Report LMSC-D553148

PERFORMANCE OF THE PHASE CORRELATOR IN IMAGE GUIDANCE APPLICATIONS

Charles D. Kuglin
Palo Alto Research Laboratory
Lockheed Missiles & Space Company, Inc.
Palo Alto, California 94304

December 15, 1976

Final Report for Period 1 January – 1 October 1976

Approved for public release; distribution unlimited

Superseded AD-B037412
mc

Prepared for

Strategic Technology Office
ADVANCED RESEARCH PROJECTS AGENCY
1400 Wilson Boulevard
Arlington, Virginia 22209

CONTROL DATA CORPORATION
Digital Image Systems Division
2800 East Shakopee Road
Minneapolis, Minnesota 55540

PERFORMANCE OF THE PHASE CORRELATOR IN IMAGE GUIDANCE APPLICATIONS

Charles D. Kuglin

LMSC-D552148 December 15, 1976

Sponsored by:

Strategic Technology Office
Advanced Research Projects Agency
1400 Wilson Boulevard
Arlington, Virginia 22209

ARPA Order 2952

Monitored by:

Control Data Corporation
Digital Image Systems Division
2800 East Shakopee Road
Minneapolis, Minnesota 55440

Contract CDC 44MK328547

Palo Alto Research Laboratory
LOCKHEED MISSILES & SPACE COMPANY, INC.
A Subsidiary of Lockheed Aircraft Corporation
Palo Alto, California 94304

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE *(When Data Entered)*

ADA955485

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE *(and Subtitle)* PERFORMANCE OF THE PHASE CORRELATOR IN IMAGE GUIDANCE APPLICATIONS		5. TYPE OF REPORT & PERIOD COVERED Final Report 1 January - 1 October 1976 6. PERFORMING ORG. REPORT NUMBER LMSC-D553148
7. AUTHOR(s) Charles D. Kuglin		8. CONTRACT OR GRANT NUMBER(s) CDC 44MK328547
9. PERFORMING ORGANIZATION NAME AND ADDRESS Lockheed Palo Alto Research Laboratory Lockheed Missiles & Space Company, Inc. 3251 Hanover Street, Palo Alto, Ca. 94304		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS ARPA Order 2952
11. CONTROLLING OFFICE NAME AND ADDRESS Defense Advanced Research Projects Agency 1400 Wilson Boulevard Arlington, Va. 22209		12. REPORT DATE 15 December 1976 13. NUMBER OF PAGES
14. MONITORING AGENCY NAME & ADDRESS*(if different from Controlling Office)* Control Data Corporation Digital Image Systems Division 2800 East Old Shakopee Road Minneapolis, Minn. 55440		15. SECURITY CLASS. *(of this report)* Unclassified 15a. DECLASSIFICATION DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT *(of this Report)*

Approved for public release; distribution unlimited

17. DISTRIBUTION STATEMENT *(of the abstract entered in Block 20, if different from Report)*

18. SUPPLEMENTARY NOTES

19. KEY WORDS *(Continue on reverse side if necessary and identify by block number)*

Image Registration
Phase Correlation
Fourier Spectrum
Terminal Homing

20. ABSTRACT *(Continue on reverse side if necessary and identify by block number)*

A study has been performed to assess the feasibility of using a phase correlation map-matching processor for terminal guidance applications which require both automatic target acquisition and high accuracy. Phase correlation is a technique for extracting highly accurate estimates of the relative displacement of two images from the phase of their cross-power spectrum. The phase correlation peak height provides a sensitive measure of the degree of image similarity and, therefore, of the correlation performance to be expected. Phase correlation is a superior method for minimizing the effects of narrow bandwidth noise or noise having a similar spectral content to the

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE *(When Data Entered)*

noise-free image. This attribute was demonstrated by computer simulations using imagery with seasonal, illumination, and atmospheric differences, as well as differences in spectral band. A parametric study has verified that, by the use of appropriate phase filters, the technique can be made to accommodate moderate amounts of geometric distortion at the price of reduced accuracy. If higher precision is required, an image rectification step must be carried out prior to the correlation computation. This procedure has been shown to significantly increase phase correlation accuracy. The possibility of correctly positioning a small sensed image within a larger reference image, as required for target acquisition, has been demonstrated using phase correlation. A compact implementation, based upon the FFT and using a novel phase quantization procedure, can provide high data rates using available components.

PROJECT RESPONSIBILITY

Title: Performance of the Phase Correlator in Image Guidance Applications

Contractor: Lockheed Missiles and Space Company, Palo Alto Research Laboratory

LMSC Program Manager: Dr. James J. Pearson
(415) 493-4411 X45845

Project Leader: Mr. Charles D. Kuglin

Analysis and Simulation: Mrs. Pauline Wong
Mr. Roy F. Lunsford

DARPA Program Manager: Major James Karam

Contract Number: CDC 44MK328547 (ARPA ORDER 2952)

Contract Period: Jan 1, 1976 – October 1, 1976

FOREWORD

This report describes the results obtained from a nine-month study of the "Performance of the Phase Correlator in Image Guidance Applications," undertaken as part of DARPA's Continuous Terminal Homing Program. The program was performed by the LMSC Signal Processing Laboratory under subcontract to the Control Data Corporation (CDC 44MK328547). The objective of this investigation has been to assess the performance of a phase correlation map-matching system under the conditions anticipated for target acquisition and homing.

CONTENTS

ILLUSTRATIONS

TABLES

Section 1
INTRODUCTION AND SUMMARY

1.1 BACKGROUND

A correlation terminal guidance system is one in which the images obtained from an onboard sensor are compared with stored reference images of the target area in order to determine the steering commands needed to move the vehicle onto its prescribed trajectory and maintain it there until it reaches its blind range. The guidance process divides naturally into two sequential phases – acquisition and homing. In the acquisition phase, the correlation system takes over from the midcourse navigation system and identifies where within the error basket the sensed image falls. Once an initial fix has been achieved, a homing phase follows in which the vehicle is steered to bring the sensed image into its proper relationship to the target and maintains it there. In its simplest form, a two-dimensional image matching technique which measures the relative displacement between the sensed image and the stored reference image is needed for both of these phases. This scene alignment process is complicated by the fact that the reference image may be constructed using a different viewing geometry with a sensor operating in a different spectral region. Additional sources of error include sensor and processor noise, illumination and seasonal changes, and atmospheric effects.

In order to improve the accuracy of guidance systems which must operate under these conditions, the Defense Advanced Research Projects Agency (DARPA) is conducting a Continuous Terminal Homing development program. This report summarizes the results obtained from a study of the "Performance of the Phase Correlator in Image Guidance Applications," which was carried out as part of the DARPA program under subcontract to the Control Data Corporation.

The phase correlation method was developed by the LMSC Signal Processing Laboratory as a part of a continuing effort to explore advanced image matching concepts.

Phase correlation is based on the fact that for images differing from each other only by a translation, all of the information about that translation resides in the phase of their two-dimensional cross-power spectrum. In its simplest form, phase correlation extracts that information by taking the inverse Fourier transform of the Fourier phase spectrum. The resulting correlation surface contains a very sharp peak of predictable shape at the location corresponding to the tip of the translation vector. Since many of the important types of noise and noncommonality between the reference and sensed images affect different regions of the Fourier spectrum differently, it is often desirable to weight the phases in a frequency-dependent manner. In almost all cases, however, it is preferable to choose the weighting function in accordance with the noise characteristics, or the type and amount of distortion which must be accommodated, rather than allowing it to be determined by the scene content as in effect happens with conventional cross correlation.

In previous work, phase correlation had been shown to yield highly accurate displacement measurements between essentially similar images. The objective of the present contract was to assess the feasibility of using phase correlation for the terminal homing problem with emphasis on the acquisition phase, where position uncertainties and geometrical distortions raise questions of both missed detections and false alarms.

Some preliminary investigations of sensitivity to various sources of noise were performed and, in addition, an estimate of implementation complexity was made.

1.2 ORGANIZATION

The report is organized in the following manner.

A description of the phase correlation concept is given in Section 2 and a general analysis is presented which relates performance to the correlation peak height and the number of sample points. A discussion of the problems which arise when images are correlated with relative perspective differences leads naturally to the concept of low-pass phase weighting functions.

A number of computer simulations are described in Section 3 which give a preliminary estimate of the effect of various noise sources on correlator performance. Particular emphasis is placed on the expected errors due to uncompensated perspective differences.

Section 4 is devoted to a discussion of the computation algorithm and its efficient implementation in digital hardware. The concept of "phase quantization" is introduced, and it is shown how the application of this idea leads to both a simplified implementation and a method for storing reference maps in compressed form.

Several topics of importance to the target acquisition and homing problem are discussed in Section 5 including false match, rectification, and reference map requirements. The highlights of each of these sections and the main conclusions are summarized in the remainder of this section.

1.3 THE PHASE CORRELATION ALGORITHM

If $g_1(\vec{r})$ and $g_2(\vec{r})$ are two images whose two-dimensional Fourier transforms are $G_1(\vec{f})$ and $G_2(\vec{f})$, their cross-power spectrum is given by $G_1 G_2^*$ and the phase of that spectrum by:

$$e^{j\phi} = \frac{G_1 G_2^*}{|G_1 G_2^*|} \tag{1.1}$$

The filtered phase correlation function d is then given by

$$d = F^{-1}\left\{H(\vec{f})\, e^{j\phi}\right\} \tag{1.2}$$

where $H(\vec{f})$ is an arbitrary function of spatial frequency, and F^{-1} denotes the inverse Fourier transform. The weighting function H is normally taken to be unity unless the form of the noise or the expected geometrical distortion dictates some other choice. (Note that if $H(\vec{f})$ is chosen equal to $G_1 G_2^*$, cyclic cross correlation results.) The "whitening" of the cross-power spectrum which results from phase correlation, removes most of the scene dependence from the analysis and yields

a narrow peak shape which has the form of a "noisy," broadened delta function. The technique is insensitive to narrow bandwidth noise and is immune to convolutional degradations applied to both images. The known peak shape facilitates highly accurate interpolation to obtain displacement measurements with subpixel precision.

1.4 PERFORMANCE MODEL

A statistical performance model has been developed based on the assumption that the Fourier phase angle differences for a pair of uncorrelated images form a set of independent random variables uniformly distributed over the range $(-\pi, \pi)$ radians. If some degree of commonality exists between the two images, the phases have a correlated component as well as a random one. The total power (normalized to 1) in the correlation plane is then divided between a coherent peak of power A^2 (where A is the peak height) and a set of incoherent peaks accounting for the remaining power, $1 - A^2$. The amplitudes of the incoherent peaks for N_s samples can be approximated by a gaussian distribution with standard deviation

$$\sigma = \left(\frac{1-A^2}{N_s}\right)^{1/2} \tag{1.3}$$

From this model follow expressions for the rms displacement error (in pixels) and the probability of false match as functions only of A and N_s. These relationships are illustrated in the contour plots given in Fig. 1-1.

The peak height A is thus seen to provide a sensitive measure of the degree of commonality between the images and consequently of the correlation performance to be expected.

1.5 NOISE AND NONCOMMONALITY EFFECTS

Computer simulations were performed to assess the expected performance of phase correlation in the presence of the various types of noise likely to be encountered in the homing application. Among the types of noise investigated were sun angle effects,

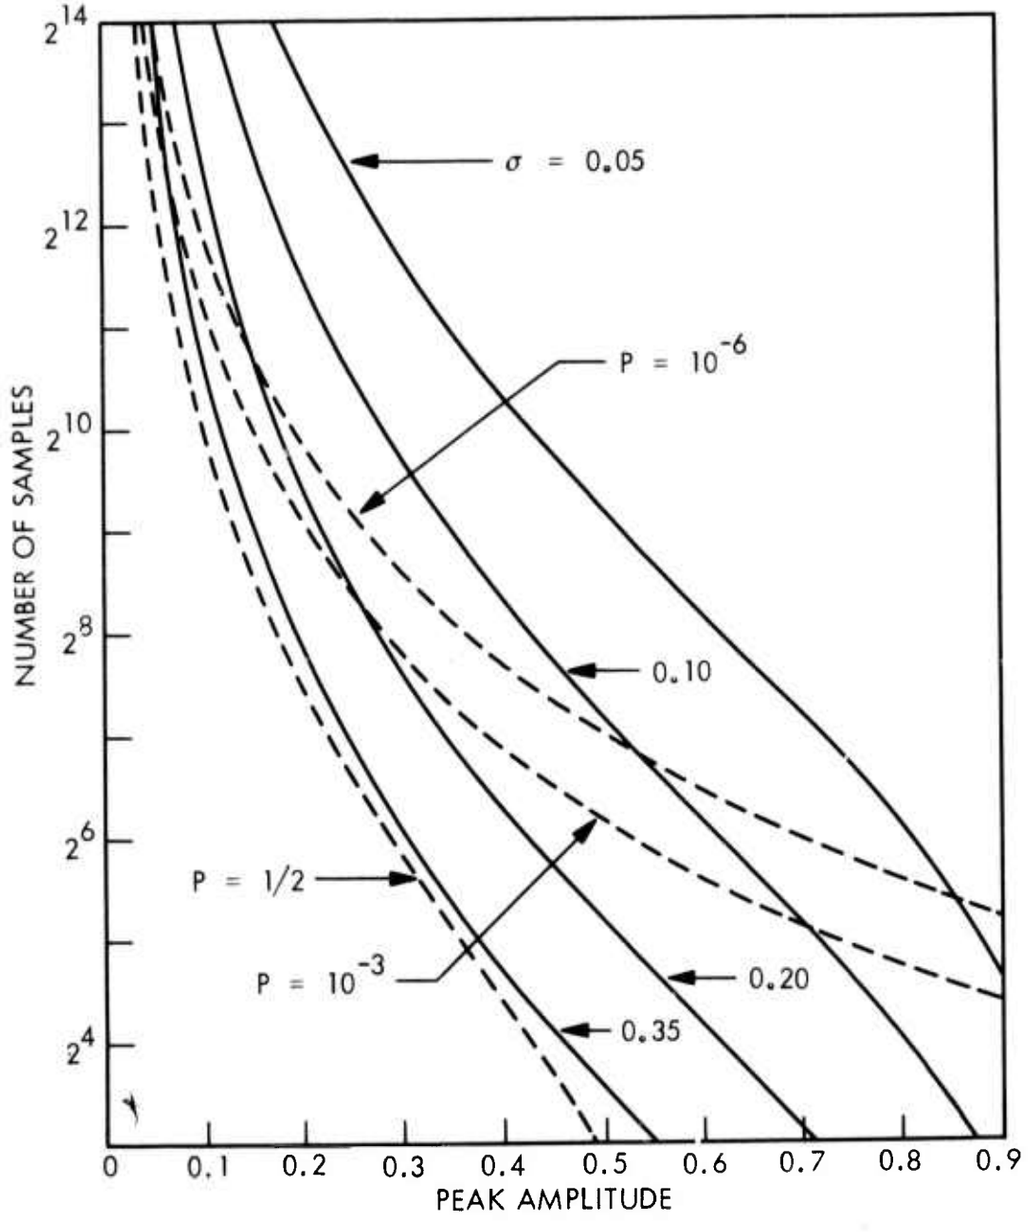

Fig. 1-1 Displacement Error and Probability of False Match Contours

seasonal changes, weather effects, spectral differences, and edge effects. Since phase correlation is well suited to minimizing the effect of either narrow bandwidth noise or noise having similar spectral content to the image (in contrast to cross correlation which is particularly adapted to the case of white noise), it would be expected to, and indeed did, perform well in the presence of all the above effects.

In addition to these studies, performance was determined for controlled amounts of independent random noise added to each image. The correct displacement was obtained to within a fraction of a resolution element for 64 x 64 element images with signal-to-noise ratios as low as -6 dB.

An example which illustrates differences due to several "noise" sources (translation, perspective variation, seasonal and illumination variations) is shown in Fig. 1-2. The phase correlation peak occurs in the right region, whereas the cross correlation result is substantially incorrect.

1.6 GEOMETRICAL DISTORTIONS

Geometrical distortions between images differ from other types of noise in that their effects are, in principle at least, largely removable. These distortions have the property that they affect the higher-frequency terms in the Fourier transform more than the lower-frequency ones. Two approaches are therefore possible. One is to ignore those transform elements which are most in error by filtering the cross-power spectrum. This has the effect of making the correlation process more tolerant to distortion, but at the cost of reduced displacement measurement accuracy, since much of the displacement information is discarded. The second approach involves correcting the distortion before performing the correlation. This method is certainly superior, since essentially all of the available information about the displacement vector can then be used (and, in some sense, it must be used if a unique displacement vector is to be defined). Its application involves prior knowledge of the distortion, however, and thus must await the development of an automatic distortion measurement technique. Work on image warping under the present contract has been limited to demonstrating that correction of a relative distortion, determined as yet using human intervention, does

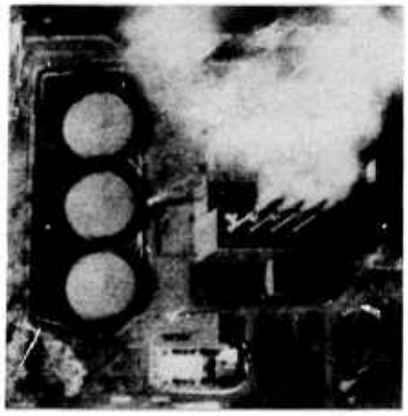

CLEAR

SNOW (RECTIFIED)

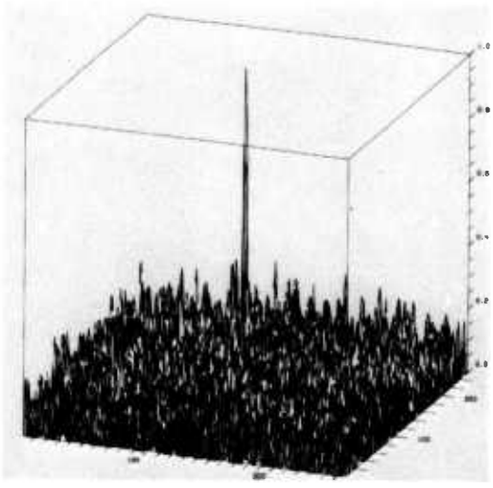

PHASE CORRELATION
($\alpha = 0$)

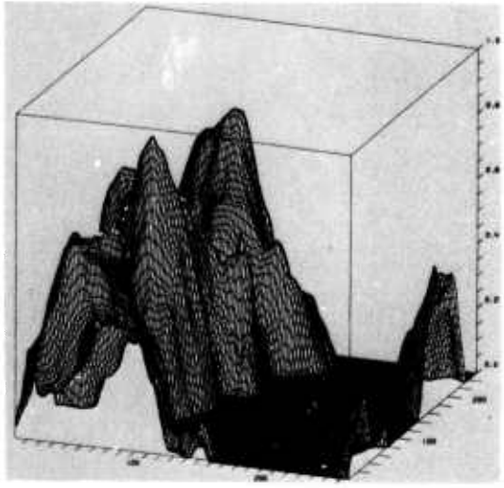

CROSS CORRELATION
($\alpha = 1$)

CORRELATION FUNCTION — $d_\alpha = F^{-1}\left\{|G_1 G_2^*|^\alpha e^{j(\phi_1 - \phi_2)}\right\}$

Fig. 1-2 Power Plant Correlations

indeed improve the displacement measurement. This work is illustrated by Fig. 1-3, which shows correlation peaks for a pair of 128 x 128 element images with and without rectification. The nearly 300% increase in peak amplitude obtained would result in a significant increase in displacement accuracy and decrease in the probability of false acquisition.

During the current program, a more detailed study of the filtering approach was performed. The results of a parametric study of filter performance consisting of some 160 experiments that use unfiltered phase correlation, cross correlation, and three types of filtered phase correlation are summarized in Figs. 1-4 through 1-7. Here, the error magnitudes are plotted as a function of the amount of rotation or scale change for two different amounts of overlap. The errors obtained are often substantially greater than predictions based upon the theoretical model. This is due to the fact that distorted images have no true match point and, thus, the correlation peak may follow some particularly strong feature within the scene as the amount of distortion is changed.

1.7 TARGET ACQUISITION

As an indication that techniques for reference map search procedures can be developed and aspect changes accommodated, the simulated acquisition computation illustrated in Figs. 1-8 and 1-9 was performed. The 128 x 128 element sensed image was correctly positioned within the 512 x 512 element reference map using 25 overlapping subarrays. The reference and sensed maps differed by a 7 1/2 degree change in azimuth and, therefore, a crude mosaic correlation computation was used to rectify the sensed map. A false match probability of less than 10^{-6} was predicted by the statistical model.

1.8 IMPLEMENTATION

The phase correlation algorithm consists of the following steps:

- The discrete two-dimensional forward Fourier transform is computed for each image.

PHASE CORRELATION – ORIGINAL SENSED MAP

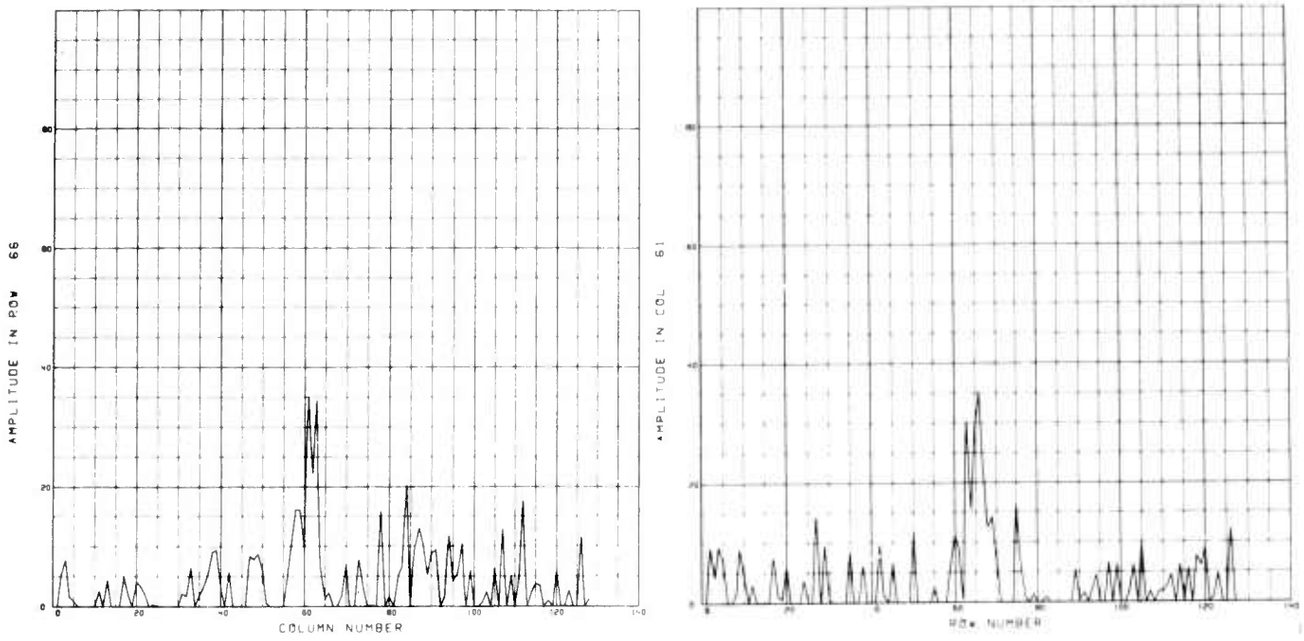

a. Column Cross Section

b. Row Cross Section

PHASE CORRELATION – RECTIFIED SENSED MAP

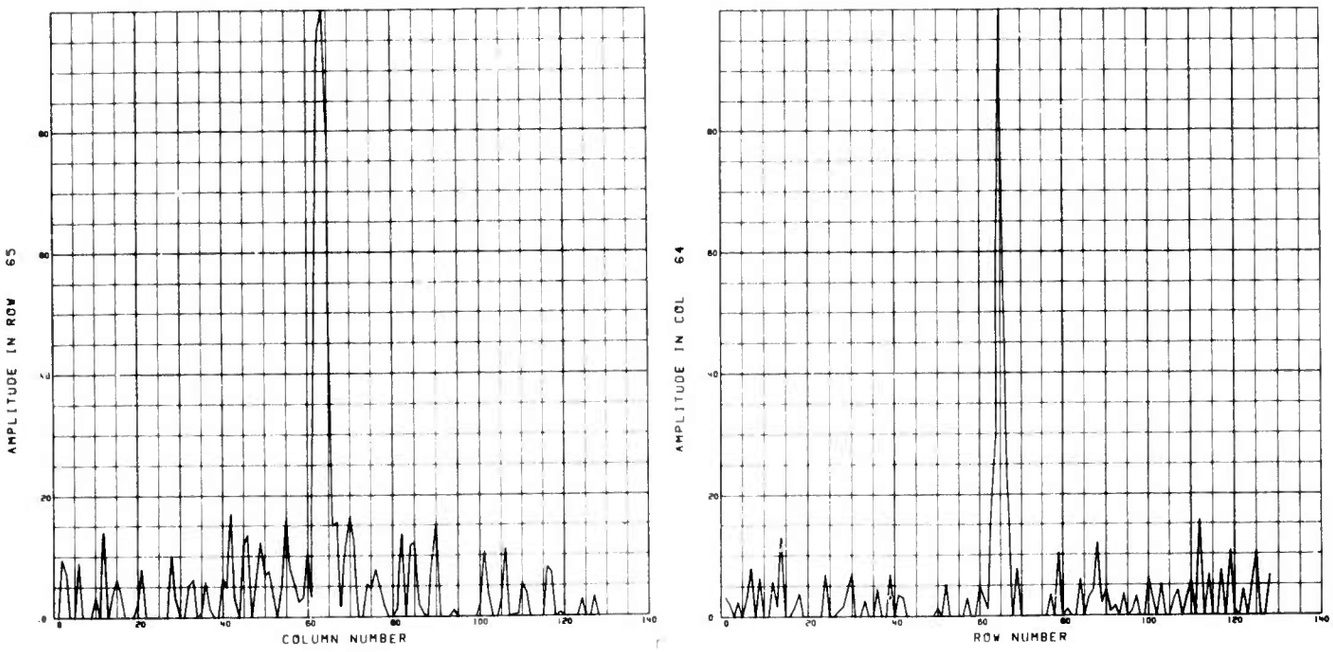

a. Column Cross Section

b. Row Cross Section

Fig. 1-3 Effect of Image Rectification on Phase Correlation Peak

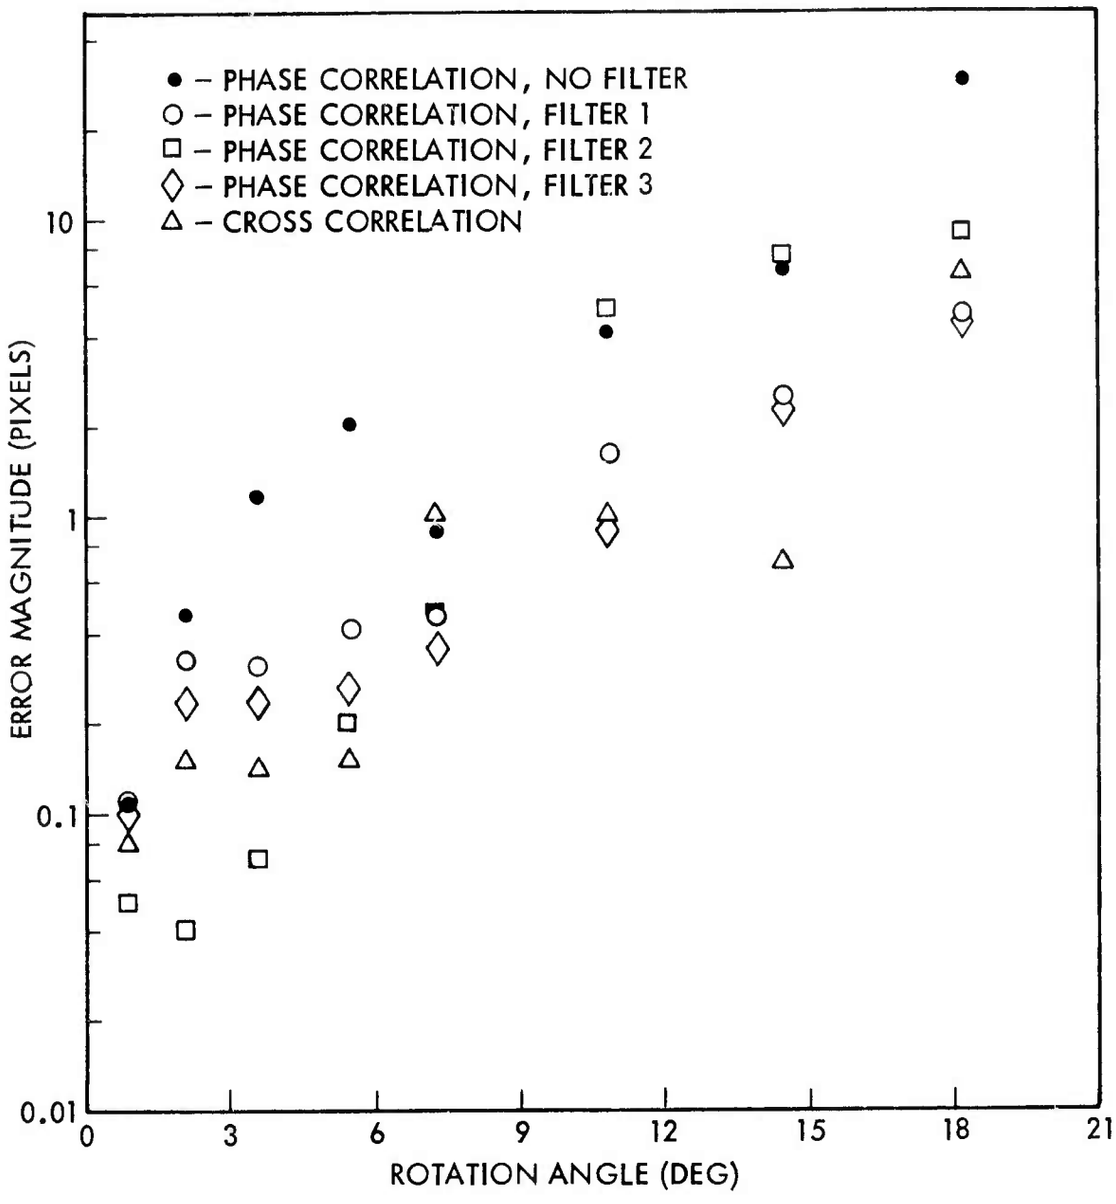

Fig. 1-4 Rotation Simulations – 100% Overlap

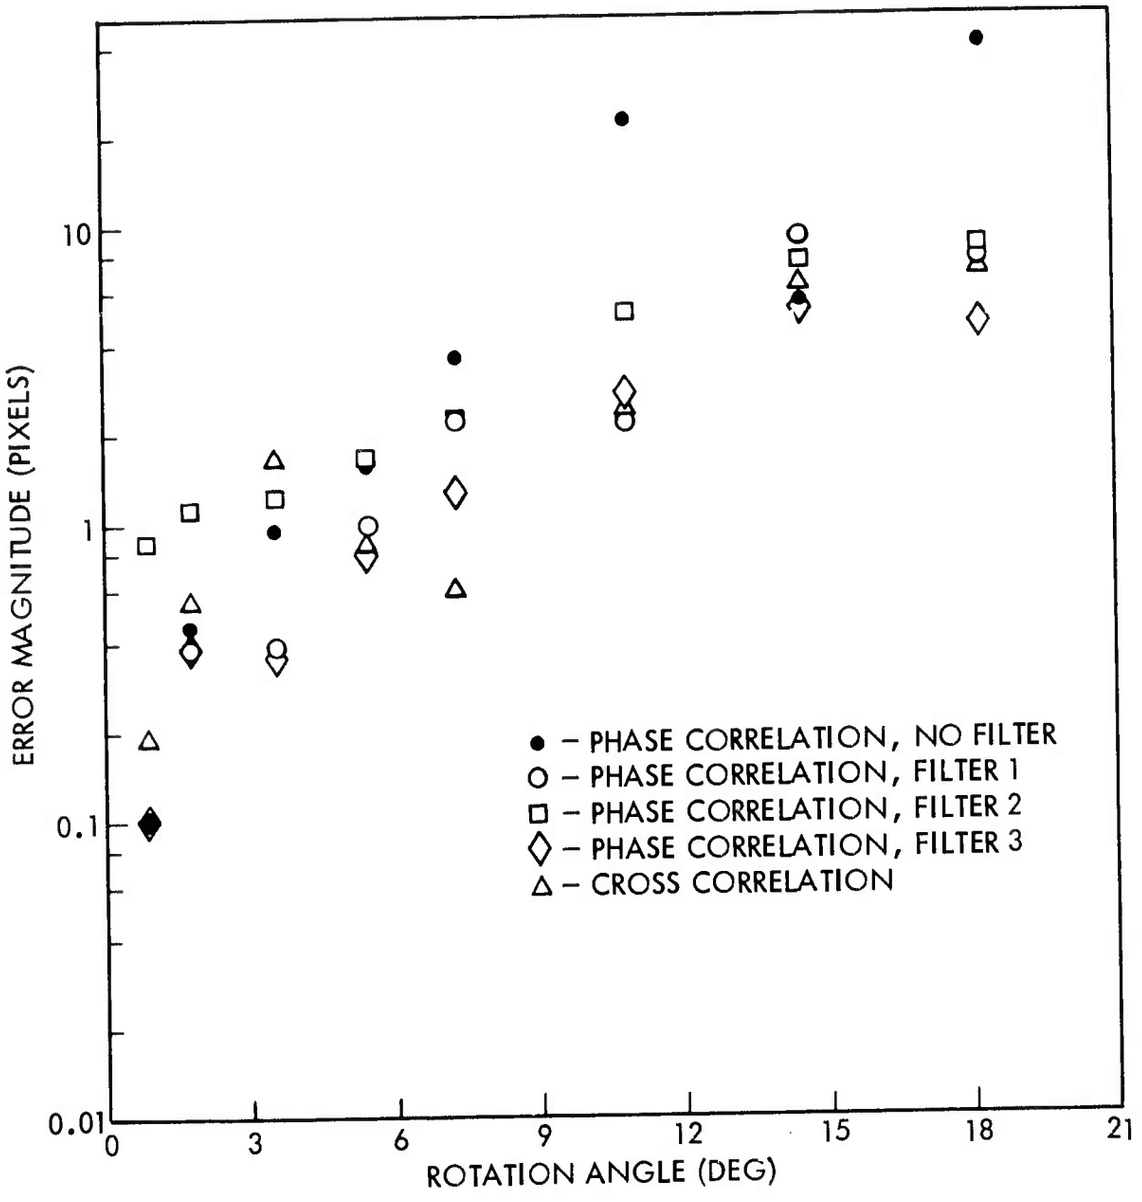

Fig. 1-5 Rotation Simulations – 80% Overlap

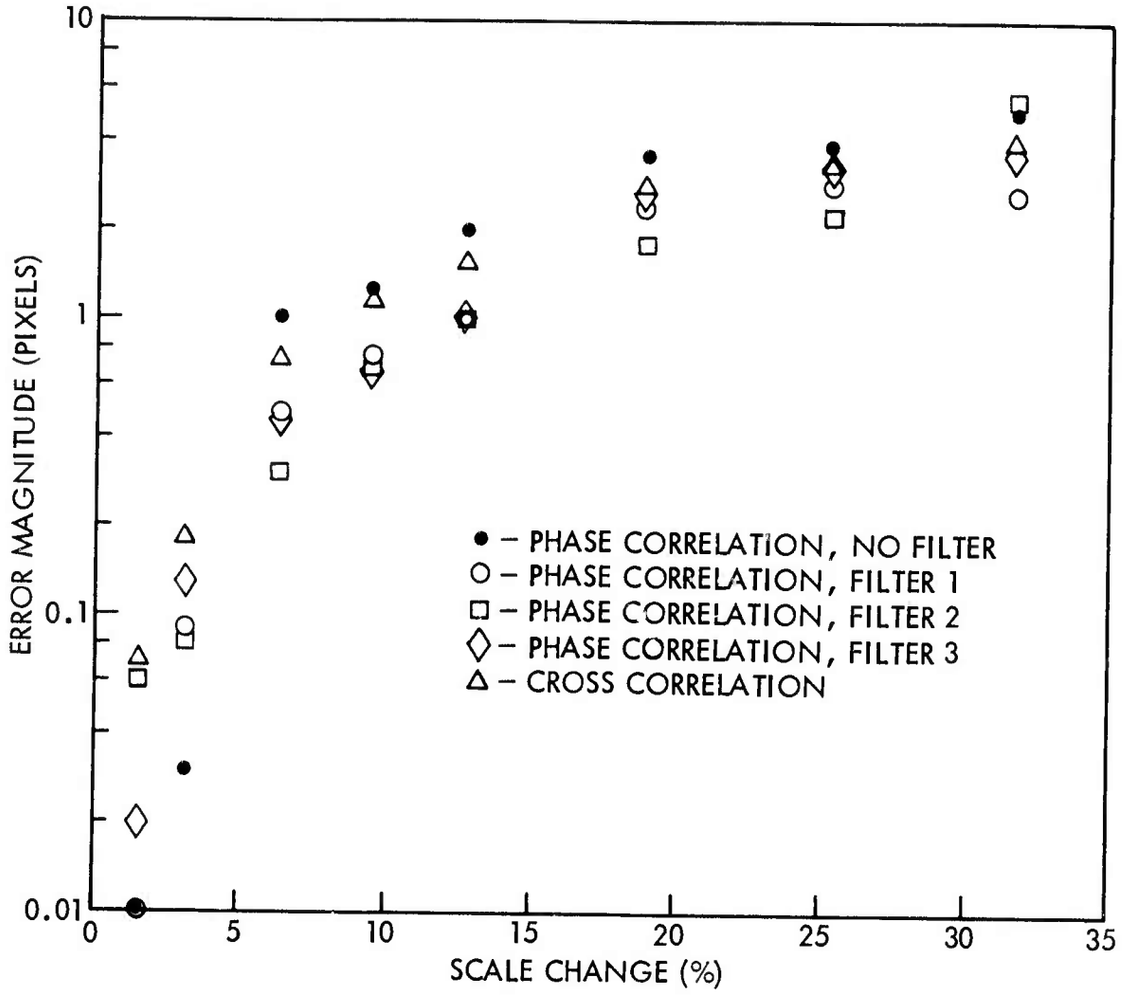

Fig. 1-6 Scale Change Simulations – 100% Overlap

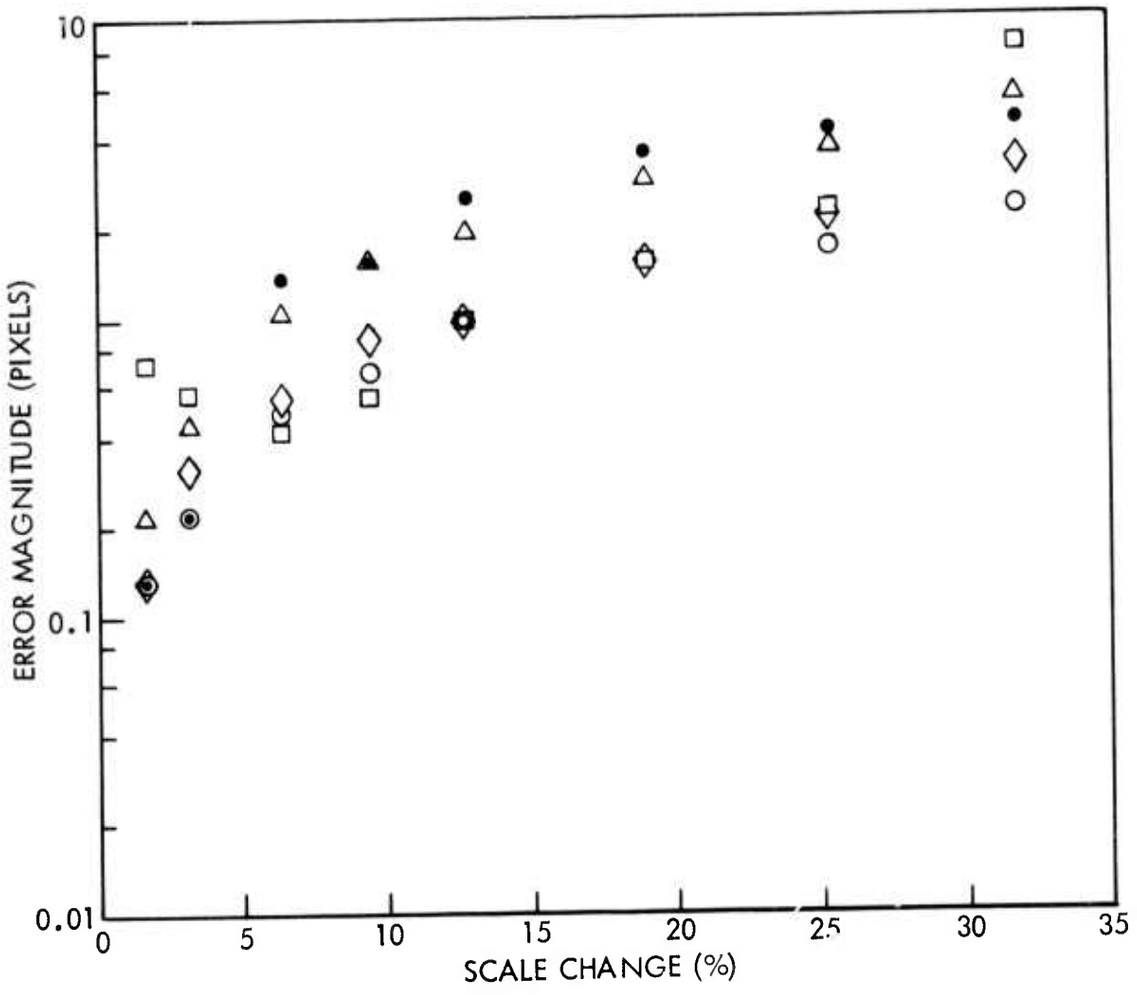

Fig. 1-7 Scale Change Simulations – 80% Overlap

a. Reference Map (512 × 512)

b. Reference Map Section (128 × 128)

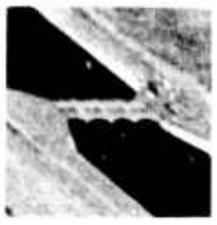

c. Sensed Map (128 × 128)

Fig. 1-8 Imagery Used for Acquisition Simulation

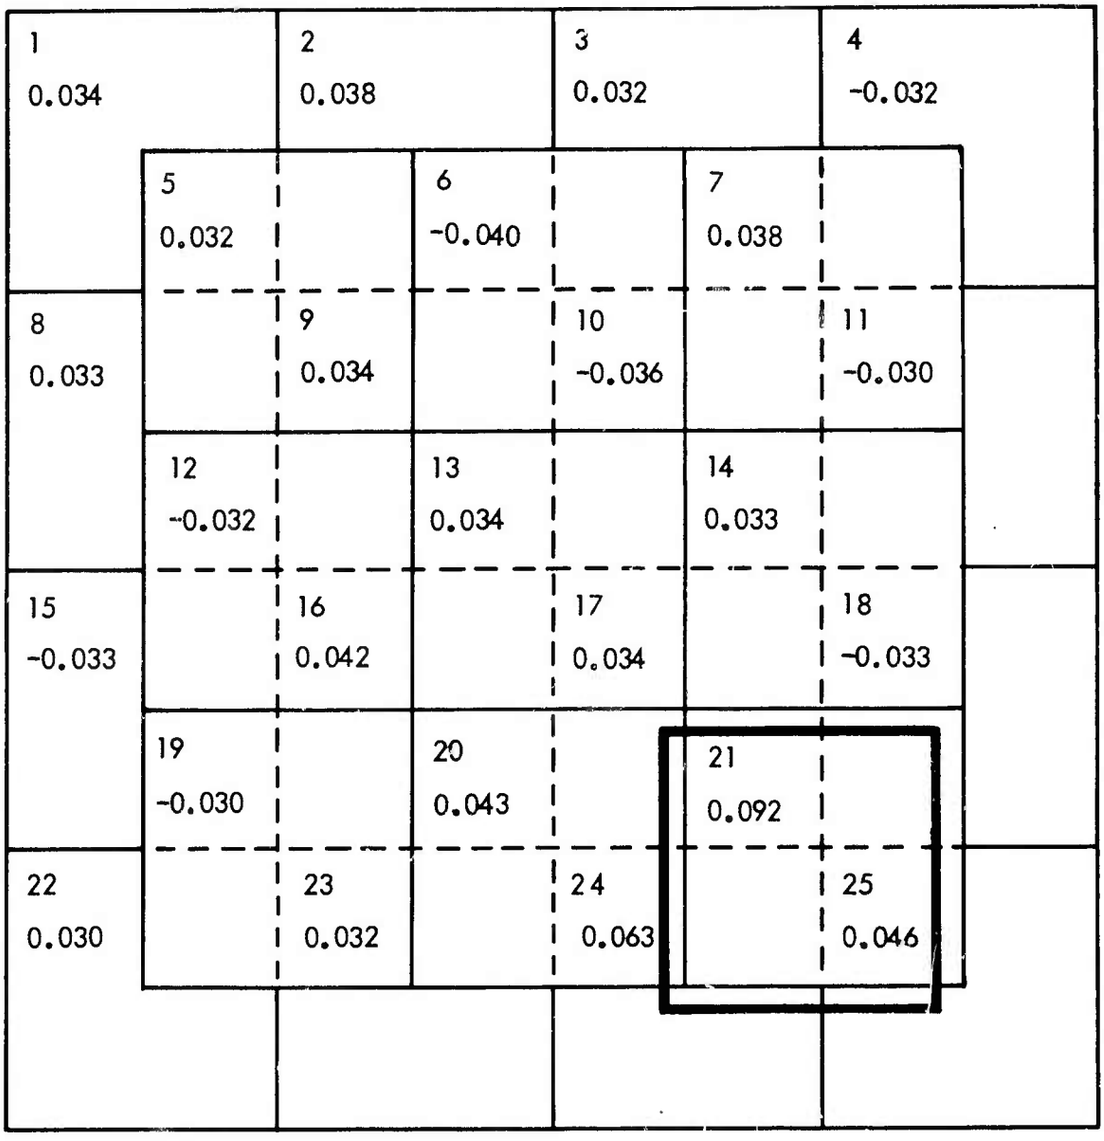

Fig. 1-9 Acquisition Simulation Results Showing Subarray Positions and Corresponding Correlation Peak Amplitudes

- The phase of each transform component is extracted and corresponding phases are subtracted.
- The phase correlation function is obtained from the inverse transform of the resulting phase matrix.
- Finally, the correlation array is searched for a maximum. The location of the maximum directly gives the displacement. Nonintegral displacements can be estimated using interpolation.

This procedure can be implemented with circuitry somewhat less complicated than that required for cyclic cross correlation. By far the greatest part of the computation is involved in taking the forward and inverse transforms, steps required by both correlation techniques. The relative advantage enjoyed by phase correlation arises from the fact that the phase matrices for sensed and reference images can be computed separately and quantized to only three bits. The phases can be obtained to this coarseness by only three binary decisions involving the real and imaginary parts of the transform elements. The phase differences are computed either by a three-bit subtraction or by a table lookup. Since the phase matrix is hermetian, only half of the phases need be stored. Thus, phase quantized reference imagery requires a storage capacity of only 1 1/2 bits/pixel.

The size and cost of the hardware required depend on the array size and the speed required. As an example, it is estimated that a 64 x 64 element image size and a processing rate of 100 frames per second would require about 300 standard TTL chips (averaging about 60 gates/chip).

1.9 RECOMMENDATIONS FOR FUTURE STUDY

In the course of the present contract, a number of problem areas have been identified as requiring further investigation. A recommended plan for future study would include the following tasks:

- Develop techniques for automatic image rectification
- Develop correlation methods suitable for three-dimensional targets
- Optimize spectral weighting functions (filters) for various distortions and construct an effective adaptive filtering algorithm

- Carry out a more detailed parametric analysis in which combinations of different noise sources are used
- Develop a theoretical model to predict phase correlator performance in the presence of different types of noise
- Determine effective image preprocessing techniques for phase correlation
- Optimize reference map search procedure

Section 2
THE PHASE CORRELATION CONCEPT

2.1 THE PHASE CORRELATION FUNCTION

The phase correlation algorithm (Refs. 1 and 2) is based upon the fact that the information pertaining to the displacement of two images resides in the phase of the cross power spectrum. This concept has been generally overlooked in previous treatments of the image-matching problem, which have usually considered cross correlation or closely related variations. Although cross correlation is an optimal procedure for additive white noise, many important sources of image noise cannot be represented by this model. It is therefore important to develop alternative correlation methods that are effective for realistic noise sources. The phase correlation method is particularly well suited to handle many of the differences that can exist between reference and sensed maps or between two sensed maps.

The application of the phase correlation technique results in a sharp peak at the point of registration on the order of one resolution element in width. The method is relatively scene-independent and is therefore very useful for aligning images taken with different sensors or under varying conditions of illumination or weather. In addition, the results are insensitive to narrow bandwidth noise and convolutional image degradations. Large relative displacements can be accurately determined without the need to have one image wholly contained within the other. Given a sufficiently high signal-to-noise ratio (SNR), it is possible to measure non-integer displacements through the use of interpolation, and it is particularly advantageous to use the phase correlation surface for this purpose since it is characterized by a very sharp, symmetric peak of known functional form. A detailed description of the phase correlation function is given below.

Consider two images g_1 and g_2 that are displaced relative to one another by a vector $\vec{L}$. Assuming the idealized case of infinite or cyclically shifted images,

$$g_2(\vec{r}) = g_1(\vec{r} + \vec{L}) \tag{2.1}$$

where the functions define brightness values at points in two-dimensional space denoted by $\vec{r}$. In terms of linear filter theory, g_1 is related to g_2 through a convolution with the shifted delta-function $\delta(\vec{r} + \vec{L})$ so that

$$g_2(\vec{r}) = g_1(\vec{r}) \circledast \delta(\vec{r} + \vec{L}) \tag{2.2}$$

where $\circledast$ denotes convolution.

The shifted delta-function, therefore, provides the exact displacement information in a manner entirely independent of the two images. The following procedure describes how the delta-function is extracted from the two images. Although these results are stated in terms of two-dimensional images, the phase correlation method can be readily applied to signals having an arbitrary number of dimensions.

When two sampled scenes are cyclically shifted relative to each other by a vector $\vec{L}$, the Fourier transform matrix of the first, G_1, is related to that of the second, G_2, by the equation (Ref. 3)

$$G_2(\vec{f}) = G_1(\vec{f})\, e^{j2\pi \vec{f} \cdot \vec{L}} \tag{2.3}$$

Here, the vector $\vec{f}$ refers to the spatial frequency of a particular transform element. If the element-by-element ratio of these two transforms is taken, the resulting matrix R is given by:

$$R(\vec{f}) = \frac{G_2(\vec{f})}{G_1(\vec{f})} = e^{j2\pi\vec{f}\cdot\vec{L}} \tag{2.4}$$

The inverse Fourier transform of the complex conjugate R^* then has the following simple form in the space domain:

$$F^{-1}\{R^*\} = \delta(\vec{r} - \vec{L}) \tag{2.5}$$

That is, the only nonzero point is located at the position of the shift vector. In the case of finite images sampled to $N_s = N \times M$ pixels, Eq. (2.5) takes the form of a sampled, periodic sinc function:

$$d_{nm}(\ell_x, \ell_y) = S_N(n - \ell_x)\, S_M(m - \ell_y) \tag{2.6}$$

where $1 \le n \le N$, $1 \le m \le M$, and ℓ_x, and ℓ_y are the components of the vector $\vec{L}$ in units of resolution elements. The S-function is defined as

$$S_K(\ ) = \sin\pi t / \tan\frac{\pi}{K}t\ . \tag{2.7}$$

Since d is a <u>known,</u> continuous function of ℓ_x and ℓ_y, it is therefore possible to determine displacements to within a fraction of a resolution element by examing the form of the peak. Note that if the displacement corresponds to an integral number of resolution elements, say (n_o, m_o), then

$$d_{nm} = N_s\, \delta_{n,n_o}\, \delta_{m,m_o} \tag{2.8}$$

where $\delta_{j,k}$ denotes a "Kronecker δ."

Although the above relations hold rigorously only for cyclically shifted scenes, they are surprisingly well satisfied for shifted images taken from a continuous scene if the algorithm is slightly modified. In general, for real imagery, $|G_1| \neq |G_2|$ because of a wide variety of effects which can cause geometrical and gray shade differences between images.

The R-matrix is then no longer a pure phase matrix and the phase correlation function d is dominated by those spectral components for which $|G_1| << |G_2|$. The solution to this problem is as follows.

Since each term in the spectrum of the shifted delta function (2.6) has magnitude unity, this condition is imposed on the ratio matrix R by defining a new matrix D called the "phase matrix" such that:

$$D = \frac{R^*}{|R|} \tag{2.9}$$

$$= \frac{G_1 G_2^*}{|G_1 G_2^*|} \tag{2.10}$$

$$= e^{j\Phi} \tag{2.11}$$

The product $G_1 G_2^*$ is known as the "cross-power spectrum." The D matrix, therefore, represents the <u>phase</u> of the cross-power spectrum, where Φ is a frequency-dependent phase angle. The correlation matrix is then given by:

$$d = F^{-1}\{D\} \tag{2.12}$$

For cyclically shifted images, $\Phi = -2\pi \vec{f} \cdot \vec{L}$. For real imagery we can, without loss of generality, define a vector function of frequency, $\vec{L}_f(\vec{f})$, so that Eq. (2.11) can be written as

$$D(\vec{f}) = e^{-j2\pi \vec{f} \cdot \vec{L}_f} \tag{2.13}$$

This, then, resembles the spectrum of a shifted delta function except that the "displacement" $\vec{L}_f$ is now a function of frequency. It might be expected that $\vec{L}_f$ has a mean value $\vec{L}$ and a standard deviation which increases with decreasing image congruence. When the inverse transform [Eq. (2.12)] is taken, the contributions from the mean value add coherently ($\sim N_s$), while the frequency-dependent variations add incoherently ($\sim\sqrt{N_s}$). The result has the form of a "noisy," broadened delta function of amplitude $\leq N_s$. In general, the correlation function cannot be expressed expressed in analytic form. It is useful, however, to fit the correlation surface in the vicinity of the peak by a function of the form:

$$p(x,y) = A \text{ sinc } \alpha_x x \text{ sinc } \alpha_y y \tag{2.14}$$

where x and y are measured from the peak maximum, and the constants α_x and α_y are peak-broadening factors which are ≤ 1.

The relationship between phase correlation and cyclic cross correlation can be readily established. The cross correlation function, c, can either be expressed as a direct convolution of the two images in the spatial domain

$$c = g_1 \circledast g_2 \tag{2.15}$$

or, using the convolution theorem, as the inverse Fourier transform of the cross-power spectrum:

$$c = F^{-1}\{G_1 G_2^*\} \tag{2.16}$$

The phase correlation function is obtained by taking the inverse Fourier transform of the phase of the cross-power spectrum. Applying the convolution theorem, the phase correlation function [Eq. (2.12)] can be written in the spatial domain as

$$d = g_1' \circledast g_2' \tag{2.17}$$

where

$$g'_k = g_k \circledast F^{-1}\left\{|F\{g_k\}|^{-1}\right\}, \quad k = 1,2 \tag{2.18}$$

The primed images are "phase images" obtained by setting the spectral amplitudes to unity and leaving the phase information unchanged. Equation (2.18) shows that the phase image is related to the original image through a convolution with the inverse Fourier transform of the reciprocal spectral magnitudes. It therefore follows that the cyclic cross correlation function and the phase correlation function are related by

$$c = d \circledast F^{-1}\{|G_1 G_2^*|\} \tag{2.19}$$

Thus, the phase correlation method can be thought of as a nonlinear processing algorithm which combines the following two operations:

(1) Each image is converted into a phase pattern which emphasizes the geometric characteristics of the scene and has a correlation length on the order of one resolution element.
(2) The phase patterns are correlated using cyclic cross correlation in order to determine the match point.

Although the individual computations described by Eqs. (2.17) and (2.18) would be extremely time consuming, the combined computation given by Eq. (2.12) can be efficiently implemented using the Fast Fourier Transform (FFT) (Ref. 4).

The phase correlation algorithm may be generalized by introducing an arbitrary weighting function $H(\vec{f})$ in the spatial frequency domain giving

$$d_H = F^{-1}\{He^{j\Phi}\} \tag{2.20}$$

For instance, if the weighting function is of the form $|G_1 G_2^*|^\alpha$, the resulting family of correlation algorithms includes both phase correlation ($\alpha = 0$) and cyclic cross

correlation ($\alpha = 1$). Immunity to various types of noise or image distortions may be provided by the proper choice of weighting function. This concept will be discussed more fully in Section 2.2.5.

2.2 GENERAL PERFORMANCE ANALYSIS

2.2.1 Phase Correlation SNR

The phase correlation signal-to-noise ratio (SNR) can be expressed quite simply as a function of the peak amplitude and the square root of the total number of sample points. This important result can then be used to estimate the probability of false match and the interpolation accuracy.

Consider first the case of two images sampled at N_s points and having no common features. The phase correlation function is then of the form

$$d = F^{-1}\{e^{j\phi}\} \tag{2.21}$$

where ϕ can be considered to be a random variable uniformly distributed over 2π radians. The N_s d-matrix amplitudes (normalized to unit total power) can then be approximately represented by gaussian random variables having zero mean and standard deviation equal to $N_s^{-1/2}$.

In the case of two images which have some degree of congruence, the power is divided between a coherent peak located at the point of image registration (signal) and the incoherent peaks (noise) resulting from the random component of the phases. If the peak amplitude is denoted by A, where $|A| \leq 1$, then the signal power is A^2 and the noise power is $1 - A^2$. The standard deviation of the noise is thus:

$$\sigma = \left(\frac{1 - A^2}{N_s}\right)^{1/2} \tag{2.22}$$

and the phase correlation SNR is

$$\psi = \frac{A}{\sigma} \tag{2.23}$$

$$= A\left(\frac{N_s}{1 - A^2}\right)^{1/2} \tag{2.24}$$

The incoherent peak amplitude distribution for a shifted pair of 64×64 element images is shown in Fig. 2-1. In this case, the coherent peak amplitude is 0.77. The predicted distribution using a gaussian model with zero mean and standard deviation computed from Eq. (2.22) is shown for comparison. The agreement is quite good.

2.2.2 Probability of False Match

A useful expression for the probability of false match can be derived if the validity of the gaussian model for the noise amplitude distribution function is assumed:

$$p(x) = \frac{1}{\sqrt{2\pi\sigma}}\, e^{-x^2/2\sigma^2} \tag{2.25}$$

where σ is given by Eq. (2.22). The probability that a noise peak will be $\geq$ some positive threshold A_o is given by

$$P = 1 - [1 - q(A_o)]^{N_s} \tag{2.26}$$

where

$$q(A_o) = \int_{A_o}^{\infty} p(x)\, dx \,. \tag{2.27}$$

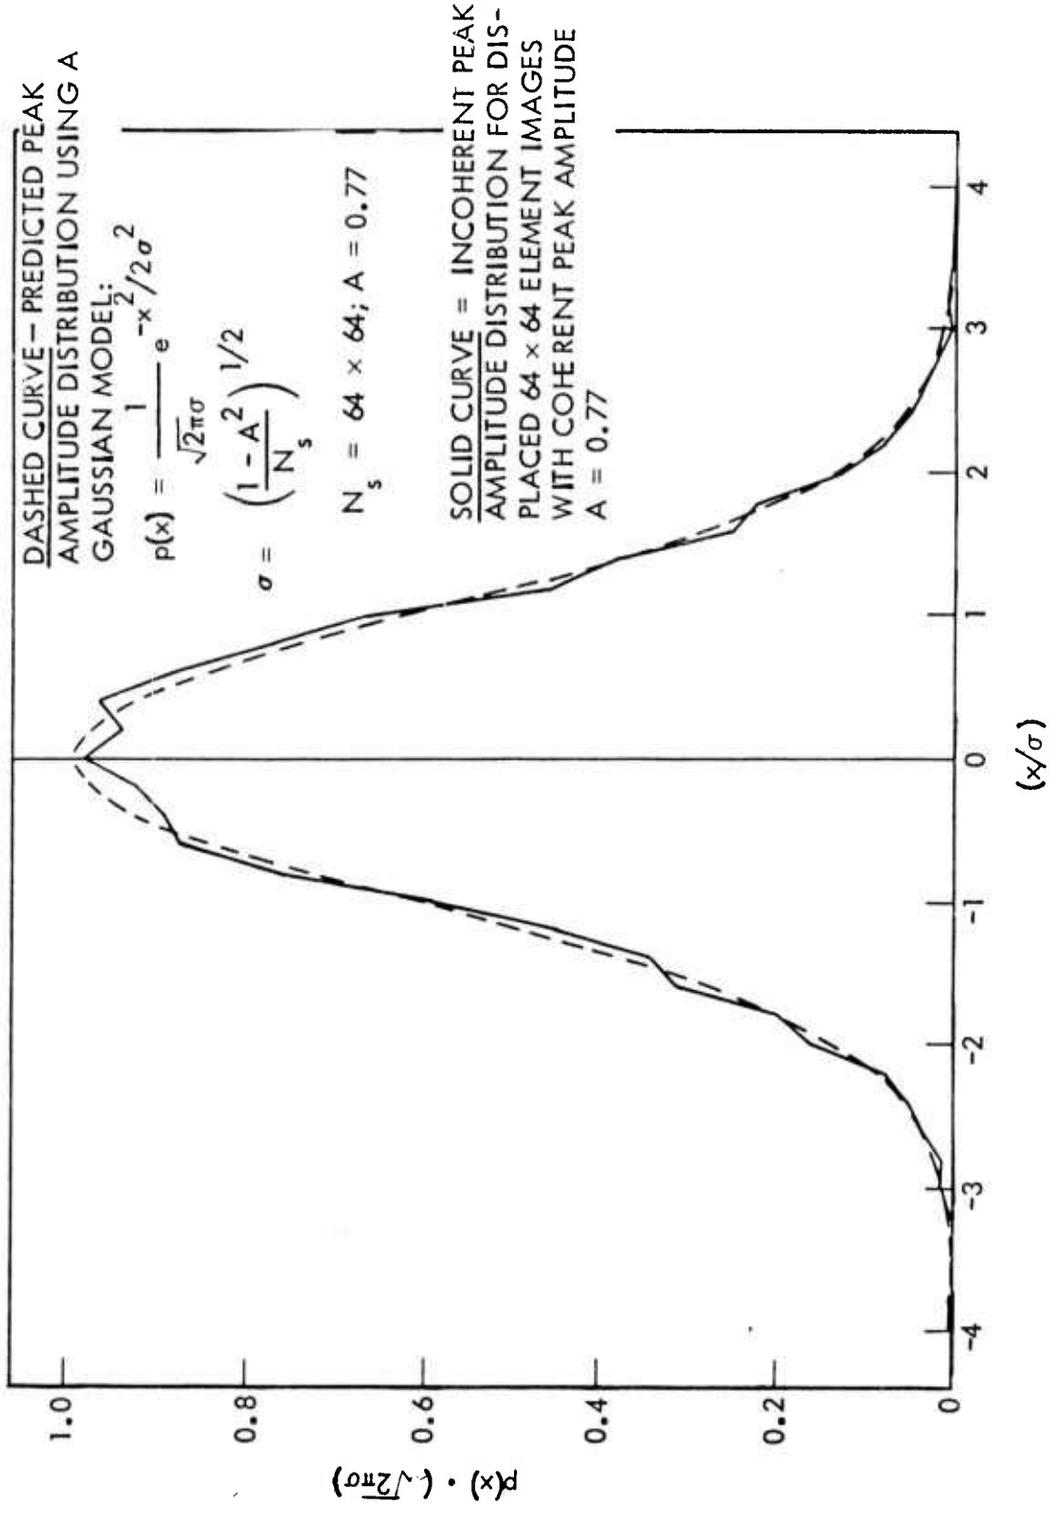

Fig. 2-1 Phase Correlation Incoherent Peak Amplitude Distribution

If it is assumed that

(1) $P \ll 1$

(2) $\sigma \approx N_s^{-1/2}$ (worst case)

(3) $A_o \sqrt{N_s} \gg 1$

the probability of false match can be written as

$$P \approx \sqrt{\frac{N_s}{2\pi}} \quad \frac{e^{-N_s A_o^2/2}}{A_o} \tag{2.28}$$

The error probability is a very strong function of peak amplitude. As an example, for a 64×64 element image ($N_s = 4096$), the error probability is 0.3 for $A_o = 0.06$ and 3×10^{-7} for $A_o = 0.1$. The phase correlation method thus allows one to distinguish answers that have a high probability of being correct from those that have a high probability of being incorrect by establishing an a priori threshold that is independent of the scene content and illumination.

The a priori prediction of the probability that a correct peak exceeding the noise threshold will actually exist is a much more difficult matter. This computation involves:

(1) The identification of all the various parameters (q_i, i = 1, K) which can produce differences between scenes

(2) An estimated joint probability distribution $p(q_i)$ for the parameters q_i

(3) An estimate of the dependence of the peak amplitude upon the scene "noise" parameters: $A(q_i)$

Some preliminary results which relate to this problem are discussed in Section 3.

2.2.3 Interpolation Precision

The correlation precision can be significantly increased through the use of interpolation with only a slight increase in overall computation time. Given a particular interpolation algorithm, the interpolation accuracy can be readily established since the noise amplitude distribution function can be closely approximated by Eq. (2.25).

In addition to the number of sample points and the peak amplitude, the accuracy will in general depend upon the peak width, the magnitude of the noninteger displacement, and the interpolation method.

A particular quadratic interpolation method has been evaluated for use with the phase correlation peak. Because of the sharpness of the peak, amplitudes at neighboring sample points usually do not have a sufficiently high SNR to provide accurate interpolated peak values. In addition, the functional form of the peak deviates substantially from a quadratic curve at distances on the order of one resolution element. These problems can be avoided by computing four nearest neighbor sample points about the peak spaced at $\frac{1}{2}$ resolution element:

$$\begin{array}{ccc} & x & \\ & p_{.5,0} & \\ x & o & x \\ p_{0,-.5} & p_{0,0} & p_{0,.5} \\ & x & \\ & p_{-.5,0} & \end{array}$$

Quadratic interpolation is then used to estimate the fractional displacement. For instance, in the horizontal direction, the noninteger part of the displacement in units of resolution elements is given by:

$$\delta = \frac{(p_{0,.5} - p_{0,-.5})}{4(2p_{0,0} - p_{0,.5} - p_{0,-.5})} \tag{2.29}$$

As the correlation peak amplitude decreases from unity, the peak tends to broaden and the phase correlation function in the vicinity of the peak can be fit by a function of the form:

$$p(x,y) = A \text{ sinc } \alpha_x x \text{ sinc } \alpha_y y \tag{2.30}$$

where x and y are measured from the peak maximum. The constants α_x and α_y are ≤ 1. A computer simulation using varying amounts of photon noise has indicated that $\alpha_x \approx \alpha_y = \alpha$, and that

$$\alpha \approx A^{1/4} . \tag{2.31}$$

If it is assumed that:

(1) Eq. (2.30) describes the correlation peak neighborhood

(2) Points separated by one or more resolution elements have independent random variations described by Eqs. (2.22) and (2.25)

(3) The peak broadening is described by Eq. (2.31)

(4) All noninteger displacements are equally likely

then the rms displacement error in δ, averaged over all x and y fractional displacements, is given by

$$\sigma_{ave} = \frac{1}{2} \left(\frac{1 - A^2}{N_s A^3} \right)^{1/2} \tag{2.32}$$

$$= 1/2 \; \psi \sqrt{A} \tag{2.33}$$

where A = maximum peak amplitude
N_s = total number of samples
ψ = SNR.

A plot of σ_{ave} versus A is given in Fig. 2-2. As an example, for $N_s = 128 \times 128$, rms displacement errors of less than 1/10 of a resolution element are obtained for $A > 0.12$. Extremely high accuracies, say $< 10^{-2}$ pixel, are probably not realizable due to inevitable small distortions, aliasing effects and computational round-off.

Equation (2.32) forms the starting point for a variety of system optimization studies. For instance, given a fixed N_s, the error is minimized by simply maximizing A. Another possible situation involves the use of an additional magnification M between

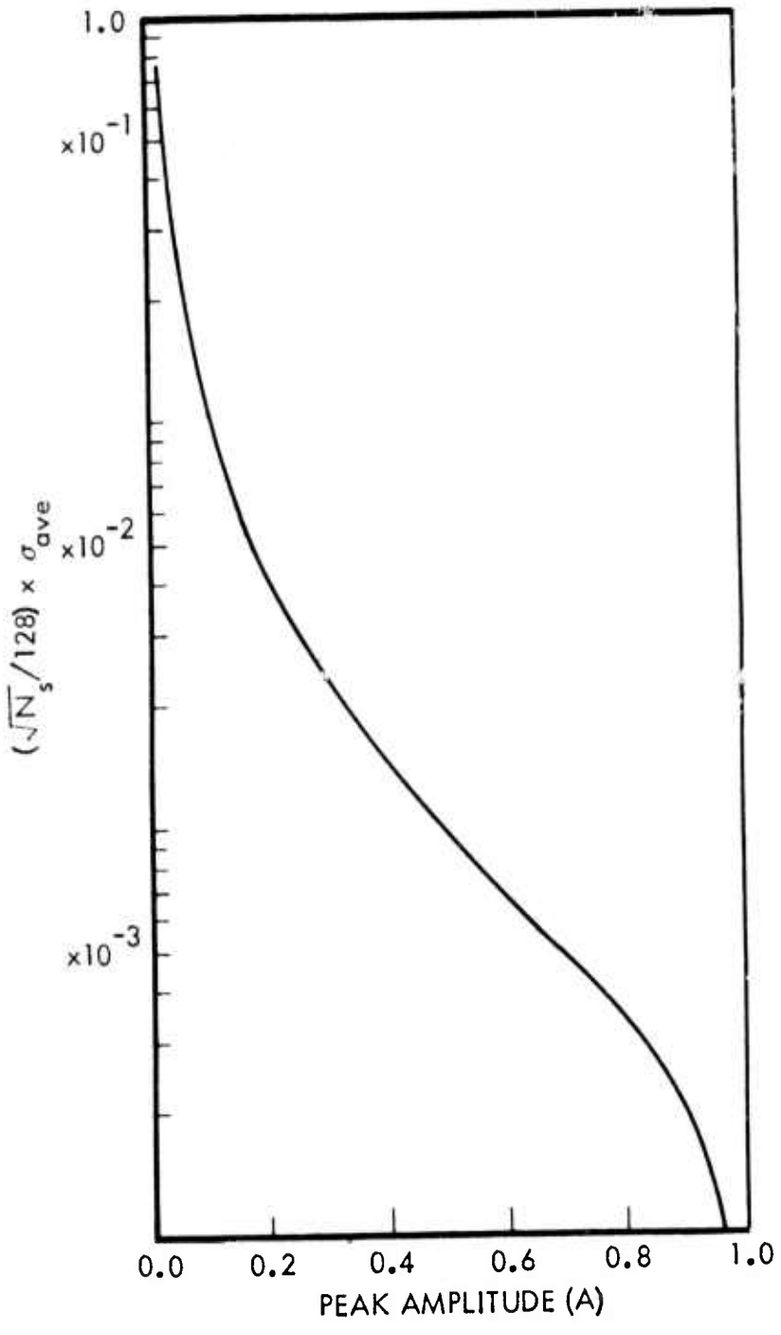

Fig. 2-2 RMS Interpolation Error Versus Peak Amplitude

the focal plane and the detector when operating in the visible or near-infrared. The focal plane precision then becomes

$$\sigma_F = \frac{1}{M} \sigma_{ave} \left[N_s, A(M)\right] \tag{2.34}$$

and the optimal operating point is the minimum of this function with respect to M. Given a fixed field-of-view, the GRD can be optimized by varying N_s so that σ_{ave} is minimized. The peak amplitude A should be relatively insensitive to the sample size as long as the scene is undersampled; however, significant oversampling will rapidly decrease the amplitude.

In this section and the previous one, it has been demonstrated that both the interpolation precision and the probability of false match can be expressed in terms of the number of samples and the phase correlation peak amplitude. Some useful curves illustrating this fact are given in Fig. 2-3. Here, three probability of false match contours corresponding to 1/2, 10^{-3} and 10^{-6} are plotted. In addition, four displacement error contours are shown corresponding to 0.35, 0.20, 0.10, and 0.05 pixels. The displacement errors correspond to rms errors in one dimension. Thus, given an estimate of the peak amplitude produced by a given set of circumstances, it is possible to determine the number of samples required to obtain a given precision or probability of false match. It is interesting to note that the probability of false match becomes significant when the displacement error is greater than a 1/4 of a resolution element.

2.2.4 Convolutional Image Degradations

Suppose that both input images are degraded by the same process (such as smear, defocusing, or antenna function) characterized by a point spread function p or transfer function $P = F\{p\}$. The input images then have the form

$$g'_k = p \circledast g_k \quad (k = 1, 2) \tag{2.35}$$

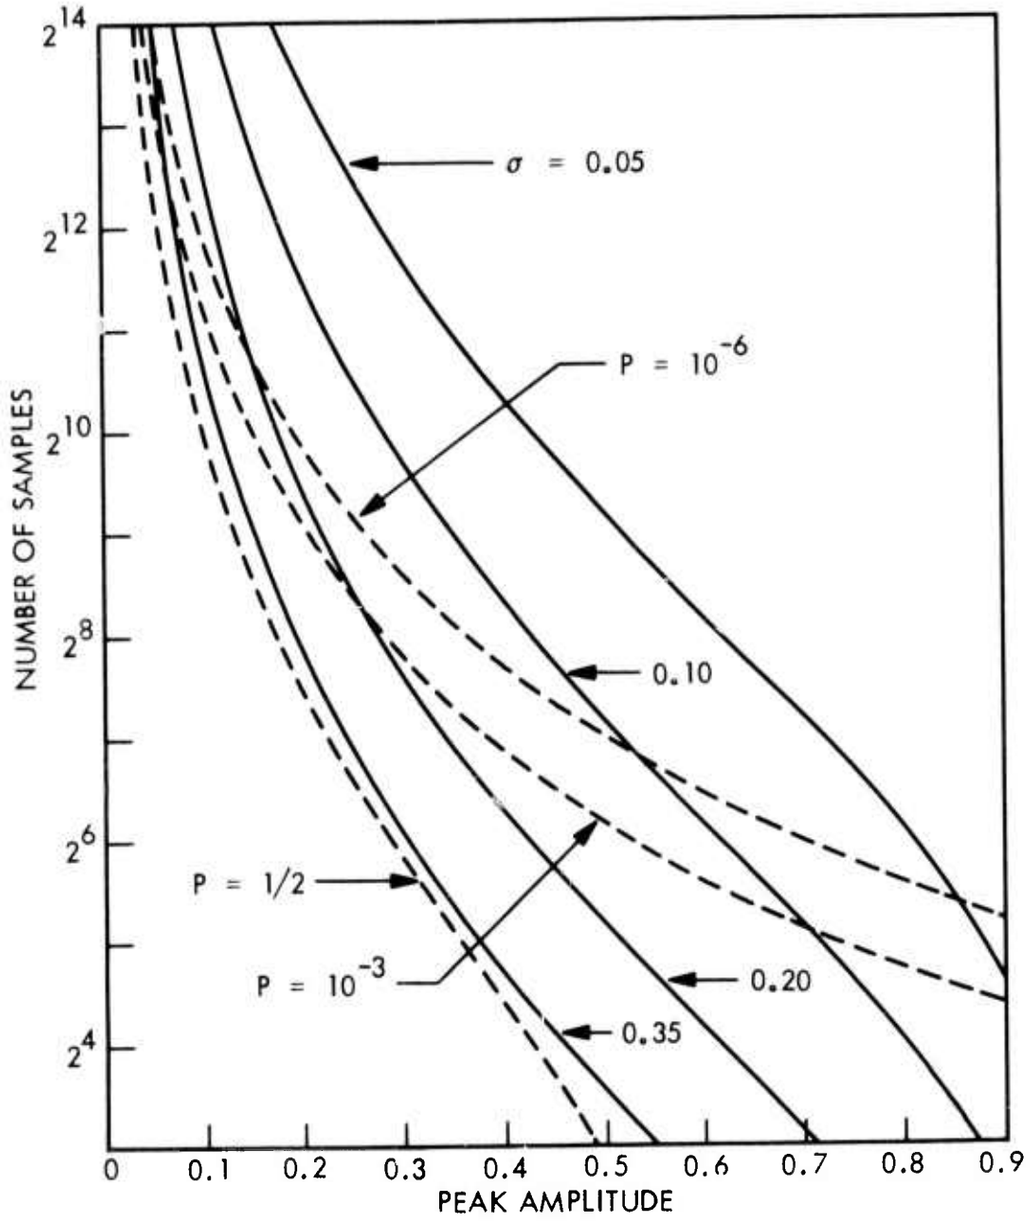

Fig. 2-3 Displacement Error and Probability of False Match Contours

and from the Fourier convolution theorem

$$G'_k = PG_k \quad (k = 1,2) \tag{2.36}$$

where $G_k = F\{g_k\}$. The cross-power spectrum is then

$$G'_1 G'^*_2 = |P|^2 G_1 G^*_2 \tag{2.37}$$

Since the phase of the degraded cross-power spectrum is the same as the phase of the undegraded cross-power spectrum, the phase correlation matrix is identical. The above comments are strictly true only in the absence of noise. In general, image degradations significantly reduce the amplitude of the high-frequency spectrum components so that noise may dominate these terms. The resulting loss of signal bandwidth gives a somewhat broadened delta-function peak of decreased amplitude. The results, however, are substantially less sensitive to image degradation than are those of other correlation techniques for high SNR imagery.

2.2.5 Correlation of Distorted Images

Perspective differences between images constitute a major problem for a high accuracy correlation map-matching system. These differences both decrease the correlation peak amplitude so that a false match becomes more probable and shift the location of the peak from the true matchpoint. The displacement errors that occur have a systematic component and thus differ from the random errors produced by white noise.

As an example, consider a function of spatial coordinates $g(\vec{r})$ whose forward two-dimensional Fourier Transform (FT) is given by

$$G(\vec{f}) = F\{g\} \tag{2.38}$$

A set of transformed coordinates is defined as

$$\vec{r}' = A\vec{r} + \vec{t} \tag{2.39}$$

where $\vec{t}$ is a vector displacement and A is a general nonsingular 2×2 matrix.

The transformed FT is then given by (Ref. 5).

$$G'(\vec{f}) = (\det J)^{-1} e^{j2\pi(A^{-1}\vec{t})\cdot\vec{f}} G[(A^{-1})^{\tau}\vec{f}] \tag{2.40}$$

where "det J" is the determinant of the Jacobian matrix and τ denotes the transpose operation. Equation (2.40) shows that aside from a scale and phase factor, a spatial coordinate transformation A results in a corresponding spatial frequency transformation which is the inverse transpose of A.

Consider now the problem of correlating two identical images which differ by a translation plus a geometrical distortion. Using the results of Eq. (2.40), it is useful to express the cross-power spectrum $G_1(\vec{f})G_2^*(\vec{f})$ as the product of three factors:

$$G_1 G_2^* = M P D \tag{2.41}$$

where

$$M(\vec{f}) = |G_1(\vec{f})\, G_1^*(B\vec{f})/(\det J)| \tag{2.42}$$

$$P(\vec{f}) = e^{j[\phi_1(\vec{f}) - \phi_1(B\vec{f})]} \tag{2.43}$$

$$D(\vec{f}) = e^{-j2\pi(B^{\tau}t)\cdot\vec{f}} \tag{2.44}$$

and

$$B = (A^{-1})^{\tau} \tag{2.45}$$

The significance of each of the three factors in Eq. (2.41) is as follows.

The first term is the magnitude of the cross-power spectrum which has a maximum at $\vec{f} = 0$. The factor, P, which is unity for B = I, introduces a random phase component which is proportional to (B - I). This results in a drop in the correlation peak amplitude and introduces some random displacement error in the correlation peak. The third term, D, is the spectrum of a shifted delta function. The correlation function (cross or phase) will thus have a peak at $\vec{r} = B^T\vec{t}$, rather than at $\vec{t}$. The error in the correlation peak is then given by

$$\vec{\Delta t} = (B^T - I)\vec{t} \tag{2.46}$$

where I is the identity matrix. This result, however, is of limited usefulness as it assumes that the scene is stationary. If this is not the case, the correlation peak may tend to follow some particularly strong feature within the scene as the relative distortion is changed. Further discussion of this point is given in Section 3.9.4. Note that the cyclic cross correlation function is obtained by taking the inverse FT of the entire cross-power spectrum, whereas the phase correlation function is computed by taking the inverse FT of the phase portion of the cross-power spectrum PD.

The Fourier phase angles can be approximated by a set of random variables having a uniform distribution between $-\pi$ and π radians. The function P thus becomes a random phasor for those spatial frequencies $\vec{f}$ such that $(B - I)\vec{f} > 1$ spatial frequency unit. This suggests the use of a spatial frequency filter in the correlation computation which has unit amplitude at $\vec{f} = 0$ and decreases to zero at the interior boundary of the random phase region. All phasors beyond this region are then set to zero. The results of a number of computer simulations using filters of this type are described in Section 3.9.

Another approach for dealing with image distortions involves the use of an image rectification step on the reference or sensed map prior to the correlation computation. This approach will be discussed in greater detail in Section 5.2.

Although the implementation of a rectification computation requires a more sophisticated processor, it does have the advantage that a significant amount of the information content of the imagery enters into the correlation computation. The low-pass phase weighting computation, on the other hand, is designed to suppress information in those regions of the transform domain for which corresponding spectral terms have been distorted by one or more resolution elements. Thus, the effective spatial frequency bandwidth is reduced, resulting in a broader, less precise correlation peak.

Geometrical rectification, on the other hand, does not ensure intensity rectification, and significant intensity variations are normally present after boundaries are properly aligned. It is likely that the best solution will be provided by a combination of filtering and rectification.

Section 3
PERFORMANCE SIMULATIONS

3.1 NOISE SOURCES

It was shown in Section 2 that the phase correlation interpolation precision and probability of false match can be estimated from a knowledge of the peak amplitude and the number of sample points. For a given set of conditions, the analysis of correlator performance can thus be reduced to the simpler task of predicting the peak amplitude.

The performance of phase correlation in the presence of several important sources of image noise will be discussed in the following sections. In some instances only a single example will be shown; in others, an extensive number of simulations have been performed.

The results of the simulations will often be given in the form of correlation surfaces. The correlation surfaces show the positive amplitude variation as a function of matrix row and column index. Peak locations are indicated by: (row #, column #). Unless indicated to the contrary, the origin of the correlation surface (corresponding to zero relative displacement) is located at the center (N/2+1, M/2+1), and the peak amplitude is normalized to 1.0. The peak location indicates the most probable location of image registration, and the ratio of the peak amplitude to the background fluctuations is a measure of the correlation function signal-to-noise ratio. Unless otherwise noted, the "cross correlation" results were obtained using unnormalized cyclic cross correlation as defined by Eq. (2.16).

The term "noise," as used here, refers to any differences that exist between the reference and sensed maps or between two different sensed maps. Since it can be shown that cross correlation is optimal with respect to additive "white" noise, this method is normally chosen for image matching applications, often without due regard to the true nature of the contributing noise sources.

A correlation map-matching system must be able to function in the presence of many sources of "noise", including:

(1) Sun angle effects: illumination gradients; reflections; shadows
(2) Weather effects: cloud cover; rain; haze
(3) Seasonal changes: snow; water; foliage
(4) Geometrical effects: scale change; aspect angle; rotation; 3-D effects
(5) Edge effects
(6) Reference map encoding errors
(7) Image degradations: sensor point spread function; vibration; smear; uneven scan rate
(8) Electronic sources: sensor; processor; external

In all of these cases, except perhaps for the last, the contributing noise source cannot be represented by white, gaussian noise. In fact, to a large extent, the first six categories can be approximated as either two-dimensional narrow bandwidth signals or signals having a spectral density that resembles the noise-free scene.

The phase correlation method is particularly suited to handle either of these types of noise. The correlation function is invariant with respect to either a scaling or a level shift of an image brightness function and phase correlation is a superior algorithm in the presence of either correlated or uncorrelated narrow bandwidth noise of unknown amplitude and spectral content. This latter property follows from the fact that all of the spectral amplitudes are normalized to unity. In addition, since the normalization procedure used to obtain the phase matrix "whitens" each image with respect to its own spectrum, the process acts as an optimum matched filter for noise having a similar spectral distribution.

3.2 NARROW BANDWIDTH NOISE

As indicated above, phase correlation is a preferred technique in the presence of narrow bandwidth noise of unknown spectral content. Such noise will dominate only a small number of (q) terms in the phase matrix. Since all of the terms in this matrix have modulus unity, the output peak amplitude will be reduced in height by a small amount on the order of q/N^2. This result is independent of the amplitude of the noise.

Figures 3-1a, b show two successive frames of imagery taken from a moving platform and scanned to a resolution of 256 × 256 pixels. The image overlap region represents about 44% of the total area. Correlated, narrow bandwidth noise has been added by computer simulation. The cross correlation surface is clearly dominated by the "noise." The phase correlation surface, on the other hand, is essentially unchanged by the addition of the noise and gives a sharp peak at the correct match point of absolute value 0.085. The origin of the correlation surface is at matrix element (1,1) in this case.

The addition of a large amplitude narrow bandwidth signal has resulted in an "infinite" correlation length along the horizontal direction. This, however, has had no significant effect on the matchpoint accuracy as determined by phase correlation.

The scene correlation length is, in general, determined by the largest spatial frequency components. The matchpoint accuracy, on the other hand, is ultimately limited by the effective scene bandwidth and signal-to-noise ratio.

3.3 WIDE BANDWIDTH NOISE

Since cross correlation is dominated by the largest spectral components, it is optimum with respect to white noise. The phase correlator performs moderately well with images degraded by noncorrelated wide bandwidth noise. Correlated wide bandwidth noise, however, presents a more serious problem, so that either an image preprocessing step or a phase matrix filter may be required. It is often possible to separate out the effects of correlated noise since the correlation results resemble those obtained from a scene composed of both a moving and a stationary component. Two main correlation peaks thus appear, one corresponding to the displaced scene and the other, located at the origin of the correlation surface, corresponding to the stationary pattern.

A typical plot of the phase correlation peak amplitude as a function of SNR is shown in Fig. 3-2. Here, the SNR is expressed in dB and is equal to 20 $\log_{10}$ (rms signal/rms noise). The computations were made by correlating pairs of images obtained by adding independent gaussian distributed noise to the 64× 64 element image shown in Fig. 3-3a. The two images used for the "0 dB" computation are shown in Figs. 3-3b, c.

a. Reference Image Plus Interference

b. Shifted Scene Plus Interference

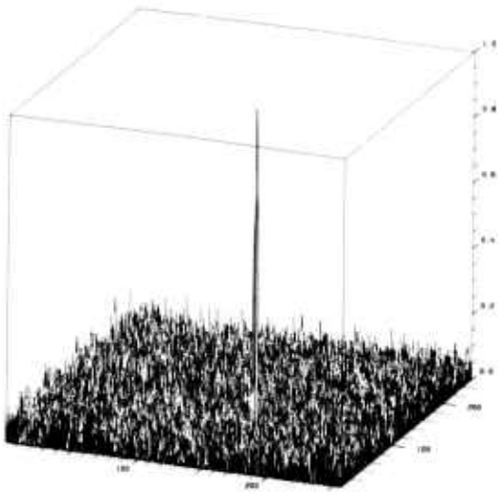

c. Phase Correlation Plot

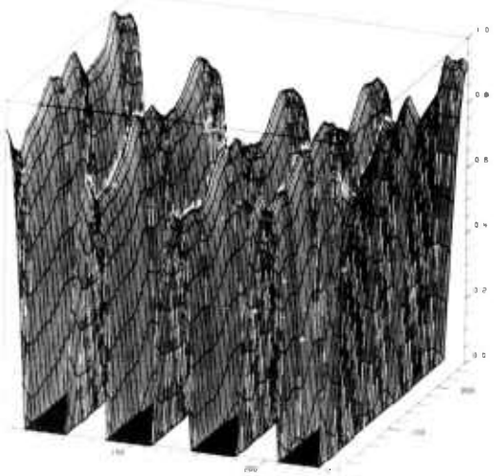

d. Cross Correlation Plot

Fig. 3-1 Effect of Narrow Bandwidth Noise on the Correlation Surface

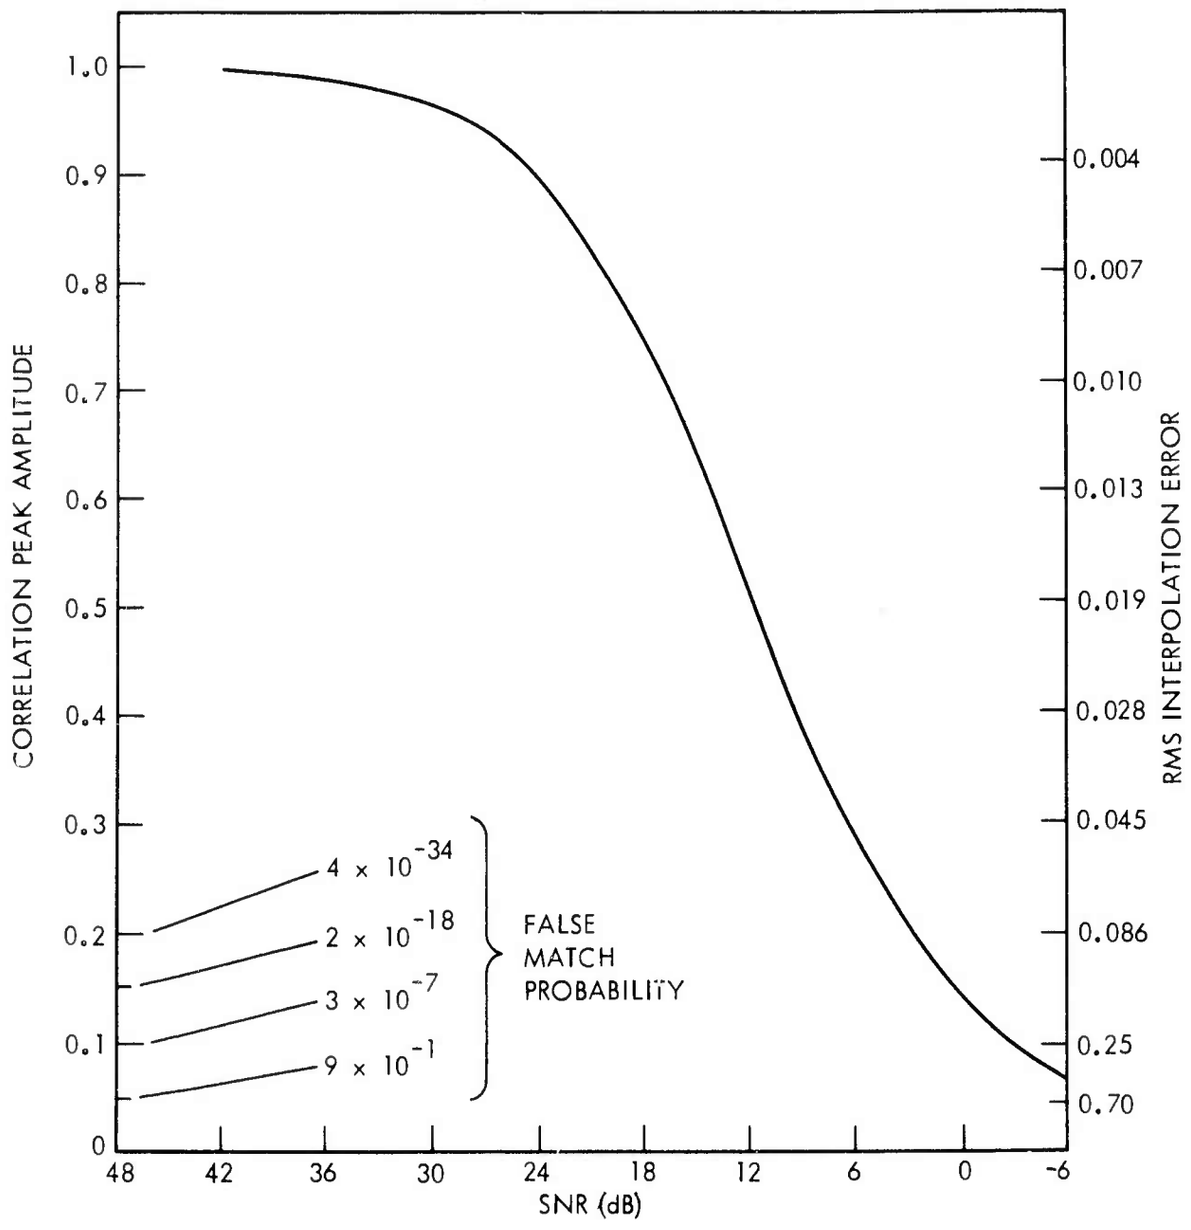

Fig. 3-2 Correlation Peak Amplitude Versus SNR

a. Noise-free Image

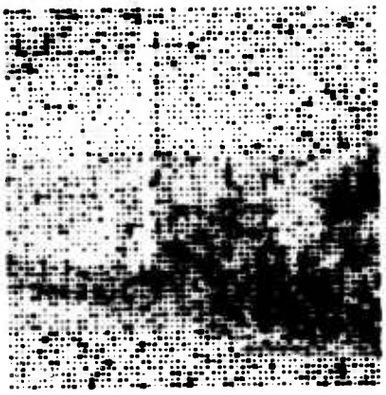

b. Image 1 (0 dB)

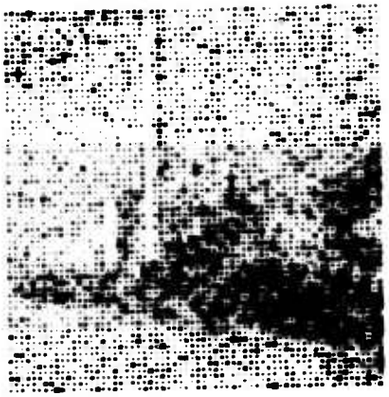

c. Image 2 (0 dB)

Fig. 3-3 Example of Imagery Used for Wide Bandwidth Noise Simulations

In all of the cases shown, the coherent peak exceeds the largest background peak. The maximum error vector occurs for the -6 dB case and has a magnitude of 0.6 resolution elements. At -12 dB, the coherent peak is lost in the noise.

The variation of the phase correlation peak amplitude with random noise is, in general, independent of the number of samples for a given scene and fixed ground resolution.

3.4 EDGE EFFECTS

The phase correlation peak amplitude is a monotonically increasing function of the common area (overlap) between two images. The form of this relationship is normally not strongly dependent upon scene content and the number of resolution elements, N_s, used in the computation.

For a sufficiently small fractional overlap, on the order of $4/\sqrt{N_s}$, the coherent peak amplitude drops to the level of the highest incoherent peak and a correct peak identification can no longer be made. In the case of 64×64 element images having an identical common region, the limiting case corresponds to a 6-percent overlap. Thus, very large relative displacements can be accurately determined without requiring that one image be contained within the other.

A series of phase correlation computations were made using overlapping 64×64 element subimages taken from the reference image shown in Fig. 3-4a. The variation of peak amplitude with "fractional overlap" is shown in Fig. 3-5. (A fractional overlap of 0.5, for instance, means that the two images compared have a 50% common area.) The data were generated by correlating "IMAGE 0" (Fig. 3-4b) with a series of displace images: "IMAGE K", where $1 \leq K \leq 50$. The K^{th} image is displaced K resolution elements to the left and K elements down with respect to the O^{th} image.

The correct displacement vector was computed in all cases with a maximum error of 0.09 resolution elements. As an example, the correlation of "IMAGE 0" with "IMAGE 50" resulted in a peak height of 0.17 and an error of 0.04 resolution elements. The two images have only a 5% common area. The peak amplitudes obtained here for

a. Reference Image

b. Sub-image 0

c. Sub-image 50

Fig. 3-4 Example of Imagery Used for Overlap Simulations

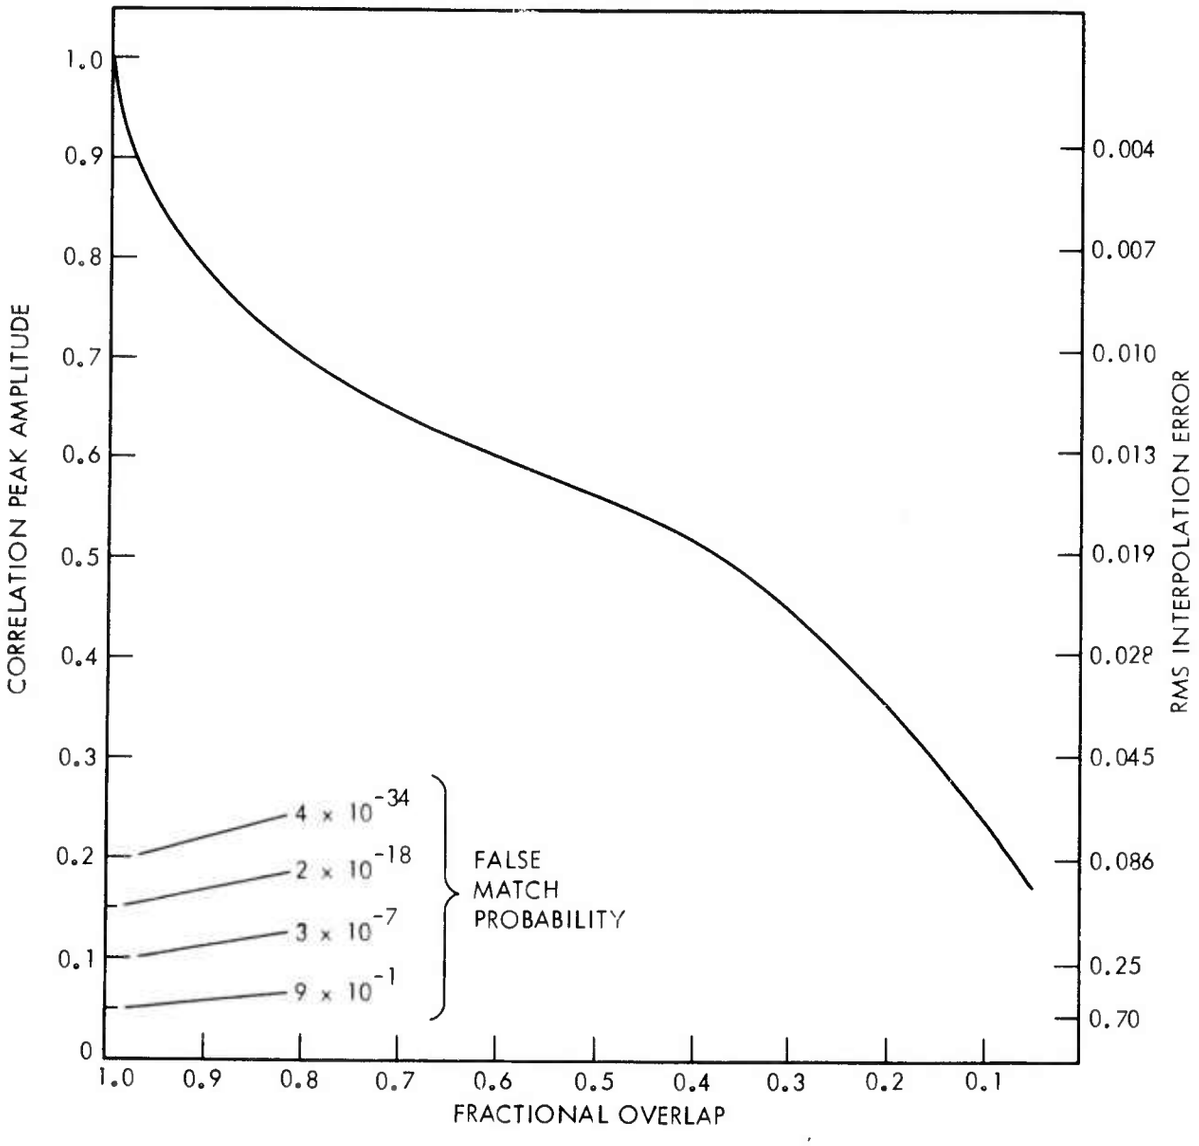

Fig. 3-5 Correlation Peak Amplitude Versus Overlap

small overlaps were higher than expected and this is probably due to the fact that the contrast in the overlap region is higher than the average image contrast.

3.5 SUN ANGLE VARIATIONS

Sun angle variations, to some extent, can be classified as narrow bandwidth noise since scene illumination functions are normally slowly varying with respect to the scene reflectivity. Therefore, the Fourier domain representation of the illumination function is concentrated at low spatial frequencies.

Figure 3-6 shows two images of the Black Dog power plant. Both images are sampled to 120×120 elements and show a sun angle variation in addition to a small relative displacement. The results of correlation experiments using aperiodic cross correlation and phase correlation are shown in Figs. 3-7 and 3-8, respectively. In both cases, the zero mean images were filled out by zeros to form 240×240 element arrays. The origin of the correlation arrays is at matrix element (121, 121). The phase correlation computation indicates a vertical and horizontal displacement of 2 and 9 pixels, respectively, whereas the cross correlation peak gives displacements of 2 and 0. The correctness of the phase correlation result can be verified by examination of Fig. 3-6. The 9-pixel error in the cross correlation computation corresponds to approximately 125 ft at the target. The phase correlation array shows some subsidiary peaks introduced by the imbedding process. These can be eliminated if cyclic phase correlation, using the original 120×120 image, rather than aperiodic phase correlation is used. The results of a cyclic computation (not shown here) were in agreement with the aperiodic case. The absolute value of the phase correlation peak was 0.085.

The error in the cross correlation calculation occurs primarily because the scene is "nonstationary"; i.e., the average intensity varies significantly over the image. In particular, the sky region is much brighter than the ground. This, of course, is a normal occurrence, rather than a pathological one.

a. Image 18 120 × 120

b. Image 38 120 × 120

Fig. 3-6 Variable Sun Angle Images of the Black Dog Power Plant

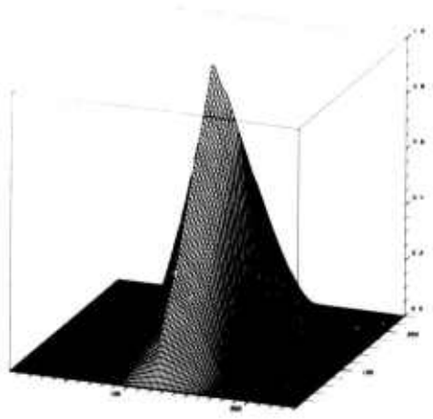

LEFT AXIS – ROW NO.
RIGHT AXIS – COLUMN NO.
VERTICAL AXIS – AMPLITUDE

a. Cross Correlation Surface

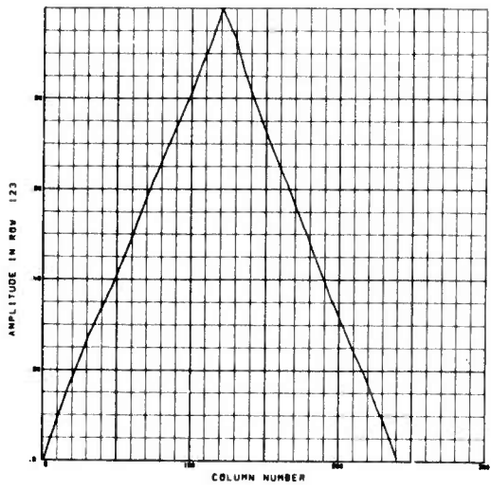

b. Column Cross Section

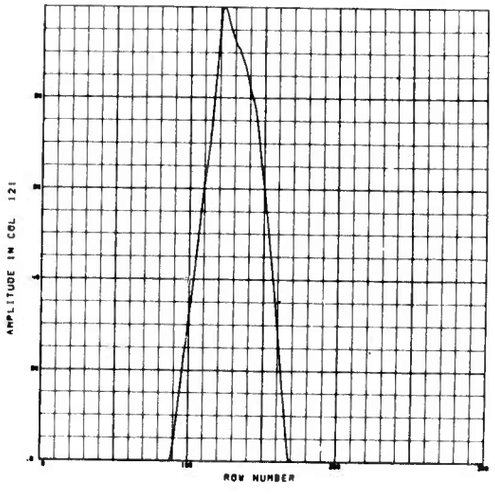

c. Row Cross Section

Fig. 3-7 Aperiodic Cross Correlation Results Using Power Plant Images

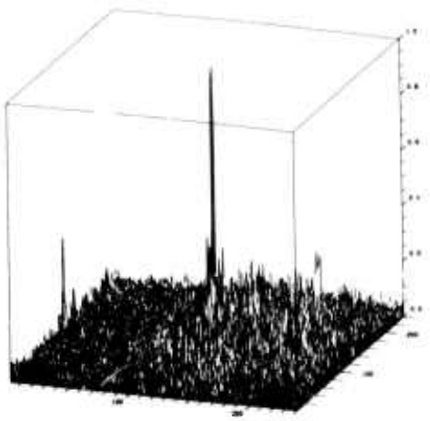

LEFT AXIS – ROW NO.
RIGHT AXIS – COLUMN NO.
VERTICAL AXIS – AMPLITUDE

a. Phase Correlation Surface

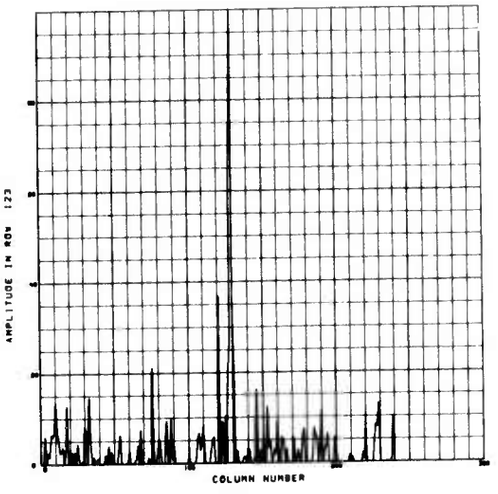

b. Column Cross Section

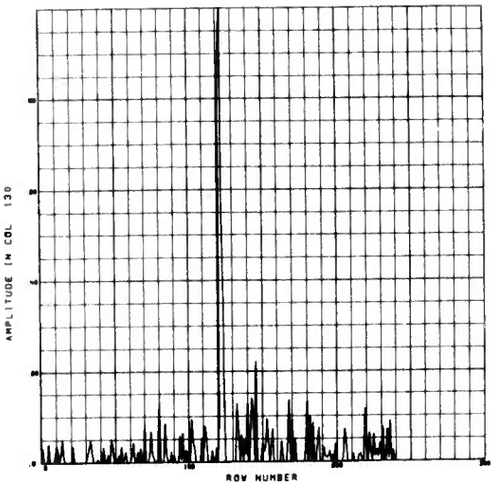

c. Row Cross Section

Fig. 3-8 Aperiodic Phase Correlation Results Using Power Plant Images

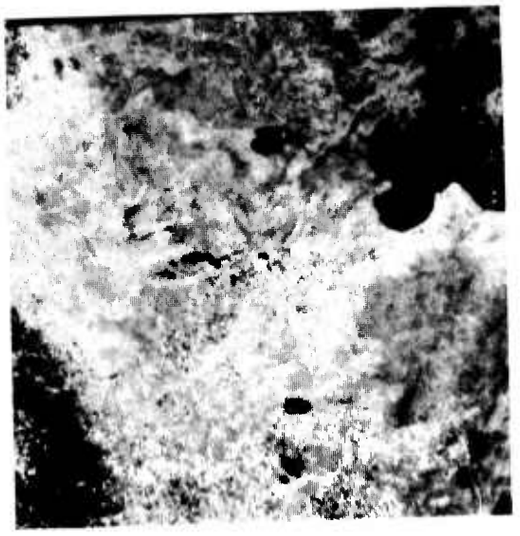

a. ERTS Scene – Clear

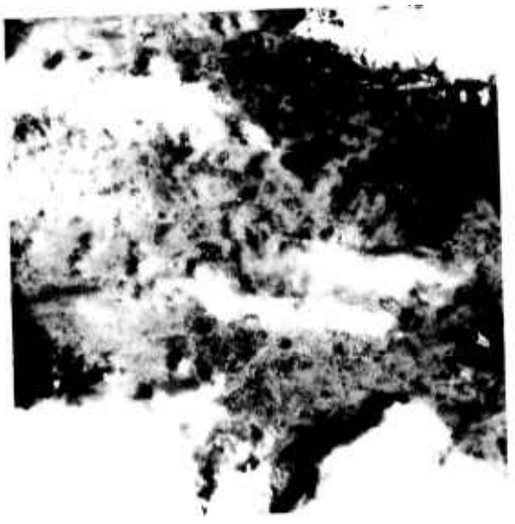

b. ERTS Scene – 50% Cloud Cover

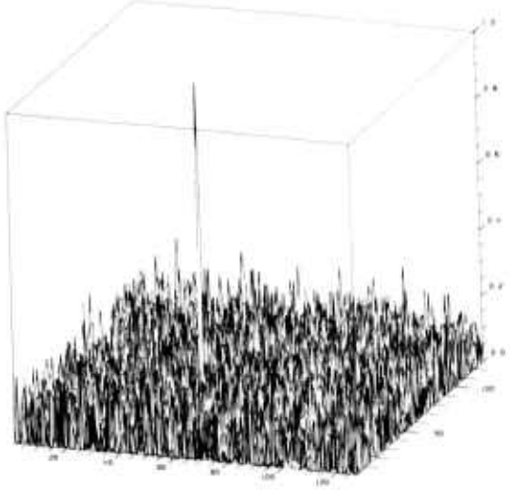

c. Plot of Phase Correlator Output

Fig. 3-9 Correlator Response to Broken Cloud Cover

3.6 WEATHER EFFECTS

Weather effects such as rain, haze, fog, cloud cover, and cloud shadows all limit the effectiveness of a correlation map-matching system. This is particularly true for sensors which operate at the shorter wavelengths. Although it is, of course, impossible to see through heavy clouds at these wavelengths, some limited capability can be achieved with broken cloud cover.

Figure 3-9a shows a cloud-free ERTS scene and a subsequent shifted image partially obscured by clouds (Fig. 3-9b). The phase correlation surface corresponding to two center sections having a 50-percent cloud cover is shown in Fig. 3-9c. The highest peak has an absolute value of 0.088.

3.7 SEASONAL CHANGES

A good example of correlation in the presence of severe seasonal changes is shown in Figs. 3-10 and 3-11. Two different views of a thermal power plant are shown, one clear and one snowy. The images, obtained from ESL, are scanned to 256 × 256 pixels and have a GRD of approximately 7 ft. In each case, the azimuth and elevation to target are unknown. There are significant differences between the two images due to translation, seasonal changes, illumination changes, and differences in scale and aspect angle. As a result, the images are very difficult to correlate, especially using conventional techniques.

Correlation computations were made using both the original scenes and a "rectified" version of the snow scene shown in Fig. 3-10c. The rectification was performed using a projective transformation which maps one plane into a second plane having an arbitrary relative orientation (Section 5.2). This transformation assumes that the entire scene lies on a planar surface, and thus objects having a significant height are not transformed correctly. As an example, although the roof area of the power plant was correctly scaled, the wall area was scaled in the wrong direction. Despite this deficiency, the use of the rectified image greatly improved the phase correlation signal to noise ratio as shown in Fig. 3-11.

a. Clear

b. Snow

c. Snow (Rectified)

Fig. 3-10 Power Plant Scenes

Phase Correlation

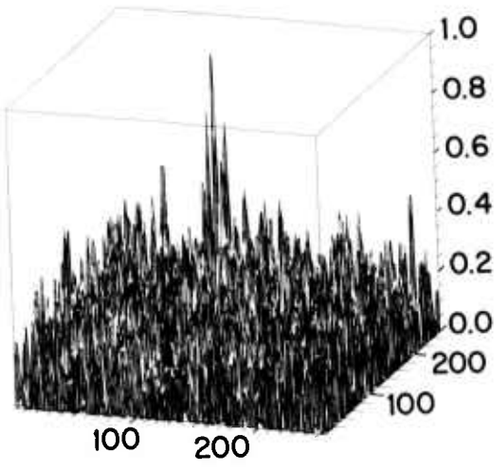

a. Clear - Snow

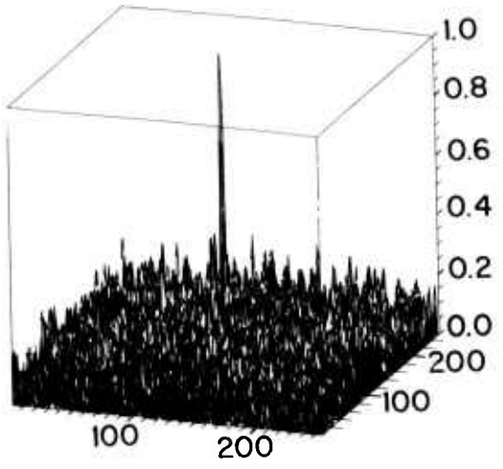

b. Clear - Snow (Rectified)

Cross Correlation

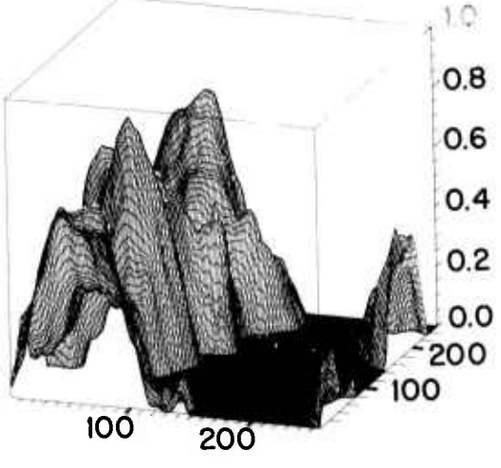

c. Clear - Snow

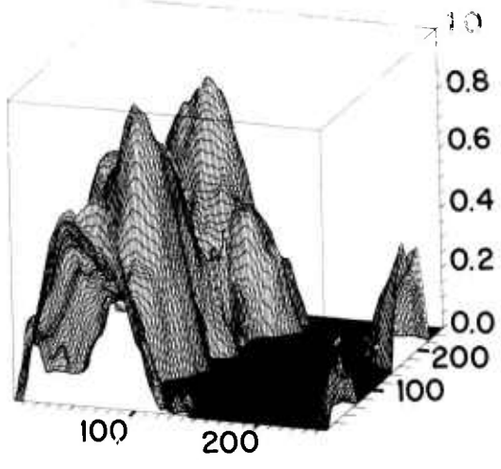

d. Clear - Snow (Rectified)

Fig. 3-11 Power Plant Correlation Surfaces

The phase correlation results were obtained using a phase-weighting function that provided some immunity to image distortions. The highest peak for the CLEAR-SNOW (RECTIFIED) case occurs at matrix element (113, 168) which indicates that the field of view for the clear image has moved 16 pixels up and 39 pixels to the right relative to the snow image. This result appears to be substantially correct. The phase correlation peak was positive in both cases, and only the positive part of the correlation surface is shown. The absolute values of the phase correlation peaks were 0.029 (original) and 0.048 (rectified).

The computations were repeated using circular cross correlation with the average image brightness subtracted out. In both cases, the largest magnitude peak indicated a negative correlation. The negative correlation peak appears to be caused by a match between the white smoke in the "clear" image with the dark power plant storage tank area in the "snow" image. The results were not significantly changed by image rectification. The highest postive peak also gives a result that is substantially incorrect. These results are shown in the lower half of Fig. 3-11 in order to give a direct comparison with the phase correlation surfaces.

3.8 SPECTRAL VARIATIONS

The properties of the phase correlation algorithm make it an attractive candidate for matching images obtained in different spectral bands, or images obtained with different sensors within the same spectral band.

Figure 3-12 compares a passive IR image of a target and a corresponding image obtained with a 10.6-μm laser radar. A rectified version of the laser radar image was correlated with the passive image, and the results are shown in Fig. 3-12d. The correlation peak is actually negative (-0.11), indicating that corresponding regions which are contrast reversed are predominant.

In general, regions which have the same contrast contribute a positive amplitude to the matchpoint peak, whereas regions that are contrast reversed contribute a negative amplitude. The danger thus exists that the two contributions may nearly cancel, resulting in a possible false peak identification. It is therefore of interest to develop some effective image preprocessing steps to minimize this effect.

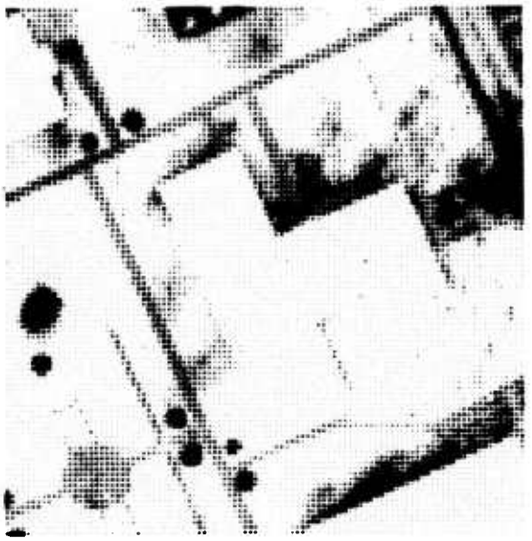

a. Passive IR

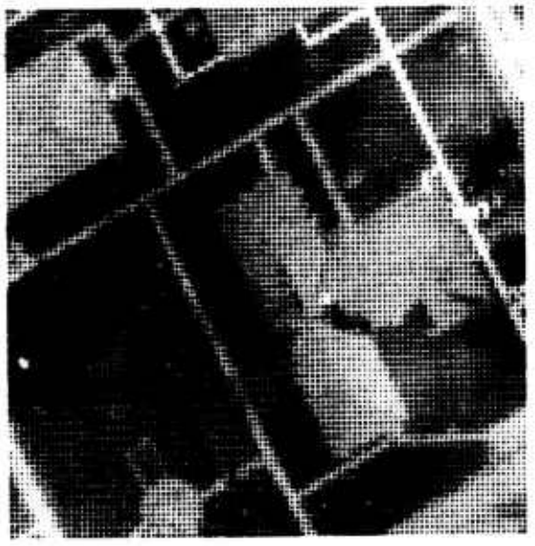

b. Laser Radar

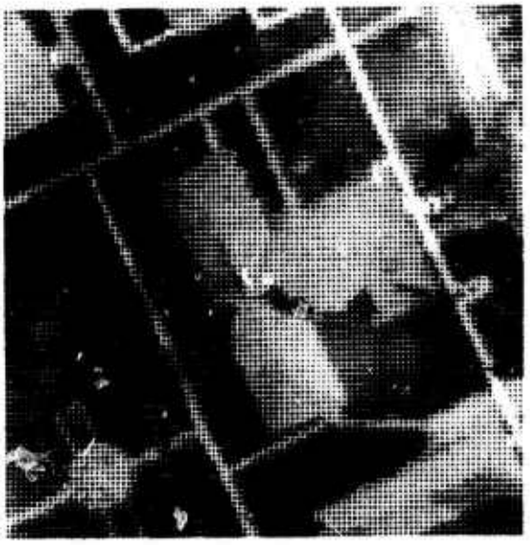

c. Laser Radar (Rectified)

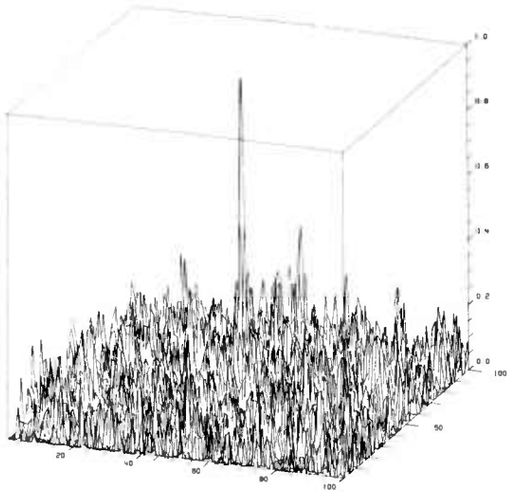

d. Phase Correlation Surface

Fig. 3-12 Correlation of Emitted and Reflected IR Images

3.9 PERSPECTIVE CHANGES

3.9.1 Approach

In this section, the results of 160 computer correlation experiments will be presented in order to assess the extent to which the performance of the phase correlation method is affected by perspective differences between the two images being correlated. Since a sufficient range of imagery with accurately known perspective differences was unavailable, the distortions were generated by computer resampling from a single image. Although the distortions were limited to rotations and scale changes for ease of analysis, the results are believed to be at least qualitatively applicable to more general distortions.

The original image used in these experiments was a high quality one in order to eliminate effects due to noise and oversampling which should be studied separately. The effect of image distortion, as was understood at the outset, is to reduce the range of spatial frequencies from which displacement information can be extracted by straightforward application of the correlation algorithm. The inclusion of information from frequencies outside this range leads to a degradation in correlation accuracy.

Two approaches can be taken to this problem. The first is to eliminate the misleading information by filtering. The second is to recover the displacement information from the entire spectrum by image rectification. The first approach is demonstrated in this section and is shown to be reasonably effective, assuming that errors of several pixels can be tolerated. The second approach is inherently superior, since the inclusion of more information results in more accurate displacement measurements and a lower probability of false match; however, its application necessitates the development of automatic rectification techniques.

Cross correlation is equivalent to a particular type of filtered phase correlation where the "filter" is the magnitude of the scene cross-power spectrum. Since the reduction in bandwidth produced by a given amount of distortion is a function only of

the distortion, it should be possible to develop a filtered phase correlation algorithm which is superior to cross correlation. Although very little work has yet been done on choosing the proper filter shape, the experiments show that the highest correlation accuracy and the lowest false match probability are achieved by cutting off the spectrum at the effective bandwidth.

3.9.2 Test Imagery

A single test image of the Black Dog Power Plant was chosen for the computer simulations in order to provide a consistent set of experiments. The original image, provided by the Control Data Corporation, had been scanned to 130×130 pixels with a gray shade range of $27 - 63$. A set of 16 replicas of the original image was produced by computer resampling corresponding to 8 values each of scale change and rotation. The resampling process used a four-point interpolation algorithm. A center section of each image, consisting of 64×64 pixels, was then chosen for the computer simulations. The "reference" image (3C) and the series of distorted images are shown in Figs. 3-13 and 3-14. In addition, a second reference image (3CA) was generated which was translated by 6 pixels in each direction from the original (80-percent overlap).

The parameters corresponding to each of the rotated (3CR1 – 3CR8) and scaled (3CM1 – 3CM8) images are given in Table 3-1. The various cases correspond to a specified "edge pixel" displacement. That is, if a pair of images are overlaid so that the centers and axes are aligned, a given rotation or scale change can be measured in terms of the displacement of an axial, edge pixel. In the case of scale change, the displacement of a given point is along a vector directed from the origin to the point, whereas a rotation produces a displacement that is perpendicular to this direction. Identical displacements are produced in the transform domain, except that a magnification in one domain will produce a minification in the other.

The Fourier power spectra corresponding to two of the test images which differ by an 18° rotation are shown in Fig. 3-15. The two spectra are also rotated by the same

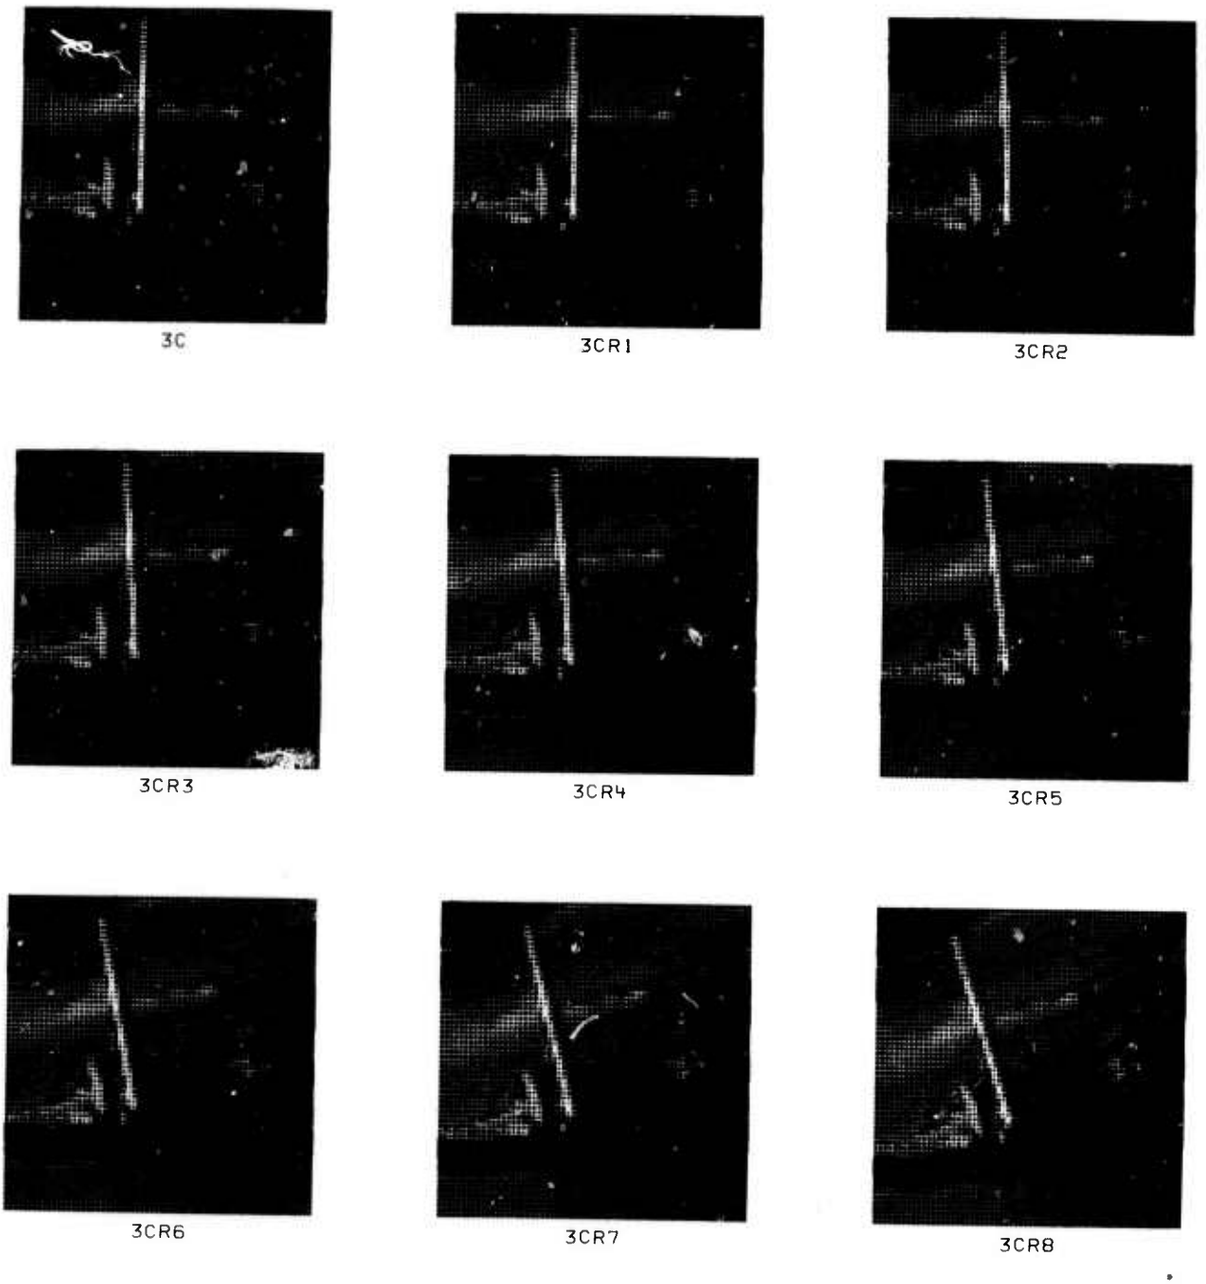

Fig. 3-13 Test Images for Rotation Simulations

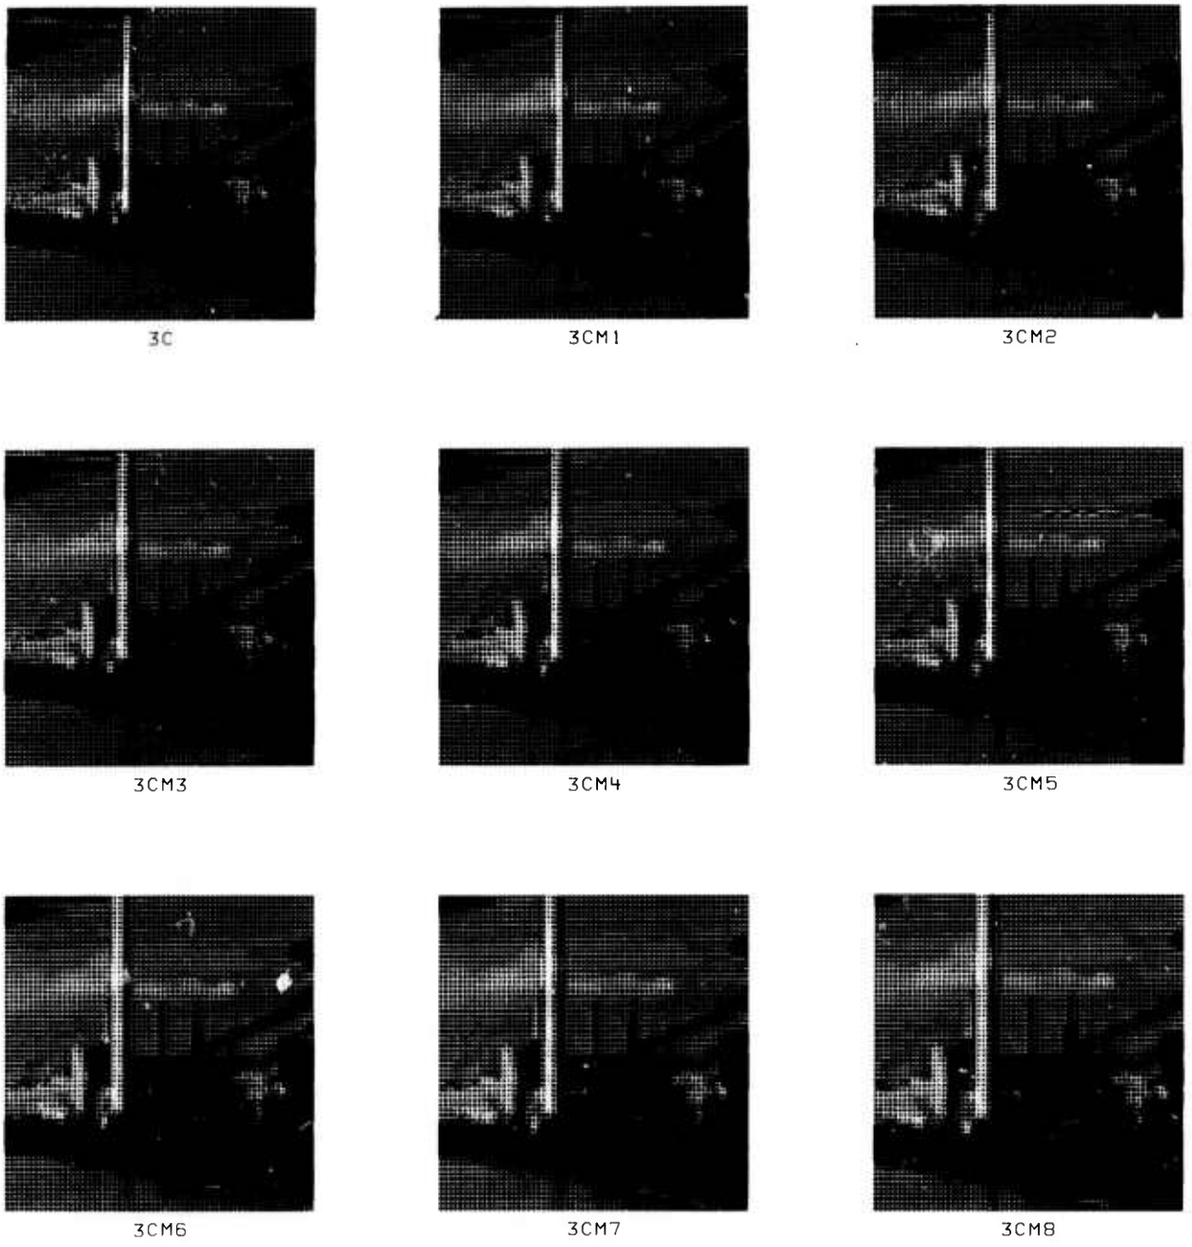

Fig. 3-14 Test Images for Scale Change Simulations

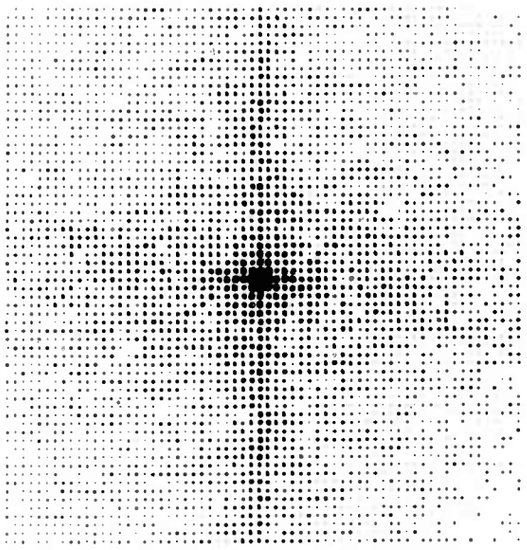

a. Image 3C

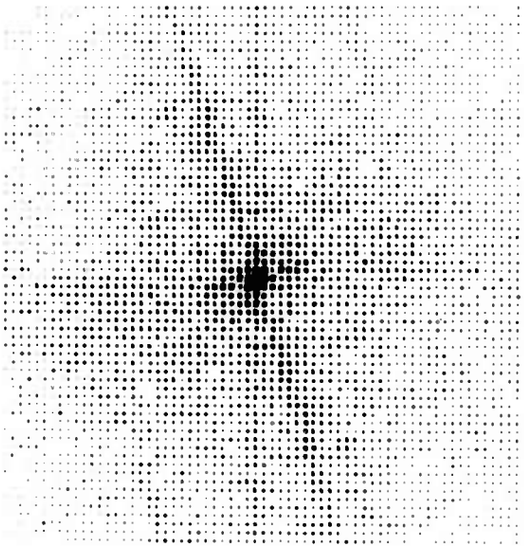

b. Image 3CR8

Fig. 3-15 Fourier Power Spectra

Table 3-1
TEST IMAGERY PARAMETERS

Case	Edge Pixel Displacement	Rotation (deg)	Scale Factor (%)
1	1/2	0.9	1.6
2	1	1.8	3.2
3	2	3.6	6.4
4	3	5.5	9.5
5	4	7.3	12.7
6	6	10.9	19.0
7	8	14.5	25.4
8	10	18.2	31.8

angle with respect to each other in the same sense as the original images. In this case, the function $\log_{10}(1+P)$ is plotted, where P is the power spectrum. The strongest feature in the spectrum is a single line which represents the combined effect of the parallel sides of the smokestacks. The identification of this strong feature could be used to rectify one image with respect to the other prior to a correlation computation in order to both increase accuracy and decrease the probability of false match.

3.9.3 Choice of Filters

The use of low-pass phase weighting functions (filters) to provide some immunity against geometrical distortions is described in Section 2.2.5. The basic idea proposed is that:

(1) Perspective changes in the image domain have a direct analog in the Fourier domain;
(2) Regions of the spatial frequency domain which are displaced by more than one resolution element produce random phase differences and therefore contribute only noise to the correlation function.

This suggests the use of a spatial frequency filter in the correlation computation which has a unit amplitude at zero spatial frequency and decreases to zero at the boundary of

the random phase region. All phases outside of this boundary are then set to zero. The choice of filter, therefore, depends upon the expected degree of distortion.

The three types of filters used for the simulations are described in Table 3-2. The first two filters have the simple shape of a pyramid. For simulations using "filter 1," the parameter "a" was matched to the distortion so that the perimeter of the pyramid base corresponded to the random phase boundary. A fixed filter ("filter 2") matched to an edge pixel displacement of 4 units was used for a second set of simulations. Simulations were also performed with a gaussian filter approximately matched to the distortion. The values of the filter parameters used in each case are given in Table 3-3. In addition to the phase correlation simulations, cyclic cross correlation computations were made for comparison. Normalized cross correlation simulations were also made (not shown) and the results were quite close to those obtained using cyclic cross correlation. Here, the term "normalized" refers to the fact that at each offset, the image elements in common are scaled so that they have a zero mean and a unit standard deviation.

The functional forms of the filters were chosen on the basis of convenience. Thus, there is no reason to believe that they are even close to the optimal shape. The optimization of the spectral weighting functions is, therefore, an important area for future study.

3.9.4 Results

The results of the 160 correlation experiments are given in Tables 3-4 and 3-5. The simulations were carried out by correlating the reference images 3C (100-percent overlap) and 3CA (80-percent overlap) with the set of rotated and scaled images using the filters described in section 3.9.2. The term "80-percent overlap" means that the distorted image was generated from an image that had 80 percent of its area in common with the reference which corresponds to a displacement of 6 image rows and columns.

For each experiment, the displacement error magnitude, $\Delta = (\Delta_x^2 + \Delta_y^2)^{1/2}$, is given in units of pixels. Zero displacement error is defined to represent the exact coincidence of the center (aimpoint) of the reference image and the identical point in the

Table 3-2

DESCRIPTION OF FILTERS*

Filter 1		
$H(u, v) = 1 - \lvert u \rvert / a$	$u \geq v$,	$u \leq a$
$= 1 - \lvert v \rvert / a$	$v > u$,	$v \leq a$
$= 0$	u ,	$v > a$
Filter 2		
$H(u, v) = 1 - \lvert u \rvert / 8$	$u \geq v$,	$u \leq 8$
$= 1 - \lvert v \rvert / 8$	$v > u$,	$v \leq 8$
$= 0$	u ,	$v > 8$
Filter 3		
$H(u, v) = e^{-(\ln 2)(u^2+v^2)/b^2}$		

*In all cases, $-32 \leq u, v \leq 31$.

Table 3-3

FILTER PARAMETERS

Edge Pixel Displacement	Filter 1 (a)	Filter 2 (a)	Filter 3 (b)
1/2	64	8	32
1	32		16
2	16		8
3	11		5.3
4	8		4.0
6	5		2.7
8	4		2.0
10	3		1.6

Table 3-4

ROTATION SIMULATION RESULTS

Angle (deg)	Error Magnitude in Pixels – 100% Overlap				
	No Filter	Filter 1	Filter 2	Filter 3	Cross Correlation
0.9	0.11	0.11	0.05	0.10	0.08
1.8	0.47	0.33	0.04	0.24	0.15
3.6	1.18	0.31	0.07	0.24	0.14
5.5	2.02	0.42	0.23	0.26	0.15
7.3	0.89	0.47	0.47	0.38	1.03
10.9	4.17	1.63	5.03	0.84	1.00
14.5	6.70	2.51	7.46	2.33	0.70
18.2	29.48	4.74	9.14	4.48	6.78

Angle (deg)	Error Magnitude in Pixels – 80% Overlap				
	No Filter	Filter 1	Filter 2	Filter 3	Cross Correlation
0.9	0.10	0.10	0.87	0.10	0.19
1.8	0.45	0.39	1.13	0.38	0.54
3.6	0.96	0.38	1.24	0.36	0.66
5.5	1.54	1.07	1.68	0.78	0.86
7.3	3.61	2.26	2.26	1.34	0.60
10.9	22.92	2.24	5.00	2.69	2.58
14.5	5.59	9.14	7.59	5.27	6.39
18.2	40.11	7.46	8.51	4.86	7.26

Table 3-5

SCALE CHANGE SIMULATION RESULTS

Scale (%)	Error Magnitude in Pixels – 100% Overlap				
	No Filter	Filter 1	Filter 2	Filter 3	Cross Correlation
1.6	0.01	0.01	0.06	0.02	0.07
3.2	0.03	0.09	0.08	0.13	0.18
6.4	1.01	0.48	0.31	0.46	0.73
9.5	1.27	0.75	0.67	0.68	1.16
12.7	2.00	1.00	1.00	1.04	1.59
19.0	3.55	2.38	1.79	2.52	2.76
25.4	3.86	2.85	2.22	3.30	3.35
31.8	4.98	2.64	5.43	3.55	4.02

Scale (%)	Error Magnitude in Pixels – 80% Overlap				
	No Filter	Filter 1	Filter 2	Filter 3	Cross Correlation
1.6	0.13	0.13	0.73	0.13	0.23
3.2	0.22	0.21	0.56	0.32	0.44
6.4	1.37	0.48	0.42	0.55	1.04
9.5	1.58	0.67	0.56	0.87	1.53
12.7	2.54	0.96	0.96	0.99	1.98
19.0	3.61	1.32	1.58	1.54	3.08
25.4	4.28	1.75	2.31	2.20	3.85
31.8	4.51	2.38	8.43	3.45	5.62

distorted image. This selection for the aimpoint was arbitrary. It is likely that improved results would have been obtained if the aimpoint had been chosen to correspond to some distinctive feature, e.g., the base of the large stack. The data in Tables 3-4 and 3-5 are plotted in Figs. 1-4 through 1-7.

Although there is a considerable amount of scatter in the data, several not unexpected conclusions can be drawn:

- Displacement errors increase with image distortion
- Errors increase with decreased overlap
- Cross-correlation is superior to unfiltered phase correlation
- Filtered phase correlation is comparable to, and in many cases superior to, cross correlation
- The "adaptive filters out-perform the "fixed" filter

The errors produced by image scale change are generally smaller than those associated with the rotations. The reason for this is not known. It may be related to the fact that the directionality of the strongest feature – the stacks – is preserved under a scale change.

The "adaptive" phase correlation accuracy is consistently better than that obtained using cross correlation for scale change. For the rotation simulations, the results are mixed. In all four cases corresponding to the greatest distortion (18-deg rotation, 32-percent scale change), adaptive phase correlation was superior to cross-correlation.

An example of two phase correlation surfaces corresponding to an 11-deg rotation and 80-percent overlap is given in Fig. 3-16. The 22.9-pixel error in the former case has been reduced to 2.2 pixels through the use of filtering. The correlation surfaces are sampled at double-resolution and therefore have dimensions of 128 × 128 pixels. The origin is located at matrix element (65, 65). The highest peak in each case has been normalized to 1.0.

A plot of the unfiltered correlation peak amplitude versus distortion is shown in Fig. 3-17. The results for rotation and scale change are quite similar when expressed

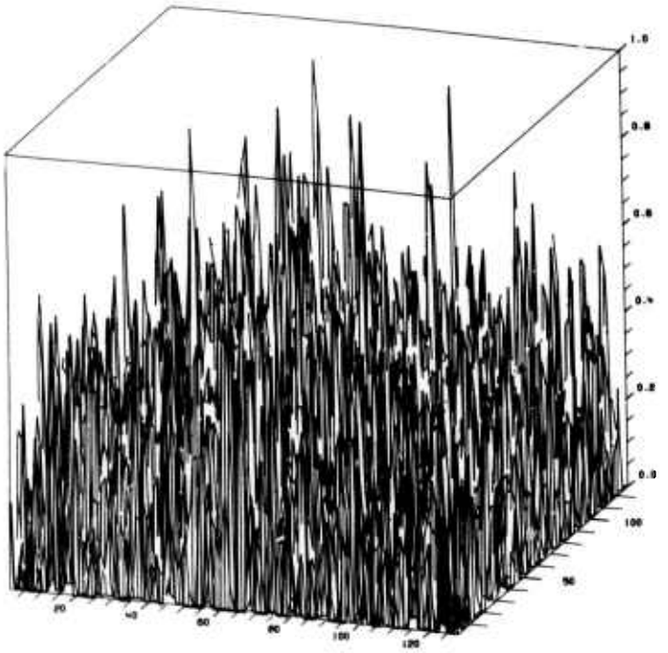

a. Unfiltered Correlation Surface

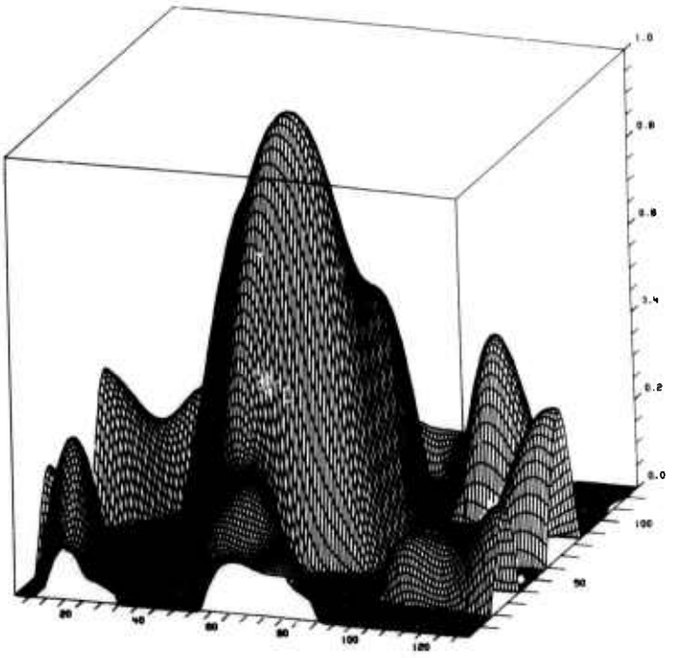

b. Filtered Correlation Surface

Fig. 3-16 Comparison Between Unfiltered and Filtered Correlation Surfaces

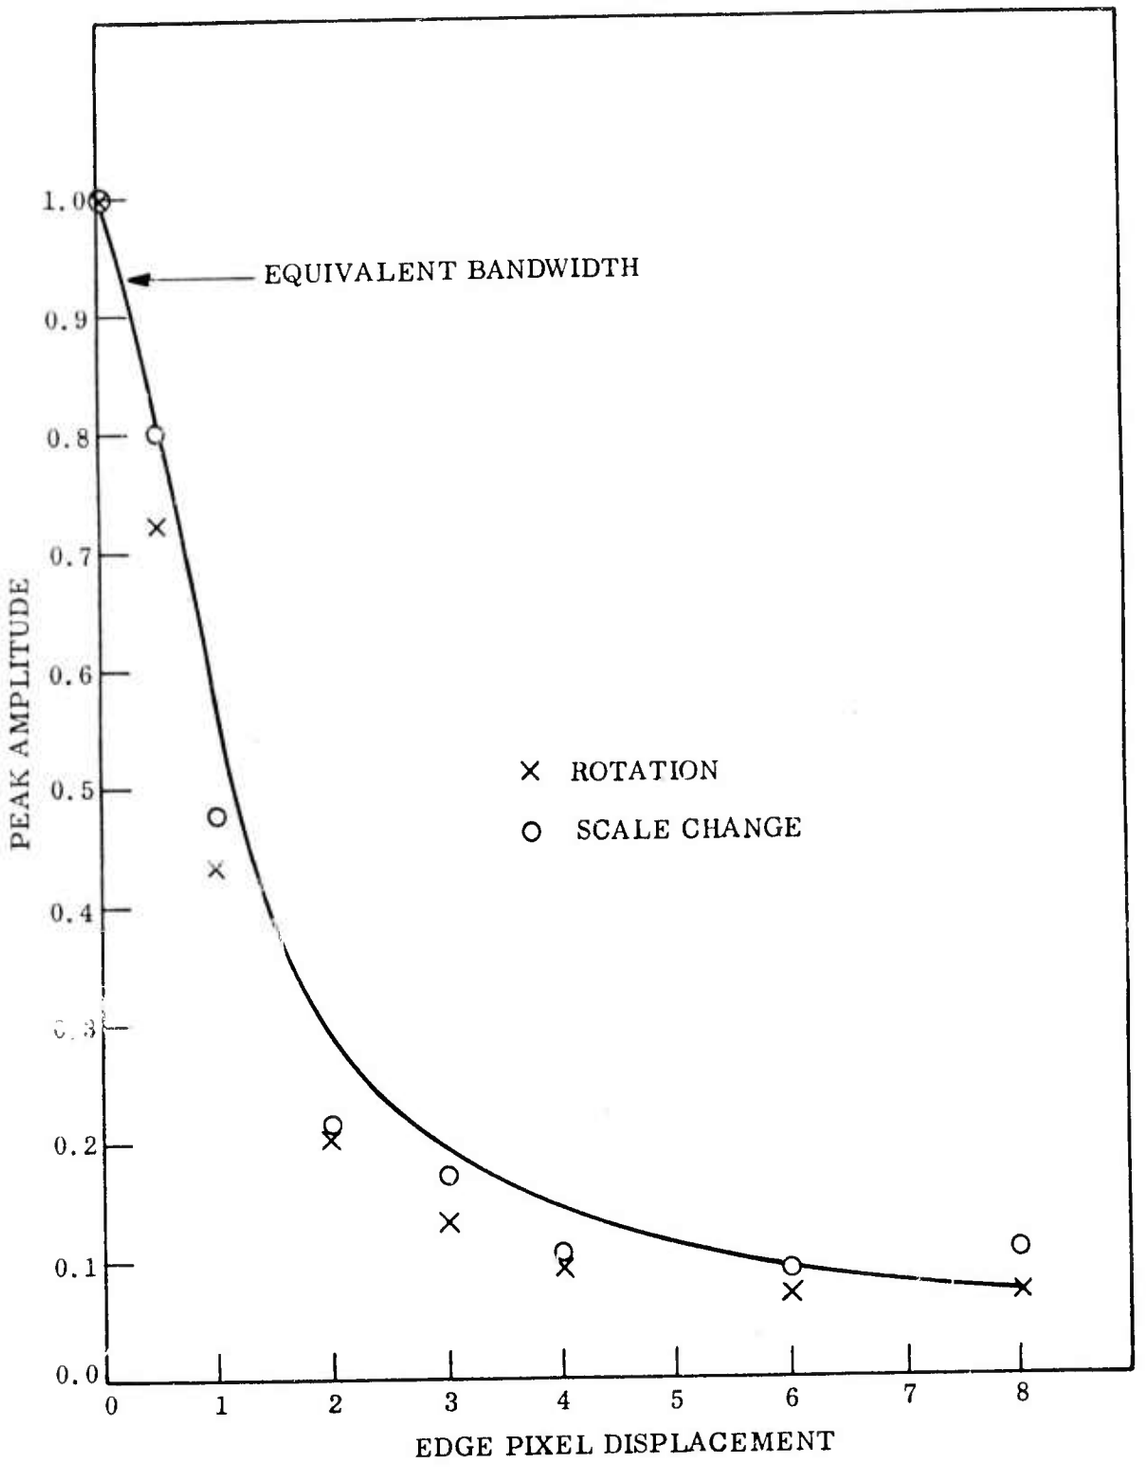

Fig. 3-17 Peak Amplitude Versus Distortion

in terms of an equivalent edge pixel displacement. The test images are degraded somewhat with respect to the reference in that they were generated using a four-point interpolation algorithm. The amplitudes should therefore be somewhat higher for actual imagery, assuming that no other sources of noise are present. Although the data were generated using 64×64 pixel images, the results can be generalized to an NXN element image by scaling the abscissa by a factor of 64/N.

The data points appear to follow a curve designated as "equivalent bandwidth." For an NXN element image, this parameter is defined as

$$\frac{1}{N}\left(\sum_i H_i\right)^{1/2}$$

where H is the amplitude of the "filter 1" function and the summation extends over all spatial frequencies. Although this definition is somewhat arbitrary, it is clear that the major effect of distortions is to effectively reduce the available bandwidth and that the equivalent bandwidth varies inversely as the edge pixel distortion.

In assessing these results, the concept of "displacement error" must be used with caution. Distorted images have no precise match point, and thus the correlation peak may wander over an increasingly large region as the relative distortion is increased. The errors determined here have a large systematic component which indicates that the correlation peak tends to follow certain image features.

This effect is illustrated in Fig. 3-18, which shows the error in the computed match-point as a function of scale change for unfiltered and filtered phase correlation, and cross correlation. In addition, the relative motion of a particular region, the base of the large smokestack, is indicated. The cross correlation peak closely follows this feature. The smokestack base, therefore, represents the "effective" matchpoint for the cross correlation computation. It would therefore be of interest to repeat these

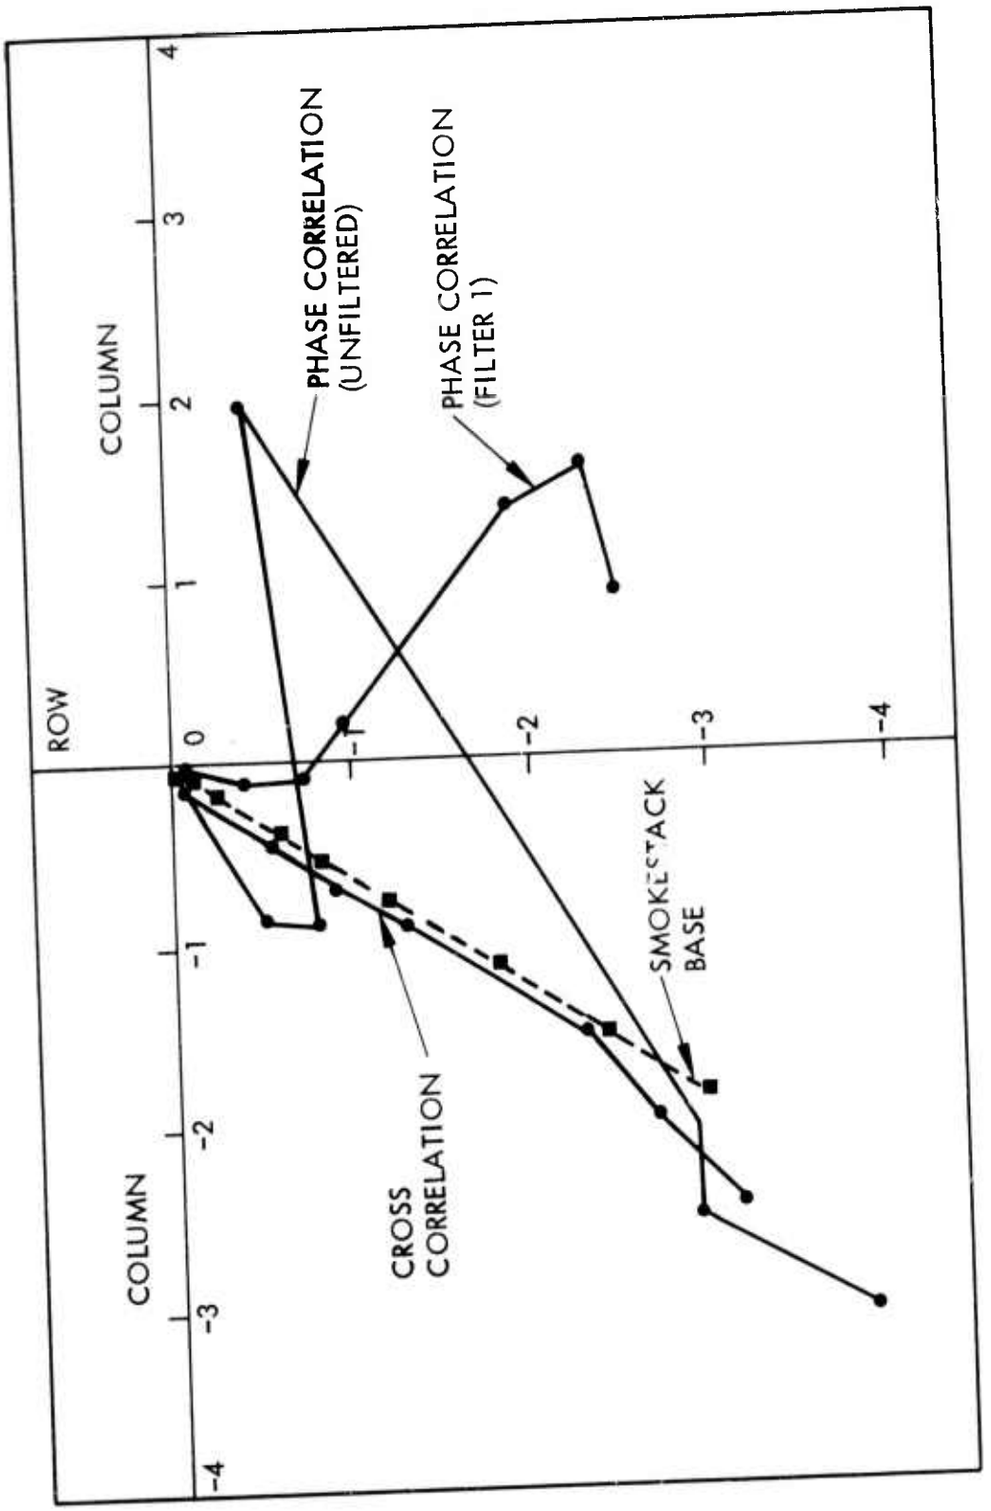

Fig. 3-18 Error in Computed Matchpoint Versus Scale Change

calculations using the smokestack base as the aimpoint in order to see whether the errors are indeed reduced.

It is of interest to estimate the optimum number of samples to use with geometrical distortions, assuming that no attempt at image rectification is made. For a fixed amount of distortion, an increase in the number of samples is beneficial on the one hand, as it represents an increased number of measurements, but detrimental on the other hand because of the loss in peak amplitude. Unfortunately, the theory developed in Section 2.2 does not readily apply to distorted images because of the lack of a well-defined matchpoint. As an example, for 64 × 64 element images, Eq. (2.32) predicts an rms displacement error of 0.09 pixels corresponding to a peak amplitude of 0.2, regardless of the mechanism which produced this particular result. At this peak amplitude, the wide bandwidth noise simulation (Section 3.3) indicated an error of 0.12 pixels, whereas the edge effect simulation (Section 3.4) gave an error of 0.06 pixels. Both values are thus in reasonable agreement with theory. On the other hand, the rotation and scale change simulations give displacement errors which are an order of magnitude larger at this peak amplitude. The discrepancy is due to the "systematic" errors discussed above which are related to the nonstationary character of the imagery. It may be possible to greatly reduce this source of error either through the proper choice of aimpoint or by appropriate image preprocessing.

If it is assumed that the systematic errors can be ignored, then the best operating point can be estimated by using the "peak amplitude versus distortion" curve (Fig. 3-17) in conjunction with Eq. (2.28) to find a minimum in the probability of false match.

For a fixed distortion, this minimum is achieved if the maximum distortion is limited to one pixel by the proper choice of the number of samples. This, of course, is exactly the idea used for choosing the filters as described in Section 3.9.3.

Section 4
PHASE CORRELATION PROCESSOR

4.1 COMPUTATION ALGORITHM

A hardware implementation of the phase correlation method can be made involving no greater complexity than that required for cyclic cross correlation. The computation consists of the following sequence of steps:

(1) The input consists of two sampled images g_1 and g_2 which have the same dimensions (say $N \times M$).

(2) The two-dimensional Fast Fourier Transform (FFT) is taken of each image, resulting in two complex $N \times M$ elment arrays, G_1 and G_2.

(3) The phase difference matrix is derived by forming the cross-power spectrum, $G_1 G_2^*$, and dividing by its modulus; i.e.,

$$e^{j\phi} = e^{j(\phi_1 - \phi_2)} = G_1 G_2^* / |G_1 G_2^*| \tag{4.1}$$

(4) If filtering is required, the phase matrix is multiplied by a weighting function $H(\vec{f})$.

(5) The phase correlation function, d, is then obtained as a real $N \times M$ element array by taking the inverse FFT of the weighted phase matrix:

$$d = F^{-1}\{He^{j\phi}\} \tag{4.2}$$

The computation can be further simplified by replacing Step (3) with a phase quantization algorithm. With this approach, the phase difference matrix $e^{j(\phi_1 - \phi_2)}$ is approximated by the following procedure:

(1) The phase angle corresponding to a transform element $G(\vec{f})$ is quantized by using three binary decisions involving the real (R) and imaginary (I) parts of G:

$$R \overset{?}{\geq} 0 \quad ; \quad I \overset{?}{\geq} 0 \quad ; \quad R \overset{?}{\geq} I \tag{4.3}$$

This generates a 3-bit index $k(\vec{f})$.

(2) The phase angle difference index is obtained through a table lookup or by a direct 3-bit subtraction:

$$\ell(\vec{f}) = k_1(\vec{f}) - k_2(\vec{f}) \pmod 8 \tag{4.4}$$

(3) The index ℓ directly references a prestored phasor, $P_\ell = e^{j\phi_\ell}$, which is used in place of the actual phasor, $e^{j\phi}$.

This computation need only be carried out on half of the transform elements. The Hermitian property of G

$$G(\vec{k}) = G^*(-\vec{k}) \tag{4.5}$$

ensures that

$$\phi(\vec{k}) = -\phi(-\vec{k}) \tag{4.6}$$

The phase quantization process is clearly easier and faster to implement than multiplying two complex quantities and dividing by the modulus of the product. The phase quantization process results in a reduction of all the peak amplitudes by about 10% and introduces a displacement error on the order of 0.02 resolution elements for 128×128 element arrays.

Although the phase correlation algorithm requires greater computational sophistication than that required for cross correlation or correlation coefficient computations performed in the spatial domain, this increased sophistication results in fewer components and computing operations in order to achieve a given processing rate. This is due to the fact that an FFT phase correlator for $N \times N$ element images requires on the order of $6N^2 \log_2 N$ complex operations, whereas conventional correlation, performed in the spatial domain requires on the order of N^4 real operations. For $N = 64$, this savings amounts to a factor of 30.

4.2 PROCESSOR DESIGN

An example of a phase correlation map-matching processor is shown in Fig. 4-1. The sensed image is digitized by the A/D converter, stored, and processed by a 2-dimensional FFT processor. The 2-D FFT consists of a 1-D FFT plus a transpose memory. The phases are then extracted from the spectrum and stored in the ϕ_1 store. The reference map is stored in the form of phase information. The data consist of a series of phase arrays corresponding to subarrays of the original image which have dimensions equal to the sensed map dimensions. The phase information is quantized to 3 bits, and only half of these data need be stored because of the Hermitian property of the transform. The reference data can therefore be compressed to 1-1/2 bits/pixel without loss of performance.

Each reference phase array is then read out of the ϕ_2 store and subtracted from the sensed phase array obtained from the ϕ_1 store. The corresponding sines and cosines are indexed in a programmable read only memory (PROM), and the resulting matrix is processed by the 2-D FFT (inverse) logic. If required, filtering can be performed by multiplying by a filter function H prior to the inverse FFT.

The result is the phase correlation matrix, which is searched for the largest value. The location and magnitude for this value are then output.

It is of interest to estimate the amount of hardware required to implement the phase correlation processor. As an example, for a 64 x 64 element image and a processing

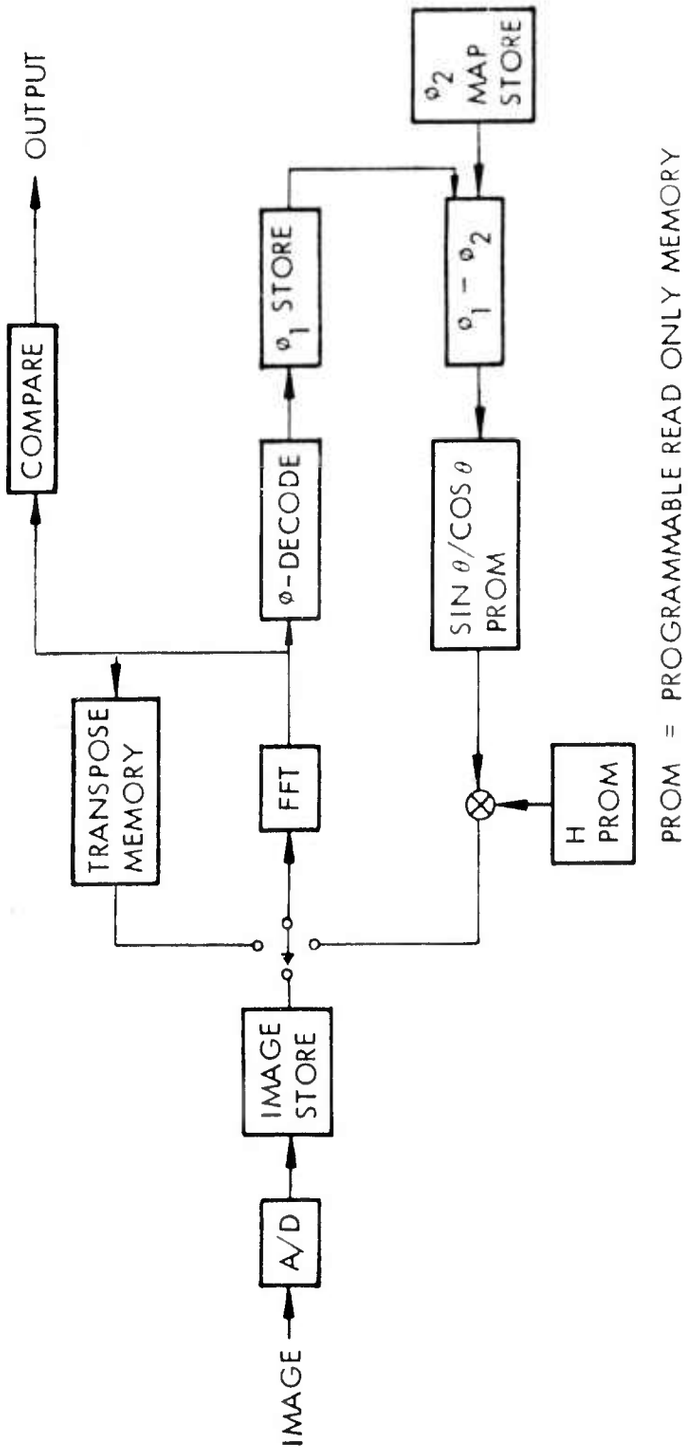

Fig. 4-1 Phase Correlation Processor Block Diagram

rate of 10 ms (100 frames/s), fewer than 300 standard TTL chips would be required. These chips average about 60 gates/chip and will fit on the two 8-in. x 9-in. circuit boards not including the power supply. Expected advances in technology should dramatically increase the available device density within the next few years, so that the processor could be constructed with only a few chips. The hardware requirements for "real-time" operation are thus fairly modest.

Section 5
CORRELATION GUIDANCE CONSIDERATIONS

5.1 TARGET ACQUISITION AND HOMING

A correlation guidance system must be able to acquire a landmark or target over some prespecified error basket which is described by a stored reference map. Ideally, the reference map has been carefully generated so as to contain a reasonable replica of the scene as obtained by the onboard sensors. As indicated in Section 3.1, many differences can occur between the reference and real-time maps such as translations, perspective changes, illumination and seasonal changes, weather conditions, sensor noise, etc. These scene differences both diminsh the correlation peak amplitude, thus increasing the probability of false match, and shift the location of the peak from the true match-point. The most serious problems appear to be caused by perspective changes, since under these conditions image alignment cannot be achieved by simple translations.

In this respect, a high precision, forward- or side-looking system presents the following dilemma. In order to achieve a fix in the presence of unknown perspective changes, a "sloppy" correlation algorithm (or reference map) is required so that correlation can be achieved using the gross details (low spatial frequencies) present in the scene. This, however, increases the probability of false fix and can lead to significant displacement errors. The initial use of a very precise correlation measure, on the other hand, may reduce the probability of target acquisition to an unacceptably low figure. The various factors that determine the probabilities of detection and false match must therefore be carefully assessed and their relationship to overall accuracy established.

There are a number of useful techniques which can be used to increase the probability of target acquisition, such as:

- Image resolution reduction
- Multiple reference maps

- Composite reference maps
- Low-pass phase filters (subsection 2.2.5)
- Correlation peak SNR enhancement using multiple looks
- Multidimensional correlations over the perspective parameter space
- Scene rectification (subsection 5.2)

Perhaps the most useful development would be an autonomous image rectification process which could remove a large fraction of the sensed map distortion prior to the correlation computation. There are several possible ways of accomplishing this including the use of pattern matching techniques in the spatial and Fourier domains, and mosaic correlation. A less precise rectification system might utilize altimeter and range data and/or input from the inertial guidance system instead of scene data to provide an estimated correction. Another possible approach is to use "blind" rectification to generate multiple versions of the sensed map in "real-time" over a prescribed set of perspective variations. The best estimate is then obtained by maximizing the phase correlation peak amplitude.

At the present time, it is not clear which of the above methods or combination of methods should be used in conjunction with the phase correlation algorithm in order to achieve a high acquisition probability. This problem should be one of the major tasks for future study.

Since it is unlikely that the initial look will provide sufficiently high accuracy to successfully complete the mission, a series of reference maps corresponding to different ranges must be stored. The initial target fix will provide both translation and perspective information which is used to generate guidance control commands. Errors in the generation and execution of these commands are then compensated for at the next fix and so on. During the time periods between fixes, it may be of interest to do frame-to-frame correlation using successive sensed maps in order to independently monitor the trajectory. The correlator would thus act as a highly accurate image velocity sensor. This computation is inherently more accurate than the map matching calculation since both images will be nearly the same.

Several topics which are pertinent to the acquisition problem are discussed in the following sections.

5.2 SCENE RECTIFICATION

Scene rectification, as discussed below, defines the process by which a central projection image of a plane object can be transformed to a second plane having an arbitrary relative orientation in space. For a down-looking sensor, image rectification can be handled reasonably well by determining the relative rotation and scale change between the reference and sensed maps. In the case of a forward- or side-looking sensor, however, a general perspective mapping must be used since even very small changes in view angle with respect to the target can produce noticeable shifts in the foreground and background detail.

The most general transformation which relates a point with coordinates (x, y) in one plane to a point (x', y') in another plane is given by:

$$x = \frac{a_{11} x' + a_{21} y' + a_{31}}{a_{13} x' + a_{23} y' + a_{33}} \tag{5.1}$$

$$y = \frac{a_{12} x' + a_{22} y' + a_{32}}{a_{13} x' + a_{23} y' + a_{33}} \tag{5.2}$$

The nine coefficients a_{ij} are constants and can be reduced to eight in number by dividing each expression by one of the coefficients in the denominator, say, a_{33}. Eight equations are necessary to determine the eight coefficients. Two equations are derived from each pair of coordinates, and a set of four corresponding points in each image is thus necessary for the computation. In order for the above equations to be valid, no three points in either plane may lie on a straight line. A simplified transformation suitable for downlooking systems is obtained by setting $a_{13} = a_{23} = 0$. Under this transformation, parallel lines are always mapped into parallel lines; e.g., a square is mapped into an arbitrary parallelogram which may be rotated and translated with respect to the original figure.

It should be noted that errors are introduced into the computation of the coefficients if the control points do not have the same elevation. Even if the control points are chosen correctly, the coefficients obtained are only appropriate for transforming regions at the control point elevation. The rectification problem for side- or forward-looking systems is thus not an easy task. Further study is required both to identify the most effective rectification methods and to incorporate them suitably into the correlation algorithm.

In spite of these problems, even a crude attempt at rectification can often result in a dramatic increase in the correlation function SNR. An example of this is shown in Figs. 5-1 and 5-2.

The upper portion of the first figure shows an electric power substation and the surrounding area. This "reference map" is digitized to 512 × 512 pixels and the GRD is 2 ft. The azimuth is arbitrarily defined as 0° and the elevation is 35°. A "sensed map" of 128 × 128 pixels taken at an azimuth of 6° is shown at the lower right. The corresponding section taken from the reference map is shown at the lower left.

Two sets of correlation surface cross-sections are shown in Fig. 5-2. In each case, the cross-sections are taken along the row and column which intersect at the peak. The results obtained using the sensed map and the corresponding reference map region are shown in Fig. 5-2a, b. A rectified version of the sensed map was generated using the four-point perspective transformation. Using this map, a dramatic increase in peak height was obtained (Fig. 5-2c, d). The absolute values of the two peak maxima are 0.043 and 0.123.

5.3 FALSE MATCH

It is important that the map-matching process be made as insensitive as possible to false match. A false match situation occurs when either the sensed map correlates most strongly with the wrong area of the reference map, or when the sensed map correlates with a similar scene appearing in the wrong reference map. As an example, two arbitrary scenes containing the horizon might cross correlate quite strongly regardless of the ground detail.

a. Reference Map (512 × 512)

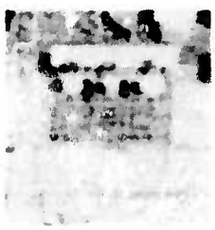

a. Reference Map Section (128 × 128)

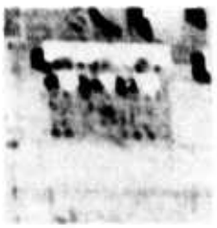

c. Sensed Map (128 × 128)

Fig. 5-1 Power Substation Reference and Sensed Images

PHASE CORRELATION – ORIGINAL SENSED MAP

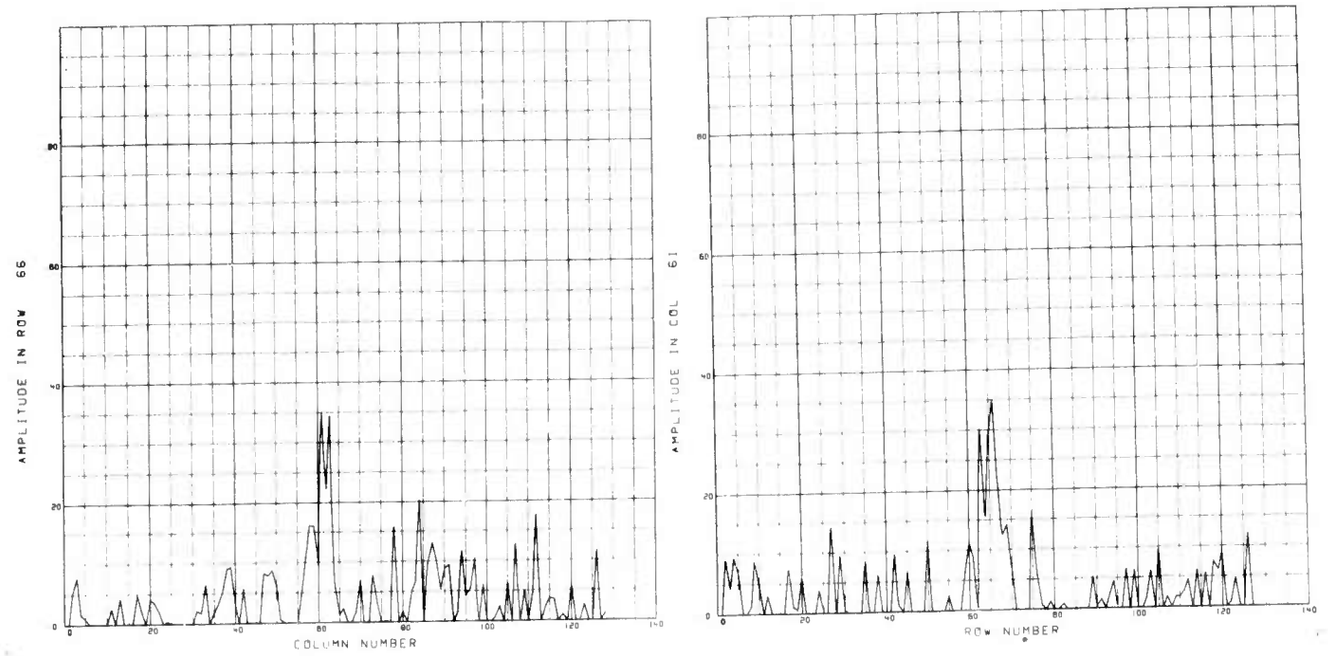

a. Column Cross Section

b. Row Cross Section

PHASE CORRELATION – RECTIFIED SENSED MAP

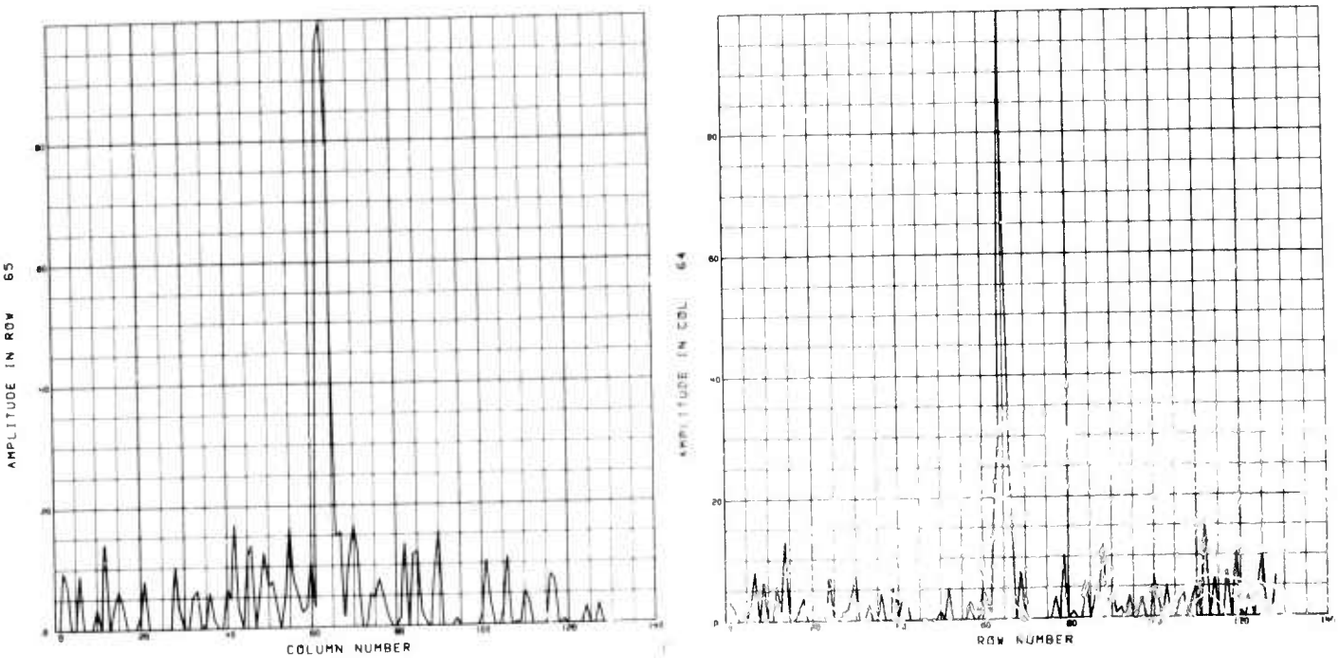

a. Column Cross Section

b. Row Cross Section

Fig. 5-2 Power Substation Correlation Peaks

Equation (2.28) indicates that given a phase correlation peak detection threshold A_o, the false peak probability can be rapidly lowered by increasing the total number of sample points, N_s. It should be emphasized, however, that the increase in sample points must represent a corresponding increase in information content. For example, oversampling a scene will not reduce the probability of false match.

The phase correlation method is less sensitive to false fix than cross correlation. An artificial, but nonetheless impressive example of this is shown in Fig. 5-3. The reference map contains four missile silhouettes embedded in a 128×128 element array of zeros. The sensed map consists of only the figure in the first quadrant embedded in a similar array. The resulting correlation surfaces were generated using cyclic cross and phase correlation. The origin of the correlation surface is located at matrix element (33, 33) which corresponds to the center of the first quadrant. The first missile image correlates with itself and the other three images to almost an equal extent, whereas the discrimination provided by phase correlation is significantly greater.

5.4 REFERENCE MAP GENERATION AND STORAGE

5.4.1 Reference Map Preparation

The basic problem of reference map preparation is to construct a two-dimensional array of samples with a reflectivity assigned to each sample such that the resulting matrix of numbers approximates as closely as possible the expected scene as obtained by the onboard sensor. The reference map should thus, if possible, be constructed using the same resolution and coordinate system used to generate the sensed map and this normally requires a detailed knowledge of the sensor characteristics, and the missile velocity, altitude, and flight path.

In general, it appears that geometrical fidelity is more important than intensity fidelity. It is quite likely that the phase correlator is not as sensitive to encoding errors as, say, cross correlation methods. As has already been noted, the phase correlation

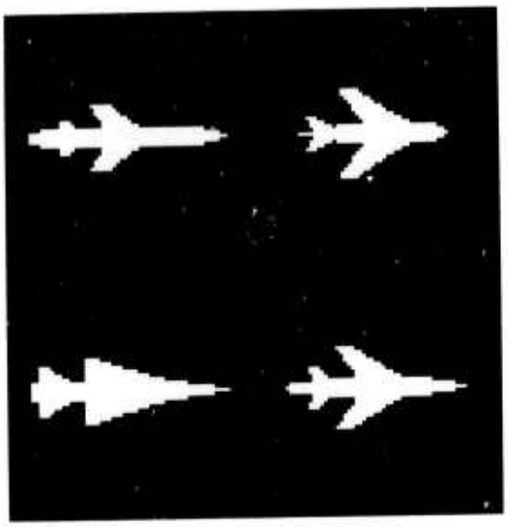

a. Missile Silhouettes

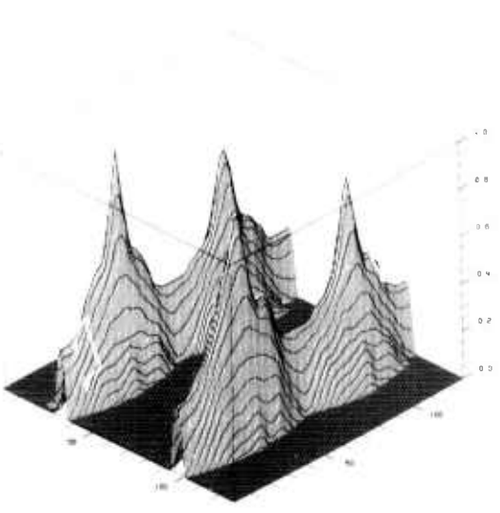

b. Cross Correlation Surface

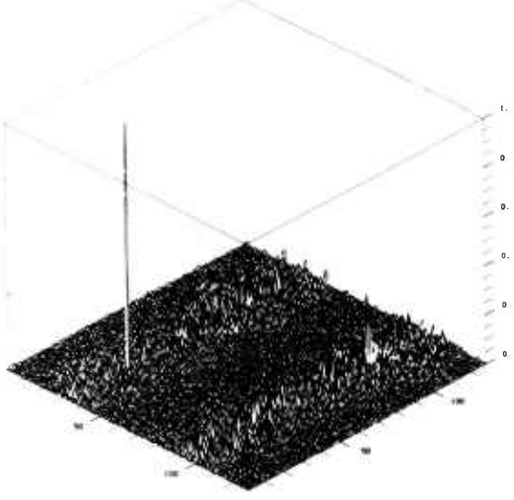

c. Phase Correlation Surface

Fig. 5-3 False Match Performance

function is invariant with respect to either a brightness scaling or a level shift. The ensemble of encoding errors can therefore be treated as an additive noise image with spectral characteristics that are similar to the true scene spectrum. The normalization procedure used to obtain the phase matrix "whitens" each image with respect to its own spectrum and hence, at the same time, acts as an optimal filter for noise having a similar spectral distribution. On the other hand, phase correlation is quite sensitive to perspective differences between images, so that this problem must be approached with some care.

Quite often it is not possible to obtain in advance a reference map of the target area using a sensor operating in the same mode or at the same wavelength as the onboard sensor. In this case it is necessary to prepare a map using data taken from a different source. For example, satellite or aircraft reconnaissance imagery in the visible can be used as a starting point to prepare maps suitable for sensors operating at longer wavelengths.

Although it is tempting to use visible maps directly, i.e., without gray shade alteration, for use with longer wavelength sensors, this usually results in a significant loss of performance, particularly when there are a significant number of contrast reversals. At the other extreme, it is possible to create a "synthetic" reference map starting with high resolution visible or near IR imagery. This process involves the identification of various objects in the reconnaissance scene and a recoding of the scene with the expected sensor return from each object. The resulting map is then convolved with the sensor point-spread function and resampled to conform to the characteristics of the onboard system. Since this process throws away so much of the scene detail, the best approach probably lies in between and involves processing to reduce the effect of the contrast reversals.

5.4.2 Image Preprocessing

As indicated in Section 2.2.4, the phase correlation algorithm is invariant under linear filters which are common to both images. As an example, the use of high pass filtering on both the reference and sensed maps does not change the phase correlation function.

There are, however, a number of nonlinear image operations which may prove to be useful. These include:

- Image clipping
- Histogram equalization
- Brightness transfer function alteration
- Grey shade reduction
- Contouring
- Noise reduction algorithms
- Space-variant filters

The evaluation of alternative preprocessing techniques is greatly simplified by using phase correlation, since the peak height is a direct measure of the degree of scene congruence. Thus, a search can be made for that preprocessing operation which maximizes peak height.

5.4.3 Reference Map Storage

As discussed in Section 4.1, the phase correlation computation can be performed using Fourier phases quantized to eight levels rather than computing the exact normalization using the magnitude of the cross-power spectrum. This process saves computation time and results in only insignificant displacement errors. The success of this method makes possible the onboard storage of reference maps in 3-bit form. Since the phase matrix is Hermitian, the required image information is reduced by 1/2 and requires a storage capability of only 1-1/2 bits/pixel. Since the phase correlation method requires that both input images have the same dimensions, the reference map will normally be divided into a series of smaller arrays. This can be accomplished in several ways and will be discussed in Section 5.5.

It should be pointed out that the quantized-phase storage of reference map subarrays restricts the computation to the selected subarray areas. It is not possible to use these data to construct a phase array corresponding to a new subarray, which does not

directly correspond to any of the original set. For instance, if a sharp correlation peak is obtained at position $\vec{L}_k$ by means of reference map M_k, it might be of interest to repeat the computation using a new reference map section offset by $\vec{L}_k$. If this feature is desired, the reference map must be stored in intensity rather than in phase form. The selected reference map section is then processed in the same manner as the sensed map.

5.5 MAP SEARCH PROCEDURE

The phase correlation algorithm requires that both input images have the same dimensions. This requirement can be satisfied in several ways in the map-matching context, including:

(1) The sensed map is filled out with an array of zeros so that it is the same size as the reference map

(2) The reference map is subdivided into a number of subarrays equal in size to the sensed map

(3) The reference map is subdivided into subarrays larger than the sensed map, and the sensed map is zero filled to the subarray size.

The goal here is to obtain a maximum probable overlap (hence, SNR) with a minimum of computation time. Further study is required to determine which approach is optimal for map-matching. Several studies using method (2) with overlapping subarrays that ensure a common area of $\geq 50\%$ have been made and the results of one of these computations is discussed below.

Figure 5-4 shows an image of the California aqueduct taken with zero azimuth and 10-deg elevation to target. The image is sampled to 512×512 pixels and the GRD is 2 ft in the plane normal to the line of sight. This image served as the reference map. The sensed map, shown at the lower right, was taken with the same elevation and a 7-1/2-deg change in azimuth. This image is sampled to 128×128 pixels. The corresponding section taken from the reference map is shown at the lower left. The reference map was divided into 25 overlapping sections.

a. Reference Map (512 × 512)

b. Reference Map Section (128 × 128)

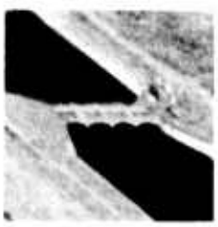

c. Sensed Map (128 × 128)

Fig. 5-4 Aqueduct Reference and Sensed Images

The sensed map was partially rectified prior to the computation using "mosaic" correlation. The map was divided into four 64×64 quadrants, and each quadrant was phase correlated with a corresponding region in the reference map. The four displacement vectors obtained provide the necessary input for the general perspective transformation [Eqs. (5.1) and (5.2)].

The results of the phase correlation acquisition computation using all 25 submaps are shown in Fig. 5-5. The highest peak was obtained at the correct location in Map 21. This peak was normalized to 1.0 and all of the other correlation surfaces use the same normalization factor.

Figure 5-6 shows the reference map partitioning and the phase correlation peak amplitudes obtained with each section. The position of the sensed map in the reference map is outlined in black. Note that the sensed map overlaps Maps 17, 18, 20, 21, 24, and 25. The greatest overlap occurs with Map 21 followed by Map 24.

Reference to Fig. 2-3 shows that a threshold of 0.05 will limit the probability of false match to $< 10^{-6}$ for a 128×128 ($N_S = 2^{14}$) element image. This threshold is exceeded only by Maps 21 and 24.

The use of cyclic correlation creates an ambiguity as to whether a given x or y displacement is a positive displacement greater than 1/2 of the corresponding array length or a negative displacement of less than 1/2 the array size. This ambiguity can usually be resolved on the basis of peak amplitude or, in the case of an image velocity sensor, by a knowledge of the general direction of motion. In the present context, an overlap of at least 50% is required with one of the reference maps so that the uncertainty does not arise. The ambiguity can be eliminated altogether by performing aperiodic correlation. In this case, the $N \times M$ element sensed map and reference submap of the same dimensions are each expanded to $2N \times 2M$ elements using the average image intensity in the added region. This, however, requires somewhat greater than four times as many operations. The increased processing time thus may outweigh the advantage obtained. An example using aperiodic correlation is given in Section 3.5.

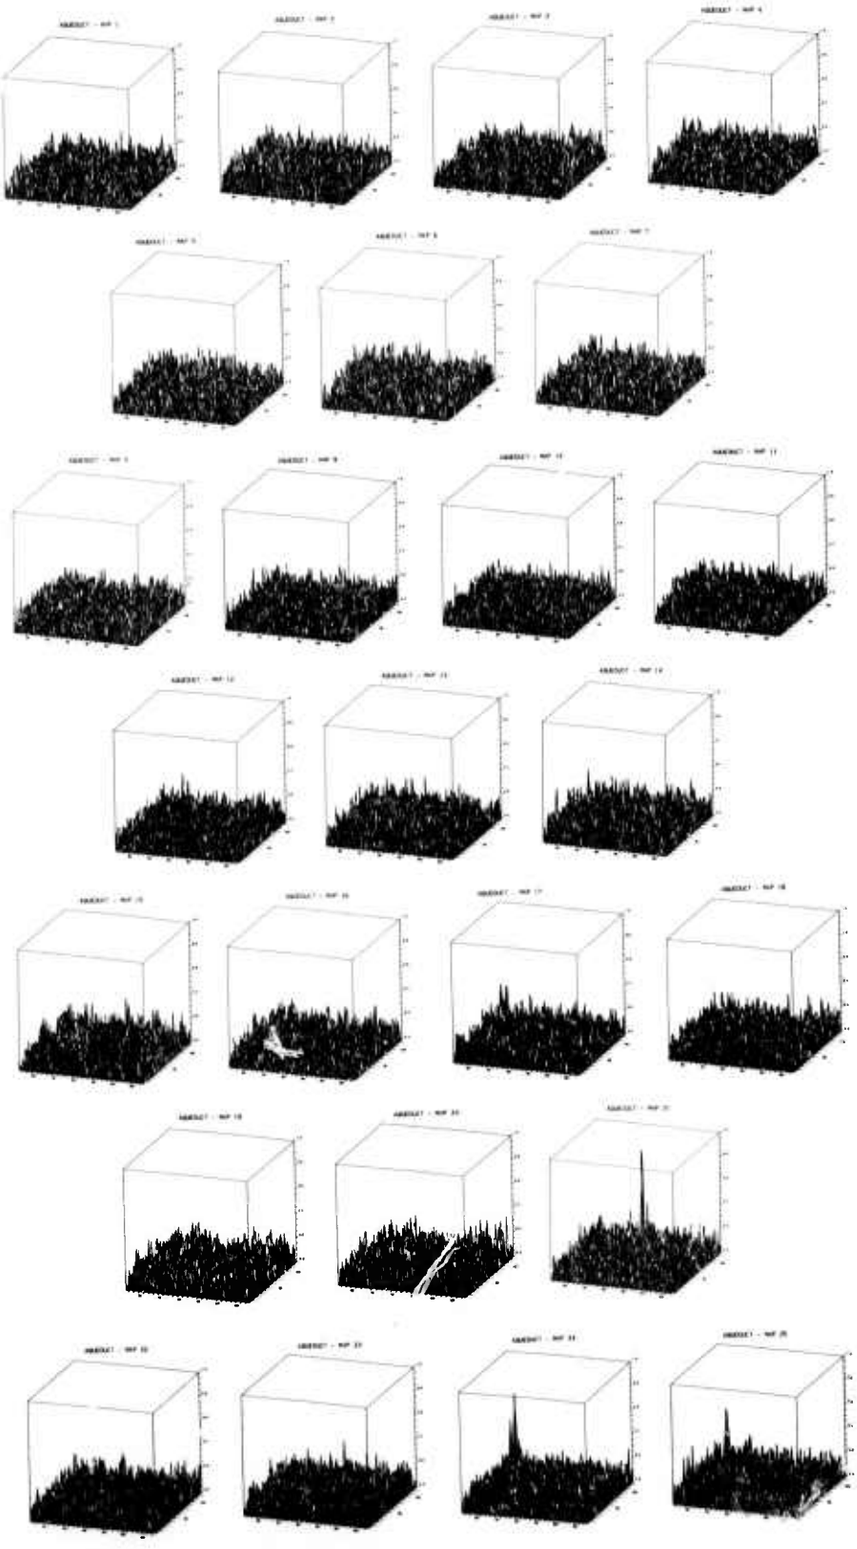

Fig. 5-5 Results of Target Acquisition Experiment Using Aqueduct Image

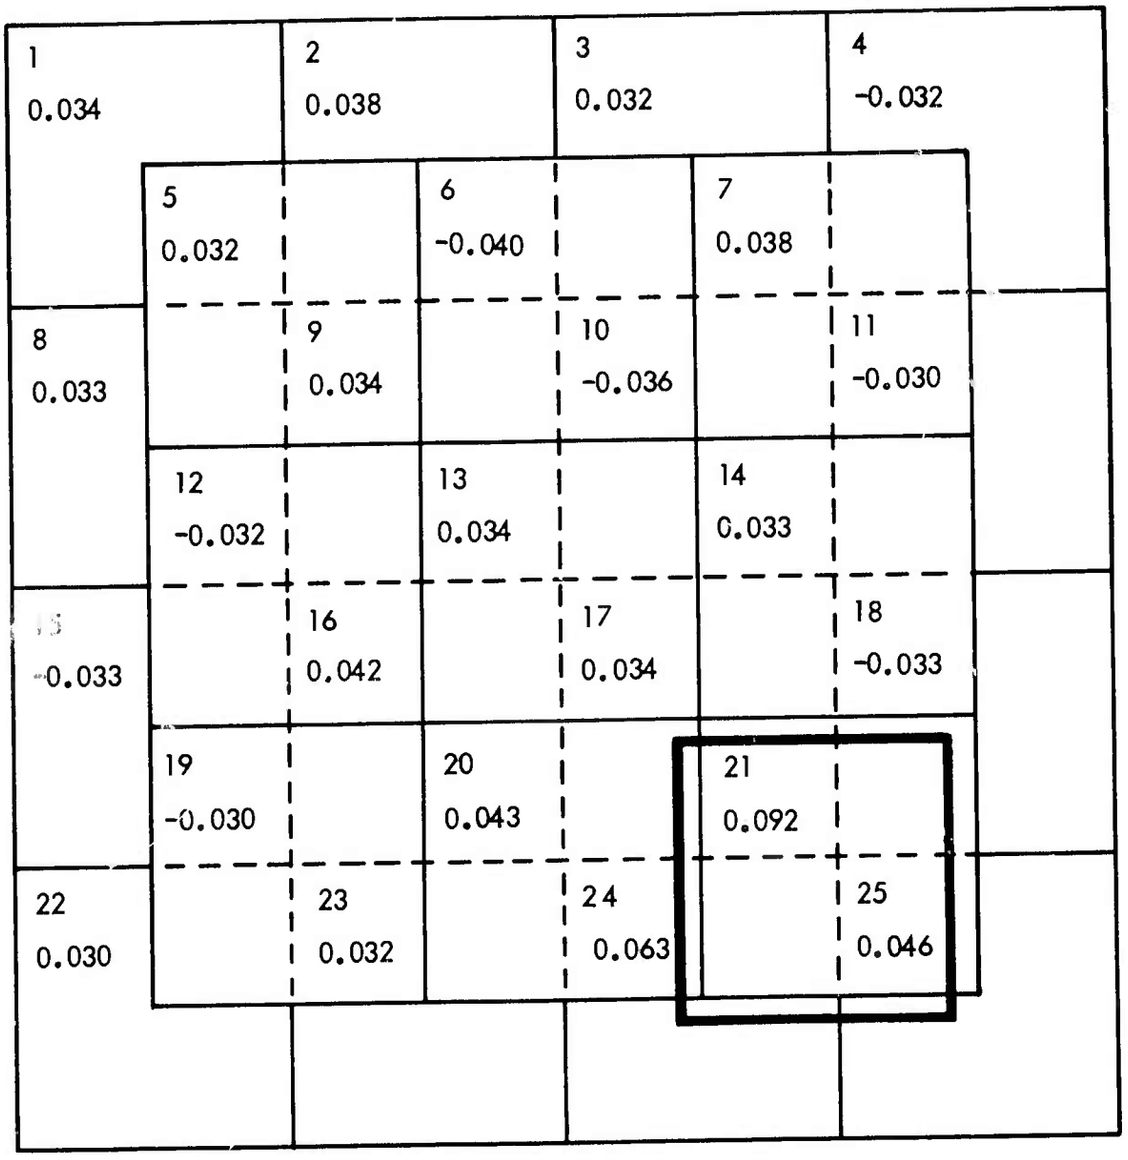

Fig. 5-6 Reference Map Positioning Showing 25 Overlapping Subarrays and Corresponding Phase Correlation Peak Amplitudes. The area corresponding to the sensed map is outlined in black.

The use of 25 overlapping 128 × 128 element reference maps in the above example requires no greater computation time than that required using two 512 × 512 element maps, since the sensed map forward Fourier transform need only be computed once. In addition, the processor data storage requirements are greatly reduced. If the reference map data are stored in phase form, the total map storage requirement is about 2.3 bits/pixel when referred to the original 512 × 512 array.

As currently envisioned, a reference map will be searched in a sequential manner beginning in that region where the probability of match is the highest. Prior to each correlation computation, the sensed map will be rectified and the transformation coefficients stored. Each correlation computation will provide a displacement vector and a peak amplitude which are stored. The highest peak obtained from each correlation computation is compared with a threshold selected to distinguish between "true" and "false" results. If this threshold is exceeded, the computation continues over a limited number of neighboring maps in order to identify that region which provides the highest peak. After this has been found, the displacement and perspective transformation data are output for guidance control.

If the search process is completed without finding a match, the search is repeated using a suitable filter to scale the phase matrix prior to the inverse FFT operation. The assumption will normally be made that target acquisition failure is the result of geometrical distortions, and hence, a low-pass filter designed to optimize performance for the most probable distortion will be chosen. If this filter fails, the process may be repeated with a more severe filter. Note that the sensed map phase matrix is stored after the initial correlation so that subsequent correlations can be performed more rapidly.

Section 6

REFERENCES

1. C. D. Kuglin and D. C. Hines, Phase Correlation Image Velocity Sensor, LMSC-D356260, Lockheed Palo Alto Research Laboratory, Nov 1973

2. -----, "The Phase Correlation Image Alignment Method," Proceedings of the IEEE 1975 International Conference on Cybernetics and Society, Sep 1975, pp. 163–165

3. J. W. Goodman, Introduction to Fourier Optics, McGraw Hill, New York, 1968, Chapter 2

4. J. W. Cooley and J. W. Tukey, "An Algorithm for the Machine Calculation of Complex Fourier Series," Math. Computation 19, April 1965, pp. 297–301

5. R. Emmert and C. McGillem, "Multitemporal Geometric Distortion Correction Utilizing the Affine Transformation," Conference on Machine Processing of Remotely Sensed Data, Purdue University, Oct 1973, pp. 1B24–32