NAVAL POSTGRADUATE SCHOOL
MONTEREY, CALIFORNIA

THESIS

NPS HIGH RESOLUTION

SYNTHETIC APERTURE

SONAR

by

Joseph Donald Welter

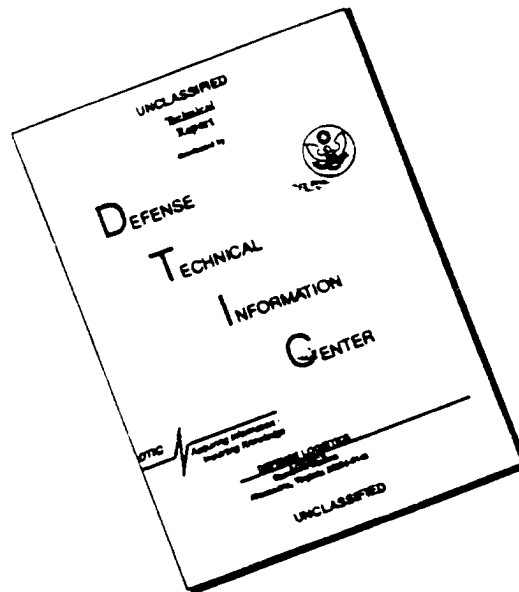December, 1995

Thesis Advisors: Donald L.Walters

Roberto Cristi

Approved for public release; distribution is unlimited.

19960411 112   DTIC QUALITY INSPECTED 1

DISCLAIMER NOTICE

REPORT DOCUMENTATION PAGE

Form Approved OMB No. 0704-0188

1. AGENCY USE ONLY *(Leave blank)*	2. REPORT DATE December 1995.	3. REPORT TYPE AND DATES COVERED Master's Thesis

4. TITLE AND SUBTITLE NPS HIGH RESOLUTION SYNTHETIC APERTURE SONAR	5. FUNDING NUMBERS
6. AUTHOR(S) Joseph Donald Welter	

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Postgraduate School Monterey CA 93943-5000	8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSORING/MONITORING AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government.

12a. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.	12b. DISTRIBUTION CODE

13. ABSTRACT *(maximum 200 words)*

This thesis investigated the use of synthetic aperture techniques to achieve a long effective aperture, high resolution, imaging sonar. The approach included a full simulation of the system using the MATLAB programming environment that provided a model for developing six data processing algorithms and a working 25KHz, 1 m baseline, air medium synthetic aperture sonar. The six azimuthal processing techniques included: 1) a normal, real aperture, 2) an unfocussed synthetic aperture, 3) a hybrid focussed-unfocussed system, 4) a fully focussed one line algorithm, 5) a limited two-dimensional, fully focussed algorithm and, 6) a limited two-dimensional, hybrid focussed-unfocussed algorithm. This thesis compared the run times, resolutions and signal to noise ratios achieved by the six techniques both in simulation and experimental measurements collected with the actual prototype.

14. SUBJECT TERMS Synthetic Aperture Sonar, Coherent Processing, High Resolution Sonar.	15. NUMBER OF PAGES 114
	16. PRICE CODE

17. SECURITY CLASSIFI-CATION OF REPORT Unclassified	18. SECURITY CLASSIFI-CATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICA-TION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UL

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. 239-18 298-102

NPS HIGH RESOLUTION SYNTHETIC APERTURE SONAR

Joseph Donald Welter
Lieutenant Commander, United States Navy
B.S.E.E. United States Naval Academy, 1982

Submitted in partial fulfillment of the
requirements for the degrees of

MASTER OF SCIENCE IN APPLIED PHYSICS
and
MASTER OF SCIENCE IN ELECTRICAL ENGINEERING

from the

NAVAL POSTGRADUATE SCHOOL
December 1995

Author: _____

Joseph Donald Welter

Approved by: _____

Donald L. Walters, Thesis Advisor

_____

Roberto Cristi, Thesis Advisor

_____

William B. Colson, Chairman, Department of Physics

_____

Herschel H. Loomis, Jr., Chairman
Department of Electrical and Computer Engineering

ABSTRACT

This thesis investigated the use of synthetic aperture techniques to achieve a long effective aperture, high resolution, imaging sonar. The approach included a full simulation of the system using the MATLAB programming environment that provided a model for developing six data processing algorithms and a working 25KHz, 1 m baseline, air medium synthetic aperture sonar. The six azimuthal processing techniques included: 1) a normal, real aperture, 2) an unfocussed synthetic aperture, 3) a hybrid focussed-unfocussed system, 4) a fully focussed one line algorithm, 5) a limited two-dimensional, fully focussed algorithm and, 6) a limited two-dimensional, hybrid focussed-unfocussed algorithm. This thesis compared the run times, resolutions and signal to noise ratios achieved by the six techniques both in simulation and experimental measurements collected with the actual prototype.

TABLE OF CONTENTS

ACKNOWLEDGMENT

I. INTRODUCTION

The proud and buried mine threat is a significant obstacle to the execution of the US Navy's mission in the littoral. One fundamental problem in this arena is the location and evaluation of mine-like objects. Conventional classification and identification methods have displayed marginal performance in this area. Synthetic aperture sonar techniques offer high azimuthal resolution (as fine as one half the size of the transducer) which does not degrade with range. Thus, these methods offer some relief in this difficult problem area.

This thesis describes azimuthal resolution techniques by comparison and contrast of the theory and results achieved by:

1. Normal real aperture processing,
2. Unfocussed synthetic aperture processing,
3. A focussed / unfocussed hybrid technique,
4. A fully focussed synthetic aperture method,
5. A limited two-dimensional fully focussed algorithm, and
6. A limited two-dimensional focussed / unfocussed hybrid algorithm.

These six azimuth resolution methods will be compared both through simulation and real data obtained from a 25KHz, 1m baseline, air medium synthetic aperture sonar built in a laboratory for this purpose.

These topics will be covered in the following sequence. Chapter II lays down the basic theoretical basis of synthetic aperture processing within the context of general detection schemes. Chapter III develops and contrasts, through simulation, these six azimuthal resolution techniques. Chapter IV describes the construction of an atmospheric test platform for the algorithms and an outline of the computer code used. (These codes were written in MATLAB and run on a personal computer for flexibility and portability.) In Chapter V, the results of that in-air test are presented and analyzed. Chapter VI details the conclusions reached and recommended directions for future research in this area at the Naval Postgraduate School including the transition to a water-borne environment. This thesis initiates synthetic aperture sonar developments at the Naval Postgraduate School.

II. BASIC DETECTION THEORY

In order to appreciate why synthetic aperture techniques are necessary, the limitations of conventional systems which led to synthetic aperture methods need to be understood. Edde (1993) and Curlander (1991) provide excellent summaries of conventional and synthetic aperture concepts applied to radar. This chapter will outline these same techniques in the context of a sonar application.

To detect, classify, and identify a mine-like object, the range and azimuth resolution of a sonar has to be fine enough to reveal unique target characteristics. A conventional sonar system is limited in angular resolution by the wavelength of the sound wave, λ, and the size of the antenna, D, by the relation $\theta=\lambda/D$. The linear azimuthal resolution of the target would be the range multiplied by this angle or $r\times\theta=r\times\lambda/D$. Improving the resolution requires either a smaller wavelength or a larger antenna aperture. In either case, resolution degrades with range.

Synthetic aperture sonar systems improve the resolution by synthesizing a large aperture, D, by moving a single smaller aperture of size d along a path and adding coherently the successive signal returns after accounting for the changing phase. The ultimate resolution turns out to be d/2, where d is the actual small antenna size, which is independent of range or wavelength. These concepts will be explored in more depth in Chapter Three.

A. INITIAL ASSUMPTIONS

To properly compare the relative benefits offered by the different methods, we first need to define the basic problem. Figure 2.1 illustrates the fundamental initial conditions of the detection problem which are:
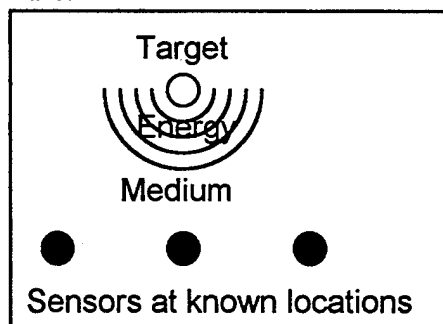
Figure 2.1 Basic Detection Problem

1. We have a series of omnidirectional sensor elements positioned at known fixed locations.

2. We have a target at some unknown fixed location.

3. There is some form of energy which passes from the target to the sensor elements.

4. There is a medium of known isotropic properties between the target and the sensors.

5. The desired knowledge is the location and an identifiable image of the target.

B. PASSIVE METHODS

The simplest method of target localization is to use the target's radiated energy passively as illustrated above. Energy from the target is sensed by the elements, relative times of arrival (phase) are recorded, and, with at least three elements, we can determine:

1. Whether or not the target is in the near field or the far field based on whether, within the sensitivity of the equipment, the time delays are all linear, hence the result of a far field plane wave, or nonlinear, hence the result of a near field spherical wave.

2. If the target is in the near field, we can calculate the range and bearing to the target by using the curvature of the received wave front.

3. If the target is in the far field, we can calculate the bearing to the target based on the time delays and a minimum range (only) to the target based on the range at which our equipment can detect the nonlinearity of a wave from this bearing.

In summary, using passive methods, we can always determine the target's bearing relative to the sensor array if it is cooperative enough to radiate sufficient energy that we can sense it with our elements. If it is close, we can gain range information also.

C. ACTIVE METHODS

1. Single Transmitting Element

To remove the restriction that the target be radiating energy, we can convert one of our sensor elements to a transmit / receive element, provide the illuminating energy ourselves to be reflected off the target, sense that returned energy and determine the target bearing in the same manner as above as illustrated in Figure 2.2.

4
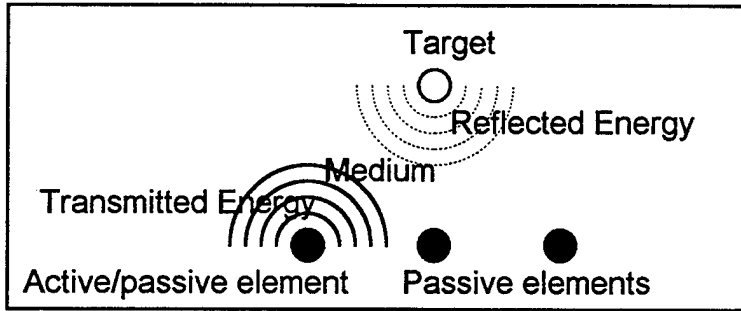
Figure 2.2 Single Active Element

Additionally, we can use our knowledge of the medium to determine the target's range to the active/passive element by recording the time difference between the transmission of the outgoing energy and the reception of the reflected energy at the element. We then multiply this time difference by the energy velocity in the medium divided by two (to account for the energy making both an outgoing and return trip in this time).i.e.

$$R=(c \times \Delta t)/2. \qquad (2.1)$$

Using all the sensors, this gives us a range and bearing to the target regardless of its cooperation but restricts the detection ranges because of the spherical spreading losses experienced by the energy both going to and coming back from the target and the limited energy we can output through one element.

2. Multiple Transmitting Elements

The easiest solution to this problem of limited energy to the target is to let all of the sensor elements transmit thereby multiplying by the number of elements the amount of energy scattered from the target. A problem that arises if all the elements transmit independently is that no element can tell if the energy it received came from its own pulse or that of another element; so, the range information can be destroyed. Additionally, if there is more than one target, no element knows from which target the energy it received returned, so bearing information is also lost. This chaos, as depicted in Figure 2.3, is worse than the passive condition.

5
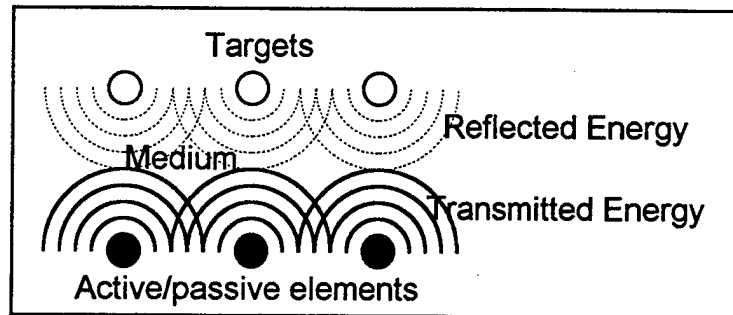
Figure 2.3 Multiple Independent Active Elements

To extract some order from this chaos, we normally resort to active beam forming, where we time-coordinate the activations of the transmitting elements such that their transmissions interfere constructively to produce a narrow "beam" traveling in one desired direction and effectively no energy in any other direction as shown in Figure 2.4. (I am intentionally ignoring side lobes in this discussion.)
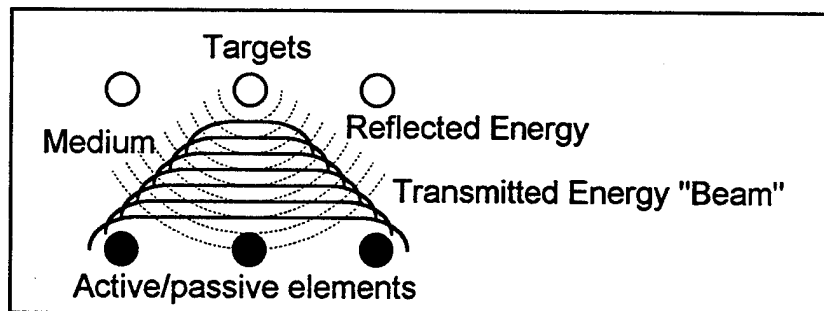
Figure 2.4 Multiple Coordinated Active Elements

In this way, we can use all of the transmitters to illuminate a single target or target direction, and then use the receivers to localize the target in range. We see less energy dissipation with increasing range because the beam does not spread as quickly as a single point, spherical wave. Energy only goes in one direction, so everything detected is on that bearing. Range data is again available because the target is now being excited by only one incident wave- the plane wave. To detect targets in other directions or steer the beam, we either physically rotate the elements as a group (as is done with a normal antenna) or adjust the timing of the excitation of the elements (as is done with a phased array antenna).

This method functions well and is adequate in most situations. As a result, it is the most popular method of target detection in cases where active transmission is appropriate. It does, however, have limitations in that the target area of uncertainty grows with range for a constant beam width as Figure 2.5 demonstrates.
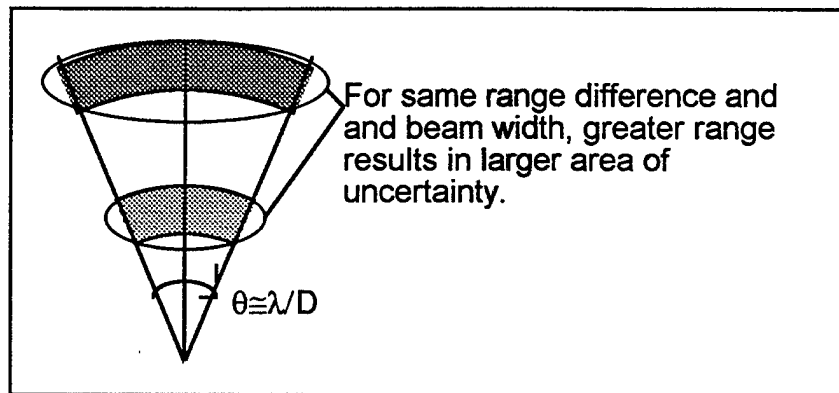
Figure 2.5 Real Beam Divergence

To reduce this uncertainty by reducing the beam width, requires a larger physical array. Additionally, if the beam is electronically steered, endfire beams are significantly broader than broadside beams because the array looks narrower to the target from an end aspect as Figure 2.6 shows.
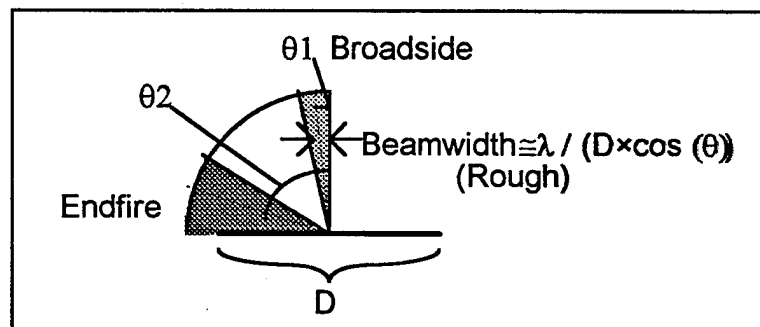
Figure 2.6 Phased Array Endfire Beam Divergence

Thus, to achieve small areas of uncertainty at long ranges, we would have to build very large arrays with elements placed in their ideal locations to within a fraction of a wavelength of the energy used (to achieve the proper phase relationships). For a long distance target location, this requirement can quickly become impractical to realize.

3. Synthetic Aperture

One innovative solution to this dilemma is to build the sensor element array over time, as illustrated in Figure 2.7, instead of requiring all of the elements to be in place for the duration of the search.
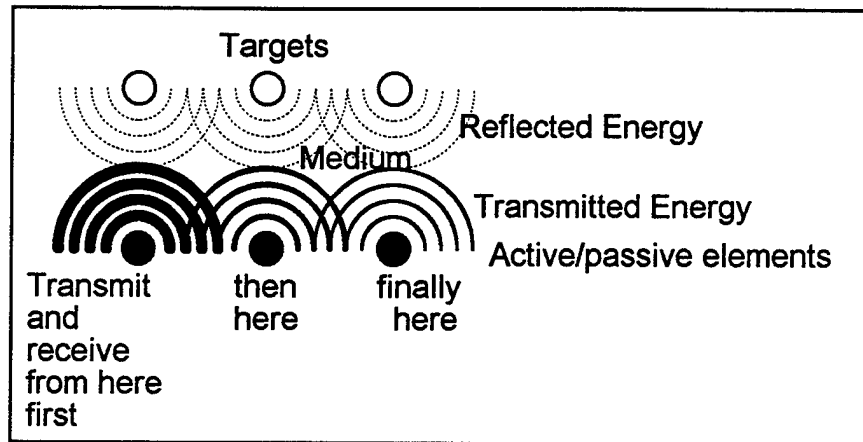
Figure 2.7 Synthetic Aperture Coordinated Active Elements

In this case, the sensor array size is limited only by the distance the platform can transport the transmit/receive element. If element positions are not ideal, their deviations from those calculated must be corrected in some manner. Now that we are no longer dealing with a highly directional plane wave, the return at any element is the result of reflections from all targets, not just those within a narrow beam. All targets are sensed on every transmission. Targets can be distinguished from one another from the viewpoint of a single element only by the time delays of their returns which are proportional to those targets' ranges from the element in question. No azimuth information is provided by the signal time history of a single element. Figure 2.8 illustrates the situation.
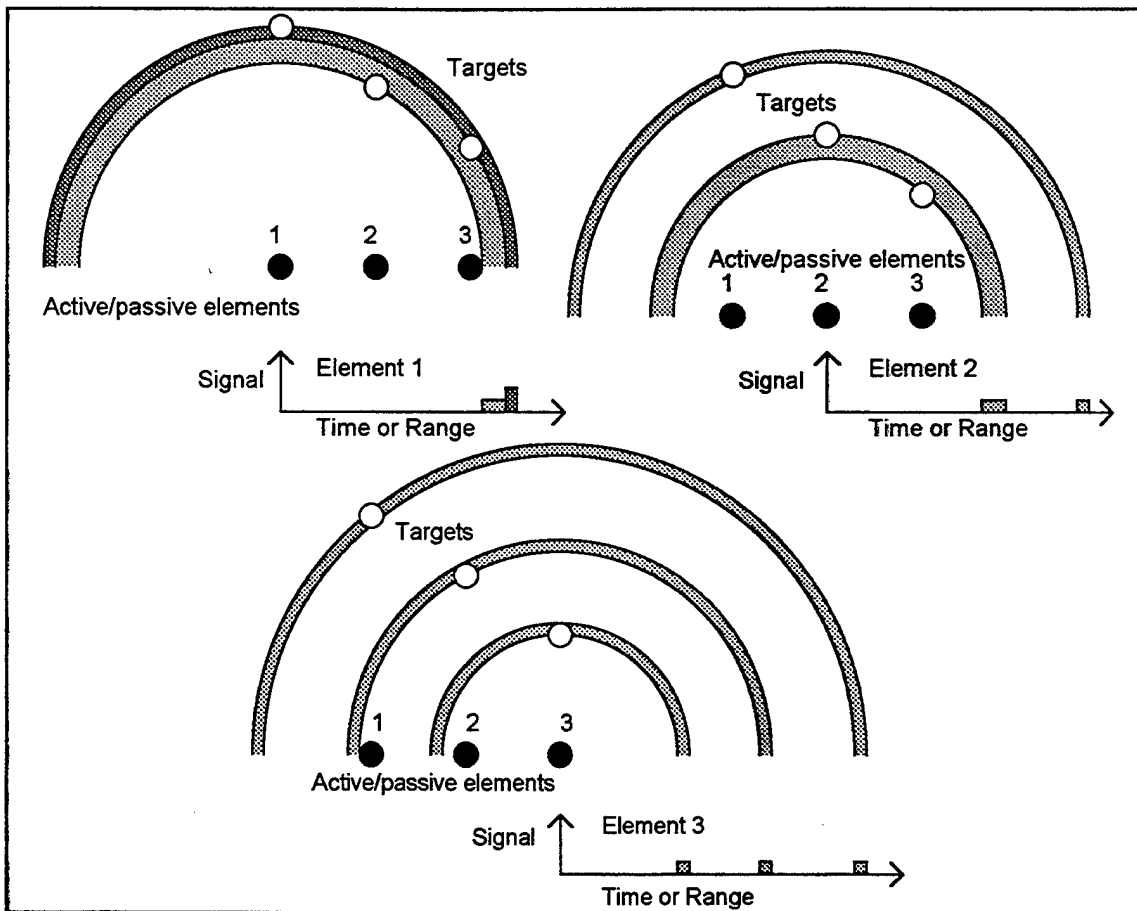
Figure 2.8 Synthetic Aperture Returns

Hence, while a normal sonar uses a large number of transmitters and receivers to locate one target at a time, a synthetic system uses a single transmitter-receiver pair to record returns from all targets simultaneously from one view point. The transmitter-receiver pair is then moved to a new location (ideally along a straight line) to view the targets from there. By coherently summing these returns (adding amplitudes, not intensities or powers) from the various elements and accounting for position offsets, we are able to reproduce the target locations, gaining azimuth information which was not present separately in the individual returns. Figure 2.9 depicts a simplified version of this concept.

9
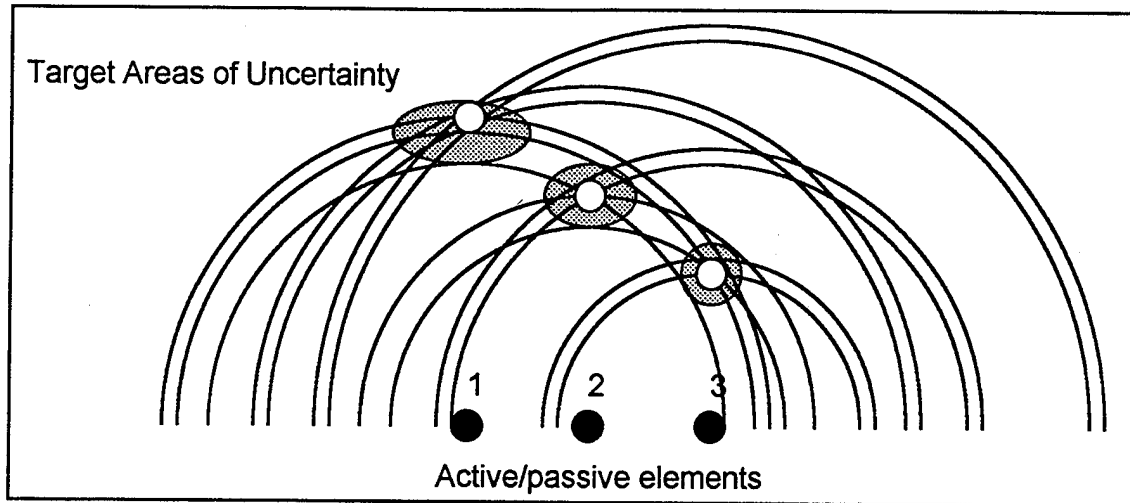
Figure 2.9 Target Localization by Range Rings

Utilizing enough elements and traveling over a wide enough "aperture" should give us any desired effective array size, hence, any desired azimuth resolution within the ranges at which we get returns from the targets.

Thus, the desire to gain a high resolution sonar without the need for a large antenna is the principle motivating force for the investigation of synthetic aperture techniques. The implementation of these techniques, as covered in the next chapter, reveals the limitations of this technique and some surprising capabilities.

10

III. IMPLEMENTATION

The goal of a synthetic aperture implementation is to take the data available at the sensors and to add the individual returns coherently to create an image of a target. To generate a satisfactory image of the target means that the target has been localized both in range and azimuth to the required resolution. Range localization is almost universally achieved in modern active radar and sonar systems (and in this paper) by means of a correlation or matched filter technique. This thesis concentrates on the azimuthal resolution techniques by comparison and contrast of the results achieved by:

1. Normal real aperture processing,
2. Unfocussed synthetic aperture processing,
3. A focussed / unfocussed hybrid technique,
4. A fully focussed synthetic aperture method,
5. A limited two-dimensional fully focussed algorithm, and
6. A limited two-dimensional focussed / unfocussed hybrid algorithm.

These six azimuth resolution methods will be compared both through simulation (in this chapter) and real data obtained from an air medium synthetic aperture sonar built in a laboratory.

A. DATA RECORDING

The first step in recreating a target via signal processing techniques is to understand the physical reality which generated those signals. In this case, a finite length pulse of energy is sent toward a point target. The target scatters the energy in all directions as depicted in Figure 3.1.
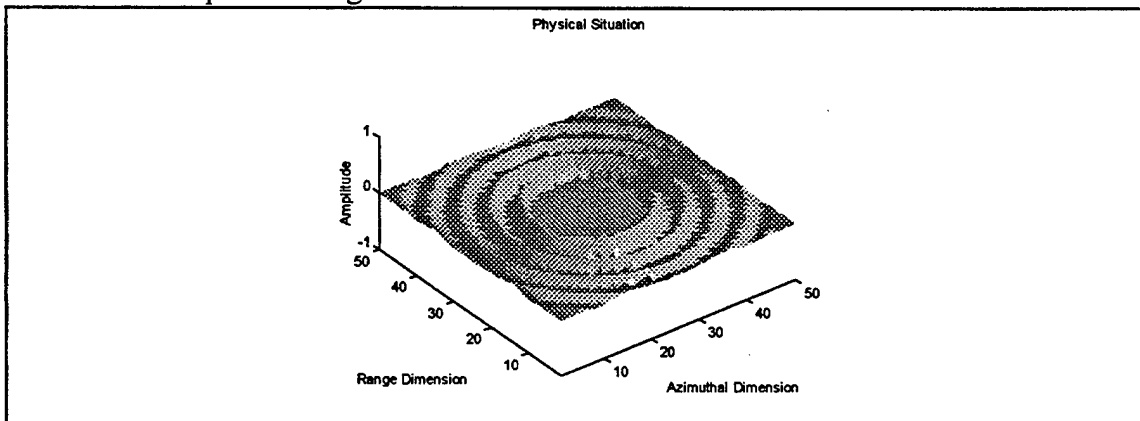
Figure 3.1 Target Scattering

If three sensor elements are placed at points along the azimuthal dimension axis, the return signals can be recorded as a function of time as shown in Figure 3.2.
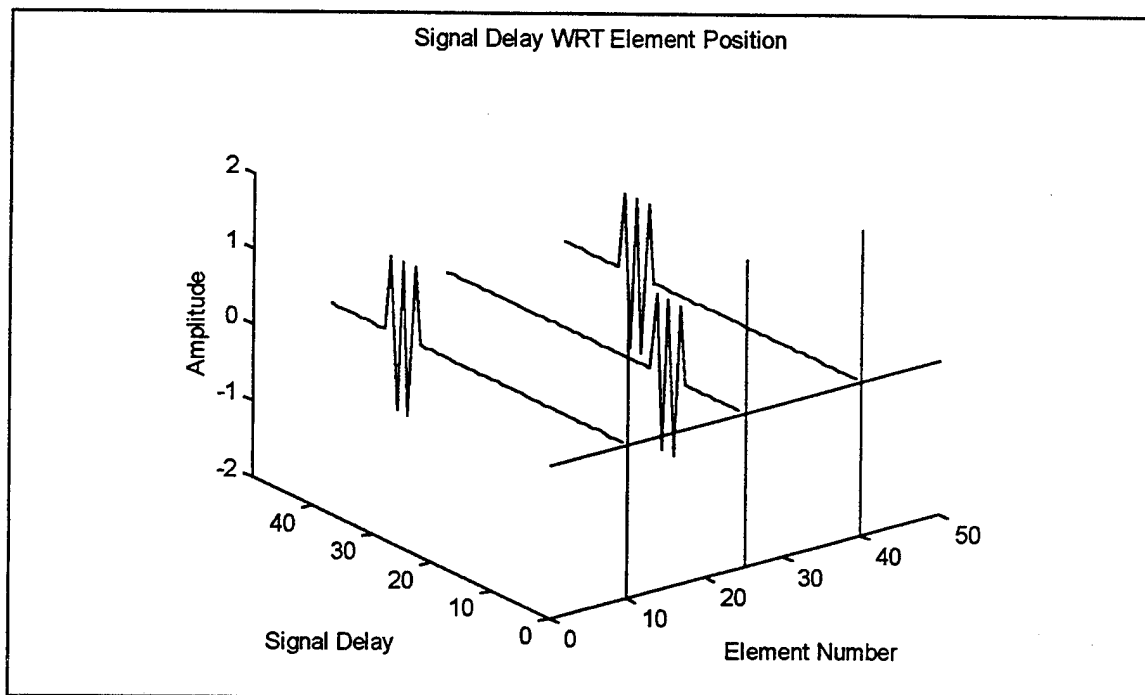
Figure 3.2 Three Element Return Signals

The greater the signal time delay between the transmission and reception of a pulse, the greater the distance between the target and the transducer. Since the center element shows the shortest time delay, the target must be closer to the center of the array. The other two elements are farther from the array and thus show longer delays.

Normally, more than three elements are used to survey a target. These plots can be represented more efficiently by packing the plots next to one another so that the individual traces combine to form a surface. The height of this surface is the amplitude of the received signal at the selected time delay for the chosen sequential element position. As noted above, the time delay is proportional to the range from the element to the target. The element positions are normally placed at equal intervals along a straight line such that the sequential element numbers are proportional to the displacement in the azimuth dimension. The data structure used to store and process this surface is a matrix where the rows correspond to discrete ranges to targets, the columns correspond to discrete azimuthal displacements, and the matrix element values correspond to the amplitude of the received signal. To display the salient features of this three dimensional surface, both a top

view as shown in Figure 3.3a and a perspective view as shown in Figure 3.3b will be used to illustrate the aspects of the topology of the signal matrices that will be exploited in this thesis.
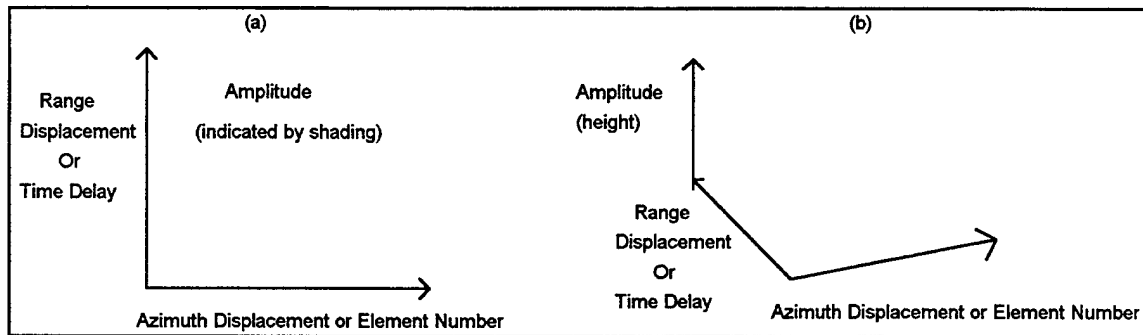
Figure 3.3 Received Matrix Formats

An example of the recording of data in these formats is that from the simulation of three point targets at different ranges illuminated by a rectangularly enveloped cosine pulse from a rectangular beam pattern source-receiver pair as illustrated in Figure 3.4.
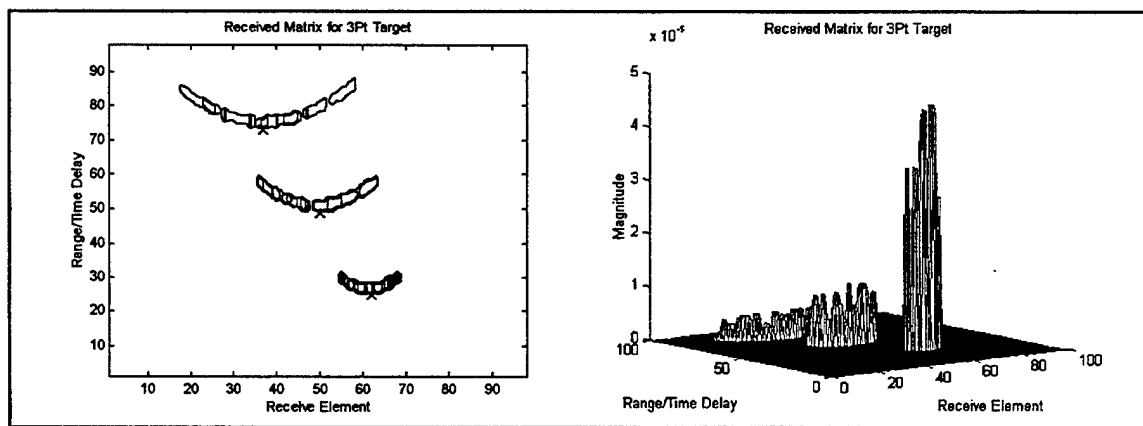
Figure 3.4 Received Signal Matrix for Three Point Targets

A casual glance at the above data yields the following observations:

    1. The targets form arcs centered and immediately farther in range than the target locations (indicated by the "x's").

    2. The arcs are of uniform thickness and uniform magnitude across that thickness.

    3. The arcs are wider and deeper but of smaller magnitude for longer range targets.

Each of these effects has implications which must be understood to be exploited so that a faithful reconstruction of the target field can be produced.

13

1. Hyperbolic Range Arcs

The arcs noted in the first observation are actually hyperbolas. Referring to Figure 3.5, we can easily explain the reason for the hyperbolas by relating the reception range, R, to the azimuthal displacement, X, and the range at closest point of approach (CPA), R(0), in the equation

$$R=\sqrt{R(0)^2+X^2}. \tag{3.1}$$
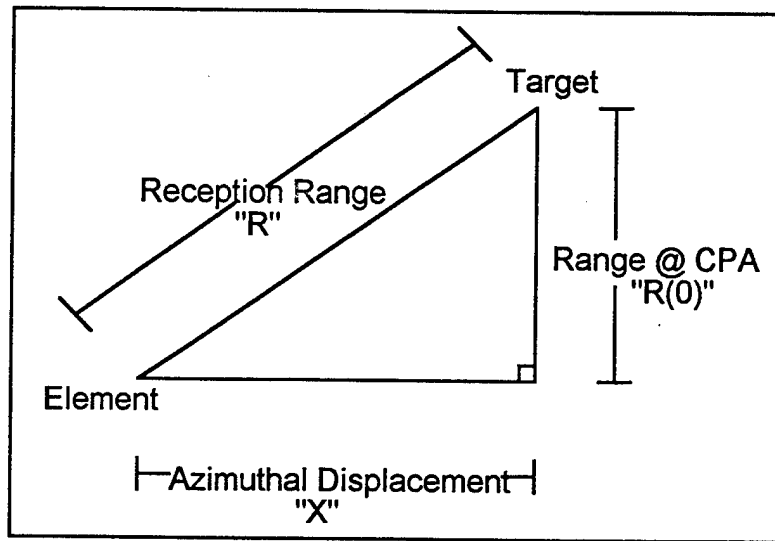
Figure 3.5 Source of Hyperbolas

In a detection or imaging situation, the receiving element spacing or azimuthal displacement as a function of element number are known. Likewise, the range resolution of the system based on pulse length is specified; so, the possible ranges of interest at CPA are known. Thus, the reception range as a function of number of elements away from CPA and CPA range bin can be calculated. Reception range is directly proportional to propagation delay time which is directly proportional to sample index in the received signal matrix as illustrated by Figure 3.6.

Hyperbolic Time Delays

Basic Shapes
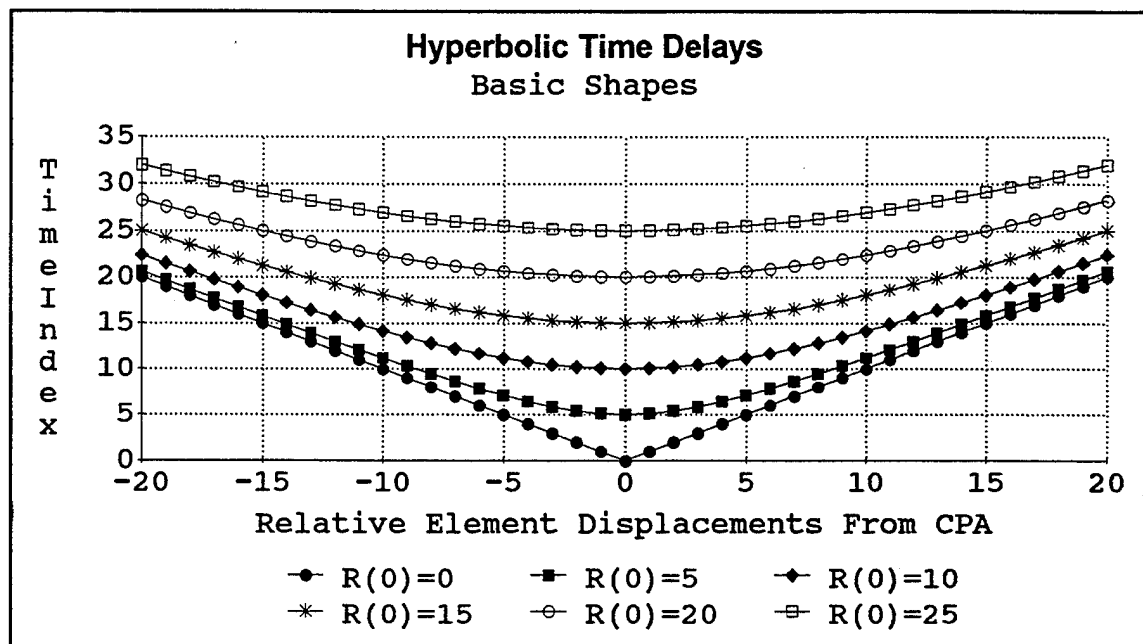
Figure 3.6 Hyperbolic Time Delays

2. Immediate Two-Dimensional Correlation

When any radar or sonar operator looks at a display similar to the one in Figure
3.4, he or she instinctively knows that the targets are located at ranges and azimuths just
short of the leading edges of the arcs on their centerlines (at the inserted "x's"). The goal is
to produce an algorithm capable of recognizing the same patterns.

The mathematical equivalent of the operator's recognition process is a two-
dimensional correlation. The operator has formed a mental picture of how the signal
associated with a target should look, and compares that prototype image to each group of
pixels in the presented signal field searching for a match. Where the signal has the same
appearance as this mental model, the correlation is high and the operator is likely to call
the result a target.

The same approach is used to design a suitable computer algorithm. In this case, a
different correlation matrix or 2-D matched filter is needed for each range bin of interest,
since shorter range hyperbolas are shaped slightly differently than longer range ones, as
illustrated previously in Figure 3.6. Since the computational power required to recognize a
two-dimensional target is directly related to the size of the template used, the first step in
creating these matrices is to discover how large they will need to be. If we were to draw
the minimum size box enclosing a complete target hyperbola, it would have the shape of

15

Figure 3.7a. The actual target path through the corresponding physical beam pattern is depicted in Figure 3.7b.
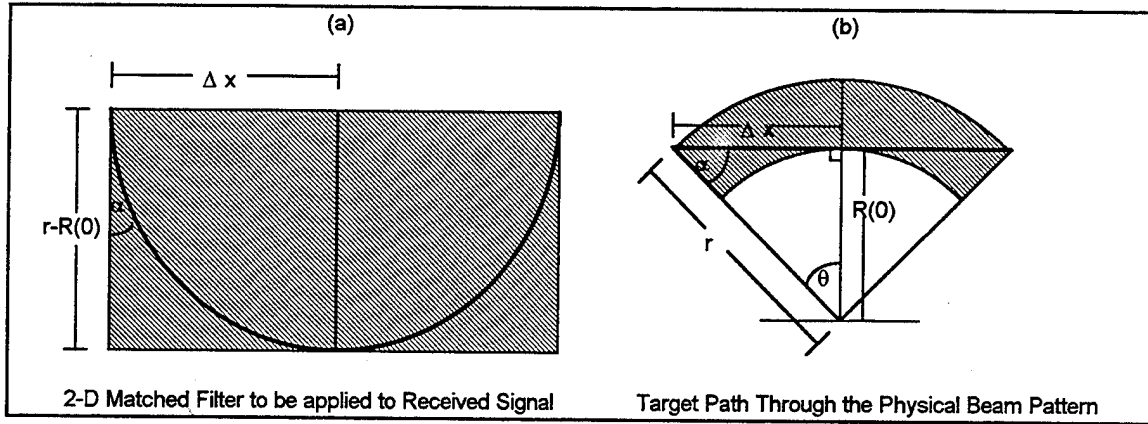
Figure 3.7 Two-Dimensional Matched Filter Geometry

The half width of the template, Δx, can be computed from the relation

$$\sin(\theta)=\Delta x \, /(R(0)^2+\Delta x^2)^{\frac{1}{2}};$$

which, when solved for Δx yields

$$\Delta x = R(0) \times \sin(\theta) \, / \, (1-\sin(\theta)^2)^{\frac{1}{2}}. \qquad (3.2)$$

Since θ is fixed by the dimensions of the transducers and the wavelength of the energy projected, the later portion of this formula is a constant. Thus, the azimuthal dimension covered by the 2-D matched filter, is directly proportional to the range of the bin being processed.

The radial range dimension to be covered by the matched filter, $r-R(0)$, can be calculated from the equation

$$r-R(0)=(R(0)^2+\Delta x^2)^{\frac{1}{2}}-R(0),$$

which, when solved for $r-R(0)$ after eliminating Δx yields

$$r-R(0)=R(0) \times [1/(1-\sin(\theta)^2)^{\frac{1}{2}}-1]. \qquad (3.3)$$

Since, θ is fixed by the dimensions of the transducers and the wavelength of the energy transmitted, the later portion of this formula is also a constant. Consequently, the radial range dimension covered by the matched filter is likewise directly proportional to the range of the bin being processed. These two facts combine to make the total size of the matched filter proportional to the square of the processing range.

This R^2 size dependence of the matched filter is a direct consequence of the conservation of energy. As the range to the target increases, the amplitude of the return

16

signal decreases as $1/R^2$ as noted in observation 3 (ignoring attenuation). Since all of the target's energy has to be recorded on the correlation template, the template size has to increase by R^2, making it wider and deeper, again as noted in comment 3, to compensate for this spreading effect.

To determine the size of the matched filter matrices, consider searching the closest 1m range swath from an air medium synthetic aperture system using a transducer full beam width of 1 radian. This makes the values of the azimuthal and down range constants of Equations 3.2 and 3.3, 0.5463 and 0.1395 respectively. These specifications along with a 100KHz sample rate, a sound velocity of 343 m/s, a 0.0125m range resolution, and an element spacing of 0.00686m result in the size requirements for the correlation matrix sizes shown in Table 3.1.

Range	Range(meters)	# Azimuth Columns	#Range Rows	Total Size
1st Bin	0.0125	2	2	4
Mid-Range	0.1125	18	10	180
Last Bin	1	160	82	13120

Table 3.1 Ideal 2-D Matched Filter Matrix Size Requirements

The range associated with the first range bin is the range resolution of the final image; hence, all ranges indicated in the final display are integer multiples of this value. The ranges noted are the distances to the point source targets whose two-dimensional received signal signatures need to be contained in the correlation matrices of the sizes specified by the next two columns. The number of azimuth columns in the two-dimensional correlation matrix or matched filter is the integer number of azimuth element spacings required to cover twice the half width calculated in Equation 3.2 for the range specified. The number of range rows in the matrix is the integer number of time delay samples required to encompass the corresponding radial distance specified by Equation 3.3. The total size or number of elements in the correlation matrix is the product of the number of rows and the number of columns. Figure 3.8 shows these results on LOG-LOG scales, demonstrating their R and R^2 dependence.

Ideal 2-D Matched Filter
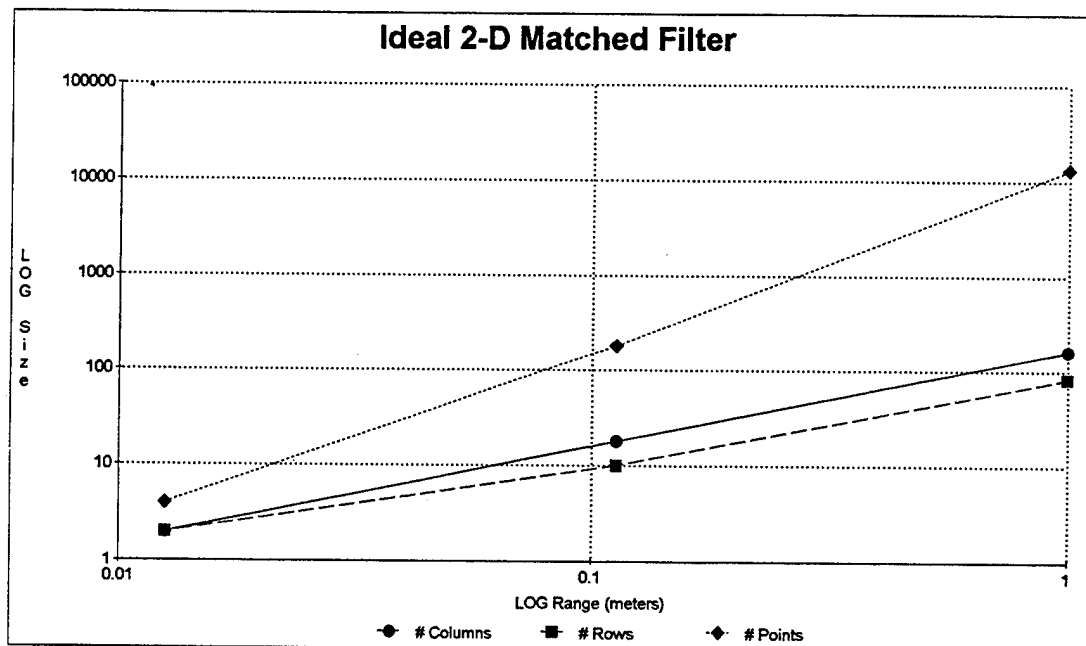
Figure 3.8 Ideal 2-D Filter Size

So, the total memory required to store the 2-D correlation matrices would be the largest matrix size times the number of matrices (or range bins) times the storage per matrix element. (The smaller matrices would be zero padded to fit in the larger matrix format.) This would require 13120 elements per matrix times 80 matrices times 8 bytes per element or over 8 Mbytes of memory just for the 2-D correlation matrices.

Other alternatives need to be considered to reduce these excessive memory requirements. Since all dimensions scale linearly with range, all the 2-D matched filters are geometrically similar - that is identical except for the change in size. In support of this assertion, the α angles at all scales are the same (61.35°), and the aspect ratios are the same (0.5+). Using this, we could store only the largest matrix and simply scale down the results for the other ranges. This would work except for one problem - the pulse width used to illuminate the target is not infinitesimal and it does not scale with range. This is the noted uniform thickness and height observed in remark 2. Assuming a five cycle burst at 25KHz sampled at 100KHz, the actual matrix sizes required are listed in Table 3.2.

18

Range	Range(meters)	# Azimuth Columns	#Range Rows	Total Size
1st Bin	0.0125	2	22	44
Mid-Range	0.1125	18	30	540
Last Bin	1	160	102	16320

Table 3.2 Actual 2-D Matched Filter Matrix Size Requirements

The numbers and the meanings of the first three columns of this table are identical to those of Table 3.1. The range row column in this table accounts for the additional twenty time delay samples which arise from the non-zero pulse length used. The resultant effects on final matrix sizes are noted in the total size column. Figure 3.9 shows these results graphically.
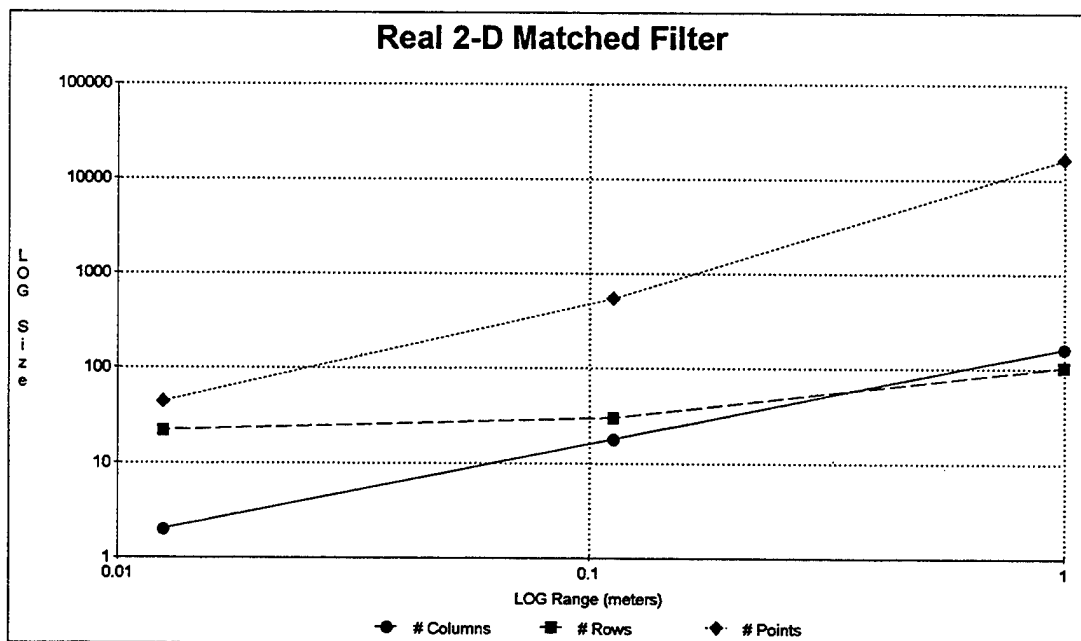
Figure 3.9 Actual 2-D Filter Size

This fact makes simple linear scaling impractical and multiplies by 1.2 the memory requirements for stored correlation matrices. By packing the matrices so as to eliminate either the zero columns or the zero rows, storage requirements can be reduced by a factor of two. However, since this example concerns only the first 1m from the platform, the next meter surveyed to the same resolution will require three times the additional memory and the next 1m after that will require another five memory blocks etc.

19

Even if a near zero length pulse, such as an impulse, could be used to force the correlation matrix to be linearly scalable, the synthesized matrices will still need to be compared to received signal data matrices of the same size. This comparing or correlation process requires computational time related to the size of the matrices compared. With prudent selection of matrix sizes and use of Fast Fourier Transform (FFT) techniques to perform the two-dimensional correlations, this calculational workload can be minimized to

$$F \cong 3 \times N_c \times \log_2(N_c) \times N_r \times \log_2(N_r), \qquad (3.4)$$

where "F" is the number of floating point computations, "N_c" is the number of columns in the matrices to be correlated, and "N_r" is the number of rows in the matrices to be compared. (Proakis, 1992, pp. 717) Since both N_c and N_r are proportional to range, F rises even more quickly than the R^2 dependence of the correlation matrix size.

Thus, the R^2 matched filter matrix size dependence quickly makes this highly accurate method impractical at reasonable search ranges. The remainder of this paper will concern itself with methods whose initial memory and computational requirements are much more modest and have a more benign range dependence.

B. RANGE LOCALIZATION

Since an immediate two-dimensional correlation to position targets in range and azimuth is impractical, some progress can be made if the targets can be located in at least one of the dimensions as a start. Invariably, the first choice is to localize the targets' ranges because this is the less difficult of the two problems to solve.

1. Envelope Processing

The easiest first approximation to the target ranges in Figure 3.4 is that their positions are known to within the thicknesses of their arcs. This is one form of envelope processing and it yields the range resolution described by Equation 3.5.

$$\Delta R = (c \times PT)/2, \qquad (3.5)$$

where "ΔR" is the range resolution, "c" is the speed of energy propagation in the medium, and "PT" is the time duration of the energy pulse sent out. Thus the time duration of the

outgoing energy pulse in this type of processing determines the minimum separation at which two targets can be differentiated in range as illustrated in Figure 3.10.
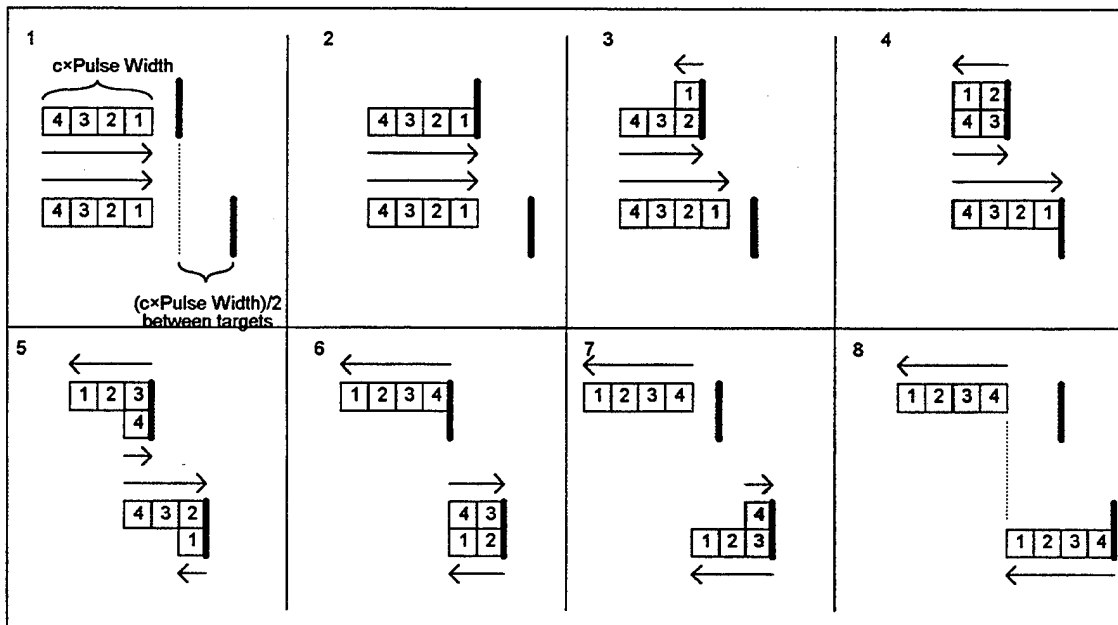
Figure 3.10 Envelope Processing Range Resolution

As can be seen above, at this separation, the pulses from the two targets return one after the other with no break or overlap. If the targets are spread in range by more than this distance, the return pulses will separate and the targets will be distinguishable. If on the other hand, the targets are closer in range than this, the return pulses will overlap and the targets will be indistinguishable.

This technique is simple and effective as long as the pulse lengths are short and the target returns are strong. If, however, the targets are weak, the transmitter may need to put more energy on the targets to illuminate them. If the transmitter cannot use a higher amplitude signal, it may use a longer pulse to provide the necessary detection power. This degrades the range resolution performance of an envelope processed system.

2. Coherent Range Processing

A technique to permit long pulse lengths and yet still obtain reasonable range resolutions is to use the phase information present in the return pulse- to use coherent processing. This is done using a correlation or matched filter process on each column (receiver element time record) from the received signal matrix above. This principle is expressed mathematically in the form the correlation integral, which is

$$R_{hx}(T) = \int_{-\infty}^{\infty} h(t) \bullet x^*(t+T) \ dt. \tag{3.6}$$

(Strum, 1994, pp. 288) The essence of this function is that it is an inner product (also called a dot-product) between h(t), the reference signal, and x*(t), the test signal (* denotes complex conjugation). As in geometry, this product is maximized when the two vectors or functions are most nearly parallel or similar except for a possible time shift. Mathematically, this is expressed in the relation

$$|R_{hx}(T)| \leq [R_{hh}(0) \bullet R_{xx}(0)]^{1/2}, \tag{3.7}$$

with equality holding and $R_{hx}(T)$ at a maximum when h(t)=x(t+T). (Strum, 1994, pp. 288) So, to optimize the output of a correlator, the reference signal, h(t), is set equal to x(t), the expected return signal from a point target. "T" then becomes the relative time offset or delay between the received signal and the reference "template". This comparison is performed with every time difference, T, to produce the function $R_{hx}(T)$. The time of maximum response of $R_{hx}(T)$ is then the propagation time delay which is proportional to the range to the target. The two dimensional extension of this correlation concept, where T is replaced by range and azimuth variables and the maximum response occurs when those variables match the location of the target, is exactly what was discussed in Section A.2. of this Chapter. This same effect can also be achieved using the relationship between convolution and correlation through the definition of the convolution integral, which is

$$y(T) = h(T) \star x(T) = \int_{-\infty}^{\infty} h(t) \bullet x(-t+T) \ dt = \text{IFT}\{H(\omega) \cdot X(\omega)\}. \tag{3.8}$$

(Strum, 1994, pp. 50) This is an inner or "dot" product, or similarity check, between the reference, h(t), and the test signal, x(-t). In a fashion analogous to that of the correlation, the peak response occurs when h(t)=x*(-t) and T equals the time delay between the transmission and reception of the return signal. But, h(t)=IFT{H(ω)}, so, the "matched filter", H(ω)=FT{x*(-t)}. This matching process between the expected return signal and the received signal is again repeated for every time shift, T. Once more, the maximum response time of this "matched filter" occurs at the time corresponding to the propagation delay time associated with the range to the target(s).

In this way, the ability to resolve targets in range is separated from the pulse length and connected to the matched filter's ability to recognize a return pulse uniquely and when

that pulse returned. Hence, the more unique and complex the pulse type used, the better the filter will be able to establish arrival time and separate targets. A direct measure of the complexity of a waveform is its bandwidth. As a result, a wide bandwidth signal used with a matched filter results in a fine range resolution as expressed in Equation 3.9.

$$\Delta R \cong c \, / \, (2 \times BW), \tag{3.9}$$

where "BW" is the bandwidth of the pulse type [$\cong 1 \, / \, (PT)$ for a rectangular constant frequency pulse].

For this simulation, we are using a rectangularly enveloped constant frequency pulse which is not optimized for this technique as are FM chirps or phase encoded signals. Thus, the matched filter will only provide a peak amplitude response at the target range at the price of widening the correlated signal thickness as illustrated in Figure 3.11.
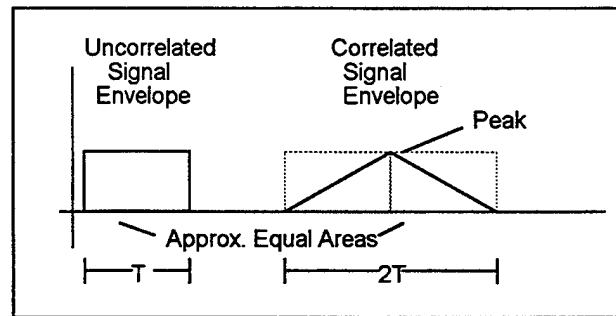
Figure 3.11 Effect of Matched Filtering on a Constant Frequency Rectangularly
Enveloped Pulse

In this application, the "matched filter" is set up as the time reversed response of the sonar to a point target, based on the assumption that most other targets can be built up as a collection of point targets. The correlation process is accomplished as a convolution between this time reversed expected response and the received signal. Because convolutions are often performed digitally via Fast Fourier Transforms (FFT's), they are computationally time efficient. The application of this method to the columns of the simulated received signal data of Figure 3.4 above is displayed in Figure 3.12.

23
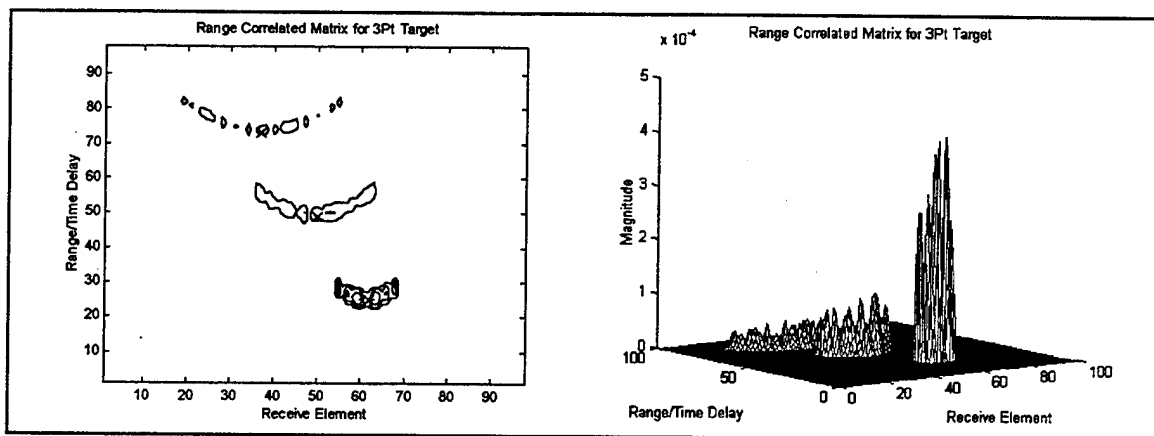
Figure 3.12 Results of Range Correlation on Received Signal Matrix

While it is difficult to discern from these plots, the center of each hyperbola "band" has a ridge or peak which was not present in the received signal. (The "x's" are located in the same place as in the raw received signal matrix.) The price of this peak is the obvious broadening of the base of the band described earlier.

If the signal processing is being performed in complex form, at this point, the signal can be modulated with the complex carrier to reduce the signal to base band. This is not required but can reduce bandwidth requirements for subsequent processing. This technique is not used here to avoid the front end overhead in terms of complex calculations.

At this point, the targets are considered to be localized in range from the bands seen in the original received signal matrix to the ridges found in the range correlated matrix. Hence, the initial two dimensional search has been reduced to a one dimensional search. The next step is to locate the target positions along these curves.

C. AZIMUTH COMPRESSION

This thesis considered six different methods of locating a target in azimuth. The six methods were:

1. Real aperture beam forming,
2. Unfocussed synthetic aperture processing,
3. A focussed / unfocussed hybrid technique,
4. A fully focussed one-dimensional synthetic aperture method,
5. A fully focussed limited two-dimensional technique, and
6. A focussed / unfocussed hybrid limited two-dimensional algorithm.

1. Real Aperture Beam Forming

The most commonly utilized method of azimuth location is real aperture beam forming, which uses the actual beam width of the source. This is used because it requires no additional signal processing. The range correlated matrix, Figure 3.12, is the final image for the real aperture. The real aperture beam forming is done before the signal comes to the electronics by the physical construction and positioning of the array elements relative to the wavelength of the energy propagated and the desired direction of detection. The array elements are wired together and physically spaced such that the signals sent out and returned coherently add in the direction of the main lobe and coherently interfere or cancel in other directions. The resultant width of this region of constructive interference or beam width is

$$\beta \cong \lambda / D, \tag{3.10}$$

where "β" is the full beam width, "λ" is the wavelength of the energy propagated, and "D" is a length characteristic of the array.

Hence, for a real array, a narrower beam can be achieved either by using a shorter wavelength relative to the array size, or building a larger array relative to the wavelength. A narrow beam width is desirable because targets within the same beam are unresolved. This means that the resolution at a range is the width of the beam at that range or

$$\Delta X \cong \beta \times R(0) \cong R(0) \times \lambda / D, \tag{3.11}$$

where "ΔX" is the azimuthal resolution at range "R(0)".

Consequently, azimuthal resolution degrades as a direct function of range. Figure 3.12 demonstrates this since the targets show azimuth swath widths directly proportional to range. While this poor azimuth resolution performance is well explained by this formula, Figure 3.12 still represents a poor final image and would not normally be considered acceptable. The reason such a wide beam width element was chosen will be explained later where this broad beam width will produce advantages for synthetic aperture processing. If a real aperture array was to be designed for this target field, it would be at least ten times as large (multiply D by 10) to produce the image of Figure 3.13 instead of Figure 3.12.
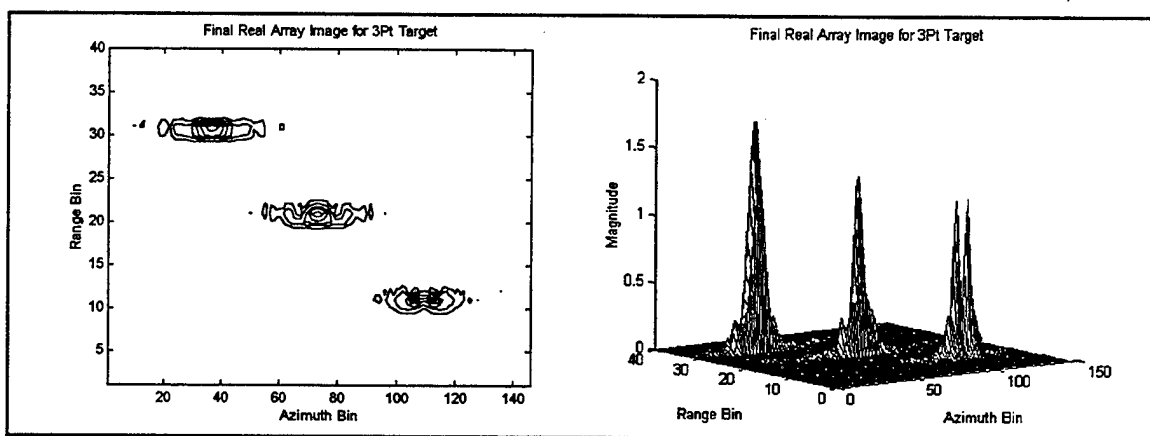
Figure 3.13 Real Aperture Image with Element Size = 10×D

As can be seen easily in Figure 3.13, the increase in size of the sensing elements significantly improved the image quality at the obvious cost and inconvenience of larger elements. While the target azimuth resolution starts out from a much better initial value, it still degrades linearly with range. This is the fundamental limitation of the azimuth resolution capabilities of a real aperture system: to gain high resolution at long ranges requires unreasonably large physical arrays. If these arrays cannot be constructed at the required size and tolerances, a different approach needs to be used.

2. Unfocussed Synthetic Aperture Processing

Since a physically larger real array is unavailable to achieve the desired resolution, the next attempt at a solution is to try to build the equivalent of a large physical array from the data processed by a smaller array. Since the larger real array coherently adds the signals it receives at each of its sensing elements, the first attempt at building this larger array should involve simply adding neighboring azimuth element returns coherently. This process is called an unfocussed array, since the spherical curvature of the waves is ignored. This strategy will improve the performance of the system as long as the phase difference among all the returns summed is small. Figure 3.14 provides the background for how this pitfall is avoided.
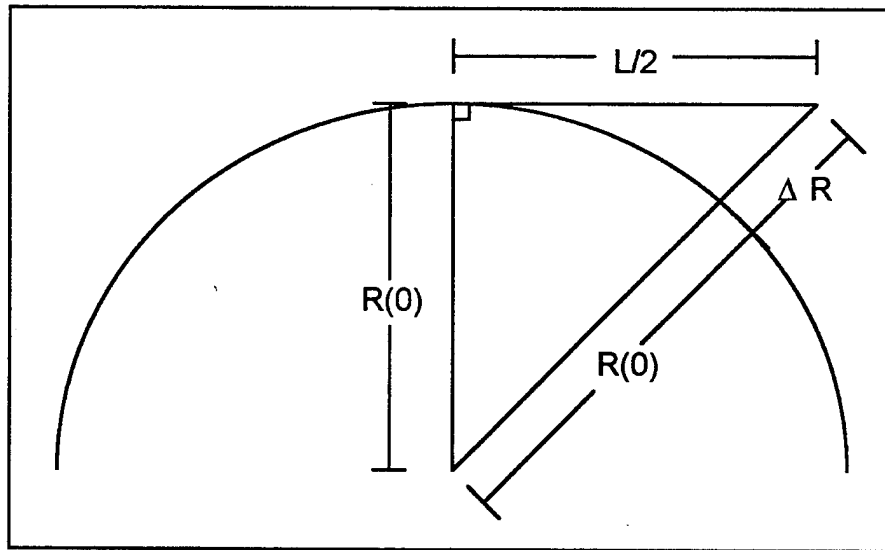
Figure 3.14 Unfocussed Plane Wave Approximation Phase Geometry

The goal of a real aperture array's coherent summation is to produce a transmit and receive plane wave in a specific direction by adding the phases of a large number of spherical waves. Thus, there are no true plane waves; but, if examined over a small enough arc, the spherical wave appears to be planar (similar to the surface of the Earth). To approximate a spherical wave front as a plane wave, a maximum tolerable two way phase error needs to be specified. In this case, a two way phase error of $\pi/4$ across a single wave front is considered acceptable. To achieve this, ΔR can be no larger than $\lambda/16$. Thus, by Pythagoras:

$$(L/2)^2 + R(0)^2 = (R(0) + \lambda/16)^2 = R(0)^2 + R(0) \times \lambda/8 + (\lambda/16)^2.$$

Since R(0)>>λ, the last term on the right can be neglected resulting in:

$$L=\sqrt{(R(0)\lambda/2)}, \qquad (3.12)$$

where "L" is the full azimuthal distance over which the element signals can be summed at range "R(0)" without exceeding a π/4 phase error, "R(0)" is the range to the target at the closest point of approach (CPA), and "λ" is the wavelength of the energy propagated in the medium.

After performing these variable width coherent sums (or moving averages) across the constant range rows, we need to determine the expected azimuthal resolution of this technique. To calculate at that value, the geometry of Figure 3.15 needs to be understood.
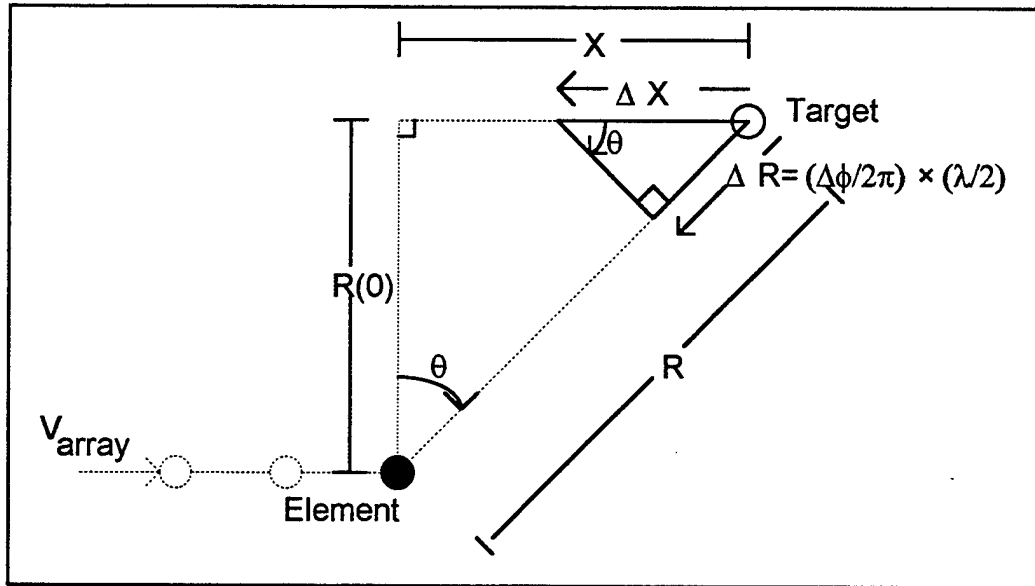
Figure 3.15 Coherent Azimuth Resolution Geometry

Here, "R" is the instantaneous range from element to target, "R(0)" is the range to the target at the closest point of approach (CPA), "X" is the target's total along track displacement to be observed, "θ" is the angular beam width where the target will be visible, "V_{array}" is the velocity of the array along its track, "ΔX" is the azimuthal resolution, "ΔR" is the change in range "R" associated with "ΔX", "λ" is the wavelength of the transmitted waveform, and "Δϕ" is the two way phase shift in the return signal associated with "ΔX". (This two way property causes ΔR=Δϕ/2π×λ/2 where Δϕ/2π is the fraction of a period phase change observed with a ΔR change in range.)

28

From the geometry of the setup:

$$\sin(\theta)=X/R= \Delta R/\Delta X,$$

which yields

$$\Delta X=\Delta\phi/2\pi\times\lambda/2\times R \ / \ X, \tag{3.13}$$

by taking into account that $\Delta R=\Delta\phi/2\pi\times\lambda/2$.

For processing, the minimum $\Delta\phi$ between wave fronts to which the system is sensitive is 2π and X, the observation distance, is L. Therefore,

$$\Delta X=(2\pi)/2\pi\times\lambda/2\times R \ / \ L,$$
$$\text{and then,}$$
$$\Delta X=(2\pi)/2\pi\times\lambda/2\times R \ / \ \{\sqrt{(R(0)\lambda/2)} \ \}. \tag{3.14}$$

However, for an element beam width of 1 radian or less R can differ from R(0) by at most 12% [1-cos(.5rad)]. So, $R\cong R(0)$. Thus Equation 3.12 can be simplified to

$$\Delta X\cong\sqrt{(R\times\lambda/2)}. \tag{3.15}$$

Comparison of this result with that of the real array in Equation 3.8 yields some interesting observations. First, there is no longer a dependence on D, the size of the sensing element. Therefore, small sensors with large beam widths can be used freely. Second, our azimuthal resolution no longer degrades linearly with the $R\times\lambda$ product, but with the square root of this quantity. Hence, an unfocussed synthetic aperture system will maintain equivalent azimuth resolution at much longer distances than a real aperture system. Since the moving average operation is computationally very simple, this capability can be added to a real aperture system for only a minor calculational cost. The result of this operation on the data of Figure 3.12 is shown in Figure 3.16.
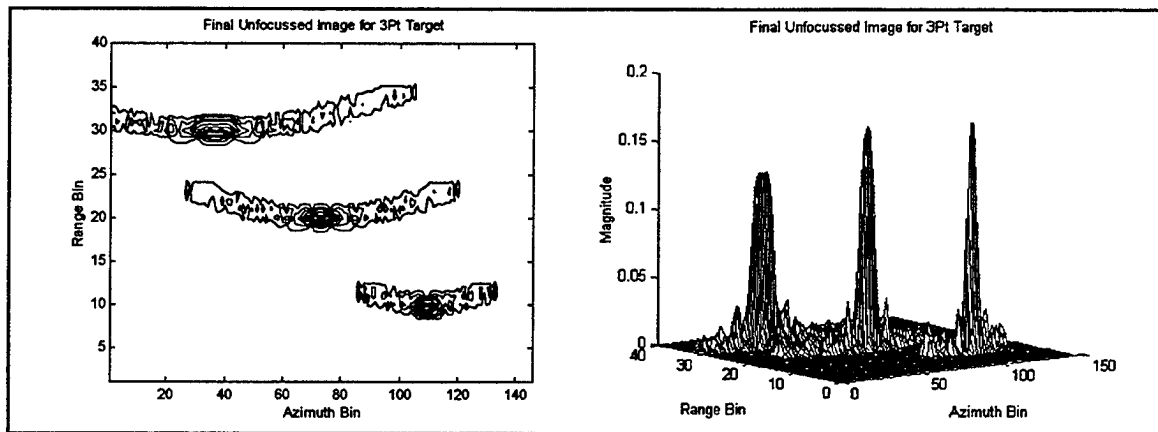
Figure 3.16 Unfocussed Synthetic Aperture Image

As expected, this image is a very significant improvement over the raw real array image of Figure 3.12 and is certainly of quality comparable to that of the real array image shown in Figure 3.13, with elements ten times as large. As anticipated, the unfocussed image does degrade with range, but at a slower rate of decay than the real array images.

3. Hybrid Focussed / Unfocussed Synthetic Aperture Processing

Thus, unfocussed synthetic aperture processing does offer improved performance over standard real aperture methods, but it is rather severely limited by the plane wave approximation. A technique is needed which does not depend on a plane wave assumption but still operates on only one range bin row per target range at a reasonable computational cost. To see how to develop such a technique Figure 3.17 provides a close up view of a target hyperbola from Figure 3.12.
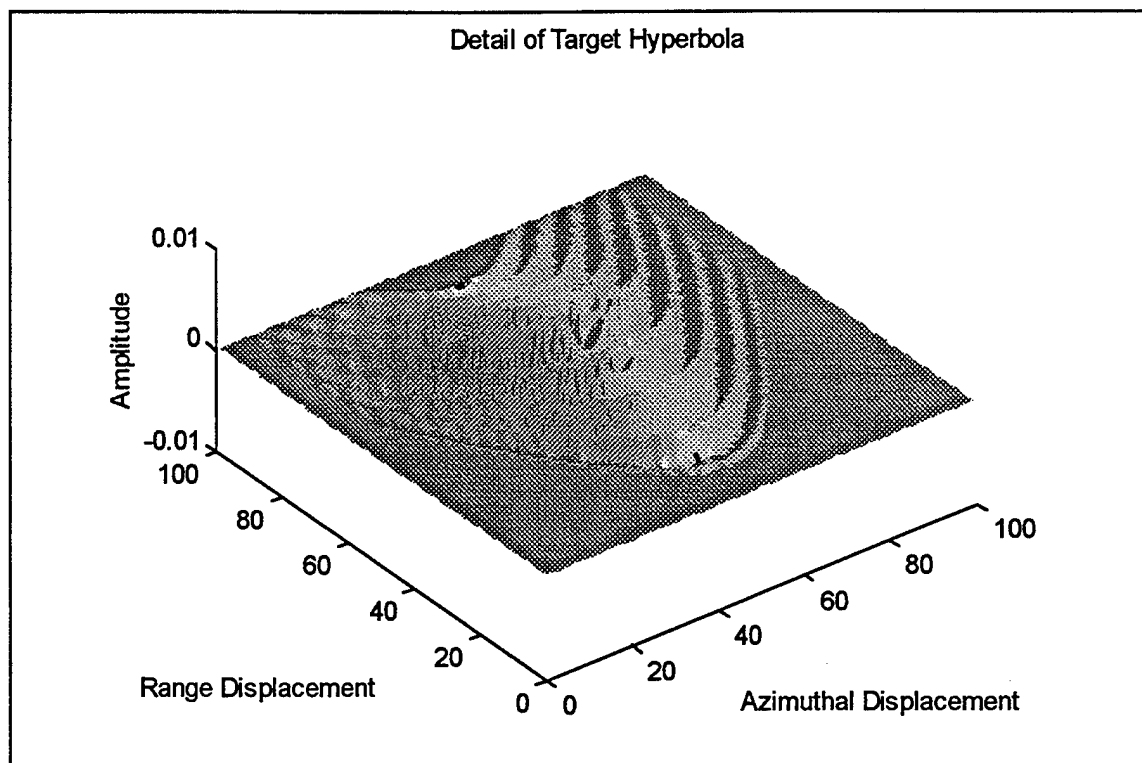
Figure 3.17 Expanded View of a Range Correlated Target Hyperbola

This Figure illustrates very effectively:

    a. The hyperbolic shape induced by the trigonometry of the situation,

    b. The peaking produced in the signal envelope by the correlation process
      in the range dimension,

    c. The constant width of the hyperbola band in the range dimension resulting
      from the non- zero pulse length,

    d. The wide though limited beamwidth in the azimuthal dimension, and

    e. The $R(0)^2 / R^2$ decay in signal amplitude away from the hyperbola centerline.

Taking a cross sectional constant range slice from this figure at the range of the highest peak in the correlation function, (the most probable location of the target) results in Figure 3.18.

31
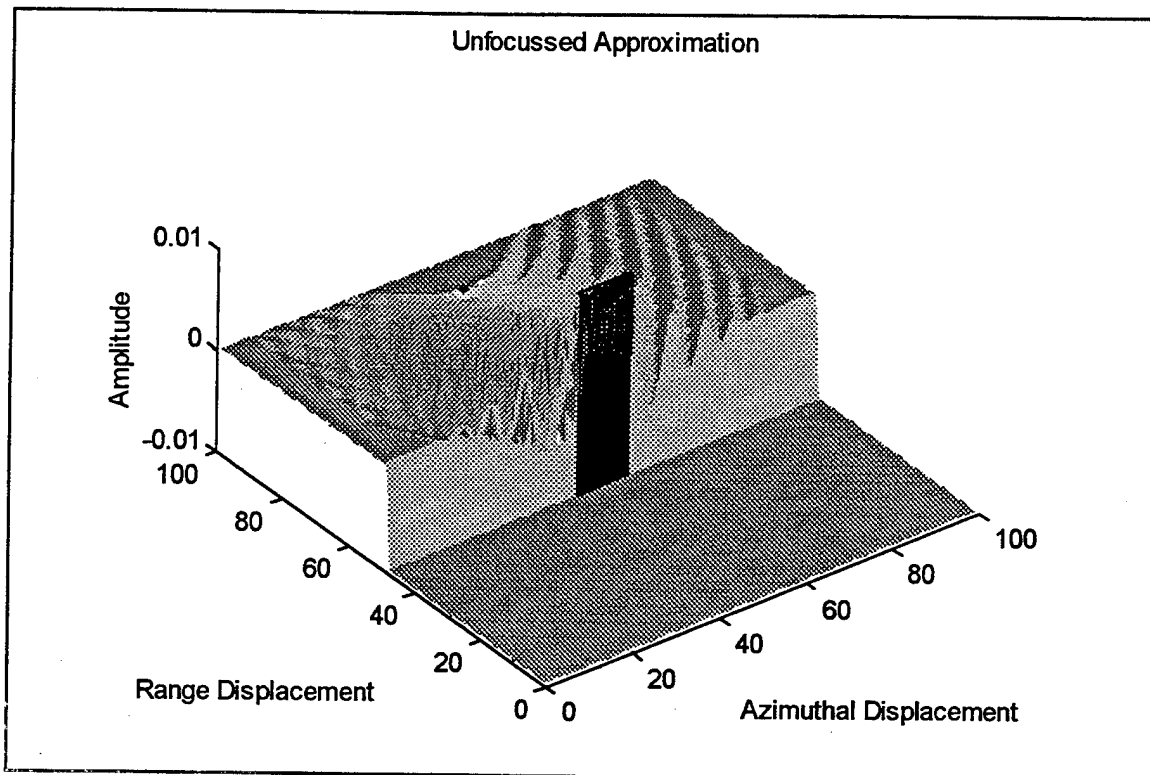
Figure 3.18 Unfocussed Approximation to Azimuthal Phase Variation

The black rectangle in the middle of the slice represents a moving average used by a pure unfocussed system as its correlator. This template is a first order attempt to recognize the location of the centerline of the hyperbolic wave pattern. However, because it is limited by phase constraints, the correlation length is very short and hence open to improvement. If, instead of this short rectangle, the entire length of the range cut was used as the correlation pattern as highlighted in Figure 3.19, significant improvements in azimuthal sensitivity (resolution) and noise rejection could be gained at the calculational expense of a nonuniformly weighted average.
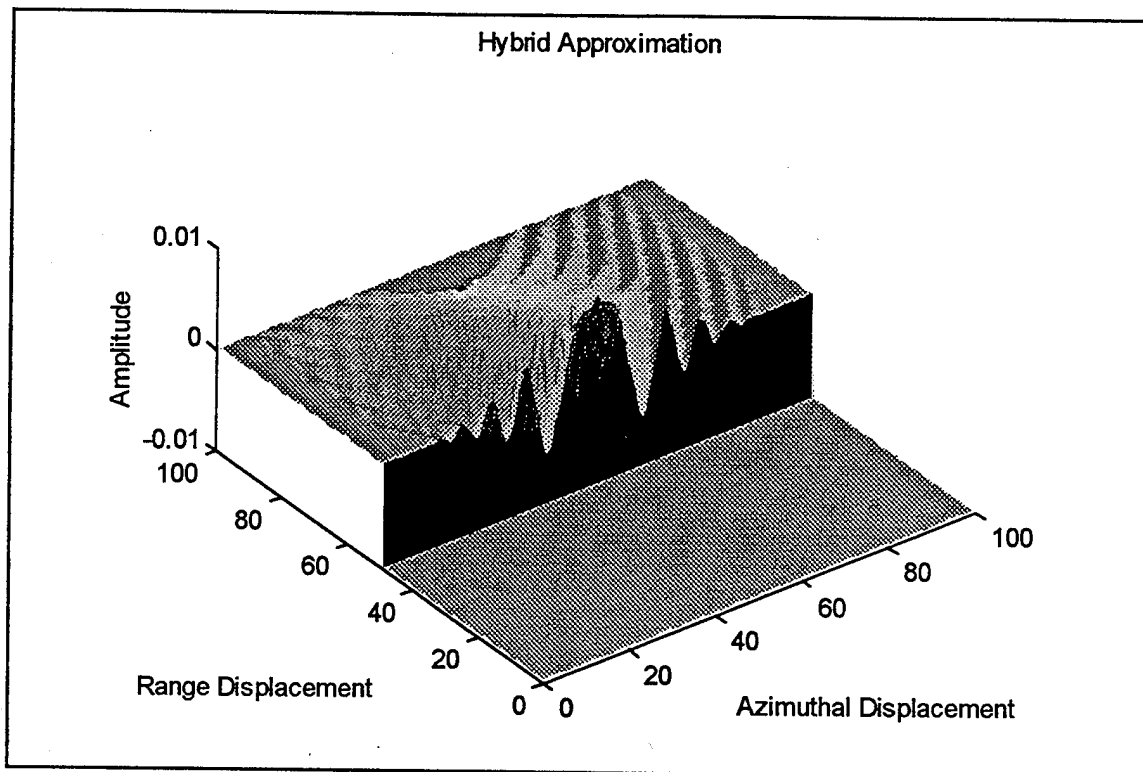
Figure 3.19 Hybrid Approximation to Azimuthal Phase Variation

This approach is a hybrid because it no longer makes the plane wave assumption that a pure unfocussed method does, but it also does not correlate along the curved range-azimuth contours for a single target the same way that a fully focussed technique does. (This will be explained in more detail later in this thesis).

Assuming that the range slice is available, we need to determine what improvement in resolution can be expected from its use. To evaluate this criteria, the width of the non-zero portion of the slice needs to be quantified. Figure 3.20 illustrates the situation.
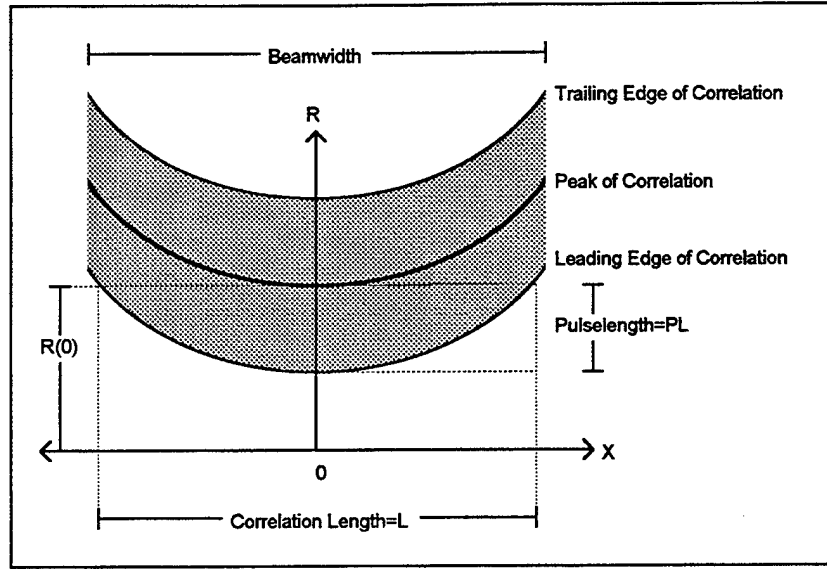
Figure 3.20 Hybrid Approximation Correlation Length

Repeating Equation 3.1, the peak of the correlation curve is located by

$$R=\sqrt{R(0)^2+X^2}. \tag{3.1}$$

The leading edge of the correlation band is described by

$$R=\sqrt{R(0)^2+X^2}-PL.$$

At X=L/2, this leading edge reaches R(0) resulting in

$$R(0)=\sqrt{R(0)^2+[L/2]^2}-PL.$$

Solving for L results in the relation,

$$L=2PL\sqrt{1+[2R(0)/PL]}. \tag{3.16}$$

Using PL = n×λ, where n is the number of cycles in a pulse, and assuming that R(0)>>PL, produces the approximation,

$$L\cong\sqrt{8R(0)\lambda n}. \tag{3.17}$$

Comparison of this result with Equation 3.12 for the pure unfocussed process reveals that this has a longer correlation length, but it is still proportional to the square root of R(0)×λ. Additionally, this correlation length grows with the square root of n. So, increasing the

34

pulse length will improve azimuth resolution. To quantify this benefit, insert this result into Equation 3.13 in the same manner as was done to produce Equation 3.14. The product of this effort is

$$\Delta X = \lambda/2 \times R / (\sqrt{8R(0)\lambda n}).$$ (3.18)

Again, for reasonable element beam widths, $R \cong R(0)$. Thus Equation 3.18 simplifies to

$$\Delta X \cong \sqrt{(R\lambda/32n)}.$$ (3.19)

This represents a significant improvement over the results of Equation 3.15 and a standard unfocussed algorithm. Additionally, the closer match between the signal and the correlation template results in a theoretically greater signal to noise ratio.

All of this correlation length and azimuth resolution analysis assumes that the outer edge of the correlation length is bounded by the leading edge of the range correlation hyperbolic band. If the hyperbolas are relatively flat and the pulse length is long enough, the correlation length could be limited by the real element beam width. This would occur when

$$L = R\lambda/D \cong \sqrt{8R(0)\lambda n}.$$

Assuming $R \cong R(0)$ and solving for R results in

$$R = 8nD^2/\lambda.$$ (3.20)

At ranges less than this, because the correlation length is now limited by real element beam width, an interesting effect occurs. Starting with Equation 3.13, which is

$$\Delta X = \Delta\phi/2\pi \times \lambda/2 \times R / X.$$ (3.13)

and using $\Delta\phi = 2\pi$, $R \cong R(0)$, and $X \cong R(0) \times \lambda / D$ from Equation 3.11 produces

$$\Delta X = 2\pi/2\pi \times \lambda/2 \times R / (R(0) \times \lambda / D),$$
which can be approximated to
$$\Delta X \cong D/2.$$ (3.21)

This is the resolution for a fully focussed system. This relation generates some interesting effects. A smaller transducer, which has a larger beam width, provides a finer azimuth resolution. This is the improvement in performance achieved by using a wide beam width

transducer which was mentioned earlier in Section III.C.1. This works in exactly the opposite manner of a real array which gains greater directivity or finer azimuth resolution from a larger element size.

What is more surprising than the D/2 factor on which focussed resolution depends is the number of factors which are missing from the relation such as range to target, total length of synthesized array, and carrier frequency. All of these factors are cancelled in essentially the same manner as illustrated in Figure 3.21.
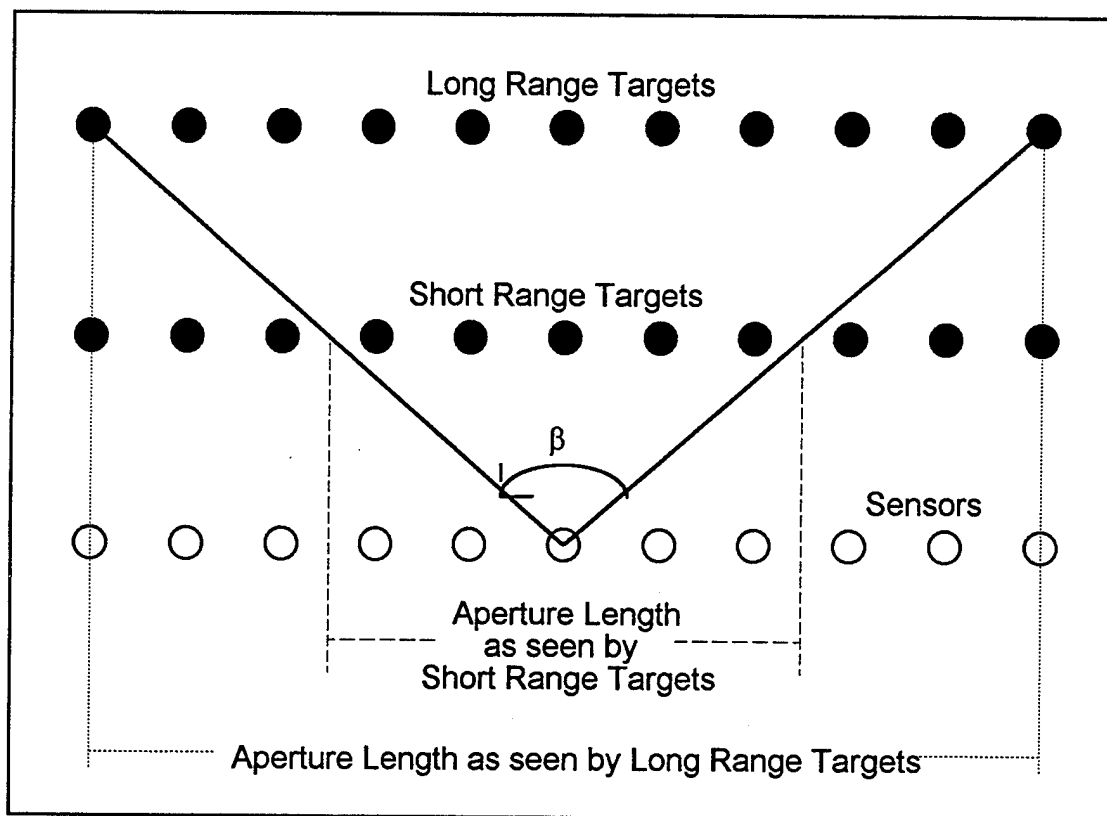
Figure 3.21 Synthetic Aperture Length as a Function of Range to the Target

While increasing range causes a real array beam to spread as a linear function of range and its resolution to degrade accordingly, a synthetic array uses this spreading effect to view the target for a distance which increases linearly with range. This makes the synthetic array appear linearly longer to a more distant target, cancelling the normal linear resolution degradation with range. As discussed above, within the conic portion of the radar or sonar footprint, the length of a synthesized array seen by a target is a function of range from the array, not the total length of the array (as long as the array keeps the target in view for the entire width of the footprint at that range). The principle effect of altering carrier frequency is to adjust the real beam width. Again, while a wider beam width

36

linearly reduces real beam resolution, it also linearly increases the length of the array which can see the target; hence, cancelling the beam divergence effect.

Thus, this result implies that infinitesimal azimuth resolution can be attained within the maximum range of the illuminator by using a small, omnidirectional sensor. This is not quite true. First, the approximations that $L \cong R \times \beta$ and $\beta \cong \lambda/D_{eff}$ introduce significant errors when $\beta > 1$ radian. Second, an omnidirectional sensor has no noise rejection capabilities from directions other than the target region making signal detection much more difficult. Finally, the energy pulse needs to travel from the near edge to the far edge and back again before the sensor can move from one azimuth resolution sample point to the next. If propagation speed is low and the two edges are widely separated, this requirement can limit the maximum speed of the sensor platform severely if full resolution capability of the system is to be exploited.

In addition to these factors, which limit all focussed systems, the hybrid method does not correlate along the maximum response range curve (hyperbola). Consequently, it has a degraded in response in low signal to noise applications when used in a fully focussed mode. However, since this procedure only requires data from a single range row, it can be run as a straight azimuthal correlation / matched filter / convolution across the azimuth columns and range rows. When implemented by Fast Fourier Transforms (FFT's), the resultant image can be produced in a time just longer than the almost instantaneous unfocussed image but also in a time orders of magnitude shorter than more robust fully focussed techniques. Hence, this hybrid process does show characteristics somewhere between those of an unfocussed and a fully focussed approach. Figure 3.22 illustrates the theoretical azimuth capabilities of the methods discussed for a five cycle, 25KHZ, rectangularly enveloped cosine wave transmitted in a rectangular shaped one radian beam pattern in air.

Theoretical Azimuth Resolution as a Function of Range

Original Real Array=V

Unfocussed Array=V

"D=10×s" Real Array=V

Azimuthal Sample Spacing=V

Focussed Array=^

Hybrid Array=V

1st Target=>

2nd Target=>

3rd Target=>

Hybrid Focussed/Unfocussed Boundary=>

Resolution (m)

Range (m)
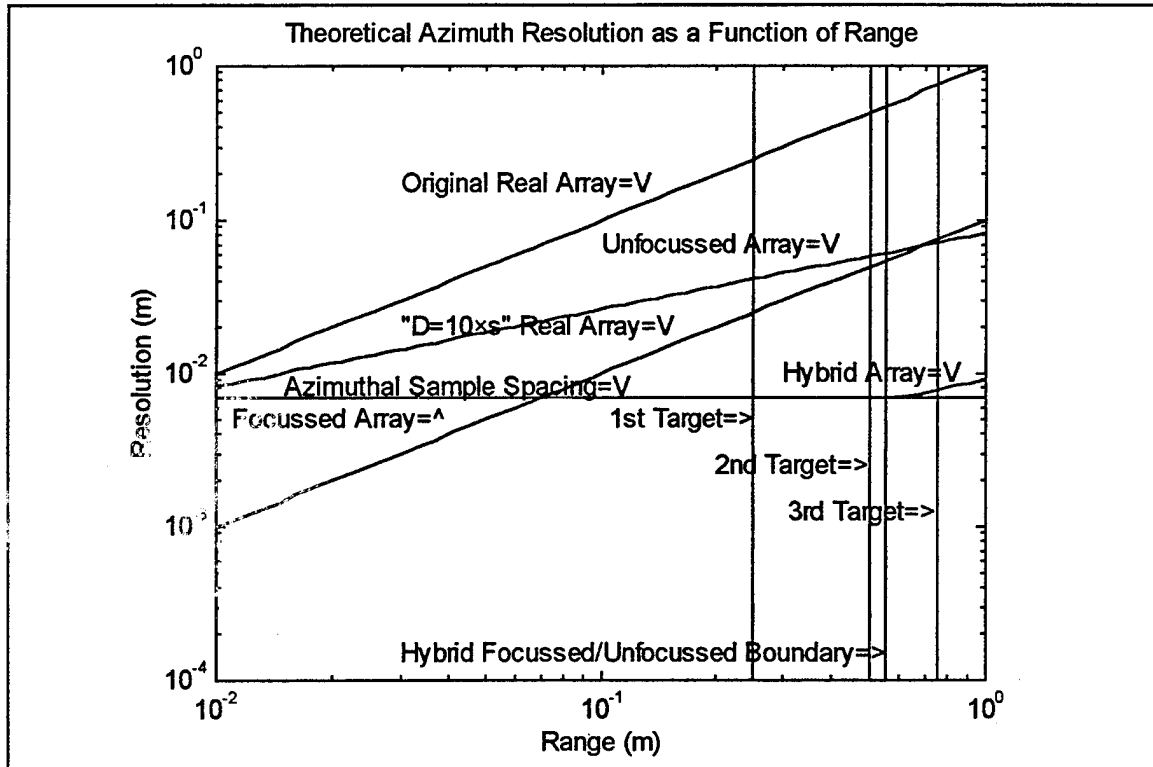
Figure 3.22 Theoretical Azimuth Resolution Capabilities

Examination of Figure 3.22 shows that the theoretical relative performances of the unfocussed and larger real array are nearly the same, as observed from the simulation results. For ranges inside 55cm, the hybrid correlation length is bounded by the beam width; hence, the process is fully focussed. Outside this range, the hybrid correlation length is pulse length limited; so, the procedure behaves in an improved unfocussed manner, reemphasizing its hybrid nature. Azimuth resolutions finer than the azimuthal sample spacing are unobservable. (The azimuthal sample spacing was chosen to coincide with the theoretical fully focussed resolution, D/2.)

Another presupposition to this entire discussion of the hybrid technique is that the applicable constant range slice is available to be used for the correlation. Curlander (1991, pp.167-169), from whose book much of the background of this paper was derived, provides a quite elegant method for calculating the form of this slice involving quadratic approximations to the theoretical Doppler shifts to be seen over the azimuth variations within a single range bin. While this closed form approximation development is intellectually satisfying, in this age of expanded computer and simulation capability, the errors introduced by the approximations are unnecessary. (Most references approximate the hyperbolic phase correction to the correlation template by a parabolic [quadratic]

38

estimate.) Additionally, these closed form approximations provide no easy way to account for the many factors which affect this cross section. Among them are the carrier frequency, sample frequency, discrete range sample spacing, discrete azimuth sample spacing, pulse shape, pulse length, beam pattern, and range. The technique used in this thesis to derive this cut was based on simulation rather than an attempt to analytically define this correlation template. This procedure had the added benefit of also providing the correlation data for two additional fully focussed techniques and a limited two dimensional version of this hybrid, from the same simulation. The basic simulation procedure was as follows:

a. Determine and record the pulse shape and beam pattern experimentally.

b. Determine the applicable initial conditions for the run.

c. Simulate a point target in the center of the acoustic field at the first increment of the range resolution of the final image.

d. Calculate the received signal matrix for a single point target based on the factors listed above.

e. Perform the same range correlation on this simulated received signal matrix as would be performed on the real signal.

f. Record the desired correlation data from the resultant range correlated matrix for the applicable range resolution bin, in this case, the amplitude of the next constant range sample row after the range bin specified above. The next sample row is used because the range resolution increments rarely fall on even signal range samples, and no signal is received until the next sample after the range resolution bin.

g. Simulate a point target in the center of the acoustic field at the next increment of the range resolution of the final image and go back to d. Continue this process until all range resolution bins in the final image have been simulated.

h. Save the results for future comparison / correlation with experimental field data.

The final product of this process for the hybrid technique is a correlation matrix with the following structure:

a. Each row of the matrix corresponds to a particular discrete range at CPA, $R(0)$, an integer multiple of the system range resolution.

b. Each column of the matrix corresponds to a particular azimuthal displacement from the center of the acoustic field, an integer multiple of the separation between along track sample points.

c. The amplitude recorded in each matrix element is the simulated amplitude value for the next integer sample after each R(0) noted at each discrete azimuthal displacement within the acoustic field.

Examination of this correlation matrix confirms many of the assertions made earlier concerning the matrix. Figure 3.23 is the correlation matrix for this hybrid technique constructed to operate on the range correlated data of Figure 3.12.
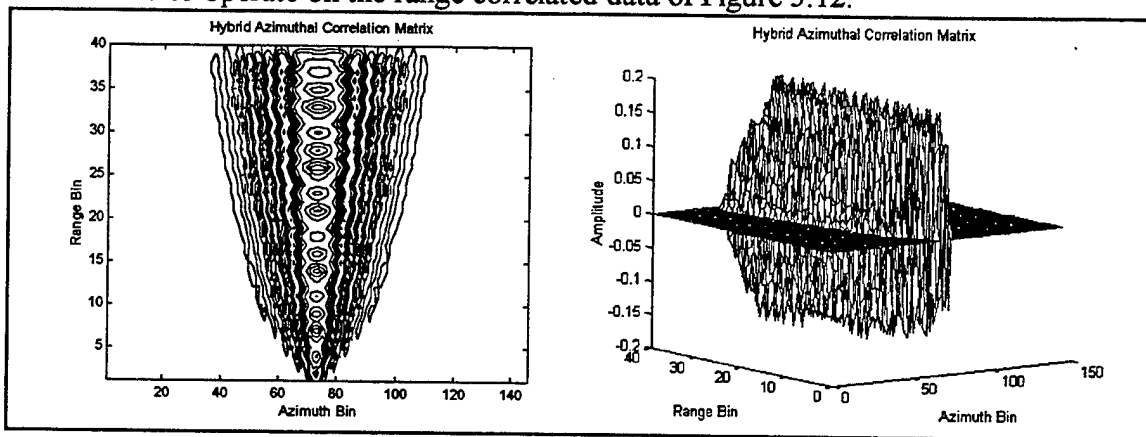
Figure 3.23 Hybrid Azimuthal Correlation Matrix

As can be visually estimated, the correlation length at short ranges is roughly proportional to the range confirming its beam width constraint while at farther distances the correlation length increases as the square root of the range where the pulse length limitation dominates. The individual rows also show the expected form as is demonstrated in Figure 3.24.
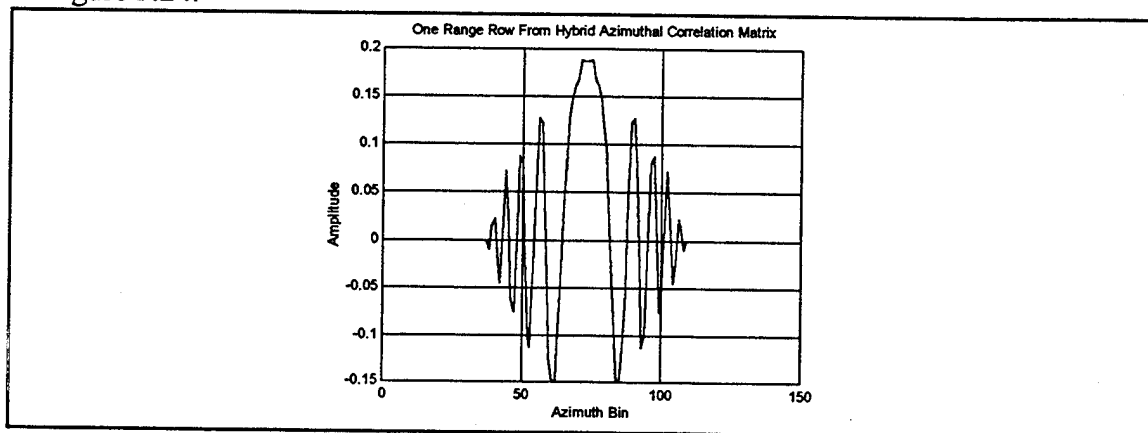
Figure 3.24 One Row From Hybrid Correlation Matrix

An interesting parallel to optical phenomena is that taking the signum function of this result is the Fresnel zone plate for a point source focussed for the given wavelength and range.

While this simulation process is computationally intensive, it can be performed at any time before or after the experimental data run. If this simulation is performed before the data run, images can be produced in near real time by utilizing the following procedure:

a. Read in pre-calculated correlation matrices and data.

b. Perform the range correlation / matched filter procedure on the experimental data as it is recorded.

c. When sufficient azimuth columns are recorded to include a complete acoustic field (based on the real element beam width and maximum range specified), begin the correlation process across the azimuth columns for all range rows using the appropriate row vectors from the simulated correlation matrix as a templates for the range bins being correlated.

The only real time calculations involved in this technique are those of the two correlations which can be performed quickly by FFT's. All other data comes from look up tables. Table look ups are faster than recomputing the analytic form of the filter template for each use (if the form were easily available). With sufficient processing capability, this technique could produce nearly focussed images in real time with only the time lag required to scan the acoustic field for the latest displayed azimuth line.

Additionally, this strategy allows the separation of the physics of the problem from the signal processing aspects. The physics is all incorporated into the simulation; then, the signal processing is handled by recording the results for later comparison with experimental data. Therefore, if the transmitter-receiver configuration is changed by physically offsetting the transmitter from the receiver or installing more receivers, the effects merely need to be simulated, the results recorded, and the signal processing handled by comparing / correlating these records with experimental reality. Hence, the quality of the final image in this algorithm does not depend on the magnitude of higher order terms left out of an approximation, but on the quality of simulation and computer round off error. Utilization of the above described method on the range correlated data of Figure 3.12 produces Figure 3.25.
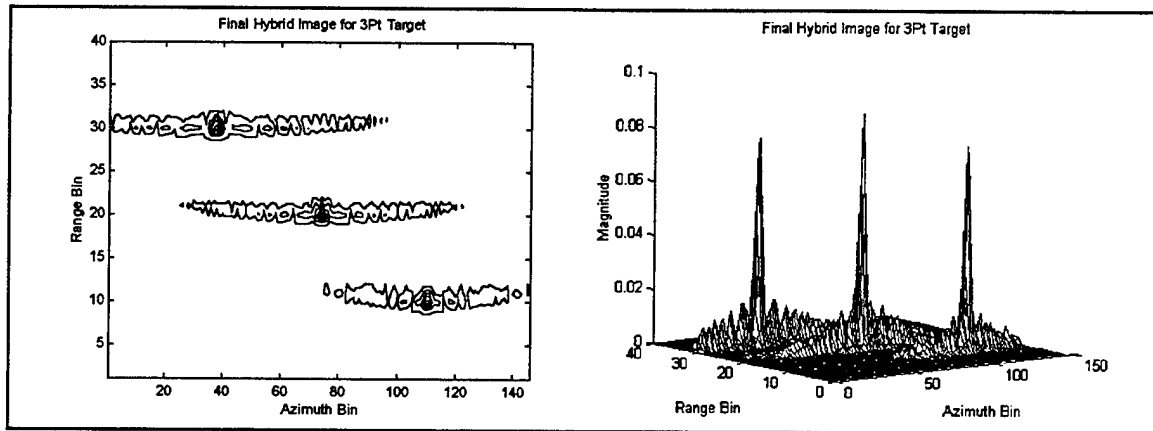
Figure 3.25 Hybrid Synthetic Aperture Final Image

This hybrid technique final image is of higher quality than either the real image of Figure 3.13 using elements ten times as large or the purely unfocussed image of Figure 3.16. As predicted, the azimuth resolution for the farthest target is only slightly worse than for the other two targets.

4. A Fully Focussed One-Dimensional Synthetic Aperture Method

This section is entitled "a technique" because it is only one of many possibilities. (Variants on the hybrid technique described previously are used much more commonly than the method described in this section since because the hybrid works so much faster [Curlander, 1991, pp.167-169].) The aspect of a focussed technique which differentiates it from an unfocussed one, is that in a focussed process, each point in the final image is the result of a correlation of some type of the received data with the hyperbolic phase variation which a target at that point would generate. This focussing for every point and the search for each associated hyperbola - not straight line - are the characteristics which both create high quality final images and require large quantities of processing effort.

To clarify the concept, just as the unfocussed and hybrid techniques could be represented as straight cuts through the range correlated target hyperbola of Figure 3.17, so too a fully focussed method can be illustrated as a cut, though not straight, as Figure 3.26 demonstrates.
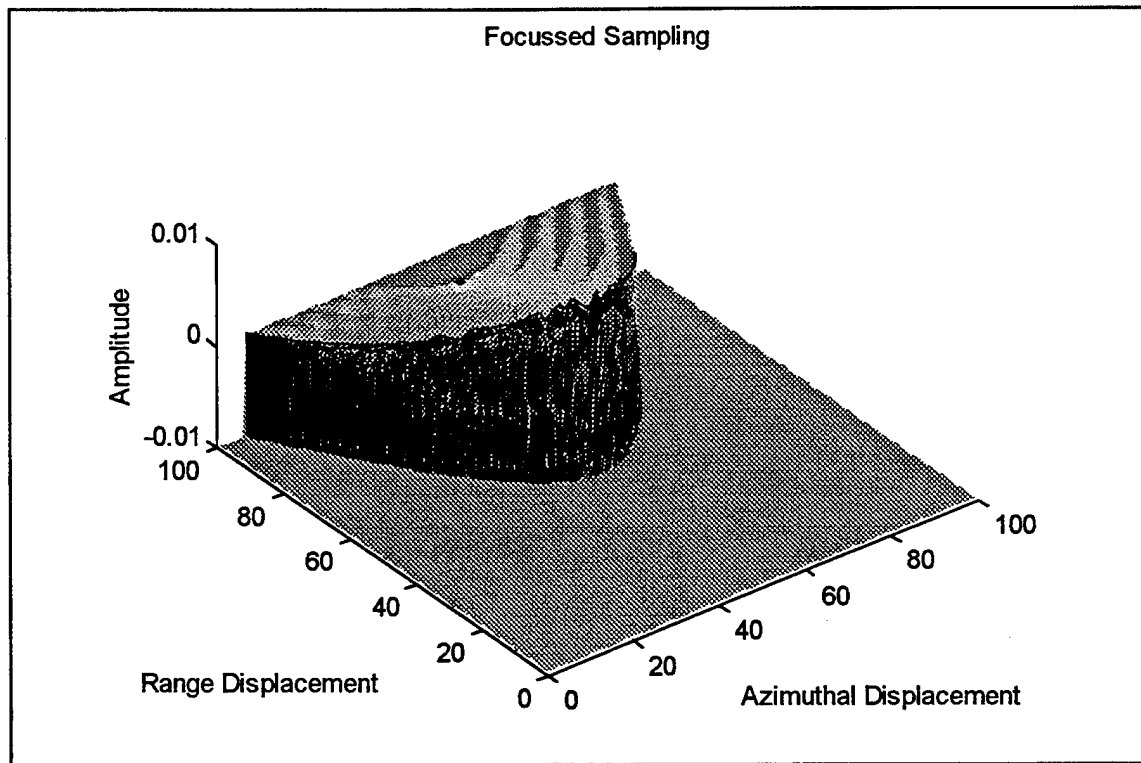
Figure 3.26 Fully Focussed Sampling of Azimuthal Phase Variation

As is apparent from Figure 3.26, the functional form of this correlation slice is greatly simplified by taking the cut along the peak of the hyperbola. It is almost the rectangular weighting function of the unfocussed algorithm. This technique uses a simple correlation function applied to a more complex data curve instead of the hybrid approach's complicated correlation template located along a straight line. Since this data path is along the maximum response curve of the signal, it attains the best possible signal to noise performance.

The first step toward implementing this strategy requires knowledge of the location of the maximum response curves for a target at any range. Equation 3.1 provides this information and Figure 3.27 repeats the basic geometry.
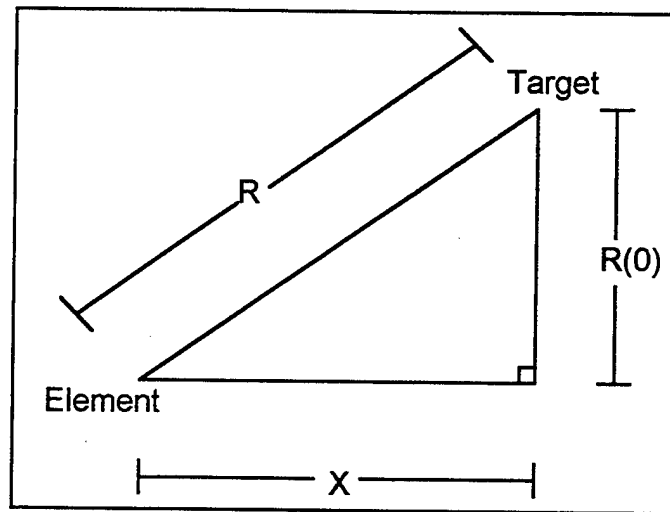
Figure 3.27 Basic Target Geometry

$$R = \sqrt{R(0)^2 + X^2}. \tag{3.1}$$

Again, this equation describes reception range, R, as a hyperbolic function of both range at CPA (closest point of approach), R(0), and azimuth displacement, X. Since the receiving element spacing is known, the azimuthal displacement as a function of element number is known. Likewise, the range resolution of the system based on pulse length is known; so, the possible ranges of interest at CPA are known. Thus, a matrix of reception ranges as a function of number of elements away from CPA and CPA range bin can be calculated. Reception range is directly proportional to propagation delay time which is directly proportional to sample index in the range correlated signal matrix. This sample index matrix was assembled in the following manner:

1. Each row vector in the matrix corresponds to a particular range at CPA [R(0)].
2. Each column vector in the matrix corresponds to a discrete number of element displacements relative to the center element where the delay corresponds to R(0).
3. The value of each element in a particular row vector is the unrounded sample delay index in the range correlated matrix which would have the maximum response for a point target at range R(0) displaced in azimuth by X element spacings from the center reference.

In graphical form, this sample index matrix looks like a cone with its apex at the zero range to the center element as Figure 3.28 illustrates.
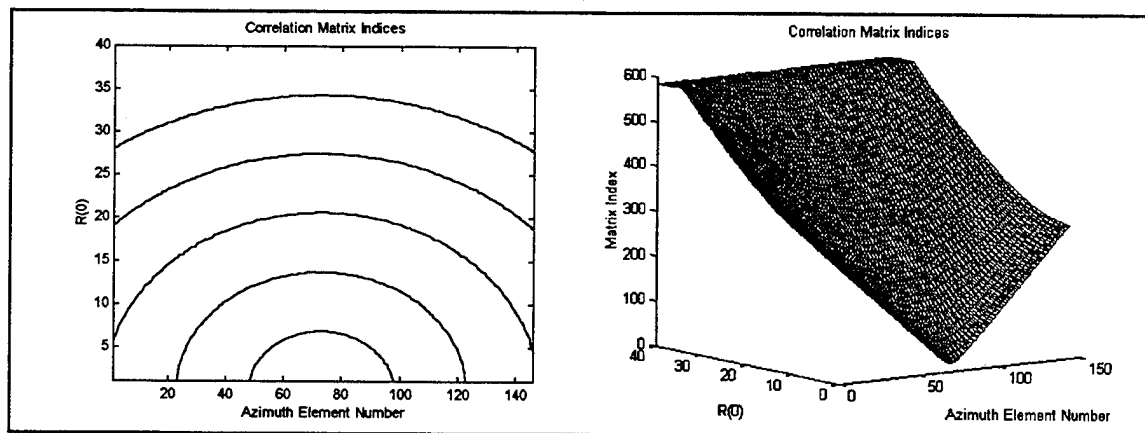
Figure 3.28 Correlation Matrix Indices

This matrix can then be used to extract focussed data from the range correlated matrix. When the center or reference element of the index matrix is aligned with the azimuthal displacement of a target, the conic / hyperbolic selection of data from the correlation matrix serves to re-map the data into a straight horizontal line - ready for azimuth correlation. Target hyperbolas located at azimuthal displacements other than this reference are still straightened but at an angle causing their signal to be distributed among several range bins. Figure 3.29 illustrates this effect on the previous range correlated matrix with the reference element set for the center target.
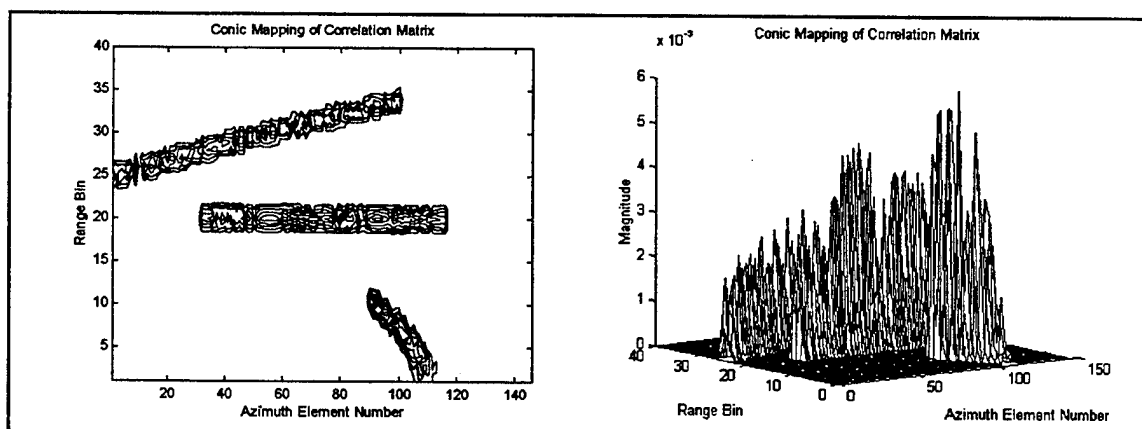
Figure 3.29 Conic Mapping of the Three Point Target Range Correlated Matrix

45

The effect of this mapping could be conceptualized as the same as reflecting the range correlated matrix off of a conic mirror onto an image plate. Because the conic mirror has a different curvature at each range, each range bin is individually focussed. As the mirror is swept across each azimuth line, each point in the image is then individually focussed. This optical analog is highly appropriate since much of the initial synthetic aperture radar data was processed using optical techniques similar to those described here.

Since the range correlated data can be focussed using this conic correlation index matrix, then the only step remaining is to find the azimuthal correlation template to be used to create the final image at each point. The method proposed here is the same as was used for the hybrid technique: pick the appropriate data from out of a simulation of the target at the range in question. The principle difference between this method and the hybrid procedure is that in the hybrid case the data was recorded from a straight, constant range line, where here, the data is retrieved from the locations specified by the conic correlation index matrix for each azimuth line. The form and function of the resultant azimuthal correlation matrix are the same except that the matrix entries for each range bin are now taken from the associated hyperbola vice a straight, constant range line. The result of this simulation and data extraction process are shown in Figure 3.30.
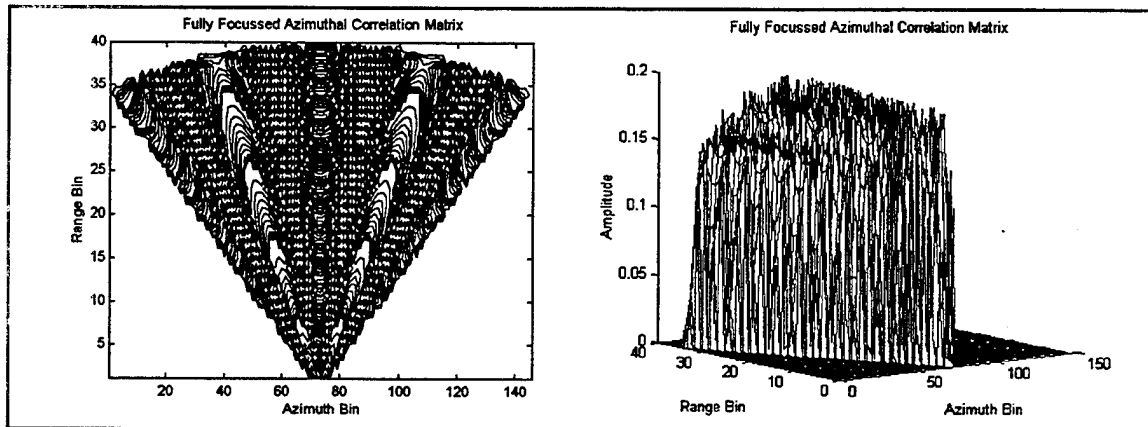
Figure 3.30 Fully Focussed Azimuthal Correlation Matrix

The one radian "pie shape" of the matrix response corresponds to the one radian beam width of the simulated illumination. This explains why the short range target had such a narrow azimuth response in the received and correlation signal matrices relative to the longer range targets - it was only in the beam for a small number of elemental azimuth displacements. The amplitude of the correlation matrix is strongly positive indicating that the samples were indeed taken at or just after the range correlated peak. The amplitudes fall off by $R(0)^2/R^2$ as the azimuth departs from the reference centerline.

46

To build a fully focussed optimum signal to noise ratio (SNR) algorithm, the following steps were executed:

a. Read in pre-calculated correlation matrices and data.

b. Read in the experimental received signal matrix.

c. Perform the range correlation / matched filter procedure on the received signal.

d. Start at the first azimuth line.

e. Create a data matrix with elements extracted from the positions in the range correlated data matrix indicated by the conic index matrix referenced to this azimuth line.

f. Create this azimuth line of the final image by performing inner products (one point correlations) between the rows of this data matrix and the corresponding rows of the fully focussed azimuthal correlation matrix. The resultant single column of these one point correlations, one element for each range bin, is the final image line for this azimuthal displacement.

g. Go to the next azimuth line and resume the procedure at step e. Continue until the full image is formed.

The application of this procedure, one azimuth line at a time, to the range correlated signal matrix of Figure 3.12 using data elements as indicated by the conic index matrix of Figure 3.28 one point correlated with the azimuthal correlation matrix of Figure 3.30 produces the image of Figure 3.31.
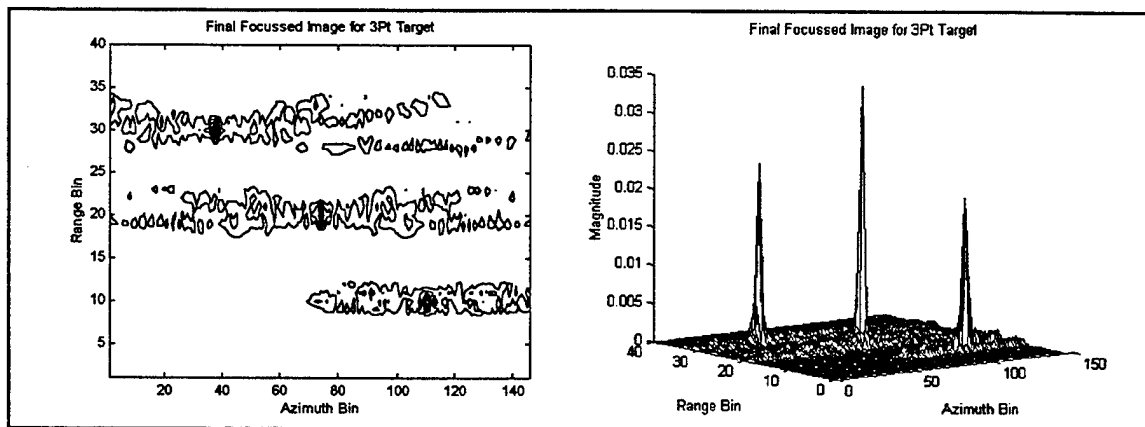
Figure 3.31 One Dimensional Fully Focussed Image

47

As can be seen in Figure 3.31, all three points are equally resolved with the best image quality seen to this point. The reduction in background noise sensitivity over the hybrid image of Figure 3.25 is also apparent.

5. A Fully Focussed Limited Two-Dimensional Synthetic Aperture Method

The fully focussed system of the last section represents the theoretical optimum in azimuthal resolution available, D/2. This resolution follows from its use of the entire real element beamwidth for an azimuthal correlation length as occurred in the hybrid case when the pulse length was long enough to produce a fully focussed image. As a review, this effect comes from Equation 3.13,

$$\Delta X = \Delta \phi / 2\pi \times \lambda / 2 \times R / X, \qquad (3.13)$$

which when used with $\Delta \phi = 2\pi$, $R \cong R(0)$, and $X \cong R(0) \times \lambda / D$ from Equation 3.9 produces

$$\Delta X = 2\pi / 2\pi \times \lambda / 2 \times R / (R(0) \times \lambda / D),$$

which simplifies to

$$\Delta X \cong D/2. \qquad (3.21)$$

Unlike the hybrid technique, the fully focussed method uses the best single data point from each azimuth line in the beam to assemble a target image. Since there are no more azimuth lines to inject more information into the process, this is the best azimuth resolution possible. Even the full two-dimensional correlation technique mentioned and rejected on practical grounds in Section III. A. 2. could provide no better azimuth resolution than this.

But, if the investment in computer power was made to enact that full two-dimensional correlation, the gain that would be made would occur in terms of signal to noise ratio (SNR). For the consecutive one-dimensional correlation techniques (one range correlation followed by one azimuth correlation of some form) discussed to this point, shadows or noisy interference could disrupt the single selected correlation line of a point target and noticeably degrade its image. In a two-dimensional process, more than just one curve contributes to the one point correlation; hence, loss of that single data line will not significantly degrade the image.

Since a full two-dimensional correlation is impractical, possibly a lesser form could provide some of the benefits of a full two-dimensional matched filter without all of the computational cost. The greatest magnitude of target response is in the vicinity of the

maximum response hyperbola. Using the two sample hyperbolas just before and the two hyperbolas just after the peak response in addition to the maximum response hyperbola itself provides the greatest return on calculational investment. Since all of these points are of similar magnitude, multiplying by five the number of points included in the one point correlation template, multiplies the amplitude of the point target response by an approximate factor of five. Noise is also increased; but since it adds incoherently or in quadrature, the noise level rises by the smaller factor of the square root of five. This raises the signal to noise amplitude ratio by a factor of the square root of five and the signal to noise power ratio by a factor of five. This represents a theoretical 7dB improvement in SNR. If greater noise immunity is required, more hyperbolas could be sampled at the obvious expense of more computational workload.

This limited two-dimensional strategy is implemented in this paper in the same manner as the fully focussed one-dimensional technique. The other four hyperbolas taken from the two range correlation matrix samples before and the two samples after the one-dimensional hyperbola form four additional correlation matrices. In the final image, the five single row one point correlations are summed together to create the single point response. Figure 3.32 graphically describes the signal topology for use in this technique.

Figure 3.32 Limited 2-D Focussed Technique

The one-dimensional focussed algorithm extracts data from the simulation and performs the one point correlation (inner product) with experimental data only along the maximum response curve as illustrated in Figure 3.32 above. The limited two-dimensional fully focussed algorithm, in addition to this correlation, executes one point correlations along the other four neighboring curves as well. The resultant amplitude of the image for the range bin in question is the sum of these five one point correlations. Thus, the target

49

signal along all five hyperbolas contributes to the final image. Hence, the limited two-dimensional fully focussed algorithm uses a two-dimensional correlation band as wide as the beam width in azimuthal extent, five samples long in range extent, and shaped to be centered in range on the maximum response curve of the point target. This is more efficient than the full two-dimensional correlation because it does not include areas where the signal is weak or not present. It is not as thorough as the full two-dimensional correlation because it ignores some of the energy received from the target outside the limited correlation band. Figure 3.33 shows the results of this approach.
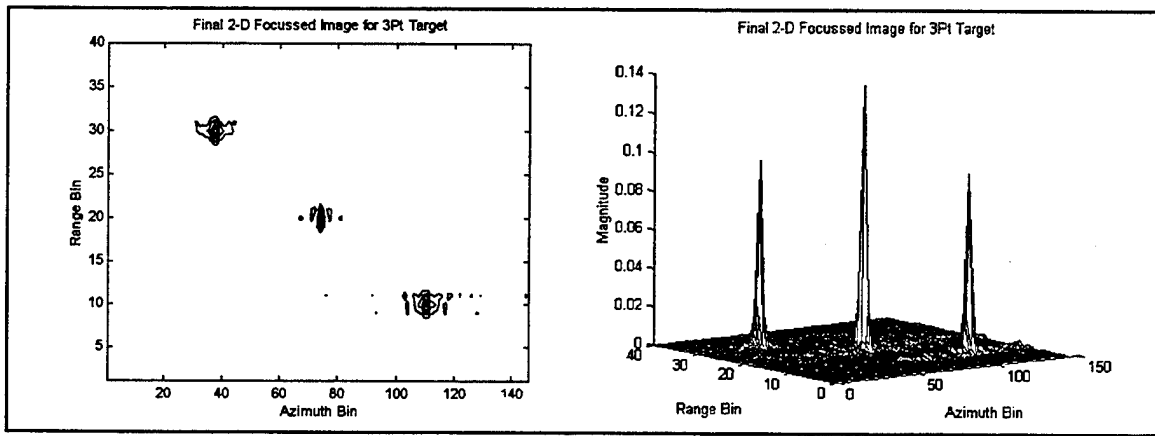
Figure 3.33 Limited Two-Dimensional Fully Focussed Image

When compared with Figure 3.31, it is apparent that the two images have identical resolution. However, Figure 3.33 displays much less background return for identical conditions. While the simulation contained no postulated noise, this lowered off peak response represents lower side lobe reaction, hence better noise rejection. Thus this technique provides a method to use processing time to improve SNR.

6. A Limited Two-Dimensional Hybrid Synthetic Aperture Method

As will be shown later, the fully focussed limited two-dimensional method, while providing the best possible resolution capability and best reasonable noise rejection capability for a moderate computational cost, is still very expensive in terms of computer time. However, the strategy used in going from the one line fully focussed technique to the limited two-dimensional five line fully focussed technique to reduce noise susceptibility by 7dB at the expense of additional computational time might be successfully applied

50

elsewhere. The hybrid technique, while demonstrating nearly the same resolution as the fully focussed algorithm in a fraction of the computer run time, suffered from poor noise sensitivity. This method of performing five one-dimensional azimuth correlations in the vicinity of the target's maximum response curve and coherently summing the results to make a limited two-dimensional correlation is easily applied to the hybrid technique. Fitch, in his book on synthetic aperture radar (1988, pp. 66-67), discusses a frequency domain version of this same concept. Figure 3.34 illustrates the signal topology used.
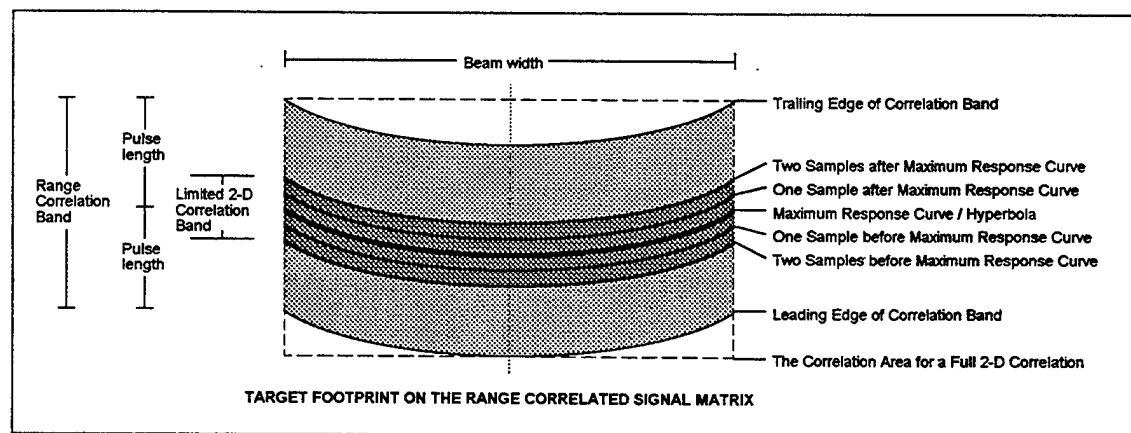

Figure 3.34 Limited 2-D Hybrid Technique

The one-dimensional hybrid algorithm extracts data from the simulation and performs the azimuthal correlation with experimental data only along the first hybrid correlation line as illustrated in Figure 3.34 above. The limited two-dimensional hybrid algorithm, in addition to this correlation, executes correlations along the next four sample lines as well. The next four sample lines are used instead of two before and two after, as was used for the fully focussed case, because, this strategy better covers the maximum response curve of the target. The resultant amplitude of the image for the range bin in question is the sum of these five correlations. Thus, the target signal along all five constant range lines contributes to the final image. Hence, the limited two-dimensional hybrid algorithm uses a two-dimensional correlation band as wide as the beam width in azimuthal extent, five samples long in range extent, and now shaped like a rectangle to be more computationally efficient.

This is still more efficient than the full two-dimensional correlation because it does not include as much area where the signal is weak or not present. It is still not as thorough as the full two-dimensional correlation because it ignores some of the energy received from the target outside the limited correlation band. It does not have quite the signal to

noise resistance or the long range fully focussed performance of the fully focussed limited two-dimensional correlation because its correlation band does not cover exclusively the maximum response zone as does that algorithm. However, this hybrid two-dimensional technique runs more than an order of magnitude faster than its fully focussed counterpart, provides better signal to noise operation than the one line focussed procedure, and exhibits fully focussed resolution over most of the acoustic field tested. Figure 3.35 shows the results of this approach.
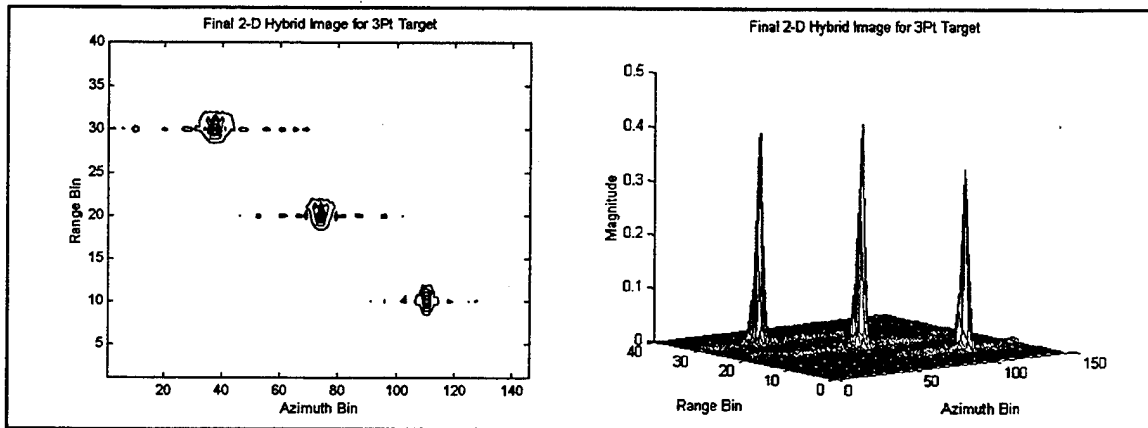

Figure 3.35 Limited Two-Dimensional Hybrid Image

When compared with Figure 3.33, it is apparent that the two images have nearly identical resolution and similar noise rejection capabilities. However, Figure 3.35 took only about 1/30th of the processing time. Thus this technique provides an efficient method to use processing time to improve SNR while maintaining acceptable resolution capability.

7. Image Degradation due to Factors Other than Noise

Up to this point, resolution and noise rejection have been the only factors considered in the pursuit of a good target image. But there are other factors which can degrade image quality.

Other limitations on system performance arise when the initial assumptions of Chapter II are violated. First, if the locations of the sensing elements are not evenly spaced on a straight line and the effects are not corrected, phase errors are introduced which deteriorate the quality of the coherently produced image.

Second, if the target is moving while imaged, any sensed Doppler shift other than that due to the motion of the array will displace the target image in the azimuthal dimension. Figure 3.36 depicts the geometry.
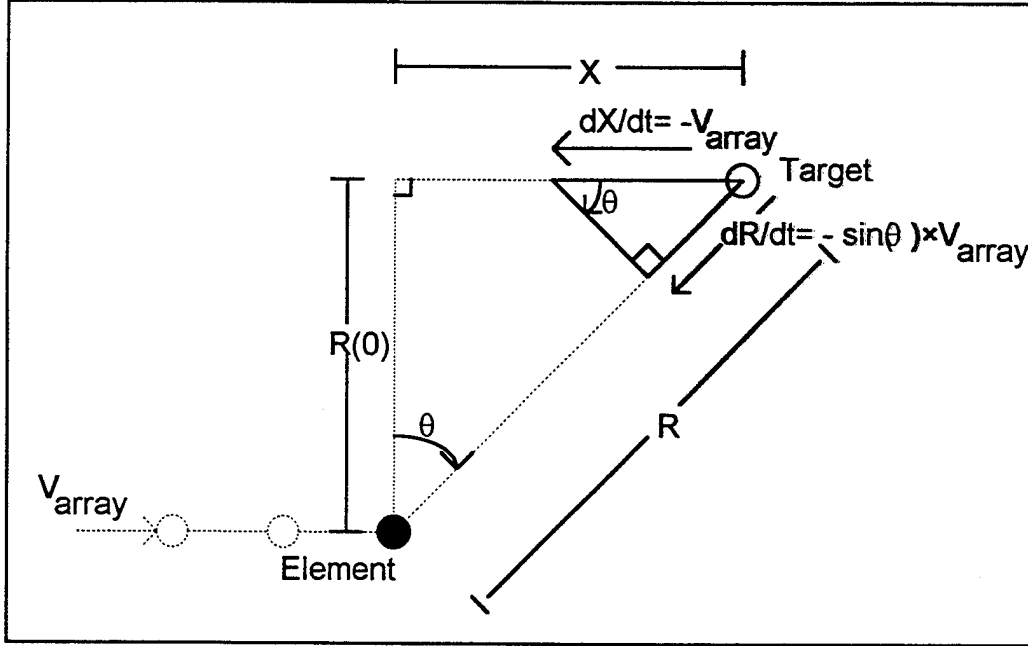
Figure 3.36 Stationary Target Geometry

The change in range, ΔR, is sensed by its phase effects, i.e.

$$\Delta R = \Delta\phi/2\pi \times \lambda/2.$$

In this case, the motions and phases are best tracked according to their time derivatives by the relation

$$dR/dt = d(\Delta\phi/2\pi)/dt \times \lambda/2 = f_D \times \lambda/2, \qquad (3.22)$$

where "f_D" is the sensed Doppler frequency and "dR/dt" is the target range rate.

The sensed azimuthal displacement, X, is given by

$$X = (R \times dR/dt) / (dX/dt),$$
$$\text{or}$$
$$X = R \times f_D \times \lambda/2 / -V_{array}, \qquad (3.23)$$

53

where "R" is measured by the time delay of the return pulse; "f_D" is the measured frequency variation from the transmitted carrier; and "λ" and "V_{array}" are known.

If the measured Doppler frequency varies from that calculated above because of target motion toward or away from the array, it will generate a Doppler frequency error according to:

$$f_{Derror} = 2/\lambda \times dR_{tgt}/dt. \qquad (3.24)$$

Applying Equation 3.24 to Equation 3.23 generates an azimuthal displacement error of

$$X_{error} = R \times f_{Derror} \times \lambda/2 / -V_{array},$$
$$\text{or}$$
$$X_{error} = (R \times dR_{tgt}/dt) / -V_{array}. \qquad (3.25)$$

If the target is rotating, dR_{tgt}/dt for any portion of the target is roughly proportional to its distance from the axis of rotation. Hence, the image of a rotating target is expanded in the azimuthal dimension. This is the principle effect sought by inverse synthetic aperture processing, where the motion of the target is used to resolve and enhance its image. For the application of this paper, this effect contributes only a distortion which degrades the final image quality.

The third violated assumption that can limit system performance is that of an isotropic medium. Variations in speed of energy propagation or path length as the pulse travels to the target and back again can produce phase variations which will degrade the final image quality if not corrected.

Image degradations caused by the first and third of these effects can be ameliorated somewhat by the use of adaptive phase compensation methods beyond the scope of this paper. The inverse synthetic aperture effect is part of the physics of the problem and can only be eliminated by a cooperative, stationary target.

IV. CONSTRUCTION

All of the theory and simulation described in Chapters II and III of this paper are little more than an academic exercise unless they can be applied to a physical result. This Chapter discusses the preparatory steps taken to create a working air medium synthetic aperture sonar. This was done initially in air to prove the concept at a very simplistic level and set the stage for more complex implementations after the basic effects have been fully investigated.

A. PHYSICAL APPARATUS

The first components of an air medium synthetic aperture sonar to be identified were the acoustic transmitters and receivers. These were the bridge between the acoustic reality to be discovered and the electronic equipment used to image that condition. The choice of these elements was governed by several considerations. First, it was desirable to operate with a carrier frequency in the vicinity of 20KHz to achieve moderate range and azimuth resolutions with a reasonable sample rate. Additionally, operating near this frequency range will provide experience helpful for the transition to water. Secondly, the transmitter-receiver combination needed a wide but finite beam width to optimize the performance of the synthetic aperture system. Third, in anticipation of possible future frequency chirping of the transmitted signal for greater range resolution, the frequency response of the combination needed to be relatively flat in the vicinity of the final carrier frequency.

The acoustic transmitter chosen was a SONY Walkman earphone. This provided a very level response from 20KHz to 30KHz. It also showed a beamwidth of about one radian at the selected operating frequency of 25KHz for an inexpensive price. The acoustic receiver chosen was a Larson-Davis ¼" wide band microphone. This microphone had a flat amplitude response from less than 1KHz to greater than 100KHz. It was also omnidirectional to a first approximation in the frequency range of interest. So, all performance limitations were essentially dictated by the transmitter which met all of the required specifications.

The next component of the physical apparatus to be determined was the means by which the speaker and microphone would be positioned to obtain their signal returns. For

this test case, a one meter by one meter square was chosen as the acoustic field of interest. To survey this region the means of transport had to meet certain specifications:

1. Minimum frame travel:

azimuthal correlation length+image width=

max range×2×sin(beam width/2)+1.0m≅2.0m.

(This requirement was later relaxed to 1.1m by zero padding the received signal matrix. This degraded the image slightly outside the center 10cm of azimuth but gained portability for the device.)

2. Maximum frame "wobble" or curvature:

Using the plane wave approximation:

$\lambda/16$=vsound/(fcarrier×16)=.86mm.

3. Maximum separation between microphone and projector vertical center lines:

Using plane wave approximation for minimum image range

$\cong \frac{1}{2}\times\sqrt{R}\times\lambda/2$=4.6mm.

Figure 4.1 shows the track and trolley system built to these specifications
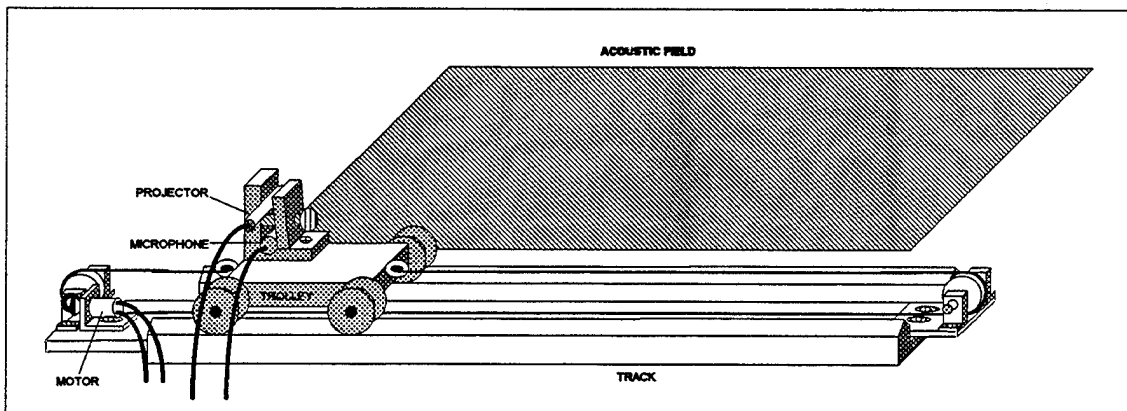
Figure 4.1 Track and Trolley System

After construction, it was noted that the motor and gear assembly were excessively noisy; however, signal processing techniques were sufficient to gain satisfactory results in spite of this interference.

B. ELECTRONIC EQUIPMENT

With the selection of transducers and the construction of the locomotive device complete, the next task was the addition of appropriate signal generation, conditioning and

56

recording equipment to make an operational system. Figure 4.2 diagrams the basic set up used.
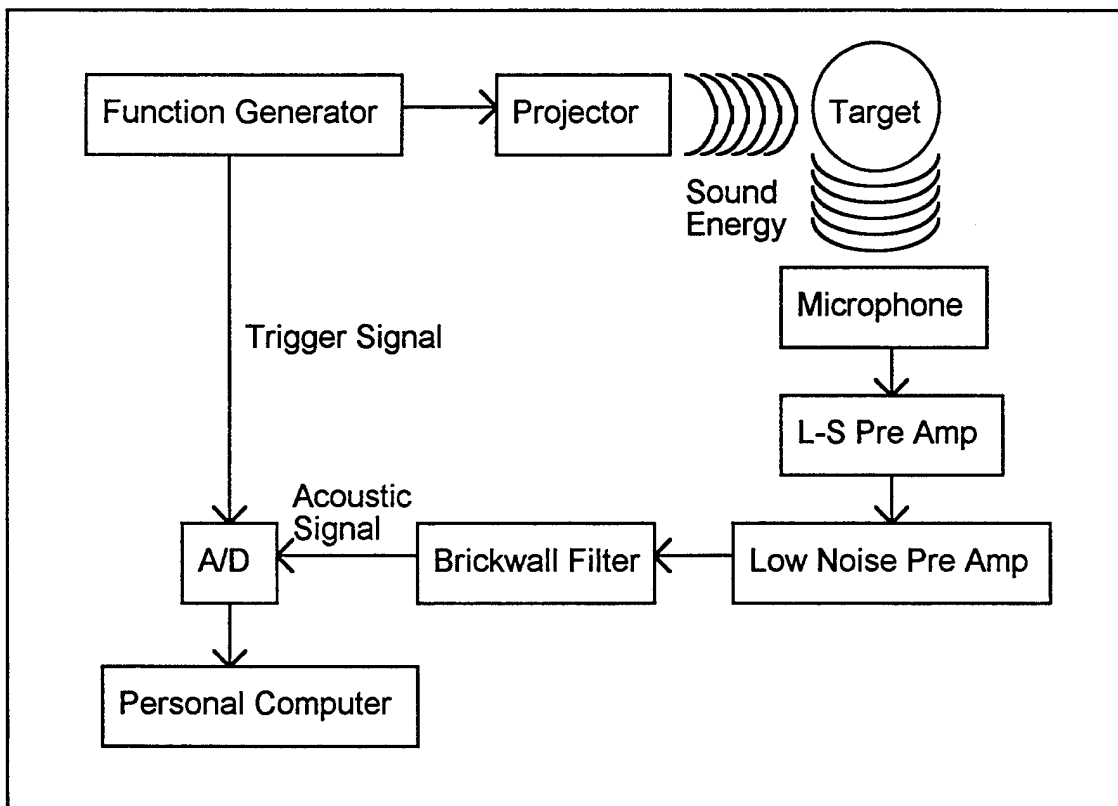
Figure 4.2 Synthetic Aperture Sonar Equipment Configuration

The specific equipment and settings used were as follows:

   A. Function generator
         Hewlett Packard Model 33120A 15MHz Function/Arbitrary Waveform
Generator
         Settings:
             ·   25kHz carrier frequency
                 5 cycle burst mode
                 13.43 Hz Pulse Repetition Rate
                 0.1, 1,10 Volt outputs on various data runs
   B. Projector / source
         SONY "Walkman" type earphone
   C. Microphone / receiver
         Larson-Davis ¼" Wide Band Microphone

57

D. L-S Preamplifier

    Larson-Davis Model 2200C Amplifier

    Settings:

        Channel 2

        Gain: 40dB

        Bias: 200mV

E. Low Noise Preamplifier

    Stanford Research Systems Model SR560 Low Noise Amplifier

    Settings:

        Gain mode: High Dynamic Reserve

        Gain: 2×10^2

        Coupling: AC

        Filter Cutoff (pass band): 10K-100KHz

        High pass filter roll off: 6dB/octave

        Low pass filter roll off: 6dB/octave

        Power: Line

F. Brickwall Filter

    Wavetek Brickwall Filter Model 753A

    Settings:

        High Pass: Gain=0dB; 15KHz Cutoff

        Low Pass: Gain=0dB; 50KHz Cutoff

G. A/D

    ADAC Model 5801MF 12 bit, 333KHz Data Acquisition Board

These filter settings were chosen to minimize extraneous noise contamination of the acoustic signal without altering the basic signal wave form (which would degrade the performance of the range correlation matched filter in the computer software portion of the system).

Later versions of this configuration may replace the function generator with the D/A portion of the A/D card presently in use to both reduce the amount of equipment necessary and to gain flexibility in the choice of output wave form.

C. COMPUTER CODE

Once the acoustic signal was passed by the A/D into the personal computer, it was the job of the software to process that data into a recognizable image. In this paper, the bulk of the programing was done in MATLAB because of the relative ease of code generation in that language and its speed at handling vectorized data. The A/D programming was done in TEST POINT based on the A/D card manufacturer's recommendation.

1. Preparatory Code

As described in the last chapter, prior to the actual data run, a simulation of targets at all range resolution bins needed to be run to assemble the correlation matrices. In order to run these simulations as accurately as possible, the beam pattern and pulse shape for the transmitter-receiver combination had to be available to the simulation.

a. Beam Pattern

First, to obtain the beam patterns of the projector and microphone, the configuration of Figure 4.3 was used.
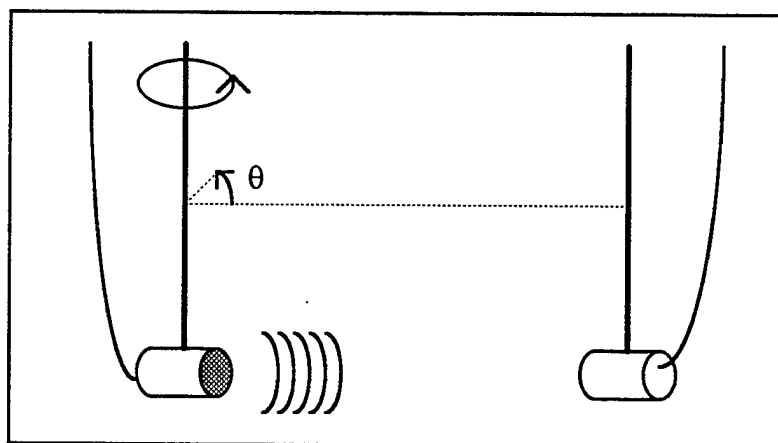

Figure 4.3 Measurement of Beam Pattern Arrangement

This measurement was conducted in an anechoic chamber to avoid contamination of the data by reflections. After the magnitude of source strength as a function of angle, θ, for the projector was recorded, the positions of the microphone and

transmitter were switched and the magnitude of response as a function of angle for the microphone was recorded in the same manner. The combined beam pattern as a function of θ was merely the product of these two patterns. However, the data form to be used in the later programs used this beam pattern as a function of $\sin(\theta)$ not θ. Therefore, this data set needed to be resampled by interpolation into the proper form.

To do this, the range of vector indices was used as follows. The old indices corresponded to a range of $-\pi/2$ to $\pi/2$. The new indices would correspond to a range of -1 to 1. Since the indices could fall only in the range 1 to the vector length (a MATLAB convention), after accounting for the shifts to eliminate non positive indices, the new indices could be extracted from the old data using the following formula:

$$\cos\left(\frac{[\text{old index-1}]\times\pi}{[\text{old length -1}]}\right) = 1 - \left(\frac{[\text{new index -1}]\times2}{[\text{new length -1}]}\right). \tag{4.1}$$

Since the new indices would correspond to fractional indices in the old vector, interpolation as described in Figure 4.4 was used.
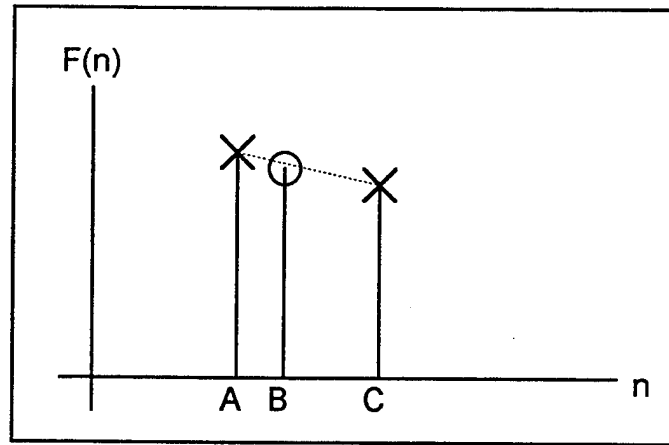

Figure 4.4 Linear Interpolation Method

Since F(A) and F(C) were known, F(B) was desired, and C-A=1,

$$F(B) \cong [\text{C-B}]\times F(A) + [\text{B-A}]\times F(C). \tag{4.2}$$

MATLAB handled this interpolation method very quickly in vector form; hence, it was used repeatedly throughout the code used in this paper. The results of these techniques applied to the data for the projector-microphone combination used here is shown in Figure 4.5.
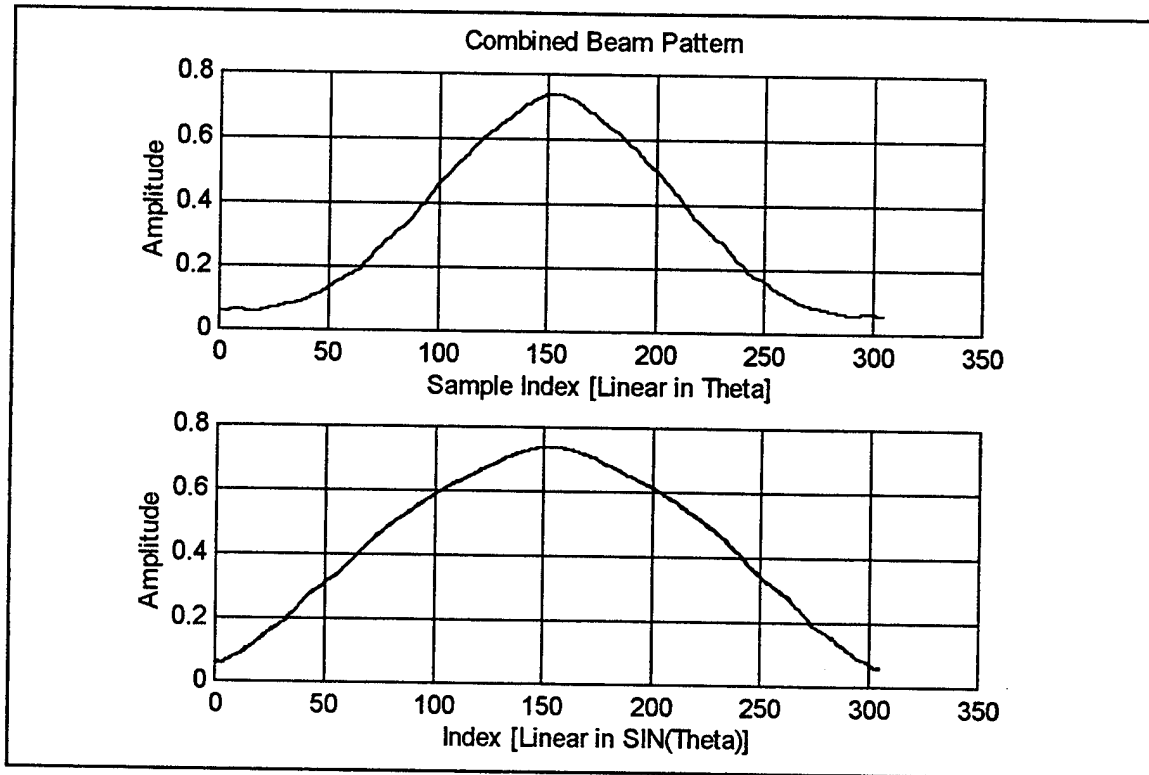
60

Figure 4.5 Combined Beam Pattern

b. Pulse Shape

With the spatial variation of the signal quantified as above, the next step was to evaluate its temporal variation. This temporal wave form would be used for the range correlation template; so, it had to match the actual received target signal as closely as possible in both phase and amplitude. This wave form was obtained using the same configuration as in Figure 4.3 except that the two transducers were pointed directly at one another and only a single pulse at the desired sample frequency was recorded. The sample frequency chosen was 100KHz to be four times the carrier frequency in anticipation of possible future adaptive phase compensation, which requires sampling in quadrature, and greater ease in early interpretation of the raw data since consecutive sample were always $\pi/2$ apart.

The carrier frequency was known; so, only a phase shift needed to be applied to make the zero crossings correct. A more interesting challenge was the extraction of the envelope. This was obtained using Hilbert Transforms in accordance with the following procedure:

1. Take the original sampled pulse.

pulse = modulation × carrier = mod×cos(ω_ct)

2. Take the original pulse's Hilbert Transform.

Hilbert = modulation × carrier all phase shifted 90°

= mod×sin(ω_ct)

3. Add the pulse to i×Hilbert.

sum = mod×cos(ω_ct) + i×mod×sin(ω_ct)

= mod ×[cos(ω_ct) + i×sin(ω_ct)]

= mod × exp(i ω_ct)

4. Demodulate the sum by mixing with exp(- i ω_ct)

result = mod × exp(i ω_ct) ×exp(- i ω_ct) = mod

This method was used because it gave a smoother (more realistic) pulse shape than simply demodulating the pulse with the carrier. It also showed phase reversals in the envelope that taking the magnitude of the sum of the pulse and the Hilbert Transform would not. Figure 4.6 demonstrates the results of these efforts using sin(2π25KHz×t +π/4) as the carrier.
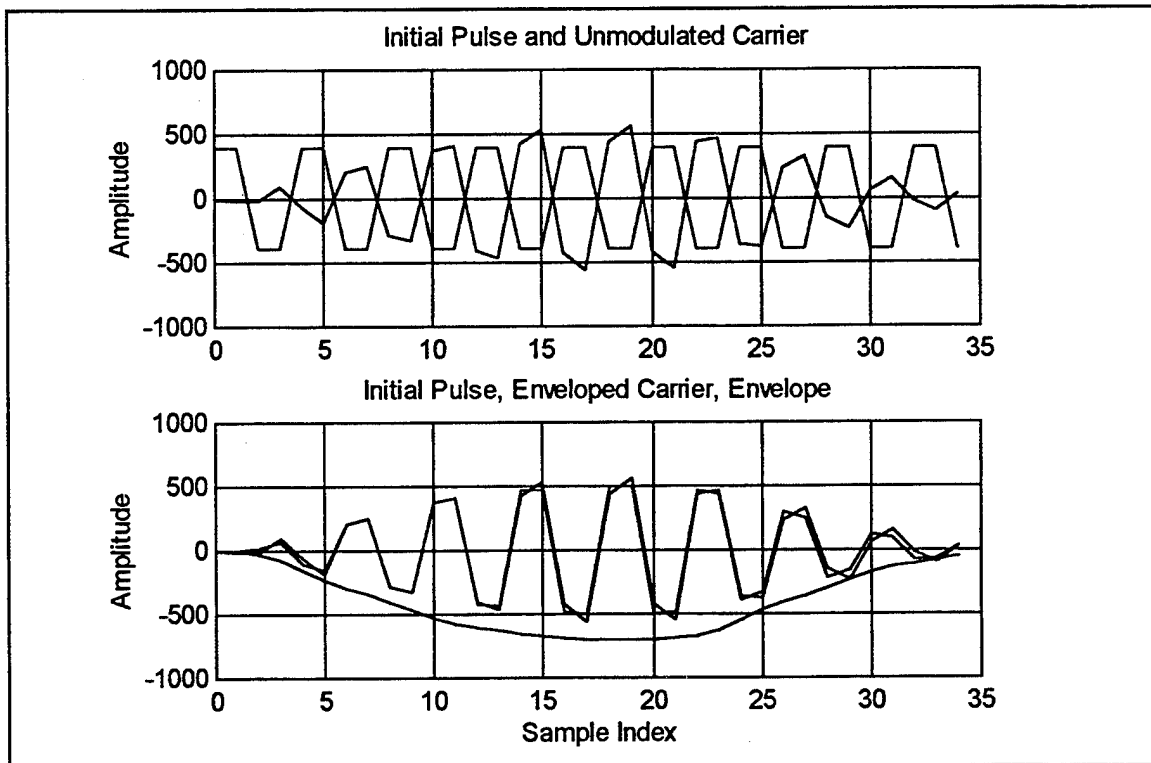
Figure 4.6 Pulse Shape and Envelope

As can be readily seen, with this carrier, the envelope was entirely negative.

c. Simulation

With the spatial and temporal characteristics of the basic pulse defined, the simulation was ready to proceed. Figure 4.7 provides an overview of the algorithm.

Figure 4.7 Simulation Algorithm

The implementing code for this algorithm is found in Appendix A. It was written in MATLAB version 4.2C for personal computer use. The program itself took 11 minutes and 43 seconds to run on a 486DX2 personal computer with 16Mbytes of RAM. Since this could be run at any time prior to the actual data run, this time expenditure was not a concern.

2. Execution Code

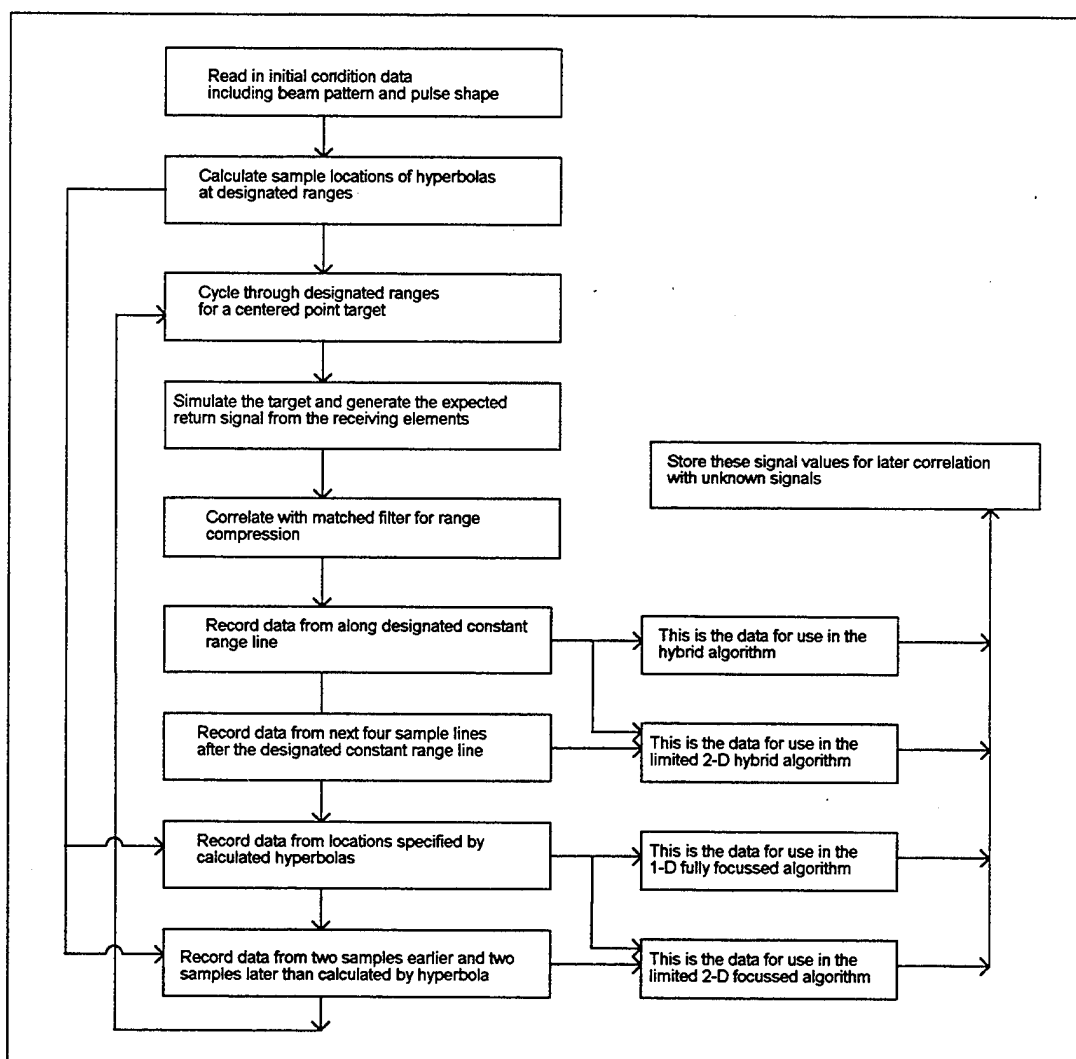With the physical apparatus assembled, the equipment connected, and the correlation matrices created, the system was ready to acquire data and form it into an image.

a. Analog to Digital Board Instructions

The first real time step in creating an image from the acoustic signal received at the microphone was to convert that signal into a digital form that the MATLAB program could manipulate. This was done by the A/D with the instructions of Figure 4.8.

Step	Action	Object	Parameters
1	Open	File 1	c:\temp\test1.dat
2	Set A/D trigger digital	A/D1	
3	Linear series	Loop 1	from 1 to 1127 step by 1
4	Acquire A/D	A/D1	#samples=584, rate=100000Hz, channel(s)=0; Raw data mode with A/D1, term.=none
5	Output to	File 1	
6	End	Loop1	
7	Close	File 1	

Figure 4.8 Test Point Program

This program told the A/D board to sample at 100KHz as discussed previously. The 584 samples corresponded to the number required for 1 meter of range coverage. With the pulse repetition rate selected, the azimuthal element spacing (dictated by D/2), an azimuthal track length of 1.1m, and the speed of the trolley set, 1127 data pulse sets gave 7 pulses per each of the 161 elements for a total data block size of 1.3Mbytes- just small enough to fit on a floppy disc for data portability. Utilizing seven pulses per azimuthal element permitted coherent summing of the seven pulses to reduce noise. Since the signal amplitude added coherently with each summation and the associated noise added in quadrature, the signal amplitude rose by a factor of seven while the noise amplitude rose by a factor of the square root of seven resulting in a net SNR power gain of 8.5dB.

b. Image Formation Program

With the data now available for MATLAB processing, Figure 4.9 summarizes the algorithm used to turn that information into an image.
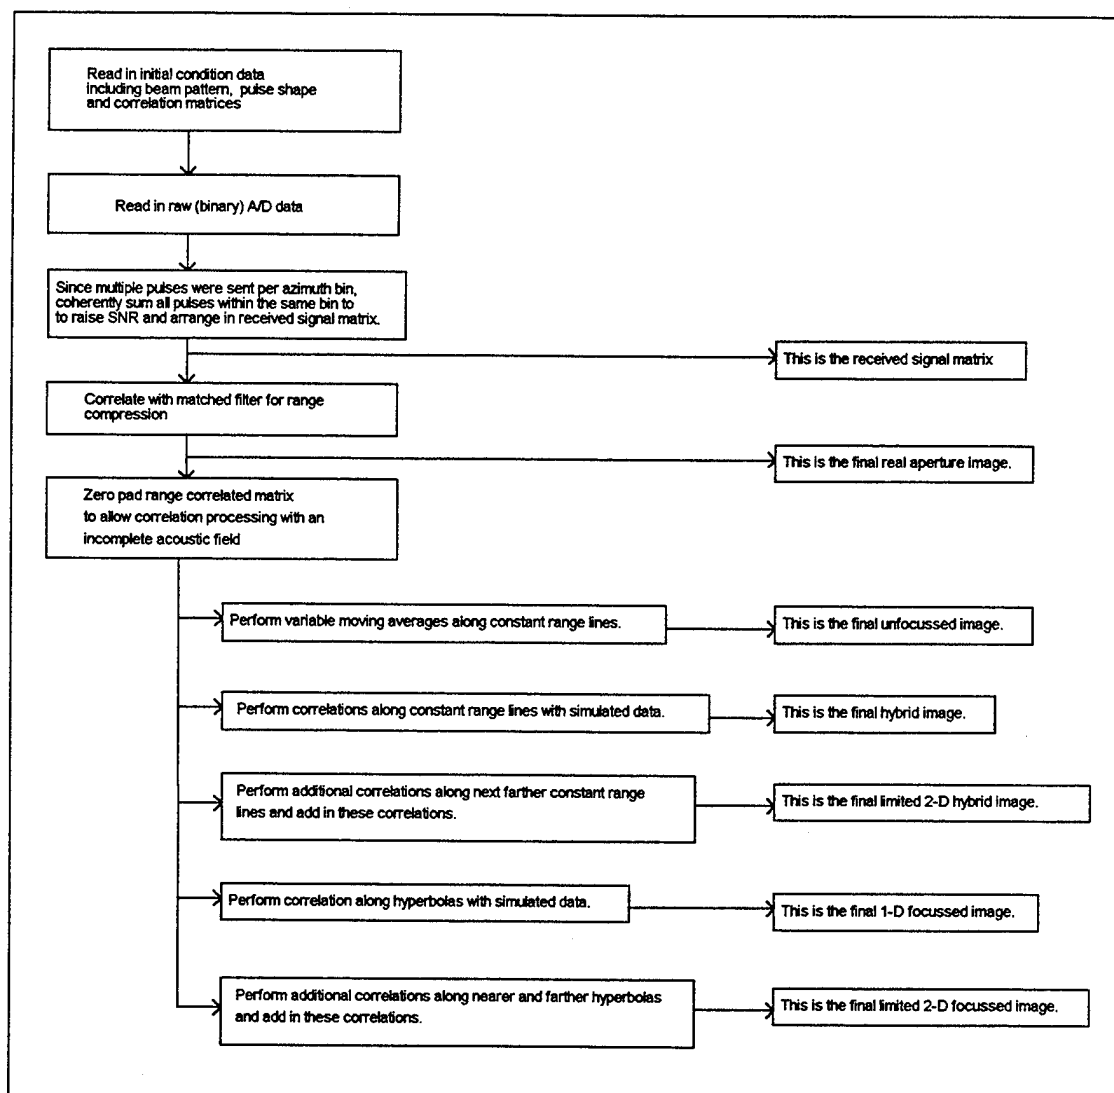
Figure 4.9 Image Formation Algorithms

The implementing code for this algorithm is found in Appendix B. It was written in MATLAB version 4.2C for personal computer use. The program itself took 28 minutes and 41 seconds to produce all seven images on a 486DX2 personal computer with 16Mbytes of RAM.

Other than correlation and sums, no other significant mathematical operations were performed in any of these algorithms. They worked primarily by extracting data from designated locations in the range correlated matrix and comparing it to the results of the previously executed simulation. This greatly enhanced speed performance. The algorithms were executed in order of speed from fastest to slowest. Figure 4.10 shows the execution times for the various azimuthal refinement methods.

Relative Processing Times
(Logarithmic Time Scale)
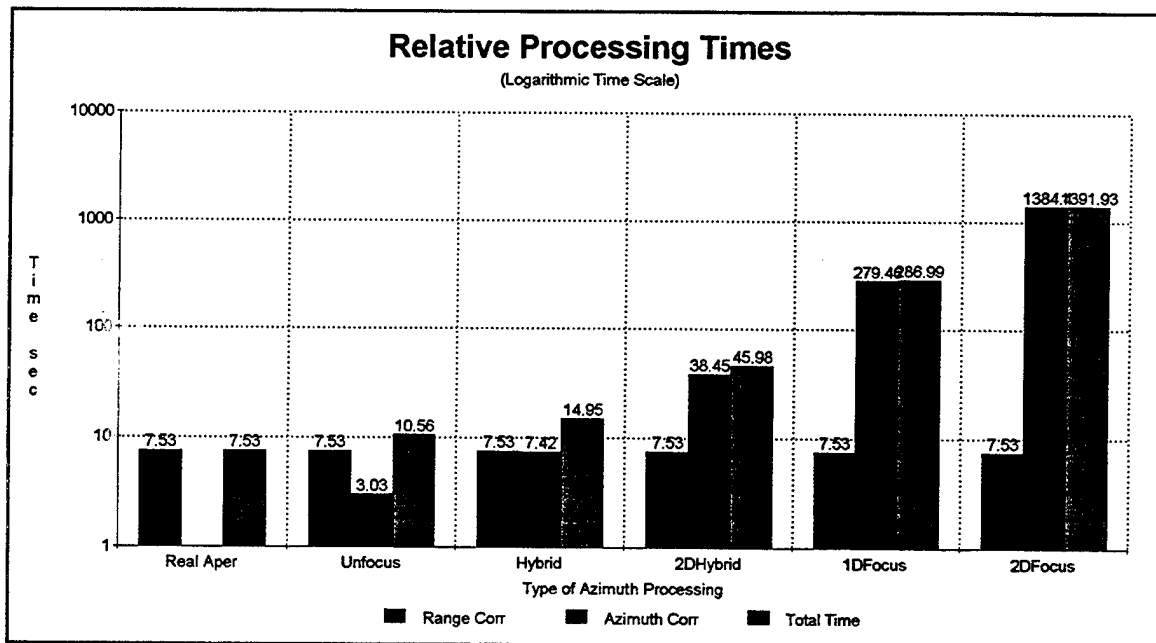
Figure 4.10 Azimuthal Resolution Algorithm Execution Times

The real, unfocussed, and hybrid methods all finish within ten seconds of one another and so are all effectively equivalent from a run time standpoint. The focussed algorithm takes a factor of almost twenty times longer than these and the limited 2-D algorithm is five times longer than the 1-D focussed technique. The limited 2-D hybrid algorithm runs between these two groupings. So, in terms of time efficiency with an eye to resolution, the hybrid algorithm is the best option. If better signal to noise ratio is required, the limited 2-D hybrid algorithm is the next best choice.

V. RESULTS

A. POINT TARGET PERFORMANCE

All of the theory, simulation, and construction were of little value unless the device worked. In this case, the product did work, and at a level very close to the theoretical predictions. To prove this assertion, a data run was completed on a 6.4 cm diameter brass sphere used as a point target. Additionally, a point target was simulated at the same range in the acoustic field using the same pulse shape and beam pattern as was used for the actual data run, but with no simulated noise. The approach of this Chapter will be to compare the two runs side by side to establish the validity of the algorithms.

The first point of comparison was, of course, the input data or received signal. Figure 5.1 shows the two data sets.
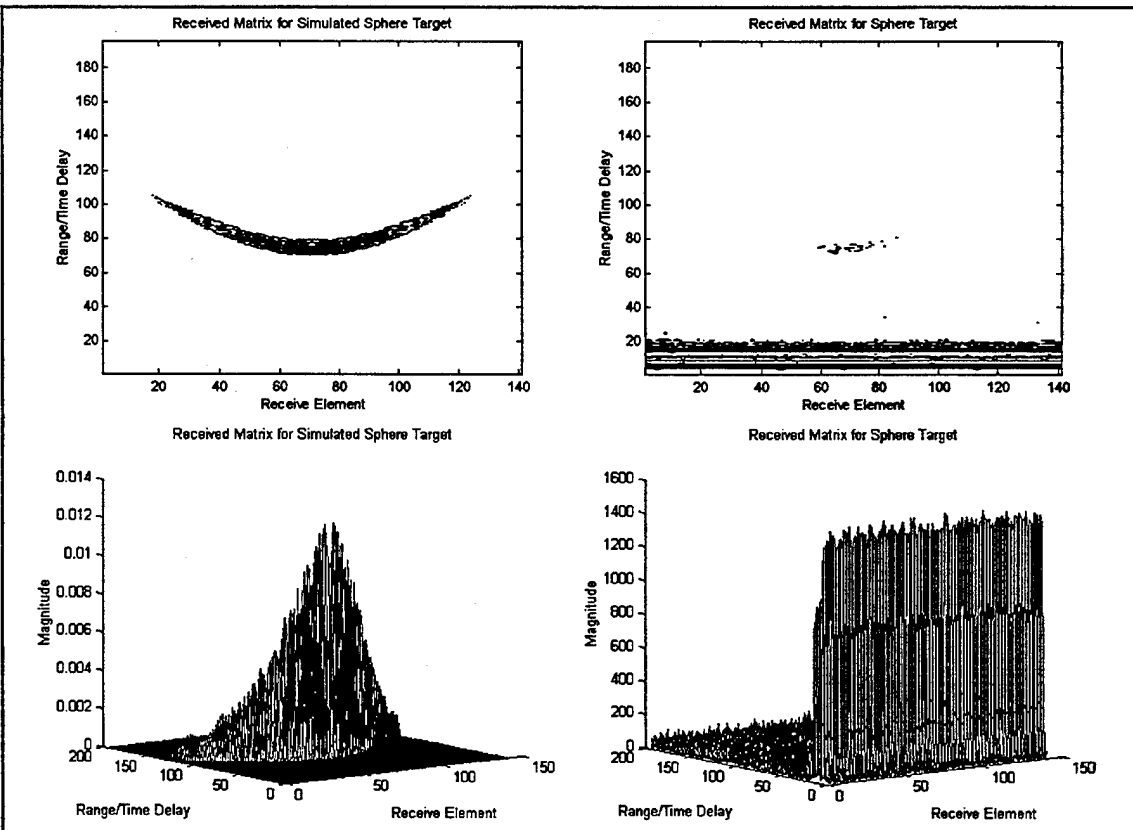

Figure 5.1 Received Signal Matrices for 1 Point Target

The presence of real noise was readily apparent in the case of the actual sphere. The wall of noise at short range was the direct cross talk from the projector to the microphone through the air. This noise was significantly reduced by later processing. This existence of the target was detectable but not blatant due to surrounding noise. This simulation showed the effect of including pulse shape and beam pattern on the received signal matrix when compared to the efforts of Chapter III.

The next step in the processing was the generation of a range correlated matrix from this received signal matrix. The results of this procedure are compared in Figure 5.2.
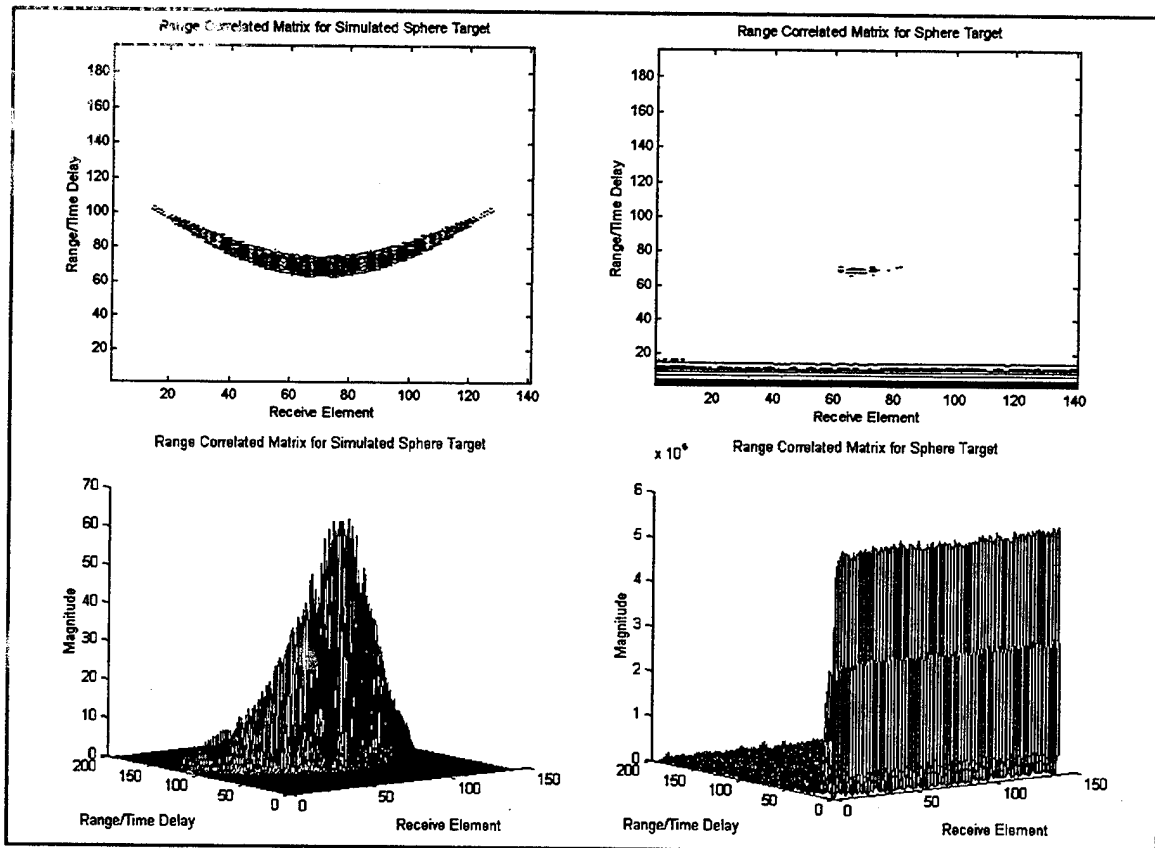
Figure 5.2 Range Correlated Signal Matrices for 1 Point Target

Again, the target was largely submerged in the noise; however, the range correlation did make a small improvement in its visibility. If this were a real array system, the processing would stop here and little further progress could be made other than to eliminate the initial cross talk noise by blanking that section of the matrix.

The first azimuthal resolution processing to be attempted was the unfocussed technique, the consequences of which are depicted in Figure 5.3.
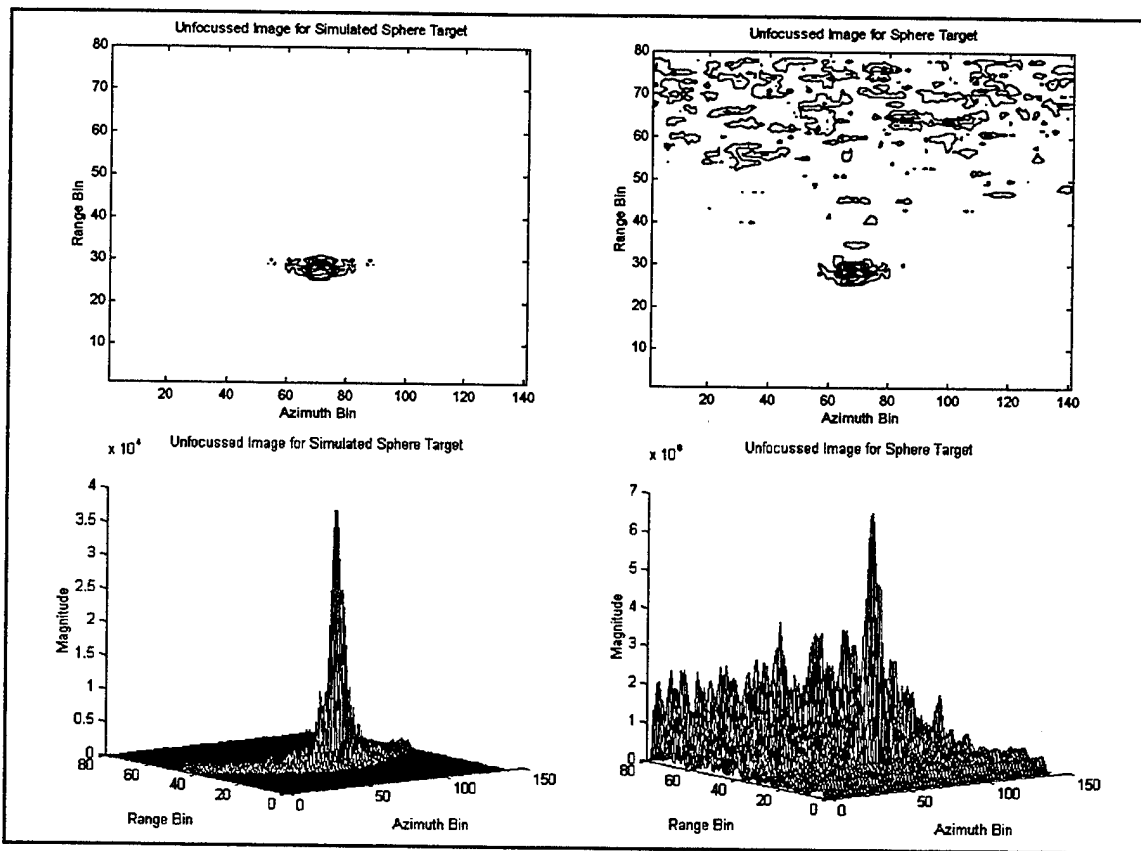
Figure 5.3 Unfocussed Images of 1 Point Target

The 1/R² compensation of this algorithm finally knocked down the short range noise wall at the cost of raising the long range noise amplitude. However, the azimuth processing produced a significant spike at the location of the target, and its resolution appeared to be close to that simulated.

The next azimuthal resolving technique to be tried was the hybrid method. Figure 5.4 demonstrates the results of this method.
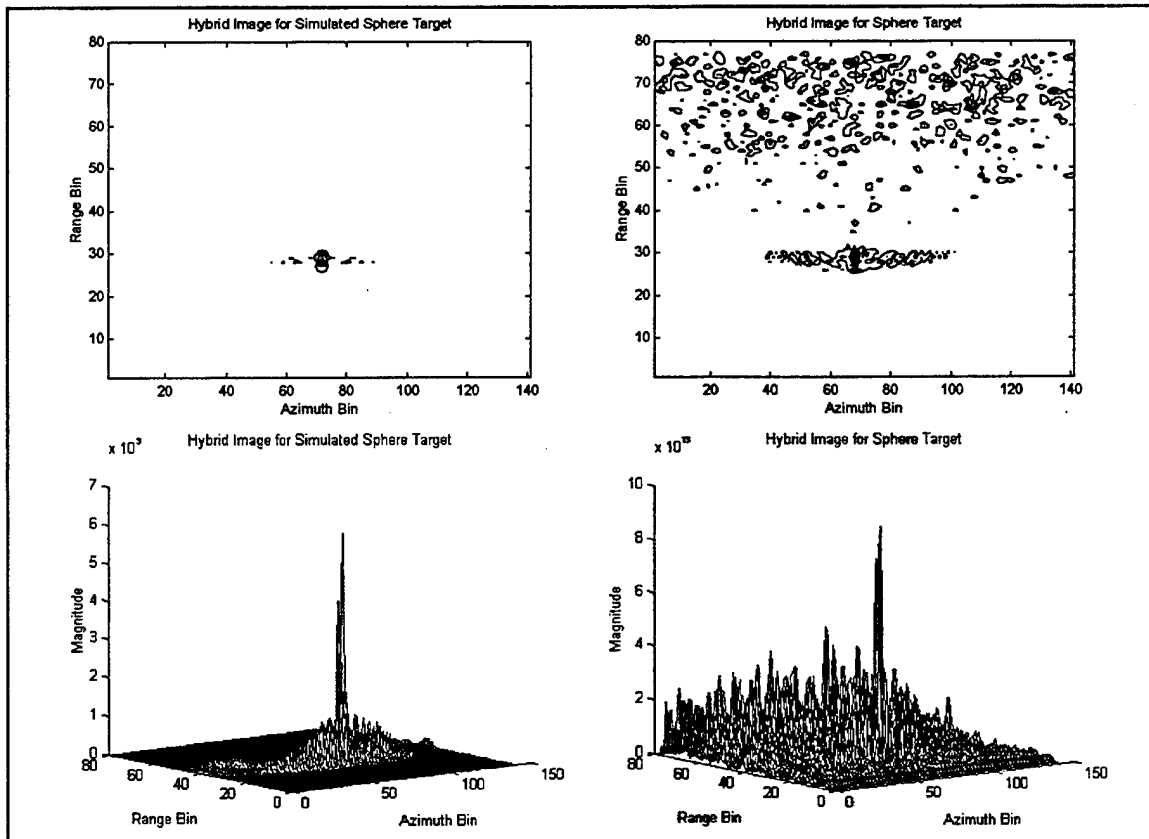
Figure 5.4 Hybrid Images of 1 Point Target

As predicted in Chapter III, the hybrid algorithm provided focussed resolution, but with poor noise resistance since its correlation path was not along the maximum response curve of the signal.

Working to correct this deficiency was the fully focussed algorithm whose effects can be seen in Figure 5.5.
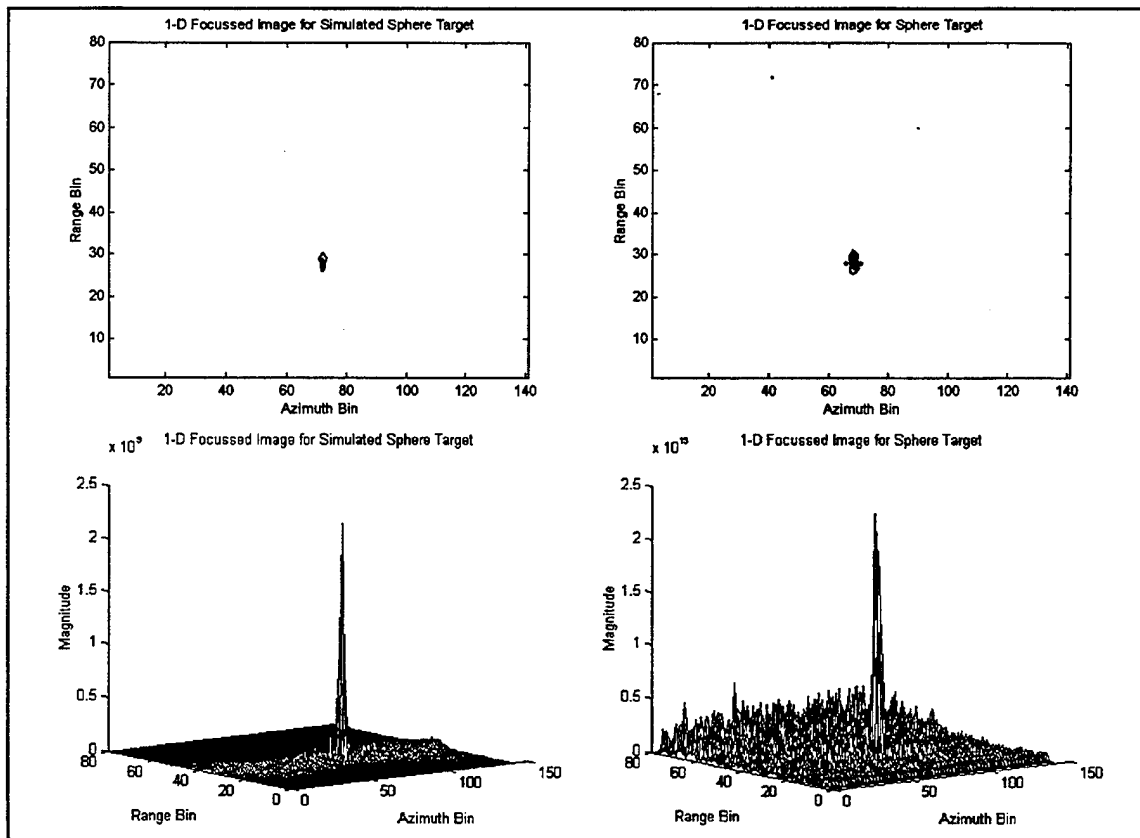
Figure 5.5 1-D Focussed Images of 1 Point Target

The correlation along the maximum response curve significantly improved the appearance of Figure 5.5 over that of Figure 5.4, although, both have the same resolution.

Another solution to the noise sensitivity exhibited by the hybrid algorithm in Figure 5.4 was to expand its correlation into the range dimension in the limited two-dimensional fashion previously described. Figure 5.6 displays the results of this expansion.
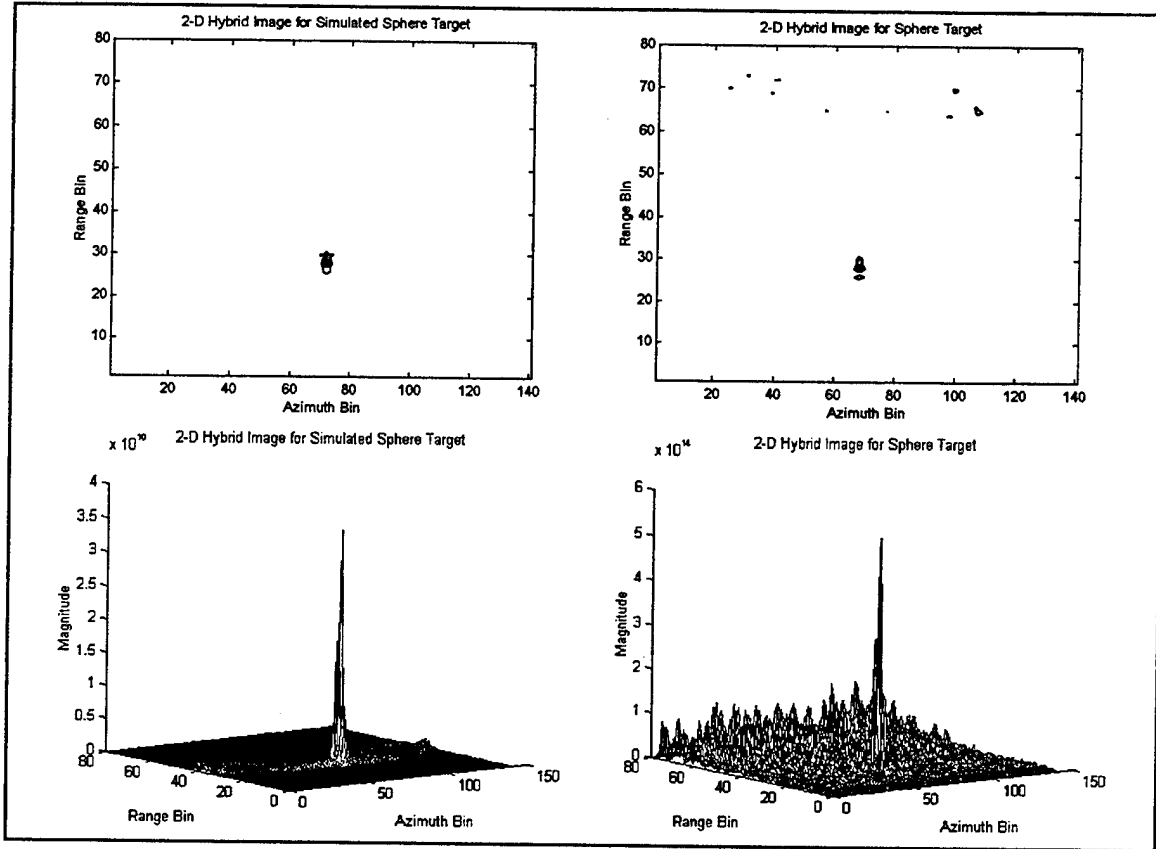
Figure 5.6 Limited 2-D Hybrid Images of 1 Point Target

The performances of this hybrid two-dimensional technique and the one line fully focussed algorithm in resolution and SNR were nearly identical from a straight visual comparison. However, this limited two-dimensional hybrid took less than one third of the computer time required for the one line fully focussed algorithm.

For the final image of the best resolution and noise performance processed, Figure 5.7 shows off the limited two-dimensional fully focussed correlation.

72
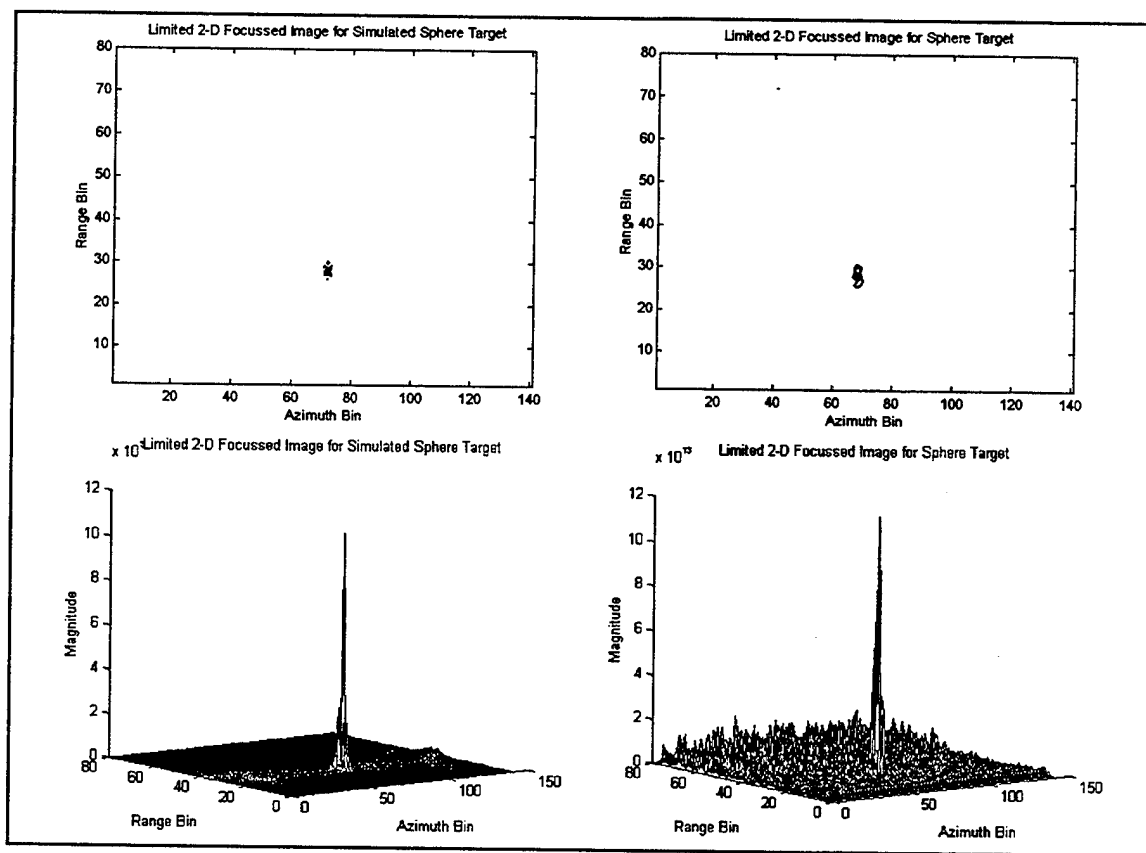
Figure 5.7 Limited 2-D Focussed Images of 1 Point Target

Comparison of Figures 5.5 and 5.7 shows that the two-dimensional process did reduce the noise for the fully focussed version just as it did for the hybrid version. Again, the 2-D fully focussed algorithm showed better SNR performance than the 2-D hybrid technique in the same manner that their one-dimensional counterparts did due to a more optimum selection of signal correlation band.

In summary, in all azimuth processed cases, the simulated and experimental resolutions were visually nearly identical to one another. The signal to noise performance of the algorithms followed the pattern predicted by theory. Thus, the basic concept of a synthetic aperture sonar has been proven in practice. However, concepts which work well in the ideal conditions of a laboratory, often fail when subjected to more realistic environments. Since these methods were coherent, using phase information from individual waves, any position errors in the location of the transmitter-receiver combination of more than a fraction of a wavelength could severely defocuss the results as shown in Figure 5.8.
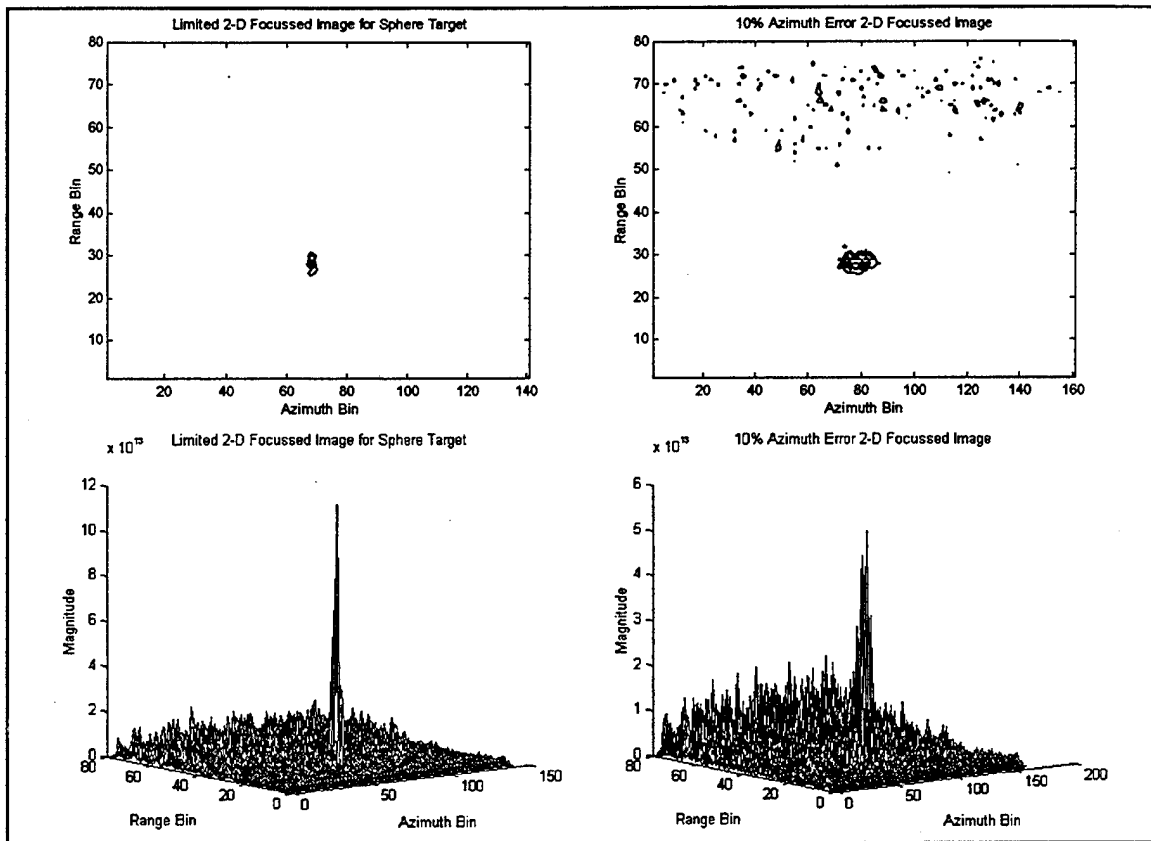
Figure 5.8 Effect of Azimuthal Travel Error

The sphere images on the left of Figure 5.8 were processed with the proper azimuthal travel distance. Those on the right were processed assuming a travel distance about ten percent too long. Otherwise, both images were processed identically with the limited 2-D focussed algorithm. The azimuthal resolution and signal to noise characteristics of this defocussed case are compared quantitatively to those of the other techniques without induced phase errors in Figures 5.9 and 5.10.
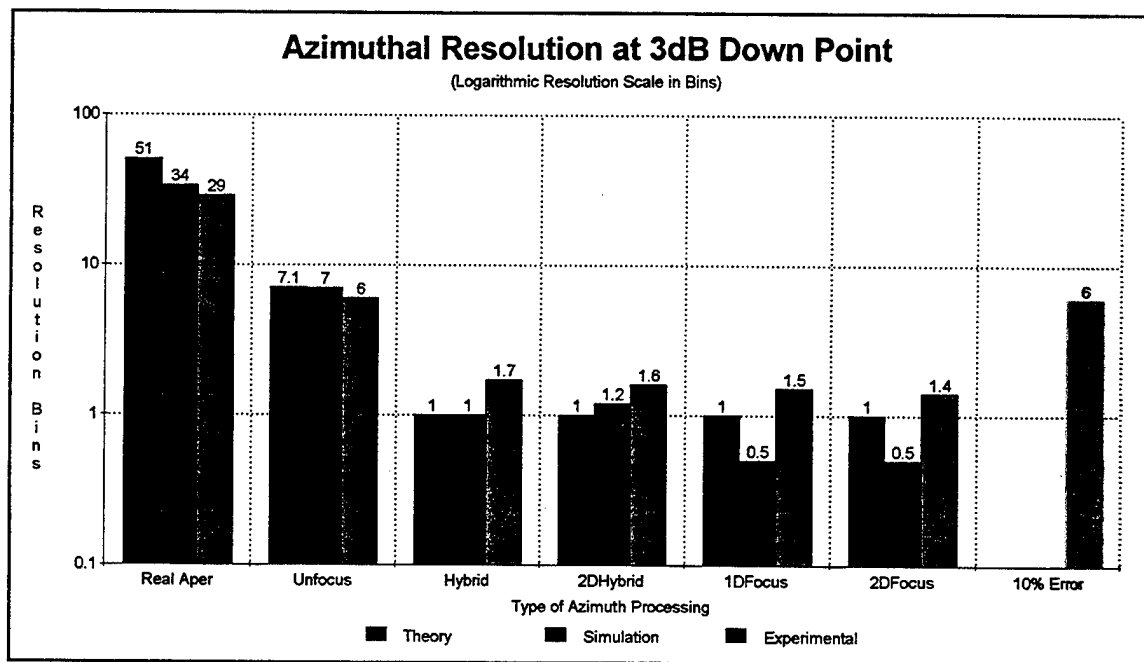
Figure 5.9 Sphere Target Azimuthal Resolution Results
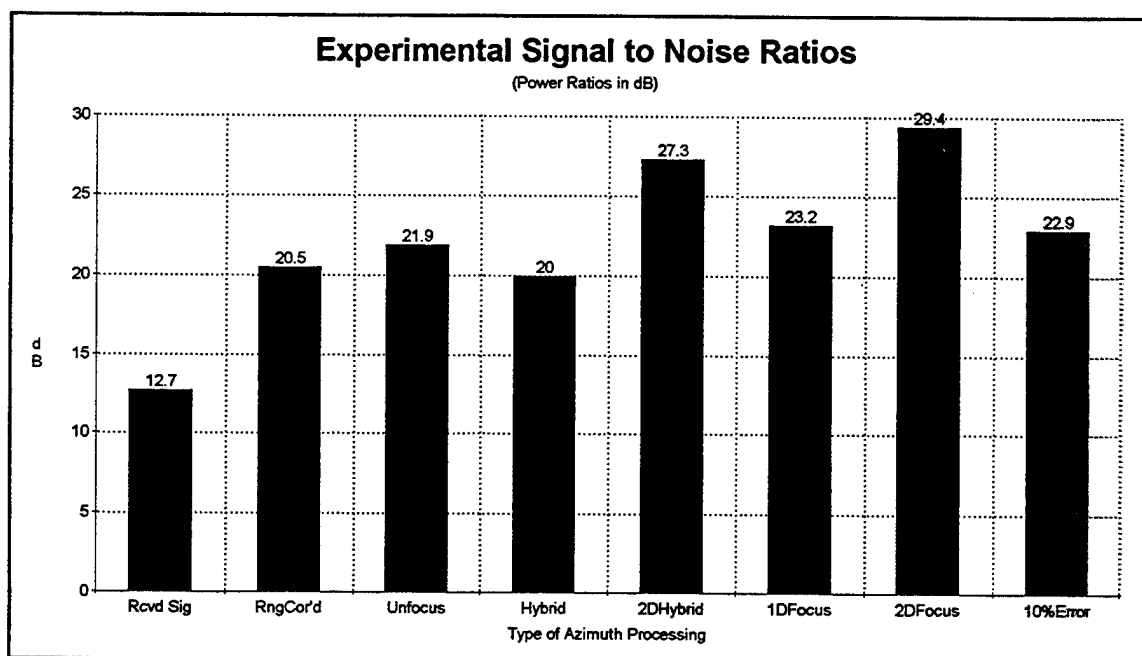
Figure 5.10 Sphere Target SNR Results

In each azimuthally processed example, the theoretical, simulated, and experimental resolutions fell within 1.1 range bins of one another, indicating good agreement. The theoretical resolution for the real aperture was coarser than that of the simulation or the experiment because the theoretical azimuth resolution calculation assumes a rectangular beam pattern. The simulation and experiment worked with the measured beam pattern which had less energy at the outer edges of the pattern, hence, a narrower 3dB beamwidth. The SNR's showed the expected trend with the exception of the hybrid case which fell slightly below the range correlated and unfocussed results. In both cases where an algorithm was generalized from one dimension to two dimensions (increasing the correlation length / area by a factor of five), about six to seven dB of SNR was gained- as was predicted by theory. Thus, the overall system performance agreed with both the theory and the simulations, not only in a subjective sense based on the appearance of the resultant images, but also on the resolution and SNR achieved.

The ten percent navigation error degraded the resolution capability of the system by a factor of four. It also reduced the noise rejection capabilities by 6.5dB. Comparison with the other data shows that the phase error made what would otherwise be a fully focussed system indistinguishable from an unfocussed system by either resolution or SNR characteristics. Wobble in the down range direction could be expected to produce even more severe degradations of the final image. Therefore, in a tactical system, either the navigation would need to be extremely precise (not very probable), or adaptive phase compensation techniques would need to be utilized.

These plots along with Figure 4.10 help to explain the popularity of the hybrid method. From Figure 4.10, it gives speed performance comparable to unfocussed methods. From Figure 5.9, it gives resolution close to that of fully focussed techniques. Finally, while slightly more susceptible to noise, its noise resistance is still quite comparable to the other methods and can be augmented easily by a limited two-dimensional technique.

B. ADVERSE SNR PERFORMANCE

Thus, the basic concept of a synthetic aperture sonar has been proven in practice. However, in some ways, the test might have been too easy. An untrained observer could have examined the experimental received signal matrix of Figure 5.1 and determined that there was a target present. The true test of the effectiveness of these techniques is to

reduce the SNR until a human operator could not decipher the received signal without help, and then, assess the recognition gain achieved by the techniques.

To test this sensitivity to SNR, a data run was configured with the target arrangement of Figure 5.11.
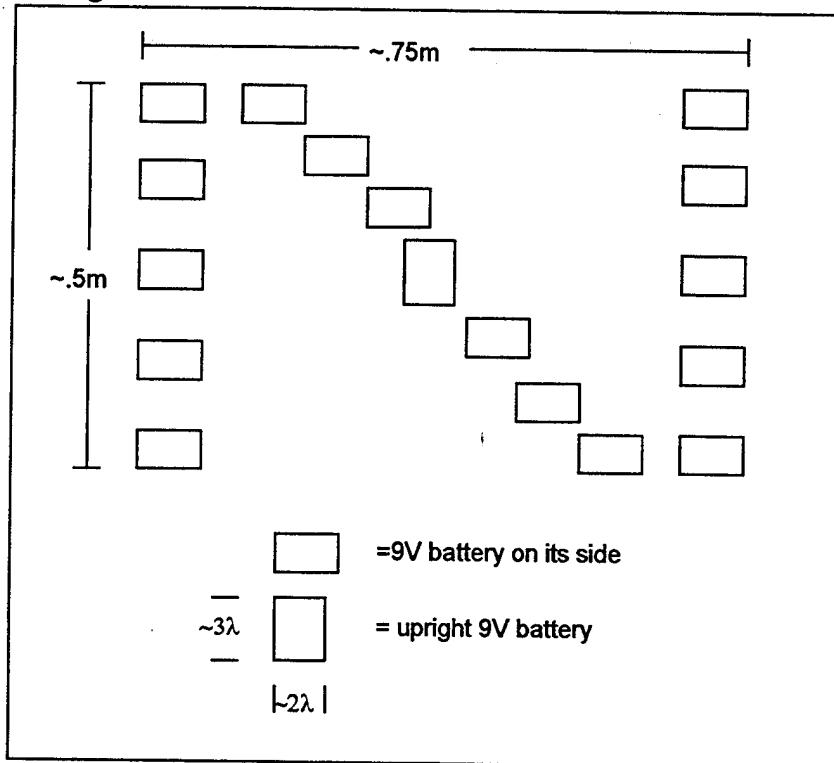
Figure 5.11 High Noise Target Configuration

To make the experiment more like the probable field use of the process, the target acoustic field was illuminated from an elevated angle as depicted in Figure 5.12.
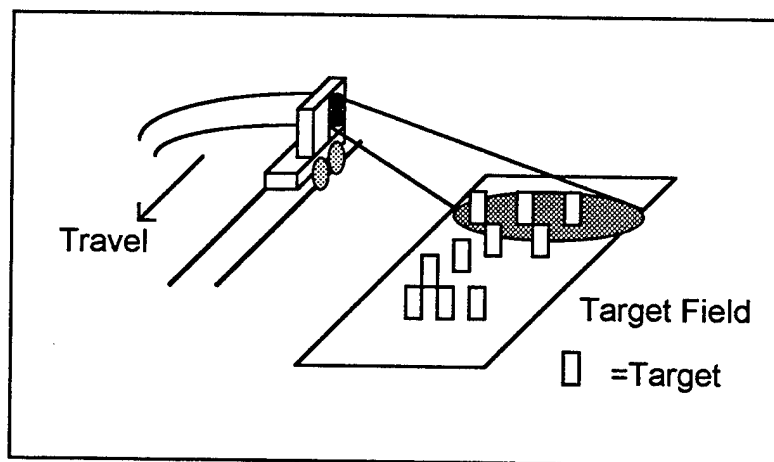
Figure 5.12 High Noise Illumination Geometry

Rather than injecting large quantities of additional noise, the signal amplitude was reduced by a factor of ten from the previously displayed results. This degraded the signal to noise power ratio by 20dB. With multiple targets, the previous display formats would quickly become unrecognizable even with strong SNR. As a result, this section of the paper uses a slightly different but more easily interpreted plotting style. Figure 5.13 shows the received signal matrix recorded for these conditions.
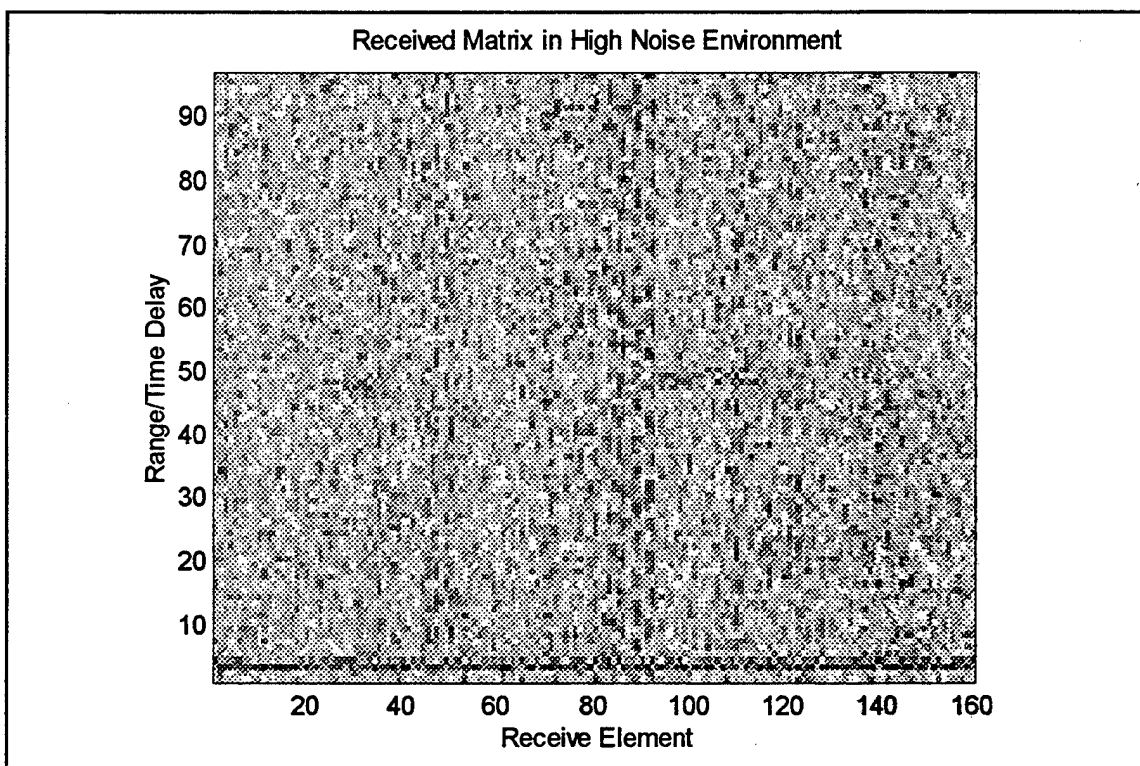
Figure 5.13 High Noise Received Signal Matrix

Even with knowledge of the target arrangement, the targets could not be visually identified from this display. Figure 5.14 depicts the next step toward identification of the targets.
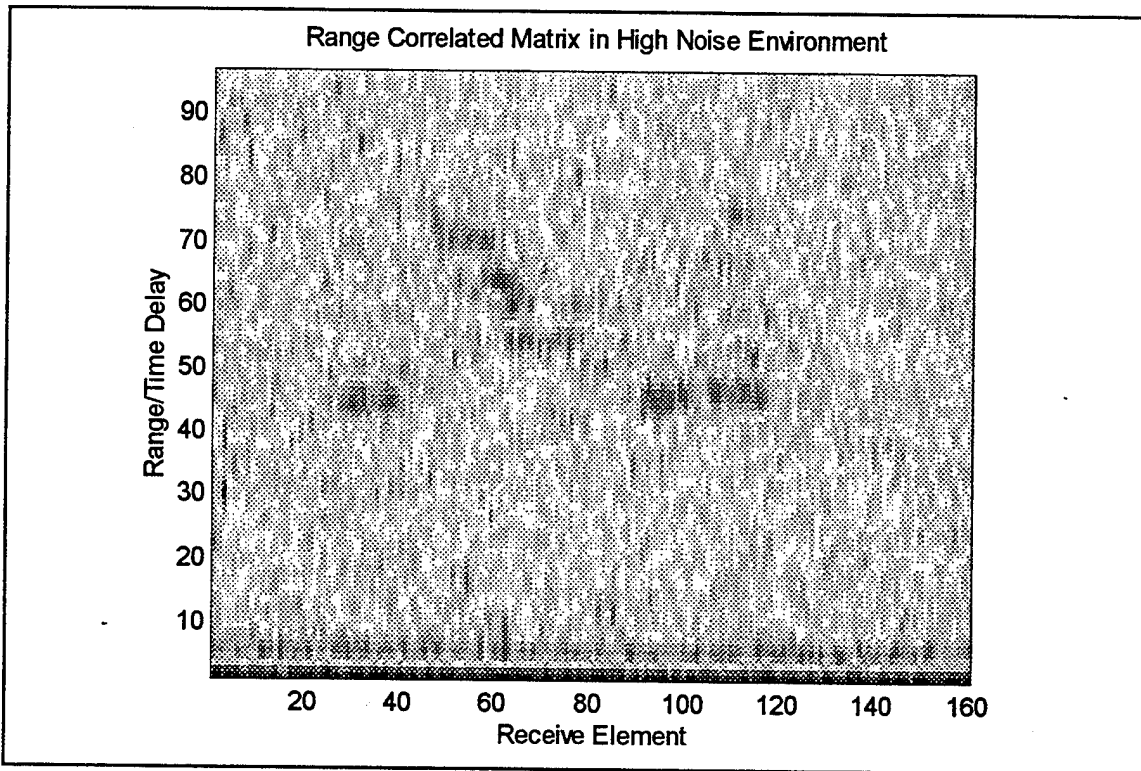
Figure 5.14 High Noise Range Correlated Signal Matrix

As is readily apparent from Figure 5.14, some SNR was gained by performing the range correlation, particularly for the nearer targets. The existence of the short range targets then became apparent, where, in the unprocessed data, it was not.

The first coherent azimuthal processing attempted on this range correlated signal was that of an unfocussed algorithm as displayed in Figure 5.15.
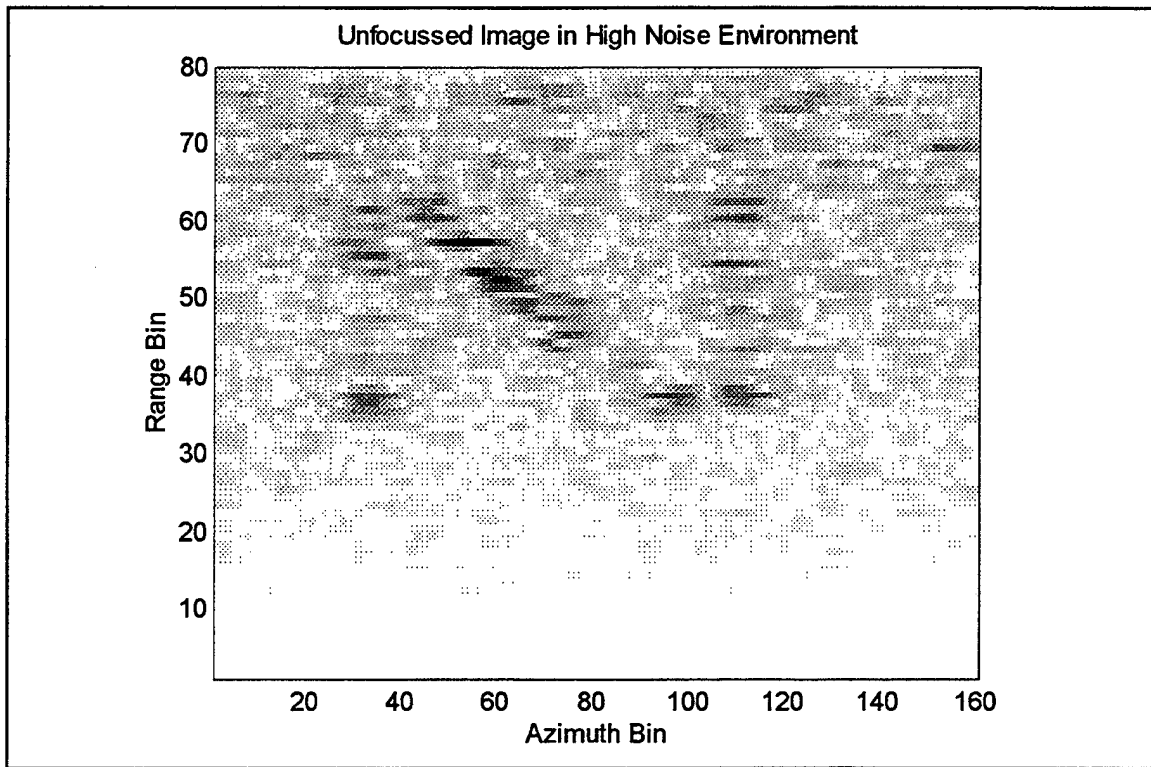
Figure 5.15 High Noise Unfocussed Image

The overall pattern of the target arrangement had become obvious, although, individual targets were difficult to resolve, particularly at long ranges.

The hybrid algorithm was the first attempt to resolve individual targets with the results shown in Figure 5.16.
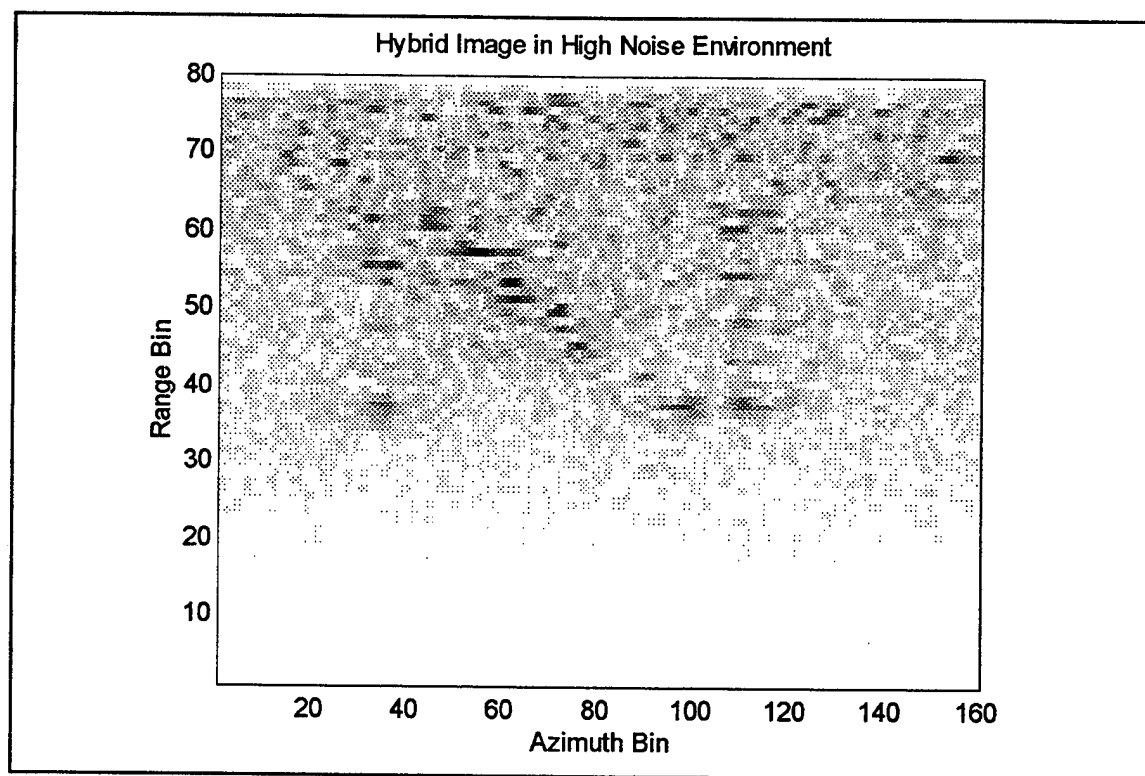
Figure 5.16 High Noise Hybrid Image

This image showed relatively good resolution, but, the low SNR made the targets appear rather faint.

The limited two-dimensional hybrid efforts displayed in Figure 5.17 were an attempt to remedy this condition.
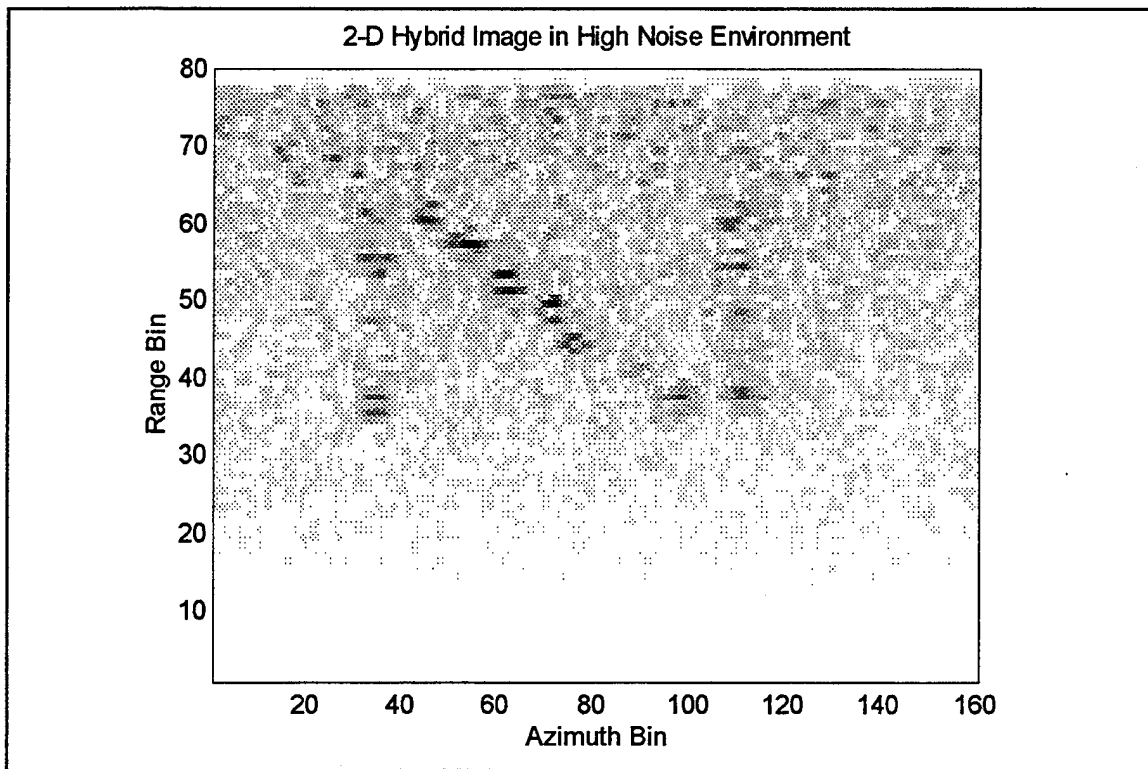
Figure 5.17 High Noise 2-D Hybrid Image

The targets in this image were a great deal more distinct, both from each other and from the background noise.

Another method used to improve the image was the use of a one line fully focussed technique, the results of which are depicted in Figure 5.18.
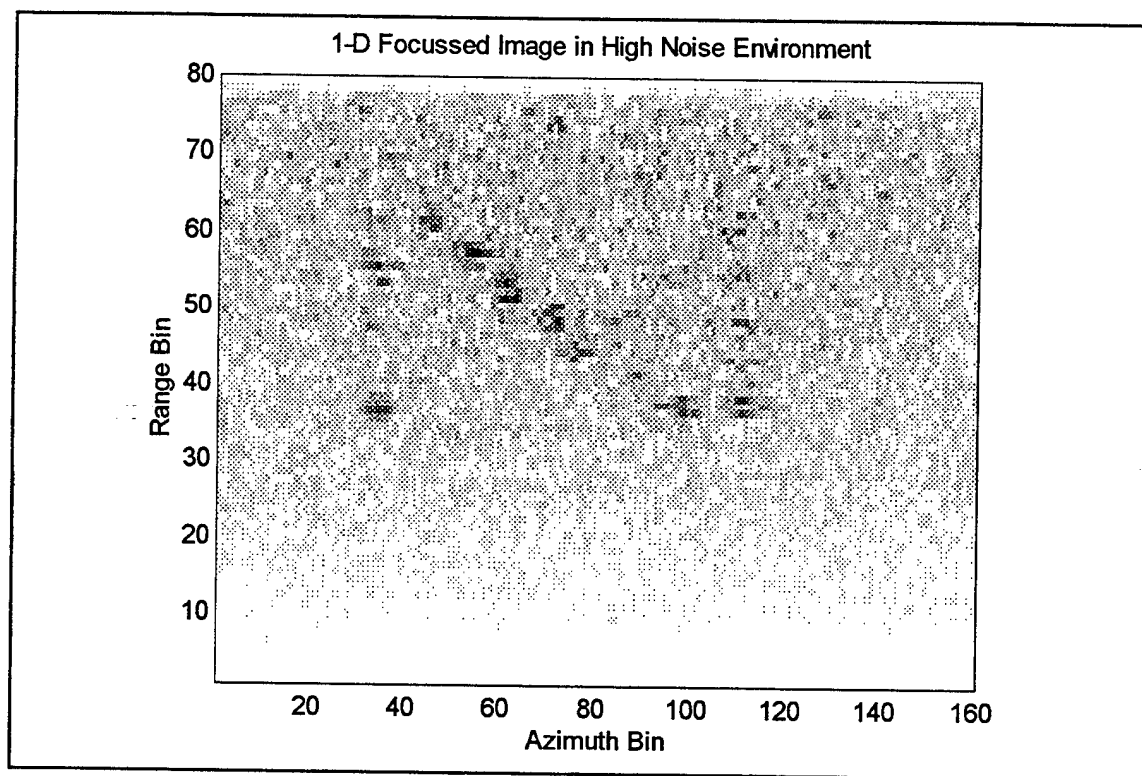
Figure 5.18 High Noise 1-D Focussed Image

It was apparent here that the individual targets were resolved, where they were not shadowed by closer targets. When compared to Figure 5.17, this image was slightly more grainy, but both showed the targets with nearly equal clarity.

Figure 5.19, which depicts the limited 2-D focussed technique, was the final attempt at gaining a good image while overcoming poor SNR conditions.
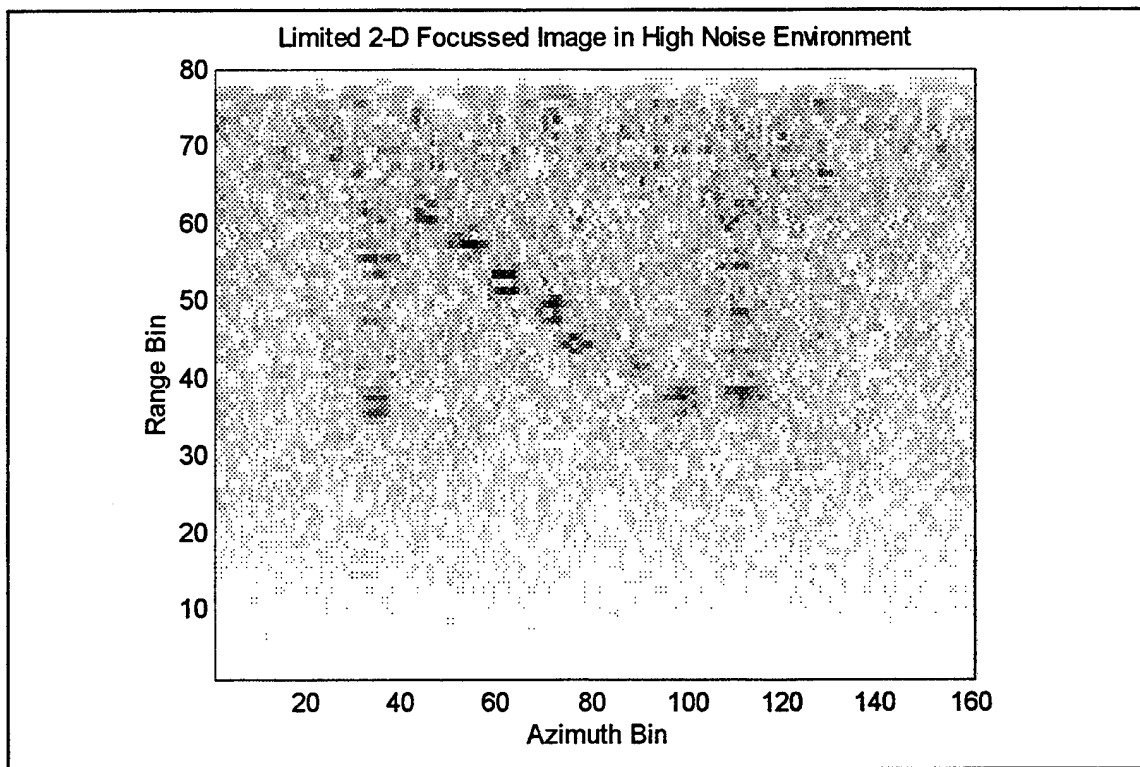
Figure 5.19 High Noise Limited 2-D Focussed Image

These results were rather impressive when considering the facts that there were no discernible targets in the received signal and that the targets were only a few wavelengths long with separations on that same order.

Hence, these techniques have proven themselves capable in a challenging SNR environment. To qualify that capability in a slightly different way, Figure 5.20 shows the RMS signal and noise plots as a function of range for each of the poor SNR displays.
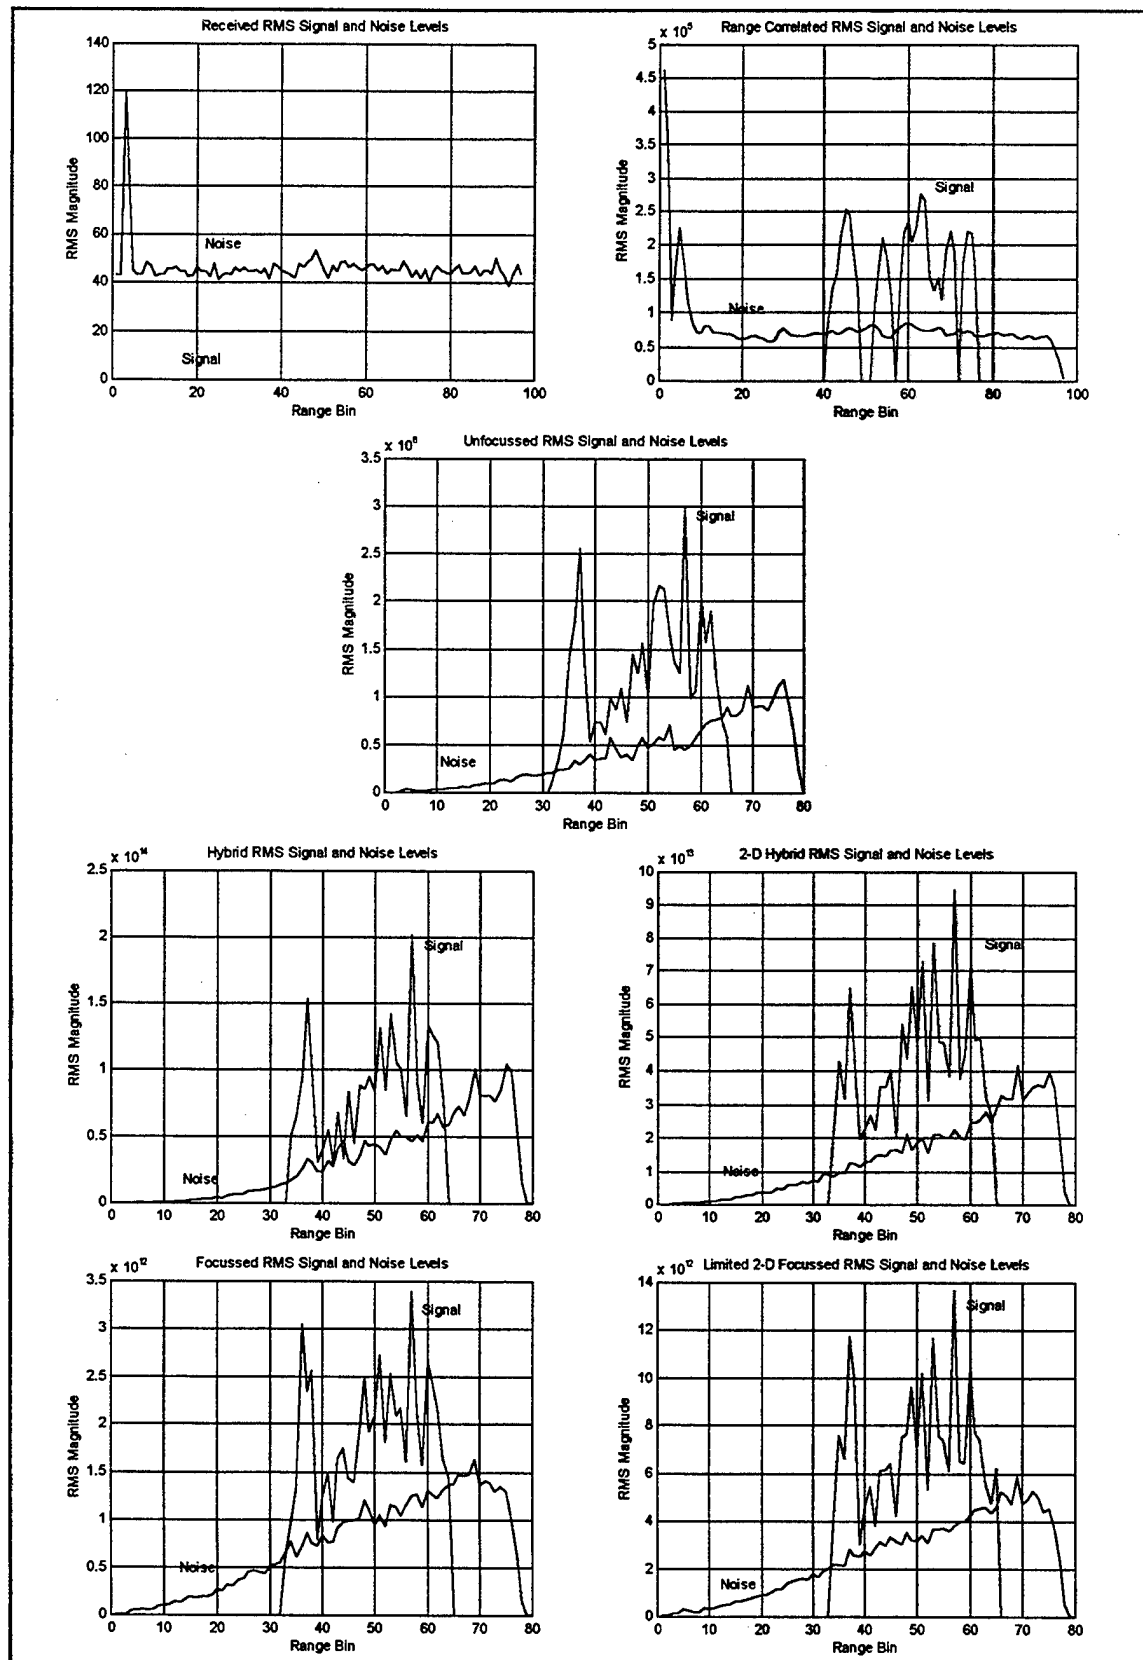
Figure 5.20 RMS Signal and Noise vs. Range

These graphs show the range dependence of the unprocessed and processed noise. In the received and range correlated signals, no $1/R^2$ compensation was used, so, the noise level was independent of range except for the direct cross talk at short range. Where the $1/R^2$ compensation was applied in processing, the noise level rose with the signal at longer ranges. The trends in noise performance of the various azimuthal processing techniques seen here were consistent with those of Figure 5.10 with the exception that the SNR differences among techniques were not as dramatic as expected.

VI. CONCLUSIONS

This paper investigated several synthetic aperture techniques from their theory, through their simulation, and finally to a working air medium synthetic aperture sonar. The results were, in summary, that the simulations confirmed the theory and that an actual working prototype confirmed the simulations. Fully focussed resolutions within a factor of two of the theoretical limit of D/2 were measured experimentally. In this instance, that attained resolution corresponded to a length of about one centimeter. Unfocussed experimental results were found to be in complete agreement with theory and simulation. Sensitivity to phase errors and their defocussing effects were experimentally observed. The noise resistance of the algorithms was found to be in general agreement with expected trends and sufficient to extract a target from received SNR's of around 0dB.

Since these techniques functioned as expected, the next logical step is to apply them. These methods are already in use in radar mapping applications. The principle reason these processes were studied at the Naval Postgraduate School is that they have the potential to be developed as sonar techniques to locate, identify, and classify mines, mine-like objects, and unexploded ordinance at sea, in the littoral, and in the ground. Additionally, this technology has the potential to be added onto a poor resolution coherent system, and without changing any transducer or pulse characteristics, improve the azimuth resolution of the system by orders of magnitude at long range. If the acoustic field can be over-sampled as was done for this paper, these techniques can be quite noise resistant. However, all applications which require the accuracy that these procedures can offer need to either navigate very precisely- to within a fraction of a wavelength- or use adaptive phase compensation.

The final question to be addressed in this paper is that of in what directions to proceed from here. The first upgrade recommended is that of FM chirping the transmitted pulse. This will improve the range resolution to make it comparable to the azimuth resolution. Second, an adaptive phase compensation technique needs to be implemented, since wavelength accurate platform position information is rarely available. This phase compensation will also reduce the errors induced by path length variance effects such as turbulence. Finally, when all effects have been adequately addressed by the air medium version, progress into the water and sand environments can be attempted.

The techniques described by this paper have the potential to make a significant contributions to the high resolution imaging necessary for the U.S.Navy to detect, classify,

and identify mine-like objects so that it can more effectively execute its mission in the littoral.

APPENDIX A. SIMULATION CODE

The simulation algorithm was implemented in the following code (Remarks are in italics.):

```
% Establish initial conditions and parameters
clear
    maxrange=1;          % Max unambiguous range in meters
    lmntspc=.00686;      % Spacing between elements in meters (chosen=D/2)
    fcarrier=25000;      % Carrier frequency in Hertz
    fsample=100000;      % Sample rate in samples/second
    pulselength=.00035;  % Pulselength in seconds
    PRR=13.4264;         % Pulse repetition rate in Hertz (chosen to fit data files on 1
floppy disk)
    vsound=343;          % Sound velocity in meters/second
    vplatform=.0131579;  % Platform velocity in meters/second (as measured)
    rangeres=.0125;      % Desired range resolution in meters (about one wavelength)


% Establish helpful conversions

    azres=lmntspc;               % azimuth resolution in meters/sample
    finerr=vsound/(2*fsample);   % fine range resolution in meters/sample
    coarserr=rangeres;           % coarse range resolution in meters/bin
    sampbin=coarserr/finerr;     % # fine samples/bin (range)
      SBW=.5;
    C=ceil(2*maxrange*SBW/lmntspc);  % total # of azimuth samps
    R=ceil(maxrange/finerr);         % total # of range samples
    Rx=ceil(maxrange/coarserr);      % total # of range bins
    corrleng=fsample*pulselength;    % # of samples per pulse
    fref=fcarrier/fsample;           % reference frequency
    lam=coarserr*vsound/(lmntspc*lmntspc*fcarrier*8);
    load 'a:\exp6bmen.mat'       % brings in bmsin, pulsenv (beam envelope, pulse shape)
                                 % (experimentally measured)
    bmwid=length(bmsin);         % for use in later interpolations


% Calculate matrix of range sample interpolation vectors for each range bin
% These are vectors of the sample #'s of maximum response
% These are not rounded to allow for later interpolation (if needed)
% Units of "rx" are "samples"
% Vectors are row vectors of constant range targets centered in azimuth
% Vectors vary along their length as a function of azimuth
% Different row vectors correspond to targets in different range bins

x=[1:C];                  % azimuth bin #'s
rx=ones(Rx,C).*(R+1);     % Unaffected locations are set out of range
```

```
for R0=1:Rx
  sq1=(x-ceil(C/2)).*(lmntspc/(R0*coarserr));
                    % Set up for Pythagorean Theorem
                    % (x-ceil(C/2))*lmntspc=1st side in meters
                    % R0*coarserr=2nd side in meters
                    % rx*finerr=hypotenuse in meters
  rx(R0,:)=(R0*sampbin).*sqrt(1+sq1.*sq1);   % rx=range in samples
end
rx(rx>R)=ones(size(find(rx>R))).*(R+1);
                    % Out of range indices are all set to a common
                    % index to later indicate a zero return
cl=ceil(rx);        % cl=1st index in corrsig matrix after peak
clear x sq1 azres


% Set up to simulate a target at each range bin

axm2=zeros(Rx,C);   % correlation values taken from correlation peak minus 2
samples
axm1=zeros(Rx,C);   % correlation values taken from correlation peak minus 1 sample
ax=zeros(Rx,C);     % correlation values taken from correlation peak
axp1=zeros(Rx,C);   % correlation values taken from correlation peak plus 1 sample
axp2=zeros(Rx,C);   % correlation values taken from correlation peak plus 2 samples

axunf=zeros(Rx,C);  % correlation values taken from target range bin
axunf1=zeros(Rx,C); % correlation values taken from target range bin plus 1 samples
axunf2=zeros(Rx,C); % correlation values taken from target range bin plus 2 samples
axunf3=zeros(Rx,C); % correlation values taken from target range bin plus 3 samples
axunf4=zeros(Rx,C); % correlation values taken from target range bin plus 4 samples

% Set up for range correlation by matched filter convolution.
t=0:corrleng-1;
carr=sin((2*pi*fref).*t+pi/4); % Carrier signal to be modulated
corrfil=pulsenv.*carr;         % Template of ideal return after modulation
cf=fliplr(corrfil);            % setup for matched filter/convolution

clear t corrfil carr scopediv T n b m

for R0=1:Rx       % The target is simulated at R0 range bin
  R0              % to track progress of the program
  tgt=ceil(C/2);  % The target is centered in azimuth
  rcvdsig1=zeros(R,C); % rcvdsig1 is the simulated signal matrix
  for lmnt1=1:C          % Cycle through each azimuthal element
    sum1=zeros(R,1);     % The focussed signals from target points will be summed here
    sq=(tgt-lmnt1)*(lmntspc/(R0*coarserr));
                    % Set up for Pythagorean Theorem
```

```
            % (tgt-lmnt1)*lmntspc=1st side in meters
            % finerr*y(tgt)=2nd side in meters
            % r*finerr=hypotenuse in meters
r=(R0*sampbin)*sqrt(1+sq*sq);
            % r=range in samples
f1=(1+(2*vplatform/vsound)*(lmntspc/finerr)*(tgt-lmnt1)/r);
            % f1=Doppler factor for fref
            % Doppler shift=(2*vplatform/vsound)*sin(theta)
            % sin(theta)=(lmntspc/finerr)*(tgt-lmnt1)/r
            % This accounts for Doppler velocity effects which turn out to be
            % 3 orders of magnitude smaller than remaining effects
n1=corrleng/f1;    % n1 is length of return pulse including Doppler effect
n2=n1;             % n2 is length of envelope including Doppler effect
ind=ceil(r);       % ind is starting index/sample of signal
  if (ind<R)&(ind+n1>=1)
            % ensures signal falls within beam pattern in range

            % These next lines allow this algorithm to be adapted for use on
            % other range zones i.e. 3 to 4m vice 0 to 1m.
  td=0;
  if ind+n1>R     % catches the head of a return at far edge
    n1=R-ind;
  end
  if ind<1         % catches tail of return inside near edge
    n1=n1-(1-ind);
    td=1-ind;
    ind=1;
  end
  ph=ind-r;        % ph compensates for ind not starting on r

            % Squeeze or stretch envelope for Doppler:
  index=[0:n2-1].*(corrleng-1)/(n2-1)+1;
            % Set up for linear interpolation
            % Index values fall between indices of env
            % Since indices of env are separated by one,
            % rem is the fractional placement of index toward
            % the next higher index of env...or the fractional
            % contribution of the next higher index of env to
            % this index of env1
            % lrem performs the same function for the next lower
            % index of env
  rem=index-floor(index);
  lrem=ones(size(rem))-rem;
  env=pulsenv(floor(index)).*lrem+pulsenv(ceil(index)).*rem;
```

```matlab
            % Account for location in azimuthal beam pattern
        sn=lmntspc/finerr*(tgt-lmnt1)/r;          % [sin(th)]
        indx=(sn+1)*(bmwid-1)/2+1;
        rm=indx-floor(indx);
        lrm=1-rm;
        bmcorr=bmsin(floor(indx))*lrm+bmsin(ceil(indx))*rm;

        n=[0:n1-1]';
        pl=sin((2*pi*fref*f1).*(n+td+ph)+pi/4).*env(td+1:td+n1).'.*(bmcorr/(ind*ind)));
        sum1(ind:ind+n1-1)=sum1(ind:ind+n1-1)+pl;
      end
    rcvdsig1(:,lmnt1)=sum1;
  end        % Go to next azimuthal element
% Now rcvdsig1 matrix is complete for a target at range bin R0


%%%%% Range correlation

corrsig1=zeros(R+5,C);          % Include five more rows of zeros so that out of
                                % range returns are zeroed


  for c=1:C                     % This loop performs range compression on each column
    cs=conv(rcvdsig1(:,c),cf);    % matched filter / range correlation
    corrsig1(1:R,c)=cs(floor(corrleng):R+floor(corrleng)-1);
  end


clear rcvdsig1


%%%%% Azimuth Correlation Data Extraction

ax(R0,:)=diag(corrsig1(cl(R0,:),:)).'.*(R0*R0);
        % Pick values out of corrsig1 from ranges / times
        % specified by rx to be stored as simulated values in ax
        % compensated for 1/R² drop off
axm2(R0,:)=diag(corrsig1(cl(R0,:)-2,:)).'.*(R0*R0);
axm1(R0,:)=diag(corrsig1(cl(R0,:)-1,:)).'.*(R0*R0);
axp1(R0,:)=diag(corrsig1(cl(R0,:)+1,:)).'.*(R0*R0);
axp2(R0,:)=diag(corrsig1(cl(R0,:)+2,:)).'.*(R0*R0);
        % Pick values out of corrsig1 from range R0, R0+1 sample, R0+2 samples, etc.
axunf(R0,:)=corrsig1(ceil(R0*sampbin),:).*(R0*R0);
axunf1(R0,:)=corrsig1(ceil(R0*sampbin)+1,:).*(R0*R0);
axunf2(R0,:)=corrsig1(ceil(R0*sampbin)+2,:).*(R0*R0);
axunf3(R0,:)=corrsig1(ceil(R0*sampbin)+3,:).*(R0*R0);
axunf4(R0,:)=corrsig1(ceil(R0*sampbin)+4,:).*(R0*R0);
end            % Go to next target range, R0
```

```
ax(rx>R)=zeros(size(find(rx>R)));
        % Ensure all out of range values are zeroed
axm2(rx>R)=zeros(size(find(rx>R)));
axm1(rx>R)=zeros(size(find(rx>R)));
axp1(rx>R)=zeros(size(find(rx>R)));
axp2(rx>R)=zeros(size(find(rx>R)));

axunf(rx>R)=zeros(size(find(rx>R)));
axunf1(rx>R)=zeros(size(find(rx>R)));
axunf2(rx>R)=zeros(size(find(rx>R)));
axunf3(rx>R)=zeros(size(find(rx>R)));
axunf4(rx>R)=zeros(size(find(rx>R)));

clear corrsig1 rcvdsig1 sum3 tgt fl cl mod lmod pulselength
clear R0 c lmnt1 ind n f1 n1 r sq cs fcarrier rangeres
clear coarserr beamwidth ph sum4 corrsig2 rcvdsig2
clear env1 rem lrem index td sn indx rm lrm bmcorr env
save 'a:\exp9corr'      % This is where the useful products of all these efforts are stored.
```

APPENDIX B. IMAGE FORMATION CODE

The imaging algorithm was implemented in the following code (Remarks are in italics.):

% Load previously calculated data, constants, and correlation matrices

```
clear
load 'a:\exp9corr'
cl=ceil(rx);        % cl=upper index in corrsig matrix for a contour
clear rx
```

% Establish additional initial conditions

```
fldwidth=.1;        % Field width to be imaged without zero padding in meters
```

% Update helpful constants

```
Ccorr=C;               % # of azimuth samples in correlation matrices
C=ceil((2*maxrange*SBW+fldwidth)/lmntspc);
                       % total # of azimuth samples
Cimage=C-Ccorr;        % # azimuth samples in final image
```

%%%% This section reads the binary data stored in 'c:\temp\testpt1.dat'
%%%% for processing.

```
repeats=round(lmntspc*PRR/vplatform);
clear lmntspc fsample PRR vsound vplatform finerr SBW fldwidth
clear x tgt sq r n1 f1 ind n ph fref

fid=fopen('c:\temp\testpt1.dat','r');
a=fread(fid,'int16');                   % Read in the binary data to "a"
fclose(fid);

b=reshape(a,R*repeats,C);               % Put data in the format of a received signal matrix
clear a
rcvdsig=zeros(R,C);

for n=0:repeats-1                       % Perform coherent sum
  rcvdsig=rcvdsig+b(n*R+1:n*R+R,:);
end
clear b
% Received signal matrix is ready for display
```

```
% Perform Range Correlation
corrsig=zeros(R+5,C);  % Include three more rows of zeros so that out of
                       % range returns are zeroed
 for c=1:C             % This loop performs range compression for each column
   cs=conv(rcvdsig(:,c),cf);
   corrsig(1:R,c)=cs(floor(corrleng):R+floor(corrleng)-1);
 end
clear rcvdsig cs cf
% Range correlated signal matrix (real aperture image) is ready for display

% Add zero padding to allow processing with an inadequate acoustic field
corrsig=[zeros(R+5,ceil(Ccorr/2)),corrsig,zeros(R+5,ceil(Ccorr/2))];
C=C+2*ceil(Ccorr/2);
Cimage=C-Ccorr;

%%%% Unfocussed Processing

Ro=[1:Rx];
nterms=ceil(2.*sqrt(Ro.*lam));       % # terms to be averaged proportional to
sqrt(Ro*lam)
image=zeros(Rx,Cimage);              % final image matrix
for R0=1:Rx                          % Cycle through one range bin at a time
 R0
 sum3=zeros(1,C+nterms(R0)-1);
 for ns=1:nterms(R0)                 % Sum the terms for the average
   sum3=sum3+[zeros(1,ns-1),corrsig(ceil(R0*sampbin),:),zeros(1,nterms(R0)-ns)];
 end
 sum3=sum3.*(R0*R0/nterms(R0));
       % Multiply by only R0^1.5 vice R0^2 because longer range
       % targets are in the footprint for a longer distance (sqrt(R0*lam))
       % This gives targets a uniform amplitude response with range.

  image(R0,:)=sum3(ceil((nterms(R0)+Ccorr)/2):ceil((nterms(R0)+Ccorr)/2)+Cimage-1);
 end          % Go to next range bin
% Unfocussed image is now ready for display

%%%% Hybrid Processing

image=zeros(Rx,Cimage);   % final image matrix
corr=zeros(1,C);          % rows from corrsig at image range bins
axunf=fliplr(axunf);       % set up for azimuth convolution

for R0=1:Rx
 R0
 corr=corrsig(ceil(R0*sampbin),:).*(R0*sqrt(R0));
```

```
        % Multiply by only R0^1.5 vice R0^2 because longer range
        % targets are in the footprint for a longer distance (sqrt(R0*lam))
        % This gives targets a uniform amplitude response with range.
  cs=conv(corr,conj(axunf(R0,:)));      % Perform azimuthal correlation
        % Perform convolution with flipped axunf which is a correlation with
        % unflipped axunf for each range bin.
  image(R0,:)=cs(Ccorr:C-1);
end            % Go to next range bin
% Hybrid image is now ready for display


%%%%% 2-D Hybrid Processing

image=zeros(Rx,Cimage);        % final image matrix
corr=zeros(1,C);               % rows from corrsig at image range bins
corr1=zeros(1,C);
corr2=zeros(1,C);
corr3=zeros(1,C);
corr4=zeros(1,C);


                               % set up for azimuth convolution
%axunf=fliplr(axunf);           % axunf was already flipped in 1-D hybrid
axunf1=fliplr(axunf1);
axunf2=fliplr(axunf2);
axunf3=fliplr(axunf3);
axunf4=fliplr(axunf4);


for R0=1:Rx
  R0
  corr=corrsig(ceil(R0*sampbin),:).*(R0*sqrt(R0));
        % Multiply by only R0^1.5 vice R0^2 because longer range
        % targets are in the footprint for a longer distance (sqrt(R0*lam))
        % This gives targets a uniform amplitude response with range.
  cs=conv(corr,conj(axunf(R0,:)));
        % Perform convolution with flipped axunf which is a correlation with
        % unflipped axunf for each range bin.

        % Perform the same operations for the next four samples after the range bin.
  corr1=corrsig(ceil(R0*sampbin)+1,:).*(R0*sqrt(R0));
  cs1=conv(corr1,conj(axunf1(R0,:)));
  corr2=corrsig(ceil(R0*sampbin)+2,:).*(R0*sqrt(R0));
  cs2=conv(corr2,conj(axunf2(R0,:)));
  corr3=corrsig(ceil(R0*sampbin)+3,:).*(R0*sqrt(R0));
  cs3=conv(corr3,conj(axunf3(R0,:)));
  corr4=corrsig(ceil(R0*sampbin)+4,:).*(R0*sqrt(R0));
  cs4=conv(corr4,conj(axunf4(R0,:)));
        % Sum the row correlations to synthesize the 2-D correlation.
```

97

```
    image(R0,:)=cs(Ccorr:C-1)+cs1(Ccorr:C-1)+cs2(Ccorr:C-1)+cs3(Ccorr:C-
1)+cs4(Ccorr:C-1);
end         % Go to next range bin
% Limited 2-D hybrid image is now ready for display


%%%% Focussed 1-D Processing

image=zeros(Rx,Cimage);    % final image matrix
corr=zeros(Rx,Ccorr);      % azimuthal correlation along hyperbolas from data
                           % to be compared with simulation
for lmnt=1:Cimage          % Build image one azimuth line at a time
 lmnt
  for R0=1:Rx              % Build data correlation matrix one range line at a time
    corr(R0,:)=diag(corrsig(cl(R0,:),lmnt:lmnt+Ccorr))'.*(R0);
        % pick values out of corrsig from ranges / times
        % specified by rx to be compared with simulated values from ax
        % Multiply by only R0 vice R0^2 because longer range targets
        % are in the footprint for a linearly longer distance
        % This gives targets a uniform amplitude response with range
  end            % Go to next range bin
  image(:,lmnt)=(sum((corr.*ax).')).';  % Create this line of the final image
end            % Go to next azimuth line
% 1-Line focussed image is now ready for display


%%%% Focussed 2-D (5 line) Processing

image=zeros(Rx,Cimage);    % final image matrix
corr=zeros(Rx,Ccorr);      % azimuthal correlation along hyperbolas from data
                           % to be compared with simulation
corrm2=zeros(Rx,Ccorr);    % a separate data correlation matrix for each hyperbola
corrm1=zeros(Rx,Ccorr);
corrp1=zeros(Rx,Ccorr);
corrp2=zeros(Rx,Ccorr);


for lmnt=1:Cimage          % Build image one azimuth line at a time
 lmnt
  for R0=1:Rx              % Build data correlation matrix one range line at a time
    corr(R0,:)=diag(corrsig(cl(R0,:),lmnt:lmnt+Ccorr))'.*(R0);
        % pick values out of corrsig from ranges / times
        % specified by rx to be compared with simulated values from ax
        % Multiply by only R0 vice R0^2 because longer range targets
        % are in the footprint for a linearly longer distance
        % This gives targets a uniform amplitude response with range
        % Perform for each hyperbola
    corrm2(R0,:)=diag(corrsig(cl(R0,:)-2,lmnt:lmnt+Ccorr))'.*(R0);
```

```
    corrm1(R0,:)=diag(corrsig(cl(R0,:)-1,lmnt:lmnt+Ccorr))'.*(R0);
    corrp1(R0,:)=diag(corrsig(cl(R0,:)+1,lmnt:lmnt+Ccorr))'.*(R0);
    corrp2(R0,:)=diag(corrsig(cl(R0,:)+2,lmnt:lmnt+Ccorr))'.*(R0);
  end             % Go to next range bin

image(:,lmnt)=(sum((corr.*ax).')).'+(sum((corrm2.*axm2).')).'+(sum((corrm1.*axm1).')).'+
(sum((corrp1.*axp1).')).'+(sum((corrp2.*axp2).')).';   % Create this line of the final image
                                          % from results of all 5 correlations
end             % Go to next azimuth line
% 5-Line focussed image is now ready for display
```

LIST OF REFERENCES

Curlander, J.C. and McDonough, R.N., *Synthetic Aperture Radar, Systems & Signal Processing*, pp. 164-170, John Wiley & Sons, Inc., New York, New York, 1991.

Edde, B., *Radar Principles, Technology, Applications*, pp. 507-651, PTR Prentice Hall, Englewood Cliffs, New Jersey, 1993.

Fitch, J.P., *Synthetic Aperture Radar*, pp. 66-67, Springer-Verlag Inc., New York, New York, 1988.

Proakis, J.G. and Manolakis, D.G., *Digital Signal Processing, Principles, Algorithms, and Applications*, pp. 717, Macmillan Publishing Company, New York, New York, 1992.

Strum, R.D. and Kirk, D.E., *Contemporary Linear Systems Using MATLAB*, pp. 50,288, PWS Publishing Company, Boston, Massachusetts, 1994.

INITIAL DISTRIBUTION LIST

1. Defense Technical Information Center 2
    8725 John J. Kingman Rd., STE0944
    Ft. Belvoir, Virginia 22060-6218

2. Library, Code 013...2
    Naval Postgraduate School
    Monterey, California 93943-5101

3. Chairman, Code EC 1
    Department of Electrical and Computer Engineering
    Naval Postgraduate School
    Monterey, California 93943-5121

4. Dr. Donald L. Walters, Code PH/WE 5
    Department of Physics
    Naval Postgraduate School
    Monterey, California 93943-5117

5. Dr. Roberto Cristi, Code EC/CX 2
    Department of Electrical and Computer Engineering
    Naval Postgraduate School
    Monterey, California 93943-5121

6. Program Executive Officer, Mine Warfare 1
    2531 Jefferson Davis Highway
    Arlington, Virginia 22242-5167

7. LCDR Joseph Welter...2
    155 April Wind East
    Montgomery, Texas 77356