

Theory and Practice of Radar Target Identification

For a complete listing of the *Artech House Radar Library*,
turn to the back of this book.

Theory and Practice of Radar Target Identification

August W. Rihaczek
Stephen J. Hershkowitz



Artech House
Boston *and* London
www.artechhouse.com

Library of Congress Cataloging-in-Publication Data

Rihaczek, August W.

Theory and practice of radar target identification / August W. Rihaczek, Stephen J. Hershkowitz.

p. cm. — (Artech House radar library)

Includes bibliographical references and index.

ISBN 1-58053-081-8 (alk. paper)

1. Radar targets-Identification. 2. Imaging systems.

I. Hershkowitz, Stephen J. II. Title. III. Series.

TK6580.R5323 2000

621.38486~21

00-059391

CIP

British Library Cataloguing in Publication Data

Rihaczek, August W.

Theory and practice of radar target identification. — (Artech House radar library)

1. Radar targets — Identification 2. Imaging systems

I. Title II. Hershkowitz, Stephen J.

621.3'848

ISBN 1-58053-081-8

Cover design by Igor Valdman .

© 2000 **ARTECH HOUSE, INC.**

685 Canton Street

Norwood, MA 02062

All rights reserved. Printed and bound in the United States of America. No part of this book may be reproduced or utilized in any form or by any means, electronic or mechanical, including photocopying, recording, or by any information storage and retrieval system, without permission in writing from the publisher.

All terms mentioned in this book that are known to be trademarks or service marks have been appropriately capitalized. Artech House cannot attest to the accuracy of this information. Use of a term in this book should not be regarded as affecting the validity of any trademark or service mark.

International Standard Book Number: 1-58053-081-8

Library of Congress Catalog Card Number: 00-059391

1 0 9 8 7 6 5 4 3 2 1

Contents

	Preface	<i>xv</i>
	Introduction	<i>xvii</i>
1	Target Identification, Measurement Requirements, and Algorithms	1
1.1	Background	1
1.1.1	Two-Dimensional Target Imaging and Complex-Image Analysis	1
1.1.2	Intensity Images Versus Complex Images	4
1.1.3	Backscattering Behavior of Man-Made Targets	6
1.1.4	Measurement Methods of Complex-Image Analysis Technology	10
1.1.5	Section Summary	28
1.2	Identification Principles	30
1.2.1	Derivation of a Practical Approach to Target Identification	30

1.2.2	Wavelength and Resolution in Identification Performance	36
1.2.3	Section Summary	40
1.3	Measurement of Range and Crossrange: The Two-Scatterer Algorithm	40
1.3.1	Resolution Principles	41
1.3.2	Implementation of the TSA for Ideal Point Targets	53
1.3.3	Application to Real Data	632
1.3.4	Section Summary	76
1.4	Special Measurements of Potential Use	77
1.4.1	Feature Extent	77
1.4.2	Cross Section of Features	80
1.4.3	Polarization Diversity	80
1.4.4	Dispersive Backscattering	82
1.5	Identification Procedure	82
1.5.1	Recognizable Target Features	83
1.5.2	Deformable Template Match	84
1.5.3	The Identification Process	86
	References	86
2	Topics Related to Target Identification	89
2.1	Ambiguity Function, Resolution, and Superresolution	89
2.1.1	Use of the Ambiguity Function	90
2.1.2	One-Dimensional Resolution	91
2.1.3	Superresolution	96
2.1.4	Two-Dimensional Resolution	98
2.1.5	Resolution of Weak Scatterers from Strong Scatterers	99
2.1.6	Effects of Weighting for Sidelobe Suppression	100
2.1.7	Resolution and Shifting Scatterers	102
2.1.8	Section Summary	103

2.2	Asymmetry of Range and Crossrange Resolution	104
2.2.1	Resolution Requirements for Target Identification	105
2.2.2	Measurement of Feature Positions	105
2.2.3	Fully Utilizing the Available Range Resolution	107
2.2.4	The Real Situation and Requirements on Crossrange Resolution	110
2.2.5	Illustration with Real Data	111
2.2.6	Checking the Limit on Crossrange Resolution	117
2.2.7	Crossrange Resolution for Targets at Broadside Aspects	121
2.2.8	Section Summary	121
2.3	Imaging Moving Targets	122
2.3.1	Imaging Principles	122
2.3.2	Motion Compensation	125
2.3.3	Motion Determination	156
2.3.4	Section Summary	158
2.4	The Need For Adaptive Processing Methods	159
	References	161
3	Aircraft Identification	163
3.1	Significance of the Radar Waveform	163
3.1.1	Section Summary	
3.2	Identification via Range Profiles	168
3.2.1	Requirements on Range Profiles for Identification	169
3.2.2	Range Profiles Without Duct Returns	170
3.2.3	Delayed Duct Returns	175
3.2.4	Range Profiles Containing Skin Returns and Delayed Duct Returns	177
3.2.5	Conclusions About Range Profile Utility	179
3.2.6	Section Summary	181

3.3	Identification via ISAR Imagery	181
3.4	Aircraft Features for Identification	182
3.4.1	Special Aircraft Features	182
3.4.2	Features Common to All Aircraft	186
3.4.3	Deriving the Database	191
3.4.4	Section Summary	198
3.5	Resolution Requirements	199
3.6	Imaging of Aircraft	203
3.6.1	Aircraft Imaging Principles	204
3.6.2	Illustrations of Aircraft Imaging Without Delayed Duct Returns	209
3.6.3	Illustrations of Aircraft Imaging With Delayed Duct Returns	223
3.6.4	Imaging of Maneuvering Aircraft	229
3.6.5	Maneuvering Combined With Vibrations	239
3.6.6	Imaging at Large Aspect Angles	245
3.6.7	The Special Case of Zero Aspect Angle	266
3.6.8	The Special Case of a 180° Aspect Angle	273
3.6.9	Imaging and Identification of Large Commercial Aircraft	277
3.6.10	Section Summary	282
3.7	Example of Positional Match for a Correct and Incorrect Aircraft	283
3.8	Procedure for Automated Aircraft Identification	285
3.8.1	Step 1: Data Collection	285
3.8.2	Step 2: Compensation and Selection of the Imaging Interval	288
3.8.3	Step 3: Image Analysis	289
3.8.4	Step 4: Identification	290
	References	290

4	Ground Vehicle Identification	291
4.1	Variability of Conditions for Identification	291
4.1.1	Application 1: Radar Platform and Ground Vehicles Stationary	292
4.1.2	Application 2: Radar Platform Stationary, Ground Vehicles Moving	292
4.1.3	Application 3: Radar Platform Moving, Ground Vehicles Stationary	293
4.1.4	Application 4: Radar Platform Moving, Ground Vehicles Moving	294
4.1.5	Section Summary	296
4.2	Basics of Ground Vehicle Identification	296
4.2.1	Radar Features Versus Vehicle Design Features	297
4.2.2	Recognizable Features	298
4.2.3	Positions of Unrecognizable Features	300
4.2.4	Section Summary	300
4.3	Identification of Stationary Ground Vehicles	300
4.3.1	Example 1: Measurements and Characterization for a Howitzer	301
4.3.2	Example 2: Measurements and Characterization for a Tank	314
4.3.3	Example 3: Measurements and Characterization for an Off-Highway Truck	320
4.3.4	Example 4: Measurements and Characterization for a Flatbed Truck	326
4.3.5	Assisting the Positional Match	331
4.3.6	Section Summary	336
4.4	Identification of Moving Ground Vehicles	336
4.4.1	Peculiarities of Moving Ground Vehicle Identification	337
4.4.2	Consequences of Different Types of Motion	338
4.4.3	General Processing Procedures	340
4.4.4	A Moving Off-Highway Truck	347

4.4.5	The Tank Moving in a Circle on Terrain	377
4.4.6	The Flatbed Truck on a Bumpy Straight Road	384
4.4.7	Recreational Vehicle Turning on a Paved Surface	389
4.4.8	Procedure for Analyzing the Survey Image and Forming a Final Image	409
4.4.9	More Sophisticated Motion Compensations	417
4.4.10	Procedure for Automated Identification of Moving Ground Vehicles	428
4.4.11	Section Summary	431
	References	432
5	Identification of Ships	433
5.1	Basics of Ship Identification	433
5.1.1	Peculiarities of Ship Identification	433
5.1.2	Imaging of Ships	434
5.1.3	Image-Time Selection	447
5.1.4	Principles of Analyzing a Ship's Motion	452
5.1.5	Motion Analysis for Real Data	467
5.1.6	Section Summary	473
5.2	Illustrations of the Measurement of Yaw and Roll Motions	474
5.2.1	Measurement Principles for Real Data	474
5.2.2	Significant Roll	477
5.2.3	Very Little Roll Motion	504
5.2.4	Turn Maneuver	510
5.2.5	Dive Boat at the End of a Turn	517
5.2.6	Small Ship in Rough Seas	523
5.2.7	Section Summary	543
5.3	Analysis of Ship Images	545
5.3.1	Measurement of Ship Length and Width	546
5.3.2	Scatterer Positions on Deck and Shape of Deck	569

5.3.3	Height of Scatterers Above Deck (Shape of Superstructure)	570
5.3.4	Section Summary	575
5.4	Processing Steps for Ship Identification/Classification	576
5.4.1	Step 1: Analysis of the Ship's Motion	576
5.4.2	Step 2: Selection of Imaging Time and Duration	577
5.4.3	Step 3: Motion Compensation	578
5.4.4	Step 4: Identification/Classification	579
6	<u>Analyzing Missiles, Rockets, and Satellites</u>	581
6.1	Overview	581
6.1.1	Section Summary	585
6.2	Basic Approach to Enhanced Sliding-Window Doppler Processing	585
6.2.1	Section Summary	589
6.3	Analysis of an Attitude Maneuver	590
6.3.1	Section Summary	610
6.4	Analysis of the Smooth Flight	611
6.4.1	Section Summary	618
6.5	Signal Sections and Response Groups	619
6.5.1	Section Summary	625
6.6	Chapter Summary	625
	Appendix A:	
	<u>One-Dimensional Two-Scatterer Algorithm</u>	627
	Appendix B:	
	<u>Transform Window Limits</u>	641
	Reference	645

Appendix C:		
Determining Scatterer Separations in Range Gates		
With Residual Uncompensated Motion		647
<hr/>		
Appendix D:		
Errors in Estimating Times of Two-Dimensional		
Motion		659
<hr/>		
Reference		663
Appendix E:		
Extracting Scatterer Locations From Range Tracks		665
<hr/>		
Reference		670
Appendix F:		
Modifications to the TSA for Interactive		
Two-Dimensional Analysis		671
<hr/>		
References		689
Appendix G:		
Tracking Interfering Scatterers by Measuring		
Phase Slopes		691
<hr/>		
G.1	Introduction	691
G.2	Choosing Intervals for Phase-Slope Measurement	692
G.3	Data Analysis	695
G.4	Summary	703
Appendix H:		
Setting an Amplitude Threshold for Analyzing		
Image Responses		705
<hr/>		
Reference		711

About the Authors	713
--------------------------	------------

Index	715
--------------	------------

Preface

In *Radar Resolution and Complex-Image Analysis*, published by Artech House in 1996, we maintained that one cannot extract enough information from radar intensity images of man-made targets to solve problems that require high resolution, such as target detection in a poor environment and target identification. To solve these problems in an operational sense (using a fully automated processor with reliable performance in a large database), it is necessary to analyze the complex image, utilizing the phase of the image in addition to the intensity. Another important point we made is that man-made targets cannot be mathematically modeled with sufficient realism that one can base processing algorithms on mathematical target models. We treated the fundamentals of the technology of extracting information from complex images, or complex processor outputs in general, by what may be called an expert system approach.

The technology of complex-image analysis is so radically different from conventional processing methods that its practical application is not readily apparent once one is familiar with the basics. For this reason we decided to publish a second book focusing on the application of the new technology to target identification. The new technology is difficult to apply because it replaces mathematically-based processing methods with an expert system approach that uses pattern recognition or interpretation. One needs a background in mathematics in order to understand conventional radar signal processing methods. In contrast, with complex-image analysis, one must

perform functions such as the interpretation of amplitude and phase patterns. Applying complex-image analysis requires insight, which we hope to provide in this book, but also extensive interactive software. This software is available in *CRISP: Complex Radar Image and Signal Processing — Software and User's Manual*, published by Artech House, 2000.

The reader may wonder why none of the practical applications treated in this book, and in fact very little on complex-image analysis technology in general, has been published in scientific and technical journals. We attempted to do so, but encountered strong resistance. In response to one reviewer's concern that our approach does not reflect the consensus of the radar community, we reply that technical journals should not only publish material on which a consensus already exists. As another example, when we claimed that mathematical modeling of a target does not approximate a real target sufficiently well to serve as a basis for processing algorithms, we were told that because real data are too diverse, one must use mathematical modeling. Such comments made it clear that we would have to seek out a different vehicle to publish our findings on the new technology.

The development of a technology that depends on extensive interactive computer software requires capital. While we had a variety of sponsors, the technology could not have been sufficiently developed were it not for Mr. James K. Hall and Dr. William J. Miceli, both with the Office of Naval Research. The former allowed us to develop the technology to a useful state; the latter is responsible for the breadth of applications. Without them there would be no technology of complex-image analysis. We also must acknowledge the indirect benefits we derived from working for White Sands Missile Range, on projects managed by Mr. Elwin C. Nunn. Lastly, there is little one can do in this field without capable interactive software, and that was developed by Dr. Richard L. Mitchell and Mr. Robert H. Mitchell of MARK Resources.

The processing algorithms discussed in this book do not lend themselves to a conventional presentation; more figures and fewer equations are required. We thank Dr. David K. Barton, radar editor of Artech House, for working through the details of the manuscript and providing innumerable and valuable suggestions for improving the presentation and organization. We also thank Ms. Bea C. Felix for effectively handling the secretarial demands associated with preparing this book.

Introduction

The decades after World War II saw the development of radar hardware so capable that (at least for stationary targets) it is now possible to generate high-quality two-dimensional images in which the range/crossrange resolution cells are much smaller than the targets. The most obvious application of this capability is target identification, but there are other applications in which the resolution of individual features on a target is extremely helpful or even necessary. However, despite the availability of excellent hardware, the development of operational systems that fully utilize this hardware has not met with much success. The main reason for this is that signal processing has been based on the conventional resolution theory. Since this is a theory for point targets, it no longer applies to man-made targets that are highly resolved.

There is a second reason for the failure of conventional signal-processing methods with respect to systems development. It has been customary to consider microwave radar as a kind of extension of an (noncoherent) optical system, even though radar is a coherent system. From the beginnings of radar resolution theory, resolution itself was wrongly defined on the basis of the intensity outputs from the processor, the criterion for resolution being that two targets generate two recognizable responses. Then it is natural to approach target identification by generating an intensity image of a target, and attempting to extract information from the positions of the response peaks. This can work satisfactorily only if the target can be

modeled by a set of point scatterers in fixed positions on the target, and if resolution is so high that every major scatterer generates a well-resolved peak. Neither requirement is met in practice.

The difference between the backscattering behavior of man-made targets at optical wavelengths and at radar wavelengths is so drastic that optics cannot provide guidance for radar, if the full potential of radar is to be realized. At optical wavelengths, the surfaces of a man-made target are rough scatterers, and hence every part of the target backscatters with more or less the same intensity. With sufficiently high angular resolution, an optical image thus shows the overall shape of the target and the positions and shapes of its individual features. Importantly, at optical wavelengths a target rarely generates highlights and spurious responses.

At radar wavelengths, a man-made target backscatters coherently. Different parts of such a target backscatter with grossly different intensities. On one hand, we have a variety of discontinuities that backscatter weakly because they are small, and smooth surfaces that backscatter weakly because by far most of the incident wave is reflected in directions away from the radar. On the other hand, we have regular and irregular corners and cavities that tend to trap the incident wave and backscatter with much higher intensities. They behave much like trihedral corner reflectors, though with less efficiency. The positions and characteristics of these wave-trapping features are a necessary ingredient for target identification, whereas the lower-level background is useless. It is much like a speckle pattern that is too crude to define feature shapes. Unless we design radar and signal-processing methods to utilize these dominant (in the image) features, we cannot reliably identify in a large database.

To make matters worse, the wave-trapping features have loosely defined positions because, by their nature, they are extended. For radar purposes, their positions are defined by their phase centers, the points from which the backscattered signal appears to emanate. Since these features are not shaped like ideal trihedrals, why should their phase centers remain at the same positions over the aspect angle sector and the bandwidth used for imaging? In fact, the phase centers do not remain fixed; the question is merely how much they move over the aspect angle utilized for crossrange resolution and the frequency band used for range resolution. When the effect is present but small, a response might be slightly widened and shifted in its position. In worse cases, we find an entire set of spurious responses instead of a single response, with the spurious responses close to the position of the scatterer. In the worst cases, the spurious responses spread much beyond the position of the scatterer. Radar wavelength and resolution must be chosen so that these

effects are small enough. In addition to the spurious responses associated with the wave-trapping features, there are more readily understood spurious responses, such as the multiple delayed returns from aircraft engines and engine ducts, or the Doppler-spread responses from rotating devices.

Conventional signal processing, derived from a definition of resolution based on the intensity output, discards the phase of the processor output and thereby sacrifices half the resolution potential. This is unacceptable. High range resolution is expensive, and high crossrange resolution introduces a variety of problems, from dwell-time requirements to difficulties with the motion compensation. In practice, for a given bandwidth and dwell time, we cannot afford to sacrifice resolution performance. Resolution was originally defined, explicitly or implicitly, as the half-power width of the point-target response (ambiguity function) of the system, but this definition does not do justice to the fundamental resolution capability of radar. Two responses may overlap to such a degree that they generate only a single peak, and yet we may recognize the presence of two scatterers and measure their positions. The resolution potential of radar can be realized only by analyzing complex responses, intensity and phase. If this is done on the basis of a mathematical target model, we arrive at the superresolution techniques, which primarily recover the factor of two lost in resolution when the phase of the processor output is discarded. If it is done without such a model, which must be the case if the method is to work on real targets, we arrive at the methods of complex-image analysis.

The imaging process itself requires rethinking. As an example, consider imaging of an aircraft. The radar collects a set of range profiles. Before these range profiles can be processed into an image, we must compensate for the aircraft motion. We use some form of range compensation, which is crude because it is based on the signal bandwidth, followed by some Doppler compensation. The latter improves the motion compensation because it is based on the wavelength. The aircraft may move smoothly, or it may execute perhaps inadvertent yaw motions, or there may be a stronger motion disturbance due to atmospheric turbulence or an intentional maneuver. If the motion happens to be smooth over the imaging interval, we will obtain an image in which the scatterer responses are well compressed (“focused”), so that scatterer positions and other information can be extracted. If the motion is not perfectly smooth, the compensation may be inadequate for this purpose. Scatterers may drift through range cells, entirely preventing meaningful Doppler measurements. Smaller drifts within a range cell introduce nonlinear phase functions, which shift and distort the responses. This is not admissible if the image is to be used to identify the aircraft in a large database.

The point is that we cannot apply a specific motion compensation procedure and blindly assume that it has worked satisfactorily. We must check how well it has worked, and if necessary correct it. This must be done automatically in an operational system. The processor also must determine whether a satisfactory motion compensation can even be performed over the chosen imaging interval, and perhaps select a subinterval within which the aircraft motion is more benign. All of this means that there cannot be a fixed motion compensation algorithm, or a fixed identification procedure. The processor must determine what the target is doing, choose the approach accordingly, and check whether or not each processing step is successful, before deciding on the next step. In other words, the processing procedure that eventually leads to the identification of the target must be adaptive.

When one considers the poor backscattering behavior of man-made targets at radar frequencies, the need to utilize the complex image rather than only the intensity image, and the absolute necessity for adaptive processing, it is clear why conventional approaches to target identification have not worked and cannot work. This understanding led us to the development of the complex-image analysis technology, which amounts to an expert system approach. It is not as elegant and straightforward a process as writing mathematical equations and programming a computer to solve them, but it works for real targets rather than mathematical target models. One can write mathematical equations only if one has a mathematical model of the target. As we have found from the analysis of innumerable images of a large variety of targets, mathematical modeling with sufficient realism is not possible, at least not when such a difficult problem as target identification is to be solved in an operational sense. Hence, mathematical equations cannot serve as the basis for the development of algorithms that are to work on man-made targets with the complexity of ground vehicles, aircraft, and ships.

The most severe impact on processing comes from the unworkability of mathematical modeling of targets. This implies that one must develop algorithms empirically, using real data from representative targets, which amounts to correctly selecting responses for analysis and interpreting amplitude and phase functions of the responses. After developing the algorithms to a point where satisfactory performance is obtained, they must be automated in a separate step. This is an expert system approach. It is not easy work, but radar problems involving highly resolved man-made targets are not easy.

We should point out that the complex-image analysis technology does not completely forgo mathematical modeling and (simple) equations. In order to implement resolution in accordance with the basic limits, we window image responses so that they are generated by either a single scatterer or

by two scatterers. These responses are modeled by two interfering point scatterers. However, aside from the simplicity of this target model, before any measurement based on it is accepted, we check the validity of the model and assign a measurement uncertainty that depends on how well the real situation approximates the model. This is very different from modeling an entire target as a set of fixed point scatterers (or using a more complicated model), writing mathematical equations based on this model, and simply accepting the results of the processor. It is the difference between a method that works in practice and one that does not.

With this book, we explain in considerable detail how the methods of complex-image analysis can be used to solve the problems that have resisted practical solution for so long. We chose the important task of target identification for this demonstration, considering the difficult problems of identifying aircraft, ground vehicles, and ships, with an indication of how these methods can be applied to missiles, rockets, and satellites. Our objective is to concentrate on the measurements that provide the inputs for target identification. We are brief on the actual identification process, the comparison of the measured data with the database, because an adequate literature already exists on the topic. We also do not discuss the automation of the measurement algorithms in great detail, because there is a wider latitude in the methods of automation than in the measurement algorithms; the automation of the algorithms is more straightforward than their development. In essence, one must automate the analyses of responses and their transforms as illustrated throughout the book. We note that the measurements discussed herein have in large part been automated, with the automated software providing a performance that approximates that of an analyst sufficiently well.

1

Target Identification, Measurement Requirements, and Algorithms

In this introductory chapter, we provide the background for two-dimensional target imaging and identification, including a summary of the most essential material from [1]. The summary is intended as an overview; important topics of this introductory discussion will be treated later in more detail.

1.1 Background

1.1.1 Two-Dimensional Target Imaging and Complex-Image Analysis

An imaging radar transmits a sequence of high-bandwidth signals at a pulse-repetition frequency that must be somewhat higher than the Doppler spread of the target. Ideally, the transmitted signals are constant-carrier pulses with a duration corresponding to the desired range resolution. Because of peak power limitations of radar transmitters, the signals are usually pulse-compression waveforms such as linear FM pulses, pulses with linear frequency stepping, or phase shift codes. Upon reception, the transmitted waveform is compressed into a spike with a duration of about the reciprocal of the bandwidth, with some energy remaining outside the compressed response in the form of range sidelobes. For imaging purposes these range sidelobes

should be as low as possible, at least as low as -30 dB. Since the chosen waveform is primarily a matter of convenient hardware implementation, in this discussion we will assume that the radar transmits a sequence of short constant-carrier pulses. For each transmitted pulse, the receiver obtains responses from the illuminated scatterers in the range gates spanning the target. This set of generally interfering responses is the so-called range profile of the target. Although the range profile is a complex function, only the intensity range profile has been used with conventional processing.

The short pulse provides the desired range resolution, but we also need Doppler or crossrange resolution. Doppler resolution is obtained by coherently integrating the sequence of pulses. Simultaneous range and Doppler resolution, or imaging, is possible only if two (not independent) requirements are met. First, each scatterer on the target must remain in the same range cell for the entire integration time. Second, the Doppler of each scatterer must be constant over the integration time. In order to meet these requirements, the received range profiles are first stored and delayed relative to one another so that the responses from a specific scatterer all appear to have arrived with the same delay. This is referred to as the alignment of the range profiles. Next, one measures the Doppler variations of the scatterers over the coherent processing time and removes them from the data. If the target spans many range cells and crossrange resolution is to be high, the target rotation (about one scatterer fixed in a range gate) needed for crossrange resolution may cause other scatterers to drift through their range gates. Removal of this drift requires a separate processing step.

The imaging process for an arbitrary transmitted waveform is shown in the flowchart in Figure 1.1. The pulse compression of the received data generates the short pulses needed for range resolution. The motion compensation of the target, consisting of range alignment followed by Doppler compensation, allows Doppler resolution to be obtained by coherent integration. However, in practice it is not possible to compensate for arbitrary target motion. As we demonstrate in detail below, *one must determine the quality of compensation possible for a given data interval and, if that quality is insufficient for the measurements to be performed on an image, choose a reduced imaging interval over which the quality is adequate*. This is a key point for generating good images: rather than proceeding with a poor motion compensation to achieve some specified nominal crossrange resolution, we must reduce the imaging interval and nominal crossrange resolution. In order to determine whether the imaging interval must be reduced, one must attempt to track and compensate individual scatterers spanning the target. If such individual scatterer tracking and compensating yields scatterers that drift by only a small

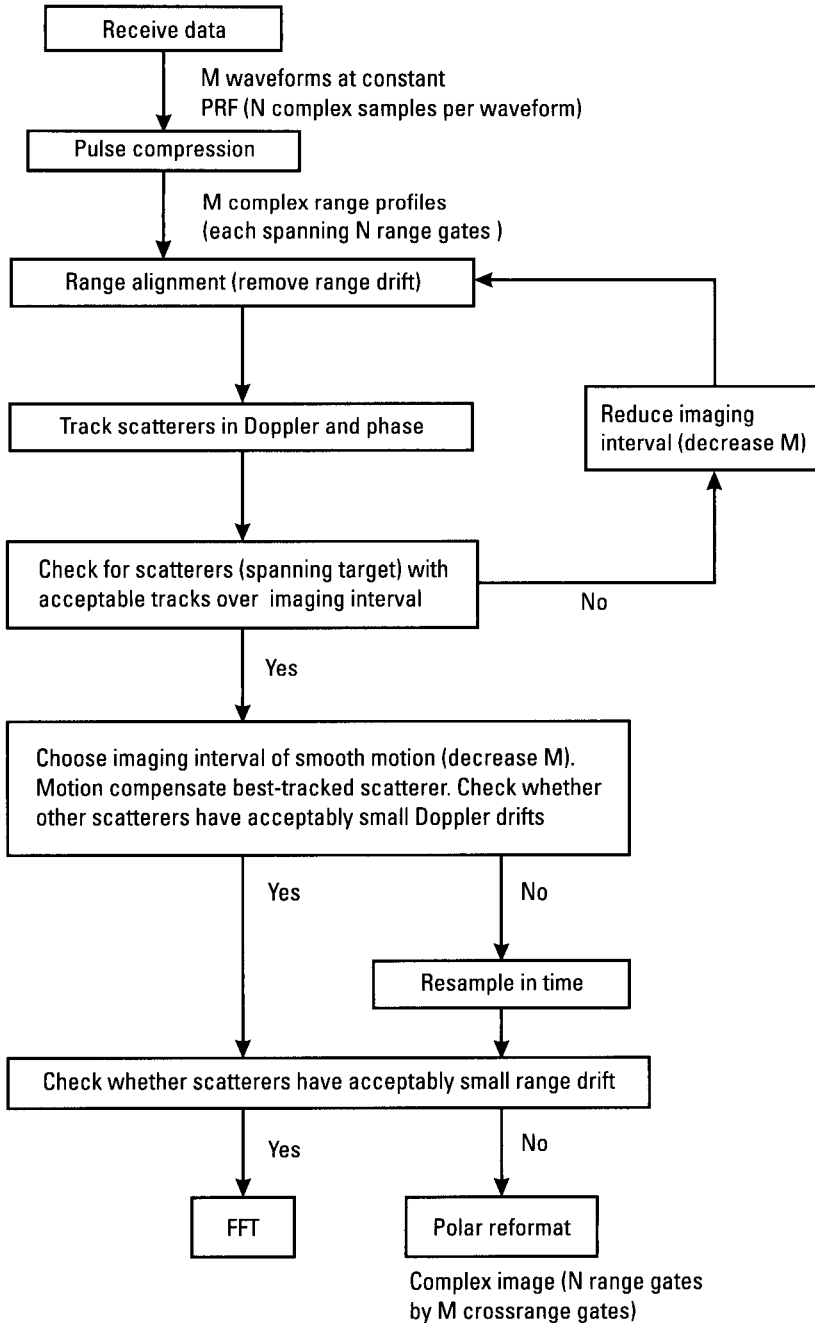


Figure 1.1 Flowchart for image generation.

fraction of a range gate over the imaging interval, the tracking is of sufficient quality to proceed with the motion compensation of the target. Otherwise, the imaging interval must be reduced, and the new interval must be evaluated in the same manner.

Having determined that individual scatterer tracks have acceptable quality, one may still need to reduce the imaging interval. Imaging over any sharp change in target motion, which will be evident from the scatterer tracks, must be avoided. Sharp changes produce pseudoperiodic errors in the process of fitting to the measured motion, and the consequent compensation errors can generate unacceptably high Doppler or crossrange sidelobes. Imaging must be restricted to an interval of smooth motion. With this restriction, the best-tracked scatterer is used to compensate the entire target. After compensation, the target effectively rotates about the compensated scatterer, at a (not necessarily uniform) rate that can be derived from the tracks of the other scatterers. If the rotation angle changes unevenly from pulse to pulse, the other scatterers will have nonconstant Dopplers and their image responses will be defocused. This can be corrected if one resamples the data (in time) so that the rotation angle changes uniformly; that is, by a constant amount from pulse to pulse. If the total rotation angle is small enough so that no scatterer drifts through a significant fraction of a range gate, the complex image can be generated by a *fast Fourier transform* (FFT) of the compensated data. Otherwise, polar reformatting [2] is required.

1.1.2 Intensity Images Versus Complex Images

When the *constant-false-alarm-rate* (CFAR) algorithm for the detection of a point target in white Gaussian noise was developed, it was found that the optimum detection procedure consists of discarding the phase by envelope detecting the receiver output, and then using a detection threshold on the intensity output. The phase of the receiver output thus was determined to be useless for this purpose. Even in applications other than target detection, it became an established practice to discard the phase of the receiver output and utilize only the intensity output. This method was carried over from the detection of a point target to the resolution of two point targets. Resolution was formulated in terms of the so-called ambiguity function [3], which is the point-target response of the transmitted radar signal. However, as with the detection of a point target, resolution was formulated in terms of the envelope of the point-target response, discarding the phase.

With the utilization of just the intensity output, it was logical to define resolution as the closest separation of two equally strong scatterers that yields

two response peaks rather than only a single peak. This minimum separation depends on the phase between the two scatterers, which would make the definition of resolution phase-dependent. Nevertheless, in an average sense we observe two response peaks from two equally strong scatterers when the minimum separation is about equal to the half-power width of the response. Hence, the resolution cell came to be defined as the half-power width of the point-target response of the system. Since a signal with a bandwidth B has a point-target response with a half-power width of about $1/B$, it was concluded that resolution in range delay is about $1/B$. Similarly, since a signal with a duration T has a point-target response with a half-power width of $1/T$, it was stated that Doppler resolution is about $1/T$. This definition of resolution implies that any processing which increases the half-power width of the point-target response degrades resolution performance in proportion to the widening. Since the necessary weighting for sidelobe suppression widens the point-target response, weighting was said to degrade resolution.

The above standard definition of resolution would be meaningful if the task of the radar were to count the number of targets, or scatterers on a target, by counting the number of response peaks. In most practical situations, however, one also must measure the range and Doppler of the target or of each scatterer. A definition of resolution that does not include this requirement is practically meaningless. Now, when we ask how far two point targets must be separated in order that one can measure their ranges or Dopplers with reasonable accuracy (a fraction of the width of the point-target response), the minimum allowable separation turns out to be about $2/B$ and $2/T$ [1, 4]. This conclusion is in conflict with the general statement that radar resolution in range delay is $1/B$ and resolution in Doppler is $1/T$.

Various workers in the field have recognized the problem that the standard definition of resolution does not take into account the need for measuring range and Doppler of the (resolved) targets. An investigation of this point leads to the conclusion that the difficulty lies with the common procedure of envelope detecting the receiver output [4]. The phase of the receiver output may be useless for the detection of a point target, but it carries essential information for the resolution of multiple targets. Use of the complex receiver output instead of the intensity output permits resolving scatterers separated by $1/B$ or $1/T$, in the sense that sufficiently accurate measurements of range or Doppler can be performed. Thus, even when the need for accurate range or Doppler measurements is taken into account, *the inherent resolution performance of radar is indeed $1/B$ and $1/T$, but in order to realize this resolution performance we must utilize the complex processor output*, not just the intensity output. If a target is resolved, one can measure its range and Doppler to a

small fraction of the resolution cell sizes in the two coordinates, but this requires use of the correct processing procedure.

When the intensity image is used instead of the complex image, resolution is degraded by a factor of about two from the inherent resolution performance of radar, both in range and crossrange. As we will discuss in more detail below, such a degradation is unacceptable in practice. Higher range resolution in order to offset this degradation is costly to implement, and may lead to spurious responses when the range resolution cell is smaller than the range extent of the scatterer. Higher crossrange resolution to offset the degradation may not be achievable on moving targets. The unacceptable degradation of resolution with intensity images is one of the reasons for the development of the complex-image analysis technology. The second reason relates to the complicated backscattering behavior of man-made targets, which does not allow us to utilize the complex receiver output via mathematically-based target models.

1.1.3 Backscattering Behavior of Man-Made Targets

If we want to formulate a viable approach to automated reliable identification of man-made targets, we must first understand their backscattering behavior. The foundation of the conventional signal-processing methods is the conventional resolution theory, which is a resolution theory for point targets. This has important implications.

If the point-target theory is to apply to imaging and identification of man-made targets such as aircraft, ground vehicles, and ships, it must be possible to model such targets as sets of point scatterers. This is not the only requirement. Even in the simplest case of a rigid target, the motion of the target is utilized to determine the crossrange positions of the scatterers. In order for this to be valid, the positions of the scatterers must not shift when the aspect angle changes during the imaging process. Similarly, the positions of the scatterers must not vary with frequency, at least not over the bandwidth of the signal. Hence, the conventional resolution theory is appropriate for radar imaging of man-made targets only if they can be modeled by sets of point scatterers in fixed positions on the targets. One need not study the problem in great detail in order to conclude that these conditions are not met by man-made targets. If they were, then the backscattering would be much like that at optical wavelengths. The intensity image would be a version of the optical image, except with the coordinates of range and crossrange instead of azimuth and elevation angle. Such an image would give an

excellent reproduction of the shape of the target. The fact is that radar images do not look like that; hence the model must be wrong.

More complicated target models have utilized so-called primitives, such as flat plates, trihedrals, dihedrals, spheres, cylinders, and so forth. Again, images produced with such target models look vastly better to the eye than the images of real targets, and allow easy identification, in contrast to real images. Again, the conclusion must be that such a model is still inappropriate. One can further increase the complexity of a target model by attempting to use electromagnetic (EM) theory to describe the backscattering of the features on real targets. However, it is questionable whether the complicated features on real targets, other than perhaps for simple targets such as missiles, can be mathematically modeled to the required fidelity. The avenue we have pursued is to use the methods of complex-image analysis technology to examine a large variety of target images for different aspect angles, with the targets stationary or moving, having a smooth motion or different degrees of erratic motion, and so forth. These extended investigations have led to an insight into the backscattering behavior of real targets, with the following conclusions.

Man-made targets of the complexity of aircraft, ground vehicles, and ships have two types of scatterer. First, there are a large number of discontinuities that act as fixed point scatterers. They produce a quasi-optical image, but at a low level. Second, there are the scatterers with shapes that tend to trap the incident wave, such as regular and irregular corners or cavities. At the short wavelengths needed for imaging, these wave-trapping features generate returns much stronger than those from the point scatterers: so strong, in fact, that they obscure the underlying quasi-optical image from the weak point scatterers. *The returns from the wave-trapping features are the only target returns that can be reliably observed, so that identification must necessarily be based on them.* Our experience shows that the number of these scatterers is fairly small. Depending on the type of target and its aspect angle, we generally observe on the order of 20 such scatterers, perhaps up to about 30 in some cases (large ships may have many more). This is discussed further in Appendix H.

The fact that the observable scatterers are extended features, rather than point scatterers, has undesirable consequences. Cavity-like features and irregular corners have complicated backscattering properties. Radar imaging is based on the assumption that the observable scatterers are in fixed positions on the target, so that their motions are determined by the target as a whole. However, what is the position of an extended feature? For radar purposes, we can define a phase center from which the return appears to emanate. If this phase

center did not shift with aspect angle (over the aspect angle sector used for crossrange resolution) or with frequency (over the bandwidth of the signal), the image would be almost quasi-optical in that single responses would appear at the true range/crossrange positions of each scatterer. The number of responses would be very much smaller than in a photograph, so that the shape of the target still would not be as well defined. However, the phase center of an extended wave-trapping feature may shift with aspect angle and sometimes also with frequency. This phase-center shifting must necessarily have an effect on imaging.

Conventionally, imaging has been performed with a correlation receiver, which compares the returned signal with the transmitted signal, and for point targets (or fixed point scatterers on a target) such a receiver is optimum. Realizing that the assumptions that led to the choice of a correlation processor for imaging are not met by complicated man-made targets, we investigated the possibility of replacing the correlation processor with another type of processor that might be more immune to the consequences of phase-center shifts of scatterers; but we concluded that for practical purposes the correlation processor is optimum. The question then becomes what consequences phase-center shifts have on the images generated by correlation processors, and how to deal with the consequences.

The answer to the first question depends on the size of these phase-center shifts relative to the wavelength and on their form: linear with time or frequency, nonlinear but monotonic, or pseudoperiodic. The specific characteristics in turn depend on the size and shape of the wave-trapping feature, with size measured relative to the wavelength, and on the motion behavior of the target. For a given shape, the larger the feature's size relative to the wavelength and the more erratic the target motion, the worse the consequences of phase-center shifts. They range from a slight widening of the response, or a slight displacement relative to the position of the scatterer, to very large displacements and a high degree of response smearing, giving the appearance of poorly resolved multiple responses. In extreme cases, these spurious "responses" may be distributed over the entire width or even length of the target. Such responses can severely distort the image, and they can mask genuine responses. (Note that these spurious responses cannot be modeled by ray-tracing methods.) In principle, such responses can be identified on the basis of their peculiar amplitude and phase patterns [1, 5], but only if they are sufficiently well resolved from other responses. Even if they can be identified, the masking problem remains. In addition to spurious responses generated by phase-center shifts of the scatterers, there are more obvious sources of spurious responses. The most problematic are the multiple delayed returns

from jet engine ducts. The spurious responses are the second major reason why quasi-optical approaches to target identification based on shape recognition generally cannot work.

How can one deal with the spurious responses in an image? With respect to such responses caused by phase-center shifts, which we designate as "sideband responses," the best approach is to avoid radar wavelengths so short that the spurious responses become a real problem. Although our experience with very high frequencies is limited, the indication is that one should avoid going above X-band if possible. If the sideband responses remain a problem at a particular carrier frequency, the problem may be severe only if the target motion is erratic or three-dimensional. However, such a target motion presents other problems that will be discussed in later chapters. One can deal with delayed duct returns by resolving them from the returns of interest, just as one can do for Doppler-spread responses from rotating features. Other delayed returns can be largely ignored by not utilizing the far (in range) parts of an image. We note again that the observable features are relatively strong, so that only strong spurious responses must be accommodated. In essence, the procedure of handling the spurious responses must be adapted to the specific application.

When a target is moving, there is yet another type of spurious response that can be even more important than the types discussed above. The motion compensation needed to form a scatterer response that is sharp in crossrange, without high crossrange sidelobes, must be performed to high accuracy. This becomes problematic when the signal-to-interference ratio for the scatterer is so low that its motion cannot be accurately measured. The problem is particularly difficult, or impossible to solve, when the motion of the scatterer is erratic. The motion compensation residuals then can generate a large number of spurious responses, so that the image quality is degraded to the point where accurate measurements of scatterer positions are impossible.

To deal with this problem, the only choice is to shorten the imaging interval, even though this means poor crossrange resolution. However, we must keep in mind the difference between resolution and the accuracy with which scatterer positions can be measured. If a scatterer is resolved in range, the crossrange position of its response can be measured to a very small fraction of the crossrange resolution cell. Hence, even when crossrange resolution is poor, position accuracy in crossrange can still be adequate. This indicates that high range resolution is much more important than high crossrange resolution. It also indicates problems when a target is viewed near broadside; if the target length is much larger than its width, range resolution

is ineffective in subdividing the target into parts, so that the burden of resolution lies with crossrange resolution. However, high crossrange resolution may not be obtainable if we cannot accurately measure an erratic motion. Also, the measurement of the scatterer motion as needed for motion compensation requires that at least some scatterers are resolved in range, and this may not be possible in a broadside geometry.

Having summarized the backscattering properties of man-made targets, we mention another advantage of utilizing the complex rather than the real image. Consider a “normal” wave-trapping feature that does not generate spurious responses, such as a good approximation to the ideal trihedral corner reflector. If it has only a small crossrange extent and is properly motion compensated, its return signal will be at a constant Doppler. Since Doppler is the derivative with respect to time of the signal phase, the implication is that the phase function will be linear. In contrast, if the crossrange extent of the feature is significant as far as the radar is concerned, its Doppler will change with aspect angle, and hence the phase function will be curved. By measuring the curvature of the phase function of the response we thus can determine the effective crossrange width of the feature. Analogously, by measuring the curvature of the response phase in the spectral domain, we can determine the range depth of the feature [1].

1.1.4 Measurement Methods of Complex-Image Analysis Technology

Earlier we discussed the fact that use of the intensity image degrades resolution in both domains by a factor of two, which is unacceptable in practice. We must therefore utilize the complex image. How should this be done? The most basic measurement is that of the range/crossrange position of a scatterer. If we use the intensity image, this measurement is performed by determining the range/crossrange positions of the response peaks. Since this procedure does not realize the inherent resolution performance of radar, we must somehow extract the scatterer positions from the complex image. The straightforward procedure would be to model the target as a set of fixed point scatterers, derive the appropriate mathematical equations, and program a processor to solve them. This is in fact the approach of superresolution, whereby the biggest part of the gain in resolution is the recovery of the factor of two lost when the phase is not utilized. However, we also concluded that such a target model is unrealistic in practice. How, then, do we achieve a resolution of $1/B$ and $1/T$ for real targets?

To answer this question, we start by considering the one-dimensional situation of two ideal point scatterers as employed in the conventional

definition of resolution. The composite signal returned from the two scatterers can be derived from the vector sum of one vector rotating about another, as will be discussed in detail later. For now we merely note that the composite return signal has an amplitude and a phase function, both modulated with a period given by the rotation period of the vector. The transform of the complex signal represented by amplitude and phase functions gives the composite response. In principle, we have the choice of analyzing this complex response or analyzing the composite return signal directly; the same information is contained in both. For two ideal point scatterers in noise, it makes no difference whether one analyzes the return signal or its transform, but this is not the real environment in which one must resolve scatterers. The scatterers are not ideal point scatterers, and in general they must be resolved in the presence of interference from other scatterers. Moreover, the changes induced by one scatterer in the amplitude and phase response of another are much more subtle than the changes in the amplitude and phase function of the signal. We have found it more practical to analyze the complex signal.

In performing such an analysis, which means determining the positions and strengths of the two scatterers, we start with the image response and take a Fourier transform to generate the corresponding composite signal. This signal is then analyzed in accordance with principles discussed below in detail. We call this basic resolution algorithm, which merely implements the inherent resolution of radar, the *two-scatterer algorithm* (TSA). Since we consider the practical rather than the ideal case, we effectively measure how well the amplitude and phase functions of the signal approximate the ideal patterns obtained from two point scatterers. The processing thus is more akin to optical pattern recognition than to conventional radar signal-processing methods. We must process in this way because the inappropriateness of a mathematical model does not permit us to assume that a good pattern is obtained. In the practical situation of interference from neighboring scatterers, we must choose the transform window so that only two scatterers are involved (checking the quality of the amplitude and phase patterns to ensure that this is the case) before we can accept the measurement results.

The preceding discussion considers one-dimensional resolution, be it in Doppler or range. In most practical applications, however, we are interested in two-dimensional resolution. How does one extend the one-dimensional TSA to two dimensions? The only practical way we have found for real targets is to take a series of image cuts through a response, at angles over 360° the image plane, then analyze the transform of each image cut in the one-dimensional manner, and combine the one-dimensional measurements into two-dimensional positions.

The one-dimensional TSA is the cornerstone of complex-image analysis, and is essential for all radar applications involving high resolution of man-made targets. However, before proceeding to the illustrations, we must make an important point concerning interpolation. With modern radar, the received signal is digitized and recorded as time samples, in each of the range gates of the entire range window. To simplify processing and recording, the minimum sampling rate is used. In accordance with the Nyquist criterion, the minimum sampling rate equals the Doppler spread of the signal (some margin might be added for practical reasons). As long as this requirement of "critical sampling" is met, the information content in the signal is preserved and conventional processing algorithms will not have any problems.

With complex-image analysis, we examine the shapes of complex (amplitude and phase) time and frequency functions. Even though it is true that digital samples of such functions retain all the information available in their analog forms, pattern-type measurements cannot be performed accurately on signals and spectra with critical sampling. For example, we often wish to derive a smooth curve from a sequence of amplitude points. The only practical way of connecting these points into a smooth curve is to draw straight lines between the points. However, when the critical sampling rate is used, the resulting curve is usually far from smooth. To generate a smooth curve with a sequence of linear segments, we must artificially increase the sample rate by interpolation. In order to obtain a smooth curve, we place additional samples between those received at the critical sampling rate.

With digital signal processing, the most convenient way to interpolate is to use the so-called zero fill. Assume a sequence of time samples at the critical sampling frequency. Let us take the FFT of this function in order to generate the Doppler spectrum. If we take the inverse transform of the spectrum, we again obtain the original time function. However, if we widen the spectral interval by some factor M , adding points with zero amplitude (zero fill), and then take the inverse transform over the widened window, the number of samples of the time function will have been increased by a factor of M . Connecting the sequence of samples by straight lines will produce a smooth function if M is large enough. Similarly, if we have a response function at the critical sampling rate, we can take the transform, zero fill, and transform back to get a smoothed version of the original function. With complex-image analysis we routinely zero fill by a factor of four or eight.

To illustrate the process, in Figure 1.2 we show a signal at the critical sampling frequency, amplitude, and phase. The actual measurement points of the amplitude and phase are connected by straight lines. The transform of this complex signal (with Hamming weighting applied for sidelobe

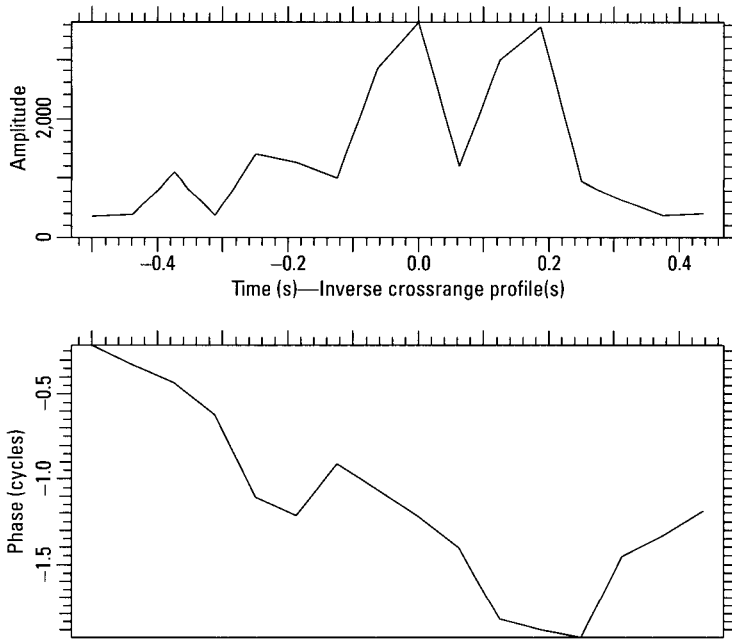


Figure 1.2 Signal at the critical sampling frequency.

suppression) is given in Figure 1.3, with the points obtained from the transform again connected by straight lines. Neither the signal shown in Figure 1.2 nor the responses in Figure 1.3 allow the required measurements. For example, the left of the two amplitude responses in Figure 1.3 gives the impression that two scatterers may be involved, but we cannot measure the two positions. In Figure 1.4 we have taken the signal of Figure 1.2, but have increased the total interval by a factor of eight through zero fill. The transform over the entire time interval (with Hamming weighting applied to the original signal) is given in Figure 1.5. The measurement points in Figure 1.5 are again connected by straight lines, but because the sample spacing was reduced to one-eighth of the original spacing, the resulting signal is smooth enough to perform the requisite measurements. Note that the left response indeed consists of two peaks. Similarly, if we want to replace the signal in Figure 1.2 by its smoothed version, we take the transform to obtain the response in Figure 1.3, widen the crossrange interval by adding points with zero amplitude, and transform back. If we again use zero fill by a factor of eight, the transform gives the signal shown in Figure 1.6. This signal is more than smooth enough for pattern measurements. The illustrations were given

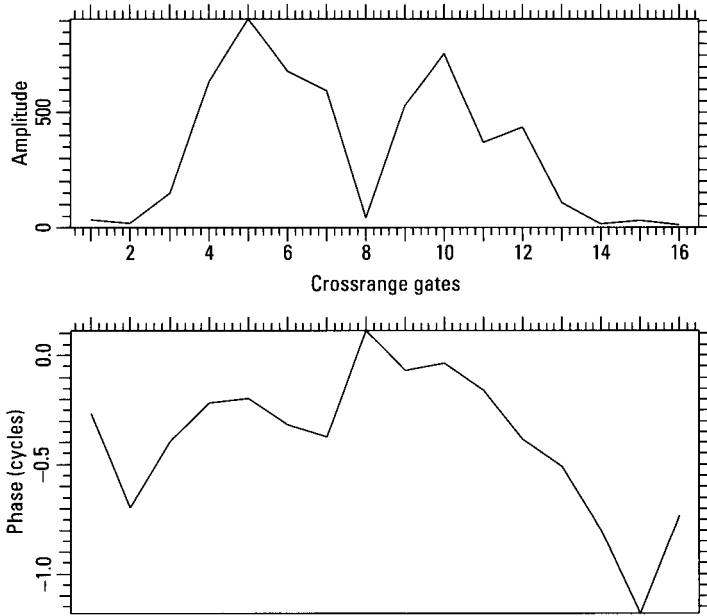


Figure 1.3 Transform of the signal in Figure 1.2.

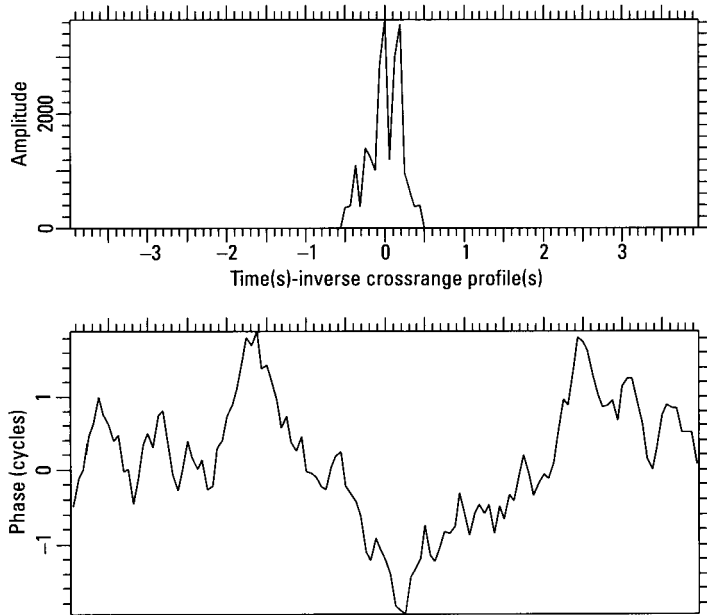


Figure 1.4 Lengthened signal interval for interpolation.

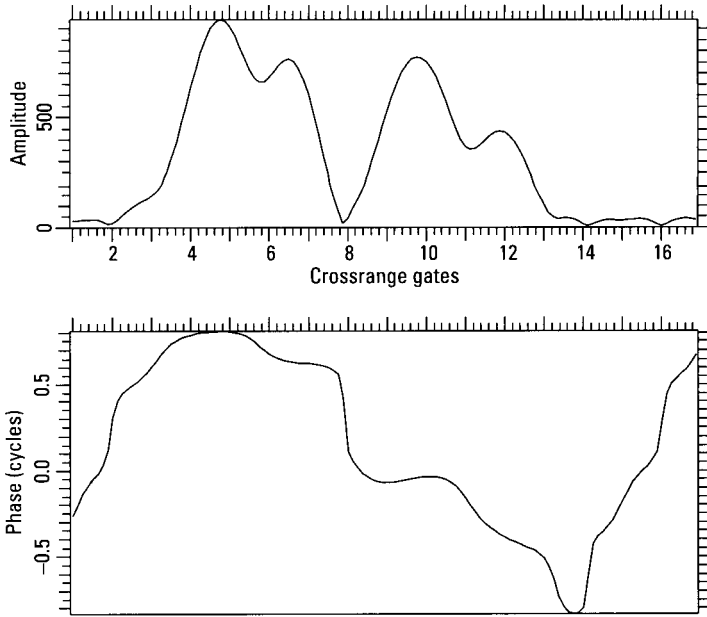


Figure 1.5 Transform over the widened time interval.

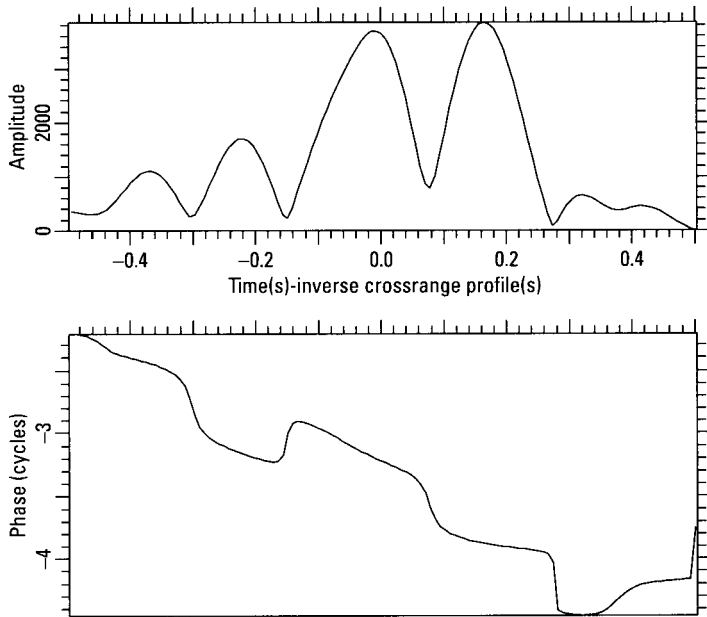


Figure 1.6 Smoothed signal.

for time signals and crossrange responses, but the smoothing method applies analogously to spectra and range responses.

The zero-fill procedure is mathematically equivalent to interpolating between the critically sampled data points, with sinc-function interpolation. The sinc-function implementation is more convenient for the analysis of two-dimensional data, which requires examining image cuts at arbitrary angles in the range-crossrange plane. These cuts may not contain any critically sampled data point. Smoothing is accomplished by using two-dimensional sinc-interpolation to generate data points along each image cut, spaced apart by one-quarter or one-eighth of a resolution cell. Having explained why the functions to be presented will all be smooth despite the fact that a practical radar always uses critical sampling, we now discuss the types of measurements made with complex-image analysis.

The following discussions are provided only for the purpose of familiarizing the reader with our way of presenting image cuts and their transforms, and the interpretation of such functions, as an introduction to complex-image analysis. Quantitative interpretations and the significance of the various measurements will be discussed in more detail later.

Figure 1.7 shows the intensity image of an aircraft. The aircraft is viewed at a relatively small aspect angle, in which case it is easy to perform a good motion compensation and generate an image of high quality. With complex-image analysis processing, we want to make measurements on the individual responses, so that a different presentation of the image is preferable (the presentation has additional advantages, as will become clear in our discussions). To obtain this modified presentation, the oversampled intensity image is scanned to find local maxima, even if they should rise only minimally above the surrounding intensity level. The intensity at each pixel is compared to the intensities at the nearest neighboring pixels in range, crossrange, and diagonally. If a pixel is stronger than all eight of its nearest neighbors, a dot is plotted at the range-crossrange coordinates of the response peak, with the area of the dot proportional to the amplitude of the peak. Such a "peaks plot" shows the individual intensity responses, regardless of whether they represent responses from resolved scatterers or the peaks of the interference pattern from unresolved scatterers.

The peaks plot corresponding to the intensity image of Figure 1.7 is shown in Figure 1.8. The fact that the dots are isolated does not imply that the associated responses are well resolved, since a local maximum is depicted even when it barely rises above the surrounding response level. Although the peaks plot gives only the intensity image, in the analysis of the responses we always utilize the image phase as well. Because we have not found a practical




Figure 1.7 Intensity image of an aircraft.

way of displaying the image phase, only the peaks of the intensity image are shown. The string of responses in about Crossrange Gate 0 are the delayed engine duct returns. The fuselage is oriented approximately along the line of large dots. Note that the number of peaks on the fuselage barely exceeds ten. The number of observable scatterers will be somewhat higher because a particular peak may be generated by more than one scatterer.

As discussed above, the analysis of the individual (complex) responses is performed via image cuts and their transforms. As a reference point, in Figure 1.9 we show the interpolated image response and signal for a fixed-range image cut through a simulated ideal fixed point scatterer. The curves on the left give the image response; the curves on the right give its Fourier transform. The upper curves show amplitude and the lower curves show phase. As this is a fixed-range cut, the image data display the crossrange variation of the response. The lower image abscissa shows crossrange position in the image. The upper abscissa will be explained below in conjunction with a general diagonal cut. For a fixed-range or fixed-crossrange cut, the upper abscissa is offset by a constant (here zero) from the lower. As crossrange is proportional to Doppler, its transform is proportional to time, as shown on the transform abscissa. (For real data, the transform of the response in a




Figure 1.8 Peaks plot for the image shown in Figure 1.7.




Figure 1.9 Fixed point scatterer.

fixed-range image cut is the time signal returned by the target in that range gate, after any pulse compression waveform was removed and the motion of the scatterer was compensated.)

The image amplitude is a symmetric peak, with a half-power width corresponding to the Hamming weighting applied to the simulated signal, in order to suppress (crossrange) sidelobes. The image phase is constant. The transform of the image data is displayed on the right of the figure. In this case, the transform was taken over the entire displayed image interval. In some figures, we use a shorter interval, which is indicated by vertical cross-hairs. The amplitude was normalized by dividing by the Hamming weighting applied to the original signal. The transform amplitude is nearly constant, with some distortion near the boundaries of the signal. This distortion corresponds to the sidelobes filtered out by our choice of image transform interval; an interval including the entire image cut would give a perfectly constant transform amplitude. The transform phase is linear, and we have normalized our presentations such that the phase slope gives the location of the image response, Crossrange Gate 0 in this case.

A note on weighting for sidelobe suppression is in order. We weight the received signal in order to form an image whose responses have low sidelobes. Then we select a response in the image, take the transform of an image cut through the response, and remove the weighting by dividing the amplitude of the transform by the weighting function. This process is unnecessary for the illustration in Figure 1.9, because only a single response is involved. Weighting is absolutely necessary in forming an image, because without weighting there would be unacceptably strong mutual interference among the various responses. In analyzing the transform of an image cut through a response, we deweight in order to generate more easily understandable amplitude patterns. Since low-level response sidelobes can rarely be included within the image transform interval, some distortions at the fringes of the transform are unavoidable.

We want to mention another side issue, which is the so-called phase unwrapping. With digitized data, we obtain values of amplitude and phase at each sample point, as derived from the analog data. However, the phase of the digitized samples has a 2π ambiguity. Since complex-image analysis requires measurements on amplitude and phase patterns, the ambiguity must be resolved. This is done by adding or subtracting full phase cycles where necessary to obtain a smooth phase function from the digital samples. This phase unwrapping process is not problematic for one-dimensional image cuts as long as the sampling rate is high enough; that is, when the data are sampled at least at the critical (Nyquist) rate.

Returning to the real data, in Figure 1.10 we show the crossrange interval over which we take an image cut in the range gate of a peak, using interactive software for our illustrations. The image cut over the interval marked in Figure 1.10 is shown on the left side of Figure 1.11. The dashed curve shows an overlaid response from a fixed point scatterer for comparison. Quantitatively, the image response has a normalized half-power width (relative to the half-power width for an ideal point scatterer) of 0.969 and a skewness (the ratio of the half-power width of the left half of the response and the half-power width of the right half) of 0.988. These indicate an almost ideal point scatterer, as is evident from the agreement between the dashed and solid curves. Since the response is strong, the scatterer must be a wave-trapping corner-like or cavity-like feature, or an antenna.

In viewing the transform of the image cut as in Figure 1.11, we must be aware of the fact that in our illustrations the transform shown is that of the entire image cut displayed in the left half of the figure, unless otherwise specified by crosshairs. In our example, the amplitude of the image cut (upper left curve) contains a minor response on the left that must not be included in the transform if we want to analyze the major response. Thus we must choose a transform window that includes as much as possible of the response of




Figure 1.10 Position of the image cut in the range gate of the peak.




Figure 1.11 Image cut in the range gate of the peak, and transform.

interest, and as little as possible of any interfering response. In easy cases this implies positioning the boundaries at the amplitude minima or the phase inflection points of the response (more difficult cases will be discussed later).

To illustrate, we choose the transform window as shown by the crosshairs in Figure 1.12. This is a transform of an image cut in the range gate of the signal, so that we obtain a time signal. As discussed earlier, it is the time signal returned by the scatterer under examination, again amplitude and phase function. If the scatterer were an ideal point scatterer, the amplitude would be constant. In this case the amplitude variation is insignificant (quite comparable to that of the point scatterer in Figure 1.9). If the motion compensation were perfect, the Doppler of the (fixed point) scatterer would be constant, and hence its phase function would be linear. This is not the case for the phase function at the lower right. However, the total variation of the phase is only 0.04 cycles, and this maximum variation occurs at the fringes, where such transforms are not very reliable. It will later be shown that phase variations below about 0.1 cycles have little consequence. Just as we conclude that the amplitude variation is acceptable, so is the phase variation. For practical purposes, the scatterer has an essentially ideal behavior in the crossrange domain.




Figure 1.12 Windowed response and its transform.

As a side remark, note that even with a perfect motion compensation, deviations from linear phase may also be caused by a shifting of the feature's phase center, which must be discriminated from motion compensation residuals by comparing the phase functions from different scatterers. If there is a motion compensation residual, scatterers with fixed phase centers will all have the same curvature of the phase function. Such a residual may be caused by inadvertently tracking and compensating a feature with a shifting phase center in the motion compensation process. The phase of the tracked scatterer will then be linear, while the phase curvature from the phase-center shifting will have been imparted to all other scatterers.

We now rotate the image cut of Figure 1.10 by 90° , so that it falls in the crossrange gate of the response peak. This image cut is shown in the left half of Figure 1.13, and its transform is shown in the right half. Since the image cut is taken in the crossrange gate of the response, the transform represents the frequency spectrum returned by the scatterer. The overlaid curve again shows a point scatterer response, for comparison. The normalized half-power width of the response is 1.272, as compared with a value of 0.969 for the image cut in the range gate of the response. The deviation from unity thus is 0.272 as compared with 0.031. Thus, in the range dimension (image




Figure 1.13 Image cut and transform in the crossrange gate of the response.

cut in the crossrange gate) the scatterer does not approach the ideal point scatterer as well as in the crossrange dimension. This is also seen from the fact that the phase function of the response, in the lower left plot, is curved rather than linear in the interval of strong amplitude. To exclude the interference, we choose a transform window as shown by the crosshairs in Figure 1.14. The periodic variation of the transform amplitude is indicative of two interfering scatterers, as will be discussed in detail below.

Figure 1.15 shows a fixed-range cut through the strongest response of the image, at Range and Crossrange Gates 0. The normalized half-power width of the response is 1.083, and the image phase function is curved rather than linear. The variation of the amplitude of the transform (upper right curve) is again acceptable for a single scatterer, since it does not show deep minima or periodic variation. On the other hand, the total phase variation of the transform (lower right curve) is about 0.3 cycles. This is an already significant deviation from the ideal point scatterer, whose phase function would be linear. The phase curvature means that the scatterer has a measurable range depth. Since the derivative of the spectral phase is delay, the total variation of the phase slope of the lower right curve gives the delay depth of the (extended) scatterer [1]. These operations can be performed for any type of




Figure 1.14 Transform window for Figure 1.13.




Figure 1.15 Fixed-range cut through strongest response of Figure 1.10.

image cut, not necessarily in a range gate or a crossrange gate. As an illustration, in Figure 1.16 we show the position of an image cut along the fuselage, with the actual image cut and its transform given in Figure 1.17. The response at the center of the displayed interval corresponds to the response marked by an x in Figure 1.16. The upper image abscissa measures distance along the image cut, in gates, from the point denoted by the x (the central point of the image cut). The lower image abscissa presents the projection of the cut on the range axis. In our illustrations, diagonal cuts will be projected on the axis, range or crossrange, for which the projection is longer. As the transform of a diagonal cut is neither signal nor spectrum, we label the transform abscissa with “Transform.”

The crosshair of Figure 1.17 also marks the response whose normalized half-power width is to be measured. The normalized half-power width is 2.1, which is about twice as large as for a resolved response from an ideal point scatterer, but the distortion is essentially all on the left side of the response, as shown by the overlaid dashed point-scatterer response. Is the marked response a composite from two scatterers, or is it merely deformed by interference from the response to the left? For a test, we position a transform window on the response as shown in Figure 1.18. As will be explained in detail




Figure 1.16 Image cut along the fuselage.

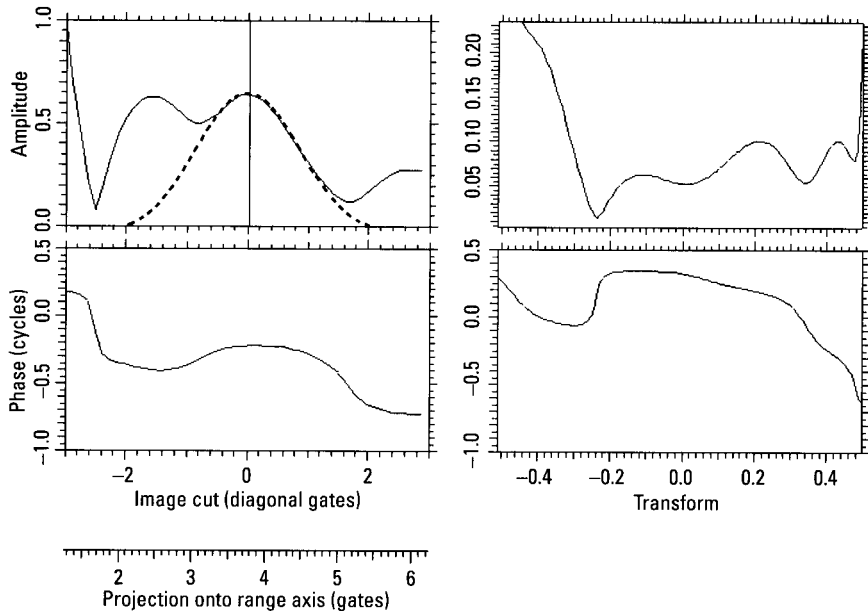


Figure 1.17 Image cut and transform, in accordance with Figure 1.16.

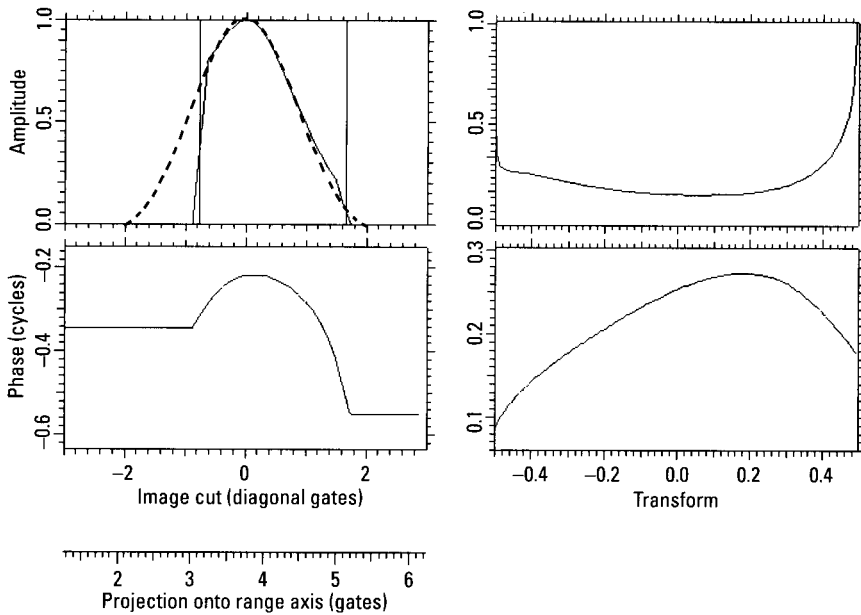


Figure 1.18 Windowed response and transform.

below, this is not the transform of two interfering scatterers of roughly comparable strengths, but the transform of a single scatterer with some low-level interference from one or more scatterers. Continuing the test, we place a transform window over both responses in Figure 1.17, as shown in Figure 1.19. Our extensive discussion of the TSA will show that the new transform is indeed a usable approximation of the interference pattern from two scatterers. Hence, the two responses essentially represent two scatterers, and we have to utilize the complex responses to obtain more accurate estimates of the positions of the two scatterers than given by the positions of the response peaks.

If we examine a response in the intensity image, whether this be a conventional intensity image or a peaks plot, we know little more than the strength of the response. If we use the procedures indicated above, we can tell whether the response is from one or more scatterers. If the response is from one scatterer, the phase function of its transform tells us whether the phase center of the scatterer is fixed or moving (and by how much), in the range domain, the crossrange domain, or any diagonal cut in the image plane. If a response is composed of two scatterers, we can determine their individual positions in range and crossrange and their strengths, and also whether they

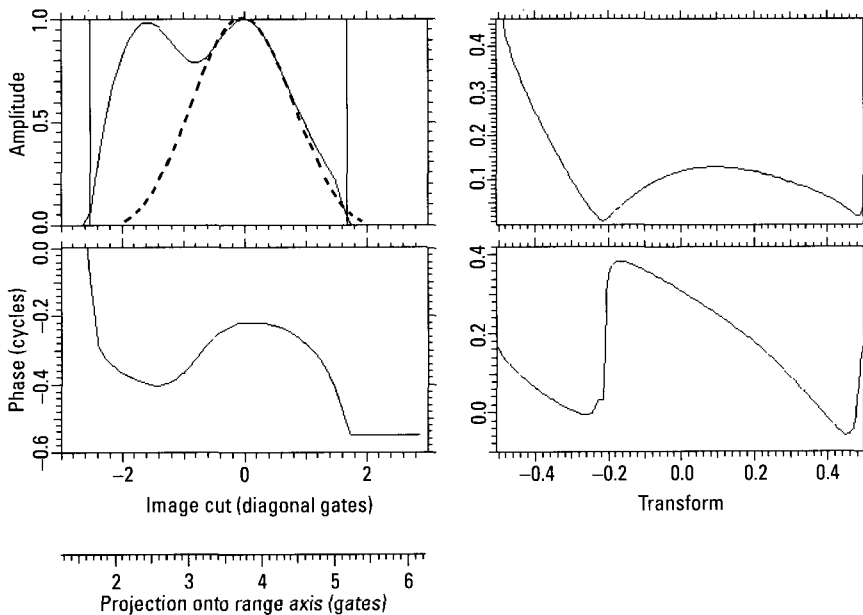


Figure 1.19 New window and transform.

have fixed or shifting phase centers for an arbitrary image cut. By comparing the transforms with those from an ideal point scatterer or the interference pattern from two point scatterers, we know how accurately we have measured the positions and properties of the scatterers. When a composite response represents more than two scatterers, the transform window must be chosen so as to include contributions from essentially only two scatterers, with minimal interference from other scatterers. This will be treated in Section 1.3. All these measurements can be performed fully automatically, as inputs to an automated identification system.

Having treated the important TSA, which implements the basic radar resolution in range and Doppler, we want to stress a point that may easily be overlooked. The TSA will give meaningful results only if the scatterers to be resolved are good approximations of fixed point scatterers and if the motion compensation has correctly compressed the two responses. When we apply the TSA to scatterers with shifting phase centers, or when the motion compensation is inadequate, the results may be meaningless. For example, the response from a single scatterer with a shifting phase center might be so distorted that application of the TSA gives a result as if two or even more scatterers were involved. This is particularly true if the TSA is applied to a set of spurious responses generated by a single scatterer, but can also happen when the response “distortions” are less extreme, whatever the cause of the distortions. This is to say that resolution works properly only if the conditions are satisfactory. This can and must be verified before the TSA is applied.

1.1.5 Section Summary

Certain issues of radar imaging and target identification govern the approach to target identification. Below we summarize the important points made above and their consequences, to set up the extended discussions that follow. This is to direct the reader’s attention to the most important issues.

1.1.5.1 Radar Backscattering of Complicated Man-Made Targets

High-quality radar images of man-made targets are dominated by a relatively small number of responses strong enough to be utilized for target identification. If an image has a large number of responses, most will be spurious.

1.1.5.2 Inputs to Target Identification

In addition to special target features (e.g., dimensions) and recognizable features (e.g., engine intakes), automated reliable target identification requires

an accurate measurement of the positions of nonrecognizable features. This is possible (to a fraction of the resolution cell) only if a scatterer can be sufficiently well resolved. In principle, it does not matter whether the scatterer is resolved in range, in crossrange, or in both range and crossrange.

1.1.5.3 Definition of Resolution

A practically meaningful definition of resolution must incorporate the need to perform reasonably accurate range and crossrange measurements. From this point of view, the basic resolution capability of a radar using a signal of bandwidth B and duration T is $1/B$ in range delay and $1/T$ in Doppler. However, realizing this resolution capability requires the use of the complex image. Use of only the intensity image degrades resolution by a factor of two in each dimension.

1.1.5.4 Requirements on Resolution

In practice, radar resolution must be high enough to ensure that most of the limited number of observable scatterers are separated by at least one resolution cell. Thus, a usable response in the intensity image must come from no more than two scatterers, or it must be possible to place a window on a response such that the window includes major contributions from at most two scatterers, with little interference from other scatterers.

1.1.5.5 Role of Range Resolution

The role of range resolution is primary. High range resolution is easy to implement and is nonproblematic. High resolution in crossrange requires dwell time and may not be achievable when the motion of a target is not smooth or when range resolution does not resolve some scatterers well enough to measure their motions for purposes of the motion compensation.

1.1.5.6 Target Identification at Large Aspect Angles

Aircraft fuselages, ships, and to some degree ground vehicles are slender, so that at large aspect angles range resolution becomes ineffective in subdividing the target into many cells. This places increased requirements on crossrange resolution; yet when range resolution is ineffective, the motion measurements needed to implement the motion compensation and high crossrange resolution may not be possible. Thus, targets near broadside may be impossible to identify reliably when the database is large.

1.2 Identification Principles

1.2.1 Derivation of a Practical Approach to Target Identification

Target identification at radar wavelengths must be approached very differently from target identification at optical wavelengths. The approach cannot be somehow to obtain an (intensity) image that allows one to recognize the shapes of the target and its features. In order to explain the drastically different approach needed at radar wavelengths, we first consider optical target identification. Because of the extremely short wavelengths in optics, the many small discontinuities on man-made targets all backscatter with similar intensities. Even the weak scatterers of the “smooth” surfaces backscatter, unless the smoothness approaches that of a mirror. The consequence is that usually all parts of a target and its features backscatter with comparable intensities, so that one can recognize the shapes of the target and its features. As long as resolution is adequate, a target can be readily recognized by inspection, even though automated identification is a difficult problem. Occasionally, a mirrorlike surface may generate a blinding flash that, if it persists over the observation interval, might prevent target identification; but usually the backscattering is not governed by a few dominant scatterers.

The situation is entirely different for radar. Radar wavelengths are so large that extended surfaces appear smooth, and most of the discontinuities backscatter very weakly. The contributions from the very large number of weak scatterers generate a pattern similar to the speckle pattern at optical wavelengths. In the absence of spurious responses, this target-generated “background” can well exceed the noise and clutter, and then would define the shape of the target in the radar coordinates of range and crossrange. A quasi-optical approach to target identification would likely work if all scatterers backscattered with similar intensities.

Our extended analyses of images of real targets have shown that complicated man-made targets contain a set of features with much stronger backscattering than the multitude of other scatterers. The strong features are those that tend to trap the radar signal, such as a variety of regular and irregular cavities and corners. Antennas also are features that backscatter relatively strongly, although we do not know whether this is due to their designed electrical properties or because they must be mounted in a way that isolates them from the surrounding metallic surface. These features are not such effective backscatterers as the ideal trihedral corner, yet their returns are strong enough to dominate over the backscattering from the large number of weak scatterers. Since the dominant features prevent us from utilizing the optical-

type background image generated by the multitude of weak scatterers, *we have no choice but to base target identification on the dominant features*. Occasionally the existing conditions might be so benign that a quasi-optical approach will work, but it will not satisfy the demands of an operational automated identification system.

The situation with real man-made targets is as follows. The typical man-made target has perhaps 20 or 30 features whose responses exceed the target-generated background sufficiently to make them usable for target identification. Although these dominant scatterers as a set are much stronger than the weak scatterers that form the quasi-optical image, their backscattering strengths vary, of course. Hence, when resolution is poor, perhaps only the 10 to 20 strongest responses may be observable. If the number of measurable responses drops further, resolution is inadequate for target identification. Thus, identification depends on a relatively small number of features, and resolution must be adequate to perform measurements on the responses from these features. If the backscattering of the features were simple, and if the number of such features were much larger than it actually is, the specific arrangement of the features might still describe the shape of the target, so that optically inspired identification methods would work. However, the backscattering of the dominant features at radar wavelengths can be rather complicated, invalidating such an approach.

As an example, consider the returns from the engine duct of a fighter aircraft, as shown in Crossrange Gates 0 to 1 of Figure 1.20. The duct intake typically is the strongest aircraft scatterer over a wide aspect angle sector, over which multiple reflections within the duct produce a string of relatively strong delayed returns. These usually extend beyond the range of any skin return (the skin returns in Figure 1.20 are located approximately within the dashed outline of the generic aircraft). As another example, any rotating device on the target generates a series of returns spread in Doppler. In Figure 1.20, the rotating compressor blades generate responses spread in crossrange and aliased by the PRF to appear both within and outside the crossrange interval of the scatterers. These responses appear at ranges greater than Range Gate 9. The spread returns in that gate are the direct returns from the blades. The spread returns at greater ranges involve both multiple reflection within the duct (as for the delayed duct returns) and reflection by the blades of other engine stages. As a less obvious fact, even more ordinary extended features can have a phase-center wander large enough to generate strong responses in false locations. The various spurious returns, in combination with the scarcity of observable features, prevent reliable shape recognition.




Figure 1.20 Aircraft image.

Although the quantitative manner in which features with shifting phase centers generate spurious responses is not fully understood at this time, the fact that these responses exist is well established [1, 5]. In its simplest form, the phase-center wander may cause a small widening of the response, leaving the response essentially in the position of the scatterer. A stronger phase-center shift may translate the response from the position of the feature, with or without a widening, depending on whether the phase-center shift is linear or nonlinear. For example, the crossrange position of the response from the engine intake of a jet fighter often differs from the actual position of the intake, which may be recognized from the crossrange position of the delayed duct responses. The nonlinear component in the phase-center wander widens the response. The “widening” may be so severe that instead of a single response, perhaps smeared, one observes an entire set of relatively strong responses that may be located quite far away from the scatterer position and that can be relatively widely spread in range and crossrange. These effects increase with carrier frequency, because the phase-center shifts become larger relative to the wavelength. Even worse, they increase as the motion of a target becomes more irregular, in particular when the target’s rotation axis changes

at the time of imaging. This complicated motion may come from inadvertent or intentional maneuvering of an aircraft, or the movement of a vehicle on a rough road or off the road, or the motion of a ship in rough seas. Such responses evidently interfere with shape recognition.

It is true that, where possible, we want to make use of target features that are at least crudely related to target shape. We would like to measure the length and wingspan of an aircraft (the latter only if we know that the aircraft is not banked), the length and width of a ground vehicle, and the length, width, and superstructure shape of a ship. However, in most instances this is not enough for reliable target identification in a large database. We often cannot measure the length and, in particular, the wingspan of an aircraft so accurately that we can reliably distinguish among many aircraft with similar dimensions. Figure 1.21 shows combined photographs of the American F-15 and Soviet Su-27 aircraft, to scale. Evidently, the sizes and shapes are similar. The F-15, on the left of the figure, has a length of 19.43m and a wingspan of 13.05m. The Su-27 has a length, including the extended fuselage stinger section, of 21.93m and a wingspan of 14.70m. Without the extended stinger, which will generally not be observable by a radar, the Su-27 length is



Figure 1.21 Combined photographs of F-15 and Su-27 aircraft, to scale.

20.18m, only 4% different from that of the F-15. These aircraft are too similar for their sizes and shapes to be used as the sole basis for discrimination, given realistic measurement errors and uncertainties in aspect and bank angles.

Similarly, we cannot rely on distinguishing similar ground vehicles on the basis of length and width measurements, because vehicles are designed to meet common road-width standards and sometimes also differ insignificantly in length. The measurement problem is much worse when the ground vehicles are moving. For a ship, the combination of length, width, and shape of the superstructure should allow classification, but not identification. In general, the measurement of shape-related target parameters will not suffice for reliable automated target identification in a large database.

For reliable identification, we must additionally measure the positions of the prominent, and hence resolvable, target features. This is particularly important for aircraft, because there are so many different aircraft and many of them are similar in size and shape. In addition, we also want to utilize any nonpositional feature properties, such as feature width and depth, that we can measure. Lastly, there are special features that can help in identification. For example, for aircraft we might be able to determine the number of engines, whether they are wing mounted or integrated into the fuselage, or similar special design characteristics. For ground vehicles, we may use the presence of a turret or a gun, or the number of wheels. The help in target identification that can be provided by overall dimensions and special features depends on the kind of target. For all targets, however, we can measure at least the range and crossrange positions of the dominant features. The need to extract these feature positions from an image becomes a driving requirement for image quality. With typical values of radar resolution, the degradation in resolution from using an intensity image is not affordable. Too many of the responses of an intensity image may be associated with more than one scatterer. This problem is greatly alleviated with complex-image analysis, because then the requirement is only that not too many of the intensity responses be associated with more than two scatterers. We point out again that improving resolution so that the intensity image is adequate is more than just a matter of cost. An arbitrarily high resolution is not achievable with real targets, particularly when they are moving, because of their backscattering properties.

An important practical question is how many dimensions a target must be resolved in, in order to permit reliable identification. An identification system would be simplest and least costly if resolution were required only in the single dimension of range. Although very high range resolution is

expensive to implement in hardware, because the hardware must be designed to accommodate a large signal bandwidth and signal processing must be performed in many range cells, the benefits of high range resolution for imaging and image exploitation are great. Also, the use of resolution only in range allows operating with very short dwell times on a target. Unfortunately, as will be shown later in more detail, in most applications the use of range resolution alone is not sufficient. Although range resolution is of primary importance, it must be supplemented by at least a crude crossrange resolution.

Several effects make the utilization of range resolution alone problematic for aircraft identification. For example, aircraft often roll in an unpredictable manner. The consequence of such a roll on the aircraft range profile is that the wing responses are translated in range relative to the fuselage responses, changing the interference conditions, thereby causing rapid variation in the profile. Wing ordnance may or may not be carried at a particular time, and if it is carried it may or may not be observable, necessitating a large profile database. Perhaps most important, over large aspect sectors, the series of delayed duct returns generated by fighter aircraft can easily be among the strongest returns from the aircraft. We need Doppler resolution to resolve the delayed duct returns from the fuselage returns, as we have illustrated in Figures 1.8 and 1.20. Furthermore, range resolution is progressively less effective on the fuselage as the aspect angle increases, and is least effective near broadside. Similar resolution requirements exist for ground vehicles, for which the ratio of length to width often is large enough to require two-dimensional resolution. Even for ships, when only range resolution is used, serious problems appear when the aspect angle becomes larger.

Thus one comes to the conclusion that reliable target identification in most situations requires at least two-dimensional resolution. However, Doppler (or crossrange) resolution is much more problematic than range resolution. This is so because moving targets typically do not move smoothly, and the implementation of crossrange resolution becomes progressively more problematic as the target's motion becomes more erratic. (The signal phase is very sensitive to even small motion disturbances.) With maneuvering aircraft, ground vehicles moving on rough roads or off the road, and small ships in rough seas, it often is impossible to implement high crossrange resolution and yet achieve sufficient focusing of the responses for accurate feature measurements. The target motion is often too erratic to permit measuring of and correcting the phase functions of the scatterers over the long periods required for high crossrange resolution. Although two-dimensional resolution is needed, with adequate range resolution there is no need for high crossrange resolution, nor may it be desirable. The latter statement refers to the

problems of the motion compensation and spurious responses. The adequacy of lower resolution in crossrange results from the ability to measure the crossrange position of a feature accurately (to a fraction of the crossrange resolution) if the feature has been resolved in range. If range resolution is adequate to resolve scatterers, their crossrange positions thus can be accurately determined even with poor crossrange resolution. Whereas resolution is approximately given by the crossrange width of a response, the accuracy of the position measurement is a small fraction of this width, corresponding to the accuracy with which the position of the resolved response peak (or of the scatterers contributing to the peak) can be measured. The general conclusion is that *reliable target identification requires at least two-dimensional resolution in range and crossrange, with high crossrange resolution less important than high range resolution.*

Should one use three-dimensional resolution? There are ways in which three-dimensional resolution can be implemented in at least some practical applications. However, here the question is whether the cost of implementing three-dimensional resolution is worth the benefits. This may be a matter of opinion, perhaps governed by what kind of approach to target identification one uses. In our opinion, which is based on our experience with complex-image analysis technology, the identification problem is solvable with only two-dimensional resolution in range and crossrange. Although the addition of one more dimension in resolution would unquestionably be beneficial, it does not appear to be either necessary or practically available in most cases.

1.2.2 Wavelength and Resolution in Identification Performance

We have pointed out that the development of a technology for target identification cannot be based on mathematical target models but only on data from real targets. One must analyze representative target data until one obtains enough insight to formulate workable signal-processing algorithms. Now, by far most of the available target data were collected at X-band, and with a range resolution of about 1 ft. Since we used this type of data to develop the technology of complex-image analysis and our identification procedures, the question arises whether we have perhaps developed an identification technology for X-band radars with a range resolution of 1 ft. As a matter of fact, we have indeed done that, but it is not a matter of data availability. As we explain in this section, X-band and an image resolution of about 1 ft not only represents an ideal combination of wavelength and

resolution, but target identification will not work under operational conditions if one deviates too much from these two parameters.

Target identification under operational conditions (automated and for a large target base) is possible only if one chooses the right carrier frequency and the right value of resolution, within relatively narrow margins. To justify this statement, let us recall the conclusion in the preceding section that the measurement of the positions and characteristics of the target features generating dominant responses is essential for target identification, because the typical target does not contain enough measurable special features. The features that give dominant responses are the wave-trapping features. However, they do not naturally generate dominant responses; *carrier frequency and resolution must be selected so that the responses from wave-trapping features will dominate the target-generated background.*

First, consider the choice of carrier frequency. A wave-trapping feature may be viewed as a kind of antenna or trihedral corner reflector, except that it will generally be less efficient in collecting and returning the radar signal. Clearly, as the carrier wavelength becomes larger, the “gain” of the feature will drop, so that its response will no longer be dominant over the background from the multitude of scatterers on the target. If the carrier wavelength chosen is too small, on the other hand, the phase-center wander of the feature will break the formerly dominant response into a set of weaker spurious responses, scattered about the location of the scatterer. The usable band of carrier frequencies thus depends on the size of the wave-trapping features. Our experience with carrier frequencies other than X-band is not so extensive that we can make accurate statements. On the high side of the carrier frequency, the phase-center wander effects can sometimes be observed even at X-band when targets such as aircraft or ground vehicles are imaged, but the effects are not yet serious. Indications are that they become serious at 18 GHz, and problematic at still higher carrier frequencies. At the lower end, C-band should be acceptable. Although it would be speculation to estimate how low a carrier frequency is usable, it probably cannot be much lower than C-band.

Next, we consider the size of the resolution cell. If the resolution cell is larger than the separation of the dominant scatterers, they will not be resolved. In this case it will be impossible to measure scatterer positions and characteristics. If the resolution cell is much smaller than the size of a feature, the feature response will be broken up in some unpredictable manner, losing its dominance without offering a practical possibility of extracting feature position and characteristics. A quantitative investigation of this question is best based on turntable data with very high range resolution, in which case

one can generate images with progressively lower values of range and cross-range resolution. Equally important, with turntable data one has perfect ground truth for the positions of features, so that it is possible to compare measured with actual feature positions and determine the way in which a specific feature response breaks up when resolution is increased.

Tests with turntable data of a ground vehicle demonstrate that the situation is as complicated as one would expect theoretically. Depending on the size and design of a specific feature, one finds all conceivable situations, even for a single ground vehicle at one aspect angle. We changed the size of the resolution cell from 30.5 by 30.5 cm in four equal-percentage steps to 12.5 by 12.5 cm. The tests showed that one response that is dominant at 30.5 cm resolution retains its dominance and even its position when resolution is improved to 12.5 cm, with the only consequence of increasing resolution being the introduction of a curvature into the phase function and a corresponding (small) widening of the response. Such a scatterer presents no problems even at the high end of resolution. Another response widens so much that application of the one-dimensional TSA indicates two scatterers. However, different from two actual scatterers, the measured separation of these does not vary with the angle of the image cut. The response remains dominant, so the task is only the correct interpretation of the measurements. With other responses, the smearing becomes so strong that dominance is lost toward the high end of resolution. Such smearing usually is very asymmetric, and can be over many gates. With still other responses, the smearing becomes so strong with increasing resolution that the response becomes part of the background. The peaks in the peaks plot are mere ripples on a very wide "response." It is not clear how much of such an extreme degree of smearing depends on the "resolution" of an extended wave-trapping feature into parts, and how much can be attributed to phase-center wander. It does not matter, because in either case we see no possibility of using such a feature for target identification.

The tests we have performed were too limited to make precise statements about the best resolution that can be used with a particular type of target. It is clear that for every extended feature there is an optimum size of the resolution cell that could be determined during the processing sequence. In fact, investigating a feature with different degrees of resolution would provide much information about the characteristics of the feature. For the perhaps more practical case in which a single resolution cell size is used on the entire target, one must not choose resolution so high that a significant part of the wave-trapping features no longer generate usable dominant responses.

Indications are that for a target such as a ground vehicle, one cannot operate with a resolution much better than 1 ft; a resolution of 0.5 ft already appears to be too high. At the other end, a resolution of 2 ft causes serious resolution problems for adjacent features. These numbers should be taken only as an indication of the range of the usable values of resolution for ground vehicles.

Of course, there are ways of utilizing higher resolution *without forming* an image with very high resolution. For example, with moving targets a range resolution better than 1 ft would be useful for the motion compensation, for which one should be able to find a scatterer that behaves well even at the high resolution. If higher range and crossrange resolution were available, one would have the flexibility of adapting the value of resolution to each of the dominant features. As a simpler way of utilizing surplus resolution, one could form multiple images with lower values of resolution, and analyze these to obtain more information than is available from a single image. However, our point is the following: The maximum resolution usable on individual features in order to measure scatterer positions and characteristics is limited. If this limiting resolution is exceeded, these essential measurements cannot be performed.

In summary, we conclude that there are relatively small margins in the choice of carrier frequency and resolution for target identification. It is obvious that resolution cannot be too low if the positions and characteristics of the dominant scatterers are to be measured, but not obvious that resolution might be chosen too high for a given image. Similarly, it is obvious that choosing the carrier frequency very low will lead to problems, because the wavelength defines the limits on resolution. It is not obvious that the carrier frequency can be too high, a fact that has unfortunate implications for missile-borne radars. Lastly, the preceding discussion might give the impression that the fact that some combination of carrier frequency and resolution offers a high performance potential is just a fortunate accident. This is not the case. It is natural for the wave-trapping features of man-made targets to be separated by more than their extents, so that one must be able to choose a resolution cell small enough for resolving the various features from each other and yet not so small that the individual features are broken up. The optimum size of the resolution cell would be expected to change from one type of target to another. Since we had turntable data available for a military ground vehicle, we were able to establish that resolution must be about one foot for this type of target. We know that this resolution is also good for aircraft and ships, but in both cases we did not investigate what value of resolution is optimum and how wide the practical margins are.

1.2.3 Section Summary

Since the dominant features prevent us from utilizing the optical-type background image generated by the multitude of weak scatterers, we have no choice but to base target identification on the dominant features.

Reliable target identification requires at least two-dimensional resolution in range and crossrange, with high crossrange resolution less important than high range resolution.

Target identification is possible only if one chooses the carrier frequency and resolution so that genuine responses from wave-trapping features dominate the image.

1.3 Measurement of Range and Crossrange: The Two-Scatterer Algorithm

Aside from spurious responses, the intensity image of a man-made target contains three types of response: (1) responses generated by a single scatterer, (2) responses generated by two scatterers, and (3) responses generated by more than two scatterers. The position of the scatterer associated with the first type of response is the position of the response peak. In order to obtain the positions of the two scatterers associated with the second type of response, we must analyze the complex response with the two-scatterer algorithm (TSA). This algorithm merely implements the basic resolution capability of radar, which is $1/B$ in range delay and $1/T$ in Doppler. With the third type of response, we must attempt to select response parts that come mainly from two scatterers and can be analyzed with the TSA. If this is not possible, and if such a failure occurs for too large a part of the main responses, radar resolution is not adequate for the task.

At the other extreme, we could try increasing resolution to the point that all of the significant responses of the intensity image come from single scatterers, so that the image phase need not be used for achieving the inherent resolution capability. Aside from the expense of such an approach, it appears unlikely that images with such extreme resolution can be generated with a quality that allows automated target identification. The reason again is that the scatterers on a man-made target rarely behave sufficiently like fixed point scatterers, for which no problems would appear as resolution is progressively increased. The basic (and important) algorithm for the resolution of two scatterers is treated in this section. We first summarize the theory of the TSA for two ideal point scatterers, and later show how the TSA is applied to real data. A full mathematical derivation is contained in Appendix A.

1.3.1 Resolution Principles

1.3.1.1 Resolution of Two Point Scatterers in One Dimension

In the development of the conventional (point target) resolution theory, resolution of two point scatterers was defined via the following question: In the envelope-detected processor output, how close can two point scatterers of equal strengths be so that one still observes two response peaks rather than a single peak? Aside from the fact that a definition of resolution that does not take into account the need for measuring range or Doppler with reasonable accuracy is not very practical, such a definition also is loose in that performance depends on the phase relation between the two scatterers as well as their relative strengths. With this definition, resolution of two scatterers that are in phase is much inferior to resolution when they are out of phase, because constructive interference tends to generate a single response peak, whereas destructive phase interference leads to two response peaks even for very close scatterer separations. Such a conclusion on the consequences of scatterer phasing is suspicious, because scatterers return more energy when they are in phase than when they are out of phase. In fact, with a meaningful definition of resolution, under limiting conditions the performance is indeed better when the two scatterers are in phase. Going back to the original definition, in an average sense (with respect to the scatterer phasing) the chance is good that one can observe two peaks when the scatterers are separated by at least the half-power width of the point-target response of the waveform. Since the half-power width of the point-target response of a waveform with bandwidth B and duration T is about $1/B$ in range delay and $1/T$ in Doppler, resolution in range delay came to be defined as the reciprocal of the signal bandwidth B , and resolution in Doppler as the reciprocal of the signal duration T . Since weighting for sidelobe suppression widens the point-target response, it was also concluded that weighting degrades resolution.

The question of resolution performance on ideal point scatterers is so basic and has such importance that it is worthwhile to reexamine it [1, 4]. In the left half of Figure 1.22, we show the composite crossrange response from two ideal point scatterers of equal strengths and zero phase difference, with a separation of 1.2 crossrange resolution cells or gates. (For a signal duration T , Doppler separation is $1.2/T$.) The intensity response of the left upper plot shows that there is a single response even when the separation of the two scatterers is 20% larger than the basic resolution cell. However, the half-power width of the response is 35% larger than that of a single point scatterer.

The transform of the response, as shown in the right half of Figure 1.22, is more revealing. The amplitude is strongly modulated, and the phase has




Figure 1.22 Response and transform of two point scatterers, equal strengths, in phase, separation of 1.2 resolution cells.

half-cycle jumps at the times of the amplitude minima. This transform of the response is, of course, just the composite signal returned from the two scatterers. Thus we find that the signal is easier to understand than the response. The general return from two scatterers is derived in Appendix A.

For our present introductory discussion, we consider the fact that the return from two interfering scatterers may be represented by the well-known phasor diagram of one vector rotating about another. The sum vector varies from the maximum when both vectors are aligned, to minimum when they are out of phase, with the sum dropping to zero for equal-strength scatterers. At this point the phase switches its sign, which implies a half-cycle phase jump. (This goes back to the definitions of amplitude and phase of a complex signal, where a real signal amplitude that drops to zero must be accompanied by a half-cycle phase jump.) These facts are represented by the amplitude and phase functions on the right. We chose the separation of the two scatterers slightly larger than one resolution cell in order to include the phase jumps in the figure.

Clearly, if we can observe the signal from one amplitude minimum to the next, or a full modulation period, we can recognize the presence of two

scatterers, and from the modulation period obtain the differential Doppler. Thus we find that *the composite signal from the two scatterers allows us to achieve a Doppler resolution of $1/T$, even though there is a single intensity peak.* From the flat phase function we conclude that the two scatterers are located symmetrically with respect to the center of the image interval, and the fact that the amplitude reaches a null implies that the two scatterers have equal strengths. Hence, the signal not only easily reveals the fact that there are two scatterers, but also their positions and strengths. The relations for arbitrary scatterer separations, relative strengths, and phase differences are derived in Appendix A. They will be used when the actual resolution algorithm is applied.

Figure 1.23 shows the response and signal of the two point scatterers, with only a changed phase relation. Instead of a phase difference of 0° we introduce a phase difference of 180° , changing from constructive to destructive interference between the two scatterers. The corresponding response is shown in the left half of the figure, with the transform or return signal in the right half. The intensity response now correctly shows two peaks. However, whereas the actual scatterer positions are at -0.6 and 0.6 crossrange gates, the positions of the peaks are at -0.9 and 0.9 gates. This is a practically




Figure 1.23 Same as Figure 1.22, except destructive interference.

unacceptable error of 0.3 resolution cells. (In practice there are too many other sources of error to allow an error of 0.3 cells in the ideal case.) The signal in the right half of the figure still allows us to measure the positions of the two scatterers relatively accurately, but not as accurately as in the case of constructive interference. The problem lies with measuring the length of the modulation cycle, in the present instance the separation of the amplitude maxima.

The difference between the two cases would not be significant if we could measure the return signal directly. In practice, we form an image, but in the process we must use weighting for sidelobe suppression. Then a response or two responses are selected by windowing, the transform is taken, and the resulting signals are dewighted to remove the weighting that is no longer needed. This entails unavoidable interference and errors, and in our present illustration of the ideal point scatterers we included the weighting and dewighting so as not to be entirely unrealistic. The weighting/dewighting process results in distortions of the amplitude function of the signal near the fringes of the display window, in this instance near the amplitude maxima, as apparent in the right top curve in Figure 1.23. Thus we cannot measure the positions of the maxima as accurately as we could measure the positions of the minima in Figure 1.22, but the resulting position error is far smaller than if one accepted the positions of the intensity peaks.

When two scatterers have a constant differential Doppler, the phase relation between them changes linearly with time. In practice, we thus cannot time our observation so that there is a specified phase relation between the scatterers. Similarly, since the Doppler difference between the scatterers is unknown a priori, the signal duration (or observation interval) likewise cannot be selected so that a specified number of amplitude maxima or minima are observed. The actual situation can be understood from a generalization of the interference pattern between two point scatterers. Figures 1.22 and 1.23 show different parts of this pattern. The general amplitude/phase pattern of the combined return from two ideal point scatterers of equal amplitudes is given in Figure 1.24.

In the previous illustrations we considered an observation interval, or signal duration, slightly larger than one modulation period. In order to extract the positions of the scatterers and their relative strengths from the composite return, we must observe at least one full modulation cycle of the repetitive signal. This requirement defines the inherent resolution limit of radar. The terms constructive and destructive interference can apply only at specific instants, and they refer to the phase relation of the two scatterers




Figure 1.24 General amplitude/phase pattern for two point scatterers of equal amplitudes.

at the center of the observation interval. By comparison, how long do we have to observe the scatterers in order to obtain their positions with adequate accuracy from the positions of their intensity responses? As is shown in [1, 4], we need about two modulation periods rather than one. This justifies the claim that utilizing the intensity image degrades resolution by about a factor of two. As was discussed above, when we work near the resolution limit of $1/T$ and have only one full modulation cycle available for the measurement, there is a performance difference between the cases of constructive and destructive interference. As the number of observed cycles increases, effects on the fringes of the displayed curves become insignificant, so that the performance difference disappears; but then resolution is no problem anyway.

In contrast, the distinction between constructive and destructive interference, or more generally phase differences of 0° and 180° , becomes rather important when the observation interval decreases below one modulation period. This is equivalent to the statement that the two scatterers are separated by less than one resolution cell, so that they cannot be properly resolved. Suppose, as an extreme, the differential Doppler between the two scatterers is so small in relation to the signal duration that the observation interval is a small fraction of one modulation period in Figure 1.24. Since the

phasing of the two scatterers depends on the accidental time of observation, this observation interval may fall anywhere on the curves of Figure 1.24. If it cuts out a window near an amplitude maximum, there will be a strong return that appears to come from a single scatterer. As the window is shifted in time (or the phasing of the scatterers changes due to the differential Doppler), when the window is centered on an amplitude minimum, the return will not be measurable in the background. Thus, we essentially have a scatterer with a fluctuating cross section that may or may not be observable, depending on the phase between the two scatterers at the time they are illuminated. (The example illustrates that the problem of a “fluctuating” target disappears with increasing resolution.)

The conventional definition of resolution considers two point scatterers of equal strength, because different relative strengths lead to different deformations of the intensity pattern. In other words, resolution becomes even less precisely defined on the basis of the intensity pattern if two scatterers can have different strengths. Nevertheless, as is shown in [1, 4], it remains true that actual resolution on the basis of the intensity response remains at about $2/T$ or $2/B$ even when the scatterer strengths differ. We now show that resolution, properly defined using the complex response, is independent of the scatterer strengths. We use the same two-scatterer configuration as for Figure 1.24, except that the amplitude of one scatterer is half the amplitude of the other. The resulting amplitude/phase pattern is shown in Figure 1.25. By comparison with Figure 1.24, the amplitude minima are less deep and the phase near an amplitude minimum changes less sharply. Nevertheless, we will retain the term “phase jump” to denote the association of the phase change with an amplitude minimum. A phasor diagram makes it clear that the modulation period is still governed by the differential Doppler, and that the phase change near the amplitude minimum is more gradual than when both vectors have equal lengths. This is seen quantitatively from the relations in Appendix A.

Even though the use of the complex response improves resolution by a factor of two over the performance obtained with the intensity response, this is not superresolution. The improvement is due merely to the utilization of the full signal, amplitude and phase, rather than only one of its two components. Superresolution starts when we try to resolve the two scatterers by observing less than a full cycle of the modulation pattern, since we then cannot directly measure the separation of the amplitude maxima or minima, or of the phase jumps. As an illustration, in Figure 1.26 we show the same case of constructive scatterer interference as in Figure 1.22, except that the




Figure 1.25 Amplitude/phase pattern when the ratio of scatterer amplitudes is 0.5.

separation of the two scatterers was reduced to 0.8 gates. Thus we cannot measure the positions of the two amplitude nulls or phase jumps. In fact, we cannot tell from the response or the signal that the amplitude minima should be null, because we do not know that the two scatterers have equal strengths.

In the absence of noise, we can compare the shape of the amplitude function at the top right of Figure 1.26 with the shapes of the amplitude functions for different relative strengths of the two scatterers, such as in Figure 1.24 for equal strengths or Figure 1.25 for an amplitude ratio of 0.5. Without noise, we could indeed determine the depth of the amplitude minima and their positions by this process, so that the scatterer positions could be determined. This would be true superresolution. The achievable accuracy evidently degrades as noise is allowed, and even more seriously if there is interference from other scatterers. Also, the fitting process between the idealized and the measured curves becomes more uncertain when a smaller fraction of a modulation cycle is observed; that is, when the scatterer separation is reduced.

The problem of superresolution varies with the phase difference between the two scatterers. In Figure 1.26 we considered constructive interference. In Figure 1.27 we show the same situation as in Figure 1.26, but with destructive instead of constructive interference. Toward the fringes of




Figure 1.26 Two scatterer returns as in Figure 1.22, except for a scatterer separation of 0.8 gates.

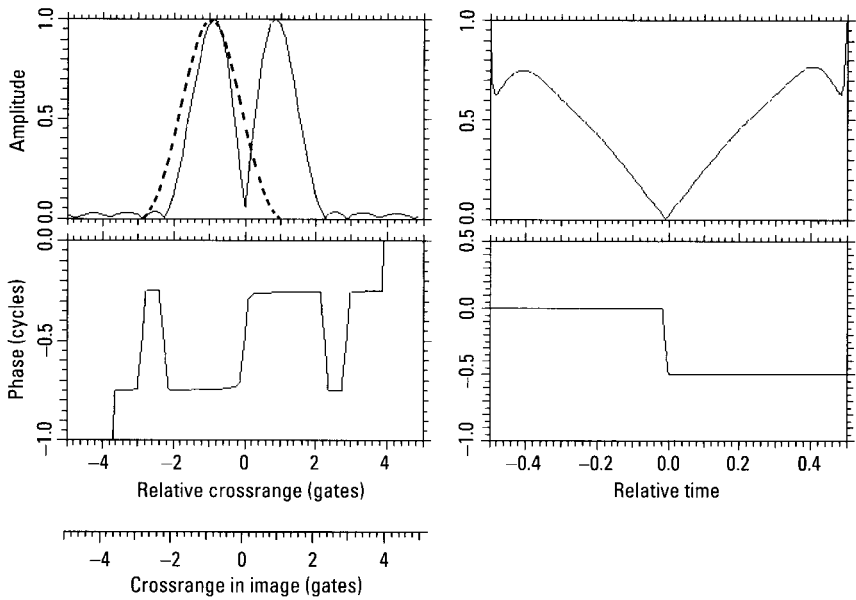


Figure 1.27 Repeat of Figure 1.26, but with destructive interference.

the displayed interval, where the composite return is strong, we have the distortions due to weighting and deweighting. In the center of the interval the signal strength is low and the amplitude changes nearly linearly. The combination of the two facts makes fitting an idealized curve much more problematic than in the case of constructive interference, so that superresolution gives less accurate results for destructive phase interference.

For ideal point scatterers as considered above, and in a benign noise and interference background, the inherent resolution performance of radar could be realized without the methods of complex-image analysis. If a target is modeled by a set of fixed point scatterers, the resolution problem can be formulated in mathematical terms, and processing algorithms can be derived on that basis. This is the approach for various superresolution methods. Indeed, the main part of the gain claimed for superresolution methods is the factor of two in resolution lost when only the intensity output is utilized. The reason that superresolution methods do not work in practice is that complicated targets cannot be modeled by sets of fixed point scatterers. Much more complicated target models have also been introduced. Our position, which is supported by the analysis of images generated by very complicated models, is that such complicated targets as aircraft, ground vehicles, and ships cannot be mathematically modeled with sufficient realism to serve as the basis for workable signal-processing algorithms.

Despite the accuracy problems with superresolution, we do make use of a small degree of superresolution in complex-image analysis. We point out, however, that this is quite different from mathematically modeling a target and simply accepting the results of superresolution processing. First, we attempt only a small degree of superresolution. Second, this is done only for the simple case of two point scatterers. Third, in each such case we check the adequacy of the two-scatterer model. This will be discussed in more detail later.

Our basic discussion of one-dimensional resolution was given in terms of Doppler resolution. It follows from the duality between time and frequency that the discussion applies equally well for range resolution. Instead of a response in crossrange, we then have a response in range. The transform now is the frequency spectrum instead of the time signal. The composite spectrum from two scatterers has amplitude and phase modulation that permits us to determine the ranges and strengths of the two scatterers. All of the preceding figures apply when the coordinates are changed appropriately.

1.3.1.2 Resolution of Two Point Scatterers in Two Dimensions

We have generally shown that the limiting resolution performance of radar can be realized by examining the composite signal returned from two

scatterers separated in crossrange, or the composite spectrum from two scatterers separated in range. In practice, the first case requires taking an image cut in the range gate of the two scatterers (they are separated only in crossrange), and taking the transform. In the second case we take an image cut in the crossrange gate of the two scatterers, and again examine the transform of the image cut. The question is how to extend the procedure to the more general two-dimensional case when the two scatterers must be resolved in range as well as crossrange. The only practical solution (in the sense that it must also apply to real targets) we have found is to examine a succession of image cuts, at different angles through a single point in the range/crossrange plane. In our implementation, we use 18 equally spaced image cuts in the image plane, thus considering an image cut every 10° . The point of intersection of the 18 image cuts is usually chosen at a response peak of the two-dimensional image, but this need not be the case.

In principle, it is possible to perform a two-dimensional analysis of an image response without examining one-dimensional image cuts. As we show below, even the one-dimensional analysis requires adaptive adjustment of a transform window about the response. A similar adjustment of the boundaries of a two-dimensional region about the response is so difficult that we have not attempted it.

As an illustration of the procedure by which two point scatterers are resolved in both range and crossrange, in Figure 1.28, Point A is the intersection point of the (18, in our implementation) diagonal image cuts, or pivot




Figure 1.28 Implementation of two-dimensional resolution.

point, and Point B is the location of a scatterer. The circles about Point B represent equal-amplitude contours for a point scatterer. In each of the diagonal image cuts, the maximum amplitude occurs at the point of the cut closest to the scatterer position. The locus of these points is traced by the heavy curve in the figure. The shape of the (complex) response generated by cutting through the response of the point scatterer is determined by the two-dimensional weighting function (for sidelobe suppression) used to form the image. For example, with circular Gaussian weighting the image cuts at the various angles generate Gaussian responses of the same width and linear phase. Thus, the actual positions of the amplitude maxima in the one-dimensional cuts, if measured accurately in the image, should fall on the heavy curve in the figure. The distance from Point A to any of these positions varies as the cosine of the angle between the image cut and the line from Point A to Point B, as given in the figure. Although we have assumed Gaussian weighting, the differences for different weighting functions are minor.

If we examine image cuts in the vicinity of a single point scatterer, we can trace out the entire cosine pattern obtained by progressively cutting through the response of the scatterer. However, if two scatterers are involved, only part of the pattern for each can be obtained. To understand this, consider two point scatterers and an image cut along the line connecting the two responses. Let this line rotate about the midpoint between the two scatterers. As is seen from Figure 1.28, the measurement point in accordance with the heavy curve is at the perpendicular projection of the scatterer position on the image cut. When the image cut is along the line connecting the two scatterers, the separation of the projected points for the two responses is thus largest. As the image cut is rotated, when it is at 90° relative to the initial connecting line, the two projected points coincide, so that the two responses fully overlap. Somewhere between the initial and the last orientation of the image cut the projected points, and hence the responses in the image cut, will be separated by the minimum resolvable distance. At higher rotation angles we cannot measure the locations of both projected points. Thus, the measurable locus of the projected point is truncated to less than the full cosine pattern. The degree of truncation depends on the relative scatterer strengths and phasing, but truncation generally occurs when the projected points are separated by slightly less than one gate. This is an example of using a small degree of superresolution to perform the measurements over an extended angular sector, with the scatterer position robustly defined by a curved segment rather than by just a point.

1.3.1.3 A Modified Procedure for Interactive Analysis

Since complex-image analysis amounts to an expert system approach where one must first analyze real data to obtain insight and develop a processing approach, and then in a separate step automate the algorithms, the TSA is used in two ways: for interactive analysis and fully automated processing. The two-dimensional TSA described above is too unwieldy for interactive processing, because it involves 18 image cuts in which scatterer measurements must be made. For this reason, we utilize a modified version in which the two-dimensional positions of the two scatterers are derived from one fixed-range and one fixed-crossrange cut through the response peak to be analyzed, possibly refined by a few additional fixed-range and fixed-crossrange cuts in the vicinity of the response peak. For real data, this modified version is less accurate than that involving diagonal cuts, but it is often good enough for the purposes of interactive processing.

Often, there may be two scatterer positions derived from the image cut in the range gate of the response, and two positions from the image cut in the crossrange gate of the response. The correct range position must be associated with the correct crossrange gate position, so as to avoid scatterer "ghosts." The association of the measurement results from the fixed-range and fixed-crossrange cuts may require additional cuts in the vicinity of the response peak. The algorithm that determines where to place these cuts incorporates several branch points. In order that the details of placing the cuts not detract from the explanation of the basic principles of implementing the inherent resolution capability of radar, we defer explanation of the modified version of the two-dimensional TSA to Appendix F.

1.3.2 Implementation of the TSA for Ideal Point Targets

In this section we discuss the basic measurement algorithm that allows one to realize the resolution performance inherent in radar. This will be partly a brief summary of material found in [1] and partly an extension to truly two-dimensional measurements. This basic resolution algorithm is indispensable to working with real data. We derive two-dimensional scatterer positions from position measurements in a series of one-dimensional image cuts through a response. As mentioned above, it is possible in principle to perform a two-dimensional analysis of each response without examining one-dimensional image cuts, but adaptive adjustment of the boundaries of a two-dimensional region about the response is so difficult that we have not attempted it.

1.3.2.1 Basic One-Dimensional TSA

The one-dimensional TSA compares a Fourier transform of a response to the signal generated by two interfering fixed point scatterers, which is characterized by a sinusoidally varying squared amplitude (power), with phase “jumps” at the times of amplitude minima and linear phase at the times of amplitude maxima, as shown in Figure 1.29. The left half of the figure shows an image cut through the two responses, separated by 5.6 gates (resolution cells). The top box gives the image amplitude; the bottom gives the image phase. The right half of the figure, which shows the Fourier transform of the left illustrates the ideal two-scatterer pattern. Again, the top box gives the transform amplitude and the bottom box gives the transform phase. Note the correspondence of these in Figure 1.25.

As discussed above in conjunction with Figures 1.24 and 1.25, we use the term “phase jumps” for the rapid changes of the transform phase at the times of the amplitude minima, even though the change is strictly a jump only for two equal-strength scatterers. The image abscissa is labeled “diagonal gates,” as a reminder that the one-dimensional procedure is not constrained to fixed-range or fixed-crossrange cuts. A separation of one gate along the abscissa corresponds to critical sampling in the image cut. The transform abscissa normalization, from -0.5 to 0.5 , is chosen so that, with phase measured in cycles, phase slope directly gives position in the image cut. With this normalization, a transform phase slope of 1.0 (accompanied by a constant




Figure 1.29 Two-scatterer interference pattern.

transform amplitude) corresponds to an image response located 1.0 gates left of the center of the image cut interval displayed in the figure.

In the example case of Figure 1.29, the positions and amplitudes of the fully resolved scatterers can be measured using the conventional image amplitude response, without phase data and without performing the transform. The two responses are separated so much that the improvement in resolution by the factor of two obtained by using the complex image is not needed. In fact, before applying the TSA to a response, we measure the normalized half-power width of the response (its half-power width divided by that of a fixed point scatterer) and its skewness (the ratio of the right and left half-power half-widths). If the normalized half-power width and the skewness are sufficiently close to unity (in our software, we require that they be within 6% and 12% of unity, respectively), we interpret the image cut as corresponding to a single well-resolved scatterer at the location of the amplitude peak and do not apply the TSA analysis. The utility of the transform plots lies in our ability to interpret the information when the scatterers are separated by less than two resolution cells or gates. With extrapolation to obtain some degree of superresolution, the interpretation of the transform is possible until the scatterer separation decreases to about half a resolution cell, depending on the relative scatterer phasing. Thus, at the most we attempt superresolution by a factor of two, but only when the pattern of the transform is clean enough.

The analysis of the two-scatterer pattern is quantitative (see Appendix A), but even inspection of the pattern allows us to extract information about the two interfering scatterers. The separation of successive amplitude minima (or successive phase jumps, or successive amplitude maxima) tells us the separation of the two scatterers. With the normalization employed in the figure, the separation in resolution cells of the scatterers is the reciprocal of the separation of the minima of the transform amplitude. In this case, a normalized separation of the minima of about 0.18 gives the correct scatterer separation of 5.6 resolution cells.

The size of the phase jumps (or, alternatively, the ratio of amplitude minima to maxima) tells us the relative strength of the two scatterers. Equal scatterer amplitudes give half-cycle jumps and zero-amplitude minima. When the two amplitudes differ, as in Figure 1.29, the amplitude minima are less pronounced relative to the maximum, and the phase jumps are smaller and less sharp. The phase slope at the time of the amplitude maximum corresponds to a scatterer location close to the position of the stronger of the two scatterers, by the phase capture effect. By combining this phase slope and the phase slope over the entire transform with the scatterer

separation and relative strengths, we can determine the positions and amplitudes of the two scatterers (see Appendix A).

Figure 1.29 shows an interference pattern of two well-resolved scatterers. In practice, the two-scatterer pattern is of interest only for unresolved intensity responses, responses separated by less than about two resolution cells. We now proceed to such a case. Figure 1.30 shows a fixed-range image cut through the intensity peak of a simulated response composed of two interfering fixed point scatterers. For such a fixed-range cut, the data points vary in crossrange. The image domain abscissa is labeled with two scales, one giving crossrange in the image, and the “relative” scale being an arbitrary translation of the first, here set to zero at the peak location. The transform of such a fixed-range cut is the signal corresponding to the response peak, so the transform abscissa is labeled “Relative time.”

The image cut itself displays a widened amplitude peak with a bulge on its right side, accompanied by a curved phase function. Rather than judge the amplitude peak width by eye, we always consider the normalized half-power width of the peak, which in this case is 1.45. The small variation of the image phase, over just about 0.1 cycles, shows that the two scatterers contributing to the response are approximately in phase. Based on interpreting the image data, we could not easily decide whether this response corresponds to two interfering point-like scatterers or to one shifting scatterer. The decision is

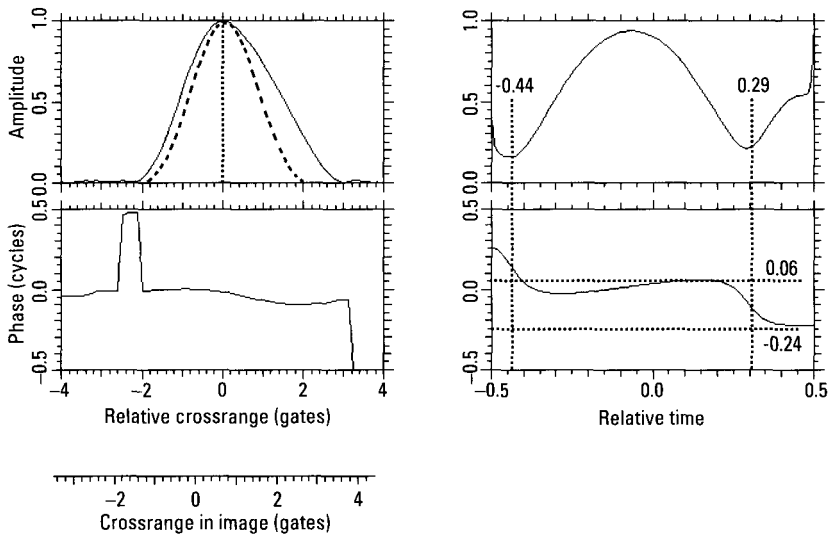


Figure 1.30 Fixed-range image cut through response from two scatterers.

much easier with the transform data, which is a prototypical example of a two-scatterer pattern (sinusoidally varying squared amplitude, with phase jumps at the times of amplitude minima and linear phase at the times of amplitude maxima).

If we measure the times of the amplitude minima (which correspond to the times of the phase jumps) and the size of the phase jump, the scatterer positions can be calculated (by the process described in Appendix A). The dashed vertical lines of Figure 1.30 show that the amplitude minima occur at relative times -0.44 and 0.29 . The phase jump appears to be 0.30 cycles, between the dashed horizontal lines tangent to the transform phase curve. However, because the phase jump is not instantaneous, the relative times when the tangency occurs will change if a linear phase function is added to the transform phase curve, corresponding to a different centering of the image response in the image cut window. If the image cut of Figure 1.30 began at Relative Crossrange Gate -3.0 instead of -4.0 , the transform phase would yield different times of tangency, and the measured size of the phase jump would change.

In order to eliminate this variability in the size of the phase jump, we remove any linear phase slope introduced by the choice of the position of the transform window. In order to do this, we perform a linear least-squares fit to the transform phase, subtract the fit, then define the phase jump. This procedure is justified in Appendix A. We note here that the size of the phase jump may also be derived from the ratio of the amplitude minima to the amplitude maxima, and that the degree of consistency between the sizes as measured from phase jumps and derived from amplitude minima and maxima provides one measure of the agreement of the data with the model of two interfering scatterers. Such measures allow one to determine when the model is applicable, and to estimate the error in the derived scatterer positions. Figure 1.31 shows Figure 1.30 after the removal of the linear fit to the transform phase. Corresponding to this removal, the image cut has been circularly shifted; the image cut peak, shown by the dashed vertical line, has moved from Relative Crossrange Gate 0.00 in Figure 1.30 to Gate 0.27 in Figure 1.31. The horizontal dashed lines tangent to the transform phase of Figure 1.31 show a phase jump of 0.25 cycles. With the phase jump relative times of -0.44 and 0.29 , the TSA (as given in Appendix A) yields scatterer positions of Crossrange Gates -0.05 and 1.33 .

For the sake of brevity, in the following illustrations that involve measurements of phase jump size, we show the image cut and its transform, but not the result of subtracting the linear fit to the transform phase. We give values for the phase jump as measured after the subtraction of the linear fit.

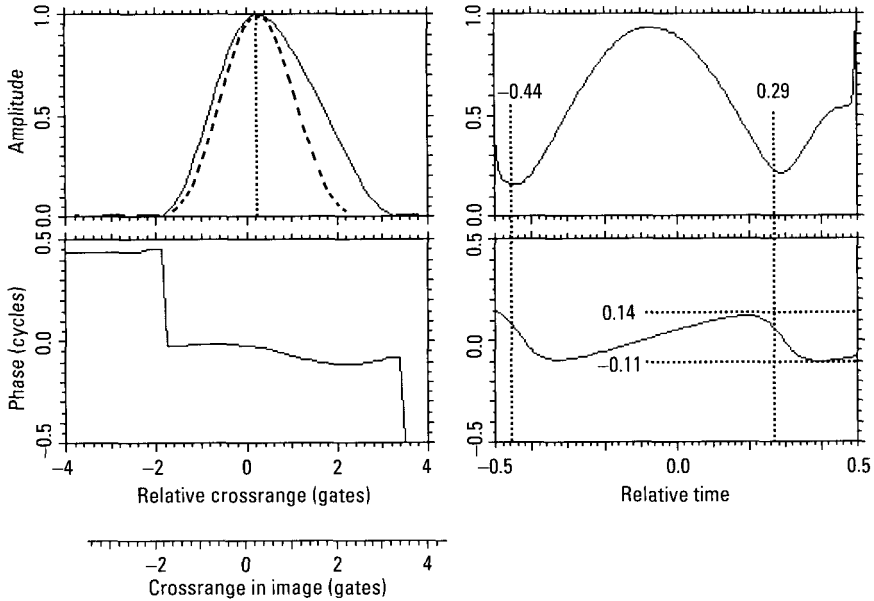


Figure 1.31 Figure 1.30 with linear transform phase removed.

Figure 1.32 shows the fixed-crossrange cut through the same two-dimensional image intensity peak as used for the fixed-range cut of Figure 1.30. For the fixed-crossrange cut the data points vary in range, and the transform is the spectrum corresponding to the response peak, so its abscissa is labeled “Relative frequency.” The response in this image cut has a normalized half-power of 1.02 and a nearly linear phase. The transform amplitude changes very slowly with frequency, and the transform phase is linear. This corresponds to a single scatterer, or to two scatterers that are so close in range relative to resolution performance that they might as well be considered a single scatterer. Thus, we must estimate a single range at the peak position in the cut, Range Gate 0.11.

We assign each measurement an uncertainty based on the degree of agreement of the measurement and the one-scatterer or two-scatterer patterns, but with a minimum uncertainty of 0.2 gates (which applies for both dimensions of our example). Details are contained in Appendix A.

A more complicated situation arises if a small amount of extrapolation, or superresolution, is required to measure the positions of two scatterers that contribute to a response. Figure 1.33 shows a fixed-range image cut through another simulated response composed of two fixed point scatterers. The

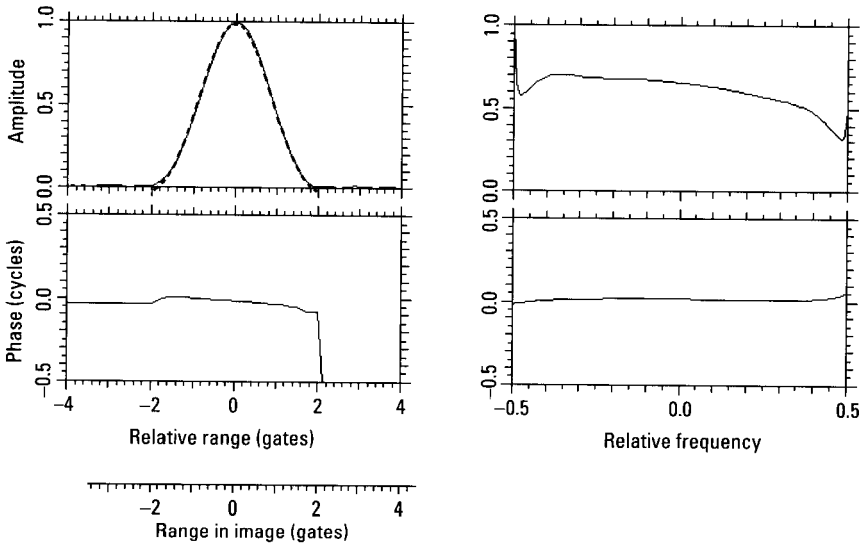


Figure 1.32 Fixed-crossrange image cut through response from two scatterers.

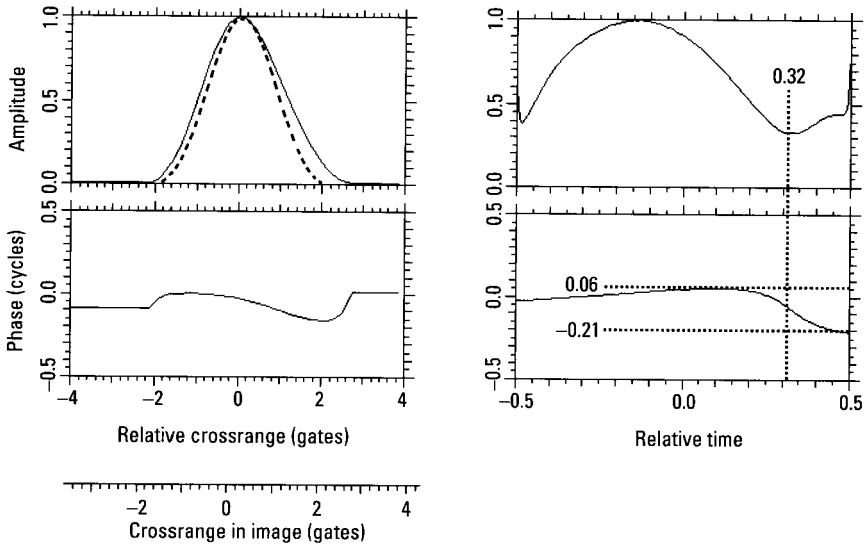


Figure 1.33 Fixed-range cut through another simulated response.

image response has a normalized half-power width of 1.21, with the right side of the amplitude response wider than the left, and a curved phase

function. The small amount of phase curvature is consistent with two interfering point scatterers that are roughly in phase, as well as with a single scatterer with a shifting phase center. The transform shows that the response corresponds to two interfering scatterers. The transform amplitude has the required variation, with the transform phase exhibiting a phase jump (indicated by the vertical dashed line) when the transform amplitude passes through its minimum, near relative time 0.32.

We observe only one phase jump in the transform, and must be suspicious about the null in the transform amplitude at relative time -0.48 , because of the smoothness of the accompanying phase function, and imperfections in deweighting the windowed image response. However, the transform amplitude is very symmetric about its peak near relative time -0.13 . We can use this symmetry to extrapolate the position of a second amplitude null (and corresponding phase jump); we substitute a symmetrical function about the peak at -0.13 , mirroring the section from -0.13 to 0.33 in order to find the true position of the null at -0.59 .

The dashed horizontal lines of Figure 1.33 show the size of the phase jump, based on the displayed transform phase function. However, as discussed above, we must subtract a linear fit to the phase function before measuring the jump. Doing so gives a phase jump of 0.17 cycles. With the jump times of -0.59 and 0.33 , the algorithm of Appendix A gives crossrange positions of Crossrange Gates -0.01 and 1.06 (the actual positions are 0.00 and 1.10).

Measurement uncertainty generally increases when such extrapolation is necessary, because fewer checks of consistency between the data and the two-scatterer model are available. For example, we cannot compare the values of different phase jumps and amplitude minima, but are restricted to examining consistency of one phase jump with its corresponding amplitude minimum. The very act of extrapolating the time of the second phase jump also contributes to the increased uncertainty. In order to minimize this contribution, our automated measurement procedure extrapolates in three different ways, and requires that at least two be consistent. The first extrapolation method is via the symmetry discussed above. The second fits a cubic polynomial to the amplitude between the transform peak and close to the transform edge, then finds the first zero or minimum of the fit. The third exploits the fact (shown in Appendix A) that the two-scatterer transform power has a sinusoidal variation whose period is the reciprocal of the scatterer separation. This means that the derivative of the transform power is a pure sinusoid with the same period. Thus, the peak of the Fourier transform of the derivative of the transform power occurs at the separation of the two scatterers, and the reciprocal of this is the separation of the transform phase jumps.

1.3.2.2 Two-Dimensional TSA

As described in Section 1.3.1.2, we analyze an image response by examining 18 evenly spaced image cuts passing through the response peak. The analytic result for perfect measurements in the one-dimensional cuts (see Figure 1.28) shows that, as the cut angle varies, the distance from the measured one-dimensional position to the intersection point of the image cuts should have a sinusoidal variation. We can describe this mathematically as

$$d(\theta) = d_0 \cos(\theta - \theta_0) \quad (1.1)$$

where θ is the polar angle of a cut, θ_0 is the polar angle of the line from the intersection point of the cuts to the scatterer, d_0 is the distance from the intersection point to the scatterer, and $d(\theta)$ is the distance from the intersection point to the measured location in the cut at polar angle θ . If two or more scatterers interfere with one another, resolution considerations dictate that we will be able to observe only part of the cosine pattern.

To perform the association of the one-dimensional measurements, we compare subsets of the one-dimensional measurements to the analytic form, solving for d_0 and θ_0 by linear-least-squares. Using standard error propagation and the estimated uncertainties in the one-dimensional measurements, we derive uncertainties in d_0 and θ_0 . Generally, several overlapping subsets correspond to a given scatterer, and we must choose one for the association. We use that subset with the smallest variance about the analytic form. In considering subsets, we begin with those one-dimensional measurements away from the pivot point (the intersection point) of the image cuts. The association results away from the pivot point dictate how we treat the measurements near it, because a scatterer away from the pivot point implies that measurements in cuts roughly perpendicular to the line between the pivot point and the scatterer will be near the pivot point, and are not indicative of a scatterer near the pivot point. Rather, they must be disregarded.

Figure 1.34 shows the results of diagonal cuts for the simulated data used in Figures 1.30 and 1.32. Line segments in the figure indicate positions measured in one-dimensional cuts, with the length of the segment showing the measurement uncertainty. The minimum uncertainty of any measurement is taken as 0.2 resolution cells. Crosses indicate two-dimensional scatterer positions derived with the association procedure described above. Curves through the crosses show the cosine dependence of the fit to the data. Annular sections show the uncertainty about the two-dimensional positions, with a minimum radial uncertainty of 0.2 resolution cells. Circles show the




Figure 1.34 Results of diagonal cuts for case of Figures 1.30 and 1.32.

actual simulated scatterer positions. The light background curves are amplitude contours separated by 3 dB. The contours show that one could not derive the scatterer positions on the basis of amplitude alone. The derived scatterer positions of Crossrange Gate 0.00 ± 0.20 and Range Gate 0.01 ± 0.38 , and Crossrange Gate 1.37 ± 0.24 and Range Gate 0.61 ± 0.47 are in excellent agreement with the actual positions of Crossrange Gate 0.00 and Range Gate 0.00, and Crossrange Gate 1.40 and Range Gate 0.60.

Figure 1.35 shows the results of diagonal cuts for the simulated data used in Figure 1.33, in the same format as in Figure 1.34. The measured scatterer positions of Crossrange Gate -0.01 ± 0.30 and Range Gate -0.01 ± 0.37 , and Crossrange Gate 1.19 ± 0.20 and Range Gate 0.72 ± 0.23 are in good agreement with the actual scatterer positions of Crossrange Gate 0.00 and Range Gate 0.00, and Crossrange Gate 1.10 and Range Gate 0.90. The difference between the derived and actual scatterer positions for the second scatterer is attributable to a mismeasurement in the position represented by the lowest line segment intersecting the uncertainty annulus. This mismeasurement is due partly to imperfections in automation and partly to the implementation of windowing procedures designed to minimize the effects




Figure 1.35 Result of diagonal cuts for case of Figure 1.33.

on one-dimensional cuts of interference from more than two scatterers. These procedures are discussed in Section 1.3.3.1.

1.3.3 Application to Real Data

We have described the one-dimensional TSA, and how the results of 18 one-dimensional image cuts are combined into the two-dimensional TSA. This was done for two ideal point scatterers. The application of the TSA to real data poses a variety of problems, and in this section we discuss the specifics of how we have solved these problems. Such details tend to detract from the main issues of target identification. Hence, at this point it might be appropriate for the reader to assume that the TSA works on real data, and to postpone studying this section in detail until the actual processing algorithms become of interest.

1.3.3.1 Properties of Real-Data Analysis

A two-dimensional image of a real target generally contains a large variety of responses. There may be a well-resolved, or isolated, response from a scatterer with a stable phase center, or from a scatterer with a significantly shifting phase center. A relatively well-resolved response may come from two interfering scatterers, both with stable phase centers, both with shifting phase

centers, or one with a shifting and the other with a stable phase center. There may be a group of poorly resolved responses. There could be isolated spurious responses or entire sets of spurious responses. It is necessary to analyze each response or each group of responses individually, and then to choose the appropriate processing algorithms.

Each response is analyzed by examining 18 evenly spaced image cuts through the response peak. In each cut, we first test whether the response is from an isolated single scatterer, in which case the amplitude of the transform must be nearly constant. If not, we next test whether the response is a composite response from two interfering scatterers. This is done by examining whether the amplitude and phase functions of the transform have the appropriate patterns. If this is not the case, we test whether some part of the response can be decomposed into two interfering scatterers.

We describe our specific implementation of these tests. The first test is that for a relatively isolated response from a single scatterer. We break this test into two checks. The first check is very simple: if the normalized half-power width and skewness of the response are within 6% and 12% of unity, respectively, we declare an isolated response from a single stable scatterer. However, a normalized half-power width outside this interval may correspond to an isolated response from a single shifting scatterer. The check for this situation is implied in the one-dimensional TSA: if the response is due to a single scatterer, the amplitude of the transform of an image cut through the scatterer will be essentially constant. This is because the backscattering strength of a scatterer does not change much over the small aspect angle sectors and the bandwidths used for imaging. The phase function of the transform is determined by the combination of the behavior of the phase center and any residual uncompensated target motion. If the motion compensation is good, the phase function describes the shifting of the phase center: a linear function in the case of a stable phase center or a curved function when the phase center shifts with aspect angle or frequency.

As a side remark, this test of amplitude constancy of the transform has an important application far exceeding its use in deciding whether a response comes from a single scatterer. It is used in fixed-range image cuts to analyze the motion behavior of scatterers in order to determine processing steps for motion compensation and image time selection. This is discussed in detail in Sections 2.3.2.3 and 2.3.3.

The second test, applied only when the response fails the first test, is for a composite response from two scatterers. In the absence of significant interference from surrounding scatterers, the amplitude function of the transform of such a response will show a good approximation of the interference

pattern from two ideal point scatterers. The phase function will also approximate that of two ideal scatterers, unless one or both of the actual scatterers have shifting phase centers. In the latter case, the phase function of the transform will be curved where it should be linear, but otherwise the amplitude and phase patterns will be those of two interfering scatterers. We will generally perform the measurements as if we were dealing with two fixed point scatterers, but depending on how well or poorly the amplitude/phase pattern of the transform approximates that of two ideal point scatterers, we will assign different uncertainties to the measurement. The basic algorithm could be extended to include the case of scatterers with shifting phase centers, but we have not done so.

The third test is applied only when the response fails the first two tests. We must try to choose a transform window such that the amplitude/phase pattern of the transform at least approximates that of two ideal point scatterers. When this can be done, we again assign measurement uncertainties that depend on the quality of the approximation of the pattern. These uncertainties will be larger than those from the second test. If too many responses fail all three tests, the radar has an inadequate resolution for the particular situation.

The only remaining point is consideration of spurious responses, which are responses appearing in positions other than those of the associated scatterers. Spurious responses can be divided into two classes. The first class, consisting of delayed returns and returns from moving parts, as shown in Figure 1.20, can be recognized by their location in the image. The second class consists of responses generated by shifting phase centers. Such responses can be readily recognized when they appear outside the target, if the clutter and noise levels are low enough, based on the response locations and phase curvatures [1, 5]. However, when the intrinsic phase curvature of the spurious responses is distorted by interference with strong target, clutter, or noise responses, these spurious responses are difficult to discriminate. They become stronger and more numerous as the carrier frequency is increased, because the sizes of the features become larger in comparison with the wavelength. For stationary targets, the best way to avoid a serious problem from spurious responses is to avoid operating at very high carrier frequencies. Indications are that X-band may be the highest practical carrier frequency in that respect. For moving targets, we should avoid imaging when the motion is erratic or three-dimensional.

The sequence of steps that we employ to extract scatterer positions and other characteristics is given by the flowchart in Figure 1.36. We locate the two-dimensional amplitude peaks and retain those stronger than an adaptively determined threshold (see Appendix H). We then discard

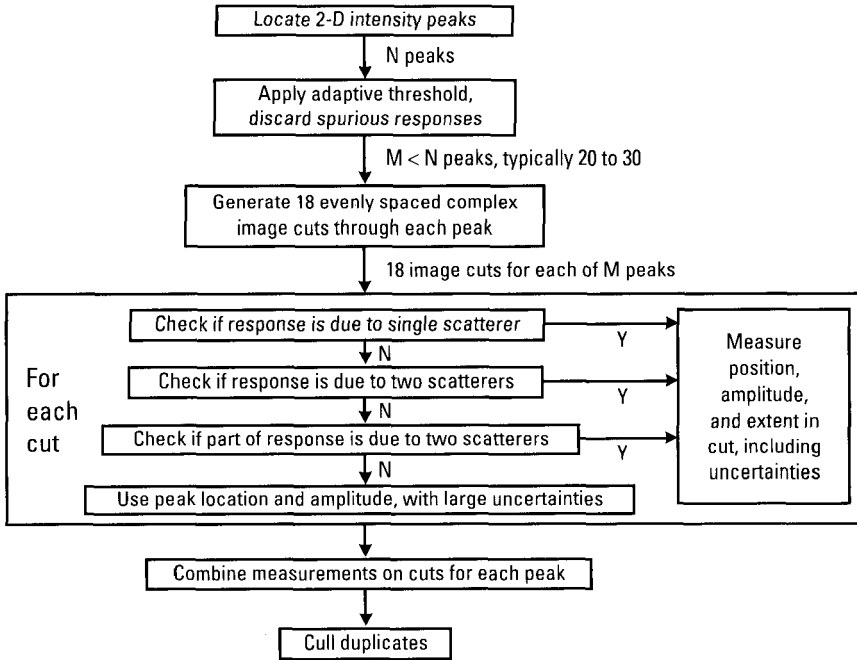


Figure 1.36 Feature extraction steps.

spurious responses. We analyze the remaining responses in the fashion described above, combining results of 18 image cuts through each response. As two interfering scatterers may generate two two-dimensional peaks, and our analysis of each peak may yield both scatterers, we cull duplicates in a final step. As well as position, we measure scatterer amplitude and extent, as discussed in Section 1.4.

1.3.3.2 Evaluating One-Dimensional Image Cuts

In most real data, the image responses we must analyze are not well resolved from other responses. We must define transform windows to exclude as much interference as possible, without excluding too much of the responses of interest. Including too much interference or excluding too much of a response of interest will cause the transform to deviate from a one-scatterer or two-scatterer pattern. Thus, choosing transform windows appropriately can be critical to measurement accuracy. We have found that it is best to place the boundaries of the transform window close to the point where the scatterer to be retained and the scatterer to be excluded have equal

amplitudes—the point at which dominance shifts. We show in Appendix B that for two interfering point scatterers this shift coincides with a phase inflection point, and that if the interference is generally destructive the shift occurs close to an amplitude minimum. The procedure that we have implemented is to examine the transforms of windows whose boundaries are phase inflection points and amplitude minima, and if a transform is a good approximation of a one-scatterer or two-scatterer pattern, to vary the window boundaries slightly in order to improve the approximation to the pattern.

Figure 1.37 shows an example of the distortion that can result from windowing out part of a response of interest. The figure shows a repeat of the image cut of Figure 1.30, with the right tail and sidelobes on both sides of the response excluded from the transform window. We note that although the shape of the transform amplitude has changed substantially, the separation of the amplitude minima (or phase jumps) has changed by just a few percent. Similarly, there is just a small change in the size of the phase jumps (or ratio of amplitude maximum to minima). These features are more robust to windowing than the overall amplitude shape, so we rely on them in our position measurements. The amplitude distortion is particularly strong near the edges of the transform. This is a general result of windowing; the first and last 5% (roughly) of the displayed transform amplitude function cannot be relied upon. One must often exclude more of a response of interest than we

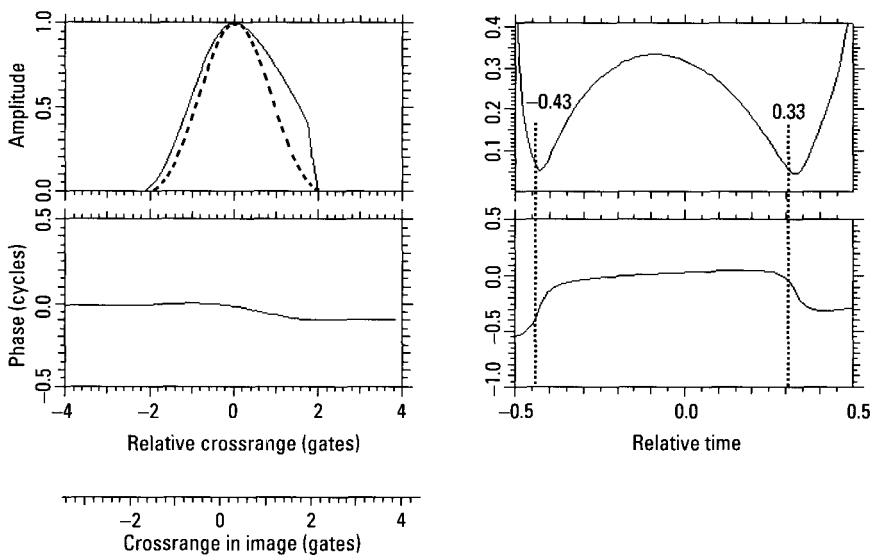


Figure 1.37 Windowed image of Figure 1.30.

have done in Figure 1.37, or place both transform boundaries to exclude parts of both responses of interest. Furthermore, the transform will be distorted by contributions from additional scatterers that could not be entirely excluded. The distortions of Figure 1.37 are relatively mild.

The phase jumps of Figure 1.37 are more widely separated than those of Figure 1.30, and will therefore yield more closely separated scatterer positions. This is consistent with the windowing of the image cut. The window cut out the right tail of the image response. We could also say that it cut out the tail of the rightmost of the two scatterers contributing to the response. This filtering shifts the centroid of that response to the left, and the transform effectively measures the separation of the scatterer centroids. This raises an important point. In practice, we apply such windowing to minimize the effects of a third scatterer on the transform. That third scatterer obscures the two scatterers of interest, making it difficult to ascertain how much of their returns are filtered out by the transform window, and by how much their centroids shift. We must ensure that we do not use a transform window that shifts the scatterer positions unacceptably.

Verifying that the shift is acceptably small requires comparing scatterer positions derived from several different windows. Once we have found the window whose transform is the best two-scatterer pattern, we vary the window boundaries by one-quarter of a resolution cell and remeasure the scatterer positions. If the window has removed only a small part of the two scatterers of interest, the centroid shift will be small, as will the differences between the scatterer positions measured for different window boundaries. We accept the scatterer positions measured in the original window only if they are reproduced in the new windows, within one-tenth of a resolution cell. This small difference is acceptable, as we assign an uncertainty of at least one-fifth of a resolution cell to each measured position.

Besides preventing mismeasurement due to interference from a third scatterer, this test also prevents us from erroneously utilizing a transform window that contains the return from one scatterer and a small part of a second, in which case the distorted pattern might be misinterpreted. Such a window could be problematic when the second scatterer is strong relative to the first. However, narrowing the boundaries of the window will usually produce an unacceptable two-scatterer pattern, and will otherwise result in unacceptably large position shifts.

1.3.3.2 Examples of TSA Analysis of Real Data

In this section, we demonstrate the measurement of scatterer positions in real data of a motion-compensated flying aircraft, with a range resolution of

about 0.3m and a crossrange resolution of about 1m. Figure 1.38 shows a fixed-range image cut through a two-dimensional intensity image response peak. The response of interest, near Relative Crossrange 0, is not fully resolved from another response near Relative Crossrange 2.4. The normalized half-power width of the response of interest is 1.42, far too high for a single scatterer. This response is likely composed of two scatterers, with its interpretation made complicated by the presence of the second response. As discussed above, we set transform window boundaries near phase inflection points and amplitude minima, then evaluate the degree to which each transform fits a two-scatterer pattern.

Figure 1.39 shows the best window position and the resulting transform. The figure also includes curves labeled “m” (for model), which are the result of extracting the two scatterer positions, amplitudes, and phases via the TSA, generating the corresponding image cut, applying the same windowing as done to the real data, and taking the Fourier transform. The modeled amplitude and phase are in excellent agreement with the real data. The amplitudes differ only near the edges of the transform, where deweighting amplifies contributions of the third scatterer to the real data. As discussed above, distortions are expected at these edges. The primary difference between the phases is a constant offset, which is insignificant. The TSA

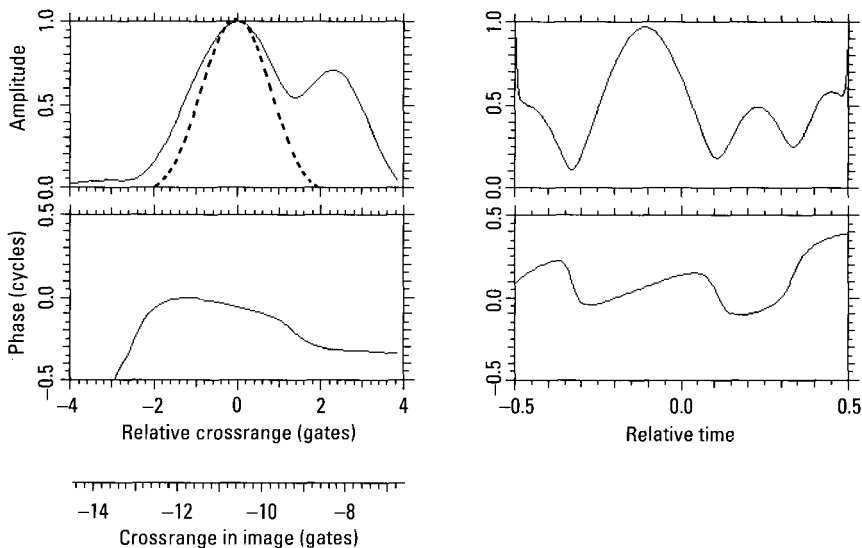


Figure 1.38 Fixed-range image cut through aircraft response.

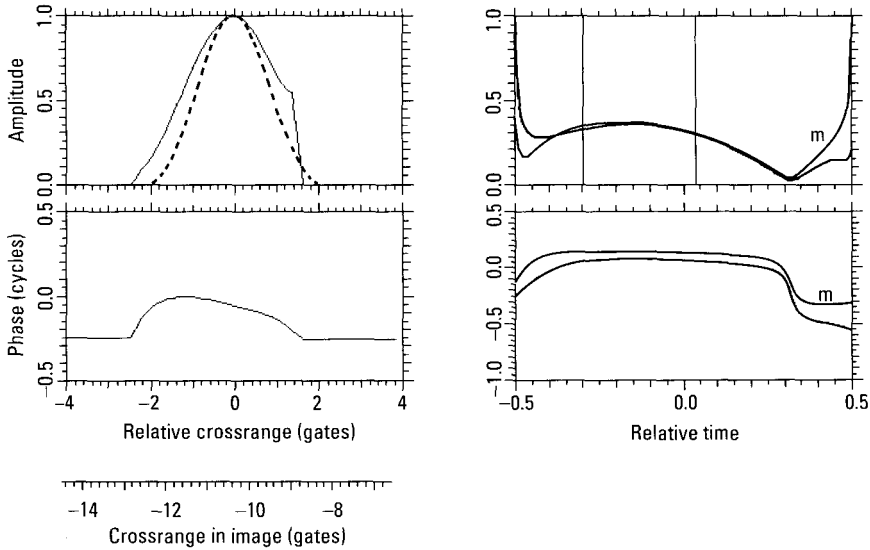


Figure 1.39 Transform and two-scatterer pattern for Figure 1.38.

gives scatterer positions (projected on the cut) of Crossrange -10.90 and Crossrange -9.70 . These positions must be, and were, verified by varying the window position as described above.

Figure 1.40 shows the fixed-crossrange cut through the response. The response amplitude has a slight bulge on its right-hand side, and the phase is curved. The response's normalized half-power width is 1.02 and its skewness (ratio of right to left half-power half-widths) is 1.11. The skewness is nearly large enough to allow measuring the positions of two interfering scatterers via the TSA, but not quite. A transform of the interval between the amplitude minima bounding the response does not allow the extraction of two scatterer positions. We use the peak location, Range 0.08, as an estimate of our projected scatterer location(s).

Figure 1.41 shows the results of automatically analyzing this response with the two-dimensional TSA. The figure is in the same format as Figures 1.34 and 1.35. The utility of the measurements in the diagonal cuts is evident. Each scatterer is clearly defined by a set of one-dimensional measurements closely fitting the required cosine dependence. We note again that the measurements near the two-dimensional intensity peak are not significant because they occur in cuts generally perpendicular to lines between the peak (the intersection point of the diagonal cuts) and the scatterers located away

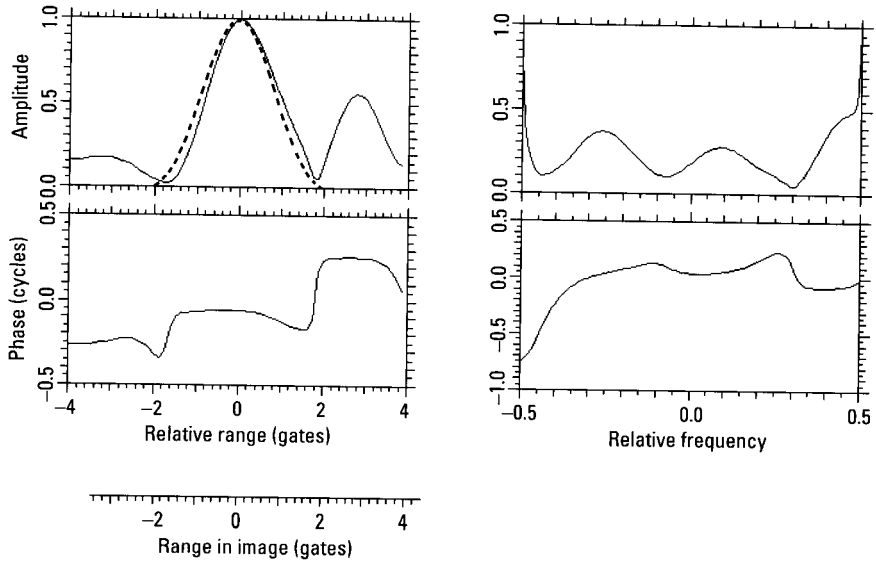


Figure 1.40 Fixed-crossrange cut through aircraft response.




Figure 1.41 Scatterers derived from diagonal cuts through aircraft response.

from the peak. We do not present any of the individual diagonal cuts because their analysis is not qualitatively different from those of the preceding cuts.

Figure 1.42 shows the results of measurements made in worse interference conditions. The automated measurement results define three scatterer positions, plus additional interference (indicated by the cluster of measurements near Crossrange -10.5 and Range -10.0). We shall discuss several of the cuts through this response that demonstrate different aspects of the measurement process.

Figure 1.43 shows the fixed-range image cut for Figure 1.42. The response has a normalized half-power width of 1.29 and a strongly curved phase function. This could correspond to a single scatterer with a shifting phase center, to two interfering scatterers with fixed phase centers, to one fixed and one shifting scatterer, or to two shifting scatterers. In the first case, the transform would have a nearly constant amplitude and a strongly quadratic phase. In the second case, the transform would be a two-scatterer interference pattern. In the third and fourth cases, the transform would resemble a two-scatterer interference pattern, but have a more strongly curved phase function when the amplitude is near its maximum. Figure 1.44 shows the




Figure 1.42 Scatterer positions derived for worse interference.

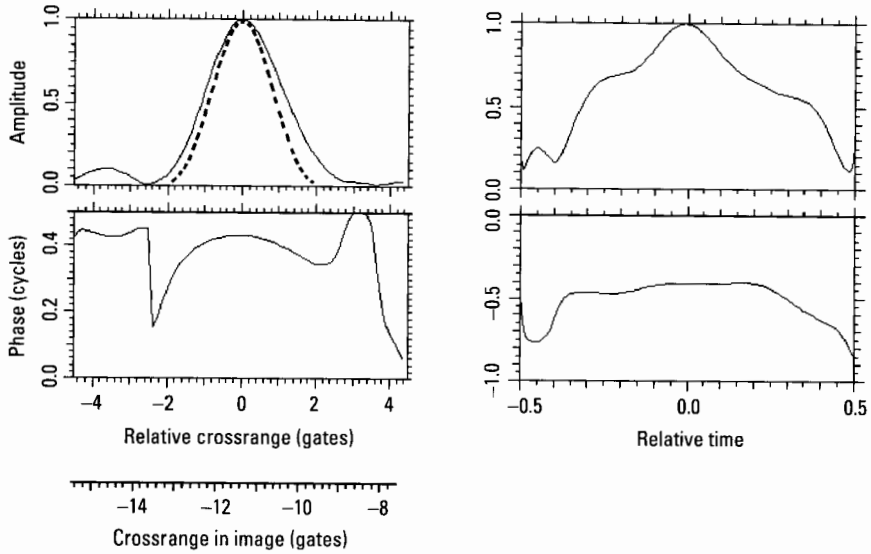


Figure 1.43 Fixed-range image cut of Figure 1.42.

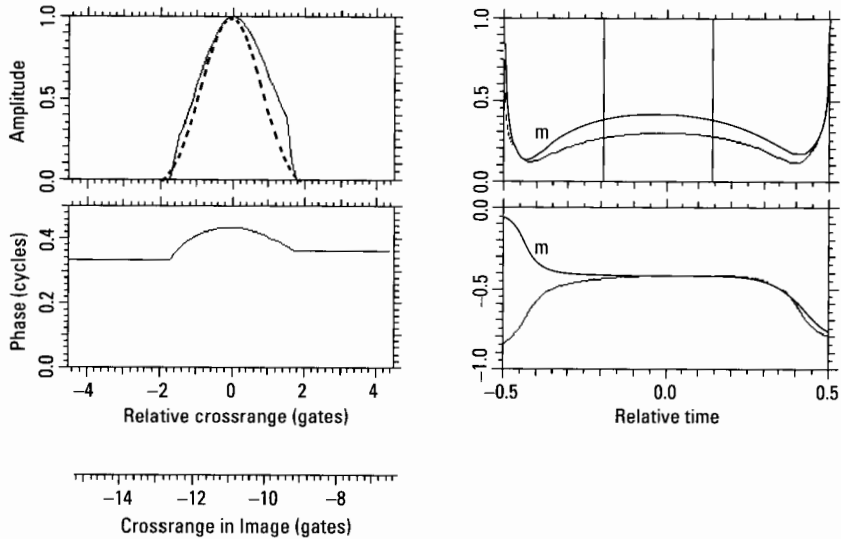


Figure 1.44 TSA analysis of Figure 1.43.

TSA analysis of Figure 1.43, with transform boundaries shifted slightly inward from the phase inflection points. The shapes of the transform data

amplitude and the TSA model amplitude are similar. Their relative scaling is caused by their different normalizations, determined in this case by their values at the very edges of the transform, and is insignificant. The apparent difference in the phase functions is due primarily to a full cycle difference in the phase unwrapping process. Accounting for this, the phases are also in good agreement. The amplitude modulation and the small phase curvature at the time of the amplitude maximum show that the response is composed of two scatterers with fixed phase centers.

Figure 1.45 shows the fixed-crossrange cut from Figure 1.42. The deformation on the right side of the amplitude indicates interference among two or three scatterers. The bulge in the left tail of the centered response indicates yet another scatterer. All this interference necessitates shifting the transform interval from that between phase inflection points. The appropriate choice of interval and the resulting two-scatterer pattern are shown in Figure 1.46. The measured and calculated data differ only in the phase unwrapping at the right edge of the transform.

Figure 1.47 shows the image cut at 120° in Figure 1.42 (fixed range being 0° and fixed crossrange being 90°). This is a case of interference worse than any yet discussed. It is clear from the response amplitude, top left curve, that at least three strong scatterers interfere. We must attempt to define a transform interval with contributions from the right and center scatterers,

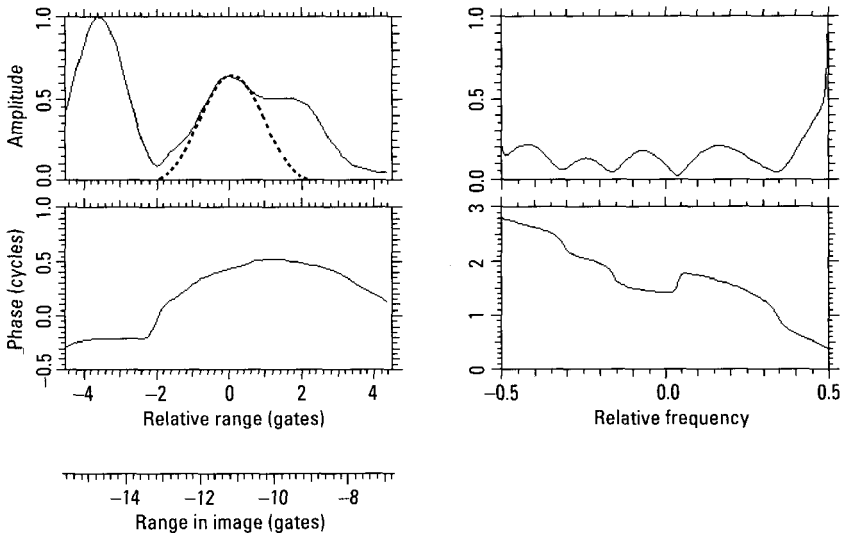


Figure 1.45 Fixed-crossrange image cut of Figure 1.42.

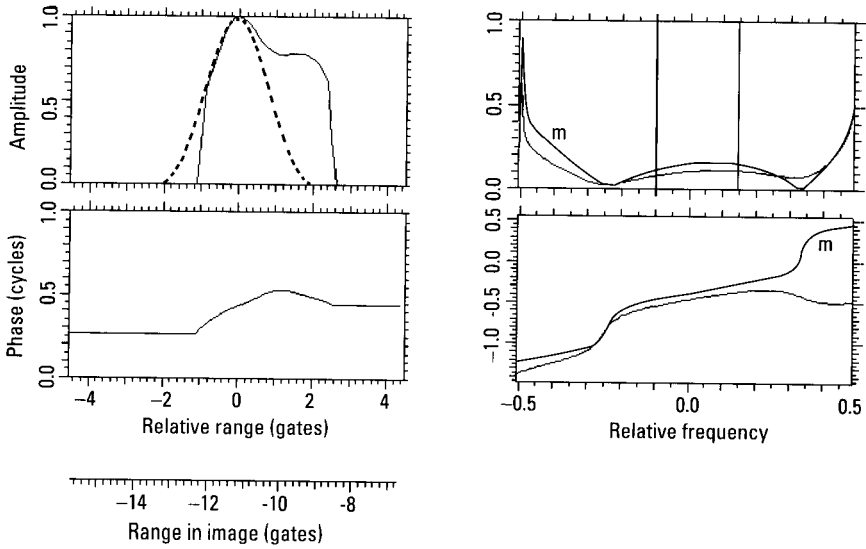


Figure 1.46 TSA analysis of Figure 1.45.




Figure 1.47 Image cut at 120° in Figure 1.42.

and another interval with contributions from the left and center scatterers. The best choices of these intervals (with respect to the quality of the TSA

patterns) are shown in Figures 1.48 and 1.49, respectively. The intervals were chosen by beginning with boundaries at the phase inflection points,

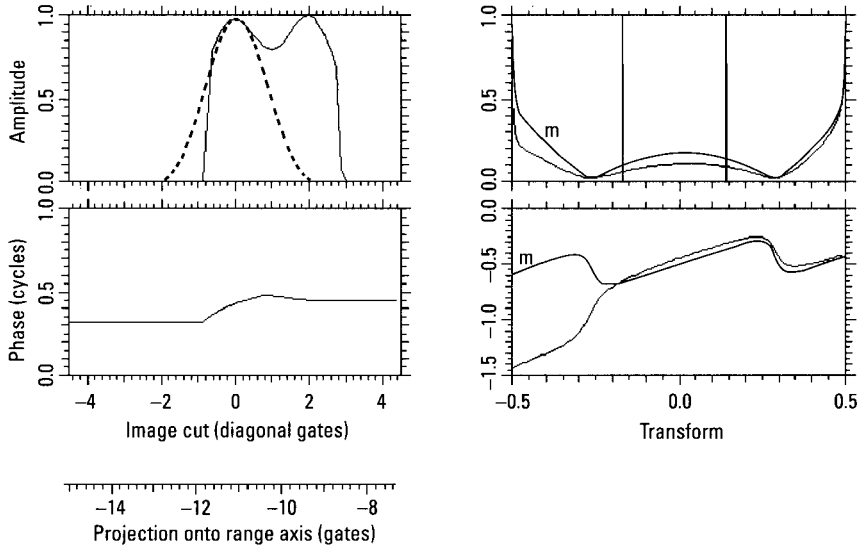


Figure 1.48 TSA analysis of the right side of the response of Figure 1.47.




Figure 1.49 TSA analysis of the left side of the response of Figure 1.47.

then varying the boundaries to produce the best two-scatterer pattern. In both figures, the differences between the measured and calculated patterns are minor, with the visual differences primarily due to different phase unwrappings. The center scatterer position should be taken as the average of the positions measured in the two figures. Unsurprisingly for such a difficult case, the automated measurements found only two locations, with some error.

Having summarized the one- and two-dimensional TSA, for ideal point scatterers and applied to real data, some clarifying remarks are in order. The one-dimensional TSA amounts to the implementation of the inherent resolution performance of radar in one dimension, range or Doppler (cross-range), on two ideal point scatterers. The need for implementing this basic algorithm via pattern recognition rather than mathematics stems from the fact that complicated man-made targets cannot be mathematically modeled with sufficient realism. For the same reason, the extension to the two-dimensional TSA must also be based on pattern interpretation. However, the method we have implemented (using 18 one-dimensional image cuts) is not basic but a choice of convenience. If there is a better way that works with real data, we did not see it. It would be very simple to apply the TSA to real data if the problem were merely to resolve two scatterers. In practice, we must contend with various forms of interference from other scatterers. This presents problems that we have solved in the manner discussed above, primarily by choosing the “correct” transform window. There is nothing basic about our particular implementation, and modifications and improvements are possible. Because of the complexity of the two-dimensional TSA as implemented via pattern recognition, it is not suitable for manual analysis. For this reason we developed a version that allows determining the two-dimensional scatterer positions from only fixed-range and fixed-crossrange image cuts, as discussed in Appendix F. This version of the TSA is more difficult to understand but less cumbersome to use with manual processing.

1.3.4 Section Summary

The one-dimensional TSA, which implements the inherent resolution capability of radar, compares a Fourier transform of a response to the signal generated by two interfering fixed point scatterers, which is characterized by a sinusoidally varying squared amplitude (power), with phase jumps at the times of amplitude minima and linear phase at the times of amplitude maxima.

We analyze each image response by examining 18 evenly spaced cuts through the response peak. In each cut, we first test whether the response is from an isolated single scatterer, in which case the amplitude of the transform must be nearly constant. If not, we use the one-dimensional TSA to test whether the response is a composite response from two interfering scatterers. If this is not the case, we test whether some part of the response can be decomposed into two interfering scatterers.

We assign each measurement an uncertainty based on the degree of agreement of the measurement and the one-scatterer or two-scatterer patterns.

1.4 Special Measurements of Potential Use

The measurement of scatterer positions, together with the determination of special features such as wing-mounted versus fuselage-integrated engines, provides the basis for target identification. One can augment target identification with other types of measurement, but at the time of this writing the questions of how useful and how necessary these measurements are was not answered. We include a brief discussion for the sake of completeness, since such measurements might become more useful in the future.

1.4.1 Feature Extent

As is shown in [1], it is possible to derive the effective width of a scatterer in range and crossrange from the phase function of an image response. This is evidently a useful capability for target identification. The problem is that the measurement is simple only when the intensity response of a feature is well resolved from the intensity responses from neighboring features. In other words, it is simple only if the separation of the response peaks is at least $2/B$ and $2/T$ rather than $1/B$ and $1/T$. It appears possible to extend the measurement to the case where the neighboring features are separated by only $1/B$ and $1/T$, but we have not developed the appropriate algorithms. We will give a brief summary of the measurements for the case where the intensity responses are well resolved.

Figure 1.50 shows a fixed-range cut through a response of a ground vehicle. The response has a normalized half-power width of 1.22 and a curved phase. *If the response is due to an extended scatterer with a shifting phase center, its transform will have an essentially constant amplitude and a curved phase function.* The instantaneous phase slope of the transform then

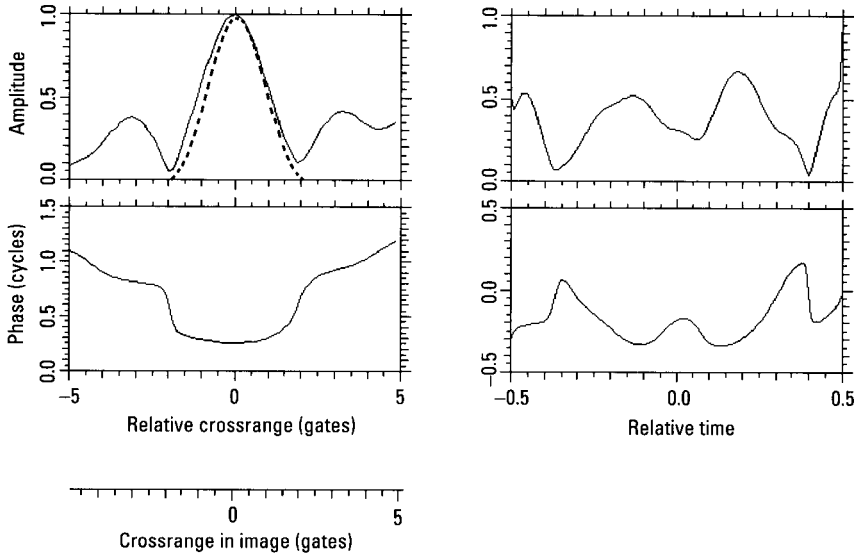


Figure 1.50 Fixed-range image cut through ground vehicle response.

gives the crossrange of the scatterer's phase center at the time of the measurement. Figure 1.51 shows a transform over the interval between phase

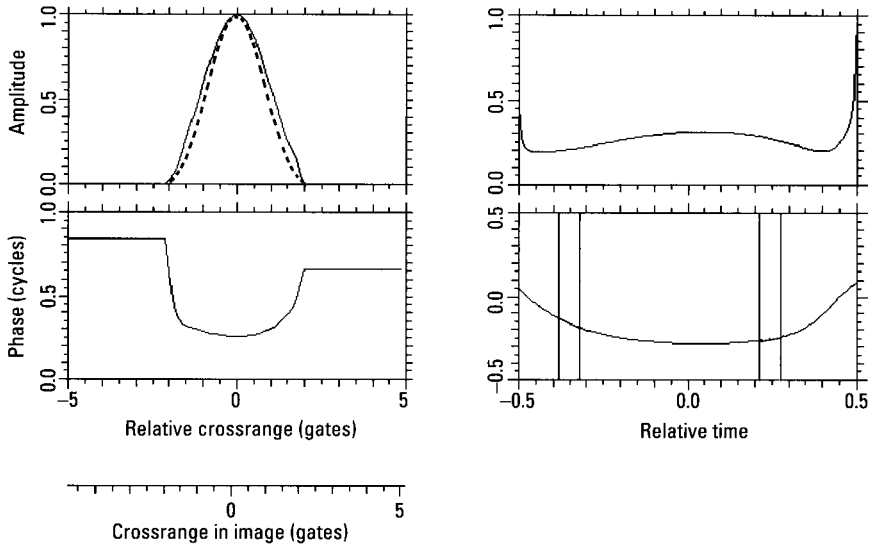


Figure 1.51 Phase slope measurement for extent determination.

inflection points bounding the response. The transform amplitude has just a small slow modulation, perhaps due to distortions induced by deweighting or perhaps due to interference between two scatterers, one much weaker than the other. Even if the modulation is due to such interference, it is small enough that the interference will introduce only negligible curvature into the transform phase, except near times of amplitude minima, which imply phase jumps.

In order to measure the feature's crossrange extent, we measure the most positive transform phase slope and the most negative transform phase slope, excluding regions of the transform that are near amplitude minima. These exclusions ensure that we measure the phase-center shift of one feature independently of interference effects from another. The difference between these two phase slopes is the amount by which the phase center has shifted, the effective extent of the feature. The intervals used for the phase slope measurements are indicated by the vertical lines in the transform of Figure 1.51. In our choice of these intervals we have excluded those parts of the transform that occur when the phase curvature changes rapidly, such as around an amplitude minimum. The regions with rapid curvature variation are more evident in Figure 1.52, which shows the scaled second derivative of the transform phase of Figure 1.51. The difference between the two phase slopes of Figure 1.51 is 1.3 crossrange resolution cells. The crossrange resolution for this data is about 0.3m, giving a feature extent of about 0.4m.

The high crossrange resolution needed to measure the width of a typical target feature is oftentimes unavailable, because dwell time is insufficient or because target motion is too severe to allow motion compensating a large enough angular interval. If crossrange resolution is about 1m or worse,




Figure 1.52 Scaled second derivative of transform phase in Figure 1.51.

because the typical features on man-made targets are then significantly smaller than a resolution cell, crossrange extent measurements are generally impossible. Similarly, range extent measurements are possible only when the range resolution cell is smaller than or comparable to typical feature extents. Since high range resolution is more generally available than high crossrange resolution, range extents can be measured more often than crossrange extents.

1.4.2 Cross Section of Features

Whenever we measure the range and crossrange positions of a target feature, we automatically also obtain the cross section of the feature. The question thus arises whether we might be able to utilize feature cross section as an additional input to target identification. We already do this in a way when we restrict our attention to the limited set of “observable” target features, utilizing only the set of 20 or 30 observable responses. However, for these responses we could add cross section to position and extent.

The problem with feature cross section is that a specific parameter of a feature is usable for target identification only if it can be predicted for each of the candidate targets, so that it may be included in the comparison database. However, feature cross section is difficult to predict over different aspect angle sectors with better accuracy than merely predicting whether or not the feature is observable. The features on actual man-made targets have designs that do not at all approximate those idealized features for which the cross sections can be readily calculated as a function of aspect angle. We will typically not have sufficient information on the designs of real features on real targets to permit predicting their cross sections to reasonable accuracy. Hence, although we feel that there should be some utility to feature cross section, we have so far not found a good way to exploit it.

1.4.3 Polarization Diversity

Much has been written about polarization diversity and its use in target identification. There is no question that fully polarimetric measurements provide more information about a target than measurements on only a single polarization. We have found this to be true for real data. However, two considerations are important. One, is the expense of designing a radar system with such a capability justified? In other words, does reliable target identification require polarization diversity? Two, if we utilize polarization

diversity, it must be from the point of view of real targets rather than simulated targets.

Although we believe that reliable target identification can be obtained with a single polarization, until an operational system has been fielded, one cannot be sure. Our concern here can be only with the second point. Most of the past work on polarization diversity is based on idealized target shapes, or so-called “primitives,” with the main emphasis on the utilization of even-bounce and odd-bounce returns to discriminate between different shapes. As we have found, such considerations have little validity for the features on real targets. For example, we have found that many dominant scatterers on ground vehicles have comparable backscattering at RL polarization and at the stronger of RR and LL, with significantly different strengths at RR and LL.

As with data at a single polarization, to make best use of polarization diversity one must exploit the image phase. Simple ratios of intensity returns at different polarizations discard much of the available information. Furthermore, one should not just make measurements on the four separate linear or circular polarizations and then combine the results. Some feature characteristics are best revealed by less conventional approaches. For example, the complicated features which dominate the backscattering from man-made targets have phase centers that shift with changes in the polarization. Hence, the image locations of their responses shift with polarization. Just as using diagonal image cuts improves the measurement of response locations in a single image, using images at slightly different polarizations improves measurement of the shift in location with polarization. We have measured this shift by selecting a strong response at, say, VV polarization and, keeping the transmit and receive polarizations equal and linear at all times, shifting them in steps to HH, and on to $-V-V$, measuring the response location at each step.

Although the shift in location with polarization is measurable and provides discrimination among features on man-made targets, predicting it is problematic for most target features, limiting its utility. We are probably better served by utilizing feature characteristics that are more easily predicted. To this end, we note that the scattering matrix that describes the polarized return from a feature may be expressed in terms of linear polarizations, circular polarizations, or an eigendecomposition [6]. The eigendecomposition gives six independent parameters which can be related to the physical structure of the scatterer. We have found one of these six parameters, the helicity, which measures the feature symmetry in the plane perpendicular to the radar line of sight, to be useful in characterizing dominant scatterers on man-made targets.

1.4.4 Dispersive Backscattering

With simple backscattering models of target features, the return from a single feature has an essentially constant backscattering amplitude over typical radar bandwidths. Real targets have features for which this is not true. The backscattering may be strong over, say, the lower half of the signal band, and very weak over the upper half. This behavior is potentially useful for target identification, because it allows relating an image response to a specific feature on the target. The utility of such a capability would be great, even if this occurs only for a single feature on the target. As with the cross section of responses, the difficulty in utilizing dispersive behavior lies with its prediction. Complicated backscattering models can predict the dispersive behavior of some shapes and substances found on man-made targets, but the composition of most dispersive target features is not readily available information. Dispersive features are likely to be electronic devices, and these play an important part at least in the identification process for military aircraft. As is the case for all the topics of this section, the utility of recognizing dispersive features is a point to be further explored. The discussions of the main part of this book are reserved for methods that we have already tested extensively.

1.5 Identification Procedure

Although this is a textbook on target identification, we emphasize the methods used to extract information from the radar return. The procedures of how to perform target identification when sufficient information has been extracted from an image have for some time been much better understood than how to extract the necessary information from an image. In this section, we summarize the other components of the identification procedure we have adopted.

Our procedure utilizes two types of features recognizable and generic. Recognizable aircraft features include the capability to carry wingtip ordnance, the location and number of engine intakes and exhausts, and a minimum target length. Recognizable ground vehicle features include the presence of a turret or gun and the dimensions of stationary targets. Recognizable ship features include length, width, the number of deck levels, and the number of superstructure blocks. Generic features are the positions and extents of scatterers composing the strong responses of the target. *The identification procedure relies primarily on comparing a template of predicted scatterer information for each candidate target with the generic information extracted from the individual responses of the measured target.* Since

knowledge of the measured target's range and crossrange location, orientation, and (for moving targets) crossrange resolution is only approximate, each template is allowed to deform in a probabilistic manner. A standard Bayes classifier [7, 8] is used to perform the target identification, comparing the chi-square probability for the template match of each candidate target against that of all others in the comparison database. Probabilities for the recognizable features are also calculated, and combined in Bayesian fashion with those from the template matching, to yield composite probabilities that are used for the identification.

1.5.1 Recognizable Target Features

Many man-made targets contain easily recognizable features. For instance, *inverse synthetic aperture radar* (ISAR) images of approaching fighter aircraft contain columns of responses at the crossrange positions of engine intakes, generated by multiple-bounce scattering in the intake. Receding fighters have similar columns at the crossrange positions of the exhausts, generated in a similar manner. Aircraft with wing-mounted engines have shorter columns of responses, offset from the fuselage. These columns of responses can be used to count and locate intakes and exhausts. Rapidly moving parts, be they engine blades or optical scanners, can be recognized by sets of responses at nearly fixed ranges. The turret of a stationary or benignly moving ground vehicle can be recognized by its shadow within the vehicle, and the vehicle's length and width can be measured to high accuracy.

The predictive database for each candidate must contain its recognizable features, with a likelihood for each feature being observed when the candidate is imaged. Because target features are difficult to recognize at low crossrange resolution, we must not penalize any candidate for possessing an unobserved feature. The predictive likelihoods for each observed feature are combined to generate Bayesian probabilities for the feature. We note that the database features must be those recognizable by radar. If a target has a fender that extends six inches beyond a large cavity, the target dimensions must be given by the position of the cavity, because the tip of the fender will usually not be observed. Even if it should be observed under benign circumstances, the weakness of its response will indicate what feature is observed.

The recognizable features have an important use beyond that already described; they may constrain the allowable template deformations in the match to the extracted generic features. Constraining an aircraft's intake position (using the crossrange gate of the duct returns) can greatly increase the discrimination capability of the template match. Conversely, an error is

disastrous. We can only apply constraints for those features recognizable beyond a doubt.

1.5.2 Deformable Template Match

The deformable template matching algorithm is based on a Bayesian model. In this model, prior information for each target candidate is available in the form of its predicted features. These consist of the dominant scatterer positions and position uncertainties, as well as the characteristics discussed in Section 1.4. Because those characteristics are only sometimes measurable, we often refer to the deformable template match as the positional match.

To allow matching of the predicted features to the extracted features, the predicted scatterer locations must be mapped into the range and crossrange coordinates of the SAR/ISAR image. The parameters necessary to perform this mapping are the range and crossrange translational position, the target orientation, and the crossrange resolution. Since these parameters are only approximately known, we generate a prototype template based on our “best guess” of the parameters. This prototype template is allowed to deform through adjustments to these parameters, as illustrated in Figure 1.53.

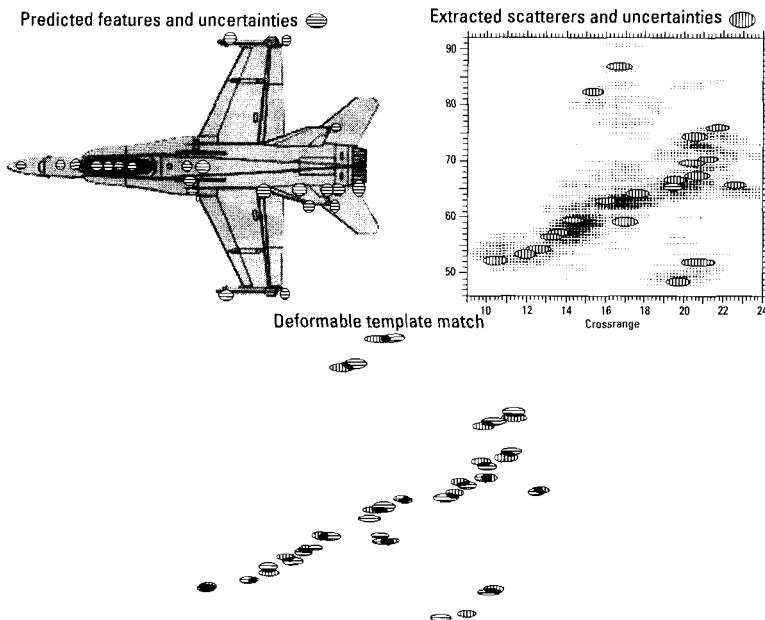


Figure 1.53 Positional match.

Each extracted scatterer is represented in the ISAR image at the upper right of the figure as a two-dimensional distribution, with the distribution widths given by the estimated positional uncertainties. Each predicted scatterer is similarly represented in the plan view at the upper left of the figure. The predicted scatterer distributions form a template that is deformed by translation, rotation, and stretching in crossrange to yield the maximal overlap with the extracted scatterer distributions.

We calculate a likelihood function for each deformed template that measures the similarity between the template and the measured positions. The likelihood function used is the average variance-normalized two-dimensional distance between individual measured and predicted scatterers (with a cap on the distance for any scatterer, so that no single scatterer drives the matching). The variance normalization underscores the necessity to estimate uncertainty in position measurements. We effectively weight the more certain measurements more heavily.

We have employed two schemes to calculate the likelihood function. The first performs a direct association of individual predicted scatterers with the corresponding measured scatterers, beginning with the best matched pair, then finding the next-best matched of the remainder, and continuing through the scatterers. This is the intuitive approach depicted in Figure 1.53, but it is not computationally efficient. The second scheme generates the likelihood function on a fixed grid in range and crossrange, computing the likelihood function for each grid element based on its variance-normalized distance to the nearest measured scatterer. As each predicted scatterer is mapped into a grid cell, the predicted scatterer is assigned a likelihood and a measured scatterer identity. The average likelihood from all the predicted features (with a normalization to account for the number of predicted and measured features [9]) is used as the likelihood function. This scheme creates an effective potential over the range-crossrange plane in which the prototype template is allowed to deform until its minimum energy configuration is achieved.

We also calculate a probability for each deformation, based on tracking information and recognizable image features (such as the location of an engine duct inlet). The tracking information constrains the allowed template orientations, and the recognizable features constrain the template centering. These probabilities limit the variability in the template shapes, in effect strongly penalizing unlikely template deformations.

We combine the likelihood function describing the goodness of the match between the predicted and measured feature values for a given deformation with the constraint-based probabilities for the deformation, to obtain

an overall probability for the deformation. We then find that deformation which maximizes the probability for each candidate target, and compare the candidates via the Bayes rule.

1.5.3 The Identification Process

The target is identified as the candidate with the maximum composite Bayes probability if that probability is above a predetermined threshold, and if the absolute agreement between measurements and predictions is sufficient. The threshold is set as a tradeoff between declaration probability and probability of misidentification, and must be determined empirically. The absolute agreement is determined by a chi-square comparison of the measurements and the predictions, and must be included in case of observation of a target not in the candidate database.

The Bayesian approach allows the specification of a priori probabilities for each candidate in the database. This permits the fusion of radar information with that from intelligence or other sensors. If such information can be used reliably to restrict the candidate database, identification will occur more often and more rapidly. However, identification on the basis of radar data alone has shown promise of meeting operational performance requirements.

References

- [1] Rihaczek, A. W., and S. J. Hershkowitz, *Radar Resolution and Complex-Image Analysis*, Norwood, MA: Artech House, 1996.
- [2] Walker, J. L., "Range-Doppler Imaging of Rotating Targets," *IEEE Trans. on Aerospace and Electronic Systems*, Vol. AES-16, No. 1, 1980, p. 23.
- [3] Woodward, P. M., *Probability and Information Theory, With Applications to Radar*, London: Pergamon, 1953.
- [4] Rihaczek, A. W., "Radar Resolution of Ideal Point Scatterers," *IEEE Trans. on Aerospace and Electronic Systems*, Vol. 32, No. 2, April 1996, pp. 842-845.
- [5] Rihaczek, A. W., and S. J. Hershkowitz, "Man-Made Target Backscattering Behavior: Applicability of Conventional Radar Resolution Theory," *IEEE Trans. on Aerospace and Electronic Systems*, Vol. 32, No. 2, April 1996, pp. 809-824.
- [6] Huynen, J. R., "Phenomenological Theory of Radar Targets," *Electromagnetic Scattering*, New York: Academic, 1978, pp. 653-712.

- [7] Duda, R. O., and P. E. Hart, *Pattern Classification and Scene Analysis*, New York: Wiley, 1973.
- [8] Winkler, R. L., *An Introduction to Bayesian Inference and Decision*, New York: Holt, Rinehart, and Winston, 1972.
- [9] Hauss, B. I., et al., "Airborne Target Identification from Low Crossrange Resolution ISAR Imagery," *SPIE Proceedings*, Vol. 2845, August 1996, pp. 120 -132.

2

Topics Related to Target Identification

Some common concepts in the field of target identification are based on the conventional resolution theory, which is valid in some cases and invalid, particularly when man-made targets are highly resolved, in others. Even though radar is a coherent system, one finds that much thinking regarding radar imaging and image interpretation has been inspired by noncoherent optics. Before treating the main topic of target identification, we will update the old concepts as necessary.

2.1 Ambiguity Function, Resolution, and Superresolution

Applying the ambiguity function in situations that require the resolution of a target into small parts creates problems. The ambiguity function does not fully describe close-scatterer resolution performance, because it is readily understood only when it is viewed as an intensity response. Intensity responses, however, do not provide the inherent resolution performance of radar. This section updates the conventional concepts of resolution performance.

2.1.1 Use of the Ambiguity Function

Since the time Woodward introduced the ambiguity function [1], it has been utilized in two ways. One use has been to determine the range/Doppler distribution of the interference generated by a point scatterer in range/Doppler resolution cells outside the one occupied by the scatterer. The practically valid assumption here is that the return from a particular waveform is received by a correlation processor. This type of interference is introduced by the radar waveform and is due to the relation between signal and spectrum as Fourier transforms of each other, and hence is basic to Fourier transform processing [2]. This specific use of the ambiguity function as a tool for waveform design to avoid mutual interference between widely separated targets and clutter remains fully valid today, because such interference must be analyzed via intensities.

The second use has been to determine close-target resolution performance from the shape of the central response of the ambiguity function. The resolution cell was defined as the area of a horizontal cut through the central spike of the ambiguity function, at about the half-power width of the response. Two targets were deemed resolved if they fell in different resolution cells. Since the necessary weighting for range and Doppler sidelobe suppression widens the central spike of the ambiguity function, it was concluded that sidelobe suppression is achieved at the cost of a loss in resolution performance. Theoretically, the conclusions that relate to the central spike of the ambiguity function were never fully correct, since in a coherent system resolution must be defined in terms of the complex receiver response. However, in view of the limited capabilities of early processors, the conclusions were practically meaningful. They lost their simple meaning when high-capability digital processors became available, because these permit the implementation of theoretically optimum processing algorithms.

Resolution performance thus was derived from the envelope of the ambiguity function, or specifically from the central spike after removal of the phase by envelope detection. This corresponds to a treatment as if radar produced an output equivalent to that of a noncoherent optical system (no phase information in the processor output). As already pointed out, the purpose of resolution is not merely to count the number of targets or scatterers, but to make at least reasonably accurate range and Doppler measurements. This becomes particularly important when the "targets" are individual features on a highly resolved extended object, where the purpose of range/Doppler imaging typically will be the identification of the object. When this practical measurement requirement is included, the resolution

performance on the basis of the envelope of the ambiguity function of a waveform with bandwidth B and duration T is about $2/B$ in delay and $2/T$ in Doppler, rather than the commonly assumed $1/B$ and $1/T$ [3]. In the two-dimensional range/Doppler plane, the resolution cell is indeed defined by the ellipse obtained by cutting through the central spike of the envelope of the ambiguity function, but in the absence of the widening caused by weighting for sidelobe suppression. Most important, however, in order *to achieve the corresponding resolution performance, we must utilize the complex ambiguity function*, or point-target response. For real targets this requires the analysis of the complex processor output, not just the intensity output. In turn, this requires use of the methods of complex-image analysis.

It is interesting to note that during the 1960s efforts were pursued to design waveforms in accordance with specified ambiguity functions. This problem appeared solvable for the complex ambiguity function but turned out not to be solvable for the envelope of the ambiguity function, which at the time was thought to be the practically meaningful specification. Thus, these attempts were considered a failure. Now we recognize that the interesting quantity is really the complex ambiguity function. Of course, even if a waveform can be designed in accordance with a specified ambiguity function, this is hardly of practical interest because the resulting waveform is unlikely to be implementable in a practical system. The ambiguity function still has great design value, but in the area of clutter suppression, be it true clutter or the self-clutter generated by the target to be identified.

2.1.2 One-Dimensional Resolution

Let us consider two point scatterers, one stationary and the other moving with constant Doppler. It is trivial to calculate the combined return from the two scatterers for a cw signal. The phasor diagram of Figure 2.1 shows the resultant vector from two scatterers at different Dopplers. In the figure, one scatterer has amplitude 1 and the second has a smaller amplitude a . The angle θ between the two varies linearly with time t as $\theta = 2\pi ft + \theta_0$, where f is the Doppler difference of the scatterers. The magnitude of the resultant is given by

$$A(\theta) = \sqrt{(1 + a \cos \theta)^2 + (a \sin \theta)^2} \quad (2.1)$$

and the phase of the resultant is given by




Figure 2.1 Phasor diagram for two scatterers at different Dopplers.

$$\phi(\theta) = \arctan[a \sin \theta / (1 + a \cos \theta)] \quad (2.2)$$

The maximum phase is

$$\phi_{\max} = \arcsin(a) \quad (2.3)$$

The “phase jump” in the composite signal has size $2\phi_{\max}$ and its duration is a fraction $(\pi/2 - \phi_{\max})/(\pi/2 + \phi_{\max})$ of the signal period. For illustration purposes we assume a Doppler of 3 Hz for the second scatterer, and an amplitude of 80% of that of the first scatterer. The return signal over one second is shown in Figure 2.2, the amplitude function on top and the phase function at the bottom. Since a constant Doppler difference implies that the differential phase changes linearly with time, the phase difference between the two scatterers goes periodically through 0° and 180° , which is responsible for the amplitude and phase modulation patterns. In practice, we cannot select the observation time in accordance with the unpredictable phasing of the two scatterers. Thus, if we use some small observation window on the




Figure 2.2 Return from two point scatterers, Doppler separation 3 Hz and amplitude ratio 0.8.

amplitude/phase functions of Figure 2.2, we cannot choose the window to observe a specific section of this general pattern. Whatever the accidental position of the observation window, it is clear that if we observe at least one full modulation cycle, we can calculate the parameters of the two scatterers. Since these parameters define the amplitude and phase pattern, but the pattern is periodic, we need observe no more than one full period to obtain all the information. The practical implication is that with the observation of a full modulation cycle we can calculate the parameters of the two scatterers, which means that the two scatterers have been resolved. *This is the valid definition of resolution*, not the half-power width of the envelope of the central spike of the ambiguity function. Most significantly, the duration of the modulation cycle equals the reciprocal of the Doppler separation of the scatterers. More details on the calculations are given in Appendix A.

As already stated, in a practical operation the observation window will fall onto an unpredictable part of the general return of Figure 2.2. It can be

centered on an amplitude maximum of the pattern of the two scatterers, on an amplitude minimum, or anywhere between a maximum and a minimum. This depends on the changing phase relation between the scatterers (moving at different Dopplers), and this phase relation is unpredictable. In practice we thus cannot choose the timing of the observation window. In addition, the observation window could be narrower than the interference cycle for the two scatterers, it could be about as wide, or it could be wider. For a given absolute width of the observation window (the transmitted signal duration), this relation depends on the Doppler difference between the two scatterers, which we also cannot choose. In practice, then, an observation window of arbitrary width relative to the interference cycle will be centered at an arbitrary point of the interference cycle.

Suppose that we accidentally observe the return over an interval centered on the time of an amplitude maximum, when the phase difference between the scatterers is zero (the two vectors are aligned at the time of the window center, which we define as constructive interference). If the observation interval is as short as indicated in the figure by the crosshairs, we evidently cannot measure the periodically repeated part of the amplitude/phase pattern, and hence the two scatterers cannot be resolved. If we extend the observation window to the first amplitude minima (or phase jumps) to the left and right, we can measure a full period of the modulation pattern, which implies that the scatterers are resolved. Since the amplitude minima are separated by $1/\Delta v$, where Δv is the Doppler separation of the scatterers, the achievable Doppler resolution is the reciprocal of the window length, or $1/T$.

Let us go to the other extreme where the observation window is centered on an amplitude minimum, which means that the phase difference between the two scatterers is 180° at the center of the window (defined as destructive interference). If the observation window is as narrow as the one indicated in Figure 2.2, not only are the scatterers unresolvable with this timing of the observation, but they probably cannot even be detected in a noise background. Because of destructive interference their combined return may be too low for detection, and with such a short window we effectively see a single scatterer, if we can detect it at all. As we shift such a short observation window in time, we merely observe a single target or scatterer with a fluctuating cross section.

If the observation window is still centered on an amplitude minimum (destructive interference) but is widened so that it extends from one amplitude maximum to the next, we observe the full amplitude/phase modulation cycle. Hence, the two scatterers can be resolved. In principle, resolution of

two point scatterers that are out of phase thus does not require a wider observation window, meaning that inherent resolution performance does not depend on the phasing of the scatterers. In practice, one must suppress additional interfering responses and noise via the processing sequence discussed in Section 1.3: apply weighting to the signal, Fourier transform, Doppler filter, inverse Fourier transform, remove the applied weighting. After this processing, the shape of the pattern will be less reliable near the fringes of the observation window, and this will be more serious when the observation window is centered on an amplitude minimum rather than an amplitude maximum. However, if there are truly only two scatterers present, this will not degrade resolution performance seriously. In effect, the two scatterers are at worst somewhat more difficult to resolve when they interfere destructively (a phase difference of 180° at the center of the observation window), where resolution again requires the ability to make accurate Doppler measurements. If a full modulation cycle can be observed, as needed if the two scatterers are to be fully resolved, the difference between constructive and destructive phase interference is not large. Thus we may also in practice define Doppler resolution for scatterers regardless of their phasing. With the use of the complex signal processor output we obtain a Doppler resolution of $1/T$, where T is the duration of the observation window or of the transmitted signal.

The very same discussion could be given for scatterers separated in delay by $\Delta\tau$ and a signal bandwidth (or width of the observation window in the frequency domain) of B . From the duality between time and frequency it is clear that delay resolution is given by $\Delta\tau = 1/B$, under the assumption that the complex processor output is analyzed. Depending on the signal-to-background ratio, delay resolution might degrade a little as the phase difference (at the center of the frequency band) between the two scatterers approaches 180° , in the sense that range accuracy is not quite as good as for constructive interference. Again, if there are only two scatterers present and if they can be resolved, the difference between the two cases is insignificant.

The preceding discussion was concerned with the resolution of two point scatterers in noise. If the signal-to-noise ratio is high, the resolution performance of $1/B$ and $1/T$ can be achieved regardless of the phase between two scatterers. This is not the case when interference comes from adjacent, incompletely resolved additional scatterers, or when the signal-to-noise ratio is low. Rather, the amplitude/phase pattern will be distorted, and resolution performance will depend on the relative phasing of the two scatterers. It is more difficult to recognize amplitude maxima at the boundaries of an observation window centered on an amplitude minimum (scatterers in phase opposition) than to recognize minima at the boundaries of an observation

window centered on a maximum (scatterers in phase). The practical significance is that under poor interference conditions it is more difficult to realize the inherent resolution performance when the scatterers interfere destructively than when they interfere constructively. It is an important problem in practice, in particular when we try (intentionally or unintentionally) to resolve scatterers a little more closely spaced than one resolution cell.

2.1.3 Superresolution

The perception that the width of the central spike of the ambiguity function defines resolution has, over the years, led to many attempts to improve resolution performance by somehow decreasing the width of the spike without increasing the signal bandwidth or the signal duration. These are the so-called superresolution methods. However, if we include in the definition of resolution the necessary requirement that one must be able to measure the position of each response with reasonable accuracy (a small fraction of the resolution cell), little remains of the utility of superresolution. For example, what is the use of sharpening the delay response by enhancing the outskirts of the signal spectrum, if in the process the position of the response is shifted by the noise? We note that in the performance claims for specific superresolution methods, one must first subtract a factor of two from the gain in resolution due to the switch from intensity responses to complex responses. This is to say that reclaiming the factor of two in resolution lost by wrongly utilizing the intensity output cannot be considered a performance enhancement due to superresolution.

The basic problem with true superresolution can be readily understood if we start from the definition of resolution, that a full cycle of the modulation pattern must be observed without significant distortions, either in time or frequency. Let us again consider Figure 2.2, with the observation window centered on a maximum of the amplitude pattern. As explained above, if the window extends from one amplitude minimum to the next, we observe the full modulation cycle and can calculate the Dopplers of the two scatterers. Now, let us go not quite to the minima to the left and right of the maximum, so that less than a full modulation cycle is observed. If the noise is low, we can extrapolate the positions of the amplitude minima and still calculate the Dopplers of the scatterers. Thus, depending on how low or high the noise is, we may not have to observe the full modulation cycle. This is true superresolution, and it gives acceptable errors only when we do not have to extrapolate the modulation cycle by much and when the noise is low.

Let us go to the other extreme where the observation window happens to be centered on an amplitude minimum, or where the two scatterers happen to be in phase opposition at the center of the window. The destructive interference between the two scatterers enhances the effects of noise because the combined return signal is weaker if less than a full modulation cycle is observed. In that case, over much of the window the interference will be destructive. Any extrapolation of the full amplitude/phase modulation pattern thus would require an extremely high signal-to-noise ratio, and such is typically not available in practice. In other words, although some degree of superresolution might work if the two scatterers happen to be in phase, the performance will severely degrade as the timing of the observation changes to when the two scatterers are in phase opposition (yet we have no influence on the timing). If two scatterers interfere destructively, we would need impractically ideal conditions to derive the scatterer Dopplers from observing only a fraction of the modulation cycle.

The actual practical conditions are even worse. Above we considered the resolution of two ideal point scatterers in a noise background, but in practice the problem typically is the interference from other scatterers rather than noise. This makes superresolution methods even less workable.

One might argue that perhaps these limitations apply for Fourier transform processing, or correlation processing, but not for more general processing methods. However, there are two basic considerations that govern performance regardless of the specific processing method. First, the performance must necessarily worsen as the total received energy decreases. There is less signal energy received when the observation window is decreased, with the decrease particularly rapid when the scatterers are in phase opposition at the center of the observation interval. In that situation, as the width of the window is reduced from coverage of a full cycle, the decrease in signal energy is at first small, but becomes very rapid below about three quarters of a cycle. Second, in some sense we must be able to recognize the interference pattern from the two scatterers, but this task becomes more difficult as the width of the observation window becomes smaller than one modulation cycle. Both problems are a matter not only of Fourier transform processing, and thus cannot be alleviated by other types of processing.

We have given a rather detailed explanation of why superresolution cannot work even for the simple model of two point scatterers. As discussed earlier, an even more serious problem arises from the impossibility of mathematically modeling complicated man-made targets with sufficient realism to serve as the basis of practical processing algorithms.

2.1.4 Two-Dimensional Resolution

We return from superresolution to ordinary resolution, but resolution in two dimensions. Generally, two scatterers that are to be resolved in Doppler will not be at exactly the same range, and scatterers to be resolved in range will not be at exactly the same Doppler. There will be a difference in range as well as Doppler. One-dimensional resolution is merely the special case in which the difference in the other dimension is so small relative to resolution performance that it does not contribute to resolution.

In accordance with the above discussion, it is justified (provided the complex processor is utilized) to consider $1/B$ the delay resolution cell and $1/T$ the Doppler resolution cell. Let us assume a Doppler separation of the scatterers of δv and a delay separation of $\delta \tau$. We normalize the separations by the resolution cell in each dimension,

$$N_v = \delta v / (1/T) \quad (2.4)$$

$$N_\tau = \delta \tau / (1/B) \quad (2.5)$$

where N_v is the Doppler separation of the scatterers in Doppler resolution cells, and N_τ is the delay separation in delay resolution cells. The separation of the two scatterers in the delay/Doppler plane is

$$d = \sqrt{N_v^2 + N_\tau^2} \quad (2.6)$$

With complex-image analysis, the resolution performance is $d = 1$. From (2.6), we then require

$$N_v^2 + N_\tau^2 \geq 1 \quad (2.7)$$

As is to be expected, Doppler resolution performance depends on the scatterer separation in range, and range resolution performance depends on the scatterer separation in Doppler. For example, if the delay separation is half of the delay resolution cell, (2.7) implies that the Doppler separation needed to resolve the two scatterers is 0.87 Doppler resolution cells. One-dimensional resolution performance is a special case of (2.7).

The relation between resolution considerations based on the width of the central spike of the ambiguity function and those based on the transforms of image cuts through the responses is a simple one. The decline in

strength of a response, as we consider image positions away from the peak, can be calculated by assuming a receiver with a progressively increasing mismatch in range rate to the return signal or spectrum. We know that the strength of a response decreases by about 3 dB when the (linearly changing) phase of the return signal or spectrum becomes mismatched to the receiver by half a cycle over the extent of the signal or spectrum, so that the 3 dB width of the response corresponds to a phase change of a full cycle. However, instead of using the 3 dB width of the response for resolution, we must use the underlying phase change of one cycle. This is the way of realizing the inherent resolution performance of radar.

2.1.5 Resolution of Weak Scatterers from Strong Scatterers

With the conventional view of resolution that two response peaks should be recognized, based on the width of the central spike of the ambiguity function, one usually considers the resolution of point scatterers of comparable strengths. In fact, resolution originally was defined for two point scatterers of the same strength. Another problem arises with the conventional definition of resolution when one scatterer is much weaker than the other. If the weak scatterer is, say, 20 dB weaker, and if they are separated by one resolution cell, the weak scatterer will merely distort the response from the stronger scatterer a little. Then one asks the question, how much must the scatterer be separated in order for its response to be detectable in the tails of the strong response? This is to say that the conventional viewpoint on resolution leads to the conclusion that resolution performance on a scatterer much weaker than a neighboring main scatterer degrades from that for two equally strong scatterers, and degrades severely if the second scatterer is much weaker.

This degradation of resolution does indeed take place if only the envelope of the correlator output is analyzed. The question is irrelevant if the complex correlator output is analyzed, in which case resolution does not depend on the relative strengths of the scatterers. This is easily recognized from the modulation patterns of Figure 2.2. As we decrease the relative strength of the second scatterer, the effect on the amplitude function will be a decrease of the modulation index. The modulation function will approach a sinusoid of progressively smaller amplitude. This is also the case with the phase function. Now, as long as we are able to measure a full cycle of the modulation (with an acceptable distortion), we can resolve the two scatterers. However, the measurability of the modulation depends only on the signal-to-background ratio of the weaker scatterer. Thus, as long as the weaker scatterer has an adequate signal-to-background ratio, it is possible to measure

the modulation pattern (assuming a minimum separation of $1/T$ in Doppler or $1/B$ in delay). This means that the weak scatterer can be resolved. The signal-to-background ratio for the stronger scatterer is of no concern because it will be much higher. We conclude that resolution performance does not degrade when the second scatterer becomes much weaker than the first, but is governed only by the background conditions for the weaker scatterer. These can be favorable even for a weak scatterer because of the huge integration gain implied by the coherent processing needed for Doppler resolution.

In the case where a weak return of adequate signal-to-noise ratio is to be resolved from a much stronger return, if a correct mathematical model were available, superresolution principles would work just as well as the analysis of the amplitude/phase patterns, but also only when the scatterers are not separated by significantly less than $1/T$ or $1/B$. In other words, superresolution methods will work as long as we only want to prevent a degradation of resolution performance due to the fact that one scatterer is much weaker than the other, but then it is not true superresolution. If we were able to model a man-made target with sufficient realism, it would not matter whether we resolve the scatterers via complex-image analysis procedures or by mathematical superresolution methods.

2.1.6 Effects of Weighting for Sidelobe Suppression

Perhaps the most popular pulse compression signal is the linear FM signal, which has the property that the shape of the spectrum is the same as the shape of the transmitted pulse. Since it is most practical to operate with constant transmitter power, the shape of the linear FM spectrum is also rectangular. However, the Fourier transform of a function with sharp edges generates high sidelobes (in the case of a spectrum, these are high range sidelobes). For this reason the spectrum is smoothed into a bell-shaped form on reception, or weighted, in order to reduce the range sidelobes. The process widens the response by a factor that is typically in the order of 1.3. This has been taken to mean that weighting degrades resolution. Analogous considerations also apply to other waveforms, such as the phase reversal code, even though the sidelobe suppression method cannot be considered as weighting. However, the result still is a smooth bell-shaped spectrum, or a weighted spectrum.

With the correct definition of resolution, implemented by utilizing the complex processor output rather than the intensity output, weighting does not affect resolution at all. Although we must use weighting under almost all circumstances in order to reduce the mutual interference between scatterers,

in the analysis of the complex image we remove the weighting after it has done its task. In other words, we take the transform of responses generated with the inclusion of weighting, but after taking the transform we remove the weighting. The weighting has no consequence whatsoever if the problem is to resolve two point scatterers, as resolution is defined. However, when we resolve two scatterers next to a third scatterer, weighting widens the responses, and the tail of the response from the third scatterer may enter the transform window for the two scatterers to be resolved. Measurement accuracy of the scatterer positions then will be degraded. Hence, although in a two-scatterer situation weighting does not matter, it does have an adverse effect when additional scatterers interfere. Analogous considerations also apply with respect to weighting in the time domain, for the suppression of the Doppler sidelobes.

There is a detail that deserves mentioning. Weighting for range sidelobe suppression works well when the Doppler difference between return signal and correlator reference is small. Similarly, weighting works well for Doppler sidelobe suppression when the delay between return and reference is small. This is not always the case, and that leads to the conclusion that in some situations the transmission of a bell-shaped time function or bell-shaped spectrum is desirable, because a weighted transmitted signal or spectrum is less sensitive to mismatches between target and reference delay and Doppler shift.

When resolution is based on the analysis of the complex processor output, the consequences of weighting the transmitted signal or spectrum are again easy to see. For matching purposes we must also use weighting of the signal or spectrum in the correlator, but this weighting is removed when the processor output is analyzed. Only the weighting upon transmission remains, and this still leaves the signal or spectrum bell-shaped. Now, in the absence of noise, we could observe the amplitude and phase patterns over the entire interval used during the transmission, regardless of the decrease of the amplitude toward the ends, because a weighted signal was transmitted. In the presence of noise, however, this amounts to superresolution. We can go only so far toward the fringes of the interval until the effects of noise start distorting the amplitude and phase functions to an unacceptable degree. The achievable resolution is limited by the decrease in the signal-to-noise ratio as we approach the fringes of the interval, since any part of the frequency band or signal with low signal-to-noise ratio is useless for resolution purposes. *It is important that we can determine when the amplitude/phase pattern becomes distorted*, so that we do not try to achieve better resolution than allowed by the noise background. This again would amount to superresolution. It is the

noise background that determines the effective width of the signal or spectrum for resolution when weighted signals are transmitted.

We could utilize some degree of superresolution, provided that in each case we check whether the interference is sufficiently low. In the case of a bell-shaped spectrum, for example, if the interference is so low that the known shapes of the spectrum tails are not significantly distorted, we can “extrapolate” the bandwidth (utilize the fringes of the bell-shaped spectrum) and improve range resolution. The essential point, however, is that the conditions must be examined for each case, which in the presence of interference from other scatterers means for each attempt to resolve a composite response into two scatterer positions. This is not feasible with mathematically formulated superresolution methods that model the entire target rather than two scatterers at a time. The practical solution is to examine and analyze a huge variety of real data in order to obtain an understanding of the actual backscattering behavior of real targets. Based on this insight, one formulates algorithms to perform the various measurements, after many tests and failures obtaining workable algorithms. The last step is to automate the algorithms so that they reproduce the manually achieved results sufficiently well. This procedure is referred to as an expert system approach.

2.1.7 Resolution and Shifting Scatterers

All of the above considerations apply for ideal point scatterers. The practical interest in regard to target identification is not so much in the case where two targets can be represented as point scatterers, but when two scatterers on a single target can be represented as point scatterers. As explained in Chapter 1 (and, in more detail, in Section 2.4 of [3] and in [4]), such an assumption often is not allowable when complicated extended target features are considered, such as cavities or irregular corners, which are the primary contributors to the target image. The phase centers of these “shifting” scatterers may be moving with aspect angle and frequency. How meaningful is the definition of resolution for this important class of scatterers?

As a first case, assume a target feature with a significant extent in cross-range but not in range, whereby we mean that the phase-center motion can be significant in crossrange but not in range. Certainly, if the range extent of a feature is small, the phase center cannot move much in range. Assume that the phase-center motion with aspect angle introduces a linearly changing signal phase, which implies that the phase center is moving with constant range rate. From the definition of resolution, it follows that if this phase-center wander causes the signal phase to go through one cycle, the response will

have shifted by one crossrange resolution cell. Since one cycle implies a range change by half a wavelength, the range of the phase center has changed by just one half-wavelength. All resolution considerations still apply and a second response does not affect the measured position of the first, but the phase-center motion has falsified the crossrange position of the response. If the phase-center motion additionally includes a nonlinear component, the response widens. However, we can resolve it from another response as long as the interference between the two returns generates recognizable amplitude and phase modulation patterns.

In the other extreme, where the extent of the feature is large in range but not in crossrange, the dual considerations apply. If the phase-center motion is proportional to frequency, the response will shift by one range resolution cell for every half-wavelength shift of the phase center over the spectral bandwidth. A nonlinearity will again cause a widening of the response, but whether or not two responses can be resolved depends on the recognizability of the interference pattern. Although we can measure the position of a response accurately despite the closeness of another response, the measured position does not represent the range position of the associated scatterer.

In the general case the target feature may be extended in range as well as crossrange, so that the phase center of the feature shifts with both aspect angle and frequency. The cases of a shift only in range or only in crossrange, in particular when the linear phase components are dominant, are just special cases of a shift of the response in both range and crossrange simultaneously. Suppose there is a component in the backscattered signal for which the instantaneous translation in delay is proportional to the instantaneous translation in Doppler. If the linear component is dominant, then the phase function will be linear in a direction in the image plane given by the proportionality factor. The response will be compressed in this direction, and it can be resolved from another response in that direction in accordance with the principles discussed above. A complicated feature may have several such components that will yield a series of responses. If they are compressed, so that peaks are formed, their phase functions will be linear in the directions in which the responses are translated in the image plane. This behavior of the phase function allows them to be recognized as a set of spurious responses, which we designate as sideband responses [3].

2.1.8 Section Summary

Resolution cannot be defined by counting the number of response peaks, but must include the capability to measure range and Doppler of the targets to be

resolved. With this practical definition, the envelope of the ambiguity function does not fully describe close-target resolution. Specifically, a widening of the central spike due to weighting for sidelobe suppression does not degrade resolution. Resolution does not degrade when one scatterer is much weaker than the other, in contrast to what the ambiguity function suggests.

Resolution of two scatterers can be understood only from the amplitude and phase patterns of two interfering scatterers. It is essential to understand these patterns in all their forms. They also apply if the scatterers have shifting phase centers.

In principle, superresolution by up to a factor of two can work because that factor represents the switch from response envelopes to complex responses. In practice, superresolution even by a factor of two will not work if the algorithm is based on a mathematical target model (e.g., a set of fixed point scatterers). With complex-image analysis, a small degree of superresolution is obtainable when the validity of the two-scatterer model is verified for each measurement.

Resolution in one dimension is aided by a scatterer separation in the other dimension.

2.2 Asymmetry of Range and Crossrange Resolution

Guided by optical thinking, one tends to assign range and crossrange resolution similar roles. This section considers resolution requirements and the different roles of range and crossrange resolution in radar imaging and target identification. High range resolution is unproblematic and necessary, whereas crossrange resolution on moving targets is so problematic that only a limited degree is achievable. The following explains these important points in more detail.

With photographs, resolution in azimuth and resolution in elevation play the same role, and it is sensible to have the same resolution performance in the two angles. This thinking appears to have migrated to radar, where one sees the tendency to choose equal range and crossrange resolutions. Such a choice would be justified if radar target identification could be performed in a way similar to optical identification, but this is not the case. Making range and crossrange resolution equal might be justified for stationary targets under some conditions, but certainly not for moving targets. In the following we will show that, in general, range and crossrange resolution must be chosen rather differently, the choice of the latter depending on the operational conditions. An exception is the special case of SAR surveillance of stationary ground vehicles.

2.2.1 Resolution Requirements for Target Identification

Once we have chosen the values of range and Doppler resolution in a particular system, if (as should be the case) the processing algorithms are based on Fourier transforms, then the required signal bandwidth and duration are fixed. As discussed earlier, the universal applicability of Fourier transforms is one of the central points of complex-image analysis technology. Even more important is the recognition that an intensity image does not contain sufficient information for fully automated reliable target identification in a large database.

One must extract specific target characteristics from radar images, but these characteristics must be the type that can be readily related to the appearance of the target. These are target parameters such as the length, width, and other special features, but this is not enough for reliable identification. We must also measure the locations of those scatterers that are not recognizable. The information to be used for identification must be easily accessible and physically meaningful. If one accepts this premise, then it becomes necessary to associate responses in an image with specific scatterers on the target, and to measure the positions of the scatterers relatively accurately. This requires range and Doppler resolution good enough to yield the required measurement accuracy for the scatterer positions, but not necessarily equal resolution in the two dimensions.

With a further pursuit of this approach, one might ask how much range resolution is needed, how much crossrange resolution, or whether range resolution by itself might be adequate. We will show that for the general class of moving targets, the appropriate questions are the following: How much range resolution is needed for the task? How much crossrange resolution is necessary to assist range resolution? Also, in those difficult cases in which range resolution is not effective because of a small range extent of the target, how much crossrange resolution is needed to replace the missing range resolution performance? Thus, we should not think in terms of one-dimensional versus two-dimensional resolution, *but consider range resolution the primary resolution mode, with crossrange resolution needed for assistance, depending on the application and conditions of the moment.*

2.2.2 Measurement of Feature Positions

As already discussed, we utilize special target features that are characteristic of a specific target whenever we can detect them. For example, if we observe a gun on a ground vehicle, this is important information. On the other hand, a gun is rarely detected in an image. Length and width measurements are quite generally applicable, yet some targets differ so little in their overall dimensions that the requirement on measurement accuracy would be too high if one were to

rely on just the length and width measurements for target identification. This is why we also must use scatterers that have the same characteristics for each target, such as corners and cavities, and for which the difference between targets lies only in differences in the scatterer positions on the targets. Thus, the measurement of scatterer positions is very important for target identification. It is the required measurement accuracy for scatterer positions that drives the necessary resolution performance.

The various scatterers on a target generate image responses, and we want to measure the range and crossrange positions of the scatterers associated with the responses. If all the scatterer responses were cleanly resolved, we could measure the range/crossrange positions of the image responses and they would correspond to the positions of the actual scatterers, but typically a radar does not provide this degree of resolution. Moreover, resolution beyond some degree does not help, because the observable features have significant extents and shifting phase centers. Now, no position measurement of reasonable accuracy is feasible unless the response of a scatterer can be resolved from the other responses. In principle it should not matter whether the resolution of a scatterer is accomplished by range resolution, crossrange resolution, or a combination of both. As long as a response is resolved by any means, the position of the associated scatterer can be measured. Such an attitude might be appropriate for a stationary target, but it is not useful for a moving target. A moving target does not permit much choice in how to resolve its scatterer responses so as to allow measuring the scatterer positions.

There are important reasons why range resolution is the primary kind of resolution, and should be as high as affordable (with a limit imposed by the effective range extents of scatterers). The obvious practical difference between range and crossrange resolution is that range resolution requires little dwell time, because it does not need an aspect angle change of the target. However, there are more important reasons as far as performance is concerned. Range resolution is straightforward to implement, which is not the case for crossrange resolution. In fact, crossrange resolution can not typically be implemented to a satisfactory degree and quality if good range resolution is not available. For example, the motion compensation needed for crossrange resolution may be difficult to implement for a moving target when the aspect angle is small and, without good range resolution, the compensation is essentially impossible to implement well enough when the aspect angle approaches broadside. Also, the spurious responses generated by cavity-type features are a much more serious problem with high crossrange resolution than with high range resolution, at least when the target motion is not very smooth.

For signal processing and image analysis, it would be most desirable if one could avoid crossrange resolution altogether, relying solely on range resolution. Unfortunately, even high range resolution does not guarantee that the scatterers will be resolved, because several still may be within a single range cell. Then we need additional crossrange resolution to separate the scatterers in the same range cell. In practice, however, we will use only the minimum amount of crossrange resolution needed to achieve this purpose, because the implementation of crossrange resolution leads to progressively worse problems as crossrange resolution is increased. The objective is to *use only as much crossrange resolution as needed* to perform the measurement of scatterer positions to the required accuracy. Since the correct crossrange resolution depends on the type of target and its behavior, the degree of *crossrange resolution must necessarily be chosen adaptively*.

2.2.3 Fully Utilizing the Available Range Resolution

Let us start with the simplest, though most unrealistic, case in which we model a target by a set of point scatterers in fixed locations on the target (no dependence of the scatterer positions on aspect angle or frequency). Suppose that there is no crossrange resolution. We begin with the unrealistic assumption that range resolution is so high that most of the resolution cells contain only one scatterer. We then can accurately measure the range positions of the scatterers without the need for crossrange resolution. We cannot measure their crossrange positions, but perhaps we can identify the target based only on its accurate range profile.

Retaining ideal point scatterers, suppose that hardware constraints prevent us from achieving a range resolution cell so small that most cells contain only a single scatterer. In each range cell the scatterers interfere with each other, and we can neither tell how many scatterers have contributed to the composite response, nor measure the ranges of the scatterers to a small fraction of the range-cell width. In other words, if we do not have sufficient crossrange resolution to separate the scatterer responses in a range cell, we also cannot measure the scatterer ranges with sufficient accuracy. Resolution in range is not enough even if we are interested only in the range positions of the scatterers.

The consequences of inadequate crossrange resolution are very serious for practical targets. We will give an illustration for the very simplest (ideal) case, and later discuss why the consequences are much worse in practice. Figure 2.3 shows the range responses from two resolved ideal point scatterers. Because of the unavoidable weighting for sidelobe suppression, the




Figure 2.3 Responses from ideal point scatterers.

half-power widths of the responses are 1.35 resolution cells (or gates). We want to measure the range of the response at the right of the figure, which in the absence of interference (as in this figure) is just the position of its peak. The response on the left is assumed to be at a different crossrange, but without crossrange resolution this is irrelevant. If the range separation of the response on the left is less than the width of the response, it will interfere with the central response. We want to determine the consequences on range accuracy; that is, on the position of the peak.

In Figure 2.4 we show the combined response when the range separation of the interfering response is decreased from that shown in Figure 2.3. As long as the range separation is at least about one gate, we can resolve the responses via analysis of the complex responses without the use of superresolution, so that an accurate range measurement can be performed. Thus, for our illustration we choose the separation smaller than one gate, selecting half a range gate for this illustration, but with different phasing of the responses. For the top response of Figure 2.4 the phase difference is 90° , for the center response it is 135° , and for the bottom response it is 180° . In all three cases, the scatterers are located at Range Gates -0.5 and 0.0 , as shown by the vertical arrows in the figure. The scatterer on the left has an amplitude of 1.3, and the one on the right has an amplitude of 1.0. The maximum amplitude for




Figure 2.4 Interference pattern for the two responses.

each interference condition is displayed in the legends on the left of the figure, and the dashed curves show the response from a fixed point scatterer, for comparison.

If we ignore superresolution for the moment, all we can do for such a close scatterer separation is to accept the peak positions as the scatterer ranges. The peak positions in the top and center traces are close to the actual scatterer positions, with the response peak displaced from the closest scatterer by 0.17 and 0.07 resolution cells, respectively. However, even though the range accuracy is acceptable, in both cases we detect a single scatterer instead of two. For the bottom trace we correctly conclude that there are two scatterers, but the peak positions are displaced from the scatterer positions by 0.45 and 0.77 resolution cells. All three receiver outputs of Figure 2.4 provide unacceptable range measurement performance (unless the radar has a bandwidth much higher than usual), with the first two cases revealing only a single rather than two scatterers. The transforms of the three complex responses

will show that two scatterers are involved, but by extrapolating the patterns we still can measure the range positions of the two scatterers only inaccurately. Nevertheless, it can be important not to interpret the composite response from two scatterers as a response from a single scatterer.

2.2.4 The Real Situation and Requirements on Crossrange Resolution

Since real targets do not backscatter like a set of point scatterers in fixed positions, the problem of position measurements is even more serious than indicated by the idealized example of Figure 2.4. In that example we assumed that a range cell contains only two scatterers, but with poor crossrange resolution there will often be more than two unresolved scatterers in a range cell. We also assumed point scatterers in fixed locations on the target. The strong scatterers tend to be extended irregular corners and cavities rather than point scatterers, and their positions may shift with aspect angle and over the frequency band of the signal. Lastly, there may be many spurious responses that indicate positions where there are no scatterers and that may mask genuine responses [3 (Section 2.4), 4]. We already discussed why the suggested remedy of superresolution cannot work.

The problems examined here mean that a nominal range resolution as defined by the signal bandwidth is meaningless for the measurement of range positions if not accompanied by a sufficient crossrange resolution. The only situation in which the nominal range resolution of the radar is meaningful is when range resolution is chosen so high that each range cell contains only a single significant scatterer. Even if we are not interested in crossrange measurements, if a range cell may contain more than one scatterer, we need crossrange resolution in order to fully resolve the scatterers, so that an accurate range measurement is possible. Again, the reason for range resolution is the need for range measurements, and a nominal range resolution of some value does not guarantee that we can measure the range of a scatterer to a fraction of the range cell. This is where we need at least some degree of crossrange resolution.

We have shown that unless adequate crossrange resolution is implemented, we will not be able to measure the ranges of the scatterers sufficiently accurately, which means that no use can be made of scatterer positions. However, having implemented some degree of crossrange resolution, we can also measure the crossranges of the scatterers to an accuracy much better than the width of the crossrange resolution cell. *Thus we can perform accurate two-dimensional position measurements even when crossrange resolution is poor.* In summary, performing accurate positional measurements

requires high resolution in only one dimension, which must be range, but it will generally have to be assisted by some degree of crossrange resolution. This situation is quite different from implementing square resolution cells, as one would in an optical system.

The reader may wonder why we are discussing in such detail an issue that should be obvious. The reason is that there have always been attempts to solve specific radar problems by providing resolution in range only, without crossrange resolution. We want to show why this is generally impossible, in particular with real data. For the following illustrations the reader should keep in mind that we are increasing crossrange resolution in order to improve measurement accuracy in range, not in crossrange.

2.2.5 Illustration with Real Data

In our illustrations throughout this book, we will not use the conventional presentations of SAR/ISAR images giving the value of the cross section in each pixel. This type of presentation is analogous to a photograph, and would be justified if one could identify a target by examining such an image. Since we strongly believe that this is not possible (particularly when target motion allows only poor crossrange resolution), we will instead use a presentation that highlights the strong image peaks, which serve as starting points for the extraction of feature positions. We will almost always use this so-called peaks plot presentation, where the positions of the local maxima of the intensity image are indicated by dots. The area of a dot is proportional to the height of the associated response peak. A dot is plotted only if the local maximum is found in the range gate of the response, the crossrange gate, and the two diagonals. One must keep in mind that the local maximum indicated in the peaks plot may exceed the surrounding level only by a minimal amount, so that separated dots by no means indicate resolved responses. Furthermore, one must not forget that the image is complex and that the phase contains useful information. However, since we have not found a practical way of presenting complex images, we will show only intensity images in peaks plot form.

The peaks plot has various advantages over the conventional presentation of SAR/ISAR images. The primary reason for its use is that we typically examine specific responses, which can be easily identified by the positions of the dots. Another reason is that the peaks plot allows one to judge the quality of an image far more easily than a conventional plot. As discussed in Chapter 1 and Appendix H, we know that a typical fighter aircraft, ground vehicle, or small ship has about 20 or 30 responses (fewer, if resolution is low) strong enough to be observable above the background from the weak scatterers

generated by the target. Hence, if we see a peaks plot with a much larger number of dots of significant size, we immediately suspect that the motion compensation is inadequate, that the image contains many spurious responses, or both. An image with many more dots almost always has a quality too poor to be of use.

We emphasize two points in connection with the peaks plot presentation. First, the peaks plot contains less information than an intensity plot; its utility lies in guiding the eye to the more useful information. Second, the number of dots serves as a convenient tool for quickly estimating crude image quality, but we do not use it for any processing decisions. Rather, regardless of whether a peaks plot shows a "reasonable" number of dots, we must check whether the individual responses are sharply "focused" (compressed is a better term because this is not optics). Some of the peaks plots contain dotted lines. These show the locations within the image of image cuts displayed in other figures. Each dotted line is labeled by the numbers of the figures to which it corresponds.

In order to demonstrate the need for at least some degree of crossrange resolution, we use a SAR surveillance image of a turning vehicle. Range resolution of the radar (the range-gate width) is 1 ft, and an automated motion compensation as described in [3] was performed. In Figure 2.5 we show the peaks plot image of the vehicle with a crossrange resolution of about 4 ft, with the approximate vehicle outline given by the dotted rectangle. Because the vehicle is turning, the imaging time is only 0.035 seconds. We choose the weak peak marked by the dotted line for our illustration. A cut through the complex image in the crossrange gate of the response is shown in Figure 2.6, with the amplitude function on top and the phase function at the bottom. The dashed line shows a fixed point scatterer response, for comparison. The half-power width of the response relative to that of the point scatterer is 1.225, which implies that the response is generated by more than one scatterer, or by a scatterer with a strongly shifting phase center. The decision between the two possibilities can be made by taking a transform of the composite response. If the response is generated by a single scatterer, the amplitude function of the transform will be roughly constant. If it is generated by two interfering scatterers, the transform will show the typical interference pattern for two scatterers, as in Figure 2.2.

This transform, taken between amplitude minima, is shown in Figure 2.7. The pattern of the transform does not have the constant amplitude associated with a single scatterer. The amplitude/phase pattern also does not approach that of two interfering scatterers well enough to measure the range positions of the two scatterers with reasonable accuracy. We could




Figure 2.5 Peaks plot image with 4 ft crossrange resolution.




Figure 2.6 Image cut in the crossrange gate of the response.




Figure 2.7 Transform of the response marked in Figure 2.6.

decide to accept the response as coming from a single scatterer with the range position given by the position of the peak. However, this would be a poor decision because mistaking a response from two scatterers as one coming from a single scatterer and then assigning a position to that scatterer is very detrimental for target identification. The better alternative is to accept two scatterers with poor accuracies in range, and not place much weight on the measured range positions. In the case of Figure 2.6, the measurement error might approach 1 gate, or 1 ft. As a test, we regenerated the image with cross-range resolution improved by a factor of four to about 1 ft. The new amplitude/phase pattern corresponding to Figure 2.7 was much improved, yet was not good enough to allow an accurate determination of the positions of the scatterers. The reason can be recognized from Figure 2.6. The sharpness of the amplitude minimum in Gate 37.2 implies that it is not a minimum between two neighboring resolved responses (as in Figure 2.3), but that it is caused by destructive interference between two adjoining unresolved responses, which is the worst case of interference (the TSA requires the most extrapolation for this case).

With a further doubling of the imaging time to 0.28 seconds we obtain an image with a crossrange resolution of 0.5 ft, with the peaks plot given in

Figure 2.8. We show the image for comparison with Figure 2.5, which has many fewer peaks. The vehicle outline is in the same location in both figures. The image cut in the crossrange gate of the response of interest is shown in Figure 2.9. By comparison with Figure 2.6 we see how the interference in range is reduced as crossrange resolution is improved. The unresolved response indicated by the bulge in about Range Gate 38 of Figure 2.6 has shifted in crossrange away from the central peak in Figure 2.9, so that the central peak more closely matches the response of a fixed point scatterer. The normalized half-power half-widths for the left side of the peak are 1.417 and 1.201, in Figures 2.6 and 2.9, respectively. Furthermore, the phase corresponding to the central peak deviates much less from linearity in Figure 2.9 than in Figure 2.6. We find that a refinement of the range position is possible, by taking the transform of the response and interpreting it in accordance with the two-scatterer model.

We again double the imaging time to 0.56 seconds, obtaining a crossrange resolution of 0.25 ft and the peaks plot image of Figure 2.10. The vehicle outline is in the same position as in the lower resolution figures. The appropriate image cut through the same response is shown in Figure 2.11. As indicated by the relative half-power width of 0.938, the response is




Figure 2.8 Peaks plot image for a crossrange resolution of 0.5 ft.




Figure 2.9 Image cut in the crossrange gate of the response.




Figure 2.10 Peaks plot image for a crossrange resolution of 0.25 ft.




Figure 2.11 Image cut in the crossrange gate of the response.

essentially generated by a single scatterer, so that the peak position may be accepted as the position of the scatterer.

The image of Figure 2.10 contains so many peaks that we suspect that the motion compensation may be inadequate (in an automatic system, the quality of the motion compensation must be checked for all images). The image indicates that we probably have attempted to achieve a crossrange resolution too high for the given situation, primarily the motion behavior of the vehicle. However, even if the motion compensation may be too poor to allow performing accurate crossrange measurements, it does improve the range measurement. As will be seen later, in situations where a target's motion behavior is so erratic that no accurate crossrange measurements can be performed, the motion still must be utilized to obtain accurate range measurements.

2.2.6 Checking the Limit on Crossrange Resolution

Typically, when too high a crossrange resolution is attempted, the residual motion compensation errors effectively generate an interference pattern from two or more images slightly displaced in crossrange, which leads to many

more responses than there are scatterers. A motion compensation of adequate quality cannot be performed over an arbitrarily long interval, in particular if the vehicle is not behaving like a rigid target. *This means we must verify the quality of the motion compensation in each case and, if necessary, to reduce the imaging time.* Even though nominal crossrange resolution is degraded in the process, we have no choice.

We will describe the important step of checking the quality of the motion compensation in considerable detail in Section 2.3. For a brief summary, to check the quality of the motion compensation for an image, we have to find a response which is sufficiently resolved from others and stands out in the background interference so well that the amplitude function of its transform does not have deep breaks. In that case the phase describes the range motion of the scatterer, which is to be checked. Finding a sufficiently resolved peak may be difficult or impossible, because the motion compensation may be poor at this stage. If such a peak cannot be found, we must reduce the nominal crossrange resolution and try again. This will also be illustrated in the chapters on target identification. As nominal crossrange resolution is reduced, the motion compensation will be more satisfactory; there is an optimum compromise between improved motion compensation and degraded crossrange resolution.

The image of Figure 2.10 contains several peaks that permit the described measurement, with the most suitable peak (as determined from the transform of the response) marked by the arrow. An image cut in the range gate of the peak is shown in Figure 2.12. The relative half-power width of 1.134 is sufficiently larger than unity to conclude that there is either a problem with interference or with the motion compensation. This also can be decided only by taking a transform. The transform over the entire interval is displayed in Figure 2.13, after subtraction of a linear fit to the phase. The rapid phase (and amplitude) fluctuations are caused by the lower level responses in Figure 2.12, whereas the roughly quadratic up-and-down trend of the phase is due either to a motion compensation residual or the phase-center shifting of the scatterer. To find the cause, we compensate the slow phase trend shown by the dashed curve, and apply the same compensation to other responses in the image. These responses will be properly compressed if the phase modulation is due to residual motion, but will be smeared if it comes from the phase-center shifting of one scatterer.

Figure 2.14 shows the image cut of Figure 2.12, after compensating the image with the slow phase trend of Figure 2.13. The compensation is essentially perfect, as indicated by the near unity (0.96) relative half-power width, the nearly linear phase, and the absence of high response sidelobes. It is




Figure 2.12 Image cut in the range gate of the peak selected for checking the motion compensation.




Figure 2.13 Transform over the interval displayed in Figure 2.12.




Figure 2.14 Image cut of Figure 2.12 after phase compensation.

always necessary to verify that the phase compensation has worked well, as done with Figure 2.14, because otherwise it becomes problematic to check other responses in the image.

When we apply the same phase compensation that compresses the response of Figure 2.11 into the high-quality response of Figure 2.14 to other responses in the image, we find that the compression also works on other responses, but not as well. From this fact we conclude that at least a good part of the slow phase variation of Figure 2.13 is due to a motion compensation residual. To determine whether all of it is a motion compensation residual or whether a part comes from a phase-center shift of the scatterer, we must compare the signal phases of the other responses (after compensation), considering that the phase centers of these scatterers also may not be stable. However, even though one can differentiate between the various effects and come to the correct conclusion ([3], Section 6.3), we should not be too ambitious if the system is to be fully automated. Thus it is simplest and practical to examine Figure 2.13, find some subinterval over which the phase function is sufficiently linear, and form an image over that subinterval. In this instance we conclude that the underlying trend of the phase is quite

linear between the normalized times of -0.2 and 0.4 seconds (shown by the dotted vertical lines), and use this particular time interval to form a new image. Actually, we could reduce the overall curvature of the phase even more, choosing the interval from 0.0 to 0.4 seconds. The image then is that already shown in Figure 2.8. The image thus is a compromise between crossrange resolution and complexity of the motion compensation. An analysis of the image responses shows that scatterer positions can be measured with sufficient accuracy (responses can be decomposed into one or two scatterer positions, with estimated uncertainties of a small fraction of a resolution cell).

2.2.7 Crossrange Resolution for Targets at Broadside Aspects

We have shown above that, for moving targets, crossrange resolution has a role subordinate to that of range resolution, even though an adequate degree of crossrange resolution is essential regardless of whether one wants to measure scatterer positions in range and crossrange or only in range. Crossrange resolution must be high enough that the image permits position measurements on enough scatterers to accomplish the task at hand. The number of scatterers depends on the task and the type of target; identification requires more than classification.

For targets such as aircraft, the important scatterers for identification are mostly on the fuselage, and the fuselage is long and narrow. When the aircraft is viewed near broadside, range resolution is not very effective on the fuselage, so that scatterer resolution depends primarily on crossrange resolution. To a lesser but still significant degree this is also true for ground vehicles. Although the ratio of width to length is not as small for ground vehicles as for an aircraft fuselage, at broadside the observable scatterers lie mostly along the illuminated edge. Hence, when the target is viewed near broadside, ground vehicle identification demands crossrange resolution comparable to range resolution. Since such a degree of crossrange resolution can seldom be obtained on moving targets, identification near the broadside aspect becomes problematic. The general conclusion is that, at least for moving targets, satisfactory identification performance is likely only when one need not place a high requirement on crossrange resolution.

2.2.8 Section Summary

Resolution requirements for target identification are driven by the need to measure the positions of target features to some accuracy. In principle, it does not matter what kind of resolution is used to resolve scatterers. In

practice, range resolution is primary, with (usually lower) crossrange resolution implemented only to assist range resolution. Problems arise when crossrange resolution replaces range resolution because the latter is not effective.

The primary function of crossrange resolution is to separate responses in the same range gate sufficiently to allow accurate range measurements. In the process, we also make crossrange measurements to a fraction of the crossrange cell. This usually allows crude crossrange resolution, with the burden of scatterer resolution carried by range resolution.

Since we address the measurement of scatterer positions, target images are presented in peaks plot form. Because a target contains a limited number of dominant responses, a peaks plot with too many peaks indicates poor image quality.

Crossrange resolution on moving targets is problematic, and the quality of the motion compensation must always be checked. This is done by finding a response dominant in its range cell, taking a transform, and verifying that the amplitude function has no deep nulls. The phase function then describes the residual motion. If such a response cannot be found, the imaging interval must be reduced and the motion compensation rechecked. The iteration of the process may lead to low crossrange resolution, which must be accepted.

Range resolution becomes ineffective near broadside aspects for targets with elongated shapes. This explains the difficulties of identifying targets near broadside.

2.3 Imaging Moving Targets

In this section we discuss imaging principles, then treat the important topic of motion compensation, providing the level of detail needed by someone working in this field. The reader only generally interested in target identification need not follow the more detailed treatments of signal-processing methods.

2.3.1 Imaging Principles

The fact that the theoretical and radar system work on target identification has been guided by quasi-optical thinking is probably related to the performance of ground surveillance SAR, which generates ground maps that approach the quality of a photograph. Optically inspired interpretation methods then are appropriate. However, metallic man-made targets do not

backscatter like trees, fields, roads, and other terrain features. Although the resolution performance of radar has been progressively improved over the years, with the hope that resolution performance might eventually be high enough to obtain images with photographic quality, this has not happened with man-made targets such as ground vehicles. Much effort has gone into focusing the intensity image, again borrowing optical terminology. The thinking was, if a SAR system performs the appropriate compensation of the platform's motion, the images of stationary ground targets will be well focused, and with sufficient resolution the targets should be identifiable. If a ground vehicle is moving, on the other hand, its image in the SAR scene will be highly smeared. Thus, we must compensate the motion of the ground vehicle in order to focus its image. For moving ground vehicles the problem has been thought to be developing the appropriate motion compensation, so that the image of a moving vehicle would be much like that when it is stationary.

We can only demonstrate that target identification is feasible with the methods described in this book; we evidently cannot prove that radar identification of a man-made target cannot be done in a way similar to the identification of a target from a photograph; that is, from the shape and details of its intensity image. However, based on our experience with real data from man-made targets, we believe that the quasi-optical approach cannot work if one is dealing with more than a small number of targets in the database; this is particularly true if reliable identification is to be fully automated. Instead of "focusing" the intensity image of a moving target, the aim must be to generate an image that allows the extraction of those parameters needed to identify the target. In an automated system we may never actually generate the final "image." After extracting the requisite target parameters from the data, there will generally be no need to form a corresponding image, the details of which are used for identification. With this understanding, in this book we will nevertheless follow conventional practice and use the term "target imaging" as if a "good" image were the final goal, but will do so merely because it allows us to simplify the terminology.

To restate an important point, *identification of a moving target cannot be considered as the identification of a stationary target after the appropriate motion compensation.* The situation is quite different. Some targets may indeed occasionally move so smoothly that a good motion compensation will generate the image of the target as if it were stationary. However, this case is so rare that no identification system can be designed on such a basis. With progressively worse forms of target motion, we lose the capability of making "focused" two-dimensional images, because crossrange accuracy may be

severely degraded. Crossrange resolution might be much worse than range resolution. Target outlines and shapes will not be recognizable by inspecting an intensity image, but then again the goal is target identification rather than imaging. In these cases we will use crossrange resolution to resolve scatterers within the same range gate to permit accurate range measurements, but will make little or no use of their crossrange positions. This is still much superior to the use of range profiles, which typically are interference patterns that do not allow accurate range measurements for more than a few strongly dominant scatterers. Crossrange resolution is needed to separate the responses in the same range gate.

The imaging of moving targets presents fundamental problems that cannot be solved with the conventional approaches. To explain the situation in more detail, we start with the simplest case of an aircraft flying at constant speed along a straight line in perfectly smooth fashion. That is to say, there is no yaw, roll, or pitch motion, nor any bouncing due to atmospheric turbulence. Unless the aircraft is headed straight toward the radar, the changing aspect angle causes the Doppler of the aircraft to change even though the aircraft may be going at constant speed. Assume that we are tracking a specific scatterer, with the signal-to-noise ratio very high. The phase function of the scatterer will be monotonically curved without being affected by anything but the translational motion of the aircraft. If we fit a sufficiently flexible polynomial or spline function to the phase function, and use the spline to compensate the scatterer response, it will be perfectly focused. We can use progressively longer processing times, so that the scatterer response becomes proportionally sharper, and yet it will remain perfectly focused.

Now, in a first step toward reality, assume that the signal-to-noise ratio of the scatterer response has more realistic values. Note that the cross section of such a scatterer typically is only a small fraction of the total cross section of the aircraft, and that radars do not employ a surplus of average transmitter power. Under these conditions, the phase function of the scatterer response will be modified by the noise. A spline function fitted to the phase thus does not provide a perfect measurement of how the phase of the scatterer response is changing, and hence how the range of the scatterer is changing. When the smeared response is compensated with this spline, it will not be perfectly compressed. If the phase measurement were affected by white Gaussian noise only, the fitting of the spline would average over the noise, effectively providing a large integration gain. In practice, however, the "noise" is mainly interference generated by other scatterers on the aircraft, because of inadequate resolution from neighboring scatterers and a background generated by all kinds of response, genuine as well as spurious. This interference has a

structure similar to that of the scatterer responses being tracked, so that it will falsify the measurement. This implies a degradation in the quality of the motion compensation and of the image.

There are other fundamental problems. Even when a target is moving relatively smoothly, there will often be disturbances in the motion. An aircraft might execute an abrupt inadvertent yaw motion, or it might be affected by atmospheric turbulence. Similar changes in motion occur with ships, and ground vehicle motion generally is even more erratic. It is impossible to compensate such abrupt changes in the motion to the degree that a good image is obtained. When a spline function is fitted to a phase function corrupted by noise and interference, the errors in the fit due to interference will generate rather high crossrange sidelobes. These will not be at the -30 dB level of the Doppler sidelobes introduced by the radar, but perhaps at -10 dB or even -3 dB. An additional problem is that a changing rotation axis tends to cause spurious responses in positions where there are no real scatterers, and these spurious responses can be strong and numerous. Flexing of moving targets is still another problem. Considering these difficulties of imaging moving targets, *our aim must be to generate images that permit one to extract those parameters that allow identification of the targets.*

2.3.2 Motion Compensation

Although the goal is not the generation of an image that looks good to the eye, but an image that allows one to extract identifying information, an appropriate motion compensation is still required. In [3] we describe a “perfect” motion compensation that will generate a high-quality image of a smoothly moving target, with high crossrange resolution. Our experience since then has shown that, more typically, the radar must operate under such severe conditions that one can use only some elements of such a sophisticated motion compensation. These elements must be used adaptively, guided by a real-time analysis of the motion behavior of the target as determined by imaging. In this section we will describe such a motion compensation in general, with more detailed illustrations given in the chapters on the identification of the various types of target.

Even in those applications where we must be satisfied with a motion compensation much simpler than the one that provides high crossrange resolution on smoothly moving targets, primarily applicable to aircraft under benign flight conditions, several motion compensation steps will normally have to be used. We generally distinguish between a motion compensation performed on the target as a whole and one that compensates a specific

scatterer on the target. The former must necessarily be much cruder than the latter. A different type of distinction depends on whether the motion compensation is based on range as measured by the pulse bandwidth or the range resolution cell, or whether the measurement is based on the signal phase. Since the range resolution cell ordinarily is much larger than the wavelength, the former compensation will be much cruder than the latter. When the range resolution cell is utilized, we speak of range compensation, and in the case of the phase we call it phase compensation. There is a third method between these two extremes, and this is Doppler tracking. Doppler tracking is based on phase, so that it is more accurate than range tracking, but since Doppler tracking involves averaging over a time window, it cannot follow time variations of the scatterer phase to the degree that is possible with phase tracking. As a last option, instead of tracking in range and in Doppler, we can track in both range and Doppler simultaneously. Such tracking is necessary when a range cell contains several scatterers of comparable strengths, in which case range tracking of a specific scatterer may not be possible. In the following we will explain how the various tracking methods are combined into a single motion compensation, given the motion behavior of a target.

2.3.2.1 Range Compensation

A radar is designed to detect a target of specified cross section at some maximum range. There are various detection schemes implemented in operational radars, and also various tracking schemes, so that one cannot make accurate statements that apply generally. However, for a rough estimate, we can reason as follows. If we use the entire target to perform the range tracking and compensation, it can be done at the maximum detection range. The subsequent step of the Doppler compensation on the entire target involves taking transforms over some time window, so that a coherent integration gain is achieved. This implies that Doppler tracking of the target could be even more readily performed at the maximum detection range than range tracking. An image with the standard motion compensation, range, and Doppler tracking of the entire target thus can be formed at the maximum detection range.

Identification depends on the visibility of individual scatterers, whose cross sections are considerably smaller than the cross section of the entire target. However, imaging involves coherent integration over many pulses, and the resulting integration gain allows the measurement of individual scatterer positions in the noise at about the maximum detection range. To illustrate, target detection might be performed with a burst of perhaps 5 to 10 ms duration. If we form an image using the same signals at the same PRF and with an

imaging time of 200 to 400 ms, we have an integration gain of 40, so that the signal-to-noise ratio will be adequate for scatterers whose cross sections are one-fortieth that of the entire target. The conclusion is that target identification should be achievable at about the maximum detection range.

The preceding reasoning is correct when the motion of a target is very smooth, so that the standard motion compensation is adequate and only low crossrange resolution is needed. As will be shown in the chapters on identification, targets usually do not move so smoothly that forming an image over some time with the standard motion compensation is adequate, in which case it is essential that one have the ability to track individual scatterers rather than only the target as a whole. This tracking is performed by taking the transform of the image cut in the range gate of a response, so that the behavior of the phase function may be examined. If the transform window extends only over the main lobe of a compensated response, the coherent gain from imaging is preserved. However, such a transform window does not include the high frequencies of the response, so that the variations of the transform amplitude and phase functions are smoothed. The phase function then does not allow determining whether or not the motion of the target changes abruptly at some time, which would adversely affect imaging and identification. Thus we must extend the transform window over at least about 10 crossrange gates, which implies that the integration gain from imaging is reduced by at least a factor of 10. For smooth target motions, we concluded that identification can be performed at about the maximum detection range. With a reduction of the integration gain by a factor of 10, from the fourth-power range law we find a reduction of the identification range by a factor of 1.78, so that identification can be performed at about half the maximum detection range.

Based on this discussion, one must adopt the following practical view. Target identification is considerably more difficult than target detection. If it is important that target identification be performed at the maximum detection range, more signal energy must be provided for identification than can be obtained by continuing to transmit the pulses used for detection at the same PRF over the imaging interval. If this is not acceptable, then target identification cannot be achieved near the maximum detection range. These are questions of policy that are beyond the scope of this book. We know that it is necessary to track individual scatterers, and in the following we will ignore questions concerning the relation between maximum detection range and maximum identification range.

To demonstrate the various steps of the motion compensation, we arbitrarily select a two-second data segment collected on an aircraft. In an

operational system, the radar would collect data over some minimum time interval and, in real time, go through the motion compensation steps and form an image while the data collection continued. The image would be analyzed in real time in the same way as will be demonstrated below. If the decision is made that more data are needed, the motion compensation is extended over a longer interval, another image is formed, and the analysis is repeated. This process is continued until an image of the desired properties and quality can be generated.

In Figure 2.15 we show a sequence of range profiles received from an aircraft over two seconds, displaying every second range profile. This particular radar uses a range window that extends over 64 range gates. The abscissa of Figure 2.15 depicts the 64 range gates that cover the range window, and the ordinate gives running time. Because of the high range rate of the aircraft, the individual range profiles drift quickly through the range window and are folded every time they move through the 64 gate range ambiguity. The first




Figure 2.15 Sequence of received range profiles.

processing step is to remove the range drift, which requires measuring the range drift. This can be done by any of several methods. We can measure the range centroid of each range profile, we can correlate one range profile with the next, or we can track a prominent response peak, in each case using smoothing to average over the measurement errors. Still other methods may be used for this crude compensation step. The important point is that *the curve fitted to the sequence of range measurements must not be too flexible*, since it might too closely follow spurious motions introduced by the measurement errors. This would be disastrous for the range compensation step, because the range changes involved are too large to be properly corrected by the following Doppler compensation. A linear fit is the correct choice for this step, even when the observation interval is as long as used for this illustration. Any nonlinear motion can be removed in the following step.

In Figure 2.16 we show the sequence of range profiles after the constant range drift was removed by correlation. The range drift has largely




Figure 2.16 Range profiles after removing the drift.

disappeared, so that one can see the range-cell wander caused by the fact that the range rate varies with time. Since in this instance each range profile is generated by a fairly long linear frequency sweep, in the order of milliseconds, the motion compensation implied in the range alignment also somewhat changes the shape of each range profile.

The same method used to remove the constant range drift will generally also be used to remove the residual nonlinear range-cell wander in Figure 2.16. It again is critical not to use too flexible a fit, with the proper choice a quadratic fit when the residual range wander is monotonic. (The adequacy of the fit, and whether too flexible a polynomial has been used, can be tested by taking the transform of a response after range and Doppler compensation, as will be illustrated later.) Actually, there is no need to split the range compensation into two steps, the first for removing the constant drift and the second for eliminating the range-cell wander due to the changing range rate. We split the process into two parts merely because we want to consider the important task of single-scatterer tracking, in which case for illustration purposes we must remove the constant range drift, at least approximately. Whereas the removal of the range drift can be done in any one of the ways mentioned above, we want to remove the residual range-cell wander by range tracking a single scatterer. This is an extremely important step in the motion compensation and imaging process, and is considered next.

In order to track single scatterers, we start with the succession of range profiles after the gross range drift has been removed. If the response peaks of a range profile truly represented the responses from single scatterers, tracking a particular peak of the range profile would be tantamount to tracking a particular scatterer. Since the range profile is an interference pattern, single-scatterer tracking is more difficult. This can be appreciated from Figure 2.17, which shows the ranges of the 17 strongest peaks of each of the received range profiles. If these peaks represented true scatterer positions, Figure 2.17 would show a series of continuous tracks over the entire two-second interval. In fact, there is not a single complete track visible in the figure. The plot is another verification that range profiles represent interference patterns that change rapidly with aspect angle. There are three different methods of dealing with the problem and establishing a scatterer track.

The crudest method is based on the peaks of the range profiles, that is, on the "peaks tracks" of Figure 2.17. We can search for a reasonably good track, perhaps with some gaps that can be bridged by fitting a polynomial, or a spline function consisting of polynomials with break points. The best candidate for such a track is the one starting in Range Gate -13.4 in




Figure 2.17 Peaks tracks corresponding to Figure 2.16.

Figure 2.17, indicated by the upper arrow. If such a relatively good track cannot be found, we can search for one that alternates between two range gates, as occurs when two primary scatterers are interfering with each other so that the response peaks are continuously shifting back and forth. An example is the pair of tracks starting in Range Gates -22 and -24 , indicated by the lower arrow. As explained earlier, resolution based on the intensity output is degraded by a factor of two with respect to the inherent resolution of radar, as provided by the complex processor output; because the peaks tracks of Figure 2.17 are based on the intensity range profiles, range resolution is degraded. Using the peaks tracks thus is a relatively crude form of scatterer tracking.

To facilitate range tracking, we should utilize the complex range profiles. The responses of the complex range profiles should be analyzed with the one-dimensional TSA to determine actual scatterer positions, and in effect the peaks tracks of Figure 2.17 should be replaced by tracks of actual scatterer positions. Such a utilization of the inherent range resolution capability will significantly improve the trackability of individual scatterers. This fact is illustrated below, but since the method has not been integrated into the interactive software used for the demonstrations in this book, we will continue this example with the simple peaks tracks. Simple tracking of the peaks of the intensity range profile has proved adequate for all demonstrations

given in this book. However, for fully automated tracking we want to use the most robust method.

Even when the peaks of the complex range profiles are analyzed to obtain actual scatterer positions, range resolution is only improved by a factor of two. The improvement will not be sufficient for certain situations, in particular when the aspect angle of a target becomes so large that range resolution is ineffective along the long dimension of the target. The problem then is that too many significant scatterers are within the same range cell. It becomes necessary in this situation to resolve the scatterers in the same range cell in Doppler, at least to some degree. Instead of range tracking we then use combined range/Doppler tracking. This tracking method was described in [3], and it is also demonstrated below. As is the case with range tracking of scatterer positions rather than peak positions, it has not been integrated into our interactive software, only into automated identification software. Thus, for our demonstrations we will use only the crudest range-tracking method, based on the peaks of the intensity range profile, and this range track will be followed by a Doppler track of the scatterer. With the third method we would not break the scatterer track into these two steps but use a combined range/Doppler track instead. Note, however, that the improved tracking methods (utilizing complex range profiles and combined range/Doppler tracking) still do not solve the problem of scatterer tracking when the aspect angle is so large that range resolution is not effective along the long dimension of the target.

For an illustration of the simple peaks tracking procedure, we select the peaks track starting in Range Gate -13.4 in Figure 2.17, and fit a quadratic polynomial as indicated in Figure 2.18, bridging one gap and smoothing the variations in the measured peak positions. These variations are not true range variations of the scatterer, but are caused by the changing interference conditions. We again point out that the fitted curve must never be so flexible that it can follow these small scale variations. It would make the next compensation step, tracking the scatterer in Doppler, impossible. However, the fact that Doppler tracking becomes impossible serves as a safeguard and prevents us from going on with the motion compensation and generating a meaningless image. When range tracking a scatterer, we must use the rule that the spline function only be flexible enough to take out the smooth trend. In cases where the situation is not as clear, we must check the success of each processing step. This will be demonstrated below.

When we remove the delay measured with the polynomial fit from each of the successive range profiles, thus motion compensating the data, we obtain the range profiles of Figure 2.19. If any of the alternative methods




Figure 2.18 Peaks track and polynomial fit.




Figure 2.19 Range profiles after range compensation.

discussed above are used to eliminate the entire range cell wander in one step, not just the linear drift, the resulting range profiles are indistinguishable from those of Figure 2.19. To illustrate another range track for Figure 2.17, in Figure 2.20 we show a peaks track for the case of two interfering scatterers (Although the range gates agree with those of Figure 2.17 in this case, in other figures they may not, because the range profiles as a whole may have been shifted.) With this lower quality track of a doublet of unresolved scatterers, the important point is that the indicated rapid range variation of the peak positions can extend only over a range interval within which two unresolved scatterers can alternately generate a single peak or two peaks. This range interval is at most about two range gates wide. When the particular range track of Figure 2.20 is used to compensate the data, the further motion compensation steps lead to an image that provides the same information as one based on the simpler peaks track of Figure 2.18.

This is generally not the case when rapid range variations like those in Figure 2.20 are larger than about two range gates, because more than two scatterers are involved, but this again would be recognized in the further processing steps. Again, this is so because of the crudeness of range measurements, where errors easily become so large that they cannot be corrected in the next motion compensation steps. Even the first processing step following range tracking—that is, Doppler tracking—might not work in the sense that no Doppler function would be obtained to which a spline function could be




Figure 2.20 Peaks track for two interfering scatterers.

meaningfully fitted. As an illustration, in Figure 2.21 we demonstrate what not to do; that is, fit through peak positions that vary over more than two range gates. There is no justification for this type of “range track.” In this instance, continuing the motion compensation by force resulted in an image with grossly smeared responses. However, although the quality of the peaks track can be verified in this manner, in critical situations where good peaks tracks cannot be found, we will use one or even both of the improved tracking methods discussed above, rather than peak tracking.

We will now demonstrate that range tracking can indeed be improved by analyzing the complex range profiles and tracking actual scatterer positions rather than the peaks of the intensity range profiles. However, since the corresponding software has not been integrated into the interactive software used for our examples of target imaging, this demonstration will be made for a different target, after which we will continue with the motion compensation for the aircraft.

In Figure 2.22 we show the peaks tracks for a ship over a time interval of 10 seconds. With a PRF of 200 Hz and an interval of 10 seconds, this peaks track plot contains the peaks of 2,000 intensity range profiles. The ship is executing a rather sharp turn, and the time of the specific interval of Figure 2.22 is shortly after the ship is at broadside aspect. This combination of strong motion during the turn and proximity to broadside makes this example one of the more difficult cases. There is not a single good track in




Figure 2.21 Wrong peaks track.




Figure 2.22 Peaks of intensity range profiles over 10 seconds.

Figure 2.22. We shall demonstrate that the use of the complex range profiles helps tracking, even though we cannot make a quantitative statement as to how much the tracking performance can be enhanced by utilizing the factor of two resolution improvement from the complex range profiles.

As a first example, in Figure 2.22 consider Range Gate 25 at the time of 321.4 seconds, indicated by the x in the left dashed vertical line. About one second later we observe a strong scatterer in that range gate, not apparent at the earlier time. The interesting section of the range profile at this particular time is shown in Figure 2.23. Comparison of the range profile with the dashed point-scatterer response clearly shows that Range Gate 25 does contain an unresolved response, which will be resolved if the TSA is used on the complex response. The same unresolved response also can be found at the time 321.8 seconds, halfway in time between the preceding measurement and the time when the scatterer becomes visible in the peaks tracks.

As a second example, consider the peaks track marked by the upper arrow, which starts in Range Gate -6 of Figure 2.22 and splits into two tracks at a time of 322.5 seconds. Arbitrarily choosing the time of 323.1 seconds, in Figure 2.22 we find a track in Range Gate -8 and a neighboring one in Range Gate -11 (indicated by the x in the right dashed vertical line), with a gap where the continuation of the earlier track should be. The range profile for this time is shown in Figure 2.24, with the dashed point-scatterer curve




Figure 2.23 Part of the range profile at 321.4 seconds.




Figure 2.24 Range profile at a time of 323.1 seconds.

overlying the response of interest. The transform of the clearly widened response, between the vertical crosshairs, is given in Figure 2.25. This is the amplitude/phase pattern from two interfering scatterers, one of them in the range gate of the lost track.




Figure 2.25 Transform of the response in Figure 2.24.

This type of enhanced range tracking will fail when it becomes impossible to define transform windows that include only two responses rather than three or more responses. In other words, compared with the intensity range profile, we cannot improve range resolution by more than a factor of two, possibly a little more than two by using a small degree of superresolution [3]. However, improving resolution by a factor of two can greatly extend the trackability of scatterers to worse motion behavior or larger aspect angles.

The examples demonstrate that range tracking of scatterers can at least be extended if one switches from tracking the peaks of the intensity range profile to tracking the scatterer positions as measured by the one-dimensional TSA. Even under relatively benign conditions, there will be cases in which two scatterers of comparable strengths are so close in range that individual tracking will not be possible. As already explained, for tracking purposes the combination of the two scatterers will act as a single scatterer with fluctuating cross section, so that it may have to be tracked through its repeated cross section minima. As long as the combination of the two unresolved scatterers can be tracked at all, it will serve the purpose of measuring the target's motion. If it cannot be tracked over the entire observation interval, another scatterer must be selected.

An illustrative example is given by the closest peaks track in Figure 2.22, starting in about Range Gate -26, marked by the lower arrow. A polynomial fit to this intermittent track (again marked by an arrow) is shown in

Figure 2.26 on an expanded scale, with the polynomial bridging the two gaps in the track. Note that the extremely long observation time chosen for this example requires the use of a polynomial of an order higher than two. The polynomial must have sufficient flexibility to follow the curved peaks track around the time of 327 seconds, but without being so flexible that it will deviate too much during the gap. The implication is that the acceptable length of a gap becomes shorter as the range variation becomes more nonlinear; that is, with a more irregular motion of the target. The acceptability of a track is determined by examining transforms of image cuts in fixed range gates after the data have been compensated. If we cannot find transform amplitudes consistent with one-scatterer or two-scatterer patterns, the track is not usable. Another track or a reduced duration must be used. In our example, the data are compensated with the fitted polynomial, and another compensation is performed in Doppler, as usual. When an image is formed over the entire interval displayed in Figure 2.26, and a transform of the image cut in the range gate of the compensated scatterer is taken, we obtain Figure 2.27. This is essentially the amplitude/phase pattern of two interfering scatterers viewed over almost two interference cycles, with a high frequency modulation from other scatterers superposed. Aside from the phase variations due to the drop in the amplitude, the phase function is essentially flat. This means that the compensation of the scatterers is very good, better than needed for measuring the motion behavior.




Figure 2.26 Polynomial fit for an intermittent peaks track.




Figure 2.27 Transform of the image cut in the range gate of the intermittent scatterer.

2.3.2.2 Doppler Compensation

After the preceding demonstration of improved range tracking on an entirely different target, we return to the motion compensation of the aircraft, where we completed the range compensation, as shown in Figure 2.19. The next compensation step is based on Doppler tracking the entire aircraft or some suitable scatterer. The usual method for Doppler-tracking the entire aircraft is to track the Doppler centroid, measuring the Doppler of the aircraft over some time window and sliding the window over the observation interval.

Range resolution is degraded for this purpose so that the aircraft is no longer resolved in range. For example, with frequency-stepped waveforms the Doppler centroid tracking can be performed on each of the frequencies (after range filtering, if the unambiguous range window is much larger than the target), and the results (actually, the range rates) are averaged. If the signal-to-noise ratio for a single frequency is too low despite the integration gain from Doppler tracking, one can Doppler track a single scatterer using the entire frequency band, instead of tracking the target as a whole. For any

waveform, one can also form short-term subimages and track the Doppler centroid of the subimages. The choice of the particular method is not critical for imaging purposes as long as it deals adequately with the system noise. In our example, the range profiles after Doppler centroid tracking are nearly indistinguishable by eye from those of Figure 2.19, so that the resulting range profiles are not shown.

The two-second image after range and Doppler tracking (both implemented by centroid track), which we might call the standard motion compensation, is shown in Figure 2.28. Even to the eye, there are too many peaks for this to be an image of acceptable quality, since the number of peaks by far exceeds the number of observable scatterers on an aircraft fuselage. As will be shown below, examination of individual complex responses shows that the image quality must be further improved. We have also considered the alternative method of range tracking a specific scatterer rather than the entire




Figure 2.28 Image after range and Doppler track of the entire aircraft.

target. In Figure 2.18 we showed the range track of a scatterer. This scatterer then is Doppler-tracked as shown in Figure 2.29. After compensation with the polynomial fitted in Figure 2.29, we obtain the image of Figure 2.30. Although this image looks rather different from that of Figure 2.28, only further analysis will show whether one is better than the other. As it happens, neither is acceptable at this stage.

We note that there is a general ambiguity problem for the Doppler track. The Doppler measurement implied in the Doppler track is ambiguous, with the ambiguity determined by the PRF of the radar. This Doppler ambiguity must be resolved by means of the range measurements based on the signal bandwidth, which is automatically done in the range compensation process. However, a range measurement based on the range resolution cell is a crude measurement that will reliably resolve the Doppler ambiguity only when the PRF is fairly high. In practice, the operational parameters of a radar can easily be such that one cannot depend on a correct resolution of the Doppler ambiguity by the range track. If the Doppler ambiguity was not correctly resolved, then after the Doppler track the range profiles will not be aligned but will drift in range at a rate corresponding to a Doppler ambiguity. In systems where this possibility exists, we must always recheck the alignment of the range profiles after Doppler compensation. Any residual drift of the range profiles due to a Doppler ambiguity must be removed by




Figure 2.29 Doppler track of a scatterer.




Figure 2.30 Image after range and Doppler track of a scatterer.

performing a new range track and compensating the data with the Doppler ambiguity closest to the range rate determined from the track.

We have now discussed the range compensation and the subsequent Doppler compensation for an aircraft. Aside from this main theme, we also demonstrated that range tracking can be improved by switching from tracking the peaks of intensity range profiles to the actual scatterer positions found by analyzing the complex range profiles. With increasing irregularity of a target's motion and the loss in range resolution as the target aspect approaches broadside, even this type of improved tracking may not be good enough. Then we must use combined range and Doppler tracking of scatterers for a further extension of performance. As already stated, at the time of this writing we could not do this for an entire imaging example. To demonstrate combined range/Doppler tracking, we will thus return to the example of the ship used to demonstrate improved range tracking.

From Figure 2.22, we select a peaks track that disappears into a mass of responses distributed over a wide range interval. To facilitate the correlation of the measurement with the peaks tracks, in Figure 2.31 we show an expanded part of Figure 2.22, after removal of the slow trend. As a consequence of a compensation, the peaks track to be examined (marked by the arrow) now is in about Range Gate 0. Combined range/Doppler tracking is equivalent to forming a sequence of short-term images in order to resolve the scatterers in both range and Doppler. The question is whether or not a particular response is trackable in the sequence of images.

Figure 2.32 gives a 0.2-second image of the ship at the time of 322.7 seconds. The track is still recognizable in Figure 2.31 during this interval, indicated by the leftmost dashed column of the figure. The response of interest in Figure 2.32, indicated by the crosshairs, is in Range Gate -0.0 and Crossrange Gate -0.4 . At this imaging time the response is fairly strong and well isolated in its range gate. The same type of image 0.2 seconds later (the central column of Figure 2.32) is shown in Figure 2.33. The response is still in Range Gate 0.0 , but has shifted slightly in crossrange, to Gate 0.0 , and has become weaker. Another image made 0.2 seconds later (the right column of Figure 2.32) is shown in Figure 2.34. The response has again shifted only slightly in range and crossrange, to Range Gate 0.3 and Crossrange Gate -0.3 . A continuation of the short-term imaging over the entire time interval of Figure 2.31 verifies that the response of interest can be tracked over the




Figure 2.31 Expanded part of Figure 2.22.




Figure 2.32 Image from 322.6 to 322.8 seconds.




Figure 2.33 Image from 322.8 to 323.0 seconds.




Figure 2.34 Image from 323.0 to 323.2 seconds.

entire interval, even though this is not possible with the intensity range profiles of Figure 2.31. Range/Doppler tracking thus is clearly superior to tracking first in range and then in Doppler, because range tracking eventually fails when a situation becomes more difficult and range gates do not contain single dominant scatterers, even when range tracking is performed on the complex range profile.

2.3.2.3 Phase Tracking of a Scatterer and Imaging Interval Selection

After the above digression into combined range/Doppler tracking, and before we return to our example of aircraft imaging, we want to discuss phase tracking in general. Phase tracking is the most critical step in imaging with a quality that allows the measurement of scatterer positions and characteristics. In order to phase track a scatterer, we must measure its phase function to determine any motion compensation residual. We also want to find any disturbances in the motion of the target that may prevent obtaining an image of sufficient quality at a particular instant. If we start with a short imaging interval, as will the processor in a fully automated system, the decision must be made whether or not to increase the imaging interval. If we start with a long

imaging interval, as is the case in our example, the decision must be made whether to improve the motion compensation in order to increase the image quality, or reduce the imaging interval in order to obtain an image of acceptable quality without further motion compensation, but with reduced cross-range resolution. This decision must be based on the motion behavior of the individual scatterers.

This type of phase tracking of a specific scatterer may not be a simple matter. We must find a range gate in which a scatterer is sufficiently well resolved so that the transform of the image cut in the range gate of the scatterer gives a reasonably constant amplitude function. Otherwise the phase variations associated with the amplitude modulation introduced by other scatterers obscure the motion of the scatterer. The transform window must be much wider than the response from the scatterer, because a narrow window cuts off the high frequencies that let us measure any abrupt changes in the phase function that indicate an erratic motion behavior of the target. The scatterer must be relatively strong, because otherwise the background interference will be strong enough to modulate the amplitude function to an unacceptable degree. Also, the motion compensation that has been performed up to this point may have particularly well compensated this one scatterer, so that an erratic behavior of the aircraft may be partly obscured. Thus we must apply the test to a second scatterer as well, and that scatterer should be as far away as possible in range from the first scatterer. Two scatterers that meet these requirements typically are not easy to find in an image.

The specific requirement on an acceptable response for the phase measurement can be stated as follows. Find a range gate with a response that has a good signal-to-background ratio, without significant responses closer than about five crossrange gates. By "response" we mean the response associated with a single scatterer, and at this stage of processing such a response might be smeared over several crossrange gates if the motion compensation is not good. A smeared response from another scatterer could be centered five crossrange gates away and yet overlap the adjacent response significantly. This would be recognized from the amplitude function of the transform over both smeared responses, in which case phase-slope tracking would be used. In any case, a crossrange interval of ± 5 gates is about the minimum needed not to cut off too much of the high frequencies. As stated, a usable response or, at most, a response doublet must be found in at least two range gates that are not too close together. We will illustrate the phase tracking of a single scatterer sufficiently frequently in our examples to forgo an illustration at this point. Instead, we will now explain the second method of phase-slope tracking, because it has not been integrated into the interactive software used in

our demonstrations and hence will not be used in many of our examples. Nevertheless, it is very important because it permits phase tracking of a specific scatterer when a second significant scatterer is interfering.

The amplitude pattern of two interfering scatterers has periodically-spaced maxima and minima. If one of the two scatterers remains stronger over the entire observation interval, at the time when both phase vectors add (time of an amplitude maximum), the phase of the combined signal is the phase of the stronger signal. Overall, the phase function of the combination of the two signals is governed by the phase of the stronger signal. This is the so called strong-signal phase capture effect. Hence, the slope of the phase function at the time of the amplitude maximum also will be governed by the phase slope of the stronger scatterer. If we *measure the phase slopes at the times of the amplitude maxima* and fit a smooth polynomial or spline function to the succession of phase slopes, we will have the phase slope of the dominant scatterer. Integration of the phase-slope function then gives the phase function itself.

The process is illustrated in Figure 2.35. On top is the amplitude function, the nearly periodic modulation of which indicates the presence of two




Figure 2.35 Phase-slope tracking.

dominant scatterers. Since these are real rather than simulated data, the amplitude function includes distortions from other than the two main scatterers. Thus, it is not entirely periodic, as it would be for only two scatterers. Nevertheless, as an approximation we measure the phase slopes at the times of the amplitude maxima. At the bottom, the curve with the phase jumps is the phase function associated with the amplitude function. It includes the phase modulation due to scatterer interference, primarily the phase jumps, as well as the slowly varying phase due to the motion compensation residual. The nonlinear function that starts at the lower left corner of the figure is the phase function of the stronger scatterer, as obtained by fitting a polynomial to the succession of phase slopes and integrating. We note that the measured phase slopes within the intervals defined by the crosshairs agree with the slopes of the phase function at the same times. The algorithm should include a correction that accommodates the effect of the weaker scatterer on the phase-slope measurement at the times of the amplitude maxima, but we have not incorporated this correction into our software.

With this extension of the phase measurement to two scatterers within the transform window, we can significantly increase the range of conditions under which phase tracking in two range gates is possible. Still, under some conditions, we may not find two range gates that allow phase or phase-slope tracking, such as when the aspect angle approaches broadside. At large aspect angles, the responses in the image after the standard motion compensation may still be quite smeared. The responses then overlap in crossrange, so that no transform window can be defined that meets the conditions for the two forms of phase tracking. The remedy is to shorten the imaging interval. With a shorter imaging interval the smearing of the responses due to the motion residual will be lower. With this discussion of the important phase-slope tracking procedure, we return to the image after standard motion compensation to demonstrate the next processing step.

In Figure 2.36 we show the transform of the tip response of Figure 2.28. For purposes of phase tracking, we consider the amplitude function “reasonably constant” over the interval between dotted vertical lines, from a normalized time of -0.4 seconds to 0.47 seconds. The phase function has a smooth trend over this interval, so that a further motion compensation would be possible. As pointed out, however, we must examine at least one other range gate, not too close to the first. The transform of the image cut in Range Gate 14.9 of Figure 2.28 is shown in Figure 2.37. In this range gate the amplitude function is reasonably constant only over much shorter time intervals. If we add the requirement that the phase function have a smooth trend over these time intervals, the longest usable time interval extends from




Figure 2.36 Transform of the tip response in Figure 2.28.




Figure 2.37 Transform in Range Gate 14.9 of Figure 2.28.

about -0.25 to 0 seconds, as shown by the dotted vertical lines. To demonstrate that there is nothing accidental about such a test, in Figure 2.38 we show the transform of the image cut in Range Gate 16.5 of the same image, with the longest usable interval again shown by the dotted vertical lines. A comparison of Figures 2.37 and 2.38 shows that the best intervals in the two




Figure 2.38 Transform of the image cut in Range Gate 16.5 of Figure 2.28.

figures overlap almost completely. The image over this shortened interval, which in real time extends from 6.5 to 7.0 seconds, is shown in Figure 2.39.




Figure 2.39 Image from 6.5 to 7.0 seconds.

The image of Figure 2.39 was obtained by range and Doppler tracking the entire aircraft, then selecting a specific imaging interval via phase tracking of a scatterer. We also generated the survey image of Figure 2.30 by range and Doppler tracking of a scatterer rather than the entire aircraft. Starting from this alternative image, we can phase track two scatterers, select a good imaging interval, and again form an image that corresponds to that of Figure 2.39. An image cut in the range gate of the tip scatterer of Figure 2.30 now shows that essentially the entire interval could be used. The problem at the beginning of the interval is less severe than in Figure 2.36. When an image cut in Range Gate 0 is taken in the image of Figure 2.30, we obtain the transform of Figure 2.40. The combination of amplitude function and phase function indicates an allowable imaging interval (between dotted vertical lines) from about -0.2 to 0.05 seconds, almost in perfect agreement with the result from Figure 2.38. The two images of Figures 2.28 and 2.30 thus are equivalent in that they lead to the same result, as one would expect.

The image over the slightly shifted interval from -0.2 to 0.05 seconds is shown in Figure 2.41. One cannot tell without examining image responses whether this or the image of Figure 2.39 has higher quality. Since it was possible to find range gates in which transforms contained intervals with constant amplitude and smooth phase, the conclusion must be that the images are equivalent in allowing the measurement of scatterer positions. Also, when an image is formed after range and Doppler tracking of a scatterer and exactly the same time interval as used for Figure 2.39 is selected (which came




Figure 2.40 Transform in Range Gate 0 of Figure 2.30.




Figure 2.41 Image from 6.6 to 7.1 seconds.

from tracking the entire aircraft rather than a single scatterer), no significant difference can be detected by eye. Thus, at least in this case, the two types of motion compensation are entirely equivalent. Moreover, if they were not equivalent, we would have discovered problems. This is to say that with proper imaging procedures we will always know whether the image has a good enough quality that one can measure scatterer positions. It is an important point for target identification. Indeed, an examination of the responses of the images of Figures 2.39 and 2.41 shows that they pass the quality test: they are sharply focused and have low Doppler sidelobes.

The preceding results lead to an important conclusion. *For moving targets, one cannot specify some desirable degree of crossrange resolution and choose the imaging interval accordingly.* We may define a desirable crossrange resolution, but if in the systematic imaging procedure it proves impossible to obtain response transforms with constant amplitudes and linear phase functions over the time interval needed to achieve the desired crossrange resolution, we must accept the achievable lower crossrange resolution. The

practically achievable crossrange resolution will be different for each type of target and will also depend on the aspect angle and the particular motion conditions. This means that crossrange resolution may on some occasions be so poor that identification must be based on the range positions of the scatterers, with crossrange resolution used only to separate scatterers in the same range gate, so that their range positions can be measured with better accuracy.

2.3.2.4 Practical Aspects of Motion Compensation

It is natural to reason that when a target is moving, its two-dimensional image will be smeared, so that we must use some kind of motion compensation to obtain a “good” image. We summarized the motion compensation methods of primary interest. However, *the way a motion compensation is structured, or whether it is or is not used, depends entirely on the type of target and its behavior.* We will now discuss these issues, starting with aircraft. More detail is given in subsequent chapters.

In the absence of any pseudoperiodic yaw motion or air turbulence, the aspect angle of an aircraft changes fairly smoothly and slowly as the aircraft proceeds along its flight path. No motion compensation (other than polar reformatting to keep all scatterers within their range gates) would be needed if the aircraft appeared to rotate at a constant rate, since that would imply constant Dopplers of the scatterers; the purpose of the motion compensation is to have all scatterers appear to move with constant Doppler. Since the apparent rotation rate of an aircraft is almost never constant, we must resample the data before an image can be formed [3], unless the imaging interval is so short that the change of the Doppler over the imaging interval is negligible. This would ordinarily imply very low crossrange resolution. Then the standard motion compensation, range tracking followed by Doppler tracking of the entire aircraft or range and Doppler tracking of a specific scatterer, will be satisfactory when the aspect angle rate is not constant but smoothly changing. On the other hand, aircraft often have an inadvertent yaw motion even when they are not intentionally maneuvering. The motion compensation must take out the combination of the Doppler variations due to the aspect angle change from the motion along the flight path with the Doppler variations due to inadvertent yawing. Any roll motion is not significant if only the scatterers on the fuselage are utilized for identification (but the vertical stabilizer may pose problems). If the Doppler variation caused by the inadvertent yaw provides insufficient crossrange resolution for imaging, the demands on the motion compensation can be rather stringent. This is particularly the case when the pseudoperiodic yaw motion is not smooth. On

the other hand, in a fraction of the imaging situations, the pseudoperiodic yaw motion generates aspect angle changes large enough for two-dimensional imaging, in which case a relatively simple motion compensation will suffice.

In summary, imaging of an aircraft requires a motion compensation that can vary from simple to very sophisticated, depending on the flight behavior of the aircraft and the desired crossrange resolution. If there is an abrupt change in the range rate, a satisfactory motion compensation may not be achievable over any interval containing the change. We must exclude such times from imaging, even when the resulting crossrange resolution becomes lower than desired. The positional match then relies on the ranges of the scatterers.

The situation is quite different with ground vehicles. When a ground vehicle moves very slowly on a smooth surface, motion compensation requirements are similar to those for aircraft, and they can also be met. On the other hand, in most cases ground vehicles move so irregularly that a motion compensation that properly compresses the responses and generates a two-dimensional image usable for vehicle identification will not be possible. As is the case with erratic aircraft motion, we must be satisfied with a motion compensation that takes out the gross motion of the vehicle but fails to compress the responses in crossrange. Again, only the ranges of the scatterers are usable for the positional match. This situation will arise almost routinely, whereas it is the exception with nonmaneuvering aircraft.

The situation is again very different with ships. Whereas the inadvertent yaw motions of aircraft often are so small that they are not usable for crossrange resolution and only make the motion compensation more difficult (or force one to form images with low crossrange resolution), the yaw, pitch, and roll motions of a ship are easily large enough to allow imaging without help from the translational motion. Let us consider a small ship with a length of 100 ft and a superstructure height of tens of feet. We might want a crossrange resolution of, say, 5 ft. At X-band the required aspect angle change then is about half a degree. However, the yaw, pitch, and roll of such a ship cause aspect angle variations that typically are much larger. Then we can select imaging intervals much shorter than the motion period, and select them so that the yaw, pitch, and roll rates are constant over the entire interval. We also do not need any motion compensation other than removing the range drift. The changes in the aspect angle may be large enough to cause the scatterers to move through range gates, in which case we must employ polar reformatting to keep them within their range gates.

In conclusion, the motion compensation requirements vary greatly from one type of target to the next, and for a given target they depend on

motion behavior and aspect angle. They vary from almost no requirement for motion compensation to one, at least in principle, for a sophisticated motion compensation. However, a sophisticated motion compensation may not be usable when the target's motion is not very smooth, because compensation may not be possible at these times without degrading the image quality by distorting responses and generating spurious responses. All these items will be extensively illustrated in the chapters on imaging of the various targets.

2.3.3 Motion Determination

Targets rarely move so smoothly that a usable image can be generated from the entire radar dwell. As discussed above, we must routinely select a subinterval of smooth motion. Then we can form a high-quality image, in which crossrange resolution at least helps resolve scatterers at the same range, and often enables us to accurately measure crossrange positions of scatterers. If the target motion permits, we also minimize spurious responses in the image, by restricting imaging to intervals when the target has a fixed rotation axis. In some applications, in order to physically interpret crossrange measurements, we image the target only when it rotates about a particular axis.

All these determinations of imaging intervals require measurement of the target's motion, which requires measurement of the motions of individual scatterers. We measure each scatterer's motion by tracking it in accordance with the motion compensation procedures discussed in Section 2.3.2. In order to judge the quality of any given track, and the residual uncompensated motion of the scatterer, we compensate the tracked motion, form an image, and examine a fixed-range cut through the scatterer. If we can choose a crossrange window about the scatterer whose Fourier transform has a constant amplitude, the corresponding phase directly gives the residual scatterer motion. This same criterion applies to other scatterers in the image: if the transform amplitude is constant, the transform phase gives the residual scatterer motion.

In practice, the transform amplitude will never be perfectly constant. This raises the questions of how much variation is acceptable and how it can be measured. As one might expect, the acceptable variation depends on the application. Perhaps surprisingly, so does the appropriate measurement procedure. The most demanding applications (primarily motion compensation) are those in which we must track the phase of the transform. Here we must be wary of phase distortions from even relatively weak interfering scatterers. The allowable amplitude modulation is small, and the amplitude cannot

approach zero at any time. The next most demanding application (determining intervals of smooth motion) requires that we recognize when large phase fluctuations or abrupt changes in phase slope occur. Then we can tolerate stronger amplitude modulation. The least demanding application (coarsely determining the relative motion of widely separated scatterers in ship identification) requires only that we estimate the size of the residual phase fluctuations. For this purpose, we can accept the amplitude so long as it does not stay near zero for an extended period.

We have found two useful measures of the amplitude variation, one dependent on the slow variation and the other on the fast, but do not claim they are ideal, in either general approach or specific implementation. We distinguish the two variations by calculating a moving average of the amplitude, using a window width equal to the lesser of one-tenth the transform duration or half the mean modulation period of the strongest significant amplitude modulation. The modulation period and strength can be determined from a transform of the amplitude function, with the phase set to zero. We can consider a peak to have significant strength if its amplitude is more than one-fifth that of the dc peak.

The measure of the fast variation is the minimum of the amplitude divided by its moving average. We next calculate a linear least-squares fit to the moving average, and subtract the linear term (but not the constant term) from the moving average. This subtraction allows for a slow variation in scatterer strength, due to the scatterer drifting slowly through its range gate and to changes in the target's orientation. The measure of the slow variation is the ratio of the minimum to the maximum of the resulting function.

For the most demanding (phase tracking) applications, the product of these two measures should be greater than about 0.5. To appreciate this threshold, first assume that the fast measure is equal to unity, and that the slow variation is entirely due to interference between two scatterers. The threshold of 0.5 would then correspond (see Appendix A) to a phase jump of 0.1 cycles, which is barely tolerable. Next, assume that the slow measure is equal to unity and the fast variation is entirely due to noise. The threshold of 0.5 then corresponds to a largest noise fluctuation of half the scatterer amplitude. If we approximate the standard deviation of the noise as one-third this largest deviation, standard noise theory [3, 5] gives a corresponding standard deviation in the phase of 0.03 cycles. Again approximating the maximum deviation as three times the standard gives a barely tolerable 0.09 cycles. We note that if all scatterers within a single range gate moved slowly through that gate, we might be better served by a measure of fast amplitude variation based on the standard deviation from the moving average. However, three-

dimensional targets may have some scatterers barely moving within a gate while others move quickly through it. Using the measure as described lets us avoid phase tracking at such a time, and also alerts us if our previous tracking was so poor that the tracked scatterer does not remain within a range gate.

For those applications in which we want to recognize abrupt changes in phase slope and large phase fluctuations, the product of the two measures can be somewhat lower. Recognizing that such problems occur is easier than phase tracking through them or through phase jumps due to interference. We have not tested a threshold extensively, but estimate that it should be in the range of 0.2 to 0.3. In practice, the more relevant question for this application is how large can acceptable phase fluctuations and phase-slope discontinuities be? Acceptable phase fluctuations do not differ by more than about 0.1 cycles from a moving average of the phase, calculated over the same windows used for the amplitude. The limits on phase-slope discontinuity depend on two criteria. First, a phase-slope difference implies a shift of the corresponding response in crossrange, which should be by less than one gate. This implies that the duration of an image containing the discontinuity can be no longer than the reciprocal of the change in phase slope. Second, the crossrange shift caused by the phase-slope difference must be a small fraction of the crossrange width of the target. This may further reduce the allowable imaging duration.

For the least demanding application, coarsely determining the relative motion of widely separated scatterers, our primary concerns are verifying that we have not lost track of a scatterer or, worse, switched our track from one scatterer to a crossing scatterer whose range at any time differs from that of the first by more than about a range gate. In either case, a loss of track is most likely to occur when the amplitude of the tracked response becomes weak and stays weak for an extended interval. Hence, only the slow measure of amplitude variation is relevant. We estimate that an appropriate threshold is in the range of 0.1 to 0.2.

2.3.4 Section Summary

The imaging of a moving target must not be considered as the problem of finding a motion compensation that generates the image of the target when stationary. Depending on the target motion, there are severe limits to the achievable motion compensation. The best obtainable image may be highly smeared in crossrange.

The goal of the motion compensation is to form an image that allows measurement of the scatterer positions.

The available motion compensation steps, in order of increasing accuracy, are:

- Range tracking of the entire target;
- Doppler tracking of the entire target;
- Range tracking of a scatterer;
- Doppler tracking of a scatterer;
- Range/Doppler tracking of a scatterer;
- Phase-slope tracking of a scatterer;
- Phase tracking of a scatterer.

The need for a motion compensation degrades the identification range relative to the maximum detection range.

For a given target, the selection of the motion compensation steps (from a very simple to a complicated motion compensation) must be performed adaptively.

2.4 The Need For Adaptive Processing Methods

In order to identify a target, one must obtain an image of the target, extract information from the image, and compare that information with the information in the database. The primary problem is to generate an image of such a high quality that one can extract enough information for identification. Extracting the information is not simple, but if the image is not good enough it becomes impossible; hence, image quality is of primary importance. It is simple to generate a good (in a radar, not optical, sense) image if the target is stationary. The typical application of this kind is SAR surveillance, and although a high-quality compensation of the platform motion is not trivial, just as the entire design of a high-quality surveillance system is not a trivial task, a satisfactory motion compensation technology has been developed. One good motion compensation provides high-quality images of all stationary targets, because their residual motions relative to the platform are smooth and known.

The situation with moving targets is much more difficult. When one images stationary targets, one attempts to fly the radar platform at constant speed along a prescribed path. With moving targets, the precise motion to be compensated varies greatly from one target to the next and from one

observation to the next on the same target, as well as between types of target. We already discussed the differences between the motions of aircraft, ground vehicles, and ships. The variability of the motion is particularly great for ground vehicles. Different types of ground vehicle have different designs—some are more rigidly constructed than others—and they have different suspensions. Vehicles may travel at a variety of speeds on roads of varying quality. Some vehicles may travel off the road, particularly such military vehicles as tanks. Some vehicles act as rigid targets and others as flexing or even vibrating targets. Even the same vehicle sometimes behaves rigidly at some times and nonrigidly at others, depending on the smoothness of the surface upon which it moves. The translational motion of a vehicle may be smooth or there might be (for radar measurement purposes) significant yaw, pitch, and roll. Depending on the vehicle design, its suspension, and the travel conditions, a bouncing motion could be important, and it could be slow or fast, regular or irregular. The motions of aircraft and ships are benign when compared with those of ground vehicles, but they are still quite variable, in particular for ships.

It is clear that, whereas one kind of motion compensation is adequate for all stationary targets, there can be no specific motion compensation that is adequate for all moving targets, even for the same kind of targets. Given that a moving target is to be imaged, the specific motion compensation that will produce an image of adequate quality depends entirely on the behavior of the target. Since that behavior is not known, the processor must measure the target behavior before a motion compensation can be applied. It then must decide which of a set of available motion compensation steps are appropriate for the particular target and in what order.

This necessary degree of adaptivity cannot be implemented in a straightforward manner, because it is intrinsically coupled with the analysis of the behavior of the target. In other words, the process of learning about the target behavior is interconnected with the various motion compensation steps. For example, the processor cannot fully analyze the target's motion without forming an image, but forming the image requires knowledge of the motion to be compensated. The solution is to apply individual motion compensation steps one at a time, and check whether the desired result has been achieved before proceeding to the next step. If the outcome of the processing step is not as desired, the processing must be modified so that the desired result is obtained. The point is that imaging of a moving target must necessarily be a highly adaptive process. This will be demonstrated by many of the imaging examples given in this book.

We want to stress that the discussions of the compensation verification steps are not meant solely for the purpose of expounding the theory of imaging and identification, but represent a process that must actually be implemented in an operational system. The many illustrations may be deceiving in that respect, but they represent processor outputs that replace the equations used in applications where mathematical target models are useful. Such models are not useful when the task is to identify man-made targets with radar.

References

- [1] Woodward, P. M., *Probability and Information Theory, with Applications to Radar*, London: Pergamon, 1953.
- [2] Rihaczek, A. W., *Principles of High-Resolution Radar*, New York: McGraw Hill, 1969. (See Reference 1 on page 290.)
- [3] Rihaczek, A. W., and S. J. Hershkowitz, *Radar Resolution and Complex-Image Analysis*, Norwood, MA: Artech House, 1996.
- [4] Rihaczek, A. W., and S. J. Hershkowitz, "Man-Made Target Backscattering Behavior: Applicability of Conventional Radar Resolution Theory," *IEEE Trans. on Aerospace and Electronic Systems*, Vol. 32, No. 2, April 1996, pp. 809 -824.
- [5] Helstrom, C. W., *Statistical Theory of Signal Detection*, Oxford: Pergamon, 1960.

3

Aircraft Identification

The most challenging aspect of aircraft identification is distinguishing among the large number of interceptor aircraft and small attack planes, most of which are similar in appearance and size. The discussion in this chapter is oriented toward the difficult problem of identifying specific military aircraft; an illustration of the imaging of large commercial aircraft will be provided for contrast.

3.1 Significance of the Radar Waveform

This section is concerned with the design of the radar used for aircraft identification, specifically its waveform. Whereas for most radar applications the specifics of the transmitted waveform are immaterial as long as the signal is coherent and has adequate bandwidth and duration, this is not so for aircraft identification. Here we have the special problem that the target to be identified carries strongly reflecting, rapidly rotating devices (jet engine blades or propellers). We show that *if an unsuitable waveform is chosen, reliable aircraft identification is impossible*. The most desirable waveform and practical compromises are discussed.

An aircraft is in one sense a more difficult target to identify than either a ground vehicle or a ship. The reason for this is that an aircraft carries parts that move with very high velocities; that is, jet engine blades or propeller

blades. In view of the relative importance of jet aircraft as compared with propeller aircraft, we will consider only jet aircraft. When a radar illuminates the rotating blades of a jet engine, these act as scatterers with much higher Dopplers than those of the skin return. The radar thus must perform measurements on an aircraft when parts of the aircraft have very high differential Dopplers with respect to the body of the aircraft. In the parlance of waveform design [1], this is the problem of operating in a "target space" with a very high Doppler extent. If the radar waveform is chosen inappropriately, the interference of the jet engine returns with the body returns can easily be so high that aircraft identification becomes impossible. A reliably operating aircraft identification system cannot be designed with a radar using such an inappropriate waveform. The problem of the jet engine returns is very serious, because these returns can be observed over wide angular regions off nose-on and off tail-on. Depending on the size of the engine intake, only relatively narrow angular sectors about broadside may be sufficiently free of these returns, and these are the very aspect angles at which aircraft identification becomes questionable in practice. Although this is not a book on waveform design, this problem is so important that we will summarize the pertinent facts.

The linear FM signal is perhaps the most popular radar waveform. For a consideration of interference within an extended target space, it does not make any difference whether the linear FM signal is a continuous signal or implemented by linear frequency stepping. Both variants of the linear FM signal have identical properties, the important one for our purposes being the so-called coupling between delay and Doppler [1]. This property implies that when a return is shifted in Doppler, the response at the receiver output will be delayed or advanced in range, depending on the direction of the FM sweep. Thus, if this waveform is used on aircraft, the jet engine returns will be translated in range; they will not remain within the range gate in which they are generated. Moreover, the jet engine returns may be concentrated in Doppler bands distributed over the entire Doppler width of the data, the PRF. For the commonly used low PRF waveforms, the engine lines thus are highly folded in Doppler. The engine lines from the different original Doppler bands fall within different range gates distributed over the entire range window. The problem is compounded by the variety of multipath returns within engine ducts and exhausts, which spread engine returns into more range gates. More than one of these range gates might be within the image of the aircraft.

The behavior of the jet engine returns does not create a problem when the PRF is higher than the maximum Doppler of the engine returns. The

first strong jet engine line, which is at the spin frequency times the number of blades, will be well outside the Dopplers of the skin return. Unfortunately, for reasons of economy, radars are usually designed with relatively low PRFs, so that the engine returns are folded many times in Doppler. If the linear FM waveform is designed so as to make the spreading in range due to delay/Doppler coupling smaller than one range gate, then we lose just the range gate in which the strongest engine returns appear (possibly plus a few more due to multipath). We also gain information about the location within the duct of the first engine stage. On the other hand, if the waveform has the wrong parameters, the engine returns may be spread over the entire range extent of the aircraft and more.

The relation between the differential Doppler Δv and the range delay $\Delta \tau$ is given by [1] as

$$\Delta \tau = \Delta v (T/B) \quad (3.1)$$

where T is the duration and B is the bandwidth of the linear FM pulse or the stepped frequency signal. If range delay is converted into range,

$$\Delta R = (c/2)\Delta \tau \quad (3.2)$$

and the range resolution cell

$$\Delta r = c/2B \quad (3.3)$$

is introduced, (3.1) becomes

$$\Delta R/\Delta r = T\Delta v \quad (3.4)$$

This relation shows that, for a given Doppler shift of an engine return, the range shift, expressed in terms of the range gate width, depends only on the duration of the linear FM signal or of the stepped-frequency pulse. With typical carrier frequencies (which govern Δv), we find that (3.4) permits pulse durations in the order of a few tens of microseconds if the engine returns are to remain within a range gate. For example, assume a Doppler spread of the jet engine returns of $\Delta v = 20$ kHz, and an allowable range shift of half a range gate, $\Delta R/\Delta r = 0.5$. From (3.4) we then obtain an allowable signal duration $T = 25 \mu s$.

As an illustration of how severe the interference from engine returns can be if the duration of the FM signal is too large, in Figure 3.1 we show the peaks plot image of an aircraft obtained by a radar with slow frequency stepping, that is, a linear FM signal with a duration that is too long. The aircraft image is near the center of the figure, as approximately indicated by the dashed generic aircraft outline. The image is surrounded by jet engine returns of about the same intensity as the aircraft returns. Since the engine responses can fall anywhere, including on top of the aircraft image, the identification problem is unsolvable with automated processing. As already stated, only the duration of the signal counts, not whether it is a linear FM signal or a frequency-stepped waveform.

The preceding discussion is based on the assumption that a frequency-stepped waveform uses pulses on each frequency that have 100% duty ratio, which reduces the required peak power. By changing to a small duty ratio, at the expense of correspondingly higher peak power, we can change the ambiguity function of the waveform in such a way that the properties of a short linear FM signal are approached. In other words, if the frequency-stepped waveform operates with low duty ratio, the spreading of the engine returns in range is avoided; but then we might as well not use frequency stepping.




Figure 3.1 Peaks plot image of an aircraft, with jet engine returns.

Using short linear FM pulses or frequency stepping with low duty ratios does not solve the problem entirely but makes it easily manageable.

We stated that the engine returns can be made to remain in the range gate in which they are generated. In addition to the first engine stage, which generates the strongest returns, there are additional engine stages that generate weaker returns in other range gates that do not appear to be much of a problem. However, just as an engine duct generates a string of multiple delayed returns from the fixed parts of the engine stages because the signal bounces around in the duct (or, to a lesser degree, within wing-mounted engines), the duct also generates multiple delayed returns from the rotating blades. Hence, the strong returns from the first compressor stage do generate observable multiple delayed returns in other range gates, but these returns are typically much weaker than the direct engine returns. They can be more readily accommodated than such strong returns as shown in Figure 3.1 for a long waveform.

The peaks plot image of Figure 3.2 shows a case in which such delayed engine returns are bad because the aspect angle is small. As in Figure 3.1, the




Figure 3.2 Aircraft image for a short linear FM pulse.

approximate aircraft location is given by the dashed generic outline. The returns from the first engine stage are as strong as the skin returns, so that the range gate with the direct engine returns may indeed be lost. Each major response is trailed by a series of responses of progressively lower magnitude, decreasing rapidly with range, much the same as with the delayed duct returns visible near Crossrange Gate 1. However, the entire front half of the aircraft image is clear because for a short linear FM signal the blade returns can be advanced only insignificantly. Although an image of the kind shown in Figure 3.2 does not prevent aircraft identification, it is clear that it would be preferable to operate with a PRF so high that there is no Doppler foldover and no interfering blade returns.

The problems of range/Doppler coupling of linear FM do not occur for another popular waveform, the phase shift code. However, the tradeoff is that the suppression of the unacceptably high range sidelobes of a phase shift code is not as easy as for a linear FM signal. Range sidelobe suppression does not work if the Doppler shift is too large compared with the reciprocal of the duration of the code element pulse. This is to say that when each subpulse of a phase shift code has a Doppler resolution cell that is not at least several times larger than the Doppler spread of the aircraft, range sidelobe suppression becomes ineffective. The consequence is a limitation on the total duration of the phase shift code. Suffice it to say that the phase shift code does not have the range/Doppler coupling problem of the linear signal, but it has problems of its own that require a careful waveform design.

Section Summary

Since a large instantaneous bandwidth is expensive to implement, a popular waveform uses linear frequency stepping over a relatively long time, a 100% duty ratio, and repetition at a low PRF. This waveform is unsuitable for aircraft identification because it translates the jet engine returns in range. They superpose on the image and make it useless. The same effect occurs with a linear FM signal with too long a duration. Because of the high Dopplers of the jet engine returns, waveforms for aircraft identification must be properly designed.

3.2 Identification via Range Profiles

Since much work has gone into attempts to develop aircraft identification systems based on high range resolution alone, or on range profiles, we

include this section on identification via range profiles. If such identification were feasible, it would have the valuable attribute that range profiles can be generated with very short dwell times on the target; radars that must detect, track, and identify perhaps a large number of targets have a severe shortage of dwell time. The approach would be of particular interest with respect to aircraft because aircraft are often free of ground and sea clutter, which interfere with target returns in the absence of crossrange resolution. In this section we show that *range profiles without crossrange resolution cannot be the basis for a reliable identification system when the database is reasonably large.*

3.2.1 Requirements on Range Profiles for Identification

Attempts to identify an aircraft on the basis of its range profile assume that the range profile of an aircraft is unique enough to allow identification. Presumably, one may identify simply by comparing a received range profile with a database containing the range profile of the target of interest. Is such an approach practical? In order to avoid having to speak in generalities, we will specifically examine whether such range-profile matching is feasible with aircraft. Here we consider the identification of similar aircraft, such as fighter aircraft, because distinguishing small from large aircraft is easy.

The first requirement for success is that the range profile of a target contain sufficient detail to allow identification of the target in a large database. It is easy to identify one out of five aircraft, but very difficult to identify one out of 200 aircraft. This implies that a fairly high range resolution is needed if small aircraft are to be identified, in particular since many fighter aircraft have similar dimensions, features, and appearances. Better range resolution is required near broadside, because range resolution loses its effectiveness on the fuselage as the broadside aspect is approached and, as already explained in Chapter 1, the features of the fuselage are far more important for identification than the features on the wings. A first conclusion is that the method will not work close to broadside.

The second requirement is that the database be manageable. The range profile of an aircraft changes with aspect angle in the plane of the wings, with bank angle, and with pitch angle. The latter is important primarily if the radar platform is positioned significantly out of the plane of the wings of the aircraft to be identified, looking up or down. If the range profile of the aircraft changes too rapidly with changing orientation, we have a twofold problem: How can we collect or calculate the range profiles of all aircraft of interest, for all orientations, and with different configurations of ordnance?

Assuming that we somehow manage to construct such an extensive database, how do we store it for use in real time?

If we use a range resolution so high that each range cell contains only one dominant scatterer, the peaks of the range profile become the responses from individual scatterers. Then we do not have an interference pattern that changes rapidly with target orientation. However, operational radars typically do not have such a high degree of range resolution, because it is expensive to implement. In addition, the dominant scatterers on an aircraft are extended, and problems will appear when the range resolution cell is made smaller than the range extent of the scatterers. The fact is that a *range profile is an interference pattern, the structure of which is extremely sensitive to changes in the orientation*. Thus we must analyze the peaks of this pattern to determine actual scatterer positions, which will not be very sensitive to orientation changes. With range resolution of critical importance, we must use the complex range profile rather than the intensity range profile in order to avoid degrading resolution by a factor of two. With optimum processing, we can generate a synthetic range profile in which the actual scatterer positions change more slowly with orientation than the peaks of the intensity range profile. As a consequence, the requirements on the database will be reduced.

Now, the improvement in resolution by using the complex range profile is only by a factor of two, and if a significant number of scatterers are separated by less than one resolution cell, the synthetic range profile will not be good enough. We need more than a few scatterers to identify an aircraft in a large database. Thus it becomes tempting to try to circumvent these limitations by superresolution. However, because of the problems of mathematically modeling real targets, superresolution techniques will give sharp response peaks whose number exceeds that of actual scatterers and whose ranges will be grossly in error. We will discuss the difficulties of identifying aircraft from their range profiles when practical processing methods are used.

3.2.2 Range Profiles Without Duct Returns

Modern fighter aircraft typically have their engines integrated into the fuselage, which leads to long engine ducts. At aspect angles that permit the radar signal to enter the ducts, they generate long series of delayed returns that cause particularly severe identification problems. We will investigate the problem of engine ducts separately, first considering the more benign but somewhat unrealistic case of range profiles that do not contain engine duct returns. The data we will use for our examples were taken on a test range, so that perfect ground truth is available. There are no duct returns because for

the measurements used in the following demonstration the engine intakes were covered.

If identification via range profiles is at all possible, the favored aspect angles must be fairly small, so that range resolution is effective on the fuselage. The worst case is near broadside. We will choose an intermediate aspect angle, not prohibitively large, but large enough to represent the problems for a realistic situation.

Figure 3.3 shows a series of range profiles taken with a signal bandwidth of 300 MHz, not overly large but still a respectable bandwidth. For the range scale used in this figure, one resolution cell is four range gates wide. The top plot gives the range profile for a yaw aspect angle of 45.00° , the next plot down for 45.05° , then for 45.10° , and the bottom plot for 45.20° . The bank and pitch angles are zero in all cases. A casual examination of the plots shows that noticeable changes in the shape of the range profile start to occur for an aspect angle change of 0.1° , and important changes for an aspect angle change of 0.2° . The changes we refer to are primarily the positions of the peaks of the range profile.




Figure 3.3 Range profiles for a 300 MHz bandwidth.

The extraordinarily rapid change in the shape of the intensity range profile is due to the interference between scatterers in the same range gates but with different crossrange positions. Although use of the complex range profile improves range resolution, the improvement is only by a factor of two, and only if most responses consist of just two scatterers. Thus, when we consider the bottom plot of Figure 3.3 and the poorly resolved peaks, resolution of the responses appears problematic even when the complex range profile is utilized. We could demonstrate the resolution problem for the case represented by Figure 3.3. However, we will show below that the problem is not even solvable when the signal bandwidth is doubled to 600 MHz, so we need not demonstrate it for a 300-MHz bandwidth. Considering this problem, we conclude that the 300-MHz bandwidth range profiles of Figure 3.3 would have to be used for aircraft identification without any further refinement. With six or seven peaks to a range profile, identification would be possible at most for a very small database.

Let us ignore this difficulty, and (wrongly) assume that we could successfully distinguish a large number of aircraft with range profiles having only such a small number of peaks. The question then becomes, how many range profiles do we need in the database to represent one target for all orientations, and how many to represent all targets of interest over all orientations?

Roughly speaking, by comparing the range profiles of Figure 3.3, it appears that a new range profile is needed when the aspect angle changes by about 0.1° . With this increment, we need 3,600 range profiles merely to cover all aspects. This number must be multiplied by the number of increments needed to cover all elevations at a given aspect. Assuming the same increment gives a factor of 1,800. In addition, we must do this for different configurations of ordnance, at least those configurations that have significant effects on the range profile. How does one obtain such a database on all aircraft of interest? If it could be obtained, how would one store such an enormous database?

As pointed out, the possible remedy for the problem of an excessive number of range profiles is to use a range resolution so high that most of the important scatterers are resolved, so that one need not deal with highly variable interference patterns. In Figure 3.4 we show range profiles at the same aspect angles as for Figure 3.3, the only difference being that the signal bandwidth was doubled to 600 MHz. For the scale of the abscissa of Figure 3.4, each resolution cell is now two gates. Taking into account that the aspect angle change from one range profile to the next is only 0.05° (0.10° for the last two range profiles), the conclusion obtained from an examination of




Figure 3.4 Range profiles for 600-MHz bandwidth.

the series of range profiles is much the same as for the lower bandwidth: The shape of the range profile changes too rapidly with aspect angle. Although we now have about twice the number of peaks as at the lower bandwidth and can better differentiate between aircraft, significant changes in the range profile start to occur when the aspect angle increases by 0.1° , and important changes occur when the increase is 0.2° . We can again ask whether the situation is sufficiently improved if we use the complex range profiles instead of only the intensity range profiles, and analyze the responses to determine actual scatterer positions.

Carrying out this analysis for the top range profile of Figure 3.4 gave the following results. First, it was not possible to extract scatterer positions from the low-level responses; a factor-of-two improvement of resolution was not enough. The measurement of the scatterer positions worked for some of the stronger responses and not for others, because even at the high signal bandwidth of 600 MHz, some responses contained significant contributions from more than two scatterers. For example, the composite response centered in about Gate -35 of the top profile can be resolved into its

constituents, despite its poor appearance. This is not possible with the wide response centered in Gate 12, and neither can good position measurements be performed on the responses in Gates 18 and 24. As an example, in Figure 3.5 we show the transform of the complex response associated with the peak centered in Gate 12. The amplitude/phase pattern is much too different from the ideal pattern of two interfering scatterers to attempt determining the number and positions of the contributing scatterers. We conclude that 600 MHz is still an insufficient bandwidth for the purposes of reducing the sensitivity of the range profiles to orientation changes by switching from the intensity range profiles to the complex range profiles. The resulting improvement in range resolution is still not adequate, despite the relatively high bandwidth of 600 MHz.

The range profiles under the same conditions but with a further doubling of the signal bandwidth to 1,200 MHz are shown in Figure 3.6. Each resolution cell now corresponds to one gate in the figure. With this very high signal bandwidth, the number of peaks in an intensity range profile appears to be high enough for distinguishing many aircraft, provided the peak positions signify scatterer positions rather than peaks of an interference pattern (if there were no other problems, such as duct returns). However, as can be seen from the plots of the figure, the changes of the shape of the range profile when the aspect angle is increased by only 0.2° remain rather severe despite the almost impractically high bandwidth. Again, the question is, can the




Figure 3.5 Transform of the peak centered in Gate 12.




Figure 3.6 Range profiles for a bandwidth of 1,200 MHz.

complex peaks be analyzed so as to determine the positions for this high bandwidth? Our tests in analogy to the example given in Figure 3.5 showed that the answer now is yes, but still only for the stronger responses. One cannot judge from the limited number of tests we have carried out whether a 1,200-MHz bandwidth would indeed reduce the variability of the range profile with aspect angle to the degree that the number of range profiles needed in the database would become manageable; there is only a good possibility. Note, however, that so far we have ignored the multiple delayed duct returns.

3.2.3 Delayed Duct Returns

For a demonstration of the difficulties introduced by the delayed duct returns, we use a different fighter aircraft under similar conditions but with the engine intakes uncovered, so that duct returns are present. The target aspect is 50° off nose-on, which is favorable with respect to the duct returns, which decrease in strength as the aspect angle approaches broadside. The bank and pitch angles are zero. In Figure 3.7 we show a SAR image of the




Figure 3.7 SAR image of the aircraft.

aircraft in peak plots form. The signal bandwidth is 1,200 MHz, the same as for the range profiles with sufficient resolution to significantly decrease the variability of the range profile with orientation. The delayed duct returns are located in Crossrange Gates -1 to 7 , between the dashed vertical lines. The aircraft lies approximately within the dashed generic outline. Although the image responses are not well focused, this is irrelevant for this particular demonstration. The figure shows that beyond about Range Gate 15 the target returns in each range gate are dominated by the duct returns. Beyond Range Gate 35 the delayed duct returns are the only returns. The sizes of the dots indicate that the duct returns are among the strongest from the target.

We showed above that a 1,200-MHz bandwidth might be just about sufficient to allow the resolution of an adequate number of scatterers on the aircraft. This is a matter of how closely in range the observable scatterers are spaced. If one goes from the two-dimensional image of Figure 3.7 to range profiles, all the responses in each range gate are folded into a single crossrange gate, aircraft responses as well as the delayed duct returns. Let us assume the possibility that somehow the delayed duct returns could be used as an additional input to aircraft identification, perhaps by relating the individual

responses to the various engine stages. This would require range resolution of the individual duct returns. Is a 1,200-MHz bandwidth large enough for this purpose? Figure 3.8 shows the range profile of the delayed duct returns, generated by collapsing Crossrange Gates -1 to 7 of Figure 3.7. Below Range Gate 15, the profiles of Figure 3.8 contain contributions from skin returns as well as duct returns.

Our examination of the (complex) duct responses of Figure 3.8 showed that a 1,200-MHz bandwidth is almost sufficient for the resolution of the strongest duct returns, and hence for the measurement of the ranges of the individual duct returns. Hence, if one could predict these positions, without interference from the skin returns one might be able to achieve a positional match for this large a bandwidth. However, even the best duct returns in Figure 3.8 do not allow a position measurement as accurate as can be obtained on the better of the skin returns, because of the complexity of the scattering process within an engine duct. All these comments still apply only to the favorable case that the duct returns could be examined without interference from the skin returns, which is possible only for large aspect angles.

3.2.4 Range Profiles Containing Skin Returns and Delayed Duct Returns

So far we can state the following conclusions. In the absence of duct returns, the range profile of an aircraft is highly variable with aspect angle, so that an enormous database would be required for identification on the basis of the




Figure 3.8 Range profile including delayed duct returns.

range profiles. The size of the database could be much reduced if range resolution were made so high that most peaks of a range profile were generated either by a single or by at most two scatterers. In this case the scatterers could be resolved without superresolution by making use of the complex range profile. The indications are that a bandwidth of 1,200 MHz might be about the minimum required bandwidth for fighter aircraft. The situation is similar for duct returns. In the absence of interference by skin returns, the duct returns could be satisfactorily resolved with a bandwidth of at least 1,200 MHz, possibly requiring a somewhat higher bandwidth. However, the real problem occurs when both skin returns and delayed duct returns are present, as will be the case for most aspect angles of the aircraft. Then the two types of return interfere with each other. The resolution of returns that require a 1,200-MHz bandwidth when only one kind or the other are present demands a still higher bandwidth when both are present. Even in that case, only half of the problem would be solved. There would still be the need to measure or predict the delayed duct returns for all aspect angles, so that they can be included in the database.

As an illustration of the practical case when both types of return are present, Figure 3.9 shows the range profiles for the combination of skin




Figure 3.9 Range profiles for the entire aircraft, including duct returns.

returns and delayed duct returns, again for the same small incremental aspect angle changes. The dotted vertical lines divide the profiles into four sections. The section at the lowest range gates includes only skin returns. The next section contains both types of return, with the skin returns dominant. The following section also contains both types of return, but now with the duct returns dominant. The last section contains only duct returns. Examining the duct returns of the last section for the different aspect angles, we see that there is some variability over the 0.2° angular change in the figure, although less than for the skin returns. We would expect this because the duct is much narrower than the fuselage, so that the variation with aspect angle should be smaller. An examination of the strong return group about Gate 20, where the duct returns are dominant, shows that the skin returns do cause considerable interference. The combined returns thus change with aspect angle. We also see significant changes in other parts of the range interval occupied by both delayed duct returns and fuselage returns. Of course, in the interval that does not contain any duct return we have the earlier situation: With a bandwidth of 1,200 MHz the changes are slow over the 0.2° aspect angle interval, and the complex responses can be analyzed to determine actual scatterer positions.

The practical question regarding an obviously impractical situation is whether the peaks of the range profile can be resolved into the contributing responses from scatterers, both on the skin and within the duct(s), in the two sections containing the most important returns, between Gates -38 and 36 in Figure 3.9. Our examination of the complex responses showed that the mutual interference between duct returns and skin returns degrades the measurement of range positions. However, the degradation was not so large that one could readily conclude that 1,200 MHz is too small a bandwidth. The result is highly dependent on the aspect angle and will change from one aircraft to another. A definite answer would require identification runs on many aircraft under many different conditions, before the adequacy of a 1,200-MHz bandwidth could be confirmed or rejected. Such a test would be of academic interest alone, because it would address only one part of the problem of identifying aircraft from their range profiles, namely, whether the size of the database could be made manageable if the bandwidth were increased to 1,200 MHz. It still would leave open the question of how to predict the delayed duct returns for all aircraft of interest and for all aspect angles.

3.2.5 Conclusions About Range Profile Utility

One cannot consider the range profile of an aircraft to be a radar signature from which the aircraft can be identified. First, the use of the intensity range profile

as a target signature would unnecessarily degrade performance. Instead, if such an approach were to be pursued at all, one should determine the scatterer positions from the complex range profile and use the scatterer positions as the target signature. The resulting improvement would still not be large enough to overcome the tremendous practical obstacles against using range profiles in a way suggested by the term "target signature"; that is, for identifying aircraft. If one utilized bandwidths in excess of 1,000 MHz, perhaps much in excess of this number, one might be able to shrink the size of the required database to practical values. This still leaves the problems of accurately predicting range profiles for the skin returns, range profiles for the duct returns, and the interference between the two. This would have to be done for all aircraft of interest, for all orientations of interest, and for the various configurations of ordnance.

Does this mean that range profiles are useless for target identification? In our interpretation it means that in applications where Doppler resolution is available, one should not try to work with range profiles. An extremely high bandwidth will facilitate identification via SAR/ISAR imaging, but it cannot provide adequate performance with resolution in range only. Resolution in the second dimension of Doppler is needed. Moreover, even with the addition of Doppler resolution, the reliable identification of man-made targets is not an easy task. Hence, when the available observation time is large enough to add Doppler resolution, this must be done in a theoretically optimum manner, not by some type of pseudo-Doppler resolution. There are applications in which Doppler resolution is unavailable. The foremost example is the identification of a stationary ground vehicle from a stationary radar platform, such as from a stationary ground vehicle. In such an application we must use the highest practical bandwidth, but even then it will not be possible to achieve the identification performance obtainable when at least one of the two objects, the ground vehicle to be identified or the one carrying the radar, is moving. It would be preferable for the vehicle with the radar to move at least during the time needed for target identification, establishing a SAR baseline.

We have discussed the problems of treating the range profile of an aircraft as a signature. Of course, one may also consider a SAR/ISAR image of a target as its signature. Here we need not discuss the extension from one-dimensional to two-dimensional signatures, because other chapters of the book address identification via SAR/ISAR imaging. We do not consider such images to be adequate as signatures. First, the necessity of utilizing the complex image makes the term somewhat irrelevant. Second, the term "signature" implies that the image is correlated with an image in the database,

which is an unworkable approach. Just as a range profile represents an interference pattern, the two-dimensional image is also an interference pattern. Instead of using images as signatures, one must extract specific information about the target to be compared with the database.

3.2.6 Section Summary

With typical values of range resolution, range profiles are so highly aspect-angle dependent that a database with a huge number of range profiles would be needed. This is true even for signal bandwidth on the order of 1,000 MHz. Obtaining and storing this database are problematic.

Since delayed duct returns are present under most conditions, range resolution would have to be high enough to resolve the delayed duct returns from the skin returns; this would allow reducing the number of range profiles in the database. However, this would still leave the problem of predicting the duct returns for all aspects.

In the case of aircraft, enough Doppler resolution is needed to resolve the delayed duct returns from the skin returns.

3.3 Identification via ISAR Imagery

The process of identification via ISAR imagery can be divided into four steps: (1) data collection, (2) compensation and selection of the imaging interval, (3) image analysis, and (4) comparison of measurements to database. The first three of these must be carried out adaptively. Data collection must continue, operational considerations allowing, until sufficient resolution is achieved for identification to be feasible. The necessary dwell varies with the aircraft type, aspect, and behavior. This is discussed further in Section 3.5. Oftentimes, it is not possible to image over the entire data collection dwell. When this is the case, we must choose an appropriate subinterval for imaging. This is done on the basis of tracking and compensating individual scatterers, and examining fixed-range image cuts to determine subintervals of smooth motion. This step is discussed and illustrated extensively in Section 3.6.

The image analysis consists of measurement of recognizable features, such as aircraft length or location of engines, as well as measurement of positions, strengths, and characteristics of nonrecognizable features. These features are described in Section 3.4, along with construction of a feature database. The feature measurement algorithms have been discussed in

the previous chapters, as has the comparison of the measurements to the database. Section 3.7 gives examples of the positional match that is used to make this comparison.

3.4 Aircraft Features for Identification

Earlier, we briefly discussed the futility of any kind of shape recognition similar to optical target identification. The conclusion was that radar target identification must be based on specific features that can be measured by radar. In this section we discuss the features usable for aircraft identification. We distinguish between two classes of feature: First, special features that are distinct for each aircraft and whose significance can be readily determined. Second, features that cannot be recognized from their responses (for example, a strong scatterer that can be located, but whose origin cannot be deduced). For this second class, we utilize only the feature positions.

3.4.1 Special Aircraft Features

When the special features discussed below can be extracted from an image, they can be very important contributors to identification, but their extraction will not always be possible. Hence, the special features are used when they can be observed, but there can be no penalty when they are not observed. As already stated, we will not consider the easy problem of distinguishing small from large aircraft, but the problem of identifying aircraft within their class. Here again, our interest is in the difficult task of identifying small military aircraft.

Aircraft Length and Wingspan

Although fighter aircraft have similar lengths and wingspans, if these overall dimensions could be measured accurately, they would be very important for identification. Unfortunately, the measurements of length and wingspan both have problems that reduce their utility.

One difficulty in measuring aircraft length is that radar can measure only range extent, and that the aspect angle of the aircraft is needed to translate range extent into length. In view of the relatively short dwell times available on an aircraft and its high maneuverability, it is generally not possible to measure the aspect angle accurately. The translation of range extent into length thus can be done only with an error that increases rapidly as the aspect angle approaches broadside. One might conclude that the length

measurement should be quite accurate at small aspect angles. However, at smaller aspect angles we might not observe scatterers on the tail of the aircraft, not necessarily because of shadowing. At a given aspect angle there may be no readily observable feature at the end of the aircraft. The situation where the range extent of an aircraft is undermeasured occurs frequently enough to conclude that one can reliably determine only some minimum length of the aircraft. Measurement of a minimum length is not as useful as measuring the actual length, but it does allow rejecting aircraft that are too short.

Because of the better quality of the images of large aircraft, the first and last responses along the centerline of the fuselage can be more reliably identified. The error in the length measurement then is determined only by the error in the aspect angle measurement provided by the tracker. This will allow at least a crude measurement of length.

The situation is worse with respect to the measurement of the wingspan. At smaller aspect angles we cannot accurately measure the wingspan because the crossrange scale will be known only inaccurately, even if one can reliably detect the wingtips (which may not be the case). At large aspect angles the inaccuracy of the crossrange scale becomes unimportant, but some aircraft have wingtip responses too weak to detect in the background. Also, if the aircraft is in a turn and hence is banked, we can only roughly estimate the banking angle even when the aircraft has been tracked over a considerable time. Hence, the radar also can measure only a minimum wingspan, which is even less useful than a minimum length. Nevertheless, it will sometimes be useful.

Since the wingtip responses are readily identifiable in the image of a large aircraft, the accuracy of the wingspan measurement also is determined by the accuracy of the tracker data. When the aircraft is turning, which is less frequently the case than for small aircraft, we estimate the bank angle from the radius of the turn, assuming that the combined vector of gravitational and centripetal acceleration is oriented perpendicular to the wing plane. We may have to use a longer observation time in order to determine whether the bank angle is changing, which will cause different degrees of shifting of the wingtips in crossrange.

Type and Number of Engines

Small aircraft typically have engines and engine ducts that are integrated into the fuselage, but not all do. Some have fuselage-mounted engines without ducts, and others have wing-mounted engines. If we can measure the type, number, and locations of the engines, we can reduce the number of candidate aircraft in a particular situation. For example, the capability of distinguishing between wing-mounted and ducted engines is extremely valuable.

For angles not too close to broadside, we will detect the delayed engine responses and Doppler-spread engine returns that allow us to determine the positions of the engines. Even if these returns are missing, we can sometimes use the good definition of the outline of a large aircraft to recognize the returns from engines (protruding beyond the leading edge of the wing, for example, or near the tail).

Size of the Opening of the Engine Intake Duct

The size of the intake of aircraft with fuselage-integrated engines can vary considerably. As a consequence, the level of the jet engine returns and of the delayed duct returns also can vary significantly. The relative size of the engine intake as judged from the level of the engine returns, both from the fixed and the rotating parts, thus can be used as a distinguishing feature, although not a very sensitive one.

Aircraft with Wingtip Missiles

The fact that an aircraft is equipped to carry wingtip missiles often can be recognized from the nature of the responses at the wingtip. This is another special feature that allows narrowing the number of candidate aircraft in a given situation.

Twin-Hull Aircraft

Most aircraft with fuselage-integrated engines have a single conventionally shaped fuselage. IN some cases, the engine ducts of aircraft with two integrated engines form two hulls reminiscent of the design of a catamaran. Measurement of the separation of the two exhausts, for example, may reveal this feature of the aircraft.

Position of Engine Intake and Exhausts

When delayed duct responses are observed in an image, which can happen up to large aspect angles, we can determine the position of the engine intake along the fuselage. This position varies among aircraft. Similarly, when an aircraft is viewed from the rear, the delayed duct returns will reveal the position of the exhaust along the fuselage. For some aircraft, the exhaust is not located at the end of the fuselage.

Angle Between Fuselage and Leading Wing Edge

For a large aircraft, we can choose the crossrange scale based on the track aspect and the centerline of the aircraft fuselage. Then we measure the angle

between the centerline and the leading edge of the wing. If the aircraft is turning, we adjust the angle in accordance with the estimated bank angle.

Wingspan-to-Tailspan Ratio

Since the tip of the horizontal stabilizer will usually be detectable for a large aircraft, we can measure the ratio of the wingspan and the span of the horizontal stabilizer, independent of any banking or an error in the aspect angle measurement.

Position of the Wings

For large aircraft, we can measure the positions of the nose, tail, and leading wing edge well enough to accurately determine the fractional position along the centerline of the intersection of the leading wing edge with the centerline. They might require extrapolation of the leading wing edge.

Relative Width of Wings

An image of a large aircraft may show such a well-defined wing that the width of the wing relative to the length of the fuselage may be measurable with useful accuracy. The lagging edge usually limits the accuracy.

Position of the Horizontal Stabilizer

Since the backscattering from the end of the fuselage is strong for large aircraft, we will often be able to measure the position of the leading edge of the horizontal stabilizer relative to the end of the fuselage.

Elevated Aircraft Tail

Elevated tails of large aircraft are distinguishable from ordinary tails, especially if the aircraft is banking. This is illustrated in Section 3.6.9.

Shape of the Fuselage

The fuselage of a large aircraft is often so well defined in the image that we can distinguish different fuselage shapes.

Other Special Features

Some aircraft have special designs that can be recognized by radar. For example, when a swing-wing aircraft operates with its wings in their extended positions, a radar can observe the cavity into which the wing folds. There are probably more of these special features that can be valuable in restricting the number of candidate aircraft. It is not our purpose here to provide a listing of

all the special features that can be utilized for aircraft identification. This is a matter of establishing the database, and we merely provide examples.

3.4.2 Features Common to All Aircraft

It is clear that special features such as the ones listed above cannot provide reliable aircraft identification in a large database. There are many features that are common to all aircraft, so that their mere presence does not help identification. For example, all aircraft have pilot seats, instrument panels, various air inlets and outlets, and a variety of antennas. When these features are observed in an image, one generally cannot recognize the feature. All that can be determined is that an observable feature is located in a particular position on the aircraft. Thus, only the positions of these features are usable for identification, not their nature. The utilization of feature positions is practical only if the same features can be observed over relatively large angular sectors; otherwise we would have the same problem of an excessively large database as with range profiles.

As explained earlier, the reason that feature positions are at all usable for target identification lies in the peculiar backscattering behavior of aircraft. All the numerous discontinuities, facets, and other small aircraft features generate a background that is too indeterminate and too sensitive to aspect angle to be utilized for aircraft identification. Only those scatterers whose responses stand out above the background are usable, and these are the various electronic devices and features shaped so as to catch the radar wave. Of course, the features must be visible over an extended angular sector. For example, an air inlet is generally visible for frontal aspects but not for rear aspects. The opposite is true for air outlets.

We also discussed earlier why there is little use to features on the wings other than the wingtips. Smoothly shaped ordnance on the wings typically cannot be seen and would be unpredictable in any case. Roll of the aircraft introduces unpredictable crossrange translation of the responses from wing features. Thus it is best to utilize only the features on the fuselage for identification.

The most conspicuous electronic device is the radar in the nose cone. It is a strong scatterer over a large aspect-angle sector centered on head-on. At these aspect angles, the radar has a peculiar backscattering behavior that allows one to recognize it. Its phase center will generally shift with aspect angle, implying that the Fourier transform of the fixed-range image cut through the response has a curved phase. The radar often generates a few delayed returns, and the return may be strong only in part of the

frequency band of the transmitted signal. However, there appears to be little use in recognizing the radar in this manner, because it can be more easily recognized from its position. Over a large aspect-angle sector centered on head-on, the radar return is the first strong return from the aircraft. In fact, the value of the return from the radar lies mainly in that it defines the effective start of the aircraft in an image. Although one might occasionally detect the return from the radome tip, it is too weak to be mistaken for the radar return. Outside the main beam of the radar antenna, one will detect the bulkhead behind the radar, or antennas often positioned there. Such returns are not as strong as the radar return within the main beam, but still much stronger than the return from the radome tip.

There are a variety of antennas, for various purposes, in various frequency bands, and in various locations distributed over the aircraft. Most of these antennas are installed on the smooth metallic surface of the aircraft, from which they must be insulated. Whether it be the cavity generated in this fashion, or the electronic properties of antennas, they can be readily observed against the background. Sometimes antennas are installed below the surface of the aircraft, but then the cover must be a dielectric material rather than metal. Since the antennas have specific functions, the possibility exists that the characteristics of the radar return might allow one to determine the type of antenna. We have made no attempt at antenna identification and have used only the locations of the antennas.

The intake of a fuselage-integrated engine is a strong scatterer over a large aspect angle sector. It cannot be observed for very small angles off nose-on, and also as broadside is approached. Between these limits the intake opening not only can be observed, but it can be identified as such from the delayed duct returns. They appear in the same crossrange gate as the intake opening, and generally provide a better measurement of the crossrange position of the intake than the intake response, which often is shifted in crossrange because of phase-center shift effects. Recognizing the intake response with the help from the delayed duct return is very important for the positional match.

The situation is similar for the exhaust, which also is a strong scatterer over a large aspect angle sector. Although it can also be recognized from the delayed duct returns, this is of more limited utility than in the case of the intake, because for most aircraft the exhaust is at the very end of the aircraft. Nevertheless, for the occasional aircraft in which this is not the case, identification of the exhaust is very important. One benefit from the exhaust return even when the exhaust is at the end of the aircraft is that it defines the end of the fuselage very well. We then can differentiate between aircraft where the

end of the fuselage signifies the end of the aircraft, and other aircraft where the vertical stabilizer protrudes beyond the exhaust.

In addition to the large intakes and exhausts, an aircraft has a variety of small air inlets and outlets, usually for cooling purposes. If there is no strong scatterer nearby, these inlets and outlets can be observed at aspect angles for which they trap the wave.

Another feature whose response is typically strong enough to exceed the background is the pilot's seat. Although the smoothly shaped cover of the instrument panel is not observable for frontal aspects, the head-up display on top of the cover generally is. For rear views, the instrument panel itself can often be detected in the background.

There are other observable features that are generally not known to someone unfamiliar with aircraft designs. Whatever their use might be, we must examine diagrams, photographs, and CAD models in order to recognize any feature that will trap the radar wave. We will not always know what a particular feature represents or what its use may be, but this is irrelevant for our purposes.

Establishing a database of observable features and their positions on the aircraft is practical only if these features can be observed over aspect angle sectors that are measured in tens of degrees rather than degrees or fractions of a degree. We shall use real data on a flying aircraft to demonstrate that *features do persist over large sectors*, showing that the primary variation is due to errors in the measured position of a scatterer caused by interference from nearby scatterers, rather than nonpersistence of the scatterer.

In Figures 3.10 to 3.14 we show matches of the same database to the scatterer positions extracted from images of a flying aircraft at different aspects. The aspect angle change over the five matches is 21° . (The aspect angles obtained from the matches themselves are somewhat inaccurate because only the fuselage scatterers were utilized; we used the tracking data to determine the aspect angles.) Range resolution is 1 ft, and crossrange resolution varies between about 2 and 3 ft. Scatterers in the database are designated by letters in the plot, and are identified in the legend at the right of the figure. Measured scatterer positions are denoted by + signs. The two sets of scatterers are associated via the deformable template procedure discussed in Chapter 1. The names assigned to the various scatterers are our own (and the letters are not fully consistent throughout the book), but they do represent specific scatterers. We again note that the range match is more heavily weighted than the less accurate crossrange match. The source of the mismatches is the interference among incompletely resolved scatterers. A minor scatterer not normally observable may sometimes shift the response from the




Figure 3.10 Positional match of the database to an aircraft at 20° aspect.




Figure 3.11 Positional match of the database to an aircraft at 26° aspect.




Figure 3.12 Positional match of the database to an aircraft at 31° aspect.




Figure 3.13 Positional match of the database to an aircraft at 34° aspect.




Figure 3.14 Positional match of the database to an aircraft at 41° aspect.

major scatterer. The number of observed scatterers (some not in all images) on the fuselage is 18.

In Figure 3.10, a significant mismatch in range exists only for two out of the 18 scatterers, G and d. In Figure 3.11, instead of the two scatterers P and U we observe only a single response. The same is true for Scatterers c and b. The response from scatterer X is not visible. In Figure 3.12, we have the pairs W,c and B,X giving a single response, with the response for scatterer k missing. In Figure 3.13, the tail responses are too weak, and the responses for scatterers W and X are missing. In Figure 3.14, the tail responses are again missing, scatterers O and W give a single response, and the response from scatterer c is missing. Nevertheless, in all five figures the great majority of the scatterers can be measured at the same locations as the aspect angle is changed over 20° . This verifies the facts that with real rather than modeled aircraft there is a limited number of scatterers that can be observed above the general background, and that these scatterers persist with aspect angle.

3.4.3 Deriving the Database

Target identification requires solution of two difficult problems: (1) extracting sufficient information from the return signal, and (2) obtaining the

database to which this information is to be compared. Much of this book is concerned with the first problem. The second problem is treated in this section, but the brevity of the treatment does not imply a lesser importance; how to construct the database is much easier to understand, even though constructing it is not simple. In this section we show how the database can be constructed from diagrams, photographs, and models of aircraft. Of course, when radar data are available for specific aircraft, the database for these aircraft should be established by extracting the aircraft features from the data.

The following discussion shows that there is no choice of how to approach assembling the database. Optically inspired approaches will not work because radar backscattering is too different from optical backscattering. Deriving the database for the recognizable features is easy, but for the unrecognizable features, on which we must perform position measurements, it is not. We must rule out data collection because it is impractical for all aircraft of potential interest. Test range measurements and measurements on scale models are also impractical, for more than one reason. Electromagnetic mathematical modeling (improved over the current state) should be able to generate a database of dominant scatterer positions, with low accuracy requirements, but this is an expensive alternative. This leaves only the method of using diagrams, photographs, and CAD models to determine the wave-trapping features that will be observed by radar. The method is illustrated by an example. It requires training and involves much work, but there is no choice.

To consider these points in more detail, if visual identification of aircraft at long ranges and under all weather conditions were practical, the derivation of the database would be simple. The database would be the collection of diagrams and photographs of the aircraft of interest, taken from all aspect angles. Of course, the automation of the identification process would still be a very difficult problem. How does one automatically distinguish many aircraft, sometimes with only slight differences, by comparing an image with the huge number of images in the database? It might be an even more difficult problem than radar identification, because at the larger radar wavelengths the information about the target appears in a more condensed form, so that it is easier to utilize. For example, only about 20 or 30 scatterers can be observed on the fuselage, whereas an optical system provides vastly more detail, perhaps too much detail for automated identification.

We showed earlier that assembling a database which can be used to identify aircraft via their range profiles is totally impractical. Assembling a database for the identification of aircraft via their special features and the positions of the observable scatterers is not impractical but is a difficult

problem. At the same time, it is unavoidable if aircraft are to be identified. Of course, it is far easier and cheaper to assemble the database from photographs and diagrams than by mathematical modeling of aircraft merely to find the prominent scatterers that a radar can observe.

The special aircraft features can be relatively easily extracted from diagrams and photographs. The problem lies with the features to be used in the positional match. The most straightforward and reliable way of assembling a radar database for these features would be to collect radar data of all aircraft of interest and in all aspect angle sectors within which they might be observed in practice, and then analyze the data to determine the observable features. The positions of the features and any measurable characteristics then would be recorded in the database. Since the observable features change only slowly with aspect angle, the number of different aspect angle sectors for which the database must be established is manageably small. There is one sector close to nose-on, and another one close to tail-on. The space in between can be subdivided into cones of a size perhaps by 30° by 30° . This is clearly the best way of establishing the database of the feature positions. It is probably also the most expensive way and, in the case of foreign aircraft, it may not be possible to collect radar data on all aircraft and for all aspect angles. However, one should extract the database from the data for those aircraft and orientations for which such data are available. Even if we have only a limited amount of real data, we should make use of it.

For those aircraft for which no radar data are available, when no funds are available to collect the data on accessible aircraft, and for aircraft where data collection is impossible, the first idea always appears to be mathematical modeling. As discussed earlier, our experience with measurements on real data versus measurements on data from modeling programs indicates that even very sophisticated mathematical modeling methods do not provide results sufficiently close to those obtained with real data to permit using these results as a database. We do not preclude using improved mathematical modeling to generate a database of dominant scatterer positions with low accuracy requirements. However, such modeling is an expensive alternative.

Another possibility for establishing a database of the feature positions would be to collect aircraft data on a test range, the only difference relative to collecting data on aircraft in flight being a possible reduction in cost. Such data collection is practical mostly for potentially friendly aircraft, and only to a limited extent on potentially unfriendly aircraft. Aside from these limitations, it is not easy to collect data at outdoor ranges that can accommodate full-scale aircraft and achieve the data quality needed for the measurement of scatterer positions. This is much easier for aircraft in flight. Scale-modeling

aircraft to reduce the size so that indoor ranges can be used has what we consider insurmountable problems. Modeling the important scatterers, such as cavities and antennas, to such a fidelity at the higher frequency of the test range that they represent with sufficient accuracy the properties of the real scatterers at the operational frequency appears to be almost as impractical as mathematical modeling.

This assessment leaves us with only one practical approach. The fact that only a certain kind of aircraft feature produces observable responses in an image suggests that the database can be derived in a way that is more practical than collecting flight test data on all aircraft of interest, and more realistic than collecting data on test ranges or mathematical modeling. Since we do understand what type of feature generates an observable response, we can examine diagrams and photographs of aircraft to find these features, so that their positions can be listed in the database. It is true that this approach requires detailed design information as provided by diagrams and photographs, which may not always be easy to obtain. However, it requires less accurate information than mathematical modeling. Also, it requires some training, but again much less so than for mathematical modeling. Obtaining the required diagrams, photographs, and CAD models and examining them closely is time consuming, but even in these respects the effort and cost compare favorably with mathematical modeling. In the following, we shall indicate how one approaches assembling the database.

Figure 3.15 shows an aircraft photograph of the type that is typically found in publications that provide general information of the various kinds of aircraft. The photograph in Figure 3.15 does not show enough detail for the purpose of assembling a database for the positional match, but can be used for illustrating the general approach. Much more detailed diagrams on various aircraft can be found in some publications, but even these detailed diagrams have to be supplemented by photographs (and CAD models, when available) because the good diagrams typically show only a particular view of the aircraft. We must include the features at the bottom as well as the top of the fuselage, yet views of the bottom are rarely shown.

In the absence of radar data, all of the other sources combined may not provide complete information, but all we need is sufficient information to distinguish one aircraft from another.

One particularly frustrating aspect of establishing the database is the proliferation of models of the same aircraft type. The changes from one model to the next may appear to be minor. For example, they might involve placing different electronic equipment, such as communications antennas, warning antennas, and so forth, on an aircraft. Yet these are the features the




Figure 3.15 Photograph of an aircraft.

radar can actually observe. Even the mere relocation of antennas to different positions on the aircraft is important. Aircraft identification depends on distinguishing differences between various aircraft, which implies that we must be familiar with these differences if identification is to be reliable.

To illustrate the process of establishing the database, we start at the tip of the aircraft shown in Figure 3.15. It is the same aircraft used for our positional matches. There is a pitot tube at the tip. This device acts as a dipole which normally will not be observable, but under some benign conditions may generate a detectable response. The feature is too unreliable to include in the database. Even if it were reliable, the presence or absence of a pitot tube would not be of much use in aircraft identification. Thus we conclude that if the radar detects a response at the nose of the aircraft that is much weaker than the average response of the image, we will ignore it.

Next comes the metallic tip of the radome. Since the radome tip is not shaped so that it can catch the radar signal, it also will be a weak scatterer. It will sometimes be detectable, only because the background generated by the aircraft is low in the area of the radome tip. As with the pitot tube, the radome tip will be observed under benign conditions, but not in general. Similarly to the pitot tube, there is little useful information in the return from the radome tip. There are some exceptional aircraft for which this statement does not apply, but in view of the relative weakness of the return from the radome tip it does not appear worthwhile to include it in the database.

Since the radar is a much stronger and hence a much more reliable scatterer, for identification purposes we should have the aircraft start with the radar. Thus we establish the policy that for frontal views, the first scatterer utilized for identification is the radar.

The radar is a strong scatterer at least within its search angle, but also somewhat beyond. The strength of the radar return decreases as the aspect angle approaches broadside and the effective position of the radar shifts toward its mechanical base. The radar may cause multiple delayed returns, which can be most easily accommodated by the policy that only the first return is utilized if several delayed returns (in the same crossrange gate) should be observed. The phase center of the radar also is not stable, so that the accuracy of the position of the radar will be lower than for other, less complicated features. Another point relevant to positional accuracy is that the radar cannot be seen in photographs and many diagrams because it is covered by the radome. However, many aircraft diagrams do show the radar, and for other aircraft we can estimate the radar position based on this information. Mechanical and electric design constraints on the radome imply that the position of the radar within the radome can vary only little.

The fuselage section immediately behind the radome typically has one or several antennas, whose positions must be determined from multiple photographs and diagrams. One such antenna is visible in Figure 3.15. Since these antennas must work over wide aspect-angle sectors, they can also be observed over wide sectors. In general, it is necessary to develop a shadowing model that shows when a particular feature cannot be observed because of shadowing even when it would give an observable response in a particular direction [2].

The next feature that might be observable is the instrument panel. However, instrument panels usually have smooth covers that deflect the radar signal, so that they can be observed only for rear views that allow the signal to penetrate into the panel, provided the angle is not one where shadowing occurs. On the other hand, there may be a head-up display on top of the instrument panel cover, and this display may be formed such that it can trap the radar wave for certain aspects. For these aspects the display will be observable, unless a much stronger scatterer such as an electronic device is too close by.

Aside from the antennas already mentioned, which could be underneath the cockpit in front of the engine intake, there is the pilot's seat. The seat contains a metallic structure of sufficient complexity to make it observable for frontal aspects, and perhaps from the side if the particular aircraft does not have a stronger feature at that position.

Next comes the engine intake. Close to nose-on, the intake opening generally is not observable because it only represents a thin metallic ring. As the aspect angle increases, the entire opening will act as a strong cavity scatterer, but because of its size relative to the wavelength, its phase center is usually not stable. This instability is primarily a function of the aspect angle, so that the crossrange position of the intake response may not be the crossrange position of the intake opening. As already mentioned, this problem is avoided if the delayed duct responses are used to determine the crossrange position of the intake opening. Starting from nose-on, the intake is observable at about the same angle at which the delayed duct responses become observable. The intake cannot be seen close to broadside, where the radar wave cannot enter the duct. Of course, the intake opening cannot be observed for aspect angles beyond broadside.

When proceeding along the fuselage to determine observable features and their positions, more detailed diagrams and photographs will show a variety of air inlets and outlets, associated with various cooling systems. We must ask, is the opening sufficiently larger than the wavelength to generate an observable return? Is it shaped such that it can trap the wave at a particular aspect angle? Air inlets are designed to catch the air, so that they are observable by radar only for frontal views. Similarly, air outlets of sufficient size can be observed only for rear views. Figure 3.15 shows the opening for a gun to the side of and slightly behind the pilot's seat. This opening will be observed for frontal views, unless the particular aircraft carries an antenna or another strong feature close enough not to be resolvable from the gun opening.

Next comes the end of the cockpit. We must examine details of the design to determine whether or not there is some wave-trapping feature, perhaps similar to a trihedral corner or a cavity. In this case we may have one or two observable responses, but only for frontal views. The shape of these features will tell us over how large an aspect angle sector they trap the radar signal; that is, over what sector they are observable. On top of the fuselage, behind the cockpit, we usually find one or more antennas. These will be widely visible, subject to shadowing only for views from below.

The photograph in Figure 3.15 shows a feature on top of the fuselage and nearly in the center of the wings. This turns out to be a cavity for in-flight refueling. It is a relatively large cavity containing the actual refueling opening. Again subject to shadowing, the combination of cavity and refueling opening will generate two observable responses.

More detailed diagrams show other antennas on top of the fuselage, which again have wide visibility. On the other hand, antennas at the end of

the vertical stabilizer that serve functions directed toward the rear are strong only for rear views; they are unreliable scatterers for frontal views.

An examination of the bottom of the fuselage shows another wave-trapping cavity associated with an arresting hook. Again, the particular angle under which the aircraft is viewed determines whether or not the cavity is observable in the image.

A problem occasionally arises in that part of the surface may not be metallic but is transparent to the radar waves. Published information shows that this is the case for the root of the vertical stabilizer for this particular aircraft. Then we must examine whether the root contains metallic wave-trapping features, again using detailed diagrams to supplement photographs.

The preceding discussion gives an indication of what is involved in establishing a database for the positional match when radar data are not available. It is a laborious process, it requires some training, and it demands that all the information that may be available on the various aircraft is accessed. It would be valuable if a substitute could be found, but the known substitutes appear even more problematic.

The difficulty of deriving the database is increased by the facts that a particular aircraft may or may not carry a variety of ordnance, and that the same aircraft may come in different models, with different positions for some of the unrecognizable features. Much of the problem with the variability of ordnance disappears when one does not utilize scatterers on the wings. (Ordnance typically is not visible to the radar.) Accommodating different models of the same aircraft requires additional information and work. However, successful identification does not demand that all features in the database be matched with the measured features, as long as substantially more features are matched for the correct rather than incorrect aircraft.

3.4.4 Section Summary

Two types of feature are used for aircraft identification: special features, whose significance we understand (length, wingspan, number of engines, etc.), and features common to all aircraft, whose nature cannot be readily recognized from their responses (air inlets and outlets, antennas, seats, etc.). Special features can usually be measured only so inaccurately that they do not suffice for identification. We must also utilize the positions of the common features, via the positional match. The use of feature positions is practical because the features persist over relatively large angular sectors, measured in tens of degrees.

In deriving the database, the special features are obtained from diagrams, photographs, and CAD models of the aircraft. If real data should be available, images are analyzed to extract the positions (and characteristics) of the dominant scatterers. When real data are not available, the database for the common features is assembled by inspecting diagrams, photographs, and CAD models to determine wave-trapping features and their positions.

3.5 Resolution Requirements

In this section, we explain why high range resolution, but only low cross-range resolution, is necessary. Range resolution should be high enough to subdivide the fuselage sufficiently well to resolve (without the help of cross-range resolution) at least a few of the observable scatterers, so that a motion compensation can be performed. Range resolution is primarily a matter of cost, whereas crossrange resolution is limited both by available dwell time and the impossibility of performing an adequate motion compensation over long dwell times unless the aircraft's motion is uncharacteristically smooth. Thus, the primary requirement on crossrange resolution is that it be sufficient to resolve the multiple delayed engine duct returns on fighter aircraft from the fuselage returns. Because range resolution becomes progressively less effective on the fuselage as the aspect angle of the aircraft increases, aircraft identification near broadside requires very high range resolution. In the absence of extensive data with high range resolution (so that range resolution can be degraded in steps), we can only estimate the resolution requirements. The discussion again points toward the requirement of adaptive processing.

In [3], we demonstrated that one can use a sophisticated motion compensation to achieve rather high crossrange resolution, perhaps 1 ft or even better. However, no implication is intended that one can do this under all flight conditions, nor that one would necessarily want to do it. Under operational conditions one would like to work with the crudest possible crossrange resolution, for a variety of reasons. The primary limitations on range resolution are the feasibility and cost of supporting the wide bandwidths required in the transmitter, antenna, receiver, and signal processor. An additional limitation applies when range resolution is made so high that serious effects come from resolving individual scatterers into parts, but for aircraft this might imply a rather high bandwidth. With crossrange resolution, there are other practical limitations.

Most importantly, crossrange resolution requires dwell time on the target. A radar that is to identify aircraft typically is carried by another aircraft,

so that the approach velocity of the aircraft to be identified can be rather high. The radar also has tasks other than aircraft identification. It will usually have to detect the aircraft, track it, and then identify it. Under war conditions, such a radar will have to deal with more than one aircraft, possibly a large number of aircraft, at the same time. Practical aircraft radars have only a single beam to detect, track, and identify all the aircraft, and hence the time available for identifying a particular aircraft will be short. Identification thus should be performed with the shortest possible dwell time, which means with the crudest possible crossrange resolution.

Even when the radar does not operate with a shortage of dwell time, using a long imaging time for very high crossrange resolution may not be practical. Aircraft do not always fly very smoothly, in particular during war-time. There may be inadvertent yaw motion, or the pilot might execute slight back-and-forth maneuvers merely in the process of examining his surroundings. There will often be intentional maneuvers, and these may involve rather erratic changes in angle, and hence in range rate. Our experience indicates that it is highly doubtful that one can achieve a satisfactory motion compensation under all of these conditions, and over times as long as might be needed to achieve high crossrange resolution. The motion compensation might be good enough to generate an intensity image that looks like an aircraft to the eye, but such an image is unlikely to allow reliable identification. One needs an image good enough to allow feature measurements, in particular the measurement of feature positions. Under some conditions, it is impossible to generate images with high crossrange resolution and such a quality that accurate feature measurements can be performed. Not only is a high-quality image with poor crossrange resolution preferable to a poor-quality image with high nominal crossrange resolution, but the latter image will not allow reliable target identification.

With expense restricting range resolution, and dwell-time requirements and aircraft maneuvers limiting crossrange resolution, the practical question is how much range and crossrange resolution one needs for reliable aircraft identification. A satisfactory answer to this question would require the following work. One would have to design an aircraft identification system with higher range and crossrange resolution than needed, test the system on real data extensive enough to represent all practical conditions, and then degrade range and crossrange resolution in steps until the identification performance was no longer satisfactory. The greatest problem with this approach is the data availability. At the time of this writing, no reasonably extensive database with a range resolution high enough for this test (probably in the order of half a foot) exists, and the databases with a range resolution of 1 ft all have

radar-related problems, such as inappropriate waveforms. Thus we can merely estimate how much resolution is needed for reliable aircraft identification.

It is clear that one needs higher resolution to identify small aircraft than large aircraft. Thus it is reasonable to consider the resolution requirements for fighter aircraft. As discussed above, such aircraft typically have between 20 and 30 observable features on the fuselage, and we explained earlier why one must use, and hence resolve, the observable features on the fuselage.

Considering the length of a fighter aircraft, the number of observable fuselage features, and their typical separation, one concludes that at small aspect angles the range resolution needed might be on the order of 1 ft. If range resolution is high enough to resolve, or almost resolve, the scatterers on the fuselage, the demands on crossrange resolution are minor. This is true only for aspect angles that are not too large. As the aspect angle becomes larger and range resolution is less effective along the fuselage, better crossrange resolution is needed in order to resolve the fuselage features. Near broadside, range resolution becomes so ineffective along the fuselage that the scatterers have to be resolved almost solely by crossrange resolution. This is very unfortunate, because near broadside the implementation of high crossrange resolution becomes problematic or even impossible. Hence, if aircraft are to be identified near broadside, a range resolution better than 1 ft is likely required. Our conclusion is that 1-ft range resolution may be sufficient if identification near broadside is not needed, and that range resolution may have to be improved to perhaps 0.5 ft if identification is to be performed near broadside.

Little crossrange resolution is needed at the smaller aspect angles if range resolution is on the order of 1 ft. We do want to resolve the wingtips from the fuselage, because responses from wingtips with the capability to carry missiles might interfere too much with fuselage responses. The same is true of any wing-mounted ordnance that backscatters strongly. Fortunately, this occurs only for small aspect-angle regions. We also need to resolve delayed duct returns from most fuselage responses. Beyond these requirements, we must keep in mind that the ultimate goal is measurement accuracy. Although a measurement with high accuracy can be performed only if a response is sufficiently well resolved, it does not matter in what manner it is resolved. Thus, if range resolution is adequate to resolve the individual features, we can perform accurate crossrange position measurements without high crossrange resolution. This is the difference between the width of the response and the accuracy with which we can measure the position of the response peak. If a response is at all detectable in the background, the peak

position can be measured with an accuracy of about 10% of the response width. Since the coherent integration gain from Doppler processing increases the signal-to-noise ratio, even for the weaker observable responses we should be able to measure the peak position to perhaps 5% of the peak width. The accuracy achievable for two unresolved scatterers is not so high, but is still a small fraction of the response width.

Our work indicates that for the kind of feature separation found on the fuselages of small fighter aircraft, a measurement accuracy of about 0.3 ft might be sufficient for reliable aircraft identification. This measurement accuracy is easily achieved in range if range resolution is 1 ft and the background is noise or noise-like. This is usually the case with resolved observable features. Potential problems with range accuracy arise from interfering scatterers, but then we must choose resolution high enough to make this interference acceptable. With respect to crossrange resolution, if the measurement accuracy is to be 0.3 ft, then the above discussion implies a crossrange resolution between 10 and 20 times larger, or between 3 and 6 ft. This is so if the burden of resolution is carried by range resolution; that is, for aspect angles that are not too large. As already stated, near broadside we have to increase crossrange resolution to replace range resolution for the fuselage, so that we might need 1 ft of crossrange resolution near broadside. But this requirement becomes academic if identification near broadside is impossible for other reasons.

The preceding discussion again points toward the need for adaptive processing when moving targets are to be identified. At smaller aspect angles we can choose the dwell time so as to implement a crossrange resolution of between 3 and 6 ft, and as the aspect angle increases we have to increase crossrange resolution, in the limiting case near broadside to about 1 ft (but only if a good motion compensation is achievable with the existing motion behavior of the aircraft). The actual numbers to be used in each situation will depend on the type of aircraft, its aspect angle, and its flight behavior. This means that the automated identification system must determine in real time how much crossrange resolution is needed, by performing a test to determine whether a particular imaging time gives satisfactory crossrange resolution, and increasing the dwell time if this is not the case.

Practically, we can implement these resolution constraints through three simple requirements: wingtip responses must be resolved from fuselage responses in the same range gates; delayed duct responses must be resolved from 90% or more of the fuselage; the fuselage must extend over at least 20 range/crossrange resolution cells. Figure 3.16 shows a flowchart of adaptive data collection, based on these requirements and the realization that an

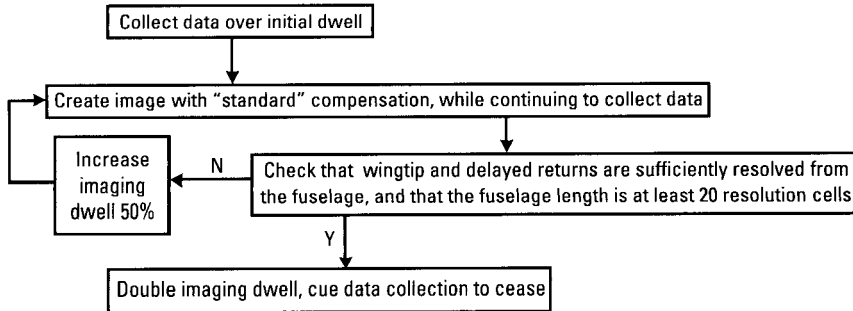


Figure 3.16 Flowchart of adaptive data collection.

operational radar should not dwell on the target longer than necessary. The “standard” compensation referred to in the figure implies any simple compensation that generally allows one to recognize the generic shape of the aircraft well enough to distinguish the fuselage, the wingtips, and delayed duct returns. Compensation on the basis of range centroid or correlation tracking, followed by Doppler centroid tracking, is sufficient.

3.5.1 Section Summary

Range resolution must be sufficient to resolve most of the dominant scatterers on the fuselage; that is, the fuselage responses must be analyzable by the TSA. Crossrange resolution must be sufficient to resolve the wingtips and the delayed duct returns from the fuselage returns.

3.6 Imaging of Aircraft

This rather extensive section treats aircraft imaging under various conditions, given the requirement that the quality of an image must be high enough to allow the measurement of scatterer positions. After a discussion of the imaging principles, we give illustrations of the imaging of aircraft with and without delayed duct returns. Next, we consider imaging under problematic conditions, such as a maneuvering aircraft and the special cases of nose-on and tail-on aspects. In order to show how much easier it is to image and identify large commercial aircraft, we include appropriate examples. Sections 3.6.1 to 3.6.3 are of general interest, whereas the treatments of the more problematic imaging conditions should be of interest to those actively working in the field.

3.6.1 Aircraft Imaging Principles

Aircraft move fast, so that little time is available for their identification, in particular since the radar must search, track, and probably identify more than one aircraft at a given time. An aircraft also does not move as smoothly as might appear to be the case, unless it is on a high altitude course from point A to point B and is a large aircraft. We cannot claim to have investigated the flight behavior of the various types of aircraft, but have seen very smooth flight only with commercial aircraft. When aircraft identification is critical, which it is for military operations, an aircraft will have inadvertent yaw and roll motions, even when the pilot is not executing a deliberate maneuver, and these motions will not necessarily be smooth. As will be demonstrated below, such flight behavior makes it difficult to generate high-quality images with sufficient crossrange resolution to separate wingtips and delayed duct returns from the fuselage returns. There is always the problem of spurious responses, and it increases as the motion becomes more erratic, in particular if the motion is in three dimensions rather than just in the horizontal plane.

In order to identify an aircraft, we thus want to generate an image that meets the following requirements: (1) the time over which the aircraft is observed is as short as possible; (2) the motion compensation is good enough to permit analyzing the image responses and determining the positions of the associated scatterers; (3) crossrange resolution is high enough to allow separating the delayed duct returns from the skin returns, but not so high that the quality of the image is degraded; and (4) the image contains few spurious responses strong enough to be mistaken for genuine aircraft responses.

We first consider the motion compensation. Its main purpose is to make one of the scatterers on the aircraft appear perfectly stationary, so that the aircraft is rotating about this reference scatterer. A secondary purpose is to generate an image of the rotating aircraft in which another scatterer (not too close to the first) can be selected, so that one can measure the precise way in which the aircraft rotates about the stationary reference scatterer. If the rotation rate is not perfectly constant, which it typically will not be, we use the measurement of rotation either to resample the data so as to make the rotation rate appear constant, or to select a shorter imaging interval within which the rotation rate is sufficiently constant. The latter is the preferred method if it results in at least adequate crossrange resolution, because we want to use the shortest possible imaging interval. If the rotation axis changes direction, we must track at least three scatterers to determine the motion. Under such conditions, choosing a shorter imaging interval within which the rotation axis and rotation rate are constant is even more crucial. We first treat the motion compensation steps needed to arrive at a stationary reference scatterer.

The first requirement of the motion compensation is that the gross range changes due to the motion of the aircraft along its flight path be removed. There are several equivalent ways, at least in principle, of measuring the range changes. One can degrade range resolution to include the entire target within one range cell. The measured range of the target is then the range of the range centroid. Alternatively, one can correlate consecutive range profiles to determine in what manner the range is changing. Another procedure is to track a prominent scatterer in the range profile. The choice will be governed by consideration of the signal-to-noise ratio. If one uses range centroid tracking or correlation, one should remove range gates beyond the range spread of the target to improve the signal-to-noise ratio. Whatever method is employed, one must fit a smooth function to the measured range values, and this function is used to remove the range drift from the data. It is important that the fitted function not be too flexible, because variations introduced by fitting to the random range changes in the measurements may not be removable in the next motion compensation steps. The consequences of such variations are distortions of the image responses and the introduction of spurious responses. Both tend to prevent the measurement of true scatterer positions.

Both tracking of the range centroid and the correlation of consecutive range profiles are different from tracking a prominent scatterer, because the former two methods utilize the entire target. From the point of view of signal-to-noise ratio, it is not advisable to rely on range tracking a single scatterer at the outset. Instead, one should follow range tracking of the entire target by Doppler tracking of the entire target, which means Doppler centroid tracking. Here again one must pay attention to the flexibility of the function fitted to the measurements, even though it is not as critical as in the case of the range compensation. Doppler tracking amounts to sliding-window tracking and hence generates an integration gain. After Doppler tracking of the entire target, the Doppler spread of the target will have narrowed sufficiently that, by eliminating the Doppler gates outside the Doppler spread of the target, one can further improve the signal-to-noise ratio (of the corresponding range profiles). Only then should the next step of range tracking a prominent scatterer be implemented. The result is a range track of the reference scatterer to an accuracy related to the width of the range cell. This is far too crude for our purposes, which require that the scatterer be tracked in phase.

One does not want to go directly from range tracking of a scatterer to tracking its phase. *It is better to include an intermediate step of Doppler tracking the scatterer.* Doppler tracking and subsequent compensation will focus the response of the scatterer sufficiently well in Doppler (or crossrange)

to resolve it from other scatterers at the same range. Then one can take a fixed-range image cut, obtain the transform of a window about the response, and use the phase of the transform to determine the motion of the scatterer. This is illustrated in detail in Section 2.3.

The described steps of the motion compensation needed to arrest a reference scatterer are critical. Since imaging conditions are continuously changing, *each step after the crude range and Doppler centroid compensation must be checked before proceeding*. Doppler tracking of the reference scatterer is used to verify the success of range tracking, because Doppler tracking will work only if the range track was sufficiently accurate. Doppler tracking of the reference scatterer is verified by taking the transform of the response and determining whether or not the amplitude function is sufficiently constant (see Section 2.3). These checks prevent one from forming an image whose responses are too poor for measuring scatterer positions.

Since there is no “standard” situation for the motion compensation, the motion compensation steps described above must be used adaptively. The minimum compensation is range centroid tracking followed by Doppler centroid tracking. When an aircraft is flying very smoothly, this simple motion compensation will produce an image with responses of such a high quality that scatterer positions can be measured. Whether or not this is the case is tested by taking transforms of some fixed-range image cuts through responses, and determining whether or not the pattern from single scatterers or two scatterers are sufficiently well approximated. If this is the case, the basic motion compensation is finished. One still must check scatterers at the extremities of the aircraft to determine whether polar reformatting might be necessary. If this is the case, the amplitudes of the transforms of the fixed-range image cuts may not stay strong over the entire imaging interval and the corresponding phase will have significant nonlinearities, with both variations increasing with distance from the compensated scatterer.

When, after range and Doppler centroid tracking, the image responses are found not to have sufficient quality for the measurement of scatterer positions, the motion compensation is continued by range tracking and Doppler tracking of a reference scatterer, and compensating the data with these tracks. Again, another response in the image (not too close to the reference scatterer) is checked. If the phase function of the transform has a kink, one should use the larger of the two subintervals divided by the kink to generate a new image, reducing nominal crossrange resolution in order to improve image quality. The quality of the responses is then checked as described above. To relate the choice of the imaging interval, and crossrange resolution, to the

practical operation of a radar, let us assume for the moment that the radar has a specific observation interval available per aircraft, such as one second. One might be tempted to conclude that in this case the image should be formed over the entire observation interval, in order to achieve the highest crossrange resolution. This would ignore a variety of serious problems, aside from the fact that in practice we should minimize the observation time per aircraft. First, we cannot expect an aircraft to fly so smoothly that a good motion compensation can be achieved over the relatively long time of one second. An erratic motion of the aircraft, or just some temporary problem in tracking a scatterer, can easily lead to a motion compensation of insufficient quality. Genuine responses from the aircraft may be distorted so that the measurement of scatterer positions becomes too inaccurate, or the deviation between the measured and actual motion may cause too many spurious responses. The image may have a high nominal crossrange resolution but totally insufficient quality. Moreover, even the measurement of scatterer ranges will be affected.

The remedy is not simply to shorten the imaging time so as to decrease the severity of the effects. A particular observation interval often contains times at which problems appear, be problems from the motion behavior of the aircraft or from inaccuracies of tracking a scatterer. At such times it is impossible to generate a high-quality image. Thus, if an observation interval of, say, one second is available, we will compensate and form an image over the entire observation interval, but treat it only as a kind of survey image needed to understand the situation. We examine the image and determine subintervals in which high-quality images can be generated. The best of these subintervals is selected for imaging, as indicated earlier in relation to the last motion compensation step.

Although identical in principle, the actual implementation under operational conditions must be different (as in Figure 3.16), because the radar should not collect more data than actually used for imaging. In a practical application, when the radar starts collecting data on an aircraft, the processor must form an image as soon as some useful amount of data has been collected. This image is examined to determine (1) whether crossrange resolution is sufficient to resolve the wingtips and any delayed duct returns from the fuselage, and (2) whether the fuselage extends over at least 20 resolution cells. As soon as the requirements for crossrange resolution are met, the radar is cued to double the sufficient dwell, then stop collecting data. The doubling of the dwell is necessary because irregular target motion often prevents the formation of an image of acceptable quality from the minimum dwell necessary for resolution.

In order to check the image quality and select an appropriate time interval for imaging, we must examine fixed-range image cuts through fuselage responses. As discussed in Section 2.3, in order to determine the motion history of a scatterer, we must be able to define a window in the image cut such that a Fourier transform of the window gives essentially constant amplitude, so that we can phase track the scatterer. If we can find such a scatterer, we examine the transform phase. Breaks in the phase slope indicate jerky target motion that is difficult to compensate, and we use such breaks to define boundaries of acceptable imaging subintervals.

If no such response exists, we attempt to improve the motion compensation, by range and Doppler tracking a fuselage scatterer. We repeat the search for an acceptable transform amplitude. If none is found, we iterate the compensation and search sequence until one is found or no trackable scatterer remains. In the latter case, we reduce the data duration by a factor of two, and repeat the compensation and search process with the reduced duration.

Having found an acceptable transform, we search for another in image cuts well-separated from the first. This second motion history will enable us to determine times of abrupt changes in target rotation, as well as abrupt changes in translation. The search for the second proceeds as did that for the first, with the difference that phase-slope tracking is acceptable for the second scatterer. As with the first scatterer, breaks in the phase slope of the second scatterer define boundaries of acceptable imaging intervals.

We next form an image over the longest interval without a phase-slope break for either response. As this is generally shorter than the total data collection time, we check that the minimum resolution requirements are still met. If not, we increase the imaging interval, adding the minimum phase-slope break possible. We continue expanding the imaging interval until the minimum resolution requirements are met, or until the imaging interval is the entire interval used for phase tracking. If the latter occurs before the former, we cannot identify from the data collected.

If our image does have at least the minimum required resolution, we then check that the quality of responses is acceptable. This is particularly of concern if we had to choose an imaging interval including a phase-slope break. The quality of responses is acceptable if fixed-range image cuts through a majority of fuselage responses yield acceptable one- or two-scatterer patterns. If the quality is unacceptable, we check whether the image has an excess of resolution (relative to the minimum). If so, we reduce the imaging interval, in proportion to the excess, and repeat the check on response quality. If not, we cannot identify from the data collected.

Once we have an image of acceptable quality, we track two scatterers near the ends of the fuselage and use their differential motion to determine the deviation from uniformity in the rotation rate and the range drift of the ends of the fuselage. If these require correction, we resample and polar reformat the data, respectively. Finally, we search for a group of responses with common phase curvature, and remove such [3].

The process of selecting an appropriate time interval for imaging, described above, is summarized in the flowchart in Figure 3.17.

3.6.2 Illustrations of Aircraft Imaging Without Delayed Duct Returns

As explained above, in practice the observation time should be varied in accordance with the existing situation. Since it would serve no purpose here to demonstrate the entire process of adaptively selecting the appropriate observation time, we will simply assume that an observation time of one second is available, and demonstrate how to select the appropriate imaging time within that observation time.

A one-second image of an aircraft after range centroid and Doppler centroid compensation is shown in Figure 3.18. With the wingtips clearly visible in this case, we can choose approximately the correct crossrange scale by making the lines connecting the wingtips be perpendicular to the centerline through the fuselage. This was done with the figure, so that the crossrange scale is approximately correct. The image shows that the aircraft is viewed at a small aspect angle, which is usually an easy case for identification. In order to check the quality of the motion compensation, regardless of what the source of any problem might be, we select a response on the fuselage, take an image cut in the range gate of the response, and examine the transform of the image cut. Evaluating the motion compensation requires an approximately constant transform amplitude. If our selected response has a variable amplitude, we must examine other responses on the fuselage. If none has approximately constant amplitude over the entire observation interval, we proceed with the response with the longest constant amplitude interval. Within this interval, we want to find a subinterval with essentially constant phase slope. (A change in the phase slope amounts to a translation of the response in crossrange, so that the scatterer response would appear in two different crossrange positions.) *A smoothly changing phase slope indicates that more compensation is necessary, but that the subinterval is usable for imaging.* Because the aircraft may be rotating irregularly about our constant amplitude response, which would not be evident from its signal phase, if the examination of its phase shows no problem we want to perform the check on a

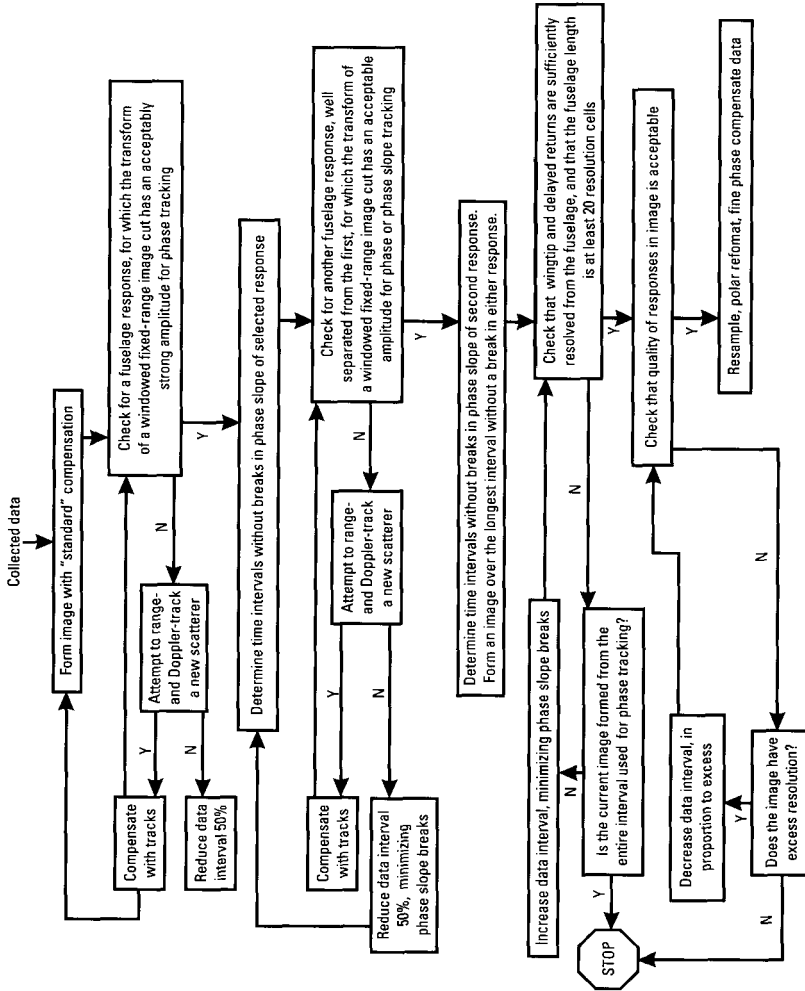


Figure 3.17 Flowchart for selecting imaging interval.




Figure 3.18 Image after standard compensation.

second response. The two responses should be separated in range as much as possible, so that the compensation cannot have corrected any kinks in both phase functions. Also, both responses must be on the fuselage.

The transform of the image cut in the range gate of the tip response is shown in Figure 3.19. The phase slopes for the first and second halves (roughly) are 0.06 and -0.21 , which with our normalization means that the responses are separated by 0.27 crossrange gates. Thus there is some smearing of the response. On the basis of the amplitude function, we might choose an imaging interval from 0.1 to 0.4 seconds of normalized time. From the transform of the image cut in Range Gate 15.5, shown in Figure 3.20, we see that this is also a good choice for the rear of the aircraft, since the amplitude is nearly constant and the phase function shows adequate linearity. The image over this subinterval is shown in Figure 3.21. We know from the way the imaging time was selected that the image has higher quality than the image of Figure 3.18, but ultimately this can be judged only when the responses




Figure 3.19 Transform of the image cut in Range Gate -22.




Figure 3.20 Transform of the image cut in Range Gate 15.5.

are analyzed. However, a closer investigation shows that the selection of the imaging time is not critical in this particular example.

We repeated the process of imaging interval selection with the alternative approach, where a prominent scatterer is first tracked in range and




Figure 3.21 Image from 6.6 to 6.9 seconds.

then in Doppler, instead of using centroid tracking on the entire aircraft. The precise tracking procedure is demonstrated in detail in Chapter 5 for ships, for which it is very important. In this case, the first check for constant amplitude is performed on the tracked scatterer. The two transforms corresponding to Figures 3.19 and 3.20 are so similar that the same imaging interval would be selected. The peaks plot of the resulting image is nearly indistinguishable from Figure 3.21. In principle, the two images need not be alike, because they use different compensations, but in this instance the tracking is easy and hence there is no difference for practical purposes.

In our next example, the selection of the precise imaging interval is more critical. In Figure 3.22 we show a one-second image of the aircraft after range centroid and Doppler centroid compensation. The multiple dots within single range gates, along most of the fuselage, already indicate that the image quality is far too poor for extracting scatterer positions. To verify this, and to select an appropriate subinterval, we must examine fixed-range cuts through strong fuselage responses. Figure 3.23 shows the image cut through the strong response in Range Gate -22.5 . Although this is the best isolated




Figure 3.22 Image after range and Doppler centroid compensation.

response on the fuselage, interference from a second scatterer is evident in the modulation of the transform amplitude and phase. The phase modulation is strong enough to obscure significant phase variations due to irregular target motion. In order to reduce the modulation, we must filter out the interfering scatterer.

Figure 3.24 shows the result of notching out the interfering scatterer to the left of the strong response in Figure 3.23, and then taking the Fourier transform. The absence of slow modulation in the amplitude and phase indicates that we have effectively removed the interfering scatterer. Clearly, we have also removed some contribution from the strong scatterer. However, because we removed just a narrow notch from the tail of the strong response, the slow variation of the phase of Figure 3.24 still corresponds to the slowly varying scatterer motion. The smooth curve of Figure 3.24 is a quadratic fit to this motion. After subtracting this fit, we can more easily visually recognize abrupt changes in the phase slope.




Figure 3.23 Image cut in Range Gate -22.5 .




Figure 3.24 Transform of filtered image cut of Figure 3.23.

However, we can do better than subtracting the fit. The fit corresponds to motion that widens both the strong response and the interfering response. If we use the fit to motion compensate the data, the two responses will be better resolved in the compensated image, so the interfering response can be more effectively filtered out. Figure 3.25 shows the result of taking a fixed-range image cut through the strong response in the compensated image, filtering out the interfering response, and taking the Fourier transform. The phase shows three intervals (relative times -0.5 to -0.15 seconds, -0.15 to 0.25 seconds, and 0.25 to 0.5 seconds, indicated by dotted vertical lines) of roughly linear phase, separated by abrupt changes of phase slope. These changes are too large to be acceptable. *If subintervals are to be combined into a single image, the phase slopes (which correspond to scatterer positions), can differ by at most a small fraction of one gate.*

Figure 3.26 shows the result of applying the same processing (filtering, compensating, and filtering) to the strong response in Range Gate 7.4 of Figure 3.22. The phase shows the same three intervals as in Figure 3.25, with the same phase-slope changes between intervals. This shows that the target is translating (rather than rotating) jerkily. The larger deviations from linearity




Figure 3.25 Transform of filtered cut in compensated image.




Figure 3.26 Motion derived from scatterer in Range Gate 7.4 of Figure 3.22.

within the intervals of Figure 3.26 are due to stronger interference in Range Gate 7.4 than in Range Gate -22.5 , but are small enough to allow phase-slope measurements.

When images are formed over the three intervals, to the eye the differences are minor. All three appear to be of good quality. However, we have repeatedly stated that the peaks plot image allows only a crude judgment of image quality, and that the actual responses must be examined. The examination of fixed-range image cuts through strong responses verifies that all three images are of acceptable quality.

To illustrate the problem of identifying an aircraft at an aspect angle near broadside, we next show three images generated by tracking the range and Doppler centroids, tracking a scatterer in range, tracking the same scatterer in Doppler, compensating the data with the tracks, forming an image over the entire one second, selecting a good imaging interval, and generating the final image. The image at an aspect angle of about 65° is shown in Figure 3.27. The aircraft is in a turn and banked, which is easily reconcilable with the image. In Figure 3.28 we show the image for an aspect angle of about 80° , and in Figure 3.29 the image essentially at broadside. With the single-scatterer tracking procedure used for these examples, the last two peaks plots do not as clearly resemble an aircraft as for smaller aspect angles. We




Figure 3.27 Image at an aspect angle of about 65° .




Figure 3.28 Image at an aspect angle of about 80° .

have not investigated whether the image details at such large aspect angles are sufficiently accurate for aircraft identification.

As the next case, we chose a one-second interval over which the aircraft is flying very smoothly, which helps the motion compensation. The




Figure 3.29 Image at broadside.

corresponding one-second image is shown in Figure 3.30. A visual examination lets one recognize the shape of the aircraft. This is of little help for automated processing, because one cannot rely on being able to generate such a high-quality image with excellent crossrange resolution. With automated processing, we must check the transforms of responses as usual. The fixed-range image cut and transform of the strongest response in the image are shown in Figure 3.31. The relative half-power width is 1.055, which means good compression. The second half of the transform amplitude droops, which has little consequence, as verified by the half-power width of the response. Although the phase function seems to have two breaks in the phase slope (ignoring the behavior at the fringe), the total phase change is only about 0.05 cycles, so the phase nonlinearities are negligible. When these parameters are (automatically) measured, for this scatterer and one near the rear of the aircraft, the conclusion is that the image is of good quality, and no shortening of the imaging time is needed.

With this image of rare quality, we can measure aircraft length with great accuracy even in the absence of an aspect angle measurement from the tracker. We adjust the crossrange scale so that a straight line connecting the first and last point along the fuselage is perpendicular to the line correcting




Figure 3.30 Image over one second.




Figure 3.31 Image cut and transform for Range Gate 0.1.

the appropriate wingtip responses. With the correct scale factor, the direction of the fuselage gives the aspect angle. Taking the range difference between the first and last points on the fuselage and dividing by the sine of the aspect angle, in this instance we obtain the length of the aircraft with an error of 2 cm. In principle, we can also measure the wingspan in this fashion, except that the aircraft is banked. This bank angle can be estimated from the track by assuming that the combined acceleration vector is oriented perpendicular to the plane of the wings. However, we do not know how reliable this assumption is. In this particular case the measurement error was 70 cm.

An image with such high resolution and good quality allows one to measure the scatterer positions very accurately, a fact that is recognized during the assessment of TSA performance on every response. To demonstrate how much more difficult the task becomes when crossrange resolution is poorer, in Figure 3.32 we show the image when only the central half of the imaging interval for Figure 3.30 is used. It is far more difficult to recognize the outline of the aircraft. Nevertheless, with respect to the measurement of scatterer positions, this is still a good image.

We want to make another general point. Figure 3.33 gives a one-second image at a later time. The image is formed over another time interval where the motion of the aircraft happens to be quite smooth, so that the image appears to be of high quality. Nevertheless, when we check image cuts in the range gates of well-resolved responses, we find that the relative half-power widths are larger than they should be. Thus we again examine the transform of the image cut of a suitable response, finding that the imaging




Figure 3.32 Image over the central half of the imaging interval.




Figure 3.33 Image from 121.0 to 122.0 seconds.

time should be restricted. Instead of the imaging interval from 121.0 to 122.0 seconds, based on the transform we select the shorter interval from 121.3 to 121.9 seconds. The corresponding image is shown in Figure 3.34. In the new image the relative half-power widths of the responses are as low as they should be, indicating good compression. This is important because a poor motion compensation that causes the TSA to result in two scatterer positions rather than one can be catastrophic for aircraft identification.

We want to point out again that the TSA can be used only when the motion compensation is of sufficient quality; otherwise the advantage of the complex over the intensity response cannot be utilized. The point we want to make is the following. The one-second image of Figure 3.33 was generated by tracking a specific scatterer. When we choose an alternative scatterer for tracking, the details of the intensity image will be somewhat different. However, if we again examine the transforms of fixed-range image cuts through responses to select a more appropriate imaging time, the differences between the images over the shorter times will be insignificant. Slight differences may appear in the intensity images, but when the responses are analyzed to determine scatterer positions, the differences will be negligible.




Figure 3.34 Image from 121.3 to 121.9 seconds.

3.6.3 Illustrations of Aircraft Imaging With Delayed Duct Returns

The examples in the preceding section are perhaps somewhat unrealistic, because they were given for an aircraft without delayed duct returns. If we examine the same aircraft when viewed from the rear, the situation is more realistic because of the delayed exhaust returns.

In Figure 3.35 we show an image with the delayed exhaust returns. The image was generated in the same way as explained above, but a short imaging time of 0.2 seconds was chosen in order to demonstrate the checks that must be performed before such a short imaging time is accepted. First, *the delayed duct returns must be sufficiently resolved from the fuselage returns not to interfere with the measurement of scatterer positions over more than about 10% of the fuselage.* This cannot be judged from the peaks plot, because the dots can represent very small local maxima of unresolved composites of responses. We must take image cuts in the range gates of responses close in crossrange to the duct returns, and determine whether or not a sufficiently accurate measurement of the crossrange position of each response can be performed. Figure 3.35 shows a fuselage response in Range Gate -10.5 and Crossrange Gate 2.5 . The image cut in Range Gate -10.5 is given on the left side of Figure 3.36, amplitude and phase function, with its transform on the right side. Although the cut shows a single response peak with a bulge on its left, there is no resolution problem. This is verified by the fact that the transform




Figure 3.35 Image from 32.1 to 32.3 seconds.




Figure 3.36 Image cut in Range Gate -10.5 and transform.

extends over about two modulation cycles. The conclusion is that crossrange resolution is sufficient with respect to the delayed duct returns.

For comparison, we now reduce the imaging interval from 0.2 seconds to 0.1 seconds, effectively halving crossrange resolution. The corresponding image is shown in Figure 3.37. In this image, when we check the resolution of the response nearest the duct returns, in analogy to Figure 3.36, resolution is still acceptable. This is to be expected because in Figure 3.36 the responses were separated by more than two crossrange gates, yet the imaging interval was reduced only by a factor of two. On the other hand, if the response doublet near Range Gate -1.5 in the image of Figure 3.37 is checked, we obtain the image cut and transform of Figure 3.38. It is an example of insufficient resolution. The transform does not have an amplitude/phase pattern from which the scatterer positions could be extracted with adequate accuracy. Checking other responses also verifies the conclusion that crossrange resolution is insufficient for this image.

As the next example, we examine the problem of imaging in the presence of duct returns when the aspect angle is near zero (at zero aspect angle there are no delayed duct returns), where crossrange resolution is particularly critical. Compensating one half-second of data in the same way as demonstrated with the earlier examples, we obtain the image of Figure 3.39. In this case the aircraft is viewed only a few degrees off tail-on, so that resolving the




Figure 3.37 Image with half the imaging time.

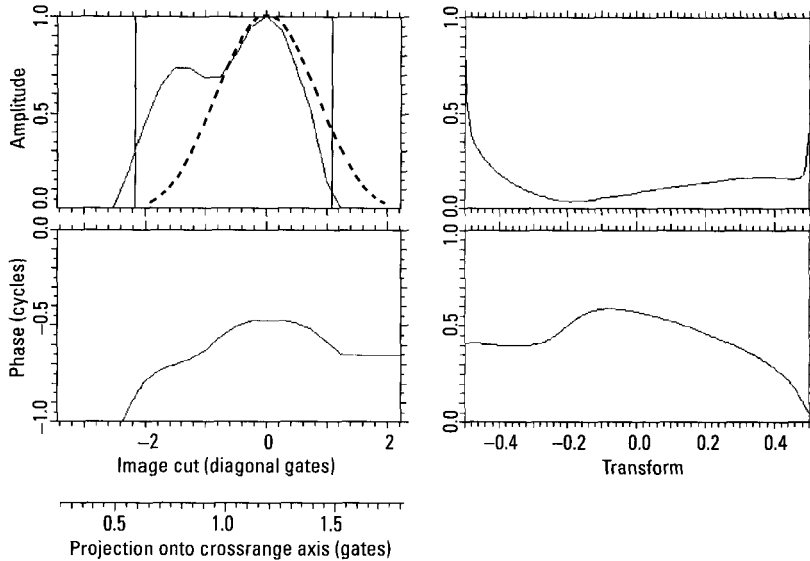


Figure 3.38 Checking resolution.

duct returns from the fuselage returns is problematic. (The crossrange scale was intentionally stretched, for a better separation of duct and fuselage returns.) We would like to use high crossrange resolution, but this often requires imaging intervals so long that the flight of the aircraft may not be sufficiently stable for a good compensation. Hence, if we increase the imaging interval, we obtain a higher nominal crossrange resolution, but the quality of the image becomes too poor for aircraft identification. Before attempting a high resolution compensation (or cueing an operational system to continue collecting data), we check whether the resolution is sufficient.

We indicated the direction of the fuselage in Figure 3.39 by the arrow (note that the crossrange position of the exhaust is best given by that of the delayed returns). Too few responses appear to be centered on this line or at least near the line. As an example of how misleading an intensity image can be, in Figure 3.40 we show the image cut and transform for Range Gate 4.75, where no response appears to be near the centerline. This would be Crossrange Gate 0.8. However, the complicated response at the left top in Figure 3.40 clearly indicates a response to the right of the peak. The positions of the interfering responses can be found with the TSA, using a transform window as shown in Figure 3.41. The left half gives the windowed response, and the right half the amplitude and phase functions of the




Figure 3.39 Image when the aircraft is viewed nearly tail-on.




Figure 3.40 Image cut in Range Gate 4.75.




Figure 3.41 TSA measurement.

transform, as usual. In addition, the idealized patterns of the amplitude and phase functions for the measured parameters are superposed (labeled “m” for model), as an indication of the quality of the measurement. In this instance the quality is not high, because the left boundary of the transform window cuts through two (destructively interfering) responses. Nevertheless, one measured crossrange position is Gate 1.2, very close to the centerline. The situation is the same with the responses in Range Gates 3.9 and 2.8, where the TSA also gives two scatterer positions, one for the delayed duct return and the other on the centerline.

The preceding results should not be taken as an indication that there is no resolution problem that cannot be solved. Compared with a visual inspection of the intensity image, the TSA can do no better than improve resolution by a factor of two, perhaps a little better with the use of a small degree of extrapolation of the amplitude and phase functions of the transform. In practice we will not use the TSA for a measurement in one dimension only, as in this illustration, but will measure the two-dimensional position in range and crossrange. Even so, if duct and fuselage responses are separated by less than about one resolution cell, we cannot accurately measure the fuselage scatterer positions. In our example, this is the case for about the rear quarter of the

fuselage. Thus, identification must be based on the front three-quarters of the aircraft. Since the aspect angle is nearly tail-on in Figure 3.39, one would expect a severe shadowing problem for the front part of the fuselage. Although there is indeed shadowing, it appears to be less than at optical wavelengths. The waves travel along the smooth surface of the aircraft and backscatter when they hit one of the wave-trapping features or an antenna. For example, we verified that the last response along the centerline in Figure 3.39 is at the position of the radar in the nose of the aircraft.

As stated above, we would like no more than 10% of the fuselage unresolved from duct returns. Our low resolution image has 25% unresolved. Therefore, we should attempt to form an image over a longer interval. However, in this case, the aircraft motion is erratic enough that no appreciably longer interval is usable. Thus, we must attempt identification on the basis of the image of Figure 3.39. Should this not lead to any target declaration, we must collect more data.

3.6.4 Imaging of Maneuvering Aircraft

Although the aircraft used in the above examples were flying in loops, the loops had a diameter exceeding 5 km, and they were flying with a steady slow turn. This is not essentially different from an aircraft flying steadily along a straight path, except that an aircraft in a loop will be banked. In this section we will consider truly maneuvering aircraft, where an aircraft is flying along a straight path and starts turning with a much smaller turn radius than 5 km. Imaging must then deal with two effects, the transition effects when the aircraft starts the turn and the turn maneuver itself.

In Figure 3.42 we show the range profiles over three seconds, starting a little before the aircraft begins the turn. For better visibility, we compensate the changing range for these range profiles and expand the range scale. The result is shown in Figure 3.43. For about the first half second the consecutive range profiles are reasonably well correlated, because the aircraft is still flying straight. Then the range profiles become almost noiselike, with some improvement in the last section. This improvement consists mainly in that the first and last peaks become better correlated over consecutive range profiles. Also noticeable is an increase in the range extent of the range profile with time, since the aspect angle of the aircraft is becoming smaller. We will separately consider one-second sections starting at 36.0 seconds, 37.0 seconds, and 38.0 seconds.

The peaks tracks for the first one-second section are shown in Figure 3.44. These are not good tracks, of course. However, as explained earlier we will




Figure 3.42 Range profiles around the beginning of the turn.

start with the crudest method of range tracking these peaks of the range profiles, rather than first analyzing the complex peaks and deriving actual scatterer positions, or using combined range/Doppler tracking. We will select what appears to be the best peaks track by eye. With automated software we can track one of the peaks, range- and Doppler-compensate the response, form an image, and search for fixed-range image cuts through two responses, whose transforms give sufficiently constant amplitude functions to use the phase functions for measuring the motions of the scatterers. If response peaks that permit these measurements cannot be found, the software can start with a different range track, and so on, until an image is found in which the transforms of two responses do give sufficiently constant amplitude functions. The success also depends on the correct choice of the order of the polynomial for fitting to the measurements.

Perhaps the best peaks track in Figure 3.44 is the one starting in Range Gate 9 (indicated by the arrow), which is separately shown in Figure 3.45,




Figure 3.43 Range profiles after compensation.




Figure 3.44 Peaks tracks from 36.0 to 37.0 seconds.




Figure 3.45 Peaks track used for the initial range track.

together with a third-order polynomial fit. It is an acceptable track because the fit is not flexible and the range variations do not exceed about one range gate. After compensating with the polynomial fit, we obtain the Doppler track of Figure 3.46 on the range-compensated scatterer, again with a third-order fit. The fact that the Doppler track is a smoothly curved function




Figure 3.46 Doppler track of the range-compensated scatterer.

already indicates that the range compensation is of sufficient quality. With a further Doppler compensation of the scatterer, we obtain the image of Figure 3.47. Although the image quality is poor, as is evident by comparing the dots to the dotted generic aircraft, this is irrelevant. The first question of interest is whether the imaging interval is sufficiently long. More specifically, is the fuselage length at least 20 resolution cells, and are wingtip and delayed returns resolved from fuselage responses over at least 90% of the fuselage? The fuselage extends from the exhaust responses near Crossrange Gate -40 and Range Gate -15 to the nose near Crossrange Gate 2 and Range Gate 1 . This is long enough that the imaging interval could be about halved. The near-wingtip responses are in Crossrange Gate -5 and Range Gate -20 , entirely resolved from the fuselage. The delayed duct returns are harder to recognize, because the compensation is so poor that they are spread over many crossrange gates. However, with a good compensation, they will be confined to one or two gates, so they will obscure an acceptably small fraction of the fuselage's 40 crossrange gates.




Figure 3.47 One-second image with the standard motion compensation.

The remaining question of interest is, can we find range gates with responses that permit a motion analysis? Despite the poor appearance of the image, it is possible to find suitable responses. The image cut in Range Gate 0, which is the gate of the compensated scatterer, and its transform are shown in Figure 3.48. The amplitude function does not have deep minima, so that the phase function is meaningful. Noting the scale of the phase function and keeping in mind that (slow) phase deviations of up to about 0.1 cycles from linear are acceptable, we conclude that one usable imaging interval extends from -0.5 to 0.0 seconds. Since these measurements have been performed on the response of the compensated scatterer, it is particularly important to go through the same analysis in another range gate. The image cut and transforms in Range Gate -11.5 are shown in Figure 3.49. Again noting the much cruder phase scale, we conclude that the imaging interval from -0.5 to 0.0 seconds, extracted from Figure 3.48, is indeed acceptable.

On the basis of these measurements we conclude that, under ordinary circumstances, using the interval from 36.0 seconds to 36.5 seconds should give a good-quality image. The actual image is shown in Figure 3.50. The dotted generic aircraft outline shows the approximate location of the aircraft fuselage. However, the wingtip responses are offset in crossrange from the wing outline; this is an indication that the aircraft is rolling. The image




Figure 3.48 Image cut and transform for Range Gate 0.




Figure 3.49 Image cut and transform for Range Gate -11.5.




Figure 3.50 Image from 36.0 to 36.5 seconds.

quality evidently is poor. The reason is that the onset of the maneuver is so abrupt in terms of radar imaging that many strong spurious responses are generated. In Figure 3.50, genuine responses (responses positioned where expected, based on ground truth) are sharply focused, because of the proper selection of the imaging interval. On the other hand, responses positioned outside the aircraft have the characteristics of spurious responses. For example, in Figure 3.51 we show the image cut and its transform for the range gate of the tip response. The relative half-power width is near unity, with some interference noticeable from nearby spurious responses. In contrast, in Figure 3.52 we show the image cut and transform in the range gate of the strong response in Range Gate -5 and Crossrange Gate 2. The relative half-power width is 1.9, and yet the transform has no resemblance to the pattern of interfering scatterers. If we examine other responses of the image, we obtain similar results; some responses are genuine and others are spurious.

In principle, we could analyze every response of the image and classify each as either genuine or spurious. This is difficult to do in practice because many responses are poorly resolved, and spurious responses may fall on top of genuine responses. However, we want to make a general point concerning spurious responses caused by a maneuvering aircraft. The aircraft data that we are using for our demonstrations were mostly collected by a radar that




Figure 3.51 Image cut and transform for Range Gate 0.




Figure 3.52 Image cut and transform in Range Gate -5.

generates a long linear FM signal by slow frequency stepping. Linear FM signals measure extrapolated ranges [1], and a long sweep implies a large degree of extrapolation. Such a radar probably enhances the effects of the phase-center wander, and may generate much stronger spurious responses than a radar using a short linear FM signal. We did not investigate this point because of lack of adequate data. In a situation in which the spurious responses represent as bad a problem as in our illustration, only a small number of genuine response positions may be readily extractable, preventing identification. Rather than trying to extract genuine responses masked by spurious ones, it appears simpler to delay identification until the aircraft exhibits a smoother flight behavior. If we utilize the imaging procedures as developed here, the resulting image will be usable when the flight is reasonably smooth, but will contain many spurious responses when the motion of the aircraft is erratic.

The quality of the peaks tracks of Figure 3.44 is so poor that one might wonder whether starting with a different peaks track would give substantially different results. Indeed, when one starts with a different track than the one chosen above, an image with somewhat different appearance is obtained. However, finding range gates with analyzable scatterers and examining the transforms of image cuts leads to the selection of the same subintervals for

imaging and to the same results as with the earlier peaks track. Also, one can reduce the consequences of the phase-center wander of scatterers by choosing a shorter imaging time, but that also reduces crossrange resolution. This is a problem when the aspect angle is close to broadside, as it is in our example. Our example thus combines the difficulties introduced by maneuvering with those from a near-broadside aspect angle.

The fuselage extends over about 40 crossrange gates in Figure 3.47, formed over one second of data, and extends over about 15 crossrange gates in Figure 3.50, formed over the first half of that interval. Therefore, the fuselage must extend over about 25 crossrange gates in an image formed over the second half of the interval. Accordingly, one might hope to obtain a usable image from the second half-second. Unfortunately, the aircraft motion during this half-second is too erratic to allow the formation of a useful image.

In examining the next second of the flight, starting at 37.0 seconds, we find the usable imaging interval in the same way as above. The image from 37.1 to 37.4 seconds is shown in Figure 3.53. It still contains many spurious responses, but the aircraft can be better recognized than in Figure 3.50. The image from 38.7 to 39.0 seconds is shown in Figure 3.54. Most of the major




Figure 3.53 Image from 37.1 to 37.4 seconds.




Figure 3.54 Image from 38.7 to 39.0 seconds.

spurious responses have disappeared, yet there are still lower-level spurious responses. We do not observe a single line of responses along the fuselage, as we do in an image of high quality. Nevertheless, the irregularities in the aircraft motion are clearly diminishing.

3.6.5 Maneuvering Combined With Vibrations

Some aircraft appear to vibrate during strong maneuvers to such a degree that it affects imaging. Since we have insufficient ground truth, we can only speculate that this might happen when a maneuver is combined with an increase in the speed of the aircraft. Whatever the cause may be, radar imaging has to deal with high-frequency motions of the scatterers. The same principles and algorithms are used to derive an image, but a higher degree of adaptivity is needed.

In Figure 3.55 we show a two-second survey image of an aircraft in a turn, after range and Doppler tracking of a scatterer. In this figure, the generic outline gives only an indication of the orientation of the aircraft. As usual, we must find responses in this image for which transforms of




Figure 3.55 Survey image of a maneuvering aircraft.

fixed-range image cuts have sufficiently constant amplitude functions to allow selecting an imaging interval on the basis of the phase function. We try to find a time interval within which all the scatterers that are tested have a linear phase function, or constant Doppler. Ordinarily, this is not difficult to do. In this instance, many intervals for different responses do not adequately overlap. Moreover, there are rapid changes in the phase that are not accompanied by corresponding amplitude variations, which means abrupt motions that we interpret as vibrations.

The image cut in the range gate of the compensated scatterer and its transform are shown in Figure 3.56. Taking into account the scale of the phase function for the transform, we recognize that one usable imaging interval extends from -0.5 to -0.3 seconds, and another one from -0.3 to -0.05 seconds. From the phase function of the transform one could at most define another interval, from 0.0 to 0.15 seconds, but this interval is clearly inferior. In Figure 3.57 we show the image cut and transform for the response in Range Gate 10.7. Disregarding the phase jumps in the early part of the transform, since these are related to the amplitude minima, we recognize a linear phase function. Thus, the first imaging interval from -0.5 to -0.3 seconds




Figure 3.56 Image cut and transform for the compensated scatterer.




Figure 3.57 Image cut and transform for another response.

is also acceptable for this response. However, the second interval (from -0.3 to -0.05 seconds) includes a strong curvature of the phase function in Figure 3.57. Hence, on the basis of these two scatterers, we must prefer the first imaging interval.

The image from 30.0 to 30.25 seconds, which corresponds to the normalized interval from -0.5 to -0.375 seconds, is shown in Figure 3.58, with the generic outline again indicating only orientation. If we, for the moment, consider only the strong responses, indicated by the large dots, it appears to be a good image. This is verified by examining the individual responses of the image, which are sharply compressed. However, the image has a relatively high level of spurious responses. To determine the source of these spurious responses, we again must select responses, take fixed-range image cuts and transforms, and examine the phase functions corresponding to nearly constant amplitudes. This is more difficult to do than when the goal is to select good imaging intervals, because we cannot tolerate significant amplitude modulation.

For a better illustration of the problem, which occurs over the entire data segment available for this aircraft, in Figure 3.59 we show the image cut




Figure 3.58 Image from 30.0 to 30.25 seconds.




Figure 3.59 Image cut and transform in the range gate of a strong response.

and transform for a strong scatterer about one second later. Note that a drop of about 0.2 cycles occurs in the transform phase function at -0.45 seconds, where the amplitude is steady. A similar rise of 0.1 cycles is found at -0.1 seconds, and another drop at 0.24 seconds. Because of the absence of a corresponding change in the amplitude function, these phase variations imply motions of the scatterer. In fact, the phase function of the transform appears to have a partly masked modulation period of about 0.1 seconds (0.2 seconds of relative time), corresponding to a modulation frequency of 10 Hz. Such phase variations can also be observed on some scatterers in different range gates, but they do not occur at the same times. This excludes the possibility that the phase variations may have been introduced by the radar. For a sinusoidal phase variation of 0.1 cycles, the sidelobe power level is at about the level found in Figure 3.58.

The spurious responses in Figure 3.58 will not prevent aircraft identification. The greater difficulty comes from the fact that the good imaging intervals for different scatterers on the aircraft do not always overlap to a sufficient degree, as discussed above. For a demonstration of the consequences, consider the phase function of the transform in Figure 3.57. We now form an image over the interval from -0.1 to 0.03 seconds of normalized time.

In accordance with the phase curve at the right bottom of Figure 3.57, this appears to be a good imaging interval. However, the same interval in Figure 3.56 indicates problems. The image is shown in Figure 3.60, with the generic outline indicating just orientation. The fact that this image has a very poor quality can be seen even from the peaks plot by considering the strong responses. Without further analysis, in Figure 3.61 we also show the image formed over an interval between those of Figures 3.58 and 3.60, from 30.4 to 30.65 seconds (relative time -0.3 to -0.175 seconds), again with the generic outline indicating just orientation. The image quality is even poorer than that of Figure 3.60. The conclusion is that the *imaging interval must be very carefully selected when the aircraft is vibrating or perhaps also flexing*. We always use the same tools for selecting the imaging interval, but the process is more critical in such a difficult situation. It is evident from these examples that correct image-time selection, even within a short observation time, makes the difference between being able to identify and not being able to identify an aircraft that is not flying smoothly.




Figure 3.60 Image from 30.8 to 30.05 seconds.




Figure 3.61 Image from 30.4 to 30.65 seconds.

3.6.6 Imaging at Large Aspect Angles

The aspect angle of a target is irrelevant for optical imaging because the imaging interval must be chosen so short that the target motion is frozen during imaging. In contrast, radar imaging uses the target motion for crossrange resolution. The motion compensation necessary for imaging can be performed only if it is possible to measure the motions of at least two scatterers, the first to remove translational motion and the second to remove rotation about the first. In both cases we must fit a polynomial or a spline function, and in the presence of interference from other scatterers these fits will not be perfect. The consequences are distortions of the responses and the generation of (relatively high) crossrange sidelobes, making the measurement of scatterer positions difficult or impossible. As the aspect angle of a slender target (the fuselage of an aircraft, since wings are of little use for identification) increases, progressively more scatterers occupy the same range cells. A motion compensation of adequate quality eventually becomes impossible because no two scatterers are sufficiently range-resolved to permit measurement of their motions with sufficient accuracy.

The statement that a range cell should not contain more than two significant scatterers needs more explanation. The practical interpretation is that a usable range gate must contain a wide enough crossrange interval with no more than two significant responses. The minimum width of the crossrange interval is defined by the requirement that it include enough of the high frequencies of the response so that the transform over the window retains any abrupt changes (which are to be measured) in the phase function of the signal. For a given range resolution, the requirement that a sufficiently wide crossrange interval contain only one or two significant responses becomes more difficult to meet when the motion of the aircraft is erratic. Such a motion smears the scatterer responses in crossrange, so that the crossrange intervals with an acceptably low background level for placing the transform window become narrower or disappear altogether. The smearing of the responses may easily lead to an overlap of all the responses in the range gate. Hence, the motion compensation at a given aspect angle becomes more difficult, or impossible, when the aircraft is not moving smoothly.

A given erratic yaw variation introduces only a small phase variation for scatterers on the fuselage when the aspect angle is small, but becomes fully effective at broadside aspects. The coarser the range resolution of the system, the smaller is the aspect angle sector about nose-on or tail-on for which the system can accommodate the combined problem of range-resolving scatterers and dealing with inadvertent phase modulations from erratic target motion. We have repeatedly pointed out the difficulties at large aspect angles, emphasizing the role of high range resolution. In view of the practical importance of these facts of radar imaging, we now consider imaging at large aspect angles in more detail.

The preceding discussion is concerned with the problem of finding two scatterers that permit measuring the aircraft's motion. There is another side of the problem that appears at large aspect angles. Because of the ineffectiveness of range resolution along the fuselage, we require a relatively high crossrange resolution, which implies a long imaging interval. The inadvertent yaw motion of an aircraft can easily be so rapid that the imaging interval needed to achieve the required crossrange resolution contains an entire yaw cycle or even more than one yaw cycle. However, to measure such a motion to high accuracy under real conditions and then fit a spline function that follows the variations accurately can be an unsolvable problem. Too many spurious responses of objectionable strengths are generated, and the genuine responses are distorted so much that the TSA cannot be applied. As was explained earlier, if the aspect angle is small enough so that a poor crossrange resolution is acceptable, we select an imaging subinterval over which the motion happens

to be one of constant Doppler. At large aspect angles this option does not exist because resolution of the scatterers on the fuselage depends on cross-range resolution. Hence, the larger and more erratic the inadvertent yaw motion of an aircraft, the smaller will be the aspect angle at which identification fails.

It is clear that the problem of motion compensation and imaging at large aspect angles can be simplified, and imaging extended to larger aspect angles, if one is willing to use a radar with very high range resolution; specifically, better than 1 ft. If the observable scatterers on an aircraft were fixed point scatterers, the choice of a sufficiently high range resolution would solve much of the problem (compensating erratic motions would still be difficult). However, the important scatterers are extended and some are dispersive, so that an increase in range resolution will cause a variety of effects such as spurious responses and dispersion. An image formed with very high range resolution might appear better to the eye, but whether it would permit accurate measurements of scatterer positions is not certain at all. In the absence of experience with very high range resolution, we can merely state that increasing range resolution will facilitate the motion compensation and identification, but probably only up to a point.

In summary, in a given situation there will be some maximum aspect angle beyond which an adequate motion compensation may not be possible, so that no image can be formed that allows accurate measurement of scatterer positions. This maximum aspect angle depends on the range resolution of the radar, and for a given range resolution it depends on the type of aircraft and its flight behavior. The limitations at large aspect angles can easily be of practical concern. Even though an aircraft viewed at a large aspect angle approaches the radar only slowly, we may still want to identify the aircraft.

In examining the task of aircraft identification at large aspect angles, we first consider the problem of definitely inadequate bandwidth. Specifically, we use data collected on a fighter aircraft by a radar with only 300 MHz bandwidth, or a range resolution of 0.5m.

We first remove the linear range drift and then the linear Doppler drift of the strongest scatterer, starting with linear compensations to ensure that no spurious motion is introduced by curve fitting in this critical situation. The resulting peaks tracks are shown in Figure 3.62. Since the removal of the linear trends in range and Doppler is not a satisfactory motion compensation in general, we must refine the compensation. An examination of the figure shows that the track of the target peaks at the closest range (Gate -9) should be usable (as is often the case), as should the strong track around Range Gate -5. These tracks are labeled by arrows in Figure 3.62. The tracks at closer




Figure 3.62 Peaks tracks after linear range and Doppler compensation.

range are sidelobes of strong target responses, which can be discerned from their amplitudes, not displayed in the figure. When the first peaks track is range- and then Doppler-compensated, we obtain the image of Figure 3.63. Let us again explain how such an image is to be interpreted.

We have compensated the tip scatterer, so that it is stationary, which implies that its response is compressed into a sharp spike in crossrange (with two low-level Doppler sidelobes). However, the other scatterers of the aircraft have residual motions, with the degree of motion increasing with distance from the compensated tip scatterer. Since this motion generally is one of varying Doppler, the responses of these scatterers are typically smeared in crossrange. The smeared responses in each range cell superpose and generate an interference pattern whose peaks are given in Figure 3.63. The higher the degree of response smearing due to a residual motion, the larger the number of peaks that will be formed by smeared responses and interfering smeared responses of scatterers in the same range cell. Thus, the number of image peaks is far larger than the number of observable scatterers, and the position of a peak does not correspond to the position of a scatterer, except in the case of a strongly dominant scatterer with minimal phase-center shifting. As we try to obtain higher crossrange resolution, the residual scatterer motions will be larger and more variable. The resulting spreading of the scatterer responses leads to an increase in the number of image peaks. Thus, the larger




Figure 3.63 Image after motion compensation of the first peaks track.

the number of significant peaks in an image, the poorer the motion compensation. For positive identification of the aircraft in a large database, we must compensate the image such that each scatterer generates at most one (properly compressed) response. Individual responses must not be superposed to such a degree that the TSA fails.

In our example, we must try to refine the motion compensation (attempt to range- and Doppler-track other scatterers, and compensate based on those tracks) until the requirements on image quality are met. This is judged from the patterns of the transforms of fixed-range image cuts through response peaks. If none of these compensations yields an image of acceptable quality, we must reduce the imaging interval by, say, a factor of two, and again evaluate image quality. Note that the crossrange spread of the fuselage in Figure 3.63, and its resolution from the wingtip and delayed duct returns, are sufficiently large that we can substantially reduce the imaging interval. The question to be considered now is whether the range resolution of 0.5m is high enough for the case at hand; that is, for identification of fighter aircraft at even the moderately large aspect angle of our example.

To refine the motion compensation, we start at the tip of the aircraft, where the scatterers are relatively sparse. Figure 3.64 gives the image cut in




Figure 3.64 Image cut in the range gate of the tip response.

the range gate of the tip response, and its transform. Examining the amplitude of the transform, the right upper curve, we conclude that the minima are too low for a high-quality compensation based on the phase function, but that the phase-slope method should work because the modulation is nearly periodic (two dominant scatterers). Next, we examine the performance of combined range/Doppler processing for this large aspect angle. Since such processing is not part of our interactive software, we implement it by first generating a sequence of short-duration images, then using the range and Doppler positions of image peaks to associate peaks in consecutive images in order to form peaks tracks in range. Next we compensate in range, and then in Doppler.

Tracks from the combined range/Doppler processing, with a Doppler window length of 7% of the displayed interval, are shown in Figure 3.65 for the closest ranges. The slope of each segment represents the Doppler of the response at the particular time. We must select those segments of the lowest peaks track whose slopes correspond to the changing range. (Note that this is impossible for the “tracks” at the larger ranges in Figure 3.65.) This leads to the range track of Figure 3.66. The following Doppler track is given in Figure 3.67. The smooth variation of the residual Doppler in Figure 3.67




Figure 3.65 Selected peaks tracks after Doppler processing.




Figure 3.66 Range track of the tip scatterer.

already is an indication that range tracking was done correctly, since otherwise no trackable Doppler history would be obtained. However, the quality of the motion compensation should always be confirmed by examining the transform of a fixed-range image cut in the new image. The image after this




Figure 3.67 Residual Doppler after the compensation in accordance with Figure 3.66.

compensation is shown in Figure 3.68. A fixed-range image cut in the gate of the tip scatterer differs insignificantly from Figure 3.64; the tip is well enough resolved that the two methods of tracking yield equivalent results.

In theory, having compensated the tip scatterer, we should select a second scatterer as far away as possible from the tip, and measure the motion of the second scatterer relative to the tip scatterer. As long as the aircraft is rigid and the direction of the rotation axis does not change, we then can sufficiently compensate the entire aircraft on the basis of these measurements. However, in this instance no response can be found in the entire image that is good enough for the motion measurement. Wherever in the rear of the aircraft we take a fixed-range image cut through a response peak, there are so many interfering smeared responses that no motion measurement is feasible. As an alternative, we might try to start at the compensated tip and gradually move the compensation along the aircraft until the tail is reached. As the motion compensation is shifted along the aircraft, we might be able to perform scatterer position measurements and compensations in a localized region around each compensated response.

Starting from the tip response, the process indeed works until we reach the strongest response, which is the intake response. The image when the intake response is compensated is given in Figure 3.69. As should be expected, the tip response is somewhat smeared in this image because the




Figure 3.68 Image after further compensation of the tip scatterer.




Figure 3.69 Image with compensated intake response.

motion differential between tip and intake is not one of constant Doppler. By comparison with Figure 3.68, we note the lower density of peaks in the region of the compensated intake response, as would be expected. We now should proceed in small steps toward larger ranges, find responses that are either dominant in their range gates or are competing at most with another strong response, and again motion compensate. It turns out that this is impossible. In whatever range gate we take an image cut, the interference situation is one that does not allow any motion compensation. It also turns out to be impossible to unscramble the responses by combined range/Doppler processing, because there are too many scatterers per range gate and the motion is not sufficiently smooth when the phase is the yardstick. The same situation occurs if we examine the image of Figure 3.63, or that image after phase-slope compensation in accordance with Figure 3.64. It is conceivable that by resampling (and possibly polar reformatting) in accordance with the measurable motions of tip and intake, one could reduce the residual motion near the tail enough so that one could track a tail scatterer. However, the small separation of the tip from the intake, and their nonpointlike backscattering, limit the achievable improvement in the tail area. In such complicated situations, the typical theoretical approach does not work on real data.

As we have discussed repeatedly, when such a situation exists, the only remedy (for a given range resolution) is to reduce the imaging interval and, with it, nominal crossrange resolution. For a given residual motion the responses then will be less smeared, so that a motion compensation is feasible. In the present instance, the fuselage extends over 16 range cells. If the significant scatterers on the aircraft were uniformly spaced along the fuselage, a subdivision of each range cell into somewhat more than one crossrange cell should be adequate to accommodate the 20 or so significant scatterers on an aircraft fuselage. This means that a crossrange resolution which spreads the fuselage over about 25 cells should be adequate. However, the important scatterers often occur in groups rather than being uniformly distributed, so that a somewhat higher crossrange resolution is needed. The requirement differs from one aircraft to another and can only be roughly estimated. The approach must be to reduce crossrange resolution in steps and try measuring the scatterer positions, trading nominal crossrange resolution against the problems of an inadequate motion compensation when the imaging interval is longer. In searching for a best compromise, identification is possible if we can find a value of crossrange resolution for which scatterer positions are measurable (properly compressed responses or responses analyzable by the TSA). If not, then range resolution is inadequate.

In our example, the aircraft fuselage spreads over nearly 90 crossrange gates, which is a higher crossrange resolution than needed. Thus one can try to reduce crossrange resolution by more than a factor of two. This raises the question of whether one should return to the data of Figure 3.62, compensated based on just linear tracks, or proceed forward from Figure 3.68 (or 3.64 or 3.69), compensated based on higher order (but still low order) fits. Use of linear corrections alone is appropriate for situations with severe enough motion that scatterers are untrackable or that their tracks have sharp kinks. If scatterer responses can be properly compressed with low-order fits (as is the case here), we make the task harder by not compensating based on actual tracks.

In this case, there is only one imaging interval, from 0.2 to 0.7 seconds, that allows measuring scatterer positions in the resulting image. In this image, the fuselage extends over 16 range gates and 21 crossrange gates, for a total extent of 26 gates, barely acceptable. On the other hand, the same aircraft did not pose any problems at smaller aspect angles, with the same 300-MHz bandwidth; images of acceptable quality could generally be formed from an arbitrary one-second or half-second dwell. Thus, it appears that a bandwidth of 300 MHz typically will be insufficient for measuring scatterer positions on a fighter aircraft at this aspect angle, which is about 60° . The general shape of the image or some special features might allow distinguishing among a few aircraft, but will not allow aircraft identification when the database is reasonably large.

We now return to data with a more reasonable range resolution of 1 ft, and consider imaging at large aspect angles with sufficient quality for the measurement of scatterer positions. Figure 3.70 shows the peaks tracks over a rather long interval of five seconds for a fighter aircraft, after removal of the constant range drift. The aspect angle is increasing from an initial value of 63° . Although the strongest peaks tracks (those in the first 20 range gates at any instant) are ill defined, there is a usable peaks track starting in Range Gate 44, indicated by the arrow. Thus, at least one response can be tracked in the intensity range profiles over the full five seconds. Although this is likely a delayed duct response, this does not matter for such an initial processing step. With a motion compensation based on range and Doppler tracking of the response, we obtain the image of Figure 3.71. As the large number of peaks of significant strength indicates, fixed-range image cuts show that the image quality is too low for the measurement of scatterer positions.

We now must try to find two fixed-range image cuts containing responses whose transforms allow analyzing the scatterers' motion. If we are successful, we can select the best (with the smoothest motion) imaging




Figure 3.70 Peaks tracks over five seconds for a fighter aircraft.




Figure 3.71 Five-second image based on range and Doppler tracking of a response.

interval long enough to provide the desired crossrange resolution. If the relative Doppler of the two scatterers varies over this best interval, we must try to

motion compensate. Success or failure of the motion compensation is evaluated by analyzing the responses of the resulting image. If the motion compensation fails to produce an image of acceptable quality, the only remedy is to select a still shorter imaging interval over which the motion is smooth. If this leads to an image of inadequate crossrange resolution, aircraft identification is not possible with a range resolution of 1 ft at this large aspect angle and for the existing motion.

In going through these steps, we try fixed-range cuts in the image of Figure 3.71, finding that the one with the lowest interference is in Range Gate -22. This image cut and its transform are shown in Figure 3.72. There is an amplitude dip and corresponding phase variation at a time of about 91 seconds. If this sharp phase variation is ignored, the phase represents the motion of the scatterer reasonably well. The polynomials used in the preceding range and Doppler compensations were far too inflexible to follow such a rapid phase variation as in the transform of Figure 3.72, so that the breaks in the slope of the phase function actually represent the changes in the range rate of the scatterer. Thus we conclude that the aircraft is undergoing a rather erratic motion, even though it may be typical.




Figure 3.72 Image cut in Range Gate -22 of Figure 3.71, and transform.

The fuselage in Figure 3.71 extends over about 130 crossrange gates, far more than necessary for measuring scatterer positions in a well-compensated image. Retaining some margin for further reducing the imaging duration, we should examine an interval of about two seconds duration, which should give a fuselage extent of about 50 crossrange gates. The interval from 92 to 94 seconds is the best available, containing slowly varying amplitude and only a single break in the phase slope. With the earlier motion compensation (based on the delayed duct return) retained, the corresponding image is shown in Figure 3.73. By comparison with Figure 3.71, the number of major responses (large dots) has decreased significantly. The aircraft fuselage indeed extends over about 50 crossrange gates, which is more than adequate. Our next step is to improve the motion compensation to where the image responses can be analyzed. Failing this, we must further reduce the imaging interval.

With the image of Figure 3.73, we can improve the compensation sufficiently. We start by measuring the range/crossrange position of a well-compensated response. (Generally, this is the reference scatterer used for compensation, but when a delayed response has been compensated, one starts with the corresponding direct response—in this case, the intake.) Then




Figure 3.73 Image from 92 to 94 seconds.

we take a fixed-range image cut through a major response in the vicinity, applying the TSA to determine whether it is a response from one or two scatterers. After finding a response from a single scatterer, we compensate in accordance with the phase function of the transform. In this manner we proceed over the entire aircraft, compensating and measuring the range/crossrange positions for the scatterers associated with the responses (incorporating the different compensations into the position measurements, to produce positions in a common coordinate system). We ignore the minor responses in the process, since they are most likely spurious responses from the imperfect motion compensation. With this type of processing, we can indeed make use of the image of Figure 3.73, which has adequate crossrange resolution. The aspect angle at the center of the imaging interval is 67° . We note that an alternative is sometimes available to a succession of compensations. If we can track two scatterers separated by a good fraction of the fuselage, and if the aircraft motion is two-dimensional, we can resample and polar reformat the data, so that all position measurements can be made in a single image.

Now we try the same imaging somewhat later in the flight, where the initial aspect angle is 77° , or only 13° off broadside. With the same processing steps and an imaging interval of two seconds, we obtain the image of Figure 3.74. In this image, it is impossible to shift the motion compensation by taking the transforms of fixed-range image cuts and compensating with




Figure 3.74 Image from 100 to 102 seconds.

their phase functions. For example, the image cut over the line in Figure 3.74 has the appearance of Figure 3.75. There is no way in which a transform window can be placed so that the transform gives a meaningful phase function. With the selection (based on transforms of other fixed-range cuts) of a suitable shorter imaging interval of 1.3 seconds, we obtain the image of Figure 3.76. The two strongest responses in this image (largest dots) both are well compensated, and so are others. However, when we take an image cut of the kind shown in Figure 3.74, it is similar to that shown in Figure 3.75. Crossrange resolution is inadequate to measure scatterer positions. Practically, it makes no difference whether accurate measurements are impossible because the responses are distorted, as in the case of the longer imaging interval, or whether undistorted responses are too closely spaced, as in the case of the shorter imaging interval. The conclusion is that we cannot form an image of adequate quality at the aspect angle of almost 80° under the existing conditions.

We have used the designation “large aspect angles,” and demonstrated that range resolution determines at how large an aspect angle one can identify a small aircraft. In our example of an image with still adequate quality, the aspect angle was 67° , but near 80° the problems proved unmanageable. How about imaging at broadside? For the type of aircraft we have used in our




Figure 3.75 Image cut and transform.




Figure 3.76 Image from 100 to 101.3 seconds.

examples, and with a range resolution of 1 ft, near broadside it has proved totally impossible to generate an image with the quality needed for identification. One might compensate some dominant scatterer rather well, but other scatterers do not stay within single range gates. Yet these scatterers cannot be properly compensated because there are too many other scatterers in the same range gates. It is clear that better range resolution than 1 ft should be available for identification of a fighter aircraft near broadside.

The preceding conclusions with respect to aircraft identification near broadside are rather negative. However, they were made on the basis of the specific processing procedures we have illustrated, which attempt to use the minimum observation time. It is dangerous to conclude generally that some problem is unsolvable. To demonstrate this fact, we now consider identification of the same aircraft almost directly at broadside, using the 500-MHz data. However, we employ a relatively long observation time. We have not processed enough data to know whether the method will give the desired results in general, because performance depends heavily on the flight behavior of the aircraft.

Figure 3.77 shows the sequence of range profiles as the aircraft approaches broadside, without any motion compensation. Broadside for the flight path, rather than the fuselage, is reached when the range no longer changes, which is at about 105.4 seconds. Since the aircraft executes inadvertent yaw maneuvers, and also because the features that generate broadside flashes are not all oriented parallel to the centerline of the aircraft, these flashes start earlier and end later than the time at which the flight path is viewed at broadside. We do not want to process in the region of these flashes. We have explained the difficulties with Doppler resolution at broadside, independent of flashes. A short flash generates a Doppler spectrum with a width equal to the reciprocal of the flash duration, which further degrades Doppler resolution of the responses. On the basis of Figure 3.77, we do not want to try to identify beyond the time of about 104.2 seconds, where the first flash starts.




Figure 3.77 Range profiles near broadside.

Let us assume that the aircraft can be observed from 102 to 104.2 seconds, a relatively long observation time of 2.2 seconds. If there were no identification problem at broadside, one could accumulate a short section of data, motion compensate, form an image, and measure the crossrange spread of the aircraft. This would be done in steps until the aircraft's spread in crossrange gates corresponded to the desired crossrange resolution. On a small fighter aircraft we might want at least about 30 to 40 crossrange cells. When the observation time became long enough to give this crossrange resolution, we would stop collecting data and identify the aircraft. Because of the problems at broadside, the process must be extended. When enough data have been collected to give the desired crossrange resolution, we try to generate an image of sufficient quality (one whose responses can be analyzed with the TSA). If the aircraft has (even small) inadvertent yaw motions, it is unlikely that we will succeed. Then we shift the starting time of the imaging interval by a fraction of the imaging interval, collect more data to offset the delayed start, and again examine whether the image quality is adequate. In this manner, we shift the imaging interval in time until we succeed. This is why the process is inefficient in observation time, which may have to be significantly longer than the imaging interval.

Figure 3.78 gives the peaks track over the interval from 102 seconds to 104.2 seconds (where the flashes begin), with the linear range drift of




Figure 3.78 Peaks tracks from 102.0 to 104.2 seconds.

Figure 3.77 removed. The lowest arrow indicates a good peaks track starting in Range Gate -1 , but even a crude image shows that the scatterer is on a wing. The highest arrow indicates another reasonably good peaks track starting in Range Gate 13 , but the scatterer also is on a wing. The peaks tracks from the fuselage scatterers are between these two tracks, and they are expectedly poor. Nevertheless, by fitting a quadratic polynomial we still can track the peaks track starting around Range Gate 8 , indicated by the middle arrow. On the other hand, in an operational system we do not want to collect data over such a long observation time, because we might obtain a good image early in our search process. However, if only the first part of the peaks track around Range Gate 8 were available, it could not be tracked with sufficient accuracy. Thus we must rely on tracking a scatterer on the wing, even though the wing motion may be different than the motion of the fuselage.

When the peaks track starting in about Range Gate -1 in Figure 3.78 is compensated, the residual motion is shown by the Doppler track of Figure 3.79. This track varies too erratically to expect a good image after the Doppler compensation. The actual image is shown in Figure 3.80. Taking a transform of the fixed-range image cut through the tip response, we find that the tip scatterer drifts so much in range that its response is observable only over a small part of the 2.2 -second interval. The responses of this image are so poorly resolved that none can be tracked. Thus, there is no way in which




Figure 3.79 Doppler track of the wing scatterer.




Figure 3.80 Image after range and Doppler compensation of the wing scatterer.

the image can be further compensated to give a final image in which scatterer positions are measurable. However, with a crossrange spread of the image of roughly 100 gates, we can reduce the imaging interval by about a factor of three, from 2.2 seconds to about 0.7 seconds.

Since the motion compensation was performed on a wing scatterer but we want to compensate the fuselage, the Doppler variation of the track in Figure 3.79 may not be significant. However, the times at which the Doppler track abruptly changes slope may be correlated with the times at which the fuselage has motion discontinuities. Indeed, one finds that the best image quality is obtained when the imaging interval is confined to the intervals where the Doppler track in Figure 3.79 is nearly linear. If one shifts the imaging interval in steps, as would be done in an automated system, one finds that the only usable image can be generated over the time interval from 103.5 to 104.2 seconds. By selecting this time interval and compensating the Doppler track of Figure 3.79 over that interval, we obtain the image of Figure 3.81.

This image is analyzable in the sense that the transforms of the responses give satisfactory approximations of the TSA patterns. Since the




Figure 3.81 Image from 103.5 to 104.2 seconds.

motion compensation was performed in the wing area, the quality of the TSA patterns decreases as one approaches the tip of the aircraft. The quality can be improved by phase-compensating near the tip, if possible. In the case of the image of Figure 3.81, such a phase compensation can be performed on the third response from the right. It is also interesting that the quality of the image, in particular near the tip, degrades noticeably if the end of the imaging interval is shifted by just 0.1 second to 104.3 seconds, where in accordance with Figure 3.77 a flash already starts building up (probably near the nose, from the shape of a fuselage). The image quality also degrades if the imaging interval starts earlier by 0.1 second, so that more of the bend of the Doppler track in Figure 3.79 is included (the section from 103.4 to 103.5 seconds).

3.6.7 The Special Case of Zero Aspect Angle

We have pointed out that aircraft identification is relatively easy at the smaller aspect angles, because then range resolution is very effective in subdividing the fuselage into cells. When the flight of the aircraft is steady

(without inadvertent yaw motions) and the aspect angle is small, we encounter crossrange resolution problems. Resolving delayed duct returns from the fuselage returns becomes difficult, but it can be done unless the aspect angle is very small. In that case, the delayed duct returns become quite weak and there is no need to resolve them. On the other hand, if an aircraft is observed at or very close to zero aspect angle, another effect appears. The various fixed engine stages then can give very strong returns, and when the aircraft is flying without significant yaw motions, crossrange resolution is unavailable for resolving the engine stage returns from the fuselage returns. One might think that this means that identification must be based solely on a range profile containing strong engine stage returns that cannot be discriminated from fuselage returns.

It is worthwhile to point out again that, more precisely, this lack of crossrange resolution occurs when the aircraft to be identified is on a collision course with the radar platform. Thus, this particular problem can be avoided if the radar platform moves off a collision course for identification purposes. The radar then effectively generates a SAR antenna. Such a circumvention of the zero-aspect-angle problem is possible when the radar is on an aircraft, but not when it is on the ground or on a ship. Even though a ship may move, its motion is too slow to generate a SAR antenna of useful length. Thus, when an aircraft is headed directly toward the radar that is to identify the aircraft, a special problem exists. Evidently, this case is of considerable practical interest.

Although, for constant aspect angle and a smooth flight, no crossrange resolution is available to discriminate the engine stage returns from the fuselage returns, crossrange resolution does allow one to identify the high-Doppler returns from the rotating blades of the engines. Even when the radar uses a waveform that suppresses the *jet engine modulation* (JEM) returns, there are sufficiently strong residuals to identify the range gates that contain the JEM returns. We can then discard the engine-stage responses in these gates of the range profile (or use them for identification if the positions of the engine stages are known), and utilize only the remaining responses for aircraft identification. Although this amounts to identification via range profiles, crossrange resolution is essential for determining which of the responses of the range profile come from the fixed engine stages. (Note, however, that the inadvertent yaw motions of an aircraft can easily be large enough to provide crossrange resolution even at zero aspect angle.)

The delayed returns from the rotating blades can be used for a different purpose. As the aspect angle increases from zero, at some point the direct returns from the fixed engine stages disappear and strong delayed duct

returns start spreading in range. When the aspect angle is large enough, these delayed duct returns can be resolved in crossrange from the fuselage returns. However, there is a transition region where crossrange resolution is not yet feasible and the delayed duct returns already extend beyond the response from the tail of the aircraft. This makes it difficult to decide where the aircraft ends, information that is not essential for aircraft identification, but is helpful. By noting the range distributions of the strength and number of delayed returns from the rotating blades, we can more reliably determine which response in the range profile is from the end of the aircraft. This topic is further treated in the following section, for tail-on aspects of an aircraft, where delayed duct returns are more difficult to deal with than for nose-on aspects.

As a side issue, an image at nose-on aspect is so easy to understand that we use one to demonstrate spurious responses due to phase-center wander of a scatterer. Figure 3.82 shows an image of a fighter aircraft at nearly exactly zero aspect angle, formed over an interval of two seconds. The image was generated by range and Doppler tracking of a scatterer. In addition, the response at the closest range was phase compensated. This produced a perfect




Figure 3.82 Nose-on image of a fighter aircraft, two-second imaging interval.

compression of not only the first response in the image, but also of the responses in Range Gates -26 and -20 . On the other hand, in the range gate marked by the line we find four response peaks. We know that the aircraft fuselage is not that wide at this point. The line is in the range gate of the radar, which is a scatterer with an unstable phase center. The corresponding image cut is shown in Figure 3.83. Note that the transform of the image cut has an amplitude function that, aside from the modulation, is strong over the entire interval. The spreading of the response thus is not due to a specular flash. Such a case is illustrated in Figure 3.84, which gives the image cut and transform through the group of responses in Range Gate 17. In this case the signal is strong only over the last third of a second, so that the image response must be correspondingly wide.

We now generate an image under the same conditions, but using only the central one second of the two seconds of Figure 3.82. This image is shown in Figure 3.85. Again, the responses in front and behind the radar are sharply focused. The image cut over the indicated interval is shown in Figure 3.86. By comparison with the top left part of Figure 3.83, the response now is much simpler and spread over many fewer crossrange gates. Note that the phase function of the transform in Figure 3.86 has a kink, even




Figure 3.83 Image cut and transform over the interval marked in Figure 3.82.




Figure 3.84 Image cut and transform for the end of the horizontal stabilizer.




Figure 3.85 Image with halved imaging interval.




Figure 3.86 Image cut and transform over the interval indicated in Figure 3.85.

though the phase functions for the responses in front of and behind the radar response were verified to be straight (aside from the inevitable minor modulations).

We generated another image, with a further halving of the imaging interval to 0.5 seconds. The new image is shown in Figure 3.87. The image cut and transform for the indicated interval are shown in Figure 3.88. As implied by the relative half-power width of 0.920, the response is almost perfectly compressed. The phase function of the transform has variations, but at about 0.1 cycles they are too small to smear the response significantly. Note, however, that the marked response in Figure 3.87 is shifted to the left from the center of the fuselage. Thus, although the nonlinear component of the phase-shift wander is too small to smear the response, a significant linear component is effective even over the short imaging interval of 0.5 seconds. This example clearly demonstrates that the achievable crossrange resolution is limited due to phase-center wander effects, even if one could compensate the motion over longer intervals without generating objectionable spurious responses.

As an interesting point, the intensity image of Figure 3.82 (at almost exactly zero aspect angle) is unique in that it approximates an optical type




Figure 3.87 Image with further halving of the imaging interval.




Figure 3.88 Image cut and transform for the interval in Figure 3.87.

image. It would allow shape recognition in the cases when aircraft do not have very similar shapes. One reason for the appearance of the image is the absence of any significant shadowing. Also, the specular flashes from the lagging straight edges of the horizontal stabilizer generate these edges in the image, and to a lesser degree we also see the lagging edges of the wings, which are not exactly at right angles to the centerline of the fuselage. Another contributing factor is that the dynamic range of the responses is not as large as at other aspect angles. For example, if some air inlet on the fuselage is illuminated at a glancing angle, its cross section is not so much larger than those of simpler discontinuities that their responses are masked. This is to say that in this special case the image of the small discontinuities is not completely masked by the responses from the cavity-type scatterers. However, it is a singular situation.

3.6.8 The Special Case of a 180° Aspect Angle

When an aircraft is viewed exactly tail-on, the situation is similar to when it is viewed exactly nose-on, but there is an important difference. In both cases the delayed duct returns become quite weak, so that they pose a lesser problem than at aspect angles not so close to nose-on or tail-on. However, a difference arises from the streamlined shape of an aircraft, in particular a fighter type aircraft. The fuselage of such an aircraft widens with distance from the tip until about the end of the cockpit. For nose-on views there is little shadowing of scatterers in the rear of the aircraft. In contrast, when the aircraft is viewed tail-on, scatterers toward the front of the aircraft will be shadowed, at least optically. We now show that this does not pose a major problem.

Figure 3.89 shows the image of a fighter aircraft viewed nearly tail-on. The aircraft is flying in a loop with a diameter of about 6 km. The image was generated by range and Doppler tracking one of several trackable scatterers. The delayed duct returns are quite weak, but so are the returns from the shadowed part of the aircraft, if they are present at all. At least by eye, we cannot distinguish the delayed duct returns from any weak scatterer returns that might be present. For a better illustration, Figure 3.90 shows the amplitude part of an image cut along the centerline of the fuselage. Starting from negative range gates, the responses are strong up to about Range Gate 0, after which they drop sharply. The question is, are the weak responses in the positive range gates true scatterer responses or weak delayed duct returns?

A comparison of the image with the actual aircraft shows that the amplitudes of the responses drop at about the end of the cockpit, beyond which the aircraft narrows so much that the fuselage is totally shadowed




Figure 3.89 Image of an aircraft viewed nearly tail-on.




Figure 3.90 Image cut along the centerline of the fuselage.

optically for a true tail-on view. However, an analysis of the individual responses and comparison with the database used for angles farther off tail-on shows that, over the entire front part of the aircraft, the same scatterers are observable in Figure 3.89 as at the other aspects. The only explanation is that, because of the smooth design of an aircraft fuselage, the radar wave travels along the fuselage and is reflected by the discontinuities. For example, we observe responses from the back of the pilot's seat, antennas under the cockpit, the radar, and the base of the metallic radome tip. Hence, there is no need to attempt identification without use of the important scatterers on the front part of the aircraft. As discussed subsequently, this is quite important.

Suppose we could not utilize the weak scatterers at the front part of the aircraft. In Figure 3.89, we would then have available only the responses between Range Gate -29 (not on the fuselage) and Range Gate -2 . However, the strong responses bunched about the tail of the aircraft are of little use, because at least some of them are spurious responses caused by the exhaust. Although an analysis shows that some of these responses are genuine, it appears to be too difficult to use these responses selectively for identification. Thus, we must exclude the set of strong responses, so that the number of usable responses decreases even further. The conclusion is that we must make use of the weak responses from the (optically) shadowed front part of the aircraft. The positions of these responses can be utilized regardless of whether or not the equally weak, or even weaker, delayed duct responses are present. Nevertheless, identification performance will improve if it is possible to discriminate the duct responses, effectively eliminating them.

We now show that such a possibility indeed exists. However, since we have not tested the method on a variety of aircraft at tail-on aspects, the discussion is to be taken primarily as an illustration of the type of expert system approach necessary to solve such problems.

One would expect that genuine responses from scatterers would have characteristics different from delayed duct responses. In what way should the characteristics differ? It appears quite hopeless to consider the scattering mechanism in the complicated structure inside an engine duct with the various compressor and turbine stages. Rather, the only practical approach is to examine the responses and try to detect a distinct difference between genuine and spurious responses. To this end, we took the usual 18 image cuts spread over 180° , and examined the structure of the responses in the image cuts. What we found was the following. If one takes an image cut at 120° through one of the delayed responses, as shown in Figure 3.89, one obtains the smeared response in Figure 3.91, instead of a sharp response. One might suspect the effect to be caused by the responses of Figure 3.89 above the one

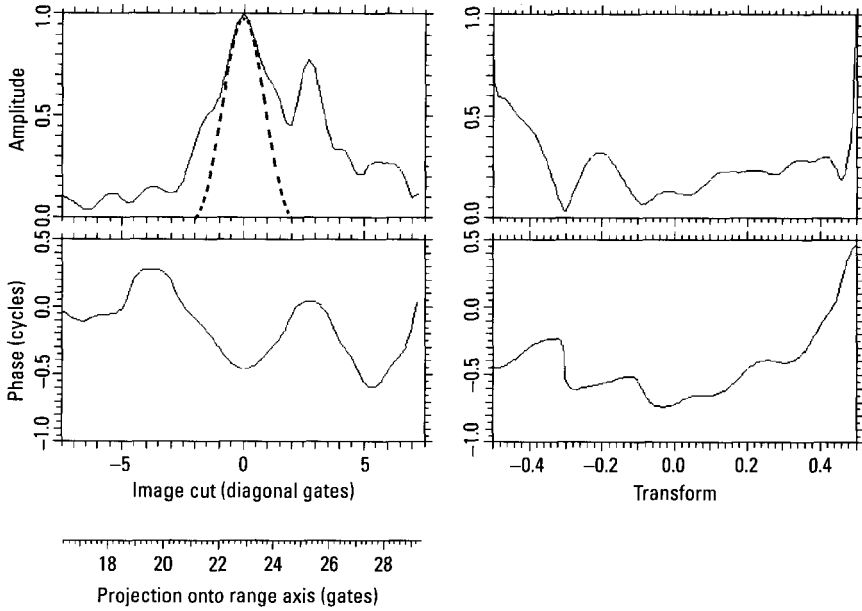


Figure 3.91 120° Image cut and transform through a delayed response.

through which the image cut was taken, but their crossrange separations of 1.5 and 3 gates from the cut are too large. In fact, when the direction of the image cut is changed from 120° to 60° , symmetrical with respect to the vertical axis, the response remains sharp. Further image cuts show that the response is smeared for image cuts within a sector of about 30° centered on 120° . Precisely the same smearing occurs for the two responses above the one marked in Figure 3.89, but not for the responses below the one marked.

The suggested conclusion is that the delayed responses can be discriminated by searching for this smearing effect. If we take the first response from the top for which the smearing does not occur, in Range Gate 19.5, and calculate the range separation from the bottom response, in Range Gate -26, we obtain the length of the aircraft with an error of 1%. Here we assumed, evidently correctly, that for this aspect angle the tip response comes from the base of the metallic radome tip. As implied in an expert system approach, these measurements would have to be repeated on a variety of aircraft, possibly with some adjustments or refinements, before they can be accepted for the discrimination of delayed responses.

3.6.9 Imaging and Identification of Large Commercial Aircraft

This chapter essentially addresses the imaging and identification of small military aircraft, primarily jet fighters, which is the most difficult part of the general problem of aircraft identification. We have shown that it is necessary to generate images of a quality that permits the measurement of at least the positions (and characteristics, if possible) of the dominant scatterers, such as wave-trapping features and antennas. Identification based on shape recognition is generally impossible. If we can measure one or the other feature related to aircraft shape, we will use it, but we cannot depend on shape features for reliable identification.

The situation is reversed for large commercial aircraft. Instead of a limited number of wave-trapping features and antennas distributed over much of the aircraft, whose strong responses mask the weak ones that would define the shapes of fuselage, wings, stabilizers, and so forth, we obtain a multitude of weaker responses that do let us recognize wing edges and the like. On the other hand, such aircraft have few of the wave-trapping features and antennas. Thus we must primarily rely on the features related to aircraft shape, including engine locations. This also implies that we must recognize the general type of aircraft from its image, and adapt our processing and identification procedures to the case on hand. It is a further indication that target identification requires a high degree of adaptivity. Although the same approach to imaging must be used for commercial aircraft as for small military aircraft, every step in the process is much simpler. Thus there is no need to illustrate the motion compensation steps that lead to a good image. Instead, we show examples of images that support our point about the identification of commercial aircraft.

Figure 3.92 shows the image of a commercial aircraft, with only the standard motion compensation. Phase modulations in the transforms of fixed-range image cuts through individual responses indicate slight yaw and roll motions, but this does not affect the peaks plot noticeably. Individual parts of the aircraft could be readily compensated if measurements were to be performed on the responses themselves. Aside from the good definition of the fuselage, one can recognize the wing as well as the horizontal stabilizer. Delayed engine returns (indicated by the vertical arrow) are apparent in Crossrange Gate -2 , with the wing-mounted engine generating a strong response in Crossrange Gate 9 (indicated by the horizontal arrow). A similar image of another commercial aircraft is shown in Figure 3.93. The same type of features can be recognized as in Figure 3.92. As indicated by the delayed engine returns, the aircraft also has wing-mounted engines. Figure 3.94




Figure 3.92 Image of a commercial aircraft.




Figure 3.93 Image of a commercial aircraft.




Figure 3.94 Image of a commercial aircraft.

shows the image of yet another commercial aircraft, with similar characteristics as for the two preceding images. The conspicuous difference is the absence of delayed engine returns in the wing area, but such returns can be seen in the rear of the aircraft, with associated strong responses.

The preceding images were taken at fairly large aspect angles, which are more problematic than the smaller angles. As another example, in Figure 3.95 we show the image of a commercial aircraft at a smaller aspect angle. As in the previous three images, the vertical arrow indicates delayed returns and the horizontal arrow indicates an engine intake. The diagonal arrow indicates Doppler-shifted returns from the engine blades. The shape-related features are most recognizable. In addition, the effects of yaw and roll are so small at the smaller aspect angle that the prominent responses stand out in the image even though again only the standard motion compensation was used. These responses represent the features that can be used for identification in addition to the shape-related features. In generating images in which the prominent responses are so well compressed that positions and characteristics of the associated scatterers can be measured, we use the same procedures as




Figure 3.95 Image of a commercial aircraft at a smaller aspect angle.

discussed for jet fighters. In particular, we select the best subinterval for imaging.

Imaging the same aircraft when it is banked gives equally good results. Such an image is shown in Figure 3.96, with the clip level set very low to enhance the outline. The most distinctive feature for this aircraft is the elevated horizontal stabilizer, but this would not be recognizable if the aircraft were not banked. We note that, in contrast to the previous images, the Doppler-spread engine returns are more useful for locating the engine than the delayed returns are. The motion of the aircraft is so smooth that the image has reasonably well compressed responses. If we investigate the phase functions of the transforms we can select a somewhat better imaging subinterval, reducing the imaging interval to 40% of that of Figure 3.96.

As a last example of how good an image of a commercial aircraft can be, in Figure 3.97 we show the image from yet another aspect, with a beautiful definition of the wing. The Doppler-spread engine returns are not confined to the range gate of the engine blades, but also appear as delayed returns. An examination of the stronger responses shows that those away from the center of the aircraft are not quite sufficiently well defined, so that we should select




Figure 3.96 Image when the aircraft is banked.




Figure 3.97 Image of the aircraft at another aspect.

another subinterval for imaging, as done with the other images. Alternatively, we could track an unfocused response, then use its differential motion relative to the compensated scatterer to resample and polar reformat the data. We note that we have never seen images from fighter aircraft that were so well defined. An identification system solely for commercial aircraft would apparently be much simpler than one that can handle military aircraft as well.

Identification of Large Planes

The class of large planes encompasses, in addition to commercial aircraft, military cargo planes and bombers. Although we did not investigate large military planes, there is little reason to expect any significant difference relative to large commercial planes. Bombers might have a few more of the prominent features, but this will only change the relation between the number of shape-related features and prominent scatterers. The shape-related features are used as discussed in Section 3.4.1, and the prominent scatterers are used as illustrated for fighter aircraft in the earlier parts of Section 3.6.

3.6.10 Section Summary

Since aircraft do not fly smoothly, achieving the desired crossrange resolution is often difficult or impossible, and each step after the crude range and Doppler centroid compensation must be checked before proceeding. Range and Doppler compensations must be done with smooth, slowly varying fits to the tracked motion. The acceptability of range and Doppler scatterer compensations is determined by measuring the constancy of the amplitudes of transforms of windowed fixed-range image cuts through the compensated and other responses.

If a compensation is acceptable for a particular response, the phase function of its transform gives its residual motion. Discontinuities in the phase slope indicate jerky motion, which must be excluded from the imaging interval.

The image should be formed by measuring the motions of two well-separated fuselage responses, then choosing the longest interval that does not contain a phase-slope discontinuity for either response. Figure 3.17 gives an adaptive procedure for determining the correct interval.

Reducing the imaging interval (and nominal crossrange resolution) in order to improve the motion compensation and the resulting image quality is necessary so that the TSA can be used, and the inherent resolution capability realized.

The section illustrates the various motion compensation steps and checks in detail, including special cases. It also demonstrates how much easier it is to image commercial aircraft than jet fighters. For large aircraft, shape-related features become more important for identification.

3.7 Example of Positional Match for a Correct and Incorrect Aircraft

Recall that we distinguish between the special features of an aircraft whose significance is known, such as the length or number of engines, and those features for which we can measure only position and perhaps extent. The inputs for the identification process thus are the results of the positional match and whatever special features can be extracted from the image. In this section, we demonstrate the power of the positional match alone to discriminate between the correct and the incorrect aircraft. We show the results of matching measurements from images of two very similar aircraft to predicted feature positions for the two aircraft. The matches of the measured feature positions to the correct aircraft are much superior to the matches to the incorrect aircraft, indicating highly effective discrimination.

We choose a relatively difficult case for this illustration. Both aircraft are jet fighters and hence so small that high resolution performance should be provided. The aspect angle is 57° off nose-on, which implies that range resolution is not very effective on the fuselage (as stated earlier, identification is most difficult at the broadside aspect). Range resolution is relatively high at 1 ft, but crossrange resolution is only 12.5 ft. Since the radar frequency is X-band, the aspect angle change utilized for crossrange resolution is only 0.24° . In addition to requiring only a short dwell, the motion compensation can be fairly simple. Despite the large aspect angle, the images have strong delayed duct returns from the fixed engine parts, so that the engine intakes can be reliably identified in the image.

To summarize the procedure, the position measurement algorithm estimates errors for the scatterer positions, depending on the interference for each response. We also introduce uncertainties in the predicted scatterer positions, depending on how well one can determine these from diagrams and photographs of the aircraft. A computer program then translates and rotates the set of measured points relative to the set of predicted points, and varies the scale of the crossrange axis until a best fit is obtained. The program calculates a probability that the two sets of positions correspond to the same aircraft, on the basis of the match.

The results of the positional matching procedure are shown in Figures 3.98 and 3.99 for the case when no response is identified as coming from a particular feature. In Figure 3.98 we match the measurements of aircraft A to the correct aircraft, and in Figure 3.99 we match them to the incorrect aircraft, aircraft B. The crosses represent the measured feature positions, and the letters give the predicted locations of the observable features on the aircraft. The calculated Bayesian probabilities differ by a factor of 10^4 . As stated above, since we can detect the delayed duct returns, we can force a match between the measured response from the intake and the intake in the predicted database. However, the match of Figure 3.98 is already so much better than that of Figure 3.99 that the improvement from fixing this point is insignificant.

The results of matching the measured positions of aircraft B (at the same aspect and resolution as the measurements for aircraft A) to the predicted feature positions of the two aircraft are shown in Figures 3.100 and 3.101. The positional match again was performed without utilizing the fact that we can identify the intake from the delayed duct returns. The Bayesian probabilities differ by a factor of over 600. When we utilize the fact that we know which of the responses represents the engine intake, the match of




Figure 3.98 Positional match when no feature is fixed, correct aircraft.




Figure 3.99 Positional match when no feature is fixed, incorrect aircraft.

Figure 3.102 is obtained (the match for the correct aircraft is essentially unchanged). The Bayesian probability has decreased by two orders of magnitude, from 1.59×10^{-3} to 1.85×10^{-5} . The latter factor thus gives the difference in the probabilities between matching to the correct and incorrect aircraft when we utilize the knowledge of the specific response that represents the intake. Even more illuminating than these numbers is the match in Figure 3.102 between the crosses and the letters, which is extremely poor.

3.8 Procedure for Automated Aircraft Identification

In this chapter we discussed the important topics related to aircraft identification, with examples of how the measurements providing the inputs to identification are performed. In this section, we will in summary fashion describe the way in which these procedures can be organized into a system that performs aircraft identification fully automatically.

3.8.1 Step 1: Data Collection

The radar starts tracking an aircraft in order to identify it, providing a sequence of range profiles to the processor for identification. The processor




Figure 3.100 Match between the data of aircraft B and the predicted positions of aircraft B.

also obtains the tracking information that gives the approximate aspect angle of the aircraft, which includes whether the aircraft is approaching or receding. As soon as range profiles have been collected over some minimum time interval, on the order of 0.1 second, the processor forms an image with the use of the standard motion compensation (range tracking of the entire aircraft followed by Doppler tracking of the entire aircraft). This image is analyzed to answer the following questions: Are the wingtip scatterers resolved from the fuselage scatterers in the same range gates? For an approaching aircraft, are the delayed engine or intake duct returns adequately resolved from the fuselage scatterers in the same range gates? (Here the criterion is that we may sacrifice up to 10% of the length of the fuselage, due to masking by delayed duct returns, but no more.) For a receding aircraft, are the multiple delayed engine or exhaust returns adequately resolved from the fuselage scatterers in the same range gates? (Again, we might sacrifice 10% of the fuselage due to masking.) Finally, does the fuselage extend over at least 20 range/crossrange resolution cells?

As long as the answer to any of these questions is no, the radar keeps collecting data on the aircraft. The processor keeps forming and evaluating




Figure 3.101 Match between the data of aircraft B and the predicted positions of aircraft A.




Figure 3.102 Repeat of the match of Figure 3.101, with the intake response identified.

images as the length of the interval over which data are collected increases in steps of 50%. As soon as an image is formed that has sufficient crossrange resolution as defined above, the radar is instructed to keep collecting data until the total interval over which data have been collected is doubled. Data thus are collected over an interval twice as long as the minimum usable interval needed to resolve wingtip and delayed duct returns, and to provide sufficient resolution on the fuselage.

3.8.2 Step 2: Compensation and Selection of the Imaging Interval

The processor forms an image from the range profiles that were collected, using the standard motion compensation. The processor selects a prominent fuselage response near the compensated point in this image, takes a transform of the image cut in the range gate of the response, and analyzes the amplitude and phase function of the transform. If the amplitude function shows an extended drop to the background level, the processor goes back to the range profiles, and range tracks and then Doppler tracks a prominent peak in the range profile. A new image is formed after this track. If the transform amplitude for the (windowed) fixed-range cut through the compensated response is not sufficiently constant that the transform phase corresponds to residual scatterer motion, the compensation is likely poor. In such a case, we switch the range and Doppler track to another scatterer. If the amplitude of the transform of a fixed-range image cut in the first (standard motion compensation) image does not show an extended drop to the background level, the peak tracking is omitted.

The processor now has found a prominent response, if necessary by trying different choices, whose transform has a sufficiently constant amplitude function for the phase function to describe the residual motion of the scatterer. The phase function then is examined to determine whether there are one or more breaks in the phase slope, and whether the phase function is linear or curved. The same analysis is performed on a second prominent fuselage scatterer that is as far separated in range from the first scatterer as possible, with the difference that for the second scatterer we can accept a transform amplitude that has the modulation introduced by two interfering scatterers. In the latter case the processor performs phase-slope tracking and generates the phase function of the dominant scatterer from the succession of phase slopes. However, if a scatterer is so far separated that the amplitude test shows that it drifts out of the range gate, a closer scatterer is selected. The processor compares the phase functions of the transforms for the two scatterers, and determines whether there is a phase-slope break within the entire

interval over which data have been collected. If not, the maximum imaging interval is the data collection interval. If there are phase-slope breaks in the two phase functions, the processor determines the longest common interval without a phase-slope break. This shorter interval then becomes the maximum imaging interval.

If the presence of phase-slope breaks forces the selection of a shorter imaging interval, the processor reduces the original imaging interval and generates an image over the shorter interval, retaining the compensation used for the longer interval. Next, the processor checks whether the wingtip and delayed responses are sufficiently resolved from the fuselage responses, and whether the fuselage extends over at least 20 resolution cells. If this is not the case, then the processor goes back to the image where a maximum imaging interval was selected by determining whether or not the phase function of the transform of fixed-range cuts has breaks. If the imaging interval was reduced because of such breaks, the longer interval is chosen for imaging, with the motion compensation executed over the break in the phase function. After iterating the image generation, the same resolution test is applied. If the test is successful, the processor proceeds to the next step of testing the quality of individual responses.

The quality of the individual responses is acceptable if most fixed-range image cuts through fuselage responses produce acceptable one- or two-scatterer patterns. If the quality is unacceptable, the processor checks whether there is an excess of resolution above the minimum required. If so, the processor reduces the imaging interval in proportion to the excess, then repeats the quality check. If there is not an excess, the radar must collect more data and try identification over the new observation interval.

In an image of acceptable quality, the processor examines transforms of image cuts in the range gates of two scatterers near the ends of the aircraft in order to determine whether the rotation rate varies enough to require resampling, and whether range drift is large enough to require polar reformatting. If so, they are applied. Next, several separated fuselage responses are examined to determine whether the phase functions of transforms in their range gates have a common curvature [3]. If this is so, then the common curvature is compensated for the entire image.

3.8.3 Step 3: Image Analysis

The image generated at the end of Step 2 is analyzed in two ways. First, the processor extracts from the image any special features (minimum length, wingtip missiles, locations of engines, etc.) that may be measurable. Second,

the processor uses the two-dimensional TSA to determine the range/cross-range positions of the scatterers on the fuselage. The processor also measures range or crossrange extent for well-resolved prominent features. The position measurements are generally initialized by taking a series of image cuts through each two-dimensional amplitude peak on the fuselage. Additionally, the processor examines a one-dimensional cut along the fuselage, uses the one-dimensional TSA to find indications of scatterers away from the one-dimensional peaks, and initializes the two-dimensional TSA at any indicated positions that are not close to a two-dimensional peak.

This analysis, as well as that of some of the earlier steps, requires approximate definition of the fuselage in the image. This implies recognition and exclusion of delayed duct returns and returns from moving devices. The position measurement also requires that a minimum amplitude threshold be set for applying the TSA. The processor does not initialize the two-dimensional TSA at any point whose amplitude is less than four times that of the noise background. Another threshold based on the amplitude distribution of the two-dimensional peaks is used to limit the analysis to the dominant scatterers of the target (typically, 20 to 30 peaks), as discussed in Appendix H.

3.8.4 Step 4: Identification

The processor performs the positional match for the features on the fuselage for all aircraft in the database. With the additional use of the measured special features it then identifies the aircraft.

References

- [1] Rihaczek, A. W., *Principles of High-Resolution Radar*, Rpt., Norwood, MA: Artech House, 1996. (See Reference 2 on page 161.)
- [2] Hauss, B. I., et al., "Airborne Target Identification From Low Crossrange Resolution ISAR Imagery," *SPIE Proc.*, Vol. 2845, Aug. 1996, pp. 120–132.
- [3] Rihaczek, A. W., and S. J. Hershkowitz, *Radar Resolution and Complex-Image Analysis*, Norwood, MA: Artech House, 1996.

4

Ground Vehicle Identification

4.1 Variability of Conditions for Identification

This section gives an overview of the various conditions under which ground vehicles must be identified. Ground vehicles undergo far more varied motions than any other type of target. Consequently, their identification requires a much higher degree of adaptivity. The radar processor must analyze the existing situation and, depending on the results, choose the appropriate approach and the specific algorithms. An overview of the general conditions for ground vehicle identification is given in Table 4.1.

The use in Table 4.1 of the terms “SAR images” and “ISAR images from SAR scenes” does not imply that our treatise applies only to ground vehicle identification by SAR surveillance systems. If identification is to be performed with a tracking radar carried by an aircraft or a missile, crossrange resolution is obtainable just as in a SAR surveillance system, provided the platform is not moving on a collision course with the target at the time of identification. After the relative motion of the platform is compensated, the conditions are the same as with a SAR surveillance system. A reference scatterer on a stationary ground vehicle or elsewhere in the scene can be used to achieve the same quality motion compensation as with a surveillance system.

Table 4.1
General Conditions for Ground Vehicle Identification

	Radar Platform Stationary	Radar Platform Moving
Target Stationary	Range profiles only	SAR images
Target Moving	ISAR images	ISAR images from SAR scenes

4.1.1 Application 1: Radar Platform and Ground Vehicles Stationary

When a stationary ground vehicle is to be identified by a stationary radar, resolution can be only in range, and only range profiles are available for identification. We have discussed extensively the unsolvable problems of reliably identifying flying aircraft from range profiles in a large database. In one sense, the situation is more benign for stationary ground vehicles, because of the absence of delayed duct returns. On the other hand, it may be more difficult if the ground vehicle is in a tree environment, where masking by trees and tree clutter may be a major problem. Although one cannot reject the possibility of identifying a specific stationary ground vehicle from its range profile, reliable identification for a large database appears to be impossible. In practice, a ground-based radar platform for identifying stationary ground vehicles will be mounted on a vehicle, which should move for at least a short time during the identification process in order to generate a synthetic aperture. Our discussions on identifying stationary ground vehicles in SAR imagery then will apply. We will not further consider the case of both a stationary ground vehicle and radar platform.

4.1.2 Application 2: Radar Platform Stationary, Ground Vehicles Moving

The difference between identifying a moving target from a moving or a stationary platform lies in the compensation of the platform motion. In the case of a SAR system, the motion compensation is usually good enough to make this difference practically insignificant. A stationary platform has a clutter advantage for most moving targets; clutter is concentrated near zero Doppler, so will only interfere with targets with low range rates. For these targets, whose unfocused images may overlap zero Doppler, the clutter is worse for a stationary platform. On the other hand, clutter cancellation is more effective for stationary platforms, because the antenna position does not change at all between consecutive looks.

The case of the stationary platform thus is included in the identification of moving ground vehicles by a SAR system, so that it need not be treated separately. Thus, we will not specifically consider identification of moving ground vehicles from a stationary platform.

4.1.3 Application 3: Radar Platform Moving, Ground Vehicles Stationary

Practical identification procedures for stationary ground vehicles depend strongly on the environment. A ground vehicle may be in a sandy or grassy area or on a paved road, all with relatively low backscattering, so that the target returns totally dominate over the returns from the terrain. Interference from the environment then is negligible, and identification requirements are similar to those for an aircraft, but simplified because an aircraft may execute small inadvertent yaw maneuvers even when flying along a straight course. If the vehicle is next to a tree but not obscured by it, the tree returns can be discriminated from vehicle returns using the methods of complex-image analysis. On the other hand, if the ground vehicle is in a forest area, is partly or fully underneath a tree, or in the shadow of a tree line, identification becomes progressively more difficult as the vehicle is more heavily obscured and the density of trees increases.

When stationary ground vehicles are camouflaged or hidden behind or underneath trees, or even are in the clear but in the vicinity of trees, conventional threshold detection will not provide sufficient detection performance. In [1] we discuss a set of algorithms that utilize the complex responses to distinguish between returns from metallic (wave trapping) features and natural objects. These algorithms give a detection performance far superior to that obtainable with any kind of threshold detection. Proceeding from detection to identification under obscurant conditions also requires the complex-image analysis technology. At the time of this writing we had not yet explored this difficult subject. Hence, we do not have an opinion concerning the performance potential. We merely point out that going to such low carrier frequencies that trees cannot effectively shield vehicles may be problematic because differentiating man-made from natural features is more difficult, perhaps impossible, when the wavelength is too large relative to the size of the design features of a vehicle. Also, the use of a large percentage bandwidth in order to achieve high range resolution, and of a wide angular sector for SAR imaging, as required at low carrier frequencies, might lead to problems. In this book, we will consider the identification of stationary ground vehicles only when they are in the clear.

4.1.4 Application 4: Radar Platform Moving, Ground Vehicles Moving

An environment in which stationary vehicles can hide underneath trees presents less of a detection problem for moving ground vehicles. If the range rate of a vehicle is large enough to shift the vehicle return outside the terrain clutter spectrum and the PRF is high enough to avoid Doppler foldover, as long as the vehicle returns are not totally attenuated, the situation is almost the same as if there were no terrain. However, ordinarily we cannot confine our attention to vehicles with such high range rates, because even fast-moving vehicles have low range rates when the aspect angle is large. Thus, we must be able to accommodate low range rates, in which case the vehicle returns remain within the terrain clutter. Furthermore, conventionally designed SAR systems have low PRFs, so returns from vehicles with high range rates also remain within the clutter.

The identification of a moving ground vehicle in a SAR scene (surveillance or otherwise) poses serious problems, because the image of such a vehicle typically is highly smeared in crossrange. As one of the consequences, the strength of the responses is much decreased relative to the case of a stationary vehicle. In addition, since the range rate of the vehicle translates the overall image in crossrange, the image of the vehicle may be superposed on nearby trees or other strong clutter. All of this implies a potentially serious detection problem, which can be satisfactorily solved only if the SAR system employs clutter cancellation. With a displaced phase center antenna, the radar effectively examines the terrain from the same spot but at two different times, canceling the returns from those objects whose ranges have not changed between the times of the two looks [2]. Although the method cannot work for vehicles with very low range rates, the returns from these vehicles also cannot be translated much—that is, not into a tree area. This means that clutter cancellation is not as important for low-range-rate vehicles. Whenever there might be problems from trees, we will assume a SAR system with clutter cancellation. This allows us to treat the terrain as noise when the ground vehicles are moving.

Even without terrain interference, the varied motions of ground vehicles require that identification be adaptive to an extraordinary degree. These motions depend on conditions that are roughly indicated in Table 4.2.

Two-dimensional imaging and identification based on ISAR images becomes difficult when the moving target does not act as a rigid structure. A heavy armored vehicle will usually act as a rigid target, but a truck or another vehicle of lighter design often will not. Whether or not flexing is a problem depends not only on the design of a vehicle but on whether it is moving on a

Table 4.2
Conditions for Moving Vehicles

Vehicle Design	Rigid	Flexible
Surface	Road	Terrain
Surface Structure	Smooth	Bumpy
Path	Straight	Curved
Speed	Low	High

smooth road, a rough road, or off the road, whether it is moving slowly or rapidly, and whether along a straight or a curved path. The same vehicle thus can act as a rigid vehicle in one situation, have significant flexing in another, and be a strongly flexing/vibrating target in a third.

Let us ignore any flexing for the moment, and consider a vehicle moving on a paved road. The requirements on and the precise approach to identification still vary greatly, depending on the other conditions. When the road is relatively smooth, vehicles at low speeds will have negligible bouncing, rolling, pitching, and yawing motions. (We must keep in mind that microwave radar is very sensitive to small motions about the center of gravity and small bouncing motions.) With increasing speed, these motion components become stronger until they govern the situation. However, this all depends on whether a vehicle is a light commercial vehicle, a heavier truck, or a very heavy military vehicle such as a tank. On poor roads, these additional motion components become significant at much lower speeds, and at faster speeds the various motion components will be as significant as they are for a ship, but much more erratic. When military vehicles move on terrain rather than on a road, it is even more important whether their speeds are low or high. There also is a significant difference between motion along a straight line and turning motion, since this will change the relations between the Doppler from the translational motion and the Dopplers due to bouncing and to motion about the center of gravity.

It should be obvious that there cannot be a single procedure for dealing with such a variety of vehicles, speeds, and motion behaviors. However, no implication is intended that identification of moving ground vehicles poses unsolvable problems. We merely must design an identification system that is sufficiently adaptive to deal with such a variety of conditions.

The variety of conditions for moving ground vehicles has an undesirable consequence for this chapter. The treatment of the identification of

moving ground vehicles cannot be simplified into a few representative cases, which makes reading this chapter more demanding than reading the others.

4.1.5 Section Summary

Identification of stationary vehicles from a large database requires that the radar platform move in order to generate a synthetic aperture. Identification of moving vehicles may be performed from stationary or moving platforms, both of which should employ clutter cancellation.

The varied motions of ground vehicles require that identification be adaptive to an extraordinary degree, depending on vehicle design, path, and speed, as well as surface type.

4.2 Basics of Ground Vehicle Identification

This section provides a general treatment of ground vehicle identification, in particular the difference between the identification of stationary and moving ground vehicles. The similarities and differences relative to aircraft identification are also discussed.

There is a basic difference between the identification of stationary and moving ground vehicles. Identification of moving ground vehicles is similar to the identification of aircraft. We will rarely obtain an image of such quality that we can perform accurate length and width measurements. Since the accuracy to which the crossrange scale of an image is known is usually poor, a reasonably accurate length measurement can be performed only at small aspect angles, but in such geometries we might not detect the scatterers at the end of the vehicle reliably. Thus we must depend on the positions of scatterers whose functions cannot be determined from an image, plus such special features as the presence of a turret, or ordinary wheels rather than a tread. The difference relative to aircraft is that as one progresses from moving vehicles under benign conditions to more serious conditions, progressively fewer of these features are measurable. In the worst case of vehicles moving on bumpy roads or terrain, only the Doppler-assisted measurement of the range positions of design features remains. This will become clear as we continue.

Identification is much simpler for stationary ground vehicles. Provided one utilizes the complex image, it is possible to measure length and width rather accurately. These two measurements alone already represent a big step toward identification. The measurement of other design features, such as presence of a turret, also is easier (or possible), and scatterer positions can be

accurately measured in both dimensions, rather than only in range. In fact, since *no adaptivity is required in the imaging process for stationary ground vehicles*, we can readily list the various processing steps.

1. Excise the complex image of interest from the SAR scene.
2. Analyze the response peaks to determine scatterer positions and characteristics.
3. Fit a straight line to the response peaks along the illuminated long edge, find the responses at the corners of the edge, and determine the length.
4. Fit a straight line through the responses along the short illuminated edge and find the far vehicle corner along this edge. Also examine cuts perpendicular to the long edge to obtain the width more reliably.
5. Examine the vicinity of the nonilluminated edges to refine their locations, and the length and width.
6. Extract special features from the image.
7. Identify on the basis of the two-dimensional positional match, length and width, and special features.

These steps involve a small degree of adaptivity, which will be explained below.

In contrast, no such simple recipe can be given for the identification of moving ground vehicles. In our treatment of moving ground vehicles, we will first summarize the consequences of identification of various types of motion, then list the general processing steps, and lastly provide the processing details with examples for representative conditions under which moving ground vehicles are to be identified.

4.2.1 Radar Features Versus Vehicle Design Features

The database needed for ground vehicle identification is generated in the same way as explained for aircraft. For those ground vehicles for which radar data are available, be it from SAR systems or turntable data, we will extract the identification features from the data. When data are not available, we use diagrams, photographs, and CAD models to determine the measurable features and the positions of observable scatterers, as explained in Chapter 3 for aircraft.

For stationary ground vehicles, the most important identification features are the length and width of the vehicle, but only if measured accurately. Too many ground vehicles have similar lengths and widths, in particular widths, to make inaccurate measurements very useful. The critical point is that *length and width of a vehicle must be defined so as to fit radar identification*. As we have stressed repeatedly, radar can generally observe only features with wave-trapping shapes. Thus, as an example, if the corner of a vehicle is defined by the corner of a fender shaped in such a manner that it cannot trap the wave at a particular aspect angle, the associated response will typically be very weak or unobservable in the background. This does not matter, however. In establishing the database, we determine which wave-trapping feature along the edge of the vehicle is closest to the vehicle corner, and this feature is used for defining the radar length of the vehicle. On occasion, a weaker feature might be observed at a vehicle corner because the interference background happens to be low. We can define the radar length of the vehicle in terms of the strong wave-trapping features, and then ignore a response that is weak compared with the typical responses of the image. Alternatively, the database could contain the weaker corner feature as well, in case a weak response is observed.

For example, for identification purposes the corner of a tank might be defined by the drive wheel, if this wheel has a wave-trapping cavity. A ground vehicle with a length of 6m and a width of 3m might have a radar length and width of 5.2m and 2.6m in the database, based on observed strong vehicle-corner responses. The measurement accuracy achieved with the methods of complex-image analysis is so high that any significant error in the length or width measurement will be due to a misinterpretation of the observed scatterer rather than a measurement error.

4.2.2 Recognizable Features

Although we have emphasized the importance of performing the measurements on the responses of the complex image, at least one special feature can be derived from the intensity image. This is the presence or absence of a turret. Because of its rounded shape, there are generally no significant wave-trapping features located on a turret, so that no strong responses are generated by turret features. In addition, a turret shadows a good part of the deck. The resulting absence of significant responses over an extended area, due to both effects, can be determined from the intensity image, since only return strength is utilized. Nevertheless, there is always some uncertainty involved.

Thus, if an extended area with weak responses is observed in the intensity image, we will conclude that the vehicle has a turret; if such an area is not reliably observed, we will not make a conclusion with respect to a turret.

Such an asymmetry is also applicable with respect to the gun of a tank or howitzer. A gun will generally not be observable, unless the radar happens to look into the barrel and the opening is much larger than the wavelength. The response from the opening of the gun barrel can be discriminated from clutter on the basis of its phase function, which will be curved because of the shifting phase center. Hence, if an isolated response is detected away from the vehicle, and the phase function of the response is curved, we conclude that the vehicle has a gun. If such a response is not detected, we do not make a conclusion one way or the other.

A third example of a feature that will be utilized when observed but ignored when absent is any rotating device with a wave-trapping design. If the device is in operation and Doppler returns distributed over some cross-range interval are observed, the detection of the feature will help identification. However, we do not require that the feature be observable if the vehicle carries it.

The number of wheels and the type of wheels are other important design features. If the wheels are cup-shaped so that they can "trap" the wave, and the wheels are solid rather than consisting of spokes, the wheels will generate observable returns. In this case we have information about the type of wheel, and we can relatively easily determine the number of wheels and their locations. Further information is obtained from the degree of difficulty in detecting the wheel returns. If the side of the vehicle is smooth, the wheel responses will be conspicuous and easily measurable. On the other hand, if the side of the vehicle has an open design that permits the radar wave to illuminate a variety of other wave-trapping features along the edge, detection of the wheel returns will be more difficult.

Certain vehicles may also have special features not found on other vehicles. For example, a vehicle might have a significant trihedral corner when viewed from certain angles, and the presence of the feature can be verified from the strength of the response and the fact that the phase function is planar (no shifting phase center). Another vehicle might have a similarly conspicuous feature, with a design that cannot be approximated by a trihedral corner. We then will find a strong response with a curved phase function. A variety of such special features may be found on only one or some of the vehicles. These features must be noted in the assembly of the database, and then used to assist identification.

4.2.3 Positions of Unrecognizable Features

Ground vehicles have many observable features that have no special significance, in which case we can utilize only their positions. This is similar to the variety of antennas, air inlets and outlets, and so on for aircraft. In the case of ground vehicles these observable features (again, of the wave-trapping kind) are often formed by metallic boxes and similar structures on the deck that create trihedral-like corners and cavities. It is easy to see that an arrangement of such features may form trihedral corners. There are also arrangements that leave small cavities between the features, and these cavities become observable features. Again, in the absence of radar data, the presence and locations of such features must be determined from an examination of diagrams, photographs, and CAD models.

Such large openings as, for example, an open access hatch of a turret are typically not observed. Although the incoming radar wave will certainly be trapped by such a feature, it will usually be so well trapped that no strong scattered signal finds its way back toward the radar.

4.2.4 Section Summary

Stationary ground vehicles can be imaged and identified via a straightforward sequence of processing steps that involve only a small degree of adaptivity. Identification is on the basis of the two-dimensional positional match, length and width, and special features. The imaging process for moving vehicles must be highly adaptive to the vehicle motion. Usually, fewer features can be extracted from their images than from those of stationary vehicles.

The features are compared to a database constructed from radar data if available, and from photographs, diagrams, and CAD models otherwise. Length and width are defined by those features observable by radar, those with wave-trapping shapes.

4.3 Identification of Stationary Ground Vehicles

This section gives specific examples of the identification of stationary ground vehicles. For the reader only casually interested in the identification of stationary ground vehicles, attention to any one of the below examples will suffice.

In this simplest of all identification cases treated in this book, the SAR processor performs the motion compensation and there are no uneven target motions, so that only the evaluation of the high-quality (complex) image

is needed. Our discussions thus will be concerned with the many measurements one can perform on the images of stationary ground vehicles. The problematic motion compensation needed for moving ground vehicles is no issue, even if the radar platform should not be a SAR surveillance system.

The real data used in our examples were collected under conditions where masking by trees was not a problem, so that the entire vehicle image is available for identification. The data were taken at X-band, with both range and crossrange resolution of about 1 ft. All images will be shown in peaks plot form, with the “gates” in range and crossrange corresponding to the inherent resolution of $1/B$ and $1/T$. We will consider the measurements on four different ground vehicles: a howitzer, a tank, an off-highway truck, and a flatbed truck. Actually, we should compare the measurements extracted from the image of one vehicle to the database of the other three vehicles, performing the positional matches for the correct and incorrect vehicles. However, with the methods of complex-image analysis, the typical military ground vehicles in our examples are sufficiently different from each other to make such a comparison useless. It is not like comparing one jet fighter with another fighter of about the same size. The problem is not to distinguish, say, a tank from a truck, but to distinguish one type of tank from another type. Although our methods are aimed at this more difficult task, a treatment of the last identification step requires working with a large database in order to obtain statistically meaningful results. This book addresses the principles and procedures of identification rather than the achievable performance. Since the same type of processing and analysis is used on all stationary ground vehicles, the various cases treated in the following are examples to make the reader familiar with the technology. For the moving ground vehicles, in contrast, the aim of the examples will be to illustrate the necessary adaptation of the processing method to the variety of possible conditions.

4.3.1 Example 1: Measurements and Characterization for a Howitzer

4.3.1.1 Measurement of the Target Outline

Measurement of the target outline is the most important step in identifying stationary ground vehicles. It serves more than just the measurement of length and width, also allowing one to determine the positions within the target outline of any special features that might be detected.

The peaks plot corresponding to the SAR intensity image of the howitzer is shown in Figure 4.1 for an aspect angle of 65° off tail-on, with the vehicle outline indicated by the dotted rectangle. An examination of the image by eye shows that the vehicle has a rectangular shape, so that the outline can be




Figure 4.1 Image of the howitzer.

determined by measuring (in the complex image) the two illuminated edges of the vehicle, finding the three vehicle corners, and drawing lines through the far vehicle corners parallel to the edges.

With automated processing, the shape of the image must be determined from those measured scatterer positions farthest left and right in range gates and lowest in crossrange gates, before a decision as to how to proceed can be made. The scatterer positions used to determine the shape of the target outline are measured by applying the two-dimensional TSA to the strong two-dimensional peaks of the target. We next find the near corner of the target, by examining the lowest-range strong response in an image cut along the long illuminated edge (both illuminated edges are determined by a linear fit to the selected scatterer positions). It is easy to find the strong response at the near corner of the image. The complex form of this response is analyzed with the one-dimensional TSA to determine whether it comes from one or two significant scatterers, as explained earlier. If it comes from a single scatterer, the peak position of the response is taken as the position of the near corner of the vehicle. If it comes from two responses, we select the measured position closer to the corner, if both responses are of roughly comparable strengths. If

one is much weaker, we use the stronger response to define the vehicle corner, in accordance with the earlier discussion on the definition of vehicle length and width.

Having found the near corner of the vehicle, we must now search along each of the two edge lines for the last response strong enough in relation to the other responses of the vehicle (see Appendix H). These last responses are taken initially to define the two far corners of the vehicle. We then examine a series of cuts near the short illuminated edge and perpendicular to the long illuminated edge, to better define the width. This establishes three corners of the vehicle, and hence the outline of a rectangular vehicle.

The determination of the short illuminated edge is somewhat more problematic than that of the long edge, because it may not be as well defined by the design of the vehicle. However, once the rectangular shape of the vehicle has been recognized and the long illuminated edge has been determined, we can draw a perpendicular to the long edge at the near vehicle corner. This assumes equal range and crossrange resolutions; otherwise we can calculate the appropriate angle. We again search for the last scatterer along the straight line that defines the shorter edge. In general, we must search for the last scatterer along each of a set of parallel lines, all perpendicular to the longest straight edge. These scatterers are used to derive the target width. In this particular example, the radar length and width of the vehicle were measured in this manner with errors of 1%, given that the correct scatterers were used in the radar definition of length and width of the vehicle.

In some cases, prominent scatterers also are observed along the nonilluminated edges. When such responses are observed, we can use them to confirm the measurement of length or width as derived by searching for the last scatterer along an edge. For example, treaded vehicles are usually designed in such a way that the radar beam has access to parts of the tread and to wheels near the far corner of the short illuminated edge. Accordingly, in Figure 4.1 we see four significant responses along the near-range part of the nonilluminated long edge, which can be used to confirm the width. On the other hand, in the case of Figure 4.1 we should not attempt such a confirmation for the length, because the responses near the nonilluminated short edge do not lie along a perpendicular to the long edge; individual responses may be delayed returns from multipath reflection within the wheel/tread assembly. This is a situation where the superstructure, here the turret and gun base, mask the short edge, so that no perpendicular line of responses is observed.

Note that the *large region with low responses in Figure 4.1 implies that the vehicle has the special feature of a turret*. Once this has been established, it is reasonable to allow a break in the series of responses defining the long

illuminated edge because at some orientations the turret may mask part of the edge. This fact can be taken into account in the process of fitting the straight line to the edge responses. However, the edge in the image is not necessarily defined by the scatterers on the top of the deck, which might indeed be shadowed by the turret. The edge can just as easily be defined by the wheels and other scatterers below the plane of the deck, and these will not be shadowed unless the turret is positioned so that its overhang is very large and the SAR system operates with an unusually large beam depression angle.

4.3.1.2 Other Distinctive Features

The scatterer positions derived from an image cut along the long illuminated edge in Figure 4.1 agree well with the peak locations. *The nearly regular spacing of the scatterers along the long illuminated edge implies that this is a case where the wheel responses dominate the edge returns.* There are no wave-trapping parts of a ground vehicle with such regular spacings along the long edge but the wheels. As one special feature, the ease with which the wheels can be detected implies that the vehicle has a smooth side, meaning flat plates. As an even more important special feature, we can determine the number and positions of the wheels. Indeed, when the complex responses along the edge are examined, one finds that they all come from single scatterers, except for the strong response at the near vehicle corner. The analysis of this response shows that it comes from the contributions of two scatterers of comparable strengths. One must be a road wheel, and the other a drive or idler wheel. Since the rightmost response along the illuminated edge is found to come from a single scatterer, we have a total of nine wheels, or seven road wheels.

As another special feature, the image of Figure 4.1 contains a very strong response near the illuminated short edge. The strength of the response implies that the feature must be fairly large. An analysis of the phase function of the response shows that it is essentially planar, with only a small curvature. The vehicle thus has a sizable feature that approximates a trihedral reflector well. A large trihedral in this position is not a feature common to many ground vehicles, and thus is a valuable help for vehicle identification.

If the strengths of the responses outside the vehicle outline are examined, we find one with coordinates C16.29/R14.10 (indicated by the arrow) that is much higher (16 dB) than the background. A closer examination of the image cut in the crossrange gate of this response shows an entire series of responses with the same crossrange positions. Thus we have a set of delayed responses. The complexity of the response provides a strong indication that this is a gun response. The facts that the response can be detected with the

gun pointed away from the radar and that multiple delayed returns exist are strong indications of a gun with a muzzle brake. In the identification process we will favor a vehicle carrying a gun with a muzzle brake, but based on this single measurement we still will not reject vehicles without a gun (there might be a wave-trapping metallic feature on the ground). The detection of a widely extended turret shadow also points toward a vehicle with a gun, but we will not reject a vehicle whose gun does not have a muzzle brake.

The analysis of the major responses of the image reveals at least one additional peculiar response at C-9.05/R0.03 that can be utilized as a special feature. The image cut in the range gate of the response shows a relative half-power width significantly larger than unity, so that the feature must have considerable crossrange width. As a representative illustration of this type of applied analysis, in Figure 4.2 we give the transform of the image cut in the range gate of the response. The amplitude is essentially constant, which implies that the response comes from a single scatterer. The phase function is strongly curved, which means that the scatterer has a sizable crossrange width. In accordance with the complex-image analysis methods, the crossrange width is obtained from phase-slope measurements over the intervals indicated by vertical lines in the figure. The phase slopes correspond to Crossrange Gates 0.7 and -0.5 , and the feature width is the difference of 1.2 gates. With a crossrange gate width of 1 ft, we conclude that the vehicle has a cavity-type scatterer with a width of about 1 ft in this position.




Figure 4.2 Transform of the image cut in Range Gate 0.

An image cut in the crossrange gate of this response shows that the response is even more complicated in the range dimension than it is in crossrange dimension; it has a large spread in range. Furthermore, the same crossrange positions are measured in different range gates over the spread response. This is indicative of unresolved delayed returns. The feature thus generates multiple delayed returns of small separations. This means that it has significant range depth, which can be measured in the image cut in the crossrange gate. The main secondary response is found to be delayed by about 1.5 range gates. The conclusion is that the feature is a cavity with a width and depth exceeding 1 ft. There are not many ground vehicles with such a conspicuous feature in this particular position.

4.3.1.3 Positional Match

Having considered a set of important special features (which, together with length and width, should suffice for vehicle identification in this relatively simple case), we next illustrate the positional match that includes those features for which position is the only usable "characteristic." In this example, the comparison template of observable scatterers was generated by examining photographs and line drawings of the vehicle. Because the vehicle's turret can be oriented in an arbitrary direction, the database does not contain any scatterers on the turret. Accordingly, having recognized that a target has a turret, we measure the positions of only the scatterers safely outside the turret area, or roughly along the edges.

Matching the template to the scatterer positions determined by analyzing the prominent complex responses of the image gives the positional match of Figure 4.3, with template features identified by letters and measurements shown by crosses. The match is very good, except for scatterer S. This is the complicated cavity investigated above. Since we know that the shifting of the phase center can translate the response in an unpredictable manner, we assign low weight to the positional match of such a complicated feature. Its value lies more in the fact that it serves as a special feature. Note that, in this fashion, the characteristics of the image responses themselves provide the information about the accuracy of the positional match for each scatterer.

4.3.1.4 Effect of Aspect Angle on Identification of the Howitzer

The image of the same vehicle for an aspect angle of 35° off nose-on is shown in Figure 4.4. Compared with the earlier image of Figure 4.1, the rear view has changed into a front view, and the orientation of the long edge is rotated by 20° . (Note that the crossrange resolution differs slightly in all the various stationary ground vehicle images we present.)




Figure 4.3 Positional match for Figure 4.1.




Figure 4.4 Image of the howitzer at a different aspect angle.

The new image shows the same nearly regular arrangement of the responses along the long edge as for the earlier aspect angle, so that the image again reveals that the side of the vehicle has a smooth design, except for the presence of the wheels. As we show below, the measurement of the wheel positions is even easier than before. The shadowing of the deck is not as extensive as for the rear view, but the presence of a turret can nevertheless be recognized. Although the gun is even better detectable at the new aspect angle, we cannot conclude from the position of the gun that the radar views the front of the vehicle, because the turret and gun could be oriented in any direction. Reliably determining the orientation of the vehicle requires the positional match.

The measurement of length and width is as easy as for the former orientation, primarily because strong scatterers define the corners of the vehicle. It is clear from Figure 4.4, which shows the weak delayed responses generated by the vehicle at larger ranges, that the responses at the vehicle corners do not have an interference problem from the terrain. The weakest corner response, the one at lowest range, is 22 dB stronger than the clutter. These responses are easily analyzed with the TSA. The separation of the two vehicle-corner responses, which is the effective radar length of the vehicle, is found to be 0.52m larger than for the rear aspect of Figure 4.1. This means that the effective length of the vehicle is different for front views and rear views, which is easily understood, because a feature that tends to trap the wave for a front view may not trap the wave for a rear view, and may not even be observable for such a radical change of the aspect angle.

As already discussed in connection with aircraft, the database must contain the observable scatterers for different angular sectors covering 360° , and hence it also must contain the effective vehicle length for the various aspect angle sectors. When no radar data are available for extracting the database, identification features such as the effective length must be determined by examining diagrams and photographs, looking for those scatterers that will be observed at the vehicle corners for a given aspect. In this instance a photograph indicates that for a frontal view the observed scatterer at the near vehicle corner should be the cavity represented by a headlight. At the far vehicle corner the radar should see the cup of a wheel. However, because the wheel is viewed at a relatively small angle of incidence, the effective position should be the far inner edge of the wheel cup rather than its center. The separation of the two features defined in this manner agrees with the measured vehicle length, within the combined length uncertainty. Quite generally, the quality of the measurement on the vehicle-corner responses will determine whether the error in the length measurement is governed by the error in the

position measurement on the scatterers, or by the error in estimating the effective positions of the vehicle-corner scatterers from photographs and diagrams. For this example, the existing interference conditions in the image indicate an expected measurement uncertainty on the order of 5 to 10 cm.

In order to illustrate the general accuracy of the TSA in a more complicated situation, in Figure 4.5 we show the image cut in the range gate of the peak that defines the far vehicle corner of the short illuminated edge. The relative half-power width of the corner peak is 0.757. This is so much less than unity that two scatterers interfering destructively must be involved. Specifically, the interference is between the responses centered in Crossrange Gates -13.3 and -11.1 in Figure 4.5, the one in Gate -13.3 being the vehicle-corner response of interest. For comparison purposes, we first take the position of the response peak as the position of the single scatterer. Since the interfering response to the right shifts the peak of the weaker response to




Figure 4.5 Image cut in Range Gate -3.85 .

lower crossrange gates, and this is the left corner of the image in Figure 4.4, the measured width will be too large. The situation is similar for the response that defines the right front corner of the vehicle, so that the total measurement error will be significant. When we proceed in this coarse manner, from the range/crossrange positions of the two scatterers we obtain a width that is too large by 0.32m.

The error in the width measurement from this crude approach is unnecessarily large. The relative half-power half-width on the right side of the response peak is 0.716, so small compared with unity that the error associated with taking the peak position as the position of the scatterer must be quite large. In this situation, for better accuracy we must determine the position of the corner scatterer with the TSA, using a transform window that includes both responses in Figure 4.5. This is simple when the second response comes from a single scatterer. In this instance, however, the transform shows that a third scatterer is involved, so that the right boundary of the window must be shifted away from the minimum in Gate -9.1 of Figure 4.5, toward the strong peak. A good compromise between cutting off too much of the response of interest and including too much of the interfering response is a boundary around Gate -10.2 . The resulting transform pattern is not good in that it does not fully represent the ideal interference pattern between two scatterers, but this type of measurement still gives better results than simply taking the peak position. The conditions are similar for the other corner response, where the TSA must also be used. When the positions of the two corner scatterers are determined in this manner, because of the unfavorable circumstances the calculation of the separation still gives a vehicle width that is too large, but only by 12 cm instead of the 32 cm when the peak positions were taken. The difference is practically significant.

Although earlier we made conclusions about the design of the side of the vehicle and the measurability of the wheel positions by examining the peaks plot image of Figure 4.4, in practice any conclusions must be based on analyzing the responses in an image cut along the long edge. This image cut is shown in Figure 4.6. Starting in Range Gate -6.3 and ending in Range Gate 5, we count seven responses from the road wheels, with the response in Range Gate 3 reduced by interference. The last response of the row is from the idler wheel. As the image of Figure 4.4 shows, the image cut along the long edge does not pass through the first two peaks, so that no analysis of the two weak responses in Range Gates -9.6 and -8 in Figure 4.6 should be performed in this image cut. When the image cut is translated so that it does go through the two peaks, the phase function of the first response is straight, whereas that of the second is curved. This agrees with the interpretation from




Figure 4.6 Image cut along the illuminated long edge.

the photograph that the first scatterer is the headlight and the second scatterer is the drive wheel. The headlight is small enough that its phase-center shift is small, whereas the drive wheel has a large cylindrical cavity.

As another example of the analysis of a response, we consider the strong response near the center of the image of Figure 4.4. Its strength already implies that it must come from a sizable corner reflector or cavity. The transforms of the image cuts in both the range and crossrange gates of the response have curved phase functions, from which one can determine the effective width and depth of the cavity. It is the “cavity” formed by the complicated design of the base of the gun.

4.3.1.5 Persistence of Scatterers

We have emphasized that target identification is impractical unless the same scatterers can be observed over a reasonably large aspect angle sector. As another demonstration of scatterer persistence over large aspect angle sectors, we perform the following test. We generate images of the howitzer for aspect angles of 31° , 56° , 105° , and 110° , all measured from nose-on. The first two images thus represent front views, and the other two rear views. The last two images are close to broadside, where all measurements are difficult. As an

illustration of this fact, in Figure 4.7 we show the vehicle image for the aspect angle of 110° , still 20° away from broadside. Compare this image with that of Figure 4.1, for which the aspect angle is only 15° farther away from broadside, or even with the image of Figure 4.4. We measured the positions of the scatterers associated with the responses near the edges of the four new images. To evaluate scatterer persistence, we matched the scatterer positions extracted from different images to each other.

In Figure 4.8 we show the match of the scatterer positions for the images at 31° and 56° aspects, with the scatterer positions for the first image rotated and translated to achieve the best match with the positions for the second image. (The first set of positions act as a template, which is adjusted to achieve the best match with the second set.) The measurements extracted from the first image are designated by letters, and those from the second by crosses. The former measurements thus act as a candidate target from the comparison database for the match. *The match appears good enough that it alone allows us to recognize the same vehicle, despite the 25° rotation.* Figure 4.9 shows the corresponding match for the two rear-view images, with the image at the 105° aspect providing the database for the image at the 110° aspect. The majority of positions again are well matched.




Figure 4.7 Image of the howitzer close to broadside.




Figure 4.8 Match between scatterer positions at 31° and 56° aspects.




Figure 4.9 Match between scatterer positions for images at 105° and 110° aspects.

Lastly, we repeat the process for images at widely differing aspect angles, rear views and front views, matching the positions extracted from the image at 31° aspect to those of the image at 105° aspect. The corresponding match is shown in Figure 4.10. The short edges must be ignored because the radar illuminates the front edge in one case and the rear edge in the other. The match for the long edge is very good. It should be noted, however, that the quality of the match along the long illuminated edge is so good because for this particular vehicle the wheel responses dominate, so that the wheel




Figure 4.10 Match between positions from the images at 31° and 105° aspects.

positions can be determined very accurately. Normally, we cannot expect excellent matches between rear views and front views, and thus must include rear views and front views in different angular sectors of the database.

4.3.2 Example 2: Measurements and Characterization for a Tank

4.3.2.1 The Tank at an Aspect Angle of 30°

The peaks plot image of a tank is shown in Figure 4.11, with the vehicle outline indicated by the dotted rectangle, and the approximate turret outline by the dotted ellipse. The gun is pointed nearly along the long axis of the tank and away from the radar. The tank is similar to the howitzer, so that the same type of measurements as already illustrated were made. Thus we shall merely point out some interesting differences. Since the right upper quarter of the vehicle image is shadowed (only weak background responses), this vehicle must also have a turret.

The responses along the long illuminated edge are not as regularly arranged as for the howitzer, which implies that the side of the vehicle is not so smooth that the wheel responses are strongly dominant. Another significant difference is recognized when the two vehicle-corner responses for the long illuminated edge are examined. The transform of the response for the far corner gives constant amplitude, so that the response comes from a single




Figure 4.11 Peaks plot image of a tank.

strong scatterer. The transform of the response from the near corner gives the modulation pattern of two scatterers; this response is either from a feature with considerable extent (so it effectively breaks up into two parts even for a 1-ft resolution) or from interference between two features at similar range and crossrange positions but different heights. We conclude that if the drive and idler wheels represent the wave-trapping features that define the corners of the target, the wheels are of very different designs. This is consistent with the tank, but not the howitzer. When length and width of the tank are measured in the same way as for the howitzer, the measurement errors turn out to be just a few centimeters.

Since an examination of some of the responses along the long illuminated edge shows single-scatterer responses with curved phase functions, we conclude that the vehicle has cup-shaped wheels without spokes. This means that *the number of road wheels can be determined* by searching for single-scatterer responses with curved phase functions. When the TSA is used to resolve a response into the contributions from two scatterers, one of which has a shifting phase center, we must determine which of the two. The phase function of the response in the image domain indicates most easily whether

the scatterer with the shifting phase center is on the left or on the right of the peak of the combined response, with curved phase indicating the shifting scatterer.

When the series of responses along the edge are examined in this manner, the responses with curved phase functions are clearly detectable. However, the situation for each wheel varies from interference so strong that the position of the wheel can be knowingly measured only very inaccurately, to the case where one wheel response (indicated in Figure 4.11 by the arrow near the long edge) is so clear of interference that the width of the wheel can be measured in the manner illustrated by Figure 4.2. The measured phase-slope difference turns out to be $(7.37 - 5.81)$ crossrange gates. For this image the crossrange gate width is 0.28m. In the case of stationary vehicles we easily obtain the aspect angle of the vehicle, which is 30° . Calculating

$$(7.37 - 5.81) \times 0.28 / \cos 30^\circ = 0.50\text{m}$$

we find the width of the road wheel. The measurement error happens to be 1 cm. Although the smallness of the error is an accident, we would expect high measurement accuracy for a response as well resolved as this particular one. (The amplitude and phase functions are an even better approximation of the ideal case than in Figure 4.2).

Another difference can be found for the response at C-11.33/R-4.23, indicated in Figure 4.11 by the arrow near the short edge. The response has a straight phase function in its fixed-range image cut, but a curved phase function in its fixed-crossrange cut. The feature thus has significant width only in range. This is clearly a special feature. Also helpful for identification is the strong response at C10.91/R2.56, indicated in Figure 4.11 by the arrow within the tank outline. Both the phase and amplitude functions of the transform of the response vary in a complicated manner with the direction of image cuts through the response. The conclusion is that the vehicle has a cavity of complicated shape at this position, which is further verified by the fact that one can identify spurious sideband responses [1] arranged along an arc and with the directions of phase linearity all pointing toward the responsible scatterer. However, the use of such spurious responses in an automated identification system does appear to be quite difficult. One might generally be satisfied with identifying and disregarding spurious responses.

There are other responses in the image that provide interesting information about the vehicle. However, length and width measurements, the special features already discussed, and the positional match, already supply more information than should be needed for vehicle identification.

4.3.2.2 The Tank at a Different Aspect Angle

The image of the tank at the new aspect angle is shown in Figure 4.12. To determine length and width, we need the three vehicle-corner responses. An examination of the image cuts along the two illuminated edges shows that there are no responses not indicated in the peaks plot sufficiently stronger (12 dB) than the background to be accepted as vehicle responses, much less strong enough to be considered wave-trapping vehicle-corner responses. Thus we analyze the three visible vehicle-corner responses (indicated by horizontal arrows) in the usual manner, taking transforms when the relative half-power width is not close to unity, and obtaining the scatterer positions and strengths via the TSA. When two scatterer positions are obtained from a single response, the position that will result in a larger length or width is chosen. However, the two near-corner responses are significantly weaker than most target responses; they come from features that might not be included in the comparison database. Therefore, we also measure the positions and strengths of the stronger responses at slightly greater range. The measured scatterer positions must be interpreted in terms of the corner scatterers the radar would observe in this aspect angle sector, as already discussed. Again, the dominant error would likely come from wrongly estimating which




Figure 4.12 Image of the tank at a different aspect angle.

scatterers near the corners can actually be observed, and whether they should be weak or strong in terms of the average strength of the image responses, rather than from measurement error.

The intensity image of Figure 4.12 also indicates the presence of a turret. Can one extract more information about the turret from the complex image? As an illustration, consider the fact that a turret has a cylindrical base on which it turns. This is indicated by the dotted circle in Figure 4.12. One would not expect such a tight fit between turret and deck that the base would not be accessible to the radar wave. With the overhang, this large base should generate a return from the point closest to the radar, and the return should be concentrated in range, having a linear phase function. When the image responses clustered about Range Gate -1 are examined, only the response at C-1.49/R-3.42 (indicated by the arrow perpendicular to the long edge) meets this requirement. We make the practical assumption that the turret is centered on the vehicle's long symmetry axis. Because of the perpendicular incidence of the wave at the point where the base cylinder backscatters, the center of the turret must be in the same crossrange gate as the point of reflection, which is Gate -1.49 . The intersection of this gate and the long symmetry axis, determined from the extracted outline, give the center of the cylinder. The range separation of the center of the cylinder and the response with linear phase gives an estimate of the cylinder's radius. When these measurements are performed in the image of Figure 4.12, this crude estimate gives a diameter of the cylinder which is 20 cm larger than the actual value.

Analysis of the strong response at C-4.13/R-8.45 (indicated by the right arrow perpendicular to the short edge) shows that it is concentrated in range, but consists of contributions from two scatterers separated in cross-range by about half a gate. This matches the design of the left headlight. The response at C-7.10/R-7.53 (indicated by the left arrow perpendicular to the short edge) is concentrated in range, but the phase curvature in cross-range indicates a wide scatterer in this dimension. This matches the design and position of the flange of the gun barrel. The responses at C-1.26/R-0.27, C0.48/R-1.09, and C0.75/R-2.70 (indicated by vertical arrows in Figure 4.13) lie along part of a circle and have generally curved phase functions that become linear for image cuts that intersect at a single point near the turret overhang. Thus we have a set of spurious sideband responses generated by the turret overhang. Conservatively, one needs at least four responses to safely identify a set of spurious sideband responses. However, two of the three responses lie within the outline of the turret cylinder and one lies very close to the cylinder, so one might be willing to rely on just




Figure 4.13 Additional features in image of Figure 4.12.

three responses. In any case, all three would be excluded from the positional match as possible turret scatterers.

When wheel responses are subject to strong interference, as they are in the case of the image of Figure 4.12, we use two criteria for determining which responses are from wheels. First, if the wheels are cup-shaped without spokes, so that they are expected to generate observable responses, the phase functions of their responses must be curved. Second, we know that the wheels of treaded vehicles are roughly equally spaced. With the two criteria, we can detect the wheel responses as marked in Figure 4.13 by horizontal arrows. Although one wheel response is masked by interference, by assuming equal spacing of the road wheels, we correctly count six road wheels in addition to the idler and drive wheels.

4.3.2.3 Persistence of Scatterers

As with the howitzer, we again test how well the observable scatterers remain the same when the aspect angle of the vehicle is changed. In this instance we formed images at aspect angles of 19° and 45° off nose-on, and extracted the scatterer positions along the illuminated edges. With the scatterer positions obtained from the image at the 19° aspect angle serving as a database, the

match with the positions from the image at 45° is shown in Figure 4.14. The crosses thus represent the positions of the scatterers observed at 19° aspect, and the letters the positions of the scatterers observed at 45° aspect. The important match along the long edge is rather good, despite the much more complicated backscattering behavior of this vehicle as compared with the howitzer. *The comparison database for this tank thus probably can be constructed with angular sectors approaching 25° .*

4.3.3 Example 3: Measurements and Characterization for an Off-Highway Truck

4.3.3.1 The Truck Viewed From the Rear

An image of an off-highway truck, viewed from the rear, is shown in Figure 4.15. The truck is of the type generally used to transport pipe for oil pipeline construction. The truck has four articulated sections and four axles. Each section has mostly smooth sides, but also small dihedral and trihedral corners where cargo braces, handholds, and steps meet the smooth side. The




Figure 4.14 Match between the scatterer positions from the images at 19° and 45° .




Figure 4.15 Image of an off-highway truck.

truck is carrying pipe, which is secured on top by a large clamp and at the rear of the truck by a large brace. Viewed from above, the truck has a rectangular outline, plus the irregularly shaped rear cargo brace. The dotted rectangle in the figure includes the brace.

As one would expect, the backscattering from the off-highway truck is very different from that of the treaded vehicles. The truck's wheels are weak scatterers, so the edge of the image is defined by the weak corner reflectors on the side of the truck. The pipe in the truck bed reflects very little energy back to the radar, so the interior of the target contains very few sizable scatterers. Most of the strong scatterers are from the rear brace. The measurement of length and width is performed in the same manner as for the other two vehicles, by starting with the analysis of the responses along the two illuminated edges and fitting straight lines. It is more critical now to search along the edges so that the furthest responses can be found, because the responses along the edges are relatively weak for this vehicle. Since the shape of the image is not a well-defined rectangle as it was for the howitzer and tank, it is critical to *analyze image cuts perpendicular to the long edge* along the length of the vehicle in order to find the widest separation of vehicle responses, and hence the width of the vehicle.

An image cut along the long illuminated edge does not reveal peaks near the far vehicle corner that are not already indicated in the peaks plot. In defining the vehicle length we thus have the choice between the weak response at C-23.15/R26.99, indicated by the arrow at furthest range, and the stronger response about one range gate closer to the radar. We know that the stronger response must come from a feature with wave-trapping design, so that vehicle length could be defined on this basis. The weak response, on the other hand, defines the actual vehicle length, because low background conditions allow the detection of even weak responses. In this instance the weak response represents the corner of the front bumper when viewed from the rear. In any case, for a given candidate vehicle we must understand the two choices from the information in the database. In this case, when we extract the length from the complex image and compare the measurement with the actual length, we have a measurement error of 0.23m.

In those cases where the image contains a sufficient number of responses along the nonilluminated long edge, we can try to refine the width estimated from cuts perpendicular to the illuminated long edge, by fitting a straight line parallel to the illuminated edge. If the quality of the fit is good in the sense that there are not too many outliers beyond the vehicle edge, with amplitudes not significantly lower than those of the responses along the edge, we will take the separation of the lines fitted to both long edges of the vehicle as the vehicle width. In this example we obtain a measurement error of 0.37m if we compare the measurement with the overall width of the vehicle. When we examine a photograph in order to estimate which scatterers on the side will be observed by a radar, we find that the measured width is only 3 cm smaller than the actual width as seen by a radar.

We briefly point out some of the characteristic special features of this vehicle. The relative weakness of the responses along the long illuminated edge as compared with the strongest responses in the image implies that the wheels do not have wave-trapping designs; they are not cup-shaped without spokes as for some armored vehicles. The intensity image of Figure 4.15 shows the top of the truck to be largely devoid of wave-trapping features. The strong response at C-9.16/R11.79, indicated by the middle horizontal arrow, has some phase curvature. Together with the strength of the response, this indicates a sizable trihedral type corner, but not a good approximation of the ideal trihedral. The response at C-4.93/R5.07, indicated by the arrow pointing left, is of the same type. The response at C1.44/R-5.66, indicated by the horizontal arrow at nearest range, has a complicated amplitude and phase function, and thus indicates a cavity of complicated design. The response at C4.43/R-2.45, indicated by the arrow pointing up, is another

complicated cavity which also has complicated amplitude/phase properties. The planar phase function of the response at C8.25/R-1.16, indicated by the arrow pointing down, marks a good trihedral corner reflector in this position on the vehicle.

4.3.3.2 The Truck at a Different Aspect Angle

Figure 4.16 shows a front view of the off-highway truck. Even from the intensity image it is fairly easy to determine that the vehicle has a rectangular shape. The main problem for outline, length, and width measurements is to find the far vehicle-corner of the illuminated long edge in the vehicle image. As before, we analyze the responses along the illuminated long edge to determine actual scatterer positions, and fit a straight line to the measured positions along the edge. We search along this line for the response that is farthest away.

This search yields a vehicle-corner in agreement with the dotted outline. However, there are strong responses further along the vehicle. We must investigate whether the strong responses are on the vehicle. Most rectangular vehicles are designed such that few strong scatterers lie beyond the last




Figure 4.16 Off-highway truck at a different aspect angle.

scatterer of the long edge. If the strong image responses are on the vehicle, it is probably not rectangular.

We find the level of the weak spurious responses (usually delayed responses) generated by the ground vehicle in the vicinity of the nonilluminated short edge. We then examine all responses in this vicinity that exceed the level of the spurious responses by 12 dB (allowing us to use the TSA), to determine whether they have the properties of responses from a single or two scatterers. Those that do not have these properties are rejected. If at least two responses are accepted by this test, the perpendicular projection of the farthest accepted response onto the long illuminated edge is compared with the position of the farthest response detected along the illuminated long edge. The end of the vehicle is taken to be along a line perpendicular to the long edge through that one of the two responses that yields a larger vehicle length.

The primary responses of interest, because they exceed the background significantly, are the ones marked by arrows in Figure 4.16. As an example, in Figure 4.17 we show the transform of the marked response in Crossrange Gate -3.84 . This is not an acceptable amplitude/phase pattern. Similarly, the transform of the marked response in Range Gate 26.54 is shown in Figure 4.18. This is an illustration of a different type of unacceptable amplitude/phase pattern. In these illustrations we used image cuts in the range gate




Figure 4.17 Transform of the response in Crossrange Gate -3.84 .




Figure 4.18 Transform of the response in Range Gate 26.54.

or crossrange gate of a response. However, as image cuts are examined in directions other than fixed range or fixed crossrange, the strange properties of the transforms of the responses become even more indicative of spurious responses.

The response at the closest range that cannot be rejected by such analysis is the one at C-10.25/R21.46. An examination of the image cuts over 360° in the image domain does not permit labeling the response as spurious, so that it must be accepted as a genuine target response. It happens that a line perpendicular to the long illuminated edge that passes through this response also passes through the response found in the earlier search. Hence, in this instance both responses define the same far edge.

The two vehicle-corners of the near illuminated edge are found in the same manner as previously illustrated. When length and width are determined from the three vehicle-corner scatterers, the errors in both length and width are 0.15m. The vehicle width was measured too large, so that we have a measurement error rather than an error in determining the observable scatterers (which can yield only a smaller vehicle length). This vehicle has a considerable height, and in estimating which scatterers determine the observable length we must take into account that scatterers at larger heights will appear

at shorter ranges when the depression angle of the SAR beam is not zero. We also note that the analysis of the responses along the long illuminated edge shows that the responses near the end are somewhat shifted from the line defining the edge toward the center of the vehicle. In other words, one can determine that the last part of the vehicle has a smaller width, or its scatterers are at a lower height. However, these details might not be needed for successful vehicle identification.

As already mentioned, from the relatively low level of the responses along the long edge we deduce that the vehicle does not have the solid cup-shaped wheels often found on armored vehicles. This is at least an indication that the vehicle is wheeled rather than tracked. All of the features discussed above can be measured regardless of the aspect angle of the vehicle. Although there are features that can be observed only within certain aspect angle sectors, as is the case with all types of targets, the most important use of these features lies in the positional match. Analyzing the stronger of the image responses so as to determine the range/crossrange positions of the associated scatterers, and finding the wave-trapping features from photographs of the vehicle in all three coordinates, leads to the positional match shown in Figure 4.19. The few discrepancies between measured and actual scatterer positions come from the inadequacy of the available photographs of the vehicle.

We do not show positional matches between images at different aspect angles for either the off-highway truck or the flatbed truck to be examined next, because the situation is similar to that illustrated for the howitzer and the tank.

4.3.4 Example 4: Measurements and Characterization for a Flatbed Truck

4.3.4.1 The Flatbed Truck Viewed at 38°

An image of a flatbed truck is given in Figure 4.20, with the truck's outline given by the dotted rectangle. By comparison with the images of the first three vehicles, the outline of the truck is ill defined in the peaks plot image. However, we again note that a response peak often comes from the contributions from two scatterers, and that the analysis of the complex responses along the edge will generally give much better defined edges than will a peaks plot image. In this instance, when the edges are determined in the manner explained for the howitzer, we obtain the solid lines shown in Figure 4.20. Since the nonilluminated edge is defined by only a few responses, the measurement of the vehicle's width cannot be expected to be very accurate.




Figure 4.19 Positional match for the off-highway truck.

In order to compare these measurements with the actual vehicle data, we must determine the effective scatterers observed by the radar along the edges. At least for the short edges these scatterers depend on the orientation of the vehicle. Although the orientation of a stationary vehicle can be determined from the image, it is more accurately found from the positional match illustrated below, so that for a vehicle with a good positional match we can also measure length and width more accurately. When a photograph of the vehicle is examined to determine where the wave-trapping features seen by the radar are located, the comparison with the length and width derived from the measured outline shows a length error of 0.10m and a width error of 0.16m.

This positional match uses only the positions of the scatterers derived via the TSA from the 15 strongest responses of the image. Since no turret is indicated in the image, the scatterers to be used for the match need not




Figure 4.20 Image of a flatbed truck.

be restricted to those close to the edges. Photographs of the vehicle were analyzed to find the scatterers that should be observable, and their three-dimensional positions on the vehicle were determined. The resulting positional match is shown in Figure 4.21. Three measurements are not matched by actual scatterer positions, and one scatterer position is not matched by a measurement; otherwise the positional match is excellent. The response corresponding to feature A can be found by analyzing the image, but it is so poorly resolved from a much stronger nearby response that it probably could not be measured in an automated system. The missing scatterers for the three unmatched responses can be explained by the lack of appropriate views in the photographs available to us.

4.3.4. The Flatbed Truck at a Different Aspect Angle

The image of the flatbed truck at a different aspect angle is shown in Figure 4.22. Among the images shown in this chapter, it has the most severe problem with spurious responses. The more significant of these responses, recognized from the amplitude/phase patterns in image cuts through their peaks, are marked by arrows. They will be briefly discussed in order to indicate the types of spurious response to be expected with ground vehicles.




Figure 4.21 Positional match for the flatbed truck.




Figure 4.22 Image of the flatbed truck at a different aspect angle.

Response 1 is one of the many delayed responses that can be observed in the image of a ground vehicle, even though weaker than the usable genuine responses. Responses of this kind are generated, for example, when the radar wave travels along a dihedral and is reflected at the end. If good photographs of the vehicle are available, such structures can be readily identified. In principle, the response can be utilized to detect the presence and to determine the location of a specific target feature. However, with automated identification this might be too difficult, in which case the policy should be to identify and ignore spurious responses. This is easy in the present instance, because the delay is large enough to place the response outside the boundaries of the vehicle. This policy will rarely dismiss genuine responses, as gun barrels oriented away from the radar usually generate only weak returns that cannot be detected in the background anyway.

Responses 2 and 3 come from rough areas on the ground. Ground responses can sometimes be strong enough to interfere with the determination of the vehicle outline, so that a mechanism is needed to distinguish a ground response from a vehicle response. This is indeed possible, because the backscattering from the ground is not dominated by wave-trapping features, as it is for the vehicle. The analysis of both Responses 2 and 3 shows amplitude/phase patterns so unusual that we cannot use the analysis methods developed for man-made targets. When a response of this kind is encountered, it should be ignored.

Responses 4 through 7, which are spurious responses generated by large cavities (again, recognized from the patterns of the transforms of image cuts), are more problematic. Since such responses can occur far away from the responsible cavity, they could degrade identification performance, primarily by simulating scatterers in positions where a particular vehicle has none. As explained in detail in [1], in principle these sideband responses occur in sets arranged along some smooth arc, with generally curved phase functions but with image cuts in which the phases are linear all meeting in one point. The responses are identifiable in principle, but in practice the responses of such a set may be poorly resolved and may be subject to interference from other responses, genuine or spurious, so that the direction of phase linearity may not be easily measurable. There are indications, not yet confirmed, that such a spurious response might be more easily discriminated by measuring the width of the angular sector in the image domain over which the phase function is essentially straight. This sector appears to be much sharper for a spurious sideband response than for a genuine response. We again note that such refinements should not be necessary for successful vehicle identification.

4.3.5 Assisting the Positional Match

In our earlier discussions we emphasized the positional match as an important tool for target identification, utilizing those scatterers for which we can *only measure positions*. With a target such as an aircraft, there are indeed many features for which we can extract no more than positions. For example, when utilizing the positions of a number of antennas on the fuselage, it is hardly possible to also determine the types of antenna. The situation is different for ground vehicles, which by their designs have a larger variety of features. We can greatly improve the contribution of the positional match to target identification *if we also utilize scatterer characteristics in the matching process*. If we are able to associate an image response with a specific scatterer, we can constrain the positional match and make it more effective. As an example, if we can identify the response from the first wheel of the vehicle, we can fix that feature in each template we match to the measurements.

More generally, if we can extract specific characteristics from any of the responses, we can compare this information with the characteristics of the associated scatterer. If the response has a planar phase function, for example, and the feature matched to the response is a trihedral corner, the correctness of the association is enhanced. On the other hand, if the feature is an extended cavity, the match to a response with a planar phase function would be penalized. We will expand on the topic in the following, using the last image of the truck, Figure 4.22. We analyzed the major responses in this image, rejecting the spurious responses as discussed above, and measuring the positions of the scatterers associated with the genuine responses. The scatterers were identified in photographs of the vehicle. The template match between the measured positions of these scatterers and the scatterer positions extracted from the image is shown in Figure 4.23. We will now examine the characteristics of various responses and relate them to the design of the associated vehicle features.

A: Flatbed Corner

Feature A of the truck is the trihedral corner at the far end of the truck bed. The phase function of the ideal trihedral is planar, and the deviation of the actual trihedral from the ideal can be determined from the curvature of the phase function, which gives the effective width and depth of the trihedral. To perform the measurement of width in crossrange, we take the transform of a fixed-range image cut through the response and make the measurement in accordance with Figure 4.2, always measuring the phase slope before any rapid change that might occur near an amplitude minimum or a sharp drop




Figure 4.23 Positional match for the flatbed truck.

in the amplitude. In this instance we measure an effective crossrange width of 6 cm, which is relatively small considering the size of the feature and the strength of the response. When we take the transform of the fixed-crossrange image cut through the response, so as to determine the range extent of the feature, we find that the phase function has an inflection point rather than being monotonically curved. This means that the range extent is too small to be measured. The two measurements imply that the feature is a good approximation of an ideal trihedral.

B: Brace of the Flatbed Sidewall

There appears to be a trihedral reflector formed at the base of the brace on the outer sidewall of the truck bed, so that the phase function of the response should be planar. An examination of the response shows that the phase function is indeed linear for an image cut in any direction in the image, signifying a planar phase function.

C: Headlight

The headlight is within a metallic enclosure, and thus acts as a cavity reflector. If the shifting of the phase center due to the changes in the aspect angle or frequency is relatively weak, the phase function of the response is curved. We then can perform extent measurements of the type already illustrated.

On the other hand, as the shifting of the phase center becomes larger, the response starts to spread into a set of sideband responses, perhaps not fully resolved. The measurement algorithm used in Figure 4.2 then does not work. If the analysis in the range gate of the response results in only two scatterer positions, but if the two scatterers are measured to have nearly the same ranges, the response is assumed to come from a single cavity. This is the situation in the present instance. Thus the predicted feature of a headlight and the measured response characteristics agree. We can take the difference between two measured scatterer positions as the feature extent.

D: Vertical Exhaust Pipe Next to the Cabin

The exhaust pipe by itself does not backscatter sufficiently strongly to be observed by the radar. For this aspect angle, the hole at the end of the pipe is turned away from the radar, so that it cannot be seen. However, the attachment of the pipe to the edge of the driver's cabin forms an irregular corner, which would be expected to have a shifting phase center. The measurement agrees with these postulates. The phase center is slightly shifting as a function of aspect angle, and produces a curved phase function. With the same method as used in Figure 4.2, we find an effective crossrange width of 16 cm. This appears reasonable in view of the specific design of the attachment. The phase curvature in range is very small, which means that the range extent of the feature is small. This likewise agrees with the design of the feature, which has the appearance of an irregular vertically oriented dihedral.

E: Inside Trihedral Corner

A trihedral corner is formed by the underside of the chassis, a vertically oriented plate along the length of the vehicle, and a vertically oriented metallic mudguard. The degree to which this feature approximates an ideal trihedral can be determined by measuring the phase linearity of the response in range and crossrange. However, there is also a metallic box nearby that introduces a complication. Measurements give a crossrange width of about 8 cm and an unmeasurably small range depth, which imply that the feature is effectively nearly an ideal trihedral.

F: Corner of the Bumper

This reflector is essentially a point scatterer, so that its response is barely strong enough to be observable even in the low-level background. If the interference allows measurement of the phase function with adequate accuracy, it should be planar. Measurements on image cuts spread over 360° in angle in the image domain show that the phase function can be measured

despite the weakness of the response, and that the phase function is indeed planar.

G: Headlight

This is similar to the first headlight, but in that case the radar sees the adjoining hood, whereas in this case it sees only the enclosure of the bulb. The analysis shows the same type of response, except that the crossrange spread is not as severe as it is for the other headlight.

H: Light

This is a small light, perhaps a turn indicator, but it is also within a metallic enclosure. Thus one would expect the behavior of a cavity-type reflector, but because of its small size the phase function of the response should be only slightly curved, with much smaller spreading of the response than in the case of a headlight. If this is the case, then we should be able to measure the size of the feature. We can indeed perform the type of measurement indicated in Figure 4.2. This measurement gives a crossrange width of 15 cm, consistent with the actual width of the feature when viewed from the front. The range extent of the feature is not measurable in this case.

I: Corner/Cavity

This feature is formed by a metallic, vertically oriented mudguard and a metallic box that leaves a small space between the box and the mudguard. The feature thus acts as a cavity. Viewed from the side of the vehicle, the cavity is relatively narrow, but it extends substantially into the vehicle, so the measurement should indicate a cavity of small crossrange extent but large range extent. Examination of the image cut in the range gate of the feature shows that the crossrange width of the cavity is unmeasurably small. The range depth, on the other hand, is found to be 51 cm, which is in agreement with the physical appearance of the feature.

J: Corner Formed by a Metallic Box and the Step to the Cabin

This feature forms a trihedral corner, so that a planar phase function should be expected. However, the interference is so heavy within an angular sector of several tens of degrees about the fixed-crossrange image cut that the phase curvature cannot be measured. Within the remainder of the angular sector the phase function is linear. Where the measurement can be performed, it thus agrees with the expectations.

K: Outside Mirror

The available photographs do not show sufficient detail to evaluate the expected backscattering characteristics of this mirror. In such situations only the scatterer position can be used.

L: Brace on the Flatbed Enclosure

The bottom of this brace acts as a small trihedral, so that the response should have a planar phase function. However, there is a metallic ring close by, which might influence the backscattering in a way that cannot be predicted at this time. For this feature we thus would not attempt to predict the backscattering characteristics. The analysis of the response shows that extent measurements cannot be performed because of interference.

M: Trihedral Corner Formed by Mudguard and Chassis

This feature is only a poor approximation of a trihedral corner, because the angle of one of the sides with respect to the other is much less than 90° . Accordingly, we would expect a feature with a strongly shifting phase center and a significantly curved phase function. This is confirmed by an analysis of the response. Not only is the phase function curved, but the amplitude function is more indicative of two scatterers than one.

N: Inside Fender Corner

This "corner" has a very complicated shape, for which we can merely predict that the amplitude and phase responses will be complicated. This is indeed verified. The response is so complicated that, at least at the time of this writing, we could not measure extents.

O: Inside Fender Feature

This feature is an extension of feature N. Again, it is so complicated that all we can predict is a complicated response. This is again verified by the actual response.

P, Q: Rear Wheels

The rear wheels have deep hubs. However, the hubs have such large diameters that they act as concave features whose effective reflection points are where the incidence of the beam is perpendicular. This point will shift slightly with aspect angle. Thus we would expect the phase curvature to be small, and this is verified from the two responses.

4.3.6 Section Summary

Stationary vehicles are the only targets of practical interest that allow a very accurate measurement of length and width, provided the quantities are defined by the wave-trapping features nearest the corners of the vehicles. They are also the only targets that always allow the use of high crossrange resolution, and the measurement of feature characteristics when resolution is sufficient.

An accurate target outline is the most important feature in the identification of stationary ground vehicles. This is supplemented by special features, such as the presence of a turret and the number and type of wheels, and the positional match of nonrecognizable features. The section describes how to measure these features.

Measurable features persist over aspect sectors on the order of 25° . The utility of nonrecognizable features can be greatly enhanced by comparing measured scatterer characteristics as well as positions to those of the predictive database.

4.4 Identification of Moving Ground Vehicles

This section is the heart of Chapter 4, treating the most difficult radar identification problem. Although the same processing steps are used regardless of the type of vehicle and its behavior, they must be used in a highly adaptive manner. Also, the measurements that can be performed on a moving ground vehicle depend very much on the type of vehicle and its behavior. A reader with a casual interest in the identification of moving ground vehicles can obtain a good understanding of the problem by reading the text without following the details illustrated in the figures.

This chapter is peculiar in that, among the targets considered in this book, it treats both the most easily identified (stationary ground vehicles) and the most difficult to identify (moving ground vehicles). The variety of conditions under which moving ground vehicles must be identified necessitates a complicated image-formation process. However, we wish to emphasize that *the same two processing and analysis steps are used under all conditions*: estimating an appropriate imaging duration and determining intervals when individual scatterers can be well imaged. The complications arise when the longest interval during which most scatterers can be well imaged is shorter than the appropriate imaging duration. Then we must investigate whether we can improve the situation by varying the motion compensation or by utilizing only part of the available dwell.

The treatment of moving ground vehicles must be quite different from that of stationary ground vehicles. With stationary ground vehicles, we obtain SAR images of high quality (even if the radar is not a surveillance system), so that the only task is to extract as much information as possible from the images. The variations from one vehicle to the next are insignificant. With a moving ground vehicle, there is the additional problem of generating a usable image. The quality of such an image will depend heavily on the type of vehicle and its behavior. Since we have given many examples of extracting information from images of stationary ground vehicles, for moving ground vehicles we will concentrate on the image-formation process and the differences in the image analysis procedures when good images cannot be formed. When we give a moving vehicle the same designation we earlier gave a stationary vehicle, we imply that it is the same vehicle observed when it is moving rather than stationary.

4.4.1 Peculiarities of Moving Ground Vehicle Identification

It is simpler to identify stationary ground vehicles than moving ones because the steady platform motion of a SAR system, rather than an erratic target motion, provides crossrange resolution for the former. This allows the same analysis and identification procedure to be used on each stationary vehicle image, regardless of the type of vehicle. The SAR system also provides images for moving ground vehicles, but these images are highly smeared because of the vehicle motions. The SAR system compensates only the motion of the platform, with the ground vehicle motion remaining. Depending on the motion characteristics and, to an important degree, also on the type of vehicle, each image may have to be formed and analyzed differently. The decision of how to proceed with the processing must be based on the analysis of the smeared vehicle image provided by the SAR system. Thus, fully adaptive procedures are needed. Since a particular vehicle may be going slowly or rapidly, straight or in a circle, and on a good road, a poor road, or off the road, the requirements on the adaptivity of the processing procedure are extraordinary.

In one particular respect, moving ground vehicles represent a special situation that requires a treatment different from that for other moving targets, such as aircraft or ships. With the latter types of target, one can either perform a good motion compensation or selectively reduce the imaging interval to such a small duration that a simple motion compensation is adequate. In these applications we can realize the inherent Doppler resolution of radar, implemented with the use of the TSA. With moving ground vehicles, we often have a situation where an adequate motion compensation is not

possible. This means that Doppler resolution (which requires constant-Doppler motions) cannot be sufficiently well implemented, and the TSA becomes inapplicable. We still may be able to resolve scatterers in crossrange and to measure the crossrange positions of the scatterers, but not by such “automatic” procedures as FFT processing and application of the TSA. Instead, we must take transforms in the range gates of responses and analyze the amplitude and phase patterns in a more basic manner, interpreting the phase function in intervals where the amplitude happens to be reasonably constant. These procedures will be illustrated where applicable in our examples. The point is that the processing and analysis for moving ground vehicles must necessarily be much more varied than for stationary ground vehicles.

The data used in all our examples were collected by SAR systems. Thus, the SAR processor compensates the motion of the radar platform and forms an image with the standard motion compensation for stationary ground vehicles. If the SAR scene contains moving ground vehicles, their motions with respect to the ground will not have been compensated. The constant range rate component of a vehicle’s motion causes the image to be shifted in crossrange, and a changing range rate smears the image, as does the varying range rate component of the vehicle’s motion about its center of gravity. The situation is much the same as if a stationary radar observed the moving ground vehicle, with the image formed without a motion compensation. The procedures developed in the following thus work equally well with any radar that tracks the vehicle over a time long enough to achieve the desired or obtainable crossrange resolution. Restricting the treatment to ground vehicles in SAR scenes does not narrow the scope of applicability of the processing techniques.

4.4.2 Consequences of Different Types of Motion

Because of the large variability of conditions under which moving ground vehicles must be imaged, we start with a description of the effects that different types of motion have on imaging and identification. The later sections will present illustrations of representative cases.

The type of information that can be extracted from an image varies from case to case. This information can be broadly categorized into three classes: (1) length and width of the vehicle, (2) scatterer positions for the positional match, and (3) special features, including scatterer characteristics. In the illustrative examples we present in the following sections, we always use a single motion compensation over an entire vehicle. As will be discussed later in more detail, there is at least the theoretical possibility of extracting

more features by varying the motion compensation over the vehicle. We will consider this topic because it might be practical when the radar uses a higher range resolution than in our examples, although we doubt this. For the situations represented by our examples, varying the compensation over the target does not appear practical, at least not for an automated system.

In the best case, we can perform all types of measurement on a moving ground vehicle that can be obtained on a stationary ground vehicle. The quality of these measurements may not be as high, but they are good enough to use in the identification process. Specifically, we can form an image good enough to allow measuring length and width of the vehicle, scatterer positions, and special features. This case arises primarily when a vehicle is turning, so that the turning Doppler dominates over the Dopplers generated by motions about the center of gravity and bouncing. The combination of the vehicle's design, the surface on which it moves, and its speed determine when this case exists. For example, a heavy and rigid armored vehicle will more likely fit this case than a lightweight vehicle. Since it rarely will be possible to measure the aspect angle of the vehicle by tracking it over a longer period, the basic requirements for a length and width measurement are that the vehicle have a rectangular outline and that both the long and the short illuminated edges can be measured accurately enough to determine the crossrange scale. This means that the width measurement will become too inaccurate for 0° and 180° aspect angles, and the length measurement for the broadside aspect.

The measurement of the short illuminated edge will be the first measurement to fail when the situation deteriorates. Except in the special case of a small aspect angle, where the vehicle length can still be measured, and the special case of near broadside, where the (less important) vehicle width might be measurable, we cannot determine vehicle length and width when only one vehicle edge is well defined. This is a very important shortcoming. We still may be able to recognize the presence of a turret, count the wheels, or measure scatterer characteristics. In general, however, the two-dimensional positional match will be most important for identification. Here the unknown crossrange scale is accommodated by stretching each template in crossrange until it gives the best match. This case occurs when the turning or translational motion does not generate sufficiently high Dopplers to dominate over the Dopplers from the irregular components of the motion about the center of gravity or from bouncing. Crossrange resolution then often cannot be made high enough for a sufficiently accurate definition of the short illuminated edge.

As the situation further deteriorates, measuring special vehicle features becomes progressively more difficult. Shadowing by a turret may not be

recognizable, for example, and the quality of the image responses may not be good enough to allow measuring scatterer characteristics. This case is likely to arise when the vehicle is moving along a straight line, and significant irregular motion about the center of gravity and bouncing occur. Again, much depends on the construction of the vehicle.

The motion of a vehicle can easily become so complicated that bending, bouncing, and vibration dominate. Essentially all that can be measured under these conditions are the range positions of the scatterers, usually of the scatterers along the long illuminated edge. The motion of the vehicle may still allow resolving the scatterers in crossrange, but the crossrange accuracy will be so poor that crossrange positions have little utility for vehicle identification. Note, however, that *this is still far better than using an ordinary range profile*, which gives the interference patterns from scatterers in the same range gates. This case will typically arise when vehicles move along straight paths on uneven surfaces, in particular at higher speeds and for lightweight vehicles.

4.4.3 General Processing Procedures

For orientation purposes, we now give a simplified summary of the processing steps that extract the information needed for vehicle identification from the radar return. The assumption is that the smeared image of the moving vehicle has been detected in the SAR scene, which may require a SAR system with clutter cancellation. In the case of SAR surveillance radar, the available imaging time will typically be much larger than needed, or usable, to image the moving vehicle. A great deal of adaptivity is required in the selection of an appropriate imaging interval. The feature extraction procedures are the same as for images of stationary ground vehicles, but the quality of the obtainable information will generally be much lower. Also, the type of information that can be extracted depends on which motion condition exists. The discussions in the following sections explain the processing steps in creating an image of a moving ground vehicle and the underlying reasons for these steps, with illustrations from representative cases. At the end of this chapter we give a brief summary of how the processing steps may be automated.

4.4.3.1 Step 1: Excise the Smeared Image of the Vehicle from the SAR Scene

Select the range gates that cover the smeared image of the vehicle. Also, select the crossrange gates that cover the vehicle image, making sure that even the lower-level returns of the smeared responses are included in the window. Choosing the range window too large merely affects the processing load. The

crossrange window should not be larger than necessary because unnecessary clutter might be included, even in the case of a system with clutter cancellation (which does not work perfectly). This is a simple processing step.

4.4.3.2 Step 2: Transform the Area Excised from the SAR Scene into the Raw Data

An FFT over crossrange of the excised data produces the sequence of range profiles over the selected range window, and over the observation time of the SAR system (with the platform motion removed). A subsequent FFT over range yields pulse-frequency data. This is another simple processing step.

4.4.3.3 Step 3: Form a Survey Image for Analysis of the Vehicle's Behavior

A standard motion compensation is applied to the excised data (range compensation based on the range centroid of the entire vehicle or a dominant scatterer, followed by Doppler compensation of the Doppler centroid of the entire vehicle or a dominant scatterer). The resulting survey image will essentially be the same as if the moving ground vehicle had been observed by a stationary radar. This is the same type of processing step as used with aircraft, except that it is only for survey purposes.

4.4.3.4 Step 4: Analyze the Survey Image and Form the Final Image

This step requires highly adaptive procedures. Ideally, one should form the survey image over the entire available observation time, estimate crossrange resolution by comparing the crossrange spread of the vehicle image with the typical size of a ground vehicle, and then form a new image over a shorter *time interval chosen to give the desired crossrange resolution (not more than needed) during the target's smoothest motion*. In practice, this requires searching for responses in the full image for which Fourier transforms of fixed-range image cuts have sufficiently constant amplitude functions to make their phase functions meaningful measures of scatterer motion, or for which the transforms correspond to two interfering dominant responses. Then one can select the subinterval with the smoothest motion. Such transforms often cannot be found in the full image. Then it is necessary to shorten the imaging interval, try finding the appropriate transforms, and iterate until one succeeds. Even when one obtains an image that permits the analysis of responses, it still may be desirable to shorten the imaging interval further, so as to select an interval in which the amplitude is very constant and the motion is very smooth. We will refer to an image formed over such an interval as the "final image."

Forming a usable image is particularly difficult for wheeled vehicles, which have a lighter construction than treaded vehicles. If the vehicle is flexing and vibrating, scatterers close to each other in the image may have significantly different motions (e.g., scatterer at different heights). Under these conditions a motion compensation cannot be performed over a sufficiently long time that the crossrange positions of scatterers can be measured with useful accuracies. Nevertheless, crossrange resolution even then is still useful because it separates responses in the same range gates, allowing higher range accuracy.

This processing step, the formation of a usable image, is difficult and very involved. It can be fully understood only through examples that illustrate the various difficulties and their solutions, and we provide these examples in much of the remainder of this chapter. The significance of these examples is difficult to grasp without a prior explanation of the issues and their solutions, yet such an explanation cannot be appreciated without the examples. For this reason, we start by giving an overview of a flowchart that describes the various processing operations, next present the illustrations (Sections 4.4.4 through 4.4.7), and then discuss the flowchart in detail (Section 4.4.8).

Figures 4.24, 4.25, and 4.26 are three versions of a flowchart that gives an overview of the processing step. Figure 4.24 shows the core operations of the processing, outlined in bold, as well as the image formation algorithms for the various types of motion. *The core operations are the estimation of the image duration required if we are to measure crossrange positions with an accuracy that is useful for identification, and the determination of time intervals during which scatterers can be well compensated.* These operations allow us to select an imaging interval of appropriate crossrange resolution and as smooth a motion as possible. The examples of Sections 4.4.4 through 4.4.7 concentrate on these core operations.

Figure 4.24 shows how the appropriate imaging interval depends on the motion conditions. If the dwell is very short or if the target and the platform are on a near-collision course, so that only low crossrange resolution is achievable, we cannot afford to reduce the imaging interval. Fortunately, motion compensation is straightforward under these conditions, so using the entire dwell is possible. When the target moves smoothly in a straight line or in a slow turn, some crossrange resolution is attainable and the translational Doppler changes generally dominate those from irregular motion. For these conditions, we choose an imaging interval when the irregular motion is small, so that a good compensation is possible. If the target moves smoothly in a rapid turn, we can achieve excessive nominal crossrange resolution.

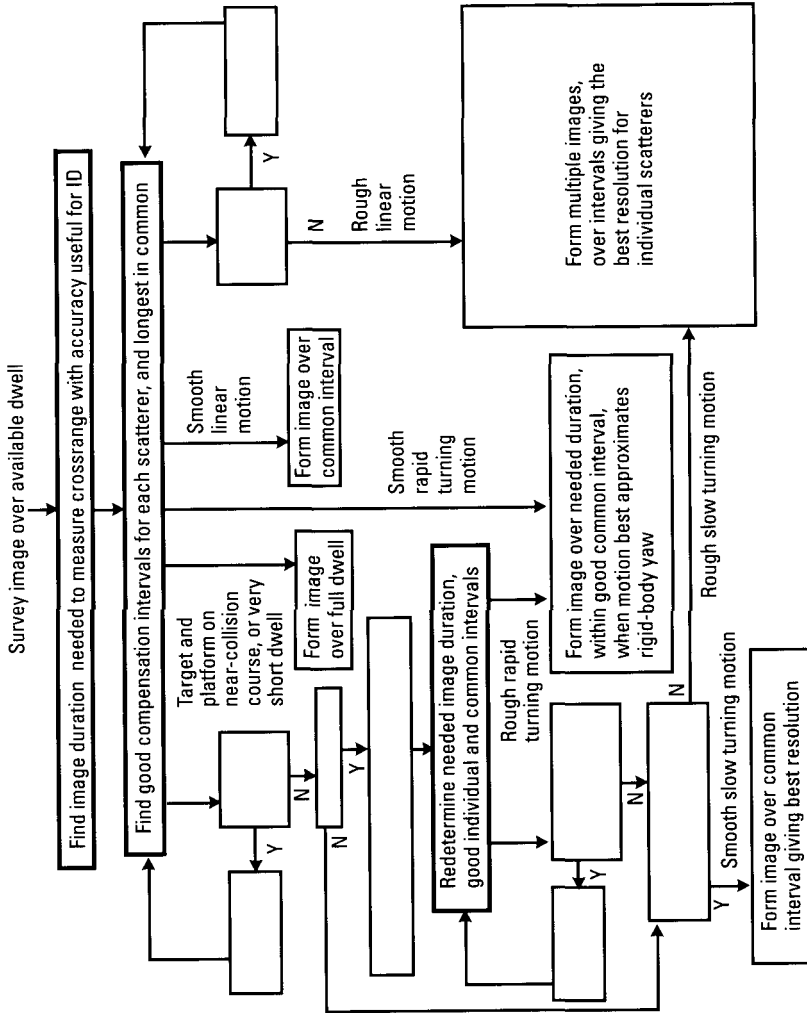


Figure 4.24 Core operations of image interval selection.

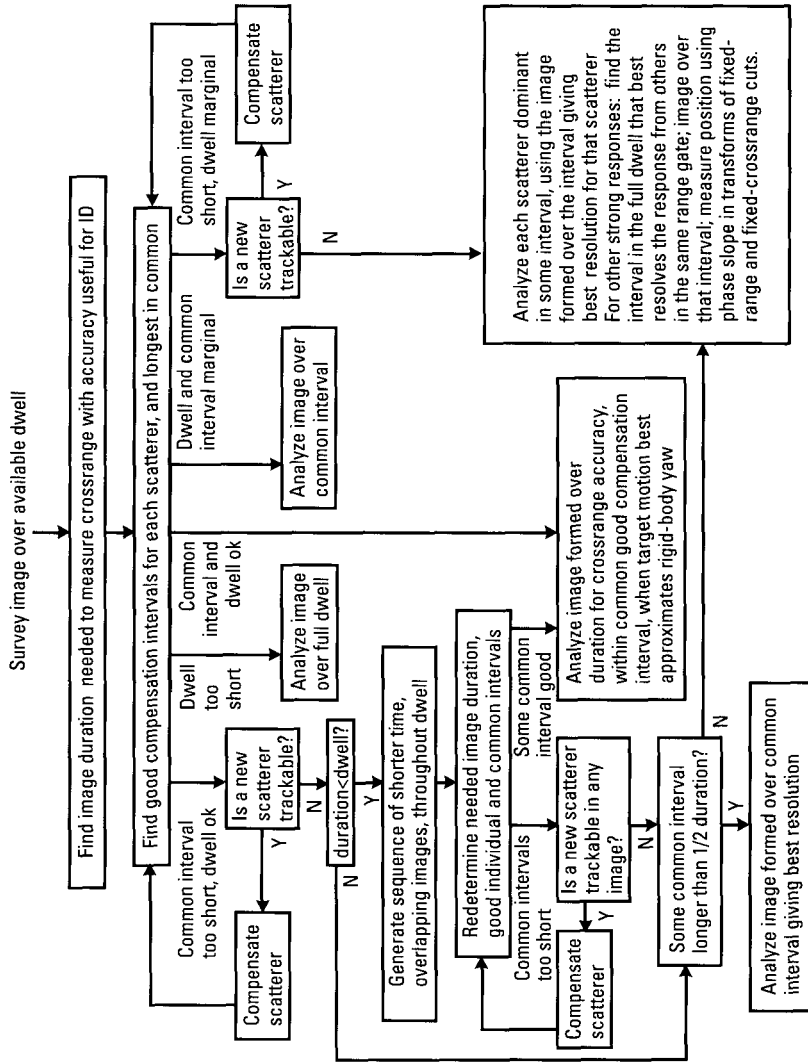


Figure 4.25 Flowchart for image interval selection.

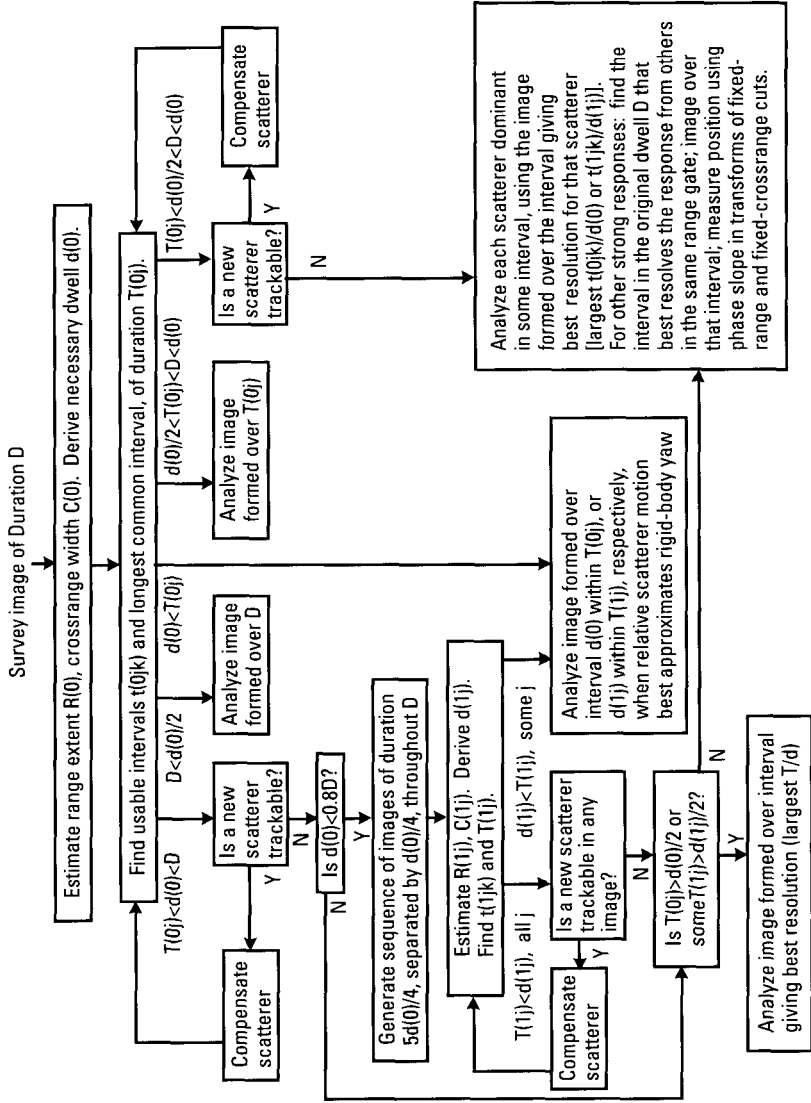


Figure 4.26 Flowchart with branching criteria.

However, even in a smooth turn, the target may be bending, flexing, or rolling. These motions complicate the interpretation of scatterer Dopplers, so we image over an interval when rigid-body yaw motion dominates and gives sufficient crossrange resolution.

When the target moves roughly, Doppler changes caused by irregular motion generally dominate those from translational motion. If the translation is linear or a slow turn, the differing motions of the scatterers, particularly scatterers within the same range gate, typically prevent the selection of a single imaging interval that gives good crossrange resolution for all the scatterers of the target. We must measure each scatterer's position in an image that gives good resolution for that scatterer. If the translation is a rapid turn, only a short interval is needed to achieve sufficient crossrange resolution. Despite the rough motion, there will often be a surplus of usable imaging time. When this is the case, we image when rigid-body yaw motion dominates.

The path followed through the flowchart is determined by the motion conditions, which are not directly measurable. However, they manifest themselves in the relative durations of the available dwell, the imaging duration necessary to measure crossrange positions with usable accuracy for identification, and the longest interval allowing a majority of strong scatterers to be well compensated. Figure 4.25 shows the flowchart in terms of these relative durations. If the translational Doppler changes are comparable to the Doppler changes from irregular motion, as is often the case for slow turns and for rough rapid turns, determining when individual scatterers can be well compensated may require us to recompensate data and examine images of varying duration, as we did with aircraft (see Figure 3.17). As shown in Figure 4.25, recompensation accounts for the three loops of the flowchart, and the use of images of varied duration accounts for the remaining complexity.

Figure 4.26 provides a more mathematical description of the branching through the flowchart. We postpone a detailed explanation of the flowchart until Section 4.4.8, because some of the rationale behind various operations and branching criteria will be better appreciated after reading Sections 4.4.4 through 4.4.7, which give illustrations and more extensive discussions of the operations involved, for four different vehicles. We also note that the factors in Figure 4.26 ($1/2$, $1/4$, $5/4$) are approximate; we do not mean to imply that precisely these values would be used in an operational system.

After creating an image, or images, with the procedure summarized in the flowchart, we must perform some additional checks before measuring scatterer positions. These checks require that we examine fixed-range cuts through strong responses, as we did when selecting the imaging interval. We

verify that scatterers can be well compensated over the imaging interval, and that nominal crossrange resolution is not excessive. If either condition does not hold, we reduce the imaging interval in the same manner as previously. In addition to the verification, we compare the scatterer motions derived from phase functions corresponding to transform intervals with constant amplitude.

If the differences in the scatterer motions scale with the scatterer separations, the target motion is dominated by rigid-body rotation. In this case, we can improve the image quality by compensating one of the rigidly rotating scatterers and then, if the differential motion is irregular enough or large enough, resampling or polar reformatting the data, respectively. Given typical ground vehicle sizes, the resampling and reformatting steps will rarely be necessary, because we have already restricted the imaging interval to one of smooth rotation through a small angle. As a final compensation step, we remove any quadratic phase variation with time common to the strong scatterers [1].

If the target motion is dominated by rigid-body rotation, we measure scatterer positions by applying the TSA, as we did with stationary targets. If not, individual responses may be poorly compensated, so that the TSA is inapplicable, and we instead determine positions by measuring phase slopes corresponding to intervals of constant transform amplitude in fixed-range and fixed-crossrange cuts through each response. In either case, we attempt to extract length, width, and special features as we did for stationary targets.

4.4.4 A Moving Off-Highway Truck

4.4.4.1 The Off-Highway Truck Moving in a Slow Circle on Flat Terrain

After detecting and excising the smeared image from the SAR scene and transforming the image data into the raw data, we arbitrarily select an imaging interval of 2.1 seconds and apply the standard motion compensation. The resulting image of the off-highway truck is shown in Figure 4.27. It is a fairly good image for a moving vehicle with only the standard motion compensation, the reason being that the relatively smooth turning motion dominates over the irregular components of the motion about the center of gravity. This survey image is the starting point for the image-formation step.

First, we decide how long a data segment is desirable for imaging. The imaging interval should be long enough to provide adequate crossrange resolution, which depends on the size and orientation of the vehicle. The range of sizes of the vehicles of interest is known, and from experience we also know that the long illuminated vehicle edge should be divided into at least




Figure 4.27 Survey image for off-highway truck in a circle.

20 to 30 resolution cells for reliable identification. In the case of the survey image of Figure 4.27, we reason as follows. The vehicle extends over about 50 range gates, which is such a large number that crude crossrange resolution should be satisfactory. We must estimate over how many crossrange cells the vehicle image is spread for the 2.1-second imaging interval.

The smeared responses appear as a series of dots in the various range gates. From the distribution of the response centers in crossrange, we can estimate the crossrange width (in each range gate) of the vehicle. For example, the response in Range Gate -9 , along the lowest dotted horizontal line in Figure 4.27, is centered in about Crossrange Gate -20 . The response in Range Gate -7 , along the next lowest dotted horizontal line, is centered in about Crossrange Gate 20 . The responses in Range Gate -3 , along the unlabeled dotted horizontal line, are also centered in about Crossrange Gates -20 and 20 . Thus, we very crudely estimate a crossrange width of 40 gates. As we demonstrate below, this estimation can be done more accurately by examining fixed-range image cuts. Assuming a practical vehicle width of 3m, the crossrange width implies that a good motion compensation would give a crossrange resolution of 7.5 cm. However, with the vehicle extending over

50 range cells, a crossrange resolution of even 0.5m (six times as wide as nominally implemented in the survey image) should be adequate. Thus we want to select a data segment roughly one-sixth as long as the one used for Figure 4.27. Note again that we do not want to implement a crossrange resolution higher than needed, even if a long observation time is available, because of the problems associated with high crossrange resolution.

In reducing the imaging interval, we want to *select the best available data segment*, with the amplitude functions of scatterers in different range gates as constant as possible and the phase functions as smooth as possible. To find a suitable data segment, we use the procedure common to all imaging and discussed many times before: Investigate fixed-range image cuts through strong or isolated responses to find transforms with sufficiently constant amplitude functions. However, because responses on moving ground vehicles are often smeared over many crossrange gates, finding acceptable transforms is more difficult than for other targets.

We start with the range gate containing the strongest response, or Gate -9.31. The image cut in this range gate is shown in Figure 4.28. A poor motion compensation smears individual responses in crossrange, producing bell-shaped response envelopes. We can use the centers of these envelopes to




Figure 4.28 Image cut in Range Gate -9.31

obtain a better estimate of the target's crossrange width than we can from the peaks plot. Only one envelope is evident in Figure 4.28, centered in about Crossrange Gate -15 . In general, the next step is to take the transform over a group of response peaks within a bell-shaped envelope (we hope that this is the smeared response from a single scatterer, but this must be verified via the transform). This is easy in the case of Figure 4.28, because there is only one such group. Thus we take the transform over the entire group, as shown in the right side of Figure 4.28. If the group of response peaks indeed represented the smeared return from one dominant scatterer, the amplitude of the transform would be essentially constant. The strong, nearly periodic amplitude modulation in Figure 4.28 implies that two scatterers of comparable strengths are the primary contributors to the smeared response.

The two scatterers in this range gate are unresolved. Thus, with the current compensation, they are not usable for selecting a good imaging interval. Should we have to recompensate the data because we cannot determine a good imaging interval from other range gates, the scatterers may become resolved and the range gate may prove usable for interval selection. In fact, if recompensation turns out to be necessary, and if we cannot find a range gate with a single dominant scatterer, we can apply the procedure for phase-slope tracking the stronger of two dominant scatterers to this gate, and use the resulting function to recompensate the data.

Although the amplitude modulation implies that we cannot currently use the range gate for imaging interval selection, *we can use the modulation pattern to measure the crossrange separation of the two scatterers*. The transform amplitude function is unaffected by uncompensated target translation. Therefore, as we explained in connection with the TSA, the period of the amplitude modulation is the reciprocal of the scatterer separation. There are about nine cycles of the dominant modulation in the transform amplitude. Thus, the two scatterers are separated by nine crossrange resolution cells. In cases where the period is less obvious, it can be obtained from an FFT of the transform amplitude, with the transform phase set to zero (see Section 4.4.9 and Appendix C).

In order to estimate the target's crossrange width and to select a good imaging interval, we analyze other range gates in the same manner as we did Figure 4.28, searching for gates that contain single dominant scatterers or two scatterers that are more widely separated than those of Figure 4.28. Despite the fact that the responses from individual scatterers are highly smeared in the survey image, and that responses from scatterers in the same range cell may overlap, it is not critical to recognize how best to choose the boundaries of the transform window in order to obtain a transform with

essentially constant amplitude. If there is uncertainty with regard to the choice of the boundaries, these boundaries can be varied until the best result is achieved. This is simple in an automated system.

One of the better image cuts through the survey image of Figure 4.27, shown in Figure 4.29, is in Range Gate -7.09 . The cut shows two bell-shaped envelopes, one centered near Crossrange Gate 5 and the other near Crossrange Gate -35 . The transform over the entire displayed interval shows a fast modulation with 35 to 40 cycles, corresponding to the separation of the two envelopes. This confirms the crossrange width estimated from Figure 4.27. The transform also shows a slower time-varying variation with about seven cycles, corresponding to two scatterers within the stronger envelope. The transform over the stronger envelope, between the vertical cross-hairs of Figure 4.29, is given in Figure 4.30. The dotted rectangles in the transform indicate sections when the amplitude modulation is small enough to make the general phase trend meaningful, and when the phase is smooth. Based on our crossrange width estimate, only the two leftmost sections are long enough (at least one-sixth of the transform length) that the corresponding images are likely to allow useful crossrange position measurements.




Figure 4.29 Image cut in Range Gate -7.09 .




Figure 4.30 Transform of the strong response of Figure 4.29.

In general, one may find that a suitable time segment selected in one range gate is not good for another range gate, which happens when flexing or vibration are dominant. We must examine other range gates and determine whether an interval exists that is good for a majority of the strong scatterers on the target. If that interval is longer than needed for useful crossrange resolution, we choose a subinterval when the motion best approximates rigid-body yaw. If the interval is only long enough to give at least marginal crossrange resolution, we form a single image over the interval. If the interval is too short for useful crossrange resolution, we must utilize different images for different scatterers.

Continuing, in Figure 4.31 we show the image cut in Range Gate 27.37 and its transform, with the acceptable section of constant amplitude and smooth phase indicated by the dotted rectangle. The gray rectangles show the acceptable sections of Figure 4.30. The longest sections that are acceptable for both scatterers (between about relative times -0.27 and -0.17 seconds and -0.13 and -0.03 seconds) are only about one-tenth the transform length, somewhat shorter than desirable.

Examining other range gates in the same manner, we find that the longest imaging interval acceptable for most scatterers on the target is

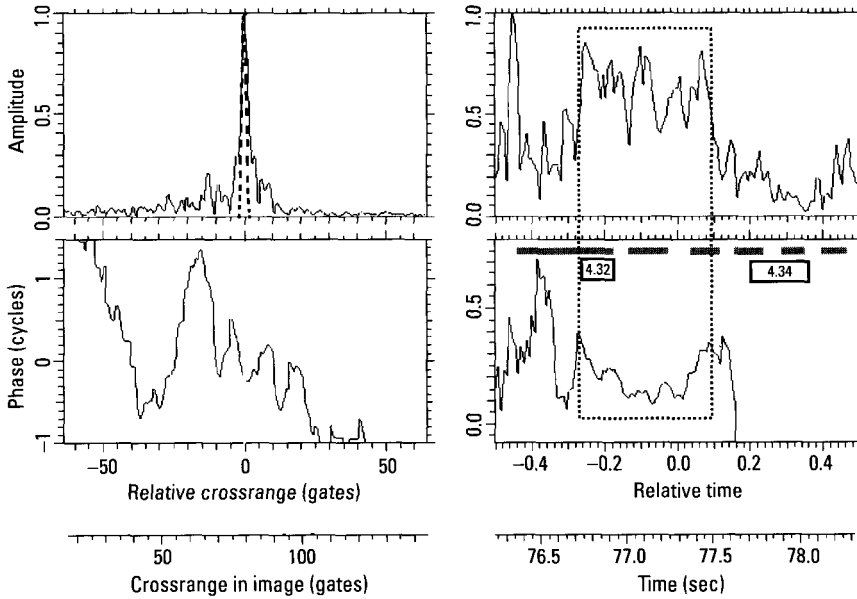


Figure 4.31 Image cut and transform in Range Gate 27.37.

between relative times -0.27 and -0.17 seconds, indicated by the box labeled 4.32 within Figure 4.31. Because this interval is shorter than our estimate of the imaging interval required for usable crossrange position measurements, we must investigate whether we can find a longer interval by recompensating the survey image (this is the leftmost branch from the second-highest box of the flowchart of Figure 4.25). Unfortunately, recompensating does not produce a longer imaging interval. Following the flowchart, we next examine a succession of reduced-duration images. These too fail to produce a longer imaging interval. The approaches did not produce a longer imaging interval in this case because they are designed to improve measurement conditions when uncompensated target motion causes strong scatterers within a single range gate to be unresolved. However, most of the strong scatterers on this target are located near its long illuminated edge. Most range gates contain just a single scatterer, in which case resolution is not an issue, or two scatterers so close together that their resolution would require tracking and accurately compensating the target's irregular motion, which is impractical.

The longest common interval allowing good compensation, between relative times -0.27 and -0.17 seconds, is about 60% of the duration that we estimated will give good crossrange position measurement accuracy for most

scatterers. Thus, we expect that an image formed over this interval will give good accuracy for some fraction of the scatterers (this puts us at the box at the lower left of the flowchart). It will likely be a fairly high fraction for this particular target, which spans 50 range resolution cells and has most scatterers along one edge. Regardless of the particular target, the alternative to using an image formed over the too short interval is forming different images for different scatterers, which generally gives lower crossrange position accuracy.

In general, even when the image-selection procedure produces a good compensation interval of sufficient duration, *compensation will not give an image equivalent to that from a stationary vehicle*. The reason is that the individual parts of the vehicle may not move in the same fashion, because the vehicle may flex and its parts may slowly vibrate in different ways. In our example, a check of the major responses shows that most move so that their phase functions are linear over the selected time segment. The image quality thus is expected to be as high as one can expect for a moving vehicle.

After compensating the motion of a single scatterer, we must again examine fixed-range cuts through strong responses in the resulting image. If the phase functions of response transforms are sufficiently linear, no further motion compensation is needed. In principle, resampling and polar reformatting may be necessary. In practice, we usually choose short enough imaging intervals that they are not. More commonly, we must compensate a nonlinear phase component common to most range gates, which we introduced by compensating a scatterer with a shifting phase center.

In our example, the short imaging time makes resampling and polar reformatting unnecessary. The final image is shown in Figure 4.32, including responses within 34 dB of the strongest, with the approximate target outline given by the dotted rectangle. We see a scarcity of dots in the range gates containing strong scatterers. This is an indication of high image quality. It is easy to analyze the responses along the illuminated edge in Figure 4.32, and fit a straight line for a good definition of the edge. Although the two dots at closest range give an indication of the orientation of the front edge of the vehicle, the peaks of a peaks plot may not be generated by single scatterers and must be analyzed with the TSA, and the edge can be better determined by the procedure we applied to images of stationary vehicles: analyze an image cut through the two responses and several slightly displaced parallel cuts. With both illuminated edges defined, we can choose the crossrange scale so that the two edges are perpendicular to each other, and then measure length and width of the vehicle as for a stationary target. The measurements of length and width will be weighted for identification in accordance with the measurement accuracy, which can be estimated from the number and




Figure 4.32 Final image.

quality of the responses used to define the two vehicle edges. The measurement accuracy is also affected by the orientation of the vehicle, which is favorable to the length measurement for this small aspect angle.

To perform the two-dimensional positional match, we must determine the positions of the scatterers associated with the strongest responses of the final image, perhaps 20 in number (see Appendix H). We also measure the scatterer characteristics for sufficiently isolated responses, to help in the identification process. In the present instance, we have the difficulty that the available photographs of the vehicle show mostly perspective views. This decreases the accuracy with which one can estimate feature positions and may preclude the use of some feature characteristics.

The template match for the scatterer positions extracted from the image of Figure 4.32 is shown in Figure 4.33. A question mark in the list of features in Figure 4.33 means that we could not determine the nature of a specific feature observed in the photographs. Since the differences in the motions of the individual scatterers are small for this vehicle, and because of the short processing time, nonrigid target motion does not shift responses appreciably in crossrange. Thus, the discrepancies between measured and




Figure 4.33 Positional match for the image of Figure 4.32.

predicted feature positions are caused by measurement errors due to interference, or by the inaccurate extraction of feature positions from photographs. Nevertheless, each extracted position agrees with a template feature, within extracted position uncertainties of 0.2 range resolution cells and 1.0 crossrange resolution cells, with template feature position uncertainties of 0.2m in each direction, and with angular uncertainties of 5° . The match of Figure 4.33 appears adequate for identifying the vehicle even if other vehicles were of a similar size, not considering that we will also utilize length, width, and scatterer characteristics.

To provide a feel for the necessity of the image interval selection process, we next perform a match for an image generated over a sufficient image duration for useful crossrange accuracy, but taken at an arbitrary time. We use the interval between relative times 0.19 to 0.36 seconds, indicated by the box labeled 4.34 in Figure 4.31. The figure shows that the new time




Figure 4.34 Image between relative times 0.19 and 0.36 seconds.

interval is not well chosen with regard to either range gate. The image over the new time interval is shown in Figure 4.34, with the approximate target location given by the dotted parallelogram (it is not rectangular, because range and crossrange are not plotted to scale). Compared with Figure 4.32, there are more responses because of the longer imaging interval (higher cross-range resolution and worse motion compensation). This problem is only partly remedied by analyzing only the stronger responses. The aspect angle change between the two imaging times is too small to expect changes in the observable scatterers. The positional match for the new image is shown in Figure 4.35, using the same comparison template as in Figure 4.33. About one-quarter of the measurements and one-quarter of the comparison features do not match.

4.4.4.2 The Off-Highway Truck Moving on a Smooth Straight Road

Figure 4.36 shows the survey image of the vehicle, again generated by range centroid and Doppler centroid compensation of the image excised from the SAR scene. The high degree of smearing of the responses in crossrange implies a significant residual motion with nonconstant Doppler. The image




Figure 4.35 Positional match for the image of Figure 4.34.

quality is poor despite the relatively smooth motion because *in the absence of a turning motion the “vibration” Dopplers are more significant relative to the Doppler due to the aspect angle change* from the translational motion. The situation is better when the vehicle is moving in a circle, despite the fact that it is moving on terrain rather than a smooth road. The first step is to determine to what degree the imaging interval may be reduced. From the crude image of Figure 4.36 we estimate that the width of the vehicle is on the order of five crossrange gates. Although we do not need high crossrange resolution when the aspect angle of the vehicle is small, five crossrange gates over the width of the vehicle is not so much that it can be greatly reduced.

The next step is the examination of image cuts in range gates spread over the range extent of the target. The transform of the image cut in Range Gate -23.41 is shown in Figure 4.37. The amplitude shows a slow modulation of slightly more than one period, plus a fast modulation corresponding




Figure 4.36 Survey image of off-highway truck on a smooth road.




Figure 4.37 Transform of the image cut in Range Gate -23.41.

to the weak peaks in the image cut. These peaks are at the noise/clutter background level of the image, so the corresponding modulation does not indicate target crossrange width. The phase function exhibits a significant behavior absent from our previous examples. The phase slope has a break at the boundary between the two dotted rectangles, and the two linear phase slopes correspond to crossrange positions that differ by about 2.3 gates. The amplitude and phase functions within each rectangle are acceptable, although the phase fluctuation near a relative time of 0.25 seconds is borderline. As explained below, the phase-slope difference is too large to accept within an imaging interval.

If all scatterers on the target had the same behavior, we could choose to image over either interval indicated by the rectangles, or alternatively correct the break in the phase function (probably not sufficiently well). The former approach was usable when the vehicle was turning, because the turning motion provides more than sufficient crossrange resolution with short imaging intervals. With a target width of only five crossrange gates, a correction of the break in the phase slope would be more in order if it could be done with sufficient accuracy, provided the vehicle is so rigid that the break occurs for all scatterers (it would generate spurious responses). We note that this check is implicit in the flowchart of Figure 4.25, in the loop between recompenation and interval determination.

The transform of the fixed-range image cut in Range Gate -6.20 lacks significant amplitude modulation and has a phase-slope break at the same time as Figure 4.37, but the break is only about 1.5 gates, implying a different motion than that of the first scatterer. The transforms of fixed-range cuts through strong scatterers in Range Gates 7.5 through 20.2 all have amplitude modulations with four to six modulation cycles, but the interference prevents determining good imaging intervals. These range gates confirm our crossrange width estimate. A last check in Range Gate 23.45 shows a transform amplitude with six modulation cycles, and a phase-slope break of 0.9 gates, at the same time as the others.

Comparison of the four phase functions shows that there is no single additional motion compensation step that would yield a good image, because the phase-slope breaks are not consistent with rigid-body translation or rigid-body rotation. In the former case, the phase-slope breaks would be equal. In the latter, they would scale with the separations of the scatterers. Thus, at the time of the phase-slope break, either the target moves nonrigidly or the target's translational and rotational motions simultaneously change jerkily. Which condition occurs is of purely academic interest. If the target moves nonrigidly, we cannot, in principle, compensate the entire target with

a single compensation. If the target moves rigidly with simultaneous translational and rotational jerks of the measured magnitude, we cannot, in practice, compensate the motion. In other words, if we wish to image over both intervals in Figure 4.37, we must vary the motion compensation over the vehicle, which in most situations is impractical for real targets. Even if it should work under some circumstances, a simpler approach is preferable. Thus we must further reduce the imaging interval, despite the loss of cross-range resolution.

To repeat the description of this important step, we take transforms over the range gates containing reasonably dominant single scatterers, such as represented by Figure 4.37, and ask: *Can we find a time segment over which the phase slopes in all range gates (spread over the range extent of the target) are reasonably continuous and the phases have reasonably small fluctuations?* As discussed in Section 2.3.3, reasonably continuous phase slope means that changes in the phase slope should be less than some threshold based on two criteria. First, the phase-slope difference implies a crossrange gate separation in the image over a reduced time segment. This separation should be less than one gate. With our normalization, this implies a phase-slope difference less than the reciprocal of the normalized duration of the reduced time segment. Second, the phase-slope difference causes a measurement error that must be a small fraction of the crossrange width of the target if the crossrange measurement is to be of use. Accepting 20% of the width gives a difference of one gate for this example. Reasonably small phase fluctuations are those that do not differ by more than about 0.1 cycles from a moving average of the phase, averaged over a duration equal to the lesser of one-tenth the dwell time and half the mean amplitude modulation period.

An examination of the phase functions in the range gates discussed above shows that the imaging interval from 0.1 to 0.4 seconds of relative time comes close to our requirements. This is not always so at the fringes of this interval, but the fringes are less important, because weighting is used for Doppler sidelobe suppression. With our estimate that the present image has a crossrange width of about six gates, use of the shorter imaging interval would reduce this number to between one and two. This appears acceptable, because range resolution performance is so high for the small aspect angle that little crossrange resolution may be needed. The crossrange positions of the scatterers still can be measured to much better accuracy than the crossrange resolution (the position of a response peak can be measured to a small fraction of the response width). We note that the reason that operating with a relatively crude crossrange resolution is practical in this instance is the tendency of the strong scatterers to be located near the illuminated edge of

the target, so that crossrange resolution becomes needed only for large aspect angles. Attempts to find a longer interval, via recompensation and the examination of reduced duration images, reproduce the interval found from the survey image. This is because the target motion is fairly smooth, except for the sudden change at the time of the phase-slope break.

An examination of the four phase functions also indicates that a second acceptable, although inferior, imaging interval is from -0.26 to 0.06 seconds of relative time. This interval is inferior because, in several of the other range gates that were examined, the interval contains phase modulations with deviations of about 0.1 cycles superposed on the slow trend, which is barely acceptable. As a test of how much selecting the inferior imaging interval deteriorates the image quality, we derive a comparison template using photographs and the first interval, and then match the template to the measured scatterer positions for the second interval. We derive the template using the first interval as well as photographs, because the photographic coverage is poor for the relevant aspect sector. Since the centers of the two intervals are separated by only about 0.7 seconds, any differences between the two positional matches should be due primarily to differences in the instantaneous motions rather than a change in the observed scatterers.

The image for the time interval from 0.1 to 0.4 seconds of relative time is shown in Figure 4.38, with the approximate target outline given by the dotted parallelogram. Reducing the imaging interval has not eliminated the spurious responses generated by the "vibrations" of the scatterers. When the vehicle moves in a circle, as in Figure 4.32, most of the responses are confined to the outline of the vehicle, while most of the spurious responses in Figure 4.38 are outside the outline. Note that in the latter figure the width of the vehicle is less than about two crossrange gates. The spurious responses are a first indication that the mere reduction of the image interval, as was used with the vehicle on a circle, will not suffice.

Quite generally, the motions of a ground vehicle and its scatterers may be so erratic that one cannot motion compensate adequately to generate an image of such quality that two-dimensional scatterer positions may be extracted. With a good motion compensation, if the transform of a response shows the amplitude/phase pattern from two interfering scatterers, we can resolve these scatterers and measure their positions. If the motion compensation is inadequate, the transform of a response may show an approximation of the amplitude/phase pattern from two scatterers, and yet there may be only a single scatterer responsible for the response. In other words, when the motion compensation is not satisfactory we cannot expect crossrange resolution to work. Hence, *before we can use the TSA to resolve scatterers*




Figure 4.38 Image from 0.1 to 0.4 seconds of relative time.

in crossrange by analyzing a response, we must be sure that the motion compensation is adequate. This will be known when the procedures developed here are used. However, even if it is not known, we can test for the adequacy of the motion compensation as follows.

Consider the image cut in Range Gate -6.5 of Figure 4.38, as shown in Figure 4.39. If the motion compensation is known to be adequate, we take a transform over the window marked in the figure, and use the TSA on the resulting amplitude and phase functions. On the other hand, when we want to test the adequacy of the motion compensation, we cannot use a narrow transform window because it cuts off the high frequencies needed to analyze the motion behavior. Thus we must take the transform over at least about 10 crossrange gates. In a case such as depicted in Figure 4.39, where the entire interval does not contain other major responses, we take the transform of the entire interval (in case other strong responses are present, we use the widest available interval that does not include these responses). Since the high frequencies are retained, the true form of the phase function is obtained. In this example the transform over the entire displayed interval is given on the right side of Figure 4.39. We conclude that a single scatterer is present,

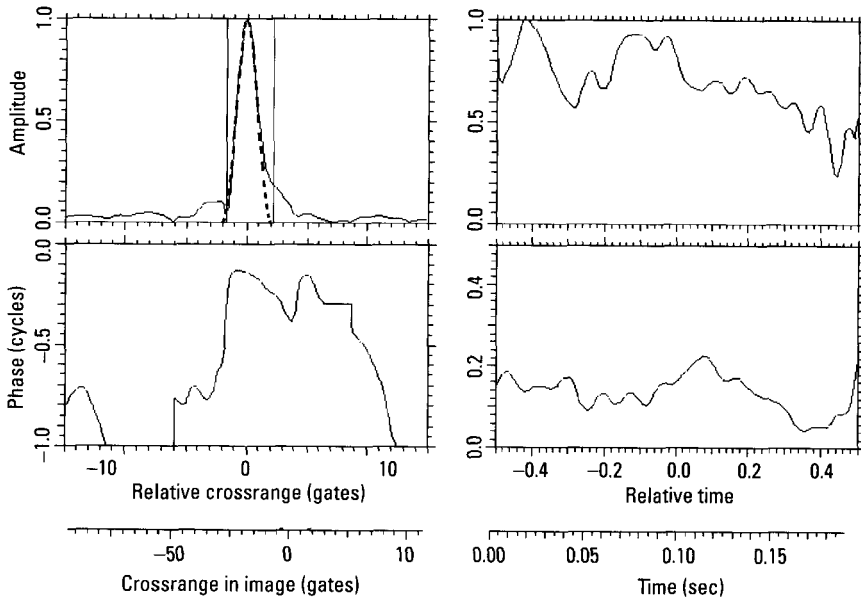


Figure 4.39 Image cut in Range Gate -6.5 .

because the amplitude does not have any significant modulation. Hence, the phase function represents the motion of this scatterer, and we observe definite breaks in the phase slope. Measurements of the phase slopes reveal that they correspond to crossrange differences of nearly one gate. Measurements of scatterer positions via the TSA, as well as any measurement of scatterer characteristics other than position and strength, would be meaningless under these circumstances.

Realizing that the image of Figure 4.38 is necessarily poorly motion compensated, with different residuals for different scatterers, *we extract the scatterer positions by transforming wide crossrange intervals and then measuring phase slopes for intervals of constant amplitude and linear phase.* In order to perform a positional match, we also examine the available diagrams and photographs of the vehicle to determine the scatterers that should be visible for such a small aspect angle, and estimate their positions in three dimensions. Because the available diagrams and photographs have poor coverage of the off-highway truck for this aspect, just a few scatterers are readily recognizable. We therefore created a feature template from these, matched it to the measurements, and then added features to the template consisting of poorly recognizable scatterers close to positions corresponding to unmatched

measurements. The match between the measured and predicted scatterer positions is given in Figure 4.40. Because of the incomplete motion compensation, we must allow a larger measurement error in crossrange than in range.

As the match of Figure 4.40 illustrates, we were (barely) able to measure a sufficient number of scatterers on the front edge to define this edge. With a good definition of the long edge, we then can adjust the crossrange scale factor in the image so that the two edges are perpendicular. This allows us to measure length and width as an input to identification separate from the two-dimensional positional match, despite the significant flexing of the vehicle. The reason is, of course, that the vehicle is moving slowly on a good road.

As already stated, we made the same measurements on the poorer image generated from the time interval between -0.26 and 0.06 seconds of relative time. This image is given in Figure 4.41. A comparison of Figures 4.38 and 4.41 makes it obvious that one cannot simply take the peak




Figure 4.40 Positional match for the image of Figure 4.38.




Figure 4.41 Image from -0.26 to 0.06 seconds of relative time.

positions as the scatterer positions, since both images then would have to be nearly identical. If the above measurement procedures are used to determine the scatterer positions from the new image, and the measured positions are matched to the feature template used in the match in Figure 4.40, we obtain the new positional match of Figure 4.42. The match is still very good, but not as good as in Figure 4.40. This is to be expected, because the motion behavior of the vehicle is not as good over the second interval. Nevertheless, even the new image permits measurement of length and width.

Moving Versus Stationary Off-Highway Truck

We have now considered the off-highway truck moving in a circle and also on a straight smooth road, showing that one can obtain length, width, and scatterer positions in both situations. This means that the different motion characteristics do not prevent us from measuring the same scatterer positions and other vehicle parameters. With respect to the scatterer positions, are they the same positions one measures on the stationary off-highway truck? They must be, if reliable target identification is to be achievable. To clarify this point, we choose an image of the stationary off-highway truck at




Figure 4.42 Match between scatterer positions from Figure 4.41 and the feature template used in Figure 4.40.

approximately the same aspect angle used for both cases of the moving off-highway truck. We extract the scatterer positions from the image and perform the positional match with the same database that was used for the preceding two positional matches. The result is shown in Figure 4.43. Because of the importance of this test, we will discuss the match in more detail.

Starting with the front row of scatterers, the stationary vehicle shows one additional scatterer, which has the position of the left corner of the bumper. It is not surprising that we can observe a weak scatterer on a stationary vehicle but not on a moving vehicle. The other matches are for the most part extraordinarily good. We do not have a response at the position of Feature B, and there is an additional response to the right of Feature S. Since we have inadequate information for the top view, we do not know the corresponding feature. The differences observable for some scatterer positions can




Figure 4.43 Match between the scatterer positions for the stationary off-highway truck and the feature template used for the moving off-highway truck.

be traced to interference from strong responses, which causes unavoidable measurement errors or even the loss of weak responses. Overall, however, the positional match between the scatterers extracted from the image of the stationary off-highway truck and the feature template of the moving off-highway truck is very good.

4.4.4.3 The Off-Highway Truck Moving on a Bumpy Straight Road

In this case the vehicle traveled on a straight road that had a variety of surface deformations to represent different degrees of bumpiness. In Figure 4.44 we show the range profiles for the vehicle over two seconds, after removing the drift due to the motion along the road, via centroid compensation. We have previously demonstrated the examination of fixed-range image cuts, and summarized the results of the recompenations that were necessary when the




Figure 4.44 Range profiles over two seconds.

longest common interval of good scatterer compensation was shorter than the imaging duration required for useful crossrange position measurement accuracy. We now demonstrate the compensation of a specific scatterer, which is found to be required when one examines fixed-range image cuts in the survey image corresponding to Figure 4.44. For this purpose, in Figure 4.45 we show the plots of the peaks of the consecutive range profiles (peaks tracks), without any refinement of the measurement. We arbitrarily select the strongest scatterer, in Range Gate -4 (indicated by the arrow), for the next compensation step, even though the peak wanders over somewhat more than one range gate. An expanded version of this peaks track is shown in Figure 4.46, together with a quadratic fit. The residual Doppler of the




Figure 4.45 Peaks tracks for Figure 4.44.




Figure 4.46 Range profile peak selected for tracking.

scatterer after the range compensation is given in Figure 4.47, again with a quadratic fit. The smoothness of the Doppler motion implies that range tracking has worked sufficiently well. When the fitted Doppler is compensated, we obtain the image of Figure 4.48.




Figure 4.47 Residual Doppler of the scatterer.




Figure 4.48 Image after range/Doppler compensation of one scatterer.

The compensated scatterer is in Range Gate 0 of Figure 4.48. The success of the motion compensation must always be verified, whether the processing is done manually or automatically, and verification is always done by examining the image cut in the range gate of the compensated scatterer. This image cut is shown in Figure 4.49. The amplitude function of the transform shows a sudden drop to a much lower level. One might consider the possibility that this is due to a tracking error, but this is unlikely because of the suddenness of the drop and the fact that the amplitude remains essentially constant after the drop. If it were due to a tracking error, slightly displaced range gates would display an amplitude increase when Figure 4.49 has the decrease. However, the displaced gates show the same behavior as Figure 4.49. Furthermore, the amplitude variation is recognizable in the range profiles of Figure 4.44, and their variation is independent of tracking. On a smooth road, such a sudden drop in scatterer strength is unusual for targets away from broadside. In this particular instance, the vehicle is moving on a road with bumps. The most likely explanation for the behavior of the scatterer is that the vehicle jerks in such a manner that the scatterer suddenly changes its aspect, causing the decrease in strength. This view is supported by the fact that from this moment on, the phase function indicates a much larger vibratory motion.




Figure 4.49 Image cut in the range gate of the compensated scatterer.

A practical question is whether the tracking performance represented by Figure 4.49 is sufficient for a usable image. To answer the question we again examine the phase function of the transform, which represents the residual motion of the scatterer, since the amplitude function shows no significant interference from a second scatterer. The high-frequency interference in Figure 4.49 is not of concern because the amplitude minima are high; the corresponding phase jump is only about 0.05 cycles (see Appendix A). We can safely attribute this modulation to interference (and thereby calculate the corresponding phase) because our earlier compensation steps cannot have introduced the modulation; we used only quadratic fits. The transform phase function at the lower right of Figure 4.49 shows that the residual phase variation is nearly one cycle, which is far too large to accept for imaging. This variation causes the crossrange smearing of Figure 4.48.

Whereas we restricted our earlier compensation steps to low-order fits, *now that we have verified that we are tracking a single scatterer, we can fit a more flexible spline to the phase.* We must take care not to follow phase variations corresponding to the high-frequency amplitude modulation and not to follow the rapid phase change associated with the sharp amplitude drop. The figure shows an appropriate fit. The deviations of the phase from the fit are much less than 0.1 cycles, except at the time of the sharp amplitude drop, hence are insignificant. After compensation of the data with this fit to the phase, the tracking of the scatterer thus is adequate. The image obtained with this compensation is shown in Figure 4.50.

The vehicle now, in effect, rotates about the compensated scatterer. If the vehicle were rigid and the roll and pitch motions were insignificant, we could track one additional scatterer and determine the necessity for resampling, polar reformatting, and shortening the duration of the data. In order to determine whether the target can be treated in this fashion, we must compare the residual tracks of at least two (and preferably more) scatterers. The target behavior can be most readily determined by examining image cuts at the extremities of the target. The effect of the motion compensation on a second scatterer, in Range Gate -8.8 , is shown in Figure 4.51. The amplitude function of the transform indicates that the scatterer drifts out of its range gate at the ends of the time interval. However, even within the time interval with a nearly constant amplitude (when the phase is a valid track of the motion), the phase varies by about ± 0.1 cycles. This already is undesirably large. The interpretation is that the new scatterer has a significant irregular motion relative to the compensated scatterer.

Residual motion is still worse with other scatterers. As an example, in Figure 4.52 we show the image cut and its transforms for Range Gate 17.9,




Figure 4.50 Compensated image.




Figure 4.51 Image cut and transforms in Range Gate -8.8.




Figure 4.52 Image cut and transform in Range Gate 17.9.

which Figure 4.50 shows to be about halfway between the compensated scatterer and the end of the vehicle. The amplitude variation indicates a strong residual range drift of the scatterer, and the phase variation is unacceptably large. We again note that the phase can be directly interpreted when the amplitude function is strong and fairly constant, during the first 40% and last 20% of the imaging time. The strong residual phase variation is the primary reason why the response of the scatterer, shown by the left upper plot, is grossly smeared. The residual phase variation also is very different from that of Figure 4.51.

The reader may wonder about some of the preceding interpretations of complicated amplitude and phase functions, but the precise interpretations are of no interest within the present context. The obvious fact is that the various scatterers on the vehicle move very differently. When one scatterer is motion compensated to give a sharp response, most of the other responses become highly smeared. By compensating the scatterers in the different range gates differently, we can properly compress their responses and, in principle, measure their crossrange positions. However, the proper measurement of the crossrange positions requires the examination of a time interval that extends over several “vibration” cycles of the scatterers, so that their average

crossrange positions may be determined. The compensations and measurements needed for this purpose will not always be successful, perhaps never fully successful. They will be successful for those scatterers that are dominant in their range gates. They may also be successful when a range gate includes two scatterers that are dominant over all others in the range gate, because we can use the phase-slope tracking procedure. When a range gate contains three or more scatterers of comparable strengths, it is unlikely that one can compensate the individual scatterers, since each of them might be moving differently.

The practical situation is not as hopeless as it may appear from this discussion. In many cases the scatterers near the illuminated edges of the vehicle are dominant in their range gates. If a scatterer away from the edges is strong, there may or may not be more than one additional strong scatterer in that range gate. The individual motion compensations in each range gate thus may be successful (which can and must be checked in each case). The practical implementation of the complicated measurement procedure is as in the rightmost branch of the flowchart of Figure 4.25. We shift the compensation from scatterer to scatterer, so long as one is trackable and we have not found an interval during which most strong scatterers can be compensated, and useful crossrange resolution can be obtained. Each new compensation is likely to improve the measurability of nondominant scatterers near the compensated one (unless they are at different heights and vibrate differently). We keep track of which compensation yields the longest usable interval for each scatterer, and eventually measure each scatterer's position in an image over its best interval, created with the appropriate compensation.

The duration over which one must track scatterers is governed by the crossrange width of the target. If the target is wide, tracking scatterers over shorter times than the full dwell may be sufficient (this would be the leftmost branch of the flowchart). In our example, the only interference evident in the transform amplitudes of Figures 4.49, 4.51, and 4.52 is a weak high-frequency modulation consistent with interference between the strong response in each gate and the background clutter/noise. Some other range gates show a slow modulation that indicates their effective crossrange widths to be less than two gates. Therefore, we must attempt to track and re-compensate over the entire two-second dwell.

Although we may be able to proceed as described above, there is an important practical question: Is such complicated processing really needed for ground vehicle identification? As a simpler alternative, we can examine each range gate on the target and determine whether there is a crossrange interval whose transform shows one scatterer to be dominant for some time

interval, then measure the range of the scatterer in an image formed over that interval. This avoids the multiple compensations, but consequently shortens the intervals used for measuring crossrange positions of nondominant responses. This degrades positional accuracy, because generally only part of a motion cycle is available for the crossrange measurement. Thus, identification depends almost entirely on the range measurements. This is still much superior to basing identification on the peak positions of an intensity range profile.

The question of whether it is necessary to perform a complicated motion compensation that varies over the vehicle, and then to determine the crossrange positions of the scatterers by averaging over several motion cycles, cannot be answered within the scope of a textbook. One would have to design a fully automated processor that can examine a large number of ground vehicles under all conditions of motion, and statistically determine whether the identification performance is sufficiently reliable. We believe that it would be, but it is not our purpose to consider such questions. We want to present the various options for ground vehicle identification, including very sophisticated methods.

We measured the positions of the well-observable scatterers in this manner, determining their crossrange positions from the phase slopes at the time of measurement, instead of determining the average crossrange position over several motion cycles. In other words, we perform an accurate range measurement but accept a possibly large error in the crossrange measurement. When these scatterer positions are matched to the database used for the off-highway truck on a smooth road, we obtain the positional match in Figure 4.53. As pointed out, in this case we must include a large uncertainty in the measured crossrange positions. Thus, visually, we must examine how good the match of Figure 4.53 is in range only. This match evidently is very good. The practical conclusion is that it might be simpler and still adequate to measure only the range positions, utilizing crossrange resolution but not crossrange positions.

4.4.5 The Tank Moving in a Circle on Terrain

4.4.5.1 The Tank at Head-on Aspect

We now consider the tank moving in a circle off the road, viewed head-on. For that aspect angle, many important scatterers visible in a sideview may not be observable, and the geometry also prevents an accurate determination of the crossrange scale. On the other hand, the rolling motion of the vehicle




Figure 4.53 Positional match for the off-highway truck on a poor road.

has little impact at this aspect angle, because it does not generate significant Dopplers. It would take a large pitch motion to complicate the situation.

Because the vehicle is in a SAR scene and is even turning, we can observe it over a longer time than needed, so the best imaging interval can be selected. The survey image generated with a two-second observation interval is shown in Figure 4.54. Because of the relative ineffectiveness of the bouncing and rolling motions, the image is only slightly smeared in crossrange. Nevertheless, it is still too smeared to permit good positional measurements.

Figure 4.54 shows strong smeared responses near Crossrange Gates 20 and 15, giving a target width of roughly 35 crossrange gates, much larger than necessary at a nose-on aspect. Fixed-range image cuts give the same width estimate, implying that we can reduce the image duration by about a factor of six. The transforms of the fixed-range cuts also show that an interval with good amplitude and phase functions which is common to all examined range gates extends over 0.30 seconds (this is the central branch from the second highest box in the flowchart of Figure 4.25). The corresponding image is




Figure 4.54 Image of the tank head-on.

shown in Figure 4.55. The relatively small number of major dots in single range gates indicates that this image is of much higher quality than the one taken over two seconds in Figure 4.54. However, because of the nose-on aspect, only the length of the vehicle can be measured accurately, not its width. From the shadowing evident in the image (the lack of responses in the central part of the rear of the vehicle), we conclude that the vehicle has a turret. Since in this example the opening of the gun barrel is clearly visible in Range Gate -20 , indicated by the arrow, we will make use of it to conclude that the vehicle also has a large gun. However, most of the time the gun will be pointed in a direction for which its response will be too weak to be observable in the background clutter/noise.

We perform the usual measurements on the image in Figure 4.55, extracting the scatterer positions. For a head-on aspect we will not see the wheels, but we will see the various metallic boxes at the side of the deck. The positional match is given in Figure 4.56. Despite the unfavorable aspect angle, the match is very good, except for two unmatched responses, which Figure 4.55 shows to be weak. We did not have photographs or diagrams of such a quality that the associated scatterers could be identified.




Figure 4.55 Image of the tank over a shorter interval.




Figure 4.56 Positional match for the tank head-on.

4.4.5.2 The Turning Tank Viewed at a Larger Aspect Angle

The rolling motion of the vehicle has little consequences near head-on aspects, and for the design of the vehicle and the beam depression angle of the SAR system, no significant pitch Dopplers are generated. We now consider the same vehicle at an aspect angle of 30° .

Figure 4.57 shows a two-second survey image of the vehicle, with the clip level for the lower responses set relatively high at -20 dB with respect to the strongest peak, so that only the major response peaks of the image are shown. Only a smooth standard motion compensation that cannot follow details of the vehicle motion was used. We observe sets of responses repeated in crossrange, the effect one obtains if a near-periodic modulation is superposed on a carrier. This is a consequence of the roll and yaw motion of the vehicle. The roll and yaw frequency is so high that the repeated responses fall outside the smeared main responses. This is fortunate because these spurious responses can effectively be ignored. However, it is conceivable that a larger degree of smearing of the main responses could combine with a smaller separation of the repeated responses, so that repeated responses might fall within the smeared main responses. This could considerably increase the severity of




Figure 4.57 Image of the tank moving in a circle, 30° aspect angle.

the processing problem, depending on how strong the spurious responses are in relation to the stronger of the genuine responses.

Figure 4.58 shows a subimage of Figure 4.57, with the lower clip level set at -40 dB from the highest peak. This figure provides a better appreciation of the degree of crossrange smearing in the image. In order to generate the final image, we use the same procedure illustrated earlier: find the major scatterers, take transforms of the image cuts in these range gates, and search for usable imaging intervals. The examination of the fixed-range image cuts shows that no single time interval allows the measurement of the crossrange positions of most strong scatterers with an accuracy useful for identification. We must proceed as we did in Section 4.4.4.3 for the off-highway truck on the bumpy road, analyzing each scatterer over its own appropriate imaging interval. As before, we made no attempt to obtain accurate crossrange positions by averaging over several motion cycles, depending primarily on the one-dimensional match in range.

Because of the arbitrary orientation of the turret, we largely ignore the responses away from the edges of the vehicle. Figure 4.59 shows the positional match between scatterer positions extracted from the final image (not shown) of the moving vehicle and a feature template that was verified to be




Figure 4.58 Subimage of Figure 4.57.




Figure 4.59 Match between template for stationary tank and the measurements for the moving tank.

in excellent agreement with the scatterer positions extracted from an image of the stationary vehicle at a similar aspect angle. The match is excellent. To illustrate the relative significance of the match in range and crossrange, in Figure 4.60 we repeat the match of Figure 4.59, with uncertainty




Figure 4.60 Positional match with uncertainty ellipses.

ellipses added. For purposes of illustrating the positional match, we assumed (arbitrarily at this point) that the uncertainties of the measured positions are 0.2 gates in range and 2 gates in crossrange. For the predicted scatterer positions we assumed an uncertainty of 0.2m in position and 5° in angle. Figure 4.60 is another illustration of the fact that, without a much more complicated procedure of measuring crossrange positions, we must rely primarily on the range measurements. Even then, coherent processing over an extended interval is necessary in order to separate the responses so much in crossrange that accurate range measurements may be performed.

There is an interesting point concerning scatterer positions that may be extracted for a moving ground vehicle, as compared with those that may be extracted for a stationary vehicle. Even though a range resolution of about one foot should be adequate for identifying ground vehicles, some of the major scatterers may not be resolved. For example, the tank has cavities formed by metallic boxes at the edge of the deck, and below these cavities are cup-shaped wheels. They may be so close in range that they cannot be resolved, or the stronger response might mask the weaker one. Such effects limit the total number of readily observable scatterers, and the changing phase relations cause slow changes in the observable scatterers as the aspect angle changes.

When a ground vehicle is moving, in particular when it is moving in an erratic manner, the phase relations between scatterers are rapidly and continuously changing. By examining the data over a longer interval, we then obtain glimpses of a larger variety of scatterers than is possible when the vehicle is stationary. At the expense of more extensive processing, we thus can detect more features when the vehicle is moving. This was found to be true with the image in Figure 4.57, although we did not try to make use of the effect in our positional matches.

4.4.6 The Flatbed Truck on a Bumpy Straight Road

The analysis of the off-highway truck has shown that movement on the bumpy road poses the most challenging identification problems. We now consider another wheeled vehicle, a truck, on a bumpy road.

A two-second survey image with only the standard motion compensation is shown in Figure 4.61, and is evidently highly smeared. Again, we must analyze several range gates in order to understand the situation and choose the appropriate approach. In Figure 4.62 we show the image cut in Range Gate 12.05, which is perhaps the cleanest of the range gates. Since the transform amplitude does not have any deep minima, the range gate is




Figure 4.61 Survey image of the flatbed truck on a bumpy road.




Figure 4.62 Image cut in Range Gate 12.05.

dominated by a single scatterer, and the transform phase function describes the motion of this scatterer. The amplitude and phase functions tell us that the vehicle motion is much worse over the second half of the observation time than over the first half. In practice, in such a situation we would choose the first half for analysis, crossrange resolution permitting. Here, for the purpose of illustration, we will analyze the flatbed truck under the worse motion conditions of the second half.

In contrast to the first half of the observation time, the second half shows a rather large cyclical motion of the scatterer. The phase deviation is about three cycles, so that the scatterer moves back and forth by about 1.5 wavelengths, or nearly 5 cm. This is a drastic motion for the purpose of radar imaging. We could compensate the motion of the scatterer and compress its response. However, the scatterers on other parts of the vehicle are found to move differently, so that they would all have to be compensated individually. Even then, depending on the phasing of the motions, the response of a scatterer could appear far away in crossrange from the true position of the scatterer. As already explained, we then would have to average over more than one motion cycle to estimate the true scatterer position, which will not always be possible, in particular not with automated processing. As long as the required identification performance is achieved, it is much simpler to rely on the range positions of the scatterers for identification.

As can be seen from the overall linear trend of the amplitude in Figure 4.62, the scatterer is drifting in its range cell. This trend is disturbed by two amplitude spikes, which coincide with a reversal of the direction of motion as seen from the phase function. The strong amplitude modulation when the scatterer is moving back and forth by only 5 cm, or about 17% of the range gate width, cannot be explained on the basis of its changing range. More likely, the scatterer's orientation changes as it moves back and forth. The subimages over the two halves of the imaging interval in Figure 4.62 are shown in Figures 4.63 and 4.64. The differences between the image quality in the two cases can be predicted from the amplitude/phase functions of Figure 4.62. Again, in practice one will work with the better image, but here we will examine the worse.

Figure 4.65 shows the image cut in Range Gate 12.09 of Figure 4.63, corresponding to Figure 4.62 for the full image. The transform in Figure 4.65 repeats the second half of that of Figure 4.62. The slope of the phase function is Doppler, with the phase slope of the left interval between vertical lines corresponding to Crossrange Gate -22.86 , and that of the right interval to Gate 15.46 . The scatterer thus shifts by nearly 40 crossrange gates in a time of about one-third second. In principle, we could measure the average phase




Figure 4.63 Image over the first half of the imaging interval of Figure 4.62.




Figure 4.64 Image over the first half of the imaging interval of Figure 4.62.




Figure 4.65 Image cut in Range Gate 12.09 of Figure 4.63.

slope over an entire modulation cycle, but is the quality of the phase function so good that the phase slope can be measured with sufficient accuracy? Moreover, the motions of the different scatterers are not synchronized, and the accuracy problem may be much worse than in Figure 4.65. As long as it appears that the positional match in range is sufficient for vehicle identification, there is no incentive for attempting to solve such complicated processing and analysis problems. At the time of this writing, we had not investigated either the practical feasibility nor the necessity of extracting accurate crossrange positions from such an image.

For this example, we analyze the range gates that contain significant responses in Figure 4.63, using the same methods as illustrated earlier in considerable detail. We then match these measurements to a feature template that earlier (Figure 4.21) was found to be in excellent agreement with positions extracted from an image of the stationary flatbed truck at a similar aspect angle. We took the high variability in crossrange position into account by allowing a crossrange uncertainty of two gates, whereas the range uncertainty was assumed to be 0.2 gates. For the predicted scatterer positions we again assumed an uncertainty of 0.20m in position and 5° in angle. The resulting positional match is shown in Figure 4.66. The quality of the match




Figure 4.66 Positional match for the truck.

in range appears completely satisfactory, despite the extreme motion behavior of the vehicle and the use of simplified processing and analysis methods. We again note the importance of crossrange resolution, even though the positional match does not much utilize the crossrange positions of scatterers.

4.4.7 Recreational Vehicle Turning on a Paved Surface

4.4.7.1 Vehicle Not Close to Broadside

We now consider data of a recreational vehicle (RV), which is not designed to be very rigid. It is turning at a low speed on a wide paved surface, which somewhat offsets the low rigidity. Whereas the radar wavelength was 3 cm for the previous four vehicles, for the RV it is only 2 cm. The data again were taken by a SAR system. Thus, since the aspect angle of the turning RV changes much faster than that of a stationary vehicle, we again have excess imaging time as long as we do not want to achieve a crossrange resolution much better than that on stationary vehicles. This excess imaging time is used to analyze the vehicle motion and to select the imaging interval for

which a good image can be generated. Since the turn is relatively slow and smooth, we do not have severe bouncing, yawing, pitching, and rolling motions, so that the crossrange smearing of the image will be relatively small. On the other hand, the large height of the vehicle and its not-so-rigid construction imply that the consequences from the changing rotation axis and flexing may be significant even for these excellent motion conditions. These consequences are magnified by the relatively short radar wavelength of 2 cm.

A survey plot of the SAR scene containing the turning RV is shown in Figure 4.67, with the vehicle image marked by an arrow. The background below -26 dB relative to the highest peak is not shown in the image. The imaging time is 2.7 seconds. As Figure 4.67 indicates, there are clutter regions with cross sections comparable to that of the vehicle. Because the vehicle image can easily be shifted into these clutter regions during the turn, and because even clear-area clutter may influence the first stages of the motion compensation, one should use a SAR system with clutter cancellation. The corresponding intensity image after DPCA clutter cancellation is shown in Figure 4.68. The strong clutter has been reduced sufficiently to make the entire smeared vehicle image clearly visible.




Figure 4.67 Survey plot of the SAR scene with the RV.




Figure 4.68 Survey of the SAR scene after clutter cancellation.

A subimage without clutter cancellation, again using an imaging interval of 2.7 seconds but three seconds later than in Figure 4.67, is shown in Figure 4.69. Both the range and the crossrange intervals have been greatly reduced. The vehicle now is viewed near broadside. As expected for a near-zero range rate, the image of the vehicle remains in the clear area where it is turning. Hence, at broadside where clutter cancellation does not work, we do not need it. This statement applies only as long as the rolling motion of the vehicle is not so great that the image is smeared in crossrange into the high-clutter region. If that should be the case, we would face a difficult problem. Although the rolling high parts of the vehicle may be retained in the clutter cancellation process, those close to the ground will not.

A turning ground vehicle sometimes offers an opportunity not usually available for moving ground vehicles, if it is observed over a significant part of the turn. The vehicle then can be tracked to determine the aspect angle at the chosen imaging time, from which we can establish the crossrange scale (and length and width) even when the illuminated edges cannot be well identified. Figure 4.70 shows a subimage of Figure 4.68, the SAR scene with DPCA clutter cancellation. During the imaging interval of 2.7 seconds, the




Figure 4.69 Survey plot three seconds later, without clutter cancellation.




Figure 4.70 Subimage of Figure 4.68.

aspect angle of the RV changes by about 40° , with aspect angles not too close to zero or to broadside. The length of the vehicle is 24 ft and the aspect angle is nonzero, yet in the image the range extent of the vehicle is more than 30 ft. This means that the vehicle is moving over a considerable number of range cells during the imaging interval. We need to apply a crude motion compensation, and examine fixed-range cuts in the resulting image, as we did for the other targets.

We apply range centroid and Doppler centroid tracking, using only linear fits. Any residual motion of nonlinear Doppler thus will be due to the vehicle motion rather than the motion compensation. The resulting image is shown in Figure 4.71. A comparison with Figure 4.70 shows the reduced range and crossrange spreads of the image. However, it is evident that the 2.7-seconds duration used in Figure 4.71 is vastly excessive. The crossrange width of the target is roughly 200 gates, whereas we require less than 10 gates. As usual, we cannot simply reduce the imaging interval, but must select the best imaging subinterval. Thus, we next take fixed-range image cuts through the major responses, and search for transform intervals when the amplitude functions are sufficiently constant to permit using the phase for the motion estimate.




Figure 4.71 Peaks plot after crude motion compensation.

The best response that can be found is that in Range Gate -2 , for which the image cut is shown in Figure 4.72. Because of the crude motion compensation, the major single-scatterer response is smeared over the entire interval marked by the two crosshairs. The center of its envelope is separated from that of the scatterer to its right by about 80 gates, or not quite half the target width. The transform of the complex response over the indicated window is given in Figure 4.73. Over the first two-thirds of the imaging interval, the amplitude is sufficiently constant to allow measuring the motion behavior. The observed phase change gives the residual range wander of the scatterer directly, with a change of one cycle corresponding to a range change of half a wavelength. However, we cannot measure the motion behavior over the last third of the interval. Moreover, we cannot find another response in the entire image that would allow any measurement of the motion of other parts of the vehicle over the entire 2.7 seconds.

On the other hand, most responses permit the measurement of scatterer motion over intervals much longer than needed for usable crossrange resolution. In fact, even the intervals that permit the measurement of most of the scatterers are much longer than needed. We have the opportunity to choose that interval with sufficient resolution that gives the most accurate crossrange measurements (this is the central branch from the second highest




Figure 4.72 Image cut in Range Gate -2 .




Figure 4.73 Transform over the window of Figure 4.72.

box of the flowchart of Figure 4.25). As explained previously, the best time for measurement occurs when the vehicle best approximates rigid body yaw. This can be determined automatically by searching for those times when the relative scatterer motions scale with the scatterer separations.

For the purpose of illustration, we reduce the imaging interval to the first 1.5 seconds, with the new peaks plot of the crudely compensated image (range and Doppler centroid tracking) shown in Figure 4.74. When the Fourier transform is taken of the same response as used for Figure 4.73, we obtain the amplitude and phase functions of Figure 4.75. Since the amplitude has no deep breaks and the motion compensation used linear fits, the phase function describes the motion of the scatterer. Although the phase fluctuates by only somewhat more than half a cycle (range motion of a little more than a quarter wavelength), this is far too large a residual motion for a usable image.

We learn from Figure 4.75 that the scatterer in Range Gate -2 moves back and forth roughly sinusoidally, by about two cycles over the imaging interval of 1.5 seconds. This could be due to the motion of the vehicle, or it




Figure 4.74 Image over 1.5 seconds.

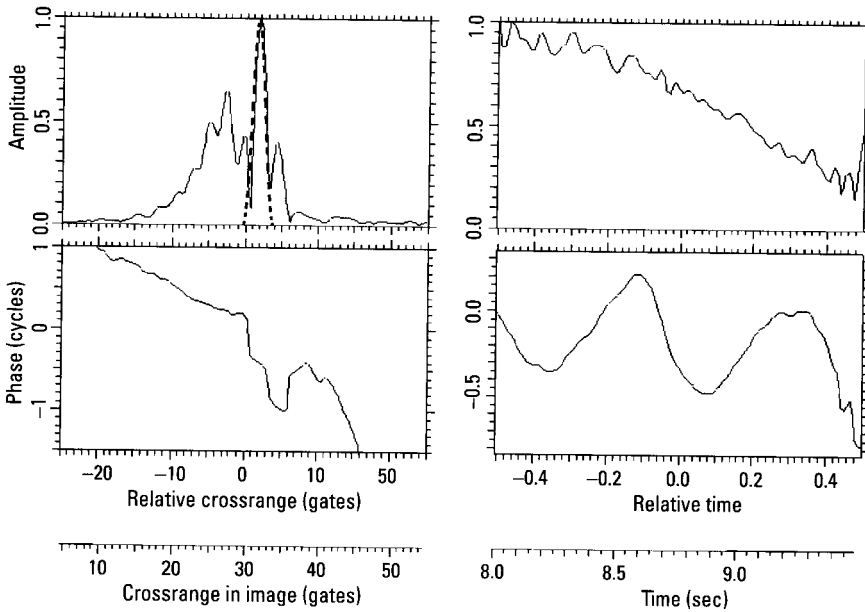


Figure 4.75 Transform of the smeared response in Range Gate -2.

could be the shifting motion of the phase center of a complicated scatterer (highly unlikely in the case of two phase cycles). We can test for this by comparing the motion to that of other dominant scatterers in nearby range gates. If nearby range gates do not contain sufficiently dominant scatterers that we can directly compare transform phase functions, we can test whether the phase center is shifting by using the phase function of Figure 4.75 to compensate the nearby range gates. The responses will be further smeared if the phase function corresponds to a phase center motion unique to one specific scatterer. If it corresponds to vehicle motion, the responses are likely to be compressed, unless the nearby scatterers have radically different motions.

We choose an image cut in Range Gate 0, about two range gates away from that of Figure 4.75. The amplitude function for this image cut is shown in the top of Figure 4.76. There appears to be a smeared response centered in about Crossrange Gate 20 in the top plot. After applying the phase compensation derived from Figure 4.75, we obtain the amplitude function at the bottom of Figure 4.76. The response clearly is much more concentrated, but far from perfectly. We conclude that the new scatterer is moving almost like that in Range Gate -2, but not quite. The implication is that the phase function of Figure 4.75 does represent vehicle rather than phase center motion,




Figure 4.76 Amplitude function in Range Gate -2 (top); after compensation (bottom).

but that the detailed motion characteristics change over the vehicle. This can be tested by finding the phase function of a scatterer far away, which should be quite different from that of Figure 4.75. However, with this type of crude motion compensation, a scatterer which remains in one range gate and whose spread response does not overlap too much with other responses (as required for this test) could be difficult to find.

Figure 4.77 shows the image cut through a response in Range Gate 7. Although deep amplitude breaks are absent only in the central part of the transform, so that the phase function gives a reliable indication of the scatterer motion only over that part, from a comparison with the phase function of Figure 4.75 it is clear that the scatterer motion is radically (for imaging purposes) different. Hence, we conclude that the vehicle has a significant motion that smears the responses in crossrange and falsifies the crossrange positions of individual scatterers when the motion compensation is refined so they are properly compressed. The vehicle motion could be due to yawing or flexing (the smooth road should minimize rolling and pitching). In order to distinguish between the two, we must examine a third scatterer, away from the first two. If the motion is rigid-body yaw, the phase differences between




Figure 4.77 Image cut in Range Gate 7.

scatterers will have the same shape, scaled by their separations. If the vehicle is flexing, this will not be the case.

Figure 4.74 shows that the target's range extent is about 20 gates. This gives a sufficient number of resolution cells along the long edge of the target, so crossrange resolution is needed primarily to provide resolution along the short edge. Five to ten cells are sufficient for this. Figure 4.74 shows the target's crossrange width to be about 100 cells. Thus, we can likely reduce our image duration by a factor of ten or more. Proceeding conservatively, we reduce it by a factor of seven, so require a duration of 0.2 seconds. After examining the resulting image, we can further reduce the duration, if necessary.

Next, we must decide when to choose our imaging interval. The 0.2-second interval must be chosen during those times that both the scatterers in Figures 4.75 and 4.77 (as well as other scatterers) are trackable; that is, when transforms of fixed-range image cuts through the scatterers have strong amplitude. Among these times, we want to choose the interval when the target behavior most resembles rigid-body yaw motion. Then, the Doppler shifts due to flexing are minimized. This means that we want the times when the flexing reverses motion. In the phase of Figure 4.75, the phase slope should be zero.

We will perform the positional match for three different images of the turning RV. The first image is formed centered at the normalized time of 0.06 seconds, where Figure 4.75 shows a stationary point for the flexing motion, so that the consequences should be minimal. We will repeat the imaging one flexing cycle earlier, at a normalized time of -0.37 seconds. This will show how quickly or slowly the observable vehicle features change with a change of the aspect angle. Lastly, we will form an image centered at a normalized time of -0.24 seconds, where in accordance with Figure 4.75 the flexing effects should be maximum. This will show whether or not these effects are serious.

To repeat an important point, in a given situation we must decide between trying to utilize crossrange resolution both for resolving scatterers and in order to obtain a high accuracy in the measured crossrange positions, or utilizing crossrange resolution only for resolving scatterers but not achieving high crossrange accuracy. We already pointed out that high crossrange accuracy requires some kind of averaging over motion or flexing cycles, which appears problematic. It will usually be preferable to choose the simpler way of selecting shorter but good imaging intervals, forgoing high crossrange resolution and perhaps accuracy. However, if a specific identification

problem should remain unsolvable with this simpler approach, an extension may be required. Section 4.4.9 will demonstrate what this involves.

Making full use of the motion characteristics of individual scatterers is rather problematic. For example, the scatterer associated with Figure 4.75 has a fluctuating phase, whereas for the scatterer of Figure 4.77 the phase varies in about a quadratic manner. An imaging interval centered on the minimum of the phase function near time zero in Figure 4.75 will also be a good choice for the scatterer of Figure 4.77, but at the time of the previous minimum the phase in Figure 4.77 is changing roughly linearly, implying a scatterer with about constant range rate. This amounts to a translation of the scatterer response in crossrange.

By selecting the described time segment about the stationary point near relative time 0.06 seconds of Figure 4.75, we obtain Figure 4.78. This image is good enough for identifying the two illuminated edges and determining special features, despite the flexing of the vehicle (but a slow turn on a smooth surface). We repeat the same imaging by shifting the center time to the previous minimum of the phase cycle of Figure 4.75, which gives the image of Figure 4.79. A comparison of Figures 4.78 and 4.79 shows little change for the relatively weak scatterers on the side of the vehicle, but more




Figure 4.78 First vehicle image.




Figure 4.79 Second vehicle image.

significant changes for the responses that are strong in Figure 4.78. This fact already indicates that these may be spurious sideband responses, because such responses change very rapidly with aspect angle. The third vehicle image is shown in Figure 4.80. We have to wait for the positional match to determine how significant the changes are; that is, to determine how important it is to select the best imaging time in this particular situation.

The tests we have so far performed show that the turning motion of the vehicle yields a good crossrange resolution in a time short compared with the motion cycle of the vehicle, which in this instance is dominated by bending rather than vibration because the vehicle is turning smoothly. Examination of fixed-range image cuts through strong responses shows that the centroid track was dominated by a single scatterer, which is well compensated, and that resampling and polar reformatting are not necessary. The only additional compensation needed before analyzing the images is to apply a fine phase compensation to each, removing phase curvature common to the strongest image peaks. We do not show the images with fine phase compensation because to the eye the (intensity) peaks with and without the fine phase compensation are indistinguishable.




Figure 4.80 Third vehicle image.

In analyzing a given peak observed in the intensity image, we must decide between the following possibilities: (1) the position of the peak does represent the position of a scatterer on the vehicle; (2) the position of a peak is determined by interference between scatterers, so that the positions of these scatterers must be derived via the TSA; (3) the peak represents a response that is somewhat smeared because of the phase-center motion of the associated scatterer (in which case the peak gives the scatterer position with a large uncertainty); and (4) the peak represents a spurious response from a scatterer that can be far removed. At the time of this writing, we did not have a good way of utilizing spurious responses for target identification. Thus, the task was to recognize and disregard the spurious responses and, to the extent possible, detect genuine responses partially masked by spurious responses. We will summarize the results of the image analysis, obtained with the procedures discussed in Chapter 1.

Our matching task is made more difficult by the fact that, as is often the case with real data, we have only imperfect ground truth. The vehicle is equipped to carry various corner reflectors and antennas on its roof, and the positions of these devices are variable and unknown for our data. Thus we

must concentrate on the equipment and features associated with the original construction of the vehicle. Since the vehicle was moving on a wide paved surface, so that ground bounces may be important, for a most detailed study we should (but did not) have information about the features on the underside.

The occurrence of spurious sideband responses depends on the design complexity of the vehicle and its motion behavior. Recreational vehicles have generally smooth shapes, but the vehicle is turning rather than proceeding along a straight path. Spurious responses thus may or may not be a problem. The analysis showed that many of the strong responses in the lower part of the image of Figure 4.78 are indeed spurious. The cause might be a feature formed by the combination of the spare wheel with the backside and the corner formed by the rear bumper. The other possibility is a complicated feature underneath the vehicle, illuminated via ground bounce. For our purposes the important point is that we want to ignore these spurious responses regardless of their origin, and want to try to detect genuine responses from the back of the vehicle.

Figure 4.81 shows the match of the feature template for the vehicle with the range/crossrange positions of the accepted responses from Figure 4.78. For most features the match is a good one, so that we will discuss only the special




Figure 4.81 Positional match for the first image.

points. The position of the left front corner of the vehicle is marked by a question mark in Figure 4.81. This corner is not visible, because for this aspect angle the start of the bumper is shaped like a bent waveguide. Thus we obtain the delayed response visible in the same crossrange gate, which is only the first in a series of unresolved delayed responses of decreasing magnitude, not shown in Figure 4.81. Features K and L are the backs of air conditioning units, whose shells cannot be readily detected at this aspect angle (51° off broadside). However, the wave penetrates into the units, so that delayed responses are observed. We do not have a matching feature for one of the readily observable responses. The feature is likely to be on the underside of the vehicle, where it can be observed via a ground bounce.

A close examination of Figure 4.81 indicates that it is unlikely that another vehicle might match the measured distribution of the features. Hence, even though one can and will augment the positional match with measured characteristics of some vehicle features and with length and width and other special features, in this case the positional match alone should be sufficient for vehicle identification. If the aspect angle is derived by tracking the vehicle, we have the vehicle length without making use of the rectangular shape of the vehicle. If we had better information on the devices on top of the vehicle, or if there were not so many strong spurious responses at the back of the vehicle, we would also obtain the vehicle width independently of the match. The good quality of the match of Figure 4.81 indicates that the flexing effects are either insignificant or that the imaging time was indeed well selected. We will see below that the latter is the case.

The positional match for the second image, presumably also formed at a good time, is shown in Figure 4.82. As seen from the legend, the aspect angle is smaller by about 9° , so that the vehicle is viewed 60° off broadside. There are several differences between the matches, even though the aspect angle change is only 9° . However, by far most of the measured responses persist from one image to the other.

Whereas the differences between the above two images are minor, the situation is much worse for the image at the intermediate time (previously judged to be a much worse imaging time), for which the positional match is shown in Figure 4.83. First, because of the relatively poor match for the features on the side of the vehicle, the measured aspect angle is in error. It should be between the angles for Figures 4.81 and 4.82. Although some of the feature positions are almost perfectly matched, the match is poor for a significant number. Moreover, there are a number of features for which there are no significant responses (higher than the background), and there are responses for which we cannot find a feature on the vehicle. These large




Figure 4.82 Positional match for the second image.




Figure 4.83 Positional match for the third image.

differences exist despite the fact that the imaging time falls between the times for Figures 4.81 and 4.82, which both show good results. The difficulties must be ascribed to the purposely poor choice of the imaging time when bending/flexing effects are significant. Note, however, that the significantly poorer results do not necessarily imply that identification is impossible. A good number of responses are still well matched, and identification will not depend solely on the positional match.

4.4.7.2 The Special Case of the Broadside Aspect

Target identification by means of radar is a difficult problem, and it is most difficult when the target is viewed at broadside. These difficulties are greatest for a moving ground vehicle viewed at broadside. Then clutter cancellation via DPCA processing is only partly successful, because the range rate of the vehicle is near zero, so that the vehicle returns get cancelled together with the clutter. Also, at the broadside aspect the visible scatterers tend to be concentrated at the single illuminated edge, so that range resolution is not very helpful. Moreover, it is difficult to perform the motion compensation when resolution is mainly in crossrange, which is the type of resolution that requires a good motion compensation. Targets also tend to generate specular flashes at broadside aspects, with their effective points of origin sometimes shifting in crossrange. It is problematic to utilize such flashes for target identification, and yet because of their strength and smearing they may hide important scatterers that could be used. We now use the RV to demonstrate that the problems can be overcome, and that target identification at broadside is indeed feasible. However, the situation is helped significantly because the vehicle is turning, not going straight.

Since the SAR system provides an observation time much longer than needed for a moving vehicle, we can track the vehicle motion in the SAR scene, in this case finding that the RV is turning roughly at 15° per second. If we choose an imaging interval of 0.3 seconds, for example, the vehicle will turn by about 5° during the imaging interval. At a wavelength of 2 cm we obtain a crossrange resolution on the order of half a foot. Thus we select a 0.3-second interval that includes the broadside aspect. If we cannot estimate a target's angular change by tracking its motion, because the target moves irregularly or stays near broadside for a long period, we must select an imaging duration by determining the target's crossrange width, as we did for the previous targets. When we compensate with the range centroid track and a subsequent Doppler centroid track over the 0.3-second interval, we obtain the image of Figure 4.84.




Figure 4.84 Image After Range and Doppler Centroid Tracks.

First, we must decide whether the image is usable for positional measurements on the scatterers. Are the strong responses in about Range Gate 0.5 meaningful, in that they allow measuring the positions of the associated scatterers, or do they represent a meaningless interference pattern? As usual, we *test by taking the transforms of individual peaks and examining the amplitude/phase patterns*. If more than two scatterers are found to contribute to many of the responses, we must try to improve crossrange resolution by improving the motion compensation. This is likely to be problematic, because many strong scatterers are within a single range gate. A poor motion compensation will generally cause their smeared responses to overlap so much that no single response can be isolated. The best hope for compensation is a scatterer in a different range gate.

As an example of the evaluation of the number of scatterers contributing to a response, in Figure 4.85 we show the transform of the (complex) peak in Range Gate 0.5 near Crossrange Gate 6, indicated by the arrow in Figure 4.84. The agreement between the two-scatterer model curves (labeled “m”) and the data curves show that the amplitude/phase pattern is a good




Figure 4.85 Transform of the peak in Range Gate 0.5 of Figure 4.84.

approximation of the ideal pattern from two interfering scatterers, so the two scatterer positions in crossrange can be measured. Similar measurements can be performed on all the major responses of the image, so the image is usable for positional measurements. Unfortunately, most of the weaker target responses must be dismissed because they do not sufficiently exceed the clutter background; there does not appear to be a general and reliable method of distinguishing between a target response and a ground clutter response. There are a few exceptions. The transforms of image cuts through some of the responses, such as those within the dotted rectangle in Figure 4.84, show backscattering flashes. Such backscattering can be taken as evidence of a target feature, particularly for a target near broadside, nose-on, or tail-on.

As already stated, because this vehicle was used for research, none of the scatterers on top of the vehicle except the air conditioners are reliable features. We measured the locations of all features on the near edge that, because of their wave-trapping shapes, should give rise to observable backscattering. This means the side of the front bumper, the four “corners” of the fenders, the wheels, the beginnings and ends of the compartments, and so forth. The resulting template match with the measured scatterer positions is shown in Figure 4.86. Considering the problems of clutter and




Figure 4.86 Positional match between measured and predicted scatterers.

backscattering at the broadside aspect, the match for the scatterers on the near edge is very good.

4.4.8 Procedure for Analyzing the Survey Image and Forming a Final Image

We reiterate that *the core operations involved in selecting an imaging time are the estimation of the image duration needed to achieve crossrange position accuracy useful for identification, and the determination of time intervals during which scatterers can be well compensated.* In order to make that determination, we may have to recompensate data and examine images of varying duration, as we did with aircraft (see Figure 3.17). The path followed through the flowchart is determined by the relative durations of the available dwell, the image duration necessary to measure crossrange positions with usable accuracy for identification, and the longest interval allowing a majority of strong scatterers to be well compensated. The relative durations are determined by the motion conditions. As discussed in Section 4.4.2 and shown in Figure 4.87, which is a repetition of Figure 4.24, different types of motion

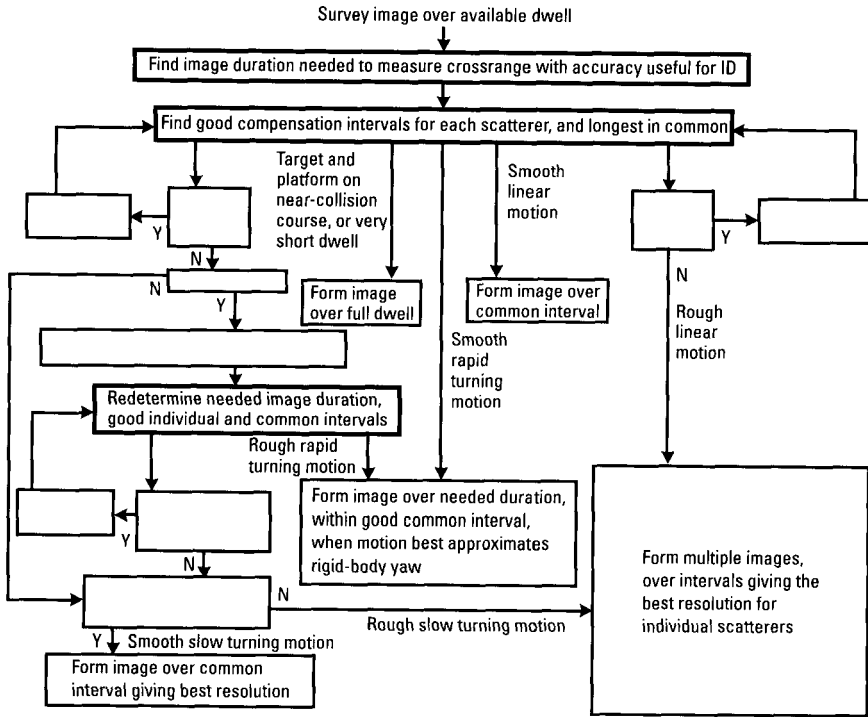


Figure 4.87 Core operations of image interval selection.

allow different measurements to be made, corresponding to the various branches that lead directly into endpoints in the flowchart. The remainder of this section is a more detailed discussion of the formation of a usable image, tracing the branches of the flowchart, as shown in Figures 4.88, and 4.89, which are repetitions of Figures 4.25 and 4.26, respectively.

The flowchart begins with our survey image of duration D . Our first action, summarized in the topmost box of the chart, is to estimate the range extent $R(0)$, the crossrange width $C(0)$, and the necessary image duration $d(0)$. Before describing how to make these estimates, we explain the notation used in the figure. Some variables have arguments, given in parentheses. The first argument, which can take the values 0 and 1, indicates whether a measurement was made on the full survey image or with a subinterval, respectively. The second argument, denoted by j , indicates that a measurement is repeated if the flowchart dictates recompensating a particular data interval. The third argument, denoted by k , indicates that a measurement must be performed on multiple scatterers in a given image.

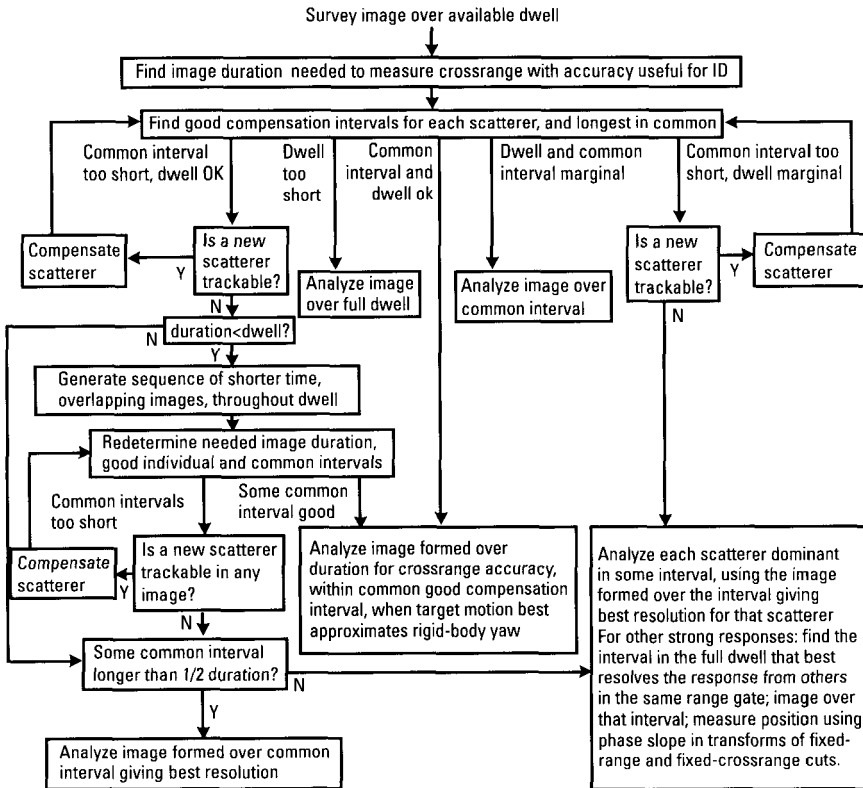


Figure 4.88 Flowchart for image interval selection.

The range extent R (in range resolution cells) is taken as the extent of range gates containing responses with amplitudes at least four times that of the noise and clutter background. If the image were perfectly compensated, this threshold would correspond to those responses to which we could apply the TSA. Our survey image will generally be poorly compensated, so responses will be smeared in crossrange, reducing their strength relative to the background. Thus, our estimate $R(0)$ is based on only some of the scatterers we hope to analyze. However, this does not imply that the estimate will be smaller than the actual target range extent. If the target is moving through range gates despite the centroid compensations used to form the survey image, we may overestimate the actual extent.

The crossrange width C is the median of target widths (in nominal crossrange resolution cells) estimated in the range gates containing strong

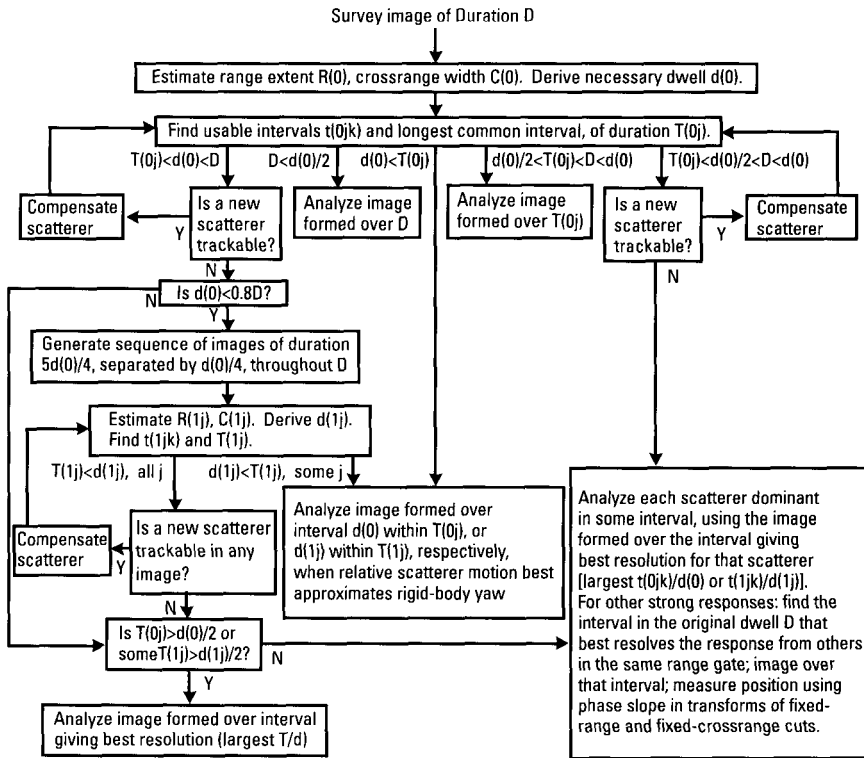


Figure 4.89 Flowchart with branching criteria.

responses. The width in each gate can be estimated in two ways. If individual responses are well enough resolved that their envelopes are bell shaped, we can take the width to be the spread between centroids of bell-shaped envelopes. If individual responses are less well resolved, we must instead take a Fourier transform over all the target responses in the range gate, and examine the amplitude of the transform. The modulation periods in the transform amplitude are the reciprocals of the scatterer separations. As is illustrated in the following sections and in Appendix C, the modulation periods can be determined directly from the transform amplitude, or by setting the transform phase to zero, taking an inverse Fourier transform, and measuring the separation of responses.

Having estimated the range extent R and crossrange width C , we can derive the image duration d necessary for us to make crossrange position measurements with useful accuracy for identification. For this purpose, we

require a minimum of about 15 resolution cells along the long edge of a ground vehicle, and about 5 along the short edge. If we assume that the target rotates uniformly during the imaging dwell, the required image duration is the dwell used for measuring R and C multiplied by the ratio of required to measured resolution cells. Allowing for measurement error, nonuniform rotation, and the possibility of further reducing the duration, we set d to twice this.

This brings us to the second highest box of the flowchart. We take fixed-range image cuts through prominent responses (those used to estimate R and C) in the survey image. We place a window around each bell-shaped response, take a Fourier transform, and search for intervals of constant amplitude without phase-slope breaks. These criteria are described quantitatively in Section 2.3.3., and are illustrated extensively in this chapter. We vary the boundaries of the window about the response, searching for the acceptable transform interval of longest duration. The longest duration interval for each scatterer k is stored in $\iota(0jk)$ (with $j=0$ for the centroid compensation). This interval is the longest one over which the scatterer can be well compensated. The longest common interval to a majority of scatterers is denoted by $T(0j)$. In order to minimize notation, we use the same symbols to denote an interval and its duration. In each instance, the meaning is clear from the context.

We have now reached our first branch point. We compare the duration of the common interval, the necessary duration, and the available dwell, in order to determine which of five branches to follow. We explain these five branches, in order of complexity. The simplest is the second from the left in the flowchart, labeled $D < d(0)/2$. This condition means that the available dwell is insufficient for us to measure the crossrange positions of responses to accuracy useful for identification. Thus, we must analyze an image formed over the full dwell D , hoping that crossrange resolution can at least help us resolve responses, thereby improving range measurements. Rather than analyzing responses in the survey image, we first attempt to track and compensate a scatterer. This branch corresponds to the vehicle and the radar platform being on a near-collision course, or to a very short dwell.

The second simplest branch is the fourth from the left in the flowchart, labeled $d(0)/2 < T(0j) < D < d(0)$. This condition means that the available dwell and the longest common interval both allow marginally useful crossrange position accuracies. We cannot afford to reduce the imaging duration very much, but are better served by imaging over the longest common interval than over the entire dwell. As explained previously, using reduced imaging intervals that allow good motion compensation is generally preferable to

using longer intervals that do not. We have already tracked several scatterers over the longest common interval. We compensate one of them and analyze the resulting image. This branch corresponds to fairly smooth linear motion (the translational Doppler change dominant over the irregular Dopplers), not too far off a collision course with the platform, or over a short dwell.

The third simplest branch is the central one in the flowchart, labeled $d(0) < T(0j)$. This is the branch likely to lead to the highest quality feature extraction. The condition means that the available dwell is sufficient for accurate crossrange position measurements, and that the necessary image duration is shorter than the duration of the longest common interval for good scatterer compensation. Although a majority of scatterers can be well compensated over $T(0j)$, a single compensation may not suffice if the target is not rigid. Also, measured crossrange positions for a flexing target will generally be a combination of scatterer height and azimuth. Thus, we search for that interval of duration $d(0)$ within $T(0j)$ that best approximates rigid-body yaw motion. This is done by determining when scatterer phase differences best scale with the scatterer separations. We compensate one of the already tracked scatterers, over the rigid-body interval, and analyze the resulting image. This branch corresponds to a smoothly and rapidly turning vehicle.

The fourth simplest branch is the rightmost one in the flowchart, labeled $T(0j) < d(0)/2 < D < d(0)$. This condition means that the available dwell is marginal for useful crossrange position measurements, but the common interval for good motion compensation is too short. However, as we discussed in Chapter 3 on aircraft identification, the intervals over which it appears possible to track and compensate individual scatterers are limited by the quality of the preceding compensation. If this compensation is too poor, scatterers may be so smeared in crossrange that it is impossible, in range gates with more than one significant scatterer, to define a window about an individual scatterer. Such a situation can be improved by tracking and compensating a scatterer that is dominant in its range gate. Thus, if we can track a scatterer that we have not previously compensated, we use its track to compensate the data (incrementing j), then repeat the determination of good compensation intervals for individual scatterers. This is represented in the flowchart by the loop back into the second highest box of the chart.

If no new scatterer is trackable, there is no common interval for good compensation that will give useful crossrange position measurements, and we move to the box at the lower right of the flowchart. This branch corresponds to rough linear target motion, with Doppler changes dominated by the irregular motion. We must measure scatterer positions in images formed over different intervals, and combine those measurements. For each scatterer k for

which we were able to determine an interval for good compensation, we form an image over the longest interval $t(0jk)$, with a compensation based on the scatterer track, and then measure the scatterer position. This interval selection will generally provide the best resolution for that particular scatterer.

We would also like to measure positions of strong responses for which we could not determine an interval for good compensation, albeit with lower measurement accuracy. The fact that we could not find a good interval suggests that each of these strong responses is likely produced by interference between two or more significant scatterers with different motions, due to flexing. As we demonstrate in Chapter 6, it is sometimes possible in such situations to iteratively track and compensate the strongest scatterer in a gate, suppress it, and move on to the next strongest. However, when the target motion is jerky, as it often is for ground vehicles, such processing will typically not yield higher accuracies than the following simpler procedure.

We take a fixed-range image cut through each strong response, define a window about the smeared multi-scatterer response, Fourier transform, select an interval of roughly constant amplitude modulation index, inverse Fourier transform, and measure the separations of the envelopes of major responses. We repeat this sequence, varying the interval chosen between Fourier transforms, to obtain the best separation of responses. This variation trades off improved nominal resolution versus response smearing due to differing motions. The latter effect is kept manageable by restricting the interval to one of roughly constant amplitude modulation index. We form an image over the interval giving the best response separation. Because at least one scatterer is not well compensated, it is unsafe to apply the TSA (see the end of Section 1.1.4). Rather, we determine scatterer positions by measuring phase slopes corresponding to constant amplitude intervals in the transforms of fixed-range and fixed-crossrange cuts through the image response of interest.

We next must combine measurements made on different images. After accounting for the different image compensations, some error in relative positions remains because the images were formed at different times. Between image times, the target may drift relative to the centroid compensation used for the survey image, and the target will rotate. However, these effects are small. Otherwise, the available dwell would be better than marginal for useful crossrange position measurements, as it is on this branch. The effects are best accommodated by increasing the minimum measurement uncertainties used in the positional match.

The most complicated branch from the second highest box in the flowchart is the leftmost, labeled $T(0j) < d(0) < D$. This condition means that

the available dwell is sufficient for useful crossrange position measurements, but the common interval for good compensation is too short. As with the fourth branch discussed, our first action along this branch is to attempt to recompensate the data. If we can successfully do so, we loop back up to the second highest box of the flowchart. If not, we try examining a sequence of images of shorter duration, hoping to reduce crossrange smearing of responses enough so that individual scatterers can be tracked over longer intervals. Note that this is the same procedure that we utilized for aircraft, as shown in the flowchart of Figure 3.17.

If the necessary image duration $d(0)$ is less than four-fifths the available dwell D , we expect that examining images of shorter duration than the full dwell may prove fruitful. If this is not the case, we check whether the longest common interval for good compensation $T(0j)$ is at least half the necessary image duration. If so, recalling that $d(0)$ may have some excess over the minimum required image duration, we proceed to the box at the lower left of the flowchart and analyze an image formed over the longest $T(0j)$, having first compensated one of the tracked scatterers. This branch corresponds to fairly smooth slow turning motion, with Doppler changes dominated by the turning. If no good common compensation interval is at least half the necessary image duration, we have the same situation as with the fourth branch discussed (rough linear target motion); we must measure response positions using multiple images, as in the box at the lower right of the flowchart. This branch corresponds to rough turning motion, with Doppler changes dominated by irregular motion.

Returning to the case where the necessary image duration $d(0)$ is less than four-fifths the available dwell D , we examine images of shorter duration than the full dwell. Specifically, we generate images of duration $5d(0)/4$, separated in time by $d(0)/4$, throughout the available dwell. We treat each of these images as we did the survey image. That is, we estimate the range extent $R(1j)$, the crossrange width $C(1j)$, the necessary image duration $d(1j)$, good compensation intervals $t(1jk)$, and good common compensation intervals $T(1j)$. Because $d(1j) < D$, only two of the five branches applicable to the survey image are relevant to these shorter duration images, the third branch discussed and the current branch, and we proceed in the same manner as we did before.

If, for any of the shorter duration images, the common interval for good compensation is sufficiently long, we proceed down the branch labeled " $d(1j) < T(1j)$, some j ," and select that image duration $d(1j)$ within $T(1j)$ that best approximates rigid-body yaw. This branch corresponds to rapid turning motion with significant irregular motion about the center of gravity.

If no common interval for good compensation is sufficiently long, we proceed down the branch labeled “ $T(1j) < d(1j)$, all j .” So long as a new scatterer is trackable in any image, we compensate that scatterer and reestimate our parameters.

When no new scatterers can be tracked, we check if some good common compensation interval is at least half the necessary image duration (some $T(0j) > d(0j)/2$ or some $T(1j) > d(1j)/2$). If so, we proceed to the box at the lower left of the flowchart and analyze an image formed over that interval with the largest ratio T/d (which should give the best resolution), first compensating one of the tracked scatterers. As already mentioned, this branch corresponds to fairly smooth slow turning motion. If no good common compensation interval has at least half the necessary image duration, we must use different images for different responses, as in the box at the lower right of the flowchart. As already mentioned, this branch corresponds to rough turning motion. In contrast to our previous cases of rough motion, the dwell available in this case may be long enough for more than marginal cross-range resolution. This implies more difficulty in combining measurements in different images. For comparable values of d and D , we can proceed as before, increasing the minimum measurement uncertainties used in the positional match. If $d \ll D$, we must register measurements in different images. This requires that, in each image, we measure the positions of some scatterers common to different images.

4.4.9 More Sophisticated Motion Compensations

On several occasions we stated that in cases of complicated vehicle motions it may be preferable in practice to forgo any attempt at accurately measuring the crossrange positions of scatterers, and to use only their range positions for the positional match. In this section, we provide an indication of the kind of processing needed if such a simplified approach is not adopted. Much can indeed be done, and a more sophisticated approach probably also can be automated, but the complexity is huge. Also, the case analyzed below is still relatively simple in that the vehicle is moving slowly on a straight road.

This particular vehicle is a refueling truck for aircraft, and it is moving on a paved road. The motion thus is relatively smooth. Moreover, the speed of the vehicle is only about 15 miles per hour, which makes its motion even smoother. On the other hand, the radar views the vehicle close to tail-on and the vehicle is moving along a straight line, so that a long dwell time is needed if a useful degree of Doppler resolution is to be achieved. If the vehicle should have even small yawing motions, they will have to be taken into

account in imaging. Under these conditions any yaw motion is unlikely to be large enough to facilitate crossrange resolution, but it may be large enough to corrupt the image if not taken into account. We will start with an imaging interval of 2.7 seconds, which in this instance is a good fraction of the available dwell time. Extending the imaging interval by perhaps a factor of two will not change the situation significantly.

The image of the vehicle after range centroid and Doppler centroid compensation is shown in Figure 4.90. The image may appear to be a two-dimensional image of a rectangular vehicle, but this impression is caused by the fact that due to the poor motion compensation the responses are about equally smeared in crossrange. It is a matter of the aspect angle. In order to estimate the target width and, hence, the necessary imaging dwell, and to select an appropriate imaging interval, we must examine fixed-range cuts through scatterers. This also will enable us to analyze the motion behavior of the vehicle to determine whether it moves rigidly (flexes by at most a small fraction of the wavelength), in which case the motion compensation is much simpler than when flexing is important. If the vehicle moves rigidly,




Figure 4.90 Peaks plot image of the vehicle after range and Doppler centroid compensation.

compensation can be based on the motion of just two scatterers. If not, we must analyze responses throughout the vehicle image.

We first consider the strongest response, the one in Range Gate -0.1 . The image cut in this range gate is shown in Figure 4.91. The question to be considered is whether the entire broad response is generated by a single scatterer, because of smearing in crossrange due to poor motion compensation, or whether these are responses from several scatterers. As always, this is judged using the transform of the image cut. The amplitude function does show significant breaks, but not deep enough to destroy the dominance of a single scatterer. The nearly regular amplitude modulation indicates the presence of a second scatterer, and from the depth of the amplitude minima one can estimate the phase jumps introduced at the positions of the amplitude minima to be on the order of 0.2 cycles (see Appendix A). If we mentally remove these phase jumps, we find that the shape of the phase function is not strongly changed by the phase jumps. The conclusion thus is that the phase slope roughly switches between four different values. We could also arrive at this conclusion by measuring phase slopes at the times of amplitude maxima. Since phase slope is Doppler or crossrange, the response of the major scatterer will appear in four parts, each part centered in one of the four




Figure 4.91 Image cut in Range Gate -0.1 .

crossrange gates determined by the phase slope. A measurement of the phase-slope values over the intervals indicated by dotted vertical lines gives, from left to right, Crossrange Gates 7.0, -3.8, 4.6, and -0.1. Figure 4.91 shows that these crossrange gates cover the main part of the smeared response.

The modulation of the amplitude of the transform shows two strong scatterers separated by about four nominal crossrange resolution cells, with a weaker scatterer separated by about twice that from the strongest. We cannot reduce the imaging interval appreciably and still hope to make crossrange position measurements with useful accuracy. Using the procedure we described in the previous sections, we would not attempt to compensate an interval containing a phase-slope break, so would have to settle for identifying on the basis of range measurements only. Here, we discuss compensating the entire interval.

Having determined the behavior of the dominant scatterer in Range Gate -0.1, we must check scatterers in other range gates. If the behaviors are all the same or they are smoothly changing with range separation, the vehicle is moving as a rigid body. If they are different, we must individually compensate the scatterers. We take image cuts, as in Figure 4.91, in the range gates of the other strong scatterers in Figure 4.90, Gates -17.6 and -11.4. The transforms are found to have very similar characteristics to Figure 4.91. In Figure 4.92, we show the phase functions of the transforms of the image cuts in the three range gates of the major scatterers. For our crude comparison we must ignore at least the major phase jumps, the one near the center of the middle plot and the one at the right side of the bottom plot. We then find that the bottom two phase functions are nearly alike, except for a displacement of the maxima of the phase functions. On the other hand, the phase function of the top plot is radically different. The expression radical is justified because the phases swing over about one cycle. This implies a large degree of response smearing, as verified by Figure 4.90.

The question now is whether or not these phase functions represent rigid body rotation. In other words, does the phase change smoothly with range? Figure 4.93 shows the phase functions of the transforms of the scatterers in Range Gates -4.4 and -7.5. The phase for Gate -7.5 behaves like that for Gate -0.1 (top of Figure 4.92), while that for Gate -4.4 behaves like that of Gates -11.4 and -17.5. Abrupt changes in the phase slopes that do not coincide in time for all parts of the vehicle, indicate nonrigid body motion. This completes the analysis of the motion behavior of the vehicle. Based on the three phase functions of Figure 4.92, which characterize the motion




Figure 4.92 Phase functions of the transforms of Range Gates -17.6 (bottom), -11.4 (middle), and -0.1 (top).

behavior over the length of the vehicle, we must now decide which specific imaging interval to use.

The simplest choice of the imaging interval would be one for which the phase functions of all the scatterers are about linear, since this implies proper compression of the responses. Let us for the moment ignore the fact that the intervals with linear phase in Figure 4.92 do not coincide, and consider one phase function at a time.

The question becomes, is any interval with linear phase long enough to allow detecting scatterers other than the dominant one, or would we have such low crossrange resolution that only the range profile of the target could be measured? In an application in which crossrange resolution is obtainable, we must not be satisfied with a range profile. We must determine whether




Figure 4.93 Phase functions of the transforms of Range Gates -7.5 (bottom) and -4.4 (top).

the amplitude functions contain the regular modulation pattern indicative of a second or perhaps third scatterer. The amplitude modulation in Figure 4.91 indicates a scatterer spread of about eight crossrange gates. The other fixed-range image cuts have similar transform amplitude modulation. We thus cannot afford to reduce the imaging duration significantly. However, the longest interval when all three phase functions of Figure 4.92 are linear is only about 10% of the imaging duration. We would obtain only the enhanced range profile of the target.

Instead of reducing the image duration, we must process over an interval extending over the flexing-induced phase variations, which must be removed (without also removing the interference-induced phase variations) by an additional motion compensation. How long a processing interval should we choose? The best approach is to start with the entire interval, and reduce it if problems appear, but only to the degree forced by the problems. For our demonstration we will process over the entire 2.7 seconds, evaluate the result, and reduce the imaging interval, if necessary.

We select the range gate considered earlier, which is Range Gate -0.1 , with the amplitude and phase functions of Figure 4.91. Figure 4.92 shows that, although the phase functions of the other two major scatterers differ

from the one in Range Gate -0.1 , all three have similar motion characteristics, and hence should be analyzable in the same way. If we wanted to generate a focused image, and if some problem should force a reduction of the imaging interval for one of the range gates with strong dominant scatterers, the reduction would have to be applied to all range gates. Since we want to “focus” responses instead of the image, and must use different motion compensations over the target extent, *we can use different imaging intervals in the different range gates*. Our association of measured responses with target features must, of course, allow for the different motion compensations and imaging intervals. They can be incorporated into the deformable template approach.

The large excursions of the phase function in Figure 4.91 are caused by the motion compensation residual. If only a single scatterer were present in this range gate, the phase function would truly represent the motion of the scatterer. Thus we could fit a flexible spline function to the phase function of Figure 4.91, and use it for motion compensation. However, the wiggles in the amplitude and phase functions imply that the response is not due to a single scatterer. For a proper motion compensation, we then must fit the spline only to the phase function due to the dominant scatterer, but not to the modulations caused by other scatterers. This is generally not possible to the desired accuracy if more than two scatterers are involved, so that phase-slope tracking cannot be used. The fitted spline will pseudo-periodically deviate from the phase function of the dominant scatterer, and these deviations will cause high crossrange sidelobes. They will generate spurious responses where there are no scatterers, and perhaps also mask weak scatterer responses.

Before we can proceed, we must again test the situation. The phase function of Figure 4.91 contains phase jumps introduced by other scatterers, but they are difficult to recognize. Thus, we try a motion compensation based on a spline fit to the phase function as it is, without correcting the phase jumps. This compensation yields a compressed dominant response. However, the secondary responses are so low that we cannot readily distinguish between genuine responses and sidelobes generated by the errors in the spline fit. When the same procedure is repeated after attempting to correct the phase jumps, the result is not significantly different.

We could attempt to use the phase-slope tracking procedure to determine the component of the phase function of Figure 4.91 introduced by the dominant scatterer. However, with amplitude and phase functions of the type shown in Figure 4.91, governed by both range cell wander and interference among three scatterers, we are unlikely to achieve better results than via

the phase tracking. Instead of pursuing such a possibility, we choose another approach in our illustration of a complicated motion compensation.

Since one scatterer is clearly dominant, we can take the FFT of the amplitude function alone, which does not contain effects from the motion compensation. This will generate responses due to Doppler differences between all scatterers; but if one scatterer is dominant, the only significant responses will be at the Doppler differences between the dominant scatterer and secondary scatterers (or secondary scatterers simulated by any periodic drift in the range gate of the dominant scatterer). Thus we obtain the separations of secondary scatterers from the dominant one, but because the FFT of the amplitude function is symmetric, we do not know whether the secondary scatterers are at lower or higher crossranges than the dominant scatterer.

The FFT of the amplitude function, or the noncoherent transform, is shown in Figure 4.94. The figure contains secondary responses indicated 3 and 4.5 gates away from the dominant response. However, the noncoherent transform of Figure 4.94 does not provide the information as to whether the genuine responses are on the left side, the right side, or on both sides of the dominant response. The solution lies in understanding the source of the problem. By examining Figure 4.91 we see that the major secondary response appears at about a time of -0.1 seconds (the periodic amplitude modulation




Figure 4.94 FFT of the amplitude function of Figure 4.91.

appears). Thus, at least in this gate, we should process only from -0.1 to 0.5 seconds.

We use this section of Figure 4.91 and remove the phase jumps occurring at the times of the amplitude minima. The result, together with a spline fit, is shown in Figure 4.95 on a renormalized scale. When the phase is compensated with the indicated spline and the transform is taken, we obtain Figure 4.96. Now we clearly see a genuine response about two gates to the left of the main response, consistent with the three-gate separation we obtained in Figure 4.94 from the entire 2.7 seconds. The low-level responses in Figure 4.96 are sidelobes generated by the motion compensation. As verification of the compensation, we apply it to Range Gate -4.4 , and obtain a perfectly focused response. As we would expect from the phase functions of Figure 4.93, the compensation does not work well for Range Gate -7.5 . An alternative approach to resolving the signs of the separations between dominant and secondary scatterers, usable when the phase jumps are more difficult to remove, is discussed in Appendix C.

We perform the same type of processing in the range gates with the other strong responses, for which the phase functions are shown in the




Figure 4.95 Spline fit for the shortened time interval after removal of phase jumps.




Figure 4.96 Transform of shortened time interval after phase compensation.

middle and lower plots of Figure 4.92. We note that the middle plot has one large phase jump (which correlates with an amplitude minimum), which must be removed before fitting a spline to the phase. The large phase jump in the bottom phase function also must be corrected before the fit. This allows us to determine the scatterers in the corresponding range gates. We will not proceed to the actual comparison with the locations of the scatterers on the vehicle because we do not have sufficient information on the vehicle.

Having established motion compensations in range gates with dominant scatterers, or just in range gates where the conditions permit good motion compensations, we can use the same motion compensations in nearby range gates. By measuring the widths of the responses we can determine whether or not a motion compensation is good enough to permit measurements on the scatterers. Beyond the range interval within which a particular motion compensation is usable we have to go through a new motion compensation process.

To show how small the range interval accommodated by one motion compensation can be, in Figure 4.97 we show the peaks plot of the vehicle when the motion compensation of Figure 4.95 is applied to the entire reduced duration image. The secondary response identified in Figure 4.96 is




Figure 4.97 Peaks plot for a compensation in accordance with the scatterer in Range Gate -0.1 .

immediately to the left of the dominant response in Range Gate -0.1 , the strongest response in the image. These two responses are well focused. However, the responses in the other range gates are drastically defocused, because of the flexing of the vehicle. We note that the responses in Range Gate -4 are also focused by the same compensation as used for Range Gate -0.1 .

The vehicle considered in this section is a slowly moving truck on a smoothly paved road, yet the processing required if the crossrange positions of scatterers are to be utilized in the positional match is extraordinarily complicated. In cases not so benign it will be even more complicated; we have not investigated any such cases at the time of this writing. The practical question is whether one would ever want to use such complicated processing in an automated operational system. Certainly, the answer is no if reliable identification of moving ground vehicles can be achieved with the much simpler procedures discussed earlier. In our opinion, it can. Aside from the question of practical need, we think that even such complicated processing can be fully automated.

4.4.10 Procedure for Automated Identification of Moving Ground Vehicles

In the preceding sections, we have given an indication of the wide range of conditions under which moving ground vehicles must be identified. Without referring to the reasons why a particular identification method must be used in a given situation, we now briefly list the processing steps that need be implemented in an automated identification system.

4.4.10.1 Step 1: Forming a Survey Image

If the moving ground vehicle is in a SAR scene, its (smeared) image is assumed to have been detected. The processor excises the image from the SAR scene, and uses the standard motion compensation (Range and Doppler tracking of the entire vehicle) to form a survey image. The processor takes fixed-range image cuts in the gates with prominent responses and over the range extent of the image. From the separation of the centroids of the smeared responses, or from the modulation periods of the amplitudes of the transforms, the processor estimates the crossrange width of the target. It also measures the range extent of the vehicle image. From the range extent and the crossrange width, taking into account the types of ground vehicle that are of interest in a given application, the processor estimates how long an imaging interval would be needed to implement a crossrange resolution that would allow measuring crossrange positions of scatterers with reasonable accuracy.

4.4.10.2 Step 2: Generating the Final Image

The goal is to find the best subinterval of the survey image that will allow forming a final image in which the crossranges as well as the ranges of the scatterers can be measured. If such a subinterval does not exist, the goal is to find a subinterval that will allow the highest accuracy in the measurement of the ranges of the scatterers. This is done as follows.

The processor takes an image cut in the range gate of a prominent (smeared) response, places a window over the bell-shaped response, and takes a transform. If the amplitude of the transform is not sufficiently constant, the boundaries of the transform window are shifted until the best result in terms of constant amplitude is achieved. The processor stores those time intervals in which the amplitude constancy is good and the phase is smooth, without a break in the phase slope (these criteria are described quantitatively in Sections 2.3.3 and 4.4.4.2). The same process is repeated for other range gates with prominent scatterers, over the range extent of the target. The processor then selects the longest common interval over which the transform

amplitudes are roughly constant and the phase functions are smooth, without a break in the phase slope. If the duration of the interval exceeds the minimum imaging duration as estimated under Step 1, the processor selects that subinterval of the minimum imaging duration during which the target motion best approximates rigid-body yaw. The processor tracks a scatterer in range and Doppler, performs a motion compensation, and forms an image over this subinterval.

If the duration of the common interval is less than the minimum imaging duration, which in turn is not much larger than the available dwell, the processor checks whether recompensating the data yields a longer common interval. If the best common interval is still too short, and if the minimum imaging duration is significantly shorter than the available dwell, the processor checks whether examining shorter duration images yields a longer common interval. Whenever a sufficiently long common interval is found, the processor selects that subinterval of the minimum imaging duration during which the target motion best approximates rigid-body yaw. The processor then range- and Doppler-tracks a scatterer, performs a motion compensation, and forms an image over this subinterval.

If no common interval was at least as long as the minimum imaging duration, but some common interval was at least half that long, the processor selects the common interval that was the largest fraction of the minimum imaging duration. The processor range- and Doppler-tracks a scatterer, performs a motion compensation, and forms an image over this interval.

The processor further examines any image formed over a selected interval, analyzing the transforms of fixed-range image cuts through prominent scatterers (presumably the same as used before). By comparing the phase functions of different scatterers, the processor determines whether resampling or polar reformatting are necessary and, if so, carries them out. It then removes any common quadratic phase component from the image.

4.4.10.3 Step 3a: Measurement of Two-Dimensional Positions

This step is employed only if Step 2 selected a common imaging interval of duration exceeding half that estimated as the minimum for measuring two-dimensional positions.

The processor measures the two-dimensional positions of scatterers associated with the major responses of the image. It then finds the leftmost and rightmost of these scatterers in each range gate, as well as the lowest-range scatterer in each crossrange gate. The scatterers that are extremal in both range and crossrange generally lie along the illuminated edges of the target. The processor notes any doubly-extremal scatterers. If the target is near

nose-on, tail-on, or broadside, there will be few doubly-extremal scatterers, but many singly extremal scatterers at the same range or crossrange. If the processor finds such a distribution of scatterers, it treats them as doubly-extremal. The processor finds the illuminated edges by fitting straight lines to the doubly-extremal scatterers.

If the range extent of the long illuminated edge is at least 15 gates, the processor estimates the crossrange width of the vehicle by finding the scatterer farthest from the long illuminated edge in each range gate. If the width is more than 10 crossrange gates, the processor reduces the imaging interval proportionally. If the range extent is less than 15 gates, the processor estimates the crossrange width of the vehicle by finding the maximum scatterer separation in each range gate. If the width is more than 20 crossrange gates, the processor reduces the imaging interval proportionally. Proportionate interval reduction is done by examining transforms of fixed-range image cuts and choosing the best subinterval on the basis of the amplitude/phase pattern.

4.4.10.4 Step 3b: Measurement of Range Only

This step is employed only if Step 2 selected a common imaging interval shorter than half that estimated as the minimum for measuring two-dimensional positions.

The processor begins with those image responses that are dominant in their range gate over some time interval, as determined in Step 2. Each such dominant response is analyzed in an image formed over the interval that gave best resolution (highest ratio of good compensation duration to minimum imaging duration) for the response, in an image formed with a compensation based on the response. The response is analyzed with the two-dimensional TSA.

Other strong responses, not dominant in their range gates, are analyzed differently. For each of them, the processor finds the imaging interval that gives the best crossrange resolution, as a compromise between insufficient resolution because of a short imaging interval and smearing of the responses because of an inadequate motion compensation. This is done by taking the transform of a fixed-range image cut through the response in the survey image, cutting out a shorter interval of roughly constant amplitude modulation index, and transforming back. When the smeared responses from major scatterers in the same range gate are best separated, the imaging interval is optimum. The processor measures each scatterer's range in an image formed over the appropriate interval for that scatterer, by measuring phase slopes during intervals of amplitude constancy in transforms of a fixed-crossrange cut through the response.

The range measurements in different subimages must be registered with one another. To this end, the processor calculates a scatterer's range as a combination of its location in an image and the compensation used to create that image. Registration also requires that some scatterers be recognizable from image to image. Therefore, the processor measures the positions of multiple responses in each image. Uncertainties in the positions of registration responses are folded into the estimated uncertainty of other responses.

4.4.10.5 Step 4: Analysis of the Final Image

If imaging with good crossrange resolution was successful, the image is analyzed in the same way as for a stationary ground vehicle, extracting scatterer positions, length, width, and other available features.

If the second imaging method had to be used, only the ranges of the scatterers are used in the positional match, and whatever few special features might be measurable.

4.4.11 Section Summary

The variety of conditions under which moving ground vehicles must be identified necessitates a complicated adaptive image-formation process. However, the same two processing and analysis steps are used under all conditions: estimating an appropriate imaging duration and determining intervals when individual scatterers can be well imaged.

Ideally, one should form the survey image by applying the standard motion compensation over the entire available observation time, estimate crossrange resolution by comparing the crossrange spread of the vehicle image with the typical size of a ground vehicle, and then form a new image over a shorter time interval chosen to give the desired crossrange resolution (not more than needed) during the target's smoothest motion. In reducing the imaging interval we want to select the best available data segment, with the amplitude functions of scatterers in different range gates as constant as possible and the phase functions as smooth as possible. The image-formation process is given in the flowcharts of Figures 4.24, 4.25, and 4.26.

In many situations, a good motion compensation that allows the use of the TSA and the measurement of accurate crossrange positions of scatterers is not possible. Crossrange position must then be determined from the phase functions derived from transforms over wider windows than with the TSA (if interference from other strong scatterers in the same range gate permits). Because the motion of the vehicle may still allow resolving the scatterers in

crossrange, this is far better than using an ordinary range profile, which gives the interference patterns from scatterers in the same range gates.

In general, even when the image selection procedure produces a good compensation interval of sufficient duration, compensation will not give an image equivalent to that from a stationary vehicle. In the best case, one can perform the same types of measurement as on stationary ground vehicles, although with lower accuracy. In the worst case, one can only measure the range positions of the dominant scatterers, with the range accuracy improved by the fact that crossrange “resolution” separates the responses in the same range gate, so that the range measurement becomes more accurate. Nevertheless, when the appropriate processing procedures are used, the same scatterers can be detected on a moving vehicle as when it is stationary.

Turning vehicles are generally easier to process, because the Doppler generated by the turn may dominate over the individual motions of the scatterers.

Since moving ground vehicles bend and “vibrate,” the scatterers have motions not fully determined by the rotation of the vehicle as a whole. The motion compensation thus should be varied over the vehicle. This is a difficult, perhaps impractical process. At least for a range resolution of (only) 1 ft, it appears more practical to use a single motion compensation as if the vehicle were rigid, and accept the fact that less information can be extracted from an image.

References

- [1] Rihaczek, A. W., and S. J. Hershkowitz, *Radar Resolution and Complex-Image Analysis*, Norwood, MA: Artech House, 1996.
- [2] Skolnik, M. I. (ed.), *Radar Handbook*, New York: McGraw-Hill, 1970.

5

Identification of Ships

5.1 Basics of Ship Identification

This section starts with a discussion of the differences between the identification of ships and aircraft or ground vehicles, and the basic problems of ship identification. The various types of image generated by the yaw, pitch, and roll motion of a ship are considered next, leading to the conclusion that ship identification must be based *primarily on undistorted topviews*, with some help from “sideviews.” This leads to a requirement that one be able to select an imaging interval that yields an image that is of the desired type and that has undistorted responses and a minimum of spurious responses. This can be done only by analyzing the ship’s motion. The principles of motion analysis are considered next, and lastly the actual measurements that must be performed on real data are discussed. The entire section should be of interest to readers wanting to gain a general understanding of ship identification.

5.1.1 Peculiarities of Ship Identification

Although the basic processing steps in ship identification remain the same as those for aircraft and ground vehicles, the details are rather different. A moving ground vehicle has yaw, roll, and pitch motions about its center of gravity, but they are generally too small and too irregular to utilize for generating a two-dimensional image of adequate quality for identification. An aircraft’s

pseudoperiodic yaw, pitch, and roll motions are much smoother than those of a ground vehicle. However, pitch motion is of little interest because an aircraft is largely a two-dimensional target and long-range identification implies a relatively small beam elevation or depression angle for the identifying radar. Only the vertical stabilizer may have sufficient height to be affected by pitch for small beam depression angles. The consequences of roll motion can be significant, but only for the wings. In fact, the potential severity of the consequences is one of the reasons wing scatterers should not be used in the positional match. The pseudoperiodic yaw motion of an aircraft, on the other hand, often is not strong enough to allow good crossrange resolution, yet modifies the translational yaw to such a degree that it limits the achievable crossrange resolution. By contrast, the yaw, roll, and pitch motions of ships are typically so large that, at such short wavelengths as X-band, *the changes in orientation due to these motions are far larger than needed for the desired crossrange resolution*. The issue with ships is not insufficient angular changes, but that ship motion can be so complicated that most images are highly distorted and degraded by spurious responses.

Another difference is that, whereas an aircraft can essentially be considered as a two-dimensional structure for purposes of a practical approach to identification, a ship must be treated as a true three-dimensional structure. The third dimension of height not only must be taken into account in imaging, but it allows one to discriminate ships on the basis of their superstructures. By comparison with moving ground vehicles, a ship not only has a much larger motion about its center of gravity, but also one that is smooth enough to utilize for good crossrange resolution. Thus we must approach ship identification differently from the identification of either aircraft or ground vehicles.

A further difference between ships on one hand and aircraft and ground vehicles on the other is that the variety of ships is far larger than that of aircraft or ground vehicles. It is relatively easy to obtain diagrams, photographs, and toy models (or even radar data) for aircraft and ground vehicles, so that a good database for identification can be generated. This is much harder to do for the huge variety of ships found in a littoral environment. Thus, we may have to be content with classification rather than identification of small ships.

5.1.2 Imaging of Ships

When a ship moves along some path, the aspect angle seen by the radar tracking the ship changes gradually, so that the ship appears to be executing a

slow and gradual yaw motion. In addition to the yaw associated with the translational motion of the ship there is generally also a pseudoperiodic yaw motion back and forth about the direction in which the ship is moving. A ship ordinarily also has pseudoperiodic roll and pitch motions. Thus, the nature of an image depends on the kind of composite motion that exists at the imaging time. We discuss the types of image that can be obtained with range/Doppler processing. The main purpose of discussing imaging under various motion conditions is to demonstrate that ship images can easily be too complicated for interpretation. Hence, when a ship is to be (automatically) identified, one of the requirements is that we must select the imaging interval so that an interpretable type of ship image is generated.

5.1.2.1 Range Profiles

When an image is formed at a time when the combined yaw, roll, and pitch Doppler is zero, range/Doppler imaging generates a one-dimensional image of the ship, its range profile. This happens when an image is formed at a time when all of the existing rotational motion components change directions, which is not a rare situation. A range profile of a real ship is shown in Figure 5.1. The responses in the image do not fall exactly along a straight




Figure 5.1 Ship image for zero yaw, roll, pitch Dopplers.

line, but the variation about a straight line amounts only to a small fraction of one crossrange gate. An image cut in the crossrange gate of the responses gives the (intensity) range profile of Figure 5.2. A test shows that all of the significant responses in the complex range profile can be readily analyzed by the TSA. However, accurate range positions alone are insufficient for identification. We note that the image dwell of 0.2 seconds used to generate the image of Figure 5.1 provides good crossrange resolution at other imaging times. In contrast, if we double the imaging dwell used for Figure 5.1, we obtain a range profile indistinguishable from that of Figure 5.2.

One should keep in mind that in order for an image to represent the range profile of the ship, the “image” of the ship must be oriented along the range axis, and the crossrange spread of the positions of the individual responses must be small compared with one crossrange gate width. This combination lets us recognize that the image represents a range profile, which implies that the ship has no yaw, roll, or pitch motion at the imaging time strong enough for forming a two-dimensional image.




Figure 5.2 Intensity range profile of a ship.

5.1.2.2 Yaw Images

Visualizing the type of radar image generated when one or more of the motion components about the center of gravity are effective, and when the beam depression angle is arbitrary, can be rather demanding. Thus we illustrate our discussion of the principles with the appropriate images for a simulated ship. Befitting this purpose, the simulated ship is rather simple, so that the essential points can easily be recognized. The ship is assumed to have three distinct height levels, the deck, the first level of a superstructure, and the second level. Point scatterers are used to define the shape of the deck. Both levels of the superstructure are represented by point scatterers at four corners. Figure 5.3 shows three views of the simulated ship, with axes labeled in feet. The simulation includes only the scatterers indicated by dots. The lines are included to show more clearly the three levels of the scatterers. In view of our negative statements concerning target simulation, we want to point out that the present simulation is intended merely to clarify the imaging conditions rather than to simulate the backscattering behavior of a target.

Before considering the types of two-dimensional ship image, we clarify some of the terms we use. By “aspect angle,” we mean the angle between the long symmetry axis of the ship hull and the projection of the radar line-of-sight (LOS) on the horizontal plane, with bow-on being zero aspect angle. Yaw motion is a change in aspect. “Depression angle” is the negative elevation of the radar LOS measured with respect to the horizontal. “Pitch” and “roll” are deck angles with respect to the horizontal. A “heeled” ship is one with nonzero roll.




Figure 5.3 Simulated ship.

We first consider the types of image a radar can form when only one of the three motions about the center of gravity is effective at the imaging time. Let us first assume a pure yaw motion. If the depression angle of the radar beam is zero, range resolution slices the ship into narrow cells in the horizontal plane. Since the Doppler of a scatterer is proportional to the distance from the center of the yaw motion projected onto a perpendicular to the line of sight, Doppler resolution gives crossrange resolution. Combined range/Doppler resolution establishes a grid of rectangular resolution cells in the horizontal plane, so that *a topview of the ship is generated by pure yaw motion*. This is the only kind of radar image that will resemble an optical image, although there will be differences in backscattering behavior. An optical sensor illuminating a ship from above establishes the same grid of resolution cells via resolution in the two angles. This is why the designation topview is justified, even though it is taken from optics.

As an example of a topview image, Figure 5.4 shows a yaw image of the simulated ship of Figure 5.3 for a beam depression angle of 0° . The eight dots near the center of the ship represent scatterers of the superstructure, with the four scatterers of the lower level connected by dotted lines, and the four scatterers of the upper level by dashed lines.




Figure 5.4 Simulated yaw image, 0° depression angle.

As the beam depression angle is increased from zero, (slant) range no longer corresponds to ground range. The plane in which the range resolution cells are established tilts upward with the line of sight. The important consequence is that scatterers at larger heights are shifted into closer range cells, with the amount of the shift proportional to scatterer height. A less important effect (for reasonably small beam depression angles) is that the yaw Doppler decreases with the cosine of the beam depression angle. This means a degradation in crossrange resolution as the beam depression angle increases. Thus, we have a gradual range shift with scatterer height, as well as a gradual widening of the crossrange resolution cell (all other factors remaining the same). The yaw image of the simulated ship of Figure 5.3 is shown in Figure 5.5 for a beam depression angle of 20° . Since $\cos 20^\circ = 0.94$, which is only 6% different from unity, the degradation of crossrange resolution is not visible to the eye. On the other hand, the responses of the superstructure are significantly shifted in range. Because of the nonzero aspect angle, the range shift means that the responses are translated closer to the edge of the ship near the radar.

For a beam depression angle of 20° , the relative translation in range of scatterers with a height difference of 5 m is $5\sin 20^\circ = 1.7$ m, which probably




Figure 5.5 Yaw image for 20° depression angle.

is not significant when a ship is viewed at a small aspect angle, but becomes significant at large aspect angles. The translation in range can be undone, at least in theory, if we can also obtain a sideview of the ship that shows the different heights of the scatterers, and if we can establish the scale factor for height. This may not be possible in practice. However, it may be adequate to correct the range shift based on some reasonable assumption about the height of the superstructure for a measured length of the ship. Since the real data used for our later illustrations were taken at very small beam depression angles, we cannot provide an actual example of the correction. In the absence of appropriate data, one cannot be sure to what accuracy a correction that appears to be straightforward can actually be performed.

At the extreme of a 90° beam depression angle, range resolution slices the ship along its height, whereas the capability for crossrange resolution disappears. We have a range profile along the height dimension. The corresponding image is shown in Figure 5.6. The image has degraded into a range profile that shows the three height levels of the ship.




Figure 5.6 Yaw image for 90° depression angle.

5.1.2.3 Roll Images

As the next case, we consider a ship with pure roll motion. At a zero beam depression angle, the roll Doppler is proportional to the height of the scatterer, so that the tendency is to generate a general type of sideview. This is illustrated by the image in Figure 5.7, which shows a pure roll image for a beam depression angle of zero and a roll angle near zero. Since the roll Doppler is proportional to the height of the scatterer, the scatterers at the same height all appear in the same crossrange gate. There is no shearing of the image in crossrange, as with yaw motion. We recognize the three height levels in Figure 5.7. The fact that one of the responses of the third height level is missing is due to inadequate resolution in the intensity image. In this instance the missing response can be found in the complex image when the TSA is used.

The roll Dopplers of scatterers on the superstructure decrease with the cosine of the beam depression angle, which is the same weak effect as with pure yaw. However, as the beam depression angle increases, scatterers at the same height no longer have the same Dopplers. A scatterer's position along the width of the ship contributes to its Doppler. The consequence is that




Figure 5.7 Roll image for 0° depression angle.

each height level is imaged similarly to the case of pure yaw motion. As an example, the pure roll image for a 20° depression angle is shown in Figure 5.8. On the right side of the image, we recognize the (somewhat distorted) outline of the deck. The four corner responses of the first level of the superstructure lie on the corners of the dotted parallelogram, and those of the second level on the dashed parallelogram. Resolution in the intensity image again is insufficient to display the fourth response of the second level. Note that the images of the three height levels are oriented along the range axis, because the scatterers at each level are at the same height. The image of Figure 5.8 can be obtained from the image of Figure 5.7 by translating the responses at each height level in crossrange. As with a yaw image, there is also the range shift of the responses due to the difference between slant range and ground range.

At the extreme of a 90° beam depression angle, we would obtain a range profile with the three different height levels if the ship had no width, just as in the case of the yaw image. Since it has width, the roll Dopplers spread the responses in each range gate in proportion to the spread in width of the scatterers at the corresponding height. Figure 5.9 gives the roll image




Figure 5.8 Roll image for 20° depression angle.




Figure 5.9 Roll image for 90° depression angle.

for this case. It can be generated by spreading the range profile of Figure 5.6 in crossrange, with the spread proportional to the width of the deck or superstructure.

In referring to images that show the different height levels, we used the term “sideview.” However, whereas the designation “topview” is justified for a yaw image, the designation “sideview” is not justified for almost all roll or pitch images. If the aspect angle of the ship is very small, range resolution slices the ship into narrow rectangles whose wide dimension extends across the width of the ship. The response of each scatterer is shifted in crossrange by an amount proportional to the height of the scatterer. Assuming that each range cell contains one dominant scatterer, we obtain a true sideview in this situation; but it is the exception. At larger aspect angles, range resolution slices the ship more along its length than its width. The dominant scatterer in each range cell can be positioned anywhere within that range cell, which means that its position can vary over a substantial part of the ship length. The image in that case has little relation to an optical type sideview image. In the extreme case of a broadside aspect, range resolution slices the ship along

its length, so that the radar image represents a kind of frontview or rearview. Hence, we will refer to images in which roll (or pitch) spread the superstructure in crossrange as “sea-level-view” images.

5.1.2.4 Pitch Images

As the third case, we consider a ship with pure pitch motion. The situation with regard to (slant) range and ground range remains the same as with yaw and roll, since range resolution does not depend on the ship’s motion. Thus, we again have the relative range shifts of the scatterers at different heights.

At a beam depression angle of zero and a pitch angle of zero, scatterers at the same height have the same Dopplers. The radar generates the range profiles of the scatterers at the same height, in this instance a range profile for the scatterers on the deck, another range profile for the scatterers at the second level, and a third range profile for the scatterers at the third level. The Dopplers are proportional to height, just as in the case of roll. For small nonzero pitch angles, *each of these profiles is sheared in crossrange*, with shearing proportional to a scatterer’s position along the length of the ship. Thus we obtain the image of Figure 5.10. It is very similar to the roll image of




Figure 5.10 Pitch image for 0° depression angle.

Figure 5.7, but is flipped in crossrange due to the phasing of the pitch motion, and has a small amount of shearing.

When the beam depression angle is increased from zero, the pitch Dopplers become significantly different even for the scatterers on the same height level, similar to the case of yaw. This means that the three range profiles of Figure 5.10 become much more sheared in crossrange. An example of how the image of Figure 5.10 changes when the beam depression angle is increased to 20° is shown in Figure 5.11. Note how the four scatterers at the second height level again define a (dotted) parallelogram with two sides oriented along the range axis, as in the case of a roll motion. Resolution again is inadequate in the intensity image to show all the corners of the analogous (dashed) parallelogram for the third height level.

At a beam depression angle of 90° , range resolution resolves the scatterers at the different height levels. For a given height level, the Dopplers of the scatterers are proportional to the distances of the scatterers from the pitch axis. The responses at each height level thus are separated in crossrange, depending on their separation along the length of the ship (rather than along the width, as in the case of roll). We obtain the image of Figure 5.12. It is the




Figure 5.11 Pitch image for 20° depression angle.




Figure 5.12 Pitch image for 90° depression angle.

same type of image as the roll image of Figure 5.9, except that length and width of the ship are interchanged.

5.1.2.5 Yaw/Pitch/Roll Images

Images generated by only one of the three motion components can be relatively easily interpreted, in particular the topview image from yaw. Image interpretation becomes much more difficult if all three motion components contribute to the image, in particular with real data. The difficulty of image interpretation increases by another level if the beam depression angle is not very small. We now consider images with all three motion components, both for zero beam depression angle and an angle of 20° . This leads to specific practical conclusions concerning imaging for ship identification.

An example of an image of the simulated ship when all three motion components are effective is given in Figure 5.13. The maximum angular deviation of the full motions about the center of gravity is 2° for the yaw, and 5° for both pitch and roll. The angular deviation over the imaging interval is 0.6° for the yaw, 1.6° for the roll, and 1.8° for the pitch. Although the beam depression angle is 0° in this example, the situation does not change




Figure 5.13 Yaw/pitch/roll image for 0° depression angle.

appreciably for moderate beam depression angles such as 20° . In Figure 5.13, we notice separate topview images of the deck, of the four scatterers of the first height level of the superstructure (dotted), and the four scatterers of the second level (dashed). When the imaging time is shifted in small steps, the essential consequence is that the responses of the superstructure move into the outline of the deck and then beyond it on the other side. Also, the direction and amount of shearing of the entire image changes.

5.1.3 Image-Time Selection

For a simulated ship as used in our illustrations, an image of the type shown in Figure 5.13 probably could be interpreted. Even with a different timing of the yaw, pitch, and roll motions, the responses still could be ordered into responses from the deck and those from the superstructure. One can form two images at different times in order to recognize the crossrange translation of the superstructure responses in relation to the deck responses, for better identification of the superstructure. Although the responses of the superstructure shift into different range cells when the beam depression angle is not zero, since the actual angle is known, it still should be possible to

unscramble the image. If one employs a somewhat more realistic ship model for which perhaps many of the scatterers of the nonilluminated edge cannot be observed and the superstructure is not divided into specific levels, the interpretation of an image formed at an arbitrary time will be at best very difficult. Now, if one deals with real data, the problem of interpreting an image that contains a combination of yaw, pitch, and roll effects appears to be unsolvable, in particular if the interpretation is to be automatic.

Normally, one would not draw such a negative conclusion without at least trying to interpret distorted ship images, which we have not done, but there is another reason one need not attempt to interpret an image of this kind. A real ship, like an aircraft or a ground vehicle, has dominant scatterers so large compared with the wavelength that there is a serious problem from the spurious responses generated by the wandering phase centers of the scatterers. This problem is particularly severe if the motion is complicated, either because it is erratic or the rotation axis is changing during the imaging interval. As will be illustrated by our examples of the imaging of real ships, spurious responses can severely degrade the quality of an image already distorted by a complicated motion about its center of gravity. In order to generate an image that allows identification or classification, it is necessary to select an imaging time at which the motion is smooth and about a fixed axis. *The central issue of ship identification thus is the choice of the imaging time*, both in order to facilitate the interpretation of the image and to avoid harmful spurious responses.

The purpose in selecting a specific imaging time is to generate an image from which one can extract the information needed to identify or classify the ship. Since ships move and maneuver much more slowly than aircraft or ground vehicles, tracking of the ship can provide a fairly accurate measurement of the ship's aspect angle. This allows one to derive the crossrange scale of a topview image, so that one can measure the actual length of the ship if bow and stern scatterers are observed. (The aspect angle cannot be supplied by the radar if the ship is docked, moored, or anchored. However, in this case the motions about the center of gravity of the ship should be so small that the SAR effect from the moving radar platform dominates. This provides the crossrange scale.) As with other types of target, the radar length is defined by the observable scatterers. If scatterers at the extremes of a candidate ship are not observable because there are no wave-trapping features at the bow or stern, we must use photographs and diagrams to define the ship's length as seen by a radar. The width of the ship also can be correctly measured after the crossrange scale is derived, provided responses from scatterers on the far side

of the ship can be recognized in the image. Also important, with the range scale known from the value of range resolution and the crossrange scale derived from the tracker data, for reasonably small beam depression angles one can measure the true positions of the observable scatterers within the outline of the ship, as with SAR images of stationary ground vehicles. Therefore, *a high-quality topview image is so useful that it should be the first objective of imaging.*

A good topview image shows the deck positions of the prominent scatterers of the superstructure, but not their heights. Because of the large number of ships in a littoral environment and the practical impossibility of obtaining and evaluating photographs and diagrams on all ships that may be encountered at a given time, *the measurement of the height "profile" of a ship also appears indispensable.* There are two ways in which this can be done. First, we can generate a sea-level-view image in addition to the topview image, and correlate the positions of the responses in the two images. This will show the heights of scatterers whose positions are measured in the topview image, at least for the prominent scatterers. In principle, one might be able to analyze the motions of a set of scatterers and from the analysis derive the height scale. We do not believe that the tracking of scatterers can be done with sufficient accuracy for such an approach to work for a real ship and most motion conditions (see Appendix E). A height profile with unknown scale factor is probably all that can be measured in practice, and it should suffice. Second, we can start from the topview image, and change the time slightly until roll/pitch effects appear on the superstructure. This will show relative height while minimizing the problem of correlating responses in two images (deghosting).

For the first approach, we need an observation interval long enough to find a time when the image of the ship is a topview and another time when it is a sea-level view. With the second approach, we only need a topview plus some time interval around it to have the topview "distorted." The second approach thus appears to be more efficient with respect to the required total observation time. This need not necessarily be so, because a relatively long observation interval, in terms of basic motion periods, may be needed in order to find a good topview image. However, small ships have relatively short motion periods, so that the overall observation time need not be too long. Whatever method is used, the primary requirement is to generate a good topview image.

Note that we have mentioned two approaches without suggesting which of the two should be preferred in practice. Our policy in this book is

to discuss those methods that can work in practice, without specifying a certain method when alternatives exist. Such decisions can be made only during the development process toward an operational system.

In principle, one could examine a sequence of images over the available observation interval in order to select a good topview image, or perhaps a topview as well as a sea-level-view image, making the selection on the basis of the target's appearance in the intensity image. However, besides the type of image, a second requirement is that *an image not have a serious problem with spurious responses*. Although the associated condition that the motion be about a fixed axis is automatically met if the image is a good topview, there might be a motion disturbance that degrades the image quality. It appears to us that examining a sequence of images is an unworkable approach in practice, if ship identification or classification is to be performed automatically. One problem is the generally complicated backscattering behavior of a target such as a ship, which may not let one recognize a high-quality radar image. Another problem is the range shift of the responses from the superstructure if the beam depression angle is not zero. A third problem is the image distortions from the composite motion about the center of gravity. A fourth comes from the spurious responses from a changing rotation axis, and, lastly, we have the spurious responses generated by a disturbance of the smooth motion. To make matters worse, the effects vary so rapidly that the sequence of images to be examined may have to be formed with overlapping imaging intervals. Given enough time, an analyst could examine these images in order to select one or more usable images of the ship, but the number of ships in a littoral environment can be so large that the entire process must be fully automated. The conclusion is that the appropriate imaging time should be selected in a manner that does not depend on the examination of the appearance of images.

To arrive at the appropriate procedure of image-time selection, we start from the requirement that *the motion of the ship must be about a fixed axis at the imaging time*. We also need an undistorted image that can be readily interpreted. As already discussed, the primary requirement is to generate the kind of topview associated with SAR images of stationary targets. This means that the responses from the scatterers of the superstructure not be translated in crossrange from their correct positions on the deck. The conclusion is that the image should be generated at a time when the rotation of the superstructure induced by roll or pitch is absent; that is, when the roll and pitch Dopplers are zero.

One might be tempted to conclude that the appropriate approach to finding this optimum imaging time would be to track as many scatterers as

needed, then to determine the ship's motion components from these tracks, and from an interpretation of the motion components derive the best imaging time. This might work with a simple model of a ship, such as we have used in our earlier illustrations. With real ships, on the other hand, accurately tracking a scatterer is a problem, and accurately tracking many scatterers is an unsolvable problem. An indication of the required tracking accuracy is obtained from the derivations in Appendix E. Even though we have not tried to solve the general motion measurement, it is our opinion that the achievable tracking accuracy is insufficient to derive usable results on the basis of simultaneous solution of equations of rigid-body rotation. A more practical approach is needed.

Such an approach exploits the fact that *one can recognize that a scatterer is located near the bow or stern, or on the superstructure*. Furthermore, scatterers in these areas are generally the best trackable on the ship. The relative Doppler of bow and stern scatterers depends only very weakly on roll motion. When this Doppler is zero, only yaw is effective and one can form a topview image. The relative Doppler of two superstructure scatterers with the same deck position but different heights depends only very weakly on yaw motion. When this Doppler is zero, only roll is effective and one can form a sea-level-view image. The weak third motion components may cause small distortions of the images, but the small refinements in imaging time needed to generate undistorted images can be estimated from the distortions themselves.

To implement the approach, we select a scatterer on the superstructure and compensate its motion so that the scatterer becomes stationary. Using the modified data, we select another scatterer of the superstructure at a different height, but as close in its deck position as possible to the first scatterer. Then we measure the changing Doppler of the second scatterer over the observation interval. When the Doppler goes through zero, we have the conditions for generating the best obtainable topview of the ship. This is a topview of the heeled, pitched ship. If the scatterers were all at the same height, it would also be a topview of the upright ship. However, differing heights will introduce some distortions, even for zero beam depression angle. If the heights are significantly different, so that the distortions become serious, we can make a first-order correction by estimating the height differences from another image that displays a sea-level view of the ship. The practical problem is to perform the required measurements with sufficient accuracy on real data.

If we want to measure the shape of the superstructure in a separate image, we must generate an image in which the roll/pitch of the

superstructure is effective. Since we want to form this image at a time when the motion is about a fixed axis, we must select a time at which the differential Doppler between scatterers at the same height and along the centerline of the ship is zero. This time is found analogously to the process of selecting the correct time for the topview image. Assuming for the moment that there is a stern scatterer on the centerline of the ship, and that the stern and bow scatterers are at the same height, we first motion compensate one of the two scatterers. In the compensated data, we then motion compensate the second scatterer, measuring the Doppler of the scatterer in the process. The image then is formed at the time at which the Doppler of the second scatterer is zero. If the two scatterers are not on the centerline of the ship and at the same height, a correction can be performed on the basis of the actual scatterer positions extracted from topviews and sea-level views, even though we will not have the height scale. There is a practical relation between the length of a ship and the height of its superstructure. Although there will be errors, we do not need the high positional accuracies for ships that we do for aircraft and ground vehicles, so that no correction or only a rough correction is needed.

The described procedure ideally utilizes tracks of two scatterers that have a small separation along the length and width of the ship, but a large height difference. It is often difficult to track two such scatterers. As a practically important alternative, we can use the tracks of the bow and stern scatterers to derive the track for any point on the line (in three-dimensional space) through the two scatterers. *We can replace one of the superstructure tracks with the derived track of a virtual scatterer whose location differs from that of the tracked superstructure scatterer primarily in height, with small differences in the other dimensions.* This alternative procedure becomes particularly important when tracking of scatterers is problematic, as it often is with real ships.

5.1.4 Principles of Analyzing a Ship's Motion

We first demonstrate the requisite measurements on the simulated data, and then proceed to real data. We want to measure the Doppler of the stern scatterer when the bow scatterer is stationary (or vice versa). This Doppler is a combination of yaw and pitch Doppler. We also want to measure the Doppler of a scatterer on the superstructure relative to another scatterer on the superstructure at a different height. This gives a combination of roll and pitch Doppler, but it is the combined Doppler that must be zero if a good topview image is to be generated. For simplicity, throughout all of what follows we will refer to the first measured Doppler as yaw Doppler, and to the

second as roll Doppler, even though both may be affected by the pitch. In other words, we assign part of the pitch effect to the yaw and the other part to the roll, since pitch generates effects analogous partly to yaw and partly to roll.

The principles of image-time selection and imaging are the same regardless of the beam depression angle, since our approach is directed toward finding the times at which the ship is not rotating about a horizontal axis (whether due to roll or pitch), and the times at which it is not rotating about a vertical axis. The value of the beam depression angle thus is irrelevant as far as image-time selection is concerned. However, at larger beam depression angles the scatterer responses are shifted in range in proportion to their heights, so that image interpretation becomes more difficult. We illustrate the procedures of image-time selection for beam depression angles of 0° and 20° . Since our real data were collected with low beam depression angles, we must use the simulated ship to demonstrate image-time selection at the higher beam depression angle.

Starting with a beam depression angle of zero, the peaks tracks for the simulated ship are shown in Figure 5.14 for maximum angular deviations of 1.5° for yaw, pitch, and roll. Even with simulated data, some of the peaks tracks are the poor tracks of unresolved scatterers. We want to select two responses on the deck at about the same height, near bow and stern, and two




Figure 5.14 Peaks tracks for simulated ship.

responses of the superstructure with close deck positions. *To recognize the height differences, we must form a (short term) sea-level-view image.* To obtain such an image, we should select two convenient scatterers near the bow and stern, determine the yaw motion, and then form an image when the yaw Doppler is zero. This will be a sea-level-view image, and hence will show whether or not the two selected scatterers are at sufficiently close heights. If not, we select another scatterer for the pair. This is the appropriate procedure for automated measurements. With the manual processing used for our illustrations, it is simpler to select an arbitrary imaging time, form the image, and check if it is a sea-level-view image. If so, we use it to estimate the relative height differences. If it is too close to a topview image, we can shift the imaging time a little and try again. The precise procedure does not matter, and in an operational system one will implement whatever method appears simpler. For our illustration, we choose the trial approach.

Since the peaks tracks indicate a motion period of two seconds, and we desire that the Doppler not vary too much over the imaging interval even for a survey image, we use an imaging interval of 0.2 seconds, or 10% of the motion period. Trying an imaging time of 1.6 seconds, indicated by the dotted rectangle in Figure 5.14, we obtain the image of Figure 5.15. The




Figure 5.15 Short-term image for scatterer selection.

strong stern and bow responses are on the long symmetry axis of the image, hence must be at nearly the same height, so that there is no question about selecting them for the yaw measurement. They are in Range Gates 42 and -42 at the chosen imaging time, indicated by solid arrows in both Figure 5.14 and 5.15. The best pair of scatterers for the roll measurement would be the strong one in Range Gate 6 (a solid arrow), and the weak one in Range Gate 4 (the dotted arrow). However, tracking the weak one presents difficulties. With a stronger scatterer in Range Gate 6 and a comparable scatterer in Range Gate 2, simple peaks tracking of the range profiles is problematic. This can be seen from the peaks track marked by the dotted arrow near Range Gate 3 in Figure 5.14. Thus, it is better to choose the other strong scatterer of the superstructure, which is in Range Gate -2 (a solid arrow). The separation of the two scatterers along the length of the ship is still only a small fraction of the separation of the bow and stern scatterers, which means that the deck positions of the two scatterers are reasonably close. We go into Figure 5.14, select the peaks tracks that fall in the above range gates at the imaging time of 1.6 seconds, and range- and Doppler-track the corresponding scatterers.

When the yaw and roll measurements are performed in accordance with the method demonstrated in detail in Section 2.3.2, using the simplest form of scatterer tracking, we obtain the yaw and roll Doppler curves of Figure 5.16, with the dashed curve giving the yaw Doppler. If we choose an imaging time when the roll Doppler is zero and the yaw Doppler is nonzero, we should obtain a topview image. At a time when the yaw Doppler is zero but the roll Doppler is not, we should obtain a sea-level-view image. From Figure 5.16 we arbitrarily select a specific time with zero roll Doppler at 9.08 seconds, with the corresponding 0.2-second image shown in Figure 5.17. Comparison with Figure 5.4, where only yaw motion exists, shows that we have indeed obtained an excellent topview image. Similarly, from Figure 5.16 we obtain an imaging time of 13.81 for an example of a sea-level-view image. The image for this time is shown in Figure 5.18. The differences relative to the pure roll image of Figure 5.7 are insignificant.

In the preceding example we selected good bow and stern scatterers, but the scatterers chosen for the superstructure do not have close deck positions. Thus, even with simulated data there can be a problem of selecting good pairs of scatterers, at least when we use the simplest method of tracking the peaks of the range profiles. The problem is worse for real data. As discussed above, there is an alternative procedure. We can track the bow and stern scatterers, and by interpolating between the two tracks, motion compensate in such a manner that the stationary point falls along the straight line




Figure 5.16 Yaw (dashed curve) and roll Dopplers for 0° depression angle.




Figure 5.17 Topview image.




Figure 5.18 Sea-level-view image.

connecting the two scatterers and within the range gate of a scatterer on the superstructure. Then there is no need to select a second scatterer for the roll Doppler measurement, because its purpose is only to serve as a stationary point. With our interpolation we replace the real stationary scatterer with a virtual stationary scatterer, so that the roll can be measured with respect to the virtual scatterer. The separation of the deck positions of the virtual scatterer and a trackable superstructure scatterer can generally be made much smaller than that of the superstructure scatterer and a second trackable scatterer. Hence, if there is an accuracy problem, the alternative approach should give more accurate results.

In Figure 5.19, the solid curve gives the roll Doppler derived with the alternative procedure, using the tracks of the bow, stern, and a high superstructure scatterer. The dashed curve is the roll Doppler curve from Figure 5.16, included for comparison. With the alternative procedure, the fourth track used to derive Figure 5.16 was replaced by a track of a virtual scatterer in the range gate of the high superstructure scatterer, generated by linearly interpolating between the bow and stern scatterer tracks. Aside from the irrelevant scale factor and a tracking problem near the fringes of the




Figure 5.19 Alternative roll Doppler, 0° depression angle, versus roll Doppler of Figure 5.16.

interval, the two roll Doppler curves are in excellent agreement. They have the same shape, with times of zero roll Doppler typically differing by less than 0.1 seconds. The fact that the two curves agree so well is an indication that choosing two scatterers of the superstructure that are not close in their deck positions did not lead to significant errors in the roll Doppler measurement, but this need not be true in general.

Figure 5.20 shows the range peaks tracks for the same simulated target motion as Figure 5.14, but with a radar beam depression angle of 20° . The bow and stern areas of Figures 5.14 and 5.20 (ranges less than about Gate -20 and greater than about Gate 20) are quite similar, because the scatterers in these regions are all at the same height. The change in depression angle compresses the range tracks in proportion to the cosine of the depression angle, only about 6% in this case. The tracks also differ slightly because roll and pitch are more effective, and yaw less effective, at the 20° depression angle. The central ranges of the figures differ more substantially, because the scatterers in this region are at three different heights. Each scatterer is shifted in range by its height multiplied by the sine of the depression angle. Therefore, the relative ranges of scatterers in the central part of the ship change as




Figure 5.20 Range peaks tracks, 20° depression angle

the depression angle varies. This change in interference conditions means that a scatterer may be trackable at one depression angle but not another.

When the motion measurement is repeated on the data for a 20° beam depression angle, with the four-scatterer procedure (using the tracks marked by arrows in Figure 5.20) we obtain the yaw and roll Doppler curves of Figure 5.21. Figure 5.22 shows an interesting section of Figure 5.21. Dotted rectangles indicate intervals used to form images shown in figures below. To begin with, the roll Doppler is zero at a time of 10.40 seconds, so that we should obtain a topview image at this time. The actual image is shown in Figure 5.23. The superstructure responses are not quite in the center of the deck, but are in the appropriate positions, as can be verified by a comparison with Figure 5.5. Thus, the combined roll/pitch Doppler is zero at this time.

The situation is similar with respect to sea-level views. Using Figure 5.22, we select a time of zero yaw Doppler at 10.9 seconds, obtaining the image of Figure 5.24. It is not a zero-yaw-Doppler image. We see separately the scatterers of the deck and of the two height levels, with the responses at each level separated in crossrange because of a residual yaw motion. If the imaging time is shifted by 0.2 seconds we obtain the image of Figure 5.25. This is clearly a (zero-yaw-Doppler) sea-level-view image. In




Figure 5.21 Yaw (dashed curve) and roll Dopplers for 20° depression angle.




Figure 5.22 Section of Figure 5.21.




Figure 5.23 Image at zero roll Doppler.




Figure 5.24 Image at 10.9 seconds.




Figure 5.25 Sea-level-view image.

practice, in order to get from Figure 5.24 to Figure 5.25, we can adjust the imaging time slightly until the crossrange spreading of the deck responses disappears. This should be automatically measurable even with real data. Imaging at other times of zero roll or yaw Doppler as determined from Figure 5.21 shows that Figure 5.24 represents the worst case of the timing error for this example.

The practical question is the following: Should one use such a simple procedure for imaging-time selection, and then *shift the imaging time slightly* (until the responses of the superstructure are roughly centered on the deck for a topview, and until the crossrange width of the deck responses is zero for a sea-level view), or should the procedure of image-time selection be refined until Doppler curves such as the ones in Figure 5.21 give the best imaging times accurately? Although we have relied on the former procedure, an answer requires extensive practical tests of identification. We will merely show that it may be possible to improve the measurement accuracy.

As a reference, consider the distorted topview image of Figure 5.26, formed at 10.15 seconds. The superstructure responses are offset from the center of the deck, but away from the side of the ship closest to the radar (the effect of the depression angle should shift them toward this side). The image




Figure 5.26 Image at 10.15 seconds.

shows the two corners of the stern, indicated by horizontal arrows. For a nonzero beam depression angle, the range difference between the two corners changes with yaw and roll, but not with pitch. Hence, if we measure the Doppler of one corner relative to the other, in the same way as for bow and stern scatterers, the Doppler describes a combined yaw and roll motion, but without pitch component. If we image at the time of zero Doppler, only the pitch motion will be effective. Similarly, consider the two responses approximately in Crossrange Gate 7.5 of Figure 5.26, indicated by vertical arrows. Their range difference is affected by pitch and roll, but only insignificantly by yaw (for small changes of aspect). Thus, the relative Doppler of these two responses will be a combination of roll and pitch, without yaw. Imaging at a time of zero Doppler will give a yaw image. The third case of no roll effects is obtained when responses along the edge are used for deriving the motion Doppler. The Doppler then is a combination of yaw and pitch Doppler, without roll. In the third case, when real data are used, we must check a sea-level-view image to verify that the scatterers are at the same level. This more general procedure allows us to form specific types of image by selecting the times at which the measured Doppler is zero. Of course, we must be able to track the relevant scatterers, which may be problematic in real data.

As an example of the procedure, deriving the Doppler of one stern corner relative to the other and selecting an image time at which the Doppler is zero gives the image of Figure 5.27. Within the measurement accuracy, there is no crossrange separation of responses at the same position along the length of the ship, because we eliminated the crossrange separation between the two corner responses of the stern. The height levels are correctly reproduced. Since we know the type of image that was generated, it can be interpreted even when real data are used.

When the two scatterers in Crossrange Gate 7.5 are used, as discussed above, and a zero-Doppler time is used for imaging, we obtain the yaw image of Figure 5.28. It is a topview image, obtained from yaw alone. In the third case, where scatterers along the edge of the ship are used, the same procedure leads to a roll image that has the same characteristics as the roll image of Figure 5.18. Again, the advantage of systematically selecting the imaging time is that we know the nature of the image generated. That may not be important for simulated data, but it is very important for real data.

The preceding results for a beam depression angle of 20° were derived by performing Doppler measurements on two scatterers at a time, regardless of what type of Doppler was to be measured. We now demonstrate that the




Figure 5.27 Pitch image.




Figure 5.28 Yaw image.

procedure in which a real scatterer is replaced by a virtual scatterer works just as well at the higher beam depression angle.

We now derive the roll Doppler from the three tracks of Figure 5.20 indicated by solid arrows, using the alternative procedure with three real and one virtual scatterer. Bow and stern scatterers were tracked in Range Gates -39 and 39 , respectively. The high superstructure scatterer was tracked in Range Gate -8 . A virtual track was generated in the range gate of the high superstructure scatterer. In this alternative derivation, we would like to use a virtual scatterer at the same position along the length of the ship as the high superstructure scatterer. However, because of the height-dependent range shift of the superstructure scatterer for nonzero depression angles, the virtual scatterer should be chosen at a greater range than that of the superstructure scatterer. In a second version of the alternative procedure, the virtual scatterer was chosen at the correct range. Figure 5.29 shows the roll Dopplers derived with the alternative procedure, with (dashed) and without (solid) taking into account the range shift of the scatterer. With the exception of the zero roll Doppler time near two seconds (poor tracking accuracy), the two Doppler curves are in very good agreement. The other times of zero roll




Figure 5.29 Alternative roll Doppler, 20° depression angle, and improved version.

Doppler agree within 0.1 to 0.2 seconds (5% to 10% of the average motion period). This agreement indicates that a correction for the height-dependent range shift need not be very precise. We can apply the correction by using a typical height for a given ship length and superstructure position.

Note that the roll Dopplers of Figures 5.29 and 5.19, the results for beam depression angles of 20° and 0° , are essentially identical. The times of zero roll Doppler are indistinguishable, and the shapes of the two curves differ only in small details. Thus, when the beam depression angle is increased to 20° , the change in the roll Doppler is minimal. Although this is to be expected on the basis of physical considerations, the differences between the peaks tracks of Figures 5.20 and 5.14 indicate that the processing for the two cases is significantly different. We have found that the three different types of measurement (four scatterers, three scatterers without range correction, three scatterers with range correction) measure the timing of the zero-Doppler crossing with good agreement, typically within 0.1 seconds (about 5% of the average motion period). However, in the context of Figures 5.23 and 5.26 we also found that a shift in the imaging time by 0.25 seconds is enough to change a highly distorted image into an undistorted image. Now, the difference between the two curves of Figure 5.29 is only a range correction of the

virtual scatterer, to take into account the range translation due to height. The differences in the timings of the Doppler zeros are very small. However, a check shows that the difference is enough to affect the distortion of the corresponding images. In other words, when the beam depression angle is still only 20° , applying the range correction gives perfect results for the simulated data. Even though the range correction cannot be perfect in practice because, in contrast to the simulated data, the height of the scatterers is not known, for a measured length of the ship we should be able to estimate the height of the superstructure with sufficient accuracy to apply a meaningful correction of the scatterer range.

5.1.5 Motion Analysis for Real Data

This section discusses the scatterer tracking accuracy needed to select imaging intervals and form well-compensated images, and the tests that should be used to judge the acceptability of a track. It also covers the type of compensation that is typically required. Section 5.2 gives illustrations of the image interval selection process.

The accuracy with which a ship's motion can be measured depends critically on how well one can track the four (or three, with the alternative procedure) scatterers needed for the measurement of the yaw and roll Doppler functions. The topic of scatterer tracking is treated in Section 2.3.2, and will be summarized here only as necessary. In a given situation we cannot simply track a scatterer; we must also know how well the scatterer was tracked, in order to be certain that the measurement of the motion behavior is sufficiently accurate or even meaningful. Thus, one must check each processing step; if a track fails any check, a new scatterer must be selected for tracking.

The following *checks must be performed on every scatterer that has been tracked* before continuing with the analysis of the motion behavior of the ship, whether the analysis is manual or automated. First, we must verify that the scatterer has at least remained in its range cell. This means forming an image on the basis of the track of the scatterer, taking an image cut in the range gate of the scatterer, and examining the amplitude and phase function of the transform of the (windowed) cut. If range tracking has kept the scatterer sufficiently well within its range cell, then the amplitude function will not remain near the background level for any extended period. Short drops in the amplitude due to interference are acceptable, as long as the minima are sharp rather than flat. In the latter case, the tracking has been grossly unsuccessful. We must improve or iterate the tracking, always performing the same test on the amplitude function.

If the amplitude function does not have any extended minima, the scatterer has remained in its range cell. We next consider whether the track contains contributions from an interfering scatterer. An amplitude function that does not have any low minima is the sign of high tracking accuracy on a single dominant scatterer. The phase function then represents residual scatterer motion. In the best case, it will be essentially linear without significant modulation. Nonlinearity is acceptable if the phase is at worst smoothly curved without any abrupt changes that exceed 0.1 cycles.

Minima of the amplitude function that reach relatively low levels imply phase jumps, meaning the Doppler track was fitted to data including variations generated by interference rather than motion. If the fit was flexible enough to follow these variations, it will have corrupted the data, perhaps too much to correct. If the amplitude minima are low, we attempt to use the phase-slope procedure of tracking the phase of the dominant scatterer, as described in Section 2.3.2.3. However, it is preferable to find a scatterer that can be better tracked, and to use the improved tracking procedures discussed in Chapter 2. Whether we employ phase tracking or phase-slope tracking, we must measure the residual phase variations and compensate any large enough to warrant it.

The criteria for how large a residual phase variation is acceptable are its effects on the motion compensation and on the measured Doppler function. Evidently, phase variations imply Doppler variations, and unless the residual Doppler variations are small compared with the yaw and roll Dopplers to be measured, these Dopplers will not have been measured with sufficient accuracy. For such an estimate, we can model the pseudoperiodic residual phase variations as a sine wave, writing

$$\phi(t) = \phi_m \sin 2\pi f_m t \quad (5.1)$$

where ϕ_m is the maximum phase deviation and f_m is the modulation frequency of the phase fluctuations. We obtain the Doppler by differentiating (5.1),

$$v(t) = -\frac{1}{2\pi} \frac{d\phi(t)}{dt} = -f_m \phi_m \cos 2\pi f_m t \quad (5.2)$$

The maximum excursion of the absolute Doppler is

$$v_m = f_m \phi_m \quad (5.3)$$

By measuring the residual phase variation ϕ_m and the cycle length $1/f_m$ we obtain v_m . This quantity must be small compared with the measured maximum yaw and roll Doppler for the measurement to be useful. For an acceptable phase compensation, the residual phase variation must be less than about one-tenth of a cycle, so that $\phi_m \leq 0.2\pi$. As an example, a short roll or yaw period of two seconds gives $f_m = 0.5$ Hz, so that $v \leq 0.3$ Hz. For the same short roll or yaw periods, the maximum roll or yaw Dopplers will typically be much larger. However, the important point is that the measurement accuracy must be checked in this manner, before accepting the roll and yaw Doppler measurements.

The actual tracking of scatterers can be a difficult problem under certain adverse circumstances. First, in case of a ship, the yaw and roll motions may be so large that the cross sections of some major scatterers might change over the motion cycle. One may observe a fluctuating scatterer whose cross section becomes too low for reliable tracking over extended periods. Then one must extrapolate through the intervals of nondetectability, which can be done with a polynomial or spline fit if the intervals are not too long. Second, scatterers that are close in range may interfere with each other, with their phase difference varying throughout the motion cycles. Thus we may observe two separated peaks in the intensity range profile over some time intervals, and a single peak over other time intervals; or if the scatterers have comparable cross sections the combined response may not be visible in the background for extended periods. In the case of two interfering scatterers, one may have to utilize the complex range profile rather than the intensity range profile, so that range resolution is not degraded. As demonstrated in Section 2.3.2, the processor should track actual scatterer positions rather than the peaks in an intensity range profile, and in the most difficult situation the tracking should be combined in range and Doppler.

With increasing aspect angle, range resolution becomes less effective, with the consequence that a single range cell will more frequently contain two or more scatterers of comparable strengths. As the aspect angle approaches broadside, tracking individual scatterers through the sequence of range profiles becomes more difficult, and sometimes impossible. This will be so even when the complex range profiles are utilized. Then it becomes necessary to use the most capable form of scatterer tracking, where one replaces tracking in range followed by tracking in Doppler by combined range/Doppler tracking, using sufficient Doppler resolution to resolve major scatterers in the same range cell. The significant improvement in tracking performance from using the complex range profiles and, even more potently, combined tracking in range and Doppler, was demonstrated in Section 2.3.2.

The purpose of measuring the motion behavior of a ship is to find times when usable images can be formed. Before generating an image at some selected time, the question is whether one must compensate the data other than taking out the range drift via centroid compensation. If the answer is yes, to what degree can this be done? In Chapter 4, we showed that under many conditions the erratic motion of ground vehicles precludes a motion compensation beyond simple range centroid and Doppler centroid tracking of the entire vehicle (or an equivalent procedure). However, such a simple motion compensation will not remove pseudoperiodic yaw and roll motions. We first consider the possibility of a more sophisticated motion compensation for ships. Of course, a sophisticated motion compensation may always be possible under some benign conditions, but in practice we are interested in what can be done under all but the most extreme circumstances.

One potential problem with complicated motions is that a particular scatterer might change its range rate (which governs motion compensation) rather abruptly. Although a motion compensation could, in principle, follow such a change, in practice the measurement of the motion will typically be degraded by interference because there often are other significant scatterers in the same range cell. An even worse problem is that the scatterers occupying the same range cell may have rather different positions in crossrange and height, and thus may be moving differently. In practice, it is generally impossible to unscramble the various returns (which are smeared before a good motion compensation is applied), and then to perform individually different motion compensations on the scatterers. Thus we need not search for a perfect motion compensation; it generally does not exist for the three-dimensional structure and changing rotation axis of a ship.

The more important question is whether we even need a motion compensation beyond the simple removal of the range drift. Ships are large objects, and even modest yaw, pitch, and roll motions lead to range changes that are large in terms of the wavelength of microwave radar. These motions are easily so large that when an image is formed over only a small and smooth part of a motion cycle, so that the distortions caused by the changes of the Doppler are insignificant, we obtain sufficient crossrange resolution. In this case no further motion compensation is required, other than perhaps polar reformatting in order to keep the scatterers at the extremes of the ship within their range gates.

Suppose that we face a situation in which the Doppler is changing rather rapidly, and good crossrange resolution would require imaging over such a large part of the motion cycle that the variation in Doppler is unacceptable. Imaging without further motion compensation will lead to a

smearing of the responses. However, the motions of the scatterers will vary with their deck positions, and also with their height above the deck. In fact, scatterers in the same range gate may have substantially different heights. This often makes the motion compensation too complicated to be practical for real data. For this case, as a practical solution it appears best to forgo any attempt at a sophisticated motion compensation, and for identification to rely on Doppler-resolved range profiles in the same manner as discussed for ground vehicles going over poor roads or terrain. However, with the large and relatively smooth motions of a ship, this is a situation far less likely to occur than with ground vehicles.

While we may not need a motion compensation beyond simple removal of the range drift, the facts remain that we have range tracked, Doppler tracked, and phase tracked four scatterers and have verified the qualities of these tracks. If a motion compensation is indeed needed, we are better served by using the bow or stern track to compensate a yaw image, and using a superstructure track to compensate a roll image, than relying on a centroid track. The improvement over the mere removal of the range drift may be small, but it costs nothing.

As a last point concerning a possible motion compensation, we want to reiterate that when the motion happens to be complicated, we do not even want to attempt a motion compensation because of the problem of spurious responses. In order to avoid excessive spurious responses, we want to image when the motion is about a fixed axis and smooth, without any disturbances. Choosing imaging times when the motion is about a fixed axis is important even if the imaging interval is so short that variations of the Doppler over the imaging interval are smaller than the reciprocal imaging interval, in which case these variations do not matter. We have made this point repeatedly, so that an illustration is in order (for which real data are needed). Figure 5.30 shows a true short-term roll image of a ship, with an imaging duration of 0.2 seconds. The fact that the deck line of the ship is oriented nearly along the range axis implies that neither yaw nor pitch is effective at the imaging time. This is also indicated by the narrowness of the deck image outside the superstructure. Figure 5.31 shows an image of the same ship, with the same short imaging duration, but at a time when all motion components are effective. Note that the deck, where recognizable, has a crossrange spread. The clip level for this image was chosen by attempting to show the far side of the deck without allowing too many spurious responses in the image. An examination of the image shows that the spurious responses have a strength comparable to that of the far deck responses. Clearly, a length measurement, for example, would have questionable accuracy.




Figure 5.30 Roll image of a ship.




Figure 5.31 Hybrid image of the ship.

5.1.6 Section Summary

The yaw, roll, and pitch motions of ships are typically so large that, at such short wavelengths as X-band, the aspect angle changes due to these motions are far larger than needed for the desired crossrange resolution.

Pure yaw motion generates topview images. Pure roll or pitch motion generates sea-level-view images. When motion includes more than one component, pitch behaves partly as yaw and partly as roll. For simplicity, we refer to the combined effects as “yaw” or “roll.”

A high-quality topview image is so useful that it should be the first objective of imaging. The measurement of the height profile of a ship also appears indispensable. A height profile with unknown scale factor is probably all that can be measured in practice, and it should suffice. It can be derived from a sea-level-view image or from the distortions of a topview image as the imaging time is shifted slightly.

Besides the type of image, a second requirement is that an image not have a serious problem with spurious responses. This means that the motion of the ship must be about a fixed axis at the imaging time. In order to obtain an undistorted image that can be readily interpreted, the axis must be vertical (a topview) or horizontal (a sea-level view).

The relative Doppler of bow and stern scatterers depends only very weakly on roll motion. When this Doppler is zero, only yaw is effective and one can form a topview image. The relative Doppler of two superstructure scatterers with the same deck position but different heights depends only very weakly on yaw motion. When this Doppler is zero, only roll is effective and one can form a sea-level-view image. The weak third motion components may cause small distortions of the images, but the small refinements in imaging time needed to generate undistorted images can be estimated from the distortions themselves.

We can use the tracks of the bow and stern scatterers to derive the track for any point on the line (in three-dimensional space) through the two scatterers, and thereby replace one of the superstructure tracks with the derived track of a virtual scatterer whose location differs from that of the tracked superstructure scatterer primarily in height, with small differences in the other dimensions.

Before continuing with the analysis of the motion behavior of the ship, we must check that every scatterer that has been tracked has remained in its range cell. The amplitude of the transform of a fixed-range image cut through the compensated response must not become weak for any extended interval.

5.2 Measurement of Yaw and Roll Motions

5.2.1 Measurement Principles for Real Data

Section 5.1 treats the principles of ship imaging and the selection of imaging times when a usable image can be formed. This should suffice for the reader who wants to obtain a general understanding of the technology. For those who are interested in the actual processing procedures, Section 5.2.1 treats a specific case in considerable detail. The remaining illustrations, considered in less detail, concern special but nevertheless practically important situations, of interest to those who want to work on ship identification. Section 5.2.2 treats the case when a ship yaws without much roll, Section 5.2.3 a turn maneuver, Section 5.2.4 a very small maneuvering ship, and Section 5.2.5 the difficult case of a small ship in rough seas. A summary is given in Section 5.2.6. We reiterate here that “yaw” and “roll” motions both may contain pitch contributions.

In order to measure yaw motion, we want to select scatterers as far apart along the length of the ship as possible. Scatterer positions along the length of the ship can be taken from an image regardless of whether there is any yaw or roll. In the absence of any motion, the positions can still be taken from the range profile. Also, since the length of a ship is much larger than its width, because the superstructure tends to be in the center of the ship and because a majority of ships have bow and stern at approximately the same heights, errors from height differences and offsets from the centerline will typically be insignificant.

For the measurement of roll motion, we need at least one scatterer on the superstructure. However, we need a roll image in order to select scatterers at different heights. In those frequent cases where the superstructure is rather concentrated, only a limited selection of scatterers will be available for the roll measurement, because the scatterers may be so poorly resolved that tracking is problematic. A slight roll translates the responses of the superstructure from their deck positions along the width of the ship, yet this cannot be recognized in the image because we do not know where the responses should be. For this reason, we must consider the measurement of the yaw and roll Dopplers in two cases, with or without significant roll. “Significant” here means that the crossrange spread of the superstructure is sufficient for one to recognize the height relations between scatterers in an image.

We want to make an important point concerning real data. The purpose of the motion measurement is to find suitable imaging times, but with real data there will be unavoidable measurement errors. These errors may be small, yet the characteristics of an image can change significantly with only

small changes in the imaging time. For example, a shift of the imaging time by only 5% of the motion cycle can have significant consequences. This does not constitute a problem in practice. We must remember *that the purpose of the motion measurement is to find imaging times at which the motion is about a fixed axis and Doppler is easily physically interpretable*. This means imaging times at which either the position of the deck or of the superstructure (relative to the deck) is fixed. In the former case the deck appears as a line oriented along the range axis in the image. In the latter case, the superstructure responses have the correct positions in the image of the deck. These two facts allow us to refine the imaging time if it was measured with an error of a small fraction of the motion cycle. Such accuracy is generally achievable, as we show in the remainder of this section and in Appendix D.

In order to refine the timing of a sea-level-view image, we first check whether the deck appears as a line or has a Doppler spread. If the latter, we form images slightly before and slightly after the measured time, using time increments that cause significant changes in the crossrange width of the deck. We measure the deck widths in the two images and determine the correct imaging time by interpolation. If the deck appears as a line that is tilted off the range axis, we form images slightly before and slightly after the measured time, using time increments that cause significant shearing of the deck line in crossrange. We measure the degree of shearing in the images, and determine the correct imaging time by interpolation. The refinement can be iterated until a deck line oriented along the range axis is obtained.

In the case of the topview image, the check of the accuracy of the measured imaging time is more difficult, in particular in an automated system, because it requires determining the position of the superstructure within the outline of the deck, which is not trivial. It is simpler to form two images as with the sea-level-view image, one before and the other after the measured imaging time. Measuring the maximal crossrange separation of the superstructure from the centerline of the deck in each image and interpolating gives the time at which the superstructure is centered along the width of the ship. This centered position of the superstructure may not be exactly the true position, but it appears satisfactory. If the depression angle is nonzero, we must shift the superstructure in range before centering it. The shift can be based on typical superstructure heights for a ship with the measured length and superstructure position.

In Figure 5.32, we present a flowchart for the analysis of a ship's motion. This chart is intended to supply an overview of the analysis and to show how the examples in the following subsections fit together. The rationale for the various branches in the chart is given in the examples contained in

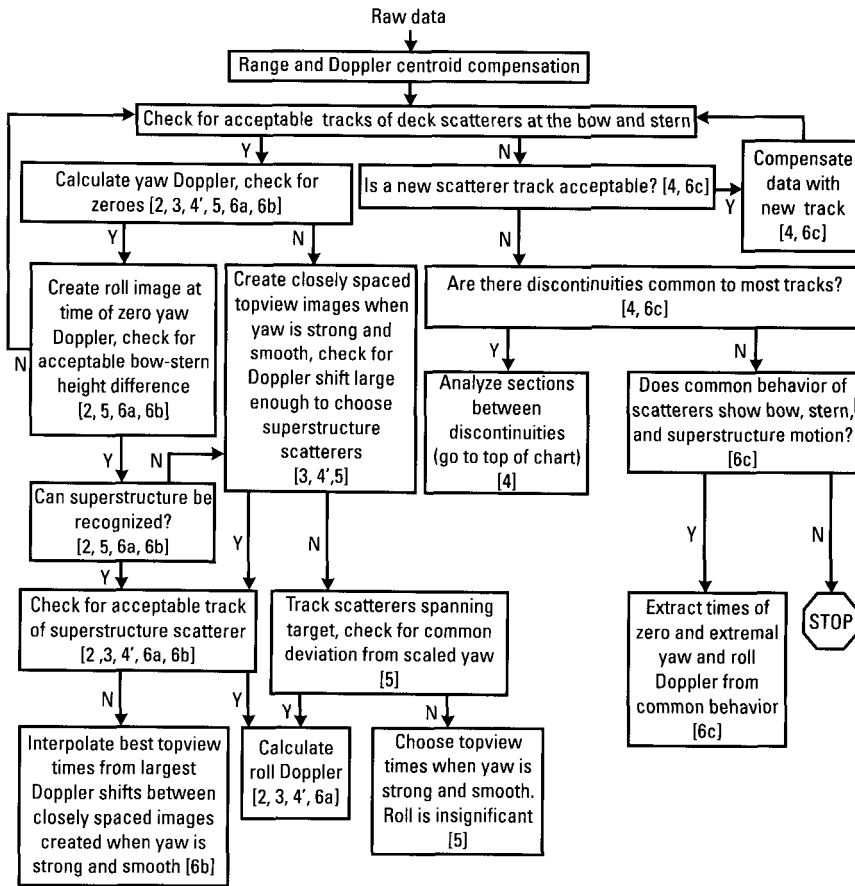


Figure 5.32 Flowchart for analysis of ship motion.

the subsections. Most boxes in the flowchart contain numerals in brackets. This indicates that the operation described in the box is employed in an example given in the corresponding subsection (e.g., the numeral 2 corresponds to Subsection 5.2.2). The data used in Subsection 5.2.4 is divided into pieces. The analysis of the entire data is denoted by 4, and that of a piece by 4'. The three examples in Subsection 5.2.6 are denoted by 6a, 6b, and 6c. The details of the flowchart will be much better appreciated after completing the subsections.

The path one follows through the chart depends on the smoothness of the ship's motion and the relative strength of yaw and roll motions. The strategy underlying the sequence of operations is first to measure the yaw

motion by tracking bow and stern scatterers, next to measure the roll motion by selecting and tracking a superstructure scatterer, and then to use the motions to select imaging intervals. Ideally, the motion measurement is done by recognizing and tracking bow, stern, and superstructure scatterers. If tracking a bow or stern scatterer is not possible, we attempt to recompensate the data, to reduce the time interval to be analyzed, or to estimate the yaw motion from common behavior of several scatterers. If tracking or recognizing a superstructure scatterer is not possible, we attempt to estimate the roll motion from common behavior of several scatterers or from the crossrange spread of images created when the yaw motion is strong and smooth.

Several steps in the flowchart mention *acceptable scatterer tracks*. As we explained above and illustrate below, a scatterer track is acceptable for the motion determination if, after compensation with the track, the amplitude of the Fourier transform of a fixed-range image cut through the scatterer response does not stay at the noise/clutter background level of the image for an extended time. When this is the case, the fluctuations of the transform phase (or equivalently, the crossrange spread of the image response) give the maximum residual scatterer Doppler. In the motion determination, we calculate differential Dopplers between pairs of scatterers. For a pair of scatterers to be acceptable for the motion measurement, their differential Doppler must be significantly larger than both residual Dopplers. Finally, the scatterers must be located in the appropriate positions on the ship. This is discussed in detail below.

We must choose imaging times when one of the two motions passes through zero Doppler and the other is strong and smooth. The Doppler rates at each such time, scaled to the full target size, allow us to determine the longest *acceptable imaging* interval duration. We must then reexamine the fixed-range image cuts used to judge scatterer acceptability, and exclude any part of the interval that includes an abrupt change in the transform phase. Because the scatterers have nonzero residual Dopplers, and because they are not ideally located on the ship, we must refine the times of the imaging intervals. These steps are discussed further and are illustrated below.

5.2.2 Significant Roll

5.2.2.1 Measurement of Yaw and Roll Dopplers

Since the measurement of the ship's motion is very important, we describe it in considerable detail for this first example of a real ship. Where appropriate, we also state how the measurement can be improved with procedures that, at the time of this writing, had not yet been integrated into our interactive

software. The improvement is meant in the sense that the measurement of the ship's motion will work for worse cases of motion (larger and more erratic movements), and at aspect angles closer to broadside.

We select a rather long data segment of 20 seconds duration, without regard to whether or not the motion conditions change. It is an arbitrary selection of a specific observation interval for an arbitrarily chosen small ship (a coast guard cutter with a length of 110 ft), shown in Figure 5.33. The sequence of range profiles, after taking out the drift due to the translational motion of the ship, is shown in Figure 5.34. Note that some scatterers at the larger ranges are seen only intermittently in the background. As the first processing step, we generate the simple peaks tracks, based on the intensity range profiles. In this example, the processor measures the range positions of the 20 strongest peaks of each intensity range profile, as depicted in Figure 5.35. In time intervals within which a response peak is reasonably strong with little interference, the peaks tracks are obvious in the figure. However, for the most part, the tracks are very poor. The situation can be somewhat improved by sliding-window Doppler processing, where a sequence of range profiles is coherently integrated to produce responses, and the integration window is shifted in steps over the entire observation time. This corresponds to Doppler processing with too low a Doppler resolution to resolve scatterers in the same range gate. When the peaks are coherently processed over 32 consecutive range profiles, we obtain the smoothed tracks of Figure 5.36. This figure has a better appearance to the eye than Figure 5.35, but in reality the situation is not much improved.

Compare the peaks tracks of Figure 5.35 for a real ship with those of Figure 5.20 for the simulated ship. One may wonder how one could possibly track scatterers in Figure 5.35 with the accuracy needed to derive meaningful yaw and roll Dopplers. Actually, this quality of peaks tracks does not yet pose



Figure 5.33 Coast guard cutter.




Figure 5.34 Sequence of range profiles.




Figure 5.35 Peaks tracks over the 20 seconds.




Figure 5.36 Smoothed peaks tracks.

a serious problem even for the simple procedure of tracking only these intensity peaks, but the performance checks we discussed in Section 5.1.5 clearly are important. The peaks tracks must be of much poorer quality before it is necessary to utilize the complex range profiles for tracking actual scatterer positions rather than peak positions of the intensity profile, and even worse before it is necessary to resort to combined range/Doppler tracking.

Even though we have demonstrated for simulated data that one of the four scatterers needed for the yaw and roll measurement can be replaced by a virtual scatterer, and we will demonstrate this for real data as well, in most examples we will track four scatterers. Our intent is to treat the principles, without suggesting which specific sequence of processing steps should be implemented in an operational system. This can be done only after working with a large variety of data from all types of ship and motion condition.

In accordance with the earlier discussion on the motion measurement, we want to establish good tracks on two pairs of the range profile peaks. Two tracks should be as far apart in range as possible, for the best measurement of the yaw motion. The two other tracks should be as close as possible in range, but the height difference for the two scatterers should be as large as possible. Following four of the peaks tracks in Figure 5.35 appears to be difficult

or impossible, but we will demonstrate that it can be done. As a firm rule, we must always check the quality of any track. The check will determine whether we have satisfactorily tracked a scatterer. Thus we will know whether or not accurate yaw and roll measurements may be performed.

In view of the poor quality of the peaks tracks of Figure 5.35, we point out that we will use the crudest form of tracking; that is, tracking of the peaks of the intensity range profile, but with manual processing. In an operational system the processor should use the improved tracking procedure discussed above. An adequate tracking performance should be achievable under much worse conditions than represented by the peaks tracks of Figure 5.35. However, since we have not integrated the improved tracking procedures into our interactive software, we cannot demonstrate under what worst-case conditions the measurement of the motion of a ship is still possible.

We first try tracking the “bow” scatterer near Range Gate 54 of Figure 5.35 (indicated by the top arrow at the left of the plot), which is seen only intermittently. The track of this peak and the associated spline fit (with one break point, indicated by the vertical line) are shown in Figure 5.37. Next, the data are compensated with the spline of Figure 5.37, which process shifts the scatterer into Range Gate 0. Figure 5.38 shows the Doppler




Figure 5.37 Range track of the scatterer near Range Gate 54.




Figure 5.38 Doppler track of the scatterer in Range Gate 0.

variations of the scatterer in Range Gate 0, together with the spline fit. Again, the spline is used to compensate the data. The selected scatterer now has been compensated in range and Doppler, and the combination of the splines in Figures 5.37 and 5.38 represents the motion of the scatterer. We must verify that the compensation is good enough for the motion measurement to be accurate. This is done by forming an image over the entire 20 seconds and analyzing the image cut in Range Gate 0. The applicable part of the image is shown in Figure 5.39. As is the case with all responses of the image, the compensated response in Range Gate 0 is evidently highly smeared in crossrange, which implies an imperfect motion compensation. However, at this point all that is needed is a *tracking accuracy good enough for the motion measurement*, not for generating properly compressed image responses. Crossrange smearing of even the compensated response is to be expected when a compensation is performed over a 20-second interval, which contains many motion cycles of this ship. The image cut in Range Gate 0 is shown in Figure 5.40, the amplitude and phase functions of the image cut itself in the left half and amplitude and phase functions of the transform in the right half.

For an accurate motion measurement, one would like to have a nearly constant transform amplitude function in order to make the phase function




Figure 5.39 Part of the 20-second image.

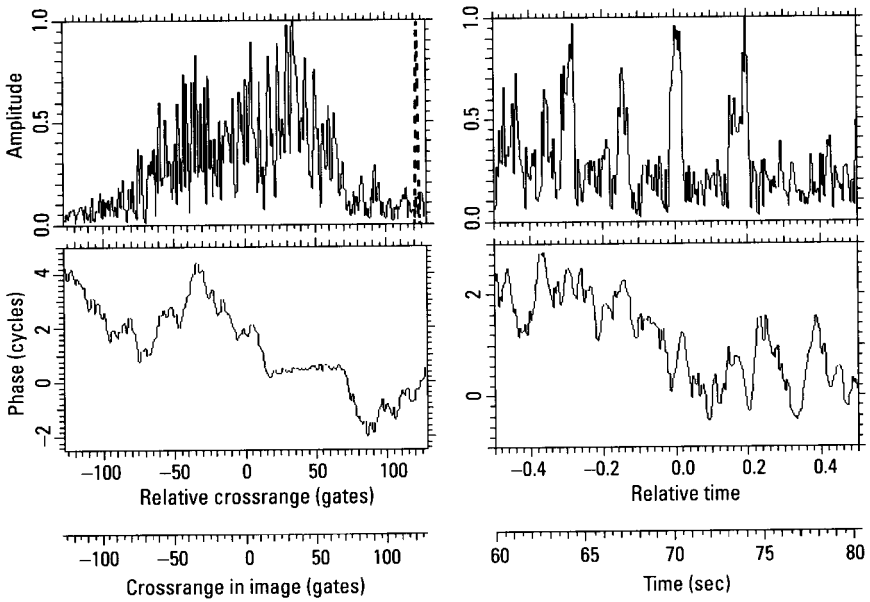


Figure 5.40 Image cut in Range Gate 0.

meaningful. By measuring the phase variations, in accordance with (5.3), we can estimate the error in the measured yaw or roll Doppler. The amplitude function at the right top in Figure 5.40 is strongly fluctuating rather than being reasonably constant. However, the amplitude is not near zero for extended periods. Moreover, if we ignore the phase jumps at the times of the amplitude minima, the variation of the phase function is quite small. Despite the poor quality of the track, it does not appear necessary to improve the tracking by using the phase-slope version, or even choosing a different scatterer.

In a situation where the amplitude/phase pattern is difficult to use for estimating the quality of the scatterer track, we can estimate it from the responses after compensation based on the measured phase function. Residual phase variations grossly smear an image even when the phase varies by no more than about one cycle over the total time interval, and the consequences of such a residual phase modulation far outweigh the widening of the responses due to an amplitude modulation. The smearing of the responses thus allows a direct estimate of the residual phase variation when the smearing is strong (in which case the amplitude function does not matter). Figure 5.40 (left top) shows that most of the smeared response energy is contained within about ± 50 crossrange gates of the response center. With an imaging interval of 20 seconds, one crossrange gate corresponds to 0.05 Hz, so that a spread over ± 50 crossrange gates corresponds to ± 2.5 Hz. From this crude estimate we conclude that, with a pseudoperiodic motion of the ship, the motion compensation of the intermittently visible scatterer near Range Gate 54 contributes an uncertainty of about 2.5 Hz to a relative bow-stern yaw Doppler. This is acceptable as long as the maximum relative Doppler is much larger than 2.5 Hz. If the results show that this condition is not met, a more accurate estimate via phase-slope tracking is needed, or one might choose a better scatterer for the yaw measurement.

This procedure estimates the residual pseudoperiodic Doppler variation of the scatterer, but neglects a potential Doppler bias, corresponding to a linear phase slope in the signal. The signal phase decreases roughly linearly by about two cycles during the 20 seconds displayed in Figure 5.40. This gives a range rate of one wavelength per 20 seconds, or a Doppler bias of 0.1 Hz. In this instance, the bias is totally negligible.

Although it is easy to estimate the Doppler width of the image response of Figure 5.40 by eye, an automatic estimate of the width of a highly smeared response is not as straightforward. Fortunately, high accuracy is unnecessary. We can automatically *estimate the response smearing, and its implied Doppler uncertainty*, as follows (perhaps with some adjustment of the specified parameter values):

1. Estimate the center of the response as the first moment of those image cut samples with amplitudes at least three times the background level of the image.
2. Moving left and right from the center of the response, note the first occurrence of three consecutive peaks with amplitudes less than three times the background level. Limit further consideration to the crossrange interval from the center to, but not containing, these peaks.
3. Fit a Gaussian function to the amplitude peaks in the interval. Take the half-amplitude width of the Gaussian as an initial estimate of the uncertainty.
4. Estimate the strengths and separations of the scatterers from an FFT of the amplitude of the transform, with the phase set to zero (see Section 4.4.9 and Appendices A and C). If any secondary scatterer has an amplitude at least half that of the primary scatterer and three times the image background, and a separation from the primary of less than the half-amplitude width of the Gaussian, our estimate of the response width should account for the secondary scatterer. Thus, our uncertainty estimate becomes the half-amplitude width of the Gaussian minus half the largest separation for a secondary response of concern.

This procedure will overestimate the tracking uncertainty if multiple scatterers in the same range gate have very different Dopplers that are rapidly changing. In that case, even if one scatterer were perfectly focused (which is unlikely), the others would create a widely smeared response overlapping that of the focused scatterer. One can recognize that overestimation has occurred by comparing the Doppler spread of the image response being used for the uncertainty estimate to the Doppler spread of the response calculated in Step 4 above. If the Doppler spreads are equal (although usually much weaker for the latter response), the widths represent the differing scatterer motions, rather than tracking errors, and the tracking uncertainty will have been overestimated. In such a situation, the best course of action is to attempt to track another scatterer.

In order to calculate the yaw motion, we need to recognize and track a stern scatterer at about the same height above the deck as the already tracked bow scatterer. We can recognize a scatterer's height by forming an image at any time when roll motion is not negligible. This can be a short-term image, so that no motion compensation beyond removal of range drift is required

on the original data. For imaging, we select a time at which the scatterer tracked above (Range Gate 54) is well visible. For example, we can use part of the 0.8-second interval starting at 70 seconds in Figure 5.35, indicated by the double-headed arrow at the top of the plot. We choose the imaging interval so that the well-trackable scatterers at lower ranges also are visible. Lastly, in the roll image we check which of these scatterers have the desired small height difference.

As an illustration, in Figure 5.41 we show the image from 70.2 to 70.4 seconds. The horizontal spread of the image is only two crossrange gates, so that the roll Doppler is low at the chosen instant; yet it is sufficient for our purpose. Since the first scatterer was chosen near Range Gate 54, its response is the top response in Figure 5.41 (indicated by the top right-pointing arrow). We utilize this scatterer as a reference for the yaw measurement. This means that we compensate the data with the tracked motion, select a new scatterer, and measure the motion of this scatterer. This scatterer should be sufficiently well trackable, and it should be at a height as little different from that of the reference scatterer as possible. The new scatterer can be selected from the range tracks of Figure 5.35 and the image of Figure 5.41. We select the scatterer in Range Gate -25 (indicated in both figures by the bottom right-pointing arrow), which is at about the same height as the reference scatterer and has a relatively good range track. The tracks of this scatterer, performed as illustrated in Figures 5.37 and 5.38, describe the yaw motion of the ship.

For the roll measurement, we must compensate a new reference scatterer, and then track a second, chosen as close in range and as far apart in height from the reference scatterer as possible. We again examine the range tracks of Figure 5.35 and the image of Figure 5.41 to make the choice. We select the strong scatterer near Range Gate -5 as a reference, and the scatterer near Range Gate 3 for the roll measurement (indicated in both figures by the left-pointing arrows). The same type of compensations and measurements as illustrated for the yaw measurement now are performed on these two scatterers, resulting in the roll Doppler.

The yaw and roll Dopplers measured in this manner are plotted in Figure 5.42, with the dashed curve giving the yaw Doppler. The estimated uncertainties for the two curves are nearly equal, although this need not be the case in general, and are indicated by the dotted horizontal lines. Note that the Doppler maxima are several times larger than the uncertainties. Accordingly, the zero crossings of the Doppler curves are sharp enough so that the Doppler uncertainties cause only small uncertainties in the times of the zero crossings, on the order of one-tenth the motion period. Thus,




Figure 5.41 Image from 70.2 to 70.4 seconds.




Figure 5.42 Yaw and roll Dopplers of the ship.

our *tracking accuracies appear acceptable*. Figure 5.42 allows us to select the imaging time so as to achieve a particular type of image, within the limitations set by the motion of the ship.

Before proceeding to examples, we also show the results of deriving the roll Doppler with the alternative procedure, using one virtual and only three real scatterers. Figure 5.43 shows the roll Dopplers derived with both procedures, with the solid curve taken from Figure 5.42 for comparison. The roll Doppler from the alternative procedure was derived using the bow and stern scatterer tracks and the track of the scatterer in Range Gate 3. The times of zero roll Doppler agree well between the two curves, with typical differences of about 0.1 seconds. The shapes of the curves are similar as well. Because the second superstructure scatterer is about as well trackable as the bow and stern scatterers, the uncertainties are also similar for the two curves.

We chose an arbitrary imaging time for the survey image of Figure 5.41, which was used for the selection of the scatterers for the motion measurement. In accordance with the Doppler curves of Figure 5.42, this happens to be a time at which the absolute roll Doppler is only a small fraction of its peak value. Nevertheless, the image still allows a good estimate of the height differences of the scatterers. Having derived the yaw and roll Doppler curves, we can check whether we can improve our choice of scatterers by choosing




Figure 5.43 Roll Dopplers from both procedures.

an imaging time at which the roll Doppler reaches a maximum, so that the height differences of the scatterers can be more accurately measured, and by selecting a time for generating a topview image, so that the scatterer positions relative to the centerline of the ship can be more accurately measured. If we observe a large height difference between bow and stern scatterers or a large difference in deck position between superstructure scatterers, we can attempt to track other scatterers that have smaller differences.

A rule of thumb, based on (D.10) and (D.14), is that scatterer positions used to measure a motion component give acceptable timing errors if the maximum value of the desired component of their relative Doppler is only a few times larger than that of the corrupting components. For motions whose components have comparable periods and maximum angular excursions, this means that the separation along the desired direction need be only a few times larger than that in other directions (see (D.11) and (D.15)), which is usually easily satisfied. Regardless of comparability of the components, the *acceptability of scatterer positions* can be determined by examining the Doppler separations of scatterers in the topview and sea-level-view images, as follows:

1. Form a topview image at a time of zero roll Doppler. Measure the Doppler separation of the stern and bow scatterers, v_{sb}^{top} , and the superstructure and nearby deck scatterer, v_{sd}^{top} .
2. Scale these by the ratio of the maximum yaw Doppler to the yaw Doppler at the time of the topview image, $v_{yaw}^{max} / v_{yaw}^{top}$, where the maximum is taken as the mean of the absolute yaw Doppler extrema bracketing the time of the topview image.
3. Form a sea-level-view image at a time of zero yaw Doppler. Measure the Doppler separation of the stern and bow scatterers, v_{sb}^{sea} , and the superstructure and nearby deck scatterer, v_{sd}^{sea} .
4. Scale these by the ratio of the maximum roll Doppler to the roll Doppler at the time of the sea-level-view image, $v_{roll}^{max} / v_{roll}^{sea}$, where the maximum is taken as the mean of the absolute roll Doppler extrema bracketing the time of the sea-level-view image.
5. Require the maximum value of the desired component to be at least a few times the maximum value of the corrupting component:

$$\frac{v_{sb}^{top} (v_{yaw}^{max} / v_{yaw}^{top})}{v_{sb}^{sea} (v_{roll}^{max} / v_{roll}^{sea})} \quad \text{and} \quad \frac{v_{sd}^{sea} (v_{roll}^{max} / v_{roll}^{sea})}{v_{sd}^{top} (v_{yaw}^{max} / v_{yaw}^{top})}$$

must both be at least about three.

Returning to Figure 5.42, the Doppler functions show that, at least within the depicted time interval, the timing of yaw and roll happens to be quite well synchronized. Both Dopplers tend to peak and go through zero nearly at the same times. This indicates that the ship's motion is dominated by pitch, which contributes to both the "yaw" and "roll" motions with the same timing, so long as the pitch angle does not cross zero. If it does cross zero, the "yaw" motion will show twice as many cycles as the "roll," with half the "yaw" peaks and zeroes synchronized with those of the "roll" motion. The synchrony resulting from the dominant pitch motion makes it a little more difficult to find imaging times where one type of Doppler is zero and yet the other type is substantial, as required for undistorted topviews and sea-level-views.

There is also the question of which of the two types of Doppler is dominant, in the sense of more easily allowing the generation of a good image of its corresponding type. This question cannot be answered by examining Figure 5.42, because the Doppler curves were not generated by measurements on the extremal scatterers of the target; they must be scaled accordingly. This can be readily done by forming an image at a time when both Dopplers in Figure 5.42 have high values. For example, Figure 5.44 shows a 0.2-second image at a time of 69.6 seconds, where in accordance with Figure 5.42 the yaw Doppler curve is about 50% higher than the roll Doppler curve. However, the crossrange spread of the image of Figure 5.44 due to the roll Doppler is 15 gates, whereas the yaw Doppler causes a crossrange spread of the base of only about 4 gates. This is so despite the fact that the yaw Doppler of Figure 5.42 is 50% higher than the roll Doppler, because the separation of the scatterers used to measure the yaw motion (indicated by right-pointing arrows) is a much larger fraction of the ship's length than the separation of the scatterers used to measure roll motion (indicated by left-pointing arrows) is of the ship's width. The much greater crossrange spread due to roll means that under the existing conditions it must necessarily be far easier to generate a good sea-level-view image than a topview image. (Note the limited utility of a sea-level-view image, in particular at larger aspect angles.)

This conclusion is easily demonstrated, and at the same time we can show that the roll and yaw Dopplers were quite accurately measured. If we want to generate a pure roll image (not sheared in crossrange), we must select an imaging time at which the yaw Doppler is zero but the roll Doppler is as high as possible. An examination of Figure 5.42 shows that such a combination of the two Dopplers occurs at a time of 64.7 seconds. The image from 64.6 to 64.8 seconds is shown in Figure 5.45. It is indeed a nearly pure roll




Figure 5.44 Image from 69.5 to 69.7 seconds.




Figure 5.45 Image from 64.6 to 64.8 seconds.

image without yaw effects, as seen from the fact that the baseline is oriented nearly parallel to the range axis. Although it is not exactly parallel, the error is too small to require correction by finding the exact imaging time via interpolation, as discussed earlier.

In order to generate a topview image, we must image at a time when the roll Doppler goes through zero yet the yaw Doppler is substantial. From Figure 5.42, the best time is at 64.0 seconds. When we form an image at this time, because of the strongly dominant roll Doppler and the unavoidable error in the measurement of yaw and roll Doppler, the responses from the superstructure are noticeably shifted by the roll. The yaw Doppler is so small that a roll Doppler residual can readily translate the responses in crossrange by amounts comparable to the crossrange spread of the ship in a given range cell. In order to obtain a pure topview, we must use the interpolation procedure discussed above. This results in a shift of the imaging time by 0.2 seconds to 64.2 seconds. The corresponding image is shown in Figure 5.46. It is a topview image, but (as expected) the definition is poor because of the low yaw Doppler of the ship. Because the ship happens to roll much more strongly than it yaws, it is unlikely that the quality of the topview image can be improved by going to longer imaging times with a motion compensation




Figure 5.46 Image from 64.1 to 64.3 seconds.

(only the centroid compensation was used for the image of Figure 5.46). A test shows that this is indeed the situation.

In the above example of selecting the imaging time from the measured yaw and roll Dopplers, we had to shift the imaging time by 0.2 seconds from the value measured with the four scatterers in order to obtain an undistorted topview image (again, the amount of shift can be found from two images and interpolation, as explained earlier). The zero crossing of the yaw Doppler thus was measured with an error of 0.2 seconds. Such errors must be expected, in particular with the low-quality peaks tracks of this example. However, we demonstrated for the simulated ship that the three-scatterer method yields more accurate results because the virtual scatterer can be moved closer to the deck position of the superstructure scatterer. From the dashed curve of Figure 5.43, which was derived with the three-scatterer method, we obtain the time of zero crossing at 64.2 seconds, which agrees with the imaging time used for Figure 5.46.

An important point regarding the selection of the imaging time concerns image quality. By this we do not mean absence of “distortions” introduced by roll or yaw, or a hybrid image, but the quality of the individual responses and the number and strengths of spurious responses. Are the genuine image responses properly compressed, with low sidelobes? When the ship goes through some erratic motion, it probably is not possible to generate a high-quality image in this respect, even with a sophisticated motion compensation. The only remedy then is to generate an image with low crossrange resolution. With slowly moving ships, where the observation time can be long enough to permit a choice in the imaging time, one of the criteria for selecting the imaging time must be the obtainable image quality. We will show that this is important and that the yaw and roll Doppler curves contain the information for the appropriate choice.

If the motion of the ship is smooth, which is a requirement for high image quality, the tracks of the scatterers will also be smooth, and so will be the spline fits and the yaw and roll Doppler curves. Hence, examining Figure 5.42, we would expect the image quality to be higher when the extrema of the curves are smooth. For example, we should obtain a roll image of higher quality around the time of 68 seconds (labeled 5.47) than around 75 seconds (labeled 5.50). We now demonstrate that *the smoothness of the yaw and roll Doppler curves is correlated with image quality*.

In Figure 5.47 we show an image centered at a time of 68.24 seconds, with an imaging interval of 0.68 seconds. The peculiar imaging interval was chosen for two reasons. First, we want a larger interval than the 0.2 seconds used earlier, because imaging problems increase with higher crossrange




Figure 5.47 Image from 67.90 to 68.58 seconds.

resolution; that is, with the length of the imaging interval. Second, for a fair comparison with the second image considered below, we want about the same degree of crossrange resolution for both images. The imaging intervals thus were scaled in accordance with the roll Dopplers at the imaging times. We now take an image cut along the baseline, even though this is in the range domain and imaging problems due to the ship's motion would be expected to occur primarily in the crossrange dimension. However, as pointed out repeatedly, range measurement accuracy is affected by crossrange resolution.

The image cut along the straight line in Figure 5.47 is shown in Figure 5.48. The responses of the image cut are well developed in the sense that when we perform TSA measurements on the individual responses, we obtain good patterns from either a single or from two interfering scatterers. The implication is that we can measure actual scatterer positions. In the left half of Figure 5.49 we show the image cut in the range gate of the major response marked by the x in Figure 5.47, Range Gate -1.6 , after we suppressed responses outside the vertical lines. The amplitude and phase functions of the transform are shown in the right half of the figure. The




Figure 5.48 Image cut along the vertical line marked in Figure 5.47.

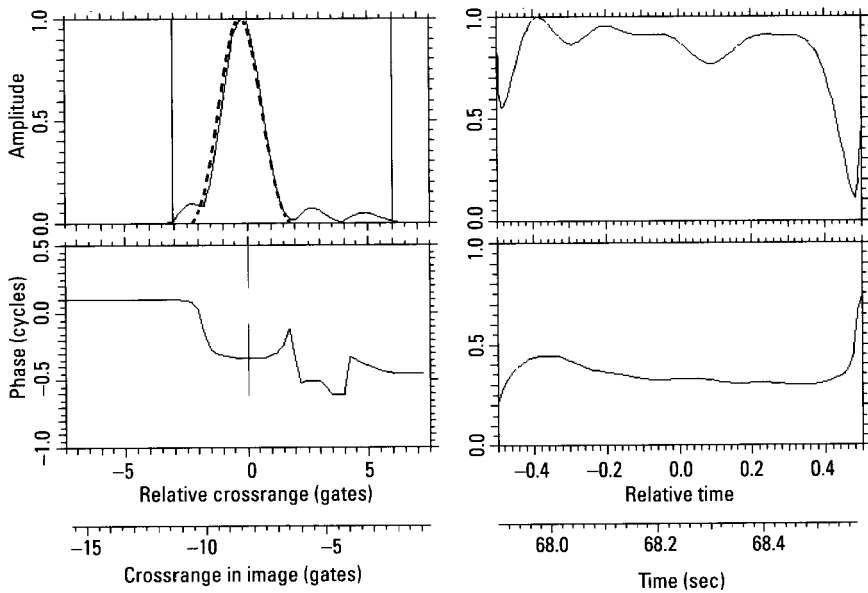


Figure 5.49 Image cut in Range Gate -1.6.

amplitude is reasonably constant, and the phase has a residual curvature of 0.1 to 0.2 cycles. Hence, this response has reasonable properties in the crossrange dimension as well. This turns out to be the case for the other responses as well. In other words, the image is good enough to determine scatterer positions.

In Figure 5.50 we show an image at a central time of 74.7 seconds, with a long imaging interval of 1 second. This is the point at which the roll Doppler curve in Figure 5.42 has a double-hump maximum. The image cut along the baseline is shown in Figure 5.51. Even to the eye, the structure of the responses compares unfavorably with that in Figure 5.48 (the minima are shallow throughout Figure 5.51). When we try to analyze individual (complex) responses, we find rather poor approximations of the amplitude/phase patterns of single and double scatterers. The image cut in Range Gate -1.6 , through the same major response as in Figure 5.49, is given in Figure 5.52. The response now is strong only within somewhat more than the first half of the imaging time, and the curvature of the phase function is much larger than in Figure 5.49. Overall, not only does the new image have poorer quality than the first image, but the quality is insufficient for the clean resolution of scatterers and the measurement of their positions. Hence, the yaw and roll




Figure 5.50 Image from 74.2 to 75.2 seconds.




Figure 5.51 Image cut along the vertical line marked in Figure 5.50.




Figure 5.52 Image cut in Range Gate -1.6.

Doppler curves also allow us to select times at which images of high quality can be generated.

5.2.2.2 Interpretation of Yaw and Roll Measurements

We have shown previously that, on the basis of the yaw and roll Doppler curves, one can predict the type of image obtainable at a given time, as well as the quality of the image with respect to focusing of the responses and the appearance of spurious responses. In this section we show more systematically how to interpret the yaw and roll Doppler curves, using a 20-second segment of data of the same ship at a different time.

In Figure 5.53 we show the peaks tracks over the new 20-second interval. In an automated system one will start with the yaw measurement in order to find good times for a roll image (zero yaw Doppler). With manual processing, it is easier to search for one or the other imaging time with significant roll effects. In our example, we start by generating short-term images (without motion compensation) at times when even the intermittent scatterer in Range Gate 55 is visible, trying different times until an approximate roll image is obtained. Such an image is shown in Figure 5.54, generated at a time of 34.9 seconds with an imaging interval of 0.2 seconds. For the yaw measurement, we need two scatterers along the baseline of the ship, with




Figure 5.53 Peaks tracks over a new time interval.




Figure 5.54 Image from 34.8 to 35 seconds.

usable peaks tracks. Thus we consider the major or isolated dots along the baseline in Figure 5.52, read their ranges, and from Figure 5.51 determine whether the peaks tracks appear usable. (If we make an error in this judgment, we must choose another scatterer after the test for track quality fails.) The selected scatterers are those in Range Gates 53 and -27 at the time of the image of Figure 5.54 (indicated by right-pointing arrows). Since the superstructure of a ship typically does not extend to bow and stern, the choice of the two scatterers for the yaw measurement usually is easier than that of scatterers for the roll measurement.

For the roll measurement, we want to select two scatterers as close in range as possible, but at very different heights. This would suggest choosing the scatterers associated with the strong responses in Range Gates -5 and -4 of Figure 5.54. However, the range separation of the two scatterers is only one gate, which makes tracking of the peaks of the intensity range profile difficult. In addition, there are other weaker scatterers near the range of the second scatterer, but at different crossranges. This is where range/Doppler tracking of the scatterers, rather than range tracking followed by Doppler tracking, would be required. Since our examples utilize the simplest tracking

method, it is preferable to select a scatterer better separated in range even if the height separation should be smaller. Based on an examination of the peaks tracks, we choose the strong scatterers in Range Gates -5 and 3 (indicated by left-pointing arrows), even though the height difference between the selected scatterers is relatively small. Note that a test on tracking the scatterer in Range Gate -4 did indeed result in a compensation of unacceptably poor quality.

The yaw measurement proceeds as before. We track the scatterer in Range Gate 53 in range and then in Doppler, and compensate. In the compensated data, we track the scatterer that used to be in Gate -27 , again compensating in range and then in Doppler. The combination of the two tracks for the second scatterer gives the yaw Doppler. The roll measurement is performed analogously, with the other two selected scatterers. As illustrated earlier, we check whether each scatterer track is adequate. Because the ship is moving rather erratically, as will be seen from the measured yaw and roll Doppler curves, none of the four scatterer tracks is good. The results from these measurements are the yaw and roll Doppler curves of Figure 5.55, with the dashed curve depicting the yaw Doppler. The uncertainty in the yaw Doppler is about ± 6 Hz, and that in the roll Doppler is about ± 4 Hz. The dotted horizontal lines are drawn at the average of these, ± 5 Hz. The large




Figure 5.55 Yaw (dashed) and roll Dopplers for the new time interval.

relative uncertainty, particularly in the center of the figure, means that we may have to improve our scatterer tracks, either by refining the existing tracks or by switching to other scatterers. Before attempting to improve the tracks, we check whether we can form good images from the time intervals with small relative Doppler uncertainty.

For orientation purposes, we would like to know how the roll and yaw Dopplers compare in strength, so that we may judge from the magnitudes of the curves in Figure 5.55 what the consequences on the image will be. Thus we generate an image at a time where both roll and yaw Doppler are strong in Figure 5.55, and compare the magnitude of the shearing of the baseline of the ship with the crossrange spread of the superstructure due to the roll. For an imaging time of 24 seconds, we obtain the image of Figure 5.56. The crossrange spread of the baseline over its length is about eight gates, and the crossrange spread of the superstructure about the same. As seen from Figure 5.55, at the imaging time of 24.0 seconds, the yaw Doppler is nearly three times as high as the roll Doppler. Hence, if the curves of Figure 5.55 are used to judge the yaw and roll effects, we conclude that the roll Doppler is nearly three times as effective as the yaw Doppler. This provides a feel for how easy or difficult it is to generate good topviews or sea-level views.




Figure 5.56 Image from 23.9 to 24.1 seconds.

Examining Figure 5.55, we conclude that a good topview image should be obtained near a time of 21.8 seconds, when the roll Doppler goes through zero. However, the absolute roll Doppler remains less than the roll Doppler uncertainty for nearly a second, centered on 21.8 seconds, so we expect to have to adjust the image time by interpolation. The image from 21.6 to 22.0 seconds is shown in Figure 5.57. This is indeed a topview image. However, we cannot be sure that it is an absolutely undistorted topview image, which means that the responses from the superstructure are in their correct crossrange positions, because some measurement error is involved in obtaining the yaw and roll Dopplers; we have not yet attempted to correct the error in the manner discussed earlier. To do so, we select one imaging time at the first roll extremum of Figure 5.55 and another imaging time at the second roll extremum (indicated by the vertical arrows at 20.7 and 22.8 seconds), both with 0.4 seconds duration. Two prominent responses of the superstructure, indicated by horizontal arrows in Figure 5.57, are separated by 5.3 crossrange gates in the first case, and by 6.3 crossrange gates in the opposite direction in the second case, with both separations measured using the complex image. By interpolation we obtain the time of zero roll Doppler as 21.66 seconds, which is shifted by 0.14 seconds from the measured value. The




Figure 5.57 Image from 21.6 to 22.0 seconds.

corresponding image is shown in Figure 5.58. The responses from the superstructure in the center of the ship differ significantly between Figures 5.57 and 5.58. The superstructure scatterers are within the outline formed by the deck scatterers in Figure 5.58, but not in Figure 5.57. The two scatterers used for interpolation form a single peak in Figure 5.58.

From Figure 5.55 we find that over the imaging interval of 0.4 seconds, the yaw Doppler changes by 20 Hz. Because of weighting for sidelobe suppression, this number is effectively halved to 10 Hz. With a duration of 0.4 seconds, each crossrange cell has a width of 2.5 Hz. Thus, the relative positions of the two scatterers vary by four crossrange cells during the imaging interval. This implies that the responses in Figure 5.58 have too much phase curvature for accurate position measurements, which a test shows to be the case. In order to measure positions and characteristics of these responses, one must either reduce the processing time or compensate the responses differently along the length of the ship by polar reformat processing. As indicated by the image of Figure 5.58, reducing the imaging interval is impractical because resolution along the width of the ship would be too low.

Instead of estimating how long an imaging interval can be chosen for a particular scatterer, we can easily use the measured yaw and roll Doppler




Figure 5.58 Image from 21.46 to 21.86 seconds.

curves to *determine how long an imaging interval is acceptable* for generating a specific topview or sea-level-view image. The yaw Doppler curve was derived from two scatterers in known range gates. From any image, we can determine the ratio of the range separation of the two scatterers and the total range extent of the observable responses of the ship. The yaw Doppler curve for the entire ship is that of the two scatterers, scaled up by this ratio. At any selected time, we then can find the maximum imaging interval usable without further motion compensation, by requiring that the change in the scaled yaw Doppler over the imaging interval be no more than about the reciprocal of the imaging interval, maximally perhaps up to twice that value. Similarly, from any sea-level-view image we can determine the crossrange separation of the two scatterers used for the roll measurement, as well as the crossrange extent of all observable scatterers. The measured roll Doppler is scaled by the ratio of crossrange extent to crossrange separation, and the allowable imaging interval at a given instant can be determined in the same way as done for the yaw Doppler.

The central part of Figure 5.55 shows a more erratic motion as well as decreased maximum Dopplers. The decrease of the absolute peak Dopplers means that longer signal-processing intervals must be used for a given crossrange resolution, yet this conflicts with the shorter motion cycles. Thus, although we can generate the same types of image, the quality will be lower. This is true even for the remainder of the interval. For example, at 36 seconds (indicated by the solid double-headed arrow), we have a similar situation as at 21.6 seconds (labeled 5.57), but with important differences. The peak yaw Doppler is only a fraction of that at the earlier time. In fact, it is so small that the errors in scatterer tracking may become significant. Second, whereas at the earlier time we could form sea-level-view images before and after the predicted instant of zero roll Doppler, the peak of the roll Doppler at 36.4 seconds is too low to permit generating a usable sea-level-view image.

The essential point of the preceding illustrations is that the yaw and roll Doppler curves permit one to select an imaging time that yields a specific type of ship image. They also allow one to judge the image quality before the image is actually generated. Lastly, the curves permit one to determine whether or not an imaging interval requires further motion compensation.

5.2.3 Very Little Roll Motion

The coast guard cutter used for the preceding illustrations tends to roll more than it yaws, so for the present illustration we select data from a cruise ship, shown in Figure 5.59, that tends to yaw more than it rolls. In fact, the roll is



Figure 5.59 Cruise ship.

so small that it is difficult to measure. The peaks tracks over a 20-second interval are shown in Figure 5.60. To again summarize the required




Figure 5.60 Peaks tracks for the cruise ship.

processing steps, the first step is to select candidate bow and stern scatterers for the yaw measurement, and evaluate the quality of the tracks associated with these responses. As discussed above, quality is evaluated by compensating the data with the track, then examining the image cut in the range gate of the response. If the quality test fails, we must track one or more additional responses until we have found an acceptable pair. The next step is to form short-term images when the yaw Doppler is near zero, until we find one that gives a sea-level-view image of usable crossrange extent. If no image is usable, we slide a short-term image window through the data, searching for an image in which the superstructure has usable extent. If we find a usable image, we verify the acceptability of the positions of the scatterers used for the yaw measurement and select candidate responses for the roll measurement. We then track these responses, verifying the quality of the tracks. In the present instance, we obtain a usable image with significant roll effects only near time 139.1 seconds. As a slightly different 20-second interval would not yield any usable image, we consider how to proceed in such a case.

If, in this imaging step, we do not obtain an image with significant roll effects, we assume that the roll is weak, so that the yaw measurement will not be falsified by the roll, even if the two scatterers should not be at the same height. One might argue that there may not be any yaw either, but this is extremely unlikely because the yaw contains the component due to the translational motion of the ship. There will be no yaw only if the ship goes absolutely stably in a straight line toward the radar. In this extremely unlikely situation, even ignoring that the beam depression angle will not be zero, we still obtain a range profile. The difference with respect to a topview image is merely that the ship image will have no width. In any case, in such an unlikely event a two-dimensional image can be obtained only by waiting until the situation is more normal.

In Figure 5.60, the most logical candidates for the yaw measurement are the peaks tracks starting in Range Gates -53 and 40, indicated by right-pointing arrows. Range-tracking and Doppler-tracking the reference scatterer, compensating, and then range- and Doppler-tracking the second scatterer in the data where the reference scatterer was compensated (in each case verifying the qualities of the tracks), presents nothing unusual to discuss. The result of the yaw Doppler measurement is shown in Figure 5.61. The conspicuous fact in this figure is the slowly decreasing Doppler bias, which implies that the ship is ending a turn. It appears to be a slow turn, because even at the beginning of the interval the periodic yaw Doppler is about as large as the turning Doppler. If we image at the time of the maximum yaw Doppler, we obtain the image of Figure 5.62. It is clearly a topview image,




Figure 5.61 Yaw Doppler.




Figure 5.62 Image from 122.2 to 122.6 seconds.

with most of the nonilluminated far side of the ship missing (as would be expected from the design of a cruise ship). However, we cannot tell from this image whether the strong responses of the superstructure, in the lower half of the image, are actually in the positions along the width as given by the image. The crossrange positions may be shifted by a slight roll, so that it may not be an undistorted topview image.

In order to place the responses of the superstructure in their correct positions on the deck, we need an image with significant roll effects. In such a situation we may try different imaging times spread over the observation interval in order to obtain a roll image (guided by the yaw measurement), but we already stated above that no roll image can be generated over 19 seconds in this case. The next step then is to *test whether the roll is at all significant*. The image of Figure 5.62 was made at the time of the first yaw Doppler peak of Figure 5.61. If we shift the imaging time in small steps, perhaps in increments of about 5% of the yaw period (0.2-second steps), if there is any significant roll at all, some of the strong responses of the superstructure will shift in crossrange, differently than the responses on the deck. If this does not happen, then we conclude that the roll is totally insignificant, and that the image of Figure 5.62 gives the responses of the superstructure in their correct positions. In this instance, however, the test does show a gradual translation of the responses, so that the roll motion is not completely insignificant. We should measure it in order to obtain the correct placement of some of the major responses of the ship, which may be important for identification. However, generally we must not assume that the roll is necessarily measurable; it may be too weak. Thus we must always test the measurability of the roll Doppler.

To perform such a test, we start with the response that was found to shift most with a translation of the imaging time, which appears in Range Gate -22 in Figure 5.62, indicated by a left-pointing arrow. The corresponding scatterer is chosen as the reference scatterer for the roll measurement. We go to the peaks tracks of Figure 5.60 and range-track the peak that appears in Range Gate -22 at the imaging time of 122.4 seconds, also indicated by the left-pointing arrow. This is the strongest response, so that range-tracking, Doppler-tracking, and the compensation are straightforward. Now we need a second scatterer at a different height for the roll measurement, but we already have determined that the roll is so small that no usable roll image can be generated over nearly the entire observation interval. Thus, we use the bow and stern tracks to generate a virtual scatterer in the range gate of the reference superstructure scatterer.

In our example we already found that the responses of the superstructure shift with respect to the edge of the ship when the imaging time is changed, so that the roll Doppler should be measurable. By performing the measurements indicated above, we obtain the roll Doppler curve for the second scatterer, as shown in Figure 5.63 by the dashed curve, with the yaw Doppler of Figure 5.61 superposed. The tracking uncertainties for the two curves are small and nearly equal, as indicated by the dotted horizontal lines. The correlation between the two curves indicates that our roll measurement is contaminated by yaw contributions. In other words, we were not able to perfectly place the virtual scatterer. However, near time 139.1 seconds, the yaw Doppler is near zero and much smaller than the roll Doppler. This is potentially a time when we can form a roll image. The 0.8-second image at 139.1 seconds is shown in Figure 5.64. This is an excellent sea-level-view image, as can be seen by comparison with the line drawing also in Figure 5.64. We can now also select an imaging time from Figure 5.63 when the roll Doppler goes through zero. If there is a choice, we select the time when the yaw Doppler is as high as possible. From Figure 5.63, a good time is 123.2 seconds.




Figure 5.63 Roll Doppler (dashed) and yaw Doppler.




Figure 5.64 Roll image, from 138.7 to 139.5 seconds.

The corresponding 0.4-second image is given in Figure 5.65. The images of Figures 5.62 and 5.65 were generated only 0.8 seconds apart, yet the latter image is of higher quality. The comparison of the two images shows that the roll motion does displace the responses of the superstructure in the image of Figure 5.62 to the right. Despite the fact that in this example the roll motion is very much smaller than the yaw motion, we have been able to measure the roll Doppler accurately enough to allow selecting the correct imaging time for an undistorted topview image.

5.2.4 Turn Maneuver

As already stated, the use of Doppler processing makes the analysis very susceptible to small but abrupt changes in the range rate, so that the analysis cannot be extended through such changes. In practice, the analysis can be performed only over *time intervals wherein the Doppler changes smoothly*. The nature of the Doppler changes thus must be recognized, and the time intervals for analysis must be chosen accordingly. The most likely occasion when abrupt Doppler changes may occur is during a turn maneuver. We analyze




Figure 5.65 Image from 123.0 to 123.4 seconds.

a section of 20 seconds duration when the cruise ship performs a rather sharp turn.

The range peaks tracks for that interval are shown in Figure 5.66. Examination of the tracks indicates that something unusual happens at 293.8 seconds and again at 297.5 seconds (the times of the dotted vertical lines). Some of the discernible peaks tracks are disrupted at these points. However, it is not necessary that one examine the peaks tracks by eye to discover these events. They can equally well be recognized when any one of the peaks tracks is followed and compensated, be it done manually or automatically. We illustrate this fact on the track indicated by the top arrow in Figure 5.66, which does not reveal anything unusual at 293.8 seconds. This track can be fairly well followed until the time of 298.0 seconds, beyond which some extrapolation is necessary. Range tracking does not reveal any problem, but a range compensation based on a polynomial fit to the track gives the Doppler track of Figure 5.67. There is a clear jump in the Doppler at 293.8 seconds (indicated by the left arrow), and no analysis can be performed across such a jump. Either the tracking must be improved so as not to generate such a jump (likely from interference between two close scatterers),




Figure 5.66 Peaks tracks during a turn.




Figure 5.67 Doppler track after range track.

or a different scatterer must be selected for tracking if this particular time is important. Note that the Doppler jump of about 2 Hz corresponds to a range rate change of only 3 cm/s at X-band.

The plot of Figure 5.67 also confirms some event around 297 seconds, although not as clearly as at 294 seconds. Since both the range track and the Doppler track are very poor in the interval beyond 297.0 seconds, we should choose a range track that can be better followed in this region. For example, over the last five seconds of Figure 5.66 we can choose the track ending in Range Gate -15, indicated by the left-pointing arrow. Neither the range track nor the Doppler track reveals problems an automated processor could reliably distinguish. Thus we complete the range and Doppler tracks over the five seconds, and (as required) check the image cut in the range gate of the tracked response. This image cut is shown in Figure 5.68, together with its transform. The amplitude function of the transform shows a prolonged drop over almost the first half of the time interval. This is the crudest criterion for failure to track properly. We can analyze only the second half of the interval, which starts near the time of 298.0 seconds. The overall conclusion is that if the interval from 280 to 300 seconds is to be analyzed, it must be broken into three segments for the analysis, from 280.0 to 293.8 seconds, 293.8 to 297.5 seconds, and 297.5 to 300.0 seconds.

To measure the yaw motion during the first segment, we select the tracks at the closest and farthest range in Figure 5.66 (indicated by the solid right-pointing arrows), which means scatterers separated by almost the full




Figure 5.68 Image cut in the range gate of the tracked scatterer.

length of the ship. This gives the yaw Doppler of Figure 5.69. Note that the average absolute yaw Doppler is over 100 Hz, because of the turn maneuver. This means that we will obtain a topview image anywhere within this interval. However, even though the roll may be small, it could make the difference between a high-quality and a poor quality image. The concern is not so much that the roll could distort the image, but that a changing rotation axis may cause too many spurious responses. Thus we again try to measure the roll motion.

Because the target is turning so rapidly, as evidenced by the large change in the range extent of the peaks tracks of Figure 5.66, we must take care that the yaw motion does not corrupt the roll measurement. The best trackable response in the center of the target is initially in Range Gate -5 in Figure 5.66 (indicated by the upper dashed arrow). The nearest well-trackable response is initially in Range Gate -19 (indicated by the lower dashed arrow). When we examine the Doppler shifts of these responses between closely spaced topview images, we find that the responses are located at different heights on the superstructure. Because the deck positions of the two scatterers are not closely spaced, we cannot derive the roll motion from their differential Doppler; the difference will be corrupted by yaw motion.




Figure 5.69 Yaw Doppler.

We must attempt to generate a virtual scatterer close to one of the two superstructure scatterers.

The range peak tracks of Figure 5.66 show that the range separation of the top track and the response initially in Range Gate -5 is initially 55% of the range separation of the top and bottom tracks, and this ratio steadily decreases to 45% at a time of 294 seconds. If the response in Range Gate -5 were along the physical line between the top and bottom responses, the ratio would be constant. The large variation indicates that it is far off the line, so that we likely will not be able to place a virtual scatterer close enough to avoid corruption from residual yaw motion. The corresponding ratio for the response initially in Range Gate -19 varies from 73% to 69%. As we will see below, this is acceptable. If it were not, we would have to search for another superstructure response.

As time passes and the ship approaches broadside, the displacement of the response from the line has a larger effect on the ratio. We therefore use the earlier value when we calculate the Doppler history of our virtual scatterer. Subtracting this from the measured Doppler of the superstructure scatterer gives the roll Doppler of Figure 5.70. The dotted horizontal lines give the uncertainty due to tracking, but do not account for uncertainty due to scatterer positions. From Figure 5.69, we see that a 1% change in the ratio




Figure 5.70 Roll Doppler.

used to set the position of the virtual scatterer will shift the derived roll Doppler by about 1 Hz, but leave the shape essentially unchanged. Therefore, the most accurate time for obtaining a good topview image is about 287.4 seconds, when the curve passes most steeply through zero Doppler.

As explained earlier, we want sufficient but not too much crossrange resolution, since the ship's motion might not be smooth, in particular during a turning maneuver. Figure 5.69 shows that the yaw Doppler varies at the rate of about 50 Hz per second at the imaging time. If we use an imaging interval of 0.2 second, the Doppler changes by 10 Hz over the imaging interval, but because of the weighting for sidelobe suppression it changes effectively by only one half this amount, or 5 Hz. This is just the reciprocal of the imaging interval, which is acceptable. The corresponding image is shown in Figure 5.71, with arrows indicating the tracked scatterers, as in Figure 5.66. Even the peaks plot image lets us recognize the shape of the ship.

For comparison, we also generate an image with the same parameters at the small double-hump peak of the roll Doppler at the time of 288.9 seconds, or a separation of just 1.6 seconds from the earlier image. The new image is shown in Figure 5.72. The comparison of the two images confirms our expectation that even relatively small roll Dopplers degrade the image quality. Tests showed that, based on the roll Doppler curve of Figure 5.70,




Figure 5.71 Image from 287.3 to 287.5 seconds.




Figure 5.72 Image from 288.8 to 289.0 seconds.

the images made at zero roll Doppler have better quality than those made at finite roll Doppler.

The peaks tracks of Figure 5.66 are greatly disturbed in the second interval, from 293.8 to 297.5 seconds. Although the yaw Doppler can be measured without difficulties, this is not so for the roll Doppler. When the roll Doppler is measured for several scatterers, in accordance with the above procedure, the resulting measurements have common elements but also differences. Thus we conclude that with the present type of tracking of scatterers, that is, simple tracking of the intensity peaks of the range profiles, a reliable roll measurement is not possible for this data segment. Unless the tracking is improved with the use of the complex range profile or, even better, with combined range/Doppler tracking, the specific imaging times for high-quality images cannot be chosen. This is not surprising, because the ship is being viewed near broadside.

5.2.5 Dive Boat at the End of a Turn

As another example, we consider a very small ship (a dive boat with a length of 47 ft) at the end of a turn, so that the motion is rather uneven. It is a

limiting case for tracking scatterers on the basis of the peaks of the intensity range profile, and tracking should be done with the enhanced methods. The ship is shown in Figure 5.73.

The peaks tracks over a 20-second interval are shown in Figure 5.74. There is only one good track, that of the strongest scatterer of the ship, starting in Range Gate -24 (indicated by the dashed arrow). There is another usable track at the farthest range (indicated by a solid arrow), provided we fit through the interleaved tracks of two scatterers in the second half of the time interval. This is easier after removing the periodic translation of the entire ship, which we do by compensating the strong scatterer. The fit to the resulting peaks track at the farthest range is indicated in Figure 5.75. The track next in quality is the one at the shortest range (indicated by a solid arrow in Figure 5.74), where we would fit through the bulge around 23 seconds.

In order to measure the yaw motion, we first compensate the strongest track in range and then in Doppler, so that it becomes the reference. Next, we range- and Doppler-track the scatterer at the largest range. Lastly, having compensated the top track in Figure 5.74, we range- and Doppler-track the bottom scatterer, in which case the combined spline gives the yaw Doppler as measured between the two scatterers at the extremes of the ship.

The Doppler measured between the top and bottom scatterers in Figure 5.74 is shown in Figure 5.76. As we will verify later, the ship has a totally insignificant roll motion within this observation interval. Hence, if the tracking of all three scatterers were perfect, the Doppler function derived between the first two scatterers should be a scaled version of the function of



Figure 5.73 Dive boat.




Figure 5.74 Peaks tracks for the dive boat.




Figure 5.75 Fitting through gaps.

Figure 5.76. The only differences actually found were that the zero crossings were off by 0.1 seconds. This is harmless because when the accuracy is in




Figure 5.76 Doppler between the scatterers at the extremes.

question we determine the exact time of zero yaw Doppler by forming short term images slightly before and after the measured value, and interpolating the times from the Doppler spread of the sheared images, as already demonstrated. Also, the average Doppler over the first four seconds was slightly rising rather than being constant as in Figure 5.76, which again is immaterial. In view of the tracking problems pointed out above and the simple type of tracking used, this excellent agreement is noteworthy. This is not an accident, because the quality of the track is always verified via the image cut in the range gate of the compensated scatterer. We always know whether or not the tracking has worked, so that if the track is found to be acceptable we can rely on the Doppler measurement.

In order to measure the roll motion, we must track a scatterer on the superstructure. We attempt to select one by forming images at the times of zero yaw Doppler given by Figure 5.76, when the crossrange spread of the responses is due to the roll motion of the ship. However, the crossrange spreads of these images are too small to recognize the superstructure. Slightly varying the image times does not improve the situation. Evidently, the roll motion is small at the same time as the yaw motion. In Section 5.2.3, we saw that when this is the case, we may be able to recognize superstructure

scatterers by forming closely spaced images when the yaw is strong and smooth, and then measuring Doppler shifts of responses from image to image. If the roll motion is also strong, the superstructure responses may have recognizably larger Doppler shifts than the deck responses. When we apply this procedure to the current data, we still are unable to recognize a superstructure scatterer.

The fact that we cannot recognize any roll effects at the time of strong yaw Doppler means that we can form a good topview image at this time. However, we cannot measure (scaled) scatterer heights by varying the image time. If possible, we would like to extract this information from the data. Although we have determined that the roll motion is too weak to allow this when the yaw Doppler is near zero and when it is strongest, the possibility still exists that the roll motion may be strong enough at other times. As we are unable, in this case, to recognize superstructure scatterers from short-term images, we are forced to try to *recognize the superstructure scatterers from their Doppler histories*.

We want to track additional scatterers, calculate their Dopplers relative to the top scatterer in Figure 5.74, and compare the relative Dopplers to the yaw Doppler of Figure 5.76. Scatterers whose motion contains a significant roll component will yield relative Dopplers that are not scaled versions of the yaw Doppler. We can apply the virtual scatterer procedure to each of these, and obtain the roll motion for the corresponding scatterer. If the roll is strong enough that the measured scatterer roll motions agree, we have found the roll motion of the ship. Unfortunately, in this instance, the combination of poor tracks and a small roll Doppler make it impossible to derive the roll Doppler with sufficient accuracy.

Because of these results, we can make a topview image at the time of maximum Doppler, which Figure 5.76 gives as 30.3 seconds. We have shown earlier that it is important to form the image when the motion is only about a fixed axis, that is, when the roll Doppler goes through zero. In this example we cannot consult the measured roll Doppler because it is small, but then it should also be irrelevant. However, the range tracks also indicate when some sudden disturbance occurs, not necessarily a roll motion. For example, in Figure 5.74 there is such a disturbance, as evidenced by the top track, at the time of 31.0 seconds (indicated by the dotted vertical line), only 0.7 seconds after the yaw Doppler goes through its peak. For comparison of the image quality at the two different times, in Figure 5.77 we show a 0.6-second image for a center time of 30.3 seconds, and in Figure 5.78 for a center time of 31.3 seconds, or 1 second later. Although the examination of intensity images does not allow one to judge the details that determine image




Figure 5.77 Image from 30.0 to 30.6 seconds.




Figure 5.78 Image from 31.0 to 31.6 seconds.

quality, the differences between the two images are large enough to let us recognize the higher quality of the image of Figure 5.77.

5.2.6 Small Ship in Rough Seas

As an example of difficult conditions, we consider an 83-foot-long coast guard cutter at Sea State 4. The new ship (actually, a different ship of the same type as the one on which data was collected) is shown in Figure 5.79. We expect scatterer tracking to be much more difficult in a rough sea, in particular for a small ship. It would be easy to state that the way to overcome the tracking problem is to use combined range/Doppler tracking rather than range tracking followed by Doppler tracking, in particular with range tracking implemented by simple peaks tracking, but this most sophisticated type of tracking also has its limitations. Thus, in this section we consider high-quality imaging under severely limiting conditions.

If a ship can be observed over a long time, then the easiest way of avoiding tracking problems is to search for a time interval when the simple tracking of the intensity peaks of the range profile is adequate. As an illustration, in Figure 5.80 we show the peaks tracks of the ship over a 20-second interval in which scatterer tracking is possible despite the rough sea. When one attempts to track the scatterers needed for the yaw and roll Doppler measurements, one finds that the quality of these peaks tracks is much poorer than that of any of the earlier peaks tracks used in our illustrations. Aside from the difficulties caused by the heavy motion of the ship, range tracking is further aggravated by the strong returns from the ocean surface, which interfere with the returns from the ship in the same range gates. Of course, this particular



Figure 5.79 Smaller coast guard cutter.




Figure 5.80 Peaks tracks for the smaller coast guard cutter.

problem largely disappears with combined range/Doppler tracking, because the ocean returns generally fall in different crossrange gates than the ship returns.

The peaks track of the strongest scatterer, indicated by the left-pointing arrow, starts in Range Gate -2 . One of the times at which even this strong return is disturbed is around 410 seconds. Figure 5.81 gives a 0.2-second image centered at a time of 409.5 seconds. The main ocean clutter is in about Range Gate -17 . Figure 5.82 shows another short-term image just 0.2 seconds later. The plotting levels in the two figures were adjusted so that the returns from the ocean in front of the ship are about equally strong in both figures. This makes a comparison of the two figures meaningful, with the conclusion that in Figure 5.82 the ocean clutter is more widely spread in range. Examination of the 20-second interval shows that it contains three short subintervals, each roughly of 1 second duration, when the ocean clutter is so weak that it does not cause any problem. It then is comparable to the clutter in our earlier illustrations. Thus, the ocean clutter varies throughout the entire 20-second interval (as it does over the other 20-second intervals of these data), from insignificant to very detrimental.

If the image of the ship occupied a specific crossrange interval, the clutter outside this interval (most of the clutter) could be suppressed before




Figure 5.81 Image from 409.4 to 409.6 seconds.




Figure 5.82 Image from 409.6 to 409.8 seconds.

generating the peaks tracks. In fact, the ship's motion about its center of gravity continuously changes the crossrange width of the ship, so that such an approach becomes impractical, at least for interactive processing. Hence, the only remedy is to use crossrange resolution for clutter suppression by tracking the scatterers in combined range and Doppler. This is the method we demonstrated in Chapter 2, but cannot yet use systematically because it has not been integrated into our interactive software. Thus, in this example we will merely demonstrate that the method of imaging time selection used on the other ships works even under such adverse conditions, and when we use the simplest procedure of tracking the peaks of the intensity range profile. Under these adverse conditions, this crude method will sometimes work with manual processing and at other times not, but automated processing may require range/Doppler tracking or the modified method described in the flowchart of Figure 5.32 and illustrated subsequently.

When three scatterers (bow and stern indicated by solid right-pointing arrows and superstructure by dashed or short right-pointing arrows in Figures 5.80, 5.81, and 5.82) are used to derive the yaw and roll Doppler curves in the manner illustrated earlier, we obtain Figure 5.83. The bow and stern scatterers are easily trackable, as indicated by the horizontal line labeled "y." However, the superstructure scatterer is much more problematic.




Figure 5.83 Yaw (dashed) and roll Doppler curves.

Figures 5.81 and 5.82 show that no range gate contains just one superstructure scatterer. The differing motions of the multiple scatterers in each gate make tracking any one difficult, and limit how well one can estimate residual motions. The horizontal line labeled “r,” which is an overestimate of the tracking uncertainty in the roll motion, reflects this limitation. Despite the poor quality of the peaks tracks of Figure 5.80 and our inability to estimate roll Doppler error, the Doppler functions of Figure 5.83 do tell us when we can generate pure sea-level-view images and pure topview images. When the precise imaging time must be adjusted by interpolation between images at an earlier and later time, the imaging time is shifted only minimally from the values extracted from Figure 5.83.

An example of a topview image is given in Figure 5.84. It is a good topview image for two reasons. First, it was generated at a time at which the ocean clutter caused by the ship is very benign. Second, at this imaging time the yaw Doppler has about the largest value within the entire 20-second interval. However, because of the severe motion of the ship, the precise imaging time is very critical. As an illustration, Figure 5.85 shows the topview image only 0.1 seconds earlier. It clearly is much inferior to that of Figure 5.84. A significantly inferior topview image also is obtained if the imaging time is delayed by 0.1 seconds. Evidently, in high sea states it is even




Figure 5.84 Image from 405.0 to 405.2 seconds.




Figure 5.85 Image from 404.9 to 405.1 seconds.

more important to select an imaging time when the ship is rotating about a fixed axis.

By forming a sequence of images throughout the entire 20-second time interval, one finds that the ocean clutter generated by the ship is correlated with the ship's roll/pitch motion (We lumped roll and pitch together for purposes of sea-level views, but for the generation of ocean clutter under the present circumstances the pitch component may be more important). Whenever the roll/pitch Doppler goes through zero starting from a negative value, the ocean clutter induced by the ship is strong. When the roll/pitch Doppler goes through zero starting from positive values, the ocean clutter tends to be weak. However, even during the periods of high ocean clutter generated by the ship, we obtain the expected images. For example, Figure 5.86 shows a sea-level-view image in strong ocean clutter from the ship. Note that the declivity of the ship has no crossrange spread.

We now select a considerably worse case from the same data file, with the worsening of the problem due to a much larger aspect angle rather than a particularly bad motion of the ship. The large aspect angle causes the scatterers to be bunched in range, so that a range cell on the superstructure is likely to contain more than one significant scatterer. The peaks tracks over a




Figure 5.86 Image from 410.4 to 410.6 seconds.

20-second interval are shown in Figure 5.87. To appreciate the severity of the tracking problem, the reader should try tracking any one of the peaks “tracks” of Figure 5.87 by eye. Even the strongest track (indicated by the left-pointing arrow), around Range Gate -10 , becomes rather poor over some intervals. (Note the crude range gate scale.) This peaks track is separately shown in Figure 5.88, with most of the outliers removed. It turns out that the most critical section for the motion measurement and imaging is around 3 to 4 seconds (indicated by the dotted vertical lines), where adequate tracking without crossrange resolution is impossible even manually. Even if it were possible to track the peaks of the range profiles, range tracking followed by Doppler tracking is not meaningful when multiple scatterers are moving within and through a particular range gate.

Some clarifying remarks concerning the peaks track of Figure 5.88 are in order. In connection with aircraft imaging we emphasized that the spline fitted to the peaks track must not be too flexible, since the following Doppler track cannot correct mistracks in range beyond some value. The optimum degree of flexibility of a polynomial or a spline depends on the type of target and its behavior, and may not be predictable for a given tracking situation.




Figure 5.87 Peaks tracks at a large aspect angle.




Figure 5.88 Strongest peaks track.

The remedy then is to try different degrees of flexibility, in each case checking the quality of the track. As pointed out repeatedly, this is done by

forming an image after range and Doppler tracking (or combined range/Doppler tracking), and checking whether the transform of the response of the tracked scatterer has a sufficiently constant amplitude. However, trying fits with both progressively more and less flexible splines than that of Figure 5.88 also did not provide usable tracks. They are good only in the regions where the peaks track of Figure 5.88 appears well defined to the eye. If straightforward tracking of the scatterers is to be used under such poor circumstances in order to generate the yaw and roll Doppler curves, it must be combined range/Doppler tracking (using a sequence of images over very short intervals). In this case, each time the response from a particular scatterer is observed in a short-term image, we obtain a range and a crossrange measurement. Since the PRF of the radar is known, the crossrange measurement can be converted into Doppler and range rate, and the latter allows one to predict the range position of the response peak in the next short-term image. This facilitates tracking the response of a specific scatterer among changing sets of responses. Taking the difference between the crossrange positions of the bow and stern scatterers, we have the yaw Doppler. Similarly, the difference between the crossrange positions of the two responses of the superstructure (or one superstructure and one virtual response) gives the roll Doppler. It is evident that, for best performance, instead of tracking the response peaks in the short-term images we should analyze each peak with the TSA and track actual scatterer positions.

Even the most sophisticated of all tracking procedures has its limitations as sea state and aspect angle increase. As an indication of the potential problems, we implement the first step of combined range/Doppler tracking of the data represented by Figure 5.87, but without associating the responses from one look to the next. We use a processing window of 0.5 seconds duration, which is about the limit for a motion cycle duration of roughly three seconds when no motion compensation is employed. The processor forms individual images over this processing window, and measures the range/Doppler positions of the stronger peaks. The result is shown in Figure 5.89, with the slope of each line segment indicating the measured range rate for an image peak at the time and range of the segment. What is missing in the figure is an indication of how, based on the range rate measurements, the short tracking sections for individual scatterers are to be associated in order to establish tracks. An attempt to do this by eye for any one of the "tracks" of Figure 5.89 fails (aside from the smallness of the details). However, it is likely that an association process that examines likely associations and chooses among them based on the constancy of the resulting signal amplitude can be automated. The performance could also be improved by




Figure 5.89 Combined range/Doppler tracks.

optimizing the length of the processing window with respect to the duration of the motion cycles, even making it adaptive. Nevertheless, with increasing sea state and aspect angle the process will eventually fail. We now describe a simpler alternative procedure.

With increasing aspect angle, the tracking difficulties for scatterers on the superstructure become severe, because of the relatively high density in range of scatterers. With scatterers moving in the same range cell, range tracking followed by Doppler tracking usually will not work. On the other hand, there are relatively few significant scatterers near the bow and the stern, so that the tracking of bow and stern scatterers for the yaw Doppler measurement will generally work. Thus we can measure the yaw Doppler, and select time sections when the yaw Doppler is strong enough to generate a usable topview image. Then we can perform a *search for the time of zero roll Doppler in a way similar to the interpolation measurement* discussed earlier, thus avoiding the direct measurement of the roll Doppler and its zeros. This will be illustrated for the case at hand.

For the yaw measurement, we select the tracks starting in Range Gates -22 and 16 in Figure 5.87, indicated by right-pointing arrows. Even though both tracks appear difficult to follow, we have two checks available, allowing

us to iterate if necessary. First, as the usual check, when we use a scatterer track to compensate the data, form an image, and take the transform of the fixed-range image cut through the tracked scatterer, the amplitude function must not drop to the noise level over extended time intervals. If that should be the case, the tracking can be repeated with an alternative branch at the point where a tracking problem appeared. Second, after deriving the yaw Doppler curve, we can check whether at the indicated times of zero yaw Doppler the declination of the ship's image is indeed vertically oriented (constant Doppler). Thus there need not be an uncertainty as to whether the tracking has worked sufficiently well in critical situations.

The yaw Doppler derived in this manner is shown in Figure 5.90. We recognize three sections with strong yaw Doppler (shown by double-headed arrows), from 0 to 6 seconds, around 16 seconds, and around 19.6 seconds. Since we always want to image when the motion is least violent, the best section for generating a topview image is the first one. We now must find the precise imaging time at which the roll Doppler is zero, but this cannot be done in the way applicable for more benign motion conditions, because we cannot track a scatterer on the superstructure.




Figure 5.90 Yaw Doppler from tracks starting in ranges -22 and 16 of Figure 5.87.

As a first step, we form a short-term image (without additional motion compensation) within the selected interval from 0 to 6 seconds. We avoid the beginning of the 20-second interval, because of possible spline fitting problems, and also the times when the yaw Doppler drops sharply. Thus we choose the center of the six-second interval, or an imaging time of three seconds. The length of the imaging interval is selected so that the crossrange spread of the superstructure due to roll is large enough to be distinguished from the spreading of the deck due to yaw. This image, together with the ocean clutter generated by the ship, is shown in Figure 5.91. The ship's image is a combination of topview and sea-level view.

In the same manner as one can recognize the responses from the superstructure by eye, one can find the response with the largest crossrange shift away from the deck automatically. This requires that the ocean clutter responses be recognized, which can be done both on the basis of their distribution and their individual properties (the signals corresponding to ocean responses are usually poor approximations to one- and two-scatterer patterns). By shifting the imaging interval slightly, one can use the shift of the superstructure responses to determine whether reversing the sign of the crossrange shift requires going to an earlier or later imaging time. In this instance




Figure 5.91 Image from 2.8 to 3.2 seconds.

it is a later time. We increase the imaging time in steps until the superstructure responses are shifted to the left of the deck. With a shift of the imaging interval by one second we obtain the image of Figure 5.92. If one measures the crossrange offsets of the response from the top scatterer relative to the centerline of the deck in both images, interpolation gives an imaging time of 3.27 seconds for a topview. However, with the uneven motion of a small ship in rough seas, such an interpolation may not be accurate. Thus we note the range gate of the response from the top of the superstructure, and fine-adjust the imaging time so that this particular response is on the centerline of the deck. This should produce the highest quality image because the top of the superstructure is likely to be near the the centerline. Note that measuring the crossrange shift of a scatterer on the superstructure relative to the centerline of the deck is essentially equivalent to the earlier use of a virtual scatterer, derived from bow and stern scatterers.

The resulting topview image, from 3.20 to 3.60 seconds, is shown in Figure 5.93. The crossrange scale was chosen so that the stern is approximately perpendicular to the centerline. Although this does not appear to be the case in the image of Figure 5.93, the reason is that response peaks rather than scatterer positions are presented in the image. We verified that the




Figure 5.92 Image from 3.8 to 4.2 seconds.




Figure 5.93 Topview image.

response from the left corner of the stern is the composite from two responses (as measured with the TSA), so that its location is not precisely that of the corner of the stern. Based on this crude adjustment of the aspect angle, we obtain a length of the ship of 90 ft, as compared with an actual length of 83 ft. One cannot expect a better accuracy under the present conditions, at least not when the range tracker does not supply the aspect angle of the ship.

In summary, the methods of image-time selection also work in rough seas, and for relatively large aspect angles, although one must adapt the precise procedures to the more difficult conditions. The tracking of individual scatterers becomes more demanding when a high sea state combines with a large aspect angle. Although combined range/Doppler tracking will give adequate results under worse conditions than range tracking followed by Doppler tracking, it is unlikely to provide these results under all conditions. Then one must modify the procedure as illustrated in this section, although combined range/Doppler tracking might sometimes be necessary for tracking the bow and stern scatterers. The other difference relative to calmer seas is that selecting the optimum imaging times is more critical for generating a usable image.

As a last illustration, we consider the most problematic situation. The task of forming a usable image becomes most difficult if a high sea state combines with a large aspect angle of the ship, in the limit the broadside aspect. It is clear that the problem would diminish if one were willing to improve range resolution of the radar, since smaller range cells facilitate tracking of scatterers. However, it is not clear that the problem would be totally solved. When the range resolution cell becomes smaller than the range extent of the dominant scatterers, the effective backscattering behavior of the scatterers changes (see Section 1.2.2). We do not have suitable real data to pursue this question for a ship. Nevertheless, even if there should be problems with very high range resolution (considerably less than 1 ft), it would not imply that identification is impossible. The heavy motion that might make tracking of scatterers impossible also causes such drastic changes in the attitude of the ship that high-quality tracking of scatterers over long periods is not needed. Images with sufficient crossrange resolution for identification can be formed over very brief intervals of smooth motion, between the drastic changes. This is illustrated subsequently.

In the long run of the coast guard cutter in Sea State 4, we select the worst 20-second data segment, when at one point the ship is viewed very close to broadside. The scatterer at the closest range typically should have a sufficiently low interference in its range gate to be at least roughly trackable under these worst-case conditions; but this is not a requirement for what we want to demonstrate. It merely facilitates the illustration. We do need to analyze scatterers in the general vicinity of the bow and stern in order to extract the yaw motion. If the ship is so close to broadside that the scatterers at closest range are from midship, identification may not be possible. We compensate this scatterer, obtaining the peaks tracks of Figure 5.94, where the tracked scatterer now is in Range Gate 0, indicated by the arrow. It is clear from Figure 5.94 that we cannot track the other two scatterers needed for a complete motion measurement. On the other hand, because the ship's motion is so strong, it lets us *recognize the motion behavior from sections of peaks tracks*. The peaks tracks can be somewhat improved by combined range/Doppler tracking, performing short-term imaging in order to obtain crossrange resolution. This processing was performed on the data of Figure 5.94, with a window length of 0.2 seconds. The tracks from this procedure are shown in Figure 5.95 for the 20 strongest peaks at every window position. A comparison of Figures 5.94 and 5.95 shows that the recognizability of pieces of scatterer tracks has improved.

We can reason that the scatterers at the far ranges of Figure 5.95 do not belong to the superstructure, unless the ship is viewed nearly exactly at




Figure 5.94 Peaks tracks of the smaller coast guard cutter under worst-case conditions.




Figure 5.95 Tracks after combined range/Doppler processing.

broadside (in which case we should probably wait a little with the identification). Since the first scatterer was compensated, the far peaks tracks indicate

the changing range difference between bow and stern. Hence, when the range difference goes through a maximum, what we call the yaw Doppler must be zero. When the slope is highest, we have the maximum yaw Doppler. Thus we can use any convenient section of an identifiable peaks track to determine the times at which we will obtain a topview image, ignoring the roll motion for the moment. Whereas the various sections of the peaks tracks at the farther ranges all vary in phase, at the ranges from about Range Gate 4 to Range Gate 15 (between the dotted horizontal lines) we observe pieces of peaks tracks that are not synchronized with the ones at the farther ranges. These must be the peaks tracks from the scatterers of the superstructure. As the tracks show that the roll motion is much more rapid than the yaw motion, the timing of the extrema and inflection points of the peaks tracks associated with the superstructure are not significantly affected by the yaw motion. Thus, we can use whatever peaks tracks sections might be convenient to determine when the roll Doppler is strong and when it goes through zero.

As an example, there is a good section of a peaks track for the yaw Doppler, from 241 to 246 seconds in the interval from Range Gate 23 to Range Gate 29 (indicated by the upper dashed rectangle). Over about the same interval we can identify a roll track around Range Gate 5 (indicated by the lower dashed rectangle). When we image at an extremum of the roll track, but within the interval when the yaw Doppler is strong, we should obtain an undistorted topview. If necessary, we can adjust the imaging time via interpolation, as discussed earlier. This approach works over the entire observation interval, but we do not want to show yet another set of images. Of course, although finding usable pieces of peaks tracks by eye in Figure 5.95 is rather simple, automating the process poses a bigger challenge than most of the automation tasks discussed throughout the book. However, another practical point is more interesting.

The conclusion from the above discussion is that undistorted images can be obtained for small ships in a rough sea even when the aspect is near or perhaps even at broadside, but the interesting question is whether even an undistorted image is useful under these conditions. For example, a heavily moving small ship can churn the sea to such an extent that the ocean clutter becomes a serious problem. As we discussed earlier, in principle we can discriminate returns from the ocean either on the basis of the distribution of the responses in the image or the properties of the responses themselves, but if the broadside aspect combines with heavy ocean clutter, perhaps this is not practical because it requires a good motion compensation of the ship in order to recognize the smeared ocean clutter responses. On the other hand, the

motion of the ship is so heavy in a rough sea that imaging need extend only over a small part of the motion cycle. A test is in order.

The tracks of Figure 5.95 show several sections of tracks with essentially linearly changing range, which implies constant yaw Doppler, as was explained above. For example, there is such a track roughly between 247 and 248.5 seconds and centered in about Range Gate 20 (the right dashed rectangle). Within this time interval, we obtain a topview image when the ship is not rolling significantly. As the peaks tracks show that the roll period is two to three seconds, we should be able to find at least one imaging time with insignificant roll. An actual test shows that undistorted topview images (images in which the superstructure responses are within the outline of the deck responses) can be generated within the interval from 246.9 to 248.8 seconds. Over this time interval the roll is effective only in degrading the image quality because it introduces a changing rotation axis. Given the constant yaw rate, we can find the best imaging time by forming images over equally long intervals, counting the number of peaks within some selected area of the ship's image, and choosing the image with the fewest number of peaks. This process leads to the image of Figure 5.96. We can readily recognize the stern in the left upper corner of the ship's image. By comparison, Figure 5.97 shows an image 0.9 seconds later. The latter image of the ship




Figure 5.96 Image from 247.2 to 247.5 seconds.




Figure 5.97 Image from 248.1 to 248.4 seconds.

has many more peaks. (In both images, the lower cutoff for responses is set 8 dB above the median clutter level.)

The imaging interval was chosen so as to provide an adequate number of crossrange cells on the ship as needed for identification, in this instance in the order of 50. An examination of the individual responses of Figure 5.96 shows that they are nearly sufficiently well focused, but not quite. Transforms of some fixed-range image cuts through major responses have amplitudes that are sufficiently constant to represent the return from a single scatterer. Compensating the entire image with the common curved phase function of these transforms gives responses on which the TSA can be used.

On the basis of the tracks of Figure 5.95, we selected an imaging time at which the yaw Doppler was strong and constant. We can use the same tracks to find times at which the range does not change, so that the yaw Doppler is zero. At these times we then can form a sea-level-view image, provided there is sufficient roll motion. However, this cannot be in doubt in a high sea state. Such an image is shown in Figure 5.98. The precise imaging interval was found by shifting the time of a short fixed-duration image in small steps until the declivity was vertically oriented in the image (zero yaw Doppler). To determine the duration of the imaging interval, we can find the center of




Figure 5.98 Image from 244.75 to 245.05 seconds.

the imaging interval as described, and then increase the duration until the height levels of the superstructure are sufficiently well defined.

If we measure the distance in Figure 5.96 between the stern responses and the response marked by the arrow, and convert by the range cell size of 34.1 cm, we obtain a length of 69 ft. At this large aspect angle the bow itself should not contain the wave trapping feature that can produce such a relatively strong response, so that the ship must be somewhat longer (it is actually 83 ft). Since the image responses are good enough to use the TSA for the measurement of scatterer positions, we should be able to identify the ship if it is in the database. The approximate length measurement and the knowledge that the actual length must be somewhat larger, together with the sea-level-view image of Figure 5.98, which gives the shape and location of the superstructure, should at least allow us to classify the ship.

Summarizing the last example, which combines a high sea state with an aspect angle close to broadside, we conclude that it still is possible to systematically select the appropriate imaging interval, rather than forming an entire

series of images and attempting to select one that meets the requirements for classification/identification. Without an extensive effort, one cannot estimate exactly how close to broadside the aspect angle can be for a given sea state (and range resolution) to allow ship classification/ identification. Identification at broadside is a problem common to all moving targets.

5.2.7 Section Summary

We have shown a straightforward and practical way of measuring a ship's motion: to tracking a pair of scatterers for the measurement of the yaw Doppler, and another pair of scatterers for the measurement of the roll Doppler. For a pair of scatterers to be acceptable for the motion measurement, their differential Doppler must be significantly larger than both residual Dopplers, and the scatterers must be located in the appropriate positions on the ship. Section 5.2.2 gives procedures for estimating residual Dopplers and determining acceptability of scatterer positions.

We must choose imaging times when one of the two motions passes through zero Doppler and the other is strong and smooth, so that the motion at the imaging time is about a fixed rotation axis. Aside from the fact that such an image is easily interpreted, a changing rotation axis generates spurious responses. When tracking and compensating a scatterer for the measurement of the ship's motion, we must check the range track and the Doppler track for discontinuities and, if the tracks are acceptable, also check the phase function of the transform of the tracked response. The times of such discontinuities must be excluded from imaging, because the image quality would be poor (poor compression of responses, plus strong sidelobes and spurious responses). The Doppler rates at each good imaging time, scaled from the tracked scatterers to the full target size, allow us to determine the longest acceptable imaging interval duration (without additional motion compensation).

To select the scatterers for the motion measurements, we form a roll image of the ship, so that two scatterers at nearly the same height and at a large separation along the length of the ship can be selected for the yaw measurement, and two scatterers at different heights and at a small separation along the length of the ship can be selected for the roll measurement. If the objective of tracking scatterers in desirable locations on the ship is not quite achievable, the measured yaw and roll Dopplers will contain errors. One of the scatterers used for the roll Doppler measurement can be replaced by a virtual scatterer derived from the two scatterers used for the yaw Doppler measurement. This alternative method gives improved accuracy because the deck

position of the virtual scatterer can be chosen close to that of the scatterer of the superstructure.

A sea-level-view image formed at a time when the yaw Doppler is indicated to be zero will be slightly sheared if the time has been measured with a small error, but this will normally be acceptable. It can be readily corrected when the baseline for the roll image is discernible. If such a slight shearing of the image is not acceptable (mainly for reasons of the associated image quality), we can form a sea-level-view image slightly before the measured time of zero yaw Doppler and another sea-level-view image slightly after that time, and by comparing the amounts of shearing of the two images determine the exact time of zero yaw Doppler.

When a topview image is formed at a time when the roll Doppler was inaccurately measured to be zero, the residual roll Doppler will shift the responses of the superstructure relative to the outline of the ship. Here again we can form an image slightly before the measured time of zero roll Doppler and another image slightly after that time, and by comparing the translations of the major responses of the superstructure find the exact time at which the superstructure will be centered within the deck outline (with adjustments for nonzero depression angle). This may not represent the actual position of the superstructure correctly, but will usually be good enough for identification or classification purposes.

In the most benign motion conditions, we can easily recognize and track bow, stern, and superstructure scatterers. However, both recognition and tracking can be difficult. The flowchart of Figure 5.32 shows the appropriate sequence of processing steps when the motion is not ideal. Section 5.2.2 demonstrates measurements under benign conditions, and Sections 5.2.3 through 5.2.6 demonstrate alternatives when motion makes scatterer recognition or tracking difficult.

In those cases when the roll Doppler happens to be very small, we will not be able to generate a sea-level-view image in which scatterers can be selected for the measurement of the roll motion, though the selection and tracking of the scatterers for the yaw motion are easy. In these cases, we form topview images over slightly shifted intervals, all when the yaw motion is strong and smooth, and determine whether the usually strong responses of the superstructure shift in crossrange relative to the ship outline (or illuminated side of the ship) with changing imaging time. This shift would be caused by roll. If the crossrange positions do not shift, the roll is declared insignificant. If such shifts are detected, we attempt to track a scatterer whose position shifts significantly. If we can track such a scatterer acceptably well, we use it and a virtual scatterer in the same range gate to derive the roll

Doppler. If we cannot track a superstructure scatterer whose position shifts with changing imaging time, and the resulting uncertainty in the crossrange positions of the scatterers of the superstructure appears objectionable, we can form a series of images that reveal a full roll cycle, and choose the imaging time halfway between the peaks of the roll Doppler.

The tracking of peaks in the intensity range profile becomes problematic when a ship is moving heavily and in an erratic manner. Even under more benign circumstances such tracking becomes difficult as a ship approaches broadside, and will often be impossible near broadside (unless the radar uses very high range resolution). The tracking performance can be improved by analyzing the responses of the complex range profiles, determining scatterer positions, and then tracking the scatterer positions. Tracking performance can be further improved by replacing the procedure of first tracking a scatterer in range and then tracking it in Doppler by combined range/Doppler tracking. Despite these improvements, severe motion may require dividing the observation dwell into separate intervals for tracking and analysis.

5.3 Analysis of Ship Images

Section 5.1 considered the principles of ship imaging, and Section 5.2 the generation of usable images. Section 5.3 addresses the task of extracting from these images the information needed to identify a ship. The material is meant primarily for the reader actively interested in ship identification. A reader more casually interested in the topic might want to read the section without attempting to follow all the detailed derivations.

With regard to effects on imaging and the analysis of images, a ship can have a large number of different motion states. It may yaw with insignificant roll, roll with insignificant yaw, or it can have considerable yaw as well as roll. In the latter case, the yaw and roll motions can have cycles of similar length, or one motion can have a cycle duration much larger than the other. When the yaw and roll cycles have about the same durations, the two motions can be in phase, out of phase, or at any phase relation. This high degree of variability of the motion state of a ship makes it impractical to organize a treatment on feature extraction in accordance with a particular motion state, such as feature extraction when the ship is only rolling or only yawing, or when both motions are significant. Rather, we will consider the measurement of important features under a variety of motion conditions.

5.3.1 Measurement of Ship Length and Width

The primary problem of determining length and width, but particularly length, is to *distinguish between real responses from the ship and spurious responses* generated by the ocean surface or by scatterers with peculiarly behaving phase centers. We assume that the ship is being tracked in angle. With conventional tracking methods, the pseudoperiodic component of the yaw motion is averaged out, so that a smooth angular track is obtained. Such a track is often not very accurate, because tracking is performed on the ship as a whole, in which case the consequence of yaw, pitch, and roll on the centroid of the ship return is smoothed rather than being taken out. For more accurate tracking, the processor should create a sequence of short-term images in real time, and from these images select a specific scatterer for tracking. Such a high-quality track would allow the derivation of a relatively accurate crossrange scale factor, which would be very useful. However, the typical radar tracker in use at the time of this writing was not designed for this type of precision tracking.

As our first example, we choose the small dive boat treated earlier. In practice, the first step is to measure the motion of the ship over some time interval, and analyze the result. If there are times at which a sufficiently high yaw Doppler coincides with zero roll Doppler, the best of these times is selected for generating a topview image. This "best" imaging time is the time when the yaw Doppler is as strong as possible, the yaw period is as long as possible, and the roll Doppler goes through the broadest available zero. Also, the strongest scatterer used for the motion analysis is checked for irregularities in the ship's motion. Strong effects of this kind can be recognized by kinks in the peaks track, or the disappearance of the peaks track. When the peaks track appears smooth, the following Doppler track must be checked for discontinuities. If the Doppler track appears smooth, the phase track must be examined. This is the most sensitive check on motion irregularities, because no averaging is involved. To repeat, this check means taking an image cut in the range gate of the scatterer, and then transforming the image cut. The phase function is examined for discontinuities at the tentative imaging time. The entire procedure of selecting the best imaging time may look very complicated, but poses no problems with automated processing. The procedure is necessary for obtaining an image of sufficient quality.

For our example, we will deliberately not choose the best imaging time, because the best imaging time in a given practical application may not give results as good as we could produce by selecting the best imaging times for the available data. In other words, by not selecting the best imaging time we

can illustrate the problems that may occur when, in a particular application, the best available imaging time is chosen. Specifically, it may not be obvious in an image which of the weaker responses are from the ship and which are from the ocean surface. The image selected for this illustration is shown in Figure 5.99. The first question is, which of the responses at close ranges is the first observable response from the ship? The accuracy of the length measurement depends on answering the question correctly, and it also might be of consequence for determining the outline of the deck. A correct outline not only provides length and width, but also allows placing the scatterers in their correct positions on the deck.

To discriminate between genuine and spurious responses, we can take image cuts in the range gates of the various candidates for the bow response, and examine their transforms in order to determine whether the responses come from real scatterers or are spurious. However, the shorter the imaging interval, the more difficult it is to distinguish between the two types of response reliably. The longer a particular response can be observed, the more information can be extracted. Also, a longer coherent integration time enhances genuine responses relative to ocean clutter because an ocean surface




Figure 5.99 Topview of the dive boat.

rough enough to generate returns that can compete with ship returns in strength is not stable. For these reasons we proceed as follows. We select an image response that undoubtedly is a genuine ship response, because of strength and location, and is as close to the bow as possible. If the selection cannot be simply based on the strength of the response, we take transforms in the range gate or crossrange gate of the response, or even along diagonals, and determine whether or not there is a good approximation of the patterns of a single scatterer, or of two interfering scatterers.

In the case of Figure 5.99 the obvious response to select is the one in Range Gate -13 , indicated by the arrow. Because of its strength, this response can be readily identified in the peaks tracks. As demonstrated repeatedly, it is tracked in range and then in Doppler, and the data are compensated with the combined spline to make the associated scatterer stationary. A new image is formed with this compensation, and the stronger of the image responses closer to the bow than the compensated response (at closer ranges) are examined to find at least one with the characteristics of a genuine response, again judged from the amplitude and phase functions of the transform. If such a response is found, we compensate the new response, and form a new image. In this image we then examine the responses still closer to the bow. If such a response is not found, the examination of the responses closer to the bow is performed in the image based on the compensation of the first response. As we shift the compensation to weaker scatterers near the bow, we image over longer times in order to obtain an integration gain of ship responses relative to the variable ocean responses. With this integration gain, ocean clutter responses may become too weak in relation to the ship responses, or they may disappear altogether below the clip level (set to eliminate most ocean responses).

For an illustration we select the strong response in Range Gate -13 of Figure 5.99, compensate, and form a new image of two seconds duration. Examination of the responses at closer ranges reveals one with the characteristics of a genuine response (although smeared because it is not properly compensated). The image cut through this response and its transform are shown in Figure 5.100, on the left side the amplitude and phase function of the response, and on the right the amplitude and phase functions of the transform. The amplitude function on the right top has only a droop of the average amplitude, indicating a slow drift in the range gate. The phase function is smoothly curved (except at the right fringe), other than the variations associated with the amplitude fluctuations. Thus we can compensate this response, and generate a new image in which the compensation has been shifted closer to the bow, if not to the bow.




Figure 5.100 Image cut in the range gate of a genuine response.

The corresponding image is shown in Figure 5.101, with the newly compensated response marked. The image cut in the range gate of this response and its transform are shown in Figure 5.102. On the left top we find a sharply focused response rising high above the background. The amplitude and phase functions of the transform confirm a genuine response. The phase variations are sufficiently small over most of the interval, and the large drop at the end of the interval is due to the low amplitude minima. The drop could be avoided by taking the transform over only the response rather than over the entire crossrange interval displayed on the left of Figure 5.102. We can now examine other responses, to determine whether or not they are genuine. In particular, we want to find any genuine response that may be at a closer range than the one compensated, because that response would determine the measured length of the ship. No such response can be found in this instance. As an example, in Figure 5.103, we show the image cut and its transform in the range gate of the response just below the compensated response in Figure 5.101. The image cut does not show a single dominant response, or perhaps two responses, as would be expected from a ship. In addition, no TSA analysis is performable on the various “responses” of the image cut.




Figure 5.101 Two-second image with the bow response compensated.




Figure 5.102 Image cut and transform of the newly compensated response.




Figure 5.103 Image cut in Range Gate -9.5 .

We have demonstrated how to sort the responses at close ranges in Figure 5.99 into genuine and spurious responses, so that the first (at closest range) response from the ship may be found. The same process must be repeated for the stern if an accurate length measurement is to be obtained. An image with a scatterer compensation analogous to Figure 5.101 is shown in Figure 5.104 for the stern. The image cut in the range gate of the compensated response and its transform are shown in Figure 5.105. The response is sharply focused. The amplitude has only a short drop from poor tracking, but for most of the interval it is rather steady and accompanied by an almost flat phase function. This is a response from a real scatterer. In contrast, in Figure 5.106 we show the image cut and its transform for the response just above the one compensated in Figure 5.104. There is no sharp response, and over the short interval where the transform amplitude function (on the right) is strong, the phase function is strongly curved. If there should be any doubt that this is a spurious response, we can try to compensate better, but we will not succeed in these cases.

As stated above, for this example we have deliberately chosen a poor imaging time, so that the example is rather demanding. We have stressed the fact that one must choose the imaging time so that the spurious responses are




Figure 5.104 Image with a scatterer near the stern compensated.




Figure 5.105 Image cut in Range Gate 0 of Figure 5.104.




Figure 5.106 Image cut in Range Gate 2.2.

at the lowest possible level. This can be done on the basis of the measured yaw and roll Doppler curves, and an examination of the tracking performance in range, Doppler, and phase. Nevertheless, even for well chosen imaging times the problem of determining the first and last responses from the ship, for an accurate length measurement, is real.

For a first demonstration of a length measurement, we choose an image of better quality, which is easy to obtain if the rules discussed above are followed. Such an image of the dive boat is shown in Figure 5.107. The aspect angle is needed in order to determine the crossrange scale, which in turn is needed to measure the ship length at aspect angles other than close to head-on. This aspect angle must generally be supplied by the radar tracker. However, if a ship has a square stern and the stern is visible in the image, we can estimate *the crossrange scale from the image by changing the scale until the lines defined by the side of the ship and the stern are perpendicular to each other*. Of main importance is the ability to define the lines established by the side of the ship and its stern. The line defining the side is drawn in Figure 5.107, in order to point out that it is defined by the series of weak responses and does not include the strong response. The side of a ship hardly ever includes a scatterer capable of producing such a strong response.




Figure 5.107 Image of the dive boat, from 26.0 to 26.3 seconds.

It is worth noting that in our examples we chose the directions on the basis of the dots of the peaks plot image, which means we are utilizing the peaks of the intensity image. As was pointed out many times, the position of a peak of the intensity response need not be the position of a scatterer, because the peak could be generated by interference between two or more scatterers. To realize the best accuracy under general conditions, each complex peak must be analyzed to obtain the positions of the scatterers, and a line must be fitted to the measured positions of the scatterers rather than to the positions of the peaks. We will not do this in our examples, because the goal of our illustrations is not to achieve maximum accuracy.

The range separation between the last response along the line and the perpendicular projection of the first response onto the line is 26.5 range gates. With a range gate size of 0.32m, the range separation becomes 8.5m. Since the aspect angle, or the angle of the line with respect to horizontal, is 43° , the length of the ship is $8.5/\sin 43^\circ = 12.5\text{m}$. At this point the measurement represents a minimum length, because the length apparent to a radar can be smaller but not larger than the actual length. However, since the difference between the measured and actual will not be large, even in the absence of any information about the ship, we now know that its length is at

least 12.5m, but not much larger; perhaps by 1m or 2m. If we have good diagrams and photographs of the ship, we should be able to estimate the radar length quite accurately, by finding the first and last wave trapping features along the side of the ship. In the present case we have a photograph with insufficient details. We will demonstrate how to measure the radar length, but our estimate of the measurable length will represent the situation where a little is known about the ship, but no details.

The photograph of the dive boat is given in Figure 5.73. At the bow we see a device that looks like an anchor, but that will hardly produce the significant response indicated by the horizontal arrow in Range Gate -19 of Figure 5.107. A more likely scatterer is what appears to be a hole between the anchor and the name of the boat. There is some confirmation of this from the fact that the response is close to the line defining the edge of the ship rather than being nearer the centerline of the ship. Moreover, we can see that at this aspect angle the part of the bow around the hole will be nearly perpendicular to the line of sight. With respect to the scatterer defining the stern of the ship, no feature that would tend to trap the radar wave can be seen when starting from the very end of the ship until slightly before the first reeling support. Thus we will take the length between this point and the hole at the bow. With an overall length of 14.3m for the ship, from the photograph we find the radar length to be 12.1m. This result differs from the measured value by 0.4m. At least part of this error is due to the inaccuracy of choosing the crossrange scale factor in Figure 5.107. Probably the larger part comes from the inability to see the construction of the features that define the radar length of the ship. Nevertheless, under operational conditions the error should be considered small. The width of the ship also can be taken from Figure 5.107, but we have no information on the actual width.

As a second example, we measure the length of the coast guard cutter shown in Figure 5.33, this time viewing the ship from the rear. A short-term image (without compensation other than removing the constant range drift) is shown in Figure 5.108. Whereas the bow response meets the test of a genuine response, at the stern of the ship it is not obvious which responses are from the ship and which are from the wake. Here, we again form an image over a longer time, so that the ship responses are enhanced through the integration gain. With automated processing, we will perform a motion compensation on a scatterer as close to the stern as possible, and proceed in the same way as in the earlier example. Trying to simplify manual processing, the precise choice of the imaging interval is dictated by the desire not to have to perform a motion compensation, so that it depends on the duration of the yaw and roll cycles. In this instance the measurement of the yaw and roll motion




Figure 5.108 Short-term image (0.2 seconds) of the 110-ft coast guard cutter.

shows shorter cycle durations than for the previous example, so that we choose an interval of one second instead of the two seconds used with the dive boat. The corresponding image is shown in Figure 5.109. The range separation between the top response and the one marked is almost exactly the separation between the top response in Figure 5.108 and the one near Range Gate -40 and Crossrange Gate -2 , indicated by the horizontal arrow. This measurement establishes where the ship responses start.

Since we do not have tracking information for our examples, we again want to determine the crossrange scale by utilizing the symmetrical design of a ship, except that in this case the side of the ship is not well defined in the peaks plot. For a more accurate definition, we should analyze the complex responses in the vicinity of the illuminated edge with the TSA, and obtain a better defined edge. However, it is simpler to use the bow response, which is strong because from the rear the radar observes the corner formed by the deck and plates that rise above it. The length of the ship thus can be measured if we use the line between the bow response and the center of the stern *to estimate the crossrange scale* (If we do not estimate the center of the stern correctly, the error is negligible). The line marked in Figure 5.108 a range




Figure 5.109 One-second image of the 110-ft coast guard cutter.

separation of 83.2 gates or $83.2 \times 0.32 = 26.6\text{m}$ and an angle with respect to the horizontal of 54° , yielding a length of 32.9m. For this particular view, where the end scatterers of the ship are well defined in the image, we know that the actual length must be close to the measured length.

We can verify this from the photograph of Figure 5.33. From the design of the stern it is clear that the radar will see the actual end of the ship. Estimating in the photograph where the deck joins the corner of the bow leads to a ship length of 33.0m, just 0.1m different from the measured value. The width of the ship could be measured by analyzing the responses along the illuminated edge and obtaining a better defined edge. Alternatively, we could analyze the responses along the stern and derive the width of the stern. In fact, considering the accuracies involved here, it appears adequate to take an image cut along the stern and use only the intensity responses as shown in Figure 5.110. A definition of the width of the stern as indicated by the crosshairs should suffice. We find a projected width of the stern of 13 range gates or 4.2m, which gives an estimated actual width of $4.2/\cos 54^\circ = 7.1\text{m}$.

As an illustration of a case where angular tracking information is needed to determine the length of a ship, in Figure 5.111 we show a one-second




Figure 5.110 Image cut along the stern.

topview of the cruise ship. When the imaging interval is extended to check for spurious responses, the first and last responses of the image remain in the same range gates as in Figure 5.111, which means that these are indeed ship responses. The stern of this ship is rounded rather than straight, so that it cannot be used as a reference for establishing the crossrange scale. In the absence of a conspicuous structure whose responses are arranged perpendicular to the centerline, we can obtain only a crude estimate of the ship's length when no aspect angle information is available from the tracker. First, we assume that the stronger responses at the bow and stern should be near the centerline, so we choose the centerline as indicated in Figure 5.111. Then there are various series of responses that must be associated with typical structures on a ship, oriented perpendicular to the centerline. We choose the crossrange scale factor so that the lines defined by the responses are perpendicular to the centerline in the figure, such as the responses around the dotted line near Range Gate 20. The length of the ship then is derived in the same manner as illustrated earlier for the case of a rectangular stern. This measurement yields a length of 183 ft, whereas the actual length is only 151 ft. Evidently, this is a very crude estimate, yet it still leaves the ship in the correct class.




Figure 5.111 Image of the cruise ship, one second.

The preceding example is simple because the ship has only significant yaw motion, so that a discussion of the problems associated with an additional roll motion was unnecessary. For the next example of length and width measurement we choose the case of combined yaw and roll motion that led to the motion measurement of Figure 5.55. The specific time interval to be considered extends from 23 to 25 seconds, covering one yaw half-cycle (dashed curve) and somewhat more than one roll half-cycle.

We have stressed that the choice of the imaging time is critical when the rotation axis is changing. We now demonstrate its effect on the length measurement. The analysis of Figure 5.55 has determined intervals during which only one motion component is effective. We must examine these intervals in more detail, checking transforms of fixed-range image cuts for small motion irregularities (fractions of a wavelength) that can spoil image quality, but that would insignificantly affect the derived yaw and roll Doppler curves.

The survey image over two seconds is shown in Figure 5.112. Because of the absence of a motion compensation, the responses are highly smeared in crossrange. Also, the responses in the farther range gates are below the clip level. Lowering the clip level so that these responses can be seen brings up too




Figure 5.112 Image from 23.0 to 25.0 seconds.

many of the low-level spurious responses. The natural choice for a compensation is the strong response in Range Gate -5 . We take an image cut in the range gate of this response, transform the image cut, and remove the slow phase curvature by fitting a smooth spline. The amplitude function of the transform and the phase function after removal of the trend are shown in Figure 5.113. Near the center we notice two deep amplitude minima. Since they do not repeat periodically and are not accompanied by correspondingly large phase jumps, they are not caused by interference from another scatterer. It is the time at which both yaw and roll Doppler go through their maxima; that is, it is the worst case of a changing rotation axis within the entire two-second interval.

In contrast, around the times of 23.2 and 24.4 seconds, in the intervals labeled “yaw” and “5.114,” the amplitude is particularly smooth. These are times at which the roll Doppler of Figure 5.55 goes through zero, so that the motion is only about a fixed axis. Thus, these are the times at which we want to generate an image. Since the yaw Doppler is much larger at the second time, this is the better time for generating a topview image. A 0.4-second image at this time is shown in Figure 5.114, again without further motion




Figure 5.113 Amplitude and phase functions for the dominant scatterer.

compensation. The phase jump of about 0.17 cycles near the beginning of the imaging interval is acceptable only because it is so close to the edge of the interval that it will be suppressed by the weighting used in the imaging process. The other phase fluctuations are just a few hundredths of a cycle, and so are acceptable. Examination of the three other scatterers used for Figure 5.55 also shows acceptable transform amplitude and phase over this time interval. Thus, it yields a good-quality image.

In this relatively favorable case, when we adjust the crossrange scale so that the centerline is perpendicular to the stern, as marked in Figure 5.114, by dividing the range separation of bow and stern by the sine of the angle of the centerline, we obtain a length of 111.3 ft, compared with an actual length of 111.0 ft. For comparison, we also generate the image at the time of maximum yaw and roll Dopplers, given by Figure 5.115. The responses for the last third of the range extent of the ship are at such low levels that they cannot be shown in the figure without also displaying a large number of spurious responses. A length measurement in the second image gives 105.2m, which is much more in error than the first measurement.

As a final example of a length measurement, we choose the dive boat at a time when the motion conditions are rather poor for identification. In fact,




Figure 5.114 Image from 24.3 to 24.7 seconds.




Figure 5.115 Image from 23.8 to 24.2 seconds.

when we attempt to measure the motion by using two scatterers for the yaw measurement and another two for the roll measurement, the simple procedure of starting with a track of the peaks of the intensity range profiles does not lead to usable measurements of the roll Doppler. This is to say that when we apply the usual tests of the track quality, we cannot find acceptable tracks for two superstructure scatterers. Retaining the simple tracking approach, it is possible to use the three-scatterer method, where the motion of one superstructure scatterer is measured relative to a virtual reference scatterer. The result from this method is shown in Figure 5.116. The uncertainty in the yaw measurement is about 6 Hz, and that in the roll about 8 Hz; the figure shows an uncertainty of 7 Hz. As always, the two Doppler functions were derived from selected scatterers, and must be scaled to the full length and height of the ship.

The correlation between the two curves indicates either that the roll Doppler was calculated incorrectly, so contains a yaw contribution, or that the ship's motion is dominated by pitch (with a nonzero pitch angle). Over the 20 seconds of the figure, examination of the range peaks tracks shows that the range separation of the superstructure scatterer and the stern varies from 39% to 43% of the range separation of the bow and the stern. This small variation indicates that the superstructure scatterer is close to the line




Figure 5.116 Yaw and roll (dashed curve) Dopplers for the dive boat.

(in physical space) between the bow and stern scatterers. Figure 5.116 was derived using a fractional separation of 41%. Removal of the common motion between the curves would require a fractional separation of 71%. The calculation cannot be in such large error. Thus, we conclude that the ship's motion is dominated by pitch, which contributes to both the "roll" and "yaw" curves. As we saw in Section 5.2.2, the curves can still be used to choose times for generating topview and sea-level-view images.

Figure 5.116 reveals a difficult situation. First, the yaw motion is so rapid that the Doppler changes at maximum rates on the order of 100 Hz per second. If we use an imaging interval of 0.2 seconds, for example, at the maximum rate of change the Doppler changes by about 20 Hz over the imaging interval, which is four resolution cells. Even with such a short imaging interval, to avoid a smearing of the responses we must image near the maxima or minima of the yaw Doppler functions. The situation is somewhat worse, because the scatterers that were used for deriving the yaw Doppler span only about two-thirds of the range extent of the ship, meaning the maximum absolute yaw Doppler will be higher for bow and stern scatterers. Furthermore, even if it were practical to compensate the Doppler variations, we would still have the problem of spurious responses.

Additional difficulty arises because the roll Doppler is also significant. The roll Doppler of Figure 5.116, even without any possible scaling to the full height of the ship, changes rapidly enough to be of concern. If this were not the case, the significance of the roll Doppler would be checked by forming a short-term image at a time when the yaw Doppler is zero but the roll Doppler is reasonably close to its maximum value, or by generating an image when both yaw and roll Doppler are significant, as was demonstrated earlier. The curves of Figure 5.116 thus indicate that the choice of imaging time and imaging interval are very critical if a high quality image is to be obtained.

In order to obtain a topview image we must choose an imaging time when the yaw Doppler is strong and the roll Doppler is zero. However, with such a rapid motion of the ship, it is particularly important, from the point of view of spurious responses, to select the correct imaging time. In order to do this, we compensate the data with one of the scatterer tracks, form the image, and examine the transform of the image cut in the range gate of the tracked scatterer. This transform is shown in Figure 5.117, for the bow scatterer. As the modulation of the amplitude function indicates, it is not a good track. This fact by itself confirms that the ship's motion makes obtaining a high-quality image difficult. Nevertheless, the track is good enough to make a choice.




Figure 5.117 Transform of the image cut in the range gate of the compensated bow scatterer.

We want to select the imaging time such that a high yaw Doppler and a small roll Doppler combine with a phase function with insignificant variations (no more than about 0.1 cycles). An evaluation of the possibilities in Figures 5.116 and 5.117 shows that such a time starts at about 171.2 seconds and ends at 171.9 seconds (indicated by the double-headed arrow), over which interval the phase varies little and has constant slope. This same interval is also acceptable for the other two tracked scatterers. In Figure 5.116 this interval starts near the peak Doppler of a major yaw cycle and ends near zero yaw.

The time at which the roll Doppler goes through zero is 171.5 seconds, but the absolute roll Doppler is less than the roll Doppler uncertainty for the entire interval from 171.2 to 171.9 seconds. We form images at the beginning and end of the interval, for derivation of the time for generating a pure topview image, and find that this time is at the beginning of the interval. This image, from 171.2 to 171.4 seconds, is shown in Figure 5.118. Although the strong motion of the ship generates high returns from the ocean surface, the outline of the ship is sufficiently well indicated. The stern responses and a bow response are clearly recognizable. We again perform a crude length measurement by utilizing the peaks of the intensity image,




Figure 5.118 Image from 171.2 to 171.4 seconds.

placing the centerline of the ship halfway between the outer stern responses and parallel to the sides. The crossrange scale factor then is chosen so that the line through the stern responses is perpendicular to the centerline. We find a range difference between the projection of the bow response on the centerline and the center of the stern of 25.1 range gates, and an aspect angle of 43° , which gives a radar length of 11.8m. This is just 30 cm shorter than the length measured earlier.

To demonstrate how critical the choice of the imaging time is, in Figure 5.119 we show the same kind of image generated just 0.4 seconds before that of Figure 5.118. Whereas the earlier image allows a very accurate measurement of ship length (and width, from the responses defining the illuminated edges), the image of Figure 5.119 is dominated by useless spurious responses, the sideband responses generated by scatterers with shifting phase centers when the rotation axis is changing. We have pointed out repeatedly that a peaks plot image allows only a crude assessment of image quality, and that one must analyze the complex responses of an image to obtain a reliable assessment. If this is done with the two images of Figures 5.118 and 5.119, the responses in Figure 5.119 do not generate edges of better quality.




Figure 5.119 Image taken 0.4 seconds before that of Figure 5.118.

To demonstrate how critical the choice of the imaging time is in general, we form a pure roll image at the time of 172.0 seconds, where Figure 5.116 shows zero yaw Doppler but a finite roll Doppler. Again, when the roll Doppler is used for generating an image, we want to have zero yaw Doppler not so much because the shearing of the image is to be avoided (the shearing is easy to deal with) but because imaging must be done when the motion is about a fixed axis rather than a changing one. The sea-level-view image is shown in Figure 5.120. The responses from the level of the deck appear in Crossrange Gate -0.7 (indicated by the dotted line), essentially arranged parallel to the range axis because the yaw Doppler is truly near zero. The responses from the superstructure are to the left of the deck responses. It is a clean image, which fact is verified by examining the individual image responses. In Figure 5.121 we show the same image generated over the contiguous imaging interval from 172.1 seconds, again with an imaging interval of 0.2 seconds. We find some resemblance to the image of Figure 5.120, in particular in the upper part, but the image is not simply a sheared version of the one at zero yaw. In Figure 5.122 we show the image over the 0.2-second interval preceding the one used for Figure 5.120. Again, this is not simply




Figure 5.120 Image from 171.9 to 172.1 seconds.




Figure 5.121 Image from 172.1 to 172.3 seconds.




Figure 5.122 Image from 171.7 to 171.9 seconds.

a sheared form of the undistorted image, but is essentially useless for identification.

5.3.2 Scatterer Positions on Deck and Shape of Deck

As for all types of target, we also want to utilize special features and the positional match for ships. Two of the special features, length and width of the ship, were treated in the preceding section. As another example of a special feature, in the topview image of Figure 5.58 we can recognize that the strong scatterers of the superstructure are concentrated near the center of the ship along its length. This is not the case for the dive boat images, such as Figure 5.77. It is even less so for the SAR images of the cruise ship, Figures 5.62 or 5.111. In fact, for this ship the lack of pronounced responses that may be associated with a concentrated superstructure might be considered a special feature. The cruise ship has a superstructure that extends over essentially the entire length of the ship. Undoubtedly, as more images of ships are examined, one can define more special features that exist for some ships but not for others.

In principle, we can use a positional match for those scatterers that have no special significance, basically in the same way as was demonstrated for

aircraft and ground vehicles. Now, a ship almost always has at least some degree of roll, perhaps not sufficient to allow generating a good sea-level-view image, but enough to displace responses in crossrange. We demonstrated earlier that it is possible to interpolate between imaging times to refine the determination of the time at which the roll Doppler is zero. However, there is a question of how accurately such a procedure can be implemented with automated processing. This suggests treating the positional match for topview images more like that of ground vehicles than aircraft, meaning that we should heavily weight the accuracy of the positional match in range but not in crossrange.

We stated at the beginning of the preceding paragraph that one can use the positional match in principle. This statement refers to the likely difficulty of obtaining good diagrams and photographs for the large number of ships of various designs. If the aim is to identify large military ships, for which the variety in designs is limited, a positional match as demonstrated for aircraft and ground vehicles appears practical, so that identification appears practical. On the other hand, for small ships in a littoral environment, as we are using for our demonstrations, this may not be practical. It may not be possible to obtain the kind of diagrams and photographs needed for estimating the positions of the stronger scatterers. Then it may only be possible to classify the small ships based on such special features as length and width, the position and shape of the superstructure, and similar; it may not be possible to identify them. Thus we will merely give an indication that the positional match would work if one had enough information on the ships.

As an example of the fact that the image quality is adequate for a positional match, the three strongest responses in the topview image of the coast guard cutter, Figure 5.58, have the positions of the back of the superstructure, the mast, and the back of the command cabin, with the mast giving the strongest response (see Figure 5.33). The strongest response of Figure 5.77, the topview image of the dive boat, is generated by the trihedral corner formed by the superstructure and some other structure, just behind the life saver (see Figure 5.73). The weaker response just above and right of the strongest comes from a structure on top of the cabin. Our photographs of the cruise ship are not good enough for identification of the various corners and "cavities" that cause strong responses.

5.3.3 Height of Scatterers Above Deck (Shape of Superstructure)

In this section, we demonstrate the processing procedure for measuring the height of the scatterers, and in the process we obtain the three-dimensional

positions of the scatterers, but without knowledge of the height scale. However, we do not perform these measurements to the full accuracy to which they are obtainable. One reason is that the effort is too large for manual processing; these measurements must be performed automatically. Another reason is that we do not have adequate information on the ship in order to determine measurement accuracies. We point out where simplifications are made and what actually should be done.

Consider the 110-ft coast guard cutter, for which Figure 5.42 gives the roll (solid) and yaw Doppler curves for one of the 20-second intervals. We must start with a topview image, which means that the image must be formed at a time when the roll Doppler is zero. We would like to find an instant when the roll Doppler is zero and the yaw Doppler is near its peak, because this choice will result in the lowest possible level of spurious responses. An examination of Figure 5.42 reveals that such an imaging time exists around 64 seconds. It is the only time where this requirement is met, because the dominant pitch motion results in the “yaw” and “roll” motions having about the same timing of extrema and zero crossings. We choose this optimum imaging time when it can be found within the observation interval. Otherwise we would choose the best approximation to the ideal imaging time, even though it implies higher spurious responses.

In order to maximize the accuracy of the scatterer height measurement, we should find the exact time for zero roll Doppler by the interpolation method discussed earlier. To this point in the processing sequence, we have merely removed from the data the linear component of the changing range of the ship, without further motion compensation. For highest measurement accuracy, one should remove any residual nonlinear slow motion by range and Doppler tracking a scatterer and compensating, but the residual in this case is found to be so small that for an imaging interval of 0.2 seconds it should not matter. Thus, we leave the data uncompensated. The image obtained at a time of 63.9 seconds is shown in Figure 5.123, and the image at an imaging time of 64.5 seconds is shown in Figure 5.124. We estimate the deck line as marked in both figures. We then measure the crossrange separations from the deck line of the scatterer near Range Gate -2.0 in both figures, indicated by the arrows. Both the deck line estimation and crossrange separation measurements should be done by analyzing responses and determining the actual scatterer positions, but we have simply used the peak positions. These measurements yield crossrange separations of 7.6 and 5.6 gates, respectively, and interpolation gives the time of zero roll Doppler as 64.25 seconds. The corresponding image is shown in Figure 5.125.




Figure 5.123 Image from 63.8 to 64.0 seconds.




Figure 5.124 Image from 64.4 to 64.6 seconds.




Figure 5.125 Image from 64.15 to 64.35 seconds, zero roll Doppler.

Next, we form an image at a time slightly different from that of zero roll, closer than the times used for interpolation. We want to avoid a change in the imaging time so large that the roll motion shifts the scatterers in range and perhaps changes the image sufficiently to require a detailed analysis of the responses. The change in the imaging time should be a compromise between measurability of the crossrange shifts of the responses caused by roll and undesirably large changes introduced by the roll. Without considering this compromise in detail, we shift the imaging time by 0.1 seconds. The resulting image is shown in Figure 5.126. We should analyze the responses and determine actual scatterer positions for both images, the one with zero roll Doppler of Figure 5.125 and the image with a roll Doppler, Figure 5.126. However, by selecting a response peak in Figure 5.125 and measuring the crossrange shift of the peak in the same range gate in Figure 5.126, we already have a good estimate of the height of the peak above an unknown point. Had the translational yaw of the ship been removed by the compensation of the range drift, the heights would be referenced to the roll axis. The location of Crossrange Gate 0 in the topview image of




Figure 5.126 Image from 64.25 to 64.45 seconds.

Figure 5.125 shows that a translational yaw motion remains in the data. If the translational yaw residual changes between the two images, the roll axis will shift Doppler position. In most cases, this shift is likely to be negligible.

The measurement of a scaled height above the roll axis is nearly equivalent to that of a scaled height above the deck. Roll gives Doppler shifts proportional to height above the plane of the line of sight and the roll axis. Pitch gives Doppler shifts proportional to height above the plane of the line of sight and the pitch axis. Stability considerations dictate that the pitch and roll axes be at essentially the same height below the deck. The typically small roll and pitch motions of ships then make the planes for the two types of motion nearly coincident. The most significant effect of the difference between them is that, rather than position along the ship's width contributing to the height as for roll, position along its length does for pitch. The shape of a ship implies that, for equal roll and pitch angles, the width contributes much less than the length. If necessary, we can correct for Doppler shifts due to length by measuring their effect in the shift of the baseline of the images used to interpolate the zero-pitch time.

We can reference the scaled height to the deck by either recognizing the centerline in both images and measuring a scatterer's height by the change in its separation from the centerline, or by recognizing a set of scatterers with the same height that are spread throughout the range extent of the target, and subtracting this height from the measured values. The described measurement procedure yields the three-dimensional distribution of the scatterers, except that it does not give absolute heights but only relative heights. Errors due to rotational yaw can be estimated from Figure 5.42. The change in the yaw, scaled to the full length of the ship, is about 1 Hz from image to image. This corresponds to 0.2 crossrange gates, which can be scaled appropriately with range. Although the entire measurement procedure could clearly be refined, there does not appear to be a payoff in achieving the highest measurement accuracy in this application.

5.3.4 Section Summary

The primary problem of determining length and width, but particularly length, is to distinguish between real responses from the ship and spurious responses generated by the ocean surface or by scatterers with peculiarly behaving phase centers (by requiring transforms of fixed-range image cuts to show good one- or two-scatterer patterns). We first select an image response that undoubtedly is genuine, because of strength and location, and is as close to the bow as possible. We then compensate this response and examine responses closer to the bow, imaging over longer times to obtain an integration gain. We iterate until we no longer find genuine responses closer to the bow. We repeat the same process for the stern.

We define the centerline of the ship to pass through the center of the stern responses and be parallel to the side of the ship. If the side is poorly defined, we instead require the centerline to pass through the bow response. We obtain the crossrange scale from tracking information or by taking the stern and side to be perpendicular.

The shape of the deck and the distribution of strong scatterers along the length of the ship are two special features that help discriminate different types of ship.

The scaled height of scatterers above the deck can be determined by generating an image at a time of zero roll, then shifting the imaging time slightly, so that the roll motion shifts the scatterers in crossrange, but not significantly in range. The differential crossrange shift of the responses gives their scaled heights.

5.4 Processing Steps for Ship Identification/Classification

We summarize the processing steps required to arrive at ship identification/classification, at the same time pointing out where (known) improved measurement algorithms should be integrated into the identification software. This summary section should be of use for the reader more than casually interested in ship identification.

5.4.1 Step 1: Analysis of the Ship's Motion

After data have been collected over some observation interval, the processor analyzes the ship's motion in order to select one or more imaging intervals.

To measure the yaw Doppler function, the processor chooses one scatterer as close to the bow as possible and another scatterer close to the stern, with as small a height difference as possible. The height determination is based on a crude short-term image without any motion compensation other than removal of the range drift. One scatterer is motion compensated by first range-tracking and then Doppler-tracking, and the compensation is applied to the data. The first scatterer thus becomes a stationary reference scatterer. In the compensated data, the processor compensates the second scatterer, again by first tracking in range and then in Doppler. The range compensation of the second scatterer is converted to Doppler and added to the Doppler compensation. The result is approximately the yaw Doppler of the ship. It is exactly the yaw Doppler if both scatterers have the same height.

To measure the roll Doppler function, the processor chooses two scatterers as close together on the deck as possible, but with different heights. The selection is based on a crude short-term image with significant roll effects. If the ship does not roll, a crude image based on yaw (topview image) is used. If the selected pair of scatterers does not produce a well-defined roll Doppler function, one of the scatterers of the pair is exchanged for a new one. If necessary, the process is iterated until the roll Doppler is measurable. Otherwise, the roll Doppler is measured in the same way as the yaw Doppler, by first tracking one scatterer, then compensating the data, and deriving the compensation function for the second scatterer. This results in the approximate roll Doppler function, if the scatterers are separated along the length of the ship by a small percentage of the distance between the bow and stern.

With the alternative method, one of the scatterers on the superstructure can be replaced by a virtual scatterer, derived by interpolating the

compensations of bow and stern scatterers. The measured yaw and roll Dopplers are more accurate with this procedure, because the virtual scatterer can usually be placed in a more favorable deck position than a real trackable scatterer.

With real data, there will be small errors in the zero crossings of yaw and roll Dopplers, causing the deck line of a sea-level-view image not to be oriented parallel to the range axis, and the superstructure in topview images not to be centered along the width of the ship. One can correct the errors in the yaw and roll Doppler measurement in one of two ways. One approach (not pursued here) is to use a roll image to determine the height differences of the scatterers used for the yaw measurement, and to use any image to determine the differences along the length of the ship for the scatterers used in the roll Doppler measurements, then use the measured yaw and roll Doppler functions in an iterative procedure for reducing the errors. This is complicated. It is simpler to estimate the exact times of zero yaw Doppler or zero roll Doppler by interpolation. The processor forms a sea-level-view image slightly before and slightly after the estimated time of zero yaw Doppler, and from the degree of shearing and spreading introduced by the residual yaw interpolates the exact time. Similarly, the processor forms a topview image slightly before and after the time of zero roll, and extrapolates the exact time from the roll shifts of scatterers on the superstructure.

For our demonstrations of imaging, we have used only the simplest range tracking procedure, which is to follow peaks of the intensity range profile. In some cases this does not allow tracking four scatterers, and the modified method for three scatterers was used. For an operational system, we must incorporate both of the improved tracking procedures demonstrated separately. A first improvement is to analyze each peak of the complex range profile to determine actual scatterer positions, and then to track the scatterer positions rather than the response peaks. A further improvement for difficult tracking situations is to track scatterers in combined range and Doppler, rather than first in range and then in Doppler.

5.4.2 Step 2: Selection of Imaging Time and Duration

Topview images must be formed at times when the rotation of the superstructure (by roll or pitch) is absent, which are the times at which the measured “roll” Doppler (which can include a pitch component) is zero. Sea-level-view images must be formed at times at which the rotation of the deck is absent, which are the times at which the measured “yaw” Doppler (which can include a pitch component) is zero. At these times the images are

least distorted and most easily interpreted, and they contain the least and weakest spurious responses.

To find these imaging times, the processor scans the measured yaw and roll Doppler functions to find the times of zero yaw Doppler and the times of zero roll Doppler. The processor determines how strong the roll Doppler is at the times of zero yaw Doppler. The processor measures the strength of the yaw Doppler at the times of zero roll Doppler. For imaging, it selects the best combination, meaning that the roll Doppler should be as strong as possible when the yaw Doppler is zero, and vice versa. To choose among the possible imaging times, the processor also measures the smoothness and width of the curve of the Doppler to be used for imaging, again selecting the most favorable case. The processor also determines the smoothness of the ship motion at potential imaging times by examining the phase functions of the transforms of compensated scatterers.

Topview images are of primary interest. They are used primarily to determine the shape of the deck, length and width of the ship, and the deck positions of the dominant scatterers of the superstructure. Aside from their need as a guide to processing, sea-level-view images can be directly utilized when the aspect angle of the ship is either small or close to broadside. In the first case one obtains a true radar version of a sideview of the ship, and in the other a frontview or rearview. For aspect angles between small and large, a sea-level-view image by itself gives only a general indication of the shape and location of the superstructure. For a better definition of the superstructure, the responses of topview and sea-level-view images must be associated (deghosting) by starting from a topview image and then shifting the imaging time so that roll effects become measurable. This allows one to derive the three-dimensional arrangement of the dominant responses of the superstructure, but without true height scale. The latter can be approximately estimated from the length of the ship. In all cases, the imaging interval is selected so as to give a sufficient number of crossrange cells on the ship, but not more than necessary. This is judged by forming images over progressively longer intervals, all uncompensated other than by removal of the range drift.

5.4.3 Step 3: Motion Compensation

There will generally be no need for a motion compensation beyond removal of the overall range drift, because of the relatively large aspect angle changes due to yaw and roll. The processor must test for range gate wander and, if necessary, use polar reformatting; but this is different from the ordinary motion compensation whose purpose is to ensure that all scatterers move

with constant Doppler over the imaging time. The situation with ships is that at those times when a motion compensation is needed the image quality will be poor even with a perfect motion compensation, because of spurious responses. At those times where good image quality is obtainable, no motion compensation is ordinarily needed.

5.4.4 Step 4: Identification/Classification

Identification is possible only if good photographs and diagrams of the ships to be identified are available. When this is not the case, one can only perform classification.

For identification and classification, one measures special features such as length and width of the ship, the distributions of the scatterers on the deck, and the distribution of the scatterers in height (not to scale). The difference between identification and classification is that the measurements can be less accurate for classification, and the positional match for the scatterers will be less detailed. For example, if detailed diagrams and photographs of a particular ship are not available, one cannot determine which scatterers define the radar-observable ends of the ship, and thus the length measurement must necessarily be inaccurate.

6

Analyzing Missiles, Rockets, and Satellites

This chapter treats the application of complex-image analysis methods to the analysis of targets that spin or tumble about some axis, generally analyzed by sliding-window Doppler (SWD) processing. This advanced type of signal processing is more sophisticated and computer intensive than conventional SWD processing, but it also gives significantly better performance. Whereas Section 6.1 is of general interest, the remaining sections of the chapter are of interest primarily to those engaged in the analysis of spinning and tumbling objects.

6.1 Overview

Although aircraft, ground vehicles, and ships behave rather differently from the point of view of radar signal processing, they have one important point in common: As long as these targets can be considered rigid, scatterers in close proximity to each other have very nearly the same motions at a given time. This allows us to perform a motion compensation that is valid for the entire target and leads to an image in which all responses are properly compressed. In the case of a strongly flexing or vibrating ground vehicle this is not so, but

then we choose to forgo any attempt at measuring the crossrange positions of the scatterers to useful accuracy.

There is another group of targets that behave similarly with respect to radar signal processing, yet very differently from aircraft, ground vehicles, and ships. These targets are missiles, rockets, artillery projectiles, and satellites. Even though all scatterers on a spinning or tumbling target have the same periodic motion, the phasing as seen by the radar is different. Hence, at a given time, the Doppler variations of the scatterers are different. If we motion compensate a specific scatterer, its response will be properly compressed. However, the responses from most of the other scatterers will be smeared. *No single motion compensation that will generate an image in which all scatterer responses are compressed.* Such targets require a modified form of the signal processing treated in the preceding chapters. We will demonstrate the application of complex-image analysis to the new class of targets, but will not consider target identification.

If a target spins, the Doppler of a scatterer off the spin axis varies sinusoidally. As long as a range resolution cell contains only one scatterer, we can measure the sinusoidally changing phase of the return signal and use this phase function to compensate the scatterer so that it becomes stationary. The compressed scatterer response can be utilized to extract the scatterer's position and characteristics in the same way as discussed in the preceding chapters. The phasing of the sinusoidal Doppler gives the position of the scatterer along the circumference of the target, and the scatterer's range is obtained from the range of the peak of the response. As a detail, most of the scatterers on a spinning target will be shadowed over part of the spin period, so that the available coherent processing time is limited.

Problems arise when a range cell contains more than one scatterer, which happens under many practical conditions. The Dopplers of the scatterers in the same range cell all vary sinusoidally but, depending on the locations of these scatterers along the circumference of the spinning target, the phasing of the sinusoids is different. Hence, at a given time the Dopplers of the multiple scatterers in the same range cell vary differently. This problematic situation led to the development and wide use of SWD processing. In a given range cell, one takes the Fourier transform of the signal over some window, but with a duration so short that the Doppler variation of a scatterer within the window is not so large that the corresponding response is smeared to an unacceptable degree. This processing window is shifted in steps, so that one obtains the Doppler of the scatterer as a function of time; that is, one generates the Doppler history of the scatterer.

In practice, SWD processing often does not produce the desired results. In many applications, the extent of a target within a given range cell is not so large that adequate Doppler resolution is achieved when the transform window is made short enough to avoid smearing of the responses. Thus *there typically exists a severe Doppler resolution problem*. The situation is aggravated by the fact that, as has been conventional, only the intensity output of the processor is utilized. The resulting degradation of resolution by a factor of two is particularly detrimental with SWD processing. It is fair to state that conventional SWD processing in many cases does not give the desired performance. What may appear to be rough Doppler histories of scatterers often are interference patterns. One significant improvement would be to present the measured Doppler positions of the scatterers instead of the positions of the intensity peaks, but resolution is still improved by only a factor of two.

As in other applications in which the resolution performance obtained by FFT processing is less than desired, this is a field where superresolution methods have been applied. Simple spinning targets do not have scatterers as complicated as is generally the case for aircraft, ground vehicles, and ships, but often some of the important scatterers are sufficiently different from point scatterers to make the applicability of superresolution methods questionable. The Doppler histories obtained with superresolution methods generally have sharper spikes (to a large extent because the phase is not discarded), but the Dopplers of these spikes may not approximate the Dopplers of the actual scatterers sufficiently well. This becomes worse when the superresolution technique is designed to recover more than the factor of two lost by discarding the phase information.

To show how the methods of complex-image analysis can be adapted to SWD processing, we first consider the analysis of the image of, say, an aircraft. Since this analysis is done for each of the range gates of the image, we might as well restrict ourselves to the analysis of a single range gate. In order to fully utilize the inherent resolution capability of radar, we want to analyze the amplitude and phase functions of the return within the range gate. However, (with the TSA) we can analyze such a return only if it comes from at most two scatterers. This forces us to start with the image cut, and choose windows about responses in such a fashion that a Fourier transform yields amplitude and phase functions that have the patterns of responses from either a single scatterer or two scatterers. The process is repeated until all responses within the range gate have been analyzed. The important point is that the analysis of the amplitude/phase pattern is preceded by the Doppler

filtering implied when we use a transform window on a response to generate the amplitude/phase pattern.

Now, consider the same processing for a spinning target, with a range cell containing more than one scatterer. The best a motion compensation can do is to compress the response from one of the scatterers, which smears the responses from the other scatterers in the range gate because the timing of their spin (tumbling) Dopplers is different. In order to analyze all the responses within the range gate, we must *motion compensate one of the scatterers and then subtract its compressed response from the signal*, then motion compensate the next scatterer and subtract its compressed response, and iterate the process until all scatterers have been analyzed.

This type of processing requires that we be able to place transform windows on the responses in the range gate in such a way that one or at most two scatterers dominate in the resulting transform. Then it is possible to measure the phase function of a scatterer, as needed for the motion compensation (treated extensively in the earlier chapters). Also, every time we subtract a compressed response, we remove the parts of the other smeared responses that overlap with the compressed response. The described iteration thus will work only up to some point. The number of iteration steps that can be successfully carried out will depend on the particular circumstances for a given range gate. Note, however, that the results will show whether or not a specific iteration has been successful.

In principle, this type of processing can also be carried out for flexing and vibrating ground vehicles. In practice, two factors make the approach less interesting for ground vehicles than for spinning and tumbling objects. First, the flexing/vibration motions of scatterers on ground vehicles generally are much more erratic than a smooth spinning motion, so that consecutively compressing different scatterer responses is much more problematic. Second, even if we successfully compress a response from a vibrating scatterer, the vibration motion is generally so complicated that deriving the correct cross-range position of the scatterer is impossible. This is far different from the smooth periodic motion of a spinning or tumbling object, and for this reason we have not considered the approach for moving ground vehicles.

Instead of demonstrating the adaptation of the complex-image analysis methods to a spinning target, we will demonstrate it in a more varied application involving the slow reorientation of an attitude-stabilized rocket, when SWD processing does not work satisfactorily. Such a target may be considered to have an instantaneous tumbling motion. To make the problem more difficult, the entire rocket is within a single range cell, so that the resolution of individual scatterers depends entirely on Doppler resolution. The purpose

is to show how standard SWD processing can be replaced by more effective methods when necessary.

6.1.1 Section Summary

In contrast to the targets considered in the preceding chapters, the Dopplers of the scatterers on spinning or tumbling targets vary sinusoidally, with different phasings. Thus, if one scatterer is motion compensated, the responses of the others will be grossly smeared and interfere with the measurements on the compensated scatterer.

The signal processing used with targets of this type has conventionally been SWD processing with only the intensity output utilized. The corresponding loss in resolution is typically unacceptable. Performance can be much improved by using the methods of the complex-image analysis to measure the behavior (and characteristics) of the scatterers in the same range cell. This requires that we iteratively compensate scatterers and subtract their compressed responses from the signal.

6.2 Basic Approach to Enhanced Sliding-Window Doppler Processing

This section extends the earlier phase-processing procedures to the case of spinning, tumbling, or rotating targets, whose scatterers have differently changing Dopplers at any given time. After we remove the translational motion of the entire target, we directly analyze the return signal in each range gate. In each gate, we process time intervals when a scatterer is dominant. This scatterer is motion compensated to compress its response, which is then removed so that another scatterer becomes dominant. The iteration is carried on as long as it gives verifiable meaningful results. The sequence of processing steps is given in the flowchart of Figure 6.1. This section presents the rationale for the sequence, as well as more explanation than contained in the figure. Sections 6.3, 6.4, and 6.5 illustrate the processing in detail.

Both the times and durations of the signal sections chosen for the analysis (Steps A, B, and C of Figure 6.1) are critical. The motion about the center of gravity smears the responses, and the degree of smearing generally increases with the duration of the section selected for analysis. If the section is chosen too long, adjacent smeared responses will overlap so much that none can be selected (windowed) for analysis. If the section is too short, adjacent responses will overlap because of inadequate Doppler resolution, so

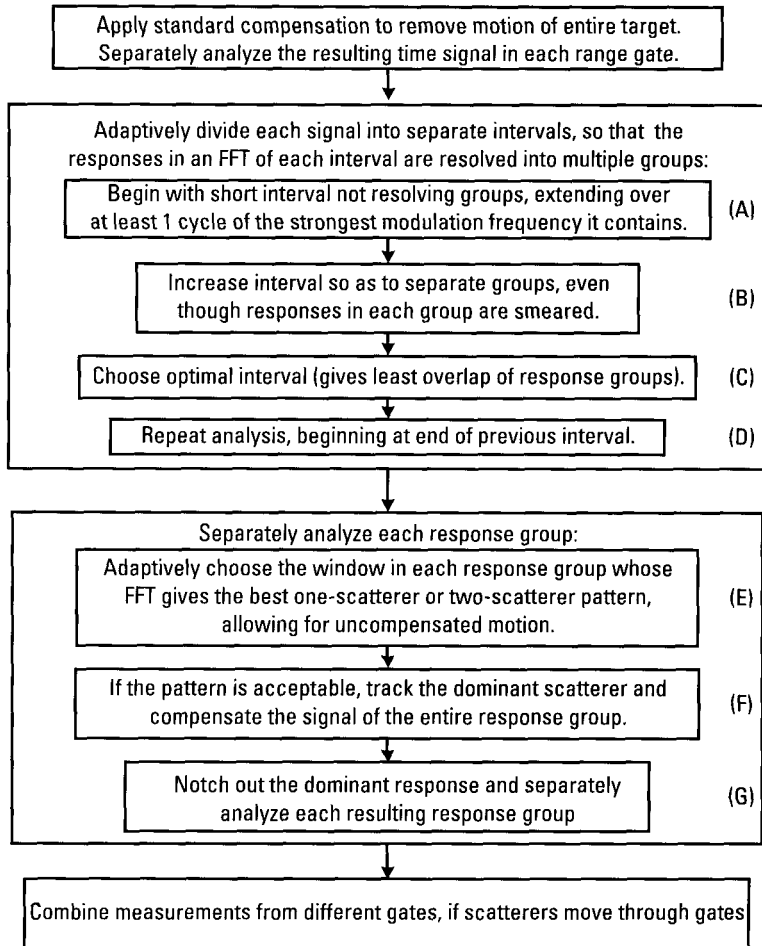


Figure 6.1 Enhanced SWD processing.

again they cannot be analyzed. Also, with spinning or tumbling targets, specific scatterers may become shadowed and reappear at certain times. This means that, in addition to the duration of the section to be analyzed, we also must select the appropriate timing.

No detailed theory of the choice of the optimum length of the signal section to be analyzed (Step C) is worth deriving. The interval must be chosen by adaptive processing: One takes transforms over successively longer intervals and investigates the degree of overlap of adjacent responses in each transform. There will be a minimum length of the signal section for which

Doppler resolution becomes good enough to eliminate a significant overlap between the responses to be analyzed. The interval can then be lengthened until the smearing of the responses (because of the absence of a motion compensation) again causes adjacent responses to overlap. We want to analyze the longest signal section for which, in the transform, a *significant overlap of adjacent responses can be avoided* by proper choice of the duration of the section. In the same manner as we select for analysis individual responses in a cut through an aircraft image, we must do this now even though many of the responses will be smeared. By systematically examining different Doppler intervals in the transform of the signal section, we select a smeared response or a group of smeared responses for analysis (Step E), and transform back into the signal domain in order to analyze the amplitude/phase pattern (Step F). The only difference compared with what is done for aircraft, ground vehicles, and ships is that the responses we transform are smeared rather than sharply compressed, and thus the signal section (comparable to the imaging interval) must be adaptively selected.

In order to proceed after selecting a smeared response, we must motion compensate a scatterer in this group (Step F), but this may pose a problem when no single scatterer is dominant over the selected signal section. As we illustrate later, the criterion for choosing the appropriate processing interval for resolving individual scatterer responses within a group is whether the responses have the *characteristics of true scatterer responses* rather than the characteristics of peaks of an interference pattern. In other words, the TSA must produce acceptable amplitude/phase patterns, allowing for phase distortions due to uncompensated motion.

It frequently happens that a single scatterer is dominant in a group of scatterers moving along roughly similar Doppler histories. By the term "dominant" we do not imply that a particular scatterer must be much stronger than all other scatterers. All we require is that the phase of the composite signal be governed by that of the dominant scatterer, at least at specific times at which the interference by the other scatterers happens to be low. In this case we can use the phase-slope measurement algorithm for deriving the phase of the dominant scatterer, so that it may be motion compensated. The compensation of the dominant scatterer compresses its response, for suppression by filtering (Step G). Transforming the resulting residual response (or part of it), we obtain the time signal minus the contribution from the dominant scatterer. Another scatterer may now be dominant, and we can perform the same measurements, compensation, and suppression. This process can be iterated as long as a dominant scatterer governs the return, at least to the degree that it can be tracked via phase-slope measurements.

If done incorrectly, this type of iterative processing could be misleading and give totally wrong results, because instead of dealing with dominant scatterer responses we might be dealing with accidentally strong peaks of an interference pattern. For this reason, we must *verify that true scatterer responses have been obtained* at each of the processing steps. In other words, instead of assuming (in Step F) that we have compensated a dominant scatterer sufficiently well, we must check whether its focused response has the characteristic of a single-scatterer response: its transform has an essentially constant amplitude function. This check is performed on each newly compensated dominant scatterer in the iteration process. Whenever the result fails the check, we must return to the previous stage where it did not fail. This is the last valid iteration.

Suppose we view a portion of a signal where a single scatterer is strongly dominant. This means that the signal amplitude function is nearly constant, showing at most a weak modulation introduced by other scatterers. The phase function will also show similarly weak modulation (the timing synchronized with that of the amplitude modulation) and it will generally be curved, because the Doppler of the scatterer is changing. The phase function thus may appear as a smooth curve with a few high-frequency wiggles superposed. In this case we can fit a spline that is flexible enough to follow the smooth curve but not so flexible that it follows the high-frequency wiggles. This is possible, of course, only if the processing interval contains some minimum number of these wiggles, perhaps four or five. Such conditions can usually be obtained with an appropriate choice of the processing interval; if they cannot, then some degradation in measurement performance is inevitable.

In those cases where a scatterer is not strongly dominant, the magnitudes of the amplitude and phase modulations introduced by secondary scatterers will be larger. At some point the phase modulation may become so large that fitting a smooth curve to the underlying variation of the (weakly) dominant scatterer will entail significant errors, which means that the response of the dominant scatterer will be incompletely compressed. It will probably also have unacceptably large sidelobes. It can readily happen that the compression of the response is so poor that we cannot continue with the iterative processing. Thus we need a different type of compensation when no single scatterer is strongly dominant.

The objectionable phase modulation from other scatterers is correlated with the amplitude modulation. Under some conditions, this allows us to remove the phase jumps and then fit a spline to the smoothly curved phase function of the weakly dominant scatterer. However, our experience has

shown that in practice the removal of the phase jumps involves errors that tend to generate a detrimental background for the responses. For this reason we have developed the alternative compensation method, variously discussed and demonstrated earlier, whereby we fit only to the slopes of the phase function when the scatterer to be compensated happens to be dominant (phase-slope tracking). We show in the following sections that this type of phase-slope compensation gives much better results than the removal of the phase jumps.

The difficult task of the phase-slope compensation method is to define the intervals when the phase slope is dominated by the same scatterer. If we misidentify some of these intervals, then phase distortions will be introduced that smear the responses and generate high crossrange sidelobes, the latter usually being the worse problem. If the distortions are too large, the responses will fail our test for true characteristics of scatterer responses. In a complicated situation we may fail to track a scatterer, but we should not track peaks of an interference pattern and assume that they represent true scatterer responses. We must always verify the accuracy of the track, as explained above.

6.2.1 Section Summary

The general procedure is to remove the range drift of the target, select a section of the return signal, track the dominant scatterer in phase, motion compensate, and remove the compressed scatterer response by Doppler filtering. The smeared responses of the other scatterers in the same range gate remain. This represents a signal in which a new scatterer is likely to be dominant, so that the processing can be iterated until no scatterer is trackable after suppression of the compressed response of the tracked scatterer.

This type of analysis can be performed only if the amplitude/phase pattern of the signal is interpretable; that is, if it comes from either one or two significant scatterers. To accomplish this, we must generally take the transform of the signal and select a suitable part of the spectrum so that its inverse transform gives an admissible amplitude/phase pattern.

The length of the signal section to be transformed must be chosen as a best compromise between Doppler resolution and the response smearing caused by the lack of a motion compensation (other than removal of the range drift). The appropriate length of the signal is found by iteration.

The criterion of success at every processing stage is whether the responses are analyzable with the TSA. In a given situation we may not

obtain all desired results, but we will not obtain false results if the checks are applied.

6.3 Analysis of an Attitude Maneuver

In this section we demonstrate in considerable detail the type of analysis discussed in Section 6.2, because it is more intricate than the analysis of returns from aircraft, ground vehicles, or ships. The section should be of interest to those who want to improve SWD processing. For this demonstration, we use data that cannot be properly analyzed with conventional SWD processing because the rotation of the target is too slow: data from a sounding rocket when it changes its attitude.

In Figure 6.2 we show the conventional Doppler history plot around the time the rocket reorients. Those intervals we analyze in more detail are labeled on the right side of the plot. For about the first 10 seconds, when only a single response is visible, we have smooth flight before the start of the maneuver. The amplitude of the response is highly clipped prior to the maneuver. The Doppler history of the response before the start of the maneuver illustrates the problems with conventional SWD processing. The processor forms a series of Doppler spectra, and then tracks the main peak. The Doppler track is converted to phase, which is subtracted from the data on a pulse-by-pulse basis. One problem is that Doppler measurements involve processing over some time window, so that Doppler tracking involves smoothing. The more serious problem comes from the choice of the flexibility and the spacing of breakpoints of the spline function fitted to the Doppler track. If the spline function is too rigid, it will only incompletely take out the Doppler wander of the dominant scatterer. If the spline function is too flexible, it will follow the disturbances in the Doppler track introduced by interfering scatterers. In the limiting case of too much flexibility, all variations of the phase function are removed, so that a noncoherent signal of low utility is generated. The capability of tracking individual scatterers is lost. The remedy is to track the phase function directly, which allows one to select a spline function that follows the phase of the dominant scatterer without being affected by the phase variations introduced by other scatterers. This is to say that one must use the phase-tracking procedures discussed in Chapter 2.

After the start of the maneuver at about 442 seconds, the poorly developed Doppler histories of Figure 6.2 provide only some general understanding of what happens. When the Doppler of a scatterer changes rapidly,




Figure 6.2 SWD plot of the sounding rocket performing an attitude maneuver.

Doppler processing smears the responses in a way that tends to generate short, vertically oriented “Doppler histories.” These are so similar to real Doppler histories of other scatterers that it often is impossible to decide whether one observes additional Doppler histories or just the consequences of Doppler processing on scatterers whose Dopplers are not constant. A Doppler history plot of the type shown in Figure 6.2 generally does not allow tracking of individual Doppler histories. The problem cannot be overcome

by judiciously choosing the length of the processing window, because of the conflict between response smearing when the window is too large, and inadequate Doppler resolution when it is too short.

In Figure 6.3 we show the amplitude and phase data from 440 to 450 seconds, where about the first two seconds represent stable flight. For a better understanding of the phase function during the stable part of the flight, in Figure 6.4 we show just the first two seconds on an expanded scale. For this strong signal there is a slight (equipment) problem with amplitude clipping, but this is of no interest here. Figure 6.4 is interpreted as follows. The fact that there are no deep amplitude minima implies that one scatterer is dominant, so that the phase function describes the motion of that scatterer. The instantaneous Doppler of the scatterer is given by the phase slope at the time of interest, but we must smoothly fit through the phase variations associated with the amplitude modulation; both are introduced by a secondary scatterer. To make the dominant scatterer stationary, we fit a polynomial to the phase function, with the flexibility chosen so that it can follow the roughly quadratic trend but not the high-frequency variations. If there is reason to present the results in the form of a Doppler history, the stationarity of the dominant scatterer after the phase compensation means that its Doppler




Figure 6.3 Data segment from 440 to 450 seconds.




Figure 6.4 Data segment from 440 to 442 seconds.

history will be aligned along the zero Doppler line, not meander as in Figure 6.2.

The simple characteristics of the modulation seen in the amplitude and phase function already indicate that, having motion compensated the dominant scatterer, we can perform measurements on the secondary scatterer. The response of the motion compensated scatterer will be sharply compressed, so that it can be suppressed by Doppler filtering. The secondary scatterer then dominates in the remaining signal, in which it has now become the dominant scatterer. Thus we can perform measurements on it. It is too simple to justify a demonstration.

In accordance with Step D, we move on to the analysis of the data section after the smooth motion. Figure 6.5 shows that the next data section, from 442.0 to 443.1 seconds, is more interesting. This interval was chosen for the purposes of illustration. As we explain below, this arbitrary choice does not affect the processing of the data. The first tenth of a second covers the end of the stable flight. For the rest of the interval we see a strong amplitude modulation, indicative of the fact that the single scatterer previously so strongly dominant (the base of the rocket) has become much less dominant. The amplitude modulation essentially disappears in the interval from 442.82




Figure 6.5 Data segment from 442.0 to 443.1 seconds.

to 443.00 seconds, which means that the secondary scatterer has become too weak to interfere. As the phase is close to linear, the transform of this interval of the amplitude/phase data would give a single strong response, and its peak position would be the true Doppler of the scatterer.

We have placed solid vertical lines in Figure 6.5 at the times when the (smoothed) amplitude is at a peak. The spacing of these crosshairs indicates that the modulation interval at first decreases and then increases. This means that the interfering scatterer has a smoothly increasing and then decreasing Doppler, which is typical for this type of application. Thus we can motion compensate either the dominant or the secondary scatterer, but not both with one compensation. The situation is similar to that of a ground vehicle with flexing or vibration, except that in the present case the motions are smoother and better defined.

To analyze the target, we must first compensate the stronger scatterer and measure its Doppler and strength. Since the response of the dominant

scatterer is sharply compressed after motion compensation, we can suppress it by Doppler filtering. Another scatterer will be dominant in the remaining signal, even though it may not be as strongly dominant. We can measure its phase function, perhaps by the phase-slope method if the scatterer is not strongly dominant, and via motion compensation compress the response of the second scatterer. Its Doppler and strength can now be measured. The compressed response can again be removed by Doppler filtering, so perhaps another scatterer becomes dominant. The process can be iterated until it stops working, which can be recognized from the fact that the compression of a response no longer works. This type of processing can be successfully done until the time of 442.82 seconds, a time at which the contributions from other scatterers become too weak for measuring more than the dominant scatterer.

In terms of Figure 6.1, Step A would begin with an interval extending over one period of the strong modulation, from 442.0 seconds to 442.12 seconds. The transform of this is a single response. Increasing the duration of the interval (Step B) improves Doppler resolution of the interfering responses, until 442.82 seconds. An FFT of this interval resolves the smeared dominant and secondary scatterers to a large degree, but not entirely. Beyond this time, extending the interval widens the response of the dominant scatterer (because it is not compensated) and leaves the secondary unaffected. The optimal interval (Step C) is the one that best resolves the responses, 442.0 to 442.82 seconds. We next (Step E) must choose the Doppler interval whose transform gives the best one-scatterer or two-scatterer interference pattern, allowing for uncompensated motion. Because the smeared dominant and secondary scatterers are not yet entirely resolved, the transform of a window about either response does not give a constant amplitude. The transform over a window about both responses, however, gives a good two-scatterer pattern, insignificantly different from that of Figure 6.5 between 442.0 and 442.82 seconds. For the sake of brevity, we track the dominant scatterer (Step F) in Figure 6.5 directly.

With the relatively strong amplitude modulation of Figure 6.5, we cannot phase track the dominant scatterer but must instead use phase-slope tracking. The crosshairs mark the times at which the (smoothed) phase slope must be measured, so that a spline can be fitted to the succession of phase slopes. A cruder approach suggested by the sharp phase jumps in Figure 6.5 is to remove the phase jumps occurring at the times of the amplitude minima, and then fit a polynomial to the phase function directly. As already indicated earlier, we have generally not been successful with this procedure because errors in the measured size of the phase jumps translate into Doppler

sidelobes, and these are particularly critical if one attempts to iterate the procedure.

As mentioned in Chapter 2, the specific phase-slope tracking procedure we have implemented contains the inaccuracy of assuming that the phase slope as measured at the times of the amplitude maxima is not affected at all by the weaker scatterer. This approximation is acceptable when the amplitude modulation is weak (relatively high amplitude minima), and when measurements are performed without iteration. In the present instance the amplitude modulation is rather strong, and we want to use an iterative procedure. Thus we should measure the changing amplitude modulation index and include a correction of the measured phase slope. The need for such a correction can be verified by examining the phase slopes at the times of the downward-pointing arrows. In the interval around 442.9 seconds, where the second scatterer is much weaker (no deep amplitude minimum), the phase slope gives the Doppler of the main scatterer. On the other hand, in the preceding interval, at about 442.8 seconds, the phase slope is rather different. If the phase slope were due to the main scatterer, the slope would have to change very quickly, which it cannot do under the present circumstances. However, note that the amplitude modulation around 442.8 seconds is very strong, so that neglecting the consequences of the secondary scatterer on the phase slope evidently is not permissible. This type of application poses increased demands on phase-slope tracking, so that the algorithm should include the correction discussed above.

The points concerning the inaccuracies in the phase-slope measurement and the Doppler sidelobes generated by errors in the measurement of the size of phase jumps apply for routine or automated analysis of targets with this type of motion behavior. However, our present purpose is the illustration of the techniques, so that we simply remove the phase jumps before fitting. The corresponding results on the basis of phase-slope tracking are given in Appendix G. We perform the processing between 442.0 seconds and the time of "disappearance" of the second scatterer, at 442.82 seconds. If the phase jumps are not removed accurately, which is usually the case, the resulting phase discontinuities limit the achievable number of iterations for analyzing progressively weaker scatterers.

The data after removal of the phase jumps and the linear trend of the resulting phase are shown in Figure 6.6 (note the change in phase scale compared with Figure 6.5). The residual slow modulation of the phase function of almost 0.3 cycles would seriously smear the responses in the transform domain. We fit a spline function to the slow modulation of the phase in Figure 6.6, and use it to correct the original phase function, which is the




Figure 6.6 Data segment after removal of phase jumps and linear trend.

phase function before removing the phase jumps (Step F). The result is shown in Figure 6.7. This is an amplitude/phase pattern as used with the TSA to determine the Dopplers of the scatterers, except that it is shown over many modulation cycles rather than only a single cycle, and the modulation period changes. This means that the observation time is much longer than needed to resolve the two main scatterers (if proper resolution were feasible when the Dopplers are not constant). One immaterial difference is the high-frequency modulation superposed on amplitude and phase, which can be readily removed by taking a transform, placing a transform window that excludes the high frequencies, and transforming back.

The important difference relative to the situation with aircraft and similar targets is that the duration of the modulation cycle changes with time. This is the consequence of the changing differential Doppler between the two scatterers. However, because the change is small over one cycle, we can measure the instantaneous Dopplers of the two scatterers, one cycle at a time. It is merely a question of pinpointing a changing Doppler precisely




Figure 6.7 Data in Figure 6.5 after phase correction.

in time. When we measure the changing Doppler from the change in the modulation cycles over the time interval displayed in Figure 6.7, we obtain the Doppler histories of the two scatterers. Any measurement errors for one cycle at a time can be smoothed by spline fitting to the sequence of measurements, if a sufficient number of measurements is available.

Having performed the Doppler history measurements on the two main scatterers, we want to suppress the dominant scatterer so that the iteration may be continued and the Doppler histories of weaker scatterers may be measured. This requires focusing the dominant scatterer without smearing the secondary scatterer so much that its smeared response overlaps with the sharp response of the dominant scatterer. The most effective way would be to motion compensate the dominant scatterer, resample the data so that the duration of the modulation cycle does not change over the interval under analysis, take the transform to obtain two sharp responses with a background from the smeared weaker responses, and suppress the response of the

dominant scatterer. However, if the modulation frequency does not change too rapidly, it is simpler to do this without resampling the data, accepting a somewhat smeared rather than sharply focused response of the secondary scatterer. The important point is that it is not a matter of believing what may be acceptable. With the simpler procedure, we can examine the transform and determine whether or not the secondary response overlaps the dominant response more than is desired. Thus we can adaptively decide how to proceed.

The time interval displayed in Figure 6.7 is of rather arbitrary length. There is no need to analyze the entire interval in one procedure. As long as we can find sections that include several modulation cycles for which the duration does not change much, we can simplify a manual analysis by examining shorter sections at a time. This may even be advisable, as we have ignored the effects of the secondary scatterer on the phase of the dominant scatterer. For this reason, we take the transform only over a section where the amplitude modulation does not change too much, as marked by the cross-hairs in Figure 6.7. This transform is shown in Figure 6.8.

The main response in Figure 6.8 is well focused, as can be seen from its sharpness and the fact that the phase is linear. Since the secondary response is not as sharply focused, when we suppress the main response we will also suppress a somewhat significant part of the tail of the secondary response. This could be avoided or reduced by resampling. We notch out the main response, as indicated in Figure 6.9 (Step G). The curvature of the phase of the remaining response indicates that this response is not focused, as expected from the fact that the various scatterers are moving with different Dopplers (we can focus only one scatterer at a time). Taking the transform (Step E) of the response in Figure 6.9 gives Figure 6.10. If the secondary response had truly become dominant after the suppression of the dominant response, the amplitude modulation in Figure 6.10 would be small, and the phase function would accurately represent the motion of the secondary response. In actual fact, some of the amplitude minima are quite deep, so that the apparent slow modulation trend of the phase function is affected by other scatterers. Nevertheless, instead of trying to decide whether the phase measurement is acceptable, with this type of amplitude function we can continue and *decide whether the continued processing was justified*, by examining whether the resulting response is sufficiently compressed. Thus, we remove the overall curvature of the phase function, as shown by the fitted polynomial (Step F). The transform after the phase compensation is given in Figure 6.11. This response is now also well focused, as seen from the linear phase function.




Figure 6.8 Transform of the corrected data.




Figure 6.9 Removal of the main response.




Figure 6.10 Transform of response in Figure 6.9.




Figure 6.11 Transform of Figure 6.10 after phase compensation.

If there were a tertiary response of significant magnitude, we could continue with this process. There is a weak response indicated by the peak around 0 Hz in Figure 6.11, and another weak response is indicated in the right tail of the main response. However, if the Doppler histories of such weak responses are to be measured, it becomes necessary to implement the procedure as accurately as possible, primarily by using phase-slope measurements rather than simply removing the phase jumps, with the phase-slope measurement taking into account the change in the phase slope of the dominant scatterer by a secondary scatterer. Otherwise, the number of iterations is limited, as indicated by the fact that beyond a certain number of iterations one does not obtain a compressed response for further suppression.

In our earlier discussions of how long a data interval should be analyzed at a time, we have given the criterion that the optimal interval is that whose transform gives the least overlap of response groups. We have implicitly assumed that an accurate motion compensation is possible over this interval. This assumption is reasonable, because complicated uncompensated motion will smear responses extensively. Our stated criterion is in effect a balancing of the need for nominal Doppler resolution versus the motion compensation difficulties and response smearing due to variable Doppler differences between scatterers. Nevertheless, it is possible that accurate scatterer tracking is feasible only over part of our selected interval. If this turns out to be the case, we must reduce the interval accordingly. The acceptability of a motion compensation can be judged by the quality of the resulting responses.

When consecutive sections of the data record of Figure 6.3 are examined, the situation with respect to processing is similar to that illustrated above. However, the interval from 448 to 449 seconds is more interesting because more than two significant scatterers are involved, and the processing can be carried further. The original signal within that interval is shown in Figure 6.12. We again simply remove the phase jumps rather than using the phase-slope procedure (the phase-slope results are in Appendix G), fit a spline function, subtract the result from the original signal, and then take the transform (Step E). These processing steps yield Figure 6.13. The compensated response lies between the crosshairs, and it is a double response, as expected from the modulation pattern of Figure 6.12. The piecewise linear phase functions verify that it is a true double response rather than the product of smearing due to an inadequate compensation. The transform over the window in Figure 6.13 gives the signal in Figure 6.14. This is an excellent approximation of the interference pattern of two scatterers with a constant differential Doppler. The amplitude/phase pattern of Figure 6.12 thus is not a true two-scatterer pattern, but a three-scatterer pattern from the three main




Figure 6.12 Data segment from 448 to 449 seconds.




Figure 6.13 Responses after motion compensation of strongest scatterer.




Figure 6.14 Transform over the window in Figure 6.13.

responses in Figure 6.13. The point is that the smoothly varying differential Dopplers create patterns different from those of scatterers with constant differential Dopplers.

In the next iteration, let us suppress the signal within the crosshairs in Figure 6.13 (Step G) and take the transform (Step E). The result is shown in Figure 6.15, where we again see an amplitude function without deep amplitude minima except for one near the end of the displayed interval, implying a new dominant scatterer. The curved phase thus may be compensated in the same manner as in the other examples (Step F). The corresponding transform after the compensation of the dominant response is shown in Figure 6.16. By comparison with Figure 6.13, the response is clearly sharper and the phase function is linear. In this fashion we can iterate the processing until all significant responses are resolved. When this is done over the entire record, we obtain the Doppler histories of the scatterers.

As another example, in Figure 6.17 we show a section during which the modulation period changes rapidly, because the differential Doppler between scatterers changes. To demonstrate the consequences, in Figure 6.18 we show the responses obtained by transforming the signal in Figure 6.17




Figure 6.15 Transform of Figure 6.13 after suppressing indicated window.




Figure 6.16 Transform of signal in Figure 6.15 after compensation.




Figure 6.17 Data segment from 452 to 453 seconds.




Figure 6.18 Transform of signal in Figure 6.17.

without any compensation. We see two responses that are smeared because the signals have curved phase functions. If we suppress the stronger signal on the left and take a transform (Step E), we obtain Figure 6.19. The nearly constant amplitude indicates that only one significant scatterer contributes to the response, so that the slope of the phase function represents the Doppler of the scatterer (Step F). The difference between the phase slopes of the two displayed intervals corresponds to a Doppler change of about 3 Hz, causing a corresponding smearing of the response. When the Dopplers of scatterers are changing, the *correct selection of the length of the processing window* is important. As stated earlier, the processing interval must be chosen as a compromise between the need for nominal Doppler resolution versus the difficulties of compensating over a longer time period and the problems of response smearing caused by a variable relative scatterer Doppler. Although the entire one-second signal of Figure 6.17 is usable in a single interval, this need not be the case.

In Figure 6.20 we show the signal from 455 to 456 seconds. The amplitude/phase pattern varies much over this interval. If we attempt to process the entire one second, rather than adaptively selecting intervals, we find that an FFT does not allow us to choose a *Doppler window* that filters




Figure 6.19 Transform of Figure 6.18 after strongest response is suppressed.




Figure 6.20 Data segment from 455 to 456 seconds.

out any of the contributing scatterers. This remains the case even after we compress the dominant response by applying phase-slope tracking in Figure 6.20. The Doppler separation of the dominant and secondary scatterer varies too much over the one-second interval of Figure 6.20. Thus, an analysis of the entire one-second interval would produce a Doppler history for just one scatterer. In contrast, by adaptively choosing the processing interval, we can easily measure Doppler histories for two scatterers.

We note that conventional SWD processing (with a fixed window), using a shorter interval than one second, can produce two accurate Doppler histories over only part of the signal. For favorable positions of the window relative to the amplitude modulation cycle, the window must be at least one full cycle long (from one minimum to the next) in order to resolve interfering scatterers. Other positions require a longer window. A window wide enough to resolve the scatterers at the beginning of the signal would yield inaccurate positions near the center of the signal, when the modulation period changes rapidly.

In Figure 6.21 we show an example where the differential Doppler between two scatterers decreases to zero and then increases again (we can tell




Figure 6.21 Data segment from 459.3 to 462.5 seconds.

from the change in sign of the phase slopes corresponding to the amplitude maxima that the value of the differential Doppler goes through zero, rather than merely going through a minimum). The crosshairs at the amplitude maxima show how the modulation period increases and then decreases. At the center of the displayed time interval we appear to have only a single scatterer. This is because the differential Doppler is so small that we cannot resolve the two scatterers. Only the continuity in processing over the entire time interval shows what is happening. Whereas the amplitude/phase pattern of Figure 6.21 is easily interpretable, the conventional output from SWD processing is worst at such times. This is because the rate of change of the Doppler is largest at these times, as seen from the strong curvature of the phase function in Figure 6.21.

In Figure 6.22 we show the signal from 465.4 to 466.6 seconds. As in Figure 6.21, Figure 6.22 shows the pattern of two interfering scatterers whose differential Doppler goes through zero. However, the displayed interval is 1.2 seconds in Figure 6.22, but 3.2 seconds in Figure 6.21, which means that the process in Figure 6.22 is almost three times as fast as in Figure 6.21. If conventional SWD processing worked at all under such




Figure 6.22 Data segment from 465.4 to 466.6 seconds.

conditions, the width of the processing window should not be the same. As a minimum, SWD processing would have to be modified to work with adaptive window widths. With amplitude/phase analysis, on the other hand, rate of change of the length of the interference cycles does not matter because the processing interval is already chosen adaptively.

6.3.1 Section Summary

We demonstrated that the complex-image analysis procedures allow one to track two interfering scatterers with different motions, under conditions when conventional SWD processing does not.

When the Dopplers of scatterers are changing, the correct selection of the length of the processing window is important. The processing interval must be chosen as a compromise between the need for nominal Doppler resolution versus the difficulties of compensating over a longer time period and the problems of response smearing caused by a variable relative scatterer Doppler. The validity of the chosen interval can be established by examining the resulting response compression.

6.4 Analysis of the Smooth Flight

In this section we demonstrate that complex-image analysis procedures also improve resolution performance when a motion is so smooth that the differential Dopplers between scatterers are very small. One can achieve far better Doppler resolution performance than with conventional SWD processing. The section is intended for those readers who are active in this field of signal processing.

In the preceding section we showed how target motions that cause the Dopplers of the individual scatterers to vary in complicated fashion also require a generally rather complicated analysis. In this section we show that the type of analysis can also be critical if the target motion is much more benign.

We examine a data segment before the reorientation of the rocket begins. The objective is to align the strong main response so well along the zero-Doppler axis that responses from much weaker scatterers are found, if possible. The detection of weak responses is a common problem for such a target if it is viewed from the rear, where the base is a very strong scatterer. We have arbitrarily chosen a signal segment of 10 seconds duration, which should enable a Doppler resolution of 0.1 Hz with complex-image analysis processing. There is no implication that even better Doppler resolution cannot be obtained when the flight is smooth. The signal is shown in Figure 6.23. The curvature of the phase function extends over about 10 cycles, which is huge by signal-processing standards. A variation of at most about 0.1 cycles is admissible if a response is to be properly focused. The amplitude modulation visible in Figure 6.23 implies the presence of a dominant and other scatterers, more than one scatterer because the modulation is not quasi-periodic despite the fact that the rocket is attitude stabilized. How well can one extract the individual scatterer Dopplers?

To compensate the dominant scatterer, we must fit a polynomial to the phase function of its return. Although not visible at the ordinate scale of the phase function in Figure 6.23, there must be phase variations associated with the amplitude variations. However, since there are seven amplitude modulation cycles, we can fit a smooth polynomial to the phase function of Figure 6.23, sufficiently flexible to follow the curved phase function but not so flexible as to be influenced by the hidden phase fluctuations. A fifth- or sixth-order polynomial meets the requirement. The signal after the compensation is shown in Figure 6.24. Since the rapid phase jumps occur at the times of the amplitude minima, the phase modulation is caused by mutual interference between scatterers rather than being a motion compensation




Figure 6.23 Data segment from 430 to 440 seconds.




Figure 6.24 Signal after compensation with fifth-order polynomial.

residual. The characteristics of the amplitude/phase pattern (only small deviations from that of two interfering scatterers) indicate the presence of a secondary scatterer that is dominant among the other weaker scatterers. The transform of the signal in Figure 6.24 is given in Figure 6.25. The normalized half-power width is 1.01, so that the compensation is indeed excellent. The other responses in Figure 6.25 are evidently rather weak. In order to determine their nature and possibly measure their Dopplers, we must employ the iterative processing illustrated earlier.

For a practical application, we would next separately examine transforms of the two response groups to the left and right of the suppressed response (Step G). For the purpose of illustration, we examine the transform of the entire spectrum after suppressing the main response. Taking the transform of the remainder gives the signal in Figure 6.26. We observe an amplitude pattern in which the detailed variations are largely synchronized with the phase pattern, so that there appears to be another dominant scatterer in the residual group of scatterers. Since the phase function has an overall curvature, it must be compensated in order to focus the response of the newly dominant scatterer. However, in contrast with Figure 6.23, the rapid phase variations now are very much larger in relation to the magnitude of the slow




Figure 6.25 Transform of Figure 6.24.




Figure 6.26 Transform of Figure 6.25 after suppressing main response.

curvature. This means that we cannot simply fit a smooth polynomial as done for the phase function of Figure 6.23, but must use the phase-slope method of measuring the phase function of the dominant scatterer. The phase slopes are measured during the times of scatterer dominance, as shown by the vertical lines of the figure. When a polynomial is fitted to the sequence of phase slopes, and the integrated fit is subtracted from the phase function of Figure 6.26, we obtain the signal of Figure 6.27. Because the phase-slope changes caused by the secondary scatterers have not been taken into account in the phase-slope algorithm, some phase curvature is left after the compensation; but it is much smaller than in Figure 6.26.

The transform of the signal in Figure 6.27 is shown in Figure 6.28. We indeed have an enhanced and sharpened response of a newly dominant scatterer, so that the objective was accomplished. The focusing is not perfect, both because of the weakness of the scatterer relative to the base of the rocket (recall Figure 6.25), and because of the approximations involved in our processing. Suppressing the strong response in Figure 6.28 and iterating the process results in Figure 6.29. From Figure 6.25 to Figure 6.28, and to Figure 6.29, the relative height of the remaining dominant response has progressively decreased, but as long as the procedure works, we obtain




Figure 6.27 Signal of Figure 6.26 after phase-slope compensation.




Figure 6.28 Transform of Figure 6.27.




Figure 6.29 Spectrum after suppressing second dominant scatterer, then compensating signal.

meaningful measurements of the Dopplers of individual scatterers. If the compensation had not worked, it would be obvious both from the poor fit of a polynomial with reasonable flexibility and because the resulting responses would be greatly smeared. Note that the main response in Figure 6.29, at -0.88 Hz, corresponds to the unresolved response on the left side of the main response in Figure 6.28. Since the strongest response in Figure 6.29 is not much stronger than the remaining responses, further iteration appears useless. However, with a careful formulation of the algorithms, continuing the iteration should be possible because *the remaining two groups of responses can be processed separately*.

Returning to Figure 6.2, one has the impression that in the lower part of the figure one observes only the main scatterer, whereas with the beginning of the maneuver many more scatterers appear. What we have shown is that, with the appropriate compensation over a sufficiently long time, we can observe these very weak scatterers even during the smooth missile flight. The combination of the strength of the base return before the maneuver and the small differential Dopplers of the weak scatterers means that something more

sophisticated than conventional SWD processing must be employed to make the weak scatterers visible.

As we have stressed repeatedly, *responses obtained with such sophisticated processing must be tested*, to ensure that they are not merely peaks of an interference pattern. For example, consider the responses in Figure 6.29. Measurement of the main response at -0.88 Hz gives a relative half-power width near unity, so that the position of the peak accurately gives the Doppler of the associated scatterer. The two responses within the interval from zero to 0.6 Hz are not fully resolved, so that the positions of the peaks may not accurately represent the Dopplers of the scatterers, if they are true scatterer responses at all. When we take a transform over the two responses, we obtain Figure 6.30. The approximation of the amplitude/phase pattern from two interfering scatterers is good enough to conclude that we indeed have true scatterer responses. A measurement with the TSA gives Doppler positions of 0.41 and 0.14 Hz, which are close to but not exactly the two peak positions.

Note that the many pattern cycles in Figure 6.30 imply that the time interval is much longer than needed to resolve the two responses (10 seconds for scatterers separated by 0.27 Hz). Since this is during the smooth part of




Figure 6.30 Transform from 0.0 to 0.6 Hz in Figure 6.29.

the flight, before the maneuver introduces highly variable Dopplers, the two responses in Figure 6.29 should be fully resolved. They are not, because imperfect phase compensations were (knowingly) used. This “resolution” problem, apparent when we examine intensity outputs, does not exist when we examine amplitude/phase patterns, as in Figure 6.30. As another example, the transform of the two responses between -1.4 and -1.0 Hz in Figure 6.29 is shown in Figure 6.31. The two responses are so close in Doppler that, with the existing phasing, we must include some extrapolation of the position of the left amplitude minimum (a small degree of superresolution). The scatterers are separated by 0.12 Hz, which (for a 10-second interval) allows accurate Doppler measurements only when the phase difference between the scatterers is not close to 180° . The fact that the amplitude minimum and phase jump occur nearly in the center of the interval in Figure 6.31 implies that the two scatterers are nearly in phase opposition at the time of the measurement.

6.4.1 Section Summary

The complex-image analysis procedures also improve resolution performance when a motion is so smooth that the differential Dopplers between scatterers are very small.




Figure 6.31 Transform from -1.4 to -1.0 Hz in Figure 6.29.

The tracking, compensation, and suppression of scatterers can be continued for several iterations. So long as compressed responses give good one- or two-scatterer patterns, further processing may be possible. This must be checked after each iteration.

6.5 Signal Sections and Response Groups

This section illustrates even more intricate details of the processing procedures than the preceding section, of interest primarily to those engaged in SWD processing.

If we want to improve resolution performance beyond what is obtained with intensity outputs, we must examine the transform of the complex response. If the transform amplitude function is essentially constant, the response was generated by a single scatterer. If the amplitude/phase pattern of the transform is that from two scatterers, we can determine the positions of the two scatterers. If the amplitude/phase pattern of the transform is more complicated than that from two interfering scatterers, we cannot analyze it. We must choose the transform window in such a way that effectively only the responses from two scatterers are contained within it (Step E). These rules are fairly straightforward to use when the scatterers on a target move in similar fashion, as is the case for aircraft, ships, and in most situations also for ground vehicles. On the other hand, the situation is more difficult with spinning targets or targets executing arbitrary maneuvers with a changing rotation axis. Here, if we focus one response, we will smear the other responses. This is why we must combine appropriate Doppler filtering with the amplitude/phase analysis, so that we need not analyze too many scatterers simultaneously. Below we will discuss these issues in more detail.

The first question to be considered is how to subdivide some flight record into sections for individual analysis. The general rule is to *choose sections within which the characteristics of the return are similar*, so as to minimize the inevitable smearing of responses with differently varying Dopplers (Steps A, B, and C) and to facilitate scatterer tracking (in Step F). We now illustrate the application of this rule.

Figure 6.32 gives the transform of the return from the rocket between 443 and 449 seconds. Comparing the level of the background at the right end of the spectral plot with the response level between the two response groups, we conclude that the separation between the two groups of Doppler spectra is sufficient to consider one group of responses at a time. In other words, around zero frequency the response level is almost as low as the




Figure 6.32 Transform of signal from 443 to 449 seconds.

background, so that the tails of the smeared responses do not overlap significantly. We first analyze the group between the two crosshairs. The transform over that window is shown in Figure 6.33. We note a change in the characteristics of the signal starting at about 447.6 seconds, when the scatterer causing the strong modulation with a cycle length of about 0.3 seconds disappears.

Using the rule that the signal should be analyzed in segments with similar characteristics, we decide to analyze the section from 443 to 447.6 seconds. If this should lead to processing problems, we can further reduce the time interval. Similarly, we could first attempt to track the dominant scatterer over the entire interval, reducing it only if we encounter difficulties. We also could begin by examining a short interval and increasing it until tracking is not possible. In order to obtain sufficient Doppler resolution, we must extend the processing interval over at least one full cycle of the lowest strong modulation frequency of the signal. This minimum interval goes from the start to about 444.5 seconds in Figure 6.33. The only reasons to use a longer interval are to improve resolution and to reduce the amount of processing that is needed to cover a particular flight interval.




Figure 6.33 Transform of interval indicated in Figure 6.32.

The signal for the longer interval is shown in Figure 6.34, after suppressing the background outside the spectral region of interest and removing the linear trend of the phase function. This signal shows a large variation in the ratio of amplitude maximum and minimum, and this causes a slight problem for the motion compensation. Around the time of 446 seconds, the amplitude minima are so deep that large phase jumps are generated. If we use a phase compensation based on the phase as presented in Figure 6.34 and fit a smooth function to the phase function, this function will be distorted in the area of the phase jumps because we cannot rely on the phase jumps alternating between plus and minus half a cycle. The result will be an incomplete focusing of the response of the dominant scatterer. On the other hand, when using the phase-slope compensation, the variability of the level of the amplitude minima implies that the measured phase slopes will significantly deviate from those of the dominant scatterer. As already stated, we have not




Figure 6.34 Signal section with similar characteristics.

developed a correction for this effect. We tried both of the approximate methods and found that fitting a smooth function through the phase jumps gives a somewhat better motion compensation than using the phase-slope fit without correction. After motion compensation with the first approach, we obtain Figure 6.35. The main response is sharply focused, and so are some of the minor responses. Even the responses around -10 Hz now have structure. Thus, although simply fitting through the phase jumps in Figure 6.34 does result in errors (the phase jumps could be removed before fitting a spline), the errors are acceptable. We can again iteratively process, as demonstrated above.

We now examine the right group of responses in Figure 6.32 over the same time interval as for the left group, for which the signal is shown in Figure 6.36. In accordance with the rule that the processing interval should be extended only over a signal segment with a similar structure, we now examine the amplitude function, and determine the sections within which the structure is similar; that is, the sections within which the primary modulation frequency is roughly the same. These sections are separated by the crosshairs in Figure 6.36. However, we point out that this preliminary examination of the signal is not critical. Manually, as well as in an automated




Figure 6.35 Figure 6.34 after compensation.




Figure 6.36 Signal from right response group in Figure 6.32.

fashion, we can start from the beginning, analyze a short section, and determine whether the result is a group of well-defined responses or an interference pattern. If the result consists of well-defined responses, we add another signal section, and similarly analyze the now longer section.

This process is iterated until the results are no longer acceptable, meaning that the responses are too much smeared. Then the last section with acceptable results is analyzed. The process is repeated starting from the end of the first section, and so forth, until the entire signal has been analyzed. The critical point is what we consider to be “acceptable” results of the processing. This is defined as follows. After a sequence of processing steps, we obtain a group of responses in which one response is perhaps dominant. If we can analyze these responses by taking individual transforms, and if the transforms show the amplitude/phase patterns associated with a single scatterer or two interfering scatterers, this is acceptable. If we observe the pattern of more than two interfering scatterers, Doppler resolution is insufficient, and we must lengthen the processing interval. The results are not acceptable if the transforms of the responses give the type of undefinable amplitude/phase patterns associated with the peaks of an interference pattern.

Although it is difficult to see on the scale of Figure 6.36, for the first section the phase slopes at the times of the stronger amplitude peaks vary so little that a correction should not be necessary. Thus we can take the transform over the first section without any phase compensation. Again, we need not be correct in this observation, because the results will confirm or contradict our assumption. In the latter case we would have to perform a phase compensation. The transform over this section does indeed produce a sharp dominant response, and its transform has essentially constant amplitude and a phase function with acceptable curvature. The same analyses on the weaker responses also confirm the acceptability of these responses as representing true scatterer positions rather than peaks of an interference pattern. Of course, the phase-slope measurement (which would be used in an automated system) also shows that there is no need for compensation in this instance. All of these tests fail when the interval for analysis is chosen too long, so that passing or failing of the tests can be the basis for choosing the processing interval, manually or automatically.

We have emphasized that one must select sections with similar characteristics for analysis. The main point that determines these characteristics is the duration of the amplitude (and phase) modulation cycle. If this duration changes, it may imply that the Doppler difference between the two major scatterers is changing, so that one cannot take a Fourier transform. Alternatively, a scatterer may have become weak, so that a different scatterer with a

different Doppler separation starts modulating the return from the dominant scatterer. One should not take a transform over a time interval within which the return from a scatterer quickly changes in strength.

6.5.1 Section Summary

At each stage of analysis, the signal should be divided into sections within which the characteristics of the return are similar, so as to minimize the smearing of responses with differently varying Dopplers and to facilitate scatterer tracking.

At each stage of analysis, the spectrum should be divided into groups of responses, so that the tails of the groups do not overlap significantly.

6.6 Chapter Summary

Spinning and tumbling objects introduce processing difficulties because at a specific time the scatterers within the same range gate are at different points in their sinusoidal Doppler variations. A general motion compensation that focuses all scatterers into a usable image is not possible.

The conventional SWD processing technology developed for this type of target has serious limitations. First, there is the critical loss of Doppler resolution from the use of the intensity output only. Second, there is a lack of adaptivity in signal processing when the behavior of the target changes. Third, there is the absence of a check as to whether the Doppler histories obtained are true Doppler histories or more or less interference patterns.

The methods of complex-image analysis can much improve the analysis results. As is the case with aircraft, ground vehicles, and ships, one analyzes amplitude and phase patterns to achieve the desired resolution performance and to perform measurements on the scatterers. The main difference is that the responses whose transforms give these patterns generally are smeared. This difficulty is overcome by focusing the response of the dominant scatterer, suppressing it after the measurements, focusing the response of the next dominant scatterer, suppressing it after the measurements, and so forth, until the iteration no longer works. Although the individual processing steps may be rather sophisticated, because they must be adapted to the existing situations, the processing is governed by rules that allow full automation.

Appendix A: One-Dimensional Two-Scatterer Algorithm

With the degree of resolution implemented in the typical “high-resolution” radar, there will be a considerable number of composite responses to which two scatterers have contributed, and other responses generated by three or even more scatterers. We can greatly improve functional resolution performance by utilizing the complex responses rather than just the intensity responses, at least in the case where a response is due to two scatterers only. In this appendix we discuss the mathematical details of our algorithm for resolving two scatterers, the TSA. We restrict the discussion to the application of the TSA to a one-dimensional complex image range cut. The combination of one-dimensional results to yield two-dimensional positions is explained in Section 1.3.2.2.

As explained previously, we take an image cut through a response of interest, consider various windows containing the response, and examine the Fourier transforms of those windows. We choose that transform best fitting the model of two interfering point scatterers. If the fit is good enough, we make measurements on the transform that allow us to extract the locations of the contributing scatterers, along with uncertainty estimates for those locations.

We begin our explanation with a consideration of the ideal case, the signal from two fixed point scatterers. By measuring the separation of amplitude extrema, or equivalently, the separation of phase jumps, we find the separation of the scatterer locations. From the size of the phase jumps at amplitude minima, or equivalently, the ratio of amplitude minima to maxima, we find the relative amplitude of the two scatterers. With this information we can use the phase slope at the time of the amplitude peak and the phase slope of the entire signal to locate the scatterers. Consider the signal from two fixed point scatterers:

$$\begin{aligned} S(t) &= A_1 \exp[i(\omega_2 t + \phi_2)] + A_2 \exp(\omega_2 t + \phi_2) \\ &\equiv A_1 \exp[i(\omega t + \phi)] \{1 + a \exp(\omega t + \phi)\} \end{aligned} \quad (\text{A.1})$$

The amplitude of this is given by

$$\text{Amplitude} \equiv A(t) = A_1 [1 + a^2 + 2a \cos(\omega t + \phi)]^{1/2} \quad (\text{A.2})$$

and is minimal when $\omega t + \phi = (2n + 1)\pi$, where n is an integer. If we measure the time between two consecutive minima to be t_{12} , it follows that the separation between the two scatterers is $\omega/2\pi = \pm 1/t_{12}$. Furthermore, having found ω , we have also determined ϕ . Similarly, if we measure the ratio of amplitude minimum to amplitude maximum to be a_{12} , it follows that the relative amplitude of the weaker scatterer is $a = (1 - a_{12})/(1 + a_{12})$.

The phasor diagram of Figure A.1 shows the resultant vector from the two scatterers, with the stronger normalized to strength 1 and the weaker to strength a , as in (A.1) and (A.2). The angle θ between the two varies with time t as $\theta = \omega t + \phi$. The magnitude of the resultant is given by $A(t)$, and the phase by

$$\psi(\theta) \equiv \tan^{-1} \{a \sin(\omega t + \phi) / [1 + a \cos(\omega t + \phi)]\} \quad (\text{A.3})$$

The maximum phase is

$$\psi_{\max} \equiv H = \arcsin(a) \quad (\text{A.4})$$

The "phase jump" in the composite signal has size $2\psi_{\max}$ and its duration is a fraction $(\pi/2 - \psi_{\max})/(\pi/2 + \psi_{\max})$ of the signal period. Therefore, a phase jump of size $2H$ has a width of $(2\pi/\omega)(\pi/2 - H)/(\pi/2 + H)$.




Figure A.1 Phasor diagram for two interfering scatterers.

Equation (A.4) relates the relative amplitude of the weaker scatterer to the size of the phase jump. Let us consider how to apply this. Figure A.2 shows the signal from two constructively interfering point scatterers of nearly equal amplitudes. Figure A.3 shows the same signal, centered differently in the transform window. This recentering is equivalent to adding a line to the phase function. Evidently, by adding a linear trend to the phase function, we can create and shift the locations of the extrema. Given a signal, how does one determine when phase extrema occur and measure a phase jump? Recall that (A.4) was derived from the phasor diagram of Figure A.1. In this diagram, we see that the phase of the resultant vector is zero at both the times of constructive and destructive interference. This gives us the linear trend we need to remove from the phase of Figure A.2; we must subtract a line that makes the phase values equal at times of constructive and destructive interference, which occur when the amplitude is extremal and the phase has an inflection point. Having done that, the difference between the phase extrema gives us the phase jump to use in (A.4).

In practice, recognizing the times of constructive and destructive interference may not be so easy as in Figure A.2. We often wish to analyze signals corresponding to two unresolved responses, so only part of a full interference




Figure A.2 Signal from two constructively interfering point scatterers.

cycle is observed. Furthermore, we may need to filter out interference in the image domain before we take a Fourier transform to generate the signal to be analyzed. This distorts the signal relative to the ideal case, so that the amplitude extrema and phase inflection points may be much less sharp, and may occur at different times than the constructive and destructive interference. Fortunately, using the times of amplitude extrema or phase inflection points in the distorted signal engenders little error in the derived scatterer positions and amplitudes. In fact, when the signal contains a single phase jump (or when all jumps are unwrapped in the same direction), removing a linear fit to the signal phase provides a sufficiently accurate approximation to the correct line. Figure A.3 shows the signal of Figure A.2 after subtraction of a linear least-squares fit to the phase (in this particular case, the linear fit agrees exactly with the correct line). The phase jumps are easily measured in Figure A.3.

We now show how to use phase slope measurements to extract the scatterer positions. First, note that the $\exp[i(\omega_1 t + \phi_1)]$ factor of (A.1) shifts the slope of the phase by $\omega_1/2\pi$ cycles and offsets it by $\phi_1/2\pi$ cycles. We thus consider the $\{1 + a \exp[i(\omega t + \phi)]\}$ factor, and add the resultant slope to $\omega_1/2\pi$. We consider a least-squares fit to the phase between the times of consecutive amplitude minima. The least-squares fit is found by minimizing the following expression with respect to α and β :




Figure A.3 Same signal as Figure A.2, except for centering of image transform window.

$$\sum_k \left[at_k + \beta - (1/2\pi) \tan^{-1} \left\{ \frac{a \sin(\omega t_k + \phi)}{1 + a \cos(\omega t_k + \phi)} \right\} \right]^2$$

$$\approx \int_{\omega t + \phi = -\pi}^{\omega t + \phi = \pi} dt \left[\alpha t + \beta - (1/2\pi) \tan^{-1} \{ a \sin(\omega t + \phi) / [1 + a \cos(\omega t + \phi)] \} \right]^2 \quad (\text{A.5})$$

where α and β are in cycles/sec and cycles, respectively.

Differentiating gives

$$\int_{(-\pi-\phi)/\omega}^{(\pi-\phi)/\omega} dt [\alpha t^2 + \beta t - (t/2\pi) \tan^{-1} \{ P(a, \omega t + \phi) \}] = 0 \quad (\text{A.6})$$

and

$$\int_{(-\pi-\phi)/\omega}^{(\pi-\phi)/\omega} dt [\alpha t + \beta - (1/2\pi) \tan^{-1} \{ P(a, \omega t + \phi) \}] = 0 \quad (\text{A.7})$$

where $P(a, \omega t + \phi) \equiv \{ a \sin(\omega t + \phi) / [1 + a \cos(\omega t + \phi)] \}$. These become

$$2\pi\alpha(\pi^2 + 3\phi^2) / 3\omega^3 - 2\pi\phi\beta / \omega^2 - (1 / 2\pi) \int_{(-\pi-\phi)/\omega}^{(\pi-\phi)/\omega} dt t \tan^{-1}\{P(a, \omega t + \phi)\} = 0 \quad (\text{A.8})$$

and

$$2\pi\beta / \omega - 2\pi\phi\alpha / \omega^2 - (1 / 2\pi) \int_{(-\pi-\phi)/\omega}^{(\pi-\phi)/\omega} dt t \tan^{-1}\{P(a, \omega t + \phi)\} = 0 \quad (\text{A.9})$$

Combining these gives

$$\alpha = (3f / \pi^3) \int_0^\pi dx x \tan^{-1}[a \sin x / (1 + a \cos x)] \quad (\text{A.10})$$

where $f = \omega / 2\pi$ and $x = \omega t + \phi$. Note that α is independent of ϕ .

Usually we will fit a segment of duration less than that between amplitude minima. Near the minima interference is relatively strong and the phase may reflect the interference more than the two scatterers of interest. Also, as shown in Figure A.4, the phase may not remain linear between amplitude minima. Say we measure the slope over an interval which is a fraction g of the time between amplitude minima and which is centered on an amplitude maximum. Then it runs over $\omega t + \phi = -\pi g$ to $\omega t + \phi = \pi g$, $0 < g \leq 1$. Equation (A.10) becomes

$$\alpha(a, f, g) = (3f / \pi^3 g^3) \int_0^{\pi g} dx x \tan^{-1}[a \sin x / (1 + a \cos x)] \quad (\text{A.11})$$

Figure A.5 is a plot of $\zeta(a, g) \equiv \alpha(a, g) / f$ versus a for values of g from 0.2 to 0.9, assuming scatterer 1 has an amplitude greater than or equal to that of scatterer 2. Curves are labeled by g . The greater the value of g , the smaller is the value of $\zeta(a, g)$. What is a reasonable value of g to use for real measurements? If g is large then ζ is sensitive to the precise value of g . If g is small there will be few data points in the measured phase slope. A good compromise is $g = 0.5$, which we have adopted in our measurements.

Where are we now? We have derived the parameters a and $|f|$ from our data, and have selected the value 0.5 for the parameter g . This allows us to calculate a value for $|\alpha|$ over the central half of the interval between amplitude minima. Let us now do a linear least squares fit to the phase over




Figure A.4 Signal of two interfering fixed point scatterers

this interval, denoting the measured phase slope by γ . We then have two possibilities for ω_1 :

$$\gamma = \omega_1 / 2\pi \pm |\alpha| \quad (\text{A.12})$$

How do we decide which of the two possible solutions to (A.12) represents the actual scatterer locations? We consider fitting a line to the phase of the entire signal, of duration T (1 second for our normalized figures). Then (A.6) becomes

$$\int_{-T/2}^{T/2} dt [\alpha_T t^2 + \beta_T t - (t / 2\pi) \tan^{-1} \{P(a, \omega t + \phi)\}]^2 = 0 \quad (\text{A.13})$$

giving




Figure A.5 Relation between phase slope and Dopplers of scatterers. Curves are labeled by the parameter g .

$$\alpha_T(a, f) = (6 / \pi T^3) \int_{-T/2}^{T/2} dt t \tan^{-1} \{P(a, \omega t + \phi)\} \quad (\text{A.14})$$

Let γ_T denote the measured phase slope over the entire signal. Analogously to (A.12), we have

$$\gamma_T = \omega_1 / 2\pi + \alpha_T(a, f) \quad (\text{A.15})$$

We have two solutions to (A.12) and two solutions to (A.15). For the correct choice of sign of f , the two solutions will agree. For the incorrect sign, they will not. Thus, by choosing the common solution, we have ω_1 . This also gives us the correct sign of ω . We have found two scatterers, one at $\omega_1/2\pi$ and the other at $(\omega_1 + \omega)/2\pi$.

For nonideal data, there will not be a common solution of (A.12) and (A.15). However, the pair of solutions corresponding to the correct sign of f will be nearly equal, while the other pair is not. We then select the nearly equal pair, determine the sign of ω , and use the value of ω_1 derived from (A.12), which avoids measurement over signal times when the two scatterers interfere destructively.

Besides measuring scatterer locations, we also wish to estimate the uncertainties in those measurements. To do so, we measure the consistency of the signal with the two scatterer model, and compare to measurements made on simulated data. We select simulated data with similar consistency measures, and use the average location error for those simulations as the uncertainty for the signal of interest.

What consistency measures are available? First, we should be able to measure two half-jumps for each amplitude minimum, and the times of the amplitude minima and the phase inflection points should be the same. Second, there should be an amplitude maximum midway between the amplitude minima. Third, the minima should be at the same levels. Fourth, the phase jumps should be symmetric about the inflection points. The phase difference between each inflection point and phase extremum should be equal. Fifth, a phase half-jump size of H implies a relative scatterer amplitude of $\sin(H)$. This should agree with the relative scatterer amplitude implied by the ratio of the amplitude minima and maxima. Sixth, a phase half-jump size of H implies a phase half-jump width of $(\pi / \omega)(\pi / 2 - H) / (\pi / 2 + H)$, which can be checked against the actual half-jump width. Seventh, measurements of the signal frequency ω and the relative amplitude a imply a particular phase function at times near the amplitude maximum. The implied phase curvature can be compared to the actual phase curvature. If we are using the TSA to superresolve scatterers, then only some of the above consistency measures are applicable.

With what simulations should the signal of interest be compared? This depends on the application. Certainly, interference between three and four fixed point scatterers is relevant, as is interference between a shifting scatterer and multiple fixed point scatterers. If target motion must be compensated, simulations should also include residual uncompensated motion. For safety, uncertainties should be derived from the worst conditions appropriate to the situation of interest.

Figure A.6 shows measurement error in the position of the stronger of the two scatterers, for the ideal case of two interfering fixed point scatterers without superresolution. The ordinate is position error in hundredths of a resolution cell, and the abscissa is deviation in the size of the phase jumps, in hundredths of a cycle. It is noteworthy that the overwhelming majority of measurements have errors of only a few hundredths of a resolution cell, and that the phase jumps are consistent within a few hundredths of a cycle.

Figure A.7 shows a histogram of the mean absolute position measurement error for the stronger scatterer as a function of phase-jump deviation. There is a nearly linear relation between the measurement error and the




Figure A.6 Strong scatterer position error versus phase-jump deviation, two scatterers.

deviation. Figure A.8 shows the corresponding histogram for the weaker scatterer. The error is somewhat larger than for the stronger scatterer, but still just a few hundredths of a resolution cell in the overwhelming majority of cases. Again, there is a nearly linear relation between the measurement error and the deviation.

Figure A.9 shows a histogram of the mean absolute position measurement error for the stronger scatterer as a function of phase-jump deviation, for a simulation of three interfering fixed point scatterers, without superresolution. Measurement error has increased relative to the ideal case, to a level of roughly 0.2 resolution cells, with only weak dependence on phase-jump deviation (over the small deviations plotted). Figure A.10 shows the corresponding histogram for the weaker scatterer. Measurement error has increased at low deviation, relative to the ideal case of Figure A.8. It has a weak dependence on phase-jump deviation, increasing slowly from about 0.2 to 0.25 resolution cells.

How does one combine error relations derived from Figures A.9 and A.10 with relations derived for other consistency measures? For N consistency measures, one could examine measurement error in an N -dimensional




Figure A.7 Histogram of strong scatterer position error versus phase-jump deviation, two scatterers.




Figure A.8 Histogram of weak scatterer position error versus phase-jump deviation, two scatterers.




Figure A.9 Histogram of strong scatterer position error versus phase-jump deviation, three scatterers.




Figure A.10 Histogram of weak scatterer position error versus phase-jump deviation, three scatterers.

consistency space. However, this would require difficult analysis and very extensive simulations. We believe it is more effective to analyze each consistency measure independently, in effect projecting the N -dimensional surface on one dimension at a time, producing histograms as in Figures A.9 and A.10. We then retain those consistency measures on which the position error depends strongly, discarding those for which there is only weak dependence.

Appendix B: Transform Window Limits

When analyzing complex radar imagery (one- or two-dimensional), one often wishes to examine the return from just one or two of several interfering scatterers. To do this analysis, one must take the Fourier transform of a data interval containing as much of the responses of interest as possible without including too much interference. As explained heuristically in [1], phase inflection and amplitude minimum locations are good choices for interval boundaries. Here we supply analytic justification for these choices. Specifically, we show that for two interfering point scatterers, the scatterer strengths are equal at the phase inflection point. Also, if the interference is generally destructive, the scatterer strengths are nearly equal at the amplitude minimum.

We consider the case of one-dimensional data, and assume Gaussian weighting, an infinite duration signal, and no sampling. Then the image domain response of N scatterers located at frequencies f_j with amplitudes a_j and phases ϕ_j is

$$I(f) = \sum_{j=1}^N a_j \exp[-b(f - f_j)^2 + i\phi_j] \quad (\text{B.1})$$

where b depends on the details of the Gaussian weighting. This has a phase Φ of

$$2i\Phi = \ln \left\{ \sum_{j=1}^N a_j \exp[-b(f - f_j)^2 + i\phi_j] \right\} - c.c. \quad (\text{B.2})$$

where $c.c.$ denotes complex conjugate. Differentiating (B.2) twice gives

$$\begin{aligned} 2i\ddot{\Phi} / 4b^2 = & \left\{ -\sum_{j=1}^N a_j f_j (f - f_j) \exp[-b(f - f_j)^2 + i\phi_j] \right. \\ & \times \sum_{m=1}^N a_m \exp[-b(f - f_m)^2 + i\phi_m] + \sum_{j=1}^N a_j f_j \exp[-b(f - f_j)^2 + i\phi_j] \\ & \times \sum_{m=1}^N a_m (f - f_m) \exp[-b(f - f_m)^2 + i\phi_m] \left. \right\} \\ & / \left\{ \sum_{k=1}^N a_k \exp[-b(f - f_k)^2 + i\phi_k] \right\}^2 - c.c. \end{aligned} \quad (\text{B.3})$$

The inflection point is $\ddot{\Phi} = 0$. This implies

$$\begin{aligned} 0 = & \sum_{j,k,m,n=1}^N a_j a_k a_m a_n \exp\{-b[(f - f_j)^2 + (f - f_k)^2 \\ & + (f - f_m)^2 + (f - f_n)^2]\} \times \{[-f_j(f - f_j) \\ & + f_j(f - f_k)] \exp[i(\phi_j + \phi_k - \phi_m - \phi_n)] - c.c.\} \end{aligned} \quad (\text{B.4})$$

Simplifying (B.4) gives

$$\begin{aligned} 0 = & \sum_{j,k,m,n=1}^N a_j a_k a_m a_n \exp\{-b[(f - f_j)^2 + (f - f_k)^2 + (f - f_m)^2 \\ & + (f - f_n)^2]\} \times f_j (f_j - f_k) \sin(\phi_j + \phi_k - \phi_m - \phi_n) \end{aligned} \quad (\text{B.5})$$

Now, consider $N = 2$. This gives 16 terms in the right-hand side of (B.5), of which only four are nonzero. These yield

$$0 = a_1 a_2 \exp\{-b[(f - f_1)^2 + (f - f_2)^2]\} (f_1 - f_2)^2 \quad (\text{B.6})$$

$$\times \{a_1^2 \exp[-2b(f - f_1)^2] - a_2^2 \exp[-2b(f - f_2)^2]\}$$

implying

$$a_1 \exp[-b(f - f_1)^2] = a_2 \exp[-b(f - f_2)^2] \quad (\text{B.7})$$

This is also the definition of the frequency f where the amplitudes of the two scatterers are equal. Hence, the same frequency f that gives the inflection point also gives the change in dominance from one scatterer to the other.

Solving for f we get the frequency at which dominance shifts:

$$f_c = (f_1 + f_2) / 2 + \ln(a_1 / a_2) / [2b(f_2 - f_1)] \quad (\text{B.8})$$

Next, we examine the location of amplitude (or power) minima for generally destructive interference. We can write the power as

$$P = \sum_{j,k=1}^N a_j a_k \exp\{-b[(f - f_j)^2 + (f - f_k)^2] + i(\phi_j - \phi_k)\} \quad (\text{B.9})$$

Differentiating (B.9) gives

$$\dot{P} = \sum_{j,k=1}^N a_j a_k \exp\{-b[(f - f_j)^2 + (f - f_k)^2] + i(\phi_j - \phi_k)\} (2f - f_j - f_k) \quad (\text{B.10})$$

At the power minimum $\dot{P} = 0$. Then at the power minimum, for $N = 2$, (B.10) becomes

$$0 = 2a_1^2 (f - f_1) \exp[-2b(f - f_1)^2] + 2a_2^2 (f - f_2) \exp[-2b(f - f_2)^2]$$

$$+ 2a_1 a_2 (2f - f_1 - f_2) \cos(\phi_2 - \phi_1) \exp\{-b[(f - f_1)^2 + (f - f_2)^2]\} \quad (\text{B.11})$$

Rewriting $f = f_c + \delta$ (so δ is the offset from the frequency of dominance change) and removing common factors from (B.11) gives

$$0 = (f_c + \delta - f_1) \exp[2b\delta(f_1 - f_2)] + (f_c + \delta - f_2) \exp[2b\delta(f_2 - f_1)] \\ + (2f_c + 2\delta - f_1 - f_2) \cos(\phi_2 - \phi_1) \quad (\text{B.12})$$

For purely destructive interference $\phi_2 - \phi_1 = \pi$, and (B.12) becomes

$$0 = \{\exp[2b\delta(f_1 - f_2)] - 1\}(f_c + \delta - f_1) \\ + \{\exp[2b\delta(f_2 - f_1)] - 1\}(f_c + \delta - f_2) \quad (\text{B.13})$$

which evidently has the solution $\delta = 0$. Thus, for purely destructive interference, the amplitude minimum and the phase inflection point occur at the same frequency, that of the change in dominance.

Now, assume generally (but not purely) destructive interference. Defining $\phi_2 - \phi_1 = \pi + \varepsilon$ and approximating $\cos \varepsilon \approx 1 - \varepsilon^2 / 2$, (B.12) becomes

$$0 \approx \{\exp[2b\delta(f_1 - f_2)] - 1\}(f_c + \delta - f_1) + \{\exp[2b\delta(f_2 - f_1)] - 1\} \\ (f_c + \delta - f_2) + (\varepsilon^2 / 2)(2f_c + 2\delta - f_1 - f_2) \quad (\text{B.14})$$

Expanding the exponentials of (B.14) in Taylor series and retaining terms linear in δ gives

$$\delta \approx \varepsilon^2 \ln(a_1 / a_2) / [4b^2(f_2 - f_1)^3] \quad (\text{B.15})$$

Equation (B.15) shows that the offset of the amplitude minimum from the phase inflection point (and the point of dominance change) is small for generally destructive interference. Hence, the two scatterers are nearly equally strong at the amplitude minimum.

Besides locating the point of dominance change, we can also find the point midway between the two scatterers. Differentiating (B.2) and multiplying by (B.9) gives

$$P\dot{\Phi} = -2b \sum_{j,k=1}^N a_j a_k (f - f_k) \sin(\phi_k - \phi_j) \exp\{-b[(f - f_j)^2 + (f - f_k)^2]\} \quad (\text{B.16})$$

For $N=2$, (B.16) becomes

$$P\dot{\Phi} = 2b_2 a_1 a_2 (f_1 - f_2) \sin(\phi_1 - \phi_2) \exp\{-b[(f - f_1)^2 + (f - f_2)^2]\} \quad (\text{B.17})$$

The derivative of this is proportional to $(2f - f_1 - f_2)$, so is zero when $f = (f_1 + f_2) / 2$. Thus, the midpoint between the two scatterers is located where $P\dot{\Phi}$ has an extremum.

Reference

- [1] Rihaczek, A. W., and S. J. Hershkowitz, *Radar Resolution and Complex-Image Analysis*, Norwood, MA: Artech House, 1996.

Appendix C: Determining Scatterer Separations in Range Gates With Residual Uncompensated Motion

Because moving ground vehicles bend and flex, a motion compensation that is based on a rigid-body assumption will generally leave parts of the vehicle with residual uncompensated motion. Typically, this residual motion will not be large enough to cause scatterers to drift from one range gate to another, but will be more than large enough to smear each scatterer response over many crossrange gates, generating multiple response peaks. If a bending/flexing region of the vehicle contains a scatterer that is sufficiently well resolved or enough stronger than its neighbors, that scatterer can be used to derive a compensation for that region of the vehicle. Unfortunately, a bending/flexing region need not be large, and such a scatterer may not be available.

This appendix describes a technique, applicable when such a scatterer is not available, for determining scatterer crossrange separations. It relies on using two nearby range gates undergoing similar bending/flexing. First, we take the Fourier transform of each range gate. The amplitude function of each resulting signal contains information determining the absolute scatterer separations in the corresponding range gate, but not the signs of the

separations. The absolute separations are most easily extracted by examining the transform of the signal amplitude, as in Section 4.4.9.

In order to determine the signs of the separations, we multiply the signal from one range gate by the complex conjugate of the signal from the other. The residual uncompensated common motion cancels out of this complex product signal. If the two range gates have different scatterer separations or relative amplitudes, the product signal allows us to determine the signs of the separations. This is shown below.

Consider the case where we have two nearby range gates containing scatterers undergoing the same unknown motion. The signals for these gates can be written as

$$S_j(t) = \sum_{k=1}^{N_j} A_{jk} \exp[i(\omega_{jk}t + \phi_{jk})] \exp[i\Phi(t)] \quad (\text{C.1})$$

where j denotes the gate, there are N_j scatterers in gate j , and $\Phi(t)$ is the common phase motion.

The power in a gate is

$$P_j(t) = S_j S_j^*(t) = \sum_{k,n=1}^{N_j} A_{jk} A_{jn} \cos[(\omega_{jk} - \omega_{jn})t + \phi_{jk} - \phi_{jn}] \quad (\text{C.2})$$

This is independent of $\Phi(t)$, and is real. An FFT will produce a symmetric spectrum, with responses at $\pm(\omega_{jk} - \omega_{jn})$.

Now multiply the signal in one gate with the complex conjugate of that in another:

$$Q_{jm}(t) \equiv S_j(t) S_m^*(t) = \sum_{k=1}^{N_j} \sum_{n=1}^{N_m} A_{jk} A_{mn} \exp\{i[(\omega_{jk} - \omega_{mn}) + \phi_{jk} - \phi_{mn}]\} \quad (\text{C.3})$$

Because the two sets $\{\omega_{jk}\}$ and $\{\omega_{mn}\}$ are different, $Q_{jm}(t)$ is not real, and its transform is asymmetric. We can use this asymmetry to measure $\omega_{jk} - \omega_{mn}$ and thereby determine the signs of $\{\omega_{jk}\}$ and $\{\omega_{mn}\}$.

We first illustrate this with synthetic data. Figure C.1 shows the image response from scatterers at Crossrange Gates 33 and 36, plus a quadratic phase drift in time. Figure C.2 shows the corresponding signal. Figure C.3




Figure C.1 Simulated scatterers at crossranges 33 and 36, including uncompensated motion.




Figure C.2 Signal corresponding to Figure C.1.




Figure C.3 Transform of the power of Figure C.2.

shows the transform of the power of Figure C.2 with peaks separated by ± 3 Hz (with our normalization, 1 Hz corresponds to one crossrange gate).

Figure C.4 shows the image response from scatterers at Crossrange Gates 33 and 28, plus the same quadratic drift. Figure C.5 shows the corresponding signal, and Figure C.6 the transform of the signal power. Peaks there are separated by ± 5 Hz.

Figure C.7 shows the product of the signal of Figure C.2 with the complex conjugate of the signal of Figure C.5. Evidently, the common quadratic phase drift is absent and the signal is complex. Figure C.8 shows the transform of Figure C.7. It contains responses at frequencies 0, 3, 5, and 8. The response at frequency 0 is generated by the combination of the strongest responses in each of the original gates. The separations and relative amplitudes of the other responses can be used to determine the response separations in Figures C.1 and C.4.

The second strongest response in Figure C.8 is separated from the strongest by 5 Hz (gates). Comparing the relative amplitude and the separation to Figures C.3 and C.6 tells us that this response is generated from the strongest scatterer of Figure C.1 and the secondary scatterer of Figure C.4.




Figure C.4 Simulated scatterers at crossranges 33 and 28, including the same uncompensated motion of Figure C.1.




Figure C.5 Signal corresponding to Figure C.4.




Figure C.6 Transform of the Power of Figure C.5.




Figure C.7 Product of signal of Figure C.2 and complex conjugate of signal of Figure C.5.




Figure C.8 Transform of Figure C.7.

Because we took the conjugate of the signal of Figure C.5, the separation of the secondary and primary scatterers of Figure C.4 is $-(5 - 0) = -5$ gates.

The response at 3 Hz in Figure C.8 is from the primary scatterer of Figure C.4 and the secondary scatterer of Figure C.1. Thus the separation of the secondary and primary scatterers of Figure C.1 is $(3 - 0) = 3$ gates. The remaining response of Figure C.8, at 8 Hz, is generated by the secondary scatterers. As expected, its position is $3 - (-5) = 8$ Hz.

Will this work with real data? The data must satisfy several conditions. First, the motion of the scatterers in the two gates must be nearly the same. Second, each gate must contain few scatterers for the unscrambling to succeed. Third, the separations or relative amplitudes in the two gates must differ. Fourth, we must have a sufficient signal to noise ratio.

Figure C.9 shows the image cut in Range Gate 0 of Figure 4.98, with the crossrange positions shifted by 128 gates. Figure C.10 shows the transform of the power of the corresponding signal. Figure C.11 shows the image cut in Range Gate 2 of Figure 4.98, with the same crossrange shift. Range Gate 2 is far enough from Range Gate 0 to contain different scatterers, but close enough to have similar motion. Figure C.12 shows the transform of the power of the signal corresponding to Range Gate 2. Figures C.10 and C.12




Figure C-9 Image cut in Range Gate 0 of Figure 4.98.




Figure C.10 Transform of the power of the signal corresponding to Figure C.9.




Figure C.11 Image cut in Range Gate 2 of Figure 4.98.




Figure C.12 Transform of the power of the signal corresponding to Figure C.11.




Figure C.13 Product of signal of Gate 0 and complex conjugate of signal of Gate 2.




Figure C.14 Transform of Figure C.13.

show that the absolute separations of the primary and secondary scatterers in Range Gates 0 and 2 are about 3 and 4.5 crossrange gates, respectively.

Figure C.13 shows the product of the signal for Gate 0 and the complex conjugate of the signal for Gate 2. Its transform is given by Figure C.14. The large peak near -1 Hz corresponds to the two primary scatterers. We must search for responses ± 3 and ± 4.5 gates from this. The second strongest response of the figure is near 2 Hz, with an indication of another near 3.5 Hz. Thus, the secondary response in Range Gate 0 is separated from the primary by $-(2 - (-1)) = -3$ gates, and the secondary response in Range Gate 2 is separated from the primary by $-(3.5 - (-1)) = -4.5$ gates. This agrees with the adaptive analysis of Section 4.4.6.

Appendix D: Errors in Estimating Times of Two-Dimensional Motion

In this appendix, we estimate the errors in the algorithms of Chapter 5 for determining the times of zero yaw and roll Doppler. To this end, we consider the rotation of a target about a motion compensated point on that target. We define a target-centered (x', y', z') coordinate system and a radar-centered (x, y, z) coordinate system, with y being the downrange direction and x horizontal. Without any loss of generality, we take the z' direction as vertical for an unbanked, unpitched target, the x' direction along the long horizontal axis, and the y' direction as defined by the right-hand rule. After motion compensation, these systems can be related by

$$\begin{bmatrix} x \\ y \\ z \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos \delta & -\sin \delta \\ 0 & \sin \delta & \cos \delta \end{bmatrix} \begin{bmatrix} \cos \alpha & -\sin \alpha & 0 \\ \sin \alpha & \cos \alpha & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos \beta & -\sin \beta \\ 0 & \sin \beta & \cos \beta \end{bmatrix} \begin{bmatrix} \cos \rho & 0 & -\sin \rho \\ 0 & 1 & 0 \\ \sin \rho & 0 & \cos \rho \end{bmatrix} \begin{bmatrix} x' \\ y' \\ z' \end{bmatrix} \quad (\text{D.1})$$

where δ is the radar depression angle, α is the target aspect angle (0° being broadside), β is the target bank angle, and ρ is the target pitch angle [1].

Carrying out the multiplications gives a range of

$$\begin{aligned}
y = & (\cos \delta \cos \rho \sin \alpha - \cos \delta \sin \rho \sin \beta \cos \alpha - \sin \delta \sin \rho \cos \beta)x' \\
& + (\cos \delta \cos \beta \cos \alpha - \sin \delta \sin \beta)y' \\
& - (\cos \delta \sin \rho \sin \alpha + \cos \delta \cos \rho \sin \beta \cos \alpha + \sin \delta \cos \rho \cos \beta)z'
\end{aligned} \tag{D.2}$$

The range and the angles in (D.2) are time dependent, while the body-centered coordinates are not. Thus, we can rewrite the equation as

$$y(t) = A(t)x' + B(t)y' + C(t)z' \tag{D.3}$$

Differentiating (D.3) gives us a relation between measured range-rate and Doppler components:

$$\dot{y}(t) = \dot{A}(t)x' + \dot{B}(t)y' + \dot{C}(t)z' \tag{D.4}$$

The primary aim of the approach described in Chapter 5 is to determine those times when $\dot{C}(t)$ in (D.4) is zero. At these times, Dopplers are independent of scatterer height. The motion is about a vertical target axis, and a good ISAR image is obtainable. The secondary aim is to determine those times when the sum of the other terms is zero. Then, the Doppler depends only on height, and the motion is about a horizontal target axis, yielding a good profile view of the target.

Under ideal conditions, we would select a scatterer at the bow and another at the stern, separated only in their x' coordinates. We would then select two scatterers (or select one scatterer and generate a virtual scatterer) separated only in their z' coordinates. Note that, even under these ideal conditions, we do not expect to determine times when yaw, roll, or pitch Dopplers are zero. We are attempting to determine times when motion is about a horizontal or vertical target axis. Under real conditions, we cannot select scatterers separated in only one coordinate. We now consider the error this introduces in time determination.

Let t_0 be a time of two-dimensional motion. For small errors, we can Taylor-expand the range rate as

$$\dot{y}(t_0 + \varepsilon) \approx \dot{y}(t_0) + \varepsilon \ddot{y}(t_0) \tag{D.5}$$

The procedure that we describe in Chapter 5 for determining times of two-dimensional motion selects those times ($t_0 + \varepsilon$) when the measured

relative range-rate between two scatterers is zero. In other words, the left-hand side of (D.5) is zero. This gives the estimate of the timing error as

$$\varepsilon \approx -\dot{y}(t_0) / \ddot{y}(t_0) \approx -[\dot{A}(t_0)x' + \dot{B}(t_0)y' + \dot{C}(t_0)z'] / [\ddot{A}(t_0)x' + \ddot{B}(t_0)y' + \ddot{C}(t_0)z'] \quad (\text{D.6})$$

As we attempt to select imaging times when the motion is smooth, as well as about a fixed axis, we can approximate the yaw, roll, and pitch motions as periodic for the error analysis. We write

$$\begin{aligned} \alpha(t) &\equiv \alpha_0 + \gamma(t) \equiv \alpha_0 + \gamma_m \sin(\omega_\gamma t + \phi_\gamma) \\ \beta(t) &\equiv \beta_m \sin(\omega_\beta t + \phi_\beta) \\ \rho(t) &\equiv \rho_m \sin(\omega_\rho t + \phi_\rho) \end{aligned} \quad (\text{D.7})$$

where α_0 is the aspect at time t_0 . Over typical observation intervals, and even more so for the purposes of Taylor expansion, γ , β , and ρ are all much less than unity. We further assume the depression angle δ is similarly small. Thus, we will estimate the timing error by expanding (D.6) in these angles, retaining only leading terms.

For the case of primary interest, error in estimating when the “roll” Doppler is zero, we have $\dot{C}(t_0) \equiv 0$. With our periodic assumption, $\ddot{C}(t_0)$ is extremal. Furthermore, we have attempted to pick scatterers with a separation just in height z' . Thus, the denominator of (D.6) is dominated by the third term. This gives

$$\varepsilon_r \approx -[\dot{A}(t_0)x' + \dot{B}(t_0)y'] / \ddot{C}(t_0)z' \quad (\text{D.8})$$

Carrying out the differentiation and Taylor expansion, retaining only leading terms gives

$$\varepsilon_r \approx \dot{\gamma}(x' \cos \alpha_0 + y' \sin \alpha_0) / [(\ddot{\rho} \sin \alpha_0 + \ddot{\beta} \cos \alpha_0)z'] \quad (\text{D.9})$$

Squaring each side and taking the expectation value over many relative phasings gives

$$\langle \varepsilon_r^2 \rangle \approx \omega_\gamma^2 \gamma_m^2 (x' \cos \alpha_0 + y' \sin \alpha_0)^2 / [(\omega_\rho^4 \rho_m^2 \sin^2 \alpha_0 + \omega_\beta^4 \beta_m^2 \cos^2 \alpha_0)z'^2] \quad (\text{D.10})$$

If the motions have comparable periods T and comparable maximum angular variations, this gives an rms timing error

$$\varepsilon_r \approx |\max(x', y')T / 2\pi z'| \quad (\text{D.11})$$

If we wish our rms error to be less than, say, 0.1 seconds, and have a period of about three seconds, we must choose our scatterers so their separation in the deck coordinates x' and y' is less than one-fifth their separation in height above the deck z' . This requirement is easily met if we use three scatterer tracks and generate a virtual scatterer near the deck position of the superstructure scatterer. For a superstructure height of 20 ft, we must match the deck position within about 4 ft (4 range gates for typical range resolution). If we wish to use four scatterer tracks, the height separation of the two superstructure scatterers will be less than the height of the superstructure, so their range separation must be reduced proportionately. Thus, when yaw, pitch, and roll have comparable motion periods and maximum angular variations, the virtual scatterer method is generally required, depending on the height of the superstructure. When roll motion is dominant, the deck spacing between the two superstructure scatterers can be larger, and the four-scatterer method can be employed. Conversely, when yaw motion is dominant, the virtual scatterer method is almost always required.

For the case of secondary interest, error in estimating when the “yaw” Doppler is zero, we have $\dot{A}(t_0)x' + \dot{B}(t_0)y' \equiv 0$. With our periodic assumption, $\ddot{A}(t_0)x' + \ddot{B}(t_0)y'$ is extremal. Furthermore, we have attempted to pick scatterers at the same height z' , with a separation primarily in x' , along the ship's length. Thus, the denominator of (D.6) is dominated by the first two terms. This gives

$$\varepsilon_y \approx -\dot{C}(t_0)z' / [\ddot{A}(t_0)x' + \ddot{B}(t_0)y'] \quad (\text{D.12})$$

Carrying out the differentiation and Taylor expansion, retaining only leading terms gives

$$\varepsilon_y \approx (\dot{\rho} \sin \alpha_0 + \dot{\beta} \cos \alpha_0)z' / [\ddot{\gamma}(x' \cos \alpha_0 - y' \sin \alpha_0)] \quad (\text{D.13})$$

Squaring each side and taking the expectation value over many relative phasings gives

$$\langle \varepsilon_y^2 \rangle \approx (\omega_\rho^2 \rho_m^2 \sin^2 \alpha_0 + \omega_\beta^2 \beta_m^2 \cos^2 \alpha_0) z'^2 / [\omega_\gamma^4 \gamma_m^2 (x' \cos \alpha_0 - y' \sin \alpha_0)^2] \quad (\text{D.14})$$

If the motions have comparable periods T and comparable maximum angular variations, this gives an rms timing error

$$\varepsilon_y \approx |Tz' / 2\pi x'| \quad (\text{D.15})$$

If we wish our rms error to be less than, say, 0.1 seconds, and have a period of about 3 seconds, we must choose our scatterers so their separation in height above the deck z' is less than one-fifth their separation along the long deck coordinate x' . Most ships have a deck height variation that is much smaller than the deck length. For these ships, the timing error is insignificant when the yaw, pitch, and roll motions are comparable. Those ships that have large enough deck height variations to be problematic will be obvious in approximate roll images. We can use such an image to choose two deck scatterers at similar heights, perhaps sacrificing some separation along the length of the ship. If yaw motion is dominant, height variation is even less significant. Conversely, if roll motion is dominant, we must be more careful in our choice of scatterers. We must use an approximate roll image to select two scatterers near the same height.

In the foregoing, we have analyzed error due to the impossibility of choosing scatterers that are separated only in height above the deck, or only along the length of the ship. Mismeasurement of scatterer range-rates also contributes to errors in determining times when motion is two-dimensional. However, so long as each track does indeed correspond to a single scatterer, rather than interference between comparable scatterers, the tracking inaccuracy is inconsequential relative to the errors due to nonideal scatterer locations.

Reference

- [1] Goldstein, H., *Classical Mechanics*, Reading, MA: Addison-Wesley, 1950.

Appendix E: Extracting Scatterer Locations From Range Tracks

In this appendix, we consider a mathematical approach to extracting scatterer locations and target motion from the tracks of four scatterers. We rewrite (D.3) as

$$R(t) = A(t)x' + B(t)y' + C(t)z' \quad (\text{E.1})$$

where $R(t)$ is the measured range of a scatterer after some other point on the target has been motion compensated. We denote the scatterer used for compensation as scatterer 0, and the others as 1, 2, and 3. We can then write

$$\begin{bmatrix} R_1(t) \\ R_2(t) \\ R_3(t) \end{bmatrix} = \begin{bmatrix} x'_1 & y'_1 & z'_1 \\ x'_2 & y'_2 & z'_2 \\ x'_3 & y'_3 & z'_3 \end{bmatrix} \begin{bmatrix} A(t) \\ B(t) \\ C(t) \end{bmatrix} \quad (\text{E.2})$$

Inverting this yields

$$\begin{bmatrix} A(t) \\ B(t) \\ C(t) \end{bmatrix} = \begin{bmatrix} x'_1 & y'_1 & z'_1 \\ x'_2 & y'_2 & z'_2 \\ x'_3 & y'_3 & z'_3 \end{bmatrix}^{-1} \begin{bmatrix} R_1(t) \\ R_2(t) \\ R_3(t) \end{bmatrix} \quad (\text{E.3})$$

We now note that

$$A^2(t) + B^2(t) + C^2(t) = 1 \quad (\text{E.4})$$

as can be verified by direct substitution from (D.2), or by observing that these three functions are generated by the rotation matrices of (D.1) and must preserve length. Substituting (E.3) into (E.4) gives

$$\begin{bmatrix} R_1(t) & R_2(t) & R_3(t) \end{bmatrix} \begin{bmatrix} x'_1 & x'_2 & x'_3 \\ y'_1 & y'_2 & y'_3 \\ z'_1 & z'_2 & z'_3 \end{bmatrix}^{-1} \begin{bmatrix} x'_1 & y'_1 & z'_1 \\ x'_2 & y'_2 & z'_2 \\ x'_3 & y'_3 & z'_3 \end{bmatrix}^{-1} \begin{bmatrix} R_1(t) \\ R_2(t) \\ R_3(t) \end{bmatrix} = 1 \quad (\text{E.5})$$

We can rewrite this as

$$\begin{bmatrix} R_1(t) & R_2(t) & R_3(t) \end{bmatrix} \begin{bmatrix} \mathbf{r}'_1 \cdot \mathbf{r}'_1 & \mathbf{r}'_1 \cdot \mathbf{r}'_2 & \mathbf{r}'_1 \cdot \mathbf{r}'_3 \\ \mathbf{r}'_2 \cdot \mathbf{r}'_1 & \mathbf{r}'_2 \cdot \mathbf{r}'_2 & \mathbf{r}'_2 \cdot \mathbf{r}'_3 \\ \mathbf{r}'_3 \cdot \mathbf{r}'_1 & \mathbf{r}'_3 \cdot \mathbf{r}'_2 & \mathbf{r}'_3 \cdot \mathbf{r}'_3 \end{bmatrix}^{-1} \begin{bmatrix} R_1(t) \\ R_2(t) \\ R_3(t) \end{bmatrix} = 1 \quad (\text{E.6})$$

where \mathbf{r}'_j is the vector (x'_j, y'_j, z'_j) . Defining $r_{jk} \equiv \mathbf{r}'_j \cdot \mathbf{r}'_k$ and inverting the matrix in (E.6) gives

$$\begin{bmatrix} R_1 & R_2 & R_3 \end{bmatrix} \begin{bmatrix} r_{22}r_{33} - r_{23}r_{23} & r_{13}r_{23} - r_{12}r_{33} & r_{12}r_{23} - r_{22}r_{13} \\ r_{13}r_{23} - r_{12}r_{33} & r_{11}r_{33} - r_{13}r_{13} & r_{12}r_{13} - r_{11}r_{23} \\ r_{12}r_{23} - r_{13}r_{22} & r_{12}r_{13} - r_{11}r_{23} & r_{11}r_{22} - r_{12}r_{12} \end{bmatrix} \begin{bmatrix} R_1 \\ R_2 \\ R_3 \end{bmatrix} = D \quad (\text{E.7})$$

where we have suppressed the time dependence of $R_k(t)$ and where

$$D = r_{11}r_{22}r_{33} + 2r_{12}r_{23}r_{13} - r_{13}^2r_{22} - r_{23}^2r_{11} - r_{12}^2r_{33} \quad (\text{E.8})$$

Changing variables, we have

$$\begin{bmatrix} R_1(t) & R_2(t) & R_3(t) \end{bmatrix} \begin{bmatrix} b_0 & b_1 & b_2 \\ b_1 & b_3 & b_4 \\ b_2 & b_4 & b_5 \end{bmatrix} \begin{bmatrix} R_1(t) \\ R_2(t) \\ R_3(t) \end{bmatrix} = 1 \quad (\text{E.9})$$

If our four tracked scatterers are not coplanar, we can solve (E.9) for the coefficients b_j by a standard polynomial least-squares fit [1]. This entails inverting a matrix that is nearly singular for four scatterers that are nearly coplanar; such a set of scatterers will produce inaccurate results. Having solved for the coefficients, we extract the dot products of (E.7) and (E.8) algebraically:

$$\begin{aligned} r_{11} &= (b_4^2 - b_3b_5)E \\ r_{12} &= (b_1b_5 - b_2b_4)E \\ r_{13} &= (b_2b_3 - b_1b_4)E \\ r_{22} &= (b_2^2 - b_0b_5)E \\ r_{23} &= (b_0b_4 - b_1b_2)E \\ r_{33} &= (b_1^2 - b_0b_3)E \\ E &= [(b_2^2 - b_0b_5)(b_1^2 - b_0b_3) - (b_0b_4 - b_1b_2)^2] / b_0 \\ &= [(b_4^2 - b_3b_5)(b_2^2 - b_0b_5)(b_1^2 - b_0b_3) - (b_2b_3 - b_1b_4)^2(b_2^2 - b_0b_5) - \\ &\quad (b_0b_4 - b_1b_2)^2(b_4^2 - b_3b_5) - (b_1b_5 - b_2b_4)^2(b_1^2 - b_0b_3) + \\ &\quad 2(b_1b_5 - b_2b_4)(b_0b_4 - b_1b_2)(b_2b_3 - b_1b_4)] \end{aligned} \quad (\text{E.10})$$

We have extracted only the dot products of the vectors between the four scatterers. This is all that can be extracted via this method. Note that we can insert a rotation matrix and its transpose in the center of the left hand of (E.5); we cannot distinguish any particular coordinate system for the scatterer locations. We can define a tetrahedron of scatterers, but not its orientation.

Uncertainty calculations for the above formulation are straightforward, but not illuminating, because of the complexity of (E.10) and the details of the linear-least-squares fit used to solve (E.9). Furthermore, such calculations reveal error propagation due to random error (noise), but ignore systematic error. In order to obtain a feel for the required accuracies, we examine a sample case.

We applied the above analysis to the simulated data of Figures 5.13 through 5.18. In this simulation, we generated 20 seconds of data at X-band with a bandwidth of 500 MHz and a PRF of 200 Hz by rotating a 100-ft-long "ship" composed of 20 fixed point scatterers. We used periodic yaw, roll, and pitch motions, all with maximum excursions of $\pm 1.5^\circ$, with a starting aspect of 35° off nose-on and zero starting bank and pitch angles. The yaw period was 2 seconds, the roll period 2.4 seconds, and the pitch period 2.7 seconds. Thus, the relative phasing of the motions produces different intervals when only one, only two, and all three motions are effective. We tracked a scatterer near the bow, another near the stern, and two at different heights and deck positions on the superstructure. We then applied the above analysis, using the bow scatterer as the reference, and examined how the results varied as we removed refinements to the scatterer tracks.

Defining z' to be height above the deck, y' to be distance along the long axis of the ship, and x' to be perpendicular to these, the (x', y', z') coordinates of the bow scatterer are $(0, -15.24, 0)$, those of the stern scatterer are $(0, 15.24, 0)$, and those of the superstructure scatterers are $(0.91, 1.52, 2.74)$ and $(-0.46, -0.61, 5.49)$, all in m. As discussed above, only the relative positions of the scatterers forming the tetrahedron can be extracted. In order to examine numerical results, we must choose a coordinate system. We choose the origin at the bow scatterer, and the natural system of y'' lying along the line between bow and stern scatterers, z'' perpendicular to this line and such that the high superstructure scatterer is in the $y''-z''$ plane, and x'' perpendicular to the other two axes. The correct answers then are (x'', y'', z'') coordinates of $(0, 30.48, 0)$ for the stern, $(0, 14.63, 5.51)$ for the high superstructure scatterer, and $(1.14, 16.76, 2.66)$ for the other superstructure scatterer. With this choice of coordinates, $\dot{B}(t)$ and $\dot{C}(t)$, which can be extracted from (E.3), correspond to the yaw and roll Dopplers of Chapter 5, respectively.

The full tracking of the scatterers consisted of a range track, followed by a Doppler track, and then a phase track. The mean error of the complete tracks was 0.0001 wavelengths, with a standard deviation of about 0.005 wavelengths; the tracks are essentially perfect. The extracted positions for these perfect tracks, shown in Table E.1, are also quite accurate, as are the

yaw and roll Dopplers. Next, we consider just range and Doppler tracking. Without the phase tracking, the mean error for three of the four tracks was about 0.05 wavelengths, with a standard deviation of about 0.009 wavelengths. The error for the fourth scatterer was about -0.19 wavelengths, with a standard deviation of 0.02 wavelengths. These deviations correspond to fitting inaccuracies, so the scatterers have pseudoperiodic residual motion on the order of the deviations. While the accuracy of the positions extracted from these tracks (shown in Table E.1) is somewhat degraded, it is still sufficient for practical application. The yaw and roll Dopplers also differ only slightly from the correct results. The small biases in the tracks are not very significant to the analysis.

Next, we consider just range tracking. The mean error and standard deviations for the four scatterers are 0.001 ± 0.19 wavelengths, 0.001 ± 0.07 wavelengths, 0.003 ± 0.20 wavelengths, and -0.04 ± 0.14 wavelengths. Now, the analysis produces dot products r_{jk} that are not self-consistent. If we use three scatterer tracks with range and Doppler tracking, and one with just range tracking, trying all four scatterers as the one without Doppler tracking, we find two cases of inconsistent dot products, one case of very inaccurate results, and one case of *inaccurate* but probably practically acceptable results. This should not be taken as an indication that one could not obtain self-consistent results from the range measurements. The requirement of self-consistency could be incorporated into the derivation of the b_j or the r_{jk} , although this might require numerical optimization rather than analytic solution of (E.7) or (E.9). However, although such an alternative formulation will produce self-consistent results, it will not produce more accurate results.

Table E.1
Variation of Extracted Positions With Tracking

Correct	Range+Doppler+Phase	Range+Doppler
30.48	30.73	31.14
14.63	14.74	14.91
5.51	5.46	5.40
16.76	16.87	17.04
1.14	1.16	1.19
2.66	2.64	2.62

This leads to the conclusion that, for motion conditions similar to those simulated, pseudoperiodic motion residuals on the order of 0.01 to 0.02 wavelengths are insignificant, while residuals on the order of 0.1 to 0.2 wavelengths can be catastrophic. The question arises, how well can we expect to track four scatterers in real data? Our experience indicates that achieving the necessary accuracy will be problematic in most cases. Rarely are four scatterers well enough isolated from interference (from true scatterers and spurious responses) that they can be tracked to this accuracy over several motion cycles. The tracking requirements may be somewhat eased because the solution of (E.9), or its equivalent in an alternative formulation, allows different weights to be assigned to different data points. However, assigning such weights reliably is not trivial. In summary, we believe that this approach to measuring the relative scatterer positions will be applicable to only a small fraction of the data of interest.

Reference

- [1] Bevington, P. R., *Data Reduction and Error Analysis for the Physical Sciences*, New York: McGraw-Hill, 1969.

Appendix F: Modifications to the TSA for Interactive Two-Dimensional Analysis

The material in this appendix treats an adaptation of the two-dimensional TSA for manual processing. To make the appendix self-contained, we repeat some of the figures and discussions of Chapter 1.

Since complex-image analysis amounts to an expert system approach where one must first analyze real data to obtain an insight and develop a processing approach, and then in a separate step automate the algorithms, the TSA is used in two ways: for interactive analysis and fully automated processing. The two-dimensional TSA described in Section 1.3 is too unwieldy for interactive processing, because it involves 18 image cuts in which scatterer measurements must be made. For this reason, we utilize a modified version in which the two-dimensional positions of the two scatterers are derived from one fixed-range and one fixed-crossrange cut through the response peak to be analyzed, possibly augmented by a few additional fixed-range and fixed-crossrange cuts in the vicinity of the response peak. For real data, this modified version is less accurate than that involving diagonal cuts, but it is often good enough for the purposes of interactive processing.

For a summary discussion of the procedure that will be illustrated in detail below, assume that an image response was generated by two resolvable

scatterers, meaning that the distance between the two scatterer positions in the image plane exceeds one gate (but is not so large that the scatterers generate two resolvable intensity responses). We also assume that the line connecting the two scatterer positions has an orientation other than along the range or Doppler axis, which would be the case of one-dimensional resolution. A fixed-range cut through the response peak then will give an amplitude/phase pattern indicative of two scatterers, and so will a fixed-crossrange cut. From the fixed-range cut and its amplitude/phase pattern we find two crossrange positions, and from the fixed-crossrange cut we obtain two range positions. In the case of the minimum resolvable separation of one gate, the separations in the fixed-range and fixed-crossrange cuts will be less than one gate, so that some superresolution is needed to obtain the scatterer positions. This is somewhat worse than for the truly two-dimensional TSA, where in the limiting case some degree of superresolution is needed to obtain part of a curve rather than a single measurement point. However, with fixed-range and fixed-crossrange image cuts, for one of the two image cuts the projected scatterer separation will be relatively small. This performance degradation is to be expected when only two image cuts are taken.

Aside from this limiting situation, suppose that the two image cuts yield two scatterer positions in range and in crossrange. We now have the problem of correctly associating one of the range positions with the appropriate crossrange position, so that no "ghost" scatterer is generated. With ideal point scatterers, this can be done on the basis of the strengths of the scatterers measured in the two image cuts. If the scatterers are so closely spaced that the strength measurement is unreliable, we can take image cuts at a slightly different range and crossrange and determine how rapidly the measured scatterer strengths decrease. This indicates how close the real scatterers are, so that the correct association can be distinguished from the incorrect.

Figure F.1 shows a fixed-range image cut through the intensity peak of a simulated response composed of two interfering fixed point scatterers. For such a fixed-range cut, the data points vary in crossrange. The image domain abscissa is labeled with two scales, one giving crossrange in the image, and the "relative" scale being an arbitrary translation of the first, here set to zero at the peak location. The transform of such a fixed-range cut is the signal corresponding to the response peak, so the transform abscissa is labeled "Relative Time." The dashed curve is a superimposed response from a fixed point scatterer, for comparison.

The image cut itself displays a widened amplitude peak with a bulge on its right side, accompanied by a curved phase function. Rather than judge the amplitude peak width by eye, we always consider the normalized half-power




Figure F.1 Fixed-range image cut through response from two scatterers.

width of the peak, which in this case is 1.45. The small variation of the image phase, over just about 0.1 cycles, shows that the two scatterers contributing to the response are generally in phase. Based on interpreting the image data, we would be hard pressed to decide if this response corresponds to two interfering pointlike scatterers or to one shifting scatterer. The decision is much easier with the transform data, which is a prototypical example of a two-scatterer pattern (sinusoidally varying power, with phase jumps at the times of amplitude minima and linear phase at the times of amplitude maxima).

If we measure the times of the amplitude minima (which correspond to the times of the phase jumps) and the size of the phase jump, the scatterer positions can be calculated (by the process described in Appendix A). The dashed vertical lines of Figure F.1 show that the amplitude minima occur at relative times -0.44 and 0.29 . The phase jump appears to be 0.30 cycles, between the dashed horizontal lines tangent to the transform phase curve. However, as discussed in section 1.3.2.1 and Appendix A, we must subtract a linear fit to the transform phase before we measure the size of the phase jump. Doing so gives a phase jump of 0.25 cycles. With the phase jump relative times of -0.44 and 0.29 , the one-dimensional TSA (as given in Appendix A) yields scatterer positions of Crossrange Gates -0.05 and 1.33 .

Figure F.2 shows the fixed-crossrange cut through the same two-dimensional image intensity peak as used for the fixed-range cut of Figure F.1.

For the fixed-crossrange cut the data points vary in range, and the transform is the spectrum corresponding to the response peak, so its abscissa is labeled "Relative Frequency." The response in this image cut has a normalized half-power width of 1.02 and a nearly linear phase. The transform amplitude changes very slowly with frequency, and the transform phase is linear. This corresponds to a single scatterer, or to two scatterers that are far from resolvable in range. Thus, we must initially estimate a single range at the peak position in the cut, Range Gate 0.11.

Two questions now arise. First, as we measured two crossrange positions and one range position, do we have one or two scatterers? In other words, did we mistakenly measure an extra scatterer, or are two scatterers unresolved in one of the dimensions? In this particular case, the two-scatterer pattern of Figure F.1 is so good (it is a prototypical pattern, by construction of the simulation) that there can be no doubt that two scatterers are involved. In the more general case, interference from additional scatterers and the non-point nature of man-made features will produce poorer two-scatterer patterns. However, we can judge the deviation of each pattern from the ideal, and we have adopted the policy of assigning two scatterers to a cut only when the deviation is small enough that the number of scatterers cannot be questioned.

Given that we have two scatterers, we have a second question as to where they are located. Can we improve our initial range and crossrange




Figure F.2 Fixed-crossrange image cut through response from two scatterers.

estimates, especially in the dimension where they were initially unresolved? In order to do so, we must be able to perform more accurate measurements in additional image cuts. This is usually doable, but may not be possible when interference from other scatterers is strong. Consider our example. We began with two image cuts through a response peak at Crossrange Gate 0.27 and Range Gate 0.11, and found two Crossrange positions of Gates -0.05 and 1.33 . If we now take additional fixed-crossrange image cuts at our measured Gates -0.05 and 1.33 , we should obtain more accurate range positions than we did from our initial fixed-crossrange cut through the two-dimensional intensity peak. This is because the contribution from the scatterer at Crossrange Gate -0.05 is greatly reduced in a measurement in Crossrange Gate 1.33 , and vice versa.

Figure F.3 shows the positioning of these and other image cuts used to refine the measurements in the initial cuts. Figure F.1 corresponds to the image cut along the horizontal line of Figure F.3, through the two-dimensional amplitude peak at Crossrange Gate 0.27 and Range Gate 0.11. The two crossrange positions found from Figure F.1, Gates -0.05 and 1.33 ,




Figure F.3 Refinements of measurements from Figures F.1 and F.2.

are represented by open circles (one circle is partially obscured by other symbols). Figure F.2 corresponds to the image cut along the central vertical line of Figure F.3, also through the two-dimensional peak. The single range position found from Figure F.2, Gate 0.11, is represented by the open hexagon.

The additional fixed-crossrange cuts at the measured crossrange positions are indicated by the vertical lines labeled "2nd estimates" and passing through the open circles. The image cuts in Crossrange Gate -0.05 and 1.33 give point-scatterer-like responses in Range Gates 0.05 and 0.47 , respectively. These are denoted by diamonds in Figure F.3, and are significant improvements upon the initial estimate of a single range of Range Gate 0.11 . As stated above, the improvement is due to performing measurements in cuts where one or the other scatterer is dominant, rather than in the initial cut through the response peak.

We can further improve our range measurements if we increase the scatterer dominance even more, by using crossrange gates somewhat less than -0.05 and somewhat larger than 1.33 , if interference conditions allow. Such cuts are indicated in Figure F.3 by the vertical lines labeled "3rd estimates," located in Crossrange Gates -0.55 and 1.83 . These cuts give point-scatterer-like responses at Range Gates 0.00 and 0.58 , respectively, denoted by squares in Figure F.3. The actual scatterer positions are denoted by the filled circles. As the figure shows, the refined ranges differ insignificantly from the actual ranges.

If we wished to refine our crossrange estimates, we could take additional fixed-range image cuts at our new range estimates. We could continue alternately refining range and crossrange positions until they converged. However, this additional refinement produces insignificant changes. Based on our initial two crossrange measurements and our refinement of the single range measurement, we have found two scatterers, one at Crossrange Gate -0.05 and Range Gate 0.00 , the other at Crossrange Gate 1.33 and Range Gate 0.58 . The scatterers are actually located at Crossrange Gate 0.00 and Range Gate 0.00 , and at Crossrange Gate 1.40 and Range Gate 0.60 .

We assign each measurement an uncertainty based on the degree of agreement of the measurement and the one-scatterer or two-scatterer pattern, but with a minimum uncertainty of 0.2 gates (which applies for our example). Thus, although additional refinements can improve the measurement accuracy, the measurements are already well within the estimated uncertainties. A more complicated situation arises if we find two crossrange positions in our initial fixed-range cut, plus two range positions in our initial fixed-crossrange cut. We must then associate the four measurements into two two-dimensional positions. This is illustrated below, for a simulated response

composed of two point scatterers, one located at Crossrange Gate 0.00 and Range Gate 0.00, the other at Crossrange Gate 1.10 and Range Gate 0.90. Figure F.4 shows part of the association sequence. The cuts through the two-dimensional peak, each labeled as an “initial cut,” have each produced two scatterer locations (open circles and hexagons). We must decide which range measurement (hexagon) corresponds to which crossrange measurement (circle). In order to make the decision, we will examine additional fixed-range cuts, shown in Figure F.4, as well as additional fixed-crossrange cuts.

The fixed-range cut through the two-dimensional response peak, at Crossrange Gate 0.18 and Range Gate 0.14, is shown in Figure F.5. The image response has a normalized half-power width of 1.21, with the right side of the amplitude response wider than the left, and a curved phase function. The small amount of phase curvature is consistent with two interfering point scatterers that are generally in phase, as well as with a single scatterer with a shifting phase center. The transform shows that the response corresponds to two interfering scatterers. The transform power has a sinusoidal variation, with the transform phase exhibiting a phase jump when the transform amplitude passes through its minimum, near relative time 0.32.




Figure F.4 Association for two positions in each initial cut.




Figure F.5 Fixed-range image cut through simulated response of Figure F.4.

We observe only one phase jump in the transform, and are suspicious about the null in the transform amplitude at relative time -0.48 , because of the smoothness of the accompanying phase function and imperfections in deweighting the windowed image response. However, the transform amplitude is very symmetric about its peak near relative time 0.13 . We can use this symmetry to extrapolate a second amplitude null (and corresponding phase jump); we substitute a symmetrical function about the peak at -0.13 , mirroring the section from -0.13 to 0.33 in order to find the true position of the null at -0.59 .

The dashed horizontal lines of Figure F.5 show the size of the phase jump, based on the displayed transform phase function. However, as discussed in Section 1.3 and Appendix A, we must subtract a linear fit to the phase function before measuring the jump. Doing so gives a phase jump of 0.17 cycles. With the jump times of -0.59 and 0.33 , the TSA algorithm of Appendix A gives crossrange positions of Crossrange Gates -0.01 and 1.06 , with corresponding amplitudes of 0.36 and 0.18 .

Figure F.6 shows the fixed-crossrange cut through the two-dimensional response peak, in Crossrange Gate 0.18 . The response amplitude has a normalized half-power width of 1.11 , with a slight asymmetry, and the response phase has a small curvature. As with Figure F.5, these are consistent with two interfering fixed point scatterers or one scatterer with a shifting phase center.

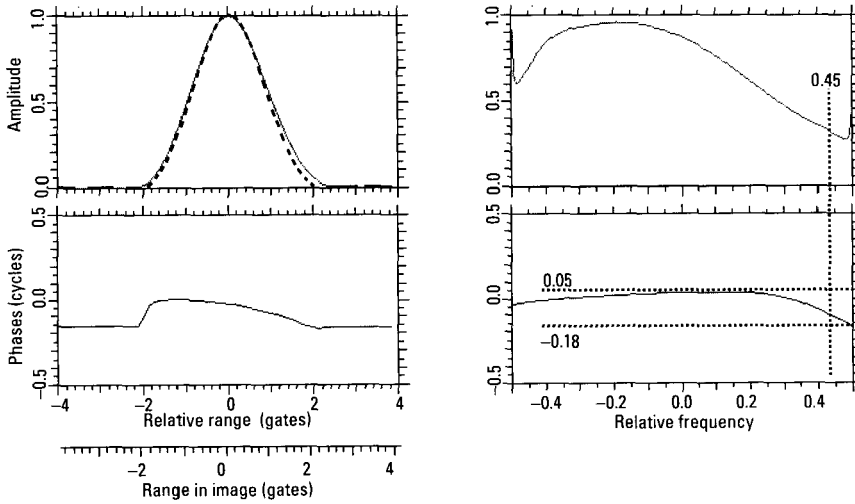


Figure F.6 Fixed crossrange image cut through simulated response of Figure F.4.

The transform amplitude shows a large variation, while the transform phase does not, except near relative frequency 0.5, where the amplitude reaches its minimum. This is inconsistent with a single scatterer with a shifting phase center, even for a truly dispersive (frequency-dependent amplitude) scatterer [1, 2]. However, it agrees with the two-scatterer interference pattern.

Position measurements on the transform shown in Figure F.6 require more extrapolation than those for Figure F.5. Figure F.6 barely contains a transform phase jump and accompanying amplitude minimum. This points out the limits of our extrapolation. The transform must contain an amplitude minimum and an amplitude maximum. This implies that, for favorable interference conditions, we can measure positions of scatterers separated by slightly more than half a resolution cell but no closer. Taking the relative frequency of the phase jump (or amplitude minimum) to be 0.45, and using symmetry about the amplitude peak at relative frequency -0.19 , gives an extrapolated phase jump at relative frequency -0.83 . With a phase jump (after subtracting the linear fit to the phase) of 0.20 cycles, we find initial scatterer ranges of Range Gate -0.04 and Range Gate 0.73 , with respective amplitudes of 0.34 and 0.20.

We must now associate these ranges with the crossrange measurements we made on the image cut of Figure F.5. One possible association method is on the basis of amplitude. The image cut of Figure F.5 gave amplitudes of

0.36 and 0.18, while the image cut of Figure F.6 gave amplitudes of 0.34 and 0.20. One might think that it is safe to associate the positions corresponding to amplitudes 0.36 and 0.34, and those corresponding to amplitudes 0.18 and 0.20. While this produces the correct association for this example, it cannot be used in general. In the general case, neither initial image cut will pass directly through either of the two unresolved scatterers. Thus, the relative contributions of the two scatterers will differ for the two cuts. If the two interfering scatterers have comparable amplitudes, especially if their separations in range and crossrange are not similar, one may be dominant in the fixed-range cut, while the other is dominant in the fixed-crossrange cut.

Having concluded that we cannot perform the association on the basis of the amplitude measurements in the initial cuts, we are left with the necessity of using additional image cuts. These additional cuts can also be used to refine the scatterer positions. As was the case with our first example, Figures F.1 through F.3, we wish to examine cuts in which one or the other scatterer is strongly dominant. This implies improving on our initial cuts, which were between the scatterers, by taking cuts at our initial estimates of the scatterer ranges and crossranges. If interference conditions permit, we can do even better by displacing our cuts from those initial estimates, in order to reduce the relative contribution of the further scatterer.

We begin with cuts at fixed range. The cut at our initial estimate of Range Gate -0.04 (the lower cut labeled "First Peak Location" in Figure F.4) gives an amplitude peak at Crossrange 0.13, with a normalized half-power width of 1.14. The peak is denoted by an up-arrow in Figure F.4. The cut at our initial estimate of Range Gate 0.73 (the upper cut labeled "First Peak Location" in Figure F.4) gives an amplitude peak at Crossrange 0.61, with a normalized half-power width of 1.37, also denoted by an up-arrow. Both these half-power widths imply that we cannot use the peak location as a scatterer location, but must apply the TSA. However, the peak location in each cut should be closer to the crossrange position of the nearby scatterer than to that of the more distant scatterer.

Graphically, the crossrange of each up-arrow should be close to the crossrange of just one open circle, and the two up-arrows should correspond to different open circles. The crossrange measurement of 0.13 from Range Gate -0.04 is much closer to the initial crossrange estimate of Crossrange Gate -0.01 than the initial estimate of Crossrange Gate 1.06, but the crossrange measurement of 0.61 from Range Gate 0.73 is nearly midway between the two initial crossrange estimates. This is likely an indication that our initial range estimate of Range Gate 0.73 is inaccurate.

We consider additional cuts at fixed range, displaced from our initial range estimates to reduce the relative contribution of the further scatterer. These cuts are labeled “Second Peak Location” and “Third Peak Location.” A cut in Range Gate -0.54 gives a peak at Crossrange 0.04 , with a normalized half-power width of 1.02 . Because this and the skewness are close enough to unity (within 6% and 12%, respectively) that the peak position is an accurate measure of the scatterer position, the measurement is denoted by a square in Figure F.4. A cut in Range Gate 1.23 gives a peak at Crossrange 0.96 , with a normalized half-power width of 1.15 , denoted by an up-arrow. A cut in Range Gate 1.73 gives a peak at Crossrange 1.09 , with a normalized half-power width of 0.98 , denoted by a square. These cuts show clearly that we should associate the initial estimates of Crossrange Gate -0.01 and Range Gate -0.04 , and those of Crossrange Gate 1.06 and Range Gate 0.73 . Furthermore, we can refine the crossrange measurements to Crossrange Gates 0.04 and 1.09 . In real data, for which interference conditions are not so benign, refinement will rely on the TSA, not measurements of peak locations.

We could have performed the association on the basis of additional fixed crossrange cuts, rather than additional fixed range cuts. In practice, both must be utilized. Interference conditions may be such that additional cuts are useful in only one dimension. Figure F.7 shows the association sequence using fixed-crossrange cuts. A fixed-crossrange cut at our initial estimate of Crossrange Gate -0.01 (labeled “First Peak Location”) gives an amplitude peak at Range 0.10 , with a normalized half-power width of 1.07 , denoted by an up-arrow. A cut at our initial estimate of Crossrange Gate 1.06 (with the same label) gives a peak at Range 0.67 , with a normalized half-power width of 1.15 , also denoted by an up-arrow. While the TSA would be required in both cuts in order to refine the range measurements, the peak locations are sufficiently different from one another, and sufficiently close to the initially estimated ranges, that we can perform the association. Graphically, the up-arrows and the hexagons can easily be paired.

We can refine the range positions by applying the TSA in these cuts, or by examining additional cuts displaced from our initial estimates. A cut at Crossrange -0.51 (labeled “Second Peak Location”) gives a peak at Range 0.02 , with a normalized half-power width of 1.00 , so denoted by a square. A cut at Crossrange 1.56 (with the same label) gives a peak at Range 0.85 , with a normalized half-power width of 1.03 , also denoted by a square. Thus, our refined range estimates are Range Gates 0.02 and 0.85 .

Figure F.8 shows the combination of the refined range and crossrange estimates. This gives estimated scatterer positions of Crossrange Gate -0.04




Figure F.7 Association sequence using fixed-crossrange cuts.

and Range Gate 0.02, and of Crossrange Gate 1.09 and Range Gate 0.85. These are in excellent agreement with the actual scatterer positions, as shown by the agreement of the triangles and filled circles in Figure F.8.

In this appendix, we have so far illustrated the application of the TSA to the ideal case of two interfering fixed point scatterers. Real data invariably contain noise and interference from additional scatterers. This necessitates windowing the image data before taking a Fourier transform and applying the TSA. We must choose the window to minimize noise and interference contributions to the transform, without excluding too much of the two scatterers of interest. Such windowing is also necessary for more complicated simulations, such as three interfering fixed point scatterers. Section 1.3.3 discusses how to choose window positions, and the distortions they cause to two-scatterer patterns. However, this affects only measurements in the one dimensional cuts, not their association.

We note that under certain rare interference conditions, three fixed point scatterers can generate a single image response peak such that two different window choices allow TSA measurement of all three positions. In such cases, one window contains contributions primarily from the leftmost and




Figure F.8 Combination of refinements.

center scatterers, and the other contains contributions primarily from the rightmost and center scatterers. The association and refinement procedure when an image cut yields three scatterer positions is similar to that described above: take additional image cuts and utilize those that provide a simpler interpretation. However, this occurs too rarely to merit detailed illustration.

The association procedure is the same for simulated and real data. The examples above were particularly simple, allowing the use of image cuts displaced by large amounts from the initial cuts. Image cuts in real data (and more complicated simulated data) cannot generally be displaced from the initial cuts by more than about the half-power half-width of a point scatterer response (0.65 gates for Hamming weighting). Also, although peak positions in the displaced cuts suffice for association of range and crossrange measurements, refinement of the measurements sometimes requires that the one-dimensional TSA be employed. Fortunately, because we are only concerned with refining the stronger scatterer in the cut, we do not need to consider even more additional cuts.

In the remainder of this appendix, we demonstrate the measurement of scatterer positions in real data, using just fixed-range and fixed-crossrange

image cuts. The data are of a motion compensated flying aircraft, with a range resolution of about 0.3m and a crossrange resolution of about 1m. Figure F.9 shows an overview of the image cuts and measurement results that we will present. Horizontal and vertical lines represent image cuts, symbols represent measured positions, and the light background curves are amplitude contours spaced by 3 dB. In the initial fixed-range cut through the response peak, we will find two scatterer positions, denoted by circles. In the initial fixed-crossrange cut, we will find one position, denoted by a hexagon. In order to refine the ranges of the two scatterers, we will examine additional cuts as indicated.

Figure F.10 shows the fixed-range image cut through the two-dimensional intensity image response peak. The response of interest, near Relative Crossrange 0.0, is not fully resolved from another response near Relative Crossrange 2.4. The normalized half-power width of the response of interest is 1.42, far too high for a single scatterer. This response is likely composed of two scatterers, with its interpretation made complicated by the presence of the second response. As discussed in Section 1.3, we set transform




Figure F.9 Overview of image cuts and measurement results.




Figure F.10 Fixed-range image cut through aircraft response.

window boundaries near phase inflection points and amplitude minima, then evaluate the degree to which each transform fits a two-scatterer pattern.

Figure F.11 shows the best window position and the resulting transform. The figure also includes curves labeled by “m” (for model) which are the result of extracting the two scatterer positions, amplitudes, and phases via the TSA, generating the corresponding image cut, applying the same windowing as done to the real data, and taking the Fourier transform. The modeled amplitude and phase are in excellent agreement with the real data. The amplitudes differ only near the edges of the transform, where deweighting amplifies contributions of the third scatterer to the real data. As discussed in Section 1.3, distortions are expected at these edges. The primary difference between the phases is a constant offset, which is insignificant. The TSA gives initial scatterer positions of Crossrange -10.90 and Crossrange -9.70 . These positions must be, and were, verified by varying the window position as described above.

Figure F.12 shows the fixed-crossrange cut through the response. The response amplitude has a slight bulge on its right-hand side, and the phase is curved. The response’s normalized half-power width is 1.02 and its skewness (ratio of right to left half-power half-widths) is 1.11. The skewness is nearly large enough to allow measuring the positions of two interfering scatterers, but not quite. A transform of the interval between the amplitude minima bounding the response does not allow the extraction of two scatterer




Figure F.11 Transform and two-scatterer pattern for Figure F.10.




Figure F.12 Fixed-crossrange cut through aircraft response.

positions. We use the peak location, Range 0.08, as an initial estimate of our scatterer location(s).




Figure F.13 Fixed-crossrange cut for refinement of first scatterer range.

With two crossrange positions and one range position, we likely have two scatterers that are poorly resolved in range. To refine their positions, we take additional fixed-crossrange cuts, beginning with cuts at our initial crossrange estimates. Figure F.13 shows the image cut in Crossrange Gate 10.9. The response has a normalized half-power width of 0.97 and a skewness of 1.03. Although the bulge on the right side of the response indicates that the response is composed of at least two scatterers, we cannot resolve them. We must rely on the response peak location, Range -0.06 .

In the beginning of this appendix, we were able to refine our measurements on simulated data by displacing our image cuts so as to reduce the relative contribution from the further scatterer. Figure F.14 shows the image cut in Crossrange Gate -11.4 . Unlike the simulated data, displacing the cut has not improved the situation. Interference with the response of interest is comparable in the two cuts. The normalized half-power width of the response is 0.93, which barely justifies attempting a two-scatterer analysis. However, no good two-scatterer pattern is obtainable. The peak location in this cut is Range -0.20 . We take the range position of the scatterer as the average of this and the previous cut, giving Range -0.13 .

Figure F.15 shows the image cut at our initially estimated crossrange for the second scatterer, Crossrange -9.70 . The response of interest, at Range 0.24, has a normalized half-power width of 1.03 and a skewness of 1.08. We must rely on the peak location. Figure F.16 shows a displaced image cut,




Figure F.14 Displaced fixed-crossrange cut for first scatterer.




Figure F.15 Fixed-crossrange cut for refinement of second scatterer range.

at Crossrange -9.20 . Interference conditions have clearly worsened, and we must rely on the image cut at Crossrange -9.70 .

The two refined range positions, Range -0.13 and Range 0.24 , are so close to one another that we are unlikely to be able to refine the crossrange

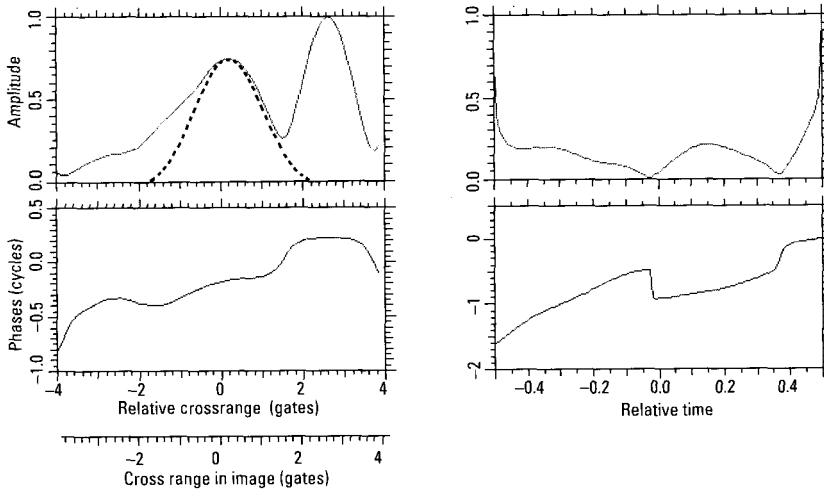


Figure F.16 Displaced fixed-crossrange cut for second scatterer.

positions by examining additional fixed-range image cuts. We accept two scatterers, one located at Crossrange Gate -10.90 and Range Gate -0.13 , and the second at Crossrange Gate -9.70 and Range Gate 0.24 . While we would like to refine these positions further, and can do so with diagonal image cuts (see Section 1.3), we have used the fixed-range and fixed-crossrange cuts to determine the by far more important fact that the response consists of two interfering scatterers. For the purpose of identification, the primary goal in analyzing a response is to determine the number of scatterers it comprises, plus their positions and reliable uncertainty estimates for those positions. Reducing inaccuracies and uncertainties to, say, 0.2 resolution cells is a desirable secondary goal, but not nearly so important as the primary. In most identification applications, such small uncertainties are not necessary for all scatterers.

References

- [1] Rihaczek, A. W., and S. J. Hershkowitz, "Man-Made Target Backscattering Behavior: Applicability of Conventional Radar Resolution Theory," *IEEE Trans. on Aerospace and Electronic Systems*, Vol. 32, No. 1, April 1996, pp. 809–824.
- [2] Rihaczek, A. W., and S. J. Hershkowitz, *Radar Resolution and Complex-Image Analysis*, Norwood, MA: Artech House, 1996.

Appendix G: Tracking Interfering Scatterers by Measuring Phase Slopes

G.1 Introduction

When two interfering scatterers have variable relative Doppler, which is typical for target tracking, it is impossible to motion compensate both scatterers with a single range history. For precision measurements, one should first compensate the stronger scatterer and perform the desired measurements on it. One should then do the following: suppress the focused response of the stronger scatterer, compensate the motion of the second scatterer, and perform measurements on it. Figure G.1 shows a signal generated by a sounding rocket performing an attitude maneuver. The changing spacing of the amplitude peaks indicates these data contain two scatterers with variable relative Doppler.

As discussed in Chapter 6, there are two ways to compensate the stronger scatterer. The first, demonstrated in detail in that chapter, is to remove the phase jumps at the times of the amplitude minima, fit a spline to the resultant phase function, and subtract the spline from the phase of the original signal. The second is to measure the phase slopes near the times of peak amplitudes, which gives the Dopplers of the strong scatterer at those times, then fit a spline function to the Doppler measurements, integrate the




Figure G.1 Signal from a sounding rocket performing an attitude maneuver.

Doppler function to get phase, and subtract the phase from that of the original signal. This appendix demonstrates the second procedure, by applying it to data intervals processed with the first procedure in Chapter 6.

Both procedures assume that the measured phase slope gives the Doppler of the stronger scatterer. This gives a significant inaccuracy when the two scatterer strengths are comparable. The weaker scatterer modifies the phase slope depending on the relative locations of the two scatterers. Thus we must incorporate information from the signal amplitude in order to correctly derive the stronger scatterer's Doppler from the signal's phase slope. This correction will not be addressed in this appendix.

G.2 Choosing Intervals for Phase-Slope Measurement

Figure G.2 shows Figure G.1 with automatically selected intervals of linear phase and strong amplitude. Each selected interval is bounded in time by left and right vertical crosshairs, and displays a linear fit to the phase function. The standard and maximum deviations of the fit are printed vertically within




Figure G.2 Figure G.1 with automatically selected intervals of linear phase.

the interval, and the slope of the fit is printed in the right margin of the plot. The remainder of this section gives details of the algorithm we have automated for interval selection, and the data analysis continues in Section G.3. We stress that this algorithm is preliminary, and must be refined through testing on additional data. More fundamentally, we have yet to develop the procedure to account for the effect of a second scatterer on the relation between a measured phase slope and the Doppler of a dominant scatterer.

The signal from a single point scatterer has constant amplitude, and its phase gives its range history. Equivalently, its phase slope gives its Doppler history. If, instead of a single scatterer, we have a dominant scatterer plus interference, the phase slopes at the times of the amplitude peaks approximately give the Doppler history of the dominant scatterer. We therefore desire to measure phase slopes at times when the amplitude is approximately constant or near the peak of the amplitude modulation cycle.

At these times, we search for intervals of linear phase. These intervals must be of long enough duration to give accurate slope measurements, but not so long that the phase slope changes appreciably within an interval. The

intervals must allow small phase fluctuations due to noise or very weak interfering scatterers, but not large enough fluctuations to corrupt the phase-slope measurement.

Given a data interval (Figure G.1 in this case), we first estimate the signal-to-noise ratio (SNR), in order to determine what level of phase fluctuations should be ignored. We use the median signal amplitude as an estimate of the amplitude of the strongest scatterer in the signal. We use the median amplitude of an FFT of the signal as the noise estimate. This assumes that the target does not occupy a large fraction of the data's Doppler spread, which should be valid for the applications of interest. These measurements allow us to estimate the expected standard deviation of the phase from noise as $1 / \sqrt{2SNR}$.

We divide the data into linear phase segments by starting with the first point, adding points to our potential segment until it is either accepted or rejected, then starting over from a point beyond the accepted or rejected segment. Each time we add a point to a potential segment we perform two tests. First, if the added point is in the seventh phase maximum within the interval, we accept the segment. This many phase fluctuations, with a small deviation, indicate that we have a long enough interval for accurate phase-slope measurement. Second, we perform a linear-least-squares fit to the phase of the potential segment. If the maximum deviation of the fit is greater than both three times the expected noise standard deviation and 0.05 cycles, we stop adding points to the segment and apply additional tests to decide whether to accept or reject it.

As the last added point caused the deviation to cross the threshold of acceptability, we reduce the segment size by dropping points off either end and redoing linear fits until the maximum deviation is less than both two times the expected noise standard deviation and 0.033 cycles. Next we check whether the remaining maximum deviation is due to phase fluctuations or phase curvature. We fit a quadratic curve to the reduced segment with the linear fit removed, then measure the phase difference between the curve at the first point of the segment and at the extremal point of the curve. If this is larger than one expected noise standard deviation, then the phase is curved, and we further reduce the segment to the interval from the first point to the extremal point.

If the reduced segment has fewer than seven points we discard it as too short. If the corresponding amplitude is not approximately constant, or if a quadratic fit to the amplitude does not have a maximum within the segment, we discard the segment as not representative of a scatterer Doppler. If the uncertainty in the phase slope, calculated by error propagation from the

least-squares fit, is greater than one Doppler resolution cell, we also discard the segment as yielding an unreliable measurement.

Having accepted or rejected the segment, we start evaluating a new potential segment. However, we must decide where to start our evaluation. This depends on whether and how the previous segment was accepted or rejected. If the previous segment was accepted on the basis of containing seven fluctuation cycles, or if it was reduced in size by the quadratic test, we start our new potential segment at the point immediately after the end of the previous segment. If the previous segment was reduced in size because of too large a deviation, and that deviation was due to phase fluctuation rather than phase curvature, then there is a good likelihood that the points dropped off are part of a phase jump. In this case, we do not want to start our new segment with these points. We begin our evaluation a few points beyond these.

G.3 Data Analysis

Figure G.3 shows a smooth curve that is the result of integrating a polynomial fit to the measured phase slopes. Note that the slope of the curve is in good agreement with the measured slopes. Thus, it appears that we have successfully extracted the motion of the stronger scatterer. If we ignore the linear component of the extracted motion, we see that there is a residual modulation of about 0.3 cycles (bottom to top). This would seriously smear the responses in the transform domain.

Figure G.4 shows the signal after the residual modulation of the curve of Figure G.3 has been used for compensation. If we had taken the effect of the second scatterer on the measured phase slope into account, the transform of the compensated signal would yield a compressed main response. Since we ignored the effect of the second scatterer, and the periodicity of the amplitude modulation changes throughout the signal (indicating a changing scatterer separation), a transform over the entire signal will not give a properly compressed response. However, a transform over a section when the amplitude modulation does not change much should give good compression, albeit with less nominal resolution than a transform of the entire signal.

The crosshairs in Figure G.4 show a section where the amplitude modulation does not change much, and Figure G.5 shows the transform of this section. The half-power width of the response, relative to that for a perfectly focused response, is 1.03 (the width in Hz of the focused response is inversely proportional to the transform interval in the signal). The phase of the response is fairly linear. Thus, we have succeeded in focusing the main




Figure G.3 Integrated fit to measured slopes.




Figure G.4 Compensated signal.




Figure G.5 Transform of the corrected data.

response. If we wished to refine the compensation, we could notch out the interference, take a transform, and measure the residual motion.

Rather than refining the motion of the stronger response, we demonstrate that we can focus the weaker response. The spread amplitude and curved phase of the weaker response in Figure G.5 indicate that it is not yet focused. Figure G.6 shows Figure G.5 with the main response notched out. Figure G.7 shows the transform of Figure G.6, along with a polynomial fit to its phase. Note that the signal duration is that of the interval chosen in Figure G.4. Figure G.8 shows the transform of Figure G.7, after the fit is subtracted from the phase. The phase of the response is linear at the location of the amplitude peak and the relative half-power width of the response is 1.11. This is commensurate with interference from a third scatterer located near 17 Hz, as evidenced by the extended amplitude tail and the flat phase function.

Figure G.9 shows the signal from 448 to 449 seconds. Figure G.10 shows the automatically selected intervals of linear phase and strong amplitude. Figure G.11 shows the smooth curve that is the result of integrating a polynomial fit to the measured phase slopes. As in Figure G.3, the slope of the curve is in good agreement with the measured slopes. After subtracting




Figure G.6 Figure G.5 with main response removed.




Figure G.7 Transform of Figure G.6




Figure G.8 Transform of compensated weaker response.




Figure G.9 Signal from 448 to 449 seconds.




Figure G.10 Automatically selected intervals of linear phase.




Figure G.11 Integrated fit to measured slopes.




Figure G.12 Transform of compensated signal.

the nonlinear component of the curve from the signal phase, then taking a transform, we obtain Figure G.12.

The compensated response lies between the crosshairs, and it is a response from two interfering scatterers. The piecewise linear phase function indicates that the response is truly from two interfering scatterers, rather than the product of smearing due to an inadequate compensation. The transform over the window in Figure G.12 gives the signal in Figure G.13. The signal after 448.15 seconds gives an excellent approximation of the interference pattern of two scatterers with constant but different Dopplers. The strong peak near 448.0 seconds is generated by a combination of a changing scatterer amplitude and imperfections in deweighting.

Next, we suppress the signal within the crosshairs in Figure G.12 and take the transform. The result is shown in Figure G.14, with a polynomial fit to the phase. Since the amplitude lacks deep minima (with an exception near 448.9 seconds), this fit represents scatterer motion rather than interference effects. Using the curve for compensation, then taking the transform, gives Figure G.15. The response phase is linear and the relative half-power width is 1.03. This response is well compensated.




Figure G.13 Transform over the window in Figure G.12.




Figure G.14 Transform of Figure G.12 with the indicated window notched out.




Figure G.15 Compensated weaker scatterer.

G.4 Summary

By measuring phase slopes at times of linear phase and strong amplitude, one can extract the motion of a dominant scatterer from a signal generated by interference among scatterers of variable relative Doppler. Having extracted that motion, one can focus that scatterer, suppress it, and repeat the procedure with a second scatterer. The intervals for phase-slope measurement can be chosen automatically.

Appendix H: Setting an Amplitude Threshold for Analyzing Image Responses

As discussed in the main text, target identification from SAR/ISAR imagery requires the measurement of scatterer positions and the comparison of those positions to a database composed of the positions of strongly backscattering features for each candidate. The question then arises, from which responses of an image should scatterers be extracted? First, we must determine which responses correspond to the target and which to clutter. The target responses must be further windowed, discarding at least those with insufficient SNR, plus range and crossrange sidelobes of stronger responses. We may also discard responses based on their locations relative to the rest of the target. For instance, as discussed in Chapter 3, for aircraft imagery we discard delayed duct responses, returns from rotating devices, and responses in the *far half* of the aircraft. This still leaves a set of responses of widely varying strengths for our consideration.

In order to accurately measure the positions of scatterers composing a response, we must be able to recognize one-scatterer and two-scatterer patterns in the transforms of image cuts through the response. This implies that, if a response is to yield usable scatterer positions, at most three scatterers can contribute significantly to the response, and that most responses must contain significant contributions from just one or two scatterers. If we attempt

to analyze a response composed of too many scatterers, we will extract a single scatterer position, with large uncertainties in range and crossrange. Too many such measurements will degrade identification performance.

There are two sources of this degradation. First, the comparison database for the correct aircraft may not contain a feature corresponding to such a poor measurement, because the database consists of wave-trapping features expected to backscatter strongly enough to stand out from their neighbors, not features composed of several closely spaced scatterers of comparable (and weak) strength. Given incomplete target information from which to construct a database, predicting the existence of a feature of the latter type is difficult. If the database does not contain the feature, its measurement will degrade identification performance. Second, predicting the feature's location is far more difficult than predicting its existence. The location shifts rapidly with target orientation, as the phase relation between the feature's constituents varies, necessitating a large uncertainty in the database. Thus, for the correct candidate, the feature will produce agreement between highly uncertain measured and predicted positions. This may yield a decrease in the relative probability of the correct candidate, compared to another candidate that has a predicted feature with a small uncertainty whose position happens to agree with the measured feature (within its large uncertainty). Additionally, the inclusion in the database of the second type of feature would have another deleterious effect. Each candidate would have features of this type, so that incorrect candidates would benefit from agreement between measurements with small uncertainties and predicted features with large uncertainties, further degrading performance.

From the above, it is clear that we should forgo extracting scatterer positions from responses that consist of too many scatterers (or too much noise). One might think that this could be accomplished by discarding any position measurements that have large uncertainties. However, many of the strong features on man-made targets are extended scatterers with shifting phase centers. These can generate responses that are not good approximations to one or two interfering fixed scatterers, and which also yield position measurements with large uncertainties. We must retain these so that the corresponding feature in the database for the correct aircraft does not go unmatched by a measurement. Another approach to avoiding responses composed of too many scatterers is to set a fixed number of responses to examine per target. However, this also has drawbacks. Any fixed number will force us either to discard usable responses for large targets, include problematic responses for small targets, or both. Furthermore, this approach

obviates the fact that the variation in the number of wave-trapping features from target to target is useful for identification.

In order to avoid the problematic responses, we recall that the strongest features on man-made targets are the wave-trapping features. In our experience, the amplitude distribution of the responses from these features is similar for many highly resolved man-made targets, such as fighter aircraft, ground vehicles, and small ships, in the following sense: if the responses are ordered by amplitude, a few responses (up to five or six) are much stronger than the rest, the next 20 or so have strengths that differ significantly from one another, and the strengths of the remainder form a slowly decreasing continuum. The problematic composite responses generally fall within the continuum, or near its upper boundary. Many spurious responses due to shifting scatterers also fall within the continuum.

Figure H.1 shows an ISAR image of an aircraft, utilizing range and Doppler centroid compensation, with a range resolution of about 1 ft and a Doppler resolution of about 1m. The image is presented in peaks plot format, where each dot represents a local two-dimensional maximum of the oversampled intensity image. The area of each dot is proportional to the amplitude of the corresponding pixel. The fact that most range gates contain only a single strong response and that there are no trails of responses at fixed ranges is an indication that the motion compensation may be acceptable. An examination of fixed-range image cuts through responses verifies that this is the case.

Figure H.2 shows the amplitude distribution of the responses of Figure H.1. The abscissa is peak number in order of descending strength. The ordinate is peak amplitude. Each peak is indicated by its corresponding numeral. For the sake of visibility, the ordinate has been truncated so that the two strongest peaks are not displayed. The distribution displays the behavior discussed above, with the continuum appearing as a nearly horizontal line of responses. The question now is how to recognize the beginning of the continuum, so as to avoid the problematic responses.

The procedure we have adopted is as follows:

1. Perform a linear-least-squares fit to the amplitude distribution as a function of peak number, for those peaks with amplitudes less than four times the background noise (or clutter) level.
2. Excluding the six strongest peaks, find a second line that minimizes the root-mean-square amplitude difference, measuring from each point to the closer of the two lines.




Figure H.1 ISAR image of an aircraft.

3. Find the crossover point between the lines, beyond which points are closer to the first line than to the second.

Figure H.2 shows both lines. The first is close to horizontal, through the continuum, most visible at its left end. The second passes near the strong responses and is close to vertical. Figure H.3 shows a blowup of Figure H.2, near the crossover region. In this case, the crossover point yields 28 responses to be analyzed. Figure H.4 shows the 28 strongest responses retained from Figure H.1.

This procedure finds the approximate transition from dominant wave-trapping features to others. Certainly, problematic features will sometimes be included and dominant wave-trapping features will sometimes be excluded.




Figure H.2 Amplitude distribution for Figure H.1.

However, the number of responses mistakenly retained or discarded has generally been a small fraction of the number retained. In the example shown, we have retained one response more than desirable. The weakest response, at closest range, is probably a return from the radome tip. As discussed in Chapter 3, this is not reliably observable and thus is excluded from the database. However, because of its location at the front of the aircraft and its weak strength relative to the next response in range (the strongest response of the target, presumably from the aircraft radar), it will not affect identification performance.

This procedure begins with a fit to peaks with amplitudes less than four times the background noise (or clutter, if applicable) level, which can be estimated by measuring the median peak level in a region a small distance from




Figure H.3 Blowup of Figure H.2.

the target (a more accurate procedure is given in [1]). As can be seen from Figure H.2, the level need not be very accurate. It is used only to set a threshold that specifies to which part of the tail of the distribution a line is fitted, and using a different part will not much affect the fit. On the other hand, if the threshold is above the tail, the procedure may not work well. However, in this case, we have an insufficient signal-to-background ratio to make accurate measurements on some of the wave-trapping features, so identification is problematic even if we correctly specify responses for analysis. (There is an exception to the foregoing statement: few responses of a stationary ground vehicle will be affected by a strong clutter background. If this is the situation of interest, the procedure should be modified.)

We have stated in the main text that good-quality images of highly resolved (in at least one dimension) man-made targets such as fighter aircraft (fuselages), ground vehicles, and small ships typically have 20 to 30 dominant responses. The procedure described in this appendix quantifies that statement.




Figure H.4 Peaks retained from Figure H.1.

Reference

- [1] Lank, G. W., and N. M. Chung, "CFAR for Homogeneous Part of High-Resolution Imagery," *IEEE Trans. on Aerospace and Electronic Systems*, Vol. 28, No.2, April 1992, pp. 370-382.

About the Authors

Dr. August W. Rihaczek has spent essentially all his professional life on research in the field of radar theory. During the 1960s, he was an active participant in the development of the theory of radar waveform design, publishing many papers and finally a textbook on the subject. This was followed by work on radar system design, which then led to a strong activity in the development and use of radar signal processing methods. It was the application of the theoretical concepts to the processing of real data that revealed fundamental problems with the conventional resolution theory, making the need for a modified approach apparent. He then spent the 1980s and 1990s on the development of the theory and the implementation of a workable method for extracting information from the returns from highly resolved man-made targets, as needed for target identification and related tasks.

Dr. Stephen J. Hershkowitz received a Ph.D. in physics from Stanford University. His early research investigated twisting and bending modes of DNA and radiative transfer through supernova atmospheres. Since joining MARK Resources, Inc. in 1986, he has worked on developing and automating the extension of conventional radar signal processing and image analysis to modern high-resolution radar, concentrating his effort on target detection and identification.

Index

- 180° aspect angle, 273–76
- Adaptive data collection flowchart, 203
- Adaptive processing methods, 159–61
- Aircraft
 - bomber, 282
 - combined photographs, 33
 - commercial, 277–82
 - Doppler of, 124
 - Doppler tracking of, 286
 - elevated tail, 185
 - feature measurement and, 34
 - fighter, 170, 249
 - horizontal stabilizer, 270, 273
 - importance scatterers, 121
 - instrument panels, 196
 - intake position, constraining, 83
 - length, 181, 182–83
 - maneuvering, 229–45
 - motion compensation and, 155
 - photograph, 195
 - pitch, 433–34
 - positional match for, 283–85
 - roll, 433–34
 - scale-modeling, 193–94
 - twin-hull, 184
 - with wing-mounted engines, 83
 - wingspan, 182–83
 - wingspan-to-tailspin ratio, 185
 - yaw, 433–34
 - See also* Jet engine returns; Jet engines
- Aircraft feature database, 188–91
 - aspect angle sectors, 188
 - assembly, 192
 - collection of diagrams/photographs, 192
 - deriving, 191–99
 - positional match at 20° aspect, 189
 - positional match at 26° aspect, 189
 - positional match at 31° aspect, 190
 - positional match at 34° aspect, 190
 - positional match at 41° aspect, 191
 - test range data collection, 193
- Aircraft features, 82, 182–99
 - air inlets/outlets, 188
 - antennas, 187, 196
 - aspect angle change, 188–91
 - common, 186–91
 - elevated aircraft tail, 185
 - engine intake and exhausts
 - position, 184
 - engine intake duct opening size, 184
 - exhaust, 187–88

- Aircraft features (continued)
 - fuselage and leading wing edge
 - angle, 184–85
 - fuselage shape, 185
 - horizontal stabilizer position, 185
 - large intakes, 187
 - length and wingspan, 182–83
 - number of engines, 183–84
 - radar in nose cone, 186–87
 - relative width of wings, 185
 - special, 182–86
 - twin-hull aircraft, 184
 - wave-trapping cavities, 198
 - wing position, 185
 - wingspan-to-tailspan ratio, 185
 - wingtip missiles, 184
- Aircraft identification, 163–290
 - automated, 285–90
 - challenging aspects, 163
 - crossrange resolution requirement, 199
 - dwelt time, 200
 - features for, 182–99
 - with ISAR imagery, 181–82
 - radar waveform significance, 163–68
 - with range profiles, 168–81
 - range resolution and, 35
 - resolution requirements, 199–203
 - time for, 200
 - See also* Target identification
- Aircraft images, 32
 - after motion compensation of first peaks
 - track, 249
 - after range and Doppler centroid
 - compensation, 214
 - after standard compression, 211
 - at aspect angle of 65° , 218
 - at aspect ratio of 80° , 218
 - banked commercial aircraft, 281
 - at broadside, 219
 - commercial, 278–79
 - cut along centerline of fuselage, 274
 - with half imaging time, 225
 - intensity, 17
 - ISAR, 707, 708
 - nose-on, 268
 - over central half of imaging
 - interval, 221
 - peaks plot, with jet engine returns, 166
 - SAR, 176
 - scatterer positions measurement
 - and, 221
 - for short linear FM pulse, 167
 - tail-on, 227, 274
- Aircraft imaging, 203–83
 - horizontal stabilizer, 270
 - illustrations with delayed duct
 - returns, 223–29
 - illustrations without delayed duct
 - returns, 209–23
 - interval selection flowchart, 210
 - large aspect angles, 245–66
 - of large commercial aircraft, 277–82
 - of maneuvering aircraft, 229–39
 - motion compensation, 204–6
 - motion compensation steps, 206
 - principles, 204–9
 - with problematic conditions, 203
 - range profiles near broadside, 262
 - requirements, 204
 - resolution check, 226
 - tip response, 252
 - zero aspect angle, 266–73
 - See also* Imaging
- Aircraft response
 - fixed-crossrange image cut through, 70
 - fixed-range image cut through, 68
 - scatterers derived from diagonal cuts
 - through, 70
- Ambiguity function, 4
 - applying, 89
 - complex, 91
 - resolution performance derived from
 - envelope of, 90
 - as tool of waveform design, 90
 - use of, 90–91
 - value, 91
- Amplitude
 - constant, 63, 77
 - cruise ship, 561
 - distortion, 66
 - distribution, 707, 709
 - extrema, 628
 - image, 19
 - interfering scatterers, 680

- maxima, 54, 76, 109, 419, 609, 673
- minima, 47, 54, 76, 112, 599, 621, 632, 673
- minimum, sharpness, 114
- minimum, threshold, 290
- peaks, 55, 64, 707
- scatterer ratio of, 47
- shape, 66
- spread, 697
- threshold, 705–11
- transform, 57, 156, 208, 288
- variations, 157, 372, 375, 679
- Amplitude function, 234
 - dive boat, 548
 - FFT of, 424
 - modulation pattern, 422
 - moving flatbed truck, 386
 - phase function combined with, 152
 - ship, 468, 494
 - See also* Phase function
- Amplitude modulation, 73, 156, 350, 419, 422, 593
 - interval, 594
 - periodic, 424
 - phase modulation correlated with, 588
 - strong, 595
 - See also* Modulation
- Amplitude/phase pattern, 619
- Antennas, 187, 196
 - location, 196
 - moving recreational vehicle, 402
 - SAR, 267
- Aspect angles
 - 180°, 273–76
 - at center of imaging interval, 259
 - change, 188–91
 - crossrange resolution and, 201
 - defined, 437
 - delayed duct returns and, 268
 - effect on howitzer
 - identification, 306–11
 - flatbed truck, 328–30
 - imaging and, 247
 - large, 29, 245–66
 - maximum, 247
 - motion compensation and, 247
 - moving off-highway truck, 357
 - moving recreational vehicle, 389
 - moving tank, 381
 - off-highway truck, 323–26
 - range resolution and, 260, 469
 - sectors, 188
 - tank, 314–19
 - zero, 266–73
- Association procedure, 682–83
 - with fixed-crossrange cuts, 682
 - for simulated and real data, 683
 - three scatterer positions and, 683
- Attitude maneuver
 - analysis, 590–610
 - sounding rocket performing, 692
 - SWD plot, 591
- Automated aircraft identification, 285–90
 - compensation/selection of imaging interval, 288–89
 - data collection, 285–88
 - image analysis, 289–90
 - See also* Aircraft identification
- Automated moving ground vehicle
 - identification, 428–31
 - final image analysis, 431
 - final image generation, 428–29
 - measurement of range only, 430–31
 - measurement of two-dimensional positions, 429–30
 - survey image formation, 428
- Backscattering, 31
 - behavior of man-made targets, xx, 6–10, 28
 - complicated, models, 82
 - dispersive, 82
 - off-highway trucks, 321
 - strength, 63
- Bandwidth
 - 300-MHz, 171, 247, 255
 - 600-MHz, 173
 - 1200-MHz, 174, 175, 178
 - instantaneous, 168
- Bayes classifier, 83
- Bayesian probabilities, 83, 284, 285
- Beam depression angle, 438–40
 - 0°, 438, 441, 444, 447
 - 20°, 439, 442, 445

- Beam depression angle (continued)
 - 90°, 440, 443, 446
 - value of, 453
- Bending/flexing region, 647
- Bomber aircraft identification, 282
- Bow scatterer, 455
 - coast guard cutter, 531, 556
 - dive boat, 565
 - relative Doppler, 473
 - tracking, 465, 473, 477
 - See also* Ships
- Carrier frequency, 37
 - choice of, 40
 - X-band, 64
- Circular polarizations, 81
- Classification, ship, 579
- Clutter cancellation
 - moving RV, 390–91, 406
 - for stationary platforms, 292
- Coast guard cutter, 498, 504, 523–43
 - automated motion
 - compensation, 555–56
 - bow scatterer, 531, 556
 - broadside, 537
 - combined range/Doppler
 - tracks, 531–32, 538
 - crossrange interval, 524
 - imaging interval, 534, 541
 - imaging interval duration
 - determination, 541–42
 - imaging time, 535, 543
 - large aspect angle, 530
 - length measurement, 556–57
 - multiple scatterers, 527
 - one-second image, 557
 - peaks tracks, 524, 529, 530, 537–38, 539, 540
 - peaks tracks at large aspect angle, 530
 - photograph, 523
 - roll Doppler, 526, 533, 571
 - roll/pitch Doppler, 528
 - in rough seas, 523–43
 - sea-level-view, 528
 - short-term image, 556
 - stern scatterer, 531
 - topview image, 527, 535–36
 - uneven motion of, 535
 - yaw Doppler, 526, 533, 541, 571
 - yaw measurement, 532
 - See also* Ships
- Commercial aircraft
 - banked, image of, 281
 - Doppler-spread engine
 - returns, 279, 280
 - identification of, 282
 - images, 278–79
 - imaging of, 277–82
 - large aspect angles, 279
 - at small aspect angle, 280
 - weaker responses, 277
 - See also* Aircraft
- Compensated scatterers, 139
 - image cut and transform for, 241
 - image cut in range gate of, 372
- Complex-image analysis, 1–4
 - adaptation to SWD processing, 583
 - defined, 12
 - as expert system approach, 671
 - improvement of analysis results, 625
 - measurement methods, 10–28
 - resolution performance, 98
 - smooth flight resolution
 - performance, 618
 - superresolution in, 49
 - zero fill and, 12
- Complex range profiles, 170
- Constant Doppler, 10, 92
- Constant-false-alarm-rate (CFAR)
 - algorithm, 4
- Correlation tracking, 203
- CRISP: Complex Radar Image and Signal Processing—Software and User's Manual*, xvi
- Crossrange
 - estimates, refined, 681
 - extent measurement, 79
 - gates, 255, 258
 - measurements, 117, 680
 - position determination, 431
 - scale, 219, 566
 - separation of superstructure, 475
 - sidelobes, 4
 - smearing, 482

- spread, 249, 490
- width, 411
- Crossrange resolution, 2, 6, 36
 - achievable, 153–54
 - additional, 107
 - for aircraft identification, 199
 - aspect angle and, 201
 - asymmetry of, 104–22
 - avoiding, 107
 - cells, 79
 - cell width, 110
 - checking limit on, 117–21
 - correct, 107
 - degradation, 439
 - high, 79–80, 226
 - imaging interval and, 9, 226
 - implementation, 106
 - inadequate, consequences, 107
 - motion compensation and, 118
 - motion measurement and, 10
 - moving ground vehicles and, 339
 - moving recreational vehicle, 399
 - moving targets and, 121, 122, 124, 153
 - peaks plot image for, 113, 115, 116
 - range interference and, 115
 - range resolution equality, 104
 - range resolution vs., 106
 - real situation of, 110–11
 - requirements, 110–11
 - subordinate role, 121
 - for targets at broadside aspects, 121
 - too high, attempted, 117–18
 - wingtips and, 207
 - See also* Resolution
- Cross section, 80
- Cruise ship, 505–10
 - amplitude function, 561
 - length measurement, 558–59
 - one-second topview, 558, 559
 - peaks tracks, 505
 - phase function, 561
 - photograph, 505
 - roll Doppler, 509
 - stern, 558
 - superstructure, 569
 - yaw Doppler, 507, 509, 560
 - See also* Ships
- Data intervals
 - analysis length, 602
 - phase slopes and, 607
 - validity, 610
- Decks, 475
 - height scatterers above, 570–75
 - scatterer positions on, 569–70
 - shape of, 569–70, 575
 - See also* Ships
- Deformable template matching
 - algorithm, 84–86
- Deformed templates
 - likelihood function, 85
 - overall probability, 86
 - unlikely, 85
- Degradation
 - crossrange resolution, 439
 - resolution, 99
 - sources, 706
- Delayed duct returns, 175–77, 181
 - aircraft imaging with, 223–29
 - aircraft imaging without, 209–23
 - aspect angle and, 268
 - imaging problem, 225
 - recognition, 233
 - resolving, 267
 - scatterer positions, 228
- Destructive interference, 644
- Diagonal cuts, 61–62
 - diagonals derived from, 70
 - results of, 61, 62
- Differential Doppler, 43, 45
 - modulation period governed by, 46
 - range delay relation, 165
 - scatterer changes due to, 46
- Dimensions, number of, 34
- Dispersive backscattering, 82
- Dive boat, 517–23
 - amplitude function, 548
 - anchor, 555
 - best accuracy, 554
 - compensated bow scatterer, 565
 - crossrange scale factor, 566
 - at end of turn, 517–23
 - image cuts, 549, 550, 551, 552, 553
 - image with scatterer near stern
 - compensated, 552

- Dive boat (continued)
 - imaging interval, 547
 - intensity image, 521
 - length measurements, 546–55
 - measurement between top/bottom scatterers, 518, 520
 - peaks tracks, 518, 519
 - phase variations, 549
 - photograph, 518
 - range gate size, 554
 - range tracks, 521
 - relative Dopplers comparison, 521
 - roll Doppler, 563, 564
 - roll Doppler uncertainty, 565
 - roll measurement, 520
 - topview, 547, 564
 - two-second image with bow response compensated, 550
 - width, 555
 - yaw Doppler, 563, 564
 - yaw measurement, 518
 - yaw motion, 564
 - See also* Ships
- Dominant scatterers, 592, 594, 595, 616
 - Doppler history of, 693
 - tracking, 620
- Doppler
 - absolute, 468
 - of aircraft, 124
 - ambiguity, 142
 - bands, 164
 - changes, 416, 510, 610
 - constant, 43, 92
 - curves, 462
 - differences between scatterers, 424
 - differential, 43, 45, 514, 602, 608–9
 - filtering, 595, 619
 - gates, 205
 - history, 515, 587, 590, 598, 693
 - instantaneous, 592
 - jump, 512
 - measurements, 95, 464, 590
 - pitch, 463
 - pseudoperiodic, 484
 - residual, 371
 - roll, 441, 442, 456, 465–66, 469
 - separation, 98, 625
 - separation of dominant/secondary scatterers, 608
 - shifts, 514, 574
 - sinusoidal, 582
 - spread, 165, 168, 485
 - uncertainty, 484
 - variable relative scatterer, 607
 - variations, 154, 468, 582, 618
 - width, 484
 - yaw, 439, 452, 454, 456, 463, 469
 - zero, 477
- Doppler compensation, 140–46, 233
 - acceptability of, 282
 - centroid, 357
 - correction by, 129
 - defined, 126
 - inflexibility, 257
 - moving off-highway truck, 370, 371
 - range profile alignment after, 142, 143
 - refueling truck after, 418
 - of wing scatterer, 265
 - See also* Motion compensation
- Doppler resolution, 5, 43
 - achievable, 94
 - at broadside, 262
 - coherent processing need, 100
 - insufficient, 624
 - nominal, 607
 - obtaining, 2, 620
 - performance, 98
 - pseudo, 180
 - SWD processing vs., 611
- Doppler sidelobes, 101, 125, 153
 - suppression, 90, 101
 - two-level, 248
- Doppler tracking, 126, 205–6, 668
 - centroid, 140, 141, 203, 205, 206, 217, 393
 - of entire aircraft, 152, 154, 286
 - of entire target, 205
 - image after, 141
 - integration gain, 140
 - measurement ambiguity, 142
 - method, 140
 - moving RV, 407
 - performance of, 126
 - of range-compensated scatterer, 232

- range tracking combination, 132, 143–46, 469, 531
- of scatterer, 142, 143, 154
- ship, 482
- ship turn, 512
- of wing scatterer, 264
- See also* Phase tracking; Range tracking; Tracking
- Dot products, 667
 - extracting, 667
 - inconsistent, 669
- Eigendecomposition, 81
- Electromagnetic (EM) theory, 7
- Enhanced SWD processing, 585–90
 - defined, 585
 - interval selection, 586
 - See also* Sliding-window Doppler (SWD) processing
- Fast Fourier transform (FFT), 4, 12
 - of amplitude function, 424
 - of interval, 595
- Features
 - aircraft, 82, 182–99
 - generic, 82, 304–6, 331–35
 - ground vehicle, 82
 - howitzer, 304–6
 - position measurement, 105–7
 - recognizable target, 83–84
 - tank, 318–19
 - wave-trapping, 707, 710
- Fighter aircraft, 170
 - large aspect angles and, 249
 - nose-on image, 268
 - peaks tracks, 256
 - tail-on, 273, 274
 - See also* aircraft
- Fixed-crossrange cuts, 58, 415
 - additional, 676
 - association sequence using, 682
 - data point range variance, 674
 - displaced, 688, 689
 - for refinement, 687, 688
 - through aircraft response, 70, 686
 - through simulated response, 679
 - two-dimensional image intensity peak, 673–74, 678
 - See also* Image cuts
- Fixed-point scatterers, 10, 635
 - illustrated, 18
 - phase function curvature, 22
 - See also* Scatterers
- Fixed-range cut, 23, 24, 63, 72, 255, 415
 - additional, 681
 - moving tank, 382
 - through aircraft response, 68, 685
 - through fuselage responses, 208
 - through ground vehicle response, 78
 - through response from two scatterers, 673
 - through simulated response, 678
 - through two-dimensional response peak, 677, 684
 - transform amplitude, 288
 - See also* Image cuts
- Flatbed trucks, 326–30
 - corner, 331–32
 - at different aspect angle, 328–30
 - enclosure brace, 335
 - image of, 328
 - moving, 384–89
 - positional match, 327, 329, 332
 - sidewall brace, 332
 - viewed at 38 degrees, 326–28
 - See also* Stationary ground vehicles; Trucks
- Fourier transform, 68, 95, 216, 583
 - fast (FFT), 4, 12, 424, 595
 - inverse, 95
 - of range gate, 647
- Frequency stepping, 166–67, 237
- Fuselage
 - direction, 221
 - image cut along centerline of, 274
 - image cuts along, 25
 - over crossrange gates, 255, 258
 - resolution cells, 207
 - shape, 105
 - two-dimensional amplitude peak, 290
 - See also* Aircraft; Aircraft features
- Gaussian weighting, 51, 642
- Generic features, 82
- Ground vehicle features, 82

- Ground vehicle features (continued)
 - brace of flatbed sidewall, 332
 - bumper corner, 333–34
 - corner/cavity, 334
 - enclosure brace, 335
 - flatbed corner, 331–32
 - headlights, 332–33, 334
 - howitzer, 304–6
 - inside fender, 335
 - measurable, 336
 - metallic box, 334
 - outside mirror, 335
 - radar features vs., 297–98
 - rear wheels, 335
 - recognizable, 298–99
 - tank, 318–19
 - trihedral corner, 299, 300, 333
 - turn indicator, 334
 - unrecognizable, 300
 - vertical exhaust pipe, 333
 - wheels, 299
- Ground vehicle identification, 291–432
 - applications, 292–96
 - automated, 428–31
 - basics, 296–300
 - conditions variability, 291–96
 - general conditions for, 292
 - of moving vehicles, 297, 336–432
 - radar platform and ground vehicle
 - stationary, 292
 - radar platform moving, ground vehicles
 - moving, 294–96
 - radar platform moving, ground vehicle
 - stationary, 293
 - radar platform stationary, ground
 - vehicles moving, 292–93
 - of stationary vehicles, 296–97, 300–306
 - See also* Target identification
- Ground vehicle images
 - flatbed truck, 328
 - howitzer, 302
 - moving flatbed truck, 385
 - moving recreational
 - vehicle, 401, 402, 403
 - off-highway truck, 321
 - tank, head-on, 379
 - tank, moving in circle, 381
- Ground vehicles
 - feature measurement, 34
 - flatbed truck, 326–30
 - flexing, 584
 - in forest area, 293
 - howitzer, 301–14
 - motion behavior, 295
 - motion compensation and, 155
 - motions, 291
 - off-highway truck, 320–26
 - pitch, 433
 - range rate, 294
 - refueling truck, 417
 - resolution and, 39
 - roll, 433
 - stationary, 104
 - tanks, 83, 298, 314–20
 - trucks, 320–30, 417–27
 - turning, 391
 - vibrating, 584
 - yaw, 433
 - See also* Moving ground vehicles; Stationary ground vehicles
- Hamming weighting, 19, 683
- Headlights, 318, 332–33, 334
- High-frequency motions, 239
- Horizontal stabilizer
 - image cut and transform, 270
 - specular flashes, 273
 - See also* Aircraft
- Howitzer, 301–14
 - aspect angle effect, 306–11
 - broadside image, 312
 - at different aspect angle, 307
 - distinctive features, 304–5
 - image, 302
 - multiple delayed returns, 305
 - peaks plot, 301
 - persistence of scatterers, 311–14
 - positional match, 306, 313–16
 - special feature, 304
 - target outline measurement, 301–4
 - See also* Ground vehicles; Stationary ground vehicles
- Ideal point scatterers, 23
 - resolution performance on, 41

- responses, 108
- TSA implementation for, 52–62
- See also* Scatterers
- Identification procedure, 82–86
 - deformable template match, 84–86
 - features, 82
 - See also* Target identification
- Image analysis
 - automated aircraft
 - identification, 289–90
 - automated moving ground vehicle
 - identification, 431
 - moving ground
 - vehicles, 341–47, 409–17
 - See also* Complex-image analysis
- Image cuts
 - 120°, 74
 - after phase compensation, 120
 - along fuselage, 25
 - circular Gaussian weighting and, 51
 - compensated scatterer, 241
 - in crossrange gate of
 - response, 23, 113, 116, 117
 - dive boat, 549, 550, 551, 552, 553
 - fixed-crossrange, 58, 70, 73, 415, 676
 - fixed-range, 55, 58, 63, 68, 72, 78, 255, 382, 415, 677
 - horizontal stabilizer, 270
 - of intermittent scatterer, 140
 - moving flatbed truck, 388
 - moving off-highway trucks, 350, 351, 353, 364, 374, 375
 - moving recreational vehicle, 394, 398
 - one-dimensional, 50, 65–67
 - in range gate of compensated scatterer, 240
 - in range gate of peak, 20, 21
 - rotating, 22
 - ship, 483, 495, 497
 - ship turn, 513
- Image generation flowchart, 3
- Image-time selection, 447–52
 - coast guard cutter, 536
 - principles, 453
 - purpose, 448
 - zero Doppler and, 477
 - See also* Ship imaging
- Imaging
 - aircraft, 203–83
 - basis assumption, 7
 - maneuvering aircraft, 229–39
 - ship, 434–47
 - two-dimensional, 1–4
 - See also* Imaging moving targets
 - Imaging intervals, 9, 146, 255
 - aspect angle at center, 259
 - coast guard cutter, 534, 541
 - common, 414, 416, 429
 - crossrange resolution and, 226
 - dive boat, 547
 - expanding, 208
 - halved, 270
 - image over central half of, 221
 - long, 146–47
 - moving flatbed truck, 387
 - moving recreational
 - vehicle, 389–90, 400
 - phase slope and, 589
 - precision, 213
 - reducing, 282, 431, 501
 - ship, 448, 494
 - two-second, 268
 - yaw Doppler and, 503
 - Imaging interval
 - selection, 146–54, 209, 213, 244
 - automated aircraft
 - identification, 288–89
 - core operations, 410
 - enhanced SWD processing, 586
 - flowchart, 210, 411
 - flowchart with branching criteria, 412
 - moving ground vehicles, 343–44
 - moving off-highway truck, 356
 - refueling truck, 421
- Imaging moving targets, 122–59
 - motion compensation, 125–56
 - motion determination, 156–58
 - principles, 122–25
 - summary, 158–59
 - See also* Moving targets
- Imaging times
 - coast guard cutter, 535, 543
 - finding, 578
 - selection of, 577–78

- Imaging times (continued)
 - shift of, 475
 - See also* Ships
- Inaccuracies
 - phase slope, 596
 - reducing, 689
- Instrument panels, 196
- Intensity images, 271–73
 - complex images vs., 4–6
 - dive boat, 521
 - peaks plot, 18
 - performance and, 34
 - resolution and, 6
 - response examination, 27
 - responses, 89
 - response types, 40
 - sample, 17
 - tanks, 318
- Intensity range profiles, 131, 136, 170
 - peaks illustration, 136
 - rapid change in shape of, 172
 - as target signature, 180
 - tracking of peaks, 131–32
 - See also* Range profiles
- Interactive analysis, 52
- Interactive two-dimensional TSA, 671–89
- Interference
 - amplitude/phase pattern, 139
 - destructive, 97, 644
 - introduction of, 90
 - pattern for two response, 109
 - poor, conditions, 96
 - range, 115
 - rare, conditions, 682
 - two-scatterer pattern, 53, 71
- Interfering scatterers, 59, 65
 - amplitude, 680
 - constructively, 630
 - fixed point, 633
 - phase curvature, 677
 - phasor diagram, 629
 - point, 630
 - tracking, 691–703
 - See also* Scatterers
- Inverse synthetic aperture radar (ISAR), 83
 - aircraft identification via, 181–82
 - aircraft image, 707, 708
 - moving targets and, 294
- Irregular motion, 416
- Jet engine modulation (JEM), 267
- Jet engine returns, 168
 - behavior, 164
 - Doppler-shifted, 279, 280
 - peaks plot image of aircraft, 166
 - rotating blades, 267, 268
 - spreading, 166
 - See also* Aircraft
- Jet engines
 - fighter aircraft, 170
 - intake and exhaust position, 184
 - intake duct opening size, 184
 - location of, 181
 - number of, 183–84
 - See also* Aircraft
- Large aspect angles, 29
 - aircraft imaging at, 245–66
 - coast guard cutter, 530
 - commercial aircraft, 279
 - fighter aircraft identification and, 249
 - limitations, 247
 - motion compensation at, 247
 - scatterer persistence over, 311
 - ships, 443
 - See also* Aspect angles
- Likelihood function, 85
- Linear FM
 - generation by slow frequency stepping, 237
 - range/Doppler coupling of, 168
 - waveforms, 165
- Linear-least-squares fit, 707
- Linear phase, 697, 700
- Linear polarizations, 81
- Line-of-sight (LOS), 437
- Maneuvering aircraft, 229–45
 - combined with vibrations, 239–45
 - defined, 229
 - Doppler track of range-compensated scatterer, 232
 - imaging, 229–39
 - imaging interval selection, 244
 - peaks tracks, 231, 232

- range profiles, 229, 230
- range profiles after compensation, 231
- range profiles at beginning of turn, 230
- survey image, 240
- See also* Aircraft; Aircraft imaging
- Man-made targets
 - intensity image responses, 40
 - wave-trapping features, 39
- Measurements
 - crossrange, 117, 680
 - cross section, 80
 - dispersive backscattering, 82
 - Doppler, 95, 464, 590
 - feature extent, 77–80
 - for feature positions, 105–7
 - length, 105–6
 - methods, 10–28
 - phase, 124, 147, 149
 - phase slope, 78, 587
 - polarization diversity, 80–81
 - range, 108, 430–31
 - roll, 455, 474–545
 - roll Doppler, 457
 - of scatterer positions, 77, 106, 173, 221
 - ship length/width, 546–69
 - ship motion, 467
 - special, of potential use, 77–82
 - target outline, 301–4
 - uncertainty, 676
 - width, 105–6
 - yaw, 455, 474–545
- Military ships, 570
- Missiles, 582
- Modulation
 - amplitude, 73, 156, 350, 419, 422, 424
 - index, 99
 - jet engine (JEM), 267
 - measurability of, 99–100
 - period, 157
 - phase, 214, 588
- Motion
 - complicated, potential problems, 470
 - determination, 156–58
 - erratic nature of, 35
 - ground vehicle, 291, 295
 - high-frequency, 239
 - irregular, 416
 - phase center, 103
 - pitch, 433–34, 444
 - relative phasing of, 668
 - residual, 373–74
 - roll, 377–78, 433–34, 474–545
 - ship, 452–67
 - two-dimensional, 659–63
 - uncompensated, 649, 651
 - yaw, 246, 263, 417, 433–34, 438, 474–545
- Motion compensation, 36, 125–56
 - aircraft, 155, 204–5
 - crossrange resolution and, 118
 - Doppler compensation, 140–46, 257
 - failure, 257
 - with fifth-order polynomial, 612
 - goal, 158
 - ground vehicles and, 155
 - ground vehicles on poor roads, 295
 - image cut for checking, 119
 - individual, in each gate, 376
 - intake response, 253
 - large aspect angles and, 247
 - moving ground vehicles and, 337–38
 - moving off-highway truck, 354, 362
 - moving recreational vehicle, 398
 - one-second image with, 233
 - phase tracking, 146–54
 - poor, 118
 - practical aspects of, 154–56
 - precise motion for, 159–60
 - purpose, 204
 - quality, checking, 118
 - range compensation, 126–40
 - range rate change and, 155
 - refueling truck, 417–27
 - requirements, 125, 155–56, 205
 - ships, 155, 470–71, 578–79
 - steps, 127–28, 159, 206
 - of strongest scatterer, 603
 - success, 257
 - tip scatterer, 252, 253
 - TSA and, 222
 - See also* Imaging moving targets
- Moving flatbed trucks, 384–89
 - amplitude/phase functions, 386
 - crossrange uncertainty, 388

- Moving flatbed trucks (continued)
 - image cut, 388
 - imaging interval, 387
 - on bumpy straight road, 384–89
 - positional match, 388–89
 - survey image, 385
 - See also* Flatbed trucks; Moving ground vehicles
- Moving ground vehicles
 - automated identification of, 428–31
 - complications, 336
 - conditions for, 295
 - core operations of image interval selection, 343
 - crossrange accuracy, 340
 - crossrange resolution and, 339
 - in erratic manner, 384
 - flatbed truck, 384–89
 - flowchart for image interval selection, 344
 - flowchart with branching criteria, 345
 - identification of, 297, 336–432
 - identification peculiarities, 337–38
 - image formation process, 337
 - length/width, 338
 - motion compensation, 337–38, 417–27
 - motion consequences, 338–40
 - moving platform, 294–96
 - off-highway truck, 347–77
 - processing procedures, 340–47
 - range gates and, 340
 - recreational vehicle, 389–409
 - rigid-body rotation, 347
 - in SAR scene, 428
 - scatterer positions, 338
 - sizes, 347
 - special features, 338
 - as special situation, 337
 - stationary platform, 292–93
 - survey image analysis, 341–47, 409–17
 - survey image formation, 341
 - tank, 377–84
 - transform area into raw data, 341
 - treatment, 337
 - turning, 391
 - two-dimensional positional match, 339
 - See also* Ground vehicles; Moving targets
- Moving off-highway trucks, 347–77
 - aspect angle, 357
 - in circle, survey image, 348
 - compensated image, 374
 - crossrange accuracy, 356
 - crossrange width, 350
 - Doppler centroid compensation, 357
 - Doppler compensation of one scatterer, 370, 371
 - feature template, 364
 - final image, 355
 - image cut in range gate of compensated scatterer, 372
 - image cuts, 350, 351, 353, 364, 374, 375
 - imaging interval selection, 356
 - individual motion compensations in each range gate, 376
 - match between scatterer positions, 367, 368
 - motion compensation, 354, 362
 - noise/clutter background level, 360
 - on bumpy straight road, 368–77
 - on smooth road, survey image, 359
 - on smooth straight road, 357–68
 - peaks tracks, 370
 - phase function comparison, 360
 - phase slope measurements, 364
 - positional match, 356, 365, 377
 - positional match, on poor road, 378
 - range gates, 349–50
 - range profile peak selected for tracking, 370
 - range profiles, 369
 - reduced-duration imaged, 353
 - residual motion, 373–74
 - rigid-body rotation, 360
 - rigid movement with simultaneous translation, 361
 - rotational jerks, 361
 - scatterer positions, 367
 - in slow circle on flat terrain, 347–57
 - stationary off-highway truck vs., 366–68
 - two-dimensional positional match, 355

- See also* Moving ground vehicles;
Moving targets; Off-highway trucks
- Moving platform
moving ground vehicles, 294–96
stationary ground vehicles, 293
- Moving recreational vehicle, 389–409
amplitude function after
compensation, 397
antennas, 402
aspect angle, 389
broadside aspect, 406–9
clutter cancellation, 390–91, 406
corner reflectors, 402
crossrange resolution, 399
Doppler centroid tracking, 393
first vehicle image, 400
image after range/Doppler centroid tracks, 407
image cut, 394, 398
imaging interval, 389–90, 400
motion characteristics of scatterers, 400
motion compensation, 398
not close to broadside, 389–406
peak analysis, 402
peaks plot after motion
compensation, 393
phase compensation, 401
positional match between
measured/predicted scatterers, 409
positional match for first image, 403
positional match for second image, 405
positional match for third image, 405
radar wavelength, 389
range centroid, 393
scatterer motion, 394
second vehicle image, 401
smeared response transform, 396
spurious sideband responses, 403
survey plot, 390
survey plot without cancellation, 392
third vehicle image, 402
turning, 399
two-scatter model curves, 407
See also Ground vehicles; Moving ground vehicles
- in circle, 30° aspect angle, 381
in circle on terrain, 377–84
fixed-range image cuts, 382
head-on aspect, 377–80
match between stationary tank
template and, 383
positional match with uncertainty
ellipses, 383
rolling motion, 377–78
scatterer positions, 384
smearing, 381
turning, viewed at larger aspect
angle, 381–84
See also Moving ground vehicles; Tanks
- Moving targets
crossrange resolution
and, 121, 122, 124, 153
flexing of, 125
identification, 123
imaging, 122–59
ISAR images and, 294
smoothly, 125
See also Moving ground vehicles
- Nyquist criterion, 12
- Observation window, 92, 93, 94, 95
- Off-highway trucks, 320–26
backscattering, 321
at different aspect angle, 323–26
end of, 324
front view, 323
image of, 321
length/width measurement, 321
moving, 347–77
moving vs. stationary, 366–68
outline, 321
positional match, 327, 358
rear view, 320–23
spurious responses, 324
vehicle corner, 323
See also Ground vehicles; Trucks
- One-dimensional image cuts, 50, 65–67
- One-dimensional resolution, 11, 91–96
defined, 98
process, 11
See also Resolution

- One-dimensional TSA, 12, 38, 53–59, 627–39
 - amplitude, 628
 - defined, 53, 76, 627
 - for ideal point scatterers, 76
 - maximum phase, 628
 - phase slope measurements, 630
 - phasor diagram, 628, 629
 - to superresolve scatterers, 635
 - See also* Two-scatterer algorithm (TSA)
- Peaks plot, 16
 - for crossrange resolution, 115, 116
 - illustrated, 18
 - moving RV, 393
 - refueling truck, 427
- Peaks tracks, 130–31
 - after Doppler processing, 251
 - after linear range/Doppler compensation, 248
 - coast guard cutter, 524, 529, 530, 537–38, 539, 540
 - cruise ship, 505
 - defined, 130
 - dive boat, 518, 519
 - for fighter aircraft, 256
 - first, image after motion compensation of, 249
 - fitting quadratic polynomial and, 264
 - illustrated, 131
 - intermittent, polynomial fit for, 139
 - maneuvering aircraft, 231, 232
 - moving off-highway truck, 370
 - with polynomial fit, 133
 - procedure, 132
 - quality, 135, 237
 - range, 20° depression angle, 459
 - ship, 453, 478, 479, 498
 - ship turn, 511, 515
 - smoothed, ship, 480
 - for two interfering scatterers, 134
 - wrong, 135
- Phase
 - differential, 92
 - distortions, 156
 - fluctuations, 158, 694
 - linear, 697, 700
 - maximum, 92
 - measurement, 124, 147, 149
 - quadratic, 71
 - relation between scatterers, 94
 - variations, 139, 147, 243, 484, 549
- Phase center, 7–8
 - motion, 103
 - radar, 196
 - shift, 32, 64, 79
 - stable, 63
 - wander effects, 37
- Phase compensation, 621
 - moving recreational vehicle, 401
 - refueling truck, 426
- Phase curvature, 23, 59, 599, 697
 - rapid change, 79
 - small, 73
- Phase function, 63, 234, 259
 - amplitude function combined with, 152
 - approximation, 64
 - behavior of, 103
 - cruise ship, 561
 - curvature, 77, 611, 613
 - fit smooth function to, 621
 - moving flatbed truck, 386
 - piecewise linear, 602
 - polynomial fit, 592, 611
 - in range gates, 361
 - refueling truck, 421, 422
 - ship, 468, 494
 - transform window placement and, 260
 - variations, 271
 - See also* Ambiguity function
- Phase jumps, 46, 56, 67, 595
 - defined, 92
 - deviation, 636–38
 - extrapolating time of, 59
 - fitting through, 622
 - measuring, 678
 - one, in transform, 59
 - rapid, 611
 - recognition of, 423
 - removal of, 596, 602
 - size of, 54, 56, 59, 678
 - transform, 59
 - values, 56
- Phase modulation, 214, 588

- correlated with amplitude
 - modulation, 588
 - objectionable, 588
 - slow, 596
 - See also* Modulation
- Phase-slope algorithm, 614
- Phase-slope breaks, 208, 289, 360, 420
- Phase slope compensation
 - method, 589
 - signal after, 615
- Phase-slope measurement, 78, 364, 587, 624, 631
 - choosing intervals for, 692–95
 - data analysis, 695–703
 - integrated fit, 696, 700
 - using, 703
- Phase slopes, 23
 - abrupt changes in, 157, 214
 - at amplitude maximum, 54
 - changes, 614
 - difference, 79, 316, 361, 607
 - inaccuracies, 596
 - instantaneous, 77–78
 - intervals and, 589, 607
 - relation between Dopplers of scatterers and, 634
 - switching between different values, 419
 - tracking, 147–48, 208, 423, 595
 - tracking interfering scatterers with, 691–703
- Phase tracking, 146–54, 205, 468, 668
 - applications, 157
 - avoiding, 158
 - of single scatterer, 147
 - slope, 147–48, 208
 - See also* Doppler tracking; Range tracking; Tracking
- Phase unwrapping, 19
- Phasor diagram, 42
 - of two interfering scatterers, 629
 - of two scatterers at different Dopplers, 92
- Pitch
 - aircraft, 433–34
 - axis, 574
 - ground vehicle, 433
 - motion, 444
 - motion period, 662
 - zero, 435
 - See also* Roll; Ships; Yaw
- Pitch Doppler
 - roll Doppler combination, 528
 - yaw Doppler combination, 463
- Pitch images, 444–47
 - at 0° depression angle, 444
 - at 20° depression angle, 445
 - at 90° depression angle, 446
 - illustrated, 464
 - See also* Ship images
- Point targets
 - half-power width response, 41
 - receiver output phase and, 5
 - theory, 6
- Polarization diversity, 80–81
 - best use of, 81
 - of real targets, 81
- Polarizations, 81
- Polynomial fit
 - compensating with, 232
 - for intermittent peaks track, 139
 - least-squares, 667
 - phase function, 592, 611
 - range compensation based on, 511
- Positional match
 - assisting, 331–35
 - for correct/incorrect aircraft, 283–85
 - flatbed truck, 327, 329, 332
 - howitzer, 306, 313–14
 - moving flatbed truck, 388–89
 - moving off-highway truck, 355, 356, 377
 - moving recreational vehicle, 403–5
 - moving tank, 383
 - off-highway truck, 327, 365
 - procedure results, 284
 - two-dimensional, 339, 355
 - when no feature is fixed, 284, 285
- Primitives, 81
- Pseudoperiodic yaw motion, 154–55
- Radar
 - aircraft diagrams and, 196
 - nose cone, 186–87
 - phase center, 196

- Radar (continued)
 - scatterer, 196
- Radar imaging. *See* Imaging
- Radome tip, 195
- Range
 - cells, 245–46
 - centroid, 203, 205, 217, 393
 - delay, 165
 - detection, 127
 - difference, 463
 - drift, 129–30, 470, 485, 589
 - extents, 80, 416
 - interference in, 115
 - interval, 134
 - measurements, 108, 430–31, 676
 - peaks tracks, 459
 - positions, refining, 681
 - rate, 294
 - rate change, 155
 - rate mismeasurement, 663
 - scale, 449
 - separation, 499
 - sidelobe suppression, 168
 - two-dimensional motion, 660
- Range compensation, 126–40
 - acceptability of, 282
 - based on polynomial fit, 511
 - range profiles after, 133
 - refueling truck after, 418
 - of wing scatterer, 265
- Range gates, 240
 - Fourier transform, 647
 - image cut, of compensated scatterer, 372
 - image cut and transform in, 243
 - image cut in, of tip response, 250
 - individual motion compensations for, 376
 - moving ground vehicles and, 340
 - moving off-highway truck, 349–50
 - negative, 273
 - phase functions in, 361
 - refueling truck, 422–23
 - superstructure scatterer, 457
- Range profiles
 - for 300-MHz bandwidth, 171
 - for 600-MHz bandwidth, 173
 - for 1200-MHz
 - bandwidth, 174, 175, 178
 - after Doppler compensation, 142, 143
 - after drift removal, 129
 - after range compensation, 133
 - at aircraft turn beginning, 230
 - comparing, 172
 - comparison with dashed pointer scatterer response, 136
 - complex, 170
 - database and, 169–70
 - with delayed duct returns, 175–77
 - detail, 169
 - for entire aircraft including duct returns, 178
 - identification via, 168–81
 - illustrated, 137
 - intensity, 131, 136, 170
 - maneuvering aircraft, 229, 230
 - moving off-highway truck, 369
 - near broadside, 262
 - peaks, 130, 132, 136, 170
 - peaks, tracking, 545
 - requirements, for identification, 169–70
 - response analysis, 131
 - sequence of, 128
 - ship, 435–36
 - ship, sequence, 479
 - with skin returns and delayed duct returns, 177–79
 - utility conclusions, 179–81
 - without duct returns, 170–75
- Range resolution, 2
 - for aircraft identification, 35, 199
 - aspect angle size and, 260, 469
 - asymmetry of, 104–22
 - cell, 165
 - coarse, 246
 - crossrange resolution equality, 104
 - crossrange resolution vs., 106
 - degrading, 205
 - high, 104
 - implementation, 106
 - importance, 9
 - of individual duct returns, 177
 - ineffective, 9–10
 - as primary resolution, 106, 122

- role of, 29
 - utilizing, 107–10
 - See also* Resolution
- Range tracking, 467, 668
 - dive boat, 521
 - Doppler tracking combination, 132, 143–46, 469, 531
 - enhanced, 138
 - entire aircraft, 152
 - facilitating, 131
 - image after, 141
 - improving, 135, 143
 - mean error, 669
 - moving RV, 407
 - scatterer location extraction, 665–70
 - scatterers, 132
 - ship turn, 512
 - tip scatterer, 251
 - See also* Doppler tracking; Phase tracking; Tracking
- Recreational vehicle (RV). *See* Moving recreational vehicle
- Refinement procedure, 681–83
 - combination, 683
 - fixed-crossrange cut for, 687, 688
 - three scatterer positions and, 683
- Refueling trucks, 417–27
 - image after range centroid and Doppler centroid compensation, 418
 - image cut, 419
 - image duration, 422
 - imaging interval reduction, 423
 - imaging interval selection, 421
 - motion compensation, 417–27
 - peaks plot, 418, 427
 - phase compensation, 426
 - phase functions, 421, 422
 - range gates, 422–23
 - spline fit for shortened time interval, 425
 - See also* Ground vehicles
- Registration, 431
- Residual motion, 373–74
 - pseudoperiodic, 670
 - size of, 647
 - uncompensated, 647–57
- Resolution
 - in azimuth, 104
 - choice of, 40
 - constraints, 202
 - crossrange, 2, 6, 9, 36, 79–80, 104–22
 - defined, 29
 - degradation, 99
 - delay, 95
 - in elevation, 104
 - ground vehicles and, 39
 - in identification performance, 36–39
 - intensity images and, 6
 - one-dimensional, 11, 91
 - principles, 41–52
 - range, 2, 9–10, 29, 35, 104–22
 - shifting scatterers and, 102–3
 - three-dimensional, 36
 - two-dimensional, 35, 50, 98–99
 - two point scatterers, 94–95
 - of weak scatterers, 99–100
- Resolution cells, 37, 109
 - crossrange, 79, 110
 - fuselage, 207
 - optimum size of, 39
 - range, 165
- Resolution performance, 95
 - close-target, 90
 - complex-image analysis, 98
 - derived from envelope of ambiguity function, 90
 - Doppler, 98
 - improvement, 123
 - inherent, realizing, 96
- Resolution requirements, 29, 121
 - aircraft identification, 199–203
 - target identification, 105
- Response groups, 619–25
- Rigid-body
 - rotations, 347, 360
 - yaw, 398–99
- Rockets, 582
 - attitude-stabilized, 584
 - performing attitude maneuver, 591, 692
 - reorientation of, 611
 - transform return, 619–20
- Roll, 474–545
 - aircraft, 433–34

- Roll (continued)
- axis, 574
 - coast guard cutter, 504
 - dominant, 662, 663
 - ground vehicle, 433
 - motion period, 662
 - range difference and, 463
 - relative strength of, 476
 - significant, 477–504
 - very little, 504–10
 - zero, 435
- See also* Pitch; Roll; Ships
- Roll Doppler, 441, 442, 465
- for 0° depression angle, 456, 458
 - for 20° depression angle, 460
 - for 20° depression angle, improved version, 466
 - coast guard cutter, 526, 533
 - cruise ship, 509
 - curves, 493, 497–98, 500, 504
 - dive boat, 563, 564
 - illustrated, 487, 488
 - image quality and, 516
 - inaccurate, 544
 - maximum, 469
 - measurability, testing, 508
 - measurement, 457, 477–98, 576
 - measurement error, 502
 - for new time interval, 500
 - pitch Doppler combination, 528
 - residual, 492
 - ship turn, 515, 517
 - strength comparison, 501
 - uncertainty, 565
 - zero, 465–66, 502, 517, 533, 571, 577
 - zero, image at, 461
 - zero crossings, 577
 - zero yaw Doppler and, 578
- See also* Yaw Doppler
- Roll images, 441–44
- at 0° depression angle, 441,
 - at 20° depression angle, 442
 - at 90° depression angle, 443
 - approximate, 663
 - finding good times for, 498
 - illustrated, 472, 510
 - sideview, 443
- See also* Ship images
- Roll measurement, 455, 474–545
- accurate, 482–84, 490
 - dive boat, 520
 - Doppler, 477–98
 - error, 502
 - function, 576
 - illustrated, 487
 - interpretation of, 498–504
 - principles, 474–77
 - reference scatterer compensation, 486
 - requirements, 499
 - scatterer positions, 489
 - ship turn, 514
- Satellites, 582
- Scale-modeling aircraft, 193–94
- Scatterer locations
- extracting, 665–70
 - variation with tracking, 669
- Scatterer positions
- center, 76
 - for delayed duct return, 228
 - derived for worse interference, 71
 - howitzer, match between, 313–14
 - matched to database, 377
 - measurement accuracy of, 101
 - measurements of, 77, 106, 173, 221
 - for motion measurement, 489
 - moving off-highway truck, 367
 - moving tank, 384
 - on ship deck, 569–70
 - two-dimensional, 429–30
 - well-observable, 377
- Scatterers
- aircraft, 121
 - amplitude, 47
 - bow, 455, 465, 473
 - compensation, 139, 241, 372
 - composite signal returned by, 11
 - dominant, 592, 594, 595, 616, 620
 - Doppler difference between, 94
 - Doppler separation of, 94, 98
 - Doppler track of, 142, 143, 154
 - Doppler variations of, 582
 - extracted, 85
 - fixed-point, 10, 18, 22, 635

- height of, above deck, 570–75
 - high-frequency motions of, 239
 - ideal point, 23, 41
 - individual, motion characteristics
 - of, 400
 - interfering, 59, 65, 629, 677
 - intermittent, 140
 - number of, 121
 - persistence of, 311–14, 319–20
 - phase relation between, 94
 - phase tracking of, 146–54
 - point midway between, 644
 - primary, 653
 - radar, 196
 - range tracking, 132
 - secondary, 593, 594, 595, 613, 653
 - shifting, 102–3
 - simulated, 649, 651
 - stern, 452, 455, 465, 473
 - strong, 99–100
 - strong, position error, 636, 637, 638
 - superstructure, 452, 457
 - three-dimensional distribution of, 575
 - tip, 252, 253
 - tracking, 130, 138
 - two point, 41–51
 - weak, 99–100
 - weak, position error, 637, 638
 - wing, 264, 265
- Scatterer separations, 647–57
- absolute, 648
 - determining, 650
 - signs, determining, 648
- Sea-level view images, 449, 454, 457, 462
- formation, 544, 577
 - timing refinement, 475
- See also* Ship images
- Shadowing, 229, 379
- absence of, 273
 - turret, 339
- Ship identification, 433–579
- basics, 433–73
 - imaging time/duration
 - selection, 577–78
 - military, 570
 - motion analysis, 576–77
 - motion compensation, 578–79
 - peculiarities of, 433–34
 - processing steps, 433, 576–79
- See also* Target identification
- Ship images
- analysis, 546–75
 - high-quality, 493
 - hybrid, 472
 - pitch, 444–47, 464
 - roll, 441–44, 472, 510
 - sea-level-view, 449, 454, 457, 462, 475, 544
 - short-term, for scatterer selection, 454
 - topview (coast guard cutter), 527, 535–36
 - topview (dive boat), 547
 - topview (ship turn), 516
 - topview, 456, 475, 492, 544, 578
 - two-dimensional, 437
 - undistorted, 450
 - yaw, 437–40, 465
 - yaw/pitch/roll, 446–47
 - for zero yaw, roll, pitch Dopplers, 435
- Ship imaging, 434–47
- interval, 448
 - range profiles, 435–36
 - time selection, 447–52
- See also* Imaging
- Ships
- aircraft vs., 434
 - amplitude function, 468, 494
 - bow scatterer, 455, 465, 473
 - centerline, 575
 - classification, 579
 - coast guard cutter, 498, 504, 523–43
 - cruise, 505–10
 - deck, 475
 - dive boat, 517–23
 - Doppler track, 482
 - erratic motion, 493, 504
 - ground vehicles vs., 434
 - heeled, 437
 - image cuts, 483, 495, 498
 - imaging duration, 577–78
 - imaging intervals, 494
 - imaging times, 577–78
 - intensity range profile, 436
 - large aspect angles, 443

- Ships (continued)
 - length, 546–69
 - military, 570
 - motion, 452–67
 - motion analysis, 576–77
 - motion analysis flowchart, 476
 - motion analysis for real data, 467–72
 - motion compensation, 155, 471–72, 578–79
 - motion measurement accuracy, 467
 - motion measurement purpose, 470
 - motion states, 545
 - outline, 544
 - peaks tracks, 453, 478, 479, 498
 - phase function, 468, 494
 - pitch, 434
 - range profiles sequence, 479
 - range scale, 449
 - roll, 434
 - scatterer positions on deck, 569–70
 - scatterer selection, 543
 - simulated, 437
 - slowly moving, 493
 - small, in rough seas, 523–43
 - stern scatterer, 452, 455, 465, 473
 - tracking performance, 469
 - turn maneuver, 510–17
 - very little roll motion, 504–10
 - width, 448, 545–69
 - yaw, 434
- Sidelobes
 - crossrange, 4
 - Doppler, 101, 125, 595–96
 - high, generation of, 100
 - low-level response, 19
- Sidelobe suppression
 - Doppler, 90, 101
 - range, 168
 - weighting for, 19, 100–102
- Signal sections, 619–25
 - after compensation, 623
 - with similar characteristics, 622
- Signal-to-noise ratio (SNR), 694
- Skewness, 69
- Sliding-window Doppler (SWD)
 - processing, 585
 - adaptive window lengths and, 610
 - complex-image analysis adaptation to, 583
 - defined, 581
 - detriment, 583
 - enhanced, 585–90
 - overview, 581–83
 - plot of attitude maneuver, 591
 - in practice, 583
 - wide use of, 582
- Smoothed peaks tracks, 480
- Smoothed signal, 15
- Smooth flight, 611–19
 - analysis, 611–19
 - differential Dopplers, 611
 - dominant scatterer compensation, 611
 - objective, 611
 - resolution performance, 618
- Smoothing, 16
- Spurious responses, 8–9, 36, 243
 - defined, 64
 - generated by cavity-type features, 106
 - minimizing, 156
 - moving target, 9
 - objectionable strength, 246
 - off-highway trucks, 324
- Stationary ground vehicles
 - flatbed truck, 326–30
 - Howitzer, 301–14
 - identification of, 296–97, 300–336
 - off-highway truck, 320–26
 - positional match, 331–35
 - SAR surveillance, 104, 337
 - simplist, 300–301
 - tank, 314–20
 - See also* Ground vehicles
- Stationary platform
 - clutter advantage, 292
 - moving ground vehicles, 292–93
 - stationary ground vehicles, 292
- Stern scatterer, 455
 - coast guard cutter, 531
 - Doppler, 452, 473
 - tracking, 465, 473, 477
 - See also* Ships
- Strong scatterer position error
 - phase-jump deviation (three scatterers) vs., 638

- phase-jump deviation (two scatterers)
 - vs., 636, 637
- Superresolution, 46, 54, 96–97, 104
 - accuracy problems, 49
 - in complex-image analysis, 49
 - phase difference and, 47
 - problem, 96
 - small degree of, 51
 - utilizing, 102
- Superstructures
 - crossrange separation of, 475
 - cruise ship, 569
 - definition of, 578
 - shape of, 570–75
- Superstructure scatterer, 452, 457
 - measured Doppler of, 515
 - range gate, 457
 - responses, 462, 502
- Synthetic aperture radar (SAR)
 - aircraft image, 176
 - antenna, 267
 - baseline, 180
 - moving ground vehicle in scene, 428
 - processor, 338
 - surveillance of stationary ground vehicles, 104
 - surveillance system, 291
- Tanks, 314–20
 - 30° aspect angle, 314–16, 381
 - aspect angles, 317–19
 - drive wheel, 298
 - features, 318–19
 - gun, 314
 - headlights, 318
 - head-on aspect, 377–80
 - head-on image, 379
 - image, in circle, 381
 - image over short interval, 380
 - intensity image, 318
 - measured phase slope difference, 316
 - moving, 377–84
 - peaks plot image, 314, 315
 - persistence of scatterers, 319–20
 - turning, viewed at larger aspect angle, 381–84
 - turret cylinder, 83, 318
 - wheels, 319
 - See also* Ground vehicles
- Target detection, 127
- Target identification
 - aircraft, 163–290
 - background, 1–29
 - derivation of practical approach
 - to, 30–36
 - ground vehicle, 291–432
 - inputs to, 28–29
 - in large aspect angles, 29
 - moving target, 123
 - performance, wavelength and resolution
 - in, 36–39
 - principles, 30–40
 - procedure, 82–86
 - quasi-optical, 30
 - reliable, 34, 35, 36, 40
 - resolution requirements, 105
 - ship, 433–579
 - stationary target, 123
 - target detection vs., 127
 - under operational conditions, 37
- Target imaging. *See* Imaging
- Target motion. *See* Motion
- Taylor expansion, 661, 662
- Taylor series, 644
- Timing error, 661
- Tip scatterer, 252–53
- Tracking
 - automated, 132
 - bow scatterer, 465, 473, 477
 - correlation, 203
 - Doppler, 126, 140, 203, 205, 668
 - interfering scatterers, 691–703
 - performance improvement, 469
 - phase, 146–54, 205, 468, 668
 - phase slope, 147–48, 423, 595
 - range, 131, 135, 138, 467, 668
 - range/Doppler, 143–46
 - reliable, 469
 - scatterers, 138
 - single scatterer, 130
 - single-scatterer procedure, 217
 - stern scatterer, 465, 473, 477
 - uncertainty, 485

- Transforms
 - after phase compensation, 601
 - after strongest response suppressed, 607
 - after suppressing indicated window, 605
 - amplitude, 57, 156, 208
 - compensated scatterer, 241
 - filtered cut in compensated image, 216
 - filtered image cut, 215
 - horizontal stabilizer, 270
 - instantaneous phase slope of, 77
 - noncoherent, 424
- Transform windows, 24
 - boundaries, 66, 228
 - choosing, 65, 619
 - distortion, 66
 - limits, 641–45
 - phase function and, 260
- Trucks, 320–30
 - flatbed (moving), 384–89
 - flatbed (stationary), 326–30
 - off-highway (moving), 347–77
 - off-highway (stationary), 320–26
 - refueling, 417–27
 - See also* Ground vehicles
- Truncation, 51
- Turrets
 - cylinder, 83, 318
 - shadowing, 339
- Twin-hull aircraft, 184
- Two-dimensional motion, 659–63
 - range, 660
 - time, 660
 - timing error, 661
- Two-dimensional resolution, 35, 98–99
 - implementation of, 50
 - separation normalization, 98
 - See also* Resolution
- Two-dimensional TSA, 52, 60–62, 290, 671–89
 - based on pattern interpretation, 76
 - complexity, 76
 - for ideal point scatterers, 76
 - interactive, 671–89
 - one-dimensional measurements
 - association, 60, 62
 - of real target, 62
 - response analysis with, 69
 - See also* Two-scatterer algorithm (TSA)
- Two point scatterers, 41–51
 - amplitude/phase pattern, 45
 - destructive interference, 43, 48
 - Doppler difference between, 44
 - in one dimension, 41–49
 - resolution, 94–95
 - response, 42
 - transform, 42
 - in two dimensions, 49–51
- Two-scatterer algorithm (TSA), 11, 40–77
 - accuracy, 309
 - amplitude/phase pattern with, 597
 - analysis example of real data, 67–76
 - application to real data, 62–76
 - applied to scatterers with shifting phase centers, 28
 - defined, 40
 - fixed-crossrange image cut analysis, 74
 - fixed-range image cut analysis, 72
 - implementation for ideal point targets, 52–62
 - for interactive analysis, 52
 - left side response analysis, 75
 - model amplitude, 73
 - motion compensation and, 222
 - one-dimensional, 12, 38, 53–59, 627–39
 - results, 28
 - right side response analysis, 75
 - two-dimensional, 52, 60–62, 290, 671–89
- Two-scatterer pattern, 69
 - fixed-range cut through aircraft response, 686
 - interference, 53, 71
 - recognizing, 705
- Uncertainties
 - large, 706
 - measurement, 676
 - reducing, 689
 - registration, 431
 - roll Doppler, 565
 - tracking, 485

- Uncompensated motion, 649, 651
- Waveforms
 design, 90–91
 frequency-stepped, 166
 linear FM, 165
 significance of, 163–68
- Wavelengths, 30, 36–39
- Wave-trapping features, 707, 710
- Weak scatterer, 99–100
 position error, 637, 638
 signal-to-background ratio, 99
- Weighting
 circular Gaussian, 51
 Gaussian, 642
 Hamming, 19, 683
 for sidelobe suppression, 19, 100–102
- Wheels, 299
 rear, 335
 tank, 319
 See also Ground vehicles
- Windowed response, 22, 26
- Windowing, 682
- Windows
 boundaries, 67
 observation, 92, 93, 94, 95
 position, best, 68
 transform, 24, 65–66, 228, 619, 641–45
- Wing scatterer
 Doppler track of, 264
 image after range and Doppler compensation, 265
- X-band, 36, 64
- Yaw, 474–545
 aircraft, 433–34
 coast guard cutter, 504
 dominant, 662
 ground vehicle, 433
 motion, 246, 263, 417, 438
 motion period, 662
 pseudoperiodic, 434
 pure, 473
 range difference and, 463
 relative strength of, 476
 rigid-body, 398–99
 translational, 573, 574
 variation, 246
 zero, 435
 See also Pitch; Roll; Ships
- Yaw Doppler, 439, 452, 454
 for 0° depression angle, 456
 for 20° depression angle, 460
 coast guard cutter, 526, 533, 541
 cruise ship, 507, 509, 560
 curves, 493, 497–98, 500, 504
 dive boat, 563, 564
 illustrated, 488
 imaging interval and, 503
 maximum, 469
 measurement, 477–98, 576
 measurement error, 502
 for new time interval, 500
 pitch Doppler combination, 463
 sea-level-view images and, 577–78
 ship turn, 514, 517
 small, 492
 strength comparison, 501
 zero, 490, 520, 533, 567, 577, 578
 zero, image, 459
 zero crossings, 577
 See also Roll Doppler
- Yaw images, 437–40
 at 0° depression angle, 438
 at 20° depression angle, 439
 at 90° depression angle, 440
 illustrated, 465
 See also Ship images
- Yaw measurement, 455, 474–545
 accurate, 482–84, 490
 calculation, 485
 coast guard cutter, 532
 dive boat, 518
 Doppler, 477–98
 error, 502
 function, 576
 illustrated, 487
 interpretation of, 498–504
 principles for real data, 474–77
 requirements, 498–99

- scatterer positions, 489
- ship turn, 513
- significant roll and, 477
- Yaw/pitch/roll images, 446–47
 - at 0 depression angle, 447
 - imaging time, 447
 - interpretation, 446
 - See also* Ship images
- Zero aspect angle, 266–73
- Zero-fill procedure, 16

Recent Titles in the Artech House Radar Library

David K. Barton, Series Editor

- Advanced Techniques for Digital Receivers*, Phillip E. Pace
- Airborne Pulsed Doppler Radar, Second Edition*, Guy V. Morris and Linda Harkness, editors
- Bayesian Multiple Target Tracking*, Lawrence D. Stone, Carl A. Barlow, and Thomas L. Corwin
- CRISP: Complex Radar Image and Signal Processing, Software and User's Manual*, August W. Rihaczek, Stephen J. Hershkowitz, Richard L. Mitchell, and Robert H. Mitchell
- Design and Analysis of Modern Tracking Systems*, Samuel Blackman and Robert Popoli
- Digital Techniques for Wideband Receivers*, James Tsui
- Electronic Intelligence: The Analysis of Radar Signals, Second Edition*, Richard G. Wiley
- Electronic Warfare in the Information Age*, D. Curtis Schleher
- High-Resolution Radar, Second Edition*, Donald R. Wehner
- Introduction to Electronic Warfare*, D. Curtis Schleher
- Introduction to Multisensor Data Fusion: Multimedia Software and User's Guide*, TECH REACH, Inc.
- Microwave Radar: Imaging and Advanced Concepts*, Roger J. Sullivan
- Millimeter-Wave and Infrared Multisensor Design and Signal Processing*, Lawrence A. Klein
- Modern Radar System Analysis*, David K. Barton
- Modern Radar System Analysis Software and User's Manual*, David K. Barton and William F. Barton
- Principles of High-Resolution Radar*, August W. Rihaczek

Radar Cross Section, Second Edition, Eugene F. Knott, et al.

Radar Evaluation Handbook, David K. Barton, et al.

Radar Meteorology, Henri Sauvageot

Radar Signal Processing and Adaptive Systems, Ramon Nitzberg

Radar Technology Encyclopedia, David K. Barton and
Sergey A. Leonov, editors

Theory and Practice of Radar Target Identification,
August W. Rihaczek and Stephen J. Hershkowitz

For further information on these and other Artech House titles,
including previously considered out-of-print books now available
through our In-Print-Forever® (IPF®) program, contact:

Artech House
685 Canton Street
Norwood, MA 02062
Phone: 781-769-9750
Fax: 781-769-6334
e-mail: artech@artechhouse.com

Artech House
46 Gillingham Street
London SW1V 1AH UK
Phone: +44 (0)20 7596-8750
Fax: +44 (0)20 7630-0166
e-mail: artech-uk@artechhouse.com

Find us on the World Wide Web at:
www.artechhouse.com
