

**SAR AND MTI PROCESSING OF SPARSE
SATELLITE CLUSTERS**

by

Nathan A. Goodman

B. Sc. (With Distinction) Electrical Engineering, The University of Kansas, 1995

M. Sc. Electrical Engineering, The University of Kansas, 1997

Submitted to the Department of Electrical Engineering and Computer Science and the
Faculty of the Graduate School of the University of Kansas in partial fulfillment of
the requirements for the degree of Doctor of Philosophy.

Dissertation Committee:

James Stiles: Chairperson

Christopher Allen

Sivaprasad Gogineni

Glenn Prescott

James Toplicar

Date of Defense: July 2, 2002

ACKNOWLEDGEMENTS

I would like to thank my Ph.D. advisor, Dr. James Stiles, for working with me throughout the last several years. Dr. Stiles taught me to look at the big picture while concurrently challenging me to be precise in my research. He has also provided invaluable friendship and advice, especially concerning the joys and trials of being a new Assistant Professor.

My undergraduate and M.S. advisor, Dr. Richard Plumb, also deserves my deepest gratitude. Ever since I have known him, Dr. Plumb has been an advocate and friend. I owe him many thanks for his professional and personal advice over the years.

I thank Madison A. and Lila Self for providing the fellowship that supported me during my Ph.D. research. Their endowment at the University of Kansas provided academic freedom, resources for additional career and professional enrichment, and personal development beyond what is normally provided in a Ph.D. program.

I thank Carl Leuschen for his friendship during all of my time at the University of Kansas. Life at the Radar Systems and Remote Sensing Lab certainly would have been less enjoyable if Carl had not been around.

I also thank my parents. Their own graduate degrees implicitly emphasized the importance of education. Moreover, they have always provided love and support, even when my graduate studies delayed their becoming Grandma and Grandpa.

Last, I thank my wife, Abby, for her support and sacrifices during the last several years. During my Ph.D. program, Abby excelled in her own job, all the while knowing that her career would likely be uprooted when I graduated. She uplifted me when I was frustrated, and applauded me when things went well. I hope I can repay her for her support and sacrifices many times over.

TABLE OF CONTENTS

ACKNOWLEDGEMENTS	II
LIST OF FIGURES AND TABLE	V
ABSTRACT.....	XII
1. INTRODUCTION.....	1
2. THE 2D SYNTHETIC APERTURE AND CHARACTERIZATION OF MULTIPLE RECEIVE APERTURE RADAR SYSTEMS ...	8
2.1. Signal-Space Representation of the Radar System	10
2.2. 2D Synthetic Aperture.....	15
2.2.1. Derivation	15
2.2.2. The Transformation Matrices and their Physical Interpretation.....	22
2.2.3. Sidelooking Case.....	24
2.3. The System Ambiguity Function.....	27
2.3.1. Resolution	27
2.3.1.1. Constant Correlation Approach to Resolution.....	30
2.3.1.2. Constant Error Bound Approach to Resolution	31
2.3.2. Resolution Simulations.....	35
2.3.3. Ambiguities	43
2.3.4. Synthetic Array and Synthetic Co-array Simulations ...	45
2.4. Summary	58
3. SAR.....	59
3.1. Introduction to Multiple Aperture SAR.....	59
3.2. Linear SAR Processing.....	64
3.3. Single-Aperture SAR.....	66
3.3.1. Time-Bandwidth Limitation on Illumination Area.....	66
3.3.2. Single-Aperture SAR Simulations.....	67
3.4. Multiple-Aperture SAR.....	71
3.4.1. Correlation Processing	71

3.5. <i>Sparse Arrays</i>	79
3.5.1. Maximum Likelihood Processing.....	85
3.5.2. Maximum a Posteriori Processing.....	89
3.5.3. Minimum Mean-Squared Error Processing.....	91
3.6. <i>Numerical Performance</i>	97
3.7. <i>Bandwidth and CPI Length versus</i> <i>Constellation Size</i>	102
3.8. <i>SNR Analysis</i>	107
3.9. <i>Summary</i>	115
4. MOVING TARGET INDICATION	117
4.1. <i>Introduction to MTI, DPCA, and STAP</i>	117
4.1.1. Space-Time Filtering and DPCA.....	119
4.1.2. Conventional Space-Time Adaptive Processing	122
4.1.3. STAP Performance Metrics	124
4.1.4. Minimum Requirements for MTI.....	126
4.2. <i>Clutter Rank</i>	129
4.2.1. Brennan's Rule	129
4.2.2. The Space-Bandwidth Product	132
4.2.3. Estimation of Sensor Size	138
4.2.4. Clutter Rank versus Constellation Diameter and Density	148
4.2.5. Synthetic Co-Array Bound on Clutter Rank.....	151
4.3. <i>MTI Filtering</i>	154
4.3.1. Derivation of Orthogonal Projection Filter.....	154
4.3.2. Sparse-Array MTI Simulations: Small Illumination	159
4.3.3. Sparse-Array, Wide-Area MTI	171
4.3.4. Angle Ambiguity Resolution	181
4.4. <i>Summary</i>	189
5. CONCLUSIONS AND FUTURE WORK	193
5.1. <i>Summary</i>	193
5.2. <i>Future Work</i>	199
5.3. <i>Conclusions</i>	201
REFERENCES	203

LIST OF FIGURES AND TABLE

Figure 2-1. Radar geometry for a constellation of radar satellites with the flat-Earth approximation.....	10
Figure 2-2. The vector, \mathbf{l} , defines the position of a differential piece of the transmitting antenna with a surface defined by S_A	20
Figure 2-3a. The 3D ambiguity function and predicted ellipse for a sidelooking scenario with resolution controlled by time and bandwidth.....	36
Figure 2-3b. The 2D ambiguity function and predicted ellipse for a sidelooking scenario with resolution controlled by time and bandwidth.....	36
Figure 2-4a. The 3D ambiguity function and predicted ellipse for a sidelooking scenario with resolution controlled by the array pattern.	38
Figure 2-4b. The 2D ambiguity function and predicted ellipse for a sidelooking scenario with resolution controlled by the array pattern.	38
Figure 2-5a. The 3D ambiguity function and predicted ellipse for a forward-looking scenario with resolution controlled by time and bandwidth.....	39
Figure 2-5b. The 2D ambiguity function and predicted ellipse for a forward-looking scenario with resolution controlled by time and bandwidth.....	39
Figure 2-6a. The 3D ambiguity function and predicted ellipse for a forward-looking scenario with resolution controlled by the array pattern.	41
Figure 2-6b. The 2D ambiguity function and predicted ellipse for a forward-looking scenario with resolution controlled by the array pattern.	41
Figure 2-7a. The 3D ambiguity function and predicted ellipse for a forward-looking scenario with resolution affected by time, bandwidth, and the physical array. Each of the sensor parameters has a uniform taper.	42
Figure 2-7b. The 2D ambiguity function and predicted ellipse for a forward-looking scenario with resolution affected by time, bandwidth, and the physical array. Each of the sensor parameters has a uniform taper.	42

Figure 2-8. The synthetic array for a dual-aperture SAR system with PRF matched to array element spacing	45
Figure 2-9. The synthetic co-array for a dual-aperture SAR system with PRF matched to array element spacing.	46
Figure 2-10. The ambiguity function calculated from the synthetic array of Figure 2-8.....	46
Figure 2-11. Doppler cut of the ambiguity function in Figure 2-10.....	47
Figure 2-12. The synthetic array for a dual-aperture SAR system with PRF unmatched to array element spacing.	48
Figure 2-13. The synthetic co-array for a dual-aperture SAR system with PRF unmatched to array element spacing.	49
Figure 2-14. The ambiguity function calculated from the synthetic array of Figure 2-12.....	49
Figure 2-15. Doppler cut of the ambiguity function in Figure 2-14.....	50
Figure 2-16a. Sparse-array ambiguity function generated numerically using a sparse-array simulator.....	52
Figure 2-16b. Sparse-array ambiguity function generated using the synthetic array projection.....	52
Figure 2-17a. The 2D synthetic array for a small microsat constellation.....	53
Figure 2-17b. The 2D synthetic co-array for a small microsat constellation.....	53
Figure 2-17c. The ambiguity function of the sensor shown in Figs. 2-17a and 2-17b.	54
Figure 2-18a. The 2D synthetic array for a large, very sparse microsat constellation.	55
Figure 2-18b. The 2D synthetic co-array for a large, very sparse microsat constellation.....	55
Figure 2-18c. The ambiguity function of the sensor shown in Figs. 2-18a and 2-18b.	56
Figure 3-1. Image of the KU campus used as the input to the radar simulator.	68

Figure 3-2. The image formed by a single aperture that satisfies the minimum SAR antenna area constraint.	69
Figure 3-3. The image formed by a single aperture that is much smaller than required by the minimum SAR antenna area constraint.	70
Figure 3-4a. The synthetic array for a single-aperture system.....	72
Figure 3-4b. The synthetic co-array for a single-aperture system.	72
Figure 3-5. An aperture satisfying the minimum SAR antenna area constraint is divided into multiple elements of an array. The large aperture was used to generate Fig. 3-2 while the array was used to generate Fig. 3-7.	73
Figure 3-6a. The synthetic array for the divided aperture of Fig. 3-5.	74
Figure 3-6b. The synthetic co-array for the divided aperture of Fig. 3-5.	74
Figure 3-7. The matched-filter result when the KU campus image is applied to the divided aperture of Figure 3-5.	75
Figure 3-8. The range-Doppler ambiguity function for a single element of the divided aperture in Figure 3-5.	77
Figure 3-9. The array pattern for the divided aperture of Figure 3-5.	78
Figure 3-10. The total ambiguity function of the divided-aperture system is the product of the range-Doppler ambiguity function in Figure 3-8 and the array pattern in Figure 3-9.	78
Figure 3-11. The locations of the apertures in a sparsely populated, irregularly spaced array.	79
Figure 3-12. The result of matched-filter processing for the sparse satellite array. ..	80
Figure 3-13a. A part of the synthetic array for the sparse array system.	81
Figure 3-13b. A part of the synthetic co-array for the sparse array system.	81
Figure 3-14. The array pattern for the sparse array formed by a constellation of microsats.	83
Figure 3-15. The total ambiguity function for the sparse array system.	83
Figure 3-16. The matched-filter result for the sparse array and varying SNR.	84
Figure 3-17. The ML-filter result for the sparse array and varying SNR.	88

Figure 3-18. The MMSE-filter result for the sparse array and varying SNR.	94
Figure 3-19. The matched, ML, and MMSE filters in a side-by-side comparison versus SNR.	95
Figure 3-20. Correlation, ML, and MMSE filter performance versus input SNR for a 12-receiver, sparse, random array.	98
Figure 3-21. Correlation, ML, and MMSE filter performance versus antenna positioning error for a sparse, random array and infinite SNR.	99
Figure 3-22. Correlation, ML, and MMSE filter performance versus number of receive apertures for a sparse, random array and moderate SNR.....	100
Figure 3-23. The array pattern for ML spatial filtering.	103
Figure 3-24. The total ambiguity function for ML spatial filtering.	103
Figure 3-25. The array pattern for ML spatial filtering of a very large microsat constellation.....	104
Figure 3-26. The total ambiguity function for ML spatial filtering of a very large microsat constellation.	105
Figure 3-27. Close-up of a range-Doppler ambiguity.....	106
Figure 3-28. Close-up of a range-Doppler ambiguity that is partially rejected by a null of the ML spatial filter.	106
Table 3-1. RADARSAT system parameters.....	108
Figure 3-29. SNR gain and processing loss in dB for varying number of receivers in elevation.	111
Figure 3-30. SNR gain and processing loss in dB for varying number of satellites in a microsat constellation.	113
Figure 4-1. Angle-Doppler clutter spectrum and moving target for a sidelooking scenario.....	121
Figure 4-2. Effective array positions for a two-element, two-pulse, DPCA system.	122
Figure 4-3. Angle-Doppler filter response for the optimum filter.	124
Figure 4-4. Output SINR vs. target Doppler for a target at zero azimuth angle.....	125

Figure 4-5. A sample clutter spectrum under the assumptions typically assumed for sidelooking STAP.....	131
Figure 4-6. A sample clutter spectrum with rank predicted by Brennan’s rule for non-integer β	131
Figure 4-7. The eigenvalues of a random signal’s covariance matrix follow the signal’s power spectral density.....	135
Figure 4-8. The 2D synthetic array and sensor boundary for a sparse, five-aperture system.....	139
Figure 4-9. Theoretical vs. actual clutter power contained within a specified clutter rank when the sensor area is approximated using a piecewise, visual approach.....	140
Figure 4-10. Theoretical vs. actual clutter power contained within a specified clutter rank when the sensor area is approximated using the synthetic array’s sampled 2 nd moments.....	143
Figure 4-11. Theoretical vs. actual clutter power contained within a specified clutter rank when the sensor area is approximated using the synthetic array’s sampled 2 nd moments. Hanning windows were used for the time and frequency tapers.....	144
Figure 4-12. Theoretical vs. actual clutter power versus clutter rank when the sensor area is approximated using the 2 nd moments of the five sensor parameters directly. Hanning windows were used for the time and frequency tapers.....	147
Figure 4-13. The eigenspectrum for a sensor with Hanning time and frequency windows. The illumination pattern was also Hanning. The average clutter power to the left of the dashed line is 85 percent of the total average clutter power.....	148
Figure 4-14. Predicted clutter rank vs. constellation size for constellations with five, eight, and 12 receivers. The dashed lines indicate the number of measurements obtained by each constellation.....	149

Figure 4-15. The projection filter is equivalent to replacing the true eigenspectrum with a uniform eigenspectrum of width determined by clutter rank.	158
Figure 4-16. Output SINR vs. target Doppler for filters with varying levels of clutter rejection.	160
Figure 4-17. Output SINR vs. target azimuth angle for filters with varying levels of clutter rejection.	161
Figure 4-18. Output SINR vs. target azimuth and elevation for filters with varying levels of clutter rejection.	162
Figure 4-19. Output SINR vs. target azimuth and Doppler for filters with varying levels of clutter rejection.	163
Figure 4-20. Output SINR vs. target Doppler for constellations of varying size. ...	164
Figure 4-21. Output SINR vs. target azimuth for constellations of varying size.....	165
Figure 4-22. Output SINR vs. target azimuth and elevation for constellations of varying size.....	165
Figure 4-23. Output SINR vs. target Doppler for constellations with varying number of satellites.....	166
Figure 4-24. Output SINR vs. target azimuth for constellations with varying number of satellites.....	167
Figure 4-25. Output SINR vs. target azimuth and elevation for constellations with varying number of satellites.	167
Figure 4-26. Output SINR vs. target Doppler for constellations with varying levels of satellite positioning errors.	169
Figure 4-27. Output SINR vs. target azimuth for constellations with varying levels of satellite positioning errors.	169
Figure 4-28. Output SINR vs. target azimuth and elevation for constellations with varying levels of satellite positioning errors.	170
Figure 4-29. The intersection of constant range lines with the spectral bandwidth contour defines the clutter subspace. The bold range lines contribute to the clutter subspace.	173

Figure 4-30. SAR map made from data that includes four moving targets.	174
Figure 4-31. MTI map for moving targets with a very low radial velocity.	175
Figure 4-32. MTI map for a radial velocity that is still lower than the slowest target.	175
Figure 4-33. Histogram and Rayleigh curve fit for the pixel magnitudes of Fig. 4-32.	176
Figure 4-34. The result of applying a threshold to Fig. 4-32. The threshold was set for a P_{fa} of 10^{-6} using the Rayleigh curve fit of Fig. 4-33.	177
Figure 4-35. MTI map for the radial velocity of the slowest target. Target location is indicated by the arrow.	178
Figure 4-36. The result of applying a threshold to Fig. 4-35. The threshold was set for a P_{fa} of 10^{-6} . The location of the moving target is indicated.	179
Figure 4-37. MTI map containing two moving targets.	179
Figure 4-39. MTI map for the highest radial velocity of the four targets.	180
Figure 4-40. The result of applying a threshold to Fig. 4-39. The threshold was set for a P_{fa} of 10^{-6}	181
Figure 4-41. The along-track-position/radial-velocity ambiguity function generated from time data only.	182
Figure 4-42. The along-track-position/radial-velocity ambiguity function generated from spatial data only.	183
Figure 4-43. The total along-track-position/radial-velocity ambiguity function.	184
Figure 4-44. Matched-filter output vs. post-detection ML-filter output for 31 detections obtained from Figs. 4-34, 4-36, 4-38, and 4-40.	186
Figure 4-45. Probability of false alarm for the matched filter vs. the post-detection ML filter for Figs. 4-34, 4-36, 4-38, and 4-40.	187
Figure 4-46. Probability of detection for the matched filter vs. the post-detection ML filter for Figs. 4-34, 4-36, 4-38, and 4-40.	188

ABSTRACT

The concept of radar satellite constellations has been proposed in the literature and is currently under research. These constellations possess the wide angular coverage necessary for performing spaceborne MTI and the spatial sampling necessary for wide-area SAR. However, the constellations form arrays that are sparsely populated, three-dimensional, and irregularly spaced. Therefore, current SAR and MTI processing algorithms, which often assume a single element or a uniform, linear array, are not well suited to this type of implementation.

In order to derive effective SAR and MTI processing techniques, I develop a novel method of characterizing the behavior of sparse-array radar. A radar's five sensor parameters-time, frequency, and 3D spatial location-are equivalently represented by a 2D synthetic aperture. This synthetic aperture is then used to characterize radar properties such as resolution and ambiguity. In addition, I use the synthetic aperture as the basis for a new method of estimating clutter rank when performing MTI. The synthetic aperture is crucial to accurately estimating clutter rank using this technique.

From the synthetic aperture, effective SAR and MTI algorithms are developed. It becomes apparent that traditional SAR matched filtering produces poor results due to the sidelobes of the sparse, irregular array; therefore, other filters are derived to handle the sparse-array case. Simulations are performed for each filtering technique, the performance of the filters versus various system parameters is compared, and impacts of the constellation concept on system design are inferred from the results.

Next, the synthetic aperture is used to estimate clutter rank and define a clutter subspace. MTI is then performed through generalized DPCA filtering, which projects the radar data orthogonal to the predicted clutter subspace. Although DPCA is a mature concept, application of DPCA to a sparse, spaceborne array is new. MTI performance is also analyzed versus system parameters, and conclusions are made concerning the design limitations on sparse, spaceborne MTI systems.

1. INTRODUCTION

The theme of the 2001 IEEE Radar Conference, *2001: Radar's Odyssey into Space* [1], indicates the current emphasis being placed on transferring radar technology to spaceborne platforms. Recent literature has also emphasized space-based radar [2-7] and its advantages for reconnaissance, remote sensing, and moving target indication (MTI). It is well known that, depending on frequency, radar performs well in a wide variety of weather conditions. While clouds and rain will undoubtedly be encountered by any spaceborne or high-altitude system, they have relatively little effect on radars operating in the lower GHz frequencies unless the rain is extremely strong [8-9]. Spaceborne radars also do not require onboard personnel. Therefore, there is no flight staff to maintain, and no personnel are placed in danger during data acquisition. Moreover, the radar satellite itself is relatively safe. During military conflicts, airborne intelligence-gathering vehicles are prime targets, but once a radar satellite is placed into orbit, it is less vulnerable to destruction through accident or military threat.

The most significant advantage of spaceborne radar, however, is its ability to provide global, 24-hour coverage. Radars on airborne platforms are limited by airspace restrictions. The best they can do is fly along the boundary of restricted airspace while transmitting into and receiving from the restricted area, and during

times of hostility, even this approach is dangerous or impossible. Radar satellites are not subject to airspace restrictions and from their high altitudes can access the polar regions and all areas in between [10]. Therefore, spaceborne radar provides the possibility for continuous imaging, moving target detection, and tracking of targets anywhere on the globe. With the importance placed in today's world on access to quality, timely intelligence and environmental information, spaceborne radar's ability to provide global, 24-hour coverage is a very important benefit.

There are also disadvantages, however, to spaceborne radar, the most significant being the tradeoff between sensor weight, size, and power. It is known that the power density of a propagating wave is inversely proportional to the square of distance; hence, the large distances involved in spaceborne radar necessitate that the main driver in the weight-size-power tradeoff be signal-to-noise ratio (SNR). In order to compensate for the large distances, spaceborne sensors must have large *power-aperture* products. Higher transmitted *power* improves SNR by increasing the initial power density. Large *apertures* improve SNR on both transmit and receive. On transmit, the power density is improved through the gain of a large aperture. On receive, SNR is improved by collecting energy over a large effective area. Therefore, the SNR problem could be overcome by designing systems with high transmit powers and/or large antennas.

Other factors requiring large apertures are ambiguity and resolution. For synthetic aperture radar (SAR), the minimum antenna size is governed by the minimum SAR antenna area constraint [11-17]. Complying with this constraint ensures that range-Doppler ambiguities are not illuminated. For MTI, high spatial resolution is necessary in order to separate moving targets, especially slow-moving ones, from ground clutter. Ground moving targets such as cars and tanks provide a difficult scenario when observed by spaceborne radar because the component of the moving target's velocity in the direction of an overhead satellite is small. Because its radial velocity is small, the moving target will not be located far from ground clutter having the same Doppler frequency. Therefore, high spatial resolution is required to

resolve targets from clutter having the same Doppler but slightly different angle of arrival. This spatial resolution can be achieved through either a large single aperture or an array, but in either case the satellite's receiving structure must be large.

Unfortunately, the requirements for high transmit power and large antennas conflict with available technology. First, the high power that is to be transmitted must come from the satellite itself, either through an onboard power supply or a power-generating system such as solar panels. Any increase in the transmitted power requirement further taxes a satellite commodity that is always in short supply, and large power supplies can add considerable weight to a satellite [18]. Second, large antennas significantly increase the difficulty of satellite launch and deployment. A large antenna may be heavy, which increases the size and cost of the launch vehicle. The area of the antenna may also exceed the capacity of most launch vehicles, which means that the antenna must be deployed once in orbit.

The large aperture requirement also conflicts with search and revisit rate requirements. Since antenna beamwidth and size are inversely related, a large aperture size translates to a relatively small illumination area. Since the illumination area at any instant in time is limited, the coverage that a single satellite can provide in a single pass is also limited. If global coverage is required and a satellite can only cover a limited area in a single pass, many passes may be required before the radar can return to a previously viewed area. This is unacceptable for remote sensing applications such as monitoring of flooding or oil spills, or for providing real-time tactical information to the military. One common method, ScanSAR [11, 16-17, 19-22], improves revisit rate, but does so at the cost of degraded Doppler resolution.

Similarly, a large aperture limits MTI search rate. As mentioned, a large effective aperture is needed in order to distinguish moving targets from ground clutter. Since a large effective aperture corresponds to a relatively small illumination area, more beams are required to cover a specified search area, and the need to scan the aperture's beam limits the area that can be searched in a given amount of time.

One approach to spaceborne radar intended to mitigate the weight-size-power tradeoff is to implement a cluster of satellites called a constellation [23-27]. Each satellite has a coherent receiver, and the constellation flies in formation to create an antenna array. The individual satellites in this concept are small, with relatively small apertures; hence, they are termed microsats. The data from each microsat and the spatial sampling of the constellation are combined to form a single, virtual radar.

Several advantages are achieved by breaking the radar system into multiple, free-flying components. First, the effective aperture of the system is still determined by its total energy-collecting area, but that area is now divided between many satellites. The sum of the aperture sizes of all the satellites determines the total effective aperture of the system, but it is more cost effective to launch and deploy a constellation of formation-flying microsats than to launch a similar monolithic satellite [20, 28-29]. Therefore, sufficient SNR can be achieved through a large effective aperture, but without the problems associated with the launch and deployment of a large, single satellite. Sufficient SNR can be achieved through increased aperture area rather than by increased transmit power, with the system's effective aperture increasing as microsats are added to the constellation. Some of the cost savings of microsat constellations come through mass production of identical microsats and by using the small size of the microsats to optimize launch vehicle payloads.

The second advantage is in the added spatial sampling. For SAR, it is possible that the spatial samples obtained from multiple receivers can be used to distinguish between range-Doppler ambiguities since ambiguous cells have different angles of arrival. For MTI, the increased illumination area results in a higher search rate without sacrificing the angular resolution necessary for detecting slow-moving targets. While the search rate will be determined by the illumination pattern of the smaller, individual apertures, the angular resolution will be determined by the size of the constellation. With proposed constellation sizes being on the order of hundreds of meters or more, the angular resolution provided by such a system would be very high.

Other advantages of satellite constellations include graceful performance degradation and ease of replacing or upgrading satellites. The performance of satellite constellations degrades gracefully compared to monolithic systems. If a major subsystem of a monolithic system fails, the satellite is rendered useless. If a microsat fails, however, it is only one of an entire constellation. Performance will be degraded due to the loss of some spatial information, but certainly not to the point of entire system failure. When microsats do fail, they can be replaced as time and cost allow. Moreover, they can be replaced by upgraded versions, and cost savings are again achieved through mass production of the replacements. Constellations can also be augmented with additional microsats for improved performance, and further cost savings can be realized by *piggybacking* replacement microsats on launch vehicles that have available payload capacity.

The disadvantage of a constellation of formation-flying satellites is that the array formed by the constellation will be sparsely populated, three-dimensional, and non-uniformly spaced. The laws of orbital physics govern microsat orbits, and a safe distance must be maintained such that microsats don't collide. It may be possible to temporarily propel the microsats into a regular, well-formed pattern, but the amount of propulsion necessary to maintain a regular formation puts this option beyond any reasonable fuel budget. Therefore, the radar designer has essentially no control of the array structure formed by the constellation except, possibly, for some input about the general, overall size of the constellation.

Unfortunately, traditional SAR and MTI algorithms assume uniform, linear arrays aligned along the radar system's direction of travel. As a result, signal processing algorithms for performing SAR and MTI on data obtained from sparse satellite clusters have not been developed. Some work has been done on multi-channel SAR [22, 30], but this work still assumes linear or planar arrays and uniform array spacing. Many advantages of satellite clusters were discussed in the preceding paragraphs, but until it is determined how best to process the data obtained by this unique radar concept, those advantages cannot be realized.

In the following chapters of this dissertation, I investigate the general question of how to process the data obtained from sparsely populated, irregularly spaced, 3D radar arrays. Because this radar concept is so unique, I begin in Chapter 2 with the radar model that is the foundation used to build effective signal processing algorithms. Then, I show that the data collected by a satellite constellation can be collected by an equivalent, 2D, synthetic aperture. Just as along-track resolution is determined in traditional SAR by the length of the 1D synthetic aperture, the 2D synthetic aperture also enables evaluation of important system characteristics such as resolution, ambiguity, and clutter width. I will show that the synthetic aperture concept is useful for evaluating sparse array performance, just as it is for traditional, sidelooking SAR.

In Chapter 3, SAR processing methods for sparse satellite constellations are investigated. The simulations begin with a uniform, planar array. Once it is shown that adding coherent apertures enables an expanded illumination area, the elements of the array are spread apart randomly to simulate a satellite cluster. I will demonstrate that matched filtering produces unsatisfactory results when applied to the satellite cluster. Consequently, maximum likelihood (ML) and minimum mean-squared error (MMSE) filtering are introduced to handle the sparse spatial sampling. Then, since ML and MMSE filtering require increased computation, a method of applying these filters to only the spatial data is discussed. Simulation results are presented to show the performance and behavior of matched, ML, and MMSE filtering. Last, simulation results are presented to show SAR performance using these methods versus various parameters such as the number of satellites in the constellation, SNR, and error in the knowledge of aperture positions.

MTI is investigated in Chapter 4. First, I present an introduction to moving target indication, space-time adaptive processing (STAP), and appropriate performance metrics. Then, the 2D synthetic aperture from Chapter 2 is used to characterize the severity of ground clutter. This characterization, also known as clutter rank, is in terms of the number of measurement degrees of freedom needed for rejecting a

predetermined fraction of the expected clutter power. Once the clutter is characterized, an MTI filter is derived that removes clutter power by projecting the measured data orthogonal to the clutter subspace. The performance of this filter versus system parameters such as the number of satellites in the constellation, size of the constellation, clutter rank, and aperture positioning error is examined. Finally, a wide-area MTI simulation is performed where the angle information provided by the satellites must be used to resolve range-Doppler ambiguities in addition to separating moving targets from ground clutter. In Chapter 5, I make my conclusions and recommendations for future work.

2. THE 2D SYNTHETIC APERTURE AND CHARACTERIZATION OF MULTIPLE RECEIVE APERTURE RADAR SYSTEMS

Radar implemented via a constellation of formation-flying satellites is a paradigm shift from the radar implementations of yesterday and today. Although multi-channel systems have been used for both MTI and SAR [22, 31-33], they have been implemented on airborne systems and have not used sparse arrays. Therefore, there are several basic issues and questions that are brought up by the constellation concept. One of the most basic issues is how to evaluate a sparse radar system's resolution. For single-aperture, sidelooking systems, there are simple relationships between bandwidth, the length of the coherent processing interval (CPI), and resolution. With large, sparse arrays, however, the beamwidth of the physical array can have an effect on resolution. Likewise, satellite clusters will scan forward and backward, but resolution is not easily determined for non-sidelooking scenarios. Hence, a method is needed for evaluating resolution for sparse arrays and various look geometries.

In addition to resolution, or the width of the ambiguity function's mainlobe, sidelobe levels, sidelobe structure, and the location of grating lobes are also important. While these features can be characterized for single-aperture, sidelooking systems based on knowledge of window functions and sampling theory, they are not easily characterized for non-sidelooking systems and systems employing sparse,

randomly sampled arrays. Therefore, a method is also needed that characterizes sidelobe levels and ambiguities of such systems.

The question of how to include sparse spatial data in SAR processing algorithms also arises. The beamforming pattern for satellite constellations will appear random and may contain high sidelobes; therefore, it is not clear if traditional matched filtering can produce satisfactory results. A succinct, convenient expression for the measured data is needed so that algorithms other than matched filtering can be derived. Also, if the new algorithms are computationally expensive, a succinct expression could aid in implementing computationally reduced versions of the algorithms.

Last, the clutter spectrum must be characterized in order to perform MTI. Traditional MTI uses a uniform, linear array and periodic pulses to detect moving targets in a single range bin. The degrees of freedom (DOFs) in the observed clutter spectrum, which are a measure of the clutter's severity, are well known for this case and have been reported many times in the literature [34-37]. For the constellation concept, however, the clutter spectrum is not known. Clearly, the clutter spectrum needs to be characterized so that it can be filtered out and moving targets can be detected.

The derivations in this chapter provide the foundation for addressing the above issues. In this chapter, the measured data are first expressed in matrix-vector notation, which leads to straightforward application of algorithms from estimation theory. These algorithms will be necessary for performing SAR and MTI using satellite clusters. Then, in order to characterize important system parameters such as resolution, ambiguity, and the spectral width of clutter, a data transformation is derived. I will demonstrate that the radar measurements obtained across the five dimensions of time, frequency, and 3D space can be equivalently represented by data obtained through a two-dimensional synthetic aperture. Therefore, system characterization can be achieved by working on data that span two dimensions rather than five.




Figure 2-1. Radar geometry for a constellation of radar satellites with the flat-Earth approximation.

2.1. Signal-Space Representation of the Radar System

The radar geometry is shown in Fig. 2-1. The spaceborne system travels in the positive x -direction at velocity, v , and the array phase reference at time zero is located at the origin of the coordinate system. Therefore, assuming a flat earth, the z -coordinate of all targets on the ground is $-h$, where h is the altitude of the array phase reference. There is a single transmitter located at the array phase reference.

Using Fig. 2-1, some variables can be defined. The vector defining the position of a spot on the Earth's surface, \mathbf{x} , is given by $\mathbf{x} = [x \ y \ -h]^\dagger$ where $(\cdot)^\dagger$ denotes the matrix or vector transpose operation. Since the transmitter is located at the array phase reference, which passes through the origin at $t = 0$, its position at any given time is $\mathbf{r}_t = \mathbf{v}t$ where \mathbf{v} is the radar system's velocity vector, $\mathbf{v} = [v \ 0 \ 0]^\dagger$. The location of a receiver at $t = 0$ is defined as $\mathbf{r} = [r_x \ r_y \ r_z]^\dagger$.

In terms of the above position and velocity vectors, the range from transmitter to scatterer at time, t , is

$$R_{tx}(\mathbf{x}, t) = |\mathbf{v}t - \mathbf{x}|. \quad (2.1)$$

Likewise, the range from the point \mathbf{x} to a receiver at time, t , is given by

$$R_{rx}(\mathbf{x}, \mathbf{r}, t) = |\mathbf{r} + \mathbf{v}t - \mathbf{x}|, \quad (2.2)$$

and the round-trip delay is given by

$$\tau(\mathbf{x}, \mathbf{r}, t) = \frac{1}{c} [R_{tx}(\mathbf{x}, t) + R_{rx}(\mathbf{x}, \mathbf{r}, t)] \quad (2.3)$$

where c is the speed of light. Let the transmitting antenna gain versus scatterer location be defined as $G(\mathbf{x})$, the complex envelope of the transmitted signal be defined as $s(t')$, and the carrier frequency be defined as ω_0 . Then, the signal measured by a radar receiver at position, \mathbf{r} , and time, t , due to unit scatterer located at \mathbf{x} is

$$\rho(\mathbf{x}, \mathbf{r}, t) = \frac{G(\mathbf{x})}{R_{tx}(\mathbf{x}, t)R_{rx}(\mathbf{x}, \mathbf{r}, t)} \int_T \exp[-j\omega_0\tau(\mathbf{x}, \mathbf{r}, t)] \delta[t - t' - \tau(\mathbf{x}, \mathbf{r}, t)] s(t') dt' \quad (2.4)$$

where T is the length of the transmit signal and $\delta(\cdot)$ is the delta function. Next, noting that the denominator of (2.4) is an amplitude component that varies little over a radar's time and space dimensions, the denominator is approximated as $R_{tx}(\mathbf{x}, t)R_{rx}(\mathbf{x}, \mathbf{r}, t) \approx R_{tx}(\mathbf{x}, 0)R_{rx}(\mathbf{x}, \mathbf{0}, 0) \approx R(\mathbf{x})^2$. Using the sifting property of the delta function, (2.4) becomes

$$\rho(\mathbf{x}, \mathbf{r}, t) = \frac{G(\mathbf{x})}{R(\mathbf{x})^2} \exp[-j\omega_0\tau(\mathbf{x}, \mathbf{r}, t)] s[t - \tau(\mathbf{x}, \mathbf{r}, t)]. \quad (2.5)$$

Equation (2.5) contains all the components of the radar physics and expresses the complex measurement that would be obtained by a radar at position \mathbf{r} and time t due to a unit scatterer at \mathbf{x} . The received signal is a delayed version of the transmit signal with a phase component that depends on the time-varying propagation delay. The amplitude of the radar response is inversely proportional to the square of the one-way range to the scatterer and proportional to the gain of the transmitting antenna. Henceforth, $\rho(\mathbf{x}, \mathbf{r}, t)$ is referred to as the *radar response* due to a unit scatterer.

The total measurements taken by a radar system are due to all illuminated scatterers. Therefore, the actual complex data samples, $d(\mathbf{r}, t)$, sampled at some point in space and time by a radar system's analog-to-digital converters can be obtained by integrating the weighted responses of all illuminated scatterers. Expressing this integration and adding noise, the complex data samples are given by

$$d(\mathbf{r}, t) = \int_A \gamma_0(\mathbf{x}) \rho(\mathbf{x}, \mathbf{r}, t) d\mathbf{x} + n(\mathbf{r}, t) \quad (2.6)$$

where $\gamma_0(\mathbf{x})$ is the complex reflectance per unit area at \mathbf{x} , A is the radar's illumination area, and $n(\mathbf{r}, t)$ is complex noise. In (2.6), the radar response due to a scatterer at \mathbf{x} is scaled by the scatterer's complex reflection coefficient. The weighted

responses for all illuminated scatterers are totaled together and noise is added to arrive at the complete, continuous-time, continuous-space received signal.

If the integration in (2.6) is approximated by a summation, the received signal becomes

$$d(\mathbf{r}, t) = \sum_i \gamma_0(\mathbf{x}_i) \rho(\mathbf{x}_i, \mathbf{r}, t) \Delta A + n(\mathbf{r}, t), \quad (2.7)$$

where the index, i , includes all differential areas of size, ΔA , that are illuminated and \mathbf{x}_i is the position vector to the center of the i^{th} differential area. If the signal is sampled by the radar in space and time, the m^{th} sample at the n^{th} receiver is

$$d(\mathbf{r}_n, t_m) = \sum_i \gamma_0(\mathbf{x}_i) \rho(\mathbf{x}_i, \mathbf{r}_n, t_m) \Delta A + n(\mathbf{r}_n, t_m). \quad (2.8)$$

The received signal is sampled in time at a rate equal to the signal's bandwidth, and the sampling interval is the length of the pulse train, T . Therefore, the number of time samples obtained by a single receiver is equal to the signal's time-bandwidth product at that receiver. Finally, the entire set of measurements can be represented using matrix-vector notation,

$$\mathbf{d} = \mathbf{P}\boldsymbol{\gamma} + \mathbf{n}, \quad (2.9)$$

where

$$\mathbf{d} = [d(\mathbf{r}_1, t_1) \ \cdots \ d(\mathbf{r}_1, t_{BT}) \ d(\mathbf{r}_2, t_1) \ \cdots \ \cdots \ d(\mathbf{r}_N, t_{BT})]^\dagger, \quad (2.10)$$

$$\mathbf{P} = [\boldsymbol{\rho}_1 \ \boldsymbol{\rho}_2 \ \cdots \ \boldsymbol{\rho}_C], \quad (2.11)$$

$$\boldsymbol{\gamma} = [\gamma_1 \ \gamma_2 \ \cdots \ \gamma_C]^\dagger, \quad (2.12)$$

$$\mathbf{n} = [n(\mathbf{r}_1, t_1) \ \cdots \ n(\mathbf{r}_N, t_{BT})]^\dagger, \quad (2.13)$$

$$\boldsymbol{\rho}_i = [\rho(\mathbf{x}_i, \mathbf{r}_1, t_1) \ \cdots \ \rho(\mathbf{x}_i, \mathbf{r}_N, t_{BT})]^\dagger, \quad (2.14)$$

and

$$\gamma_i = \gamma_0(\mathbf{x}_i). \quad (2.15)$$

Also, N is the number of satellites in the constellation, BT is the received time-bandwidth product, and C is the number of differential areas of size ΔA in the summation of (2.8). Equation (2.9) expresses the radar physics in a concise, linear form that will be useful for deriving SAR algorithms in Chapter 3. It shows that the measurements obtained by the radar are a result of the scattering from many resolution cells on the ground. Equation (2.9) also demonstrates how to model the radar system in my simulations. A set of complex scattering coefficients is assigned to a grid of resolution cells, and the radar response from each resolution cell is calculated and weighted by the corresponding scattering coefficient. The sum of all the scaled responses becomes the simulated data vector. However, the size of each cell and the number of cells required to properly model the radar measurements have yet to be determined.

The vector described by (2.14) is an important vector. Previously, the scalar value, $\rho(\mathbf{x}, \mathbf{r}, t)$, was derived as the radar response at any point in space and time due to a unit-reflectance scatterer at \mathbf{x} . The *response vector*, $\boldsymbol{\rho}_i$, is the full set of measurements obtained by a radar at a discrete set of space and time locations, again due to a unit reflectance scatterer. If these response vectors are weighted by their appropriate scattering values and added together, the result is the set of measurements obtained by a radar due to all scatterers within the illuminated area.

2.2. 2D Synthetic Aperture

2.2.1. Derivation

When the behavior of the radar response is evaluated over the complete range of radar measurement parameters, it is possible to characterize important system parameters such as resolution and ambiguity. In this section, the time-varying frequency response of a radar system is developed because it is beneficial to investigate system behavior in both the time and frequency domains. Then, the radar response is expressed in terms of a synthetic aperture, which demonstrates that the five sensor parameters of space, time, and frequency can be projected into an equivalent two-dimensional sensor. This 2D sensor then leads to convenient characterization of radar resolution and ambiguity.

Equation (2.4) represents a linear system with an input function, $s(t')$, and a system impulse response given by

$$h_g(\mathbf{x}, \mathbf{r}, t, t') = \frac{G(\mathbf{x})}{R(\mathbf{x})^2} \exp[-j\omega_0 \tau(\mathbf{x}, \mathbf{r}, t)] \delta[t - t' - \tau(\mathbf{x}, \mathbf{r}, t)]. \quad (2.16)$$

The time-varying frequency representation [38-39] of the radar system impulse response is defined as

$$h_{gf}(\mathbf{x}, \mathbf{r}, \omega', t) = \int_{-\infty}^{\infty} h_g(\mathbf{x}, \mathbf{r}, t, t') \exp[-j\omega'(t-t')] dt'. \quad (2.17)$$

From the previous section, the radar impulse response is approximately

$$h_g(\mathbf{x}, \mathbf{r}, t, t') = \frac{G(\mathbf{x})}{R(\mathbf{x})^2} \exp[-j\omega_0 \tau'(\mathbf{x}, \mathbf{r}, t)] \delta[t - t' - \tau'(\mathbf{x}, \mathbf{r}, t)]. \quad (2.18)$$

Substituting (2.18) into (2.17), the time-varying frequency response of the radar system is

$$h_{yf}(\mathbf{x}, \mathbf{r}, \omega', t) = \frac{G(\mathbf{x})}{R(\mathbf{x})^2} \exp[-j(\omega_0 + \omega')\tau'(\mathbf{x}, \mathbf{r}, t)]. \quad (2.19)$$

Equation (2.19) describes the received signal at the radar due to a transmit signal of the form $\exp(j\omega't)$ that is mixed to a center frequency of ω_0 and reflected off a unit scatterer at \mathbf{x} .

To simplify notation, let $\omega = \omega_0 + \omega'$ and

$$\Psi(\mathbf{x}, \mathbf{s}) = \omega\tau'(\mathbf{x}, \mathbf{r}, t) \quad (2.20)$$

where \mathbf{s} is the vector of sensor parameters, $\mathbf{s} = [\mathbf{r}^\dagger \quad \omega \quad t]^\dagger$. The time-varying frequency response is then

$$h_{yf}(\mathbf{x}, \mathbf{s}) = \frac{G(\mathbf{x})}{R(\mathbf{x})^2} \exp[-j\Psi(\mathbf{x}, \mathbf{s})]. \quad (2.21)$$

From (2.19) and (2.21), it is apparent that several variables affect the phase of the radar response. The phase varies with time, space, frequency, and the location of the scatterer. The only scatterer-dependent parameters that affect how the received phase varies over the sensor are the x and y coordinates of the scatterer. Since the scatterers only have two variables that affect their response at the radar, it is reasonable to assume that only two sensor dimensions are needed for representing the SAR data. Therefore, the hypothesis is that although the phase expression of (2.21) varies versus five sensor parameters, those sensor parameters can be projected into the coordinates of two independent *eigensensors*. The projection of a sensor's time, space, and frequency parameters onto these two eigensensors forms a 2D synthetic sensor that can be used to characterize SAR and MTI performance.

I derive the synthetic sensor by performing two, first-order Taylor expansions of the phase of the radar response. First, Ψ is expanded around the radar sensor parameters, \mathbf{s} . The result is

$$\Psi(\mathbf{x}, \mathbf{s}) \approx \Psi(\mathbf{x}, \bar{\mathbf{s}}) + \left(\nabla_s \Psi \Big|_{\mathbf{x}, \bar{\mathbf{s}}} \right)^\dagger \Delta \mathbf{s} \quad (2.22)$$

where

$$\left. \begin{aligned} \bar{\mathbf{s}} &= [r_{x0} \quad r_{y0} \quad r_{z0} \quad \omega_0 \quad t_0]^\dagger \\ \nabla_s &= \left[\frac{\partial}{\partial r_x} \quad \frac{\partial}{\partial r_y} \quad \frac{\partial}{\partial r_z} \quad \frac{\partial}{\partial \omega} \quad \frac{\partial}{\partial t} \right]^\dagger \\ \Delta \mathbf{s} = \mathbf{s} - \bar{\mathbf{s}} &= [r_x - r_{x0} \quad r_y - r_{y0} \quad r_z - r_{z0} \quad \omega - \omega_0 \quad t - t_0]^\dagger \end{aligned} \right\}. \quad (2.23)$$

In (2.22), $\bar{\mathbf{s}}$ is the set of sensor parameters around which the expansion is performed. Using the array phase reference, mean sensor time, and mean frequency, the sensor parameters around which the expansion is performed are given by

$$\bar{\mathbf{s}} = [0 \quad 0 \quad 0 \quad \omega_0 \quad 0]^\dagger. \quad (2.24)$$

The first term on the right hand side of (2.22) does not vary with changes in the sensor parameters or over the antenna. Therefore, it contains no useful information and can be absorbed into $\gamma_0(\mathbf{x})$ of (2.6) as part of the random scattering phase. Then, writing out the remaining part of (2.22) using the derivatives of the Taylor expansion, the phase is

$$\begin{aligned} \Psi(\mathbf{x}, \mathbf{s}) &= \frac{\partial \Psi}{\partial r_x} \Big|_{\mathbf{x}, \bar{\mathbf{s}}} r_x + \frac{\partial \Psi}{\partial r_y} \Big|_{\mathbf{x}, \bar{\mathbf{s}}} r_y + \frac{\partial \Psi}{\partial r_z} \Big|_{\mathbf{x}, \bar{\mathbf{s}}} r_z + \frac{\partial \Psi}{\partial \omega} \Big|_{\mathbf{x}, \bar{\mathbf{s}}} (\omega - \omega_0) + \frac{\partial \Psi}{\partial t} \Big|_{\mathbf{x}, \bar{\mathbf{s}}} t \\ &= k_x r_x + k_y r_y + k_z r_z + \tau(\omega - \omega_0) + \omega_D t \end{aligned} \quad (2.25)$$

where k_x , k_y , and k_z are spatial frequencies, τ is the delay due to range, and ω_D is the Doppler frequency. By performing the Taylor expansion around sensor parameters, the phase becomes a summation of derivatives where each derivative already has a traditionally accepted physical interpretation.

Some important approximations are made when expanding the phase in terms of single derivatives as shown in (2.22) and (2.25). Range, Doppler, and spatial frequency are assumed constant across the sensor parameters because the first-order approximations are evaluated at a single point in the sensor space. Therefore, it is assumed that there is no range or Doppler walk and that the satellite constellation is in the far field of the Earth's surface. For example, consider the first term on the right hand side of (2.25), interpreted as the x -component of a pixel's spatial frequency. The derivative varies with the evaluation point, \mathbf{x} , but does not vary with the sensor parameters since it is evaluated only at the mean sensor parameters, $\bar{\mathbf{s}}$. But if an array is not in the far field, the radius of curvature of the propagating field can be large enough for this approximation to be invalid. Consequently, the arrays formed by satellite constellations could be large enough that this first-order approximation is invalid. Similar arguments apply to the variation of Doppler over time, frequency, and space, the variation of range over time and space, and the variation of spatial frequencies over time and frequency.

Some of the above approximations already have names associated with them. The far field approximation was already mentioned. Likewise, assuming that spatial frequency is constant over frequency and assuming that range is constant over space are known as narrowband assumptions. These are common assumptions, but they must be taken into account when considering the span of sensor parameters over which the synthetic aperture is valid.

Returning to (2.22), which was the phase function after expanding around the sensor and antenna parameters, the sensor-dependent component of phase is written as

$$\Psi(\mathbf{x}, \mathbf{s}) = \mathbf{k}_s(\mathbf{x})^\dagger \Delta \mathbf{s} \quad (2.26)$$

where $\mathbf{k}_s(\mathbf{x}) = \nabla_s \Psi|_{\mathbf{x}, \bar{\mathbf{s}}}$ is used to emphasize similarity with the standard wavenumber vector. Next, the first-order Taylor expansion of $\mathbf{k}_s(\mathbf{x})^\dagger$ around the position of the scatterer on the ground is performed. Defining $\bar{\mathbf{x}} = [x_0 \ y_0 \ -h]^\dagger$ as the central point of illumination on the ground and the point around which the expansion is performed, the expansion is

$$\mathbf{k}_s(\mathbf{x})^\dagger \approx \mathbf{k}_s(\bar{\mathbf{x}})^\dagger + \Delta \mathbf{x}^\dagger \left[\nabla_x \mathbf{k}_s(\mathbf{x})^\dagger \Big|_{\bar{\mathbf{x}}} \right] \quad (2.27)$$

where

$$\Delta \mathbf{x} = [x - x_0 \ y - y_0 \ 0]^\dagger \quad (2.28)$$

and

$$\nabla_x = \left[\frac{\partial}{\partial x} \ \frac{\partial}{\partial y} \ \frac{\partial}{\partial z} \right]^\dagger. \quad (2.29)$$

The received phase then becomes

$$\begin{aligned} \Psi(\mathbf{x}, \mathbf{s}) &\approx \mathbf{k}_s(\bar{\mathbf{x}})^\dagger \Delta \mathbf{s} + \Delta \mathbf{x}^\dagger \left[\nabla_x \mathbf{k}_s(\mathbf{x})^\dagger \Big|_{\bar{\mathbf{x}, \bar{\mathbf{s}}} \right] \Delta \mathbf{s} \\ &= \left(\nabla_s \Psi \Big|_{\bar{\mathbf{x}, \bar{\mathbf{s}}} \right)^\dagger \Delta \mathbf{s} + \Delta \mathbf{x}^\dagger \left[\nabla_x \left(\nabla_s \Psi \right)^\dagger \Big|_{\bar{\mathbf{x}, \bar{\mathbf{s}}} \right] \Delta \mathbf{s} \\ &= \left(\mathbf{k}_s^0 \right)^\dagger \Delta \mathbf{s} + \Delta \mathbf{x}^\dagger \Lambda_s \Delta \mathbf{s} \end{aligned} \quad (2.30)$$

where $\mathbf{k}_s^0 = \mathbf{k}_s(\bar{\mathbf{x}})$ and Λ_s is coined the *sensor transformation matrix*. Using (2.30) and (2.21), the time-frequency response of the radar is




Figure 2-2. The vector, \mathbf{l} , defines the position of a differential piece of the transmitting antenna with a surface defined by S_A .

$$h_{yf}(\mathbf{x}, \mathbf{s}) = \frac{G(\mathbf{x})}{R(\mathbf{x})^2} \exp \left\{ -j \left[(\mathbf{k}_s^0)^\dagger \Delta \mathbf{s} + \Delta \mathbf{x}^\dagger \Lambda_s \Delta \mathbf{s} \right] \right\}. \quad (2.31)$$

Next, attention is turned to the radiation pattern of the transmitting antenna, $G(\mathbf{x})$. The antenna pattern is given by

$$G(\mathbf{x}) = \int_{S_A} w_l(\mathbf{l}) \exp[-j\Psi_l(\mathbf{x}, \mathbf{l})] d\mathbf{l}. \quad (2.32)$$

In (2.32), $w_l(\mathbf{l})$ is the antenna's current or field distribution, \mathbf{l} is a vector to a point on the antenna's conducting structure or aperture, $\mathbf{l} = [l_x \quad l_y \quad l_z]^\dagger$, S_A is the surface of the conducting structure or aperture, and $\Psi_l(\mathbf{x}, \mathbf{l})$ is the relative phase shift over the antenna due to a slightly varying range to the scatterer location, \mathbf{x} . The phase shift, $\Psi_l(\mathbf{x}, \mathbf{l})$, is given by

$$\Psi_l(\mathbf{x}, \mathbf{l}) = \frac{\omega_0}{c} |\mathbf{l} - \mathbf{x}|, \quad (2.33)$$

and the vector \mathbf{l} is shown graphically in Fig. 2-2. The antenna distribution, $w_l(\mathbf{l})$, has a magnitude and phase associated with it. Therefore, it can be written as

$$w_l(\mathbf{l}) = |w_l(\mathbf{l})| \exp[-j\Psi_a(\mathbf{l})]. \quad (2.34)$$

The phase in (2.32) is similar to the phase used previously to derive the sensor response. The range to a point, \mathbf{x} , on the scattering surface varies slightly over the antenna structure. The radiation pattern is obtained by integrating this variation over the antenna structure as indicated by (2.32). Performing Taylor expansions similar to those performed earlier for the sensor parameters, the phase of the antenna distribution can be expressed as

$$\Psi_l(\mathbf{x}, \mathbf{l}) = (\mathbf{k}_l^0)^\dagger \Delta \mathbf{l} + \Delta \mathbf{x}^\dagger \Lambda_l \Delta \mathbf{l} \quad (2.35)$$

where

$$\mathbf{k}_l^0 = \nabla_l \Psi_l |_{\bar{\mathbf{x}}, \bar{\mathbf{l}}}, \quad (2.36)$$

$$\Lambda_l = \nabla_x (\nabla_l \Psi_l)^\dagger |_{\bar{\mathbf{x}}, \bar{\mathbf{l}}}, \quad (2.37)$$

$$\nabla_l = \left[\frac{\partial}{\partial l_x} \quad \frac{\partial}{\partial l_y} \quad \frac{\partial}{\partial l_z} \right]^\dagger, \quad (2.38)$$

$$\Delta \mathbf{l} = \mathbf{l} - \bar{\mathbf{l}}, \quad (2.39)$$

$\bar{\mathbf{l}}$ is the point on the antenna structure around which the first expansion is performed, and Λ_l is termed the *antenna transformation matrix*. Using (2.35), the transmit gain is

$$G(\mathbf{x}) = \int_{S_A} |w_l(\mathbf{l})| \exp \left\{ -j \left[\Psi_a(\mathbf{l}) + (\mathbf{k}_l^0)^\dagger \Delta \mathbf{l} + \Delta \mathbf{x}^\dagger \Lambda_l \Delta \mathbf{l} \right] \right\} d\mathbf{l}. \quad (2.40)$$

Last, the transmitting antenna can be focused on the mean scatterer location, $\bar{\mathbf{x}}$, by forcing a phase taper of $\Psi_a(\mathbf{l}) = -(\mathbf{k}_l^0)^\dagger \Delta \mathbf{l}$. Then, the transmit gain is

$$G(\mathbf{x}) = \int_{S_A} |w_l(\mathbf{l})| \exp(-j\Delta \mathbf{x}^\dagger \Lambda_l \Delta \mathbf{l}) d\mathbf{l}. \quad (2.41)$$

2.2.2. The Transformation Matrices and their Physical Interpretation

The matrix, Λ_s , is called the sensor transformation matrix because of the way its function can be interpreted. To better make this interpretation, Λ_s is written in terms of its singular value decomposition (SVD). The derivative with respect to the target's z location is ignored here because all scatterers are located at the same altitude; therefore, Λ_s is a two by five matrix. When expanded in terms of its SVD, Λ_s becomes

$$\mathbf{U}\mathbf{S}\mathbf{V}^\dagger = \Lambda_s \quad (2.42)$$

where

$$\mathbf{U} = [\mathbf{u}_1 \quad \mathbf{u}_1], \quad (2.43)$$

$$\mathbf{S} = \begin{bmatrix} \sigma_1 & 0 & 0 & 0 & 0 \\ 0 & \sigma_2 & 0 & 0 & 0 \end{bmatrix}, \quad (2.44)$$

and

$$\mathbf{V} = [\mathbf{v}_1 \quad \mathbf{v}_2 \quad \mathbf{v}_3 \quad \mathbf{v}_4 \quad \mathbf{v}_5]. \quad (2.45)$$

Then, with the SVD explicitly expanded, the received phase is

$$\begin{aligned}\Psi(\mathbf{x}, \mathbf{s}) &\approx (\mathbf{k}_s^0)^\dagger \Delta \mathbf{s} + \sigma_1 (\Delta \mathbf{x}^\dagger \mathbf{u}_1) (\mathbf{v}_1^\dagger \Delta \mathbf{s}) + \sigma_2 (\Delta \mathbf{x}^\dagger \mathbf{u}_2) (\mathbf{v}_2^\dagger \Delta \mathbf{s}) \\ &= (\mathbf{k}_s^0)^\dagger \Delta \mathbf{s} + (\tilde{\mathbf{k}}_s)^\dagger \tilde{\mathbf{r}}\end{aligned}\tag{2.46}$$

where $\tilde{\mathbf{k}}_s = [\sigma_1 \Delta \mathbf{x}^\dagger \mathbf{u}_1 \quad \sigma_2 \Delta \mathbf{x}^\dagger \mathbf{u}_2]^\dagger$ and $\tilde{\mathbf{r}} = [\mathbf{v}_1^\dagger \Delta \mathbf{s} \quad \mathbf{v}_2^\dagger \Delta \mathbf{s}]^\dagger$. The sensor transformation matrix can be interpreted using (2.46). The basis vectors for the rows of Λ_s , \mathbf{v}_1 and \mathbf{v}_2 , project the five sensor parameters into two independent dimensions of a 2D synthetic aperture. Based on their interpretation using eigenanalysis, the two dimensions are termed eigensensors. The first eigensensor is obtained through the inner product of the sensor parameters with \mathbf{v}_1 , and the second eigensensor is obtained through the inner product of the sensor parameters with \mathbf{v}_2 . Consequently, the coordinates of an element in the synthetic array are given by $\mathbf{v}_1^\dagger \Delta \mathbf{s}$ and $\mathbf{v}_2^\dagger \Delta \mathbf{s}$. When these two dimensions are used as the measurement dimensions, they are the only two dimensions needed because they are the only two dimensions with nonzero singular values. The other three dimensions: \mathbf{v}_3 , \mathbf{v}_4 , and \mathbf{v}_5 , are associated with zero singular values, and a change in the position of a scatterer produces no change in the measurements obtained in these three dimensions. Therefore, only the first two dimensions provide information about stationary scatterers, and these two dimensions preserve all the information collected by the sensor in space, time, and frequency.

In addition to preserving all the sensor information in just two dimensions, the two eigensensors are also independent. Therefore, the two eigensensors provide information about two orthogonal spatial frequencies where these spatial frequencies are obtained through inner products with the basis vectors for the columns of Λ_s , \mathbf{u}_1 and \mathbf{u}_2 . The values of the new spatial frequencies are given by $\sigma_1 \Delta \mathbf{x}^\dagger \mathbf{u}_1$ and $\sigma_2 \Delta \mathbf{x}^\dagger \mathbf{u}_2$. One reason that sidelooking scenarios are so widely used is that they are easy scenarios to understand. Along-track location of the scatterer is completely

measured by time and the along-track position of the receiver. Cross-track location is completely measured by frequency and cross-track receiver position. The eigensensors provide the opportunity to characterize radar behavior with this same simplicity. The sensor dimension \mathbf{v}_1 provides information about scatterer position along dimension \mathbf{u}_1 , and the sensor dimension \mathbf{v}_2 provides information about scatterer position along dimension \mathbf{u}_2 .

The antenna transformation matrix, Λ_l , performs for an individual antenna a function similar to what Λ_s performs for the sensor samples. Λ_l takes any point on the structure of an antenna and projects that point onto a plane. The plane is perpendicular to boresight of the antenna as defined by $\bar{\mathbf{x}}$. Therefore, the antenna transformation matrix, Λ_l , takes a 3D antenna and transforms it into a 2D antenna that has an equivalent illumination pattern for the region surrounding the point $\bar{\mathbf{x}}$. In addition, Λ_l projects the original spatial frequencies corresponding to the x and y positions of the scatterer to new spatial frequencies that are measured by the primary axes of the equivalent 2D antenna.

2.2.3. Sidelooking Case

It is interesting to look at the sidelooking case because radar behavior for this scenario is known, and (2.30) can be checked to see if it predicts the proper eigensensors. In terms of derivatives, Λ_s is given by

$$\Lambda_s = \left[\begin{array}{ccccc} \frac{\partial^2}{\partial x \partial r_x} & \frac{\partial^2}{\partial x \partial r_y} & \frac{\partial^2}{\partial x \partial r_z} & \frac{\partial^2}{\partial x \partial \omega} & \frac{\partial^2}{\partial x \partial t} \\ \frac{\partial^2}{\partial y \partial r_x} & \frac{\partial^2}{\partial y \partial r_y} & \frac{\partial^2}{\partial y \partial r_z} & \frac{\partial^2}{\partial y \partial \omega} & \frac{\partial^2}{\partial y \partial t} \end{array} \right] \Psi \Bigg|_{\bar{\mathbf{x}}, \bar{\mathbf{s}}} . \quad (2.47)$$

Substituting (2.20) and (2.1)-(2.3) into (2.47) and evaluating at $\bar{\mathbf{s}}$ and $\bar{\mathbf{x}}$, the sensor transformation matrix becomes

$$\Lambda_s = \frac{\omega_0}{c} \begin{bmatrix} \frac{-(h^2 + y_0^2)}{R_0^3} & \frac{x_0 y_0}{R_0^3} & \frac{-x_0 h}{R_0^3} & \frac{2x_0}{\omega_0 R_0} & \frac{-2v(h^2 + y_0^2)}{R_0^3} \\ \frac{x_0 y_0}{R_0^3} & \frac{-(h^2 + x_0^2)}{R_0^3} & \frac{-y_0 h}{R_0^3} & \frac{2y_0}{\omega_0 R_0} & \frac{2vx_0 y_0}{R_0^3} \end{bmatrix} \quad (2.48)$$

where $R_0 = \sqrt{h^2 + x_0^2 + y_0^2}$. For sidelooking, $x_0 = 0$, and Λ_s becomes

$$\Lambda_s = \begin{bmatrix} \frac{-1}{R_0} & 0 & 0 & 0 & \frac{-2v}{R_0} \\ 0 & \frac{-h^2}{R_0^3} & \frac{-y_0 h}{R_0^3} & \frac{2y_0}{\omega_0 R_0} & 0 \end{bmatrix}. \quad (2.49)$$

Substituting (2.49) back into (2.30), the phase response is

$$\Psi(\mathbf{x}, \mathbf{s}) = (\mathbf{k}_s^0)^\dagger \Delta \mathbf{s} + \frac{\omega_0}{c} (y - y_0) \left[\frac{2y_0}{\omega_0 R_0} (\omega - \omega_0) - \frac{h^2}{R_0^3} r_y - \frac{y_0 h}{R_0^3} r_z \right] - \frac{\omega_0}{c} \frac{x}{R_0} (r_x + 2vt), \quad (2.50)$$

and (2.21) becomes

$$\begin{aligned} h_{rf}(\mathbf{x}, \mathbf{s}) &= \frac{1}{R(\mathbf{x})^2} \exp \left[-j (\mathbf{k}_s^0)^\dagger \Delta \mathbf{s} \right] \\ &\times \exp \left\{ -j \frac{\omega_0}{c} (y - y_0) \left[\frac{2y_0}{\omega_0 R_0} (\omega - \omega_0) - \frac{h^2}{R_0^3} r_y - \frac{y_0 h}{R_0^3} r_z \right] \right\} \\ &\times \exp \left[j \frac{\omega_0}{c} \frac{x}{R_0} (r_x + 2vt) \right] \\ &\times \int_{S_A} |w_l(\mathbf{l})| \exp(-j \Delta \mathbf{x}^\dagger \Lambda_l \Delta \mathbf{l}) d\mathbf{l} \end{aligned} \quad (2.51)$$

Looking at (2.51), it is seen that the along- and cross-track components contribute independent components. In (2.51), the first line contains the attenuation due to

spreading and the measured phase due to the mean target frequencies. The second line describes the cross-track component, the third line describes the along-track component, and the fourth line accounts for the radiation pattern of the transmitter.

Isolating the along-track component and noting that $\frac{\omega_0 x}{cR_0} \approx k_x$, the along-track component is

$$\rho_x = \exp\left[k_x(r_x + 2vt)\right]. \quad (2.52)$$

Equation (2.52) is the along-track response commonly used in SAR and MTI analysis. It shows that the along-track spatial frequency component is sampled by a synthetic array with elements located at $(r_x + 2vt)$. Consequently, (2.30) is awarded a degree of confidence based on proper prediction of the along-track component for the sidelooking case. Simulations presented later in this chapter demonstrate that (2.30) is valid for a wide span of scenarios.

2.3. The System Ambiguity Function

Now that (2.30) provides a method for projecting all sensor parameters into a synthetic 2D aperture, it is straightforward to evaluate the radar system's resolution and ambiguity function. The ambiguity function is the pattern resulting from the synthetic aperture. Resolution corresponds to the width of the pattern's mainlobe. Likewise, sidelobes in the ambiguity function correspond to sidelobes in the synthetic aperture pattern, and ambiguities correspond to grating lobes in the synthetic aperture pattern. In the first part of this section, the equation of the elliptically shaped contour that defines resolution is derived. Then, simulations are presented that compare the size of the mainlobe of numerically generated ambiguity functions to theoretical predictions. The comparison is performed for several constellation sizes and look geometries.

2.3.1. Resolution

There are many ways to define resolution. For example, in [40-41] resolution was defined statistically based on the Cramer-Rao lower bound (CRLB) for the given data. However, when dealing with radar, resolution is traditionally considered in terms of the correlation between two targets, which implies use of the matched filter. Since the CRLB for the given data is independent of the estimation procedure used [42], the approaches in [40-41] cannot be used to derive the resolution due specifically to the output of the matched filter. Instead, an expression must be obtained for the output of the matched filter, and then a resolution criterion based on that output must be defined.

The output of a matched filter in the presence of two targets is derived. The filter is matched to the first signal, which has $\mathbf{x} = \bar{\mathbf{x}}$, and, therefore, $\Delta\mathbf{x} = \mathbf{0}$. If the transmit aperture is assumed to have a 3D Gaussian amplitude taper with width described by the matrix \mathbf{J}_t , then

$$|w_l(\mathbf{l})| = \frac{1}{(2\pi)^{3/2} \sqrt{|\mathbf{J}_l|}} \exp\left[-\frac{1}{2} \Delta \mathbf{l}^\dagger \mathbf{J}_l^{-1} \Delta \mathbf{l}\right], \quad (2.53)$$

and the signal received from the first scatterer is

$$\begin{aligned} y_1 &= \frac{\gamma_0(\bar{\mathbf{x}}) \Delta A}{R(\bar{\mathbf{x}})^2} \exp\left[-j(\mathbf{k}_s^0)^\dagger \Delta \mathbf{s}\right] \int_{-\infty}^{\infty} \frac{1}{(2\pi)^{3/2} \sqrt{|\mathbf{J}_l|}} \exp\left[-\Delta \mathbf{l}^\dagger \mathbf{J}_l^{-1} \Delta \mathbf{l}\right] d\mathbf{l} \\ &= \frac{\gamma_0(\bar{\mathbf{x}}) \Delta A}{R(\bar{\mathbf{x}})^2} \exp\left[-j(\mathbf{k}_s^0)^\dagger \Delta \mathbf{s}\right] \end{aligned} \quad (2.54)$$

The second signal is given by

$$\begin{aligned} y_2 &= \frac{\gamma_0(\mathbf{x}_2) \Delta A}{R(\mathbf{x}_2)^2} \exp\left[-j(\mathbf{k}_s^0)^\dagger \Delta \mathbf{s}\right] \exp\left[-j \Delta \mathbf{x}^\dagger \Lambda_s \Delta \mathbf{s}\right] \\ &\quad \times \int_{s_A} \frac{1}{(2\pi)^{3/2} \sqrt{|\mathbf{J}_l|}} \exp\left[-\frac{1}{2} \Delta \mathbf{l}^\dagger \mathbf{J}_l^{-1} \Delta \mathbf{l}\right] \exp\left[-j \Delta \mathbf{x}^\dagger \Lambda_l \Delta \mathbf{l}\right] d\mathbf{l} \end{aligned} \quad (2.55)$$

By completing the square in the integrand of (2.55), y_2 is

$$y_2 = \frac{\gamma_0(\mathbf{x}_2) \Delta A}{R(\mathbf{x}_2)^2} \exp\left[-j(\mathbf{k}_s^0)^\dagger \Delta \mathbf{s}\right] \exp\left[-j \Delta \mathbf{x}^\dagger \Lambda_s \Delta \mathbf{s}\right] \exp\left[-\frac{1}{2} \Delta \mathbf{x}^\dagger \Lambda_l \mathbf{J}_l \Lambda_l^\dagger \Delta \mathbf{x}\right]. \quad (2.56)$$

The filter that is matched to y_1 is given by

$$h_c = \frac{R(\bar{\mathbf{x}})^2}{\Delta A} w_s(\Delta \mathbf{s}) \exp\left[j(\mathbf{k}_s^0)^\dagger \Delta \mathbf{s}\right] \quad (2.57)$$

where $w_s(\Delta \mathbf{s})$ is the window function, or amplitude taper, placed on the sensor data.

The window function describes weighting such as windowing of pulses or an amplitude taper on the physical array. Let the sensor parameters have a 5D, jointly Gaussian taper with width described by the matrix \mathbf{J}_s , then

$$w_s(\Delta \mathbf{s}) = \frac{1}{(2\pi)^{5/2} \sqrt{|\mathbf{J}_s|}} \exp\left(-\frac{1}{2} \Delta \mathbf{s}^\dagger \mathbf{J}_s^{-1} \Delta \mathbf{s}\right). \quad (2.58)$$

The output, a , of the correlation filter is

$$\begin{aligned} a &= \int_S (y_1 + y_2 + n_i) h_c d\Delta \mathbf{s} \\ &= \gamma_0(\bar{\mathbf{x}}) \int_S w_s(\Delta \mathbf{s}) d\Delta \mathbf{s} \\ &\quad + \gamma_0(\mathbf{x}_2) \frac{R(\bar{\mathbf{x}})^2}{R(\mathbf{x}_2)^2} \exp\left[-\frac{1}{2} \Delta \mathbf{x}^\dagger \Lambda_l \mathbf{J}_l \Lambda_l^\dagger \Delta \mathbf{x}\right] \int_S w_s(\Delta \mathbf{s}) \exp[-j \Delta \mathbf{x}^\dagger \Lambda_s \Delta \mathbf{s}] d\Delta \mathbf{s} + n_o \end{aligned} \quad (2.59)$$

where n_i and n_o are additive white Gaussian noise at the input and output of the filter, respectively, and the integrations are performed over the full range of each parameter. The first term in (2.59) is simply an integration over a 5D Gaussian pdf; therefore, the integral goes to one and the first term becomes the desired output, $\gamma_0(\bar{\mathbf{x}})$. The second term, a_2 , is

$$a_2 = \gamma_0(\mathbf{x}_2) \frac{R(\bar{\mathbf{x}})^2}{R(\mathbf{x}_2)^2} \exp\left(-\frac{1}{2} \Delta \mathbf{x}^\dagger \Lambda_l \mathbf{J}_l \Lambda_l^\dagger \Delta \mathbf{x}\right) \exp\left(-\frac{1}{2} \Delta \mathbf{x}^\dagger \Lambda_s \mathbf{J}_s \Lambda_s^\dagger \Delta \mathbf{x}\right). \quad (2.60)$$

Noting that

$$\frac{R(\bar{\mathbf{x}})^2}{R(\mathbf{x}_2)^2} \approx 1 \quad (2.61)$$

when the points $\bar{\mathbf{x}}$ and \mathbf{x}_2 differ only by a resolution cell, the output of the correlation filter is now

$$a = \gamma_0(\bar{\mathbf{x}}) + \gamma_0(\mathbf{x}_2) \exp\left[-\frac{1}{2} \Delta \mathbf{x}^\dagger \left(\Lambda_s \mathbf{J}_s \Lambda_s^\dagger + \Lambda_l \mathbf{J}_l \Lambda_l^\dagger\right) \Delta \mathbf{x}\right] + n_o. \quad (2.62)$$

Equation (2.62) shows that the output of the correlation filter depends on three terms. The first term is the desired output: the reflectance due to the pixel of interest. The second term is the amount of the second target that leaks into the filter output. It depends on the exponential term, which is the correlation between the two targets. The last term is the noise that passes through the matched filter. The last two terms, therefore, are error terms.

2.3.1.1. Constant Correlation Approach to Resolution

Two approaches will be taken to deriving resolution. The first uses the traditional approach that says two targets are resolved when their correlation drops to a specified level. The required level of correlation is chosen by an individual and may depend on the current situation and application. If the first target is held at its position while the second target's position is allowed to vary, then using (2.62), the equation for the contour of constant correlation is

$$\kappa_c = \exp \left[-\frac{1}{2} \Delta \mathbf{x}^\dagger \left(\Lambda_s \mathbf{J}_s \Lambda_s^\dagger + \Lambda_l \mathbf{J}_l \Lambda_l^\dagger \right) \Delta \mathbf{x} \right] \quad (2.63)$$

where κ_c is a constant that depends on how much correlation is chosen as the definition of resolution. Taking the natural logarithm and multiplying out the constant,

$$-2 \ln \kappa_c = \Delta \mathbf{x}^\dagger \left(\Lambda_s \mathbf{J}_s \Lambda_s^\dagger + \Lambda_l \mathbf{J}_l \Lambda_l^\dagger \right) \Delta \mathbf{x} \quad (2.64)$$

which is the equation for an ellipse. The size of the ellipse depends on the value of κ_c and on the term in parentheses. The term in parentheses is a matrix, and the size of the ellipse actually depends inversely on that matrix's determinant. For example,

if the sensor's measurement parameters are widened, the variances in \mathbf{J}_s become larger, and the determinant of the matrix in parentheses increases. Hence, the resolution ellipse decreases in size, which is expected since the sensor's size was increased. Typically, the effect of an individual antenna's illumination pattern on resolution can be ignored and (2.64) is approximated as

$$-2 \ln \kappa_c = \Delta \mathbf{x}^\dagger \Lambda_s \mathbf{J}_s \Lambda_s^\dagger \Delta \mathbf{x}. \quad (2.65)$$

2.3.1.2. Constant Error Bound Approach to Resolution

The second approach to resolution is based on the Cramer-Rao lower bound (CRLB) for the estimation error of a scatterer in the presence of noise and a second scatterer. This allows a radar engineer to define resolution based on error rather than on correlation coefficient. To calculate the CRLB for the estimation error, the output of the matched filter in (2.62) is rewritten as

$$\xi = \gamma_0(\bar{\mathbf{x}}) + \kappa_c \gamma_0(\mathbf{x}_2) + n_o \quad (2.66)$$

where κ_c was described by (2.63). The reflectivities are assumed to be independent and complex Gaussian with $E\left[|\gamma_0(\bar{\mathbf{x}})|^2\right] = E\left[|\gamma_0(\mathbf{x}_2)|^2\right] = \sigma_\gamma^2$ and $E\left[\gamma_0(\bar{\mathbf{x}})\gamma_0(\mathbf{x}_2)^*\right] = 0$. Likewise the output noise is white and complex Gaussian with $E\left[|n_o|^2\right] = \sigma_n^2$. Therefore, the pdf of the matched-filter output is

$$p_\xi\left[\xi/(\gamma_0(\bar{\mathbf{x}}), \gamma_0(\mathbf{x}_2))\right] = \frac{1}{2\pi\sigma_n^2} \exp\left\{-\frac{1}{2\sigma_n^2}\left[\xi - \bar{\xi}(\gamma_0(\bar{\mathbf{x}}), \gamma_0(\mathbf{x}_2))\right]^2\right\}. \quad (2.67)$$

The CRLB is found by taking the inverse of the Fisher information matrix [43]. The Fisher information matrix has two components: one due to the data and the other due to a priori knowledge. Each component is found by taking the second-order derivatives of the logarithm of the appropriate pdfs with respect to all parameters being estimated. For example, the data in this case are the output of the matched filter. The Fisher information matrix for this component is found by

$$\mathbf{J}_\xi = \begin{bmatrix} \frac{\partial^2}{\partial^2 \gamma_0(\bar{\mathbf{x}})} & \frac{\partial^2}{\partial \gamma_0(\bar{\mathbf{x}}) \partial \gamma_0(\mathbf{x}_2)} \\ \frac{\partial^2}{\partial \gamma_0(\mathbf{x}_2) \partial \gamma_0(\bar{\mathbf{x}})} & \frac{\partial^2}{\partial^2 \gamma_0(\mathbf{x}_2)} \end{bmatrix} \ln p_\xi. \quad (2.68)$$

Then, using (2.67), the Fisher information matrix for the filter output is

$$\mathbf{J}_\xi = \begin{bmatrix} \frac{1}{\sigma_n^2} & \frac{\kappa_c}{\sigma_n^2} \\ \frac{\kappa_c}{\sigma_n^2} & \frac{\kappa_c^2}{\sigma_n^2} \end{bmatrix}. \quad (2.69)$$

Likewise, the Fisher information matrix from the a priori information about the targets is

$$\mathbf{J}_\gamma = \frac{1}{\sigma_\gamma^2} \mathbf{I} \quad (2.70)$$

where \mathbf{I} is the identity matrix. The total information matrix is the sum of (2.69) and (2.70),

$$\mathbf{J} = \mathbf{J}_\xi + \mathbf{J}_\gamma = \begin{bmatrix} \frac{1}{\sigma_n^2} + \frac{1}{\sigma_\gamma^2} & \frac{\kappa_c}{\sigma_n^2} \\ \frac{\kappa_c}{\sigma_n^2} & \frac{\kappa_c^2}{\sigma_n^2} + \frac{1}{\sigma_\gamma^2} \end{bmatrix}. \quad (2.71)$$

The CRLB is contained in the inverse of the information matrix,

$$\mathbf{J}^{-1} = \frac{1}{\sigma_n^2 + \sigma_\gamma^2 (1 + \kappa_c^2)} \begin{bmatrix} \sigma_\gamma^2 (\sigma_n^2 + \sigma_\gamma^2 \kappa_c^2) & -\kappa_c \sigma_\gamma^4 \\ -\kappa_c \sigma_\gamma^4 & \sigma_\gamma^2 (\sigma_n^2 + \sigma_\gamma^2) \end{bmatrix}. \quad (2.72)$$

The lower bounds for the estimation error of the two targets are given on the main diagonal of (2.72). The error bound for the first target is the upper-left component of (2.72) while the error bound for the second target is the lower-right component. Therefore, the CRLB for the error in the estimation of $\gamma_0(\bar{\mathbf{x}})$ is

$$\kappa_e = \frac{\sigma_\gamma^2 (\sigma_n^2 + \sigma_\gamma^2 \kappa_c^2)}{\sigma_n^2 + \sigma_\gamma^2 (1 + \kappa_c^2)}. \quad (2.73)$$

For a given error variance and a specified level of acceptable estimation error, κ_c^2 is constant. Since κ_c^2 is constant, and it represents the amount of correlation between the two targets, the constant error bound approach to resolution produces an ellipse with the same orientation and ellipticity as the correlation approach. Using (2.63),

$$\kappa_c^2 = \exp \left[-\Delta \mathbf{x}^\dagger \left(\Lambda_s \mathbf{J}_s \Lambda_s^\dagger + \Lambda_l \mathbf{J}_l \Lambda_l^\dagger \right) \Delta \mathbf{x} \right], \quad (2.74)$$

or in terms of κ_e ,

$$-\ln \frac{\sigma_\gamma^2 \sigma_n^2 - \kappa_e \sigma_\gamma^2 - \kappa_e \sigma_n^2}{\sigma_\gamma^2 (\kappa_e - \sigma_\gamma^2)} = \Delta \mathbf{x}^\dagger \left(\Lambda_s \mathbf{J}_s \Lambda_s^\dagger + \Lambda_l \mathbf{J}_l \Lambda_l^\dagger \right) \Delta \mathbf{x}. \quad (2.75)$$

It is also interesting to note the behavior of (2.72) for some special cases. First, let the noise variance and the correlation coefficient go to zero. In this case,

$$\lim_{\kappa_c \rightarrow 0} \lim_{\sigma_n^2 \rightarrow 0} \mathbf{J}^{-1} = \begin{bmatrix} 0 & 0 \\ 0 & \sigma_\gamma^2 \end{bmatrix}, \quad (2.76)$$

and the bound shows that for zero noise and zero target correlation it is possible to estimate the reflectance of target one without error. The lower bound for the estimation error on target two goes to the a priori knowledge of target two. This is because the matched filter completely rejects energy due to target two; hence, the matched filter output provides no information toward estimating target two. With no signal energy at the output of the filter attributed to target two, its estimate must be based on any previous knowledge that was available.

If the noise is held at zero, but the two targets are completely correlated, $\kappa_c = 1$ and

$$\lim_{\kappa_c \rightarrow 1} \lim_{\sigma_n^2 \rightarrow 0} \mathbf{J}^{-1} = \begin{bmatrix} \sigma_\gamma^2 / 2 & \sigma_\gamma^2 / 2 \\ \sigma_\gamma^2 / 2 & \sigma_\gamma^2 / 2 \end{bmatrix}. \quad (2.77)$$

It is seen in (2.77) that since the two targets are completely correlated, signal energy at the output of the matched filter is evenly divided between them, and they have equal estimation errors. The last special case is when the noise power is infinite. The inverse of the Fisher information matrix in this case is

$$\lim_{\sigma_n^2 \rightarrow \infty} \mathbf{J}^{-1} = \begin{bmatrix} \sigma_\gamma^2 & 0 \\ 0 & \sigma_\gamma^2 \end{bmatrix}. \quad (2.78)$$

Therefore, when the noise is strong enough that the measurements are useless, the estimation error is the same as the variance of the a priori knowledge. Note that (2.78) is the inverse of the a priori information matrix given in (2.70), meaning that the filter output provided no additional information.

2.3.2. Resolution Simulations

The accuracy of (2.64) in predicting resolution will now be demonstrated through simulation. In this section, the ambiguity functions for several hypothetical radar systems with sparse arrays are numerically generated using the radar data model from Section 2.1. Then, the resolution ellipse is predicted using (2.64) and compared graphically with the mainlobe of the simulated ambiguity function.

The first simulated scenario is sidelooking, and the array, although sparse, is relatively small in size. The limited size of the physical array means that pixel resolution is controlled by the length of the coherent processing interval (CPI), or integration time, and by the bandwidth of the signal. Therefore, for a sidelooking, signal-dominated case, the main axes of resolution should be aligned with the along-track and cross-track dimensions. The length of the CPI determines along-track resolution, and bandwidth determines cross-track resolution. The simulation result and the theoretically predicted resolution ellipse, which is shown as a dark ring in the 3D plot, are shown in Fig. 2-3a. As expected, the axes of both the ambiguity function and the resolution ellipse do align with the along- and cross-track dimensions. Furthermore, the size of the ellipse fits the numerically generated ambiguity function perfectly. The resolution definition was chosen to be a correlation level of $\kappa_c = 0.707$. The predicted ellipse intersects the ambiguity function at this level. The size and orientation of the ambiguity function and predicted ellipse are more easily seen in the two-dimensional plot of Fig. 2-3b. In the 2D plot, the resolution ellipse is plotted in white in order to enhance contrast.

Another detail to notice in Figs. 2-3a and 2-3b is that there are no sidelobes in the ambiguity function. This is a result of the Gaussian tapers assumed for the dominant sensor parameters: time and frequency. The frequency spectrum is Gaussian, and since a cross-track cut of the ambiguity function is proportional to the Fourier transform of the frequency spectrum, the cross-track cut is also Gaussian. The same goes for the along-track cut of the ambiguity function and the window placed on the




Figure 2-3a. The 3D ambiguity function and predicted ellipse for a sidelooking scenario with resolution controlled by time and bandwidth.




Figure 2-3b. The 2D ambiguity function and predicted ellipse for a sidelooking scenario with resolution controlled by time and bandwidth.

received pulses. The next several simulations were all performed with Gaussian frequency spectra and time windows, but uniform array weighting.

The next scenario is still sidelooking, but the size of the constellation, or array, is increased significantly. As the size of the array gets larger, the size of its mainlobe decreases. Therefore, one would expect that eventually the mainlobe of the array would become smaller than a range-Doppler resolution cell. When this occurs, overall resolution improves, and the resolution cell decreases in size. In addition, the axes of the resolution cell can rotate, depending on the structure of the array. There is no reason to expect that an array, which has a shape and structure defined by orbital physics, would produce a radiation pattern with its main axes aligned with the along- and cross-track dimensions. Figures 2-4a and 2-4b show the rotation and reduced size of the resulting resolution cell. In Fig. 2-4a, the predicted ellipse intersects the ambiguity function at the specified correlation level. In Fig. 2-4b, it is obvious that the mainlobe of the ambiguity function has rotated and decreased in size.

In Figs. 2-4a and 2-4b, sidelobes have appeared even though the time and frequency tapers are still Gaussian. The reason for this is that the finer details of the ambiguity function are now controlled by the array because of its small beamwidth. The array is uniformly weighted and sparsely populated; therefore, there are sidelobes in the array pattern. Since the array pattern is dominant and it has sidelobes, the ambiguity function has sidelobes as well.

The third scenario is, once again, for a smaller constellation such that time and bandwidth control resolution, but now the geometry is forward-looking. The size of the resolution cell, as seen in Figs. 2-5a and 2-5b, is large again because the array beamwidth does not contribute to resolution. The resolution cell has rotated because of the forward-looking geometry, and no sidelobes are present because the time and frequency samples control the ambiguity function, and they are Gaussian weighted.




Figure 2-4a. The 3D ambiguity function and predicted ellipse for a sidelooking scenario with resolution controlled by the array pattern.




Figure 2-4b. The 2D ambiguity function and predicted ellipse for a sidelooking scenario with resolution controlled by the array pattern.




Figure 2-5a. The 3D ambiguity function and predicted ellipse for a forward-looking scenario with resolution controlled by time and bandwidth.




Figure 2-5b. The 2D ambiguity function and predicted ellipse for a forward-looking scenario with resolution controlled by time and bandwidth.

To round out the set, the fourth scenario, shown in Figs. 2-6a and 2-6b, is for a very large array such that beamwidth is a factor in resolution, and the scenario geometry is forward-looking. Sidelobes are present due to the uniformly weighted, sparsely sampled spatial information. The predicted resolution ellipse again intersects the ambiguity function at the correct level.

Equation (2.64) is also useful for non-Gaussian tapers and window functions. This has already been demonstrated somewhat by the fact that the array was uniformly weighted in the previous examples. The key to accurate prediction of resolution for any weight function is proper evaluation of the covariance matrix, \mathbf{J}_s , which can be obtained through the sample variance. In the simulations above, the entries corresponding to the array were obtained through the sample variance. For example, the (1,2) entry of \mathbf{J}_s corresponding to the covariance of the r_x and r_y components of the array, is given by

$$J_s^{1,2} = \frac{1}{N} \sum_{n=0}^{N-1} (r_x^n - \bar{r}_x)(r_y^n - \bar{r}_y). \quad (2.79)$$

Other components of \mathbf{J}_s can be evaluated in a similar manner.

The results of a final simulation for demonstrating the accuracy of (2.64) are shown in Figs. 2-7a and 2-7b. The geometry is forward-looking with a moderately sized physical array that has some effect on resolution. In this simulation all sensor parameters had uniform tapers. Figure 2-7 demonstrates that even though (2.64) was derived using Gaussian assumptions, it still has utility for other amplitude tapers. The important component in determining resolution is the variance of the taper, which can be obtained, even for sparsely sampled sensor parameters, through the sample variance.




Figure 2-6a. The 3D ambiguity function and predicted ellipse for a forward-looking scenario with resolution controlled by the array pattern.




Figure 2-6b. The 2D ambiguity function and predicted ellipse for a forward-looking scenario with resolution controlled by the array pattern.




Figure 2-7a. The 3D ambiguity function and predicted ellipse for a forward-looking scenario with resolution affected by time, bandwidth, and the physical array. Each of the sensor parameters has a uniform taper.




Figure 2-7b. The 2D ambiguity function and predicted ellipse for a forward-looking scenario with resolution affected by time, bandwidth, and the physical array. Each of the sensor parameters has a uniform taper.

2.3.3. Ambiguities

In addition to resolution, a fundamental performance characteristic of a sparse radar array is the sidelobe structure exhibited in the sensor ambiguity function. At a minimum, no grating lobes, which result from perfect measurement correlation between dissimilar targets, should occur within the sensor ambiguity function. In addition, the sidelobes of the ambiguity function should ideally be small. Again considering the traditional sidelooking, single-aperture SAR, significant sidelobes within the sensor ambiguity function can be directly determined from knowledge of the transmit pulse repetition frequency (PRF). For a sparse radar array, however, the concept of the synthetic aperture must be implemented to determine sensor ambiguity. The complexity of this sensor: five sensor measurement parameters, forward- and backward-looking geometries, and the large but sparse spatial array, make a more direct analysis problematic.

For the analysis in this section, we assume the radar transmits a coherent pulse train at a constant PRF. As a result, the sensor weight function for the time parameter, $w_t(\Delta t)$, is sampled periodically $1/\text{PRF}$ across the sensor's CPI. Similarly, the weight function for the frequency parameter, $w_f(\Delta\omega)$, is sampled periodically by $2\pi\text{PRF}$ across the sensor bandwidth. Since the receive array formed by the satellite constellation likewise represents sampling in three dimensions of space, the entire sensor weight function, $w(\Delta\mathbf{s})$, represents a five-dimensional array. This array can then be projected into a two-dimensional *synthetic array*, using the projection vectors described in Section 2.2.2.

As stated earlier, the sensor ambiguity function for a sparse radar array is determined from the two-dimensional synthetic aperture or array. As with any aperture or array, the sidelobe levels of the resulting pattern depend on the aperture function since the two are related by a Fourier transform. As a result, the 2D synthetic aperture function of a sparse radar array can provide direct insight into the behavior of this sensor. For example, the size of the aperture in each orthogonal

measurement direction determines sensor resolution. Likewise, to avoid large sidelobes, the aperture or array must be filled—that is, measurements must be made across the entire aperture extent. Additionally, grating lobes can occur if the synthetic aperture or array has periodicities.

Specifically, the concept of a co-array will be used to evaluate the sidelobe performance of a synthetic array. The co-array is the autocorrelation of the array and is a measure of the spatial *lags* sampled by the array. The far-field power pattern produced by an array is simply the Fourier Transform of its co-array. Regularly spaced arrays and co-arrays are defined as having spacing that can be laid out on an underlying, evenly spaced grid. Therefore, the distances between samples are always multiples of each other. If samples are missing such that the Nyquist criterion is not satisfied, the array is said to be sparsely populated. If no underlying grid can be found on which to locate the samples, the array is said to be randomly, or non-uniformly, spaced.

There are two cases when true grating lobes do not occur: when a regularly spaced co-array is Nyquist sampled or when the co-array is irregularly spaced. If a regular co-array has grating lobes, their separation from the mainlobe are determined by the minimum separation between co-array samples. Therefore, it is possible for a regularly spaced, sparse array to not have true grating lobes if some of its co-array samples are closely spaced. In this situation, however, there may be high sidelobes that nearly approach the level of grating lobes. For an irregular co-array, high sidelobes are possible, but the lack of a periodic structure prevents true grating lobes. Therefore, a sparse array radar system based on the satellite constellation concept will not have any true grating lobes, but high sidelobes could prove difficult to handle. For randomly placed arrays such as we have assumed for microsat constellations, the number of apertures in the array strongly controls sidelobe levels. A more detailed discussion of co-arrays is available in [42].




Figure 2-8. The synthetic array for a dual-aperture SAR system with PRF matched to array element spacing.

2.3.4. Synthetic Array and Synthetic Co-array Simulations

In order to demonstrate how the synthetic array and synthetic co-array can aid in ambiguity analysis of a radar system, some simulated results are presented. First, a simple dual-aperture system is modeled. The areas of the individual antennas are much smaller than required by the minimum antenna area constraint, and the PRF is low; therefore, there are Doppler ambiguities illuminated but no range ambiguities. The elements of the synthetic array are due to sampling in time, frequency, and the two spatial locations. For a sidelooking scenario, the synthetic array dimensions are the same as for traditional SAR. The synthetic array is shown in Fig. 2-8 and is seen to have a regular structure. The corresponding synthetic co-array is also regular and is shown in Fig. 2-9. The regular sampling of the synthetic co-array causes Doppler ambiguities in the ambiguity function as seen in Fig. 2-10. All range ambiguities are outside the illuminated area. The minimum distance between along-track samples in the synthetic co-array determines the locations of the Doppler ambiguities. An along-track cut provided in Fig. 2-11 shows the Doppler ambiguities more clearly.




Figure 2-9. The synthetic co-array for a dual-aperture SAR system with PRF matched to array element spacing.




Figure 2-10. The ambiguity function calculated from the synthetic array of Figure 2-8.




Figure 2-11. Doppler cut of the ambiguity function in Figure 2-10.

In the prior example, the spacing between the two elements was half of the distance traveled in a single PRI. Hence, the synthetic array and synthetic co-array were both regular on the same rectangular grid. The distance between elements in the synthetic array was the same as the distance between elements in the synthetic co-array; therefore, either could be used to assess the location of ambiguities. In addition, the synthetic array was fully filled so the sidelobes between ambiguities followed the usual form, which is the Fourier transform of the uniform taper used.

In other situations, the co-array reveals more subtle information about the array it represents. Consider the previous example with a single modification: let the PRF be reduced by a factor of $\frac{4}{5}$. The corresponding synthetic array is shown in Fig. 2-12. In the cross-track dimension, the spacing is still regular and more closely spaced, so range ambiguities will certainly not appear in the ambiguity function. In the along-track dimension, however, the synthetic samples are not regularly spaced and the




Figure 2-12. The synthetic array for a dual-aperture SAR system with PRF unmatched to array element spacing.

overall density of samples has decreased. On first intuition, one may think that the Doppler ambiguities would become more closely spaced, but because of the slightly irregular spacing, the closest true Doppler ambiguity is actually twice as far away. In the synthetic co-array for this situation, shown in Fig. 2-13, it is seen that some of the co-array samples in the along-track dimension are twice as close as the minimum co-array spacing of Fig. 2-9. As a result, the closest true Doppler ambiguity in the ambiguity function of Fig. 2-14 is twice as far away as in Fig. 2-10. A Doppler cut of the ambiguity function in Fig. 2-14 is shown in Fig. 2-15. It can be seen in Fig. 2-15 that some of the peaks are high sidelobes rather than ambiguities. The along-track dimension of the synthetic co-array in Fig. 2-13 is actually based on a grid twice as small as in Fig. 2-9, resulting in more separation between Doppler ambiguities. However, some of the samples on that smaller grid are obviously missing, and very high sidelobes appear in the ambiguity function.




Figure 2-13. The synthetic co-array for a dual-aperture SAR system with PRF unmatched to array element spacing.




Figure 2-14. The ambiguity function calculated from the synthetic array of Figure 2-12.




Figure 2-15. Doppler cut of the ambiguity function in Figure 2-14.

Depending on the situation, the sidelobes in Fig. 2-14 may be just as damaging as the true ambiguities of Fig. 2-10. This is important because the constellation concept discussed in this paper is similar to the scenario just presented. There will likely be gaps in the synthetic co-array meaning high sidelobes are inevitable. Fortunately, the synthetic co-array spacing will be even less regular than in the previous example, due to the random satellite placement rather than adjusting the PRF by a rational number. There should be fewer extreme sidelobes in the ambiguity function for the satellite constellation, with the average sidelobe level decreasing as the number of apertures in the constellation increases. Nevertheless, the potential for high sidelobes exists and must be considered by any suitable SAR or MTI algorithm.

Now that I have introduced the co-array concept, I demonstrate that the synthetic array does accurately predict the sensor's ambiguity function. I generated ambiguity functions for a forward-looking scenario with both a numeric sparse-array radar

simulator and the synthetic array. The ambiguity function obtained from the simulator is shown in Fig. 2-16a, and the ambiguity function obtained from the synthetic array is shown in Fig. 2-16b. There are some differences in the ambiguity functions since the approximations made in deriving the synthetic array do not apply to the numerical simulation. Specifically, we see that the differences are more apparent at the edges of the ambiguity function. This is because these regions are further from the center of the illuminated area where the second Taylor expansion was performed. However, although there are some differences, the ambiguity functions in Fig. 2-16a and 2-16b are generally in good agreement.

Figures 2-17 and 2-18 demonstrate application of the synthetic co-array. In Fig. 2-17a, the 2D synthetic array for a system with a five-element satellite constellation is shown. In this example, the size of the constellation is relatively small. Hence, many of the synthetic array samples overlap, and bandwidth and CPI rather than the size of the constellation primarily determine the overall size of the synthetic array. By taking the autocorrelation of the synthetic array, we generate the synthetic co-array shown in Fig. 2-17b. Because of the random nature of the spatial sampling, there is no underlying sample grid for either the synthetic array or co-array. Therefore, the corresponding ambiguity function shown in Fig. 2-17c has no true grating lobes.

The random-array ambiguity function, however, does show effects of both the periodic sampling of the radar signal and the sparse, non-periodic sampling of the satellite constellation. Since there is periodic sampling in both time and frequency, there are hints of range-Doppler ambiguities in the resulting ambiguity function. The potential range-Doppler ambiguities are multiplied by sidelobes of the radar constellation's radiation pattern in order to arrive at the sidelobes seen in Fig. 2-17c. Since the array pattern is random in nature, the expected sidelobe levels depend on the sample density in the synthetic array and co-array [42]. This density, in turn, depends on the system PRF, the overall size of the constellation, and the number of satellites in the constellation.




Figure 2-16a. Sparse-array ambiguity function generated numerically using a sparse-array simulator.




Figure 2-16b. Sparse-array ambiguity function generated using the synthetic array projection.




Figure 2-17a. The 2D synthetic array for a small microsat constellation.




Figure 2-17b. The 2D synthetic co-array for a small microsat constellation.




Figure 2-17c. The ambiguity function of the sensor shown in Figs. 2-17a and 2-17b.

A final experiment is presented in Fig. 2-18 in order to demonstrate the effect of sample density on the system ambiguity function. The constellation used to generate the results shown in Fig. 2-17 was relatively small. As a result, many elements of the synthetic array overlapped, the size of the synthetic array was primarily determined by the transmit signal, the density of samples in the synthetic co-array was relatively high, and the sidelobes immediately surrounding the ambiguity function's mainlobe were small. In Fig. 2-18, however, the satellite constellation is much larger. As a result, the synthetic array shown in Fig. 2-18a has wide regions without samples. The overall size of the synthetic array has increased because of the huge size of the constellation, but the sample density seen in the synthetic co-array of Fig. 2-18b is much less than the density seen in Fig. 2-17b. Hence, although resolution has improved through the wider extent of the synthetic array, the reduced sampling density results in high sidelobes immediately surrounding the ambiguity function's mainlobe. The ambiguity function for this case is shown in Fig. 2-18c.




Figure 2-18a. The 2D synthetic array for a large, very sparse microsat constellation.




Figure 2-18b. The 2D synthetic co-array for a large, very sparse microsat constellation.




Figure 2-18c. The ambiguity function of the sensor shown in Figs. 2-18a and 2-18b.

The compromise between resolution and sample density is a primary issue for processing of sparse radar arrays. Of course, we desire a high sample density over a wide extent, resulting in good resolution and low sidelobe levels, but this may not be achievable. We can, however, use the simulations in Figs. 2-17 and 2-18 to make some judgments about when sparse arrays can be effectively processed for SAR and MTI. The situation in Fig. 2-18 is one where resolution has improved through the large size of the constellation. In essence the number of resolution cells in the SAR map has increased, but the number of measurements for estimating the scattering from those cells has remained constant. Therefore, we cannot expect to be able to perform SAR processing effectively. If, however, we design a system such that bandwidth and CPI are the dominant factors in determining resolution, we ensure a sufficient sampling density such that quality SAR processing is achievable. In other words, the synthetic co-array analysis shows that since the time and frequency measurement dimensions are more sufficiently sampled than the spatial measurement

dimensions, they should be the dimensions that largely determine resolution. For extremely wide satellite constellations, this may require very long coherent processing intervals and wide signal bandwidths.

2.4. Summary

In this chapter, I have developed methods for analyzing multiple aperture radar systems. I began with a representation of the physics of the radar process. Then, the radar data were derived by sampling that response. The expression for the radar data, however, does not necessarily lend itself to straightforward interpretation of the radar's behavior. Therefore, in order to analyze and interpret radar performance better, a two-dimensional synthetic aperture was derived using first-order Taylor expansions of the phase of the received signal. The synthetic aperture takes measurements equivalent to the radar's measurements in time, frequency, and spatial location, but the synthetic aperture information is collected in only two dimensions. Therefore, the synthetic aperture is a useful concept for characterizing important system attributes such as resolution or the presence of ambiguities.

The importance of characterizing a radar's resolution and ambiguities will become more apparent in the succeeding chapters. In the next chapter, SAR processing of microsat constellations is considered. The importance of resolution for a SAR system is obvious, but analysis of the ambiguity function is important as well. The algebraic radar model derived in the first part of this chapter will also be put to use since it describes how to simulate radar data and can be used to derive processing filters.

Resolution and the system ambiguity function are also important for MTI. The resolution of the system affects the radar system's ability to separate and detect slowly moving targets from stationary ground clutter. In addition, the synthetic aperture is crucial to the evaluation of clutter rank, which is an important parameter in MTI performance.

3. SAR

3.1. Introduction to Multiple Aperture SAR

It is desirable for SAR to have high resolution and a large spotlight area or wide swathwidth. Instantaneously illuminating a large area is advantageous for applications that require immediate access to the information that SAR provides. For dynamic remote sensing applications, such as soil moisture measurement, that require wide coverage and monitoring over time, increased illumination area translates to fewer orbits necessary for imaging a wide area. Therefore, increased illumination area means decreased revisit time, and, for these applications, short revisit times are crucial to maintaining valid and current information. For military applications, accurate and up-to-date situational awareness is critical. Tactical information provided by SAR can be useless if it is only just a few hours old. Consequently, wide-area, high-resolution SAR can be crucial for both military and civilian applications.

Unfortunately, a commonly known design requirement that limits SAR illumination area is the minimum SAR antenna area constraint [11-17]. The requirement arises because the area that a SAR illuminates on the ground must be restricted so that the radar does not receive ambiguous returns in range or Doppler. In

spotlight mode, therefore, there is a maximum area that can be illuminated during a dwell. In stripmap mode, the width of the SAR swath is restricted. Since illumination area and aperture size are inversely related, the maximum allowable illumination area corresponds to a minimum allowable aperture size. The minimum SAR antenna area is also proportional to range, making the constraint especially restrictive for spaceborne systems.

There have been many proposed solutions for increasing illumination area; however, most of them do so at the cost of degraded azimuth resolution. ScanSAR has been a popular method [11, 16-17, 19-22] that makes this compromise. In ScanSAR, the synthetic aperture size is limited, thereby reducing azimuth resolution but also allowing time for the beam to be scanned in elevation for multiple range swaths. A technique proposed in [15] changes the elevation dimension of the antenna in order to optimize the illumination area for different grazing angles, but the fundamental area limitation is still present. Another proposed method [44] uses multiple along-track beam positions and a larger along-track aperture. The method provides better signal-to-noise ratio (SNR) but does not improve revisit time for large areas. Last, several papers in the literature divide illumination time of a single receiver into smaller segments called subapertures, but this is done to null large sidelobe targets, improve processing efficiency, or reduce speckle [17, 21, 45-49].

Multiple coherent receive apertures can be used to improve illumination coverage while maintaining azimuth resolution [13, 22, 30, 50-51]. In this case, the illumination area is determined by the size of the individual apertures, and the spatial information provided by those apertures is used to resolve any range-Doppler ambiguities that are illuminated. The sum of the aperture areas must still satisfy the minimum area constraint, but illumination area can be improved over the single-receiver case by a factor equal to the number of receivers.

Intuitively, the use of a multiple-aperture array works as follows. Each aperture is smaller than required by the minimum SAR antenna area constraint. Depending on the shape of the antenna and the PRF of the transmit signal, range ambiguities,

Doppler ambiguities, or both are illuminated by the system. After SAR processing, the size of a pixel and its ambiguities are limited to the range-Doppler resolution of the system. Since the desired pixel and its ambiguities are limited to finite, fixed locations, the array pattern can be used to distinguish the desired pixel from its ambiguities. A phase taper is placed on the array such that the mainlobe of the array pattern falls on the desired pixel. If there are enough apertures in the array, nulls can be placed on all range-Doppler ambiguities. Therefore, the array can be used to pass the desired pixel while rejecting any energy due to its range-Doppler ambiguities. The process can be repeated for all range bins, Doppler bins, and ambiguous regions to complete the final SAR image.

Some important issues arise when considering the intuitive explanation of the preceding paragraph. The first issue concerns the number of apertures necessary for synthesis of the mainlobe and nulls in the array pattern. The number of constraints on the array pattern is the sum of the mainlobe constraint and the null constraints. Consequently, the minimum number of receive apertures in the array is $N_{amb} + 1$, where N_{amb} is the number of illuminated range-Doppler ambiguities. However, while this minimum number of apertures is enough to ensure that the array passes energy from the desired pixel while nulling range-Doppler ambiguities, constraining the array pattern so tightly may produce poor results. In order to form the required number of nulls, the mainlobe of the array pattern may become distorted. The mainlobe may broaden, distort, and shift away from the desired pixel, resulting in SNR loss. Also, the peak of the mainlobe could fall on range-Doppler sidelobes. Although the ambiguities will have been rejected, decreased SNR and increased sidelobe leakage may result in a poor SAR image.

Another issue that goes along with the number of apertures required to synthesize an array pattern is the array's quiescent, or beamforming, pattern. If the beamforming pattern already has nulls near the range-Doppler ambiguities, it will be relatively easy to synthesize an array pattern that shifts these nulls slightly. However, if a major sidelobe of the beamforming pattern falls on a range-Doppler ambiguity, it will be

more difficult to generate a null on that ambiguity because the null goes against the natural beamforming pattern.

The quiescent pattern and the number of apertures in the array combine to determine how much the pattern mainlobe will distort when nulls are formed. Since the beamforming pattern is determined by placement of the apertures, and since aperture placement is determined by the physics governing the satellite orbits, controlling the original beamforming pattern is not likely an option for improving pattern synthesis. The working assumption is that the beamforming pattern will be a random mess of sidelobes, some of which will inevitably fall on range-Doppler ambiguities. Therefore, the remaining option for improving pattern synthesis is to add more apertures to the satellite constellation. It will be shown in this chapter that adding extra apertures is an effective and important method of improving performance of a SAR satellite cluster.

The last issue concerns the overall size of the satellite constellation. If the physical array formed by the constellation is large compared to the size of the synthetic array formed by the time and frequency samples, problems will result. First, a large physical array can improve resolution. While this may seem like an advantage, in this case it is actually a detriment. The basis for improving swathwidth through the use of multiple apertures is that the apertures must add spatial samples without improving resolution. Instead, the spatial samples must be used to resolve range-Doppler ambiguities. If the constellation is large enough to improve resolution through the beamwidth of its array, the number of pixels within the illuminated area increases. If the number of illuminated pixels exceeds the number of measurements obtained, ambiguities are unavoidable.

The size of the satellite constellation compared to the size of the time/frequency synthetic aperture is also important because of the width of the nulls placed by the satellite array. In order to work, multiple aperture SAR must use the array pattern to reject range-Doppler ambiguities. If the physical array is too large, its nulls will be

too small to null the entire pixel of a range-Doppler ambiguity, and only a small part of the ambiguity's energy will be rejected.

3.2. Linear SAR Processing

In the beginning of Chapter 2, the expression for the sampled radar data was approximated with a summation and written in matrix-vector notation,

$$\mathbf{d} = \mathbf{P}\boldsymbol{\gamma} + \mathbf{n}. \quad (3.1)$$

The SAR problem is one of estimating the values in the reflectance vector, $\boldsymbol{\gamma}$, using noisy, complex data samples, \mathbf{d} , and knowledge of the radar parameters, geometry, and propagation physics, \mathbf{P} . As seen in (3.1), the radar process is approximated as a linear summation of vectors, with each vector being weighted by the appropriate element of the reflectance vector. Since the radar process is linear, the estimator used to estimate $\boldsymbol{\gamma}$ should also be linear. A weight vector, or filter, which can be applied as a linear process, will be determined for each resolution cell. When the inner products between the received measurements and each of the weight vectors are taken, the estimated reflectance vector, $\hat{\boldsymbol{\gamma}}$, is

$$\hat{\boldsymbol{\gamma}} = \mathbf{W}\mathbf{d} \quad (3.2)$$

where

$$\hat{\boldsymbol{\gamma}} = [\hat{\gamma}_1 \quad \hat{\gamma}_2 \quad \cdots \quad \hat{\gamma}_C]^\dagger, \quad (3.3)$$

$$\mathbf{W} = [\mathbf{w}_1 \quad \mathbf{w}_2 \quad \cdots \quad \mathbf{w}_C]^\text{H}, \quad (3.4)$$

\mathbf{w}_i is the weight vector for the i^{th} pixel, and $(\cdot)^\text{H}$ denotes the complex conjugate transpose operation. It is also important to note that this is the same type of processing traditionally done in SAR where \mathbf{w}_i is typically the matched filter, which is the complex conjugate of the i^{th} pixel's response vector.

Although some algorithms presented in this paper will require more computation to find each \mathbf{w}_i , the total size of \mathbf{W} remains constant; consequently, the data-

dependent process of calculating inner products is equivalent for all linear estimators. This is important because the numerically intensive algorithms will not need to be computed in real time on the spaceborne platform. The coefficients of the weight vectors can be calculated offline based on the projected orbit of the constellation. Then, the coefficients can either be transmitted to the satellites for onboard processing or the data can be transmitted from the satellite to Earth for ground-based processing. Even the more complex algorithms can be implemented through onboard processing because the actual filtering process is computationally equivalent regardless of the estimation algorithm used to compute the filter coefficients.

3.3. Single-Aperture SAR

3.3.1. Time-Bandwidth Limitation on Illumination Area

Ambiguities occur when the number of illuminated resolution cells exceeds the number of independent measurements collected. In spotlight mode, SAR focuses on a particular area for some time, T . During that time, independent, complex samples can be collected at a maximum rate determined by the signal bandwidth, B . Therefore, the maximum number of independent, complex samples that can be collected is equal to BT , also known as the time-bandwidth product. The problem is that bandwidth and observation time are fixed by the resolution requirements of the system. The range resolution requirement determines bandwidth, and the azimuth resolution requirement determines observation time. Since it is only possible to unambiguously image as many targets as there are independent samples, and since the number of independent samples is related through bandwidth and time to particular range and azimuth resolutions, the maximum image area is fixed. A simple example for a sidelooking, spotlight SAR is presented in the following. Suppose the resolution requirements imposed on the radar are δx and δR in azimuth and range, respectively. Range resolution is given by

$$\delta R = \frac{c}{2B}. \quad (3.5)$$

For a 90° sidelooking geometry where the azimuth extent is small compared to the range, the azimuth resolution can be approximated as

$$\delta x = \frac{\lambda_0 R_0}{2v} \frac{1}{T} \quad (3.6)$$

where λ_0 is the wavelength at the center operating frequency and R_0 is the average pixel range. The area per pixel is then the product of the azimuth and range resolutions,

$$\delta x \delta R = \frac{c\lambda_0 R_0}{4v} \frac{1}{BT}. \quad (3.7)$$

The maximum area is the time-bandwidth product multiplied by the area per pixel,

$$A_{max} = BT\delta x\delta R = \frac{c\lambda_0 R_0}{4v}. \quad (3.8)$$

As can be seen from the right side of (3.8), the maximum area that can be imaged is determined by range, wavelength, and platform velocity.

In the middle of (3.8), increasing the time-bandwidth product appears to be one solution to the problem. However, (3.5) and (3.6) clearly show that the resolution dimensions are inversely proportional to bandwidth and time. For example, it is possible to obtain four times as many independent samples by increasing bandwidth and time each by a factor of two, but the resolution pixels would be half the original size on each side. Although there would be four times as many samples, the area of each pixel would be four times smaller, and the total area would remain the same. Radar designers sometimes try to solve the problem by staggering the PRF or by applying other signal coding methods. However, the total energy in the ambiguity function must remain constant [52]. Therefore, these methods can rearrange ambiguity, but cannot make it disappear. It is impossible to get around the fact that only as many resolution cells can be illuminated as there are measurements available to distinguish them.

3.3.2. Single-Aperture SAR Simulations

Figure 3-1 shows the magnitude of the image used as the input to the radar simulator. It was taken from a photograph of the campus of the University of Kansas in Lawrence, Kansas. The photograph was scanned, changed to a 256-level grayscale image, and cropped to be 512 by 512 pixels. Each pixel intensity was given a random phase, and the resulting set of complex pixel values was used as the vector of



Figure 3-1. Image of the KU campus used as the input to the radar simulator.

scattering coefficients, γ , in the radar simulations. The image was chosen because it realistically represents what could be seen in a SAR scenario. The image has a wide variety of green spaces, buildings, parking lots, cars, roads, and landmarks.

The simulations presented in this paper were designed to have equal resolution in the along- and cross-track dimensions. Except for the first simulation in this section, the sizes of the individual apertures were chosen such that the 3 dB boundaries of their illumination patterns approximately coincided with the edges of the 512 by 512 image. The time-bandwidth product of the received signal was 36864; therefore, only about $\frac{1}{7}$ of the original 512 by 512 input can be unambiguously imaged by a single-aperture system.

The first simulation is for a single aperture that meets the minimum SAR antenna area constraint. The antenna's size and orientation are such that its 3 dB beamwidth approximately covers the middle ninth of the original image. There is no noise added




Figure 3-2. The image formed by a single aperture that satisfies the minimum SAR antenna area constraint.

to the data in this simulation, and the scenario is sidelooking with a boresight elevation angle of 45 degrees. For an aperture that meets the antenna area constraint, the area illuminated on the ground is smaller than what is spanned by the 512 by 512 image. The resulting matched-filtered image, shown in Fig. 3-2, is good only over a limited area. Outside the beamwidth of the aperture, the received signal energy rapidly approaches zero.

If, in order to meet swathwidth or revisit rate requirements, a wider area of the ground is illuminated, the result is similar to the image shown in Fig. 3-3. The aperture used for the simulation that produced Fig. 3-3 was smaller than the antenna used for Fig. 3-2 by three times in both the along- and cross-track dimensions. Therefore, there are approximately three ambiguous range swaths and three ambiguous Doppler swaths illuminated in the simulation. With ambiguities present, and no spatial data available to resolve them, the entire SAR image is of poor quality.




Figure 3-3. The image formed by a single aperture that is much smaller than required by the minimum SAR antenna area constraint.

The results shown in Figs. 3-2 and 3-3 demonstrate the challenge presented by the minimum SAR antenna area constraint. For any given range and Doppler resolutions, the integration time and signal bandwidth are set. These, in turn, define the maximum number of independent measurements that can be collected, and, therefore, the maximum number of resolution pixels that can be illuminated. The illuminated area in Fig. 3-2 may not be enough to satisfy revisit rate or real-time data requirements, but any attempt to illuminate a wider area with a single aperture ends in disaster as seen in Fig. 3-3.

3.4. Multiple-Aperture SAR

A necessary requirement for increasing SAR map area is to increase the number of independent samples, or amount of information, that is collected without modifying resolution cell size in the process. It was shown that this is not possible by increasing bandwidth or integration time. However, it was shown in Chapter 2 that the additional data samples can come from using multiple apertures. If N is the number of apertures with coherent receivers, the number of independent data samples available to the radar system is NBT . If the physical array formed by the satellite constellation is not so large that its beamwidth affects resolution, the area per SAR pixel, repeated here for convenience, is still expressed by (3.7),

$$\delta x \delta R = \frac{c\lambda_0 R_0}{4v} \frac{1}{BT}. \quad (3.9)$$

The maximum allowable illumination area is again the product of the number of independent samples and the area per pixel,

$$A_{max} = NBT\delta x \delta R = N \frac{c\lambda_0 R_0}{4v}, \quad (3.10)$$

which is N times larger than was possible with a single receiver.

3.4.1. Correlation Processing

Assume the antenna of a single-aperture radar system satisfies the minimum SAR antenna area constraint. Portions of the synthetic array and synthetic co-array formed from time and frequency samples are shown in Fig. 3-4. Because the synthetic array is regularly spaced on a rectangular grid, the synthetic co-array is also regularly spaced. The minimum distance between co-array samples determines the location of range-Doppler ambiguities, which are rejected by the antenna pattern because it satisfies the minimum area constraint.




Figure 3-4a. The synthetic array for a single-aperture system.




Figure 3-4b. The synthetic co-array for a single-aperture system.




Figure 3-5. An aperture satisfying the minimum SAR antenna area constraint is divided into multiple elements of an array. The large aperture was used to generate Fig. 3-2 while the array was used to generate Fig. 3-7.

Next, assume that the single antenna satisfying the antenna area constraint is divided into several sections as seen in Fig. 3-5, and a coherent receiver is placed behind each section. Because each aperture is three times smaller in both dimensions than the original, an area three times wider in both the along- and cross-track dimensions will be illuminated. However, the spatial sampling provided by the array will resolve range-Doppler ambiguities.

Portions of the synthetic array and synthetic co-array for the divided aperture scenario are shown in Fig. 3-6. The portions shown in Fig. 3-6 cover the same area as the portions shown in Fig. 3-4. Therefore, it can be seen that the divided aperture has increased the sampling density by three times in each dimension of the synthetic array. The minimum distance between co-array samples has decreased by the same amount. Since this distance has decreased by three times in each dimension, ambiguities are pushed beyond the wider illumination area.




Figure 3-6a. The synthetic array for the divided aperture of Fig. 3-5.




Figure 3-6b. The synthetic co-array for the divided aperture of Fig. 3-5.




Figure 3-7. The matched-filter result when the KU campus image is applied to the divided aperture of Figure 3-5.

When a simulation using the image of the KU campus is applied to the divided aperture of Fig. 3-5, the result of matched filtering is the image shown in Fig. 3-7. The vector representation of the correlation, or matched, filter for the i^{th} resolution cell is a weighted version of its response vector,

$$\mathbf{w}_i^{corr} = \frac{\mathbf{p}_i}{|\mathbf{p}_i|^2} \quad (3.11)$$

where the superscript, *corr*, denotes that the filter is a correlation filter. If the matched filters for all targets are placed into a matrix, the matched-filter estimator, \mathbf{W}_{corr} , is given by

$$\mathbf{W}_{corr} = \mathbf{D}^{-1} \mathbf{P}^H \quad (3.12)$$

where \mathbf{P} is given in (2.8) and the diagonal matrix, \mathbf{D} , is defined as

$$\mathbf{D} = \begin{bmatrix} \boldsymbol{\rho}_1^H \boldsymbol{\rho}_1 & 0 & \cdots & 0 \\ 0 & \boldsymbol{\rho}_2^H \boldsymbol{\rho}_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \boldsymbol{\rho}_C^H \boldsymbol{\rho}_C \end{bmatrix}. \quad (3.13)$$

When the matched filter is applied to the data, the estimate of the i^{th} target is

$$\hat{\gamma}_i = (\mathbf{w}_i^{corr})^H \mathbf{d} = \gamma_i + \sum_{j \neq i} (\mathbf{w}_i^{corr})^H \boldsymbol{\rho}_j \gamma_j + (\mathbf{w}_i^{corr})^H \mathbf{n}. \quad (3.14)$$

As seen in (3.14), error in the matched filtering operation comes from two terms. The first error component is represented by the summation term. When estimating the i^{th} resolution cell, other cells are considered clutter. Hence, the error represented by the summation is error due to clutter because of correlation between the matched filter and other targets. The second error component is shown in the last term of (3.14) as the error component due to noise. The matched-filter vector, however, has the smallest magnitude of any filter vector that gives γ_i as its expected result. Therefore, the matched filter minimizes noise power at its output. While the matched filter maximizes SNR, it does not account for clutter in any manner. In cases that are clutter limited rather than noise limited, the matched filter does not provide optimal estimates. However, it is this lack of dependence on clutter that also makes the matched-filter vectors the least computationally expensive to generate.

In the generation of Fig. 3-7, no noise was added to the simulated data vector. Any error in the image must be due to the clutter term of (3.14), but the excellent result indicates that this error component was small as well. The co-array of the divided-aperture system explains why the error due to the clutter component was small; however, an alternative way to understand the small error is to look at the range-Doppler ambiguity function versus the radiation pattern of the physical array. The range-Doppler ambiguity function is shown in Fig. 3-8. From this figure it is apparent that for most pixels there are eight other pixels that are ambiguous in the




Figure 3-8. The range-Doppler ambiguity function for a single element of the divided aperture in Figure 3-5.

time and frequency domains. Figure 3-9, however, shows the array pattern when the array is focused on the center pixel. This is the spatial-domain ambiguity function for the center pixel. The eight pixels that are ambiguous with the center pixel are located in the nulls of the array pattern. The total ambiguity function for the center target is, approximately, the product of the range-Doppler and spatial ambiguity functions. This product is shown in Fig 3-10, which shows that the total ambiguity function approaches the ideal *thumbtack* shape. Energy from ambiguous targets has been eliminated rather than rearranged because of the addition of spatial information.




Figure 3-9. The array pattern for the divided aperture of Figure 3-5.




Figure 3-10. The total ambiguity function of the divided-aperture system is the product of the range-Doppler ambiguity function in Figure 3-8 and the array pattern in Figure 3-9.




Figure 3-11. The locations of the apertures in a sparsely populated, irregularly spaced array.

3.5. Sparse Arrays

The microsat concept calls for placing each receive aperture on its own, small satellite. Furthermore, the orbital dynamics of formation flying restrict the satellites from having small, regular spacing. Figure 3-11 shows the random 3D array to be used for the following SAR simulations. The aperture locations were generated using a 3D uniform probability density function. This represents control of the array structure by orbital dynamics rather than by radar design. Candidate aperture locations were generated using the uniform pdf. If the proposed location was less than a specified minimum separation, the point was thrown out and a new point was generated. If the candidate location was farther from the other satellites than required by the minimum distance, the location was retained, and the process was repeated until locations were obtained for all apertures. The mean of the aperture




Figure 3-12. The result of matched-filter processing for the sparse satellite array.

locations was then subtracted off so that the locations could be scaled to simulate differing sizes of satellite constellations. The rectangular apertures are all the same size and are oriented such that they illuminate the same area.

The result of correlation processing for the sparse satellite constellation is shown in Fig. 3-12. From this result, it is obvious that the sparse, randomly sampled array significantly impacts SAR processing. The image in Fig. 3-12 is considerably worse than the image in Fig. 3-7, and since no noise has been added to the data, the degradation must be a result of the sparse, irregular nature of the satellite constellation.

The poor result shown in Fig. 3-12 can be explained as before through the synthetic co-array or through the range-Doppler and spatial ambiguity functions. First, a part of the synthetic array and synthetic co-array for a sparse radar system are shown in Fig. 3-13. Because of the random nature of the spatial sampling, there is no




Figure 3-13a. A part of the synthetic array for the sparse array system.




Figure 3-13b. A part of the synthetic co-array for the sparse array system.

clear, underlying sample grid for either the synthetic array or co-array. Therefore, the total ambiguity function does not have true ambiguities. However, because of the random nature of the synthetic co-array, the behavior of the sidelobes in the total ambiguity function is also random. Some areas of the ambiguity function have high sidelobes and some have low sidelobes. The high sidelobe areas are the source of the degradation seen in Fig. 3-12.

Likewise, the poor result can also be explained using the range-Doppler ambiguity function and the array pattern of the satellite constellation. The range-Doppler ambiguity function is the ambiguity function for a single aperture resulting from the time and frequency data. Since the transmit signal has not changed, the range-Doppler ambiguity function is the same as shown in Fig. 3-8. The array pattern, however, has changed considerably and is shown in Fig. 3-14. The total ambiguity function for the sparse system is the product of the ambiguity functions in Figs. 3-8 and 3-14. This product is shown in Fig. 3-15. The range-Doppler ambiguities can still be seen in the total ambiguity function. This is because the ambiguities are multiplied by sidelobes in the spatial pattern. The sidelobe level that multiplies the ambiguities is a varying quantity, and, therefore, different ambiguities bleed through into the total ambiguity function by differing amounts. Figure 3-15 shows this effect as the ambiguities have varying intensities.

The advantage of the matched filter is that the ratio of desired pixel energy to white noise energy is maximized at the filter output. Therefore, the matched filter minimizes the error due to noise, and it is expected that performance should degrade gracefully as the SNR of the measurements is decreased. Figure 3-16 demonstrates this behavior. The results of matched-filter processing for various SNRs are shown. The SNR is defined as the ratio of the average signal power in a single data sample to the average power of a single noise sample. In Fig. 3-16, the SNRs are -5 dB, 5 dB, 20 dB, and ∞ dB. It can be seen that, although the image clearly worsens with decreased SNR, the degradation is relatively slow. In the low-SNR, noise-limited




Figure 3-14. The array pattern for the sparse array formed by a constellation of microsats.




Figure 3-15. The total ambiguity function for the sparse array system.




Figure 3-16. The matched-filter result for the sparse array and varying SNR.

case, the matched filter provides optimal results by maximizing output SNR. The matched filter is not optimal in the high-SNR case, however, because it does not sacrifice SNR in order to mitigate clutter. In the next section, a filter is derived which is optimal in the high-SNR case. When results from other filters are shown later in this chapter, it will be clear that the matched filter provides the most graceful degradation versus SNR.

3.5.1. Maximum Likelihood Processing

If the vector of measurements, \mathbf{d} , is defined in (2.7) and the vector of noise values, \mathbf{n} , is jointly Gaussian, complex noise with zero mean and a covariance matrix given by \mathbf{K}_n , the conditional pdf of the measurement vector is

$$p_{\mathbf{d}/\gamma} = \frac{1}{\sqrt{|2\pi\mathbf{K}_n|}} \exp\left[-\frac{1}{2}(\mathbf{d} - \mathbf{P}\gamma)^H \mathbf{K}_n^{-1} (\mathbf{d} - \mathbf{P}\gamma)\right]. \quad (3.15)$$

The maximum likelihood (ML) estimator is obtained by maximizing the natural logarithm of (3.15). This maximum is found by taking the derivative with respect to γ and setting to zero,

$$\frac{\partial \ln p_{\mathbf{d}/\gamma}}{\partial \gamma} = (\mathbf{d} - \mathbf{P}\gamma)^H \mathbf{K}_n^{-1} \mathbf{P} = \mathbf{0}. \quad (3.16)$$

Solving (3.16) yields the ML estimate of γ ,

$$\hat{\gamma}_{ml} = \left(\mathbf{P}^H \mathbf{K}_n^{-1} \mathbf{P}\right)^{-1} \mathbf{P}^H \mathbf{K}_n^{-1} \mathbf{d} \quad (3.17)$$

where $(\cdot)^{-1}$ denotes the pseudoinverse operation that is based on the SVD. The matrix that performs the ML estimation is, therefore, given by

$$\mathbf{W}_{ml} = \left(\mathbf{P}^H \mathbf{K}_n^{-1} \mathbf{P}\right)^{-1} \mathbf{P}^H \mathbf{K}_n^{-1}. \quad (3.18)$$

Next, if a white noise assumption is made such that the noise samples are independent, the noise covariance matrix is diagonal,

$$\mathbf{K}_n = \sigma_n^2 \mathbf{I} \quad (3.19)$$

and the ML estimator reduces to

$$\mathbf{W}_{ml} = \mathbf{P}^{-1}. \quad (3.20)$$

The estimate of the i^{th} pixel's reflectance due to the ML filter is

$$\hat{\gamma}_i = \left(\mathbf{w}_i^{ml} \right)^H \mathbf{d} = \gamma_i + \left(\mathbf{w}_i^{ml} \right)^H \mathbf{n}. \quad (3.21)$$

Comparing (3.21) and (3.14), it is seen that the clutter term is absent in (3.21). The last term in (3.21), however, becomes important. If \mathbf{P} is ill conditioned, the magnitude of the weight vectors that make up \mathbf{W}_{ml} can be very large and any noise present in the measurements will be significantly amplified. Thus, SNR and the condition of \mathbf{P} become critical factors in determining the quality of the SAR image obtained through the ML filter.

The condition of \mathbf{P} can be improved by adding more apertures. The size of \mathbf{P} is the number of measurement samples by the number of SAR pixels. If adding more apertures can be done without improving resolution, the condition of \mathbf{P} can be improved. This is because the number of rows in \mathbf{P} is increased without increasing the number of columns. In linear algebra terms, by making the measurement vectors longer, the dimension of the subspace that the vectors lie in becomes larger. However, since the number of pixels remains unchanged, there are the same number of vectors that exist within that subspace. With the same number of vectors in a larger subspace, there is more free space available for the pseudoinverse operation to work with. From an ambiguity function and array pattern point of view, the addition of more apertures means there are more DOFs available to the array for synthesizing a spatial pattern. Therefore, range-Doppler ambiguities can be nulled without sacrificing as much SNR.

Equation (3.20) is difficult to implement in practice. It is likely that an operational system would illuminate thousands of pixels and store thousands of data samples. The size of \mathbf{P} would be thousands of rows by thousands of columns and calculating its inverse would be impractical. Fortunately, the ML filter can be applied to a subset of the data as long as the subset is carefully chosen. The data can be

divided up between the matched and ML filters in any manner as long as the matched filtering process limits the number of non-zero pixels to less than the number of measurements available for ML filtering. For best results, the number of pixels passed by the matched filter should be as small as possible compared to the number of measurements available for ML filtering.

For example, applying the matched filter in the time and frequency domains and the ML filter in the spatial domain leads to the interpretation that the array pattern is being used to null range-Doppler ambiguities. After matched filtering, there will be a finite number of range-Doppler ambiguous pixels. The ML filter applied in the spatial domain can be interpreted as placing array pattern nulls on those ambiguities.

A new matrix of response vectors, \mathbf{P}_s , must be defined in order to apply the ML filter to the spatial data. Since the ML filter is being applied to the spatial data only, the number of rows in \mathbf{P}_s is the number of apertures, N . Also, it was previously assumed that the output of range-Doppler filtering contains energy from only the desired pixel and its ambiguities. Therefore, range-Doppler sidelobes are ignored, and the number of columns in \mathbf{P}_s only needs to be $N_{amb} + 1$. Hence, \mathbf{P}_s is a matrix of beamforming vectors. Each column of \mathbf{P}_s contains the beamforming weights for either the desired pixel or one of its ambiguities. In order to obtain an accurate estimate of the pixel reflectances, the amplitudes of the beamforming weights, a_i , for a target or ambiguity at \mathbf{x}_i is

$$a_i = \frac{\Delta A}{R(\mathbf{x}_i)^2} G(\mathbf{x}_i) |\mathbf{w}_{tf}| \quad (3.22)$$

where \mathbf{w}_{tf} is the vector of time-frequency, matched-filter coefficients. If Φ_i is the vector of beamforming phases shifts for the i^{th} pixel or ambiguity, \mathbf{P}_s is

$$\mathbf{P}_s = \left[a_1 \exp(j\Phi_1) \quad a_2 \exp(j\Phi_2) \quad \cdots \quad a_{N_{amb}+1} \exp(j\Phi_{N_{amb}+1}) \right]. \quad (3.23)$$




Figure 3-17. The ML-filter result for the sparse array and varying SNR.

Results for the spatial ML filter for four different SNRs are shown in Fig. 3-17. The SNRs are the same as for the matched filter: -5 dB, 5 dB, 20 dB, and ∞ dB. Since the matrix estimator depends only on the pseudoinverse of \mathbf{P}_s , it does not vary with SNR. The pseudoinverse operation calculates filter vectors that are orthogonal to the response vectors from other pixels; consequently, the resulting error due to correlation with other targets is zero. In high-SNR cases this behavior produces optimal results. However, some of the weight vectors that result from the

pseudoinverse are very large in magnitude in order to yield accurate estimates of the pixel reflectances in high-SNR scenarios. When the input noise increases, those large vectors amplify the noise, and the output becomes unstable. This behavior, typical of inverse-type solutions, is clearly apparent in the low-SNR results shown in Fig. 3-17.

The ML filter maximizes the signal-to-clutter ratio (SCR). In high-SNR cases where the number of measurements significantly exceeds the number of pixels, the ML filter produces an excellent result. One disadvantage of the ML filter includes unstable results, or pixel estimates that increase rapidly with input noise power, when the input SNR is low. Another disadvantage is increased computation required to calculate \mathbf{W}_{ml} when compared to the matched filter. This is because the matched filter requires a matrix transpose while the ML filter requires a pseudoinverse operation.

3.5.2. Maximum a Posteriori Processing

While the ML estimate is obtained from the a priori density, the maximum a posteriori (MAP) estimate is obtained by maximizing the conditional a posteriori density. Using the a priori density described in (3.15), Bayes' rule gives the a posteriori density,

$$p_{\gamma/\mathbf{d}} = \frac{p_{\mathbf{d}/\gamma} p_{\gamma}}{p_{\mathbf{d}}}, \quad (3.24)$$

and the MAP estimate is obtained by maximizing

$$p_{\mathbf{d}/\gamma} = \frac{p_{\gamma}}{p_{\mathbf{d}}} \frac{1}{\sqrt{|2\pi\mathbf{K}_n|}} \exp\left[-\frac{1}{2}(\mathbf{d} - \mathbf{P}\boldsymbol{\gamma})^H \mathbf{K}_n^{-1} (\mathbf{d} - \mathbf{P}\boldsymbol{\gamma})\right]. \quad (3.25)$$

For a given data vector, $p_{\mathbf{d}}$ is constant and does not affect the location of the maximum. Hence, it can be ignored. Next, assuming that the elements of $\boldsymbol{\gamma}$ are

uncorrelated, complex, zero-mean, Gaussian random variables of variance

$E[|\gamma_i|^2] = \sigma_\gamma^2$, the scattering-coefficient covariance matrix, \mathbf{K}_γ , is

$$\mathbf{K}_\gamma = E[\boldsymbol{\gamma}\boldsymbol{\gamma}^H] = \sigma_\gamma^2 \mathbf{I}, \quad (3.26)$$

and the pdf of the scattering-coefficient vector is

$$p_\gamma = \exp\left[-\frac{1}{2}\boldsymbol{\gamma}^H \mathbf{K}_\gamma^{-1} \boldsymbol{\gamma}\right]. \quad (3.27)$$

Substituting (3.27) into (3.25), and noting that the estimate that maximizes (3.25) also maximizes its logarithm, the MAP estimate is the solution to

$$\frac{\partial \ln[p_{\gamma/d} p_\gamma]}{\partial \boldsymbol{\gamma}} = \frac{\partial}{\partial \boldsymbol{\gamma}} \left[-\frac{1}{2}(\mathbf{d} - \mathbf{P}\boldsymbol{\gamma})^H \mathbf{K}_n^{-1} (\mathbf{d} - \mathbf{P}\boldsymbol{\gamma}) - \frac{1}{2}\boldsymbol{\gamma}^H \mathbf{K}_\gamma^{-1} \boldsymbol{\gamma} \right] = 0. \quad (3.28)$$

This solution is

$$\hat{\boldsymbol{\gamma}}_{map} = \left(\mathbf{P}^H \mathbf{K}_n^{-1} \mathbf{P} + \mathbf{K}_\gamma^{-1} \right)^{-1} \mathbf{P}^H \mathbf{K}_n^{-1} \mathbf{d}. \quad (3.29)$$

Equation (3.29) uses a priori information about the scattering coefficients in order to make an estimate of $\boldsymbol{\gamma}$ that may be better than the previous ML estimate. In fact, estimation theory shows that ML estimation is the same as MAP estimation in the limit as a priori information about the value being estimated goes to zero. In the above, this is the same as letting the variance of the scattering coefficients go to infinity, which leads to

$$\lim_{\sigma_\gamma^2 \rightarrow \infty} \hat{\boldsymbol{\gamma}}_{map} = \left(\mathbf{P}^H \mathbf{K}_n^{-1} \mathbf{P} \right)^{-1} \mathbf{P}^H \mathbf{K}_n^{-1} \mathbf{d} = \hat{\boldsymbol{\gamma}}_{ml}. \quad (3.30)$$

When a priori information about γ is not zero, this information and information about the noise statistics amount to an a priori estimate of SNR, which can be used to obtain a better estimate of γ . For the Gaussian assumptions that have been made, however, the MAP estimate is equivalent to the MMSE estimate. Therefore, I will demonstrate use of a priori SNR information in a modified form. This form, the MMSE estimate, is derived in the next section.

3.5.3. Minimum Mean-Squared Error Processing

The minimum mean-squared error (MMSE) method uses statistical properties of the targets and noise to calculate the filter that achieves the best compromise between SCR and SNR in terms of mean-squared estimation error. The MMSE filter, therefore, is the mathematically optimum compromise between the correlation and ML filters.

The derivation of the linear MMSE filter begins with the orthogonality principle, which states that all linear combinations of the data must be orthogonal to the estimation error. This can be expressed as [42]

$$E\left[(\hat{\gamma}-\gamma)^H \mathbf{M}\mathbf{d}\right] = 0 \quad (3.31)$$

where \mathbf{M} is used to represent all linear combinations of the data, \mathbf{d} . Substituting for the reflectance estimate, (3.31) is

$$E\left[(\mathbf{W}_{mmse}\mathbf{d}-\gamma)^H \mathbf{M}\mathbf{d}\right] = 0, \quad (3.32)$$

which can be equivalently expressed as [42]

$$E\left[(\mathbf{W}_{mmse}\mathbf{d}-\gamma)\mathbf{d}^H\right] = \mathbf{0} \quad (3.33)$$

where $\mathbf{0}$ is a matrix of zeros. Using (3.1) for the data vector and rearranging terms changes (3.33) to

$$\mathbf{E} \left[\mathbf{W}_{mmse} (\mathbf{P}\boldsymbol{\gamma} + \mathbf{n})(\mathbf{P}\boldsymbol{\gamma} + \mathbf{n})^H \right] = \mathbf{E} \left[\boldsymbol{\gamma}(\mathbf{P}\boldsymbol{\gamma} + \mathbf{n})^H \right] \quad (3.34)$$

or

$$\mathbf{W}_{mmse} \mathbf{E} \left[\mathbf{P}\boldsymbol{\gamma}\boldsymbol{\gamma}^H\mathbf{P}^H + \mathbf{P}\boldsymbol{\gamma}\mathbf{n}^H + \mathbf{n}\boldsymbol{\gamma}^H\mathbf{P}^H + \mathbf{n}\mathbf{n}^H \right] = \mathbf{E} \left[\boldsymbol{\gamma}\boldsymbol{\gamma}^H\mathbf{P}^H + \boldsymbol{\gamma}\mathbf{n}^H \right]. \quad (3.35)$$

Next, assuming that the noise samples are independent of the scattering values and that $\mathbf{E} \left[\mathbf{n}\mathbf{n}^H \right] = \mathbf{K}_n$ is the noise covariance matrix, (3.35) becomes

$$\mathbf{W}_{mmse} \left(\mathbf{P} \mathbf{E} \left[\boldsymbol{\gamma}\boldsymbol{\gamma}^H \right] \mathbf{P}^H + \mathbf{K}_n \right) = \mathbf{E} \left[\boldsymbol{\gamma}\boldsymbol{\gamma}^H \right] \mathbf{P}^H. \quad (3.36)$$

The MMSE matrix estimator is then

$$\mathbf{W}_{mmse} = \mathbf{E} \left[\boldsymbol{\gamma}\boldsymbol{\gamma}^H \right] \mathbf{P}^H \left(\mathbf{P} \mathbf{E} \left[\boldsymbol{\gamma}\boldsymbol{\gamma}^H \right] \mathbf{P}^H + \mathbf{K}_n \right)^{-1}. \quad (3.37)$$

With the further assumptions that the reflectance values from pixel to pixel are uncorrelated and that the mean scattering value is zero, $\mathbf{E} \left[\boldsymbol{\gamma}\boldsymbol{\gamma}^H \right] = \sigma_\gamma^2 \mathbf{I}$ where σ_γ^2 is the expected value of the squared reflectance magnitude for each target. Also, assuming white noise with power given by σ_n^2 , the filter is

$$\mathbf{W}_{mmse} = \sigma_\gamma^2 \mathbf{P}^H \left(\sigma_\gamma^2 \mathbf{P}\mathbf{P}^H + \sigma_n^2 \mathbf{I} \right)^{-1}. \quad (3.38)$$

Examining (3.38) provides important insight into the behavior of the MMSE filter. First, in a low- or zero-noise case, σ_n^2 will be negligible and \mathbf{W}_{mmse} becomes

$$\mathbf{W}_{mmse} \approx \mathbf{P}^{-1}, \quad (3.39)$$

which is the same as the ML filter. In the low-noise case, therefore, the MMSE filter maximizes the SCR. In the high-noise case, σ_n^2 dominates and \mathbf{W}_{mmse} becomes

$$\mathbf{W}_{mmse} \approx \frac{\sigma_\gamma^2}{\sigma_n^2} \mathbf{P}^H. \quad (3.40)$$

The filter represented by (3.40) is very interesting. It is a vector in the same direction as the matched filter. An important difference, however, is that the MMSE filter becomes inversely proportional to the noise variance. Hence, as the noise variance approaches infinity, the magnitude of the MMSE filter approaches zero. Therefore, inherent in the equation for the MMSE filter is the concession that in the presence of overwhelming noise, it is best to estimate the reflectance values not by the measurements, but by the a priori statistical properties of the targets.

It is impractical to apply the MMSE filter to an entire data set just as it was for the ML filter. The reason is the same: an extremely large matrix inverse would be required. However, the MMSE filter can also be applied to a subset of the data. A new matrix of pixel response vectors must be defined as in the ML case, but in the MMSE case, the effect of the matched filter on the noise power at the input of the MMSE filter must also be taken into account. If the vector containing the matched filter weights is \mathbf{w}_{tf} , the noise power expected at the input of the MMSE filter, $\sigma_{n,mmse}^2$, is

$$\sigma_{n,mmse}^2 = \sigma_n^2 |\mathbf{w}_{tf}|^2 \quad (3.41)$$

where σ_n^2 is the noise power at the input to the matched filter. If the matched filter is applied as before to the time and frequency data, the spatial MMSE filter is given by

$$\mathbf{W}_{mmse} = \sigma_\gamma^2 \mathbf{P}_s^H \left(\sigma_\gamma^2 \mathbf{P}_s \mathbf{P}_s^H + \sigma_n^2 |\mathbf{w}_{tf}|^2 \mathbf{I} \right)^{-1} \quad (3.42)$$

where \mathbf{P}_s was defined in (3.23).




Figure 3-18. The MMSE-filter result for the sparse array and varying SNR.

Figure 3-18 shows the same SNR cases: -5 dB, 5 dB, 20 dB, and ∞ dB, for the spatial MMSE filter as were shown before for the matched and spatial ML filters. Figure 3-19 shows results of the three filters side by side for low-, moderate- and infinite-SNR cases. From Figs. 3-18 and 3-19, it can be seen that the MMSE filter provides the best of both worlds in terms of performance. In the high-SNR situation, the MMSE filter minimizes correlation with other pixels in order to reduce the error




Figure 3-19. The matched, ML, and MMSE filters in a side-by-side comparison versus SNR.

due to clutter. In this case the MMSE filter produces the same output as the ML filter, but the matched filter is seen to be clutter limited. In the low-SNR case, the MMSE filter maximizes SNR but also scales the filter in order to rely on target statistics rather than noisy measurements. The matched filter maximizes SNR but does not reduce its magnitude regardless of the input noise power. The ML filter goes unstable for the low-SNR case. The magnitude of the ML filter is very large for some pixels because it must compensate for loss of gain due to strict enforcement of

the null constraints. In the moderate-SNR case, the MMSE filter is the optimum compromise between the matched and ML filters.

The primary advantage of the MMSE filter is that it, by definition, provides the minimum mean-squared error in all noise and clutter scenarios. The filter accounts for the statistical properties of both the pixels and noise. The cost of implementing the MMSE filter is increased computation for calculating \mathbf{W} compared to both the ML and matched filters.

3.6. Numerical Performance

Thus far, simulation results have been presented pictorially in order to demonstrate the advantages of arrays in SAR processing, but it is also informative to assess performance of the different algorithms numerically. Several error curves are presented that demonstrate the performance of the three algorithms versus different variables. The error criterion is the mean-squared error (MSE) of the pixel magnitudes normalized by the image's mean pixel magnitude

$$\text{MSE}_n = \frac{(\hat{\gamma} - \gamma)^H (\hat{\gamma} - \gamma)}{\gamma^H \gamma}. \quad (3.43)$$

First, the effect of SNR on performance is investigated. Figure 3-20 shows the normalized mean-squared error as a function of input SNR, and clearly validates the conclusions that have been stated about the performance of the three filters. Most important, the MMSE filter has the lowest error at every SNR. Also of importance is the rapid increase in error as SNR decreases for the ML filter and the flattening of the matched filter curve for high SNR. The flat curve for the matched filter as SNR increases shows that the matched filter is clutter limited; therefore, improving SNR does not improve the results. Last, the ML and MMSE curves begin to flatten at lower error and higher SNR than the matched filter. The error floor is due to range-Doppler matched filtering. The level of range-Doppler sidelobes due to matched filtering in time and frequency determines the best achievable performance when the SNR is high. If the ML and MMSE filters were applied to the entire data set in time, frequency, and space, the error would continue to decrease for increasing SNR. Application of the ML and MMSE filters in the spatial domain only, however, does yield significant improvement over the matched filter as shown by the right side of Fig. 3-20.




Figure 3-20. Correlation, ML, and MMSE filter performance versus input SNR for a 12-receiver, sparse, random array.

Another important factor that has not yet been investigated is the accuracy of information about the physical scenario. Since the physics must be known in order to calculate the response vectors, any deviation from the expected scenario introduces error in the results. Figure 3-21 shows the behavior of the estimation error versus one particular deviation: receiver positioning error. In this simulation, the filters were calculated using assumed receiver positions. When the radar measurements were simulated, however, the receiver locations were randomly deviated from the assumed locations according to a Gaussian distribution of standard deviation, σ_p , in each




Figure 3-21. Correlation, ML, and MMSE filter performance versus antenna positioning error for a sparse, random array and infinite SNR.

dimension. The standard deviation of the total positioning error for each receiver was then

$$\sigma_e = \sqrt{3}\sigma_p. \quad (3.44)$$

Figure 3-21 shows estimation error versus σ_e . The simulations were performed at infinite SNR for a different input image. The matched filter results vary little versus amount of positioning error. However, since placing nulls requires accurate phase information, the ML and MMSE results are very sensitive to positioning error.




Figure 3-22. Correlation, ML, and MMSE filter performance versus number of receive apertures for a sparse, random array and moderate SNR.

It was mentioned earlier that forming nulls in the spatial pattern requires enough degrees of freedom. More DOFs allow forming the required nulls without sacrificing gain on the target. Hence, it is expected that estimation error will decrease as the number of apertures increases. Figure 3-22 demonstrates this result. The estimation error versus the number of receivers in Fig. 3-22 only improves slightly for the matched filter because it only has one constraint: that the gain on the target be maximized. The entire range from nine to 20 receivers provides enough degrees of freedom to satisfy this one constraint. The slight improvement with increasing receivers is due to more signal energy being collected by the increasing combined aperture area. The ML and MMSE filters, however, have nine constraints (the target

and eight ambiguities). Therefore, as the number of receive apertures is increased from nine to 20, the degrees of freedom increase significantly, and the estimation error improves more rapidly than for the matched filter. This improvement is especially noticeable in the nine to 13 receiver range where the extra degrees of freedom in the array pattern increase from zero to four. Beyond 13 receivers, the improvement is mostly due to the increased signal energy collected.

3.7. Bandwidth and CPI Length versus Constellation Size

The relationship between the beamwidth of the physical array and resolution as determined by bandwidth and the length of the integration interval has been mentioned a couple of times up to this point. In Chapter 2, the size of the constellation versus bandwidth and CPI length was discussed in terms of the density of the sampled synthetic co-array. Simulations were presented that demonstrated high sidelobes in the case where the co-array density was decreased due to large satellite spacing. This case was also interpreted as a scenario where the beamwidth of the physical array is small enough to improve resolution, thereby increasing the number of resolution pixels illuminated by the radar system. The problems associated with a small physical array beamwidth were also mentioned in the introduction of this chapter, but have not been elaborated on. In this section, I provide results that further explain the problem.

The ML and MMSE SAR filters have been presented as techniques to reject range-Doppler ambiguities that pass through the first step of traditional SAR processing. After time and frequency processing, the ML and MMSE filters place array pattern nulls on range-Doppler ambiguities to cancel them out. The assumption has been that the ambiguities are completely rejected by the radiation null. However, if the width of the null is smaller than a range-Doppler resolution cell, only a fraction of ambiguity will be rejected, leaving some of the ambiguous energy to leak through to the filter output.

Figure 3-8 showed the range-Doppler ambiguity function for a single satellite. Since traditional SAR processing is applied before spatial processing, this ambiguity function does not change for different constellation sizes. If the satellite constellation is small enough that bandwidth and CPI length determine resolution, a sample array pattern for ML filtering is shown in Fig. 3-23. When this spatial pattern is multiplied by the range-Doppler ambiguity function of Fig. 3-8, the total ambiguity function for ML spatial filtering is shown in Fig. 3-24. In Fig. 3-24, it is seen that




Figure 3-23. The array pattern for ML spatial filtering.




Figure 3-24. The total ambiguity function for ML spatial filtering.




Figure 3-25. The array pattern for ML spatial filtering of a very large microsat constellation.

some of the range-Doppler sidelobes leak through sidelobes of the array pattern, but the range-Doppler ambiguities have been largely rejected.

If, however, the microsat constellation is large enough that the mainlobe of the physical array is smaller than a range-Doppler resolution cell, the ML spatial filtering pattern may look like the pattern shown in Fig. 3-25. When the array pattern in Fig. 3-25 is multiplied by the range-Doppler ambiguity function of Fig. 3-8, the resulting total ambiguity function for ML processing is shown in Fig. 3-26. In Fig. 3-26, the resolution of the desired pixel in the middle has clearly improved, but there is more residue apparent in other regions of the ambiguity function. If I show a closeup of one of the range-Doppler ambiguous areas in finer detail, the reason for the extra residue becomes obvious. A close-up of one of the range-Doppler ambiguities before




Figure 3-26. The total ambiguity function for ML spatial filtering of a very large microsat constellation.

spatial processing is shown in Fig. 3-27. In Fig. 3-28, the same region of the ambiguity function is shown after ML spatial processing. The null applied by the ML filter should significantly reject the ambiguity. We see in Fig. 3-28 that a narrow null passes through the middle of the ambiguity, but the null is not wide enough to cancel the entire ambiguity. Hence, range-Doppler ambiguous pixels will leak into the filter output causing a degraded image.




Figure 3-27. Close-up of a range-Doppler ambiguity.




Figure 3-28. Close-up of a range-Doppler ambiguity that is partially rejected by a null of the ML spatial filter.

3.8. SNR Analysis

A constellation of radar satellites must overcome two sources of loss in SNR. First, since the size of the transmitting antenna is reduced, there is a loss of antenna gain on transmit. Another way of looking at this is that the same amount of transmit power must be spread over more pixels within the increased illumination width. Second, there will be a processing loss compared to the matched filter when the ML and MMSE filters are used. This is due to the loss of processing gain that is necessary in order to reject range-Doppler ambiguities with spatial nulls. On receive, there is no effective loss of antenna gain since the total effective aperture is the sum of all the microsat antennas. In fact, since Fig. 3-22 showed that the number of receivers should surpass the number of illuminated range-Doppler ambiguities in order to provide extra DOFs to the pattern synthesis process, the effective aperture on receive will likely be larger than for traditional single-aperture SAR. In other words, the constellation concept suffers from reduced transmit gain, but for each satellite added to the constellation, the radar's effective aperture is increased and spatial processing loss is decreased.

Since SNR is an important factor in successfully employing the ML and MMSE filters presented in this chapter, this section investigates the SNR of a constellation-type radar system. First, a baseline computation was performed based on the RADARSAT satellite. Since we are interested in wider swathwidths without degrading resolution, the center beam for RADARSAT's fine resolution mode was used [53-56]. RCS values were taken from [57]. After the baseline was computed, I calculated the output SNR on a per pixel basis for a system such as RADARSAT that attempted to increase its swathwidth by decreasing the size of its antenna by three times in elevation. Of course, a single-aperture system like this would produce poor results due to illumination of range-Doppler ambiguities, but the calculation gives an indication of the received SNR for a single-aperture system that spreads its transmit

Table 3-1. RADARSAT system parameters.

f	5.3 GHz
λ	5.62 cm
Peak Power, P_T	5 kW
Average Power, \bar{P}	300 W
Antenna Along-Track Length, L_x	15 m
Antenna Cross-Track Length, L_y	1.5 m
Altitude, h	800 km
Bandwidth, B	30 MHz
FM Chirp Length, τ_p	40 μ s
PRF	1330 Hz
Looks	1
Elevation Angle, θ_{el}	37.1°
System Noise Temperature	290 K
System Noise Figure	2.8 dB
Boltzmann's Constant, k_b	1.38×10^{-23} J/K

power over a wider illumination area. Next, I computed the total output SNR including MMSE processing loss as satellites were added in a sparse configuration and compared the output SNR to the RADARSAT output. Finally, I performed the same computations for a constellation-type system with antennas that are reduced in size in both cross-track and along-track.

Table 3-1 shows the parameters obtained [53-56] for the center beam of RADARSAT's fine-resolution mode, except that the system noise temperature and noise figure were obtained from [57]. Using equations in [57], RADARSAT's

altitude and elevation angle lead to a grazing angle of $\theta_{gr} = 47.2^\circ$ and a slant range of $R_s = 1.043e6$ m. The range resolution on the ground is given by

$$\delta y = \frac{c}{2B \cos \theta_{gr}} = 7.3 \text{ m}, \quad (3.45)$$

and the along-track resolution for SAR is half the length of the physical aperture,

$$\delta x = \frac{L_x}{2} = 7.5 \text{ m}. \quad (3.46)$$

The along-track beamwidth is

$$\theta_h = \frac{\lambda}{L_x} = 0.00375 \text{ rad}, \quad (3.47)$$

resulting in an along-track illumination width of

$$W_h = \theta_h R_s = 3910 \text{ m}. \quad (3.48)$$

For a satellite traveling at a ground speed of $v = 7800 \text{ m/s}$, the illumination width means that a pixel stays within the radar beam for an integration time of

$$T = \frac{W_h}{v} = 0.5 \text{ s}, \quad (3.49)$$

which for the stated PRF, the number of integrated pulses is

$$M = T \times \text{PRF} = 665 \text{ Pulses}. \quad (3.50)$$

The gain of the RADARSAT antenna is

$$G = \frac{4\pi A}{\lambda^2} = 90160 = 49.55 \text{ dB}. \quad (3.51)$$

The final parameter needed for computing the RADARSAT baseline is the RCS per unit area, which is obtained from [57] by interpolating between the given grazing angles of 40° and 50°. This gives an RCS per unit area of $\sigma_0 = 0.063 = -12$ dB .

The expected SNR for a single resolution cell after pulse compression but before Doppler processing is

$$SNR_{in} = \frac{P_T \tau_p G^2 \sigma_0 \delta x \delta y \lambda^2}{(4\pi)^3 R_s^4 k_b T_{sys} F_{sys}} = 1 = 0 \text{ dB} . \quad (3.52)$$

Last, the SNR after Doppler processing without windowing is

$$SNR_o = M \times SNR_{in} = 665 = 28.2 \text{ dB} . \quad (3.53)$$

The output SNR in (3.53) is the expected value of the power at the output of SAR processing due to a single pixel of a given resolution related to the expected output white noise power. Therefore, the value does not include leakage of other pixels into the output through sidelobes of the range-Doppler matched filter. It should also be noted that the 28.2 dB mark is not a magic number, but merely the output SNR for one mode of RADARSAT, which will be used for comparison throughout the rest of this section.

Next, assume that a system has the same baseline parameters as the previous RADARSAT mode, except that the antennas are one-third the size of RADARSAT's antennas in elevation. The only parameter in (3.52) that changes is the antenna gain. In this case, for a 15m by 0.5m antenna, the gain is $G = 30053 = 44.78$ dB . Therefore, the SNR before Doppler processing would be -9.54 dB and the SNR after Doppler processing would be

$$SNR_o = M \times SNR_{in} = 18.66 \text{ dB} . \quad (3.54)$$




Figure 3-29. SNR gain and processing loss in dB for varying number of receivers in elevation.

The output SNR in (3.54) is approximately 9.5 dB below the original RADARSAT output SNR due to the loss of antenna gain. In addition, range-Doppler ambiguities will be illuminated, which will ruin the SAR processing. If additional apertures are added in elevation, the range-Doppler ambiguities can be rejected using the array pattern, and increased SNR can be obtained through the added effective aperture. A minimum of three apertures is needed for resolving range-Doppler ambiguities. Figure 3-29 shows the SNR gained through added apertures and the SNR loss due to ambiguity nulling as the number of apertures is increased from three to eight. The SNR loss was computed numerically by averaging the loss due to the MMSE filter over a wide range of pixels in cross-track. It is seen in Fig. 3-29 that SNR improves with added apertures through both increased area and reduced processing loss. For example, for five satellites in elevation, there is a 7 dB improvement due to added effective aperture and a 2 dB loss due to ambiguity

nulling. The result is an improvement of 5 dB compared to the single aperture case. Therefore, the output SNR per pixel is

$$\begin{aligned} SNR_{mmse} &= SNR_o + 10 \log_{10} N_r - Loss_{dB} \\ &= 18.66 + 7 - 2 = 23.66 \text{ dB} \end{aligned} \quad (3.55)$$

which compares favorably with the RADARSAT output SNR.

Last, I present calculations for a full microsat scenario with increased illumination in both along- and cross-track. Assume, again, the same baseline parameters except for antennas that are $\frac{1}{3}$ smaller in both along- and cross-track. With smaller antennas in along-track, the system could coherently integrate over an interval three times longer than before, resulting in an along-track resolution that is improved by the same factor of three. However, since we are discussing output SNR on a per pixel basis, and the output SNR certainly depends on the pixel's RCS, I will keep the coherent integration the same as before. This means that along-track resolution stays the same, but the extra illumination time can be used to create three looks that are incoherently averaged.

The gain for the antennas that now have $\frac{1}{9}$ the area of the original RADARSAT antenna is $G = 10018 = 40$ dB. Keeping all other parameters the same, the input SNR is

$$SNR_{in} = -19.1 \text{ dB}. \quad (3.56)$$

Integrating over the same number of pulses, the SNR after Doppler processing is

$$SNR_o = M \times SNR_{in} = 9.1 \text{ dB}. \quad (3.57)$$

Again, the total output SNR includes the effect of increased effective aperture and processing loss due to use of the MMSE or ML filter rather than the matched filter. This loss for the MMSE filter was estimated numerically and is shown in Fig. 3-30 along with the gain from added apertures. The total output SNR includes an




Figure 3-30. SNR gain and processing loss in dB for varying number of satellites in a microsat constellation.

additional improvement equal to the square root of the number of looks. Using the aperture gain and processing loss from Fig. 3-30 for 16 apertures, a sample output SNR for the MMSE filter is

$$\begin{aligned}
 SNR_{mmse} &= SNR_o + 10 \log_{10} N_r - Loss_{dB} + 10 \log_{10} \sqrt{N_{looks}} \\
 &= 9.1 + 12 - 5 + 2.4 = 18.5 \text{ dB}
 \end{aligned}
 \tag{3.58}$$

Although the result in (3.58) is several dB less than the RADARSAT result of (3.53), the output SNR for this sample microsat scenario should be sufficient for a good SAR image. Also, it must be kept in mind that the sample scenario takes an immediate hit of 9.55 dB due to a loss of transmit antenna gain. It should be expected that if the same amount of transmit power is spread out over nine times as many pixels, that the output SNR per pixel should decrease. That is not to say that the 18.5

dB mark in (3.58) is insufficient or that it is in any way a final number. Certainly, output SNR could be improved by decreasing satellite altitude, adding more satellites, or transmitting more power. It has even been suggested that each satellite in the constellation transmit its own unique signal that could be received and identified by other satellites. For example, each satellite could transmit at a slightly different frequency band. Then, the signals due to each transmitter could be integrated for further SNR improvement. If the 18.2 dB mark is sufficient, the currently assumed transmit power could be spread over several transmitters, resulting in a lower peak and average power requirement for each satellite. In the end, the SNR analysis shows that coherent processing of the signals from multiple apertures more than makes up for the reduced gain on receive, and the flexibility of a satellite constellation provides many options for dealing with the reduced gain on transmit.

3.9. Summary

I have demonstrated the utility of multiple receive apertures for SAR processing. The spotlight area, or swathwidth, of a SAR system is fundamentally limited by the amount of information that can be collected at a given resolution, and multiple receive apertures mitigate this situation by adding independent angle-of-arrival information. A sparsely populated array with randomly placed elements was presented, and the motivation for such an array in space was discussed. I also presented additional algorithms for processing the angle-of-arrival information and applied them to multiple receive aperture simulations. I showed results produced by the three different algorithms for varying signal-to-noise ratios and investigated numerical performance versus the factors of SNR, receive-aperture positioning accuracy, and number of receive elements.

The results presented in this chapter demonstrate both the needs for and methods of applying multiple receive apertures to obtain wide-area SAR images. Future spaceborne systems of the microsat concept are being studied, and this chapter demonstrates how the measurements from such systems can be processed effectively. Several factors affect this processing. One of the most important factors is the number of receivers. In order for a system such as the one proposed to work, there must be enough receive apertures to null all range-Doppler ambiguities as well as keep the array's mainlobe on the target. A sufficient number of receivers ensures that the ML and MMSE algorithms will not produce unstable results in the presence of noise. This sufficient number was determined to be the minimum number of receivers needed to satisfy the minimum SAR antenna area constraint with their combined aperture, plus at least an additional 50 percent. The SNR received by a sparsely populated, spaceborne array also affects the processing that should be used. In situations where the received SNR is high enough, the improved performance of the ML and MMSE filters justify their added computational expense.

Another important factor is the ability to position the receive apertures accurately. As with any pattern synthesis problem where nulls are desired, the relative phase shifts between elements are the dominant components in the algorithm. Ability to estimate these phase shifts accurately is crucial. Although the results presented in Figure 3-21 are for only one particular microsat constellation, the improvement shown by the ML and MMSE solutions for absolute positioning errors of one-tenth of a wavelength or less is a good rule of thumb. Again, the benefit of having extra receive apertures becomes apparent, as increasing the number of apertures will ease positioning accuracy requirements for a given error level.

The final factor presented in this chapter was the size of the satellite cluster relative to the 2D synthetic aperture formed by the time-frequency data. In order for the proposed technique to work, the size of the full 2D synthetic aperture must be predominantly controlled by bandwidth and the length of the coherent integration interval. This ensures that nulls formed by the physical array will be wide enough to completely reject range-Doppler ambiguities. Examples of systems that both did and did not meet this final criterion were discussed in detail and demonstrated graphically.

The MMSE solution is the most robust solution. It maximizes signal-to-clutter ratio in clutter-limited cases and signal-to-noise ratio in noise-limited cases. Its computational expense is significantly more than for the matched filter, but only slightly more than for the ML filter. Furthermore, the added computation is not data dependent, and the MMSE filter can be calculated before data are collected. In moderate- to high-SNR cases with low positioning error, the results produced by the MMSE filter justify its added computational burden.

4. MOVING TARGET INDICATION

4.1. Introduction to MTI, DPCA, and STAP

The detection of moving targets using radar has evolved significantly over the years. In this section, I will describe the basic steps taken toward implementation of MTI on airborne platforms. Once I describe the filtering concepts and performance metrics most commonly used today to perform and analyze MTI, I will be able to discuss the unique challenges that apply to MTI applied via a spaceborne platform. I will also be able to assess performance of a sparse, spaceborne system using standard performance metrics.

There is a predictable relationship between clutter's angle and Doppler frequency when observed by a moving radar platform. This relationship makes it possible to distinguish between moving and stationary targets when space and time filtering are jointly applied, because moving targets have a different angle-Doppler relationship than ground clutter. For a spaceborne platform, however, the component of a moving target's speed in the direction of the radar will be small. Therefore, high Doppler and angle resolutions are required in order to separate targets from ground clutter. While high Doppler resolution can be obtained through long integration intervals, the angular resolution must be obtained through a large aperture or array. The difficulties

of large apertures have already been discussed in Chapter 1 and Chapter 3; moreover, the high angular resolution required for spaceborne-based MTI will likely require an even larger aperture than is required by the minimum SAR antenna area constraint. Clusters of microsats are attractive in this regard because of the wide angular coverage and high angular resolution they can provide without the difficulties associated with large satellites.

In order for space-time filtering to detect moving targets, however, the number of measurement degrees of freedom used up by clutter, also known as clutter rank, must be less than the number of measurements obtained. Also, in order to effectively apply DPCA or STAP to a satellite cluster and in order to design effective MTI radars, the clutter rank needs to be estimated. Unfortunately, current methods of estimating clutter rank are not well suited to the satellite cluster concept. They apply only to sidelooking, linear-array scenarios with wide illumination widths and Nyquist spatial sampling. They also require strict relationships between the spatial element spacing, the speed of the radar platform, and the transmit signal's PRF.

In this chapter, after introducing DPCA and STAP in more detail, I present a method of estimating the clutter rank observed by sparse satellite clusters. I show that the current rule for predicting clutter rank, Brennan's rule, is based on the time-bandwidth product. Based on this result, 2D clutter rank, in general, depends on a product called the *space-bandwidth* product, and I show that this product can be obtained through the 2D synthetic aperture from Chapter 2. The resulting method for estimating clutter rank is more general than Brennan's rule and can be applied to the design and processing of the MTI modes of satellite clusters.

Next, I use the 2D clutter rank estimation procedure to derive and analyze DPCA performance for satellite clusters. I show that a clutter subspace can be defined. The dimension of the clutter subspace is the clutter rank. Then, the measurements obtained by a satellite cluster can be stripped of clutter energy by projecting the measurements orthogonal to the defined clutter subspace. I show performance of this method versus system parameters such as the size of the satellite constellation, the

number of satellites in the constellation, and error in the knowledge of the satellite positions.

Next, I present some wide-area MTI simulations. I show that it is possible, using clutter rank prediction and projecting orthogonal to the clutter subspace, to detect moving targets using a sparse satellite cluster, even when range-Doppler ambiguities are illuminated. There is, however, an interesting problem that arises in terms of estimating the locations of the detections. The reason why this problem occurs and a method for mitigating the problem are also presented. Finally, the results from this chapter are summarized, and their implications on system designed are discussed.

4.1.1. Space-Time Filtering and DPCA

When MTI is performed with a stationary radar, the Doppler frequency of ground clutter is centered at zero with a small spectral width caused by things such as wind-induced intrinsic clutter motion (ICM). A high-pass filter that rejects the low-frequency clutter can detect moving targets, or those targets with a nonzero velocity relative to the radar. When MTI is performed by a moving radar platform, however, the Doppler spectrum of reflections from the Earth spreads significantly. The Doppler frequency of a clutter return observed by a moving radar can be very positive for clutter directly in front of the radar, or it can be very negative for clutter directly behind the radar. In addition, a low PRF can cause clutter to fill the entire unambiguous Doppler spectrum, making it impossible to perform MTI through Doppler filtering alone.

Fortunately, there is a predictable relationship between clutter angle and Doppler frequency, f_D , given by

$$f_D = \frac{2v}{\lambda} \sin \alpha \quad (4.1)$$

where α is the angle between the clutter and the nearest point perpendicular to the radar's flight direction. It is this relationship between angle and Doppler that is

exploited by SAR processing to obtain azimuth resolution. For MTI, however, the ground reflections are considered clutter since they can obscure reflections due to moving targets; consequently, the angle-Doppler relationship is exploited for rejection of ground reflections.

Consider the sidelooking case for a single range bin. For relatively small beamwidths, the small-angle identity can be used to approximate clutter's angle-Doppler relationship as

$$f_D \approx \frac{2v}{\lambda} \alpha. \quad (4.2)$$

A simulated azimuth-Doppler clutter spectrum for this scenario is shown in Fig. 4-1. In the spectrum, a structure that represents ground clutter can be seen following a linear relationship between azimuth angle and Doppler. If the position and velocity of a moving target are such that its angle and Doppler are located in the spectrum at the arrow, the necessity of space-time filtering becomes apparent. There is clutter at the same Doppler frequency as the moving target and at the same angle as the moving target. However, there is not clutter at the same *combination* of Doppler and angle as the moving target. If spatial filtering is employed, the moving target competes with clutter at the same angle. If Doppler filtering is employed, the moving target competes with clutter at the same Doppler. Jointly filtering in angle and Doppler, however, provides a method for rejecting clutter without rejecting the desired moving target.

The first approaches to space-time filtering were developed intuitively. Stationary radars detected moving targets using high-pass filters implemented as single- or multi-pulse cancellers. The idea was that stationary interference produced identical responses in successive received pulses. Therefore, clutter could be rejected by subtracting the measurements obtained from successive pulses, and any energy that remained after subtraction should be due to either noise or moving targets.




Figure 4-1. Angle-Doppler clutter spectrum and moving target for a sidelooking scenario.

The displaced-phase-center-antenna (DPCA) technique was an attempt to apply this approach to a moving platform. If multiple antenna elements were displaced along the radar's direction of travel at a specific distance determined by the speed of the radar and the PRF, antenna phase centers for successive pulses could be made to have the same *effective* location. Then, measurements obtained at co-located effective phase centers could be applied to the same pulse-canceller filters as were used for stationary radars.

The approach is depicted in Fig. 4-2, which is adapted from [34]. In Fig. 4-2, a radar is moving at velocity, v , and has two array elements offset along the direction of travel. The two elements are separated by a distance, d . Now, let the PRI be chosen




Figure 4-2. Effective array positions for a two-element, two-pulse, DPCA system.

such that the distance traveled by the radar in a single PRI, $v \cdot \text{PRI}$, is equal to half of the element spacing, $d/2$. Because both the transmit and receive elements move, the effective distance moved by the array is d , and as Fig. 4-2 shows, the effective phase center of the leading element on pulse #1 is the same as the effective phase center of the trailing element on pulse #2. In DPCA, these two measurements are subtracted in order to subtract out stationary clutter. When the PRF, element spacing, and radar velocity are constrained to have this relationship, the radar is said to satisfy the DPCA condition. This is often measured with a coefficient, β , defined as

$$\beta = \frac{2v\text{PRI}}{d}. \quad (4.3)$$

4.1.2. Conventional Space-Time Adaptive Processing

While DPCA is both intuitively satisfying and relatively simple to implement, it suffers from some considerable limitations. Specifically, DPCA's clutter cancellation is sensitive to array element mismatch errors and provides no mechanism for dealing with jammers [34]. Fortunately, space-time adaptive processing (STAP) provides a method of space-time filtering that is both more robust than DPCA and is well grounded in statistical theory.

Many adaptive array techniques were developed and studied starting in the late 1960s [58-63]. In 1973, however, Brennan and Reed [64] described the optimum space-time filter and opened the door for the last 30 years of research into space-time adaptive processing. It was shown in [64] that the optimum space-time filter is given by

$$\mathbf{w}_{opt} = a\mathbf{M}_i^{-1}\boldsymbol{\rho}_t \quad (4.4)$$

where a is a constant, and $\boldsymbol{\rho}_t$ is the response matched to the desired target. The vector, $\boldsymbol{\rho}_t$, is analogous to the responses defined for clutter in Chapter 2. The matrix \mathbf{M}_i is the covariance matrix of the measurements due to the interference:

$$\mathbf{M}_i = E\left[\left(\mathbf{d}_c + \mathbf{d}_j + \mathbf{d}_n\right)\left(\mathbf{d}_c + \mathbf{d}_j + \mathbf{d}_n\right)^H\right] \quad (4.5)$$

where \mathbf{d}_c , \mathbf{d}_j , and \mathbf{d}_n are the measured data due to clutter, jammers, and white noise, respectively. The filter in (4.4) optimizes probability of detection (P_d) for a given probability of false alarm (P_{fa}). In practice, of course, the interference statistics are not known a priori and must be estimated. This is where the adaptive part of STAP appears. The interference covariance matrix is estimated from secondary support data that presumably have the same statistics as the interference in the data to which the filter will be applied. Most of the research concerning STAP in recent years has dealt with the implications of this approach, including how much secondary support data are needed for a good estimate of \mathbf{M}_i , how to deal with other targets in the support data, how to deal with non-stationary secondary and primary data, and how often to update \mathbf{M}_i [31, 65-67].

Figure 4-3 demonstrates the behavior of the optimum filter in (4.4) and of space-time filtering in general. In Fig 4-3, the response of the optimum filter versus Doppler and azimuth is shown in dB scale. It can be seen that the mainlobe of the filter response is centered at the target parameters: a normalized Doppler of 0.25 and




Figure 4-3. Angle-Doppler filter response for the optimum filter.

a normalized k_x of 0. The sidelobes of the filter response, however, are interrupted by a null running diagonally through the response. This null rejects ground clutter that has a linear relationship between angle and Doppler.

4.1.3. STAP Performance Metrics

Any processing technique must have a measure with which to assess performance. In the previous chapter on SAR processing, the performance metric was the average squared error between the true and estimated scattering coefficients. In MTI, we are most concerned with the detectability of moving targets; therefore, the most commonly used performance metric is the output signal-to-interference-and-noise ratio, SINR. This metric provides the output signal power to the output interference power including white noise. This metric is useful since SINR is a direct measure of target detectability.




Figure 4-4. Output SINR vs. target Doppler for a target at zero azimuth angle.

In this chapter, SINR will be shown versus various target parameters. For example, output SINR may be shown versus target Doppler with the angular position of the target held fixed. An example of this type of plot for a sample, sidelooking scenario satisfying the DPCA condition is shown in Fig. 4-4. In the figure, target Doppler is normalized to the transmit PRF. Beginning at the left side of the figure, we see that the output SINR for a target at zero Doppler is very low. This makes sense because a target that is directly to the side of the radar that has no Doppler relative to the radar cannot have a ground velocity in the direction of the radar. This is essentially a stationary target and can't be distinguished from clutter at the same location. As the target begins to move relative to the ground and in the direction of the radar, however, SINR increases quickly and the target can be detected. Eventually, as the target's relative Doppler approaches the PRF, the Doppler measurements become ambiguous with a zero Doppler target, and the target again appears identical to clutter.

Another important piece of information that can be obtained from Fig. 4-4 is the minimum detectable velocity (MDV). If for a given probability of false alarm and probability of detection, an output SINR of 10 dB is needed, it can be seen from Fig. 4-4 that targets directly to the side of the radar with a Doppler of more than approximately $0.05 \cdot \text{PRF}$ can be detected. An exception, as mentioned above, is a target whose Doppler is a multiple of the PRF.

4.1.4. Minimum Requirements for MTI

From the preceding discussion on DPCA and STAP, it can be determined that there are two prerequisites that must be met in order for MTI to be possible. First, there must be measurements taken by the radar that are different for moving and stationary scatterers. For traditional MTI, the space and time dimensions are used. While a moving target may have the same Doppler frequency as one or more stationary target pixels, part of that Doppler is due to the velocity of the target. Therefore, the moving target and stationary pixels must be at different angular locations. Likewise, a moving target and a stationary pixel at the same location must have different Doppler frequencies because one scatterer is moving and the other is not. For the constellation concept, this idea does not change. Any moving target must have a different Doppler frequency than a stationary scatterer at the same location. Consequently, it can be distinguished when Doppler and spatial frequency filtering are applied jointly. However, for the constellation concept and non-sidelooking scenarios, Doppler and azimuth spatial frequency may not map directly to the space and time measurement dimensions, and no pre-defined filtering dimensions should be assumed. For example, in forward-looking scenarios, all five sensor parameters: time, frequency, and space, provide information about target range. It does not make sense to filter the frequency samples, time samples, or spatial samples independently. Hence, the filtering dimensions are yet to be determined.

The second item that must be present in order to perform MTI is a low-clutter region in the received spectrum. In other words, measurement degrees of freedom must be available for applying toward moving target detection; the subspace spanned

by the clutter must not fill the entire measurement space. The basis for this is contained in the solution for the ideal detection filter. If the response vector for a moving target is given by \mathbf{p}_t , the filter that maximizes probability of detection is given by

$$\mathbf{w}_{opt} = a\mathbf{M}_i^{-1}\mathbf{p}_t. \quad (4.6)$$

For our purposes, the matrix \mathbf{M}_i is the sum of the clutter and noise covariance matrices,

$$\mathbf{M}_i = \mathbf{M}_c + \mathbf{K}_n. \quad (4.7)$$

Equation (4.6) is the colored-noise matched filter. If the interference is completely white, \mathbf{M}_i is a diagonal matrix and the optimum filter is exactly matched to the response vector. If \mathbf{M}_i contains both white and non-white interference, the optimum filter lies in a region where there is less signal energy but also less interference energy. The optimum filter maximizes output SINR in the presence of colored interference by operating in regions of lower interference.

One can see that if the clutter covariance matrix, \mathbf{M}_c , is full rank, then there are no dimensions where the optimum filter can avoid clutter energy. If \mathbf{M}_c is less than full rank, then subspaces exist where a filter can attempt to detect moving targets in the presence of reduced clutter. In linear algebra terms, if \mathbf{M}_c is full rank, then there is no low-noise subspace in which the optimum filter can operate. In spectral terms, there is no low-clutter region in the received spectrum. Consequently, in order for MTI to work, the number of measurements needed to fully define the clutter spectrum, or subspace, must be less than the number of measurements taken.

Clutter rank is the radar community's term for the degrees of freedom in the received data that are used up by interference other than white noise. It is termed clutter rank because it is also the rank of the clutter covariance matrix, \mathbf{M}_c . If clutter

rank is reduced, more DOFs are available for moving target detection. Today's MTI achieves reduced clutter rank through implementation within a very strict architecture such as the spacing requirements of DPCA. The clutter covariance matrix, \mathbf{M}_c , is a square matrix with the number of rows and columns being equal to the number of data samples. The DPCA condition causes measurements to be repeated, which in turn reduces the rank of \mathbf{M}_c through redundant rows and columns.

It is not necessary, however, to reduce rank through periodic sample spacing. The important thing is that the number of measurements must exceed the degrees of freedom used up by clutter. Restricting DOFs can also be achieved by controlling the illuminated clutter bandwidth and/or the size of the observing sensor. Currently, STAP is applied to each range bin individually. The result is that the space-time data applies to a single range bin, which is essentially one-dimensional. Clutter rank could be controlled by causing redundant measurements, but if not, rank depends on the azimuth extent of the synthetic array and the illuminated azimuth bandwidth. In other words, clutter rank for a single range bin depends on the time-bandwidth product.

For the constellation concept, MTI filtering may need to take place on the entire data set, not just the space-time data subset. Therefore, the total number of measurements must exceed the degrees of freedom used up by a 2D area of illuminated clutter. I will show in the next section that clutter rank for a 2D area depends on the product of the sensor area and the illuminated area, a product known in optics as the *space-bandwidth* product [68-74]. As long as the total number of measurements exceeds the space-bandwidth product, there is a set or subset of data with which moving targets can be detected.

4.2. Clutter Rank

Clutter rank is important because it characterizes the severity of ground clutter. A high clutter rank relative to the measurement dimension means that most of the measurements are required for representation of the clutter, leaving few measurements available for moving target detection. A low clutter rank relative to the measurement dimension, however, means that there are many measurements available for moving target detection. Consider the problem in the following way. If the radar collects N independent measurements, then moving targets can exist anywhere in N -dimensional space. It is already known that there are some regions of that space where moving targets compete with ground clutter: for example, the region that represents zero-azimuth, zero-Doppler targets in a sidelooking scenario. We haven't, however, characterized how much of the N -dimensional space is occupied by clutter, or how powerful the clutter is in those regions. If the interference covariance matrix is known, the optimum solution naturally attempts to detect moving targets in low-clutter-power regions of the measurement space. But how do we know if there will be a low-clutter region of the measurement space, and if so, how large will that region be? These questions are answered by understanding the rank and eigenspectrum of the clutter covariance matrix. The answers to these questions determine whether or not MTI can be performed. In addition, once clutter rank is understood, a high-clutter-power subspace can be defined. Then, an appropriate filter can be designed that rejects the high-clutter-power subspace and searches for moving targets in the remaining low-clutter data.

4.2.1. Brennan's Rule

There is one particular set of assumptions under which clutter rank is well understood. If the array elements are spaced by a half-wavelength precisely along the flight line of the radar, the clutter azimuth angles range from $-\pi/2$ to $\pi/2$ with

equal power, and the PRF is chosen such that the spacing coefficient, β , is an integer, then the clutter rank, r_c , follows Brennan's rule [34]:

$$r_c = N + (M - 1)\beta \quad (4.8)$$

where N is the number of array elements and M is the number of transmitted pulses. Equation (4.8) can be proven by counting the number of unique measurements obtained by the radar [34]. Because β is an integer, some measurements are repeated, as in DPCA, and it is easy to count the number of unique measurements. Unique measurements contribute nonzero eigenvalues to the eigenanalysis of \mathbf{M}_c while redundant measurements each produce a zero eigenvalue. Therefore, if the eigenvalues are plotted, there is a sharp transition between nonzero and zero eigenvalues that can be interpreted as the rank of the interference covariance matrix. Any detector intending to detect moving targets should lie, at least partially, outside of the subspace spanned by the eigenvectors corresponding to nonzero clutter eigenvalues. Figure 4-5 shows a sample *eigenspectrum* for a system satisfying the appropriate assumptions for Brennan's rule to apply. It is seen that the rank for this system is $r_c = 105$. Klemm [35-37, 75] has also studied clutter rank under similarly restrictive circumstances.

Brennan's rule is the most commonly used method of clutter rank prediction. Unfortunately, the set of assumptions necessary for Brennan's rule to be strictly valid rarely apply in practice. In [34], the effects of relaxing some of the assumptions are studied, and it is shown that Brennan's rule is still approximately correct for small crab angles and non-integer values of β . However, Brennan's rule is deemed correct in that it approximately predicts the knee of the now-rounded clutter eigenspectrum. A sample eigenspectrum and clutter rank as predicted by Brennan's rule for $\beta = 1.83$ are shown in Fig. 4-6.




Figure 4-5. A sample clutter spectrum under the assumptions typically assumed for sidelooking STAP.




Figure 4-6. A sample clutter spectrum with rank predicted by Brennan's rule for non-integer β .

4.2.2. The Space-Bandwidth Product

While Brennan's rule is useful for roughly estimating the clutter severity for sidelooking scenarios, the rigid circumstances under which Brennan's rule is valid make a more general clutter characterization desirable. In addition, it was previously mentioned that sidelooking, single-range-bin scenarios are not the only scenarios of interest for a spaceborne system. Specifically, forward-looking scenarios and scenarios that illuminate range and/or Doppler ambiguous clutter are of interest. In some of these scenarios, full space-time-frequency filtering may be required, which will also require a more general clutter characterization.

In [76], the number of unique signals observed by an array over multiple pulses was shown to be dependent on the spatial bandwidth of the signals impinging on the array and the total observation time from the time the first pulse arrived at the first array element to the time the last pulse arrived at the last element. Buckley called this length of time the *temporal aperture*. In other words, the rank of the measurements depended on the observed time-bandwidth product. This theorem, known as the Landau-Pollak theorem, was also observed by Van Trees who said "when a bandlimited process $[-W, W \text{ Hz}]$ is observed over a T -second interval, there are only $(2TW + 1)$ significant eigenvalues" [43].

In fact, Brennan's rule turns out to be the time-bandwidth product for the special case of Nyquist spatial sampling. For element spacing of d and full azimuth illumination, the spatial bandwidth, B_s , is $1/d$. The length of the synthetic aperture, L_s , is $(N-1)d + 2v(M-1)\text{PRI}$. Applying the spatial bandwidth and the synthetic aperture to the Landau-Pollak theorem yields

$$\begin{aligned} r_c &= L_s B_s + 1 = \left[(N-1)d + 2v(M-1)\text{PRI} \right] \frac{1}{d} + 1 \\ &= N + (M-1) \frac{2v\text{PRI}}{d} \end{aligned} \quad (4.9)$$

Recognizing the definition of β , we arrive at Brennan's rule,

$$r_c = N + (M - 1)\beta. \quad (4.10)$$

Therefore, Brennan's rule is a veiled form of the received time-bandwidth product, which for Nyquist spatial sampling and integer β also happens to be the number of unique measurements collected.

Still, these theorems predict the knee of the eigenspectrum but do not give a firm foundation for what that knee represents. Moreover, experimental results show that the knee of the eigenspectrum may be ambiguous. Even Van Trees [43] does not strictly define what is meant by a *significant* eigenvalue. It would be better to be able to predict a specific point in the eigenspectrum, such as the 3-dB point or a point that represents a specific amount of the clutter energy. A series of papers by Slepian, Landau, and Pollak [77-79] and work in the area of optics [68-74] provide a basis for just such a definition of clutter rank.

Slepian, Landau, and Pollak performed a rigorous mathematical analysis of the number of eigenfunctions needed to represent a bandlimited function over a finite time interval. They used strict definitions for time- and bandlimited functions to show that the number of eigenfunctions necessary for representing a signal over a finite time interval to within a small error in the signal's energy was $BT + 1$ where BT was the time-bandwidth product according to their definitions. These authors were also able to show that the optimum functions for representing the signal were the prolate spheroidal wave functions.

The prolate spheroidal wave functions are difficult to work with analytically. Fortunately, for large time-bandwidth products, the prolate spheroids become sinusoidal functions with evenly spaced frequencies [43]. Since a high-resolution radar will inevitably collect thousands of measurements, we can conclude that the appropriate time-bandwidth product is large and that the asymptotic behavior of the prolate spheroids applies. This makes the analysis much simpler, and we can define a method for determining the amount of clutter power represented by a finite number of eigenfunctions.

Van Trees [43] notes that the prolate spheroids are the eigenfunctions of the Karhunen-Loeve expansion of random signals; therefore, the prolate spheroids are the optimal functions for representing the PSD of clutter. Since the prolate spheroids approach sinusoids for large time-bandwidth products, and the PSD of clutter gives the expected clutter power as a function of varying frequency sinusoids, the eigenvalues of the clutter covariance matrix are equal to the height of the clutter's power spectral density (PSD), $S_c(f)$, at the frequency of the corresponding eigenfunction [43]. Furthermore, if the eigenfunctions span a finite bandwidth, B , the amount of average clutter power represented within that bandwidth is the integral of the PSD over the specified bandwidth. Mathematically, the total amount of average power in the clutter is

$$P = \int_{-\infty}^{\infty} |S_c(f)| df, \quad (4.11)$$

the amount of average power within a finite bandwidth, B , is

$$P_B = \int_B |S_c(f)| df, \quad (4.12)$$

and the ratio of powers is defined as

$$\varepsilon_P = P_B / P. \quad (4.13)$$

The number of eigenfunctions, r_c , necessary to span the specified bandwidth is the time-bandwidth product plus one:

$$r_c = BT + 1 \quad (4.14)$$

where T is the length of the observation interval and the interval over which the average clutter power must be approximated. Hence, the average power within a specified band can also be written as a summation of the first r_c eigenvalues of \mathbf{M}_c :




Figure 4-7. The eigenvalues of a random signal's covariance matrix follow the signal's power spectral density.

$$P_B = \Delta f \sum_{i=1}^{r_c} \lambda_i . \quad (4.15)$$

where Δf is the separation between eigenfunction frequencies.

Figure 4-7 compares the eigenspectrum of a random signal to the signal's single-sided PSD. The signal's two-sided PSD is the mainlobe of a Sinc function. The frequencies of the eigenfunction sinusoids are separated by $1/T$. It is seen in Fig. 4-7 that the eigenspectrum does, indeed, follow the singled-sided PSD.

Now that the theory of an eigenfunction representation of ground clutter has been discussed, the theory must be extended to two dimensions in order to account for the 2D spectrum that will be received from the Earth's surface. It has been shown that the number of degrees of freedom in an optical image is equal to the space-bandwidth product [68, 71-74], which is the 2D extension of the time-bandwidth product. In the

space-bandwidth product, the bandwidth is an area of the 2D spectrum, and the space is an area of a 2D observation region. The space-bandwidth product provides a direct extension of the eigen-representation just discussed for 1D signals. In the 2D case, however, the observation and spectral dimensions are measured as areas rather than widths, the eigenfunctions are two dimensional, and the eigenvalues take on values of the 2D PSD. If the spectral area is defined as A_B and the observation area is defined as A_S , the power in the eigen-representation is

$$P_B = \iint_{A_B} |S_c(f_\alpha, f_\beta)| df_\alpha df_\beta, \quad (4.16)$$

and the number of eigenfunctions needed to represent the power is

$$r_c = A_B A_S + 1. \quad (4.17)$$

Now we can finally relate the clutter statistics to the DOFs used up by clutter and use the eigenfunctions to define a high-clutter region and a low-clutter region of our measurement space. If, for example, we desire to reject 80 percent of the expected clutter power, we use (4.16) and (4.17) to find the necessary DOFs. First, we find the smallest possible sub-region of the clutter PSD that represents 80 percent of the total expected clutter power. Then, we compute the size of the 2D sensor, and the 2D clutter rank is the product of the two areas.

Consider the following example. Let a sidelooking radar's illumination pattern on the Earth be a 2D Gaussian given by

$$g(f_\alpha, f_\beta) = \exp\left(-\frac{1}{2} \Delta \mathbf{f}^\dagger \mathbf{J}_f^{-1} \Delta \mathbf{f}\right). \quad (4.18)$$

The expected power received from a clutter patch is

$$S_c(f_\alpha, f_\beta) = \frac{P_p g(f_\alpha, f_\beta)^2 \sigma(f_\alpha, f_\beta) \lambda^2}{(4\pi)^3 R(f_\alpha, f_\beta)^4}. \quad (4.19)$$

If we assume that the percent variation in the clutter spectrum due to changes in range, $R(f_\alpha, f_\beta)$, are small and that the expected value of the clutter RCS, $\sigma(f_\alpha, f_\beta)$, is constant over the illuminated area, then the two-dimensional PSD of clutter is

$$S_c(f_\alpha, f_\beta) = \frac{P_p g(f_\alpha, f_\beta)^2 \sigma \lambda^2}{(4\pi)^3 R^4}, \quad (4.20)$$

and the variation is seen to depend only on the radiation pattern. Since the radiation pattern is Gaussian, the 2D PSD is also Gaussian, and any spectral region of maximum power will be an ellipse oriented with the Gaussian. Such an ellipse with variable size defined by a constant, a , is

$$\Delta \mathbf{f}^\dagger (2\mathbf{J}_f^{-1}) \Delta \mathbf{f} = a^2, \quad (4.21)$$

the ratio of the average power within the elliptical area to the total average power is [80]

$$\frac{P_B}{P} = 1 - \exp\left(-\frac{1}{2}a^2\right), \quad (4.22)$$

and the area of the ellipse is

$$A_B = \frac{1}{2} \pi a^2 |\mathbf{J}_f|^{1/2}. \quad (4.23)$$

Next, let's further assume that the observing sensor is a planar, well sampled array of widths L_α and L_β . The area covered by the sensor is

$$A_S = L_\alpha L_\beta. \quad (4.24)$$

If the elliptical bandwidth was chosen to define an area of the clutter PSD that encompasses 80 percent of the total average power, the number of eigenfunctions needed for representing 80 percent of the clutter power is

$$r_c = A_S A_B + 1 = \frac{1}{2} L_\alpha L_\beta \pi a^2 |\mathbf{J}_f|^{1/2} + 1. \quad (4.25)$$

4.2.3. Estimation of Sensor Size

The space-bandwidth product is the product of two areas. The bandwidth is a region of the 2D PSD of clutter, and the space is the 2D area of the sensor. Most radars, however, collect measurements over the aforementioned five sensor parameters. How is a 2D area obtained from a 5D sensor parameter set? The 2D area of the sensor is obtained from the five measurement parameters through the synthetic aperture projection that was derived in Chapter 2. The sensor transformation matrix projects the sensor into two dimensions and projects the 2D illumination pattern into a 2D PSD with the proper units for measurement by the synthetic aperture. The space-bandwidth product predicts the amount of average signal power represented by a finite number of eigenfunctions when the sensor is a continuous, uniform, two-dimensional aperture. The radar, however, only collects data at specific sampling points in time, space, and frequency, and those points project to sample points in the 2D synthetic aperture. Therefore, the radar makes observations at specific points in a 2D synthetic *array* rather than over a continuous 2D synthetic aperture. If the samples are close enough in two dimensions, we know from Nyquist that we don't need the samples in between. The fact that the synthetic aperture is sampled, though, makes the notion of a sensor area somewhat difficult to define. In the following, I will provide some ideas and examples on how to approximate the 2D sensor area.

If the tapers on the sensor parameters are uniform, the first potential method for estimating sensor area is to project the sample locations into their respective




Figure 4-8. The 2D synthetic array and sensor boundary for a sparse, five-aperture system.

locations in the 2D array and estimate the overall area through visual inspection. In Fig. 4-8, a radar system with five receivers has been projected into the appropriate synthetic array. A boundary has also been drawn around the synthetic array. The sensor area can be estimated using a piecewise summation of areas. Although this approach is awkward and must be repeated for any change in the sensor, the results are usually very good. In Fig. 4-9, the theoretical fraction of clutter power contained within a finite number of eigenfunctions is compared with the true power fraction obtained numerically. The theoretical curve represents, for each bandwidth and PSD fraction, the number of required eigenfunctions according to the space-bandwidth product. The numerical curve represents the number of required eigenfunctions obtained by summing the computed eigenvalues of the clutter covariance matrix. In Fig. 4-9, the PSD is a Hanning-squared function.




Figure 4-9. Theoretical vs. actual clutter power contained within a specified clutter rank when the sensor area is approximated using a piecewise, visual approach.

A second method of estimating the sensor area is to use the sampled second moments of the sensor. This is similar to how the second moments were used in Chapter 2 to estimate the width of the sensor for resolution purposes, but in this case, the second moments will be used to estimate an effective area of the sensor. First, the second-moment equation is calibrated to an elliptical, uniformly weighted sensor. Then, the calibrated equation is used to predict the effective sensor area for more general sensor shapes.

For calibration purposes, assume that an elliptically shaped 2D sensor is uniformly weighted, and that the boundary of the ellipse is described by

$$\frac{x^2}{\sigma_x^2} + \frac{y^2}{\sigma_y^2} = a^2. \quad (4.26)$$

The true area of the elliptical sensor is

$$A_e = \pi a^2 \sigma_x \sigma_y. \quad (4.27)$$

From (4.26), the calibration ellipse is oriented along the x - and y - axes. Therefore, the off-diagonal elements of this sensor's mean-squared-width matrix, \mathbf{J} , will be zero. The mean-squared width in the x -direction is

$$J^{1,1} = \frac{1}{A_e} \iint_{A_e} x^2 dy dx = \frac{a^2}{4} \sigma_x^2. \quad (4.28)$$

Likewise, the y -component is

$$J^{2,2} = \frac{1}{A_e} \iint_{A_e} y^2 dy dx = \frac{a^2}{4} \sigma_y^2, \quad (4.29)$$

giving a sensor second-moment matrix of

$$\mathbf{J} = \frac{a^2}{4} \begin{bmatrix} \sigma_x^2 & 0 \\ 0 & \sigma_y^2 \end{bmatrix}. \quad (4.30)$$

Now, since \mathbf{J} is a measure of width squared, it is hypothesized that the sensor area is approximately related to the square root of the determinate of \mathbf{J} . The sensor area equation is proposed as

$$A_S = \pi a^2 \sigma_x \sigma_y \approx b |\mathbf{J}|^{1/2} \quad (4.31)$$

where b is a constant. Substituting (4.30) and solving (4.31) for b gives $b = 4\pi$, and the sensor area estimate is

$$\hat{A}_S = 4\pi |\mathbf{J}|^{1/2}. \quad (4.32)$$

Equation (4.32) can be used to estimate the area of a 2D sensor when the sensor's size is characterized in terms of second moments. The matrix, \mathbf{J} , is the sensor covariance

matrix that is analogous to the second-moment matrices used for resolution in Chapter 2.

As an example and to test (4.32), consider a rectangular aperture of widths L_x and L_y . The true area is $A_S = L_x L_y$. The sensor's second-moment matrix is

$$\mathbf{J} = \begin{bmatrix} \frac{L_x^2}{12} & 0 \\ 0 & \frac{L_y^2}{12} \end{bmatrix}. \quad (4.33)$$

Using (4.32), the area estimate is

$$\hat{A}_S = \frac{\pi}{3} L_x L_y = \frac{\pi}{3} A_S, \quad (4.34)$$

which is in error by about five percent.

A significant advantage of this approach to estimating sensor area is that the integrals that define the elements of \mathbf{J} can be approximated for sampled synthetic arrays using the sample variance. In addition, any sensor tapers such as time weighting can be included. For example, if a synthetic array has been created through projection using the sensor transformation matrix, the component of \mathbf{J} corresponding to the α synthetic dimension, $J^{1,1}$, can be approximated as

$$J^{1,1} = \frac{1}{N_{meas}} \sum_{i=1}^{N_{meas}} w_s(\tilde{r}^i) (\tilde{r}_\alpha^i - \bar{\tilde{r}}_\alpha)^2. \quad (4.35)$$

where $w_s(\tilde{r}^i)$ is the sensor taper at the point, \tilde{r}^i , in the synthetic aperture. Figure 4-10 demonstrates the prediction of clutter rank using the sample variance estimate of \mathbf{J} in (4.32). In Fig. 4-10, the radar's illumination pattern was the mainlobe of a Sinc




Figure 4-10. Theoretical vs. actual clutter power contained within a specified clutter rank when the sensor area is approximated using the synthetic array's sampled 2nd moments.

function and the sensor taper was uniform. In Fig. 4-11, the same situation applies except that the time samples and frequency samples were given Hanning tapers. Figures 4-10 and 4-11 show that the estimated area gives a good prediction of the number of eigenfunctions necessary to represent clutter for sampled systems with various sensor tapers.

The last method of estimating sensor area makes use of the full synthetic aperture notation developed in Chapter 2. Therefore, this method approximates effective sensor area when Gaussian sensor tapers are used for each of the time, frequency, and space dimensions. Using the results of Chapter 2, the signal measured at a set of sensor parameters, $\Delta \mathbf{s}$, due to a scatterer at location, $\Delta \mathbf{x}$, when the sensor tapers are Gaussian, is

$$d(\Delta \mathbf{x}, \Delta \mathbf{s}) = b \exp\left(-\frac{1}{2} \Delta \mathbf{s}^\dagger \mathbf{J}_s^{-1} \Delta \mathbf{s}\right) \exp\left(-\frac{1}{2} \Delta \mathbf{x}^\dagger \Lambda_l \mathbf{J}_l \Lambda_l^\dagger \Delta \mathbf{x}\right) \exp(-j \Delta \mathbf{x}^\dagger \Lambda_s \Delta \mathbf{s}) \quad (4.36)$$




Figure 4-11. Theoretical vs. actual clutter power contained within a specified clutter rank when the sensor area is approximated using the synthetic array's sampled 2nd moments. Hanning windows were used for the time and frequency tapers.

where b is a constant. In (4.36), the first exponential is the sensor amplitude taper, the second exponential is the illumination pattern, and the third exponential is the sensor-target phase response. Equation (4.36) can be transformed into synthetic aperture coordinates using the SVD of the sensor transformation matrix,

$$\mathbf{USV}^\dagger = \mathbf{\Lambda}_s . \quad (4.37)$$

Let

$$\tilde{\mathbf{r}} = \mathbf{SV}^\dagger \Delta \mathbf{s} , \quad (4.38)$$

and

$$\tilde{\mathbf{k}} = \mathbf{U}^\dagger \Delta \mathbf{x} . \quad (4.39)$$

The first exponential in (4.36) is the 5D sensor taper, which is set equal to the sensor taper in synthetic coordinates. Defining a matrix, \mathbf{J}_v , that contains the widths of the 2D Gaussian taper in synthetic coordinates and setting the tapers in 5D-real and 2D-synthetic coordinates equal:

$$\exp\left(-\frac{1}{2}\Delta\mathbf{s}^\dagger\mathbf{J}_s^{-1}\Delta\mathbf{s}\right) = \exp\left(-\frac{1}{2}\tilde{\mathbf{r}}^\dagger\mathbf{J}_v^{-1}\tilde{\mathbf{r}}\right), \quad (4.40)$$

yields

$$\mathbf{J}_v = \mathbf{S}\left(\mathbf{V}^\dagger\mathbf{J}_s^{-1}\mathbf{V}\right)^{-1}\mathbf{S}. \quad (4.41)$$

Likewise, setting the antenna radiation patterns in real and synthetic coordinates equal and defining an illumination width matrix in synthetic coordinates, \mathbf{J}_u , yields

$$\mathbf{J}_u = \mathbf{U}^\dagger\left(\Lambda_l\mathbf{J}_l\Lambda_l^\dagger\right)^{-1}\mathbf{U}. \quad (4.42)$$

Using (4.41) and (4.42), the measurements can be written, approximately, in synthetic sensor coordinates as

$$d(\tilde{\mathbf{k}}, \tilde{\mathbf{r}}) = b \exp\left(-\frac{1}{2}\tilde{\mathbf{r}}^\dagger\mathbf{J}_v^{-1}\tilde{\mathbf{r}}\right) \exp\left(-\frac{1}{2}\tilde{\mathbf{k}}^\dagger\mathbf{J}_u^{-1}\tilde{\mathbf{k}}\right) \exp(-j\tilde{\mathbf{k}}^\dagger\tilde{\mathbf{r}}). \quad (4.43)$$

Equation (4.43) says that the synthetic sensor has a bivariate Gaussian taper with widths described by \mathbf{J}_v . It also shows that the illumination pattern is bivariate Gaussian with widths described by \mathbf{J}_u . The second-moment approximation previously described can now be used to estimate the effective sensor area. Noting that $\tilde{\mathbf{k}}^\dagger\tilde{\mathbf{r}}$ is in radians, I factor out a $(1/2\pi)^2$ to get the approximate sensor area from (4.32),

$$\hat{A}_S = 4\pi \left| \frac{1}{4\pi^2} \mathbf{J}_v \right|^{1/2} = \frac{1}{\pi} |\mathbf{J}_v|^{1/2}. \quad (4.44)$$

Next, since the illumination is Gaussian, the spectral area is elliptical and oriented with the Gaussian pattern. An elliptical bandwidth can be defined as

$$\tilde{\mathbf{k}}^\dagger (2\mathbf{J}_u^{-1}) \tilde{\mathbf{k}} = b^2, \quad (4.45)$$

which has an area of

$$A_B = \frac{1}{2} \pi b^2 |\mathbf{J}_u|^{1/2} = \frac{1}{2} \pi b^2 \left| \mathbf{U}^\dagger (\boldsymbol{\Lambda}_l \mathbf{J}_l \boldsymbol{\Lambda}_l^\dagger)^{-1} \mathbf{U} \right|^{1/2}. \quad (4.46)$$

Substituting (4.41), (4.44), and (4.46) into the space-bandwidth product, the corresponding clutter rank is

$$\begin{aligned} r_c &= \hat{A}_S A_B + 1 \\ &= \frac{1}{2} b^2 \left| \mathbf{S} (\mathbf{V}^\dagger \mathbf{J}_s^{-1} \mathbf{V})^{-1} \mathbf{S} \right|^{1/2} \left| \mathbf{U}^\dagger (\boldsymbol{\Lambda}_l \mathbf{J}_l \boldsymbol{\Lambda}_l^\dagger)^{-1} \mathbf{U} \right|^{1/2} + 1, \end{aligned} \quad (4.47)$$

and using the formula for the integral of a bivariate Gaussian over an elliptical region [80], the fraction of power represented is

$$\frac{P_B}{P} = 1 - \exp\left(-\frac{1}{2} b^2\right). \quad (4.48)$$

The advantage of (4.41), (4.44), and (4.47) is that the sensor area can be approximated directly from the widths of each of the five real sensor parameters. Instead of first projecting into a 2D synthetic array and finding the effective area of the projection, the area is obtained from the five real measurement parameters.




Figure 4-12. Theoretical vs. actual clutter power versus clutter rank when the sensor area is approximated using the 2nd moments of the five sensor parameters directly. Hanning windows were used for the time and frequency tapers.

Therefore, rank can be characterized using rms bandwidth, integration time, and physical array size, and it is likely that these real parameter widths are better understood than the equivalent 2D widths. For example, if the frequency spectrum of the transmit pulses and the time windowing of the pulses are known, they can be used directly. Figure 4-12 shows the results of rank prediction using the widths of the five sensor parameters directly. The rms widths of the Hanning time and frequency tapers and uniform array taper are plugged directly into (4.41), and (4.44) is used to approximate the effective sensor area. Figure 4-12 shows that the theoretical clutter rank again agrees with the clutter rank found through numerical simulation.

Figure 4-13 shows the eigenspectrum from the simulation of Fig. 4-12. The eigenvalues of the clutter covariance matrix for the tapered sensor are shown. According to the second-moment area approximation using the five sensor widths




Figure 4-13. The eigenspectrum for a sensor with Hanning time and frequency windows. The illumination pattern was also Hanning. The average clutter power to the left of the dashed line is 85 percent of the total average clutter power.

directly, the clutter rank corresponding to 85 percent clutter rejection is 24. This cutoff is marked by the dashed line in Fig. 4-13. Note that there is no obvious knee in the eigenspectrum of Fig. 4-13; therefore, there is no obvious cutoff point to use in defining the clutter bandwidth. However, the eigenfunctions associated with the eigenvalues to the left of the dashed line can be used to define a clutter subspace. On average, this clutter subspace would contain 85 percent of the total clutter power.

4.2.4. Clutter Rank versus Constellation Diameter and Density

In the previous chapter, the importance of the density of the synthetic array was discussed in terms of SAR processing. For SAR, requiring a sufficient sampling density was equivalent to requiring that resolution be determined by bandwidth and integration time, rather than by the beamwidth of the constellation. This allowed the




Figure 4-14. Predicted clutter rank vs. constellation size for constellations with five, eight, and 12 receivers. The dashed lines indicate the number of measurements obtained by each constellation.

physical array to be used for nulling range-Doppler ambiguities. In this section, it is shown that the density of the synthetic array is also important for MTI processing.

It was just demonstrated that clutter rank for a 2D illuminated area depends on the space-bandwidth product of the synthetic sensor. This directly affects the size of allowable constellations and the number of satellites that must fill the constellation. If the constellation size is increased, rank goes up according to the respective increase in sensor area. If the number of satellites is kept constant while this size increase takes place, the total number of measurements obtained remains the same; consequently, increasing constellation diameter without increasing the number of satellites amounts to reducing the DOFs available for moving target detection.

Figure 4-14 demonstrates the constellation-size/number-of-receivers dilemma. In Fig. 4-14, the predicted clutter rank using the second-moment sensor approximation technique is plotted versus a scale factor for the constellation size. The scale factor is

a constant that multiplies the receiver positions relative to the array phase reference. In Fig. 4-14, the constellation scale factor is normalized to the smallest constellation size simulated. Note that a scale factor of two does not necessarily mean that the synthetic aperture has twice the diameter or area as for a scale factor of one. The three data curves are for simulations with five, eight, and 12 receivers. First, the simulation for five receivers resulted in a total number of measurements collected equal to 400, which is shown as a dashed, horizontal line. The clutter rank starts at approximately 100 and increases slowly at first as the synthetic array samples gradually become non-overlapping. As the synthetic array samples due to each satellite become separated, the predicted rank increases more quickly with the array scale factor. Similar plots are shown for eight and 12 receivers, which collected 640 and 960 measurements, respectively.

The significance of Fig. 4-14 comes when considering the available DOFs for moving-target detection versus the MDV. Remembering that wide angular coverage improves MDV and is a driving force behind the constellation concept, let's assume that a scale factor of 20 results in a spaceborne physical array large enough to meet an MDV specification. For the constellation of five receivers, this scale factor results in a clutter rank of approximately 390, which means that 97.5 percent of the measurements must be used just to reject clutter. Clearly, the remaining 10 measurement dimensions will be insufficient for detecting moving targets.

For the constellations of eight and 12 receivers, the situation improves dramatically. In both of these cases, a scale factor of 20 results in a clutter rank of approximately 425. For the constellation of eight receivers, this clutter rank corresponds to approximately 66 percent of the measurement space, which leaves 34 percent of the measurement space open for moving target detection. For the constellation of 12 receivers, however, a clutter rank of 425 corresponds to 44 percent of the measurement space, which leaves more than half of the measurement space open for moving target detection.

4.2.5. Synthetic Co-Array Bound on Clutter Rank

I began this section on clutter rank by introducing Brennan's rule, which is the clutter rank prediction method used today for sidelooking STAP. As mentioned, Brennan's rule is strictly correct under certain assumptions that lead to the number of unique measurements collected being equal to the product of the synthetic aperture length and the Doppler bandwidth. This section discusses the foundation for the *unique-measurement* bound on clutter rank that is found in the structure of the synthetic co-array.

The optimum solution given in (4.4) can detect targets in two situations. The first situation is where the target's RCS is much larger than anything nearby including clutter. In this case, just about any detection filter somewhat matched to the target will work. Unfortunately, it is not a very likely scenario. The more probable situation is where the measurement space contains both low- and high-clutter power regions. Targets in the high-power regions appear similar to ground clutter and will probably be missed. Targets in the low-power regions, however, can be detected by exploiting the angle-Doppler relationship of clutter. As has been discussed, the second scenario occurs when the rank of the clutter covariance matrix is much lower than the matrix's dimension.

Going further, it is noted that each element of the clutter covariance matrix, \mathbf{M}_c , is a measure of the expected correlation between two measurements where the signal at the two sample locations is due to clutter. If \mathbf{d}_c represents the vector of data samples due to clutter only, each element of \mathbf{M}_c is given by

$$M_c^{i,j} = \mathbb{E} \left[d_c^i (d_c^j)^* \right]. \quad (4.49)$$

The i and j components of the data vector, moreover, are collected at the i^{th} and j^{th} synthetic array locations: \tilde{r}^i and \tilde{r}^j , and (4.49) is more explicitly written as

$$M_c^{i,j} = E \left[d_c(\tilde{r}^i) (d_c(\tilde{r}^j))^* \right]. \quad (4.50)$$

Next, assuming that the clutter statistics are stationary, we can write (4.50) in terms of a 2D lag between sample locations, $\boldsymbol{\tau}$, rather than each individual sample location,

$$M_c(\boldsymbol{\tau}) = E \left[d_c(\tilde{r}) (d_c(\tilde{r} - \boldsymbol{\tau}))^* \right]. \quad (4.51)$$

It is seen in (4.51) that each element of \mathbf{M}_c represents the correlation between two measurements separated in synthetic coordinates by a 2D lag, $\boldsymbol{\tau}$. Therefore, estimating the clutter covariance matrix is equivalent to estimating clutter's 2D autocorrelation function. In addition, we know that a random signal's power spectral density function is the Fourier Transform of the signal's autocorrelation function. This not only confirms that each eigenvalue of the clutter covariance matrix represents a portion of the power spectral density, but also validates use of the word *eigenspectrum* and the discussion of MTI in both algebraic and spectral language.

The lags sampled by the radar system are represented graphically by the co-array presented in Chapter 2. Each dot in the co-array represents a lag that is represented in the clutter, noise, and interference covariance matrices. When lags are repeated in the co-array, a better estimate of that lag's statistical correlation is obtained through averaging, but this improvement comes at the expense of sampling a different, unique lag. Furthermore, since clutter's power spectral density is the Fourier Transform of its autocorrelation function, missing lags in the autocorrelation estimate can cause ambiguity in the corresponding estimate of the power spectral density function. On the other hand, if a radar system is designed such that lags are repeated while still satisfying the Nyquist criterion, this must come at the expense of estimating clutter's correlation function over a wider area. This effectively windows the autocorrelation estimate, which, in turn, smooths the power spectral density estimate. Reducing the resolution of clutter's PSD through smoothing causes the clutter spectrum to collapse into a smaller dimension of the measurement space, thereby opening up regions of the

measurement space for moving target detection. This leads to the conclusion that unique values of the estimate of the PSD of clutter, also known as clutter rank, can be limited by redundant lags in the synthetic co-array.

DPCA is the situation that minimizes PSD resolution while maximizing averaging without causing ambiguity. The repeated measurements that DPCA produces show up as redundantly sampled lags in the space-time co-array, and for this situation, the number of unique measurements obtained is also the number of uniquely sampled space-time lags. Therefore, regardless of the Doppler bandwidth arriving at the radar, the number of unique space-time lags is an upper bound to clutter rank, and limiting clutter rank in this manner can open up DOFs for moving target detection.

For the constellation concept, there are both periodic and random spacings in the 2D synthetic array and co-array. The periodic samples appear because a periodic pulse train has been assumed for a transmit signal. The random spacings occur because of the random separations between satellites. While the periodic samples will produce some redundant lags in the 2D synthetic co-array, it is unlikely that a randomly spaced constellation will ever produce enough redundant lags to actually limit clutter rank. This implies that the size of the constellation, the number of satellites in the constellation, and the illumination width will always be the critical components that determine clutter rank. If the constellation concept is to work for MTI, the number of receivers must be high enough to avoid saturation of the measurement space with clutter, as was pointed out in the previous sub-section. It is possible that a different, non-periodic transmit signal could cause enough lags to be repeated such that clutter rank could be limited in this manner. However, this will be left to ideas for future work.

4.3. MTI Filtering

4.3.1. Derivation of Orthogonal Projection Filter

In Chapter 3, the SAR problem was described as one of estimating a vector of reflectance magnitudes through linear filtering of the data vector, \mathbf{d} . The data vector was made up of noise as well as the desired signal due to scattering from the Earth's surface,

$$\mathbf{d} = \mathbf{d}_c + \mathbf{n} . \quad (4.52)$$

In this Chapter, the data are assumed to contain one or more moving targets in addition to noise and ground reflections,

$$\mathbf{d} = \sum_i \mathbf{d}_t^i + \mathbf{d}_c + \mathbf{n} . \quad (4.53)$$

A normalized response vector, $\boldsymbol{\rho}_t^i$, is defined for each moving target. The response is also known as the normalized steering vector. Then, the component of the data due to moving targets is

$$\sum_i \mathbf{d}_t^i = \sum_i \zeta_i \boldsymbol{\rho}_t^i \quad (4.54)$$

where the ζ_i are coefficients that account for the correct received signal power.

In MTI, we not interested in estimating the RCS of the targets as much as we are interested in maximizing the ratio between the output power due to a moving target and the output power due to ground reflections and noise. Maximizing this ratio maximizes P_d for a given P_{fa} . There are literally thousands of possibilities for moving targets at different angles and Dopplers, however, so it is commonly assumed that only one moving target at a time is in the data. With this assumption, the filter that maximizes the stated ratio is (4.4), repeated here:

$$\mathbf{w}_{opt} = a\mathbf{M}_i^{-1}\mathbf{p}_t. \quad (4.55)$$

The matrix, \mathbf{M}_i , in (4.55) is the interference covariance matrix that was discussed in detail in Section 4.2. Therefore, an *a priori* guess about the structure of \mathbf{M}_i is available based on radar geometry, the antenna illumination pattern, synthetic sensor area, and so on. The true structure of \mathbf{M}_i , however, differs from the assumed structure. While some of the difference is due to discrepancies between the assumed and actual radar system parameters and geometry, much of the error in the expected structure of \mathbf{M}_i is due to using the expected value of the RCS of clutter, which inevitably is different from the actual RCS.

Adaptive MTI processing attempts to account for differences in system parameters and clutter scattering by estimating \mathbf{M}_i from secondary data that hopefully have the same statistical properties as the primary data that are being tested for moving targets. Therefore, adaptive MTI attempts to get as close as possible to the optimum solution given in (4.55). The adaptive approach is important, and any final implementation of MTI on a spaceborne platform will need to be adaptive in nature, but the advantages and issues associated with adaptive processing are not the focus here. Instead, I wish to demonstrate that, given a sparse satellite system designed within certain parameters, the information needed for moving target detection is available. Hence, I will not be presenting any adaptive methods.

Instead, I present an orthogonal projection approach to MTI filtering. This method uses the research on clutter rank presented in Section 4.2 to calculate a generalized DPCA filter. First, the desired amount of clutter rejection is chosen in terms of expected clutter power. Then, the clutter rank, r_c , is determined for the specified clutter rejection, and a clutter subspace is defined that is spanned by the first r_c eigenfunctions. The data vector is then projected orthogonal to the clutter subspace, leaving noise, moving target energy, and a fraction of the clutter power. This method of clutter rejection is analogous to the subtraction of clutter performed in DPCA through multi-pulse cancellers, except that the orthogonal interpretation can be

applied to more general antenna arrays. Last, the projected data vector is tested for the presence of moving targets. The orthogonal projection filter is derived mathematically in the following.

In a previous example, the PSD of clutter was shown to be

$$S_c(f_\alpha, f_\beta) = \frac{P_p g(f_\alpha, f_\beta)^2 \sigma(f_\alpha, f_\beta) \lambda^2}{(4\pi)^3 R(f_\alpha, f_\beta)^4}. \quad (4.56)$$

Although presented as an example, this result is general since no antenna pattern, clutter model, or radar geometry are inherent in the equation. We assume that the percent variation in range over the illuminated area is small, leading to a slightly simplified clutter PSD of

$$S_c(f_\alpha, f_\beta) = \frac{P_p g(f_\alpha, f_\beta)^2 \sigma(f_\alpha, f_\beta) \lambda^2}{(4\pi)^3 R_0^4}. \quad (4.57)$$

If the desired fraction of rejected clutter power, P_B , to total clutter power, P , is defined as χ , the next step is to find the smallest spectral region, A_B , such that

$$\chi = \frac{P_B}{P} = \frac{\iint_{A_B} S_c(f_\alpha, f_\beta) df_\alpha df_\beta}{\iint_{\infty} S_c(f_\alpha, f_\beta) df_\alpha df_\beta} = \frac{\iint_{A_B} g(f_\alpha, f_\beta)^2 \sigma(f_\alpha, f_\beta) df_\alpha df_\beta}{\iint_{\infty} g(f_\alpha, f_\beta)^2 \sigma(f_\alpha, f_\beta) df_\alpha df_\beta}. \quad (4.58)$$

Next, the clutter subspace must be represented mathematically so that the data can be projected orthogonal to it. The first step is to determine the number of eigenfunctions necessary for representing the subspace. The sensor area is estimated using an appropriate method, and the predicted number of eigenfunctions necessary for representing a fraction, χ , of the clutter power is

$$r_c = \hat{A}_S A_B + 1. \quad (4.59)$$

Equation (4.59) gives the number of eigenfunctions. Now, the eigenfunctions need to be calculated. The spectral region, A_B , is sampled uniformly at r_c points. Then, the radar response due to each of the spectral points is calculated. When the synthetic sensor is sampled, each spectral response can be represented as a vector, \mathbf{u}_i , which is the eigenfunction for the i^{th} spectral sample. When the responses from each of the spectral points are put into the matrix, \mathbf{U}_{\parallel} , the clutter subspace is defined as

$$\mathbf{U}_{\parallel} = [\mathbf{u}_1 \quad \mathbf{u}_2 \quad \cdots \quad \mathbf{u}_{r_c}]. \quad (4.60)$$

Using (4.60), the projection, \mathbf{d}_{\perp} , of the data vector, \mathbf{d} , orthogonal to the clutter subspace is [42]

$$\mathbf{d}_{\perp} = \left[\mathbf{I} - \mathbf{U}_{\parallel} (\mathbf{U}_{\parallel}^H \mathbf{U}_{\parallel})^{-1} \mathbf{U}_{\parallel}^H \right] \mathbf{d}. \quad (4.61)$$

Equation (4.61) has some interesting properties. First, each of the \mathbf{u}_i have the same magnitude because the amplitude factors are accounted for in the PSD magnitude, not the eigenfunctions. Hence, the orthogonal approach defined by (4.61) treats the first r_c eigenfunctions equally, and it is like replacing the true eigenspectrum with a uniform eigenspectrum of width, r_c . Figure 4-15 demonstrates this step graphically. Uniformly rejecting clutter in this manner is indicative of DPCA approaches. Rather than estimating the clutter spectrum and calculating the optimum detection filter, DPCA techniques attempt to reject clutter completely even if the true clutter scattering at a given Doppler frequency is much less than anticipated. Eliminating the need to estimate clutter adaptively, however, makes the DPCA approach of (4.61) simpler to implement.

Also of interest is that the inverse in (4.61) is smaller than the inverse required by the optimum solution. The optimum solution calls for an inverse of the interference covariance matrix, which is square with the number of rows and columns equal to the




Figure 4-15. The projection filter is equivalent to replacing the true eigenspectrum with a uniform eigenspectrum of width determined by clutter rank.

number of measurements collected. The number of rows of the clutter subspace matrix, \mathbf{U}_{\parallel} , is also equal to the number of measurements; however, the number of columns is equal to the number of spectral samples, r_c . Therefore, the inverse in (4.61) is of a square matrix of dimension r_c . Since the clutter rank should be much smaller than the number of measurements, the inverse in (4.61) should be smaller than required for the optimum solution. The fact that the inverse is smaller is also apparent in Fig. 4-15 since the width of the approximated eigenspectrum is smaller than the full width of the true clutter spectrum.

After the data are projected orthogonal to a fraction of the clutter, they are tested for the presence of moving targets. Since the assumption is that the clutter power has been significantly reduced, the remaining interference is largely white noise, and the matched filter can be used. Remembering that the matched response for the i^{th}

potential combination of a moving target's location and Doppler was defined earlier as $\boldsymbol{\rho}_t^i$, the test statistic for this moving target is

$$\zeta_i = \left| \left(\boldsymbol{\rho}_t^i \right)^H \mathbf{d}_\perp \right| = \left| \left(\boldsymbol{\rho}_t^i \right)^H \left[\mathbf{I} - \mathbf{U}_\parallel \left(\mathbf{U}_\parallel^H \mathbf{U}_\parallel \right)^{-1} \mathbf{U}_\parallel^H \right] \mathbf{d} \right|. \quad (4.62)$$

4.3.2. Sparse-Array MTI Simulations: Small Illumination

The effect of a sparse array on the MTI filter expressed in (4.62) is studied in this section. The performance of (4.62) is investigated in terms of output SINR versus several system parameters such as the number of satellites in the constellation, the size of the constellation, fraction of clutter energy rejected, and satellite positioning error. Because of the high sidelobes inherent in the pattern of the sparse physical array, SINR performance is not expected to be as well defined in the sparse case as in the traditional STAP example demonstrated in Fig. 4-4. The sample simulations will demonstrate this phenomenon.

A baseline scenario is defined for comparison of performance versus particular system parameters. This baseline scenario includes 10 receivers and a constellation size that corresponds to a normalized array scale factor of one. The scale factor is the same as in Section 4.2.4. The scenario is sidelooking, and the baseline filter is designed for 80 percent clutter rejection. The illuminated area is small so that clutter rank is kept small for manageable processing time over a wide set of simulation parameters. Wider illumination areas that include range and/or Doppler ambiguities will be presented later.

First, SINR versus varying clutter rejection is studied. Figure 4-16 shows output SINR versus target Doppler frequency for various levels of clutter rejection. The target position was held constant at a zero azimuth angle. Because the input noise power in this simulation was relatively low compared to the clutter power, the filter that produces the best output SINR is the filter designed to reject the most clutter




Figure 4-16. Output SINR vs. target Doppler for filters with varying levels of clutter rejection.

energy. If the clutter power did not exist, the output SINR would be at approximately 23 dB, as shown in the figure. Therefore, the difference between white-noise performance and colored-noise performance, also known as SINR loss, is seen to range from only about two dB at the edges of the plot to more than 30 dB in the center where target Doppler corresponds to ground clutter. SINR loss can be interpreted as the cost of rejecting the colored clutter power.

If the output SINR required for meeting P_d and P_{fa} specifications was 10 dB, the 90-percent rejection filter is seen to meet that specification for a zero-azimuth target at all Dopplers except for a small region around zero Doppler. The 50-percent and 60-percent rejection filters, however, are seen to exceed 10 dB over less than half of the possible Doppler frequencies. The MDV for the 50-percent and 60-percent cases are greater than $\frac{1}{4}$ of the transmit PRF. Last, the 70-percent and 80-percent rejection filters have multiple regions that fail to meet the SINR requirement. As usual,




Figure 4-17. Output SINR vs. target azimuth angle for filters with varying levels of clutter rejection.

performance is poor for a zero-Doppler target because a zero-azimuth, zero-Doppler target corresponds to ground clutter. There is another region, however, around $\pm 0.25 \cdot \text{PRF}$ where the output SINR also fails to meet requirements. This is due to the high sidelobes of the physical array. When the physical array is focused on a target at zero azimuth, high sidelobes pop up at a different elevation/azimuth angle that corresponds to clutter at a Doppler of approximately $\pm 0.25 \cdot \text{PRF}$. When the clutter coming through these sidelobes is rejected, it causes a loss of output SINR. The loss, however, is not as severe as for clutter competing through the mainlobe or grating lobes, as evidenced by the fact that output SINR in the $\pm 0.25 \cdot \text{PRF}$ regions is not as bad as around zero Doppler.

Figures 4-17 and 4-18 show output SINR versus other target parameters for filters with various clutter rejection percentages. Figure 4-17 plots SINR versus target




Figure 4-18. Output SINR vs. target azimuth and elevation for filters with varying levels of clutter rejection.

azimuth for a target at a constant Doppler of zero. The notch at zero azimuth (sidelooking), therefore, coincides with a zero-azimuth, zero-Doppler target that has the same parameters as clutter. Figure 4-18 shows output SINR as a grayscale image versus two parameters: target azimuth and elevation. Each of the grayscale images corresponds to a different clutter rejection percentage. As the percentages are increased from 50 percent to 90 percent, the images become slightly whiter, which indicates improved SINR. Another feature seen in Fig. 4-18 is the decreased SINR at the edges of the image. This is due to a reduction in the received signal power at the edges of the illumination pattern. The effect of the illumination pattern on received signal power was not factored into Figs. 4-16 and 4-17 in order to isolate performance of the clutter rejection filter relative to a constant input signal power. The color scale is shown in the bottom, right-hand corner of Fig. 4-18.




Figure 4-19. Output SINR vs. target azimuth and Doppler for filters with varying levels of clutter rejection.

As a final comparison between filters of various clutter rejection percentages, Fig. 4-19 shows output SINR as a grayscale image versus azimuth and Doppler. The images again become progressively whiter for increased clutter rejections. The common feature among all the images in Fig. 4-19 is the deep null running diagonally along a linear azimuth-Doppler relationship. This diagonal is known as the clutter ridge.

The next set of simulations compares output SINR performance versus the size of the constellation. The number of satellites is maintained at the baseline level of 10. The size of the constellation formed by those 10 satellites, however, is varied according to a scale factor normalized to the baseline constellation size. The varying constellation size, of course, changes sensor area and clutter rank. For each scale factor, the rank is chosen such that 80 percent of the clutter is rejected.




Figure 4-20. Output SINR vs. target Doppler for constellations of varying size.

In Fig. 4-20, output SINR is shown versus target Doppler for various scale factors. It is seen in Fig. 4-20 that the width of the notch at zero Doppler decreases with increasing constellation size. This is because as the constellation increases in size, the angular resolution improves and the radar is better able to resolve slowly moving targets from nearby ground clutter at the same Doppler frequency. This is precisely why large angular coverage is needed for performing MTI from a spaceborne platform. As the constellation gets too big with the same number of receivers, however, the physical array becomes more sparse and sidelobes of the physical array pattern degrade performance. In Fig. 4-20, the largest two scale factors produce a narrow null at zero Doppler, but the loss compared to the white-noise-only case is at least 10 dB over most of the unambiguous Doppler spectrum.

Figures 4-21 and 4-22 again show performance versus target azimuth and versus target azimuth and elevation, respectively. As in Fig. 4-20, the clutter notch gets




Figure 4-21. Output SINR vs. target azimuth for constellations of varying size.




Figure 4-22. Output SINR vs. target azimuth and elevation for constellations of varying size.




Figure 4-23. Output SINR vs. target Doppler for constellations with varying number of satellites.

smaller as increased constellation size improves angular resolution, but eventually the constellation gets too sparse, and performance degrades over most azimuth angles. In the azimuth-elevation images, the clutter notch narrows with increasing constellation size, and most of the illuminated area whitens, implying that output SINR improves. The largest scale factor, however, goes beyond a sustainable sampling density, and output SINR degrades.

Figures 4-23 through 4-25 repeat the previous three figures except that each curve or image is now for a different number of receivers rather than a different constellation size. In these simulations, the approximate area of the synthetic sensor was maintained constant so that clutter rank would remain constant. The result is that as more receivers are added, more measurements are available for detection of




Figure 4-24. Output SINR vs. target azimuth for constellations with varying number of satellites.




Figure 4-25. Output SINR vs. target azimuth and elevation for constellations with varying number of satellites.

moving targets. In other words, the proportion of the clutter subspace to the entire measurement space decreases, and performance improves. This performance improvement as receivers are added is seen in each of Figs. 4-23 through 4-25.

Both the constellation size and the number of receivers relate to sampling density. A large constellation is desired for angular resolution, but there must be enough satellites in the constellation to maintain sufficient sampling density. The sampling density can be evaluated visually using the projection of data samples into the 2D synthetic array, or it can be evaluated numerically by comparing clutter rank to the number of measurements collected.

The last set of simulations in this section evaluates SINR performance versus deviation of the satellite receivers from their assumed positions. For this simulation set, the clutter subspace and moving target responses were calculated using incorrect, assumed satellite positions. However, the output SINR was calculated using accurate positions. The incorrect satellite positions differed from the correct positions according to a Gaussian distribution of standard deviation, σ_p , in each dimension. Therefore, the standard deviation of the total positioning error for each receiver was

$$\sigma_e = \sqrt{3}\sigma_p, \quad (4.63)$$

just as was done for SAR in the previous chapter.

The previous chapter showed that the benefits of multiple apertures could be realized for total satellite positioning errors of less than approximately $\frac{1}{10}$ of a wavelength. Figures 4-26 through 4-28 demonstrate that $\frac{1}{10}$ of a wavelength is a good rule of thumb for MTI as well. In Fig. 26, the degradation in output SINR is minimal for positioning errors up to approximately $\frac{1}{20}$ of a wavelength. Some degradation is seen for nearly $\frac{1}{10}$ of a wavelength, and considerable degradation is seen for errors over $\frac{1}{10}$ of a wavelength. Similar conclusions can be made when viewing the positioning error analyses versus target azimuth and versus target azimuth and elevation in Figs. 4-27 and 4-28, respectively.




Figure 4-26. Output SINR vs. target Doppler for constellations with varying levels of satellite positioning errors.




Figure 4-27. Output SINR vs. target azimuth for constellations with varying levels of satellite positioning errors.




Figure 4-28. Output SINR vs. target azimuth and elevation for constellations with varying levels of satellite positioning errors.

The simulations in this section demonstrate the principles that have been discussed in this paper. The first set of simulations investigated performance versus the amount of clutter rejected. It was seen that output SINR generally improved as the amount of clutter power rejected was increased. This increase, however, comes at the expense of increased computation due to increased clutter rank.

The second and third sets of simulations investigated performance versus the size and density of the 2D synthetic array. It was shown that the clutter ridge narrows as the size of the constellation increases due to improved angular resolution of the physical array. This, in turn, improves minimum detectable velocity. If the constellation becomes too sparse, however, SINR performance degrades due to rapidly increasing clutter rank. The way to mitigate the sparse spatial sampling is, of course, to add more satellites. In the end, the simulations demonstrated in Figs. 4-20 through 4-25 show that the size of the constellation and the number of receivers go

hand-in-hand to determine MTI performance. The constellation must be large enough to resolve slowly moving targets from stationary clutter, but while the constellation can be sparsely populated, the simulations show that the sparseness of an operational system is limited.

The last set of simulations demonstrated performance versus satellite positioning errors. As in the SAR case, it was seen that knowledge of the satellite positions to within $\frac{1}{10}$ of a wavelength is needed for quality performance. Beyond $\frac{1}{10}$ of a wavelength, output SINR degrades significantly due to improper clutter subspace specification and due to mismatched target steering vectors.

In this section, the area illuminated by the satellite constellation was kept small in order to reduce computation time and memory requirements. This facilitated the evaluation of performance over a wide set of target parameters and system specifications. One of the stated advantages of a satellite system, however, is to be able to search for moving targets over a wide area due to the increased illumination width of the microsats. In the next section, wide-area MTI is demonstrated.

4.3.3. Sparse-Array, Wide-Area MTI

A wide-area MTI scenario, where range-Doppler ambiguities exist and the 2D clutter rank is in the thousands, is more difficult to process than the limited-illumination scenarios of the previous section. Because clutter rank is so large, a reduced-rank method must be employed. This means that a subset of the data must first be matched filtered in order to reduce the number of measurements and restrict the possible target parameters. In the following simulations, for example, matched filtering is first applied in frequency. This leaves only space-time data remaining and limits moving targets to a specific range and range ambiguities. The space and time dimensions of the radar are then used to resolve range ambiguities, reject ground clutter, and detect moving targets.

The process for the wide-area simulation is as follows. First, the SAR data for an image similar to the one used in Chapter 3 are computed. Next, the responses from

four moving targets are added to the data. These moving targets are within the illuminated area but have non-zero radial velocities relative to the ground and in the direction of the radar. The four moving targets have three distinct radial velocities, which means that two of the targets share the same radial velocity. The RCS of the targets are equal to the average RCS of a single SAR pixel. This agrees with the relative target and clutter models of [57].

As mentioned, the MTI filtering for this simulation begins with range processing. After filtering in frequency, any remaining moving targets must be at the appropriate range or one of its range ambiguities. Next, the space-time data that remain after frequency filtering are projected orthogonal to the clutter subspace. In order to define this subspace, ground clutter at the intended range must be characterized as described in Section 4.2 as well as the ground clutter in ambiguous ranges. The PSD of this clutter depends on both the along-track and cross-track illumination patterns, which for this simulation take the shape of Hanning windows.

Figure 4-29 shows a sample of how the intended range and its ambiguities shape the clutter subspace. In Fig. 4-29, the intended range is at the center of the illuminated pattern. Above and below the intended range are two range ambiguities, and the elliptical curve is a constant-level contour of the illumination pattern. This contour defines the bandwidth of the clutter spectrum that will be used for defining clutter rank. The range lines are in bold where they fall within the spectral bandwidth. Therefore, the range lines in bold contribute to the clutter subspace while the range lines that are not in bold fall outside the specified clutter bandwidth. The contour is chosen such that the expected clutter power in the three bold lines is 95 percent of the total expected clutter power contained within the intended range and its ambiguities.

Once the clutter subspace is defined, the space-time measurements are projected orthogonal to the subspace. Then the remaining data vector is matched filtered for




Figure 4-29. The intersection of constant range lines with the spectral bandwidth contour defines the clutter subspace. The bold range lines contribute to the clutter subspace.

the presence of moving targets at all possible radial velocities and along-track positions. Once the data vector is tested for moving targets, the process is repeated for the next range bin.

The next several figures graphically demonstrate the results obtained through the wide-area MTI simulation. First, the corresponding SAR map is shown in Fig. 4-30. Even though moving targets have been injected into the data vector, they are not apparent in Fig. 4-30 because the ground clutter obscures them. Once the ground clutter is rejected, however, moving targets can be seen. Each of the following figures represents the output of MTI filtering versus range and along-track position for a constant radial velocity.

The first MTI map is the result of testing for moving targets with a very low radial velocity that borders on the MDV limit. Since the radial velocity is so low, the



Figure 4-30. SAR map made from data that include four moving targets.

moving-target matched filters are very similar to the responses of ground clutter. Moreover, since the data have been projected orthogonal to all ground clutter within an elliptical bandwidth, and the matched-filter target vectors are similar to the ground clutter responses within that bandwidth, the output of MTI filtering for such a low radial velocity is negligible within the bandwidth. A moving target at this velocity would have to be very bright in order to show up in the MTI map since much of its energy is stripped away in the clutter rejection process. The MTI image for this low radial velocity, shown in Fig. 4-31, demonstrates the rejection of all low-speed energy within the specified bandwidth.

As the radial velocity increases, however, less moving target energy is stripped away in the orthogonal projection process. In addition, some of the low-power clutter that was outside the clutter bandwidth begins to leak into the bandwidth through sidelobes of the sparse physical array. An example of this is shown in Fig. 4-32. In




Figure 4-31. MTI map for moving targets with a very low radial velocity.




Figure 4-32. MTI map for a radial velocity that is still lower than the slowest target.




Figure 4-33. Histogram and Rayleigh curve fit for the pixel magnitudes of Fig. 4-32.

Fig. 4-32, the radial velocity is still lower than any of the four moving targets that were injected. It is apparent that, as we move away from the zero-velocity clutter subspace, noise and some low-power clutter reappear within the bandwidth.

In order to detect moving targets, a detection threshold must be established. First, a histogram of the pixel magnitudes is computed. This histogram, shown in Fig. 4-33, appears to form a Rayleigh pdf. With this in mind, a Rayleigh curve with the same mean and variance as the histogram is also shown in Fig. 4-33. This curve can be used to find a detection threshold for a desired P_{fa} . Applying a threshold for a P_{fa} of 10^{-6} to Fig. 4-33 yields the image shown in Fig. 4-34.

Figure 4-34 shows six detections, which is far more than expected for a P_{fa} of 10^{-6} and the number of pixels in the image. Most of the extra detections are false alarms due to noise and remaining clutter, and the number of false alarms is higher than expected because of deviation of the pixel magnitudes from a true Rayleigh pdf at the



Figure 4-34. The result of applying a threshold to Fig. 4-32. The threshold was set for a P_{fa} of 10^{-6} using the Rayleigh curve fit of Fig. 4-33.

highest output levels. This deviation probably comes from high-RCS ground clutter just outside the defined clutter bandwidth. The only way to reduce these false alarms is to increase the clutter bandwidth, which results in improved clutter rejection but also increased clutter rank.

A few of the detections, however, are due to moving targets at radial velocities other than the velocity being tested. To explain this, note that there are many combinations of along-track position and radial velocity that combine to create the same Doppler frequency. The beamforming pattern of the constellation is supposed to resolve these possible combinations, but sidelobes of the physical array pattern can cause different along-track-position/radial-velocity combinations to look similar in both space and time. A maximum-likelihood approach to resolving these moving target ambiguities will be discussed in the next section. For now, it is important to




Figure 4-35. MTI map for the radial velocity of the slowest target. Target location is indicated by the arrow.

note that a few of the detections shown in this section are actually multiple detections of the same target and that the multiple detections occur because of the high sidelobes of the sparse array.

The radial velocity tested in Fig. 4-35 is the first velocity that is supposed to have a moving target. One of the four moving targets does show up in Fig. 4-35 at the proper location. The detections obtained based on the same P_{fa} of 10^{-6} are shown in Fig. 4-36. Figures 4-37 through 4-40 demonstrate the same process for two more radial velocities. There are two moving targets at the radial velocity shown in Figs. 4-37 and 4-38. Both of the detections are indicated while other detections are due to either false alarms or moving target ambiguities. The location of the true moving target in Figs. 4-39 and 4-40 is also indicated.



Figure 4-36. The result of applying a threshold to Fig. 4-35. The threshold was set for a P_{fa} of 10^{-6} . The location of the moving target is indicated.




Figure 4-37. MTI map containing two moving targets.



Figure 4-38. The result of applying a threshold to Fig. 4-37. The threshold was set for a P_{fa} of 10^{-6} .




Figure 4-39. MTI map for the highest radial velocity of the four targets.




Figure 4-40. The result of applying a threshold to Fig. 4-39. The threshold was set for a P_{fa} of 10^{-6} .

4.3.4. Angle Ambiguity Resolution

In the previous section, it was mentioned that a single moving target can appear in several different velocity maps. This is because there are many different combinations of along-track position and radial velocity that can combine to create the same Doppler frequency. If a sidelobe of the sparse array falls on one of these combinations where there is actually a target, the target can appear to come from a different location and velocity. This can cause false alarms and confusion by causing multiple detections of the same target. In this section, I describe the angle ambiguity problem in more detail and discuss a ML approach to resolving the ambiguities. The ML approach is similar to the ML filtering that was performed for SAR.

In order to demonstrate that multiple along-track/radial-velocity combinations are ambiguous and that array sidelobes may not completely mitigate the ambiguity, the




Figure 4-41. The along-track-position/radial-velocity ambiguity function generated from time data only.

following along-track/radial-velocity ambiguity functions were calculated. First, the ambiguity for a single range bin was calculated using only the time samples. This time-data-only ambiguity function is shown in Fig. 4-41. The axes of the ambiguity function are along-track position and radial velocity. The desired location and radial velocity are at the center of Fig. 4-41 and the black lines represent all combinations of along-track position and radial velocity that produce an ambiguous Doppler frequency. From Fig. 4-41, it is seen that even for a single radial velocity, there are multiple along-track positions where the target could be located. This is because of the wide antenna patterns and finite PRF that cause Doppler ambiguities to be illuminated.

Next, the ambiguity function for the same target parameters is shown, but the ambiguity function was generated from only the spatial samples. In other words,




Figure 4-42. The along-track-position/radial-velocity ambiguity function generated from spatial data only.

Fig. 4-42 shows the along-track/radial-velocity ambiguity function corresponding to the physical array pattern. Since the array pattern does not vary with a target's radial velocity, the ambiguity function shows no change horizontally. The vertical change is the same as an along-track cut of the array pattern.

The total ambiguity function resulting from both the time and spatial samples is shown in Fig. 4-43. The total ambiguity function is the product of the spatial ambiguity function and the time ambiguity function. Wherever sidelobes in Fig. 4-42 fall on the Doppler ambiguities of Fig. 4-41, the total ambiguity function has high sidelobes. It is these sidelobes that cause moving-target angle ambiguities.

As in the SAR case, the total ambiguity function has high sidelobes rather than true grating lobes. If the SNR is high enough, it should be possible to place nulls on along-track/radial-velocity ambiguities. However, in the SAR case it was known that




Figure 4-43. The total along-track-position/radial-velocity ambiguity function.

ambiguous energy could only come from a small number of compactly located range-Doppler ambiguities. From Fig. 4-43, we see that ambiguous energy can still come from many angle-velocity combinations that have varying levels of correlation with the desired angle and velocity. It would be impossible to require the array pattern to null all angle-velocity ambiguities. Instead, the following approach is proposed. Since all angle-velocity combinations cannot be rejected, the first step should be to detect those angle-velocity combinations that appear to contain energy. Some of these combinations will be true detections while others will be leakage from other angle-velocity combinations. However, if the detection is due to leakage, the true location should be detected elsewhere in another velocity map.

Next, the space-time responses for all detections are grouped by range and range ambiguities since angle-velocity false alarms due to angle ambiguity can only occur at the same range or range ambiguities. Finally, instead of using the space-time matched filter, the space-time ML filter is used. As in the SAR scenario, this requires

the space-time filter to be orthogonal to all other detections in the same range group. If the detection is a true detection, the output of the ML filter should not change significantly from the matched filter. If the detection is due to leakage from another detection in the same range group, however, rejection of that other detection by the ML filter will cause the filter output to drop significantly compared to the matched filter. Note that this approach can work only because the ambiguities of Fig. 4-43 are not true grating lobes, and, just like the SAR ML filter, performance of this approach will depend strongly on SNR.

The approach just described was applied to the MTI simulation of the previous section. The range, along-track position, and radial velocity of 31 detections were taken from Figs. 4-34, 4-36, 4-38, and 4-40. Then, these detections were re-tested using the ML filter. The ML filter was based on the detection being tested and all other detections at the same or ambiguous range. Therefore, the ML filter is given by

$$\mathbf{W}_{ml} = \mathbf{P}^{-1} \quad (4.64)$$

where

$$\mathbf{P} = [\boldsymbol{\rho}_t \quad \boldsymbol{\rho}_{r.a.d.} \quad \cdots], \quad (4.65)$$

$\boldsymbol{\rho}_t$ is the space-time response of the desired target, and $\boldsymbol{\rho}_{r.a.d.}$ is the response from another range-ambiguous detection.

Figure 4-44 shows the output of the ML filter compared to the matched filter for these 31 detections. The solid line shows the magnitude of the matched-filter output while the dashed line shows the ML-filter output. The dotted line shows the detection threshold that was used. Note that the detection threshold changed slightly for each velocity map. The detections corresponding to true targets are further marked with a circle.




Figure 4-44. Matched-filter output vs. post-detection ML-filter output for 31 detections obtained from Figs. 4-34, 4-36, 4-38, and 4-40.

From Fig. 4-44, it is seen that the ML approach successfully reduced many of the false alarms below the detection threshold. Specifically, detections two, nine, 11, and 28 were judged to have come from the four moving targets showing up in the wrong location. All four of these false detections were reduced below the detection threshold using the ML approach. On the negative side, the output magnitude of the true detections can also decrease. The first two true target detections are relatively unaffected by the ML approach. The third true detection is reduced but had a high enough output magnitude originally such that it was not in danger of falling below the detection threshold. The ML filter reduces the last true target, however, to a level dangerously close to the detection threshold. The reason that the last target's matched-filter magnitude is smaller than the other true targets is because the last




Figure 4-45. Probability of false alarm for the matched filter vs. the post-detection ML filter for Figs. 4-34, 4-36, 4-38, and 4-40.

target suffered a straddling loss by being evenly located between two along-track bins. Nonetheless, a slightly lower RCS could have resulted in the last target being missed by the ML filter.

In order to compare the abilities of the matched and ML filters to detect true targets while rejecting false alarms, P_{fa} and P_d curves were numerically generated. As the detection threshold was varied, the number of false alarms and number of detections were noted. The probability of false alarm was computed as the number of false alarms divided by the total number of pixels tested. The probability of detection was computed as the number of true detections divided by the four moving targets known to be in the data. Figure 4-45 shows the P_{fa} curves for both the matched and ML filters while Fig. 4-46 shows the P_d curves. In Fig. 4-45, it is seen that the P_{fa} is lower for the ML filter than for the matched filter for all threshold levels. This is due




Figure 4-46. Probability of detection for the matched filter vs. the post-detection ML filter for Figs. 4-34, 4-36, 4-38, and 4-40.

to the rejection of multiple detections of the same target. The additional constraints used to form the ML filter, however, can also result in a loss of gain on the desired target. Hence, ML-filter output for true targets may be lower than their corresponding matched-filter output, resulting in lowered P_d . This is seen in Fig. 4-46, which shows that the ML filter misses its first true target at a lower threshold than the matched filter. The ML filter also misses a second target within the threshold range covered by Fig. 4-46 while the matched filter never misses more than one target.

4.4. Summary

In this chapter, I have presented the fundamentals and performance metrics of traditional moving target indication. From those fundamentals, it was determined that, in general, there are two primary requirements that must be met in order to perform MTI. First, there must be a set of measurements capable of distinguishing between moving and stationary targets. Since moving targets always differ from stationary targets in the joint Doppler-angle domain, this requirement is easily met by ensuring that space and time samples are jointly available, and the ability to resolve moving targets from stationary targets will depend on the size of the sensor in both the space and time dimensions.

The second requirement was that there must be measurement degrees of freedom available for moving target detection after rejection of ground clutter. There are many ways of interpreting this requirement. For example, this requirement states that clutter cannot occupy the entire measurement spectrum, or that the number of significant eigenvalues of the clutter covariance matrix must be less than the matrix's dimension, or that the dimension of the vector subspace spanned by the clutter response vectors must be smaller than the number of space-time measurements.

Whether or not this second requirement is met depends on a value called clutter rank. Since a constellation of satellites does not adhere to the strict spacing requirements of traditional DPCA and STAP, current methods of clutter rank prediction do not apply. Therefore, I presented a new, more general method of estimating clutter rank. In one dimension, clutter rank depends on the clutter bandwidth and the length of the observing aperture. In two dimensions, clutter rank depends on the area of the 2D clutter bandwidth and the area of the observing sensor. Since a radar generally collects measurements over the five dimensions of time, frequency, and 3D space, the synthetic aperture of Chapter 2 was used to find an equivalent 2D aperture. Then, several methods of estimating the area of the 2D

synthetic dimensions were presented. These estimation methods were applied to simulations in order to analyze how well they can be used in estimating clutter rank.

Once clutter rank was derived, the performance of MTI filtering could be investigated. First, a filter was derived that projected the measured data orthogonal to all clutter within the specified clutter bandwidth and then detected moving targets using a matched filter. Then, simulations were performed where the illuminated area was relatively small. This kept the clutter bandwidth small, which, in turn, kept clutter rank down and enabled performance evaluation over a wide range of variables including the amount of clutter rejection, the size of the constellation, the number of satellites in the constellation, and error in the knowledge of the satellite positions.

Finally, a wide-area MTI simulation was performed by injecting a SAR simulation with moving target data. Although obscured by clutter in the SAR map, the moving targets were readily detectable after the data were projected orthogonal to clutter. False alarms also appeared above the detection threshold, and these were determined to have come from noise, remaining clutter energy, and moving target ambiguities. The problem of moving target angle ambiguity was presented, and an ML method intended to reject moving target ambiguities was proposed and analyzed.

Just as for the SAR case, a sparse array can perform MTI as long as the array is not too sparse. For SAR, the array could not be too large because nulls in the array pattern had to be wide enough to reject the full width of range-Doppler ambiguities. For MTI, the array cannot be too large because then it will dramatically increase clutter rank, which leaves fewer measurement DOFs available for moving target detection. If the size of the 2D synthetic aperture is predominantly determined by a radar system's bandwidth and coherent integration time, the satellites can be used for distinguishing between moving and stationary targets in the joint angle-Doppler domain. If the size of the physical constellation increases the area of the 2D synthetic aperture, however, it is likely that spatial measurements obtained will be used up by clutter before they are ever applied to moving targets.

On the other hand, the need for wide angular coverage in distinguishing between stationary and slowly moving targets is a significant driver in the constellation concept. Therefore, the need for wide angular coverage conflicts with the requirement that the physical array not be so large that it increases clutter rank. Fortunately, the one variable that is able to relieve the tension between these requirements is the number of satellites in the system. As more satellites are added while maintaining the same overall size of the constellation, more measurement DOFs are added without increasing clutter rank. The results of Section 4.3.2 dramatically demonstrated the tradeoff between constellation size and number of receivers. Any system that uses the constellation approach must balance these conflicting requirements.

For example, consider the sample system used for the SNR analysis of SAR in Section 3.8. In this system, the coherent processing interval was 0.5 s, leading to a single-satellite synthetic aperture of 7820 m. Assuming that in order to maintain low clutter rank, it is desired to increase the along-track dimension of the full, multi-satellite synthetic aperture by less than 15 percent, the along-track diameter of the satellite cluster could be approximately 1.2 km.

Next, the SAR scenario was such that the combined aperture of nine microsattellites was equivalent to the aperture of a single RADARSAT antenna. However, it was seen that approximately 16 apertures were needed in order to reduce the SNR loss due to ML filtering. For the MTI scenario, 16 apertures still may not be enough. For a two-to-one ratio of measurements to clutter rank, the number of apertures within the 1.2 km diameter would need to be between 18 and 20. Of course, the required ratio of measurements to clutter rank and the number of receivers needed for achieving this ratio depend heavily on satellite antenna radiation patterns, target SNR, and other factors.

For a sparse array, one of the most difficult issues that must be solved in order to perform MTI is the angle ambiguity problem. A traditional system with linearly, closely spaced elements typically has one blind spot at zero radial velocity. As

targets increase in speed, the radar's ability to detect them generally increases monotonically. If the transmit PRF is not high enough to avoid Doppler ambiguities, there will be multiple blind zones, but the performance of the system between these blind zones is still well defined. For a sparse array, however, performance is less well behaved. Because of the high sidelobes associated with a large array, targets at multiple angular-position/radial-velocity combinations can appear to come from the same location, or a single target can appear to come from many locations. Therefore, just as for SAR, matched filtering is insufficient. For MTI, matched filtering may produce many false alarms that could overwhelm the tracker. Other filtering methods, such as the post-detection ML method presented in this chapter, must be used to reduce the number of false alarms.

Similarly, high sidelobes of the sparse array can cause detection performance to be reduced for some angle/radial-velocity combinations. For example, if a particular angle/radial-velocity combination results in a particular Doppler frequency, and the physical array has a high sidelobe where ground clutter has the same Doppler, then rejection of the clutter will likely reject much of the moving target energy as well. This behavior was also demonstrated in Section 4.3.2 where output SINR was seen not to improve uniformly as velocity was increased from zero. Some higher velocities were seen to have worse performance than some lower velocities. Again, both this problem and the angle ambiguity problem can be mitigated by reduced the average sidelobe level. The average sidelobe level, in turn, directly depends on the number of receivers in the constellation.

5. CONCLUSIONS AND FUTURE WORK

5.1. Summary

SAR and MTI radar modes, when implemented via spaceborne platforms, both require large apertures or arrays. For traditional SAR, the antenna must be large such that range-Doppler ambiguities are not illuminated. For MTI, high angular resolution is needed to distinguish moving targets from ground clutter since the component of a moving target's velocity in the direction of a satellite will be small. But while spaceborne radars can provide global, 24-hour coverage, placing and operating large apertures in space is difficult. For example, large apertures are heavy and bulky, making them inefficient and expensive to launch. In addition, the limited illumination area provided by large apertures conflicts with the desire for large swathwidths and high search rates.

The satellite cluster concept is intended to mitigate the disadvantages of placing large, heavy satellites into space while providing the advantages of increased swathwidth, decreased revisit times, and improved search rates. The cluster concept does this by adding angle-of-arrival information through multiple, free-flying microsattellites. For SAR, the array formed by the microsats can be used to resolve range-Doppler ambiguities. For MTI, the wide extent of the satellite cluster provides

the high angular resolution necessary for detecting moving targets. Other advantages such as graceful degradation of the system and cost savings through mass production were also discussed in Chapter 1.

These advantages, however, do not come easily. Since an unrealistic amount of fuel would be required to constantly propel the cluster into a closely and uniformly spaced array, the constellation must be allowed to fly freely according to the laws of orbital physics. Factoring in the desire not to have the satellites collide, the array formed by the satellite cluster will be sparse, irregularly sampled, and three-dimensional. Moreover, current SAR and MTI algorithms assume uniform, linear arrays aligned along the radar's direction of travel, and if the cluster concept is to be applied, new radar signal processing algorithms must be developed.

The development of signal processing techniques applicable to the satellite cluster concept is the focus of this work. To that end, I have presented a radar model that facilitates system analysis, developed a synthetic aperture technique that simplifies system characterization, and derived and tested filtering algorithms for both SAR and MTI. The synthetic aperture characterization is a novel approach to characterizing system behavior due to all of the sensor's dimensions-time, frequency, and spatial position. Likewise, the development of sparse-array SAR and MTI techniques is also a new contribution, as is the analysis that brought out the important parameters in determining under what conditions this unique radar concept performs satisfactorily.

In Chapter 2, I presented the radar model that was used throughout this dissertation. It was shown that the radar data could be approximated as a summation of radar responses from each of the illuminated resolution cells on the ground. Then, the summation was written in matrix-vector notation, which later facilitated the derivation of appropriate SAR and MTI algorithms. Using the linear algebra notation, the filters that were derived were commonly applied as the inner product between the data and a weight vector.

The new 2D synthetic aperture was also derived in Chapter 2. This synthetic aperture is completely analogous to the synthetic aperture interpretation that gives

SAR its name. In traditional SAR, a single aperture that moves over a series of radar pulses is equivalent to a long aperture with many elements. In the 2D synthetic aperture, the additional sensor parameters of frequency and multiple spatial dimensions were considered. Instead of synthesizing a long, linear aperture, however, the five sensor parameters were shown to synthesize a planar aperture. Of course, the synthesis is dependent on far-field and narrowband approximations, but this is no different than for traditional SAR.

Using the synthetic aperture, I showed that, since the Earth's surface is two-dimensional, the size of the 2D synthetic aperture could be used to predict the size and orientation of resolution cells. Again, this is analogous to the way along-track resolution is predicted for traditional SAR using the length of the synthetic aperture. Simulations were performed in order to demonstrate the effectiveness of the 2D synthetic aperture in predicting the system ambiguity function.

Finally, the sampled synthetic aperture, or the 2D synthetic array, was discussed in terms of its sampling density. The synthetic co-array was shown to be irregularly sampled, which means that the overall system ambiguity function will not have true grating lobes. Instead, the radar system will have sidelobes with heights determined by the density of the 2D synthetic array. This density, in turn, depends on the overall size and number of satellites in the cluster.

Signal processing for SAR was discussed in Chapter 3. First, the minimum SAR antenna area constraint was presented and demonstrated through simulation. With this constraint, the instantaneous illumination area of a SAR is limited since range-Doppler ambiguities cannot be illuminated. If a SAR that satisfies the constraint, however, is divided into multiple, adjacent antenna elements, the illumination area can be increased. While range-Doppler ambiguities may be illuminated by the smaller size of each antenna element, the angle-of-arrival information provided by adjacent elements can be used to resolve the ambiguities.

This divided-but-contiguous aperture can increase allowable illumination area, but does not resolve the weight-size-power issues associated with putting large

apertures into space. In order to mitigate these issues through the cluster concept, the adjacent apertures must be spread out into a random, sparse array such as would be formed by a satellite constellation. When this occurs, it was shown that the total ambiguity function is unacceptable due to range-Doppler ambiguities lining up with high spatial sidelobes. Consequently, other algorithms that force array pattern nulls on the locations of range-Doppler ambiguities were derived. These algorithms were the maximum-likelihood and minimum-mean-squared-error algorithms. The behavior of these two algorithms as well as matched filtering was analyzed versus SNR, number of satellites, constellation size, and satellite positioning error. Specifically, it was shown that the satellite constellation must not be so large that the spatial nulls placed by the ML and MMSE algorithms are smaller than the size of range-Doppler ambiguities. Finally, an SNR analysis of the SAR operating mode of a sparse satellite cluster was presented. The SNR analysis was compared to a baseline scenario taken from RADARSAT.

Signal processing for MTI was discussed in Chapter 4. After presentation of current DPCA and STAP techniques and performance metrics, it was determined that two requirements must be met in order to perform MTI. First, there must be a set of measurements that are able to distinguish between moving and stationary targets. For both traditional and sparse-array systems, this distinction occurs when filtering is applied in the joint angle-Doppler domain. The second requirement is that there must be measurement degrees of freedom remaining after clutter rejection. In other words, the dimension of clutter must be much less than the dimension of the measurements.

The 2D synthetic aperture was shown to be essential in predicting the dimension of clutter. In the general 2D case, clutter rank depends on the space-bandwidth product. This is the product of the spectral area of clutter and the area of the observing sensor. Since a radar sensor collects information over five parameters: time, frequency, and 3D space, the 2D synthetic aperture was essential for projecting these five sensor parameters into an effective 2D aperture that could be used in the space-bandwidth product. The final analysis showed that the amount of clutter power

rejected, on average, depends on the amount of power in the clutter PSD within a specified bandwidth. The number of dimensions needed for representing that clutter power is the product of the specified bandwidth and the sensor area. Several techniques were presented for approximating the area of non-uniformly weighted and abnormally shaped 2D synthetic sensors.

Next, a filter was derived that projected the measured data orthogonal to all clutter within the bandwidth specified for clutter rank. The remaining data vector was then tested for targets at various locations and radial velocities using a matched filter. This approach demonstrates that clutter can be rejected and moving targets can be detected using sparse arrays. Although the orthogonal projection approach is not optimal, it uses the fact that clutter rank is less than the measurement dimension to save computation. Therefore, not only is the projection approach intuitive, but it is also computationally efficient.

The orthogonal projection filter was first tested using a simulation with a small clutter rank. Results were presented that compared filter performance versus sensor parameters such as the amount of clutter rejected (which is determined by the definition of the clutter bandwidth), constellation size, number of satellites in the constellation, and satellite positioning error. As in the SAR case, the combination of constellation size and the number of receivers was seen to be a key factor in determining performance. Also similar to the SAR result was the conclusion that satellite positioning knowledge to within one-tenth of a wavelength is needed in order to obtain acceptable results.

Finally, a wide-area simulation was performed. For this simulation, the data had to be reduced before the orthogonal projection. This reduction was achieved by forming range bins. The remaining data were then projected orthogonal to the clutter within each range bin and its range ambiguities, and the projected data were tested for moving targets. The SAR data from a previous simulation were injected with four moving targets, and simulations showed that these four targets could be detected. However, other false alarms occurred as well, some of which were the moving targets

being detected in incorrect locations due to angle ambiguity. An approach to resolving these ambiguities was proposed and tested. The approach called for a second pass over all targets detected using the matched filter. The second pass used the ML filter to reject other detections at the same range or range ambiguity. If the target was real, the ML output was similar to the matched-filter output. If the target was due to angle ambiguities, however, the ML output was reduced compared to the matched-filter output.

5.2. Future Work

While some technologies, such as accurate satellite positioning necessary for implementation of this type of SAR may not yet be ready, I think the approach to SAR signal processing presented in this dissertation is ready for a full system analysis. The next step should be to perform higher fidelity simulations that include real waveforms, true satellite orbits, a curved Earth, and so on. The one signal-processing issue that may still need to be addressed is how to control placement of spatial nulls over long integration times and wide bandwidths. Since radiation patterns vary with frequency and will shift as the constellation moves, it is likely that the ML and MMSE filters will need to be implemented multiple times over a set of sub-bands in time and frequency. However, these issues would be best analyzed in a full system simulation.

Improved, high-fidelity simulations should also be performed for the MTI mode. The clutter rank characterization presented in this paper could then be tested against simulations with accurate radiation patterns and angle-dependent clutter models. Also, performance parameters such as MDV, P_d , and P_{fa} could be taken as true indications of system performance rather than just measures of system behavior. While the analysis in this paper represents a proof of concept and provides bounds on system design, MDV and other performance descriptions need to be obtained from full-fidelity simulations.

More detection filters also need to be investigated. Specifically, there will be a strong need for adaptive detection techniques. This means secondary data that can be used to estimate interference statistics must be identified and analyzed. Of course, even when the source of secondary data is identified, prior STAP research will need to be leveraged in order to deal with non-stationary and non-homogeneous clutter.

Last, the angle ambiguity problem should be studied. The post-detection ML approach provided in this dissertation shows some promise, but significant research still needs to be performed in this area. Other approaches should be identified and

tested. For example, logical methods of deciding when the same detection shows up multiple times within the same or several MTI maps need to be developed. If multiple detections are attributed to the same target, the tracker should be able to use multiple looks to determine which, if any, of the detections behave in agreement with their detected characteristics of range, angle, and speed.

5.3. Conclusions

The recurring theme throughout this dissertation has been the interplay between the overall size of the satellite cluster and the number of satellites in it. Several results have pointed to the relationship between these two factors as being key in determining performance of a cluster-concept system.

In the theory of random arrays, the expected value of array pattern sidelobes depends on the number of elements in the array. As more elements are added, the expected value of the sidelobes decreases. Consequently, the ideal situation would be a fully filled array of the same size as a satellite constellation, but this possibility has been presented as impractical. The main question, therefore, is how sparse can a constellation be and still provide good results?

For SAR, two factors affect the allowable sparseness of a satellite cluster. First, the constellation cannot be so large that the width of nulls in its radiation pattern are smaller than the width of range-Doppler ambiguities. As shown in Chapter 3, if the constellation is too large, its radiation nulls will only be able to reject a fraction of range-Doppler ambiguities. The second factor is SNR. The nulling of range-Doppler ambiguities through the ML or MMSE algorithms causes a loss in SNR for each pixel. Therefore, there must be enough available signal power to overcome this processing loss, and the SNR requirement depends heavily on the number of receivers in the constellation. Not only do additional apertures collect more signal power, but they also reduce processing loss by increasing the degrees of freedom available for pattern synthesis.

Another way of looking at the constellation-size/number-of-receivers tradeoff is in terms of pixel resolution. If the constellation is large compared to the synthetic aperture formed by the transmit signal, resolution improves due to the small beamwidth of the physical array. In this case, the spatial measurements obtained through the satellite constellation are used for additional resolution rather than for ambiguity resolution. Instead, the size of the constellation should be limited such that

bandwidth and the length of the coherent processing interval are the primary factors in determining resolution. This leaves the spatial measurements for ambiguity resolution, or, in other words, this ensures that the number of measurements collected is significantly more than the number of resolution pixels illuminated.

For MTI, the desire to have a large constellation is even stronger since high angular resolution is necessary for resolving slowly moving targets from ground clutter. However, the physical constellation must not be so large that clutter rank approaches the number of measurements collected. If a minimum-detectable-velocity (MDV) requirement forces a large constellation, there must be enough receivers in the constellation to make sure that measurements far exceed clutter rank.

Additional satellites again provide a twofold benefit. First, additional apertures increase SNR through additional energy-collecting area. Second, additional apertures reduce processing loss through improved conditioning of the problem. Average sidelobe levels are decreased, the number of measurements is increased, and sensitivity to positioning errors is reduced. Furthermore, not only do reduced sidelobes improve target detectability, but they also reduce angle ambiguity since targets can only show up in the wrong location if they leak through a spatial sidelobe.

REFERENCES

- [1] *Proc. IEEE 2001 Radar Conference*, Atlanta, Ga., May 1-3, 2001.
- [2] W. B. Scott, "Testbeds wring out technologies," *Aviation Week and Space Technology*, pp. 52-53, April 5, 1999.
- [3] D. A. Whelan, "DISCOVERER II program summary," *Proc. IEEE 2000 Int. Radar Conf.*, Washington DC, pp. 7-8.
- [4] J. Entzminger, C. Fowler, and W. Kenneally, "JointSTARS and GMTI: past, present and future," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 35, no. 2, pp. 748-761, April, 1999.
- [5] M. E. Davis, "Technology challenges in affordable space based radar," *Proc. IEEE 2000 Int. Radar Conf.*, Washington DC, pp. 18-23.
- [6] K. Overman, K. Leahy, and R. J. Fritsch, "The future of surface surveillance – revolutionizing the view of the battlefield," *Proc. IEEE 2000 Int. Radar Conf.*, Washington DC, pp. 1-6.
- [7] T. J. Nohara, P. Weber, and A. Premji, "Space-based radar signal processing baselines for air, land and sea applications," *Electronics & Communication Engineering Journal*, vol. 12, no. 5, pp. 229-239, Oct., 2000.
- [8] F. T. Ulaby, R. K. Moore, and A. K. Fung, *Microwave Remote Sensing: Active and Passive*. Norwood, MA: Artech House, 1982, vol. 1, ch. 5.
- [9] R. E. Collin, *Antennas and Radiowave Propagation*. New York: McGraw-Hill, 1985, ch. 6.
- [10] V. A. Chobotov, "Orbital considerations for space-based radar," chap. 2 in *Space-Based Radar Handbook*, L. J. Cantafio, Ed. Norwood, MA: Artech House, 1989.
- [11] C. Elachi, *Spaceborne Radar Remote Sensing: Applications and Techniques*. New York: IEEE Press, 1988.

- [12] F. K. Li and W. T. K. Johnson, "Ambiguities in spaceborne synthetic aperture radar systems," *IEEE Trans. Aerosp. Electron. Syst.*, vol. AES-19, no. 3, pp. 389-396, May, 1983.
- [13] S. W. McCandless, "SAR in space - the theory, design, engineering and application of a space-based SAR system," chap. 4 in *Space-Based Radar Handbook*, L. J. Cantafio, Ed. Norwood, MA: Artech House, 1989.
- [14] K. Tomiyasu, "Image processing of synthetic aperture radar range ambiguous signals," *IEEE Trans. Geosci. Remote Sensing*, vol. 32, no. 5, pp. 1114-1117, Sept., 1994.
- [15] K. Tomiyasu, "Performance of a proposed spaceborne synthetic aperture radar with variable antenna height," *IEEE Trans. Geosci. Remote Sensing*, vol. 28, no. 4, pp. 609-613, July, 1990.
- [16] K. Tomiyasu, "Conceptual performance of a satellite borne, wide swath synthetic aperture radar," *IEEE Trans. Geosci. Remote Sensing*, vol. GE-19, no. 2, pp. 108-116, April, 1981.
- [17] F. T. Ulaby, R. K. Moore, and A. K. Fung, *Microwave Remote Sensing: Active and Passive*. Norwood, MA: Artech House, 1982, vol. 2, ch. 9.
- [18] A. K. Hyder, R. L. Wiley, G. Halpert, D. J. Flood, and S. Sabripour, *Spacecraft Power Technologies*. London: Imperial College Press, 2000.
- [19] R. K. Moore, "Scanning spaceborne synthetic aperture radar with integrated radiometer," *IEEE Trans. Aerosp. Electron. Syst.*, vol. AES-17, no. 3, pp. 410-421, May, 1981.
- [20] D. L. Evans, Ed. *Spaceborne Synthetic Aperture Radar: Current Status and Future Directions*. Washington DC: NASA, 1995.
- [21] A. Moreira, J. Mittermayer, and R. Scheiber, "Extended chirp scaling algorithm for air- and spaceborne SAR data processing in stripmap and ScanSAR imaging modes," *IEEE Trans. Geosci. Remote Sensing*, vol. 34, no. 5, pp. 1123-1136, Sept., 1996.

- [22] A. Currie and M.A. Brown, "Wide-swath SAR," *IEE Proceedings, Part F: Radar and Signal Processing*, vol. 139, no. 2, pp. 122-135, April 1992.
- [23] P. Zetocha et al., "Commanding and controlling satellite clusters," *IEEE Intelligent Systems*, vol. 15, no. 6, pp. 8-13, Nov., 2000.
- [24] C. Kitts et al., "Emerald: a low-cost spacecraft mission for validating formation flying technologies," *Proc. 1999 IEEE Aerospace Conf.*, Aspen, CO, pp. 217-226.
- [25] D. Massonnet et al., "A wheel of passive radar microsats for upgrading existing SAR projects," *Proc. IEEE 2000 Int. Geosci. Rem. Sens. Symp.*, Honolulu, HI, pp. 1000-1003.
- [26] D. Massonnet, "Capabilities and limitations of the interferometric cartwheel," *IEEE Trans. Geosci. Rem. Sens.*, vol. 39, no. 3, pp. 506-520, March, 2001.
- [27] Air Force Research Lab. TechSat 21 Next Generation Space Capabilities. Available: www.vs.afrl.af.mil/TechProgs/TechSat21/NGSC.html.
- [28] M. Martin and M. Stallard, "Distributed satellite missions and technologies – the TechSat 21 program," *Proc. of the AIAA Space Technology Conference and Exposition*, Albuquerque, NM, AIAA 99-4479, September, 1999.
- [29] A. Das, R. Cobb, and M. Stallard, "TechSat 21: a revolutionary concept in distributed space based sensing," *Proc. of the AIAA Defense and Civil Space Programs Conference and Exhibit*, Huntsville, AL, AIAA 98-5255, October, 1998.
- [30] M. Younis and W. Wiesbeck, "SAR with digital beamforming on receive only," *1999 IEEE Geoscience and Remote Sensing Symposium Proceedings*, Hamburg, vol. 3, pp. 1773-1775.
- [31] W. L. Melvin, "Eigenbased modeling of nonhomogeneous airborne radar environments," *Proc. of the 1998 IEEE National Radar Conference*, Dallas, TX, pp. 171-176.

- [32] J. J. Goldstein, I. S. Reed, P. A. Zulch, and W. L. Melvin, "A multistage STAP CFAR detection technique," *Proc. of the 1998 IEEE National Radar Conference*, Dallas, TX, pp. 171-176.
- [33] C. Covault, "Joint-Stars patrols Bosnia," *Aviation Week and Space Technology*, pp. 44-49, February 19, 1996.
- [34] J. Ward, "Space-time adaptive processing for airborne radar," MIT Lincoln Laboratory, Lexington, MA, Tech. Rep. 1015, Dec. 1994.
- [35] R. Klemm, *Space-Time Adaptive Processing*. London, UK: Institute of Electrical Engineers, 1998.
- [36] R. Klemm, "Adaptive airborne MTI: an auxiliary channel approach," *IEE Proceedings: Part F*, vol. 134, no. 3, pp. 269-276, June, 1987.
- [37] R. Klemm, "Adaptive clutter suppression for airborne phase array radars," *IEE Proceedings: Part F*, vol. 130, no. 1, pp. 125-132, Feb., 1983.
- [38] A. R. Stubberud, *Analysis and Synthesis of Linear Time-Variable Systems*. Los Angeles: University of California Press, 1964.
- [39] H. D'Angelo, *Linear Time-Varying Systems: Analysis and Synthesis*. Boston: Allyn and Bacon, 1970.
- [40] A. Satish and R. L. Kashyap, "Maximum likelihood estimation and Cramer-Rao bounds for direction of arrival parameters of a large sensor array," *IEEE Trans. Antennas Propagat.*, vol. 44, no. 4, pp. 478-491, April, 1996.
- [41] S. T. Smith, "Accuracy and resolution bounds for adaptive sensor array processing," *Proc. Ninth IEEE Signal Processing Workshop on Statistical Signal and Array Processing*, Portland, OR, Sept. 14-16, 1998.
- [42] D. H. Johnson and D. E. Dudgeon, *Array Signal Processing: Concepts and Techniques*. Englewood Cliffs, NJ: Prentice-Hall, 1993.
- [43] H. L. Van Trees, *Detection, Estimation, and Modulation Theory, Part I*. New York: John Wiley and Sons, 1968.
- [44] V. Mrstik, "Agile-beam synthetic aperture radar opportunities," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 34, no. 2, pp. 500-507, April, 1998.

- [45] R. K. Moore, "Trade-off between picture element dimensions and noncoherent averaging in side-looking airborne radar," *IEEE Trans. Aerosp. Electron. Syst.*, vol. AES-15, no. 5, pp. 697-708, Sept. 1979.
- [46] R. L. Fante, "Adaptive nulling of SAR sidelobe discretets," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 35, no. 4, pp. 1212-1218, Oct., 1999.
- [47] A. Moreira, "Real-time synthetic aperture radar (SAR) processing with a new subaperture approach," *IEEE Trans. Geosci. Remote Sensing*, vol. 30, no. 4, pp. 714-722, July, 1992.
- [48] J. Mittermayer, A. Moreira, and O. Loffeld, "Spotlight SAR data processing using the frequency scaling algorithm," *IEEE Trans. Geosci. Remote Sensing*, vol. 37, no. 5, pp. 2198-2214, Sept., 1999.
- [49] M. A. Richards, "Synthetic aperture processing," chap. 10 in *Airborne Pulsed Doppler Radar*, G. Morris and L. Harkness, Eds. Boston, MA: Artech House, 1996.
- [50] J.M. Stiles and N. Goodman, "Processing of multi-aperture SAR to produce fine-resolution images of arbitrarily large extent," *Proc. of the 2001 IEEE Radar Conference*, Atlanta, GA, pp. 451-456.
- [51] N.A. Goodman, S.C. Lin, D. Rajakrishna, and J.M. Stiles, "Processing of multiple-receiver spaceborne arrays for wide-area SAR," *IEEE Trans. Geosci. Remote Sensing*, vol. 40, no. 4, pp. 841-852, April, 2002.
- [52] N. Levanon, *Radar Principles*. New York: John Wiley & Sons, 1988.
- [53] Alaska SAR Facility. RADARSAT's Beams – Signal Parameters. Available: www.asf.alaska.edu/radarsat/radarsat_beam_signal_parameters.html.
- [54] Canada Centre for Remote Sensing. RADARSAT Technical Specs – Summary. Available: www.ccrs.nrcan.gc.ca/ccrs/tekrd/radarsat/specs/radspece.html.
- [55] Canada Centre for Remote Sensing. RADARSAT Technical Specs – Details. Available: www.ccrs.nrcan.gc.ca/ccrs/tekrd/radarsat/specs/radtable.html.
- [56] Canadian Space Agency. RADARSAT-1. Available: www.space.gc.ca/csa_sectors/earth_environment/radarsat/default.asp.

- [57] R.A. Coury, M. Gruber, A.E. Filip, and S.M. Kogon, "Discoverer II performance analysis tool (DPAT) modeling report," Lexington, MA, MIT-Lincoln Laboratory, July, 1999.
- [58] O.L. Frost, "An algorithm for linearly constrained adaptive array processing," *Proc. of the IEEE*, vol. 60, pp. 926-935, Aug., 1972.
- [59] S.P. Applebaum, "Adaptive arrays," *IEEE Trans. Antennas Propagat.*, vol. AP-24, no. 5, pp. 585-598, Sept., 1976.
- [60] W.F. Gabriel, "Adaptive arrays – an introduction," *Proc. of the IEEE*, vol. 64, no. 2, pp. 239-272, Feb., 1976.
- [61] B. Widrow, P.E. Mantey, L.J. Griffiths, and B.B. Goode, "Adaptive antenna systems," *Proc. of the IEEE*, vol. 55, no. 12, pp. 2143-2159, Dec., 1967.
- [62] J.W.R. Griffiths, "Adaptive array processing – A tutorial," *IEE Proc.: Parts F & H*, vol. 130, no. 1, pp. 3-10, Feb., 1983.
- [63] I.S. Reed, J.D. Mallett, and L.E. Brennan, "Rapid convergence rate in adaptive arrays," *IEEE Trans. Aerosp. Electron. Syst.*, vol. AES-10, pp. 853-863, Nov., 1974.
- [64] L.E. Brennan and I.S. Reed, "Theory of adaptive radar," *IEEE Trans. Aerosp. Electron. Syst.*, vol. AES-9, no. 2, pp. 237-252, March, 1973.
- [65] M.C. Wicks, W.L. Melvin, and P. Chen, "An efficient architecture for nonhomogeneity detection in space-time adaptive processing airborne early warning radar," *Proc. 1997 Radar Edinburgh Int. Conf.*, Edinburgh, UK, Oct. 14-16, pp. 295-299, 1997.
- [66] P. Antonik, H.K. Schuman, W.L. Melvin, and M.C. Wicks, "Implementation of knowledge-based control for space-time adaptive processing," *Proc. 1997 Radar Edinburgh Int. Conf.*, Edinburgh, UK, Oct. 14-16, pp. 478-482.
- [67] D.J. Rabideau and A.O. Steinhardt, "Improving the performance of adaptive arrays in nonstationary environments through data-adaptive training," *Proc. 30th Asilomar Conf.*, Pacific Grove, CA, pp. 75-79, Nov. 3-6, 1996.

- [68] A. Papoulis, "Limits on bandlimited signals," *Proc. of the IEEE*, vol. 55, no. 10, pp. 1677-1686, Oct., 1967.
- [69] B.R. Frieden, "Evaluation, design and extrapolation methods for optical signals, based on use of the prolate functions," chap. VIII in *Progress in Optics, Vol. IX*, E. Wolf, Ed. Amsterdam: North-Holland Publishing Company, 1971.
- [70] D. Gabor, "Light and information," chap. IV in *Progress in Optics, Vol. I*, E. Wolf, Ed. Amsterdam: North-Holland Publishing Company, 1961.
- [71] J.W. Goodman, *Introduction to Fourier Optics*. New York: McGraw-Hill, 1996.
- [72] A. Papoulis, *Systems and Transforms with Applications in Optics*. New York: McGraw-Hill, 1968.
- [73] F. Gori and G. Guattari, "Degrees of freedom of images from point-like-element pupils," *J. Opt. Soc. America*, vol. 64, no. 4, pp. 453-458, Apr., 1974.
- [74] G. Toraldo Di Francia, "Degrees of freedom of an image," *J. Opt. Soc. America*, vol. 59, no. 7, pp. 799-804, July, 1969.
- [75] R. Klemm, "Introduction to space-time adaptive processing," *Electronics & Comm. Eng. Jour.*, vol. 11, no. 1, pp. 5-12, Feb., 1999.
- [76] K.M. Buckley, "Spatial/spectral filtering with linearly constrained minimum variance beamformers," *IEEE Trans. Acoustics, Speech, and Sig. Proc.*, vol. ASSP-35, no. 3, pp. 249-266, March, 1987.
- [77] D. Slepian and H.O. Pollak, "Prolate spheroidal wave functions, Fourier analysis and uncertainty – I," *Bell Syst. Tech. J.*, vol. 40, pp. 43-64, Jan., 1961.
- [78] H.J. Landau and H.O. Pollak, "Prolate spheroidal wave functions, Fourier analysis and uncertainty – II," *Bell Syst. Tech. J.*, vol. 40, pp. 65-84, Jan., 1961.
- [79] H.J. Landau and H.O. Pollak, "Prolate spheroidal wave functions, Fourier analysis and uncertainty – III: The dimension of the space of essentially time- and band-limited signals," *Bell Syst. Tech. J.*, vol. 41, pp. 1295-1336, July., 1962.

- [80] A. Abramowitz and I.A. Stegun, Eds. *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*. New York: Wiley, 1972.