interested in determining whether the phase of the received signal in these channels offers improvements for the detection of growlers in sea clutter. As opposed to the one-dimensional processing that was described in Chapter 4, in this chapter we will be performing three-dimensional processing. The three dimensions we make use of are the I and Q or amplitude and phase

dimensions, as well as the polarization dimension.

A natural place to begin a coherent analysis is with the characterization of the phase statistics. We saw in the last chapter that the amplitude statistics of sea clutter are K-distributed. We shall see that the phase statistics are uniformly distributed. Then we will proceed by examining the time-varying spectra of growlers in sea clutter. We will describe our discovery of the *winking phenomenon* which has already been reported in the literature [7, 8]. Then we will examine in some detail characteristics of Doppler spectra of sea clutter and growlers based upon the processing of a fraction of a second of data. This type of analysis is performed in order to develop *medium dwell-time* Doppler techniques for improved growler detection that will satisfy the six goals outlined on page 11. We will develop a couple of different ways to parameterize these Doppler spectra. Finally, we will consider ways to use Doppler spectral parameters to estimate growler to clutter ratios (GCR) and growler to noise ratios (GNR), as well as radar cross sections (RCS).

5.2 Phase Statistics of Sea Clutter

In this section, we are interested in analyzing the phase statistics of sea clutter data, that is, the distribution of $\arctan Q/I$. It is generally accepted that the phase statistics of sea clutter are uniform, although we have seen no verification of this *fact* with real measurements in the literature. The uniformity of sea clutter statistics seems correct because phase is associated with the absolute distance to the sea clutter scatterers from the radar, modulo π/λ where λ is the operational wavelength. It is reasonable to expect that nature does not favour one location for sea clutter scatterers over another, with respect to the radar. However, if care is not taken in the estimation of the phase statistics from digital I and Q samples, strange looking results can occur. In fact, it is not difficult to end up with a Gaussian looking distribution for the phase statistics!

Since analog/digital (A/D) converters are used to quantize the I and Q video outputs of the radar in order to get the digital I and Q samples, care must be taken that some phases are not *preferred* over others, as a result of the quantization process. For example, any dc offsets on the I and Q video signals will force the point (0,0) to fall away from the center of the square A/D quantization grid causing certain phases to be favoured. A simple example will demonstrate the point. Consider the case where the I video signal has a zero dc offset but the Q video signal has a positive dc offset. A pure complex sinusoid with an amplitude that just falls short of saturating the I channel A/D converter will necessarily cause saturation of the Q channel A/D converter, but only for positive values of the Q video signal. If the Q axis is taken as a phase of 90 degrees, then there will be a band of phases about 90 degrees that will be impossible to achieve! Thus, the phase statistics will not be uniform. Other problems arise due to the quantization effect. Since the quantization of the A/D converters is independent of the amplitude of the I or Q video signals, smaller amplitudes will result with large quantizations in phase whereas larger amplitudes will have a much finer phase quantization. Therefore, it is important to limit the I and Q samples that will be used for phase statistic calculations to lie within a circle that is completely contained within the square A/D converter quantization grid.

In Figure 5.1, we show a typical phase distribution obtained from the sea clutter file with file ID B100. The radar parameters for this data are found in Table 3.1 and Table 3.2. Indeed, the phase statistics are uniform. Any deterministic-looking variations in the phase statistics are a result of the quantization effects of phase measurements on a square A/D converter measurement grid.

5.3 Time-Varying Spectra of Growlers in Sea Clutter

During the IPIX field trials, we discovered that the I and Q video returns from growlers were distinctively different than those of sea clutter. We have coined the term *winking*




Figure 5.1: Typical phase statistics of sea clutter data (B100). Both the HH and HV channel statistics are shown. The HH channel is indicated by 'o' and the HV channel is indicated by 'x'.

phenomenon to describe this characteristic of growler returns because the I and Q traces displayed on an oscilloscope appeared to wink at the location of a growler. That is, the rate at which zero crossings would occur in either the I or Q signal varied as a function of time. The zero crossing rate would start off slowly at first, increase to a maximum, decrease to a minimum rate and then return back to its starting rate. This pattern seemed to repeat itself in a somewhat regular pattern, and the period of the cycle was on the order of several seconds (5–10s). The *winking phenomenon* was also observed for other small floating targets and some of these analyses are reported in [7, 8].

Recall that a point target moving with a radial velocity of v_r knots (where v_r is assumed positive) away from the radar results in a negative Doppler frequency shift on the RF carrier. If the target was moving toward the radar so that the velocity is $-v_r$ knots, a positive frequency shift would result on the RF carrier. This is known as the Doppler effect, and the frequency shift is called the Doppler shift. The amount of Doppler shift for velocity v_r knots is derived in [20] and is given by:

$$f_d = \frac{1.03v_r}{\lambda}, \quad (5.1)$$

where f_d is the Doppler shift in Hz. and λ is the operational wavelength in metres. The wavelength λ used in the IPIX radar is 3.19 cm so that a velocity of 1 knot results in a Doppler shift of 32.2 Hz. We would expect that sea clutter returns would have a band of Doppler shifts due to the fact that ocean waves advance with varying speeds depending on their wavelength, and there is a spectrum of wavelengths of ocean waves. We might expect the growler on the otherhand to have a smaller band of Doppler shifts due to the fact that it is like a point target, and it is rigid. Applying spectrum estimation techniques to sea clutter or growler baseband time series would therefore characterize the distribution of Doppler shifts. As a result, a spectrum estimate formed from a baseband radar signal is often referred to as the Doppler spectrum. Technically speaking, Doppler spectra should be referred to as either Doppler Frequency spectra or Doppler Velocity spectra so that no

confusion arises as to which direction the target is moving from the spectra. However, this practice is often not used in the literature so that confusion does arise. In order to avoid any confusion in this thesis, we will use the convention that negative Doppler frequencies and negative Doppler velocities correspond to a target approaching the radar, and positive Doppler frequencies and positive Doppler velocities correspond to targets moving away from the radar. In this way, any Doppler spectrum can have its axis relabelled easily as either frequency or velocity without reflecting the spectrum about zero.

Doppler spectrum analysis should be able to validate the *winking phenomenon* that was observed during the IPIX field trips. However, since the spectrum of the growler returns varies in almost a cyclical pattern in a period on the order of 5-10 s, a temporal/spectral representation of the growler data is required in order to characterize the data properly. Temporal/spectral representations were studied by Nohara and Haykin in [15], and it was determined that the spectrogram [53] was best suited to characterize radar returns from growlers.

By performing spectrograms, time-varying Doppler images are formed from growler and sea clutter data. A typical set of images are given in Figure 5.2, Figure 5.3, Figure 5.4, and Figure 5.5. The first two images are the result of spectrogram processing on file B98, at the range cell containing the growler. Recall that the corresponding amplitude image of file B98, for the entire range swath, is given in Figure 4.14. Figure 5.4, and Figure 5.5 are the result of spectrogram processing at a sea clutter range cell in front of (closer in range than) the growler location. The radar parameters associated with file B98 are given in Table 3.3 and Table 3.4.

The spectrogram images show the first 26 s of data in file B98. Each frequency raster was calculated by performing a hamming-windowed FFT (Fast Fourier Transform) on 512 time samples. Since the effective PRF for each polarization channel is 1 KHz, each frequency raster is determined by processing 0.512 s of data. Adjacent frequency rasters are calculated by sliding the 512-point window 50 time samples through the data. The images

Growler in Clutter HH Time-Frequency Plot




Figure 5.2: This image shows the time-varying Doppler spectrum of a growler in sea clutter. The dark, narrow, continuous line represents the changing growler Doppler shift as a function of time. The remaining spectral contributions are due to the sea clutter in the growler range cell as well as receiver noise. Black represents large intensity and white represents small intensity. A logarithmic scale is used to code the Doppler spectrum value at each frequency bin onto the greyscale. This image is from the HH channel in file B98.

Growler in Clutter VV Time-Frequency Plot




Figure 5.3: This image shows a typical time-varying Doppler spectrum of a growler in sea clutter. The dark, narrow, continuous line represents the changing growler Doppler shift as a function of time. The remaining spectral contributions are due to the sea clutter in the growler range cell as well as receiver noise. Black represents large intensity and white represents small intensity. A logarithmic scale is used to code the Doppler spectrum value at each frequency bin onto the greyscale. This image is from the VV channel in file B98.

Sea Clutter HH Time-Frequency Plot




Figure 5.4: This image shows a typical time-varying Doppler spectrum of sea clutter. The darker regions correspond to the significant band of Doppler frequencies associated with sea clutter returns. The remaining spectral contributions are due to receiver noise. Black represents large intensity and white represents small intensity. A logarithmic scale is used to code the Doppler spectrum value at each frequency bin onto the greyscale. This image is from the HH channel in file B98.

Sea Clutter VV Time-Frequency Plot




Figure 5.5: This image shows a typical time-varying Doppler spectrum of sea clutter. The darker regions correspond to the significant band of Doppler frequencies associated with sea clutter returns. The remaining spectral contributions are due to receiver noise. Black represents large intensity and white represents small intensity. A logarithmic scale is used to code the Doppler spectrum value at each frequency bin onto the greyscale. This image is from the VV channel in file B98.

are greyscale coded logarithmically with respect to the magnitude of each FFT bin. The darker the color is, the larger the magnitude is.

By comparing Figure 5.2 and Figure 5.3 with Figure 5.4 and Figure 5.5 respectively, one can immediately see that the *winking phenomenon* has been captured by these images. The sea clutter Doppler spectra are fairly *wideband* in the sense that there is significant clutter power across a frequency band of about 100 Hz which for this data set is centered at about -100 Hz. The result of a growler in the sea clutter cell is quite dramatic as shown by the dark, narrow, snakelike line that runs down the images in Figure 5.2 and Figure 5.3, centered at approximately -30 Hz. This snakelike contribution from the growler returns represents the *winking phenomenon*. The growler is moving at mean translational velocity of about 1 knot towards the radar, however its velocity accelerates and decelerates in a cyclical fashion, causing a cyclical variation in the instantaneous Doppler frequency. It was this cyclical behaviour in the instantaneous velocity of the growler that caused the I and Q video signals to *wink* as described above.

One thing that is obvious in comparing the Doppler spectra images shown here with the corresponding amplitude images given in Figure 4.14, is the much larger separation that exists between the growler data and clutter data when spectrogram processing is used. This improvement is a direct result of the *winking phenomenon* and has prompted the study of *long dwell-time* Doppler techniques at the CRL, as described in Chapter 1. In the next section, we will examine the underlying physics responsible for the *winking phenomenon*.

5.3.1 Dynamic Modeling of Growler Motion

The *winking phenomenon* described in the last section is very important for improving the detection performance of growlers in sea clutter. Most of our effort concerns the exploitation of this phenomenon. It is critical that we are able to justify the presence of the *winking phenomenon* because our measurements represent only a subset of the possible sea state and growler combinations. In this section, we will examine the underlying dynamics of the

problem in order to explain the *winking phenomenon*. Since Doppler processing describes the radial motion of a target under examination, we will develop the equations that describe the motion of a floating target in the ocean. Clearly this will involve the development of wave and particle dynamic models.

Many wave theories exist as a result of assumptions made in the solution of fundamental differential equations that govern the motion of water particles and hence, waves. We begin by examining the basics that are fundamental to any wave theory, and then we describe simplifications that can be made which are suitable for our problem. Our development follows those given in [54, 55]. Most wave theories are based on a few basic parameters: water depth d , peak to peak wave height H , and wave period T . A flat bottom is assumed. If the wavelength of the wave is λ , and the period T is the time for the crest of the wave to travel a distance of λ , then the wave velocity is given by $c_{ph} = \lambda/T$. The wave velocity is often called the phase velocity or celerity. Now, assuming that the water is incompressible, the conservation of volume results in a continuity differential equation given by

$$\frac{\partial u}{\partial x} + \frac{\partial v}{\partial y} + \frac{\partial w}{\partial z} = 0, \quad (5.2)$$

where u , v , and w are the three components of a fluid particle velocity in the rectangular Cartesian coordinate system spanned by (x, y, z) . The origin of the coordinate system is at the still water level, the x direction is the direction of wave propagation, and the y direction represents the height of the fluid particle. Wave theories usually assume a wave model in the two dimensions x and y . It is assumed that the wave height is independent of the z direction. Assuming a sinusoidal waveshape, a single wave is therefore described by the equation

$$h(x, t) = A \cos(kx - \omega t), \quad (5.3)$$

where $h(x, t)$ is the height of the wave at position x and time t , $k = 2\pi/\lambda$ is the spatial wavenumber, and $\omega = 2\pi/T$ is the temporal frequency of the wave.

If the water particles are purely translational (no rotation), then a velocity potential

function can be shown to exist [54] which satisfies the relation,

$$\mathbf{V} = -\nabla\Phi \quad (5.4)$$

where \mathbf{V} is the velocity vector and Φ is the velocity potential function. Substitution of Equation 5.4 into Equation 5.2 results in a new continuity equation given by

$$\nabla^2\Phi = \frac{\partial^2\Phi}{\partial x^2} + \frac{\partial^2\Phi}{\partial y^2} + \frac{\partial^2\Phi}{\partial z^2} = 0. \quad (5.5)$$

This equation is known as Laplace's equation and its solution describes the kinematics of wave particles (velocity and acceleration).

As well as the kinematics, the kinetics (pressure and force) of wave particles are also required. Beginning with the unsteady Navier-Stokes equation for irrotational flow, the Bernoulli equation is derived and is given by

$$p + \rho gy + \frac{1}{2}\rho(u^2 + v^2 + w^2) = C, \quad (5.6)$$

where p is the pressure, g is the acceleration due to gravity, ρ is the mass density of the wave, y is the vertical coordinate, and C is a constant.

Wave theories are derived by solving the Laplace and Bernoulli equations for particular boundary conditions and simplifying assumptions. Two boundary conditions of particular importance assume that a particle of water lying on the surface at some time will continue to remain there, and that the atmospheric pressure outside the fluid is constant. These boundary conditions constrain the Bernoulli equation and are called the *free surface boundary conditions*. The Bernoulli equation is nonlinear and therefore satisfying the free surface boundary conditions is difficult, since the surface is continually changing. An approximation to this problem is achieved by linearizing the free surface boundary conditions so that they are satisfied only at the mean water level. In order for this linearization to be valid, it is assumed that the wave height is small with respect to λ or d .

The wave theory that results from this assumption and simplification is appropriately called the Small Amplitude Wave Theory or the Linear Airy Wave Theory (also called

Sinusoidal Wave Theory) and it is perhaps the most useful theory available today. It gives solutions to the Laplace and Bernoulli equations for the wave model described in Equation 5.3. These solutions describe the horizontal and vertical velocities, u and v , in the x and y directions respectively. With the deep water assumption and position $x = 0$, these equations simplify to:

$$\begin{aligned} u(t) &= A\omega \cos(\omega t) \\ v(t) &= -A\omega \sin(\omega t). \end{aligned} \quad (5.7)$$

Furthermore, the Linear Airy Wave Theory gives a relationship between the temporal frequency ω and the spatial wavenumber k . That is,

$$\omega^2 = gk. \quad (5.8)$$

This equation is known as the Dispersion Relation.

Equation 5.7 describes the velocity of a water particle at a particular range, as a function of time. These equations describe a cyclical pattern that the water particle follows as the energy of a wave passes through the particle. Tank studies have verified this cyclical motion that water particles undergo. Furthermore, tank studies have shown that floating ice targets such as growlers undergo the same particle motions, so long as their size is small ($< \lambda/13$) compared to the wavelength of the driving wave [56]. Therefore, Equation 5.7 is suitable for describing the kinematics of growlers in the ocean for a single wave as described in Equation 5.3. If the sea could be described by a single wave, Doppler processing would therefore show a sinusoidal frequency variation as a function of time, for a growler floating on a wave that is travelling in the direction of the radar. For small grazing angles, the frequency variation would follow $u(t)$ in Equation 5.7. The period of this frequency variation is therefore controlled by the wave period T . In a very coarse sense, the development that we have described thus far, models the *winking phenomenon*. The dynamic model of the growler motion can be improved upon significantly by using a more realistic wave model.

The wave model given in Equation 5.3 is for a single wave description of the sea. In reality, the sea surface is made up of a continuum of waves of various wavelengths. Assuming a linear medium, the growler motion will be the result of a superposition of the motion induced by each wave. Therefore, accurate modeling of growler motion requires a good model for the distribution of waves present in the sea. A maximum wavelength can be supported in a given sea, beyond which breaking waves result. In addition, a continuum of smaller and smaller waves exist and their relative power typically drops off at a rate of ω^{-5} with respect to their frequency ω .

Many mathematical models have been proposed in order to describe the wave spectrum of the sea. These models are usually based on fetch, wind and wave height parameters. The most common single-parameter spectrum in use today is the Pierson-Moskowitz (PM) model which is based on wave height [57]. This model is given by

$$S(\omega) = \alpha g^2 \omega^{-5} \exp(-1.25(\omega/\omega_o)^{-4}) \quad (5.9)$$

where the peak of the spectrum occurs at ω_o and $\alpha = 0.0081$. The frequency ω_o is related to the significant wave height H_s by

$$\omega_o^2 = 0.161g/H_s. \quad (5.10)$$

Now, extending our wave model to include M wave scales, we have

$$h(x, t) = \sum_{i=1}^M A_i \cos(k_i x - \omega_i t + \theta_i) \quad (5.11)$$

where k_i and ω_i are the wavenumber and frequency of the i^{th} wave in our model, and A_i and θ_i are the amplitude and phase of the i^{th} wave. By the principle of superposition (valid for linear wave model), the horizontal velocity of this model at $x = 0$ is given by

$$u(t) = \sum_{i=1}^M A_i \omega_i \cos(\omega_i t + \theta_i). \quad (5.12)$$

We should emphasize that A_i and θ_i are random variables since the sea surface itself is random. The PM spectrum given in Equation 5.9 can be used to specify the expected value

of A_i for each ω_i . In order to create a typical sample function of $u(t)$, the PM spectrum is discretized into M frequency bins centered on frequencies ω_i . The area under $S(w)$ in each of these bins specifies the expectation of the power, P_i , associated with each of the M waves in Equation 5.11 and Equation 5.12. A sample value for A_i is obtained by drawing an independent sample, p_i , from a zero-mean Gaussian process with variance P_i and assigning $A_i = p_i$. The random variables θ_i are selected independently from a uniform distribution on 2π . Realizations of $u(t)$ determined by the procedure just described result in time-varying velocity functions that are very similar to the time-varying Doppler characteristics of growlers that we have measured.

In this section, we have shown that the *winking phenomenon* is the result of the kinematics associated with small floating bodies in water waves. This is very exciting for a couple of reasons. First, it validates our measurements of the time-varying Doppler signatures of growlers in sea clutter, allowing for realistic conclusions to be made based on the measured data. Second, it allows us to model and simulate growler returns under conditions for which there is a lack of real measurements.

5.4 Doppler Spectrum Estimates of Sea Clutter Returns

In Section 5.3, we saw that the *winking phenomenon* associated with the growler returns was captured by spectrogram processing of the data. Having spectrogram images like those given in Figure 5.2, Figure 5.3, Figure 5.4 and Figure 5.5 made it quite easy to distinguish between clutter cells and growler cells. However, we are concerned in this thesis in distinguishing between the two using only a fraction of a second of data. This amounts to making a decision based upon a single raster in the spectrogram images. Each of these rasters corresponds to a *medium dwell-time* Doppler spectrum estimate. In this section, we examine sea clutter Doppler spectra in detail and develop models to describe a typical clutter spectrum estimate. These models allow us to characterize the *medium dwell-time* Doppler spectra of sea clutter

returns with just a few parameters.

5.4.1 Nonparametric Spectrum Estimators

Sea clutter data collected with the IPIX radar have been analyzed in the Doppler domain. The data files that were used for Doppler processing are described in Table 3.3 and Table 3.4. Typical Doppler spectra of sea clutter data are shown in Figure 5.6. There are three different types of spectra that are typical of sea clutter returns. The first and most common spectrum is bell-shaped and fairly broadband with the base spanning several tens of Hz. The second case occurs less frequently, but still is common. In this case, the spectrum looks bell-shaped as well but much narrower in width (less than 10 Hz). Finally, there often appears to be little or no return from the clutter. The Doppler spectrum estimate seems to model just the receiver noise with a small amount of additive clutter.

The spectrum estimates shown in Figure 5.6 are nonparametric in that no model is assumed for the spectrum. Although this type of analysis is useful to study the data with, it does not help in reducing the dimensionality of the detection problem, or in describing the features of the spectra that may be of interest. For example, it is often of interest to determine the mean velocity and the spectral spread of the velocities of clutter. In fact, these parameters may be of direct value to oceanographers and meteorologists. We too are very interested in these parameters because they describe the motion of clutter which is quite different from the characteristic motion of growlers described by the *winking phenomenon*. Therefore, we might expect these parameters to provide the basis for improved detection performance.

We consider a few different approaches in order to estimate parameters such as mean Doppler frequency (velocity) and Doppler spectral spread from clutter data. In the first method, we borrow classical results from probability theory. The mean frequency and spectral spread parameters mathematically correspond to the first and second normalized moments of the Doppler spectrum function. If m_n is the n^{th} moment of the Doppler

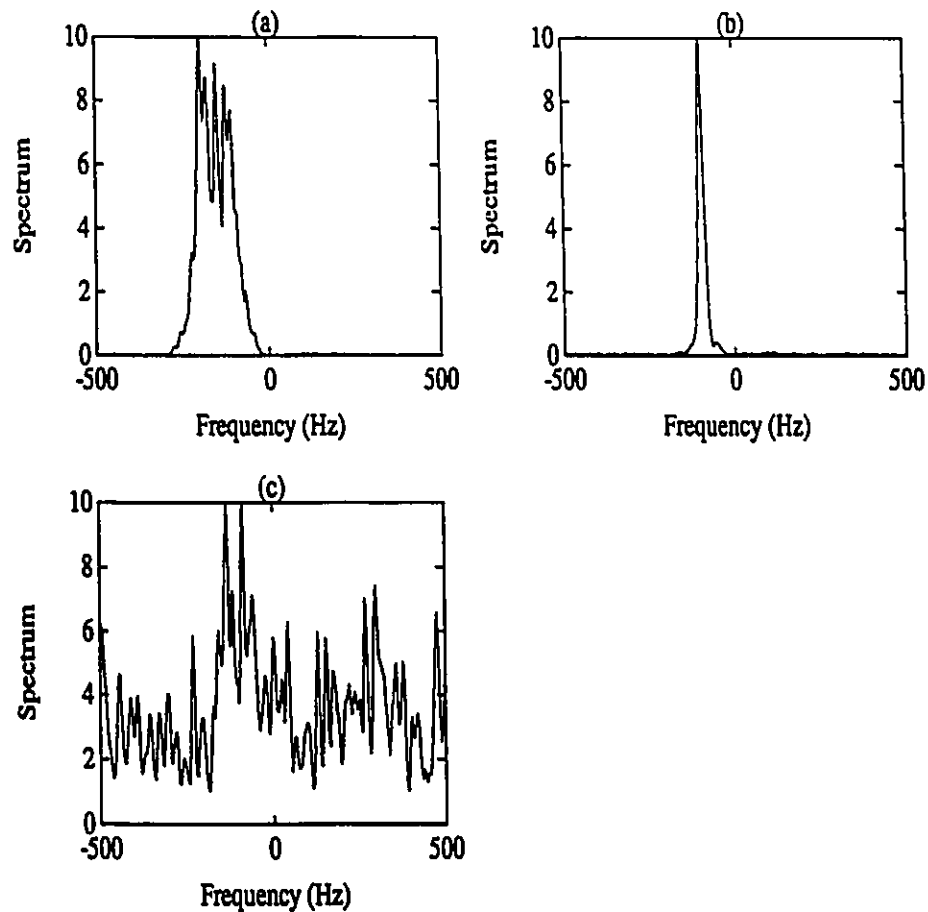


Figure 5.6: This figure shows typical Doppler spectrum estimates of sea clutter based on Fourier analysis of about 0.5 s of data. Most of the time, clutter spectra appear like that in (a), bell-shaped with a fairly broad base (tens of Hz). Often, clutter spectra can look like the estimate given in (b) which is much more narrowband than the spectrum estimate shown in (a). Often, the CNR is so small that the clutter spectral estimate is buried in noise as shown in (c). This is usually due to shadowing. The data shown here is from file B98. Note: all spectra have been normalized to a maximum value of 10.

spectrum $S(f)$, then

$$m_n = \int f^n S(f) df. \quad (5.13)$$

Therefore, the mean frequency f_c is simply given by m_1/m_0 . In a similar manner, the central moment definition can be used to calculate the Doppler spectral spread (width) σ_f :

$$\sigma_f^2 = \frac{\int (f - f_c)^2 S(f) df}{m_0}. \quad (5.14)$$

In practice, we have to estimate $S(f)$ and therefore Equation 5.13 and Equation 5.14 can be used only to estimate the values of f_c and σ_f . It has been shown that this direct classical approach is inefficient and suffers from noise bias problems [58]. We do not use these estimators here.

5.4.2 Pulse Pair Estimators

Other methods have been proposed that work on the autocorrelation function of the data rather than on the spectrum and hence have been called covariance type estimators. Probably the oldest and most widely used of these methods are the pulse pair estimators introduced by Rummler in 1968 [59, 60, 61]. These estimators have been used widely in pulse Doppler weather radars and have been studied and modified by many other researchers [58, 62, 63, 64, 65, 66]. The pulse pair estimators were derived by exploiting the relationship between the probability density function and the moment generating function or characteristic function. Applying this relationship gives the following:

$$m_n = R^{[n]}(0)/(j2\pi)^n \quad (5.15)$$

where m_n is the same as defined in Equation 5.13 and $R^{[n]}(0)$ is the n^{th} derivative of the autocorrelation function at lag zero. In order to derive the pulse pair estimators (as well as other poly-pulse type estimators), the autocorrelation function is written in terms of a magnitude and phase function and each of these are expanded as a McLaurin series. Then Equation 5.15 is used to replace the derivatives of the autocorrelation function in each of

the series. Finally, the series is truncated and the pulse pair estimators result. The pulse pair estimators for mean frequency and spectral width of sample data are given by

$$\begin{aligned}\hat{f}_c &= (2\pi T_s)^{-1} \arctan \left[\frac{\text{Im} \hat{R}(T_s)}{\text{Re} \hat{R}(T_s)} \right] \\ \hat{\sigma}_f &= (\sqrt{2}\pi T_s)^{-1} \left[1 - \frac{|\hat{R}(T_s)|}{\hat{R}(0)} (1 + \text{SNR}^{-1}) \right]^{1/2}.\end{aligned}\quad (5.16)$$

The units for \hat{f}_c and $\hat{\sigma}_f$ are Hz, $\hat{R}(T_s)$ is an estimate of the autocorrelation function of the sample data at lag T_s seconds where T_s is the sampling period, and SNR is the signal to noise ratio. The SNR can be estimated as $(\hat{R}(0) - P_n)/P_n$ where the noise power P_n must be provided or estimated independently. Autocorrelation estimates are calculated from the N data samples $z_k, k = 1, \dots, N$ by using the formula

$$\hat{R}(mT_s) = (N - m)^{-1} \sum_{k=1}^{N-m} z_k^* z_{k+m}, \quad (5.17)$$

where only the lags $m = 0$ and $m = 1$ are required to evaluate the pulse pair estimates.

If the pulse pair estimates are applied to the typical clutter spectra given in Figure 5.6, Figure 5.7 results. Figure 5.7 indicates that the pulse pair mean frequency estimator is quite good. The spectral width estimator however seems to overestimate the spectral width when the clutter is narrowband. The dependence of the accuracy of the pulse pair spectral width estimator on the actual spectral width is discussed in [62]. In order to overcome this dependence, different poly-pulse estimators should be used for different ranges of spectral widths.

5.4.3 Autoregressive Spectrum Estimators

The next method that we describe for estimating Doppler spectra and spectral moments involves the use of parameter spectrum estimation techniques. The Autoregressive (AR) Model provides the basis for a parametric spectrum estimator that has been used successfully in the past [67] for radar detection and classification. Haykin [9] also describes the wide use of the AR model as a tool for fitting real time series data in fields other than

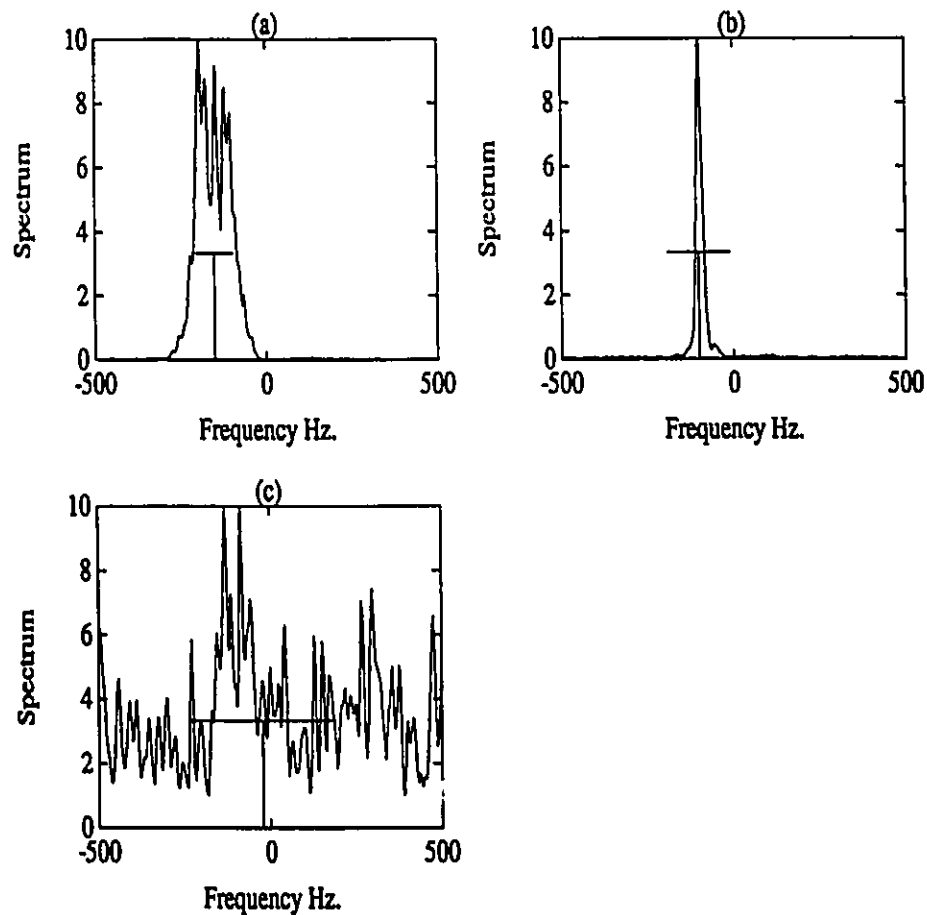


Figure 5.7: This figure shows the results of applying pulse pair estimators on typical sea clutter data. The estimated Doppler spectra are shown as well as the pulse pair estimates in order to demonstrate the performance of pulse pair estimators. A 'T' bar is drawn on each spectrum to indicate the pulse pair mean frequency estimate as well as the spectral width estimate. The height of the 'T' bar is arbitrarily set to one third of the maximum height.

radar. We examine the usefulness of this model here for both estimating and modeling the Doppler spectra, as well as providing spectral moment estimates. A time series, $u(n)$, is said to arise from an AR process of order M if it satisfies the following difference equation:

$$u(n) + a_1^* u(n-1) + \dots + a_M^* u(n-M) = v(n) \quad (5.18)$$

where $a_1^*, a_2^*, \dots, a_M^*$ are the AR parameters and $v(n)$ is a white noise process with zero mean and variance σ_v^2 .

Examination of the AR difference equation indicates that an AR process is generated by driving white noise through an all-pole filter. The filter imposes correlation on the output samples. Therefore, fitting time series data to an AR model is equivalent to designing the best all-pole filter of a given order that models the spectrum of the data most closely. The derivation of this optimization is described in [9]. The solution of the AR parameters is given by the matrix equation

$$\mathbf{a} = -\mathbf{R}^{-1}\mathbf{r} \quad (5.19)$$

where the AR parameters are given by

$$\mathbf{a}^H = [a_1^*, a_2^*, \dots, a_M^*], \quad (5.20)$$

the correlation matrix \mathbf{R} is given by

$$\mathbf{R} = \begin{bmatrix} r(0) & r(1) & \dots & r(M-1) \\ r(-1) & r(0) & \dots & r(M-2) \\ \vdots & \vdots & \ddots & \vdots \\ r(-M+1) & r(-M+2) & \dots & r(0) \end{bmatrix}, \quad (5.21)$$

and the correlation vector \mathbf{r} is given by

$$\mathbf{r}^T = [r(-1), r(-2), \dots, r(-M)]. \quad (5.22)$$

These equations are known as the Yule-Walker equations and are used extensively in time series analysis. In practice, the autocorrelation function values $r(k)$ are not known and need

to be estimated for lags $0, 1, \dots, M$. These estimates can be made using Equation 5.17. The negative lags can be calculated by $r(-k) = r^*(k)$.

Once the AR parameters have been calculated by Equation 5.19, the AR spectrum estimate can be computed using the following equation:

$$S(w) = \frac{\sigma_v^2}{\left|1 + \sum_{m=1}^M a_m^* e^{-jm w}\right|^2} \quad (5.23)$$

where $-\pi \leq w \leq \pi$, and

$$\sigma_v^2 = r(0) + \sum_{k=1}^M a_k r(k). \quad (5.24)$$

We will now apply AR spectrum estimation to the typical sea clutter data already presented in Figure 5.6. We use the Yule-Walker equations to calculate the parameters of an AR model of order six, for each of the time series that were used in the generation of Figure 5.6. (A sixth order AR model was determined to be suitable empirically). Then the AR spectrum estimates are evaluated and displayed along with the Fourier spectrum estimates. These spectra are shown in Figure 5.8. Figure 5.8 shows that the AR model is capable of tracking the variations in clutter spectra that are indicated by the nonparametric Fourier estimates.

In order to use the AR model to provide spectral moment estimates, the denominator in Equation 5.23 needs to be factorized to solve for the M poles that govern the AR all-pole transfer function. These M poles can then be plotted within the unit circle in the z -plane (we assume stability). It is easy to show that the AR spectrum described by Equation 5.23 has a geometric interpretation with respect to its poles in the z -plane. The unit circle in the z -plane represents the complex frequency axis where one revolution around the circle corresponds to w going from $-\pi$ to π . The AR spectrum is evaluated by setting z on the unit circle, and letting it make a single revolution. At each frequency, the spectrum is equal to the square of the inverse of the product of the distances from z to each pole, normalizing σ_v^2 to unity. The closer a pole is to the unit circle then, the larger will be the peak of the spectrum at the frequency determined by the phase of the pole, and

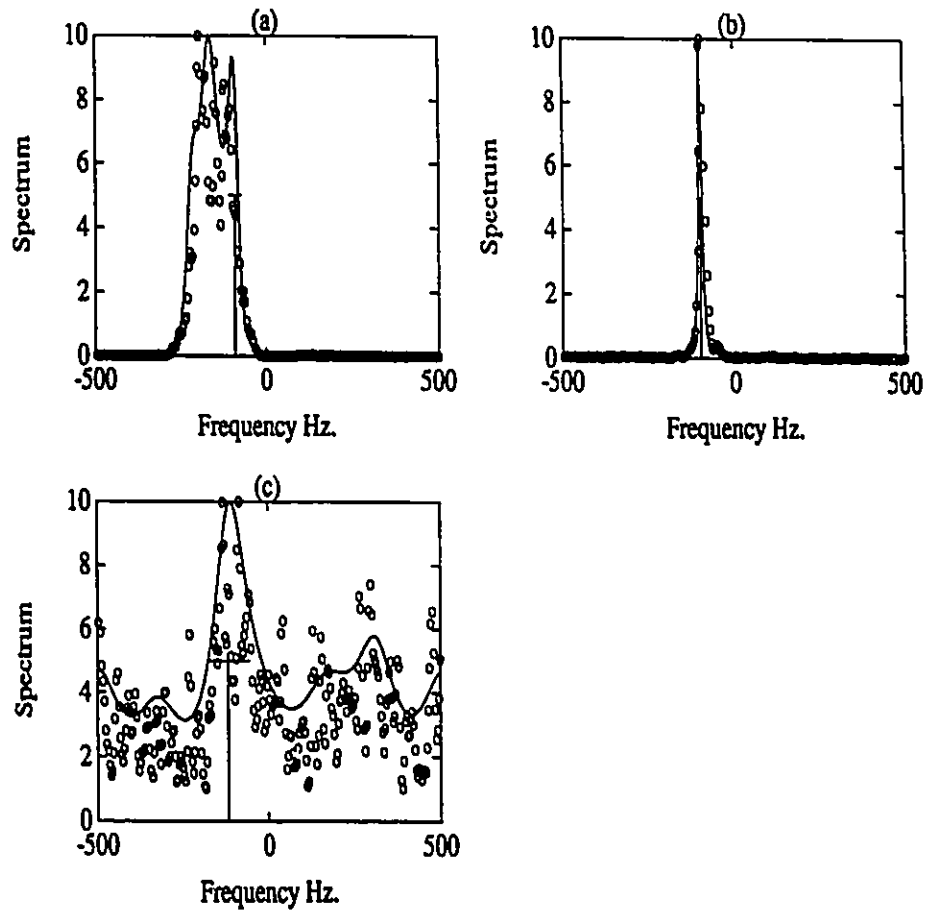


Figure 5.8: This figure shows both AR (order=6) Doppler spectrum estimates and Fourier Doppler spectrum estimates for typical sea clutter data. The Fourier estimates are indicated by 'o' and the AR estimates by solid lines. Also shown are AR mean frequency and spectral width estimates indicated by a 'T'. The 'T' is drawn at the mean frequency estimate and the width of the 'T' is the half-power spectral width estimate.

the narrower in width that peak will be. In other words each pole will cause a peak in the AR spectrum at the frequency specified by the argument of the pole and with a height and width controlled by the magnitude of the pole. In Figure 5.9, we show the AR poles that correspond to the spectral plots given in Figure 5.8.

The argument of the pole with the largest magnitude serves as the mean frequency estimate of the clutter Doppler spectra. If we denote the pole closest to the unit circle as $pe^{j\phi}$ where p is the magnitude of the pole, and ϕ is the argument of the pole, then the mean frequency estimate is given by

$$\hat{f}_{AR} = \frac{\phi}{2\pi T_s} \text{ Hz.} \quad (5.25)$$

An estimate of the spectral spread about the mean frequency can be derived if we assume that the other poles are negligible in their contribution to the spectrum near the mean frequency. From the geometric arguments given above, $S(\phi) \approx 1/(1-p)^2$. If we define the spectral width as the separation in frequency between the points on either side of the maximum, where $S(\phi)$ falls to one half of its maximum value, these points will occur along the unit circle a distance $2(1-p)^2$ from the pole. A triangle is formed with indices at the origin, the pole, and the each half-power frequency on the unit circle, and all three lengths are known (the third length is of course unity). The angle between the two sides originating at the origin represents one half of the spectral width. It can be solved for simply by using the Cosine Rule from trigonometry. The solution for this spectral width estimate denoted by $\hat{\sigma}_{AR}$ is

$$\hat{\sigma}_{AR} = \frac{2}{2\pi T_s} \arccos\left(\frac{4p - (1+p^2)}{2p}\right) \text{ Hz,} \quad (5.26)$$

and T_s is the sampling period. This relationship is plotted in Figure 5.10. Evaluating Equation 5.26 for the limiting case of $p = 1$ gives a spectral width of 0 Hz for a 1 kHz sampling rate. Equation 5.26 gives a spectral width of 1 kHz (the sampling rate) for $p = 0.1716$. For $p < 0.1716$, Equation 5.26 no longer provides meaningful estimates as the sampling rate is exceeded and the argument of the arccos function has a magnitude greater than one. The AR mean frequency and spectral width estimators defined in Equation 5.25




Figure 5.9: This figure shows the AR (order=6) poles associated with typical sea clutter data. The poles closest to the unit circle model the clutter spectral components, while the other poles are mostly modeling noise.

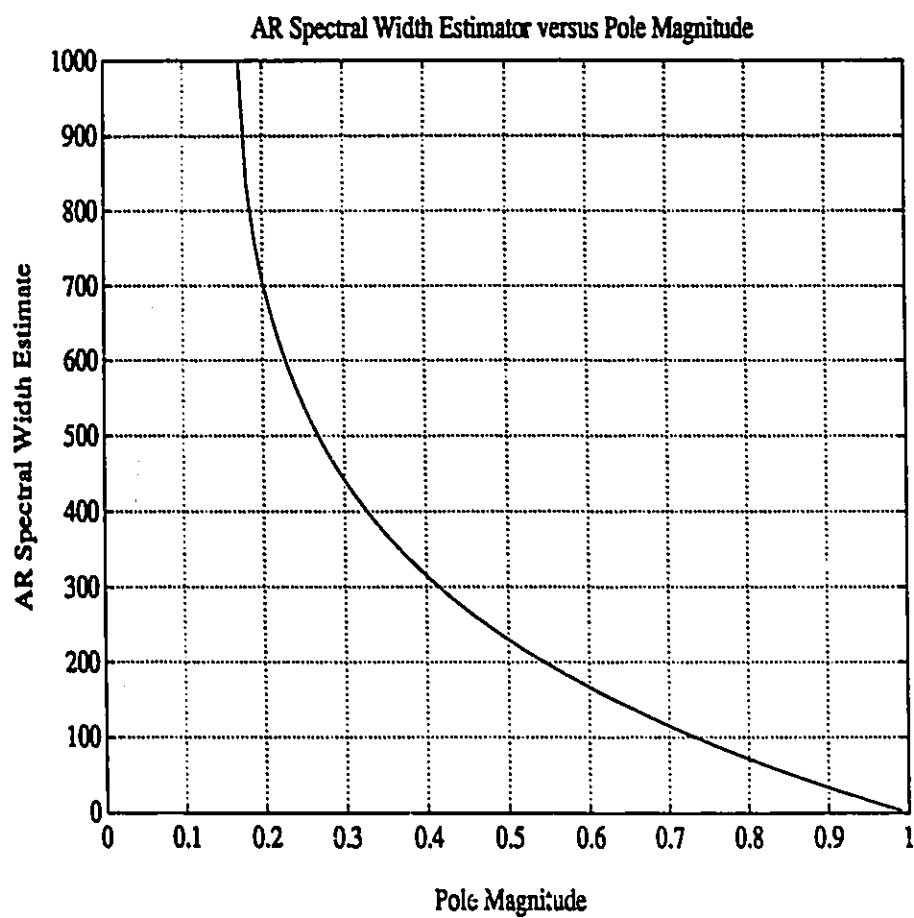


Figure 5.10: This figure shows the AR spectral width estimate as a function of the pole magnitude. The larger the pole magnitude, the smaller the spectral width. The sampling rate or PRF is assumed to be 1 kHz.

and Equation 5.26 are used to estimate mean frequency and spectral width for the typical clutter spectra given in Figure 5.6. These estimates are indicated in Figure 5.8 by a ‘T’ drawn with a height of 50 % of the peak spectral value.

5.4.4 Gaussian Spectrum Modeling

The final method that we consider for modeling the Doppler spectra as well as providing spectral moment estimators takes a more direct approach. Since we are after modeling the clutter Doppler spectra, as well as making spectral moment estimates, a spectral model which incorporates these parameters directly would seem ideal. The Gaussian model is parameterized appropriately for our needs as the mean frequency and spectral width are two of parameters that define the model. Furthermore, a Gaussian spectral model is justified from the form of the typical clutter spectra presented in Figure 5.6. Therefore, we propose to model typical sea clutter spectra as having a Gaussian spectral component plus a flat component which models the white noise floor of the receiver. The noise floor is estimated independently from the data or is provided, and it is denoted by N_o . Therefore, our problem reduces to *fitting* a one-Gaussian (1G) spectrum model to the data. Let $\mathbf{S}(\Theta)$ denote the Gaussian spectrum model vector, which is an N -vector evaluated at the N discrete frequencies $\mathbf{f} = [f_1, f_2, \dots, f_N]^T$. The parameter vector Θ is given by $\Theta = [a_1, f_1, \sigma_1]^T$ where a_1 , f_1 , and σ_1 are the amplitude, mean frequency and spectral width respectively of the model. Thus,

$$\mathbf{S}(\Theta_k) = a_1 e^{\frac{-(f_k - f_1)^2}{2\sigma_1^2}} + N_o \quad (5.27)$$

where $\mathbf{S}(\Theta_k)$ is the k^{th} element of $\mathbf{S}(\Theta)$ and f_k is the k^{th} element of \mathbf{f} , $k = 1, \dots, N$. We will assume that we have an estimate of the clutter Doppler spectrum given by $\hat{\mathbf{S}}(\mathbf{f})$, and evaluated at the same set of discrete frequencies defined above. The squared error between our model and the spectrum estimate is

$$\epsilon(\Theta) = |\hat{\mathbf{S}}(\mathbf{f}) - \mathbf{S}(\Theta)|^2. \quad (5.28)$$

The minimization of Equation 5.28 with respect to the parameter vector Θ yields the parameter estimates. This is a nonlinear, least squares optimization problem and the solution is treated in Appendix D. We have applied this optimization to the typical clutter data presented in Figure 5.6, and the results are shown in Figure 5.11. The 1G spectrum model provides a good fit to the clutter spectra, and thus provides good spectral moment estimates as well.

In this section, we have seen that the Gaussian spectral model outperforms the AR spectrum model and the pulse-pair model in providing the best estimates of the clutter spectral moments. This should not come as a great surprise as the Gaussian model has as two of its parameters, the first and second spectral moments. In the next section, we will perform a similar analysis on growler spectra.

5.5 Doppler Spectrum Estimates of Growler Returns

In the last section, we examined *medium dwell-time* Doppler spectrum estimates of sea clutter. As well as performing spectrum estimation, we also developed several spectral moment estimators. We follow a very similar procedure in this section, only we focus on radar resolution cells that contain growlers.

Consider the typical growler in sea clutter Doppler spectrum estimates shown in Figure 5.12. Just like sea clutter, growler spectra have a few common forms. The first and most common growler spectrum contains two components, one from the sea clutter in the radar cell, and the other from the growler. We have already seen various forms of the clutter component, and these are described in Figure 5.6. The growler component on the other hand has a very narrow spectral width of only a few Hz (4–6). For most of the data sets that we have examined, the amplitude of the clutter component is considerably smaller than the amplitude of the growler component, although as we shall see later, the growler to clutter ratio is not very large. There are several occasions however when the clutter

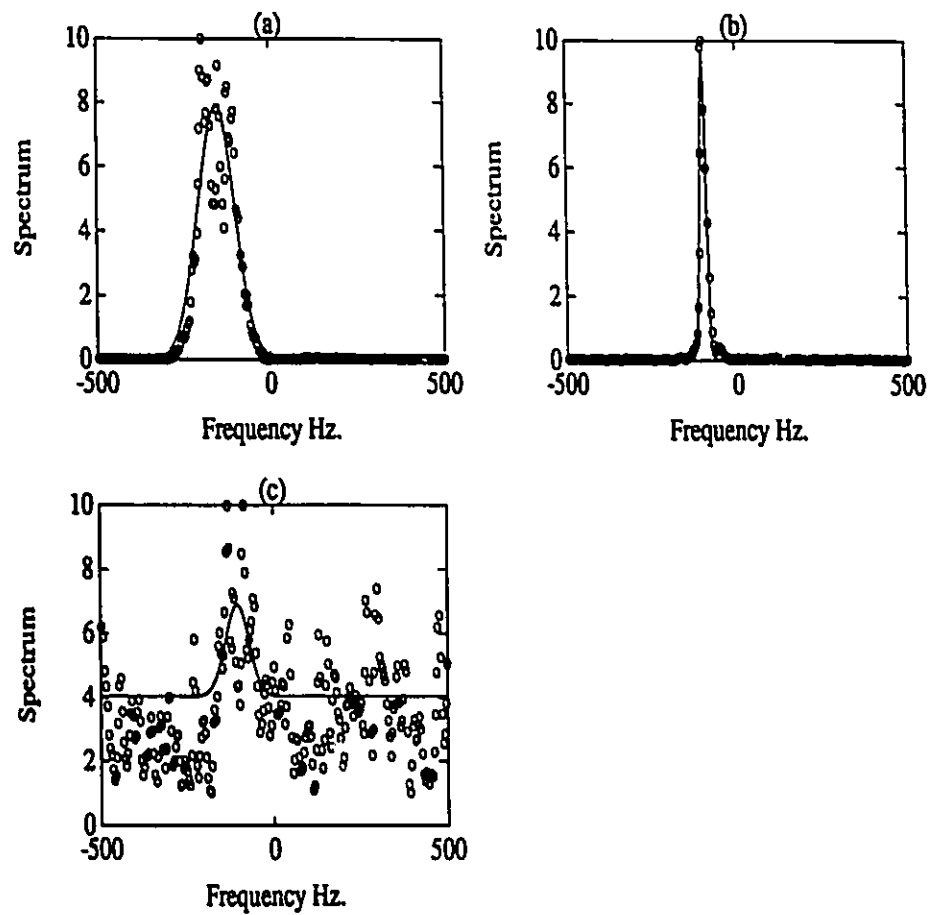


Figure 5.11: This figure shows Gaussian spectrum models fitted to the Fourier Doppler spectrum estimates for typical sea clutter data. The Fourier estimates are indicated by 'o' and the fitted Gaussian models by solid lines. The first and second spectral moments are given by the corresponding parameters of the Gaussian models.

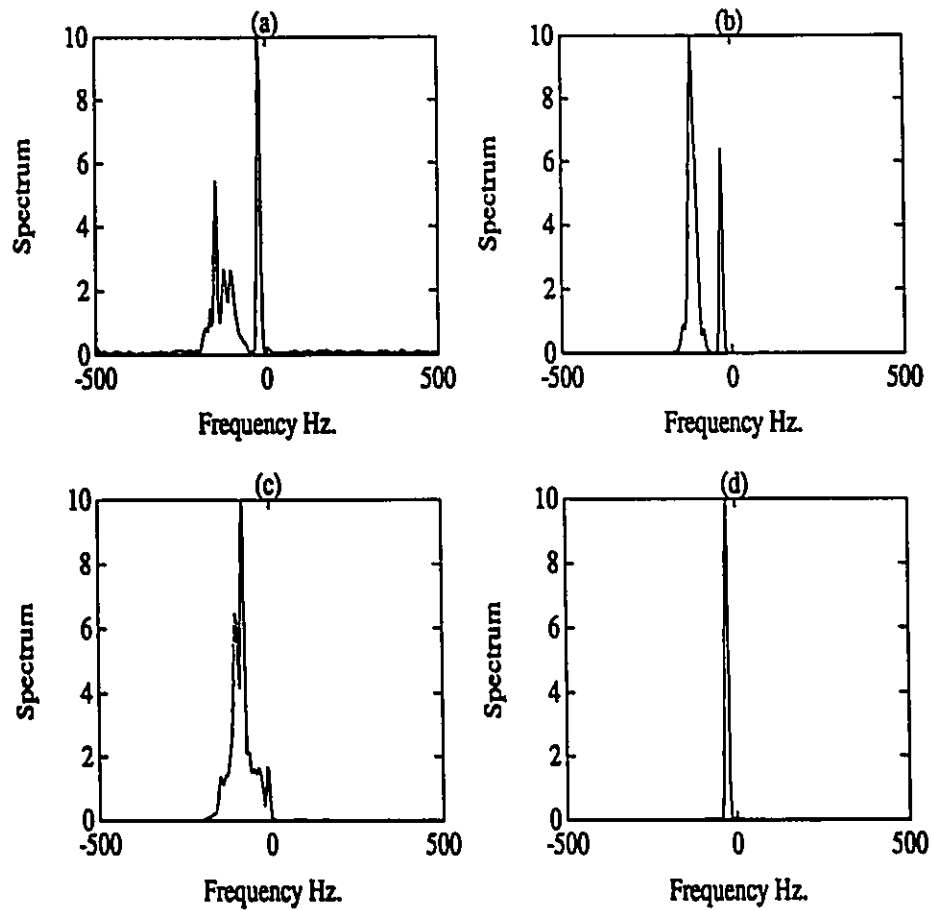


Figure 5.12: This figure shows typical Doppler spectrum estimates of a growler in sea clutter based on Fourier analysis of about 0.5 s of data. Most of the time, growler spectra appear like that in (a), with a wider bell-shaped component due to the sea clutter, and a taller narrower bell-shaped component due to the growler. Sometimes, the peak of the clutter component is larger than that of the growler as shown in (b). On other occasions, a growler component may not be present at all due to shadowing or submersion of the growler as shown in part (c). Finally, sometimes the clutter component can be nonexistent as shown in (d). The data shown here are from file B98.

component overpowers the growler component, or the growler component vanishes entirely due to shadowing or submersion. All of these cases are described in the nonparametric spectra presented in Figure 5.12.

5.5.1 Pulse Pair Estimators

We apply the same pulse pair estimates described in Equation 5.16 to the typical growler spectra given in Figure 5.12. The results of these estimates are shown in Figure 5.13. Figure 5.13 indicates the poor performance of the pulse pair estimators in the presence of growlers. When both a clutter and growler component are present in the spectra, the pulse pair mean frequency estimator tends to give a value somewhere midway between the mean values of each of the components. The spectral width estimator tends to overestimate the spectral width of either component due to the extended width of the two components. The poor performance of the pulse pair estimators is to be expected here because they were derived for unimodal spectra. Clearly, we must look at other spectral moment estimators for the growler in sea clutter case.

5.5.2 Autoregressive Spectrum Estimators

Now we discuss the application of the AR spectra and spectral moment estimates that were derived in Section 5.4.3, to the growler in clutter case. For this case, the same Yule-Walker equations described in Equation 5.19 are used. Then, the spectrum estimate is evaluated by substituting the AR coefficients into Equation 5.23. Again, by using a sixth order AR model, AR spectra are computed for each of the time series that were used in the generation of Figure 5.12. These spectrum estimates are displayed in Figure 5.14, along with the nonparametric Fourier spectrum estimates that were given in Figure 5.12. Figure 5.14 shows that the AR model is capable of tracking the variations in growler spectra that are indicated by the nonparametric Fourier estimates.

In order to use the AR model to provide spectral moment estimates, we factorize

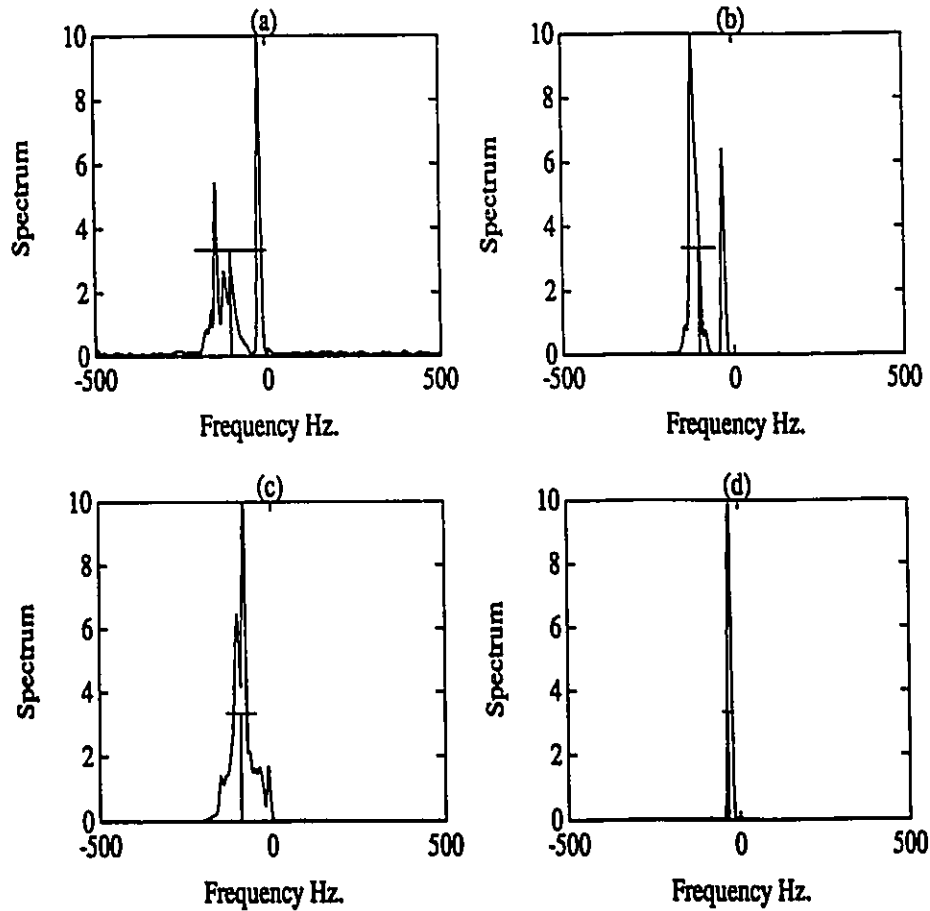


Figure 5.13: This figure shows the results of applying pulse pair estimators on typical growler in sea clutter data. The estimated Doppler spectra are shown as well as the pulse pair estimates in order to demonstrate the performance of the pulse pair estimates. A 'T' bar is drawn on each spectra to indicate the pulse pair mean frequency estimate as well as the spectral width estimate. The height of the 'T' bar is arbitrarily set to one third of the maximum height.

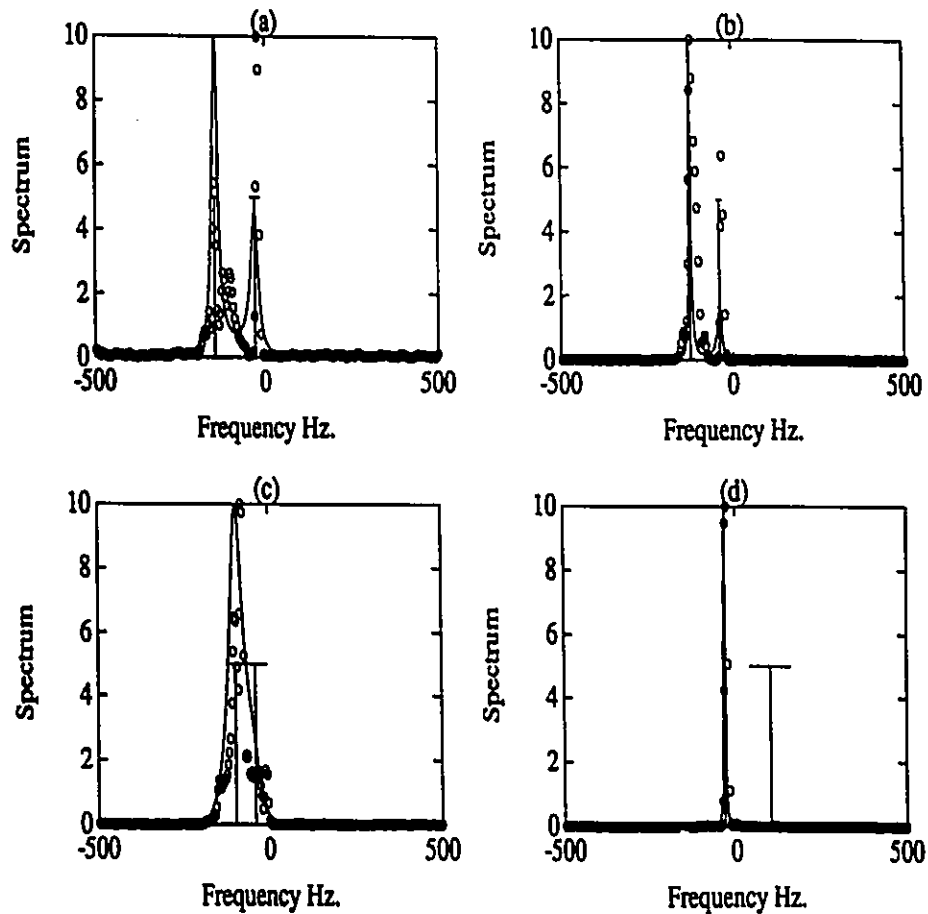


Figure 5.14: This figure shows both AR (order=6) Doppler spectrum estimates and Fourier Doppler spectrum estimates for typical growler data. The Fourier estimates are indicated by 'o' and the AR estimates by solid lines. Also shown are AR mean frequency and spectral width estimates indicated by a 'T'. The 'T' is drawn at the mean frequency estimate and the width of the 'T' is the half-power spectral width estimate. Since in general there are two components in the growler spectra, the mean frequency and spectral width estimates for both components are shown.

the denominator in Equation 5.23, as we did for the clutter-only case, and solve for the M poles that govern the AR all-pole transfer function. In Figure 5.15, we show the AR poles that correspond to the spectral plots given in Figure 5.14.

These M poles are plotted within the unit circle in the z -plane. As we saw in Section 5.4.3, the closer a pole is to the unit circle, the larger will be the peak of the spectrum at the frequency determined by the phase of the pole, and the narrower in width that peak will be. Therefore, we identify the pole closest to the unit circle as the growler pole, and the pole that is the second closest as the clutter pole. Here we are assuming that both a clutter and growler component are present in the spectra.

Once a clutter pole and growler pole have been identified, the same equations derived in Section 5.4.3 for spectral moment estimation are used with each pole. The AR mean frequency and spectral width estimators defined in Equation 5.25 and Equation 5.26 are used to estimate mean frequency and spectral width for the typical growler spectra given in Figure 5.12. These estimates are indicated in Figure 5.14 by a 'T' drawn with a height of 50 % of the maximum spectral value.

5.5.3 Gaussian Spectrum Modeling

The Gaussian model approach that was described in Section 5.4.4 for modeling sea clutter spectra and giving spectral moment estimates can be extended to the growler in sea clutter case. Since, for this case, a typical spectra has two bell-shaped components, a two-Gaussian (2G) model is used to model the growler-in-clutter spectra.

Let $S(\Theta)$ denote the Gaussian spectrum model vector, which is an N -vector evaluated at the N discrete frequencies $\mathbf{f} = [f_1, f_2, \dots, f_N]^T$. The parameter vector Θ is given by $\Theta = [a_1, \bar{f}_1, \sigma_1, a_2, \bar{f}_2, \sigma_2]^T$ where a_1 and a_2 , \bar{f}_1 and \bar{f}_2 , and σ_1 and σ_2 are the amplitudes, mean frequencies and spectral widths respectively of the 2G model. Thus,

$$S(\Theta_k) = a_1 e^{-\frac{(f_k - \bar{f}_1)^2}{2\sigma_1^2}} + a_2 e^{-\frac{(f_k - \bar{f}_2)^2}{2\sigma_2^2}} + N_o \quad (5.29)$$




Figure 5.15: This figure shows the AR (order=6) poles associated with typical growler in sea clutter data. The pole closest to the unit circle models the growler spectral component and the second closest pole to the unit circle models the clutter spectral components, while the other poles are mostly modeling noise.

where $S(\Theta_k)$ is the k^{th} element of $S(\Theta)$ and f_k is the k^{th} element of \mathbf{f} , $k = 1, \dots, N$. Like before, we assume that we have an estimate of the clutter Doppler spectrum given by $\hat{S}(\mathbf{f})$, and evaluated at the same set of discrete frequencies defined above. The squared error between our model and the spectrum estimate is

$$\varepsilon(\Theta) = |\hat{S}(\mathbf{f}) - S(\Theta)|^2. \quad (5.30)$$

The minimization of Equation 5.30 with respect to the parameter vector Θ yields the parameter estimates. This is a nonlinear, least squares optimization problem and the solution is treated in Appendix D. We have applied this optimization to the typical growler data presented in Figure 5.12, and the results are shown in Figure 5.16. The Gaussian spectrum model provides a good fit to the growler in sea clutter spectra, and thus provides good spectral moment estimates as well.

5.6 Spectral Moment Data for Sea Clutter and Growler

In Section 5.4 and Section 5.5 several estimators were derived for evaluating sea clutter and growler Doppler spectra, as well as their spectral moments. The pulse pair estimates, although they work reasonably well for mean-frequency estimation in clutter-only, they do not perform well for spectral width estimation or for growler in sea clutter data. The AR mean frequency and spectral width estimators work reasonably well in high growler to clutter ratio environments; however, they lose track easily when the growler component is weak. This is easily seen by way of an example. If the AR model is used to parameterize the spectra shown in the spectrogram image in Figure 5.2 on a raster by raster basis, then the AR mean frequency and spectral width estimates can be plotted for each raster as shown in Figure 5.17. Examination of Figure 5.17 shows that the AR estimator loses the growler track several times. This is due to the coupling between power and spectral width in the AR spectral width estimator. When the growler component is small relative to the clutter spectral component, the AR estimator tracks the clutter instead of the growler. Comparing

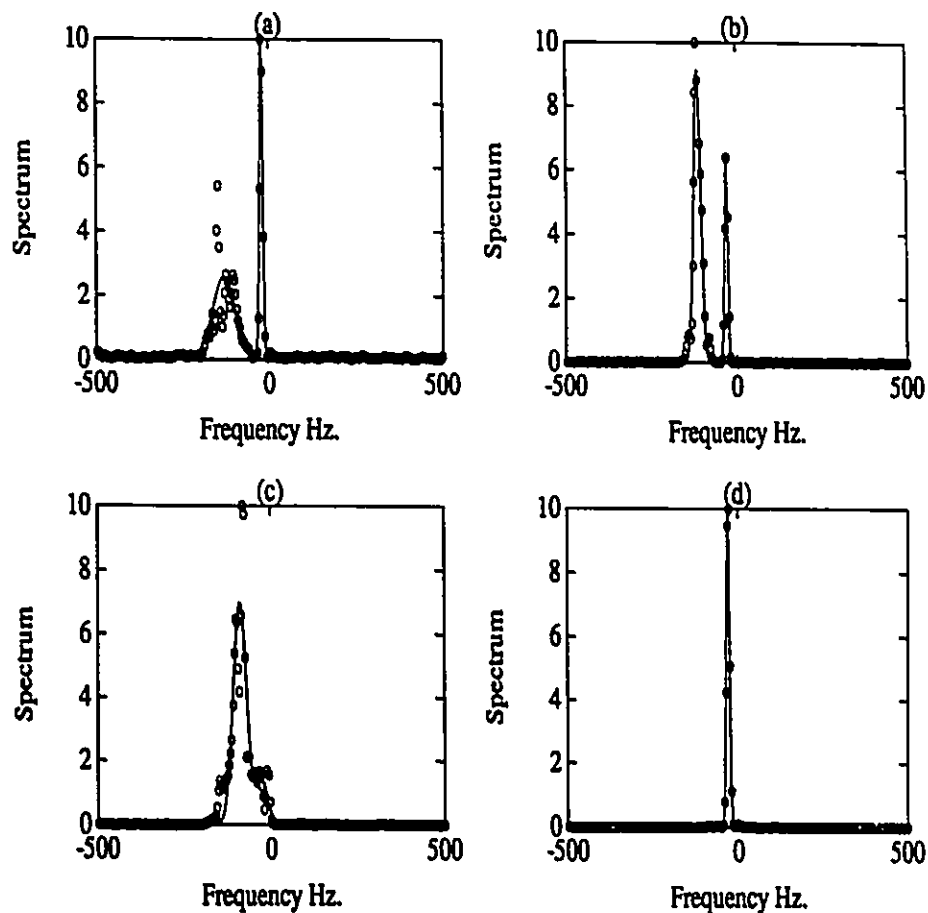


Figure 5.16: This figure shows Gaussian spectrum models fitted to the Fourier Doppler spectrum estimates for typical growler in sea clutter data. The Fourier estimates are indicated by 'o' and the fitted Gaussian models by solid lines. The first and second spectral moments are given by the mean and standard deviation parameters of the Gaussian models.




Figure 5.17: This figure shows the AR mean frequency and spectral width estimates for the growler spectrogram shown in Figure 5.2. The upper graph shows the mean frequency estimates as a function of time and the lower curve gives the corresponding spectral width estimates. The AR mean frequency estimator loses track when the growler spectral component is small relative to the clutter. At these times, the clutter pole is closer to the unit circle and hence the frequency and spectral width of the clutter component is tracked by the estimator, rather than those of the growler.

Figure 5.17 with the growler to clutter ratios shown in Figure 5.22 confirm that the AR estimators lose track whenever the growler to clutter ratio is small.

The Gaussian models on the otherhand perform very well in both clutter-only and growler in sea clutter cells. The 1G model is used for the clutter-only case and the 2G model is used for a growler cell. Clearly, the 1G model is just a special case of the 2G model, where one of the Gaussian amplitudes is zero. In this section, we will use the 2G model to study the spectral moments of growlers and sea clutter in some depth.

Reconsider the growler cell spectrograms shown for the HH and VV channels in Figure 5.2 and Figure 5.3 respectively. For the most part, these images show both a clutter component as well as growler component in the radar cell. We will use the 2G spectrum model to evaluate the spectral moments for each raster shown in the images. Therefore, we should be able to quantify and track the variations in mean frequency and spectral width of the growler and clutter as a function of time, for both polarizations HH and VV.

If Figure 5.18 and Figure 5.19, we present the clutter mean frequency and spectral width estimates for the HH and VV channel respectively. Figure 5.20 and Figure 5.21 give the corresponding growler moments. (Note: The 2G model measures amplitude as well as mean frequency and spectral width. Sometimes, the growler or clutter component is virtually nonexistent in the data due to shadowing and/or submersion. Therefore, if the estimated amplitude of either the growler or clutter component is comparable to the estimated noise level, the 2G algorithm disregards the estimates. This explains the gaps of data points that are seen in Figure 5.18 to Figure 5.21.)

Several comments are in order. First, we see that the 2G model is capable of providing good estimates of spectral moments of growler in sea clutter data. Both growler components and clutter components are modelled well. Second, the 2G model has successfully tracked the mean frequency variation of the growler that describes the *winking phenomenon*. Third, the small spectral widths of the growler component contrast the larger spectral widths associated with the clutter component. This spectral width difference is




Figure 5.18: This figure shows mean Doppler frequency and spectral width estimates for sea clutter as a function of time. A 2G spectrum model was fit to Fourier spectral estimates in order to estimate the spectral moments. The upper graph shows the mean frequency estimate as a function of time, and the lower graph shows the corresponding spectral width estimates. These estimates are determined from the HH channel of file B98.




Figure 5.19: This figure shows mean Doppler frequency and spectral width estimates for sea clutter as a function of time. A 2G spectrum model was fit to Fourier spectral estimates in order to estimate the spectral moments. The upper graph shows the mean frequency estimate as a function of time, and the lower graph shows the corresponding spectral width estimates. These estimates are determined from the VV channel of file B98.




Figure 5.20: This figure shows mean Doppler frequency and spectral width estimates for a growler as a function of time. A 2G spectrum model was fit to Fourier spectral estimates in order to estimate the spectral moments. The upper graph shows the mean frequency estimate as a function of time, and the lower graph shows the corresponding spectral width estimates. These estimates are determined from the HH channel of file B98.




Figure 5.21: This figure shows mean Doppler frequency and spectral width estimates for a growler as a function of time. A 2G spectrum model was fit to Fourier spectral estimates in order to estimate the spectral moments. The upper graph shows the mean frequency estimate as a function of time, and the lower graph shows the corresponding spectral width estimates. These estimates are determined from the VV channel of file B98.

easily seen in the Figure 5.2 and Figure 5.3. Fourth, the variation in the mean frequency estimate of the clutter component is much greater and less deterministic looking than the corresponding variation in the mean frequency estimate of the growler component. This difference is related to the *winking phenomenon*. Finally, in comparing the HH and VV spectral moments of sea clutter given in Figure 5.20 and Figure 5.21, a few relationships seem apparent. First, the mean frequency shift in the VV channel is smaller than the corresponding shift in the HH channel. Second, the spectral widths associated with the clutter component in the VV channel seem to be larger than those in the HH channel. That is, although the mean translational velocity of the sea appears to be smaller in the VV channel, the sea appears to be more *confused* implying a broader spectral width. These results agree with those described in [68]. In Table 5.1, the mean clutter frequency shifts and mean clutter spectral widths are given for each of the files that we have analyzed. For all of these files, the spectral width estimated in the VV channel exceeds the spectral width estimated in the HH channel. However, it is difficult to see a general relationship between the mean frequency shifts in the two channels.

5.7 Growler and Sea Clutter Power Ratios

Power ratios of growlers and sea clutter are very important for characterizing the relative size of target and clutter. In this section, we calculate these quantities and we provide a means to determine absolute powers.

Three different ratios are important for characterizing the relative power levels of target and clutter signals. We consider the growler return to be the desired signal. We call the ratio of average growler power to noise power the growler to noise ratio (GNR). The ratio of the average clutter power to noise power is called the clutter to noise ratio (CNR). Finally, the ratio of the average growler power to average clutter power is called the growler to clutter ratio (GCR). We are interested in providing estimates for all of these

ID	POL (h,v)	CLUTTER MEAN FREQUENCY (Hz)	CLUTTER MEAN SPECTRAL WIDTH (Hz)
B97	h	-100.3	31.5
B98	h,v	-107.1, -79.9	26.1, 37.4
B99	v	-77.6	34.7
B110	h	-86.9	25.4
B111	v	-76.3	30.5
B112	h,v	-20.1, -61.7	21.5, 32.8
B113	h	-50.6	21.6
B114	v	-38.2	23.5
B115	h,v	-14.7, -24.3	13.3, 23.4
B123	h	2.4	9.8
B124	v	-3.5	14.9
B125	h,v	5.0, -1.7	8.9, 14.3

Table 5.1: This table shows the mean clutter frequency and spectral width estimates for each of the files that were used in the coherent analysis of sea clutter and growler returns. These estimates were calculated from the parameters of a 2G model that was applied to estimated spectra of the data.

ratios. Clearly, this involves estimating the growler and clutter powers.

In this chapter, we have examined ways to model the spectra of clutter and growlers in sea clutter. One of the benefits of Doppler spectra estimation and modeling is the separation of growler and clutter components in the received signal. We have seen that the 2G model accurately models the growler-in-clutter spectra, and that a separate Gaussian component in the model is used to characterize the growler and clutter portions of the spectra. If the first Gaussian component is identified with the clutter and its parameters are a_1 , \bar{f}_1 , and σ_1 , then the clutter power is given by the area under the Gaussian curve which is easily shown to be $a_1\sqrt{2\pi}\sigma_1$. Similarly, the growler power is given by $a_2\sqrt{2\pi}\sigma_2$. If the average white noise level is N_o , and the pulse repetition frequency is denoted as PRF, then the following power ratios result:

$$\text{CNR} = \frac{a_1\sqrt{2\pi}\sigma_1}{N_o\text{PRF}} \quad (5.31)$$

$$\text{GNR} = \frac{a_2\sqrt{2\pi}\sigma_2}{N_o\text{PRF}} \quad (5.32)$$

$$\text{GCR} = \frac{a_2\sigma_2}{a_1\sigma_1} \quad (5.33)$$

Using the 2G model, each raster shown in Figure 5.2 is parameterized and Equation 5.33 is used to estimate the GCR. These GCR estimates are shown plotted in Figure 5.22. The histogram of these GCR estimates is also shown in the figure. The average GCR is computed to be 7.3 dB for this growler. Similar calculations were carried out for each of the files specified in Table 3.3. The GCR for each of these files is given in Table 5.2.

By making use of the power ratio calculations described above, it is possible to estimate the RCS of growlers and sea clutter. Since the data that we describe in this thesis were collected while operating in the linear region of the receiver, we can use the simple procedure described in Appendix B, Section B.2, to map the output powers onto the input powers. We use an input noise power at the antenna of -92 dBm, as given in Section B.2. Therefore, adding the output CNR and GNR to this noise level gives the input clutter

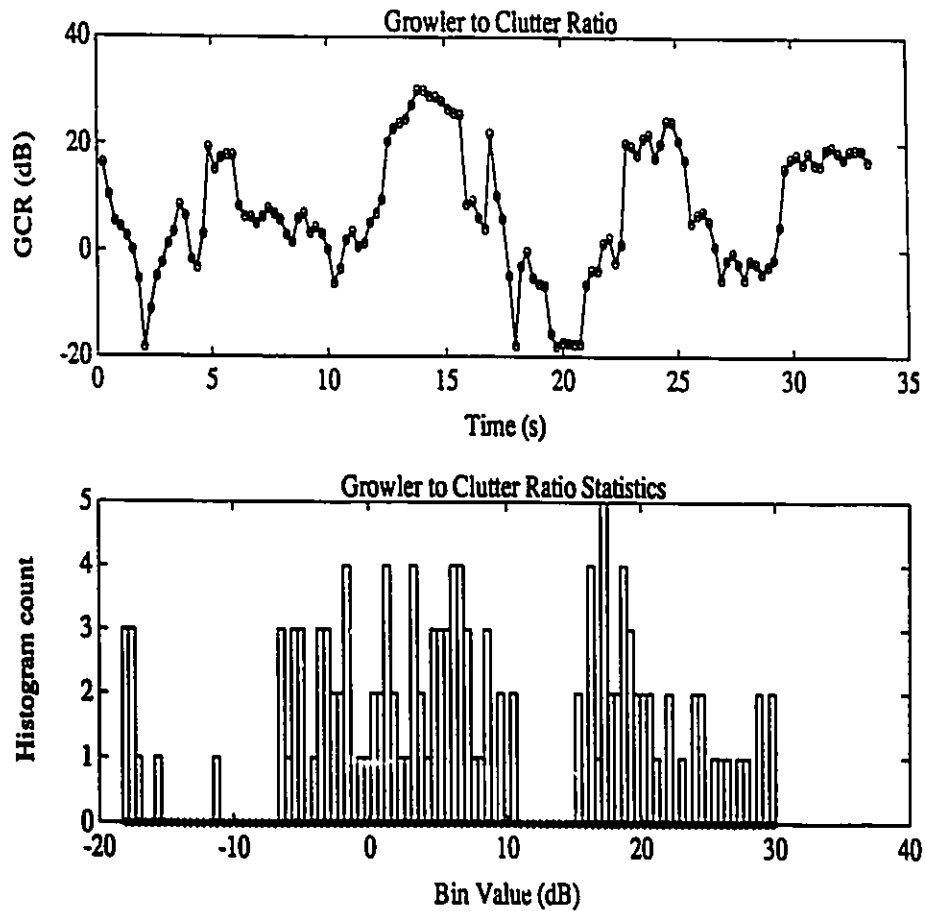


Figure 5.22: This figure shows GCR estimates as a function of time, calculated from the parameters of a 2G model applied to growler-in-clutter Fourier spectral estimates. In the upper graph, the GCR estimate is given as a function of time, and in the lower graph, the corresponding histogram of the estimates are given. These estimates are determined from the HH channel of file B98.

ID	GCR (dB)	POL (h,v)	GROWLER RCS (m^2)	CLUTTER RCS (dB- m^2/m^2)	GRAZING ANGLE (degrees)
B97	8.9	h	7.05	-41.6	0.188
B98	7.3,8.8	h,v	6.52,5.40	-38.8,-41.9	0.193
B99	14.7	v	7.75	-45.9	0.196
B110	3.6	h	1.25	-38.0	0.295
B111	2.7	v	0.14	-45.5	0.296
B112	18.1,3.9	h,v	2.85,0.17	-45.1,-45.4	0.298,0.300
B113	7.1	h	0.36	-45.0	0.304
B114	14.3	v	1.63	-48.0	0.302
B115	17.2,9.9	h,v	4.92,0.71	-40.3,-47.4	0.301
B123	20.5	h	9.77	-39.6	0.275
B124	10.0	v	0.67	-47.5	0.274
B125	11.9,6.8	h,v	4.72,0.35	-36.4,-46.9	0.272

Table 5.2: This table shows the GCR, growler RCS, and sea clutter RCS estimates for each of the files that were used in the coherent analysis of sea clutter and growler returns. The estimates were calculated from the parameters of a 2G model that was applied to estimated spectra of the data.

power and growler power, respectively. Then Equation 4.6 and Equation 4.5 can be used to determine the clutter RCS and growler RCS, respectively. This procedure has been carried out with each of the files listed in Table 3.3. The corresponding RCS estimates are given in Table 5.2. These data will be called upon in Chapter 7 when the performance of several detectors on these data files are compared.

5.8 Summary

In this chapter, we have examined some multidimensional characteristics of sea clutter and growler data. In particular, we have focussed on the coherent properties of sea clutter and growler radar data received from the HH and VV like-polarized channels. We began by examining the time-varying Doppler characteristics of growlers in sea clutter, and were able to verify the discovery of the *winking phenomenon* by generating time-frequency maps of the growler return signal. An analysis of the hydrodynamics of a floating growler in the sea showed that the *winking phenomenon* is the result of the cyclical motion patterns of water fluid particles in waves. Floating bodies that are much smaller than the wavelength of the underlying waves, were shown to behave in the same manner as water particles, thus explaining the time-varying Doppler signatures that our coherent measurements indicate.

We proceeded by examining the characteristics of sea clutter and growler spectra estimated from a fraction of a second of data, with the goal of developing *medium dwell-time* Doppler detection techniques. Several Doppler spectrum estimation and moment estimation techniques were considered in order to parameterize the spectral data. Both an AR based technique and a Gaussian modeling technique were shown to represent well the second order characteristics of sea clutter and growler data. These techniques were used to parameterize the data and the parameters were used to estimate the mean frequency and spectral width of the data spectra, as a function of time. The very small spectral widths associated with the growler spectra, and the larger spectral widths associated with the clutter data offer the

potential of reliably classifying received data as being either clutter or growler-in-clutter. Detection techniques based on exploiting the spectral width differences between sea clutter and growler returns will be dealt with in the next chapter.

Finally, we used the spectral modeling described above to separate the energy associated with growlers and sea clutter returns from the same radar resolution cell. Equations were developed that allowed GCR, GNR, CNR and RCS estimates to be made from the spectral parameters associated with the data. These estimates allow different radar data files to be compared by virtue of the relative and absolute strengths of growler and clutter, and will be important in evaluating the performance of a detector on a given file.

Chapter 6

Detector Philosophy and Design

6.1 Introduction

In the last two chapters, we studied the noncoherent and coherent properties of radar returns from growlers in sea clutter and sea clutter alone. We are now ready to consider the design of growler detectors. These detectors fall into two basic categories, noncoherent detectors and coherent detectors. The noncoherent detectors operate on the amplitude of the radar echo signal from a given radar resolution cell, and decide whether the received signal is most likely due to clutter or growler plus clutter. The coherent detectors operate jointly on the I and Q samples in order to decide whether or not a growler is present. We will make use of the detection and estimation theory presented in Chapter 2 in our development of the various detectors.

As we saw in Chapter 5, the growlers are more visible in the Doppler domain than they are in the amplitude domain. This should come as no surprise since we expect improved growler detection with coherent detectors. However, to say that coherent detectors are better for detecting growlers is a bit naive. It is important to understand the engineering tradeoffs that are associated with the detector designs that are proposed. Clearly, the most obvious tradeoff in using coherent detectors is the increase in cost and complexity of the

radar transmitter and receiver/processor. Whether or not this added cost is justified will depend on the application at hand and on the added performance that results from the increased cost and complexity. In Chapter 7, we examine in detail the growler detection performance that results from each of the detector designs proposed in this chapter.

Recall that we are committed to developing *medium dwell-time* techniques that offer improved detection over conventional noncoherent detectors. Therefore, our processors will operate on a fraction of a second of data before making a decision. In Chapter 5, we integrated about 0.5 s of data for each coherent Doppler spectrum estimate. This integration time was no accident. As we have seen, typical growler spectral widths are about 4-6 Hz. With a PRF of 1 kHz which is typical of our measurements, 0.5 s of integration results in a frequency resolution of about 2 Hz per FFT frequency bin. Therefore the growler spectral component will span over a few frequency bins. We can in fact get away with an integration time of about 1/8 s but Doppler spectral estimates will have higher variances. Using integration times of about 0.5 s is a reasonable compromise that is still considered to be *medium dwell-time*.

In order to compare the performance of different detectors, it is necessary to compare them on an equal basis. This can be difficult because of the number of different factors that any comparison can be based upon. We have decided that the fairest basis for comparison is the amount of information that the detector processes before making a decision. For example, information can be gathered from a given resolution cell by collecting several consecutive radar pulses (dwelling) and then stepping the antenna to the next azimuth position, or equally by scanning the entire region of interest and processing pulses from a given range cell on a scan-to-scan basis. Assuming the same PRF in either case, the total number of pulses integrated from a given radar resolution cell will then determine the amount of information processed. Therefore the detectors developed in this chapter will all use the same number of pulses for each detection. With this common basis established, several noncoherent and coherent detector designs will now be considered.

6.2 Noncoherent Detectors

Although we expect noncoherent detectors to perform considerably worse than coherent detectors, they are still very useful to examine. They will allow us to quantify conventional marine radar performance in detecting growlers and they will provide us with a benchmark with which to compare the improved performance of coherent detectors. Before rushing into the design of noncoherent detectors that operate on the amplitude (or power) of the radar echo signal, we should develop a basic understanding of their operation.

Fundamentally, noncoherent detectors are based on a very simple premise: the received radar echo from a given resolution cell is the sum of the returns from the clutter in the cell and the target in the cell. The clutter power signal (we discuss power since powers add by superposition) by itself can be thought of being composed of a mean clutter level m_{cl} and a varying component v_{cl} . Similarly, the growler echo can be thought of being composed of a mean growler power level m_{gr} and a varying component v_{gr} . Therefore, the radar echo from a sea clutter cell containing a growler has a mean level of $m_{cl} + m_{gr}$ and the varying component is $v_{cl} + v_{gr}$. A noncoherent detector forms an estimate of the mean signal level and decides whether the estimate is closer to m_{cl} or $m_{cl} + m_{gr}$. If the estimate is closer to m_{cl} , then clutter is declared; otherwise a growler is declared.

The performance of noncoherent detectors is solely a function of how close the estimate of the mean signal level is to the true value. If we could make perfect estimates of the mean signal level, then regardless how small the growler return m_{gr} is, we could pick a threshold between m_{cl} and $m_{cl} + m_{gr}$ and perfect detection would result. In practice, we cannot make perfect estimates unless we have an infinite amount of data. Therefore, noncoherent detection performance is limited by the accuracy of the mean signal level (power or amplitude) estimates. The accuracy of these estimates can be described by the estimate bias and variance. These, in turn, are functions of the statistical distributions that govern v_{cl} and v_{gr} and the manner in which the processor performs its mean level

estimation. Better estimates are always formed when more samples are integrated by the detector (so long as the underlying statistical processes haven't changed), independent of the detector design. Furthermore, depending on the behaviour of the varying terms v_{cl} and v_{gr} , one strategy for combining the received samples may be superior to another. It is therefore natural to consider different ways in which the received samples can be integrated or combined for noncoherent detection.

6.2.1 Amplitude Detectors

The simplest form of noncoherent detection uses the amplitude information of a single echo pulse to compute an estimate of the mean amplitude level. If the size of the varying terms v_{cl} and v_{gr} are small relative to the mean level components, the performance of single-pulse amplitude detectors is quite good. The noncoherent analysis of sea clutter data given in Chapter 4 showed that sea clutter is K-distributed. We saw that the K-distribution models the spiky nature of sea clutter backscatter. That is, very large values of sea clutter returns are common, relative to the mean amplitude level. In fact, examination of Figure 4.15 reveals that sea clutter samples that are greater than five times the mean amplitude level are common. It is these large sea clutter levels that cause conventional marine radars to perform poorly in detecting growlers. Furthermore, the compound nature of the K-distribution implies that the clutter varying term v_{cl} is actually composed of two varying terms, each with a different correlation time: a fast speckle component and a much slower underlying local mean level component.

Consider the radar amplitude echos shown in Figure 6.1 and Figure 6.2. Figure 6.1 shows the received amplitude level as a function of time from a range gate in file B98 (see Table 3.3 and Table 3.4 for radar parameters). The *spiky* nature of sea clutter returns is clearly visible. The mean amplitude level of the clutter return is much less than 20 while spiky bursts persist for several seconds throughout the record, with amplitudes exceeding 100. In Figure 6.2, a similar description is given of the growler returns from file B98.




Figure 6.1: This figure shows the amplitude of the radar echo as a function of time, from a clutter cell in file B98. The upper graph shows the time history of the received clutter amplitude. The spiky nature of the sea clutter returns is immediately visible. The lower graph shows the histogram of amplitude data given in the upper graph. The clutter amplitude statistics are K-distributed.




Figure 6.2: This figure shows the amplitude of the radar echo as a function of time, from the growler cell in file B98. The upper graph shows the time history of the received growler amplitude. The lower graph shows the histogram of amplitude data given in the upper graph.

The mean amplitude level from the growler cell is only about 35, implying that amplitude thresholding will result in high false alarm rates. By varying an amplitude threshold from zero to the maximum received amplitude, a receiver operating characteristic (ROC) can be calculated which quantifies the detection performance that would be achieved with this data. The ROC calculated from the data presented in Figure 6.1 and Figure 6.2 is shown in Figure 6.3. The poor performance that results using single pulse amplitude detection is evident. For example, to achieve a probability of detection of about 90 %, a false alarm rate of 50 % must be endured. That is, one out of every two clutter cells will be misclassified as a growler! This kind of performance is completely unacceptable for any practical problem such as ship navigation in ice infested waters.

Since our coherent processors will be given about 0.5 s of data upon which to base a decision, we now examine the improvement afforded by the integration of amplitude pulses over a similar period. As discussed above, integration of amplitude pulses should reduce the variance of the mean amplitude level estimates and therefore improve detection performance. If 512 (0.512 s at 1 kHz PRF) consecutive pulses are integrated (averaged) to form mean level estimates from the amplitude data presented in Figure 6.1, then Figure 6.4 results. Recall our discussion in Chapter 4 of the compound nature of K-distributed sea clutter. It is clear from the integrated amplitude plot in Figure 6.4 that we have only averaged out the speckle component of the clutter signal, but the underlying component with the longer correlation time still persists. The overall variance of the clutter amplitudes relative to the mean amplitude level is reduced, so some improvement in detection performance can be expected.

The integrated amplitudes for the growler data shown in Figure 6.2 are presented in Figure 6.5. Here too, the variance in the amplitude of the growler return has been reduced. In Figure 6.6, the ROC curve is given for the integrated amplitude data shown in Figure 6.4 and Figure 6.5. The ROC curve from Figure 6.3 is also redrawn here for comparison. The integrated amplitude ROC is given by the dashed line and the single pulse amplitude ROC

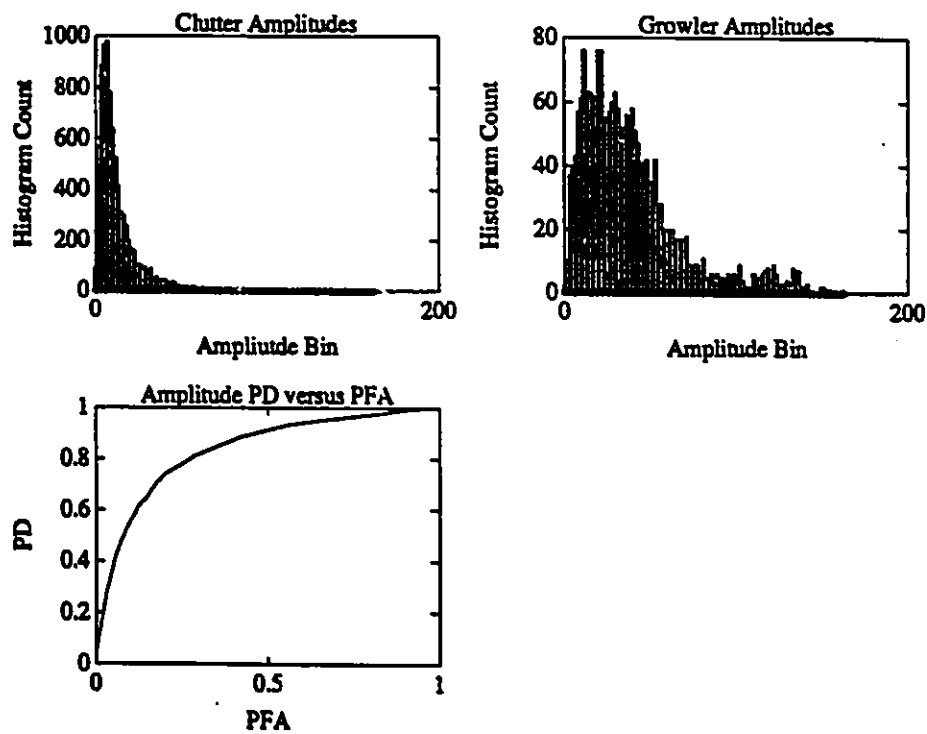


Figure 6.3: This figure shows the single pulse clutter and growler amplitude statistics for the file B98. Also, the ROC for this data is given. The probability of detection PD is shown plotted against the probability of false alarm PFA.




Figure 6.4: This figure shows the 512-pulse integrated amplitude of the radar echo as a function of time, from a clutter cell in file B98 (HH). The upper graph shows the time history of the integrated clutter amplitude. The local underlying mean level component described by the compound K-distribution is clearly visible. The lower graph shows the histogram of amplitude data given in the upper graph.




Figure 6.5: This figure shows the 512-pulse integrated amplitude of the radar echo as a function of time, from the growler cell in file B98 (HH). The upper graph shows the time history of the integrated growler amplitude. The lower graph shows the histogram of amplitude data given in the upper graph.

by the solid line in Figure 6.6. A modest improvement in the probability of detection can be observed for a probability of false alarm greater than about 4 %. For smaller probabilities of false alarm, the probability of detection estimates are not statistically significant due to the small number (130 for the clutter and 130 for the growler) of integrated amplitude measurements that were used in the calculation. We call this detector the IA (Integrated Amplitude) detector since it simply integrates amplitude pulses and compares the sum to a threshold.

The modest improvement afforded by the IA detector is a result of the reduction of the speckle component of the sea clutter amplitude echo. Since the speckle component has a decorrelation time of about 10 ms, integrating over 0.512 s reduces the speckle component significantly since many independent samples (approximately 50) are involved in the integration. The underlying local mean level component is not reduced though. This is because the underlying component has a correlation time on the order of several seconds (see Chapter 4). As discussed briefly above, the nature of the varying component v_{cl} will have a strong effect on the detection performance possible with a given detector. In order to achieve further gains in noncoherent detection performance, it is necessary to reduce the variance of v_{cl} further. Since the speckle component of v_{cl} has already been reduced, further reduction in the variance of v_{cl} can come about only by reducing the variance of the underlying local mean level component. It seems reasonable then that further improvements can be made by integrating amplitude pulses whose underlying local mean levels are independent.

A detector that uses scan-to-scan integration seems ideal for reducing the underlying local mean level of sea clutter backscatter. For example, if the scan rate of the antenna is 30 rpm, then each radar resolution cell would be revisited every 2 s. Integration of amplitude samples separated 2 s in time should reduce the variance caused by the underlying local mean level component of v_{cl} since the correlation between consecutive samples would be smaller. Furthermore, within each scan of the antenna, about 10 pulses will illuminate each radar




Figure 6.6: This figure shows the ROC for integrated amplitudes of clutter and growler data. The ROC was determined from the growler and clutter data in file B98 (HH) by integrating 512 amplitude pulses to form each test statistic. The solid curve corresponds to the single pulse amplitude ROC, and the dashed curve corresponds to the integrated amplitude ROC.

resolution cell assuming a PRF of 1 kHz and an antenna beamwidth of 2 degrees. These 10 pulses can be noncoherently integrated to reduce the speckle component of v_{cl} . Therefore a scan-to-scan integrating amplitude detector should be capable of reducing the v_{cl} term significantly. Unfortunately, our data base on growlers is limited and integration over tens of seconds to form a single test statistic results in a small number of test statistics. These test statistics are too few in number to make reasonable estimates of detection performance for small probabilities of false alarm. Although at this time we are unable to test a detector such as this empirically, we can make some comments on the order of improvements that one could expect with scan-to-scan type processors, by looking at the long term temporal correlation properties of sea clutter and growlers.

The temporal correlation for a sea clutter and growler cell from file B100 have been evaluated and are shown in Figure 6.7. The graphs plot the temporal correlation of the growler and clutter returns over time lags up to 60 s. These plots show several relationships between m_{cl} , m_{gr} , v_{cl} , and v_{gr} . The speckle component and the underlying local mean level component in v_{cl} are both visible in clutter correlation plot. After a fraction of a second in time, the clutter correlation has been reduced significantly, largely due to the decorrelation of the speckle component. Although there continues to be a reduction in the correlation as the temporal lag increases, the rate of correlation reduction is small. This indicates that scan-to-scan integration will provide better estimates of the clutter mean level, but the improvement will be small with increased scan time. The growler correlation plot shows a similar two component behavior. The correlation is characterized by a steep reduction over a period of a fraction of a second, which is largely due to the correlation reduction of the clutter in the same resolution cell as the growler. For longer lag times, the correlation of the growler return continues to fall but at a very slow rate. Again, these characteristics indicate that improvements afforded by scan-to-scan integration over periods of tens of seconds will offer some additional improvement, but the improvement will not be drastic with increased integration time.




Figure 6.7: This figure shows the temporal correlation of a growler plus clutter return and a clutter return from file B100. The polarization channel is HH. The correlations are calculated up to a temporal lag of 60 s. The growler plus clutter correlation at lag zero is about 5600 units, while for the clutter, the correlation at lag zero is about 1600 units. Therefore the GCR is about 3.5 in this case.

If scan-to-scan strategies are to be used, Figure 6.7 deserves further mention. The growler correlation is seen to be periodic with lag time. This period is controlled by the period of the dominant sea waves. Figure 6.7 indicates a period of about 10 s which agrees with the peak period reported by the nondirectional waverider, during the collection of file B100. This periodicity occurs as a result of the dynamics of a floating growler (see Section 5.3.1). When a large wave lifts and drops the growler, scintillation and shadowing results. Therefore, scan-to-scan strategies should incorporate the effects of beating with the wave period when trying to detect floating targets.

6.2.2 CFAR Detectors

In Section 6.2.1, we examined a few noncoherent integration strategies. Our goal was to provide good estimates of m_{cl} and m_{gr} , the mean signal level of sea clutter and growler returns respectively, so that a detection threshold could be set accordingly. These levels are the mean levels observed over long periods of time. The actual signal level observed over short periods of time will vary above and below the mean level. Figure 6.4 and Figure 6.5 show these *local* variations as a function of time for a sea clutter and growler cell, respectively. During times when the local clutter level is small compared to the mean level, low probabilities of false alarm will be observed; when local clutter levels are large with respect to the mean, high probabilities of false alarm will be observed. Therefore, the PFA rate will vary as a function of range and time. In Section 2.4, we examined constant false alarm rate processors that might be suitable for growler detection in a sea clutter environment. We take a closer look at these CFAR processors now.

The purpose of a noncoherent CFAR processor is to track the local variation in sea clutter so that a threshold can be set relative to the local clutter amplitude (or power) level, as opposed to the long term, mean amplitude level. In this way, the false alarm rate will not vary as a function of range and time. Recall that the local level of sea clutter in a test cell is estimated by averaging the local levels from neighbouring cells. The compound

K-distribution described in Chapter 4 provides the key for setting the local threshold, once the local clutter level estimate is obtained. If the CFAR processor estimates the local underlying mean amplitude level y for the cell under test, then Equation 4.11 can be used to set the local threshold with a prescribed false alarm rate. This approach should give CFAR operation providing that all of the cells in the CFAR window have the same local statistics as the cell under test.

To gain some insight into the detection performance of a CFAR processor in K-distributed sea clutter, the underlying mean levels of the growler cell given in Figure 6.5 and a neighbouring sea clutter cell given in Figure 6.4 are shown replotted together in Figure 6.8. The solid curve corresponds to the growler cell returns, and the dotted curve corresponds to the clutter cell returns. For the moment, consider the growler cell to be the cell under test; the clutter cell represents the other cells in the CFAR window that are used to estimate the local clutter level. The local detection threshold is typically set higher than the estimated clutter level to achieve a prescribed PFA. Therefore, Figure 6.8 indicates that the growler will not be detected around the 2 s, 10 s, and 30 s mark, due to the spiky clutter at these times.

The probability of false alarm performance of a noncoherent CFAR processor in K-distributed clutter can be demonstrated by considering a CFAR window that contains only clutter cells. Assume that the sea clutter curve shown in Figure 6.8 represents the returns from the test cell. Another neighbouring clutter cell (within the CFAR window) is required to represent the mean level of clutter in the neighbourhood of the test cell. In Figure 6.9, the test cell returns (dotted curve) and the returns from a neighbouring clutter cell (solid curve) are plotted as a function of time. The two cells are separated in range by only five resolution cells and would most certainly be contained in the same CFAR window. The detection threshold for the test cell is set relative to the clutter level measured from the reference cells in the CFAR window. Thus, the data presented in Figure 6.9 indicate that false alarms will occur in the test cell around the 2 s, 10 s, 22 s, and 30 s mark, due to




Figure 6.8: This figure shows the 512-pulse integrated amplitude signal from the growler cell and a neighbouring clutter cell in file B98 (HH). These are the same curves drawn in Figure 6.5 and Figure 6.4 respectively. The solid curve corresponds to the growler cell return, and the dotted curve corresponds to the clutter cell return. A spatial CFAR processor sets the test cell detection threshold according to the mean level in the neighbouring clutter cells. These data indicate that missed detections will occur around the 2 s, 10 s, and 30 s mark due to the spiky clutter returns at these times.




Figure 6.9: This figure shows the 512-pulse integrated amplitude signals from two clutter cells in file B98 (HH), separated in range by five radar resolution cells. The dotted curve represents a clutter cell under test, and the solid curve corresponds to neighbouring clutter data that are used by a CFAR processor to estimate the clutter level in the test cell. These data indicate that false alarms will occur around the 2 s, 10 s, 22 s, and 30 s mark.

the relatively low clutter levels that are found in the clutter reference cell at these times.

The surprisingly poor behaviour of a CFAR approach for the data presented above can be explained by considering the range correlation of sea clutter. Let's reexamine Figure 4.22 which shows spatial (range) and temporal correlations of sea clutter data. The spatial correlation curve shows that sea clutter echo from a given cell is less than 60 % correlated with sea clutter echo 30 m in range away. Since the pulse length associated with our data offers a range resolution of 30 m, the spatial correlation results imply that the correlation between sea clutter in adjacent radar resolution cells is not great. The small spatial correlation is explained by the K-distribution. The underlying mean level y is understood to be associated with the overall *tilt* of a patch of clutter, in response to the sea swell. Since the swell wavelength is typically between 100-150 m, adjacent 30 m pulses sample the wave at approximately quarter-wave intervals. Therefore, the underlying mean level component y can be expected to change considerably from range cell to range cell. Figure 4.7 and Table 4.1 show the large variability of mean clutter levels as a function of range.

A CFAR window combines measurements from adjacent range cells under the assumption that all of the cells have the same clutter statistics. We have seen however, that the statistics in adjacent clutter cells are in fact quite different, due to the variation in the underlying mean level component y described in the K-distribution. Therefore, CFAR processors will perform poorly in K-distributed sea clutter unless the pulse width is reduced substantially, allowing for a large CFAR window of similarly distributed clutter variates.

Although the analysis described above is for noncoherent CFAR processors, the same conclusions are drawn for coherent CFAR processors. The Doppler CFAR method of Larsson *et al.* [5] will not perform well on our data because of insufficient, similarly distributed clutter cells surrounding the test cell. Our growler data base at present does not offer the CFAR window lengths and range resolution necessary to provide meaningful performance estimates of traditional CFAR methods, applied to our problem. Therefore, these techniques are not considered further in Chapter 7.

6.3 Coherent Detectors

In Section 6.2, we discussed noncoherent detection strategies that are useful for detecting growlers in sea clutter. The IA detector is analyzed in the next chapter to quantify the performance of noncoherent detectors and to provide a benchmark with which to measure the performance of coherent detectors. In this section, we use the analysis carried out in Chapter 5 to design two *medium dwell-time* coherent detectors that satisfy the requirements enumerated on page 11. We are interested in designing detectors that offer improved growler detection. Since the detector performances are evaluated empirically from IPIX data, it is important that the detectors we design exploit characteristics that are representative of growler and clutter returns in general, and not simply of our data base of measurements. For example, in most of our data files, the growler Doppler velocity is limited to a small band of frequencies. Therefore, using Doppler frequency as a discriminant between growler and clutter data will improve the detection performance on our data base; however, these results will be of no general use because the range of growler frequencies encountered in practice is variable. Furthermore, our data base represents a finite range of growler and clutter radar cross sections. As a result, the sensitivity of proposed detectors to changes in GCR is also considered.

In Chapter 5, we saw that improved growler visibility in a sea clutter background could be realized by the difference in their spectral width statistics. Spectral width is independent of mean frequency and GCR, and is therefore an ideal statistic upon which to base a detector, as it can be generalized to cases beyond our data base of measurements. Furthermore, for *medium dwell-time* coherent processing, spectral width seems to be the only characteristic of the Doppler spectrum that uniquely distinguishes between clutter and growler spectra. Therefore, following the discussion given in Section 2.3.5, we call the spectral width a *sufficient* statistic. We now describe two different detectors that make use of the spectral width of the radar echo signal in order to detect growlers in sea clutter.

6.3.1 Autoregressive Largest Pole Magnitude Detector

In Chapter 5, we saw that AR parameters were suitable for modeling the Doppler spectra of growler and clutter returns. The AR poles, which model the second order statistics of the radar data, provide a useful discriminant between clutter returns and growler-in-clutter returns. In particular, the AR poles lie within the unit circle in the z -plane, and the magnitude of the pole closest to the unit circle can be associated with the most narrowband component in the radar signal (see Equation 5.26). In Figure 6.10, a scatter plot of the poles from a 6th order AR model of radar returns from a growler-in-clutter cell is shown. Since the growler returns are in general more narrowband than the clutter returns, the cluster of poles nearest the unit circle are associated with the growler component; and the next closest cluster of poles to the unit circle is associated with the sea clutter component of the radar returns.

As shown earlier, the magnitude of the pole closest to the unit circle is related to the spectral width of the most narrowband component in the data, and can be used as a statistic for detection. The detector can be summarized as follows:

1. The AR parameters (we use a 6th order model) are calculated according to Equation 5.19.
2. The denominator of Equation 5.23 is factorized giving the AR poles.
3. The pole with the largest magnitude is selected.
4. This magnitude is compared to a threshold. If the threshold is exceeded, a growler is declared; otherwise clutter is declared.
5. If desired, Equation 5.25 is evaluated to give a mean frequency estimate and Equation 5.26 is evaluated to give a spectral width estimate, based on the pole with the largest magnitude.




Figure 6.10: This figure shows a scatter plot of the poles from a 6th order AR parameterization of growler-in-clutter data in file B98 (VV). The cluster of poles closest to the unit circle model the narrowband growler component of the radar returns. The next closest cluster of poles to the unit circle model the wideband sea clutter component of the radar returns. Finally, the remaining four pole clusters model the receiver noise in the data.

Since this detector uses the largest pole magnitude for detection, we call it the ARLPM (AutoRegressive Largest Pole Magnitude) detector. The performance of this detector is analyzed in detail in the Chapter 7.

6.3.2 Gaussian Model Spectral Width Detector

A more direct approach to estimating the spectral width(s) associated with growler and clutter radar data involves the use of a 1G or 2G spectrum model as described in Section 5.4.4 and Section 5.5.3. As we saw in Section 5.4.4, fitting a 1G model to a sea clutter spectra provides good mean frequency and spectral width estimates. In a similar manner, Section 5.5.3 showed that good mean frequency and spectral width estimates of both the growler and clutter components result by fitting a 2G model to estimated growler-in-clutter data. Knowing whether to use a 1G or 2G spectral model is simple if you know in advance whether your test data is from a clutter cell or a growler-in-clutter cell. In practice though, it is the detector that must determine whether a given test cell contains a growler. Therefore, some procedure must be available in order to decide which model (1G or 2G) to use.

It might seem like a good idea to always use the 2G model. After all, the 1G model is just a special case of the 2G model. The error in this line of thinking is easily seen by an example. Imagine that the current test cell is a clutter-only cell and a 2G model is going to be fitted to the estimated spectrum. A 2G model will always provide a better fit than a 1G model since the extra degrees of freedom can be used to model some of the error associated with using just a 1G model. This over-modeling can cause a lot of problems to a spectral width detector, because the second Gaussian component can easily end up having a very narrow spectral width. Since the clutter peak typically spans tens of Hz and the estimate of the spectrum may be noisy, a single clutter spectral component can actually be over-modeled as two components, one of them being very narrow in spectral width. This kind of over-modeling can confuse the detector into thinking a growler is present.

A solution to the problem of selecting the appropriate model order, 1G or 2G, can be found by considering the *generalized likelihood ratio test* (GLRT) described by Equation 2.20. The GLRT applied to our problem of model order selection can be stated as follows:

Assume that the 2G model is the correct model and make an ML estimate of the 2G parameter vector described by Equation 5.29. Then evaluate the 2G likelihood using this ML estimate. In a similar manner, assume that the 1G model is the true model and evaluate the ML estimate of the 1G parameter vector described by Equation 5.27. The 1G likelihood is evaluated using this ML estimate. Finally, the correct model (2G or 1G) and associated parameters are selected by comparing the ratio of the 2G likelihood and the 1G likelihood to a threshold.

In our problem, maximizing the likelihood for the 2G and 1G models is equivalent to minimizing the modeling error described by Equation 5.30 and Equation 5.28, respectively. The residual modeling errors are analogous to the likelihoods associated with the optimum parameter vectors. The ratio of the residual modeling errors for the 2G and 1G optimizations is compared to a threshold to select the model that most likely fits the data. Comparing the ratio of the residual modeling errors to a threshold is equivalent to saying that the 2G model must reduce the residual modeling error by a certain percentage over the 1G residual modeling error, in order for it to be considered valid, and not just over-modeling the data. Once the model order has been selected, its corresponding parameters are then considered to be representative of the data. The operations that we have just described pertain to the *feature extractor* stage of Figure 2.1. The feature (parameter) vector is then passed on to the *threshold detector*. At this stage, we recognize that the spectral width (the smaller of the two if the 2G model is used) is the sufficient statistic and the threshold device simply compares the spectral width to a threshold. Since the detector we have described makes use of a Gaussian spectrum model to describe the data, we call this detector the GSW (Gaussian Spectral Width) detector. The operations performed by the GSW detector are

now summarized:

1. A Fourier spectrum estimate is formed from 512 (0.512 s) complex samples of time series data from a given test cell.
2. The 2G model specified by Equation 5.29 is applied to the spectrum estimate and Equation 5.30 is minimized using the nonlinear optimization procedure outlined in Appendix D. The 2G least squares modeling error is calculated.
3. The 1G model specified by Equation 5.27 is applied to the spectrum estimate and Equation 5.28 is minimized using the nonlinear optimization procedure outlined in Appendix D. The 1G least squares modeling error is calculated.
4. The ratio of the 2G and 1G modeling errors (2G/1G) is formed and compared to a threshold (empirically determined to be 0.4). If the ratio exceeds the threshold, the 1G model is selected; otherwise, the 2G model is used.
5. The spectral width from the chosen model is selected. If the 2G model is used, the smaller spectral width is selected.
6. Compare the spectral width to a threshold. If the threshold is exceeded, sea clutter is declared; otherwise, a growler is declared.

This detector is also evaluated in detail in Chapter 7. In the next section, we consider some of the practical requirements associated with the IA, ARLPM and GSW detectors.

6.4 Practical Considerations

In order to appreciate the performance evaluations of the IA, ARLPM, and GSW detectors given in Chapter 7, it is important to understand the costs associated with each detector. The radar hardware requirements associated with coherent radars are significantly more demanding than those of noncoherent radars. For example, the inexpensive and efficient

magnetron used in noncoherent radars must be replaced with a stable klystron or TWT (Travelling Wave Tube) amplifier in order to make coherent measurements of the receive signal. Furthermore, stable global system clocks and parallel I and Q receivers are required by the coherent radar whereas local clocks and a single receiver path are required by the noncoherent radar. These hardware requirements translate to increased cost. Clearly, the IA detector has a large cost advantage over the ARLPM detector and the GSW detector!

Putting radar hardware cost requirements aside, it is important to compare the relative signal processing requirements of the three detectors. A fundamental requirement of the radar is real time operation which places constraints on the types of operations that can be performed with available technology, and the type of processors required to do the job. With real time operation in mind, we examine the computing requirements associated with each of the three detectors identified above. We assume that calibrated I and Q digital samples are available to the signal processor (detector). Since we are concerned with *medium dwell-time* techniques, we assume that the radar operates in a stare-and-step mode. That is, it will stare along a given radial and collect 512 pulses (0.512 s at 1 kHz PRF) and then step to the next azimuth position. Therefore, each radial needs to be completely processed in 0.512 s. This type of operation places minimal storage requirements on the signal processor. Whether a single processor is used, or an array of processors (one for each range cell), 512 data samples place only a modest storage requirement on the processor. With the inexpensive price of memory today, the difference in storage requirements between the IA, ARLPM, and GSW detectors can be neglected. Assuming a 200 ns pulse (30 m) and a range coverage of 6 km, 200 radar resolution cells span each radial. Therefore, a single processor must process each resolution cell in 2.56 ms, whereas an array of 200 processors has 0.512 s to process each cell of data. If a 90 degree azimuth cover is required and the antenna has a 1 degree beamwidth, the full scan update time will be about 46 s. We now consider the computational complexity of each of the three detectors.

6.4.1 IA Detector Complexity

The IA detector simply sums the 512 samples to form a test statistic that is compared to a threshold. Therefore, only 511 addition operations are required. With a single processor, the 511 additions must be computed in 2.56 ms. Therefore, a single processor capable of 0.2 Mflops is required. A personal computer would be capable of handling this job in real time!

6.4.2 ARLPM Detector Complexity

The ARLPM detector requires the following operations: the evaluation of autocorrelation estimates specified in Equation 5.17, the evaluation of the Yule-Walker equations given in Equation 5.19, and the solution to the roots of the polynomial in the denominator of Equation 5.23. The autocorrelation estimates require about MN complex operations (CO's) (ie. complex multiply and adds) where $M = 6$ is the AR order and $N = 512$ is the number of data samples. The evaluation of the Yule-Walker equations requires on the order of $M^3 + M^2$ CO's. The number of operations required to solve for the roots of the polynomial in the denominator of Equation 5.23 is difficult to specify, because the computation is performed using iterative techniques such as Laguerre's Method [69]. The convergence properties of these methods are highly sensitive to the nature of the roots of the polynomial. Furthermore, once roots are found, they can be *polished* in several different ways to improve their accuracy. As a result, it is difficult to quantify the number of operations that are required.

In order to get a realistic measure of the computational complexity of the ARLPM detector, it was coded in Matlab [70] and run on 100 different windows, 0.512 s in length, of real radar data. The average number of floating point operations (additions or multiplications) required by Matlab was reported to be 460000. If these operations are to be performed on a single processor, then a processor capable of 180 Mflops is required, otherwise, parallel processors each capable of 0.9 Mflops are required.

6.4.3 GSW Detector Complexity

The GSW detector requires the following operations. First, a Fourier spectrum estimate is performed requiring $N \log N$ CO's. Then a nonlinear least squares optimization is performed for each of the 1G and 2G models. Each optimization is described in Appendix D and requires the following operations for each of I iterations. A Jacobian matrix is formed requiring approximately $10N$ CO's and N TFE's (transcendental function evaluations) for the 1G model and twice that for the 2G model. A psuedo-inverse operation is then performed. If the calculation is performed according to Equation D.12, then the parameter vector update requires approximately $L^3 + 2L^2N + LN$ CO's where $L = 3$ for the 1G model and $L = 6$ for the 2G model. The model evaluations required by Equation D.3 require an additional $2N$ CO's for the 1G model and $3N$ CO's for the 2G model. Therefore, for $N = 512$, each 1G iteration requires about $42N$ CO's and N TFE's; each 2G iteration requires about $110N$ CO's and $2N$ TFE's. Our analysis has indicated that convergence is usually achieved within 20 iterations. If the number of iterations is fixed at 20, the GSW detector would then require on the order of $3040N$ CO's and $60N$ TFE's. Assuming that the TFE's can be neglected (available from table lookup), a single processor would need to be capable of performing on the order of 1556480 CO's in 2.56 ms or 608 million CO's per second. On the other hand, a parallel bank of processors would each need to be capable of 3.04 million CO's per second. Although these processing requirements are very demanding, they can be reduced substantially. For example, since the amplitude parameters in the 1G and 2G models appear linearly, the nonlinear optimization can be rewritten to exploit this property. This results in a significant reduction in computational expense. (Details are given in Appendix D). Furthermore, rather than evaluating Equation D.12 directly, the SVD solution described in Appendix D can be used and even implemented using a systolic architecture. This however will require specialized hardware.

In order to make a realistic comparison between the number of computations required by the GSW detector, as compared to the ARLPM detector, the GSW detector was

also evaluated in Matlab. The average number of operations reported by Matlab from 100 independent runs of the GSW algorithm on real IPIX data was 1.12 million. Therefore, a single processor would need to be capable of 438 Mflops whereas a parallel bank of processors would require 2.10 Mflops. These Matlab figures compare quite well with the predicted requirements considering that the Matlab implementation used an SVD operation to calculate the pseudo-inverse, and that the number of iterations actually required was usually less than the maximum number of iterations set at 20.

6.5 Summary

In this chapter, we described three different detectors that will be evaluated in Chapter 7. The IA detector noncoherently integrates 0.512 s of data in order to come up with a test statistic. This detector will be representative of the performance that can be expected from noncoherent marine radars in use today, and can be implemented very inexpensively in real time. Two coherent detectors have been designed to take advantage of the different spectral width characteristics of clutter and growler returns. The ARLPM detector fits an autoregressive model to the received data and uses the poles of this model to classify growler and clutter returns. This detector requires significantly more computation than the IA detector but is still implementable in real time with currently available technology. Finally, the second and most computationally expensive coherent detector is the GSW detector. This detector uses a nonlinear least squares optimization to decide between a 1G or 2G spectrum model. Once the model and its associated parameters are solved for, the spectral width parameter is passed on to the threshold device for detection. The GSW detector has a very large computational cost associated with it. It requires more than twice the computing capabilities that the ARLPM detector requires. Although its processing requirements can be met with currently available computing hardware, it remains to be seen whether or not the added computational expense is justified.

Chapter 7

Detector Performances

7.1 Introduction

In this chapter, the three detectors developed in Chapter 6 are tested on real IPIX radar data, and their performances evaluated. These three detectors are the Integrated Amplitude detector, the AutoRegressive Largest Pole Magnitude detector, and the Gaussian Spectral Width detector. We begin in Section 7.2 by describing the data files that are used in the performance evaluations, along with the manner in which the detector testing is carried out. We refer to this selection of IPIX growler and clutter data files as the *data testbed*.

In Section 7.3, each file specified in the data testbed is operated on by the IA, ARLPM, and GSW detectors, and performance curves are generated for each file. Furthermore, composite performance curves are also generated that are based on the analysis of the entire data testbed. Since each file in the data testbed has different radar parameters associated with it, and each file was collected at different times, the composite performance curves are representative of the kind of performance that could be expected over the entire viewing area seen by a marine radar. The composite performance curves will also show the performance advantages, if any, of using one linear polarization configuration, say HH, over the other, VV.

Finally, in Section 7.4, the performances of the IA, ARLPM, and GSW detectors are compared in detail. In particular, the basis for the performance of the ARLPM and GSW detectors is re-examined so that realistic predictions of their performances can be made in cases where the GCR is smaller than that represented by the data testbed.

7.2 Data Testbed

The data testbed is composed of selected data from the files specified in Table 3.3. (The associated weather and waveheight data are given in Table 3.4 and Appendix A). In order to identify which range cells contain only sea clutter returns and which range cells contain growlers, a time-frequency map was performed for each range cell in a given file as described in Section 5.3. Once each range cell in the file was labelled as being either growler or clutter, the growler cell along with a few clutter cells were matched filtered and extracted for analysis. The matched filtering was carried out according to Equation 2.5 in Section 2.2. Since the range swath of each file in Table 3.3 is about 200 m and the range resolution of the 200 ns pulse is 30 m, there are only about six independent range cells in each file. Depending on the range position of the growler(s) in the file, the number of independent clutter-only cells selected from each file varied. Typically, a matched filtered growler cell along with three or four matched filtered clutter cells were extracted from each file in Table 3.3. These extracted clutter and growler data form the data testbed used to test each of the IA, ARLPM, and GSW detectors.

Since each original data file was collected under different conditions, the extracted cells are referred to by the name of the original data files from which they were extracted. Whenever we refer to a data file name in this chapter, it is to be understood that we are actually referring to the extracted growler cell data and clutter cell data from the named file. The performance evaluations are first carried out on the extracted cells for each file. In this way, the relative performance of a given detector, from file to file, can be compared against

the power and radar cross section estimates associated with each data file. These estimates are given in Table 5.2. Following the performance evaluations on a file by file basis, all of the extracted cell data in the data testbed are used to generate composite performance curves. These performance curves are representative of the expected performance over the entire viewing area covered by the marine radar, since many different aspect angles are represented by the data testbed.

The data testbed provided labelled range cells of clutter and growler radar data against which the IA, ARLPM, and GSW detectors could be tested. Each extracted range cell of data, (clutter or growler), contained 34,000 samples (pulses) of data with 0.001 s between each sample. (Note: if required, original data was undersampled so that an effective PRF of 1 kHz was obtained for each file). The detector performances were evaluated in the following manner. Each range cell of data was divided into windows of 512 (0.512 s) consecutive pulses, with an inter-window spacing of 256 pulses (0.256 s). Each range cell therefore provided 130 windows of data, with 50 % overlap between adjacent windows. Since Hamming weighting was applied to each window of data before processing, the data overlap was approximately zero between windows. Each detector operated on the 130 windows for each range cell of data, and made 130 independent detection decisions. Since the data were labelled as clutter or growler, the detector decisions were tabulated as being correct or in error, and the performance of each detector was thus evaluated.

As described above, three or four clutter cells along with a growler cell were extracted from a given IPIX data file. Therefore, each file offered approximately 500 clutter windows and 130 growler windows for testing. False alarms occur when a clutter cell is mistakenly detected as a growler. Since each file offers about 500 different clutter windows, estimates of probabilities of false alarm smaller than 2 % will have high variances and thus are not meaningful. Similarly, estimates of probabilities of detection greater than 90 % will have high variances. However, when performances are evaluated on the entire data testbed, each polarization channel (HH and VV) offers about 4000 clutter windows and 1000 growler

windows. Therefore, probabilities of false alarm as low as 0.003 and probabilities of detection as high as 0.99 can be estimated with reasonable confidence!

In the next section, we describe the file by file performances as well as the composite performances of the IA, ARLPM, and GSW detectors.

7.3 Detector Performances

7.3.1 Performance on a File by File Basis

In this section, we examine the detector performance of the IA, ARLPM, and GSW detectors for each of the files that make up the data testbed. For each data file, the clutter and growler data windows are operated on by each of the detectors, as described in the last section. As described in Chapter 6, each detector operates by first performing a feature extraction operation, followed by a threshold comparison of the test statistic. The feature used by the IA detector is the integrated amplitude of the windowed data. The ARLPM detector uses the largest AR pole magnitude determined from the windowed data as its feature. The GSW detector uses the spectral width of the windowed data as its feature. The performance of these detectors depends solely on the statistical separation of the test features under the clutter-only hypothesis, and the growler-in-clutter hypothesis. Therefore, for each file that we analyze, the histograms of the test features will be given for both hypotheses.

Detector performance is quantified as follows. A feature threshold is set to divide up the feature space into two decision regions corresponding to the two hypotheses. Then, the PFA and the PD for the given threshold are estimated from the clutter and growler histogram data respectively. Varying the decision threshold causes the PFA and PD to vary in a corresponding way. Since it is impossible to determine in advance the operational requirements (and hence the appropriate threshold), the threshold is varied over the range of the feature space, and the PFA and PD are estimated for each threshold. A plot of the PD versus the PFA is referred to as the ROC [23] and is given for each file along with the

FILE ID	FIGURE NUMBER
B97 HH	7.1
B98 HH	7.2
B98 VV	7.3
B99 VV	7.4
B110 HH	7.5
B111 VV	7.6
B112 HH	7.7
B112 VV	7.8
B113 HH	7.9
B114 VV	7.10
B115 HH	7.11
B115 VV	7.12
B123 HH	7.13
B124 VV	7.14
B125 HH	7.15
B125 VV	7.16

Table 7.1: This table lists the data files and polarization channels that make up the data testbed. For each file, the figure number is given where the performance results of the IA, ARLPM, and GSW detectors can be found.

histograms of the test features.

For each of the files in the data testbed, the performances of the IA, ARLPM, and GSW detectors are presented. For files that contain both HH and VV data (see Table 3.3), performance evaluations are done for each of the two polarization channels. Thus, there are eight different performance evaluations for the HH channel, and eight more performance evaluations for the VV channel. These detector performances are given in Figure 7.1–Figure 7.16. Each figure indicates the file and polarization channel in the data testbed on which the results are based. Table 7.1 indicates which data file in the data testbed each figure is based upon.

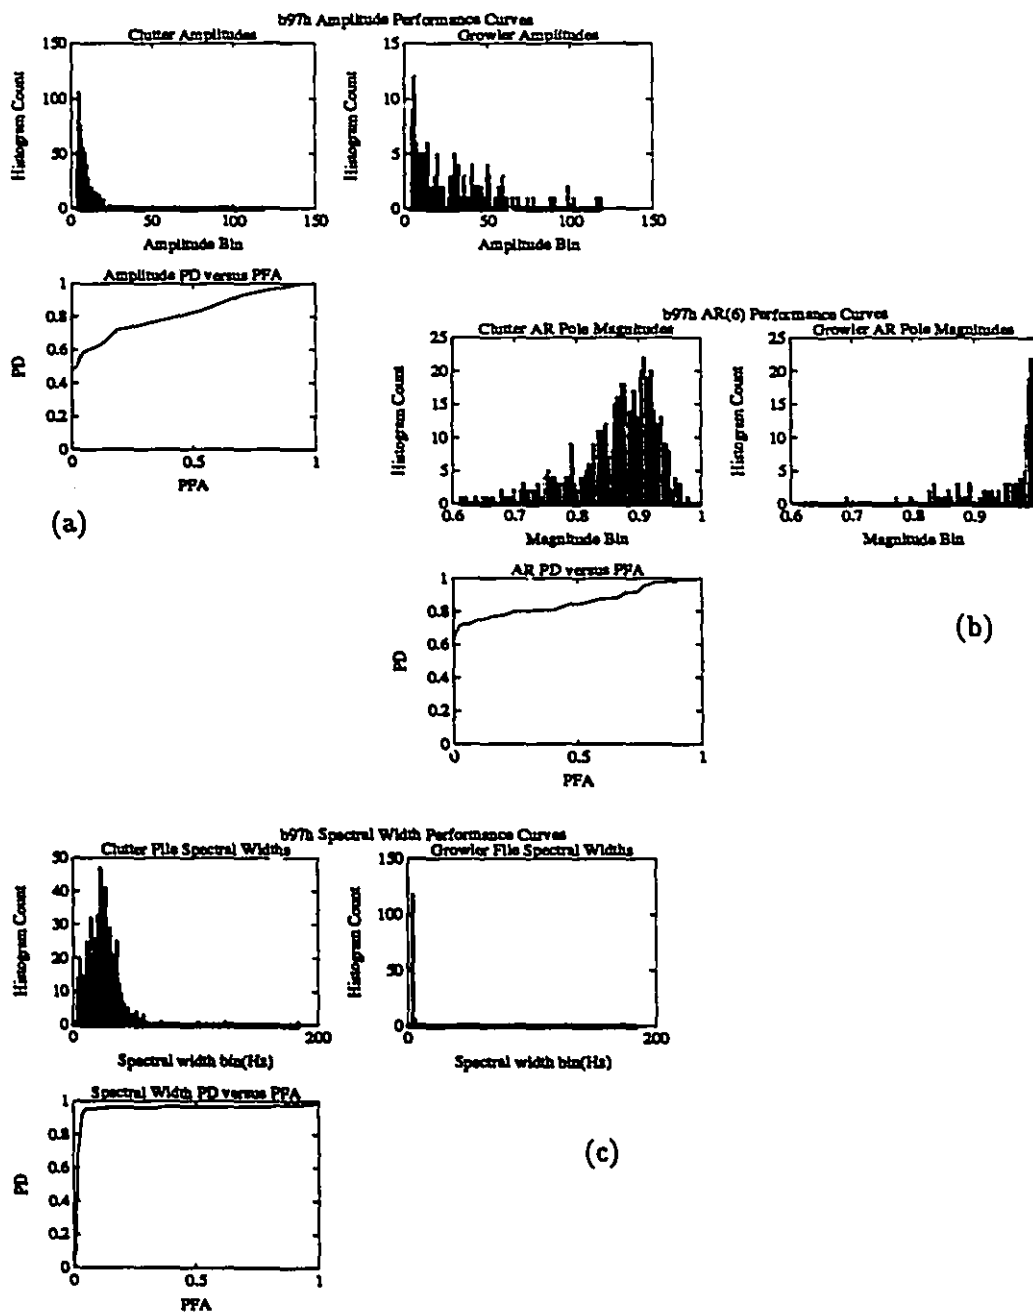


Figure 7.1: This figure gives the detector performance curves generated from file B97 (HH). (a) The IA detector performance: the clutter and growler integrated amplitude histograms are shown along with the corresponding ROC curve. (b) The ARLPM detector performance: the clutter and growler largest pole magnitude histograms are shown along with the corresponding ROC curve. (c) The GSW detector performance: the clutter and growler spectral width histograms are shown along with the corresponding ROC curve.

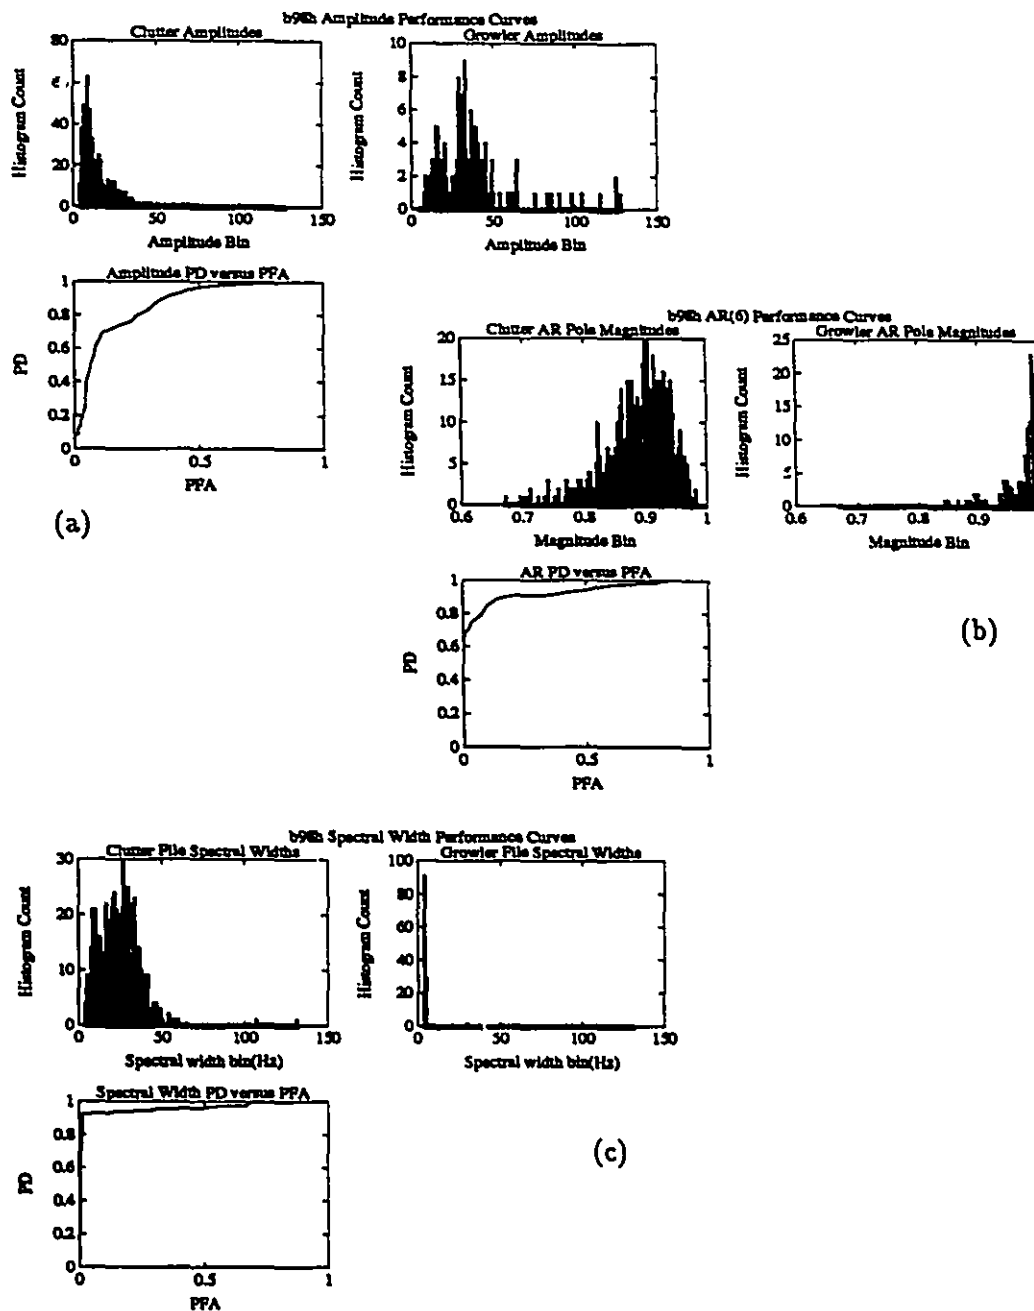


Figure 7.2: This figure gives the detector performance curves generated from file B98 (HH). (a) The IA detector performance: the clutter and growler integrated amplitude histograms are shown along with the corresponding ROC curve. (b) The ARLPM detector performance: the clutter and growler largest pole magnitude histograms are shown along with the corresponding ROC curve. (c) The GSW detector performance: the clutter and growler spectral width histograms are shown along with the corresponding ROC curve.

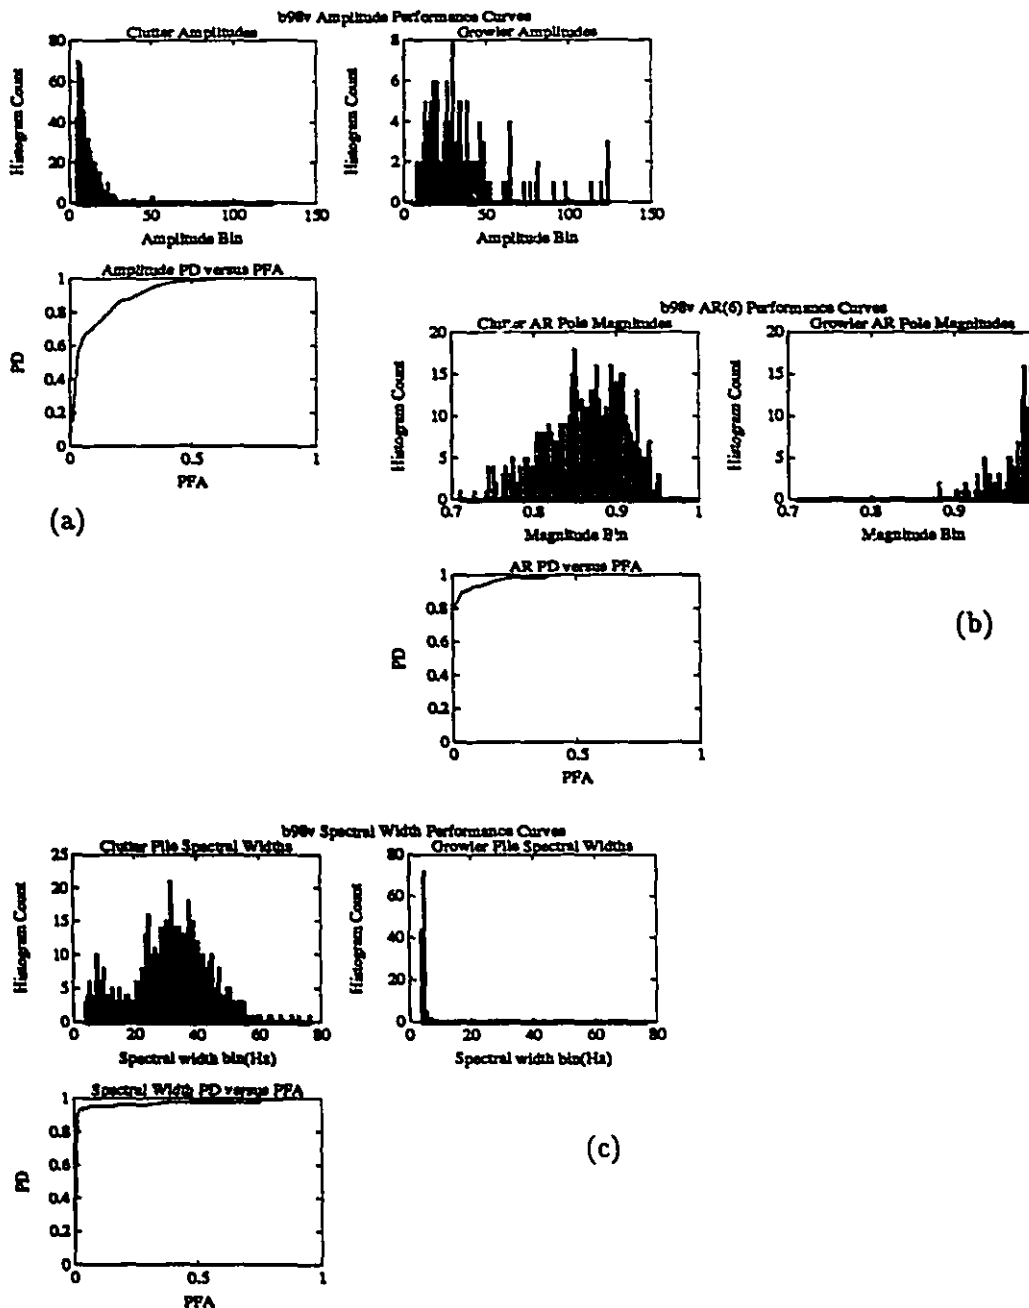


Figure 7.3: This figure gives the detector performance curves generated from file B98 (VV). (a) The IA detector performance: the clutter and growler integrated amplitude histograms are shown along with the corresponding ROC curve. (b) The ARLPM detector performance: the clutter and growler largest pole magnitude histograms are shown along with the corresponding ROC curve. (c) The GSW detector performance: the clutter and growler spectral width histograms are shown along with the corresponding ROC curve.

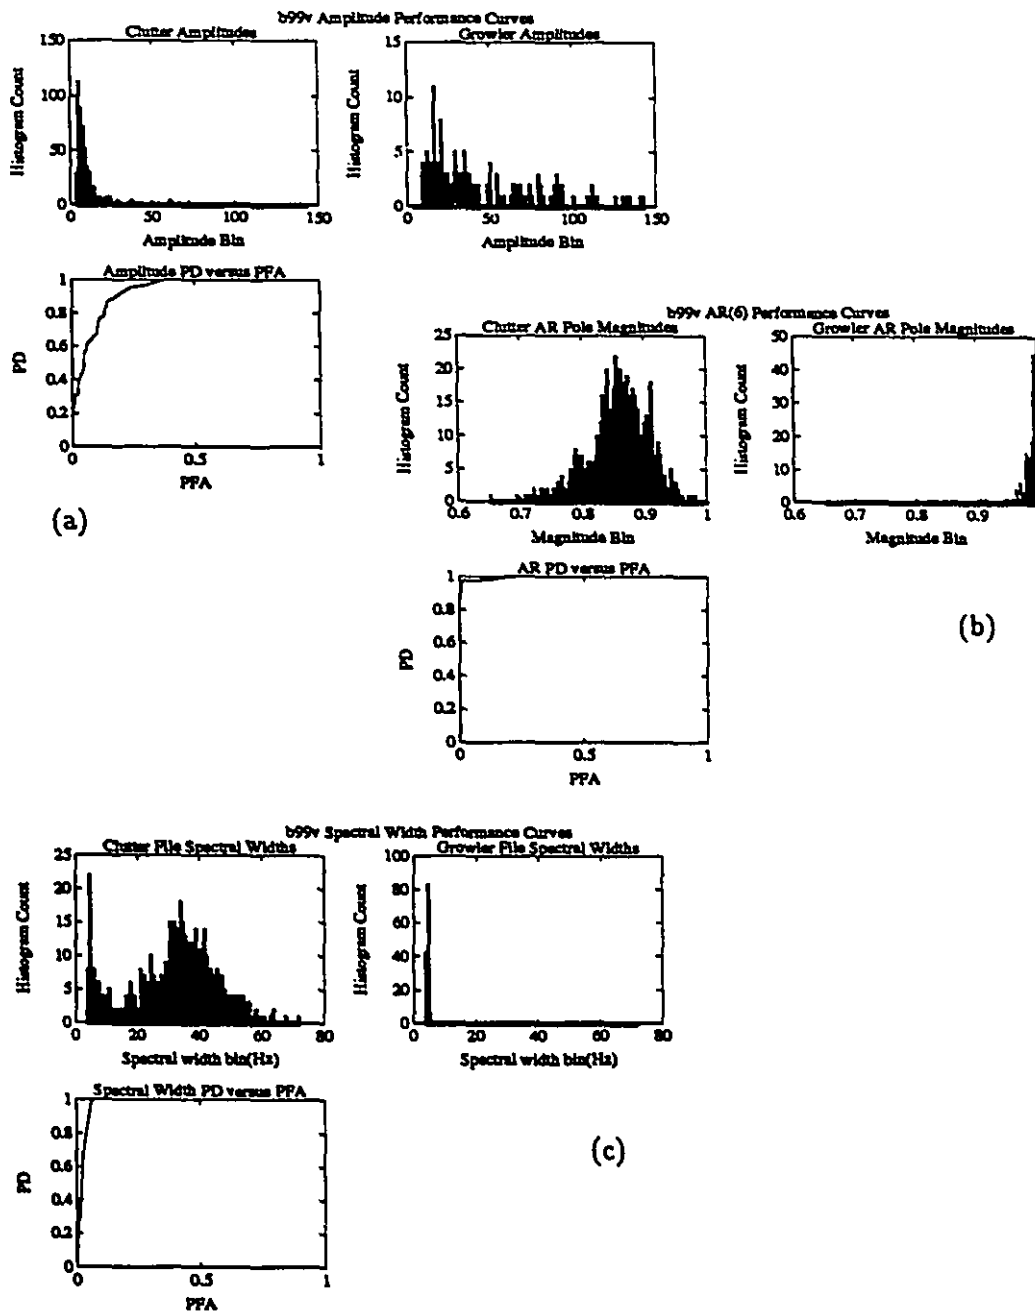


Figure 7.4: This figure gives the detector performance curves generated from file B99 (VV). (a) The IA detector performance: the clutter and growler integrated amplitude histograms are shown along with the corresponding ROC curve. (b) The ARLPM detector performance: the clutter and growler largest pole magnitude histograms are shown along with the corresponding ROC curve. (c) The GSW detector performance: the clutter and growler spectral width histograms are shown along with the corresponding ROC curve.

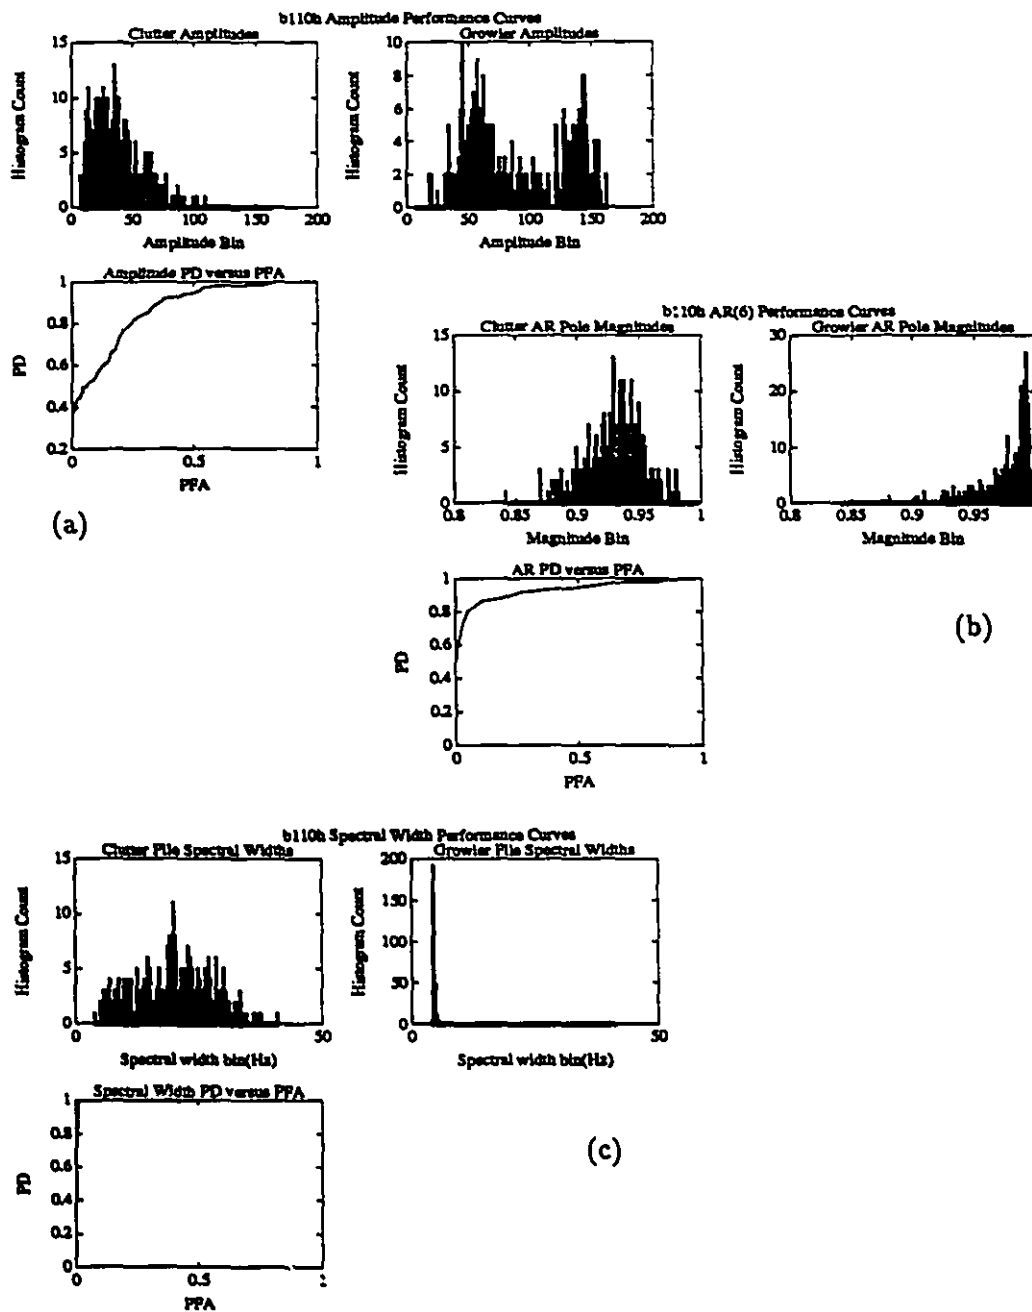


Figure 7.5: This figure gives the detector performance curves generated from file B110 (HH). (a) The IA detector performance: the clutter and growler integrated amplitude histograms are shown along with the corresponding ROC curve. (b) The ARLPM detector performance: the clutter and growler largest pole magnitude histograms are shown along with the corresponding ROC curve. (c) The GSW detector performance: the clutter and growler spectral width histograms are shown along with the corresponding ROC curve.

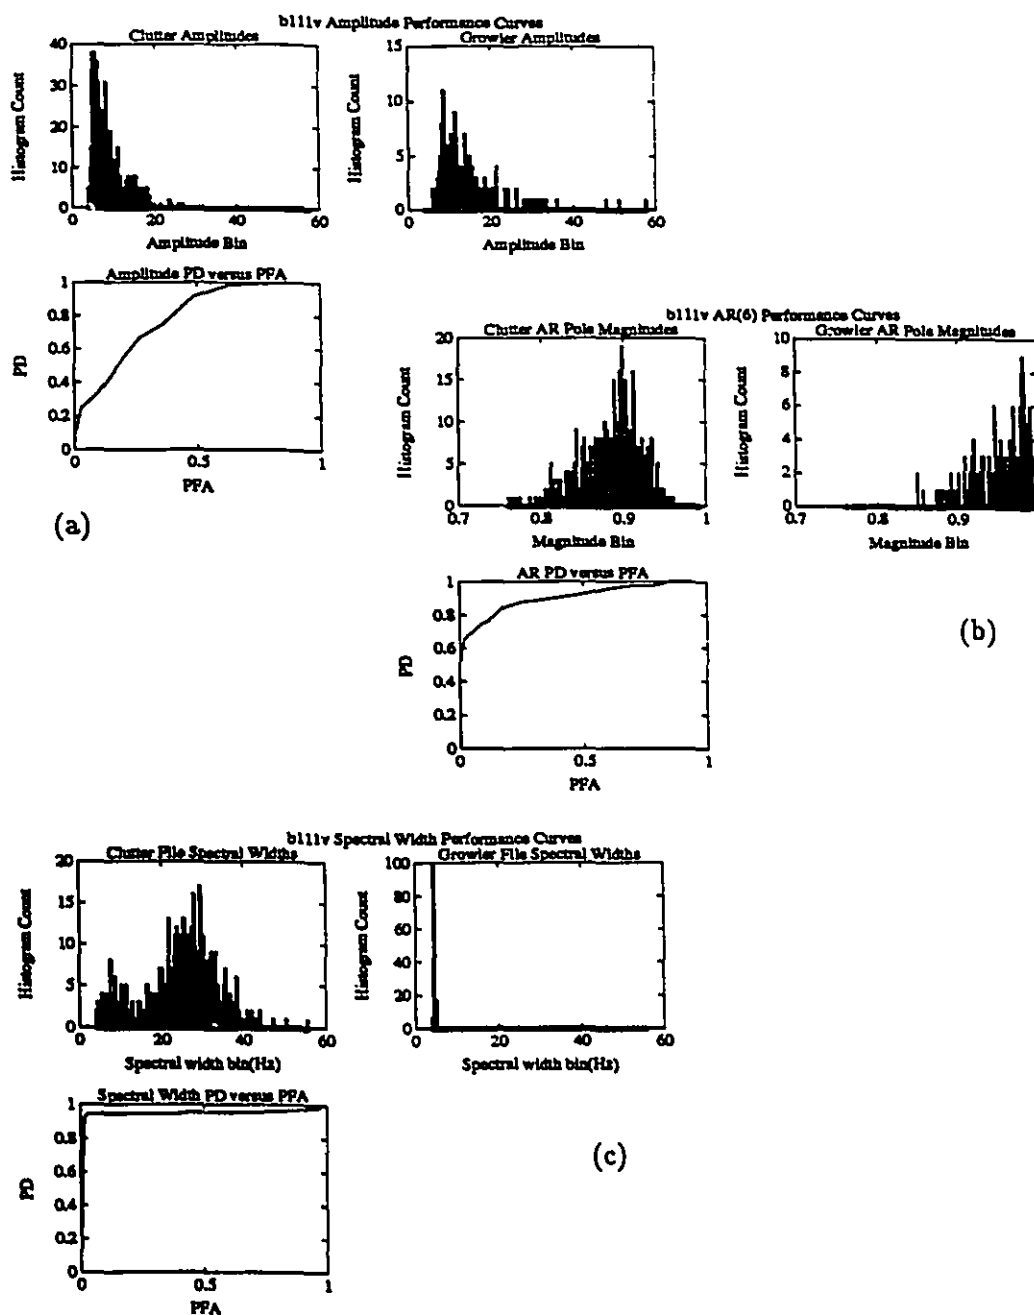


Figure 7.6: This figure gives the detector performance curves generated from file B111 (VV). (a) The IA detector performance: the clutter and growler integrated amplitude histograms are shown along with the corresponding ROC curve. (b) The ARLPM detector performance: the clutter and growler largest pole magnitude histograms are shown along with the corresponding ROC curve. (c) The GSW detector performance: the clutter and growler spectral width histograms are shown along with the corresponding ROC curve.

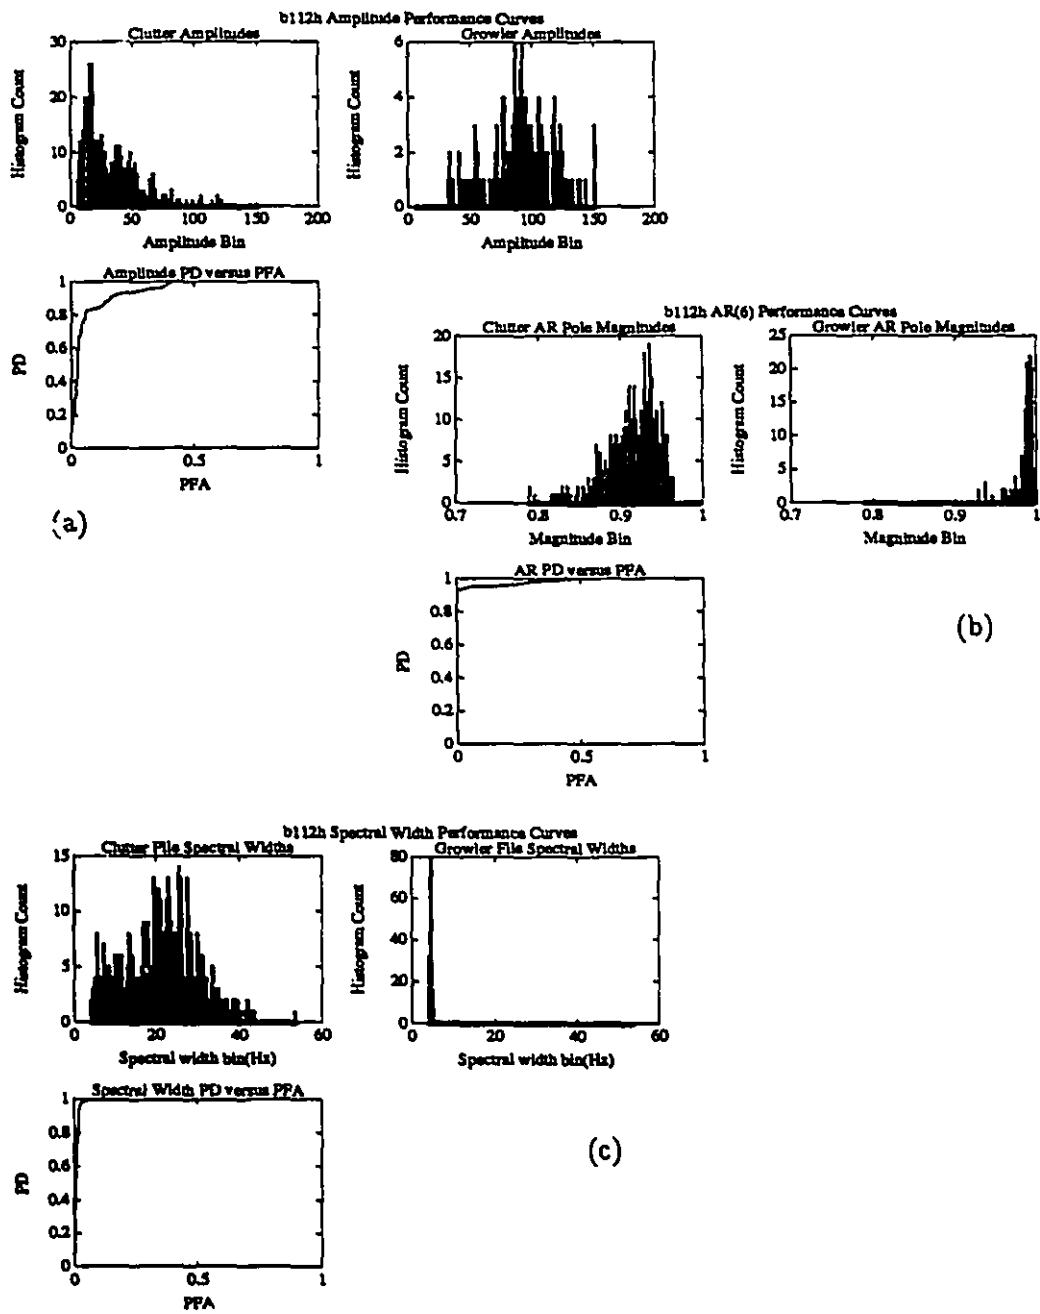


Figure 7.7: This figure gives the detector performance curves generated from file B112 (HH). (a) The IA detector performance: the clutter and growler integrated amplitude histograms are shown along with the corresponding ROC curve. (b) The ARLPM detector performance: the clutter and growler largest pole magnitude histograms are shown along with the corresponding ROC curve. (c) The GSW detector performance: the clutter and growler spectral width histograms are shown along with the corresponding ROC curve.

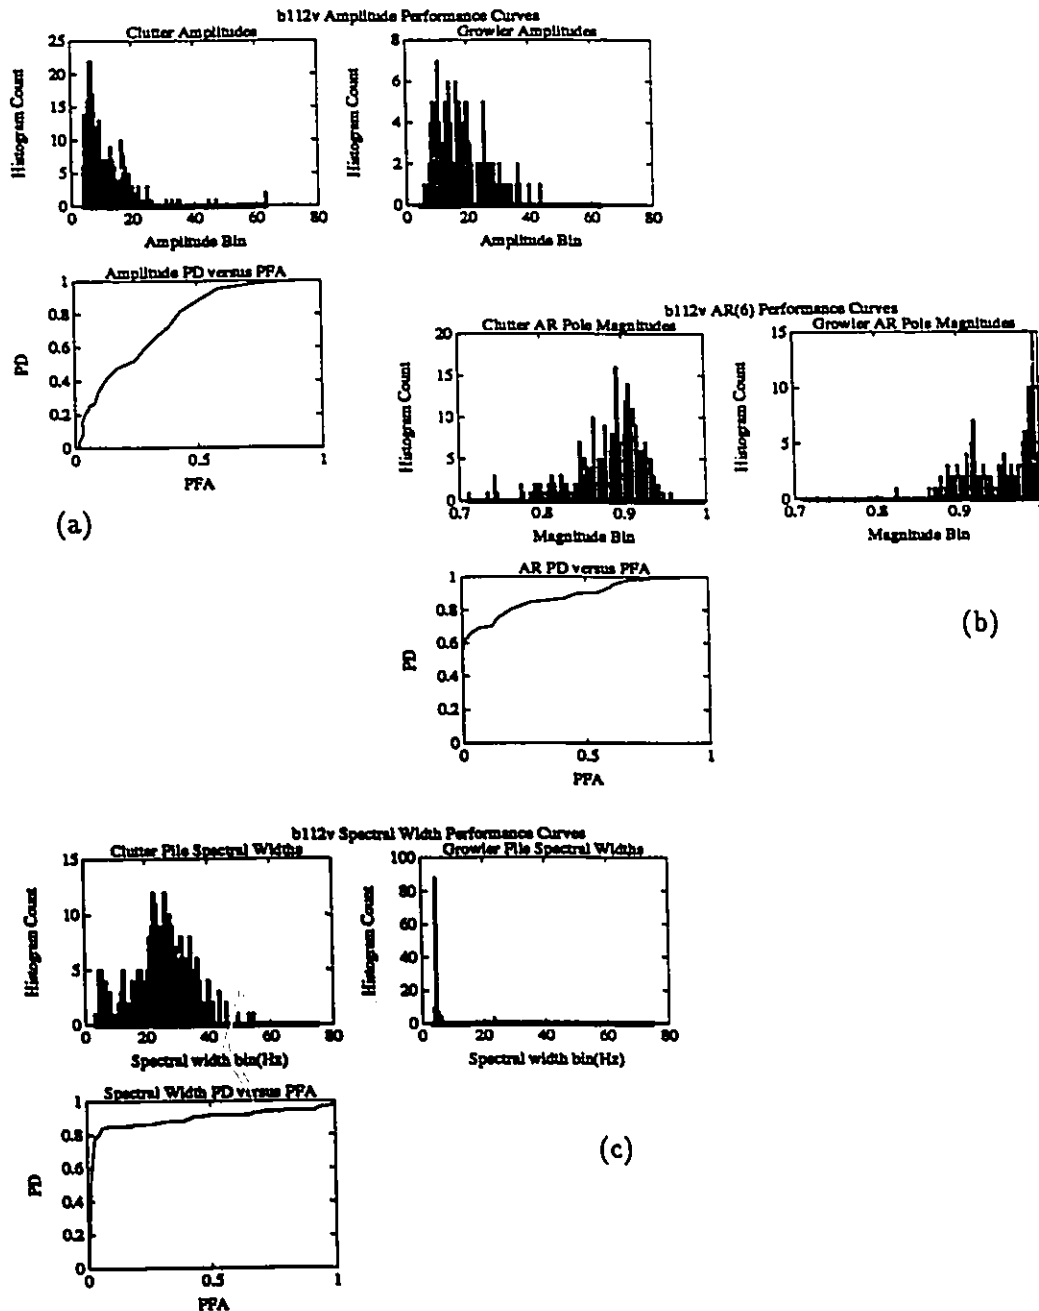


Figure 7.8: This figure gives the detector performance curves generated from file B112 (VV). (a) The IA detector performance: the clutter and growler integrated amplitude histograms are shown along with the corresponding ROC curve. (b) The ARLPM detector performance: the clutter and growler largest pole magnitude histograms are shown along with the corresponding ROC curve. (c) The GSW detector performance: the clutter and growler spectral width histograms are shown along with the corresponding ROC curve.

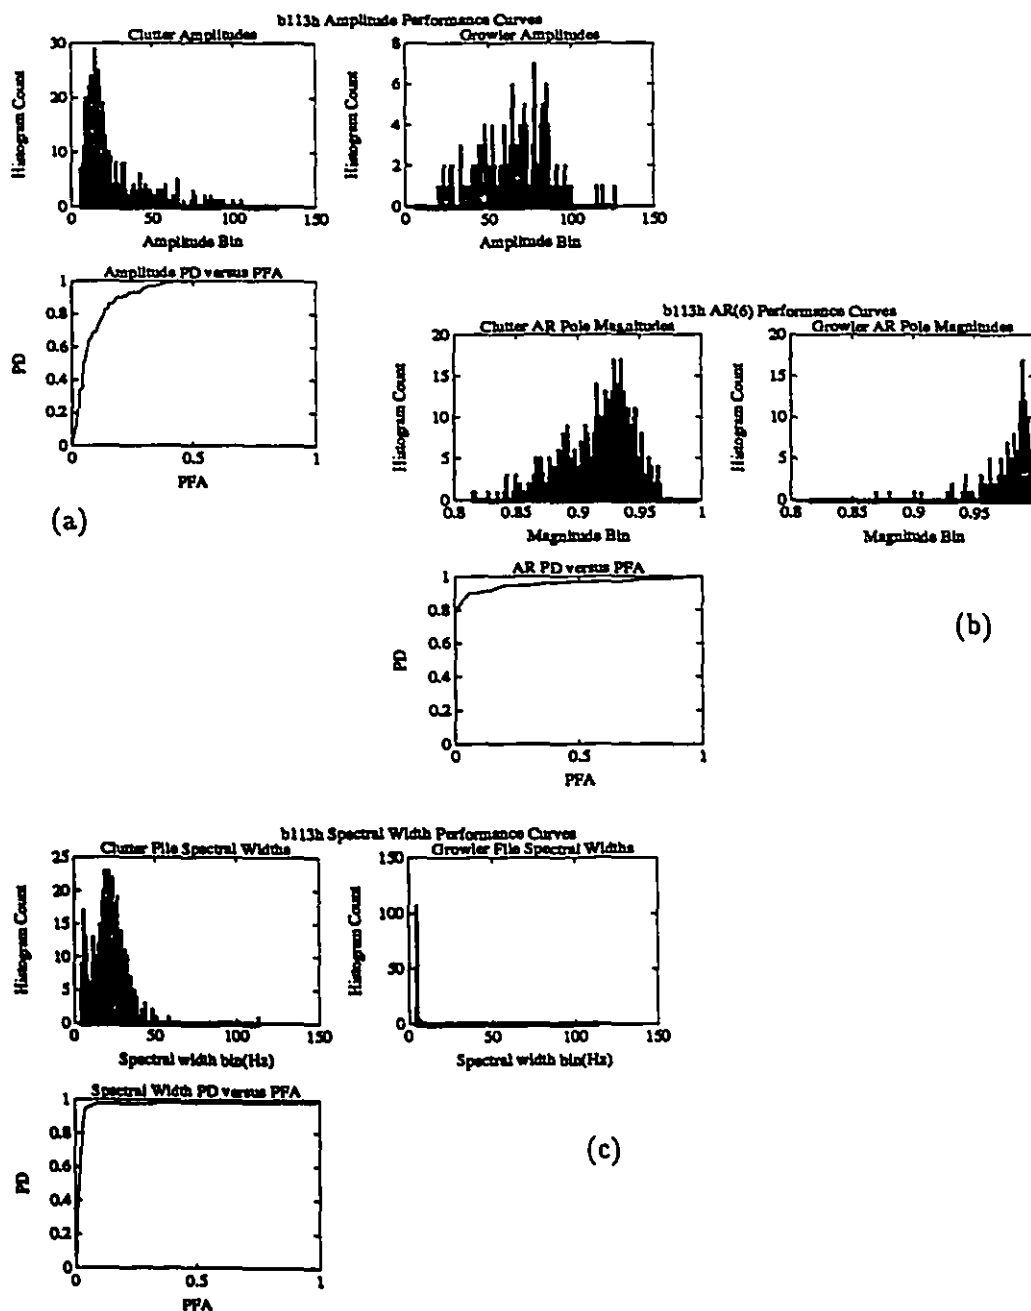


Figure 7.9: This figure gives the detector performance curves generated from file B113 (HH). (a) The IA detector performance: the clutter and growler integrated amplitude histograms are shown along with the corresponding ROC curve. (b) The ARLPM detector performance: the clutter and growler largest pole magnitude histograms are shown along with the corresponding ROC curve. (c) The GSW detector performance: the clutter and growler spectral width histograms are shown along with the corresponding ROC curve.

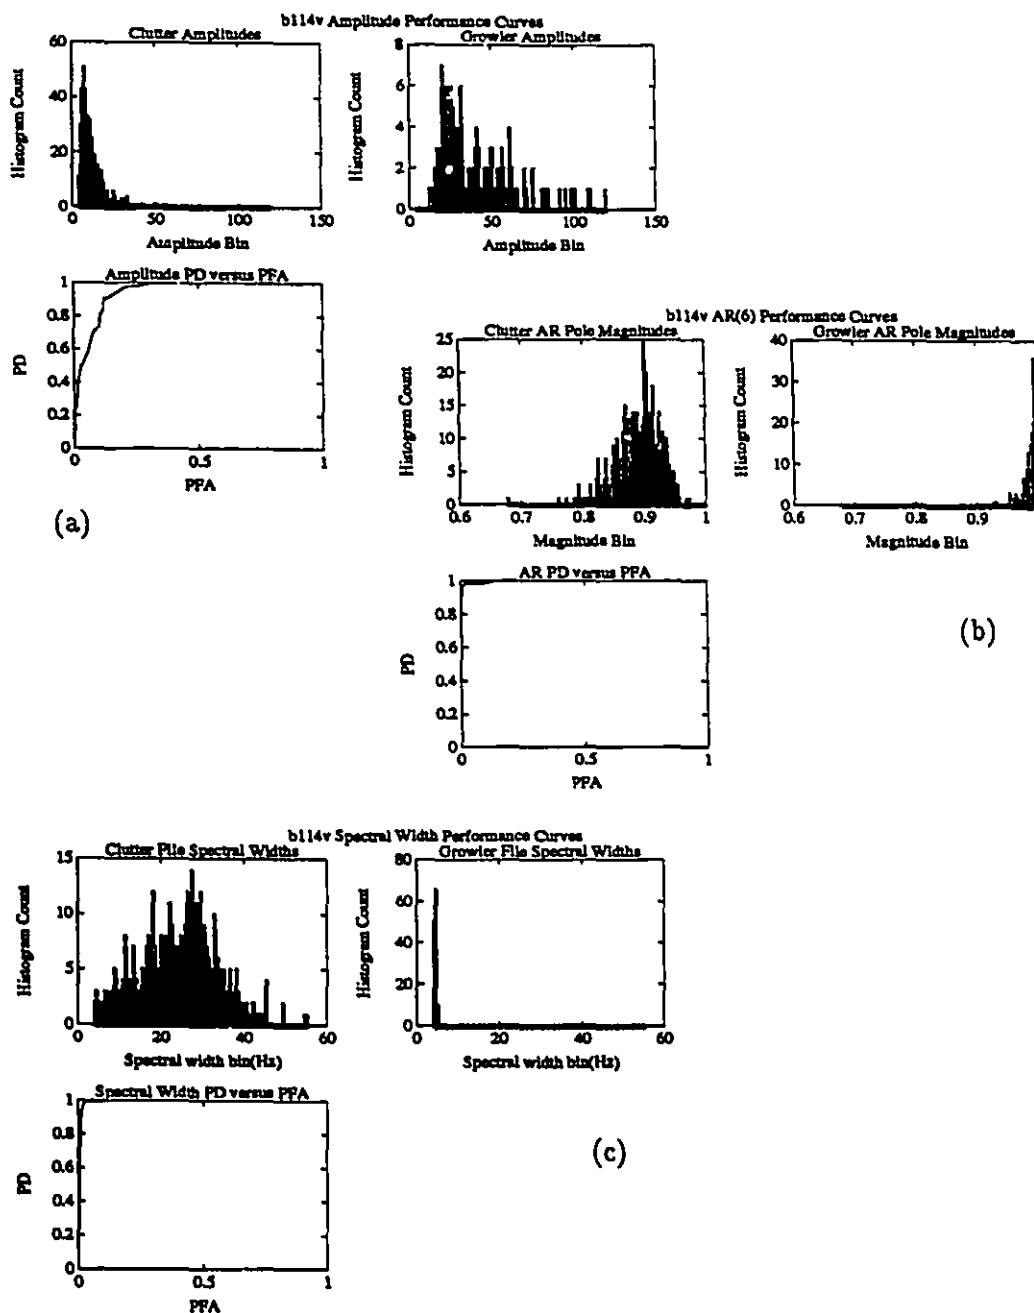


Figure 7.10: This figure gives the detector performance curves generated from file B114 (VV). (a) The IA detector performance: the clutter and growler integrated amplitude histograms are shown along with the corresponding ROC curve. (b) The ARLPM detector performance: the clutter and growler largest pole magnitude histograms are shown along with the corresponding ROC curve. (c) The GSW detector performance: the clutter and growler spectral width histograms are shown along with the corresponding ROC curve.

The relative performance of the IA, ARLPM, and GSW detectors varies from file to file as shown in Figure 7.1–Figure 7.16. Detectors like these, that are designed to separate growler returns from clutter returns, are expected to detect better in an environment where the GCR is larger. Table 5.2 provides GCR estimates for each of the data files described here. By comparing the detector performances in Figure 7.1–Figure 7.16 with the corresponding GCR entries in Table 5.2, several observations can be made and are now described.

Generally speaking, the IA detector has better detection performance with larger GCR. The trend is not as simple though when it comes to false alarm rate performance. This is because GCR represents an average quantity. For a fixed GCR, the nature of the clutter amplitude variation can be quite different. When clutter amplitudes are characterized by a small variance about their mean value, the false alarm performance is good; however, when the clutter amplitudes are characterized by a larger variance about their mean value, the false alarm rate performance is worse.

The ARLPM detector shows definite improvements as the GCR increases. These improvements are accentuated by the fact that the ARLPM detector performs coherent integration on the data. As a result, the detector operates with a larger effective GCR due to the integration gain of the coherent processing.

The GSW detector also shows detection performance improvements as the GCR increases. That is, the PD reaches 1.0 for a smaller value of PFA. The improvement in detection performance of the GSW detector, though, is not as drastic as the improvement shown by the ARLPM detector. The reasons behind this are twofold. First, the GSW detector has better detection performance than the ARLPM detector to begin with, so the maximum possible improvement that the GSW detector can make is smaller than the corresponding improvement of the ARLPM detector. Second, the GSW detector does not explicitly use amplitude information, whereas the ARLPM detector does (see Section 7.4). Therefore, the ARLPM detector has direct gains due to increased GCR.

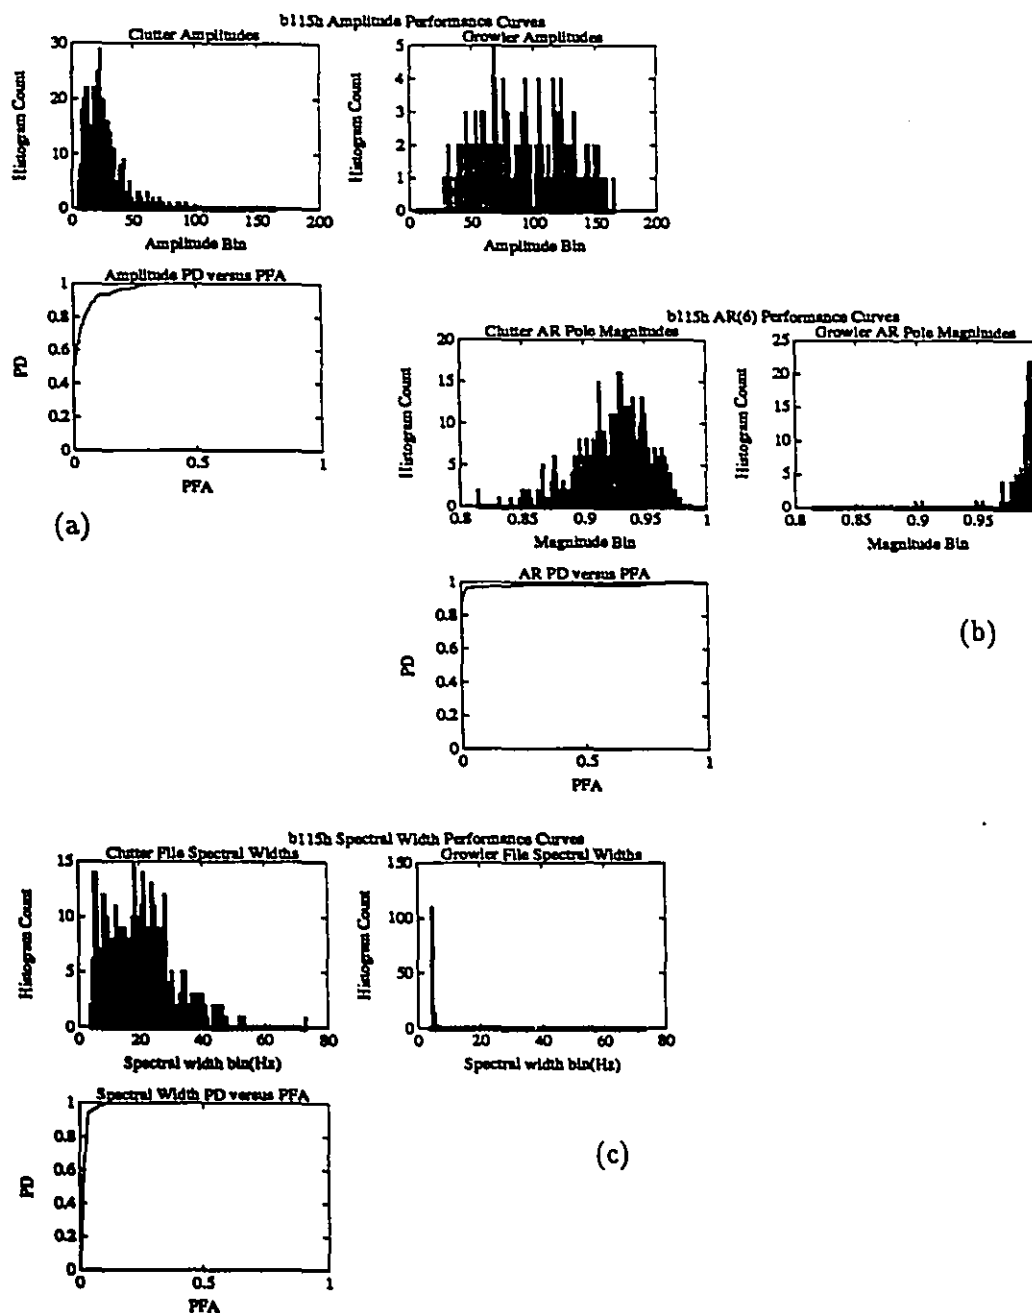


Figure 7.11: This figure gives the detector performance curves generated from file B115 (HH). (a) The IA detector performance: the clutter and growler integrated amplitude histograms are shown along with the corresponding ROC curve. (b) The ARLPM detector performance: the clutter and growler largest pole magnitude histograms are shown along with the corresponding ROC curve. (c) The GSW detector performance: the clutter and growler spectral width histograms are shown along with the corresponding ROC curve.

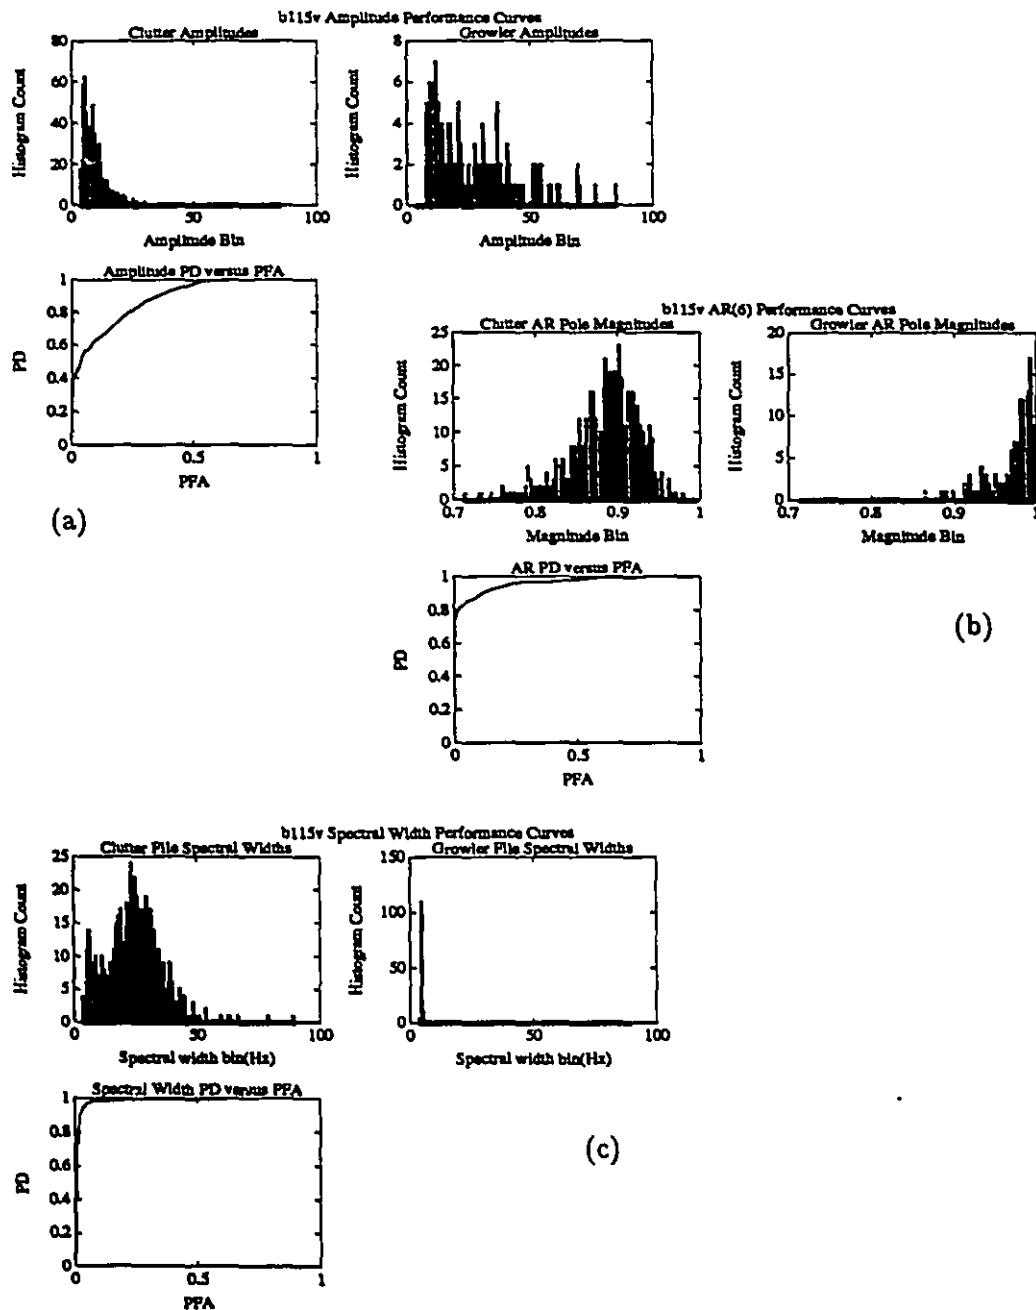


Figure 7.12: This figure gives the detector performance curves generated from file B115 (VV). (a) The IA detector performance: the clutter and growler integrated amplitude histograms are shown along with the corresponding ROC curve. (b) The ARLPM detector performance: the clutter and growler largest pole magnitude histograms are shown along with the corresponding ROC curve. (c) The GSW detector performance: the clutter and growler spectral width histograms are shown along with the corresponding ROC curve.

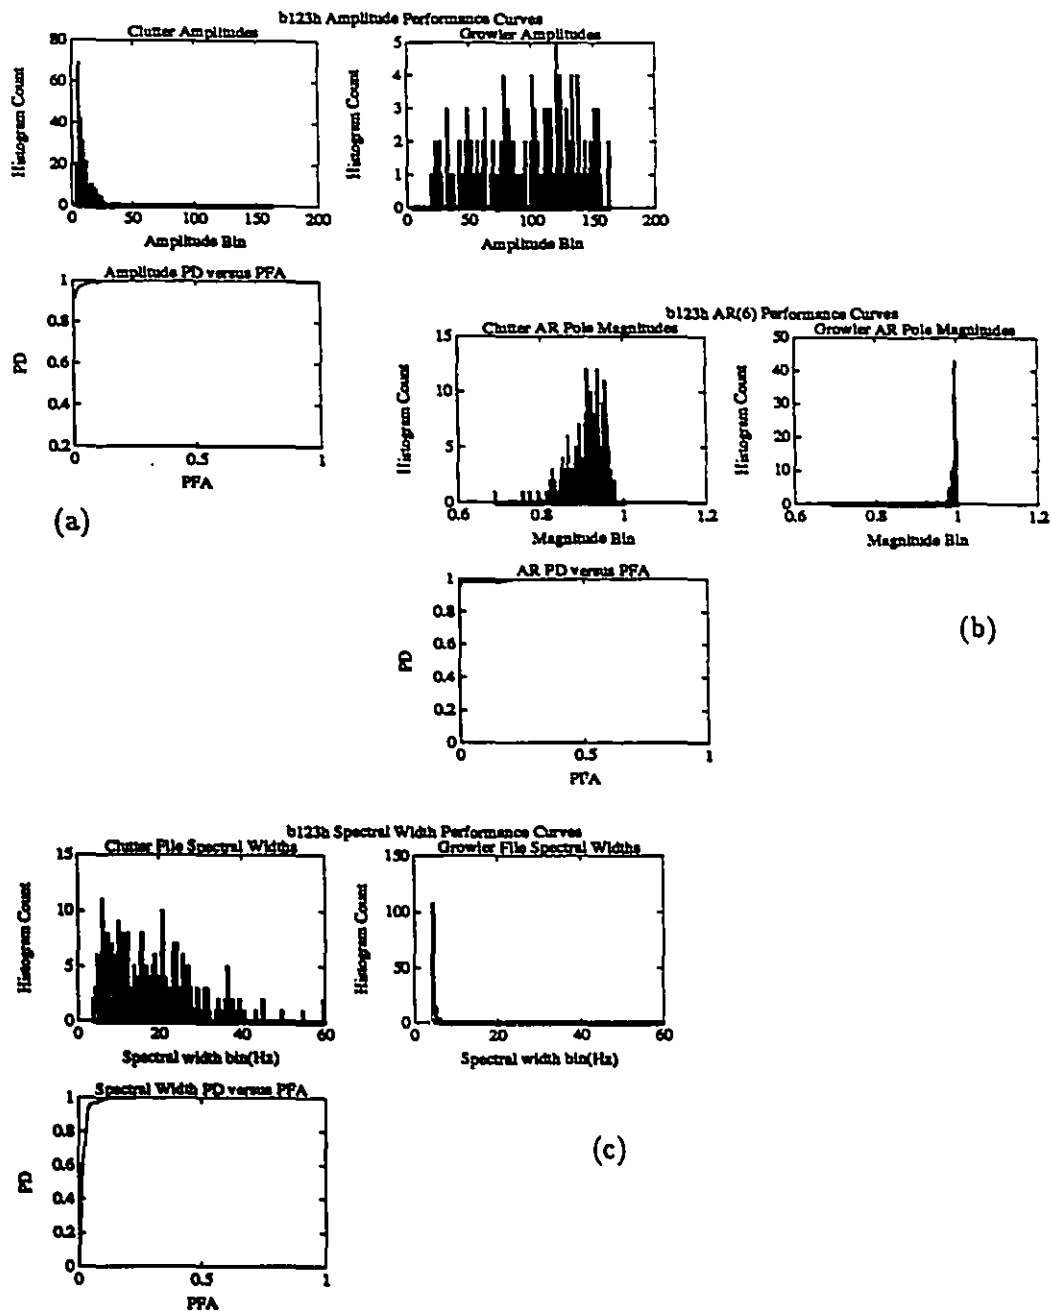


Figure 7.13: This figure gives the detector performance curves generated from file B123 (HH). (a) The IA detector performance: the clutter and growler integrated amplitude histograms are shown along with the corresponding ROC curve. (b) The ARLPM detector performance: the clutter and growler largest pole magnitude histograms are shown along with the corresponding ROC curve. (c) The GSW detector performance: the clutter and growler spectral width histograms are shown along with the corresponding ROC curve.

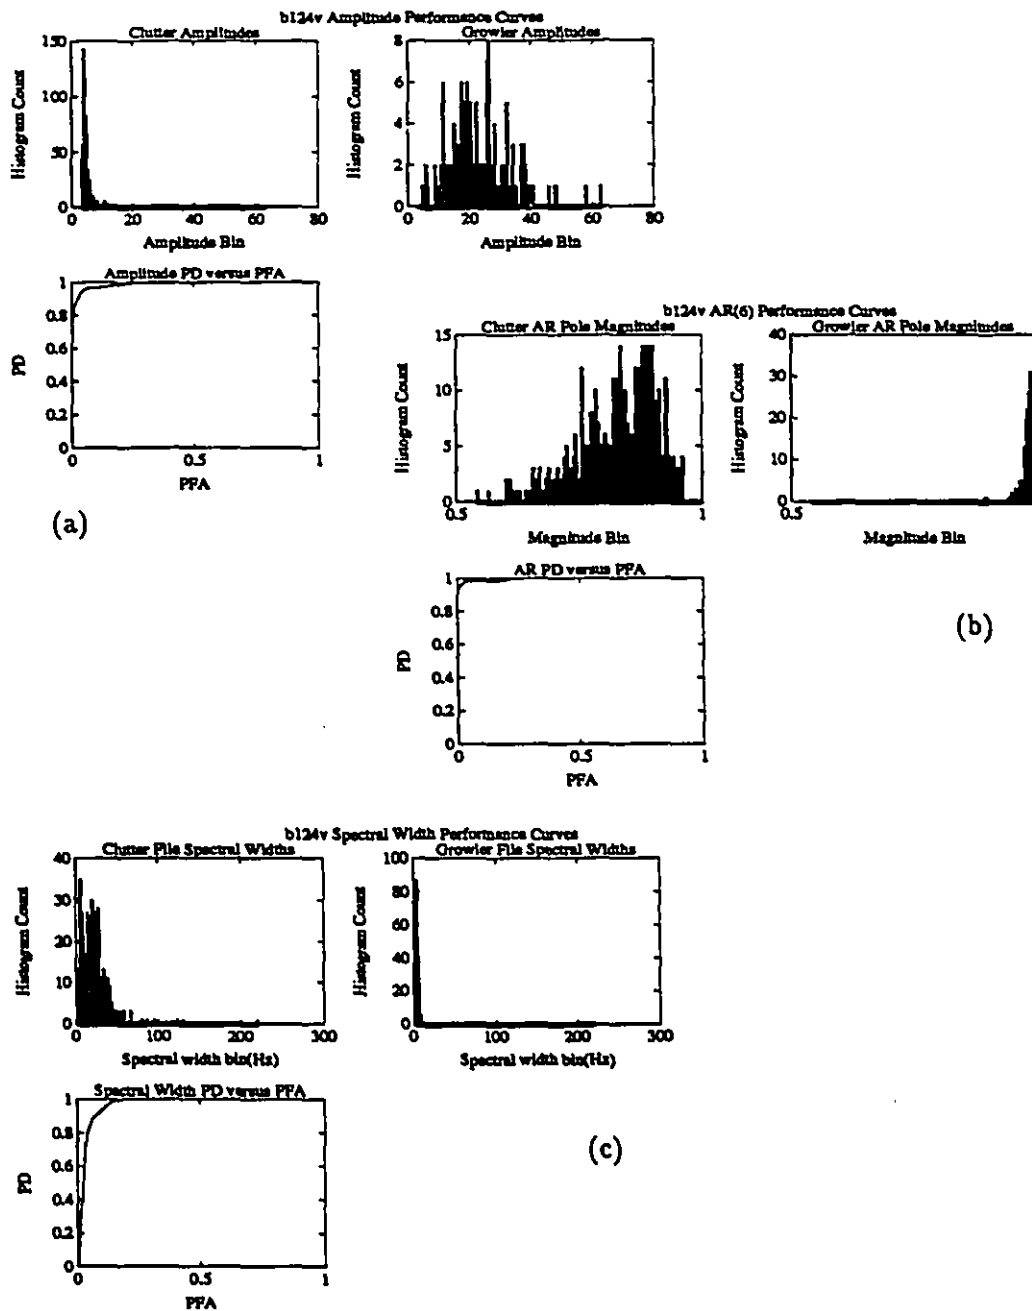


Figure 7.14: This figure gives the detector performance curves generated from file B124 (VV). (a) The IA detector performance: the clutter and growler integrated amplitude histograms are shown along with the corresponding ROC curve. (b) The ARLPM detector performance: the clutter and growler largest pole magnitude histograms are shown along with the corresponding ROC curve. (c) The GSW detector performance: the clutter and growler spectral width histograms are shown along with the corresponding ROC curve.

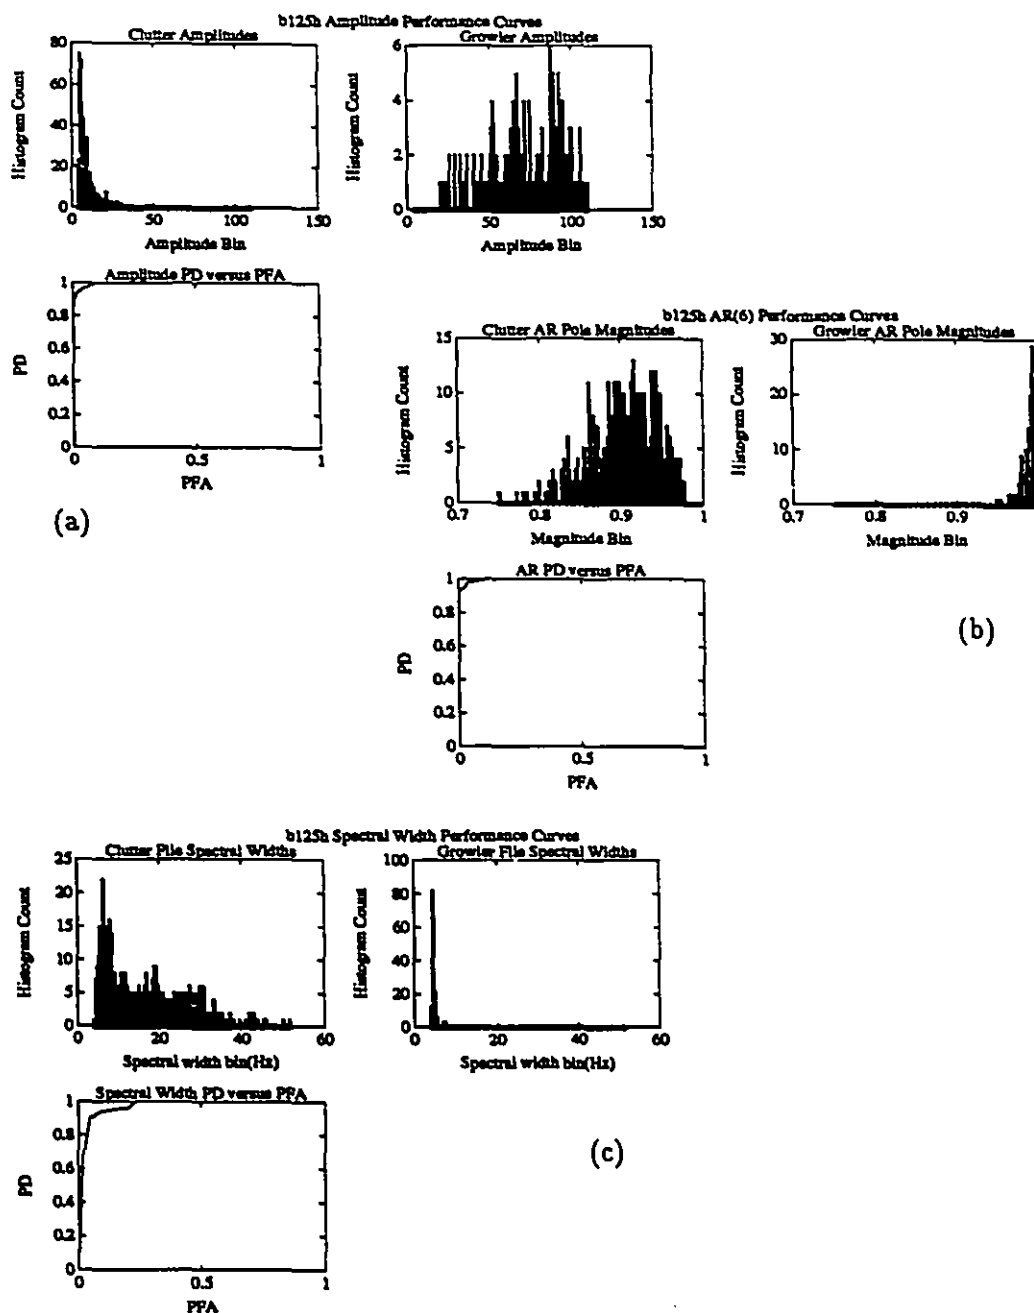


Figure 7.15: This figure gives the detector performance curves generated from file B125 (HH). (a) The IA detector performance: the clutter and growler integrated amplitude histograms are shown along with the corresponding ROC curve. (b) The ARLPM detector performance: the clutter and growler largest pole magnitude histograms are shown along with the corresponding ROC curve. (c) The GSW detector performance: the clutter and growler spectral width histograms are shown along with the corresponding ROC curve.

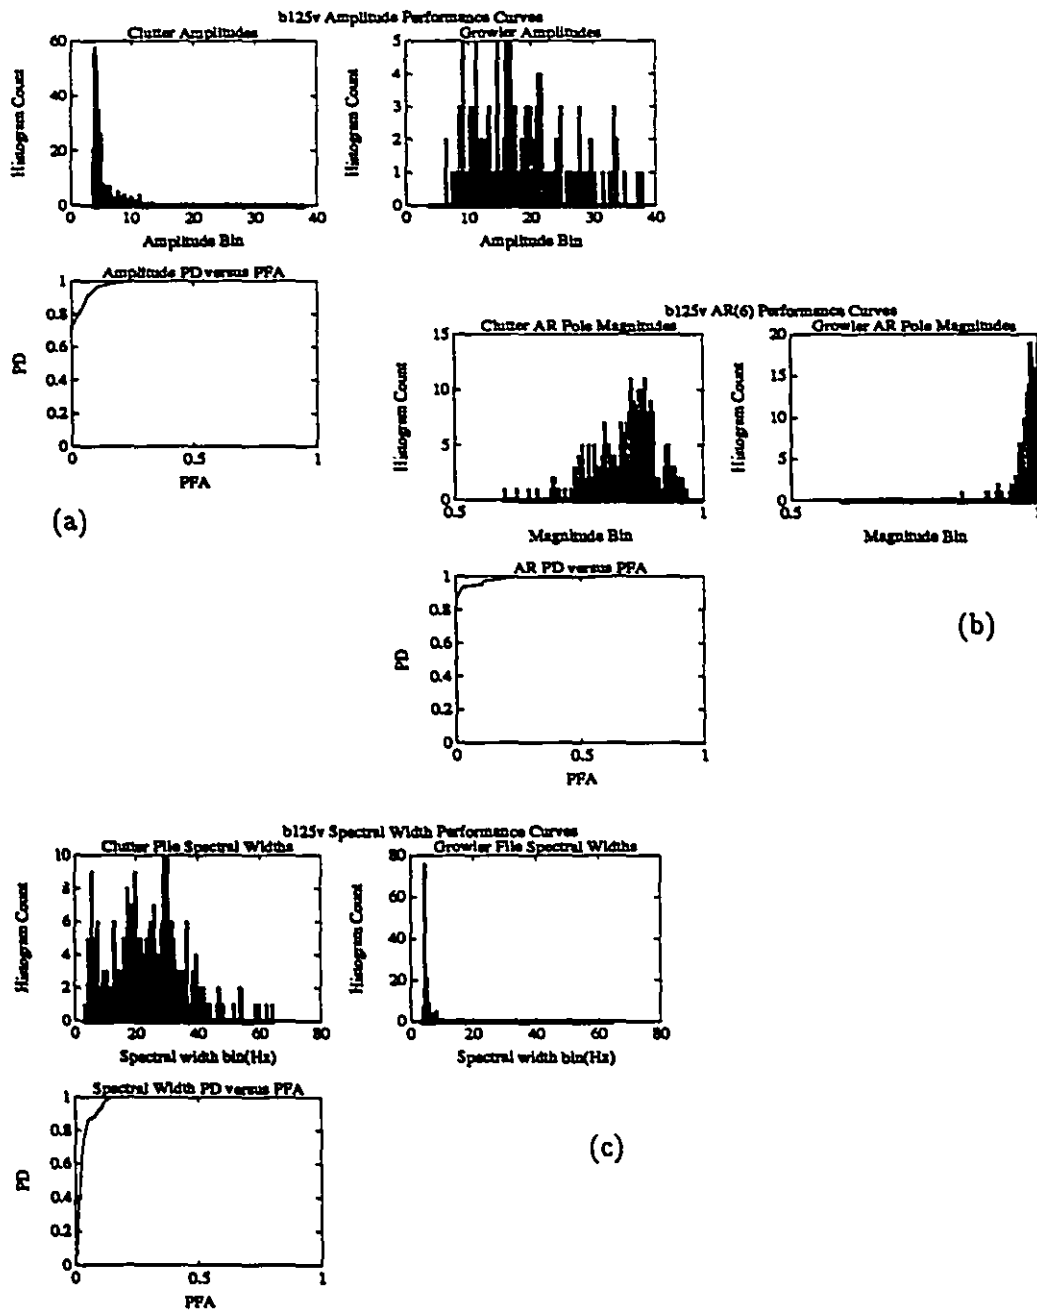


Figure 7.16: This figure gives the detector performance curves generated from file B125 (VV). (a) The IA detector performance: the clutter and growler integrated amplitude histograms are shown along with the corresponding ROC curve. (b) The ARLPM detector performance: the clutter and growler largest pole magnitude histograms are shown along with the corresponding ROC curve. (c) The GSW detector performance: the clutter and growler spectral width histograms are shown along with the corresponding ROC curve.

The performance evaluations given in this section demonstrate the individual abilities of the IA, ARLPM, and GSW detectors on a file by file basis. We have seen that detection performance improves with CCR as we expect. However, we have not made any comments on the relative performance of the three detectors with respect to each other. In the next section, the data presented here on a file by file basis will be combined so that composite performance evaluations over the entire data testbed can be made. These composite evaluations will allow us to better compare the relative merits of each of the three detectors.

7.3.2 Detector Performance Summary

We now examine the overall performances of the IA, ARLPM, and GSW detectors. In order to do this, we combine all of the HH channel data in the data testbed (all of the HH files in Table 7.1) into one data set, and all of the VV channel data in the data testbed into another data set. These composite data sets will be used to measure the composite detector performances in the HH and VV channels, respectively. Since each of the data files that make up the two composite data sets have different aspect angles or look directions associated with them, and since they were collected under comparable conditions (see Table 3.3 and Table 3.4 for details), the composite files are representative of the radar returns over the region of coverage of a typical navigation marine radar. Performance evaluations based on the composite data sets will not only provide more independent windows of data for detector testing, which allow for finer measurements in detection and false alarm probabilities, but they will also be more representative of the kind of detector performance that would be observed in a real operational scenario.

In Figure 7.17 and Figure 7.18, the composite performance of the IA detector is given for the HH and VV channels, respectively. Following an identical format to the performance curves given on a file by file basis in the last section, the integrated amplitude histograms of the labelled clutter and growler data are given, along with the ROC, for both

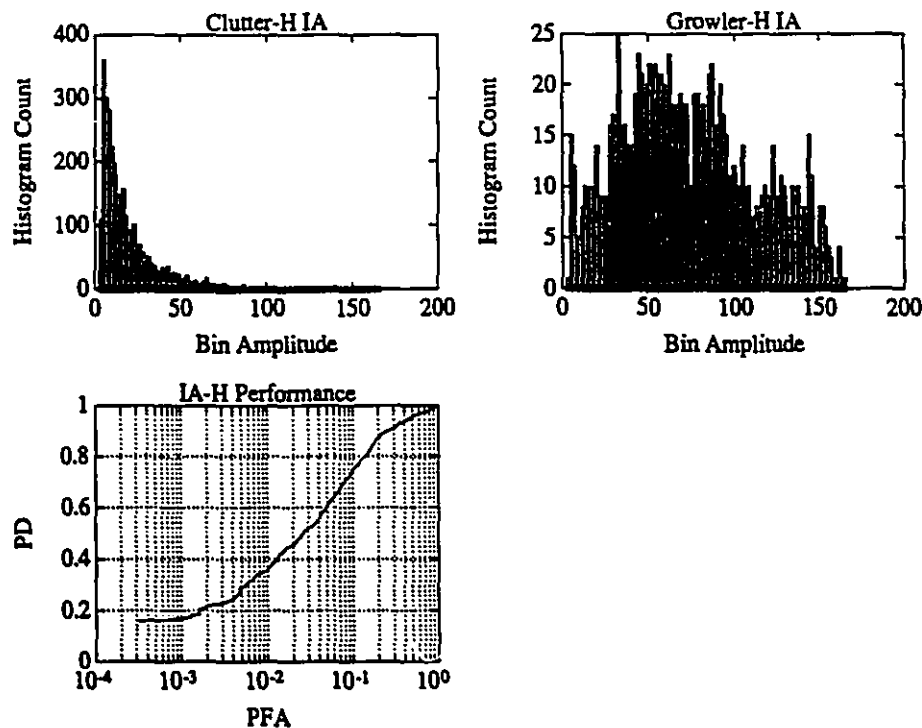


Figure 7.17: This figure gives the composite detector performance curves for the IA detector in the HH channel. The composite clutter and growler integrated amplitude histograms are shown along with the corresponding ROC curve. These composite statistics are based on the entire data testbed.

the HH and VV channels. The PFA data are plotted on a logarithmic scale so that small values of PFA are easily visible.

In Figure 7.19 and Figure 7.20, the corresponding performance curves are given for the ARLPM detector. Again, these curves are based on the composite data sets that are described above. The clutter and growler histograms of the largest AR pole magnitude are given for each of the HH and VV channels, along with the corresponding ROC curves. The PFA data is plotted on a logarithmic scale for the same reasons described above.

Finally, the performance curves for the GSW detector are given in Figure 7.21 and Figure 7.22. The clutter and growler spectral width histograms are given along with the

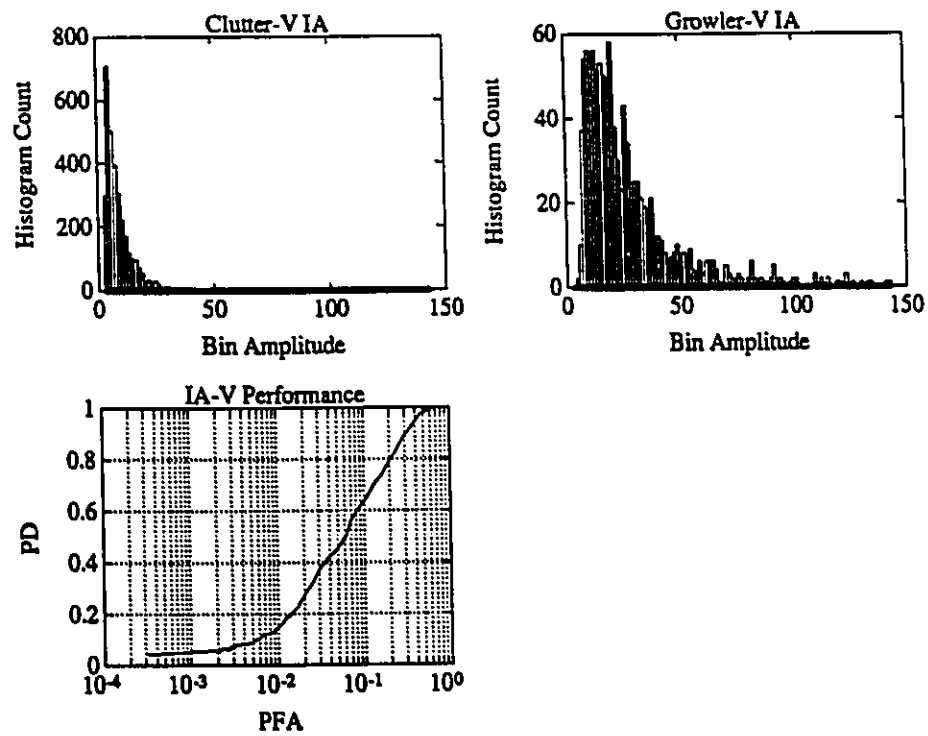


Figure 7.18: This figure gives the composite detector performance curves for the IA detector in the VV channel. The composite clutter and growler integrated amplitude histograms are shown along with the corresponding ROC curve. These composite statistics are based on the entire data testbed.

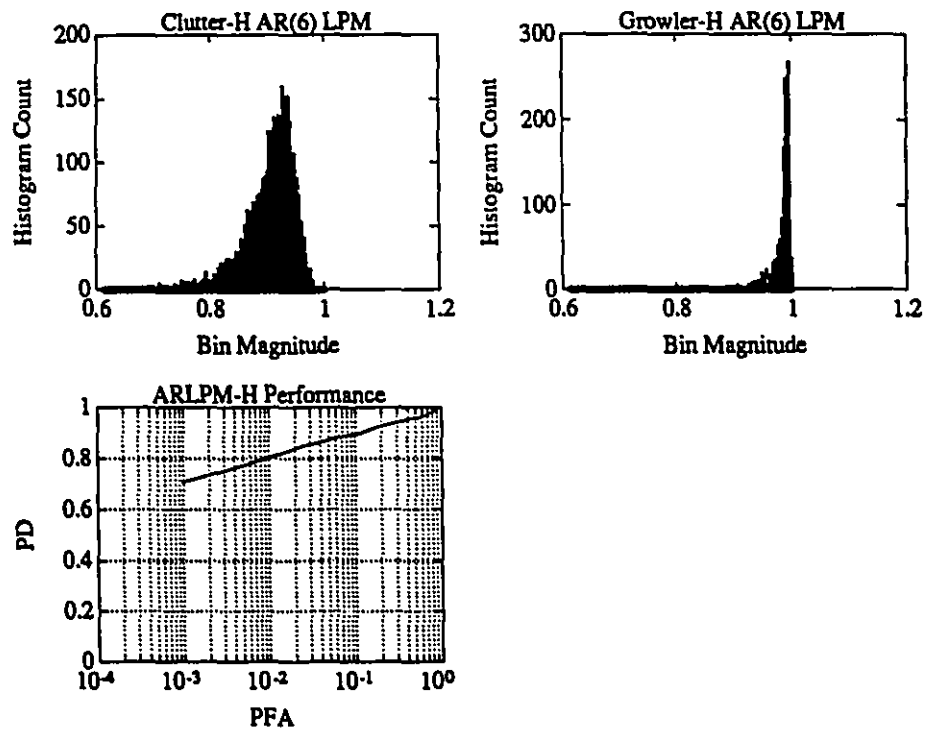


Figure 7.19: This figure gives the composite detector performance curves for the ARLPM detector in the HH channel. The composite clutter and growler largest pole magnitude histograms are shown along with the corresponding ROC curve. These composite statistics are based on the entire data testbed.

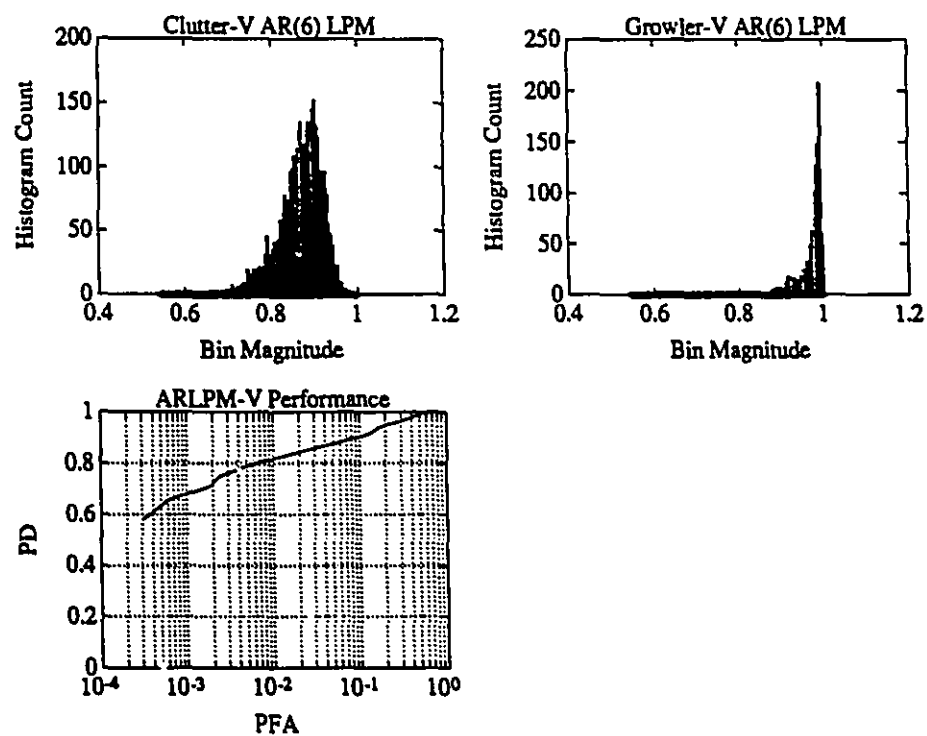


Figure 7.20: This figure gives the composite detector performance curves for the ARLPM detector in the VV channel. The composite clutter and growler largest pole magnitude histograms are shown along with the corresponding ROC curve. These composite statistics are based on the entire data test bed.




Figure 7.21: This figure gives the composite detector performance curves for the GSW detector in the HH channel. The composite clutter and growler spectral width histograms are shown along with the corresponding ROC curve. These composite statistics are based on the entire data testbed.

resulting ROC for the HH and VV channels, respectively. Here too, we have chosen to plot the PFA data on a logarithmic scale in order to accentuate the small values of false alarm rates.

In order to compare the ROC's of the three detectors for the HH polarization channel, the IA, ARLPM, and GSW ROC's shown in Figure 7.17, Figure 7.19, and Figure 7.21, respectively, have been replotted together in Figure 7.23. Similarly, the corresponding VV ROC's are replotted together in Figure 7.24.

In the next section, the performance curves of the IA, ARLPM, and GSW detectors are compared, and the reasons behind their relative performances are discussed.




Figure 7.22: This figure gives the composite detector performance curves for the GSW detector in the VV channel. The composite clutter and growler spectral width histograms are shown along with the corresponding ROC curve. These composite statistics are based on the entire data testbed.




Figure 7.23: This figure shows the composite ROC for the IA, ARLPM and GSW detectors, all on the same graph for comparison. These ROC curves are based on processing of the HH channel in the entire data testbed.




Figure 7.24: This figure shows the composite ROC for the IA, ARLPM and GSW detectors, all on the same graph for comparison. These ROC curves are based on processing of the VV channel in the entire data testbed.

7.4 Detector Performance Comparisons

Examination of Figure 7.23 and Figure 7.24 leads to the following observations:

1. The ARLPM and GSW detectors perform equally well in both the HH and VV channels. However, the IA detector definitely performs better in the HH channel. It is no wonder that currently available noncoherent marine radars operate using the HH polarization. The IA detector performs better in the HH channel because the growler to clutter ratio is larger in the HH channel than it is in the VV channel.
2. The IA detector performs worse than both the ARLPM and GSW detectors. This should come as no surprise since the IA detector is noncoherent, whereas the ARLPM and GSW detectors are both coherent detectors.
3. For probabilities of false alarm less than about 2.5 %, the ARLPM detector performs better than the GSW detector. That is, the ARLPM detector offers a higher PD for each PFA less than 2.5 %. The trend reverses though for probabilities of false alarm greater than about 2.5 %. In this region, the GSW detector has larger probabilities of detection.

It appears as though the ARLPM and GSW detectors each have their own region of superior performance, and one method would be chosen over the other simply on the basis of the required PFA. However, this simple approach neglects other important considerations. For example, the computational cost associated with the ARLPM and GSW detectors should be considered as described in Chapter 6. Furthermore, the relative performance of the two detectors in an environment different from that described by the data testbed should also influence the choice of which detector to use. After all, our evaluations have been based on a finite amount of data, over a limited range of conditions. Clearly, some care must be taken in drawing far reaching conclusions on the relative performance of the detectors in general.

As discussed in Section 7.3.1, the GCR has a tremendous effect on the performance of the ARLPM and GSW detectors. Under different weather and sea conditions, the GCR can be expected to be quite different from the average GCR represented in the data testbed used here. Therefore, in order to properly characterize the performance of the ARLPM and GSW detectors in general, their sensitivity to the GCR should be examined more closely. Since the GSW detector bases its decision solely on the spectral width of the test data, and it makes no use of the amplitudes of the fitted Gaussian spectral components, it is expected that its performance will not degrade quickly with small reductions in the GCR. The ARLPM detector, on the other hand, does make explicit use of the relative amplitudes of the clutter and growler spectral components and hence depends directly on the GCR. This dependence is due to the coupling that exists between the magnitude of an AR spectral peak, and the spectral width associated with that peak. This coupling is explicitly indicated in Equation 5.26. As a result of this dependence, we expect the ARLPM detector to be very sensitive to changes in GCR. That is, if the average GCR were to drop considerably lower than the average GCR that is representative of our data testbed, we would expect the performance of the ARLPM detector to degrade more than the performance of the GSW detector.

In order to test the validity of these performance expectations as the GCR drops, we performed a simple test using a little bit of reverse logic. Since our expectations suggest that a drop in GCR would affect the ARLPM detector more severely than the GSW detector, then clearly a rise in GCR would cause a larger improvement in the ARLPM detector performance over the GSW detector. Stated a little differently, if the GCR is reasonably large, the relative amplitude information between growler and clutter can be used to improve the performance of a detector not already using the amplitude information. That is, the GSW detector performance could be improved if amplitude information were used in addition to the spectral width information.

In Figure 7.25, a scatter plot is shown of the joint amplitude and spectral width parameters for growler data and clutter data taken from file B98 (HH). These parameters were determined by the feature extractor stage of the GSW detector, applied to the growler and clutter data. The growler parameters are indicated by 'x' and the clutter parameters by 'o' in Figure 7.25. This scatter plot indicates that the amplitude parameter as well as the spectral width parameter are useful for separating growler and clutter data. For example, the simple decision region (although not the optimal decision region) shown by the box in Figure 7.25 would result in improved growler detection over the standard GSW detector. That is, specifying that the spectral width parameter must be less than a specified threshold and the amplitude must be greater than a specified threshold before a growler is declared will provide improved growler detection when the GCR is reasonably large. We will call the detector that operates in this manner the GSW/A detector (the 'A' is for amplitude). The performance of the GSW/A detector on the data used in Figure 7.25 is indicated in Figure 7.26 by the solid line. The performance of the standard GSW detector on the same data is indicated by the dashed line in the same figure. The increase in performance by using amplitude information is very clear, supporting our expectation that the ARLPM detector will suffer more than the GSW detector in lower GCR environments. Figure 7.26 also shows the detection loss that results for PFA's greater than 1 %. This loss is due to the fact the amplitude threshold discards as clutter any growler returns with amplitudes smaller than the threshold. This loss can be reduced somewhat by choosing an optimal decision region in the amplitude-spectral width space.

The analysis just given explains why the ARLPM detector had better performance than the GSW detector for small PFA's, and poorer performance (more detection loss) for larger PFA's. The ARLPM detector took advantage of the reasonably large GCR in our data testbed by explicitly using amplitude information as part of its decision making process. However, by the arguments just given, the ARLPM detector will not be as robust to GCR variations as will the GSW detector. It is for this reason that the GSW detector was designed




Figure 7.25: This figure shows a scatter plot of the amplitude and spectral width parameters determined by the GSW detector on file B98 (HH). The growler amplitude/spectral width pairs are indicated by 'x' and the clutter amplitude/spectral width pairs are indicated by an 'o'. This scatter plot shows that there is information in the joint amplitude/spectral width statistics that is useful for classifying the growler and clutter cells in file B98. A simple (not optimal) decision boundary is shown as a box in the figure. Simply stated, a growler is declared if its amplitude/spectral width parameters fall inside of the boxed region. Otherwise, clutter is declared.




Figure 7.26: This figure shows the detector performance of the GSW and GSW/A detectors on the HH channel of file B98. The GSW detector performance is indicated by the dashed line, and the GSW/A detector performance by the solid line. The GSW bases its detection solely on the spectral width. If the spectral width is less than the specified threshold, a growler is declared; otherwise, clutter is declared. The GSW/A detector uses both an amplitude and spectral width threshold comparison as the basis for its operation. The amplitude and spectral width thresholds specified by the boxed region in Figure 7.25 are used here. Growlers are declared if the spectral width is smaller than the spectral width threshold and the amplitude is greater than the amplitude threshold; otherwise, clutter is declared.

to use spectral width as its sufficient statistic! When the GCR is expected to be quite small at times, or vary over a large range, the GSW detector is preferred as its performance will be more robust with respect to the GCR. It is expected that in situations where detection of very small growlers (smaller than Growler-1 and Growler-2) is required, or when the significant waveheight of the sea is larger than for our measurements (see Table 3.4), the GSW detector will be the detector of choice. If however the operational conditions are quite moderate as represented by our data testbed, then the ARLPM detector is still a contender if low PFA rate performance is essential. (The GSW/A detector would perform as well, however the computational cost is about double that of the ARLPM detector).

7.5 Meeting The Goals

In the previous section, we looked at composite detection performance curves in order to compare the relative performance of the IA, ARLPM, and GSW detectors. In this section, we reconsider the requirements of a *medium dwell-time* detector that we defined on page 11 of the opening chapter. Since our analysis has indicated that the GSW detector is the best performing detector under the most general conditions, we will check to see if it meets all of our stated requirements. These requirements are restated below for convenience and will be treated one at a time. The detector will:

1. provide faster detections which translate to faster updates.
2. offer improved growler detection performance over *short dwell-time* approaches.
3. be able to provide meaningful results based on the current data-base.
4. require less computation and less memory than *long dwell-time* approaches which translates to cost savings.
5. indicate the kind of performance improvements that can be expected from *long dwell-time* methods.

6. be modular in design in that it will provide the necessary features needed by a *long dwell-time* system.

The first requirement refers to *long dwell-time* techniques. This requirement is satisfied by definition since the GSW detector makes a detection decision every 0.512 s. As described in Chapter 1, long dwell time techniques coherently process data for durations on the order of the peak wave period (5-12 s). Rather than requiring several seconds before a detection decision is made, the GSW detector decides in a fraction of a second, resulting in faster update times.

The second requirement specifies that the GSW detector must provide improved growler detection over *short dwell-time* approaches. Conventional marine radars are short dwell-time processors, and they make a decision by simply comparing the received signal amplitude (or power) to a threshold. These noncoherent detectors have performance curves that are no better than the ROC of the IA detector which noncoherently integrates over 0.512 s. Clearly, Figure 7.23 and Figure 7.24 show the superior performance offered by the GSW detector for the detection of growlers. For example, Figure 7.23 indicates that for a PFA of 5 %, the IA detector offers a PD of 60 %, while the GSW detector offers a PD of 96 %!

The third requirement states that the detector should be such that meaningful results can be derived from the current IPIX data base. The emphasis here is on using real measurements to validate detector performance. The summary of empirical results and discussion given in Section 7.4, and indeed throughout this chapter, are testimony to the fact that meaningful and explainable results have been obtained from the IPIX data base, in characterizing the performance of the GSW (as well as the IA and ARLPM) detector.

The fourth requirement states that the GSW (and the IA and ARLPM detectors as well) detector should require less memory and computational requirements than long dwell-time detection techniques. The memory requirements are certainly less than long dwell-time techniques because only 0.512 s of data (512 samples) need to be stored and

operated on for each resolution cell with the GSW detector, whereas several seconds of data (several thousand samples) need to be stored and coherently operated on for long dwell-time techniques. As far as computation goes, without knowing the kind of processing that would be used for long dwell-time detection, a comparison of the computational complexity cannot be made. However, we certainly can say that if the long dwell-time detector uses sophisticated processing techniques such as those used by the GSW detector, the computational expense will be much higher.

The fifth requirements states that the GSW detector should be able to provide an indication of the amount of performance improvement that can be expected by long dwell-time Doppler methods. Whereas the GSW detector focusses on the spectral width characteristics associated with a fraction of a second of clutter and growler radar returns, long dwell-time processors would track the *winking phenomenon* by coherently processing many seconds of data. An indication of the kind of performance that can be expected by such a method is given by noncoherently integrating consecutive GSW detections. Although noncoherent integration may result in a loss with respect to coherent integration, it will certainly provide a lower bound on the performance of long dwell-time Doppler processors. In order to simplify the performance prediction procedure, we use a dual binary integration technique that is common for noncoherent processors. We simply perform five consecutive detections with the GSW detector, where each detection is based on 0.512 s of data. Therefore, with a 50 % overlap between windows, 1.536 s of data is processed. The final decision is based on a majority rule of the five individual GSW detections. That is, if at least three out of five of the GSW decisions say a growler is present, then we say a growler indeed is present; otherwise, clutter is present. We call this detector the DBI-GSW detector (**Dual Binary Integration-GSW**). Clearly, since the DBI-GSW detector processes much less data than a true long dwell-time detector, and since it uses noncoherent integration of even smaller windows of data, it can only provide a lower bound to the performance of optimal long dwell-time Doppler detectors.

In Figure 7.27, the ROC for the DBI-GSW detector is given, based on processing the entire data testbed for the HH channel. The composite GSW ROC is also shown for comparison. The performance improvement achieved simply by using a 3/5 rule on the GSW detector outputs is remarkable! For PFA's as small as 0.001, the PD is almost perfect! Similarly, the ROC for the DBI-GSW detector applied to the VV channel is given in Figure 7.28. Again, the improvements are incredible. As these figures clearly show, the GSW detector has satisfied the requirement of providing an indication of the kind of performance that can be expected from long dwell-time Doppler techniques.

Before leaving the DBI-GSW detector and going on to the sixth and final requirement, a few more comments will bring light to the significance of these results. Although the DBI-GSW ROC's were generated by processing 1.536 s of data, only durations of 0.512 s were processed coherently. Furthermore, the duration of 1.536 s was a contiguous time segment. At least the same performance would result if the five 0.512 s windows were staggered in time. This is because the time difference between windows would be larger, removing what little correlation that may have existed between consecutive windows. Therefore, the five GSW detections for a given range cell could be obtained from five consecutive scans of the radar; one GSW detection per range cell per scan. Hence, with a five-bit shift register to store the five GSW detection decisions (one from each scan) for each radar resolution cell, the DBI-GSW detector could be implemented. Furthermore, excluding the five initial scans that are required to initialize the decision shift register for each resolution cell, decisions can be updated every scan, simply by shifting out the oldest decision and shifting in the GSW detector decision from the current scan. As far as the time for updating goes then, the DBI-GSW detector behaves just like the GSW detector; the scan to scan time and update time are the same, as well as the computational requirements. In any case, the DBI-GSW detector still requires a longer time interval between independent output detection decisions. Therefore, it is approaching a long dwell-time Doppler technique. From Figure 7.27, we have been able to achieve a probability of detection of about 98 %, while




Figure 7.27: This figure shows the composite GSW detector ROC curve for the HH channel, along with the corresponding DBI-GSW ROC curve. The DBI-GSW detector shown here uses a dual binary integration procedure to improve growler classification. First, five consecutive GSW detector decisions are made. It then uses a majority rule on these five decisions. If three (or more) out of the five decisions say 'growler', then a growler is declared; otherwise clutter is declared. The improvement gained by a procedure such as this is substantial.




Figure 7.28: This figure shows the composite GSW detector ROC curve for the VV channel, along with the corresponding DBI-GSW ROC curve. The DBI-GSW detector shown here uses a dual binary integration procedure to improve growler classification. First, five consecutive GSW detector decisions are made. It then uses a majority rule on these five decisions. If three (or more) out of the five decisions say 'growler', then a growler is declared; otherwise clutter is declared. The improvement gained by a procedure such as this is substantial.

maintaining a false alarm rate of 0.1 %. Compare this detection performance to the 16 % probability of detection achieved by the IA detector (see Figure 7.23), for the same PFA. This indicates the remarkable improvement obtained over the performance of conventional marine radars!

Finally, the sixth requirement states that the GSW detector should be modular in design. That is, it should provide the necessary parameters required by a long dwell-time Doppler processor. Indeed, the GSW detector does satisfy this requirement as well. The parameters provided by the GSW detector: namely, the spectral width, the mean frequency, and the spectral amplitude, can all be passed on to a tracking operator in order to test for the presence of the winking phenomenon as described in Chapter 5. The combined GSW detector front-end and tracker backend can be viewed as a long dwell-time Doppler processor. We have already demonstrated this tracking ability as shown in Figure 5.20, where the time-varying mean frequency of the growler signal is being tracked by the parameters of the GSW detector.

In this section, we have demonstrated the superior properties of the GSW detector. We have shown that this detector satisfies all of the requirements that were specified for a *medium dwell-time* detector at the beginning of this section. Furthermore, the performance improvement of this detector over traditional noncoherent detectors is truly outstanding.

7.6 Summary

In this chapter, we have evaluated the detection performances of the IA, ARLPM, and GSW detectors that were developed in Chapter 6. We began by defining the IPIX data testbed upon which our analysis was based. In Section 7.3.1, the performance of the three detectors was shown to have a strong dependence on the GCR. The performance of the detectors improved with increasing GCR, as is expected from detection theory. In Section 7.4, the GSW and ARLPM detectors were shown to be far superior to the IA detector,

demonstrating the performance improvements that result from coherent processing. These performance evaluations were based on the entire data testbed, so they are considered to be representative of the average performance one could expect, from a given look direction. It was also shown that the ARLPM detector was more sensitive at very low PFA's, while the GSW detector was more sensitive at higher PFA's. This difference was shown to be the result of the fact that the GCR was reasonably large for the data testbed, and the ARLPM took advantage of this fact. It was also shown though that the GSW detector would perform equally well (in fact better) for small PFA's, if it too took advantage of the large GCR. However, as the GCR falls, the ARLPM performance is expected to degrade much more quickly than the GSW detector. Therefore, the GSW detector was singled out as the best performing and most robust of the three detectors. Finally, in Section 7.5, the GSW detector was checked in order to see that it satisfied the six goals originally outlined for a *medium dwell-time* detector. Indeed, all six goals were satisfied by the GSW detector. In fact, in evaluating one of the goals, the DBI-GSW detector was introduced to indicate the performance improvements that would result by processing more data. The improvement attained is incredible! The DBI-GSW detector, operating in a scanning mode, was shown to provide a PD of 98 %, while maintaining a PFA of only 0.1 %. Compare this performance to the IA detector which is representative of current marine radar technology. The PD of the IA detector was only 16 %, for the same PFA.

Chapter 8

Summary and Recommendations

8.1 Summary of Results

In this thesis, we have investigated techniques for the improved detection of growlers in sea clutter. This investigation was based on the analysis of real radar data collected with the IPIX radar, at a radar site at Cape Bonavista, Newfoundland. Our study of improved growler detectors began with a detailed characterization of the noncoherent and coherent properties of sea clutter and growler radar returns.

Several results emerged from the noncoherent analysis of growler and clutter radar returns. First and foremost, we showed from a statistical point of view that the sea clutter amplitude statistics are K-distributed. One of the key properties of the K-distribution is its ability to model *spiky* clutter. As confirmed by our analysis, sea clutter amplitude returns often appear as target-like spikes that persist for several seconds. This causes great confusion for conventional radars that assume Rayleigh distributed clutter statistics which are known to have a very small probability of observing large amplitudes, and which are expected to be fast-varying so that integration offers considerable clutter rejection. The K-distribution is a compound model that has a fast-varying Rayleigh speckle component; but this component is modulated by a slow-varying root-Gamma component which models the underlying mean

level of the sea clutter backscatter. It is this slow-varying clutter component that causes havoc to conventional marine radars, producing target-like spiky returns. Furthermore, it explains the spatial dependence on the performance of fixed threshold detectors. Since the slow-varying component is partially associated with the mean tilt of the ocean surface which is controlled by the underlying sea swell, the mean level has a spatial dependence. (Tilt contributes to the modulation, but is not the whole story; this area is the subject of active research at the moment.) Therefore, fixed threshold detectors will have false alarm rates that vary spatially.

The amplitude statistics of growlers were also examined and were shown to be quite similar to clutter amplitude statistics, indicating the poor performance of noncoherent growler detectors. Furthermore, both the like-polarized and cross-polarized channels were studied to see whether there was improved growler/clutter separation in the cross-polarized channels. No improvements were evident. Due to the undesirable spatial variation of false alarms associated with fixed threshold detections, CFAR methods were examined. Temporal CFAR detectors were ruled out because of target masking problems associated with those methods. A spatial correlation analysis of sea clutter data was done in order to determine the potential of spatial CFAR processors. However, our analysis indicated that the spatial correlation of sea clutter dropped off substantially over just a few radar resolution cells. This was due to the fact the radar resolution cell was 30 m in range. Typical swell wavelengths are between 100 m and 150 m. Therefore, each clutter cell was significantly displaced along the swell wavelength resulting in a different underlying mean level component from cell to cell. As a result, the sea clutter cells spanning a CFAR window did not have similar statistics, making traditional CFAR methods unusable with our data.

Turning to the task of quantitatively evaluating noncoherent detection performance of growlers, a *medium dwell-time* detector called the IA detector was devised and tested on the IPIX data testbed. The IA detector based its decision on about 0.5 s of data, by simply integrating the amplitude returns over this period. It was determined that the HH channel

offered better performance over the VV channel, as the clutter was shown to vary less in the HH channel. Nevertheless, the detection performance was shown to be quite poor. For example, with a PFA of 5 %, the LA detector offered a PD of only 60 %.

The coherent studies were just as fruitful as the noncoherent studies. Our analysis verified the presence of the so called *winking phenomenon* that we discovered during our field trials with the IPIX radar. The Doppler returns from a floating growler were characterized by an almost periodic variation in center frequency with time. As further study showed, the growler motion can be modeled by the hydrodynamics associated with ocean waves. This discovery laid the foundations for improved Doppler detection of growlers in sea clutter.

In order to develop *medium dwell-time* Doppler detectors that process a fraction of a second of data before making a decision, the Doppler spectra of sea clutter and growlers were examined in detail. It was determined that a Gaussian-shaped (1G) spectrum model closely matches the estimated spectra from sea clutter data. Also, a two-Gaussian (2G) model is a good fit to the growler-in-clutter spectra, where one of the Gaussians is associated with the clutter returns, and the other Gaussian is associated with the growler returns. Our studies indicate that the sufficient statistic that distinguishes the growler and clutter spectra is the spectral width associated with the Gaussian components. Growlers have spectral widths that are typically less than 6 Hz, while sea clutter tends to have much larger spectral widths. With the basis for a Doppler detector decided upon, two different detectors were designed to exploit the spectral width differences between growler and clutter data. The first detector is based on comparing the largest pole magnitude from an AR modeling of the data to a threshold. The closer the pole is to the unit circle, the more likely it is due to a growler. This detector is called the AutoRegressive Largest Pole Magnitude detector. The second coherent detector is based directly on the spectral width sufficient statistic. A spectrum estimate is formed from the data, and a 1G and 2G spectrum model are both fitted to the spectrum estimate in the nonlinear least squares sense. The model with the *best* fit is chosen as the spectrum model, and its parameters (amplitude, mean frequency, and spectral

width) are passed on to the detector. If a Gaussian component in the model has a spectral width less than a preselected threshold, a growler is declared; otherwise, clutter is declared. This detector is known as the Gaussian Spectral Width detector.

Both the ARLPM and GSW detectors were evaluated on the IPIX data testbed. Both detectors performed equally well in the HH and VV channels. Although both the ARLPM and GSW detectors are far superior to the IA detector, the GSW detector is the more robust of the two. The GSW detector is recommended in low GCR environments. For the GSW detector, a PD of 96 % was realized for a PFA of 5 % (The ARLPM detector had a PD of 87 %). It was also shown that the PFA could be reduced drastically by performing several GSW detections consecutively, and only declaring a growler if the majority of the detections indicated a growler. For example, if a growler was declared only when three or more out of five consecutive GSW decisions indicated the presence of a growler, a PD of 98 % was achieved while suffering a PFA of only 0.1 %! Compare this performance with the 16 % PD for the IA detector when the PFA equals 0.1 %. There is no doubt about the performance improvements that we have made with the GSW detector.

A few more comments are noteworthy on the GSW detector. Although the GSW detector as defined in Chapter 6, makes use of only the spectral width parameter, the amplitude and mean frequency parameters are also available if required. For example, we have shown the amplitude and spectral width parameters to be perfectly suited for making GCR, GNR, CNR, RCS, and absolute power measurements which allow estimates of growler RCS and clutter σ_o . Also, all of these parameters could be passed on to a tracking system which could track the mean frequency variation in order to identify the target based on its dynamics. The combination of the GSW stage and the tracking stage may be viewed together as a long dwell-time Doppler detector. Data reduction has also been achieved since in our case, 512 samples have been reduced to three or six parameters. Thus, the GSW detector can be considered to be very modular and versatile, and although it is computationally an expensive operator, many added benefits are bundled *for free*.

8.2 Contributions of this Research

The contributions of this research may be summarized as follows:

1. The collection of real, coherent, dual-polarized, X-band sea clutter and growler-in-sea clutter radar backscatter data. This database has been catalogued and summarized in the thesis. Associated radar parameters are indicated, along with extensive surface truthing data. Furthermore, supplemental calculated data such as GCR, growler RCS, and sea clutter σ_o data are also given. As well, a complete chapter is devoted to the experimental effort so that all of the pertinent information is put together in an organized fashion. This arrangement is valuable to other researchers who may wish to build on the results presented herein; especially since data and extensive ground-truthing information are rarely found together in the literature. Furthermore, there is very little low grazing angle X-band sea clutter data published in the literature, let alone growler data.
2. Verification of the applicability of the K-distribution for modeling the amplitude statistics of very low grazing angle sea clutter, and its use in performance prediction.
3. Discovery of the *winking phenomenon* associated with growler radar returns.
4. Characterization of the time-varying Doppler spectra of sea clutter and growler data and identification of spectral width as a sufficient detection statistic.
5. Development of two detection algorithms and ROC curves based on the IPIX data base, that offer substantial improvements over conventional noncoherent techniques. These are the AutoRegressive Largest Pole Magnitude detector and the Gaussian Spectral Width detector. In particular, the GSW detector offers performance, robustness, and modularity giving the system designer maximum freedom. Furthermore, the timeliness and scope of this study suit well the needs of the recent Canadian effort for the Hibernia development.

6. Development of calibration, data validation, data handling, and analytical software library that is portable and runs directly on the radar system computer. The software is designed to operate on the data, in place, as they are collected. This library will be very helpful to graduate researchers who are working on the IPIX radar project.
7. Publication of several conference papers and journal papers which report on the IPIX radar, the data collection effort, and results obtained [7, 8, 10, 11, 12, 13, 14, 15].

8.3 Conclusions

In the last two sections, highlights of the results of our study on the detection of growlers in sea clutter, along with the contributions to research that resulted from this work were presented. As always, any attempt to steer a study along all possible paths, while at the same time exhausting all courses of action is futile, especially with limited time and resources. It is important therefore, to focus on a few issues, with enough depth that meaningful results emerge, leaving all the other issues for another time, or another researcher. This approach has been taken in this thesis. At this time, it is appropriate to discuss the conclusions that we have arrived at, as a result of this work. These in turn will point to some of the areas in the growler detection problem that still require investigation.

Our results indicate that improved growler detection is achievable using the coherent properties of growler and clutter radar returns. Although medium dwell-time Doppler detectors such as the GSW detector offer a very significant improvement in growler detection performance over conventional techniques, it is necessary to integrate larger amounts of data to reduce the probability of false alarm rate even further. Fortunately, our research indicates that by integrating GSW detections on a scan-to-scan basis (like the DBI-GSW detector), the performance continues to improve. Indeed, our results show that long dwell-time Doppler techniques (such as the DBI-GSW detector with many scan-to-scan integrations) are required to meet stringent operational requirements. This author feels that a continued

effort to exploit the Doppler characteristics identified in this thesis will be fruitful, resulting in the development of a modern marine radar capable of reliable growler detection in sea clutter. This conclusion justifies a continued effort in the development of long dwell-time Doppler methods.

Our studies were limited largely to Doppler processing. Although we made use of independent polarization channels, polarimetric processing based on scattering matrix measurements has not been considered here, as it requires a very large effort for it to be treated properly. Preliminary polarimetric studies at the CRL have indicated that there is additional discrimination between growler and clutter returns, supporting polarimetric studies as an area of future research.

Other radar discriminants such as multi-frequency illumination offer the potential to distinguish rigid targets (e.g. growlers) from fluid targets (e.g. sea). IPIX is in fact in the process of being upgraded to multi-frequency operation (one fixed, one agile). The addition of multi-frequency measurements offer the promise of further improvements to growler detectability in sea clutter.

Another open area for future research concerns the manner in which various discriminators should be combined in an overall detection system. From a theoretical as well as practical standpoint, the joint processing of all radar variables is optimal, and would be highly desirable. There is a lot of room for the development of strategies for combining various measurements in a fashion that will be practically achievable, while at the same time offering large performance gains.

Finally, the last area for future research involves expanding the database upon which the results presented herein are based. It is important to have growler and clutter measurements under a wide variety of environmental conditions in order that conclusions are realistic. Although our database covers a limited range of growlers, sea conditions, and wind conditions, efforts were made to make our results as generally applicable as possible. We have indeed been able to achieve a remarkable improvement in growler detection using

medium dwell-time Doppler techniques. However, more data are required to adequately evaluate the performance of long dwell-time Doppler techniques. Our analysis began in the field where the problem lives, with an experimental effort to collect data. Following data collection, the analysis phase took centre stage, closing with recommendations. However, it is only after we come full circle, and our methods are validated in the field, that we may be invited to take a bow.

Appendix A

Weather and Wave Height Records

This appendix contains weather and wave height records that cover the period of the IPIX field trials, between the dates of June 2, 1989 and June 14, 1989. This data is presented in the form of graphs. The weather information plotted on the graphs was obtained from the AES Weather Office in Cape Bonavista, Newfoundland, and the wave height and wave period information were recorded by a nondirectional waverider that was moored 6.75 km offshore at a bearing of 72 degrees. The data from these two sources was then compiled by Brian Currie and are plotted here in the form of graphs.

A.1 Weather Records

We begin with a plot of the Experimental ID for each collected data file, versus the date and time the data file was collected. This plot is given in Figure A.1 and will enable the matching of a particular data file with the recorded weather and wave conditions. This information is then followed by wind speed, wind direction, visibility, temperature, and pressure data plotted in Figure A.2, Figure A.3, Figure A.4, Figure A.5, and Figure A.6




Figure A.1: This figure shows the date and time that each experiment was performed. The experiments are identified by an Experiment ID, and this is plotted against the corresponding date and time. The date May 29, 1989 at 12:01 am is normalized to day 1.

respectively. For simplicity of notation, a day number is plotted instead of a date and a time label. Day 1 corresponds to May 29, 1989 at 12:01 am. For example, a day number of 2.5 corresponds to May 30, 1989 at 12:00 noon.

A.2 Wave Height Records

In this section, wave height and wave period information is presented. Figure A.7 shows the significant and maximum wave height measurements recorded during our field trials with IPIX. This is followed in Figure A.8 with a plot of the wave period as compared to the significant wave height data. In Figure A.9, we plot a comparison of the wind speed data given in Figure A.2 and the significant wave height data given in Figure A.7. The significant wave height represents that height exceeded by 30 % of the wave crests, while the maximum wave height is simply the maximum measured waveheight reported by the waverider during each observation interval.




Figure A.2: This figure shows the wind speed in knots as a function of time during the period May 29, 1989 to June 13, 1989 as measured by the AES Weather Office at Cape Bonavista, Newfoundland.




Figure A.3: This figure shows the wind direction in degrees as a function of time during the period May 29, 1989 to June 13, 1989 as measured by the AES Weather Office at Cape Bonavista, Newfoundland. A bearing of 0 degrees is north.




Figure A.4: This figure shows the visibility in miles as a function of time during the period May 29, 1989 to June 13, 1989 as measured by the AES Weather Office at Cape Bonavista, Newfoundland.




Figure A.5: This figure shows the temperature in degrees C as a function of time during the period May 29, 1989 to June 13, 1989 as measured by the AES Weather Office at Cape Bonavista, Newfoundland.




Figure A.6: This figure shows the sea-level pressure in kPa as a function of time during the period May 29, 1989 to June 13, 1989 as measured by the AES Weather Office at Cape Bonavista, Newfoundland.




Figure A.7: This figure shows the significant and maximum waveheight in metres, plotted as a function of time during the period May 29, 1989 to June 13, 1989. This data was measured by a nondirectional waverider.




Figure A.8: This figure shows the wave period in seconds along with the significant wave height in metres, plotted as a function of time during the period May 29, 1989 to June 13, 1989. This data was measured by a nondirectional waverider.




Figure A.9: This figure shows the significant waveheight in metres along with the wind speed in knots, plotted as a function of time during the period May 29, 1989 to June 13, 1989. This data was measured by a nondirectional waverider and the AES Weather Office.

Appendix B

IPIX Calibrations

In this appendix, the I and Q calibrations and the input/output calibrations are described. The I and Q calibration involves correcting amplitude and phase imbalances that arise from mismatches in the quadrature demodulator section of the IPIX receiver. The input/output calibration involves calculation and removal of the effective gain applied by the receiver to the input signal.

B.1 I and Q Calibration

In a coherent radar such as the IPIX radar, the received signal at RF is mixed down to IF and then is passed through a *quadrature demodulator* stage. The quadrature demodulator multiplies the IF signal with two coherent IF sources shifted in phase by 90 degrees (I and Q sources), and then lowpass filters the result. The effect of this operation is to mix the IF signals down to baseband, retaining only the difference frequency components. If the quadrature demodulator performs ideally, a pure baseband tone on the IF carrier would result with a complex exponential baseband signal at the baseband tone frequency. Plotting the resulting I and Q baseband signals against each other would result in a perfect circle.

Quadrature demodulators in practice do not perform ideally. The coherent I and

I and Q sources are never perfectly orthogonal and gain imbalances between the two channels are common. The effect of the phase and amplitude imbalance on a pure baseband tone is to create an image component at the negative of the tone frequency, along with the normal baseband tone. The larger the imbalances, the larger is the magnitude of the image component.

Along with the amplitude and phase imbalances, dc offsets also result from various components in the receiver. The point of I and Q calibrations is to correct for the amplitude and phase imbalances as well as the dc offsets that are found in the system. In this section, we consider two different approaches that provide these corrections.

The first method that we consider is a generalization of the method described by Churchill *et al.* [71]. The output of the nonideal quadrature demodulator is given by

$$\begin{aligned} I(t) &= (1 + \epsilon)A \cos \omega t + d_i \\ Q(t) &= A \sin(\omega t + \phi) + d_q \end{aligned} \quad (\text{B.1})$$

where ϵ , ϕ , d_i , d_q are the amplitude imbalance, phase imbalance, and dc offsets respectively, and ω is the baseband frequency. The dc offsets d_i and d_q are easily corrected by removing the average value of $I(t)$ and $Q(t)$ respectively. Therefore, correcting these in Equation B.1 and dropping the time dependence for simplicity of notation, we get

$$\begin{aligned} I &= (1 + \epsilon)A \cos \omega t \\ Q &= A \sin(\omega t + \phi). \end{aligned} \quad (\text{B.2})$$

If I and Q are viewed as components of a 2-dimensional vector, then a rotation and scaling operation will correct for the phase and amplitude imbalances. Such a correction can be written in matrix form as

$$\begin{bmatrix} I_c \\ Q_c \end{bmatrix} = \begin{bmatrix} E & 0 \\ P & 1 \end{bmatrix} \begin{bmatrix} I \\ Q \end{bmatrix}, \quad (\text{B.3})$$

where I_c and Q_c are the corrected I and Q components and E and P are given by

$$\begin{aligned} E &= \frac{\cos \phi}{(1 + \epsilon)} \\ P &= -\frac{\sin \phi}{(1 + \epsilon)}. \end{aligned} \quad (\text{B.4})$$

In order to make use of Equation B.3, we need to determine the values of E and P for our system. One way to do this is to inject a known test signal into the receiver and measure the image component that results from the imbalances in the system. For example, consider a sampling system that samples the I and Q components with a frequency f_s . Furthermore, assume that N samples of the test tone are collected, and that the test tone frequency is given by $\omega = 2\pi f_s u/N$, where u is an integer between 1 and $N/2$ so that the baseband test tone and its image both fall on a DFT (Discrete Fourier Transform) bin. It is easy to show that in the absence of noise, E and P are exactly given by

$$\begin{aligned} E &= -\text{Real} \left[\frac{2F_{N-u}}{F_u^* + F_{N-u}} \right] + 1 \\ P &= -\text{Imag} \left[\frac{2F_{N-u}}{F_u^* + F_{N-u}} \right], \end{aligned} \quad (\text{B.5})$$

where F_k represents the k^{th} point in an N -point DFT. In other words F_u is the DFT value at the tone frequency and F_{N-u} is the DFT value at the image frequency. Therefore, any test tone can be used so long as its frequency falls on a valid DFT bin. Another way of saying this is that a full number of cycles of the tone must be processed by the DFT. This is a generalization of the method described by Churchill *et al.* in [71]. They derive equations for the case where the test tone has a single cycle processed by the DFT.

The i-q correction procedure described above works very well when the frequency of the test tone can be controlled precisely. A test source such as this is often not available and so an alternate procedure is desirable. We have derived such a procedure based on maximum likelihood estimation theory.

In cases where the test tone is not known accurately, a reasonable approach would

be to estimate this tone from the output data. Consider an imbalanced receiver with zero-mean AWGN which is typical of IPIX (and other radar receivers). With a test tone as the input source, the digital output samples are modelled as:

$$r(nT_s) = s(nT_s) + e(nT_s) \quad (\text{B.6})$$

where

$$s(nT_s) = a_1 \exp(-j\omega nT_s) + a_2 \exp(j\omega nT_s) + d. \quad (\text{B.7})$$

T_s is the sampling period, ω is the baseband frequency of the tone, d is the complex dc offset (ie. the real part is the I channel dc offset and the imaginary part is the Q channel dc offset), and $e(nT_s)$ represents the complex AWGN. We recognize a_1 as the DFT value at the image frequency and a_2 as the DFT value at the tone frequency. Not knowing the baseband frequency ω of the test tone, use of Equation B.3 requires estimation of a_1, a_2, d , and ω , based on the received data $r(nT_s)$. Because of the excellent properties of ML estimates, we compute ML estimates of these unknown parameters. Since the noise model is zero-mean AWGN, the ML estimates are equivalent to the *minimum least squares* estimates of the parameters [9, section 7.9]. In order to simplify the computational requirements of these estimates, we recognize that only the parameter ω appears nonlinearly in the model described by Equation B.7. Since the parameters of a linear model can be solved for in a single step, we assume for the moment that an estimate of ω given by $\hat{\omega}$ is available. If N is the number of samples of the received signal available, the signal model defined in Equation B.7 can be rewritten in vector notation as

$$\mathbf{s} = \mathbf{M}\boldsymbol{\alpha} \quad (\text{B.8})$$

where the matrix M is given by

$$M = \begin{bmatrix} \exp[j\hat{\omega}(N-1)/2] & \exp[-j\hat{\omega}(N-1)/2] & 1 \\ \vdots & \vdots & \vdots \\ 1 & 1 & 1 \\ \vdots & \vdots & \vdots \\ \exp[-j\hat{\omega}(N+1)/2] & \exp[j\hat{\omega}(N+1)/2] & 1 \end{bmatrix} \quad (\text{B.9})$$

and the parameter vector α is

$$\alpha = [a_1, a_2, d]^T. \quad (\text{B.10})$$

The problem now is to solve for the minimum least square or ML estimates of the parameter vector α . If we define the error function as

$$\epsilon^2 \triangleq \|r - s\|^2, \quad (\text{B.11})$$

and the conditional parameter vector α_c as

$$\alpha_c = [\alpha|\hat{\omega}], \quad (\text{B.12})$$

then minimizing ϵ^2 with respect to α_c provides the ML estimates of α . It can be shown [9] that this minimization results with the solution:

$$\alpha_{ML} = (M^H M)^{-1} M^H r. \quad (\text{B.13})$$

Since the result given in Equation B.13 assumes that $\hat{\omega}$ is available, in practice, the optimization would be performed as follows:

1. Pick a starting value for $\hat{\omega}$. This should be easy since the test tone will be known approximately.
2. Evaluate Equation B.13 to solve for the parameter vector α_{ML} .
3. Evaluate the error function given by Equation B.11 using the current parameter estimates.

4. Use the error function value and the current estimate of $\hat{\omega}$ along with the corresponding previous values to update the value of $\hat{\omega}$. For example, use the Golden Section Method described in [69].
5. Repeat the last three steps until the error function defined in Equation B.11 is at a minimum within the required tolerance.

A reduction in the computation of α_{ML} can be realized if N is odd. For this case, $\mathbf{M}^H\mathbf{M}$ reduces to a 3x3 real matrix given by

$$\mathbf{M}^H\mathbf{M} = \begin{bmatrix} N & \frac{\sin(\hat{\omega}NT_s)}{\sin(\hat{\omega}T_s)} & \frac{\sin(\hat{\omega}NT_s/2)}{\sin(\hat{\omega}T_s/2)} \\ \frac{\sin(\hat{\omega}NT_s)}{\sin(\hat{\omega}T_s)} & N & \frac{\sin(\hat{\omega}NT_s/2)}{\sin(\hat{\omega}T_s/2)} \\ \frac{\sin(\hat{\omega}NT_s/2)}{\sin(\hat{\omega}T_s/2)} & \frac{\sin(\hat{\omega}NT_s/2)}{\sin(\hat{\omega}T_s/2)} & N \end{bmatrix}, \quad (\text{B.14})$$

where we have made use of the identity

$$\sum_{n=-(N-1)/2}^{(N-1)/2} \exp(j\beta nT_s) = \frac{\sin(\beta NT_s/2)}{\sin(\beta T_s/2)}. \quad (\text{B.15})$$

We now give an example of this ML calibration procedure using real IPLX data. A test tone of approximately -1 MHz was injected into the IPLX receiver and the I and Q outputs were recorded. A sample of this data is shown plotted in Figure B.1. The phase imbalance is easily seen. The ML calibration procedure described above is applied to the data, and the data is corrected as described in Equation B.3. This corrected data is shown plotted in Figure B.2.

B.2 Input/Output Calibration

The input/output calibration allows us to map the digital I and Q samples at the output of the receiver onto the real powers that were received at the antenna. In the IPLX radar, the built-in calibration equipment (BICE) can be used to evaluate the input/output characteristic automatically. By varying a known internal test source and a digitally controlled




Figure B.1: Uncalibrated I and Q test tone plotted as I versus Q. The test tone of approximately -1 MHz in frequency was injected into the receiver, and sampled at 30 MHz.




Figure B.2: Calibrated I and Q test tone plotted as I versus Q. The test tone of approximately -1 MHz in frequency was injected into the receiver, and sampled at 30 MHz. Then the amplitude and phase imbalances were corrected using the ML calibration procedure.

attenuator, several different power levels can be injected into the receiver and the outputs measured. In order to determine the actual input power levels, a power meter is used at a convenient reference port. Then nominal losses are used for waveguide and directional couplers that precede the reference port. In this way, approximate input powers are derived. By adding the sum of the squares of the I and Q components at the output of the receiver, we arrive at the output power. Then a plot of the input/output power curve can be made.

An alternative procedure is to use an external, variable, RF power source to inject known powers into the IPIX receiver. In this way all losses in the system are included. Again output powers are calculated and the input/output characteristic is plotted. Using this procedure with the IPIX STC's in the maximum gain position, the input/output characteristic shown in Figure B.3 was produced. This characteristic is valid for all of the experiments described in this thesis. For the most part, the received data fell within the linear region of this characteristic. This allows for a simple mapping of output to input powers. That is, providing that we are operating in the linear region of the receiver, a doubling in the output power implies a doubling in the input power. Furthermore, if a convenient point such as the noise floor of the receiver were known at the input, output signal powers could be mapped onto the input simply by adding the output signal power to output noise power ratio to the noise power at the input of the receiver (the power units are in dB's). The input/output characteristics given in Figure B.3 did not take into account the waveguide losses between the receiver input and the antenna. These losses were measured to be 4 dB one way. From Figure B.3, the noise floor is seen to be around -96 dBm at the receiver input. This is therefore equivalent to -92 dBm at the antenna. This input noise level agrees well with that predicted in theory. The bandwidth of the IPIX receiver was determined to be approximately 25.1 MHz (3 dB). We will use this figure as the effective noise bandwidth B as described in [20]. The theoretical noise power at the receiver input is given by

$$N_i = kT_oBN_f \quad (\text{B.16})$$

where N_i is the input noise power in Watts, k is Boltman's constant equal to $1.38\text{e-}23$ Watts/Hz, T_o is room temperature at 290 K, and N_f is the noise figure of the receiver which is approximately 4.17 dB for IPIX. Substitution into Equation B.16 gives a theoretical input noise power of -95.8 dBm for IPIX.




Figure B.3: The input/output characteristic of the IPIX radar with the STC's at maximum gain. The horizontal channel is indicated by 'o' and the vertical channel by 'x'. This characteristic is valid for the experiments described in this thesis.

Appendix C

Confidence on Empirical Amplitude Statistics

In this appendix, the question of confidence intervals on empirical estimates of amplitude statistics is addressed. The confidence intervals on the empirical estimates of the cumulative distribution function $F_X(x)$ for the observable X will be solved, rather than for the probability density function of X , $p_X(x)$, due to simplicity.

C.1 Confidence Limits

Consider a set of independent, identically distributed variates $X_i, i = 1, \dots, N$ with known cumulative distribution function $F_X(x)$. An empirical estimate $\hat{F}_X(x)$ of $F_X(x)$ is to be determined from the variates X_i . Suppose that we choose to determine $\hat{F}_X(x)$ at M equally spaced points $x = jb$ where $j = 1, \dots, M$ and b is the bin size or sampling interval of $\hat{F}_X(x)$. Now, imagine that it is known that the true distribution function $F_X(x)$ evaluated at the point $x = kb$, where k is some constant, is given by

$$F_X(kb) = \frac{c}{N} \tag{C.1}$$

and c happens to be an integer. Further suppose that the estimate of this value based on N variates X_i just happens to yield the same result. That is,

$$\hat{F}_X(kb) = \frac{c}{N}. \quad (\text{C.2})$$

The question then arises, what confidence can be placed on the estimate $\hat{F}_X(kb)$? Let's examine just how the estimate $\hat{F}_X(kb)$ is determined empirically. Clearly, this value is estimated by counting the number of X_i 's that were less than or equal to kb in value. The estimate of c/N implies that c out of the N measured X_i 's were less than or equal to kb . Stated in another way, this is just the probability of c successes in a set of N binary trials, where the probability of success is simply $F_X(x = kb)$, and is given by,

$$\text{Prob}[\hat{F}_X(kb) = \frac{c}{N}] = \binom{N}{c} F_X(x)^c [1 - F_X(x)]^{N-c} \quad (\text{C.3})$$

where $\binom{N}{c}$ is the number of different ways c items can be chosen out of a set of N .

Equation C.3 is the desired result. However, we may generalize this result to an interval rather than a single value by writing

$$\frac{r}{N} \leq \hat{F}_X(kb) < \frac{s}{N} \quad (\text{C.4})$$

where r and s are integers. It is easily seen that the probability of measuring $\hat{F}_X(kb)$ in the interval $\{r/N, s/N\}$ is given by the summation of Equation C.3 over each point in the interval:

$$\Pi(r, s, N, F_X(kb)) = \sum_{m=r}^{s-1} \binom{N}{m} F_X(x)^m [1 - F_X(x)]^{N-m}. \quad (\text{C.5})$$

Now let us apply Equation C.5 to the data presented in this thesis. A question that we might ask is: What is the confidence in our measurements of the cumulative distributions (we present probability densities instead) for a cumulative probability of 0.999? The IPLX radar PRF was 200Hz, and an average of about 60000 samples was used for each amplitude

distribution curve. Using a 10ms decorrelation time based only on the speckle component, there are about 30000 independent samples in each calculation. Using Equation C.5, error bars between 0.997 and 0.9993 give a 95 % confidence for the 0.999 cumulative probability. If, on the other hand, we take into account the underlying component which can persist for a couple of seconds or more, we will have much fewer independent samples. Consider a decorrelation time of 1 second. Then 60000 samples give rise to only 300 independent samples. Now using Equation C.5, error bars between 0.967 and 0.983 give a confidence of only 47 % for the 0.98 cumulative probability.

Appendix D

Nonlinear Least Squares Spectrum Estimation

In this appendix, the solution to a nonlinear least squares minimization is derived. The problem we consider is that of fitting a spectrum model to a spectrum estimate determined from some measured data. We will develop the optimization equations for a growler-in-clutter spectrum. The spectrum model that we use to model a growler in sea clutter consists of two Gaussian components, each with unknown amplitude, mean frequency and spectral width. (A clutter-only model would just include a single Gaussian component). Also, there is another parameter, N_o , which describes the receiver white noise level, and can be independently estimated to save computation.

D.1 Nonlinear Least Squares Solution

Let $S(\Theta)$ denote the N -vector of the deterministic spectrum model evaluated at the N discrete frequencies $\mathbf{f} = [f_1, f_2, \dots, f_N]^T$ Hz. The parameter vector of the model, Θ , is a 6-vector (we have assumed that the white noise power, N_o , is provided by some other means) given by $\Theta = [a_1, \bar{f}_1, \sigma_1, a_2, \bar{f}_2, \sigma_2]^T$ where a_1 and a_2 , \bar{f}_1 and \bar{f}_2 , and σ_1 and σ_2 are

the amplitudes, mean frequencies and spectral widths, respectively, of the two Gaussian model. Thus,

$$S(\Theta_k) = a_1 e^{\frac{-(f_k - f_1)^2}{2\sigma_1^2}} + a_2 e^{\frac{-(f_k - f_2)^2}{2\sigma_2^2}} + N_o \quad (\text{D.1})$$

where $S(\Theta_k)$ is the k^{th} element of $S(\Theta)$ and f_k is the k^{th} element of \mathbf{f} , $k = 1, \dots, N$. We assume that we have an estimate of the clutter Doppler spectrum given by $\hat{S}(\mathbf{f})$, and evaluated at the same set of discrete frequencies defined above. The squared error between our model and the spectrum estimate is

$$\varepsilon(\Theta) = |\hat{S}(\mathbf{f}) - S(\Theta)|^2. \quad (\text{D.2})$$

The minimization of Equation D.2 with respect to the parameter vector Θ yields the parameter estimates. This is a nonlinear, least squares optimization problem and its numerical solution is now described. The numerical solution that we derive is a Gauss-Newton approach. It is described in more depth with other valid approaches in [72]. We begin by defining a residual function $\mathbf{R}(\Theta)$ given by

$$\mathbf{R}(\Theta) = \mathbf{S}(\Theta) - \hat{\mathbf{S}}(\mathbf{f}). \quad (\text{D.3})$$

Therefore, the squared error that we wish to minimize can be written as

$$\varepsilon(\Theta) = \frac{1}{2} \mathbf{R}(\Theta)^T \mathbf{R}(\Theta) \quad (\text{D.4})$$

where the factor of 1/2 is introduced for convenience. A linear expansion of $\mathbf{R}(\Theta)$ around the point Θ_c can be written as

$$M_c(\Theta) = \mathbf{R}(\Theta_c) + \mathbf{J}(\Theta_c)(\Theta - \Theta_c), \quad (\text{D.5})$$

where $\mathbf{J}(\Theta_c)$ is the Jacobian matrix of the model evaluated at Θ_c . The ij^{th} element of the Jacobian matrix is given by

$$\mathbf{J}(\Theta) = \frac{\partial r_i(\Theta)}{\partial \theta_j}, \quad (\text{D.6})$$

where in general, the residual function $\mathbf{R}(\Theta)$ is written as

$$\mathbf{R}(\Theta) = [r_1, r_2, \dots, r_N]^T, \quad (\text{D.7})$$

and the parameter vector is written as

$$\Theta = [\theta_1, \theta_2, \dots, \theta_6]^T. \quad (\text{D.8})$$

For the model specified in Equation D.1, the elements of the Jacobian matrix are given by:

$$\begin{aligned} \frac{\partial r_i(\Theta)}{\partial \theta_1} &= e^{-\frac{(f_i - \bar{f}_1)^2}{2\sigma_1^2}} \\ \frac{\partial r_i(\Theta)}{\partial \theta_2} &= \frac{\alpha_1}{\sigma_1^2} (f_i - \bar{f}_1) e^{-\frac{(f_i - \bar{f}_1)^2}{2\sigma_1^2}} \\ \frac{\partial r_i(\Theta)}{\partial \theta_3} &= \frac{\alpha_1 (f_i - \bar{f}_1)^2}{\sigma_1^3} e^{-\frac{(f_i - \bar{f}_1)^2}{2\sigma_1^2}} \\ \frac{\partial r_i(\Theta)}{\partial \theta_4} &= e^{-\frac{(f_i - \bar{f}_2)^2}{2\sigma_2^2}} \\ \frac{\partial r_i(\Theta)}{\partial \theta_5} &= \frac{\alpha_2}{\sigma_2^2} (f_i - \bar{f}_2) e^{-\frac{(f_i - \bar{f}_2)^2}{2\sigma_2^2}} \\ \frac{\partial r_i(\Theta)}{\partial \theta_6} &= \frac{\alpha_2 (f_i - \bar{f}_2)^2}{\sigma_2^3} e^{-\frac{(f_i - \bar{f}_2)^2}{2\sigma_2^2}}. \end{aligned} \quad (\text{D.9})$$

It is easy to show that the gradient vector of the squared error function described in Equation D.2 is given by

$$\nabla \varepsilon(\Theta) = \mathbf{J}(\Theta)^T \mathbf{R}(\Theta). \quad (\text{D.10})$$

Now, desiring at each step that the error $M_c(\Theta)$ goes to zero, we get the linear model

$$\mathbf{J}(\Theta_c)(\Theta - \Theta_c) = -\mathbf{R}(\Theta_c). \quad (\text{D.11})$$

We would like to solve for the optimum weight vector Θ so that the squared error defined in Equation D.2 is minimized. The solution at each step, from linear systems theory, is given by [9]:

$$(\Theta - \Theta_c) = -(\mathbf{J}(\Theta_c)^T \mathbf{J}(\Theta_c))^{-1} \mathbf{J}(\Theta_c)^T \mathbf{R}(\Theta_c). \quad (\text{D.12})$$

Recognizing the psuedo-inverse of the Jacobian matrix in Equation D.12, namely

$$\mathbf{J}(\Theta_c)^\# = (\mathbf{J}(\Theta_c)^T \mathbf{J}(\Theta_c))^{-1} \mathbf{J}(\Theta_c)^T, \quad (\text{D.13})$$

and isolating Θ in Equation D.12, we arrive at an update equation for the numerical solution of the nonlinear least squares problem:

$$\Theta_+ = \Theta_c - \mathbf{J}(\Theta_c)^\# \mathbf{R}(\Theta_c), \quad (\text{D.14})$$

where we have replaced Θ with Θ_+ to emphasize that this is an update equation.

The psuedo-inverse given in Equation D.13 can be solved for efficiently by performing a singular value decomposition on the Jacobian matrix. The singular value decomposition can be written as

$$\mathbf{U}^T \mathbf{J}(\Theta) \mathbf{V} = \begin{bmatrix} \mathbf{S} \\ \mathbf{0} \end{bmatrix} \quad (\text{D.15})$$

where \mathbf{U} is the $N \times N$ matrix of left singular vectors, \mathbf{V} is the 6×6 matrix of right singular vectors, and \mathbf{S} is the 6×6 matrix of singular values. Then, it is easy to show that the psuedo-inverse is given by

$$\mathbf{J}(\Theta_c)^\# = \mathbf{V} \begin{bmatrix} \mathbf{S}^{-1} \\ \mathbf{0} \end{bmatrix} \mathbf{U}^T. \quad (\text{D.16})$$

Further simplifications result by recognizing that the amplitude parameters, a_1 and a_2 occur linearly in the model. These can therefore be replaced with their functional solutions as described in [73], resulting in fewer computations.

The optimization procedure can be summarized as follows:

1. Pick suitable starting values for the parameter vector Θ and denote these by Θ_c .
2. Then use Equation D.14 to solve for the next value of Θ , Θ_+ .
3. Replace Θ_c with Θ_+ in Equation D.14 and solve for new parameter vector Θ_+ .
4. Repeat Step 3 until required accuracy is achieved.

Bibliography

- [1] S. Haykin, B. W. Currie, E. O. Lewis, and K. A. Nickerson, "Surface-based radar imaging of sea ice," *Proceedings of the IEEE*, vol. 73, no. 2, pp. 233–251, 1985.
- [2] E. O. Lewis, B. W. Currie, and S. Haykin, "Effect of polarization on the marine radar detection of icebergs," in *IEEE International Radar Conference*, (Washington, D.C.), pp. 253–258, 1985.
- [3] E. O. Lewis and B. W. Currie, "Arctic ship trials of a dual-polarized ice navigation radar," in *IEEE International Radar Conference*, (London, England), pp. 404–408, 1987.
- [4] E. O. Lewis, B. W. Currie, and S. Haykin, *Detection and Classification of Ice*. Research Studies Press, 1987.
- [5] B. Larsson, A. Nelander, and G. Stenstrom, "Icerad a high resolution pulse doppler radar for detection of small targets in sea clutter," in *International Radar Conference, RADAR87*, (London, England), 1987.
- [6] K. D. Ward, "Compound representation of high resolution sea clutter," *Electronics Letters*, vol. 17, no. 16, pp. 561–563, 1981.
- [7] B. W. Currie, S. Haykin, and C. Krasnor, "Time-varying spectra for dual-polarized radar returns from targest in an ocean environment," in *IEEE Conference Proceedings RADAR90*, (Washington, D.C.), 1987.

- [8] S. Haykin, C. Krasnor, T. J. Nohara, B. W. Currie, and D. Hamburger, "A coherent dual-polarized radar for studying the ocean environment," *IEEE Transactions on Geosciences and Remote Sensing*, vol. 29, pp. 189–191, January 1991.
- [9] S. Haykin, *Adaptive Filter Theory*. Prentice-Hall, 1986.
- [10] C. Krasnor, S. Haykin, B. W. Currie, and T. J. Nohara, "A coherent dual-polarized radar for ice surveillance studies," in *International Conference on Radar*, (Paris, France), pp. 438–443, 1989.
- [11] C. Krasnor, S. Haykin, T. J. Nohara, and B. W. Currie, "A new coherent marine radar for ocean-related studies," in *Canadian Conference on Electrical and Computer Engineering*, (Vancouver, Canada), pp. 711–714, 1988.
- [12] D. Hamburger, C. Krasnor, S. Haykin, B. Currie, and T. J. Nohara, "A computer-integrated radar system for research use," in *Canadian Conference on Electrical and Computer Engineering*, (Montreal, Canada), pp. 594–596, 1989.
- [13] T. J. Nohara, S. Haykin, B. W. Currie, and C. Krasnor, "Towards the improved detection of small ice targets in k-distributed sea clutter," in *International Symposium on Noise and Clutter Rejection*, (Kyoto, Japan), pp. 66–71, 1989.
- [14] T. J. Nohara and S. Haykin, "Canadian east coast radar trials and the k-distribution," *IEE Proceedings, Part F, Special Issue on Clutter*, February 1991.
- [15] T. J. Nohara and S. Haykin, "The wavelet transform: A tool for time/frequency signal processing with applications in radar," in *Canadian Conference on Electrical and Computer Engineering*, (Ottawa, Canada), pp. 77.4.1–77.4.5, 1990.
- [16] D. O. North, "An analysis of the factors which determine signal/noise discrimination in pulsed-carrier systems," *Proceedings of the IEEE*, pp. 1014–1027, 1963.

- [17] G. L. Turin, "An introduction to matched filters," *IRE Transactions on Information Theory*, pp. 310–329, 1960.
- [18] D. Middleton, "On new classes of matched filters and generalizations of the matched filter concept," *IRE Transactions on Information Theory*, pp. 349–360, 1960.
- [19] P. Bello, "Joint estimation of delay, doppler, and doppler rate," *IRE Transactions on Information Theory*, pp. 330–341, 1960.
- [20] M. I. Skolnik, *Introduction to Radar Systems, Second Edition*. McGraw-Hill, 1980.
- [21] E. C. Westerfield, R. H. Prager, and J. L. Stewart, "Processing gains against reverberation (clutter) using matched filters," *IRE Transactions on Information Theory*, pp. 342–348, 1960.
- [22] P. M. Woodward, *Probability and Information Theory with Applications to Radar*. McGraw-Hill, 1953.
- [23] H. L. Van Trees, *Detection, Estimation, and Modulation Theory (Part I)*. John Wiley and Sons, 1968.
- [24] H. Cramér, *Mathematical Methods of Statistics*. Princeton University Press, 1946.
- [25] C. R. Rao, "Information and accuracy attainable in the estimation of statistical parameters," *Bull. Calcutta Math. Soc.*, vol. 37, pp. 81–91, 1945.
- [26] R. A. Fisher, "On the mathematical foundations of theoretical statistics," *Phil. Trans. Roy. Soc.*, vol. 222, 1922.
- [27] D. Dugué, "Application des propriétés de la limite au sens du calcul des probabilités à l'étude des diverses questions d'estimation," *Ecol. Poly.*, vol. 3, no. 4, pp. 305–372, 1937.

- [28] T. J. Nohara, "Collection and analysis of noncoherent radar returns from the sea," *DSS Contract No. 07SE.FP921-8-0164*, 1989.
- [29] A. Farina and F. A. Studer, "A review of cfar detection techniques in radar systems," *Microwave Journal*, p. 115, 1986.
- [30] D. C. Schleher, "Radar detection in log-normal clutter examined statistically," *Electronic Engineering*, pp. 53-58, 1976.
- [31] G. Green, "Radar detection probability with log-detectors," *IEEE Transactions on Information Theory*, vol. IT-4, 1958.
- [32] H. M. Finn, "A cfar design for a window spanning two clutter fields," *IEEE Transactions on Aerospace and Electronic Systems*, vol. AES-22, no. 2, pp. 155-169, 1986.
- [33] J. F. Walker, "Performance data for a double-threshold detection radar," *IEEE Transactions on Aerospace and Electronic Systems*, vol. AES-7, no. 1, pp. 142-146, 1971.
- [34] H. Rohling, "Radar cfar thresholding in clutter and multiple target situations," *IEEE Transactions on Aerospace and Electronic Systems*, vol. AES-19, no. 4, pp. 608-621, 1983.
- [35] S. Tantaranta, "Two-stage cfar detectors for multiple-range-bin radars," *IEEE Transactions on Aerospace and Electronic Systems*, vol. AES-19, no. 5, pp. 688-703, 1983.
- [36] R. Lefferts, "Adaptive false alarm regulation in double threshold radar detection," *IEEE Transactions on Aerospace and Electronic Systems*, vol. AES-17, no. 5, pp. 666-675, 1981.
- [37] R. Nitzberg, "Clutter map cfar analysis," *IEEE Transactions on Aerospace and Electronic Systems*, vol. AES-22, no. 4, pp. 419-421, 1986.
- [38] T. Bucciarelli, F. Caltagirone, and G. Picardi, "Cfar problems in a non-gaussian clutter environment," *Electro-Technology*, pp. 36-41, 1984.

- [39] M. Sekine, T. Musha, Y. Tomita, T. Hagsawa, T. Irabu, and E. Kiuchi, "Suppression of weibull-distributed weather clutter," in *IEEE International Radar Conference*, pp. 294–298, 1980.
- [40] M. Sekine, S. Ohtani, T. Musha, T. Irabu, E. K. T. Hagsawa, and Y. Tomita, "Suppression of ground and weather clutter," *IEE Proceedings Part F*, vol. 128, no. 3, pp. 175–178, 1981.
- [41] M. Sekine, T. Musha, Y. Tomita, E. Kiuchi, and T. Irabu, "Suppression of clutter and detection of targets," in *International Symposium on Noise and Clutter Rejection*, (Tokyo, Japan), pp. 279–284, 1984.
- [42] R. A. Johnson and E. G. Baxa Jr., "On characterizing the performance of modified cfar radar processors in the presence of non-homogeneous interference," in *International Symposium on Noise and Clutter Rejection*, (Tokyo, Japan), pp. 285–290, 1984.
- [43] Y. Yamada, H. Yanagisawa, and Y. Ohhasshi, "Sea clutter analysis for adaptive cfar design," in *International Symposium on Noise and Clutter Rejection*, (Tokyo, Japan), pp. 117–122, 1984.
- [44] L. G. Cole and P. wei Chen, "Constant false alarm rate detector for a pulse radar in a maritime environment," in *IEEE National Aerospace Electronics Conference*, pp. 1110–1113, 1978.
- [45] E. Jakeman and P. N. Pusey, "A model for non-rayleigh sea echo," *IEEE Transactions on Antennas and Propagation*, vol. 24, no. 6, pp. 806–814, 1976.
- [46] S. Watts, "Radar detection prediction in k-distributed sea clutter and thermal noise," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 23, no. 1, pp. 40–45, 1987.

- [47] S. Watts, "Radar detection prediction in sea clutter using the compound k-distribution model," *IEE Proceedings Part F.*, vol. 132, no. 7, pp. 613–620, 1985.
- [48] S. Watts and K. D. Ward, "Spatial correlation in k-distributed sea clutter," *IEE Proceedings Part F.*, vol. 134, no. 6, pp. 526–532, 1987.
- [49] K. D. Ward, "A radar sea clutter model and its application to performance assessment," in *IEE Conference Proceedings RADAR82*, (London, England), pp. 203–207, 1982.
- [50] E. Jakeman, "On the statistics of k-distributed noise," *J. Phys. A: Math. Gen.*, vol. 13, pp. 31–48, 1980.
- [51] C. J. Baker, K. D. Ward, and S. Watts, "The significance and scope of the compound k-distribution model for sea clutter," in *IEE Conference Proceedings RADAR87*, (London, England), pp. 207–211, 1987.
- [52] C. J. Oliver, "Representation of radar sea clutter," *IEE Proceedings Part F.*, vol. 135, no. 6, pp. 497–500, 1988.
- [53] R. Koenig, H. K. Dunn, and L. Y. Lacy, "The sound spectrograph," *Journal of Acoustic Society of America*, vol. 18, pp. 19–49, 1946.
- [54] S. K. Chakrabarti, *Hydrodynamics of Offshore Structures*. Springer-Verlag, 1987.
- [55] B. LeMéhauté, *An Introduction to Hydrodynamics and Water Waves*. Springer-Verlag, 1976.
- [56] J. H. Lever, E. Reimer, and D. Diemand, "A model study of the wave-induced motion of small icebergs and bergy bits," *CCORE Report*, pp. 282–290, 1987.
- [57] W. J. Pierson and L. A. Moskowitz, "A proposed spectral form for fully developed wind seas based on the similarity theory of s. a. kitaigorodskii," *Journal of Geophysical Research*, vol. 69, no. 24, pp. 5181–5203, 1964.

- [58] R. J. Keeler and R. E. Passarelli, "Signal processing for atmospheric radars," Technical Report NCAR/TN-331+STR, National Center for Atmospheric Research NCAR/TN-331+STR, Boulder, Colorado, May 1989.
- [59] W. D. Rummler, "Introduction of a new estimator for velocity spectral parameters," Technical Memo MM-68-4121-5, Bell Labs, Whippany, N.J., 1968.
- [60] W. D. Rummler, "Accuracy of spectral width estimators using pulse pair waveforms," Technical Memo MM-68-4121-14, Bell Labs, Whippany, N.J., 1968.
- [61] W. D. Rummler, "Two-pulse spectral measurements," Technical Memo MM-68-4121-15, Bell Labs, Whippany, N.J., 1968.
- [62] R. E. Passarelli and A. D. Siggia, "The autocorrelation function and doppler spectral moments: Geometric and asymptotic interpretations," *Journal of Climate and Applied Meteorology*, vol. 22, pp. 1776–1787, 1983.
- [63] K. S. Miller and M. M. Rochwarger, "A covariance approach to spectral moment estimation," *IEEE Transactions on Information Theory*, vol. IT-18, no. 5, pp. 588–596, 1972.
- [64] K. S. Miller and M. M. Rochwarger, "Estimation of spectral moments of time series," *Biometrika*, vol. 57, no. 3, pp. 513–517, 1970.
- [65] D. S. Zrnić, "Estimation of spectral moments for weather echoes," *IEEE Transactions on Geoscience Electronics*, vol. GE-17, no. 4, pp. 113–128, 1979.
- [66] P. R. Mahapatra and D. S. Zrnić, "Practical algorithms for mean velocity estimation in pulse doppler weather radars using a small number of samples," *IEEE Transactions on Geoscience and Remote Sensing*, vol. GE-21, no. 4, pp. 491–501, 1983.
- [67] W. Stehwien, *Radar Clutter Classification*. PhD thesis, McMaster University, Hamilton, Ontario, November 1989.

- [68] D. B. Trizna, "A model for doppler peak spectral shift for low grazing angle sea scatter," *IEEE Journal of Oceanic Engineering*, vol. OE-10, no. 4, pp. 368-375, 1985.
- [69] W. H. Press, B. P. Flannery, S. A. Teukolsky, and W. T. Vetterling, *Numerical Recipes in C*. Cambridge University Press, 1988.
- [70] T. M. Inc., "Matlab software package."
- [71] F. E. Churchill, G. W. Ogar, and B. J. Thompson, "The correction of i and q errors in a coherent processor," *IEEE Transactions on Aerospace and Electronic Systems*, vol. AES-17, no. 1, pp. 131-136, 1981.
- [72] J. E. Dennis and R. B. Schnabel, *Numerical Methods for Unconstrained Optimization and Nonlinear Equations*. Prentice-Hall, 1983.
- [73] J. A. Cadzow, "Signal processing via least squares error modeling," *IEEE ASSP Signal Processing Magazine*, pp. 12-31, October 1990.