

AD-A142 421

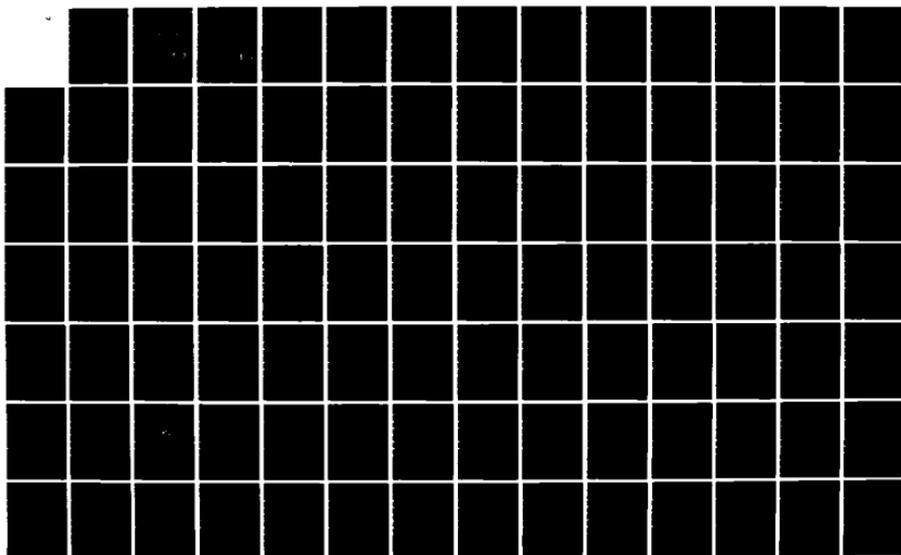
THE MISSING CONE PROBLEM IN COMPUTER TOMOGRAPHY AND A
MODEL FOR INTERPOLA. (U) ILLINOIS UNIV AT URBANA
COORDINATED SCIENCE LAB D A HAYNER JUL 83 R-989
N00014-79-C-0424

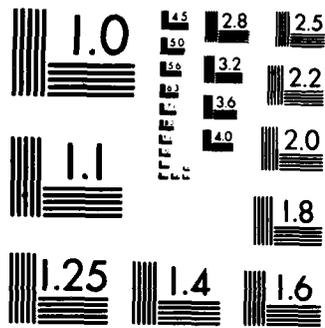
1/2

UNCLASSIFIED

F/G 17/9

NL





MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

12

REPORT R-989 JULY 1983

UICU-ENG 83-2210

AD-A142 421

THE MISSING CONE PROBLEM IN COMPUTER TOMOGRAPHY AND A MODEL FOR INTERPOLATION IN SYNTHETIC APERTURE RADAR

DAVID ALAN HAYNER

DTIC FILE COPY

DTIC
ELECTE
JUN 26 1984
S B

APPROVED FOR PUBLIC RELEASE. DISTRIBUTION UNLIMITED.

84 06 26 027

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER	2. GOVT ACCESSION NO. AD-A142421	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) THE MISSING CONE PROBLEM IN COMPUTER TOMOGRAPHY AND A MODEL FOR INTERPOLATION IN SYNTHETIC APERTURE RADAR		5. TYPE OF REPORT & PERIOD COVERED Technical Report
7. AUTHOR(s) DAVID ALAN HAYNER		6. PERFORMING ORG. REPORT NUMBER R-989; UILU-ENG-83-2210
9. PERFORMING ORGANIZATION NAME AND ADDRESS Coordinated Science Laboratory University of Illinois at Urbana-Champaign Urbana, Illinois 61801		8. CONTRACT OR GRANT NUMBER(s) N00014-79-C-0424
11. CONTROLLING OFFICE NAME AND ADDRESS Joint Services Electronics Program		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)		12. REPORT DATE July 1983
		13. NUMBER OF PAGES 158
		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) MISSING CONE COMPUTER TOMOGRAPHY SYNTHETIC APERTURE RADAR		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) The first part of this thesis considers the missing cone problem in computer tomography. In this problem, an incomplete set of projection data is available from which an image must be reconstructed. The object of the algorithms presented in this thesis is to reconstruct a higher quality image than that obtainable by treating the projections as the only source of information concerning the image to be generated. This is accomplished by treating the problem in terms of spectral extrapolation. With this interpretation, various assumptions concerning the image and other forms of a priori information can be included in the data set to increase		

DTIC
ELECTE
S JUN 26 1984 D

B

20. (continued)

the total information content available.

In order to understand the subtleties of these enhancement algorithms, the spectral extrapolation techniques employed must be well understood. A result of studying the Gerchberg and Papoulis extrapolation techniques is that either can be characterized as a contraction mapping for any realizable discrete implementation. Further more, it is theoretically derived and experimentally verified that these algorithms will in general obtain an optimal solution prior to converging to the unique fixed point. ← 7

In the second part of this thesis, nearest-neighbor interpolation for image reconstruction in spotlight mode synthetic aperture radar is analyzed. A model reflecting the effects of nearest-neighbor interpolation is derived and simulations are provided to support these results. It is also shown that an adaptive presumming operation, performed on the collected azimuth data, will significantly reduce nearest-neighbor interpolation noise.

THE MISSING CONE PROBLEM IN COMPUTER TOMOGRAPHY
AND A MODEL FOR INTERPOLATION IN
SYNTHETIC APERTURE RADAR

BY

DAVID ALAN HAYNER

B.S., University of Illinois, 1979
M.S., University of Illinois, 1980

THESIS

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Electrical Engineering
in the Graduate College of the
University of Illinois at Urbana-Champaign, 1983

Urbana, Illinois

ACKNOWLEDGMENT

I would like to express my sincere appreciation for the constant guidance and encouragement Prof. W. Kenneth Jenkins has provided over the last few years. His insight and subtle prodding has been of immeasurable help. I would also like to thank Prof. David C. Munson, Jr., Prof. Thomas S. Huang and Dr. Jorge L. C. Sanz for their critiques and assistance through the course of these studies. This work was supported under the Joint Services Electronics Program, contract number N00014-79-C-0424.

Author	
Title	
Project	✓
Year	
By	
Distribution	
Availability Codes	
Avail and/or	
Dist	Special
A-1	

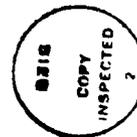


TABLE OF CONTENTS

	Page
1. INTRODUCTION	1
2. ITERATIVE BAND-LIMITED EXTRAPOLATION	6
2.1 Some Iterative and Non-iterative Techniques	7
2.2 Fixed Point Analysis	14
2.2.1 Notation and definitions	16
2.2.2 Fixed point analysis for Papoulis' algorithm	17
2.3 Error Analysis of Papoulis'' Algorithm	24
2.3.1 Notation and definitions	26
2.3.2 Derivation of error equations	28
2.3.3 Discussion of theoretical results	31
2.4 Experimental Results	35
2.4.1 Periodic case	35
2.4.2 General case	41
2.5 Conclusion	51
3. THE MISSING CONE PROBLEM IN COMPUTER TOMOGRAPHY	52
3.1 Introduction to Computer Tomography	54
3.1.1 Projection data	54
3.1.2 Reconstruction techniques	57
3.1.3 Relative image quality and computational requirements	69
3.2 The Missing Cone Problem and Some Solutions	73
3.3 Algorithms for the Solution of the Missing Cone Problem	78
3.3.1 Some algorithms	80
3.4 Computational Details and Experimental Results	87

	Page
3.4.1 Computational details	87
3.4.2 Experimental results	96
3.5 Conclusion	115
4. ANALYSIS OF A JITTER MODEL FOR COORDINATE TRANSFORMATION IN SYNTHETIC APERTURE RADAR	117
4.1 Problem Description and Background	118
4.2 Jitter Analysis	128
4.3 Experimental Results	135
4.4 Conclusion	147
5. CONCLUSION	150
REFERENCES	153
VITA	158

1. INTRODUCTION

Consider the problem of generating the three-dimensional structure of a biological cell by the use of an electron microscope. Any single electron micrograph reveals the cell in only two dimensions. In order to generate a three-dimensional picture, the depth of various structures in the cell has to be determined. The method by which information in depth is obtained is by use of relative perspective. The cell is tilted with reference to the electron beam in order to generate parallax. However, due to the mechanics of electron microscopy, the relative tilt angle over which this information can be collected is limited, typically less than 90 degrees. The problem is how to generate a three dimensional structure from incomplete information in one direction. This is a typical example of an inverse problem. More generally, inverse problems are characterized by the collection of incomplete information or observations concerning a signal and partial constraints on the class of solutions containing the original signal. In recent years inverse problems have become increasingly important; not because they are new problems, but because both the analytical tools and computational means to solve these problems have become available. As a second example, consider the following generic problem. Given an experiment that allows only the partial observation of a data set, say an interval of a signal, extrapolate this signal to obtain data values outside the initial observation interval.

In [1], Gerchberg considers a problem in which the observation consists of an interval of the Fourier transform of a signal. In order

to increase the spatial resolution beyond what the observational band-limit would apparently support, an iterative technique that combines implicit, or 'a priori information,' with the original observations is employed to extrapolate the signal. In this way, he was able to extend the band-limit of the observations thus increasing the resolution of the data set (super-resolution). Since that time, the idea of incorporating a priori information in order to improve the information content of a signal has proliferated. Similar problems exist in a variety of disciplines: radio astronomy, remote sensing and electron microscopy are examples. The specific problem considered in this work is one that arises in computer-aided tomography (CAT) [2]. In brief, tomography consists of reconstructing an image from a set of projections collected over 180 degrees, sometimes referred to as a complete perspective. A projection is an integral transformation from two-dimensional image space to one-dimensional function space. An x-ray photograph is a common example of a projection. The problem considered here is how to reconstruct and/or enhance an image when only a subset of these projections is available, perhaps those spanning only 45 degrees instead of 180 degrees. The limited perspective provided by the incomplete projection data is akin to the problem of estimating the range of distant targets with only a short baseline over which to triangulate. Note, however, that relative cross-range or azimuth positions are easily obtained from the observations on a short baseline. In this work, the reconstruction/enhancement problem is posed as a spectral extrapolation problem. By improving the spectral information content through the inclusion of a priori knowledge, and combining this with the original observation data, final image quality is improved.

A goal of this work is to develop algorithms that will generate higher quality images than those obtained by employing only the observed data. The key problem is to develop a technique that allows the incorporation of various sources of information from different domains. Examples of this type of information include band or spatial limits, specific function values or averages, non-negativeness and shape or size restrictions. The two algorithms developed in this work incorporate these various constraints by iteratively transforming between different domains in which the information can be included. After each transformation, constraints are imposed reducing some measure of error in the data set. In Chapter 3 the reconstruction/enhancement techniques are derived and discussed. In the examples provided the relative trade-offs between techniques are shown and the degree of image recovery possible is illustrated. One of the important features of these algorithms is their relative insensitivity to noise. In one example, with only 35% of the data and a 20 dB signal-to-noise ratio, a reconstruction is obtained that allows nearly complete identification of the image.

A key component in these reconstruction/enhancement techniques is spectral extrapolation. A major problem with the iterative extrapolation techniques employed in this work is their peculiar convergence behaviour. Since these problems have a direct influence on the quality of the reconstructions obtained in the tomographic algorithms, it is important to understand their behaviour. In Chapter 2 some iterative extrapolation techniques are studied in order to obtain a qualitative understanding of their convergence properties. It is shown

that as a consequence of finite length processing intervals, i.e., finite number of samples of a signal or filter, these iterative techniques obtain a distinct minimum error point, and then with more iterations, converge to a fixed point that represents a larger error than the already passed optimal value. In Chapter 2 some results are established for determining when the best solution is obtained and what factors affect the quality of this solution.

The fourth chapter considers an interpolation problem present in synthetic aperture radar (SAR). Although this work is somewhat disconnected from the previous two chapters, it served as the original motivation for much of this work. The conclusion of this thesis will comment on the relationships between SAR and the material presented in chapters two and three. In (SAR) [3], the objective is to generate an image of a scene, usually of terrain, by illuminating the scene with microwave radiation and coherently processing the reflected signals. One of the key problems in developing a real-time digital processor to accomplish the processing task is a data reformatting operation. This operation usually takes the form of a polar-to-rectangular interpolation. In this work, a method is proposed for circumventing polar-to-rectangular interpolation. A 'smart' sampler is used to obtain samples on a keystone [4] raster instead of a polar raster. A nearest-neighbor interpolation scheme is then used to obtain the samples on a rectangular grid. In Chapter 4, this technique is discussed and a mathematical model is proposed. The last half of this chapter is concerned with the verification and testing of this model.

In this thesis, each chapter begins with an introduction of the subject, followed by a presentation of the fundamental concepts as they apply to the problem under consideration. After the concepts and notation have been established, a review of the pertinent literature is presented and briefly discussed. The last portion of each chapter presents new work and ideas, followed by experimental verification and examples.

2. ITERATIVE BAND-LIMITED EXTRAPOLATION

In this chapter the effects of finite processing intervals on iterative techniques for deterministic spectral extrapolation are discussed and two specific methods are analyzed. The term 'finite length processing interval' refers to the finite number of data samples or filter coefficients that can be stored and/or manipulated by a realizable machine in finite periods of time. Since in any practical application these restrictions apply, it will be seen that these effects determine the performance of certain types of algorithms. These techniques are deterministic in the sense that known data are considered to be an observation of a unique, deterministic signal. The extrapolation attempts to approximate the original signal in the sense of a norm rather than with some statistical measure. It will be shown that in some cases an exact extrapolation is possible and in other cases a minimum-norm least-squares solution is either obtained or approached. The goal of this chapter is to qualitatively characterize the properties and numerical behaviour of various extrapolation techniques under the influence of finite length processing intervals. This knowledge is then used as an aid in determining the optimal manner in which to apply a given technique in order to obtain the 'best' solution.

Three topics will be discussed in this chapter. First, Papoulis' algorithm will be analyzed and shown to be a contraction mapping for any finite length processing scheme. Theoretical results of this work will be used to characterize the properties of the fixed point solutions, i.e., the iterative solution. The next section will take a slightly

different track in analyzing these algorithms. Here, the errors introduced by finite length processing intervals (finite records) will be examined in detail to derive an equation bounding the induced error. In the last section, some computer experiments are provided to support and demonstrate the theories presented in this work. It is convenient to start with a general review of modern deterministic spectral extrapolation techniques.

Although this chapter considers iterative techniques almost exclusively, it must be noted that there exists a large class of non-iterative methods. Of these non-iterative techniques only a few will be specifically discussed in this chapter. A comprehensive comparison between iterative and non-iterative methods is available in a paper by Huang et al. [5].

2.1 Some Iterative and Non-iterative Techniques

Gerchberg [1] presents an iterative algorithm for deterministic spectral extrapolation. The object of this extrapolation is to improve the spatial resolution obtainable from Fourier observations that are diffraction limited (in frequency). By extrapolating the spectrum, frequencies above the diffraction limit are recovered and then used to improve the resolution of the target -- i.e., super-resolution. The basis for this technique is that a spatially limited object has an analytic Fourier transform; in fact, the FT is an entire function. A basic theorem of complex variables states any finite interval of an analytic function uniquely determines the whole function [6]. Since the diffraction limited observations provide an interval of this entire

function, it is theoretically possible to recover any desired portion of the complete function.

Gerchberg's algorithm is illustrated in Figure 2.1. Starting with $F_0(j\omega)$, the available observation of $F(j\omega)$, and τ the known spatial limit, an initial approximation is made to the unknown part of the spectrum. Denote this guess as $F'_0(j\omega)$. Next the inverse Fourier transform of $F'_0(j\omega)$ is found generating $f_1(t)$. Clearly $f'_1(t)$ is a better approximation to $f(t)$ than $f_1(t)$ because the erroneous signal outside τ has been removed. By use of a Fourier transform, $F'_1(j\omega)$ is obtained from $f'_1(t)$. The process of substituting $F_0(j\omega)$ into $F_1(j\omega)$ generating $F'_1(j\omega)$ reduces the error in the spectrum a second time. This iterative process is repeated, reducing the error in two steps until a satisfactory result is obtained.

The dual of spectral extrapolation is spatial or time extrapolation in which an interval of the time domain signal is known and a band-limit in the frequency domain is available. The basis of solution in this problem is that a band-limited, finite energy signal has a uniformly convergent Taylor series approximation. Given any interval of the time signal, in theory it is possible to calculate all the derivatives around some point in the known interval, to generate the Taylor series and to calculate the unknown function to arbitrary accuracy for any point in time. Papoulis [7] discusses this variation of Gerchberg's algorithm and presents some theoretical results including a proof of convergence to the unique solution. This proof is based on the repeated application of Parseval's relation. Essentially the same proof can be used to show convergence of Gerchberg's algorithm. Either of these techniques can be

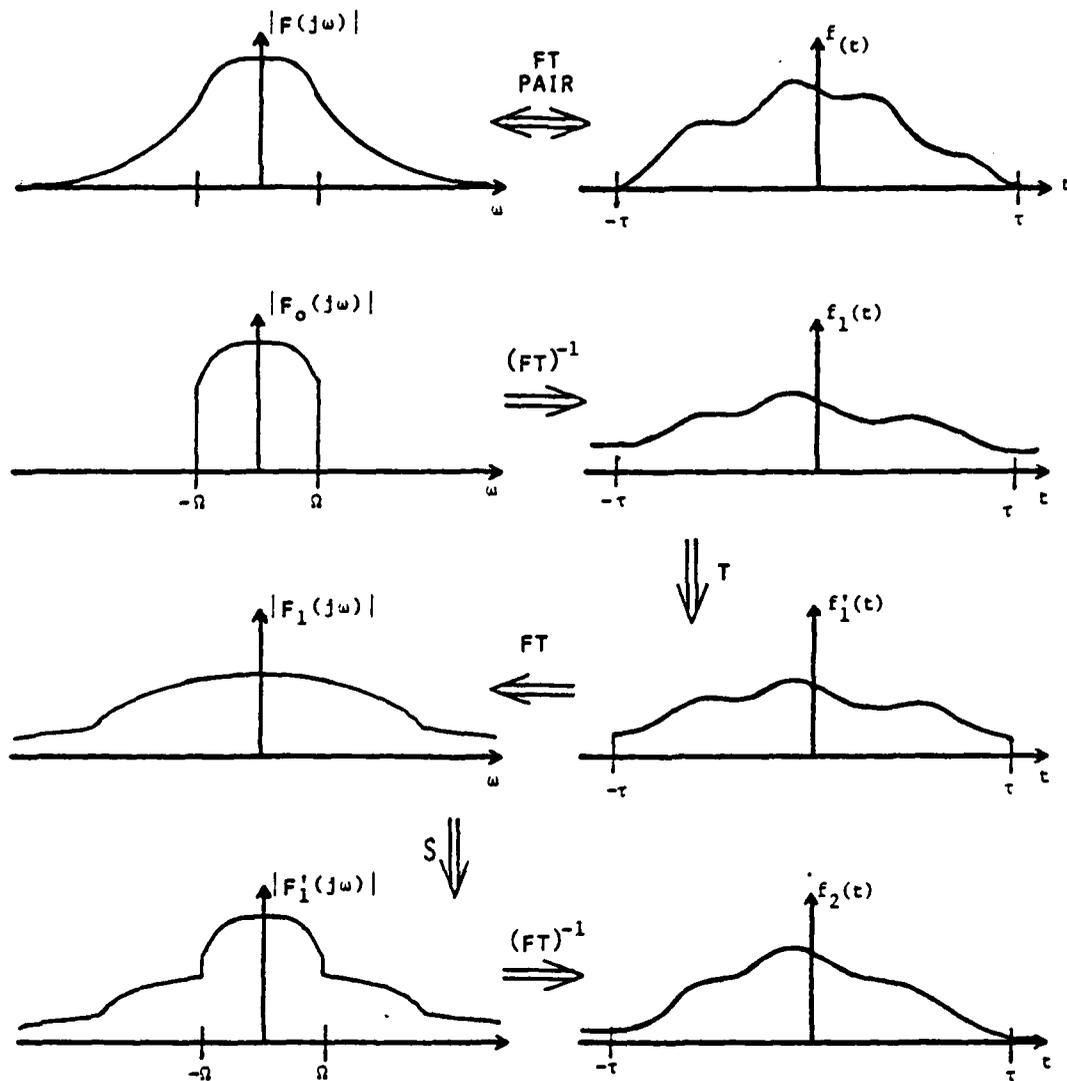


Figure 2.1 Gerchberg's algorithm.

implemented in one domain by observing that the process of transforming to the opposite domain, truncating to a known time or band-limit and transforming back is equivalent to convolving with a low-pass filter in the original domain. This is illustrated for Papoulis' algorithm in Figure 2.2 where the switches realize the substitution procedure. It should be pointed out that both Papoulis' and Gerchbergs's algorithms are special cases (see Sanz and Huang [8]) of an iterative method for the solution of Fredholm integral equations of the first kind. This technique was first proposed by Landweber in [9].

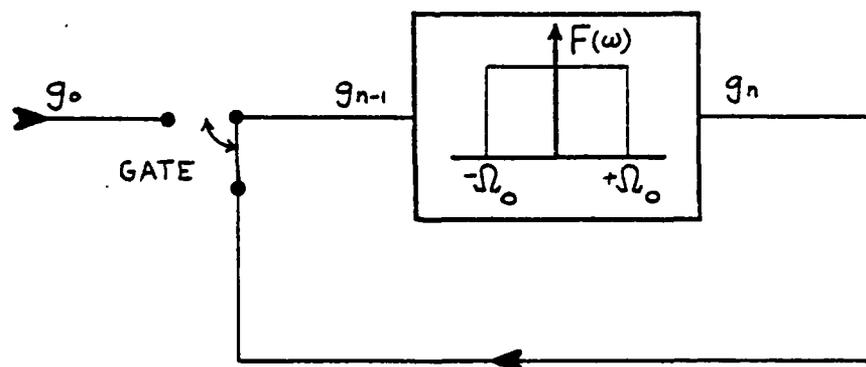


Figure 2.2 Papoulis' algorithm.

Using operator notation, Papoulis' algorithm can be expressed as

$$g_n = g_0 + \bar{D}^L B g_{n-1} \quad (2.1)$$

where B represents the low-pass filter and operator \bar{D}^L zeros, or time limits, the appropriate interval of the low-pass filtered signal so g_0 can be substituted. The subscript n denotes the iteration count.

A non-iterative extrapolation technique proposed by Sabri and Steenaart [10] involves the use of an extrapolation matrix. This matrix, E_n , is obtained by solving the difference equation (2.1) for g_n in terms of g_0 , i.e.,

$$g_n = A_n g_0 \quad (2.2)$$

where

$$A_n = \sum_{i=0}^{n-1} H^i, \quad H = \bar{D}^L B. \quad (2.3)$$

As Sabri and Steenaart point out, because Papoulis' algorithm converges and equations (2.2) and (2.3) are an alternate realization of (2.1), the solution obtained from the A_n matrix will also converge as n approaches infinity.

A second non-iterative technique was proposed by Cadzow [11]. This technique is essentially a two-step algorithm. The first step, which Cadzow recognizes as the most difficult step, involves the solution of a Fredholm integral equation of the first kind (see above discussion). In the second step, the solution to the Fredholm integral equation is low-pass filtered to obtain the final result. That Sabri's and Cadzow's methods are very similar has been the topic of significant debate in the literature [12,13,14].

There is, however, a significant problem with either the iterative techniques of Papoulis and Gerchberg or the non-iterative methods of Sabri and Cadzow. The assumption of a continuous model for the Fourier and time (spatial) domains is the primary cause of these difficulties. The first problem is that the data used and the calculations employed must be discrete. Because the data are sampled, the analyticity or Taylor series arguments presented earlier are no longer valid. As a result, a unique solution no longer exists (in general). In fact, an infinity of solutions is available. A second problem is that the ideal filtering implied in (2.1) cannot be implemented. This is because only a finite number of samples (either of the signal or filter coefficients) can be stored and manipulated. These finite record length effects further degrade the performance of the algorithm. These comments are also relevant to non-iterative techniques because an extrapolation matrix of only a finite size can be manipulated, and only a finite number H^i terms could be calculated and included. The point of this work is to qualitatively characterize the effects of finite records and finite iterations on the performance of Papoulis' and Gerchberg's algorithms, and consequently, on similar techniques.

Jain and Ranganath [15] were the first to assume a discrete model for the extrapolation problem. One of their results was to show that that discretization of Papoulis' or Gerchberg's algorithms results in convergence to the minimum-norm least-squares solution (MNLS). This result further implies that this solution is obtained from Papoulis' algorithm only in the infinite record case. In order to circumvent the finite record length problem, Jain and Ranganath phrased the

extrapolation problem in terms of linear algebra. The $2M+1$ point extrapolation solution vector g_* is found in terms of a generalized matrix inverse and the $2L+1$ point observation vector g_0 . Denoting the signal from which the observation is obtained as f_0 , the observation g_0 can be represented as

$$g_0 = S f_0 \quad (2.4)$$

where the S matrix is $2L+1 \times \infty$, with $s_{ii} = 1$, $-L \leq i \leq L$ and $s_{ij} = 0$ for all other i, j . This operator selects a subsequence of f_0 as the observation. Assuming that f_0 is a band-limited sequence, $B f_0 = f_0$, then

$$g_0 = S B f_0. \quad (2.5)$$

In (2.5), B consists of samples of the low-pass operator where

$$b_{ij} = \frac{\sin[\Omega_0(i-j)]}{\pi(i-j)}, \quad -\infty \leq i, j \leq \infty. \quad (2.6)$$

Denote by S^L the operator that extrapolates the $2L+1$ long vector with an infinite number of zeros, i.e., S^L is $\infty \times 2L+1$ and $s_{ii}^L = 1$, $-L \leq i \leq L$. Then solving (2.5) for f_0 via a generalized inverse, we obtain

$$g_* = B S^L (B^L)^{-1} g_0, \quad (2.7)$$

which of course is the MNLS solution to (2.5) or (2.4). In (2.7) the truncated operator B^L is as defined in (2.6) for i, j such that $-L \leq i \leq L$, $-L \leq j \leq L$.

An iterative algorithm Jain and Ranganath propose is

$$z_n^L = (1/\gamma) z_0 + z_{n-1}^L - (1/\gamma) B^L z_{n-1}^L, \quad z_0 = g_0 \quad (2.8)$$

$$g_n^M = B^M z_n^L \quad (2.9)$$

which also converges to the MNLS solution. In (2.8), z_n^L is a dummy vector $2L+1 \times 1$, that the algorithm iterates on. The final solution (to the desired number of samples $2M+1$) is obtained after (2.8) converges, by low-pass filtering z_n^L . As with equation (2.7), this is a well-known linear algebra result. Details of the derivation can be found in [16]. The point of (2.7), (2.8) and (2.9) is that the best obtainable solution is the MNLS result. This is a consequence of the discrete nature of the data and implementations. As pointed out earlier, an infinite number of solutions exist in the sampled data case. These various algorithms (Papoulis, Sabri, Cadzow and Jain) simply converge to different solutions. A reason that these different techniques converge to different solutions results from the underlying models assumed for the various techniques. This is explained in more detail by Huang and Sanz in [17]. Huang and Sanz [18] also point out that the discrete and continuous models are conceptually well-unified. A major emphasis of this chapter is to identify the effects of finite record lengths on the solutions generated by discrete realizations of Papoulis' and Gerchberg's algorithms.

2.2 Fixed Point Analysis

In this section Papoulis' algorithm will be analyzed in terms of a contraction mapping. It will be proved that the finite record implementation of Papoulis' algorithm is a contraction mapping and consequently has a unique fixed point. The behaviour of the algorithm with respect to this fixed point will then be studied to establish some bounds on the error and convergence properties. The underlying model is

the continuous-continuous model discussed in [17]. This implies that g_0 is a subset, i.e., an observation, of samples of a continuous signal. The objective of the extrapolation is to generate samples of the continuous signal outside the observation interval. This is the assumed model throughout this chapter. The system to be analyzed is represented by the difference equation

$$g_n^M = g_0 + \bar{D}^L B^M g_{n-1}^M \quad (2.10)$$

where M represents the record length, n the number of iterations, g_n^M is a $2M+1 \times 1$ vector and the matrix $\bar{D}^L (= [I - D^L])$ is $2M+1 \times 2M+1$ where the matrix D^L consists of $d_{ii} = 1$, for i , $-L \leq i \leq L$ and $d_{ij} = 0$ for all other i, j . The matrix B^M was described earlier and is $2M+1 \times 2M+1$. Lastly, g_0 is extended with sufficient zeros to be consistent with the rest of the equation. By defining a substitution operator T and interval I_T (where needed, k will denote individual elements of a vector)

$$T[x(k)] = \begin{cases} g_0(k) & \text{for } -L \leq k \leq L, \text{ i.e., } k \in I_T \\ x(k) & \text{for } k < -L \text{ or } k > L, \text{ i.e., } k \notin I_T \end{cases} \quad (2.11)$$

equation (10) can be written as

$$g_n^M = TB^M g_{n-1}^M, \quad g_{-1}^M = 0. \quad (2.12)$$

The objective now is to prove that the mapping TB^M is a contraction mapping for all values of M . If this is true, then a unique fixed point will exist for all finite values of M .

2.2.1 Notation and definitions

In the following discussion, the symbol $||*||$ denotes any valid norm in the space under consideration. The spaces will be subsets of R^M or R^∞ for vectors and $R^{M \times M}$ or $R^{\infty \times \infty}$ for operators. For the operator spaces, $||*||$ will denote the induced norm [19]. Clearly for the R^∞ and $R^{\infty \times \infty}$ cases this norm must exist in the sense of convergence in order to have a complete space. In all cases a linear space, a Banach space, will be assumed. Much of the literature [19,20] on contraction mappings and fixed point theorems use the symbol $d(*,*)$ to denote a measure of distance (or energy). To be consistent with this notation, $d(*,*)$ will be used in conjunction with $||*||$ where

$$d(x,y) \equiv ||x-y||, \quad x,y \in R^M \quad \text{or} \quad R^\infty. \quad (2.13)$$

and

$$d(x,0) = ||x||, \quad x \in R^M \quad \text{or} \quad R^\infty. \quad (2.14)$$

Two definitions needed pertain to contraction and non-expansive mappings. Let A be a normed linear space and C contained in A . The mapping $G: C \rightarrow A$ is a contraction mapping if there exists a constant γ , $0 < \gamma < 1$ such that for all $x,y \in C$

$$d(Gx,Gy) \leq \gamma d(x,y). \quad (2.15)$$

Based on this definition is the Contraction Mapping Theorem: If G is a contraction mapping on a closed subset C contained in a complete linear space A , then there exists a unique fixed point $x_* \in C$, i.e. $Gx_* = x_*$. In addition, the sequence $x_n = Gx_{n-1}$ converges to x_* for any initial point $x_0 \in C$ and

$$d(x_*, x_n) \leq \frac{\gamma^n d(x_1, x_0)}{1-\gamma}. \quad (2.16)$$

A less restrictive mapping which allows γ to be equal to 1 is known as a non-expansive mapping. These mappings may have multiple fixed points as opposed to the uniqueness property of contraction mappings. Letting A be a normed linear space and C contained in A , the mapping $H: C \rightarrow A$ is non-expansive if for all $x, y \in C$

$$d(Hx, Hy) \leq d(x, y). \quad (2.17)$$

If strict inequality holds in (2.17) for all $x, y \in C$, $x \neq y$, then H is strictly non-expansive in which case one or more fixed points may exist. In order to guarantee that only one exists, the image of C under H must be compact [20]. For this case H has a unique fixed point and $x_n = Hx_{n-1}$ converges to x_* for all $x_0 \in C$ [20].

2.2.2 Fixed point analysis for Papoulis' algorithm

Denote the Banach space containing finite energy sequences as S . These finite energy sequences may be infinite in extent, $x \in S^\infty$ which is contained in R_∞ , or of finite extent, $x^M \in S^M$ contained in R^M . Denote the subsets of S^∞ and S^M containing finite energy band-limited sequences by S_Ω^∞ and S_Ω^M respectively where Ω represents the band-limit. The first step in proving B^M to be a contraction mapping is to show that B is non-expansive.

Theorem 1: $B: C \subset S^\infty \rightarrow S^\infty$ is a non-expansive mapping for all $x, y \in S^\infty$.

Proof: Let $x, y \in S_{\Omega}^{\infty}$, $x \neq y$. Clearly $Bx=x$ and $By=y$. Thus

$$d(Bx, By) = d(x, y) \leq d(x, y) \quad (2.18)$$

which is the definition of a non-expansive mapping. Now let $x, y \in S^{\infty}$, $x \neq y$. In this case $Bx=x'$ and $By=y'$, implying that $x', y' \in S_{\Omega}^{\infty}$. Consider the Fourier transforms of the sequences x and x' (the following argument also holds for y and y')

$$x = x(k) \longleftrightarrow X(e^{j\omega})$$

$$x' = x'(k) \longleftrightarrow X'(e^{j\omega}).$$

Since the signals x and y are low-pass filtered, x' is band-limited to $\Omega < \pi$ and

$$\int_{-\pi}^{\pi} |X(e^{j\omega})|^2 d\omega > \int_{-\pi}^{\pi} |X'(e^{j\omega})|^2 d\omega \quad (2.19)$$

because $X'(e^{j\omega})=X(e^{j\omega})$ for $-\Omega \leq \omega \leq \Omega$ and $X(e^{j\omega})$ is non-zero outside this interval. By Parseval's relation

$$\|x'\| < \|x\| \quad (2.20)$$

or

$$\|Bx\| < \|x'\| \quad (2.21)$$

and the same holds for y and y' . For $x, y \in S^{\infty}$ there exists a $z \in S^{\infty}$ such that $x-y=z$. Substituting into (2.20) and (2.21)

$$d(Bx, By) \equiv \|Bx-By\| = \|B(x-y)\| = \|Bz\| < \|z\|. \quad (2.22)$$

Since

$$\|z\| = \|x-y\| \equiv d(x, y)$$

we obtain

$$d(Bx, By) < d(x, y). \quad (2.23)$$

Thus for all $x, y \in S^\infty$, B is a non-expansive mapping. QED.

Theorem 2: $\|B^M x^M\| < \|Bx^M\|$, $x^M \in S^M$, $B^M \subset R^{M \times M}$ and $B \subset R^{\infty \times \infty}$ for all $x^M \neq 0$

Proof: $B^M x^M = y^M \in S^M$, $Bx^M = y \in S^\infty$ and $y = y^M$ over the $2M+1$ points on the central interval denoted by I_M and define y^M equal to zero outside I_M . Further, y is non-zero over some interval outside I_M . Therefore $\|y^M\| < \|y\|$ or

$$\|y^M\| = \|B^M x^M\| < \|Bx^M\| = \|y\|. \quad (2.24)$$

QED.

Theorem 3: $\|B\| = 1$.

Proof: Let $x, y \in S^\infty$. Since B is non-expansive (Theorem 1)

$$d(Bx, By) \leq \gamma d(x, y) \quad \text{for } \gamma \leq 1 \quad (2.25)$$

or

$$\|B(x-y)\| \leq \gamma \|x-y\|, \quad \gamma \leq 1. \quad (2.26)$$

Since we also know that

$$\|B(x-y)\| \leq \|B\| \|x-y\| \quad (2.27)$$

implies that $\|B\| \leq 1$. Next consider $x \in S_\Omega^\infty$, then $Bx = x$ or $\|Bx\| = \|x\|$. Since $\|Bx\| \leq \|B\| \|x\|$ and $\|Bx\| = \|x\|$, the implication is that $\|B\| \geq 1$. This coupled with $\|B\| \leq 1$ implies that $\|B\| = 1$ for all $x \in S^\infty$. QED.

Theorem 4: $\|B^M\| < 1$.

Proof:

$$\|B^M x^M\| \leq \|B^M\| \|x^M\| \quad (2.28)$$

From Theorem 2:

$$\|B^M x^M\| < \|B x^M\|, \quad (2.29)$$

and

$$\|B^M x^M\| < \|B x^M\| \leq \|B\| \|x^M\| \quad (2.30)$$

implies that

$$\|B^M x^M\| < \|x^M\|. \quad (2.31)$$

The only way for (2.29) and (2.30) to be consistent is for $\|B^M\| < 1$.

QED.

The next to last step in proving that Papoulis' algorithm as described in (2.11) is a contraction mapping is to prove that B^M is a contraction mapping.

Theorem 5: $B^M: S^M \rightarrow S^M$ is a contraction mapping for all $x^M \in S^M$.

Proof: Let $x^M, y^M \in S^M$. There exists a z^M such that $x^M - y^M = z^M$.

From

$$\|B^M z^M\| \leq \|B^M\| \|z^M\|, \quad (2.32)$$

and substitution of $z^M = x^M - y^M$,

$$\|B^M z^M\| \leq \|B^M(x^M - y^M)\| \leq \|B^M\| \|x^M - y^M\| = \|B^M\| \|z^M\| \quad (2.33)$$

is obtained. Applying the definition of $d(*, *)$ and using $\gamma = \|B^M\| < 1$ (Theorem 4),

$$\|B^M(x^M - y^M)\| \equiv d(B^M x^M, B^M y^M) \leq \gamma d(x^M, y^M) \equiv \|B^M\| \|x^M - y^M\|. \quad (2.34)$$

The above meets the definition of a contraction mapping, therefore, B^M is a contraction mapping. QED.

It is quite easy to show that the substitution operator T , equation (2.11), is a non-expansive mapping. The proof will be omitted from this work. Combining this fact with the contraction mapping properties of B^M it is easy to show that TB^M is a contraction mapping.

Theorem 6: TB^M is a contraction mapping for all $x^M \in S^M$.

Proof: Let $x^M, y^M \in S^M$ and $x_1^M, y_1^M \in S^M$. Denote $B^M x^M$ as x_1^M and $B^M y^M$ as y_1^M . Then

$$d(TB^M x^M, TB^M y^M) = d(Tx_1^M, Ty_1^M) \leq d(x_1^M, y_1^M) \quad (2.35)$$

because T is non-expansive. By definition of the contraction mapping property of B^M , (Theorem 5)

$$d(x_1^M, y_1^M) = d(B^M x^M, B^M y^M) \leq \gamma d(x^M, y^M), \quad \gamma < 1. \quad (2.36)$$

Clearly then

$$d(TB^M x^M, TB^M y^M) \leq \gamma d(x^M, y^M), \quad \gamma < 1 \quad (2.37)$$

and TB^M is a contraction mapping. QED.

Since the mapping TB^M describing the discrete implementation of Papoulis' (and equivalently Gerchberg's) algorithm is a contraction mapping, a unique fixed point exists and the following theorem can be stated.

Theorem 7: The iterative algorithm

$$g_n^M = TB^M g_{n-1}^M \quad (2.38)$$

will converge to a unique fixed point g_*^M for any $g_1 \in S^M$. Further, the error between g_*^M and g_n^M is bounded by

$$d(g_*^M, g_n^M) \leq \frac{\gamma^n d(g_1, g_0)}{1-\gamma}. \quad (2.39)$$

Proof: Since TB^M is a contraction mapping, by the Contraction Mapping Theorem, a unique fixed point exists; further, the error is determined by equation (2.16) or equivalently by (2.39). QED.

Some comments are in order concerning the above statements. First, (2.39) is only a conservative bound on the error. In the next section a tighter bound is derived. Second, the fixed point to which (2.38) converges is a function of g_0 , not of any initial guess over the unknown intervals in the extrapolation. This property is a necessary requirement of the fixed point theorem. Two factors that affect the solution are the band-limit, Ω , and the record length employed, M . In section 3 of this chapter, effects of the band-limit will be examined and discussed. To conclude this section, the effects of M on the solution and behaviour of Papoulis' algorithm will be qualitatively studied.

That the value of M determines whether B is a contraction or non-expansive mapping would indicate M has some effect on the fixed point itself. Since Papoulis' algorithm is guaranteed only to converge to the MNLS solution with infinite records, the implication is that for finite records, the solution is sub-optimal (in the sense that the MNLS solution is optimal). The next theorem states that $\|B^M\|$ approaches $\|B\|$ as the record length goes to infinity. This fact is used to infer that g_*^M approaches the MNLS solution (g_*).

Theorem 8:

$$\lim_{M \rightarrow \infty} \|B^M\| = \|B\| \quad (2.40)$$

Proof: Since $\|B\| = 1$, $\lambda_{\max} = 1$ together with $\|B^M\| < \|B\|$ implies $\lambda_{\max}^M < 1$ because $\lambda_{\max}^M \leq \|B^M\|$. By definition of norm equivalence, two norms $\|\cdot\|_a$ and $\|\cdot\|_b$ are said to be equivalent if and only if there exist two positive numbers α and β such that

$$\alpha \|x\|_a \leq \|x\|_b \leq \beta \|x\|_a \quad \text{for all } x \in S^\infty. \quad (2.41)$$

Consider now the l_2 induced norm on S^∞ :

$$\|B\|_{l_2} = [\lambda_{\max}(A^*A)]^{1/2} \quad (2.42)$$

For

$$b_{ij} = \frac{\sin \Omega(i-j)}{\pi(i-j)} \quad (2.43)$$

$\|B\|_{l_2}$ exists. This could be inferred also from the fact that B is non-expansive. Further, since for any ϵ there exists an M such that

$$\frac{\sin \Omega i}{\pi i} < \epsilon, \quad \text{for some } i \geq M \quad (2.44)$$

$\|B^M\|_{l_2}$ defined by (2.42) for i , $-M \leq i \leq M$ converges uniformly. By employing the concept of norm equivalence, $\|B^M\|$ also converges uniformly to $1 = \|B\|$. Because B^M is positive definite, all the eigenvalues, λ^i for $i=1, \dots, M$ are distinct and $0 < \lambda^i \leq \|B^M\|$ for all $i=1, \dots, M$. Further, λ_{\max}^M converges uniformly to $\lambda_{\max} = 1$. QED.

Consider equation (2.37) with $\gamma = \|B^M\|$. Clearly, as M approaches infinity $\|B^M\|$ approaches $\|B\|$ and γ goes to 1 implying that the mapping becomes non-expansive as opposed to being a contraction. Theorem 8 also implies that g_*^M approaches g_* as M goes to infinity.

The result is that the error associated with the fixed point decreases as the record length, M , increases. Further, the error bound established by (2.39) indicates that as M gets large and consequently γ approaches 1, the convergence is slower. Combining these two facts, qualitative error curves such as those illustrated in Figure 2.3 can be sketched. These error bounds are very conservative and do not identify specific sources of error. What this bound does establish is that as M gets large, the fixed point for equation (2.10) or (2.12) approaches the fixed point for equation (2.1), i.e., the MNLS solution. In the next section a tighter and more descriptive error bound is derived that identifies specific sources of error.

2.3 Error Analysis of Papoulis' Algorithm

In order to derive a more descriptive error bound, the sources of error need to be identified and accounted for in some manner. Clearly, one source of error is the fact that an infinite number of iterations can never be realized. Consequently, an error term accounting for this component can be identified. As discussed in the previous section, a second obvious source of error is caused by the finite records employed. These two error terms will be defined and bounds found on their respective contributions to the total error.

To achieve this analysis, two cases will be considered. In the first case, the vector g_0 is considered to be a subset, i.e., an observation, of samples of a continuous, band-limited, periodic signal $f(t)$. In the second case, g_0 is an observation of samples of a finite

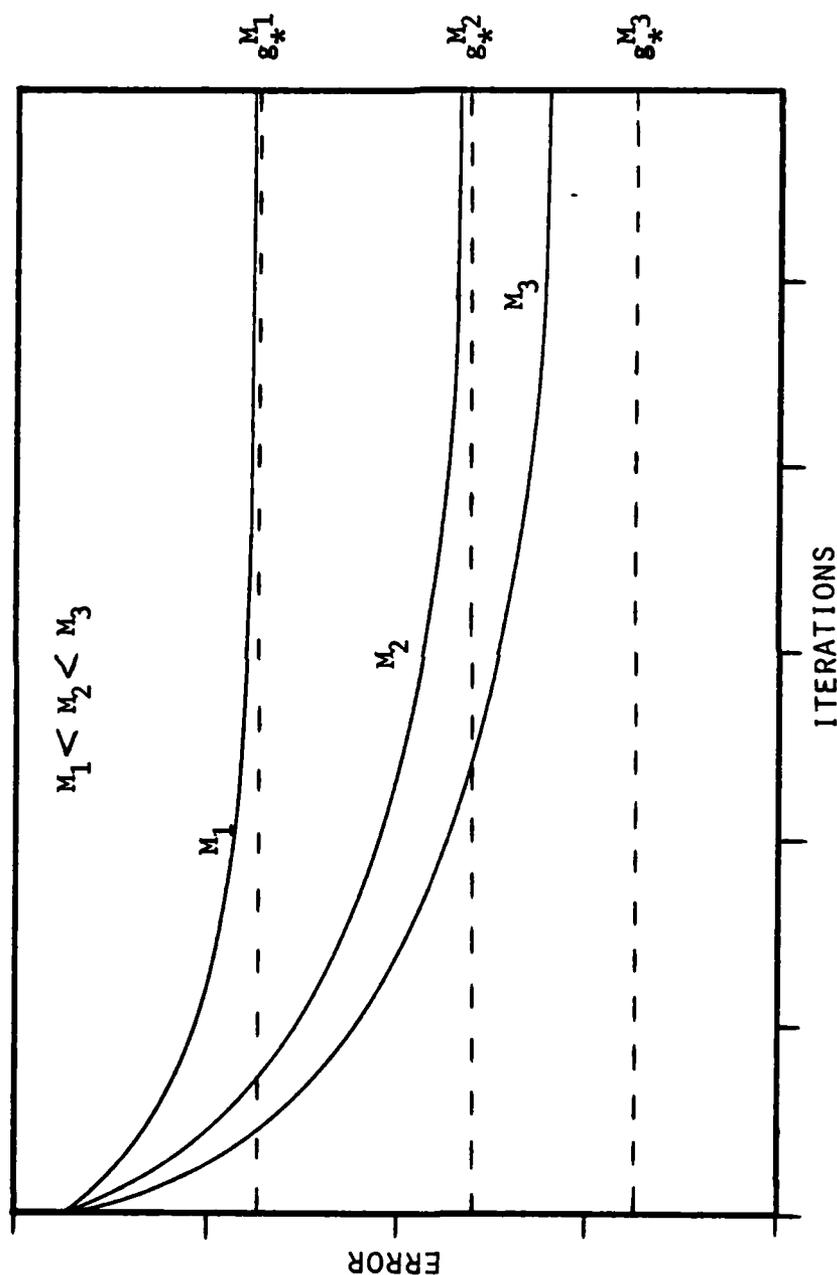


Figure 2.3 Convergence of g_n^M to the associated fixed point, g_*^M , for different record lengths, M .

energy, band-limited continuous signal $f(t)$; this case will be referred to as the general case. As discussed earlier, this model is the continuous-continuous case presented in [17]. An equation will be derived that bounds the error for the general case. The error bound for the periodic case will be shown to be a special case of the non-periodic solution.

2.3.1 Notation and definitions

Define TE_n^M as the magnitude of the total error between g_n^M of (2.10) and the MMLS solution obtained from (2.7). As before, superscripts will denote record lengths and subscripts will represent iterations. The error TE_n^M is bounded by the sum of two terms. Referring to Figure 2.4, let E_n^M represent the error between g_n^M and g_n , i.e., the difference between equations (2.10) and (2.1). Another interpretation of E_n^M is that it represents the error caused by a record length of M after n iterations have been performed.

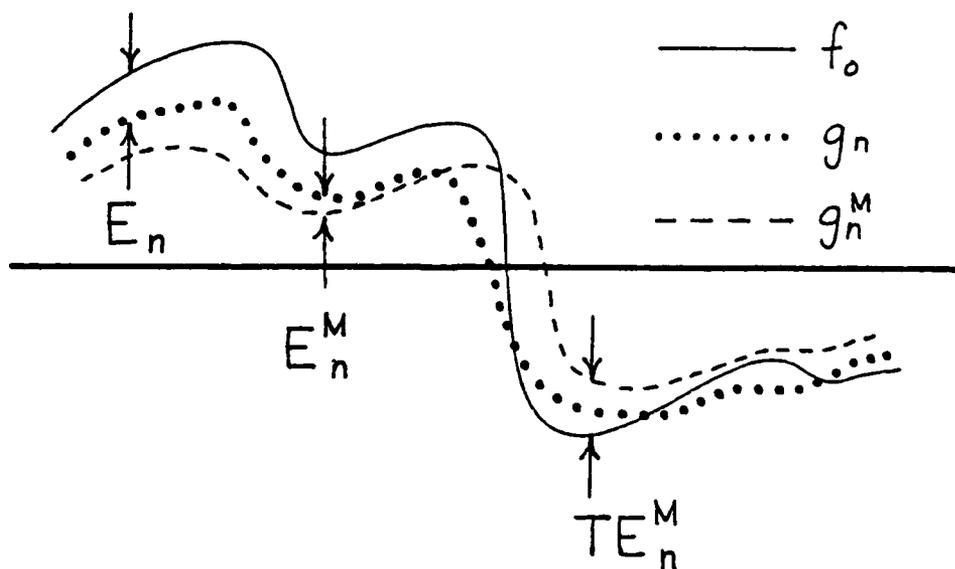


Figure 2.4 Definition of error terms.

Error remaining by performing only a finite number of iterations with infinite record lengths is denoted by E_n . This term represents the error between g_n and g_* (the MNL solution). It should be noted that this is not the same as the error bound of equation (2.39). In (2.39) the bound is between the fixed point for (2.12) implemented with a record length of M and the result of (2.12) after n iterations. Summing the above two error terms:

$$TE_n^M \leq E_n^M + E_n. \quad (2.45)$$

Some asymptotic properties of these terms are now considered. Since E_n^M is the error caused by using finite records, as the record length approaches infinity, E_n^M should approach zero, i.e.,

$$\lim_{M \rightarrow \infty} E_n^M = 0. \quad (2.46)$$

If E_n denotes the error remaining after completing a finite number of iterations (with infinite records), then

$$\lim_{n \rightarrow \infty} E_n = 0. \quad (2.47)$$

Equations (2.46) and (2.47) imply

$$\lim_{n \rightarrow \infty} \lim_{M \rightarrow \infty} TE_n^M = 0. \quad (2.48)$$

These equations represent necessary requirements that any bound on the total error must satisfy.

2.3.2 Derivation of error equations

Restating Sabri and Steenaart's extrapolation matrix solution to (2.1)

$$g_n = A_n g_0 \quad (2.49)$$

where A_n is given in equation (2.3). This represents the infinite record case and, consequently, A_n is $\infty \times \infty$ and the g_n are $\infty \times 1$. For finite records, (2.1) becomes (2.10), and (2.3) is modified to account for the finite records as

$$g_n^M = \sum_{i=0}^n (H^M)^i g_0 = A_n^M g_0. \quad (2.50)$$

In this case, the extrapolation matrix A_n^M is $M \times M$ and the vectors g_n^M are $M \times 1$. Since the extrapolation matrix

$$A_\infty = \lim_{n \rightarrow \infty} \sum_{i=0}^n (H)^i \quad (2.51)$$

does not exist (Jain [15]), the symbol A_∞ will instead be used to represent the pseudo-inverse extrapolation matrix that obtains the MNLS solution (equation (2.7)). Using this notation

$$g_* = A_\infty g_0 \quad (2.52)$$

where the dimensions of A_∞ and g_* are dependent upon the number of points desired in the extrapolation.

First, a bound on the error term E_n^M will be found. As in the previous section, error will be defined as the norm of the difference between the two terms under consideration. From the previous definition, E_n^M is expressed as

$$E_n^M = \|g_n^M - g_n\|. \quad (2.53)$$

If g_n^M is padded with sufficient zeros after M terms, then

$$E_n^M \leq \|D^M[g_n^M - g_n]\| + \|\bar{D}^M g_n\|. \quad (2.54)$$

The motivation for this decomposition is to isolate the error in the extrapolation interval (record length) from the error outside this interval. Substituting (2.49) and (2.50) into (2.54) gives

$$E_n^M \leq \|D^M[A_n^M g_0 - A_n g_0]\| + \|\bar{D}^M A_n g_0\| \quad (2.55)$$

or

$$E_n^M \leq \left[\|D^M[A_n^M - A_n]\| + \|\bar{D}^M A_n\| \right] \|g_0\|. \quad (2.56)$$

It is easy to demonstrate that (2.56) meets the necessary requirements of (2.46). As M approaches infinity the first term on the right-hand side of (2.56) goes to zero. By the properties of \bar{D}^M , as M approaches infinity the norm of $\bar{D}^M A_n$ goes to zero.

Next, a bound on E_n is found. This term represents the error remaining after completing a finite number of iterations with infinite records. By definition,

$$E_n = \|g_n - g_*\|. \quad (2.57)$$

Substituting (2.49) and (2.52) into (2.57) gives

$$E_n = \|[A_n - A_\infty]g_0\| \quad (2.58)$$

or

$$E_n \leq \|A_n - A_\infty\| \|g_0\|. \quad (2.59)$$

Since Jain proved that (2.1) goes to the MNLS solution as n approaches infinity, then E_n must go to zero, thereby satisfying the necessary

requirement of (2.47).

An equation bounding the total error can now be formed. Combining the results of (2.56) and (2.59) into (2.45) generates:

$$TE_n^M \leq \left[\|D^M[A_n^M - A_n]\| + \|\bar{D}^M A_n\| + \|A_n - A_\infty\| \right] \|g_0\|. \quad (2.60)$$

Since the components of (2.60) satisfy their respective necessary conditions, then the above inequality must satisfy the requirements of (2.48).

Inspection of the terms in (2.60) indicates that the error is a function of M , Ω and n . An important feature of (2.60) is that the relationships between TE_n^M and the various parameters are independent of the signal from which g_0 is an observation. This is important in the sense that properties established for one $f(kt)$ (and thus g_0) are essentially the same for any $f(kt)$. It will be assumed that $f(kt)$ is scaled such that the norm of g_0 is equal to one. However, it should be noted that TE_n^M is not independent of the observation length. This is represented by L and by reference to the definitions of H and H^M (eq. (2.3)), its inclusion in TE_n^M can be identified.

An error bound for the periodic case is easily derived from (2.60). In most cases involving periodic sequences, the record lengths available for processing are substantially longer than the periods involved. Consequently, since ideal filtering can be accomplished (or simulated), A_n^M and A_n are identical. Since there is no information lost outside the processing interval due to truncation, the second term of (2.60) is also zero. The error equation is then simply

$$TE_n^M \leq ||A_n - A_\infty|| = E_n. \quad (2.61)$$

The error in this case is due solely to the finite number of iterations performed and is the E_n term of (2.60); E_n^M is zero for the periodic case. Equation (2.61) indicates that with infinite records the error will approach zero as n goes to infinity. Clearly, infinite record length processing is only possible with periodic signals. Therefore, when periodic signals are being extrapolated, a solution as accurate as desired is obtainable. But this result has been well-known for quite some time and, reasonably enough, considers extrapolation of a periodic signal as an interpolation problem.

2.3.3 Discussion of theoretical results

The motivation behind (2.60) and (2.61) is to obtain some qualitative information concerning the error generated by finite records. Equation (2.60) contains two competing factors. Error remaining after a finite number of iterations with infinite records (E_n) decreases to zero monotonically with increasing n . In fact, because Papoulis' algorithm is a gradient technique, this error is decreasing with at least linear convergence [15]. Opposing this factor is the E_n^M term. As the number of iterations increases, error due to finite records is increasing from zero. Some comments about the rates of increase and decrease will be made later. Another interpretation concerning E_n^M is that the iterates g_n and g_n^M are diverging from each other. Increasing the record length would slow this divergence, i.e., cause E_n^M to increase at a slower rate. Note that E_n is independent of

M. In Figure 2.5 the dotted line represents E_n and the dashed lines represent E_n^M for different M. The sum of these curves is plotted as the solid lines.

The most distinctive feature of these curves is the minimum. These curves indicate that the best solution in terms of a norm is obtained with a finite number of iterations. A larger record length results in a better solution at the cost of more iterations. Another prediction of (60) is that a larger observation interval will uniformly shift the TE_n^M curves downward. Assume for the moment that E_n^M is constant with respect to L. It is easy to show that for any n, increasing L forces E_n to be uniformly shifted down. Adding this lower E_n curve to the assumed constant E_n^M curves results in a downward shift of the TE_n^M curves. If E_n^M is not assumed to be constant, it seems reasonable in light of the effects of M on E_n^M , that a larger L will cause E_n to grow more slowly. This would apply an additional downward shift to the TE_n^M curves. The effects of altering the extrapolation bandwidth are discussed with the experimental results.

Some additional comments are in order concerning the rate of divergence between g_n^M and g_n . In Papoulis' algorithm, a sequence is filtered in a non-ideal manner thus generating some error. This sequence plus error is then returned to the input and processed a second time generating error on both the original sequence and the feedback error. A difference equation can model this as:

$$s(n) = \alpha s(n-1) + \beta, \quad s(-1)=0 \quad (2.62)$$

where the coefficients α and β determine the stability and limits respectively. The general solution to this equation is

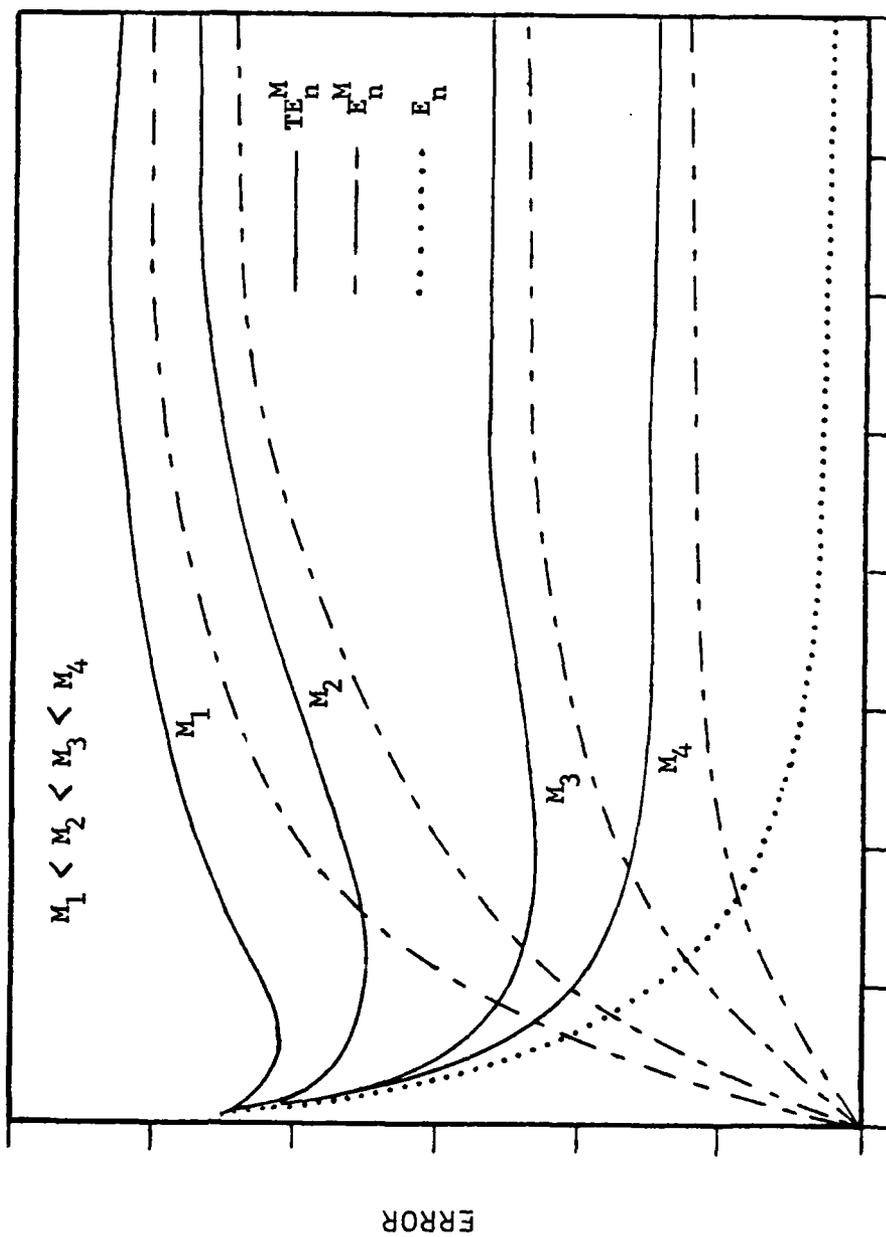


Figure 2.5 Theoretical error curves.

$$s(n) = \frac{\beta}{1-\alpha}(1-\alpha^{n+1}). \quad (2.63)$$

Obviously if α is greater than 1, $s(n)$ diverges to infinity, otherwise $s(n)$ approaches $\beta/(1-\alpha)$. Since the process representing α in Papoulis' algorithm is the low-pass operation, which in the previous section was demonstrated to be at least non-expansive, and in the truncated implementation, a contraction mapping, α is less than 1. The coefficient β represents the amount of noise injected into the processed sequence with each iteration (assumed constant for a given M). The curves plotted in Figure 2.5 are obtained with this model of E_n^M : the error between g_n^M and g_n . As M approaches infinity, both the constant amount of noise injected with each pass is reduced (β gets smaller) and g_n^M diverges from g_n at a slower rate (α goes to one). These combine to generate a lower asymptotic error. The last comment is consistent with the theory put forth in the previous section concerning the effects of M .

The sum effect of M on the total error as a function of n follows. As M increases, a distinct minimum error value is obtainable which decreases with increasing M . At some, perhaps large, value of M the rate of error increase of E_n^M is matched by the decreasing error of E_n and the distinct minimum is no longer present. For M even larger than this value, g_n^M asymptotically approaches g_n . The iterates, g_n^M , for this large value of M approach g_n^M , not g_n .

2.4 Experimental Results

In this section, experimental evidence is presented to support the theory discussed in the two previous sections. To this end, two cases are considered. The first case is the periodic case and the second is the non-periodic situation. All error plots are the l_1 norm of the difference between either g_n and $f(kt)$ or g_n^M and g_* for the periodic and non-periodic cases respectively. In the non-periodic situation, g_* is the MNL solution.

The ideal signal, $f(t)$, is illustrated in Figure 2.6. A periodic or non-periodic interval of this sequence is used depending upon the case under study. From this interval, a sub-interval is used as the observation. For example, in the periodic case 512 points of the sequence are used as the ideal waveform. An observation of length $2L+1$ is constructed by zeroing all points except those between $-L$ and L . In the general case, 600 points of $f(kt)$ are the original sequence and the observation is constructed in the same way as for the periodic observation.

2.4.1 Periodic case

In the periodic case considered, the sequence used is that described above with $L=200$ and $\Omega = 0.115$ radians. An issue is whether the algorithm implemented with circular convolution and a version realized with an FFT will produce the same results. The reason for questioning this fact is that it was felt that the cumulative effects of

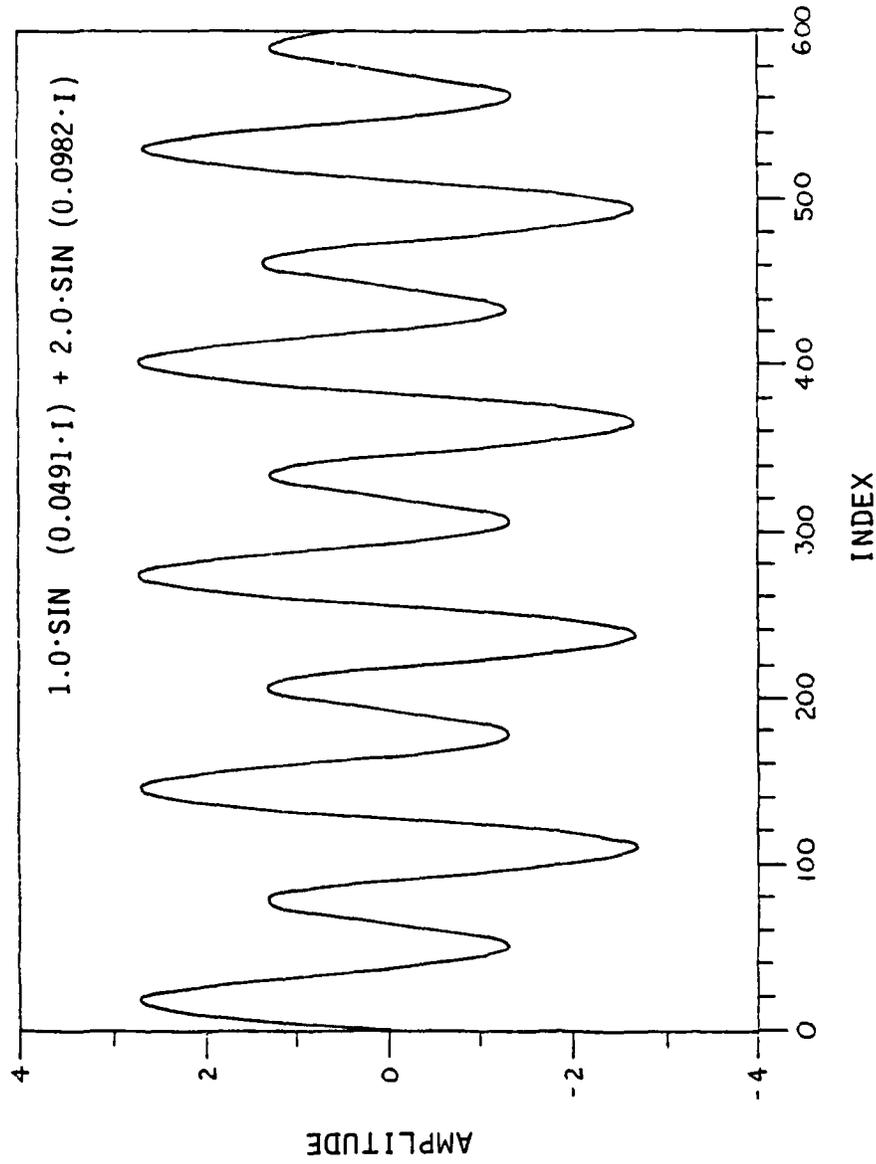


Figure 2.6 The ideal signal.

round-off error after a large number of iterations may cause some divergence in the results generated by the two implementations. Figure 2.7 illustrates the error curves for these two realizations; they are identical. To achieve these results, the same filter must be implemented with both circular convolution and the FFT method. The periodic impulse response coefficients for the convolutional technique are obtained from the inverse discrete Fourier transform of an ideal discrete low-pass filter. Circularly convolving with these coefficients is then the same as the implicit periodic processing achieved by the FFT. Both techniques simulate infinite record length processing if the periods are chosen correctly. Consequently, there are no finite record length errors and the remaining error is due solely to the lack of performing an infinite number of iterations, E_n . This supports equation (2.61) and the assumed behaviour of infinite record length processing.

A consequence of finite register length effects, sampling and the chosen Ω , the E_n curve must ultimately level off, enter a limit cycle or even start to increase. If the algorithm is correctly implemented this will not occur until after a large number of iterations. The result illustrated in Figure 2.8 indicates that this leveling off does not occur until at least 3000 iterations. Further, this curve demonstrates the linear convergence assumed for this type of algorithm, i.e., Papoulis' with infinite records. Figure 2.9 illustrates the results after 60 and 3000 iterations. The improvement is quite evident, although the cost may be prohibitive.

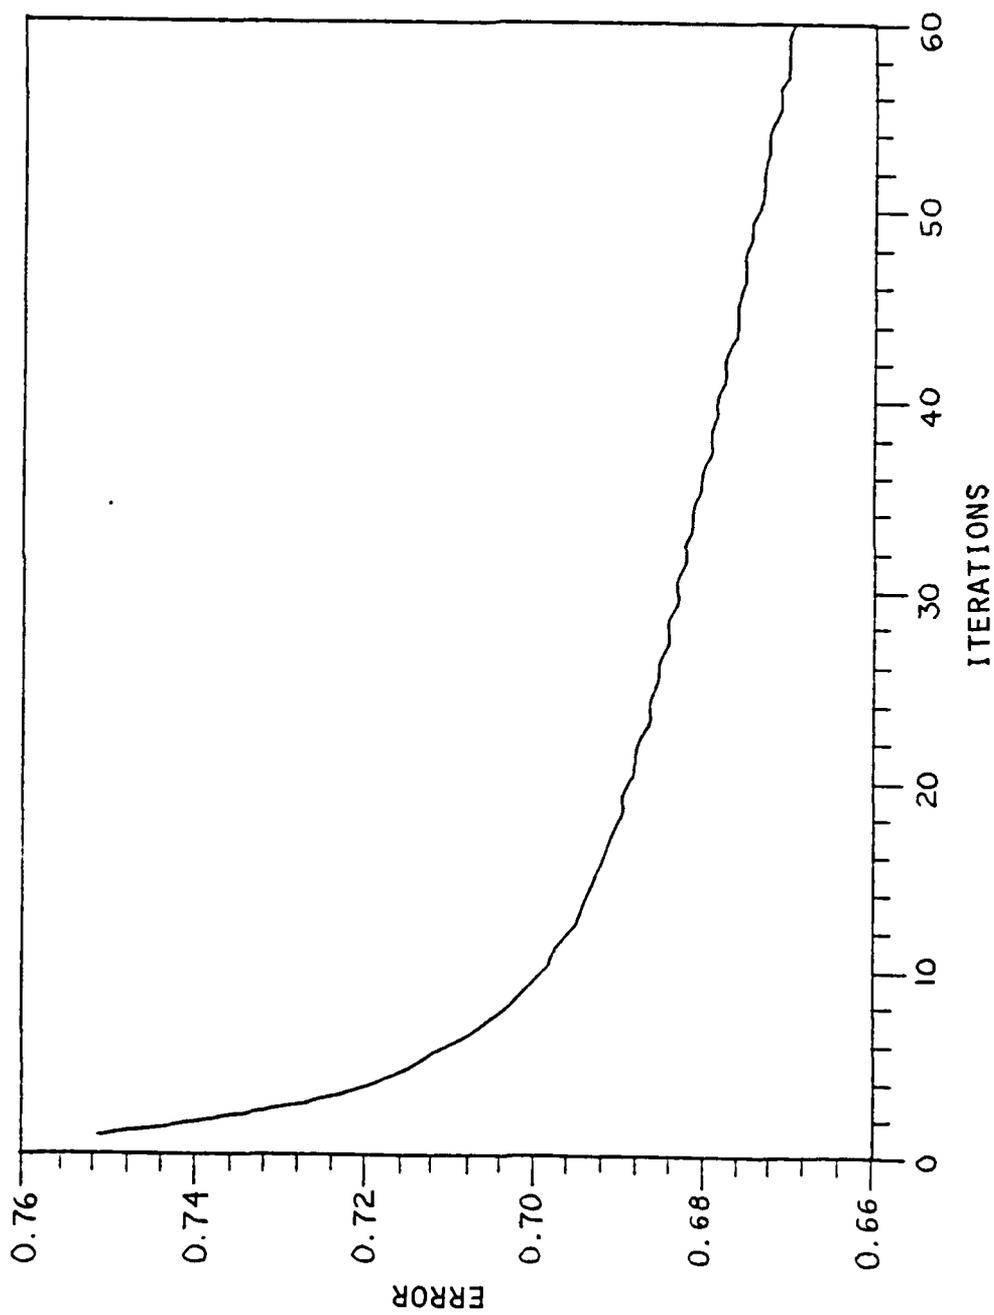


Figure 2.7 FFT and convolution results.

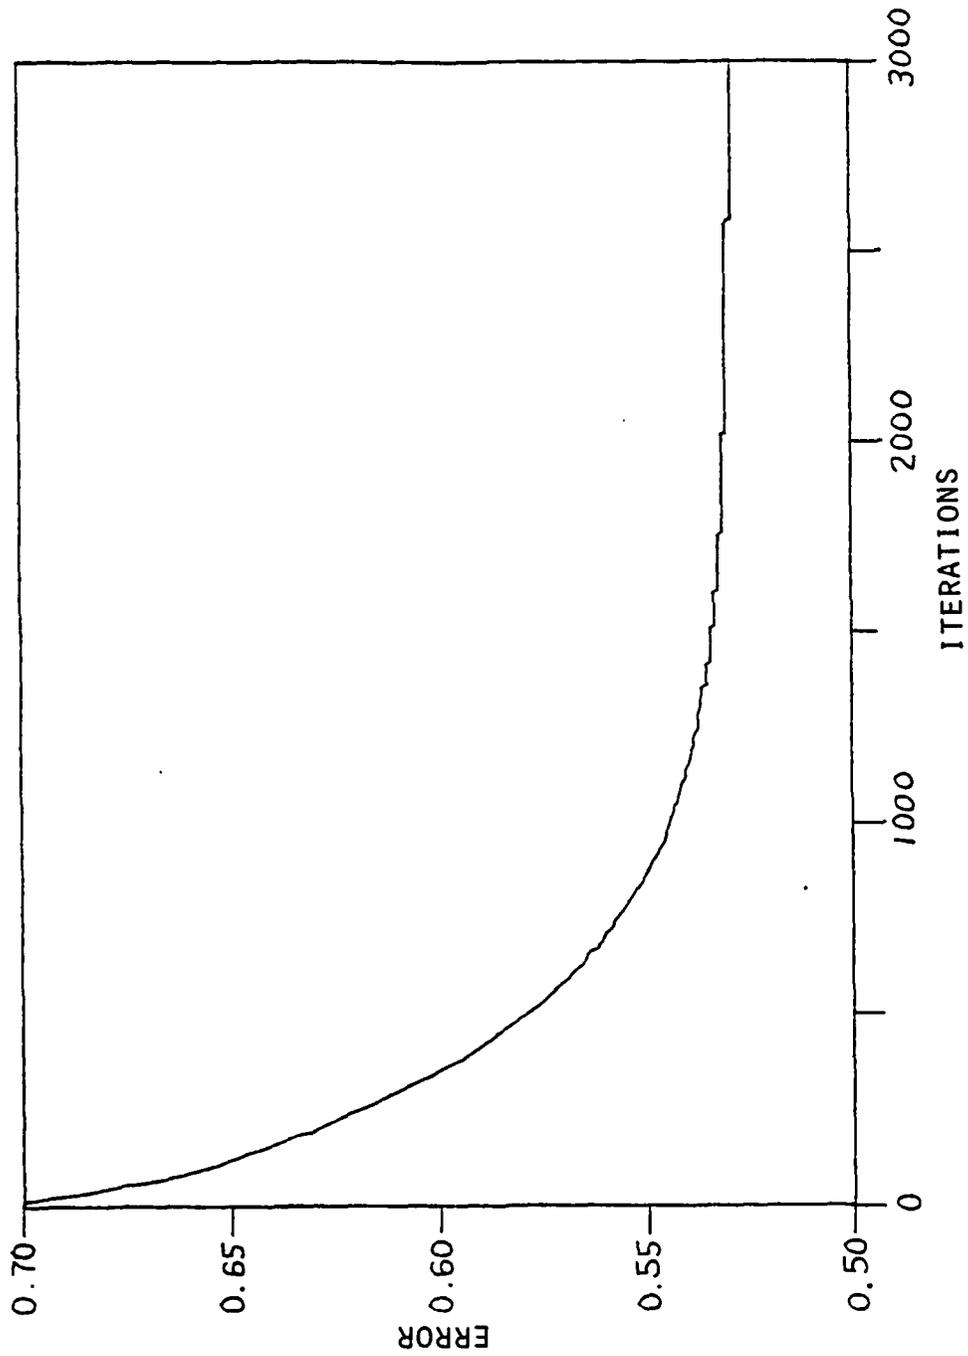


Figure 2.8 Error for the periodic case.

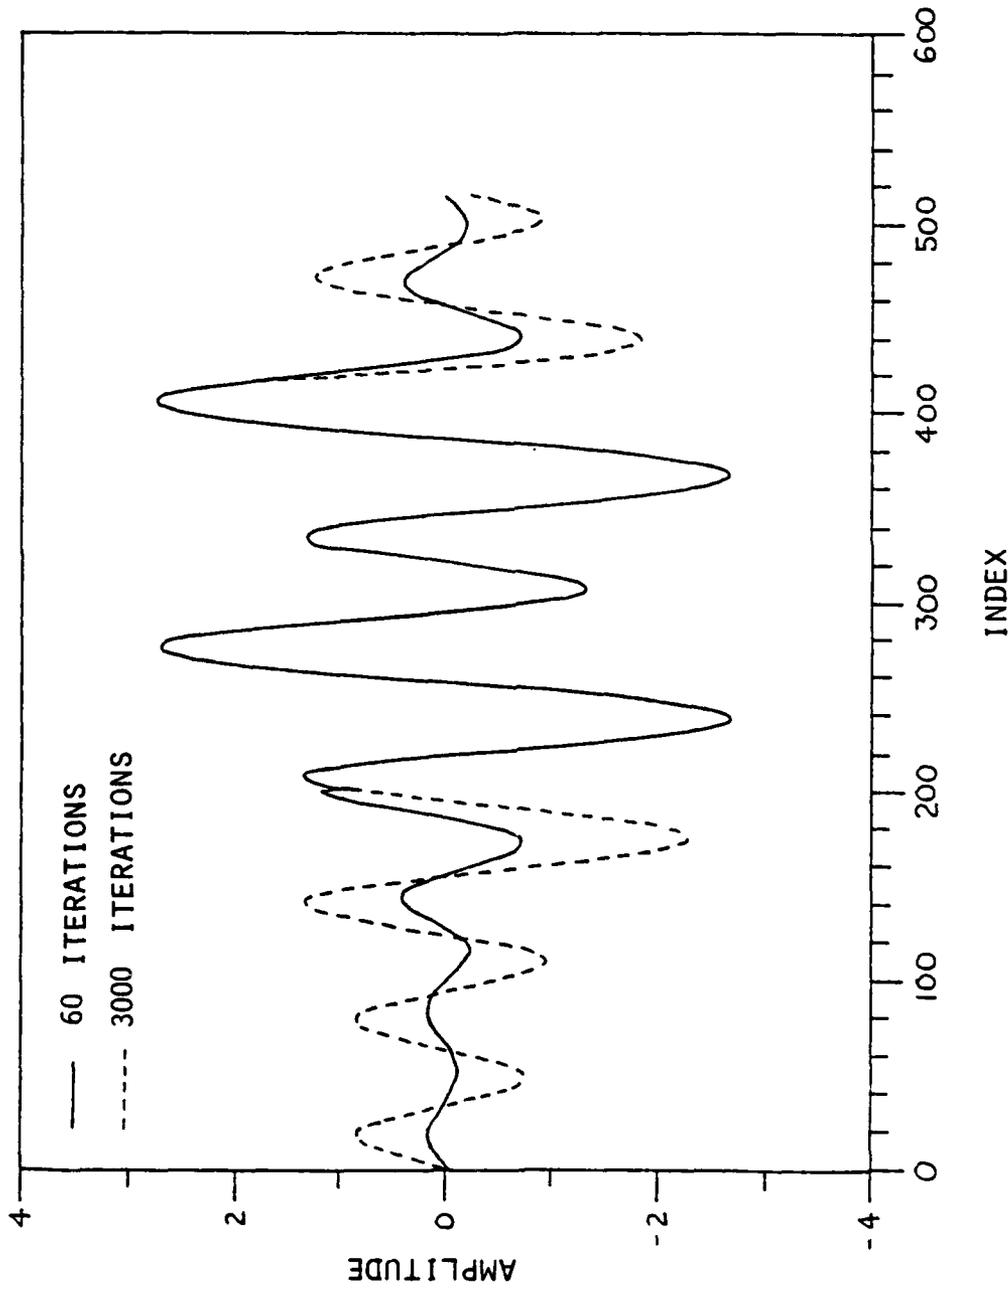


Figure 2.9 Results for periodic extrapolation after 60 and 3000 iterations.

2.4.2 General case

Experimental results for the general case will now be presented. Since evidence has been presented supporting the E_n curve and the related theory, if further evidence supports the expected TE_n^M curves, then it would be reasonable to assume correctness for the proposed E_n^M theory and curves. Previously demonstrated was the expected fact that the convolutional and FFT approach would yield identical results. Therefore, for manipulative ease, a convolutional approach will be used in the examples for the general case.

First, the effects of changing the observation interval will be examined. As discussed in section 3, longer observation intervals (L) will uniformly shift the TE_n^M curves downward. The result is that the optimal solution is obtained with the same number of iterations and the optimal total error is lower. In the cases illustrated in Figure 2.10, the observation window length is: 150, 125, 100, 75 and $\Omega = 0.115$ radians. The plots of Figure 2.10 verify the predictions of (2.60).

The critical test of (2.60) is verification of the expected behaviour of TE_n^M with respect to the record length M . Observation length experiments have supported the general shape of the curves, but further evidence is needed to verify the effects of M . It was predicted that as M increases, the error in the optimal solution would decrease but more iterations would be required to obtain this lower error. Values of M used are: 150, 175, 200, 225 for $L = 100$ and $\Omega = 0.115$ radians. The error curves illustrated in Figure 2.11 verify the theoretical predictions of equation (2.60). In Figure 2.12 is the

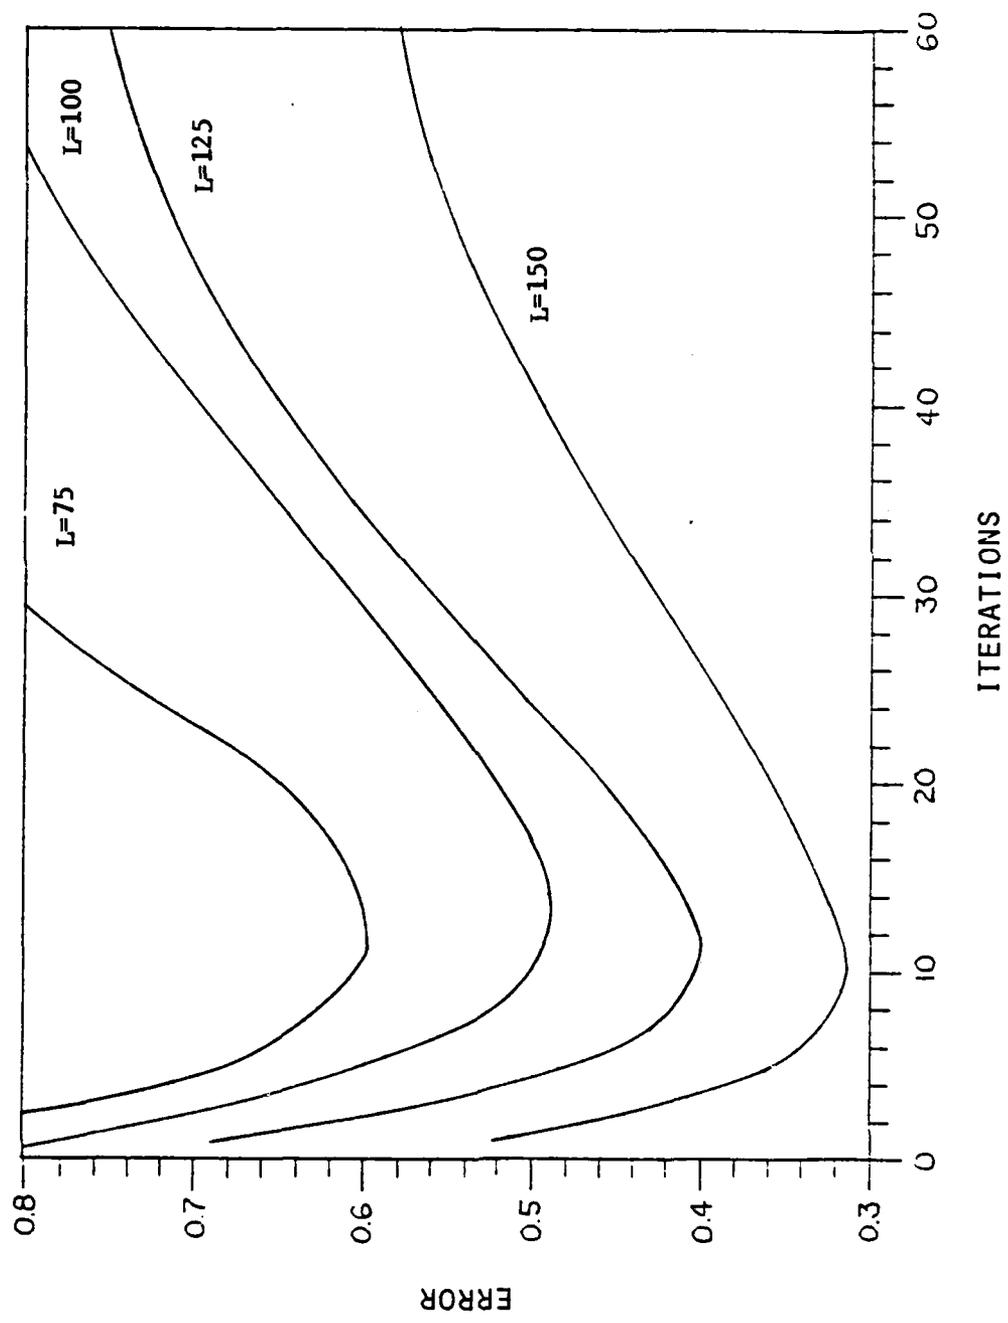


Figure 2.10 The affects of the observation length, L, on convergence.

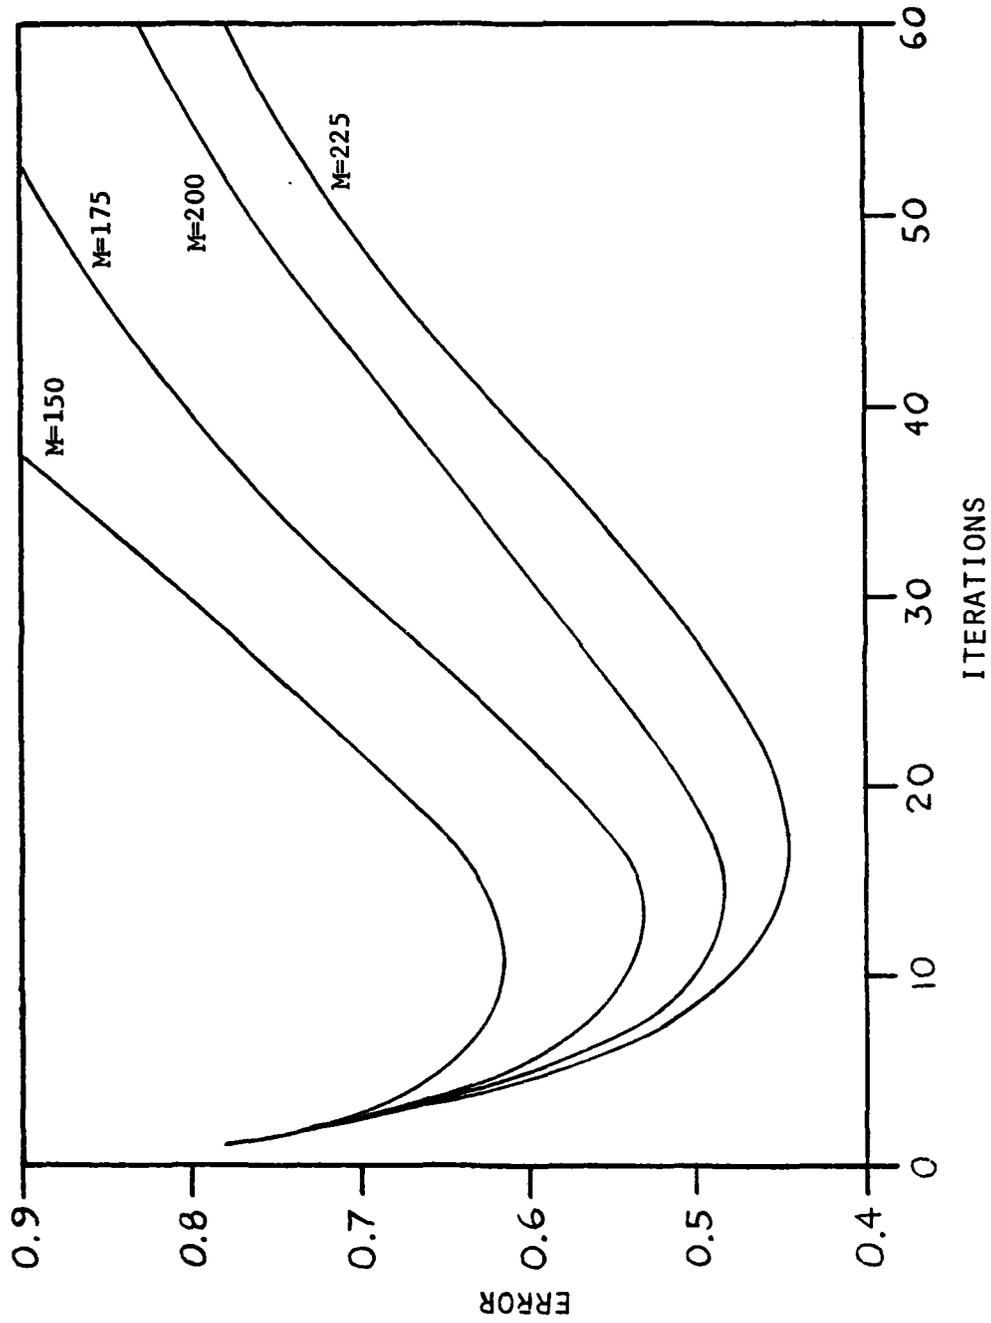


Figure 2.11 The affects of record length, M, on convergence.

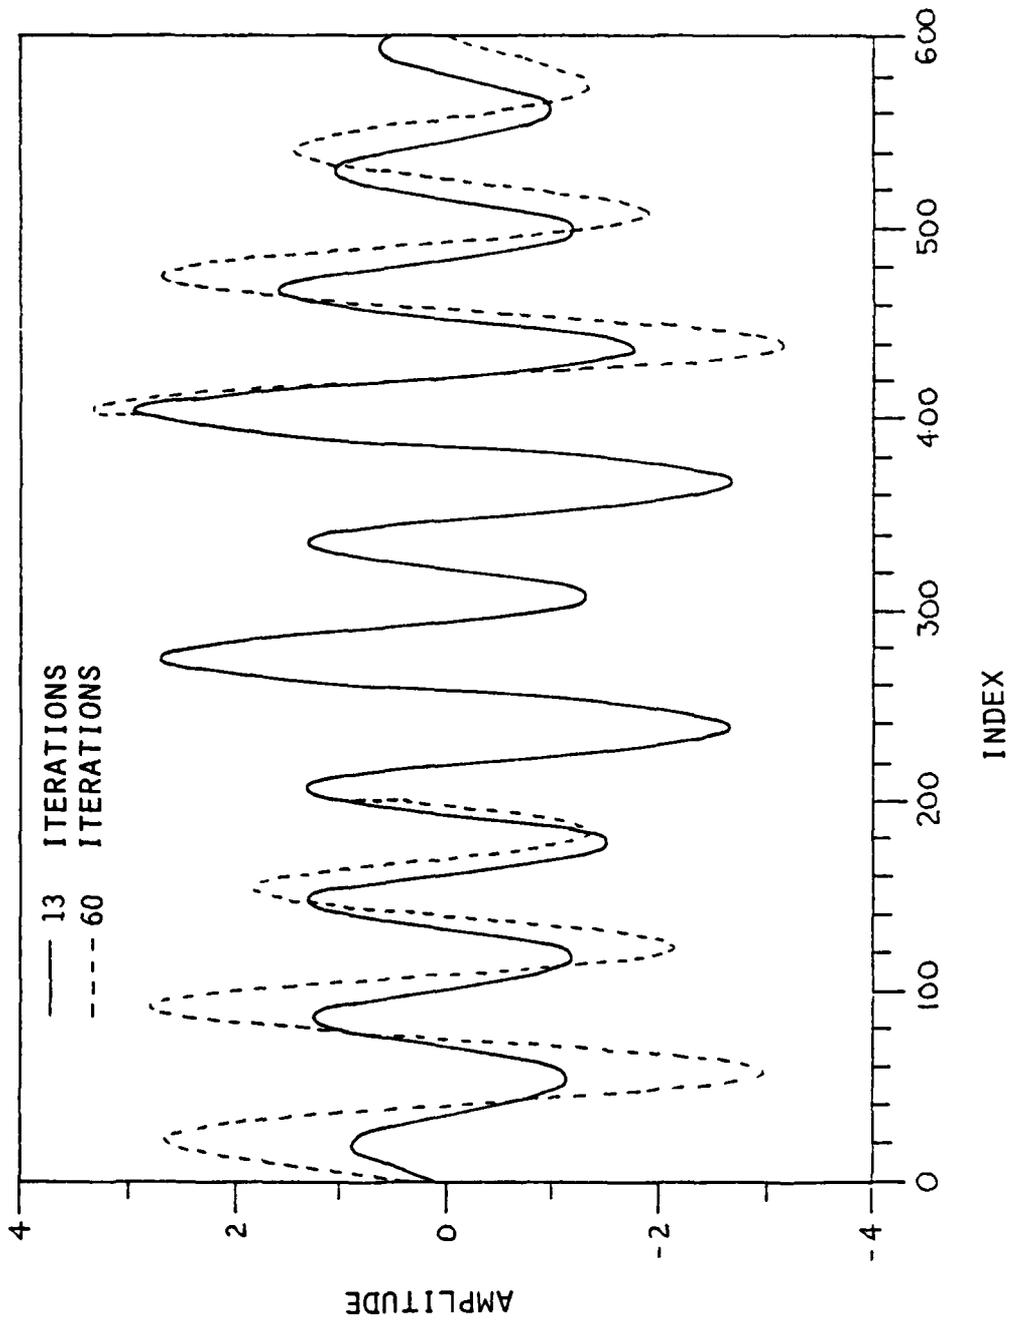


Figure 2.12 Extrapolated results after 13 (optimal) and 60 iterations.

extrapolated signal at the optimal error point, 13 iterations, and near the fixed point obtained after 60 iterations. While the signal obtained after 60 iterations appears 'bigger' than that after 13 iterations, the solution at 13 iterations is the best in terms of a norm. Small phase shifts that our subjective evaluation is insensitive to are identified and illustrated by calculating the norm. If either were to be employed as an estimate of an unknown signal, that obtained with 13 iterations would have to be selected. Additional experiments were performed to further verify the relationship between record length and error. The specific issue of concern is that the error curves of Figure 2.11 should level out (because of the existence of a fixed point) and further, that as M increases, the error associated with this fixed point should decrease. For this experiment the values of M employed are: 150, 175, 200, 225 and 250 for $L=150$ and $\omega=0.115$ radians. Error plots for this case are illustrated in Figure 2.13. A number of features need to be identified in this set of curves. First, the error curves between 0 and 60 iterations are simply shifted (down) versions of the curves in Figure 2.11. This agrees with the theoretical predictions that longer observation lengths will produce lower error. Second, as M increases, the optimal error decreases and the number of iterations required to obtain this optimal error increases. This result is identical to both theory and the results of the previous experiment. Third, the error curves are converging to a constant error which corresponds to the fixed point for that value of M . As M increases, the error associated with this fixed point decreases and the convexity of the error curves decreases. It was predicted that for a sufficiently large value of M , there is no minimum error point, i.e., the convexity of the error curve

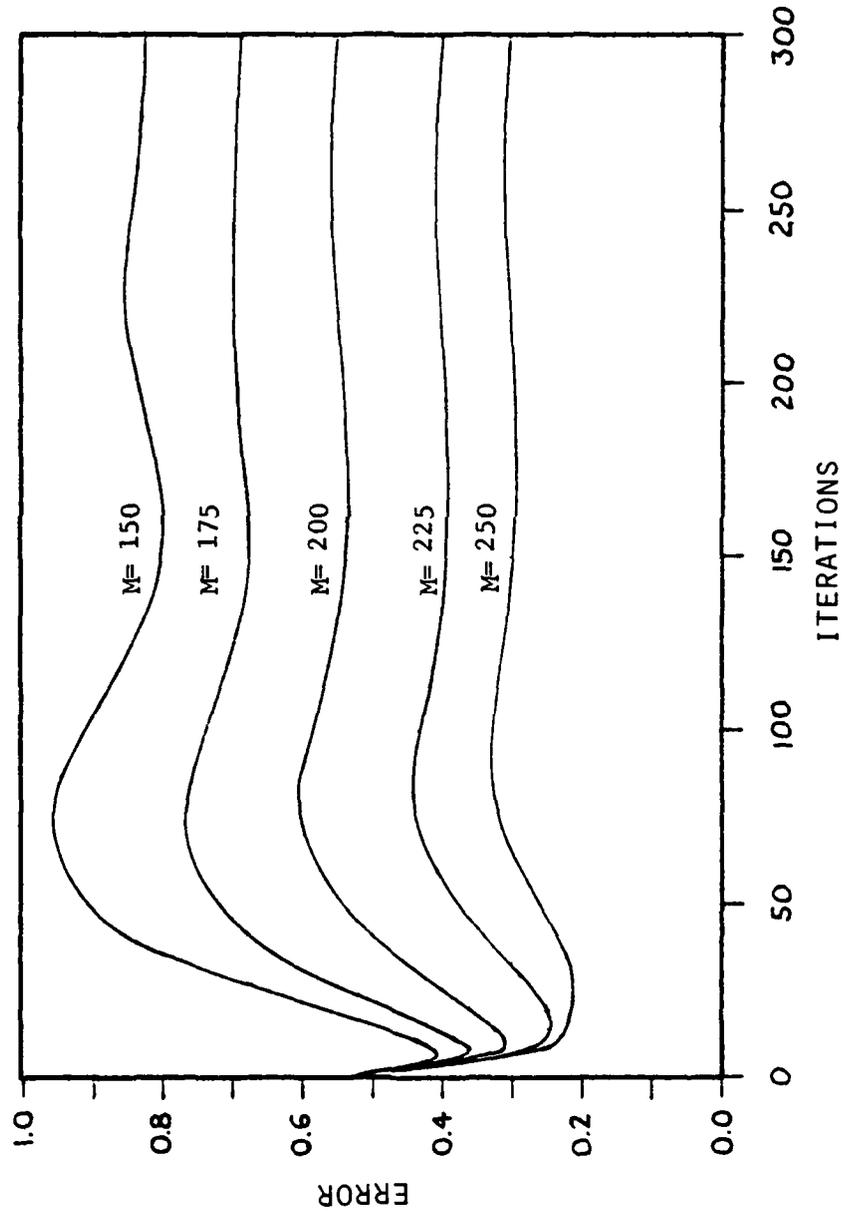


Figure 2.13 The affects of record length on convergence for a large number of iterations.

vanishes. This feature has not yet been observed experimentally. Some further comments concerning this will be made later. Lastly, while the oscillatory nature of the curves is not predicted in this analysis, it is not unexpected, and further, the distinct maximum error point for each M is anticipated as illustrated in Figure 2.5.

Curves in Figures 2.11 and 2.13 indicate that a distinct minimum error point does exist and that this error is obtained after a relatively small number of iterations. Small is used in the sense of the number of iterations required to approach the fixed point.

Next the effects of varying the extrapolation band-width, Ω , will be studied. After some thought it becomes apparent that Ω is not necessarily the least upper bound of frequencies contained in the observation. There are two reasons for this. First, it would not be desirable to choose an Ω so tight as to disturb sidelobe behavior around the higher frequencies. If the sidelobes are sufficiently affected, the mainlobes could be seriously distorted. This distortion could result in an inconsistent data set and consequently, a non-existent solution. The effect on the iterative algorithms would be to cause a constantly increasing error (increasing without bound).

An explanation for the second reason is more involved. The model employed here was presented earlier by Schaefer et al. in [21] and lately discussed by Sanz and Huang in [22]. Denote the observation of $f(n)$ with the vector $x(n)$ for $n=0, \dots, L-1$. Because any extrapolation of $x(n)$ approximating $f(n)$ must be of finite extent and can therefore be treated as periodic (implicitly at least) there exists a discrete Fourier series such that

$$x(n) = \sum_{k=-K_0}^{K_0} X(k) e^{-j2\pi nk/N}, \quad K_0 = \frac{\Omega N}{2\pi} \quad (2.64)$$

where $X(k)$ represent the Fourier series coefficients. The extrapolation problem can now be rephrased as: Given an observation set $x(n)$, calculate the Fourier series coefficients $X(k)$ that generate $x(n)$. Thus by knowing the $X(k)$, $f(n)$ can be recovered. Rewriting (2.64) in vector notation,

$$x = WX \quad (2.65)$$

where x is $L \times 1$, W is $L \times 2K_0+1$ and X is $2K_0+1 \times 1$. The desired solution of (2.65) is X given x . Clearly, for an observation length of L , if $2K_0+1$ is strictly less than L , then (2.65) is an underdetermined system of equations and a solution may not exist. The implication of this for Gerchberg-Papoulis type algorithms is that if for a given observation set, Ω is chosen too low, then the iterations will immediately diverge. An obvious problem is that the minimum extrapolation bandwidth required by (2.64) may be too high to be of any practical value. A method of circumventing this is to apply various decimation/interpolation [23] schemes to the observation set to alter L as well as to vary N in order to achieve a more useful band-limit.

Independent of the factor determining the lower bound for Ω , it is desirable to get as close (from above) to this value as possible. If Ω is significantly higher than necessary, either too much spectral energy is allowed to leak out of the mainlobes (of interest), or the system of equations is too unconstrained allowing too many solutions. In either case, the result will be a solution inferior to that obtainable with a

tighter band-limit.

The point of the above arguments is to suggest that an optimal extrapolation bandwidth exists. This value of Ω will be a function of many factors, L , M , N and the properties of the observation for instance. The net effect of Ω on the error curves should be as follows. If Ω is below some critical value, an inconsistent data set is generated and a solution may not exist implying that the iterations will diverge. As Ω is raised above some critical value, the resultant optimal solution will become gradually worse as the number of possible solutions increases. The experimental results in Figure 2.14 are for $M=200$, $L=100$ and varying Ω between 0.1075 and 0.1200. These curves support the existence of an optimal extrapolation bandwidth.

With reference to a previous remark concerning the inability to achieve a constantly decreasing error plot for sufficiently large M , the above analysis supplies a possible reason for this problem. Since L was held constant while M was increased (for the experiments in Figures 2.11 and 2.13), a larger class of solutions was admitted, thus degrading the quality of any one solution. In order to correct this problem, the sampling density in the known observation interval would have to be increased, thus reducing the degrees of freedom for the system of equations and improving the solution. Since this was not done, the errors caused by the lack of constraint swamped the benefits of longer record length.

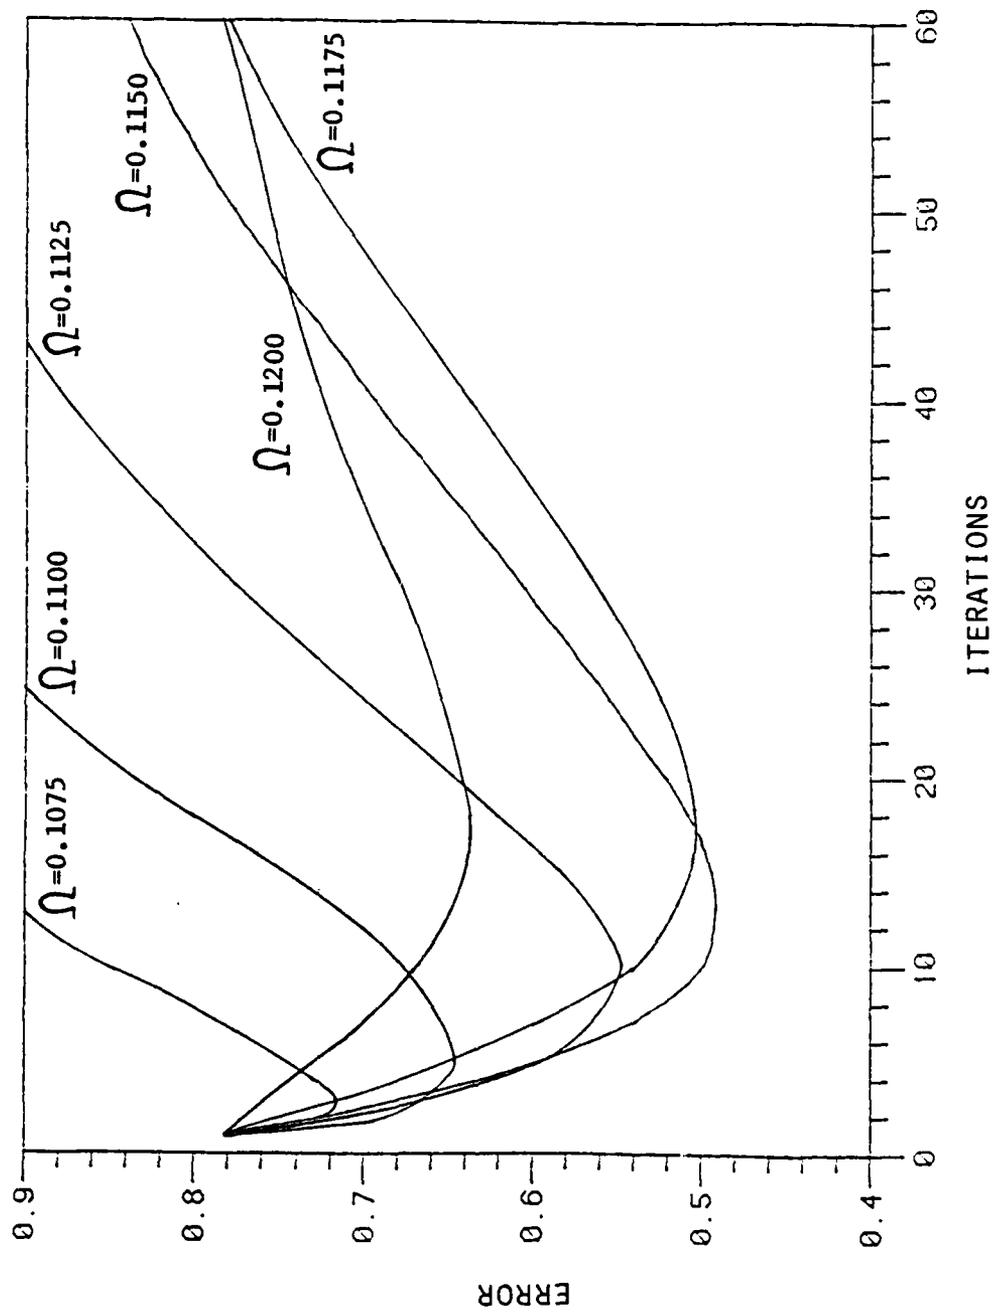


Figure 2.14 The affects of the extrapolation bandwidth, Ω , on convergence.

2.5 Conclusion

For the sake of brevity, only one example has been illustrated in this chapter. To date, over 20 different signals have been employed and tested in varying detail. Results of these experiments have in every case supported the theory presented in this chapter.

The purpose of the analysis in this chapter is to study the effects of finite length processing intervals on iterative signal extrapolation techniques. The key results are: 1) The discrete implementations of Gerchberg's and Papoulis' algorithms can be rephrased as fixed point problems in which the effect of employing finite length processing intervals in these algorithms is to make the mapping a contraction mapping. A property of a contraction mapping is that it has a unique fixed point, implying that these algorithms converge to a specific solution. 2) A second consequence of finite length records is that in general, these algorithms attain the best solution after a relatively small number of iterations. As the examples have shown, the number of iterations required to obtain this solution is approximately an order of magnitude less than the number required to approach the fixed point. In the next chapter, these results will be employed in analyzing the performance of two-dimensional image reconstruction/enhancement schemes.

3. THE MISSING CONE PROBLEM IN COMPUTER TOMOGRAPHY

Computer-aided tomography (CAT) is a well-known technique for obtaining high resolution cross-sectional images of an object from many different angular views. X-ray tomography is probably the best known type of CAT system due to its remarkable success in the field of medical diagnostics. A typical x-ray CAT scanner may record as many as 600-800 projections taken at equally spaced angular increments over a total viewing angle of 360 degrees. Data from the x-ray sensors are digitized, stored in a digital medium and later processed into a final image by one of several popular reconstruction techniques. Since exposure to x-rays should be minimized, there is considerable interest in generating high-quality images with a minimal amount of projection data.

Reconstruction algorithms fall into three general categories: 1) algebraic reconstruction techniques (ART), 2) convolutional-backprojection (CBP) techniques and 3) direct Fourier (DF) domain techniques. Although all of these have been used in commercial CAT scanners with varying degrees of success, it appears that CBP is currently the most popular in the present generation of machines. A number of these algorithms are discussed in more detail in Section 1.

Many other remote sensing systems share the common problem of attempting to reconstruct a high resolution object function from a limited set of data recorded in the frequency domain, the spatial domain or projection space. Examples are found in synthetic aperture radar, beamforming sonar, electron microscopy and radio astronomy. In these

systems it is often impossible to collect projections over an entire 360 degree viewing angle. For example, in a synthetic aperture radar, which images an area on the earth's surface with a microwave radar carried in a satellite, the total viewing angle may be quite small, perhaps on the order of 15-30 degrees. In electron microscopy, the viewing angle is limited by the extent that a specimen can be tilted with respect to the electron beam. In some cases where x-ray CAT is used for non-destructive testing of manufactured items, physical limitations prevent the collection of projections over a complete 360 degrees of angle. Therefore, the problem frequently arises as to the best way to reconstruct the image when an angular interval of projection data is missing. This constitutes the missing cone problem which is addressed in Section 2. Section 2 also discusses how maximum entropy methods and algebraic reconstruction techniques have been used in the past to deal with the missing cone problem.

Section 3 describes two recently proposed iterative reconstruction algorithms that estimate the data in the unknown missing cone region and then employ this data to improve the resolution of the reconstruction. These algorithms are called the projection-slice algorithm (PSA) and the angular iteration method (AIM). Essentially they are iterative band-limited (space-limited) extrapolation algorithms which have been modified to take into account problem dependent relationships that result from the missing cone geometry.

Section 4 presents a number of computer generated examples to illustrate the salient features of the PSA and AIM algorithms. Throughout this chapter, the missing cone problem is discussed within

the context of x-ray tomography. However, since the missing cone problem naturally arises in other imaging systems as mentioned above, it is hoped that new solutions to this problem will have important applications in a variety of different disciplines.

3.1 Introduction to Computer Tomography

The purpose of this section is to review basic concepts, establish notation and create the framework in which the missing cone problem is discussed. An excellent general reference for tomography is provided by G.T. Herman in [2].

3.1.1 Projection data

Projections, alternately referred to by some authors as shadowgraphs, are the format of collected data in tomographic systems. As illustrated in Figure 3.1, a projection is created by illuminating an object from a source of penetrating radiation, typically x-rays. The magnitude of the transmitted radiation is recorded on film or with sensors and from these data a projection can be calculated. The recorded signal intensity at a point on x' of a projection is related to the incident radiation intensity $I_0(x')$ and the two-dimensional (2D) attenuation $f(x,y)$ of the object by

$$I(x') = I_0(x') e^{-\int_L f(x,y) dy'} \quad (3.1)$$

where the y' direction is perpendicular to the projection and the

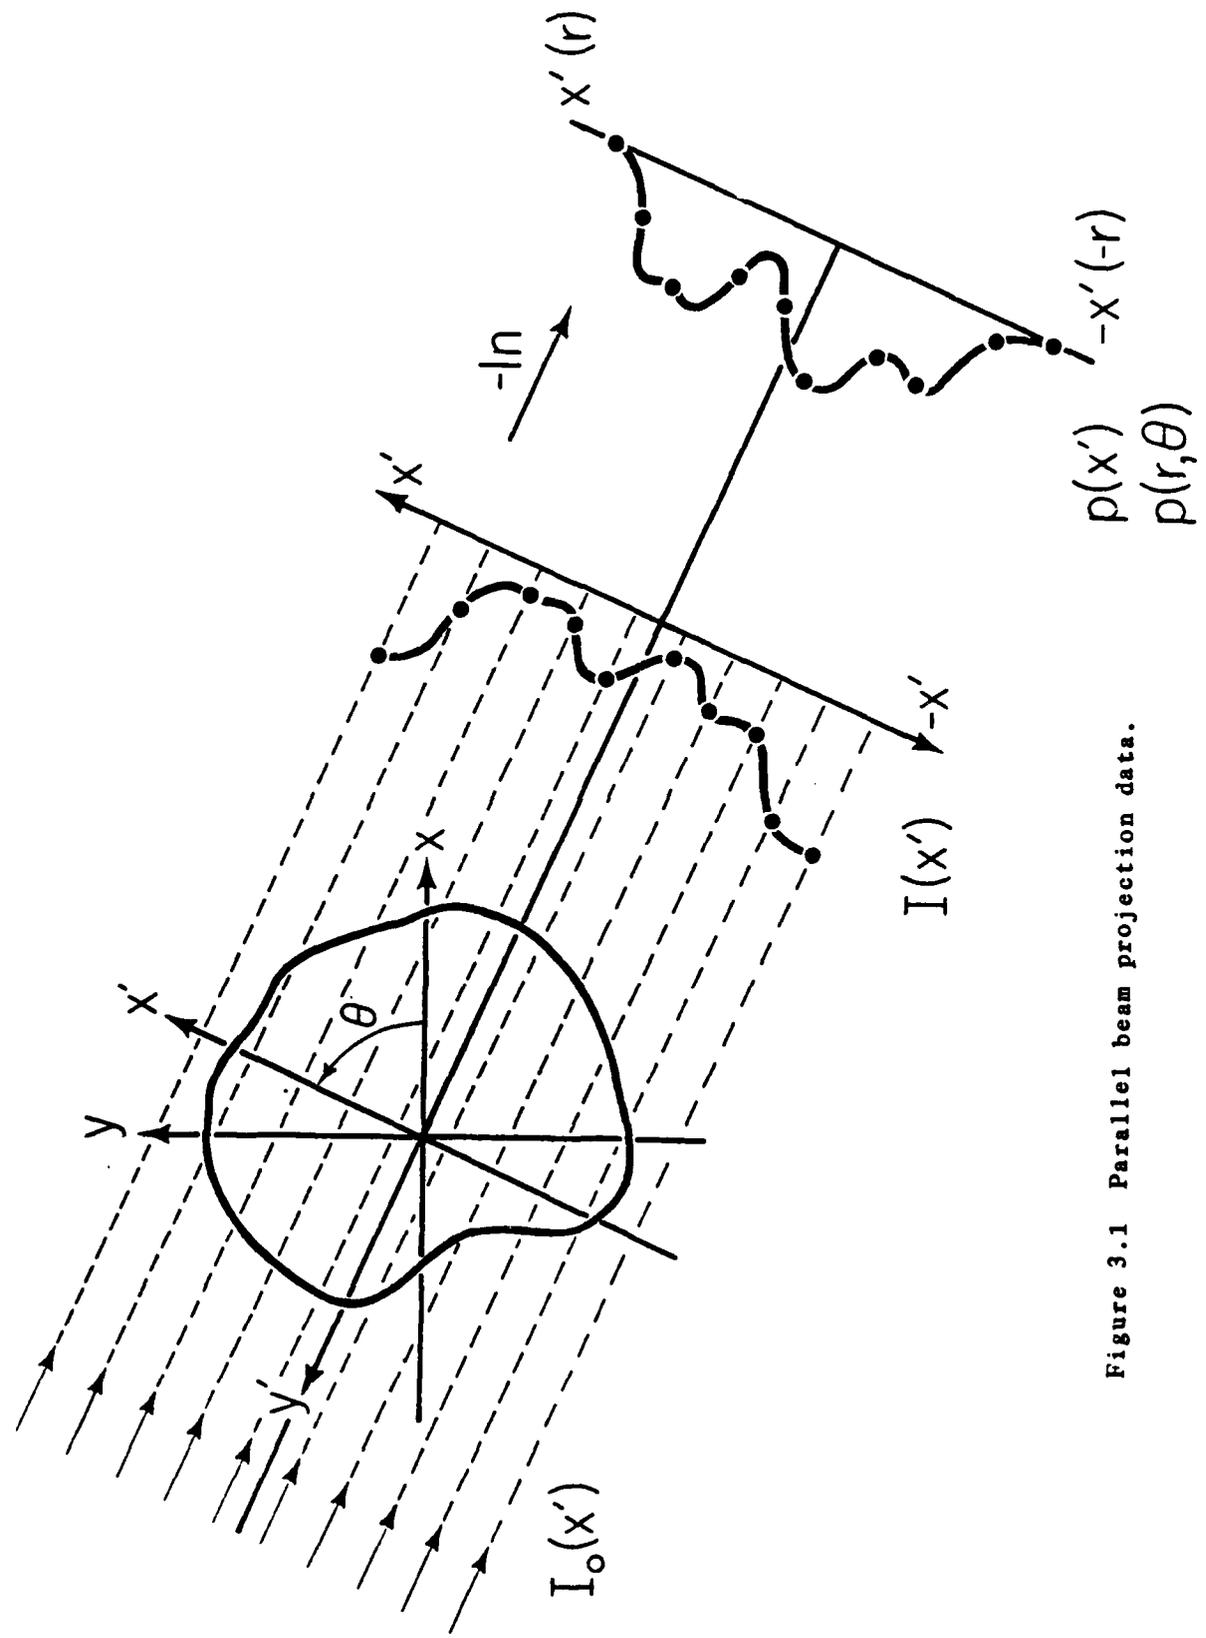


Figure 3.1 Parallel beam projection data.

integration is performed over a line parallel to y' . The projection, $p(x')$, is then defined as the line integral over $f(x,y)$,

$$p(x') = \int_L f(x,y) dy' \quad (3.2)$$

This value can be obtained from $I(x')$ through the relation

$$p(x') = \int_L f(x,y) dy' = -\ln \left[\frac{I(x')}{I_0(x')} \right] \quad (3.3)$$

In order to obtain the perspective necessary to uniquely reconstruct $f(x,y)$, projections are taken over a continuum of angle, typically π radians. Generalizing, a projection is a bivariate function of r and θ where

$$p(r,\theta) = p_\theta(r) = \int_{-\infty}^{\infty} f(r\cos\theta - v\sin\theta, r\sin\theta + v\cos\theta) dv \quad (3.4)$$

and the rotated projection coordinates x',y' are replaced by a more natural polar coordinate system. This projection operation is also referred to as the Radon transform [24].

$$p(r,\theta) = p_\theta(r) = R[f(x,y)] \quad (3.5)$$

The two slightly different notations of $p(r,\theta)$ and $p_\theta(r)$ will be used to emphasize the difference between operations on the two-dimensional (2D) data set $p(r,\theta)$ and calculations involving a specific projection $p_\theta(r)$. Throughout this chapter, use of parallel beam projection data as illustrated in Figure 3.1 will be assumed. It is shown later that this assumption does not restrict the applicability of techniques to be discussed here.

From a mathematical standpoint, the projections over the first π radians are identical to those over the second π radians, i.e.,

$$p(r, \theta) = p(r, 2\pi n + \theta) = p((-1)^n r, n\pi + \theta) \quad (3.6)$$

for n an integer. Since Radon's transformation (3.4) is uniquely invertible [25], projections over any continuous interval of π radians unambiguously describe the image. In practice however, the effects of beam hardening, beam spreading and other non-symmetrical anomalies disrupt the relationship of (3.6). Consequently, high-performance tomographic scanners obtain projections over 2π radians in order to reduce non-symmetrical effects. Since these non-symmetrical anomalies are not considered here, projections over π radians will suffice. Projections available over π radians will be referred to as a complete set of projections (a complete data set).

In any practical implementation, projections are sampled in both the angular and radial components: θ and r of (3.4). These sampled projections will permit an unambiguous reconstruction only to a finite resolution which is determined by both the sampling density and sampling geometry. The sampling geometry assumed for this work employs uniform angular and uniform radial sampling intervals. The sampling of $p(r, \theta)$ in the radial direction is indicated in Figure 3.1.

3.1.2 Reconstruction techniques

In order to derive reconstruction techniques for the case where the projections are sampled, the mathematics of the continuous case will be examined first. Given a complete set of projections, many reconstruction techniques are available. Some of these methods are: Algebraic Reconstruction Techniques (ART) [2], Rho-filtered layergram

[2,24], Convolutional-Backprojection (CBP) [2,24] and Direct Fourier (DF) [2,26,27]. ART, CBP and DF methods will be discussed in some detail while the Rho-filtered layergram method will only be identified in passing. A key concept in tomography is the Projection Slice Theorem (PST) which forms the basis for both the CBP and DF reconstruction techniques.

The PST provides a relationship between projections of $f(x,y)$ and center cross-sections of the two-dimensional Fourier transform (2D-FT) of $f(x,y)$. Denote the 2D-FT of $f(x,y)$ by $F(u,v)$. The PST theorem states that the one-dimensional Fourier transform (1D-FT), over r , of a given projection $p_\gamma(r)$ denoted by $P_\gamma(R)$ is identical to a function which is the center cross-section (slice) of $F(u,v)$ at the same angle (γ), i.e.,

$$P_\gamma(R) = F(u,v) \quad (3.7)$$

where $F(u,v)$ is evaluated along the line $u=R\cos\gamma$, $v=R\sin\gamma$. To prove this, consider the projection at a given angle γ which can be written as

$$p_\gamma(x') = \int_{-\infty}^{\infty} f(x',y') dy' \quad (3.8)$$

in the x',y' coordinate system, where x',y' is equivalent to the original x,y coordinate system rotated by angle γ . The same relationship exists between u,v and u',v' . The 1D-FT of $p_\gamma(x')$ is

$$P_\gamma(u') = \int_{-\infty}^{\infty} p_\gamma(x') e^{-ju'x'} dx' \quad (3.9)$$

and the 2D-FT of $f(x',y')$ is

$$F(u',v') = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(x',y') e^{-j(u'x'+v'y')} dx' dy'. \quad (3.10)$$

A slice of $F(u',v')$ at angle γ is defined as $F(u',v')$ with $v'=0$. Substituting (3.8) into (3.9) and equating with (3.10) results in

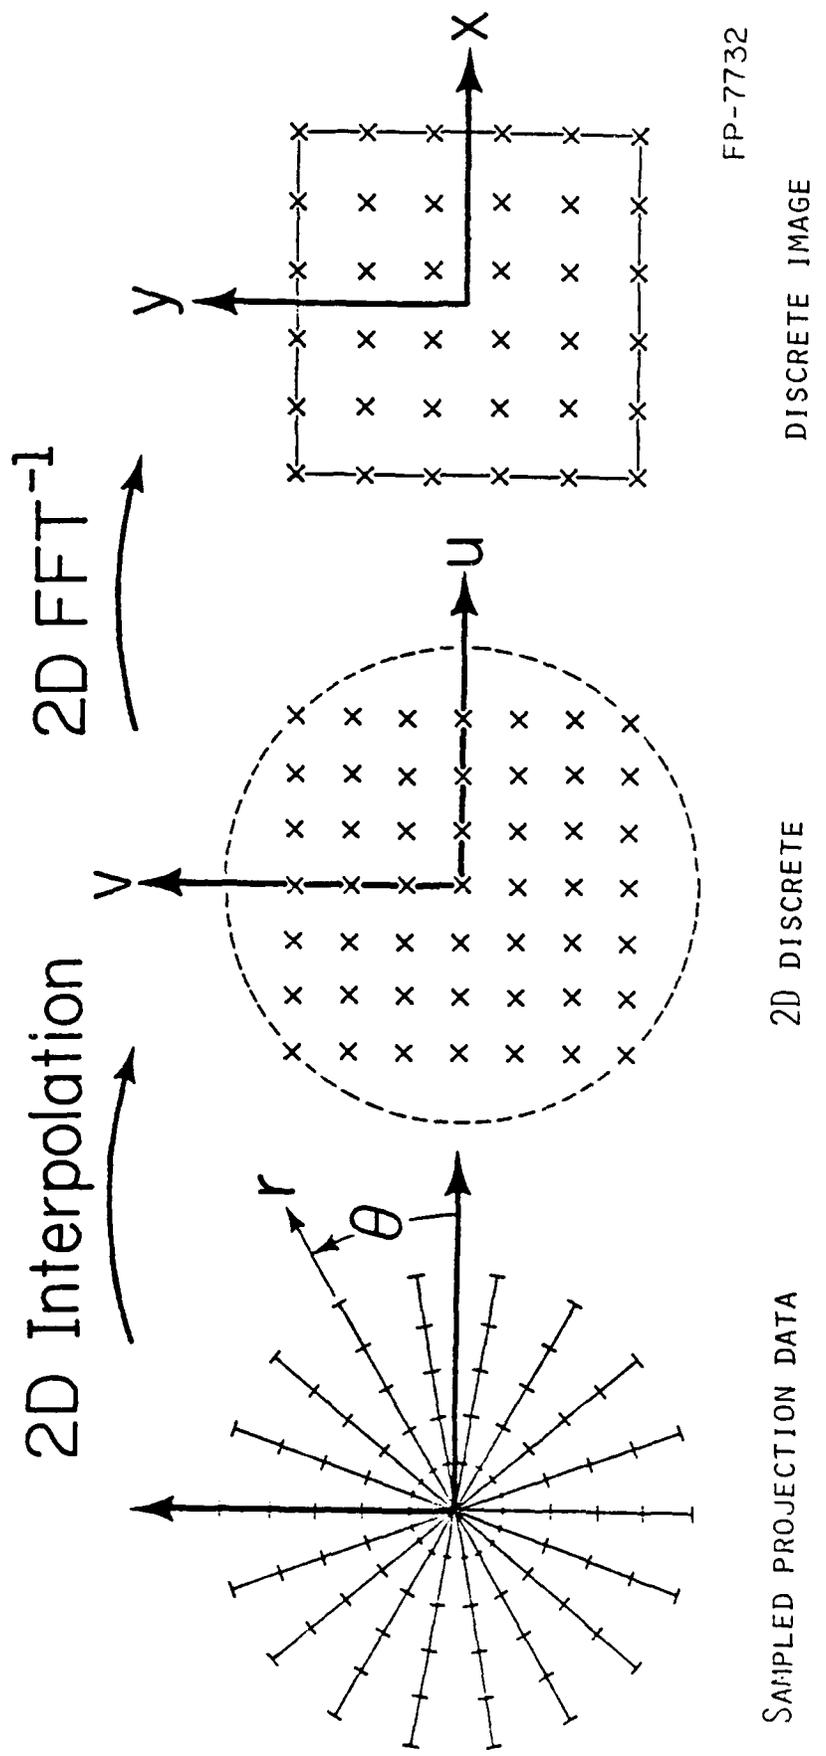
$$P_{\gamma}(u') = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(x',y') e^{-jx'u'} dx' dy' = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(x',y') e^{-j(u'x'+v'y')} dx' dy' \quad (3.11)$$

when $v'=0$. Since γ was arbitrary, the above is true for any angle of rotation (projection). This proves the PST.

One possible reconstruction technique is based on a direct implementation of the PST. Starting with projection data, a 1D Fourier transform is applied to $p(r,\theta)$ to generate $P(R,\theta)$. These transformed projections, $P(R,\theta)$, completely describe $F(u,v)$. An inverse 2D Fourier transform is next applied to $F(u,v)$ resulting in the image $f(x,y)$. This is the direct Fourier (DF) reconstruction method. A major difficulty with this technique is that in any practical case the data are sampled as described earlier. Since projections are sampled in angle, information concerning $F(u,v)$ will only be available along a discrete set of radial lines passing through the origin. The angular interval corresponds to the angular sampling of projections. Compounding this problem is the radial sampling of each projection. These data, can at best, only be used to find approximate values of $F(u,v)$ along the radial lines defined by the angular sampling rate. The reason is, with discrete data, a discrete Fourier transform (DFT) must be used which can only provide approximate samples to the FT of the continuous signal from which the original samples were taken. Consequently, these samples are

only approximate values of $F(u,v)$ on a polar raster, i.e., erroneous samples are available for discrete values of R and θ . This is illustrated in Figure 3.2. In order to efficiently generate the final image, a fast discrete Fourier transform (an FFT) is implemented. To employ an FFT, samples of $F(u,v)$ must be available on a rectilinear raster, implying a polar-to-rectangular interpolation is required. The steps involved in the DF reconstruction method are: 1) Calculate the 1D-FFT of each projection. 2) Using this transformed data set that describes $F(u,v)$ on a discrete polar raster, a two-dimensional interpolation is performed to obtain samples of $F(u,v)$ on a discrete rectilinear raster. 3) Finally, a 2D inverse FFT is employed to calculate samples of $f(x,y)$ on a rectilinear grid. Variations on this method include filtering the data in the radial or angular directions [27], or using Hankel transforms to calculate the inverse polar Fourier transform from a discrete polar grid [28,29,4]. This last technique generates samples of $f(x,y)$ on a polar grid directly from the transformed projections.

Convolutional-backprojection is the next reconstruction technique to be discussed. Backprojection can be viewed as the reverse of the projection operation. Instead of integrating over the image to generate the projections, each projection is "smeared" across the region of support for the image. The idea of smearing is to evenly distribute over the entire image, information contained in each projection. Since range information has been integrated out by the projection process, smearing is the most unbiased action to take. Let S denote the region of support for the image in the x,y plane. Consider the image of Figure 3.3 and



FP-7732

SAMPLED PROJECTION DATA

2D DISCRETE

DISCRETE IMAGE

Figure 3.2 Direct Fourier reconstruction method.

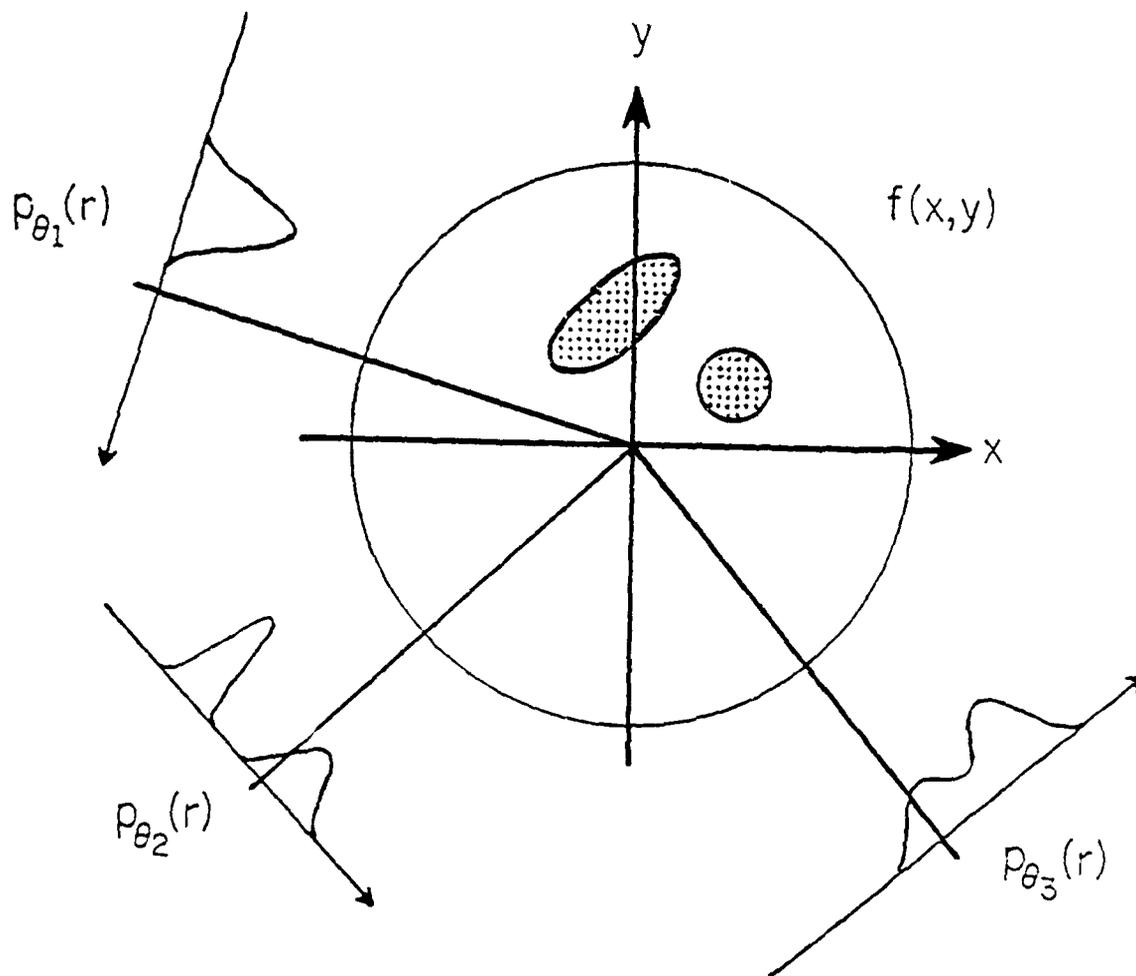


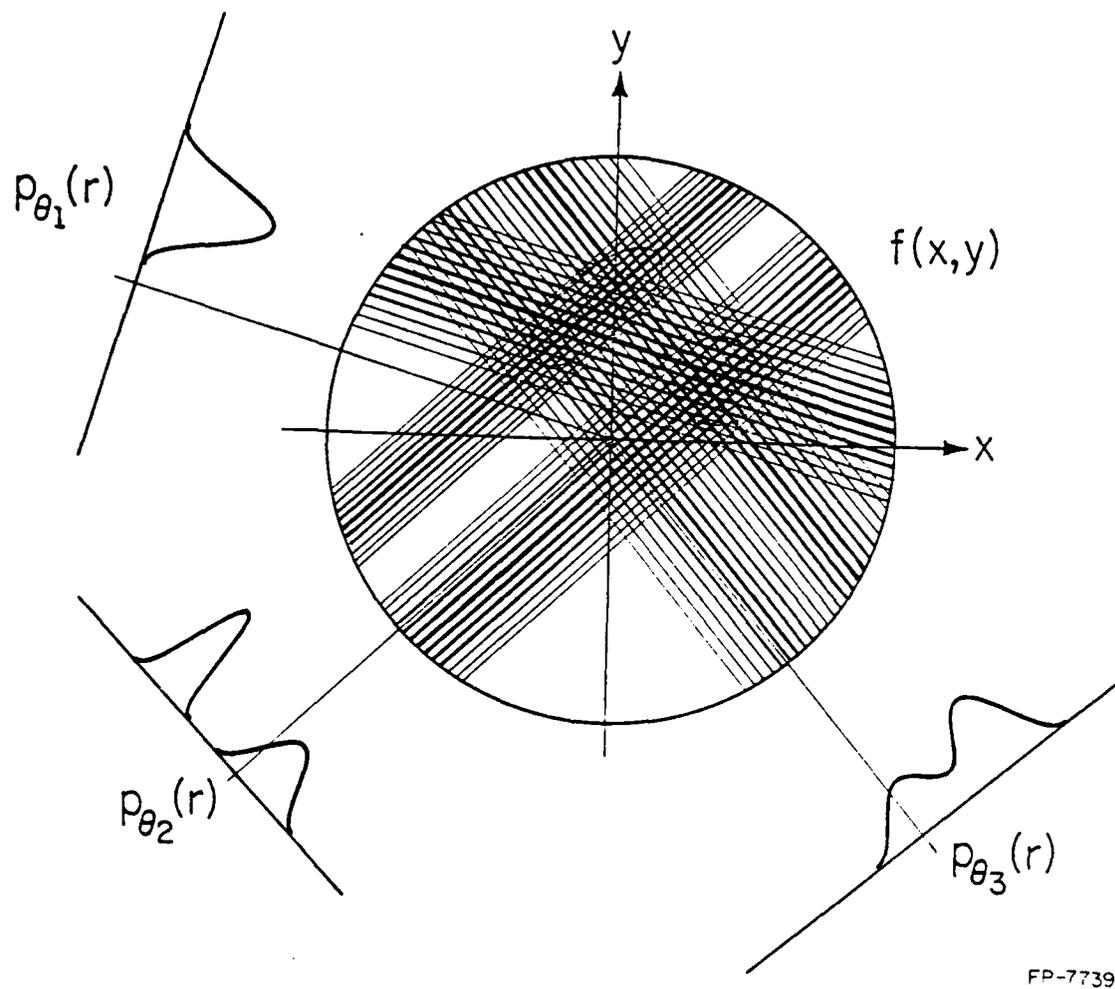
Figure 3.3 Examples of projection data.

three arbitrary projections of that image $p_{\theta_1}(r)$, $p_{\theta_2}(r)$ and $p_{\theta_3}(r)$. Smearing is the process in which the projection value for each r for a given θ is uniformly divided along the line of integration in S defined by the point (r, θ) . Since each projection contains relative spatial information only in the direction perpendicular to θ (azimuth), in order to be maximally fair, the data must be uniformly divided (smeared) along the line of integration over S . This is illustrated in Figure 3.4. Regions in S where this smearing overlaps to the greatest extent are objects in the image. Little information is present in any given projection concerning the position, magnitude and quantity of specific targets in the range direction. These data are degraded in a given projection because the projection process has integrated this information together. However, since projections are obtained over a range of angles, in most cases π radians, range information not present in one projection is available in others as cross-range (azimuth) data. It is this property of having perspective that allows an unambiguous reconstruction of the image from projections. This is also illustrated in Figure 3.4. In a complete data set, each projection and its orthogonal complement is present.

With some thought it can be seen that "smearing" is the operation called backprojection as defined by

$$f(x,y) = \int_0^{\pi} p(x\cos\theta + y\sin\theta, \theta) d\theta. \quad (3.12)$$

In (3.12), instead of dividing each projection value among all the points along the original line of integration, i.e., among each point in x,y for which it may possibly contain some information, each point in



FP-7739

Figure 3.4 Backprojection.

x, y is given all of every projection value it may be associated with. The operation described in (3.12) is an integration along circular contours in the projection domain. The contour used is determined by the specific values of x and y , thus selecting which projection values are possibly associated with that image point. A result of backprojecting is a substantial low-frequency offset in the reconstruction. Consider the effects of significantly more projections in the example of Figure 3.4. In order to compensate for this offset, some sort of filtering procedure must be included. In Rho-filtered layergram reconstruction, a 2D filter is applied after backprojection to correct for this offset. In the CBP method each projection is filtered prior to backprojection, consequently the name, convolutional-backprojection.

To derive this filter function, let $p(r, \theta)$ be a complete set of continuous projections of $f(x, y)$ and let $P(R, \theta)$ represent the FT of $p(r, \theta)$ in r . Then

$$P(R, \theta) = \int_{-\infty}^{\infty} p(r, \theta) e^{j2\pi Rr} dr \quad 0 \leq \theta < \pi. \quad (3.13)$$

Via the PST, the image in polar coordinates, $f(r, \phi)$, can be expressed as

$$f(r, \phi) = \int_0^{2\pi} \int_0^{\infty} P(R, \theta) e^{-j2\pi Rr \cos(\phi - \theta)} R dR d\theta. \quad (3.14)$$

Considering the geometry of the projection scheme illustrated in Figure 3.1, i.e., $0 \leq \theta < \pi$ implying that $-\infty < R < \infty$, (3.14) can be rewritten as

$$f(r, \phi) = \int_0^{2\pi} \int_{-\infty}^{\infty} P(R, \theta) e^{-j2\pi Rr \cos(\phi - \theta)} |R| dR d\theta. \quad (3.15)$$

The integral with respect to R in (3.15) can be considered a filtering

operation where each $P_{\theta}(R)$ is multiplied by $|R|$ and then inverse transformed. Letting $p'(*,*)$ denote the filtered projections

$$p'(r\cos(\phi - \theta), \theta) = \int_{-\infty}^{\infty} P(R, \theta) |R| e^{-j2\pi R r \cos(\phi - \theta)} dR. \quad (3.16)$$

Substituting $p'(*,*)$ into (3.15) results in

$$f(r, \phi) = \int_0^{\pi} p'(r\cos(\phi - \theta), \theta) d\theta, \quad (3.17)$$

which is equivalent to

$$f(r, \phi) = \int_0^{\pi} p'(r\cos\phi\cos\theta + r\sin\phi\sin\theta, \theta) d\theta. \quad (3.18)$$

Applying some trigonometry, (3.18) can be rewritten for $f(x, y)$ as

$$f(x, y) = \int_0^{\pi} p'(x\cos\theta + y\sin\theta, \theta) d\theta \quad (3.19)$$

which is the backprojection operator defined in (3.12). The filtering operation defined in (3.16) can also be implemented in the spatial domain via convolution. In this case, the original projections $p(r, \theta)$ are first filtered by a function with an impulse response given by

$$k(r) = \int_{-\infty}^{\infty} |R| e^{-j2\pi R r} dR. \quad (3.20)$$

After filtering the projections either in the frequency domain by multiplying by $|R|$ or in the spatial domain by convolving with $k(r)$, $p'(r, \theta)$ is backprojected with (3.12) to generate the final image $f(x, y)$.

As with the DF method, discretization of (3.19) required by the sampled data poses a variety of problems. Consider first the filtering of samples of $p(r, \theta)$ to obtain a close approximation to $p'(r, \theta)$ at the sample points. With non-periodic data it is generally not possible to

discretely filter data without introducing aliasing errors. The object then is to design a filtering scheme that minimizes these aliasing errors. This can be accomplished in various ways, some of which include: increasing the number of filter coefficients used in convolution calculations, using longer FFT's or modifying the kernel $k(r)$ as to maximize or minimize certain reconstruction features. In reference to the last statement, many authors have proposed assorted approximations to $k(r)$ for specific applications [24,30,26].

A second issue is the discretization of the backprojection integral in (3.12). With discrete filtered projection data, (3.12) will have to be implemented as a finite sum. A trapezoidal approximation has been shown to be in some sense optimal for performing backprojection [24]. With a trapezoidal approximation, (3.12) becomes

$$f(x,y) = \delta \sum_{n=0}^{M-1} p'(x \cos \delta n + y \sin \delta n, \delta n) \quad (3.21)$$

where δ is the angular sampling increment and there are M projections over π radians, i.e., $\delta = \pi/(M-1)$. Inspection of (3.21) shows that interpolation will be required between available samples of $p'(r, \theta)$ to obtain values of the projection at the point $x \cos \delta n + y \sin \delta n (=r)$. Since the integration in (3.19) is over θ , the summation in (3.21) is over δn , which implies that interpolation is only required in r . Contrast this with the 2D interpolation required by the DF method.

A significantly different reconstruction approach to those previously described above is the Algebraic Reconstruction Techniques (ART). The major way in which these techniques differ from convolutional-backprojection or direct Fourier methods is that ART

assumes by design a discrete data set and a discrete reconstruction grid. The object is to reconstruct sample values of the image to an accuracy dictated only by sampling constraints (density and geometry). Compare this to CBP or DF where approximations are made to the continuous model in order to employ sampled data. These approximations involve error prone processes such as filtering and interpolation which at best can only lead to corrupted sample values of the image. While this type of corruption may be small in the CBP method, it is even smaller or non-existent in ART even if the same initial data set is used. For a good general reference, again see [2]. It is sufficient for the purposes of this chapter to only briefly introduce ART.

Two basic algebraic reconstruction techniques are direct multiplicative ART and direct additive ART, both of which are iterative. Denote the sampled projections by $p(\varepsilon k, \delta m)$, $k=0, \dots, K-1$ and $m=0, \dots, M-1$, where ε is the radial sampling increment and δ is the angular sampling interval. The image to be reconstructed is represented by an $N \times N$ array of pixels represented by d_{ij} $i, j=1, \dots, N$. Using the superscript q to denote the iteration number, let $p^q(\varepsilon k, \delta m)$ represent the projection calculated from the q 'th image estimate d_{ij}^q . These reconstruction techniques update each pixel value by a factor related to the discrepancy between the original projections $p(\varepsilon k, \delta m)$ and the previously calculated projections $p^q(\varepsilon k, \delta m)$. An initial value assumed for all pixel values can be the average pixel density η calculated from the projections by

$$\eta = \frac{1}{MN^2} \sum_{m=0}^{M-1} \sum_{k=0}^{K-1} p(\varepsilon k, \delta m) = d_{ij}^0 \quad i, j=1, \dots, N. \quad (3.22)$$

In direct multiplicative ART, the updating scheme is

$$d_{ij}^{q+1} = \left[\frac{p(\varepsilon k, \delta m)}{p^q(\varepsilon k, \delta m)} \right] d_{ij}^q \quad (3.23)$$

for all $k=0, \dots, K-1$ and $m=0, \dots, M-1$, where only those pixels possibly involved with the specific projection value defined by the values of k and m are modified. For direct additive ART, the modification scheme is

$$d_{ij}^{q+1} = \max[d_{ij}^q + (p(\varepsilon k, \delta m) - p^q(\varepsilon k, \delta m))/T, 0] \quad (3.24)$$

for all $k=0, \dots, K-1$ and $m=0, \dots, M-1$. In (3.24), T is the number of pixels in the projection ray defined by k, m and again, only those pixels in the specified projection value updated. In both of these algorithms, the process is iterated to produce the final result.

3.1.3 Relative image quality and computational requirements

Two factors by which a specific reconstruction algorithm can be judged are image quality and processing requirements. Some qualitative relationships will be stated here for reference in later sections.

Image quality is at best difficult to measure and, as such, it is hard to compare this aspect of different methods. It would appear from the literature that in most cases convolutional-backprojection generates the highest quality images, ART the second best and direct Fourier third. There are, however, some situations in which algebraic reconstruction techniques perform better than CBP [2]; we will not comment any further on this. Some of the reasons for this ordering of CBP, ART and DF follow. Interpolation is an error-producing operation because it must approximate an unknown value with a finite number of

calculations on possibly noisy data. Since ART is based on a discrete model and employs only the available data, no interpolation is required. Convolutional-backprojection requires only 1D interpolation in the radial direction and can be performed quite accurately. Since the DF method requires 2D interpolation, given the same amount of computational time as a 1D method, it can only result in an inferior result. That 2D interpolation is poorly defined and a less studied problem further complicates the issue. Secondly, the fact that the density of polar format Fourier data decreases for increasing R (frequency) implies that more widely separated polar points are used to calculate one rectangular point at higher frequencies than at lower frequencies. This larger separation (lower density) will cause a larger uncertainty and consequently more error in the interpolated value. It should be pointed out that interpolation error is in most cases the single largest source of processing error in these reconstruction techniques. From these comments, ART would appear to be the superior method; however, the recursive and consequently asymptotic nature of ART tends to limit the obtainable quality.

Another factor affecting image quality is the sensitivity of the reconstruction method to noise in the data. To our knowledge, no definitive statement has been made establishing one of these techniques as the superior method; however, some comments are in order.

Algebraic reconstruction techniques are recursive in the sense that the old image is processed in order to generate the new image. Since recursive techniques are generally less sensitive to noise than non-recursive methods [31], it is reasonable to expect ART or recursive

techniques to perform relatively well with noisy data. In reference to the direct Fourier and convolutional-backprojection methods employed here, both can be characterized with linear operators. The direct Fourier technique generates the image via an inverse 2D Fourier transform which is of course a linear operator. Prior to this, the projection data are transformed with a FT into the Fourier domain and a 2D interpolation is performed (both of these operations are linear). In CBP, the 1D filtering of each projection is linear and backprojection is also a linear operation. If an additive noise model is assumed, then in either the DF or CBP case the resultant image can be treated as a sum of the reconstruction technique operating on noise-free data and the reconstruction technique applied to the noise only. The DF method does not contain a single step in which the noise can be reduced. In CBP the filtering operation can also be used to help reduce the noise in the data. Details will be discussed in Section 4. Thus in terms of sensitivity to noisy data, CBP may be expected to perform slightly better than DF. Further, both the 1D and 2D interpolators generate error and in this sense can be treated as noise sources. Since the error introduced by a 2D interpolator is generally larger (as discussed earlier) than for a 1D interpolator, the CBP image should again be somewhat superior to the DF image in terms of noise and image quality. In some of the examples provided these features can be identified.

The second issue of processing requirements will be discussed with reference to computational and memory demands. Computational needs will be accessed by the most operationally demanding process in the algorithm in terms of multiplications and additions. In DF methods, if a 2D

interpolator requiring less than $\log_2 N$ operations per interpolated value is employed, then the most demanding process is the calculation of 2D inverse DFT's. This operation requires $O(N^2 \log_2 N)$ calculations. For both CBP and ART, $O(N^3)$ operations are required by the image generation procedures. In CBP this image generation (reconstruction) is performed only once while in ART this procedure is repeated many times. Additionally, ART has the secondary expense of calculating projections prior to each iteration. Including these additional factors, ART requires more operations than CBP but still approximately $O(N^3)$. As expected, the higher quality images are more expensive in terms of computations, with ART and CBP imposing the largest computational burden and DF the least.

Lastly, the storage requirements of the various techniques will be reviewed. In order to meet the sampling requirements of the DF method, approximately $8N^2$ memory locations are needed to store the complex DFT of the image. For both CBP and ART the largest memory demand is the KM locations needed to store the projection data. Because of sampling requirements, KM is usually about twice the size of the N^2 locations required to store the image. This brings the total storage needs for CBP or ART to approximately $3N^2$ locations which is still less than half that needed for the DF technique.

3.2 The Missing Cone Problem and Some Solutions

In this section the missing cone problem is explained and some of the solutions other researchers have proposed are discussed. The missing cone problem belongs to the class of so-called inverse problems. These problems are characterized by having incomplete observations of a scene and partial constraints on the reconstructed solutions. For the missing cone problem the incomplete observations are represented by the set of projections known only over a limited angle $\beta < \pi$; the complete observation set would include projections over π . Some of the partial constraints that could be included are: non-negativeness of the reconstruction and observations, spatial or frequency bounds, signal magnitude limits and specific known structural features such as shape, size or position of objects in the image. Examples of more subtle sources of knowledge, often called a priori information, include a known degree of smoothness in one domain as a consequence of a spatial or frequency limit in the other domain, certain symmetries in Fourier space or restrictions on phase. These are referred to as a priori information because while they can sometimes be derived as being a consequence of the constraints, rarely are they explicitly exploited or employed in reconstruction/enhancement schemes.

As discussed in the introduction, the missing cone is a problem of limited or restricted perspective. Consider Figure 3.5. Projections available over β provide resolution in x but little information in y . To improve the resolution in y , information (projection data) is needed over the $\pi - \beta$ ($= \alpha$) region. Another interpretation provided by the PST is that spectral information is missing over the α region in the

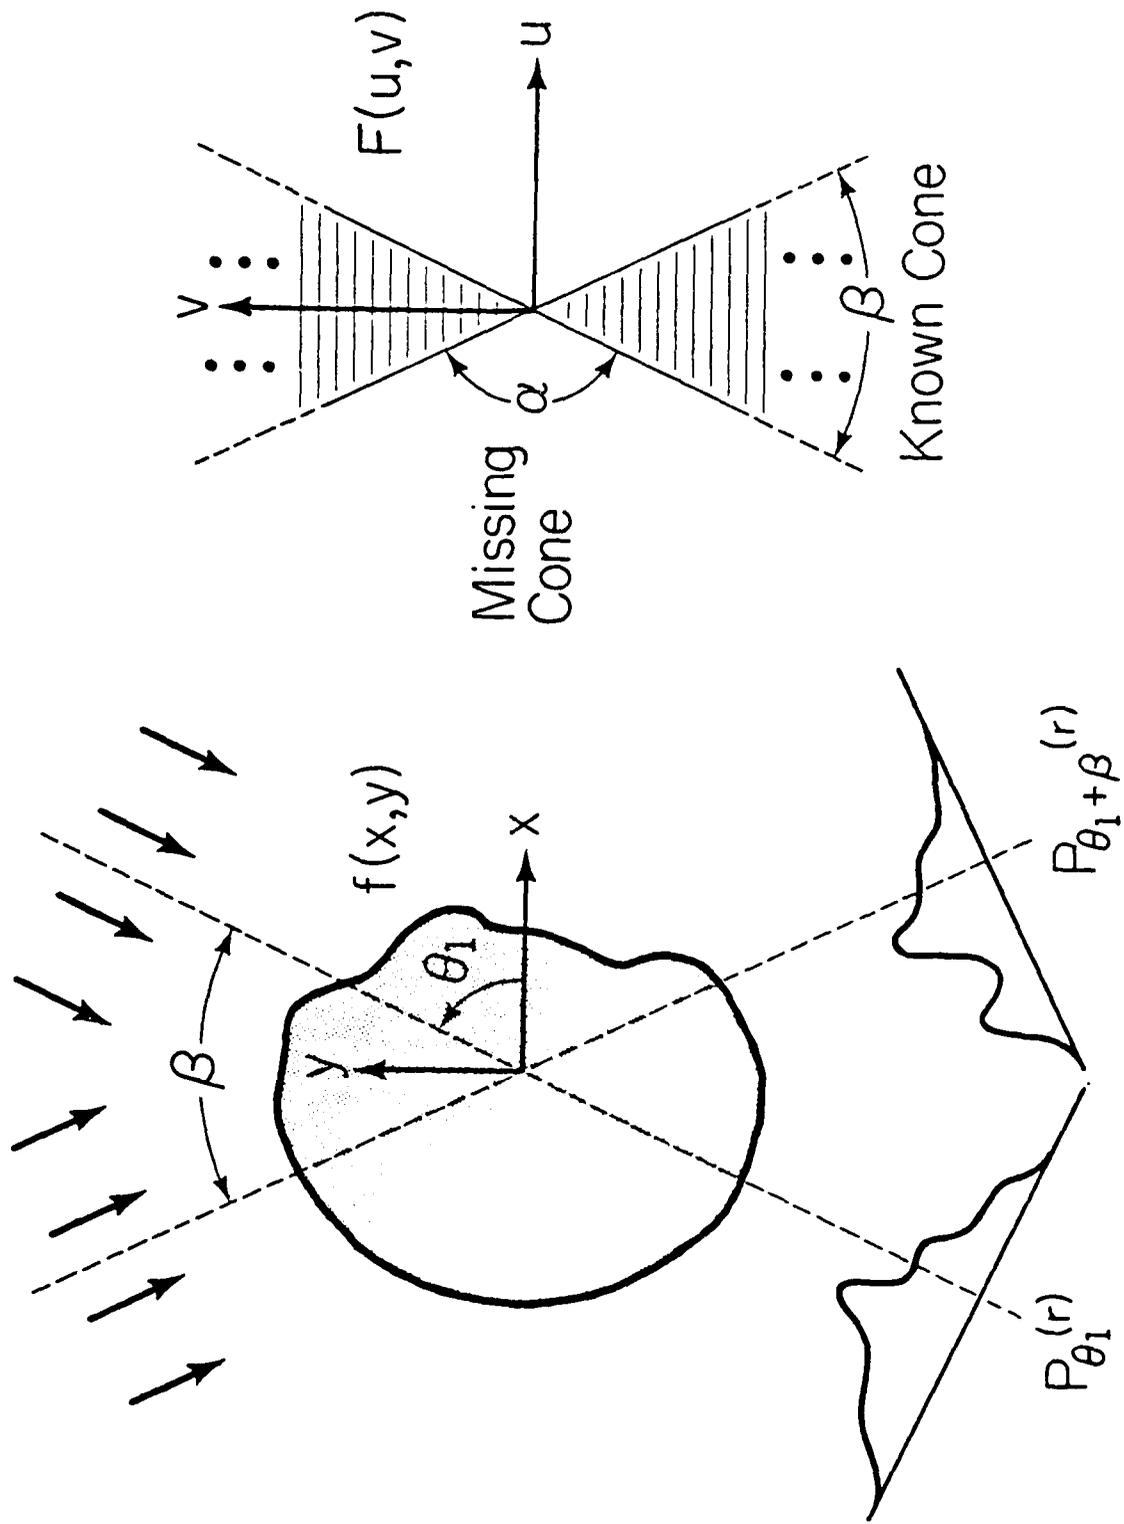


Figure 3.5 Missing cone problem.

Fourier domain. With this view, the goal of an enhancement scheme is to generate approximate information over the Δ region of the Fourier domain. The missing cone situation can now be interpreted as a spectral extrapolation problem. One extrapolation technique is the maximum entropy method (MEM) of Burg [32,33]. In this technique, the idea is to choose from the infinity of possible extrapolations the one that maximizes some measure of entropy in the solution. By maximizing the entropy, the most unbiased solution that is consistent with the known data and constraints is chosen. This solution is purported to be the most reliable, based on the principle that no information has been added that is inconsistent with the given data. In the Fourier domain, one proposed measure of entropy is

$$E = \iint_{\Delta} \log F(u,v) du dv \quad (3.25)$$

where $F(u,v)$ is the 2D Fourier transform of the image [34]. The symbol Δ represents the region of effective support for $F(u,v)$ in Fourier space. It should be noted that any measure of entropy is artificial in the sense that a deterministic signal has zero entropy. Since the known and unknown data are assumed, at least implicitly, to represent a unique and deterministic object or scene, measures of entropy are clearly somewhat artificial. Placing these theoretical notions aside, the source of rationalization for (3.25) is the definition of entropy rate (ER) of a random process. The entropy rate for a stationary, band-limited random process is

$$ER = \int_{-\Omega}^{\Omega} \log S(\omega) d\omega + O(\omega) \quad (3.26)$$

where $S(\omega)$ is the power spectrum of the process, Ω is the cutoff

frequency and $O(\omega)$ represents higher-order terms. While Fourier MEM techniques have been applied in other imaging situations [34] with apparent success, to our knowledge they have not been used in the tomographic situation.

Maximum entropy techniques can also be applied in the spatial domain if an appropriate measure of entropy can be defined. In [35], Baba et al. discuss the details and present some results of a spatial domain MEM where the measure of entropy is defined as

$$E = \sum_i \sum_j \log d_{ij}, \quad (3.27)$$

and where d_{ij} are the pixel elements of the image. These results were clearly superior to those of an ART technique that these authors previously considered [36]. Gordon et al. [37] discuss a similar ART technique applied to the missing cone problem in which the entropy is claimed to be maximized as a result of the technique although it is not a specific objective. The published results in these two papers [36,37] are nearly identical.

In the application of ART to the missing cone problem, only a minor modification is needed. Referring to the discussion in Section 1, instead of updating the pixels based on a function of all the original and all the calculated projections, only the original projections over β and calculated projections over β are used to modify the pixels as in (3.23) or (3.24). Because the updating procedure does not employ or calculate projections over the α region, no constraints can be imposed on these projections. This means that no restrictions can be imposed on the image as a result of constraints applied to the projections. The

point here is that ART leaves an entire segment of the data set unconstrained and accordingly ignores a valuable source of information that could be used to enhance the image.

The reason that MEM is superior to ART is now obvious. Since MEM generates data in the α region and applies solution restricting constraints both on these data and as a result of this application on the image, some additional implicit information is gleaned from the known data. This additional information is then used to generate an enhanced image. In ART the only constraint used is the raw data provided. No a priori knowledge or additional constraints are included in order to enhance the solution. Finally, since MEM actually attempts to extrapolate the spectrum prior to reconstruction and ART simply reconstructs the image, MEM should produce the superior result.

From the results of MEM and ART as applied to the missing cone problem, much improvement in image quality can be obtained by the imposition of constraints and the inclusion of a priori information. Unfortunately, the improvement MEM realizes over ART is obtained at a significant computational cost. While ART requires $O(N^3)$ operations per iteration, MEM consumes $O(N^6)$ operations per iteration for the same $N \times N$ image. It is desired to find a technique that is computationally comparable to ART but which can also incorporate constraints and a priori information in some manner to achieve results superior to MEM.

Some other solutions to the missing cone problem include a Bayesian approach [38], which maximizes an a posteriori conditional probability density. This conditional probability density is related to assumed statistical measures of either the data or of the required image. In

another approach, a set of eigenfunctions are defined for typical tomographic images. It is proposed that any image (in the class) then can be reconstructed with linear combinations of these eigenfunctions. The problem then reduces to one of determining an appropriate set of eigenfunctions for the class of images under consideration, and then estimating the coefficients from the available data [39].

The last technique to be discussed was proposed by Lent and Tuy in [40]. In this method, the various sources of constraints are used to define an intersection of convex sets that the solution must lie in. The algorithm iteratively applies these various constraints by the use of orthogonal projections on the convex sets in which constraints are available. Since with each application of constraints, the error must be reduced, or the distance to the intersection reduced, the algorithm will converge (assuming that the intersection is non-empty).

3.3 Algorithms for the Solution of the Missing Cone Problem

The previous section characterized the missing cone problem as a spectral extrapolation problem with constraints on the solutions and partial information in both the spatial and Fourier domains. In this section, the techniques of one-dimensional spectral/spatial extrapolation are generalized into two-dimensions. This two-dimensional technique is then modified for the missing cone problem by the incorporation of problem specific relationships into the algorithms.

An unstudied application of Gerchberg's algorithm is for the case in which multiple unconnected intervals of the signal are known. In this situation the recursive filtering and substitution scheme of Figure 2.2, although still theoretically applicable, is no longer easy to implement. Instead, the transform-and-substitute (or constrain) technique originally described is used. It is conjectured that this multiple interval case possesses convergence properties similar to the single interval situation. In the theoretical case with no noise or other signal degradations, any one interval uniquely describes the complete signal. As a result of either the analyticity or Taylor series argument, one interval is sufficient and the rest are redundant. Practically, with noise and other data collection degradations present, the multiple intervals can be considered as further constraints on the solution. The net effect of multiple intervals is to further constrain the system with a better solution as a result. For the periodic case, the multiple intervals would simply reduce the size of the extrapolation/interpolation region and the iterates still approach the unique solution. For the non-periodic case, since the cause of the divergence phenomena is the inability to perform ideal filtering and this problem is still present, the convergence behavior will be the same as in the single interval case. It will be seen, in later examples, that the convergence phenomena present in the 1D case carries over to the 2D situation.

3.3.1 Some algorithms

Generalizing the multiple interval 1D algorithm for the 2D missing cone problem is straightforward. Consider the case where regions of the 2D Fourier transform of an image are missing and some constraints are present in the spatial domain. The 2D Gerchberg algorithm (2DGA) parallels the 1D algorithm exactly. First, an approximation is made to the unknown parts of the 2D Fourier transform. Next, a first estimate for the image is generated via a 2D inverse Fourier transform and the constraints are imposed on the spatial image, thereby reducing the error. This estimate of the image is then Fourier transformed to obtain the 2D Fourier transform. The error is again reduced by the substitution of the known spectral data and the algorithm is repeated. See Figure 3.6. Typical methods for determining convergence are: to monitor small changes in either the image or Fourier domains or to test for correlation between the known observations and the calculated values that correspond to those observations. In [41] Renjen and Huang discuss the details of this algorithm and present some preliminary results. As other authors [41,42] have observed with similar techniques, it was noted that the error in the 2DGA decreased for a few iterations and then increased. For the 2DGA it is reasonable that the cause of increasing error is the same as that for the 1D case.

The 2DGA is by no means restricted to the missing cone problem. It is sufficiently general to accommodate nearly any combination of regions and constraints. Variations of the iterative technique of transforming between domains where information or constraints are available have been applied to a wide class of problems. As an example, consider the case

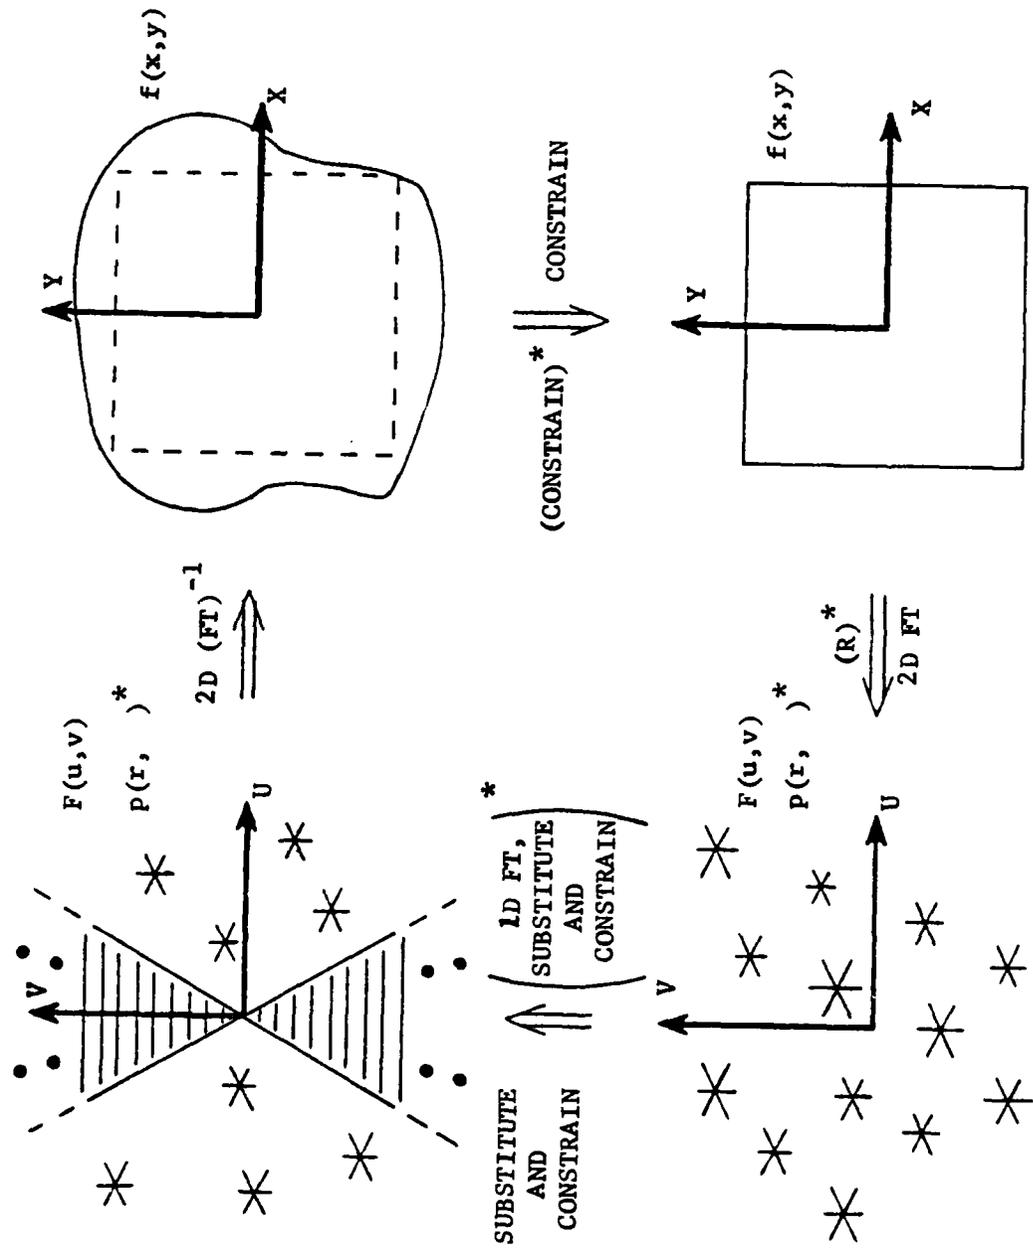


Figure 3.6 2D Gerchberg or Projection-Slice* algorithm.

where the magnitude of the Fourier transform and some image information is known but phase information is missing [43].

To best exploit the power of these iterative techniques, problem dependent relationships should be identified and incorporated into the algorithm. For the missing cone problem a unique relationship that can be exploited is the link between projections, the image and slices of the 2D Fourier transform of the image. The first algorithm to be discussed involves calculating the projections of successive image estimates and then using these data to reconstruct the image. This algorithm is also illustrated in Figure 3.6. Starting with the original incomplete data set (the known cone), an approximation is made to the unknown data. From this estimated but complete set of projections, an image is reconstructed using either a convolutional-backprojection or direct Fourier method. This image is the first approximation to the result. Next, constraints are applied to the image, reducing the total error. From this modified image a complete set of projections is calculated and additional constraints are imposed on the projections. The original projections, those over β , are substituted into this data set reducing the error a second time. This now complete data set is next used in reconstructing a better image estimate and the process is repeated until some convergence criterion is met. This algorithm will be referred to as the Projection Slice Algorithm (PSA) [44].

The only difference between the PSA when implemented with direct Fourier reconstruction and the 2D Gerchberg algorithm is that the 2D Fourier transform of the 2DGA is replaced by the calculation of projections and the application of the Projection-Slice Theorem, i.e., a

direct Fourier reconstruction. Note that in both cases a 2D inverse Fourier transform is used to obtain the image from the Fourier domain data. In the case of using convolutional-backprojection to generate the image, a totally spatial domain version of the PSA is realized which is structurally identical to the direct Fourier method and consequently quite similar to the 2D Gechberg algorithm. While the DF and CBP realizations of the Projection Slice Algorithm may be very similar structurally, they are substantially different in the results they produce. Reasons and examples are provided in the next section.

A second type of algorithm is based on the periodicity of projections, see equation (3.6). Functions $s(r, \phi)$, periodic in ϕ can be constructed for each r , $0 \leq r < K/2$, from the projection data. The available data provide the known intervals in each of these periodic functions. These functions are constructed in the following manner. Let K denote the number of samples in each projection. Consider the signal $p(r, \theta)$ as a function of θ , $0 \leq \theta < \pi$ for $0 \leq r < K/2$. By concatenating the signal $p(-r, \theta)$ for $0 \leq \theta < \pi$ to $p(r, \theta)$ forming $s(r, \phi)$, a signal periodic in 2π is generated as illustrated in Figure 3.7. Assuming a reliable band-limit is known for this periodic function, the signal can be extrapolated/interpolated to arbitrary accuracy as previously mentioned. Practical considerations affecting the accuracy of this extrapolation are: accuracy of the band-limit, width of the known intervals, the number of iterations performed and the specific implementation of the algorithm. By performing this extrapolation for each r , $0 \leq r < K/2$, the unknown intervals of each periodic function can be recovered. Taken as a set, these extrapolated functions uniquely determine the missing cone.

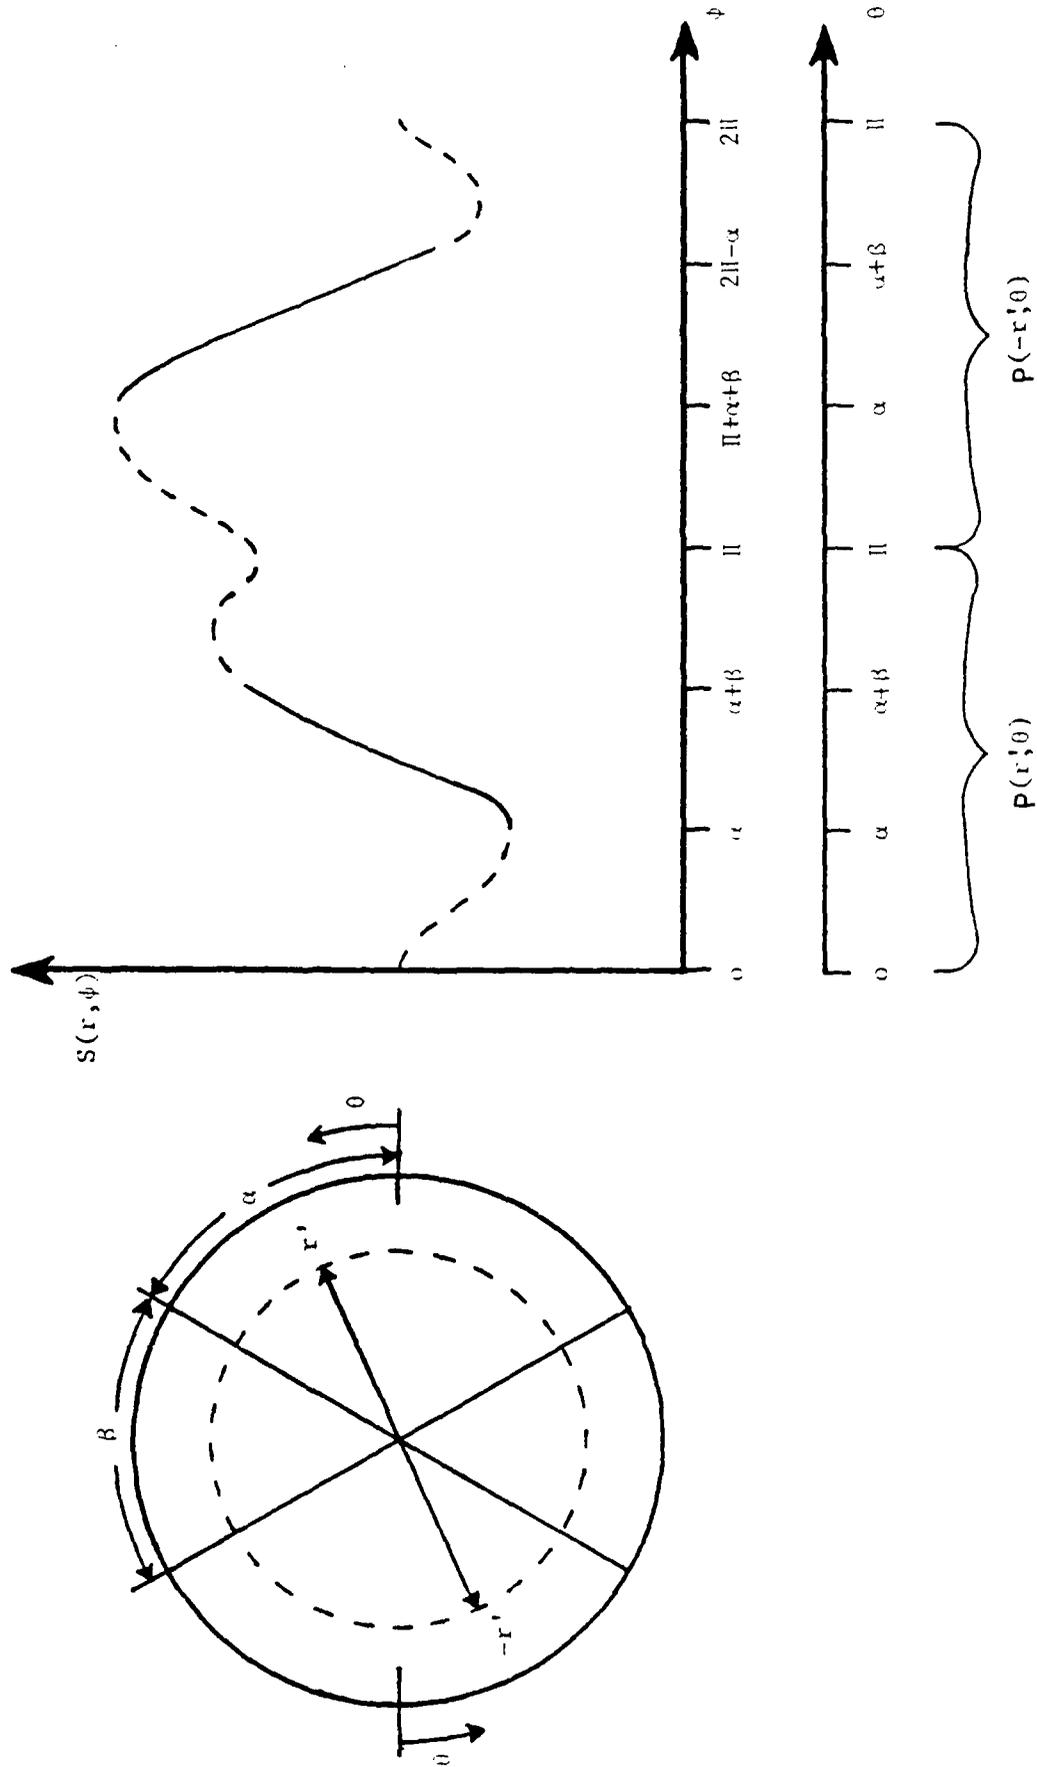


Figure 3.7 Angular iteration method.

The now complete data set is used to generate the image with either direct Fourier or CBP reconstruction. This algorithm is called the Angular Iteration Method (AIM) [44].

Studies of the mechanics of tomography indicate that the angular bandwidth for these periodic signals is a function of the spatial extent of the image. Needed is an accurate method of calculating this band-limit. While it has been shown that these periodic functions are not strictly band-limited in ϕ even for a function $f(x,y)$ that has a bounded region of support in the (u,v) Fourier plane, these functions can be considered to be effectively band-limited in ϕ [45]. This effective band-limit contains 98% of the spectral energy. Rattey and Lindgren [45] also supply the required relationship between spatial extent r_T and ω_ϕ . Denoting by ω_B the radius of support for $f(x,y)$ in the (u,v) Fourier plane, then

$$\omega_\phi = 1 + r_T \omega_B. \quad (3.28)$$

Using the above equation, a reasonable approximation can be found for ω_ϕ as a function of the radial extent of the image.

The rate of convergence for Papoulis' algorithm can be significantly improved if a good initial estimate is used. In AIM this translates to having an initial approximation to the missing cone. A good source of these data is the PSA. An alternative algorithm is a two-step method where one or more iterations of the PSA are performed in order to obtain a good estimation for the missing cone data. Known data augmented by this approximation to the unknown projections, which are available after one pass through PSA, are used as the starting point for AIM. This initial estimate to the unknown intervals of the periodic

functions will improve the extrapolation obtained after a finite number of iterations on any of the $K/2$ signals. The image can now be reconstructed from the complete data set or further iterations of PSA and AIM could be performed.

While other variations on these algorithms and concepts are possible, those described above exemplify the key idea of iterating between various domains in which constraints and information are available. As in Papoulis' and Gerchberg's algorithms, the imposition of constraints forces convergence. Another interpretation concerning this type of algorithm is that the observed data set describes a class of solutions. The effect of applying constraints is to narrow this class of solutions. If more constraints can be imposed, the class of solutions will be smaller and consequently, the resultant image will be of higher quality (less uncertainty). The purpose of iterating between various domains is to supply a means of applying various constraints and imposing additional sources of information in order to force consistency between the iterative solutions and the available information. Convergence is obtained when observations and constraints agree to some specified tolerance. If, however, the applied constraints contradict or force an inconsistency in the iterations, then the algorithm may diverge or become stuck in a oscillatory loop. Therefore, it is important to insure that the applied constraints do not contradict each other.

3.4 Computational Details and Experimental Results

In this section some experimental results from the projection-slice algorithm and the angular iteration method are presented and discussed. Additionally, many of the computational details concerning implementation of these techniques are reviewed. These topics include interpolation, filtering, integration and methods for calculating projections. Since interpolation is necessary in both of the reconstruction techniques, this problem will be reviewed first.

3.4.1 Computational details

One-dimensional interpolation is an extensively studied field with many significant results. It will be sufficient for the purposes of this work to simply state and employ some of these facts. An important concept is the idea that 1-D interpolation can be considered a linear filtering problem [46]. This view is provided via a frequency domain analysis of the interpolation process. The key result is that 1-D interpolation in a band-limited signal can be performed with arbitrary accuracy by a linear finite impulse response (FIR) filter [46]. As an example, consider 1-D linear interpolation between two points $x(n)$ and $x(n+1)$ separated by a distance Δ . The desired interpolated value $y(n)$ is at a distance r measured from $x(n)$. A FIR filter has the form

$$y(n) = \sum_{k=-M}^M x(k+n)b(k) \quad (3.29)$$

where M is finite. For the linear interpolator,

$$y(n) = \frac{x(n)}{2(1-r/\Delta)} + \frac{x(n+1)}{2(r/\Delta)} \quad (3.30)$$

which is the same form as (3.29) with $M=1$, $b(-1)=0$ and $b(0)$ and $b(1)$ as given above. In order to obtain a more accurate interpolation, more terms are required implying a larger M and more non-zero coefficients. If the position of the interpolated point changes as a function of n , then the values of $b(*)$ must also change if linear interpolation is desired. In the implementation of convolutional-backprojection, the linear interpolator of (3.30) is used with variable coefficients.

Two-dimensional interpolation is significantly less well-defined and accordingly much more difficult to perform. In this work a bilinear interpolator is employed. This interpolator is of the same form as (3.28) as it is essentially a FIR filter where the four known values surrounding the desired point are involved in the calculations. Denoting these known values as $x(1)$ through $x(4)$, the equation for interpolating y is

$$y = b(1)x(1)+b(2)x(2)+b(3)x(3)+b(4)x(4). \quad (3.31)$$

The coefficients $b(1)$ through $b(4)$ are functions of the position of y . Ordering the $x(n)$ clockwise around the point y , let Δ_{ij} denote the distance between two corner points $x(i)$ and $x(j)$. Let r_i represent the length of the line from $x(1)$ to y projected onto the line defined by $x(i)$ and $x(j)$. The value of Δ_a is the distance between the point r_1 along the line $x(1)$ to $x(2)$ and the point r_3 which is along the line joining $x(3)$ to $x(4)$. Variable r_a represents the distance from r_1 to y . The values for Δ_b and r_b are similarly calculated from the points r_2 and r_4 . Denoting the ratio of r_i/Δ_{ij} by δ_i , one possible set of coefficients for (3.31) is:

$$b(1) = 0.5[(1-\delta_a)(1-\delta_1) + (1-\delta_b)(1-\delta_4)]$$

$$b(2) = 0.5[\delta_1(1-\delta_a) + \delta_b(1-\delta_2)] \quad (3.32)$$

$$b(3) = 0.5[\delta_3\delta_a + \delta_b\delta_2]$$

$$b(4) = 0.5[\delta_a(1-\delta_3) + \delta_4(1-\delta_a)].$$

With this set of coefficients, when y is on any line joining corner points, the above b 's reduce to a 1-D linear interpolator as in (3.30).

For the convolutional-backprojection method a filtering operation is necessary along each projection. In the simulations presented in this chapter, the filtering operation (equation 3.16) is performed in the Fourier domain by use of FFT's. Reasons for this choice include computational efficiency and a more flexible means for altering the approximation made to $|R|$. The approximation used is $|R|$ multiplied by an appropriate window [47]. Some sort of approximation is necessary because $|R|$ is not a realizable filter. To achieve a realizable filter, a function is used to window $|R|$ such that it is of finite extent and closely approximates $|R|$ at frequencies where useable spectral information is present. A second purpose of the window is to ameliorate the effects of Gibb's phenomena by introducing a smooth transition from $|R|$ to zero. Lastly, by properly choosing the window and the width of the window, some of the noise present in the data can be filtered out. The window chosen is a Hamming window [47]. The filter function, $F(n)$, implemented in this work is given by

AD-R142 421

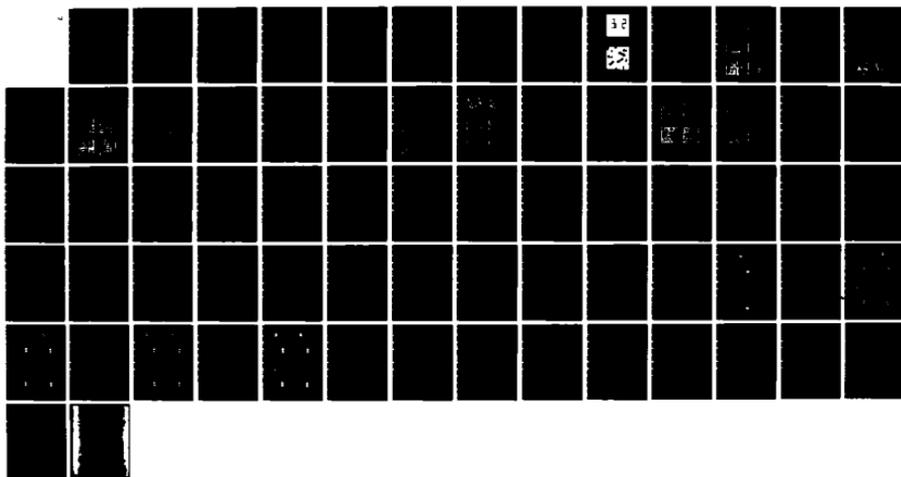
THE MISSING CONE PROBLEM IN COMPUTER TOMOGRAPHY AND A
MODEL FOR INTERPOLA. (U) ILLINOIS UNIV AT URBANA
COORDINATED SCIENCE LAB D A HAYNER JUL 83 R-989
N00014-79-C-0424

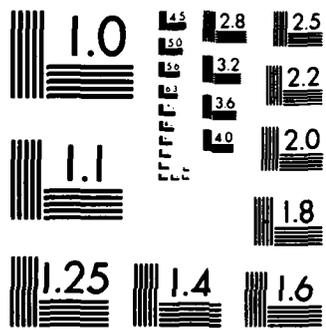
2/2

UNCLASSIFIED

F/G 17/9

NL





MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

$$F(n) = \begin{cases} 2\pi(n-1)/Q & n=1\dots L-1 \\ [\pi(n-1)/Q][1 + \cos \frac{\pi(n-L)}{Q/2-L}] & n=L\dots Q/2 \end{cases} \quad (3.33)$$

where L represents the damping factor (or window length) and Q is the length of the FFT used. Figure 3.8 illustrates the effects of L on the filter function $F(n)$. The damping factor denotes the value of n at which the window starts to modify $|R|$. Through experimentation a value of L that passes 78% of $|R|$ unmodified was chosen. For an FFT length of 256 this corresponds to L equal 100.

In the direct Fourier implementation of the projection-slice algorithm, a 2-D low-pass filter is employed to aid convergence as in the 2-D Gerchberg technique and to reduce the effects of noisy data. This filter is realized by low-pass filtering each projection prior to the 2-D interpolation step. Since the data will be in the Fourier domain for the interpolation step, implementation of this filtering (windowing) operation is easily performed by truncation (or by multiplying point by point with the appropriate window). The cutoff for this filter is primarily determined by the desired resolution in the final image although other factors may influence the choice. To employ a filter with a cutoff that passes frequencies representing detail finer than needed would be ignoring a possible source of constraint on the solution. This idea also can be a motivating factor in the choice of the damping factor, L ; see (3.33).

A third filtering operation that can be performed is angular filtering of the projection data to apply further constraints and reduce the effects of noise. Since the projection data can be formatted as signals periodic in angle, (see equation 3.6) it is possible to ideally

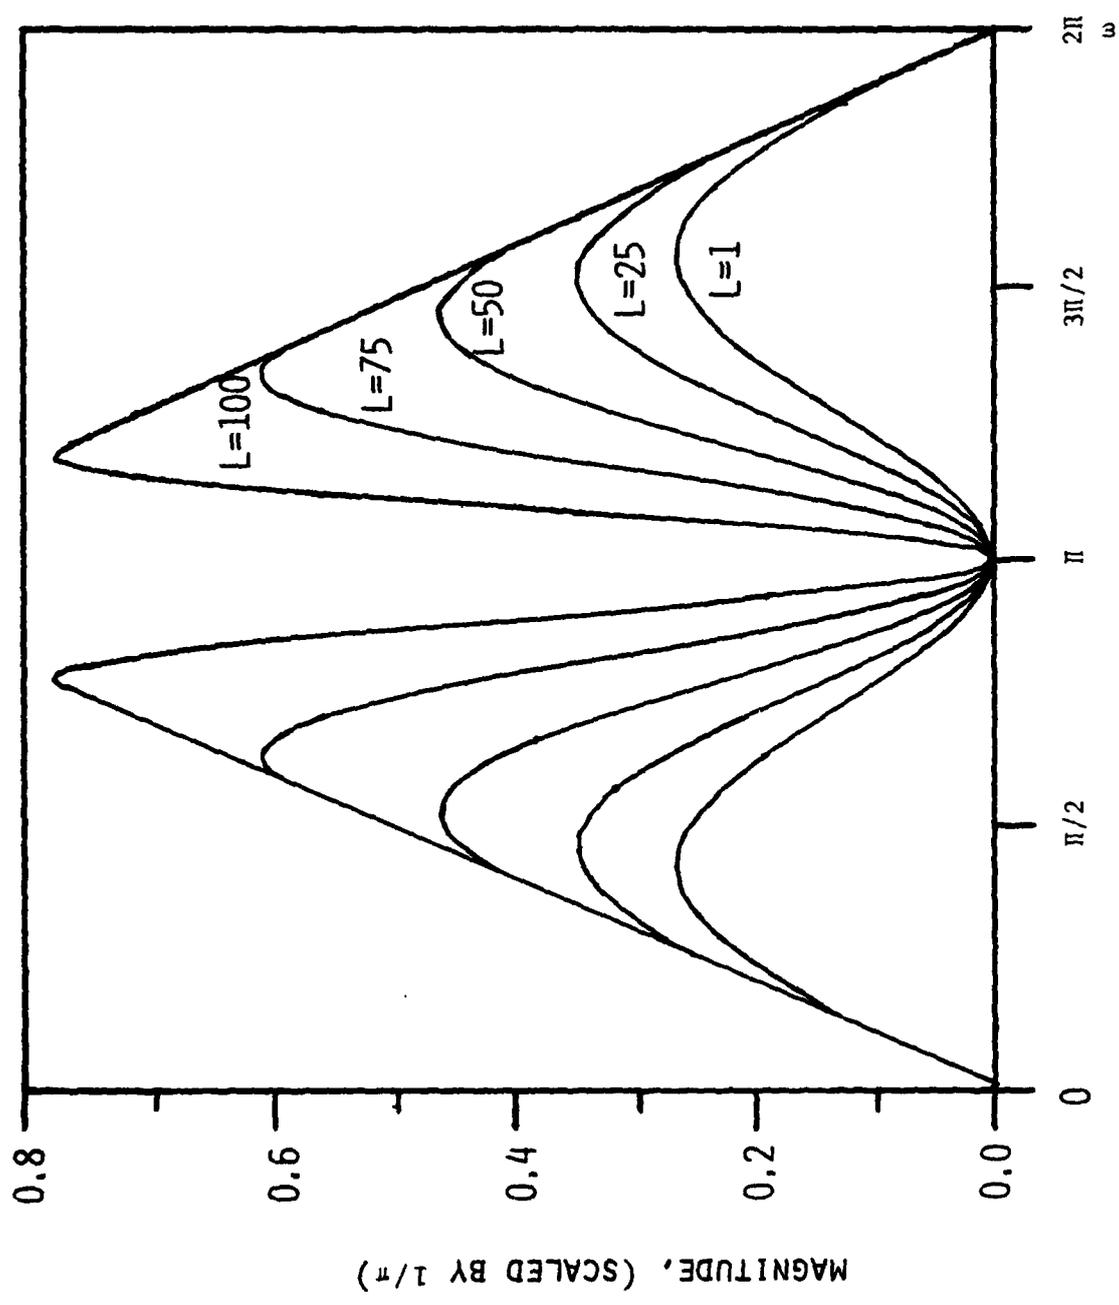


Figure 3.8 |R| modified by a Hamming window.

filter these data. What is implemented is a low-pass filter passing only spectral information that is consistent with a desired resolution. Windows may be used to modify the effects of this filtering to obtain specific results. As in the radial filtering, the choice of a band-limit higher than the needed resolution would squander resources. It must be noted that this angular filtering operation is independent of the filtering employed in the angular iteration method. In AIM, the purpose of angular filtering is to extrapolate the signal. The angular filtering is employed as the primary method of enhancing the reconstruction. While these two ideas are closely related, in AIM this angular filtering is the primary source of extrapolation. In the projection-slice algorithm it is not necessary; it is simply used as another source of constraint.

An important step in both the PSA and AIM is the calculation of projections. As mentioned earlier, the various non-symmetrical and non-linear effects present in real projection data are not considered in this work. The reason is simple. Many of the aforementioned effects are peculiar to the specific problem at hand. Since the motivation of this work is to provide some general concepts and ideas, detailed consideration of these various problems could severely limit the applicability of this work and certainly obscure some of the issues. For similar reasons, only parallel beam projection data are employed. This is not particularly restrictive because other data sets, such as fan beam, can be employed by either modifying the backprojection operator [2] or by calculating parallel beam projection data from the fan beam data -- a process called rebinning [2].

Two different projection models have been used in the algorithms. Thus far not enough difference has been noted in the results to justify considering one method as superior. The first and simplest technique called the nearest point method, models a projection scheme in which very narrow beams of radiation are used to collect data. Let the image be represented by a set of pixels d_{ij} for $i, j=1 \dots N$ (see discussion concerning ART). Referring to Figure 3.9, consider the projection at angle γ ($=m\delta$). For each value of r ($=zk$), a summation is performed along the line defined by γ and r . In order to insure that a relatively consistent number of points are included in each projection value, the summation over the image is indexed by the variable corresponding to the direction of longer intersection between the image and line. In Figure 3.9 indexing is over i for the projection at angle γ . If the projection were at angle $\gamma+90^\circ$, then indexing would occur over j . For each i or j , the value of the point nearest to the line is included in the sum.

In the second technique, all the points that lie within a specified distance of the projection line are included in the sum. This technique is more flexible because both narrow beam, as above, as well as wide-beam data collection scheme can easily modelled. If the defined beamwidth is sufficiently wide, this scheme will not have to account for the relative projection angle, as in the first case, since the width of the beam will naturally include a fairly consistent number of points.

A third technique, not used in this work, models each pixel as a square (or other regular shape). In the calculation of a projection value, a beam model is used in which the area of each pixel square intersected by the beam is calculated and a corresponding portion of

that pixel's value is summed. This technique is computationally expensive in comparison to the previous methods and for this reason was not used in the simulations.

One problem common to nearly all tomographic imaging systems is projection noise. In the process of collecting data, noise will invariably either already be present in the data or will enter the system via the data collection scheme. The latter is often called sensor noise and can be modelled; the first is more difficult. Noisy data are to some extent naturally generated by the artificial projection schemes just discussed. Clearly, neither method will generate perfect projection data, and hence this deviation from ideal can be considered as signal noise. Sensor induced noise is easily simulated by either multiplying or adding a different random noise vector with each projection. In this work, an additive Gaussian noise process is assumed.

The goal of this work is to obtain an improvement in image quality over that which is provided by directly reconstructing the image from the available data. In order to quantify this gain, some type of objective measure must be applied. With the original image available, as it will be in all the examples considered, measures that relate the reconstructed to the original can be calculated. Those measures are the average error (AE) and the variance of the error (VE).

$$AE^q = \sum_{i=1}^N \sum_{j=1}^N \frac{|d_{ij}^i - d_{ij}^q|}{N^2 \eta} \quad (3.34)$$

$$VE^q = \frac{\sum_{i=1}^N \sum_{j=1}^N (d'_{ij} - d^q_{ij})^2}{\sum_{i=1}^N \sum_{j=1}^N (d'_{ij} - \eta)^2} \quad (3.35)$$

In the above d'_{ij} represent pixels of the original image, d^q_{ij} is the reconstruction after q iterations and η is the average pixel value found from (3.22). These are the same measures used by Baba et al. [35].

3.4.2 Experimental results

In this section, some experimental results are presented to demonstrate the properties of these reconstruction/enhancement algorithms. All of the programs are written in FORTRAN VII and run on a VAX 780 under a UNIX operating system. Two significantly different pictures are used in these examples. The first one, Image #1, is the number 32; the second, Image #2, is a picture of chromosomes, Figures 3.10 and 3.11 respectively. Both pictures are 64 x 64 pixels in size with 64 gray levels. In all of the following examples, the nearest point projection method is employed in which 64 equally spaced projections taken over π radians represent a complete data set. From each projection, 128 equally spaced samples are obtained. A beamwidth method was also employed, but since no significant differences were noted the nearest point method is used in all of these examples. The missing cone situation is constructed by simply calculating projections over some restricted angle and supplying only this incomplete data set to the algorithms. The term 50% of the data implies that 32 equally spaced projections are taken over $\pi/2$ radians, with each sampled 128

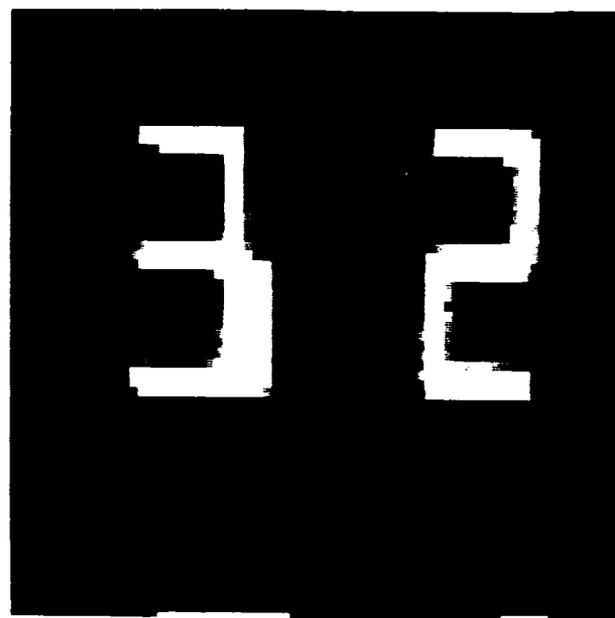


Figure 3.10 Original Image #1.

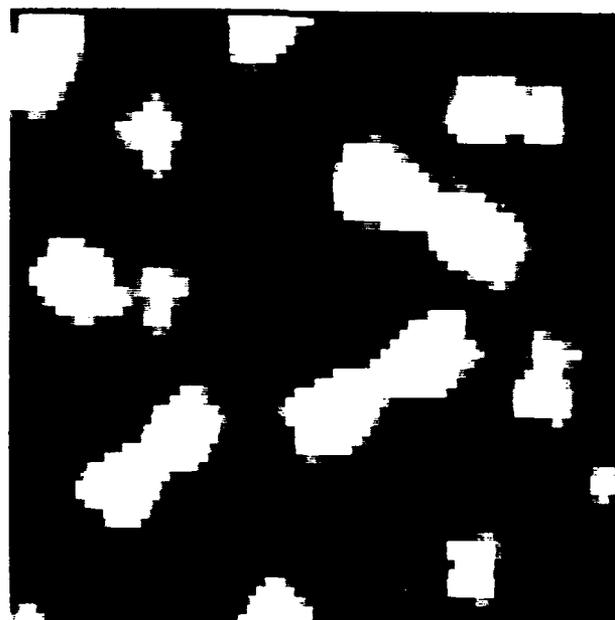


Figure 3.11 Original Image #2.

times. In all cases, projections are calculated over the same relative angle with respect to the picture. A line drawn from the center of the picture to the right-hand edge is the reference line. Projections are always calculated starting from this line, i.e., zero radians is the initial projection angle. Consequently, a 50% data set calculates the first projection at zero radians and the last one at $\pi/2$ radians -- middle of picture to top edge. In all cases, zero is the initial approximation to the unknown Fourier or projection domain data.

As discussed earlier, the projection-slice algorithm can be implemented with either direct Fourier (DF) or convolutional-backprojection (CBP) reconstruction. Examples of the DF reconstruction will be presented first. Figures 3.12a) and 3.12b) show DF reconstructions for Image #1 with 100% and 50% of the data. Figure 3.12c) shows the result after 3 iterations of the PSA algorithm. A visual inspection indicates that PSA has improved image quality to the extent that the 32 can be recognized in Figure 3.12c), whereas it could not be identified in Figure 3.12b). The mean and the variance of the error are shown in Table 3.1 to illustrate the quantitative improvement. In this case, as in all the following cases, the statistics are obtained by comparing the original and reconstructed pixel values over the entire 64 x 64 image array. The same example is shown in Figure 3.13 with the addition of enough Gaussian noise to degrade the average signal to average noise ratio to approximately 20 dB. A similar sequence of examples is shown in Figures 3.14 and 3.15 for Image #2, which is a microscopic image of a collection of chromosomes. In the final images of Figures 3.14c) and 3.15c), it may be difficult to identify these as

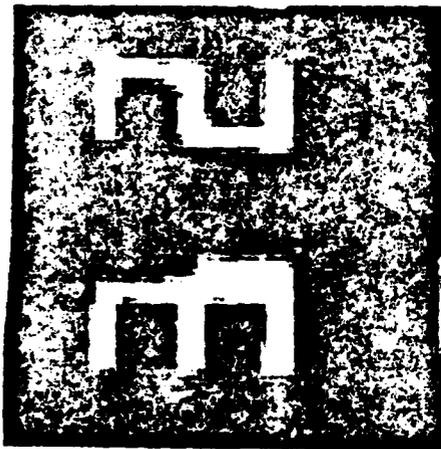


Fig. 3.12 a) Image #1, DF with 100% data and no noise.



Fig. 3.12 b) Image #1, DF with 50% data and no noise.

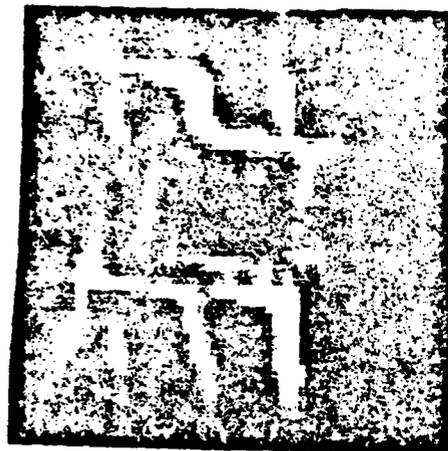


Fig. 3.12 c) Image #1, DF with 50% data, no noise and 3 iterations of the PSA.

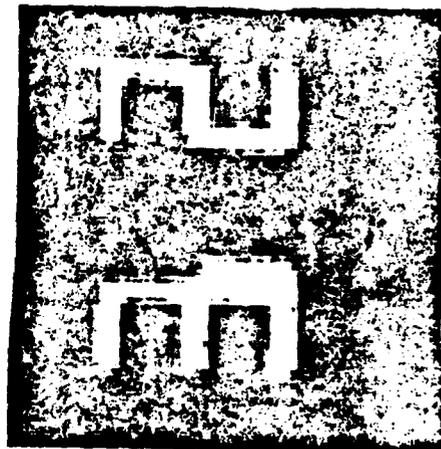


Fig. 3.13 a) Image #1, DF with 100% data and Gaussian noise.

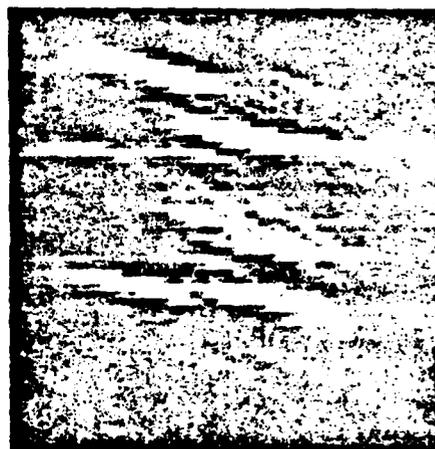


Fig. 3.13 b) Image #1, DF with 50% data and Gaussian noise.

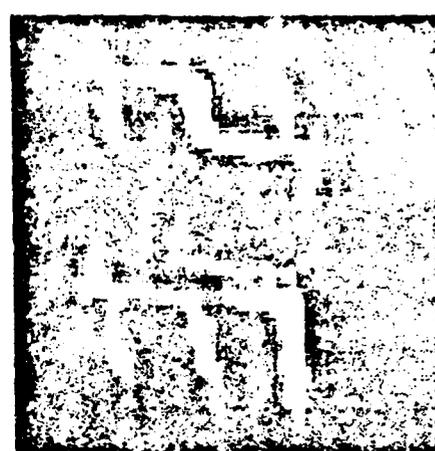


Fig 3.13 c) Image #1, DF with 50% data, Gaussian noise and 3 iterations of the PSA.

Table 3.1 Projection-Slice Algorithm: Direct Fourier reconstruction.

	<u>no noise</u>		<u>noise</u>	
	mean	variance	mean	variance
picture 1				
100%	0.388	0.538	0.382	0.557
50%	0.465	1.265	0.483	1.323
best	0.435	0.958	0.446	1.009
picture 2				
100%	0.364	0.285	0.380	0.358
50%	0.676	1.863	0.692	1.768
best	0.645	1.728	0.555	1.558



Fig 3.14 a) Image #2, DF with 100% data and no noise.



Fig. 3.14 b) Image #2, DF with 50% data and no noise.



Fig. 3.14 c) Image #2, DF with 50% data, no noise and 3 iterations of the PSA.

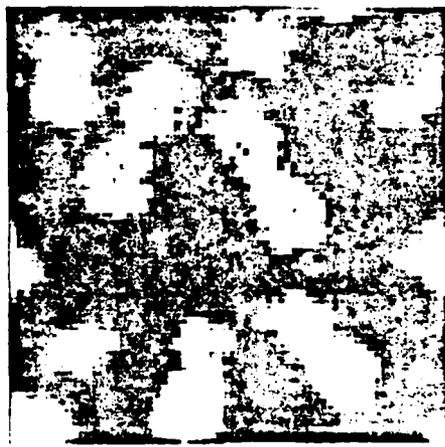


Fig. 3.15 a) Image #2, DF with 100% data and Gaussian noise.



Fig. 3.15 b) Image #2, DF with 50% data and Gaussian noise.



Fig. 3.15 c) Image #2, DF with 50% data, Gaussian noise and 3 iterations of the PSA.

pictures of chromosomes. However, if it is known a priori that the image represents a collection of chromosomes, then the PSA enhancement is sufficient to identify individual chromosomes and to determine their quantity and orientation, something not possible with the initial reconstructions shown in Figures 3.14b) and 3.15b). The effect of additive Gaussian noise is also illustrated by these examples. Since the initial reconstructions result in unrecognizable images, the added noise does not really affect our interpretation. However, in the PSA enhanced results for both images, the noise does not seem to seriously affect our ability to visually interpret the results. The statistics for Image #2 are also included in Table 3.1. Note that in both the noise-free and noisy cases, the algorithm achieves a reduction in the mean and variance of the error. Further, the improvement in the statistics is quite small and does not reflect the noticeable improvement observed by visual inspection.

Figures 3.16-3.19 show a similar sequence of examples using the PSA implemented with convolutional-backprojection (CBP) reconstruction. The corresponding statistics are summarized in Table 3.2. In comparing the CBP examples with the previous DF examples, two noticeable differences are evident. First, the initial reconstructions with 50% data both with and without noise are considerably better than the corresponding examples with DF reconstruction. In Figures 3.16b), 3.17b), 3.18b) and 3.19b), image identification is possible without further enhancement. Second, the PSA-enhancement improves the statistics more than it appears to improve subjective image quality. Perhaps the improvement in visual quality is not as significant as the statistical measures of Table 3.2

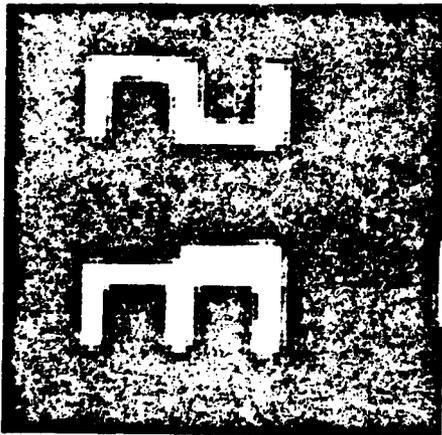


Fig. 3.16 a) Image #1, CBP with 100% data and no noise.

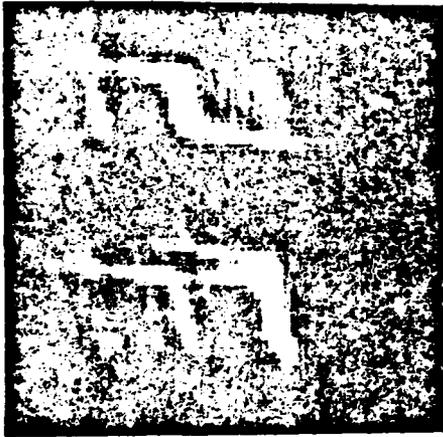


Fig. 3.16 b) Image #1, CBP with 50% data and no noise.

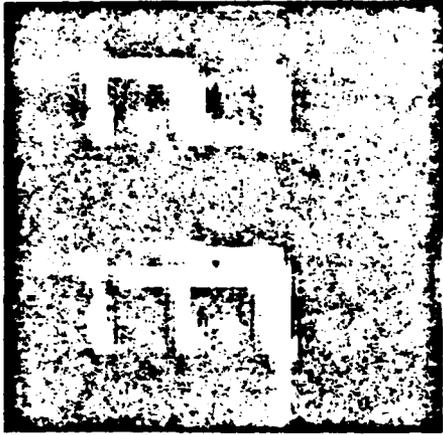


Fig. 3.16 c) Image #1, CBP with 50% data, no noise and 3 iterations of the PSA.

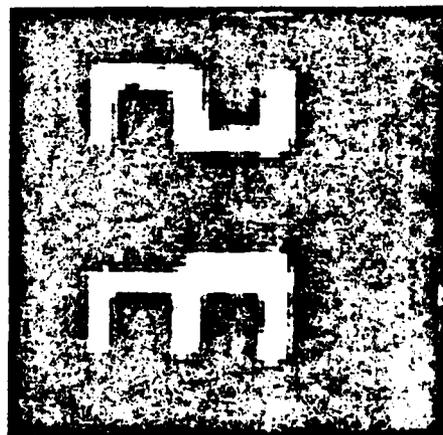


Fig. 3.17 a) Image #1, CBP with 100% data and Gaussian noise.

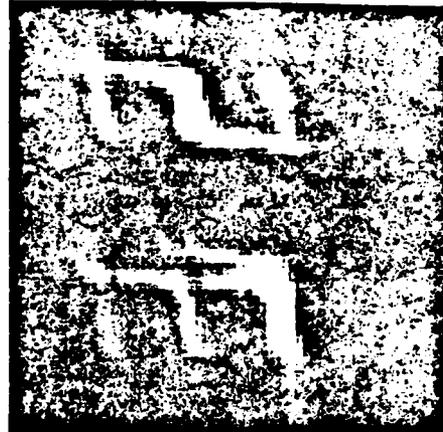


Fig. 3.17 b) Image #1, CBP with 50% data and Gaussian noise.

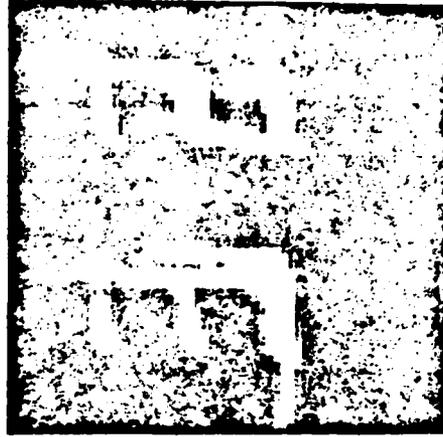


Fig. 3.17 c) Image #1, CBP with 50% data, Gaussian noise and 3 iterations of the PSA.



Fig. 3.18 a) Image #2, CBP with 100% data and no noise.



Fig. 3.18 b) Image #2, CBP with 50% data and no noise.



Fig. 3.18 c) Image #2, CBP with 50% data, no noise and 3 iterations of the PSA.



Fig. 3.19 a) Image #2, CBP with 100% data and Gaussian noise.

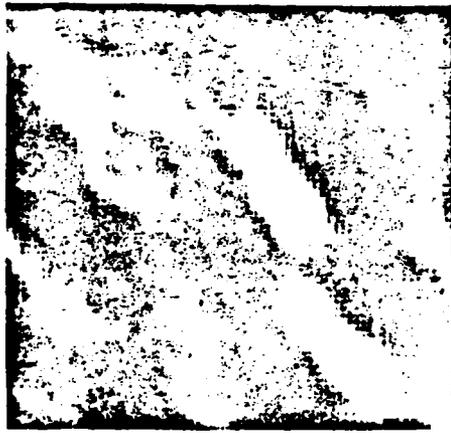


Fig. 3.19 b) Image #2, CBP with 50% data and Gaussian noise.



Fig. 3.19 c) Image #2, CBP with 50% data, Gaussian noise and 3 iterations of the PSA.

Table 3.2 Projection-Slice Algorithm: Convolutional-Backprojection.

	<u>no noise</u>		<u>noise</u>	
	mean	variance	mean	variance
picture 1				
100%	0.144	0.110	0.152	0.121
50%	0.530	0.935	0.523	0.930
best	0.327	0.617	0.346	0.732
picture 2				
100%	0.297	0.114	0.314	0.127
50%	0.599	0.982	0.595	0.981
best	0.516	0.418	0.509	0.505

may imply. It should be noted that it is not possible to directly compare the numbers of Table 3.1 with those of Table 3.2, because in CBP reconstruction the DC value of the data is removed by the filtering, causing the reconstructed image to be normalized differently than in the DF reconstruction. Also, the actual values of the statistics are not very meaningful, but the trend from one case to the other is quite significant. A study of Table 3.2 indicates that the added noise may be masking some of the image degradation that is a result of the limited data. In fact, this trick is sometimes used in image processing as a means of reducing the effects of data irregularities or processing inaccuracies. As in the DF reconstruction, the addition of noise does not seem to harm the effectiveness of the PSA in improving image quality.

The remarkable superiority of the initial convolutional-backprojection reconstruction over the initial direct Fourier reconstruction deserves some explanation. Part of the degradation in the DF reconstruction is undoubtedly a result of the polar-to-rectangular interpolation, which was carried out in these examples with a first-order 2D inverse distance algorithm. However, this is not the primary reason for the differences. Assume for the moment, implementations of both DF and CBP in which no processing errors occur, i.e., there are no interpolation, filtering or finite record errors. Then the process of calculating the inverse 2D-FT of the Fourier domain containing known data and zero padding as the initial guess, is equivalent to calculating the inverse 2D-FT of the entire correct Fourier transform, and then convolving this image with a distorting

filter function. This filter function is the impulse response of the Fourier domain function that assumes a value of unity in the known conical regions and zero outside. Clearly, the impulse response will be non-symmetrical and of significant relative spatial extent. The effect of convolving the correct image with this function is to seriously distort the image. This degradation is what is observed in the initial DF reconstructions. Consequently, the effect of making an initial guess of zero degrades the image in two ways: 1) it is clearly incorrect which must degrade the image, and 2) the discontinuities introduced into the Fourier domain enter the spatial domain by means of distorting convolving functions.

As the angular region of the missing cone approaches zero, the impulse response tends to an impulse function, which, when convolved with the image will cause no degradation. In the limiting case, the two reconstructions would be identical. The reason that CBP reconstruction is better is that only the lack of data harms the result. The presence of zero data over a region of backprojection (integration) and the necessary discontinuous edges do not affect the reconstruction.

Figures 3.20-3.23 show reconstructions using a combination of the PSA and AIM algorithms. The version of PSA employed is the convolutional-backprojection implementation. The combination PSA-AIM algorithm is a two-step technique in which one or more iterations of PSA provide an initial guess for the nested iterations that AIM performs on the periodic functions constructed from the projections. In the examples shown, one pass of PSA supplies the initial data for 10 iterations on each of the 32 periodic functions. These two steps are

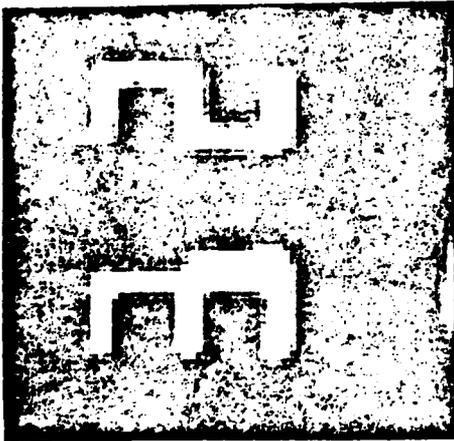


Fig. 3.20 a) Image #1, CBP with 100% data and no noise.



Fig. 3.20 b) Image #1, CBP with 50% data and no noise.

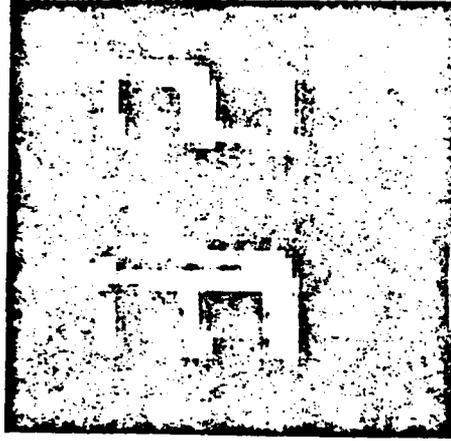


Fig. 3.20 c) Image #1, CBP with 50% data, no noise and 4x10 iterations of PSA and AIM.

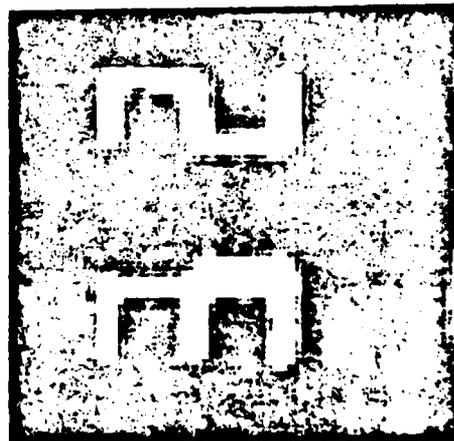


Fig. 3.21 a) Image #1, CBP with 100% data and Gaussian noise.

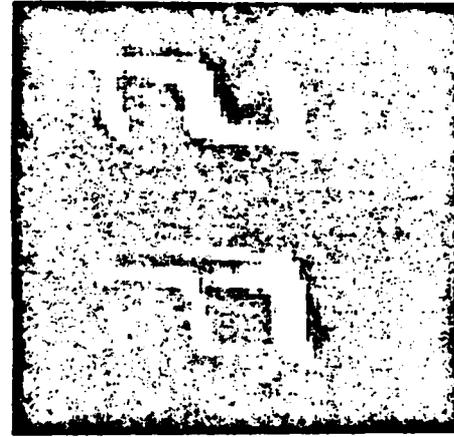


Fig. 3.21 b) Image #1, CBP with 50% data and Gaussian noise.

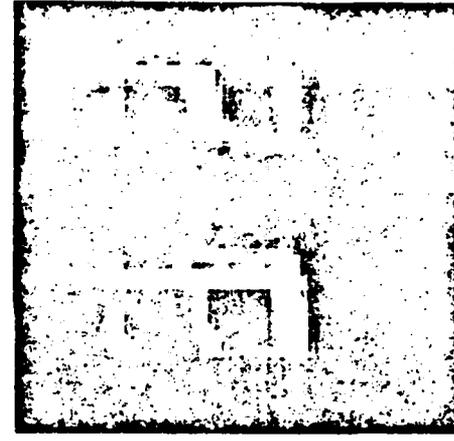


Fig. 3.21 c) Image #1, CBP with 50% data, Gaussian noise and 4x10 iterations of PSA and AIM.



Fig. 3.22 a) Image #2, CBP with 100% data and no noise.



Fig. 3.22 b) Image #2, CBP with 50% data and no noise.



Fig. 3.22 c) Image #2, CBP with 50% data, no noise and 4x10 iterations of PSA and AIM.



Fig. 3.23 a) Image #2, CBP with 100% data and Gaussian noise.



Fig. 3.23 b) Image #2, CBP with 50% data and Gaussian noise.

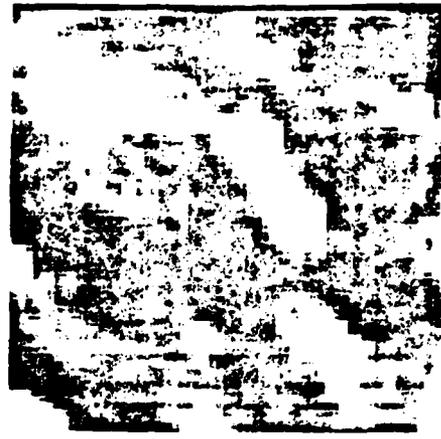


Fig. 3.23 c) Image #2, CBP with 50% data, Gaussian noise and 4x10 iterations of PSA and AIM.

repeated four times (thus the 4×10 notation in the captions). Figures 3.20 and 3.21 are the noise-free and noisy reconstructions for Image #1, while Figures 3.22 and 3.23 are the corresponding reconstructions for Image #2. The statistics for the PSA-AIM algorithm are supplied in Table 3.3. Although the statistical measures for AIM are slightly inferior to those for PSA in the noise-free case, it appears that AIM does resolve detail that is not present in the PSA results. It is not unexpected that AIM performs statistically better with noisy data, considering the intense filtering of the projections during the angular iteration process.

As a final example, PSA-AIM reconstructions with 35% and 65% data and a 20 dB SNR are presented. Figures 3.24a) and 3.26a) contain the initial reconstructions from 35% of the data and Figures 3.25a) and 3.27a) contain the 65% data case. The 'best' reconstructions are in Figures 3.24b), 3.25b), 3.26b) and 3.27b) for the 35% and 65% cases respectively. Table 3.3 includes the statistics for these examples. As in the 50% data case employing AIM, the statistical measures suggest greater improvement than observations would indicate. This is particularly true of Figures 3.25 and 3.27. The result in Figure 3.26 demonstrates how considerable detail can be recovered from a fairly narrow cone of data when the appropriate constraints are imposed.

In Section 3 some comments were made concerning the expected convergence behavior of these techniques. It was pointed out that these techniques generally appear to reduce the error for a few iterations and then diverge. An explanation for this is provided by generalizing the 1D analysis for the 2D case. That both the Projection-Slice algorithm

Table 3.3 Angular Iteration Method.

	<u>no noise</u>		<u>noise</u>	
	mean	variance	mean	variance
picture 1				
100%	0.144	0.110	0.152	0.121
65% initial best			0.409 0.252	0.729 0.380
50% initial best	0.530 0.336	0.935 0.679	0.523 0.336	0.930 0.669
35% initial best			0.649 0.399	1.111 0.968
picture 2				
100%	0.297	0.114	0.314	0.127
65% initial best			0.481 0.446	0.442 0.247
50% initial best	0.599 0.513	0.982 0.422	0.595 0.527	0.981 0.416
35% initial best			0.693 0.615	1.705 0.906

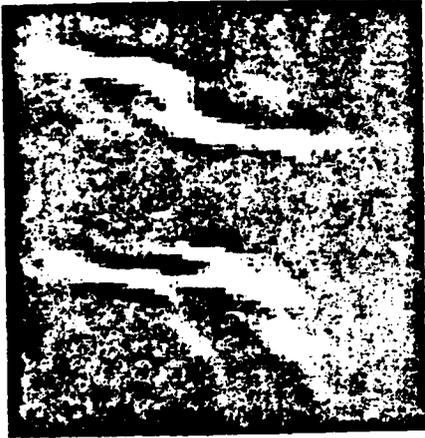


Fig. 3.24 a) Image #1, CBP with 35% data and Gaussian noise.

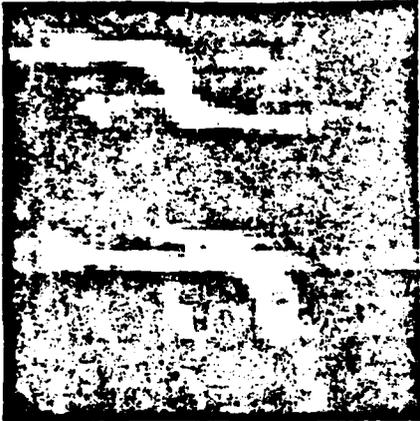


Fig. 3.24 b) Image #1, CBP with 35% data, Gaussian noise and 4x10 iterations of PSA and AIM.

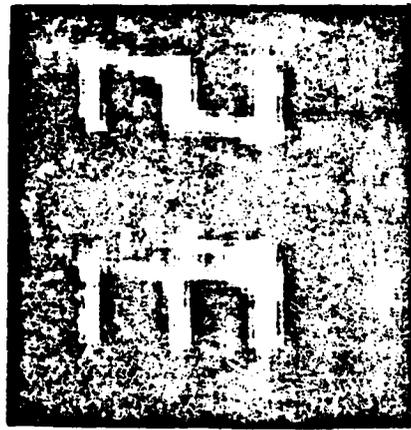


Fig. 3.25 a) Image #1, CBP with 62.5% data and Gaussian noise.

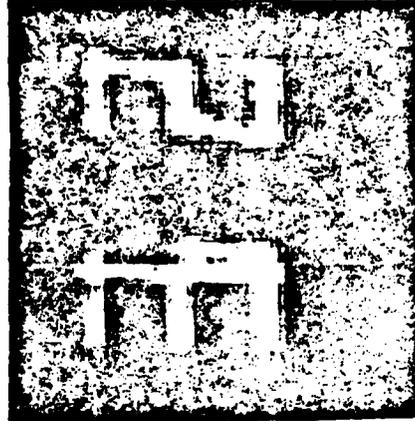


Fig. 3.25 b) Image #1, CBP with 62.5% data, Gaussian noise and 4x10 iterations of PSA and AIM.



Fig. 3.26 a) Image #2, CBP with 35% data and Gaussian noise.



Fig. 3.26 b) Image #2, CBP with 35% data, Gaussian noise and 4x10 iterations of PSA and AIM.

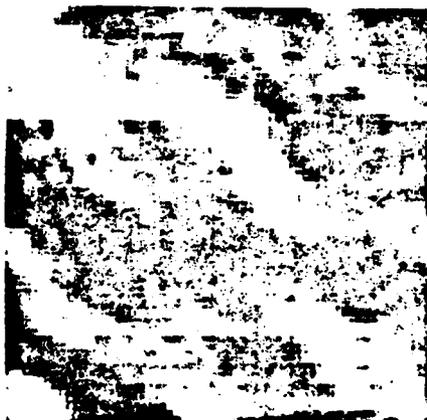


Fig. 3.27 a) Image #2, CBP with 62.5% data and Gaussian noise.



Fig. 3.27 b) Image #2, CBP with 62.5% data, Gaussian noise and 4x10 iterations of PSA and AIM.

(DF or CBP implementation) and AIM demonstrated similar tendencies in all the examples was not, therefore, unexpected. The cause of this phenomenon is the inability to implement ideal filtering (in either the spatial or Fourier domains). The introduction of more error with each iteration and the compounding effects of recursively processing this error eventually overcomes the converging nature of these algorithms and causes divergence. Iterative techniques may be less sensitive to noisy data than non-iterative methods. However, any noise (or error) remaining after one iteration will still be present at the start of the next iteration and add to the error generated by this pass, thus increasing the total error in the data. Consequently, convergence and solution quality are related to initial quantities of error and the relative rates at which these may be reduced and generated.

Some final comments can be made concerning processing times. The Projection-Slice algorithm requires approximately 90 seconds per iteration when implemented with direct Fourier reconstruction, and about 240 seconds per iteration when convolutional-backprojection reconstruction is used. For AIM, approximately 360 seconds are required for each iteration. This time includes the 10 iterations on each of the 32 $s(r, \theta)$ functions. The actual quoted times are not very significant because they can vary greatly depending upon the computer used and operating environment of that computer system. These times are for a VAX 780 under a UNIX operating system. The important point of relative time measurements is that they supply a good indication of the comparative cost of these algorithms.

3.5 Conclusion

In this chapter, the missing cone problem in computer-aided tomography has been discussed. a number of state-of-the-art techniques for dealing with this problem were reviewed, and two new experimental algorithms. the projection-slice algorithm (PSA) and the angular iteration method (AIM), were presented and illustrated with examples. The goal of these algorithms is to produce a higher quality image than can be obtained by directly reconstructing from the limited data. The algorithms achieve this improvement by combining two related concepts: 1) spectral and/or spatial domain extrapolation techniques and 2) the inclusion of a priori information. The PSA and AIM algorithms specialize the Gerchberg-Papoulis iterative extrapolation techniques by incorporating characteristics of the projection data into the algorithms. That this enhancement requires computations on the order needed for the original reconstructions is significant, particularly when compared to the computational requirements of 2D maximum entropy methods.

Throughout this chapter, efforts are made to present sufficient detail to allow the reader to implement these techniques in practical problems. One issue that has not been discussed is a convergence criterion for indicating when the iterative process should stop. It has been observed that, in virtually all the experiments carried out with finite record length processing, the iterative procedure tends to improve the image quality to a point, after which the algorithm begins to diverge from the best solution. This phenomenon is a function of the block lengths used in the FFT computations, and to some extent on the

sampling rate of the observed data. Therefore, it is important to monitor convergence closely, and to stop the iterations when the minimum error solution is obtained.

The crudest but perhaps the most widely applied method of monitoring convergence is to visually inspect each iterative result, and choose the most appealing one (in some ad hoc way). A second, more quantitative method involves monitoring the statistics for a minimum point. There are some arbitrary decisions to be made about this technique, e.g., which statistics to use and what relative weights should be assigned to these measures. It was observed in these experiments that a different number of iterations are performed to obtain the best solution, depending on whether the mean or variance is considered more important. A third technique is to compare the original data to the calculated data and iterate until they agree within a prescribed tolerance. For example, the original projections can be compared to the synthetically calculated projections over the angular interval where these are both known.

Although the missing cone problem was discussed in this chapter within the context of computer-aided tomography, it has become increasingly apparent that many other inverse problems in the fields of synthetic aperture radar, beamforming sonar, radio astronomy, electron microscopy and geophysical exploration can benefit from new solutions and better algorithms to deal with regions of missing data. Therefore, the missing cone problem represents an important generic problem which will very likely receive increased attention in future years.

4. ANALYSIS OF A JITTER MODEL FOR COORDINATE TRANSFORMATION IN SYNTHETIC APERTURE RADAR

In the generation of spotlight mode synthetic aperture radar (SAR) images from digitally recorded data, one of the most computationally demanding tasks is the two-dimensional interpolation from a polar raster to a rectangular raster [48]. This chapter analyzes a simple interpolation scheme that takes advantage of the significant oversampling of data in the azimuth direction and a 'smart' A/D converter (sampler). By 'smart' it is meant that the sampler can perform at varying and controllable rates and that these rates can be altered dynamically. The interpolation scheme proposed significantly reduces the computational requirements of a digital SAR processor [3,49,50].

This chapter is organized in the following manner. First, the basic spotlight mode SAR geometry relevant to the sampling issues involved will be presented and discussed. At this point, sufficient background material and terminology will have been covered to allow a discussion of other interpolation schemes proposed for this problem. Next, the jitter model for nearest-neighbor interpolation is presented, analyzed and discussed. Lastly, some computer experiments are presented to support the theoretical results and to illustrate the affects of nearest-neighbor interpolation.

4.1 Problem Description and Background

The geometry for data collection in a typical spotlight mode synthetic aperture radar system is illustrated in Figure 4.1. As the vehicle carrying the radar passes to the side of the terrain to be mapped, a large aperture antenna is steered so as to keep the beam of the antenna pointed at some reference point in the terrain. The angle over which the antenna illuminates the terrain is termed the look angle, α . At equidistant intervals along the flight path (assumed to be a straight line), the radar set transmits a linear FM wavepulse with center frequency f_0 . The form of this signal is:

$$f(t) = \begin{cases} \cos(2\pi f_0 t + \gamma t^2/2) & \text{for } |t| \leq T/2 \\ 0 & \text{for } |t| > T/2 \end{cases} \quad (4.1)$$

where γ is the FM rate. The value of T (pulse duration) is determined by ambiguity and resolution requirements. Clearly, T also determines the bandwidth swept by the transmitter. The received signal, which is the transmitted signal convolved with the complex reflectivity of the scene, is mixed with quadrature reference signals [3] and demodulated, bringing the collected information down to baseband. Consider the two demodulated quadrature channels to represent the real and imaginary components of a complex signal. With this interpretation, the data can be shown to be a portion of the 2D FT of the complex reflectivity of the target [48].

Refer once again to the SAR geometry illustrated in Figure 4.1. Due to the finite look angle of the data collection scheme, 2D Fourier domain data are recorded only over a finite angular region. Further,

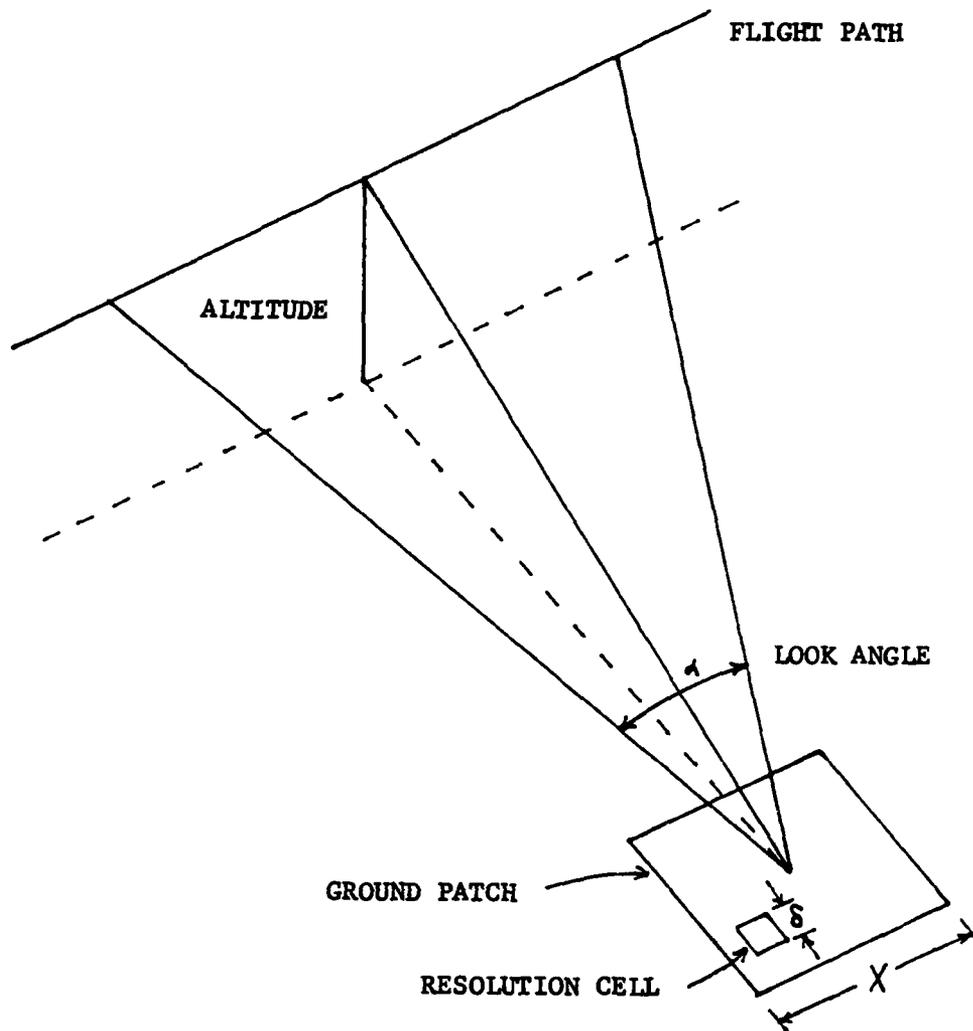


Figure 4.1 SAR data collection geometry.

the bandwidth (due to the finite T) of the transmitted pulse determines the span of frequencies for which data are available in the range direction. The combination of these effects is to constrain samples to a conical region as illustrated in Figure 4.2.

The center of this conical region is at a spatial frequency W_0 , which is equal to $4\pi f_0/C$ where C is the speed of light. Minimum and maximum frequencies in the range direction are determined by the desired resolution, thus affecting T . Physical design constraints on the bandwidth of the transmitter and receiver may be the determining factor, thus determining T and the resolution. In practice this conic section is a relatively thin strip, the ratio of azimuth bandwidth to range bandwidth is often greater than 5:1 [51]. Since the data are recorded in the Fourier domain, a 2D FT must be performed in order to generate an image.

One very fast method for transforming from the Fourier domain to the spatial domain is by the use of lenses. Since Fourier transforms can be calculated optically with lenses, the FT inversion can be accomplished almost instantaneously. Additionally, that the data occupy a conical region is of little significance to optical processors. It is easy to compensate for this data format by the proper design of the optics and recording techniques. Harger [3] presents a theoretical discussion of optical methods. In [51,52] some hardware details are presented that illustrate this type of SAR processing.

Because the exposing of film and the later development requires an off-line and time-consuming process, a real-time processor incorporating this technique is not feasible. Since the motivation of this work is to

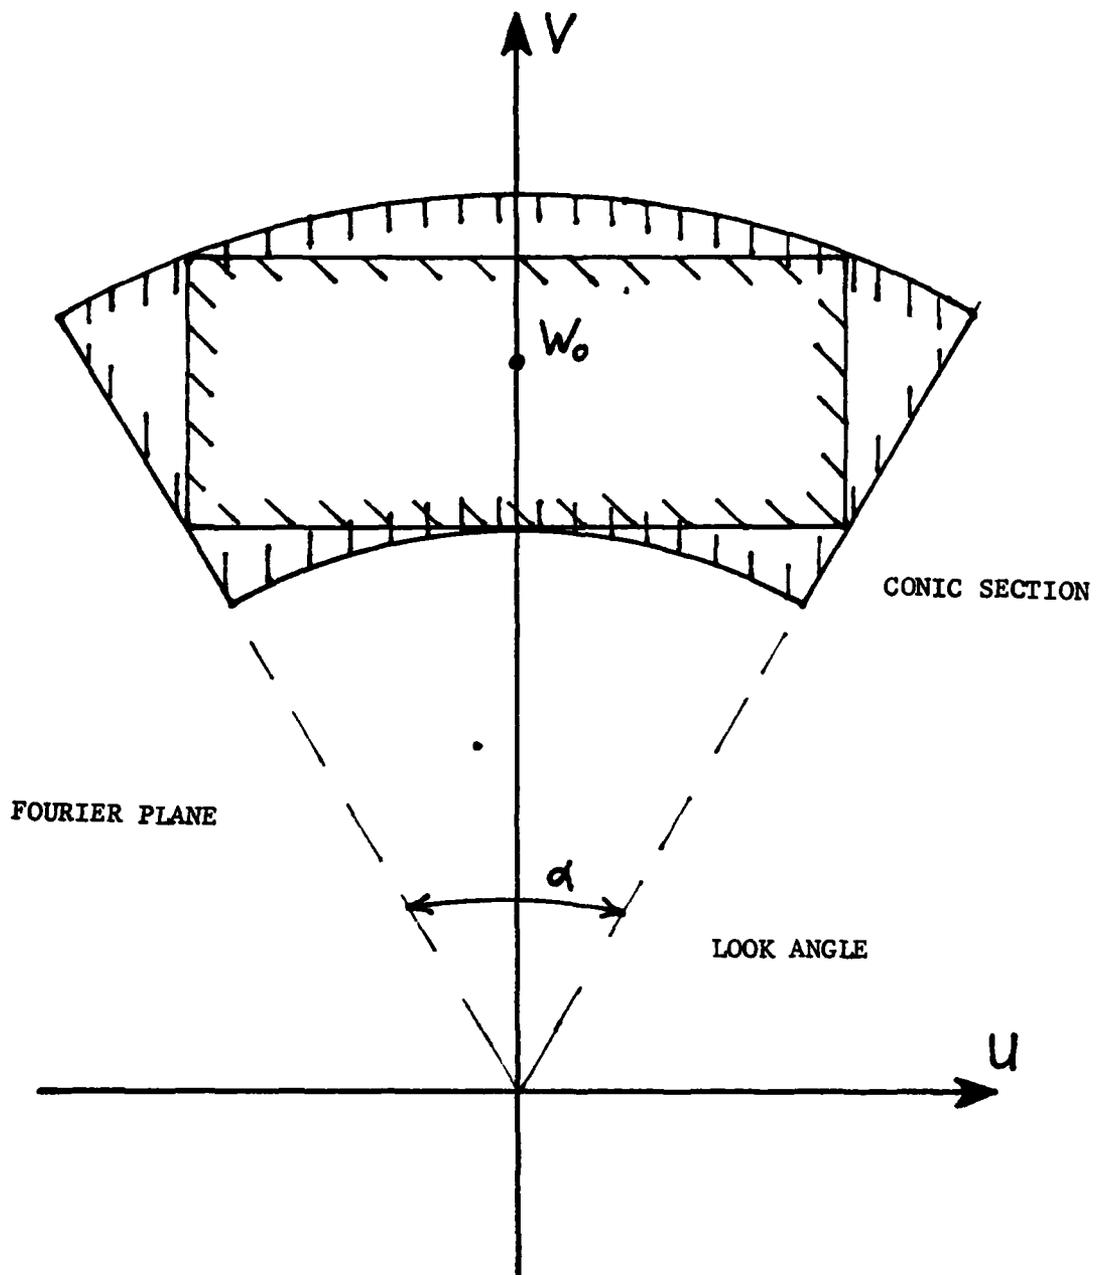


Figure 4.2 Format of collected SAR data.

develop techniques for real-time SAR processing, further attention is centered around digital techniques for SAR imaging. From the above discussion, an important process in generating the image is the transformation from the 2D Fourier domain to the 2D spatial domain. One efficient method of accomplishing this transformation is to employ FFT's (or other fast realizations of the DFT). Further, in order to employ digital techniques the data must be sampled. If the two data channels representing the real and imaginary parts of the complex Fourier domain data are sampled at uniform rates, then samples of the 2D FT are obtained on a polar raster [53]. Since FFT's and most similar techniques require samples on a rectilinear raster, some sort of interpolation is required to change the format of the collected data. The region of interpolation is typically taken to be the largest rectangular region that can be contained wholly in the conic region over which data are available. This is illustrated in Figure 4.2. In general, 2D interpolation is computationally expensive, and in SAR processing it is a primary data processing bottleneck. The aim of this chapter is to investigate an efficient method for simplifying the interpolation problem. Next, a review of present methods as well as some recently proposed techniques will be discussed.

One of the more obvious methods for transforming from a polar to a rectangular format is to perform a 2D interpolation. A fairly simple algorithm would employ a first-order inverse-distance technique. In [53], Schwartz discusses this technique and points out that sensitivity to noise is a major drawback. Further, since this technique must employ spatially varying coefficients, a constant coefficient FIR filter

implementation is not possible. As a result, this simple interpolation scheme is computationally expensive. As illustrated in Chapter 3, polar-to-rectangular interpolation is both ill-posed and, for spatially varying interpolation points, quite expensive. Consequently, significant effort has been spent attempting to circumvent one or both of these problems (ill-posedness and/or spatial variation).

In perhaps the simplest technique for polar-to-rectangular interpolation, the polar raster is assumed to closely approximate a rectangular raster. Interpolation to a rectangular grid is performed by selecting the polar sample nearest the rectangular point. In this method, called 'nearest-neighbor' interpolation, the only calculations involved are those for determining the nearest neighbor. Further, if the look angle of the polar grid is small, then the polar raster is a good approximation to a rectangular grid and the error introduced by nearest-neighbor interpolation is quite small. As pointed out in [53], this technique is rather insensitive to noisy data. However, if the look angle is not small, this technique can cause severe misregistration of targets, as well as significant loss in resolution.

A method that avoids the spatially varying problem is to interpolate from the available polar grid to a finer polar grid. Since the new grid points are regularly spaced with reference to the original raster, a FIR filter with constant coefficients can be designed to perform the interpolation. While the resulting polar grid points are not on a rectangular raster, on the average, these polar points lie closer to a rectangular grid point than prior to interpolation, and the error induced by performing a nearest-neighbor interpolation will be

smaller than with the original data. Mersereau and Oppenheim discuss this idea in [4], where they design a specific interpolator to optimize some feature in the final image. In [25], Stark et al. proposes an exact scheme for interpolating to a finer polar grid. This technique assumes two facts: 1) Data are band-limited in angle for each circular arc and 2) A complete polar grid of data are available. While the first requirement may be essentially met (as discussed in Chapter 3), the second requirement is clearly not satisfied by the conic region.

A different technique employed for polar to rectangular interpolation is to perform two 1D interpolation steps. In one realization, interpolation is performed along radial lines to obtain samples on a keystone format (see Figure 4.3) followed by interpolation along each horizontal line to the desired rectangular grid point. While this technique requires spatially varying coefficients, the computations involved are generally fewer than required by a non-separable 2D spatial varying interpolator. In some implementations, the interpolation is only performed in one direction with nearest-neighbor interpolation performed in the other direction.

So far in this discussion it has been assumed that the data are originally sampled on a polar raster. In the technique to be discussed next, the data are either sampled on a keystone raster or interpolation has been performed to transform the data to a keystone raster. From a practical standpoint, sampling data on a keystone raster is not difficult. With data in a keystone format, either 1D or nearest-neighbor interpolation can be performed to obtain the data on a rectangular raster. A significant, but rarely exploited aspect of SAR

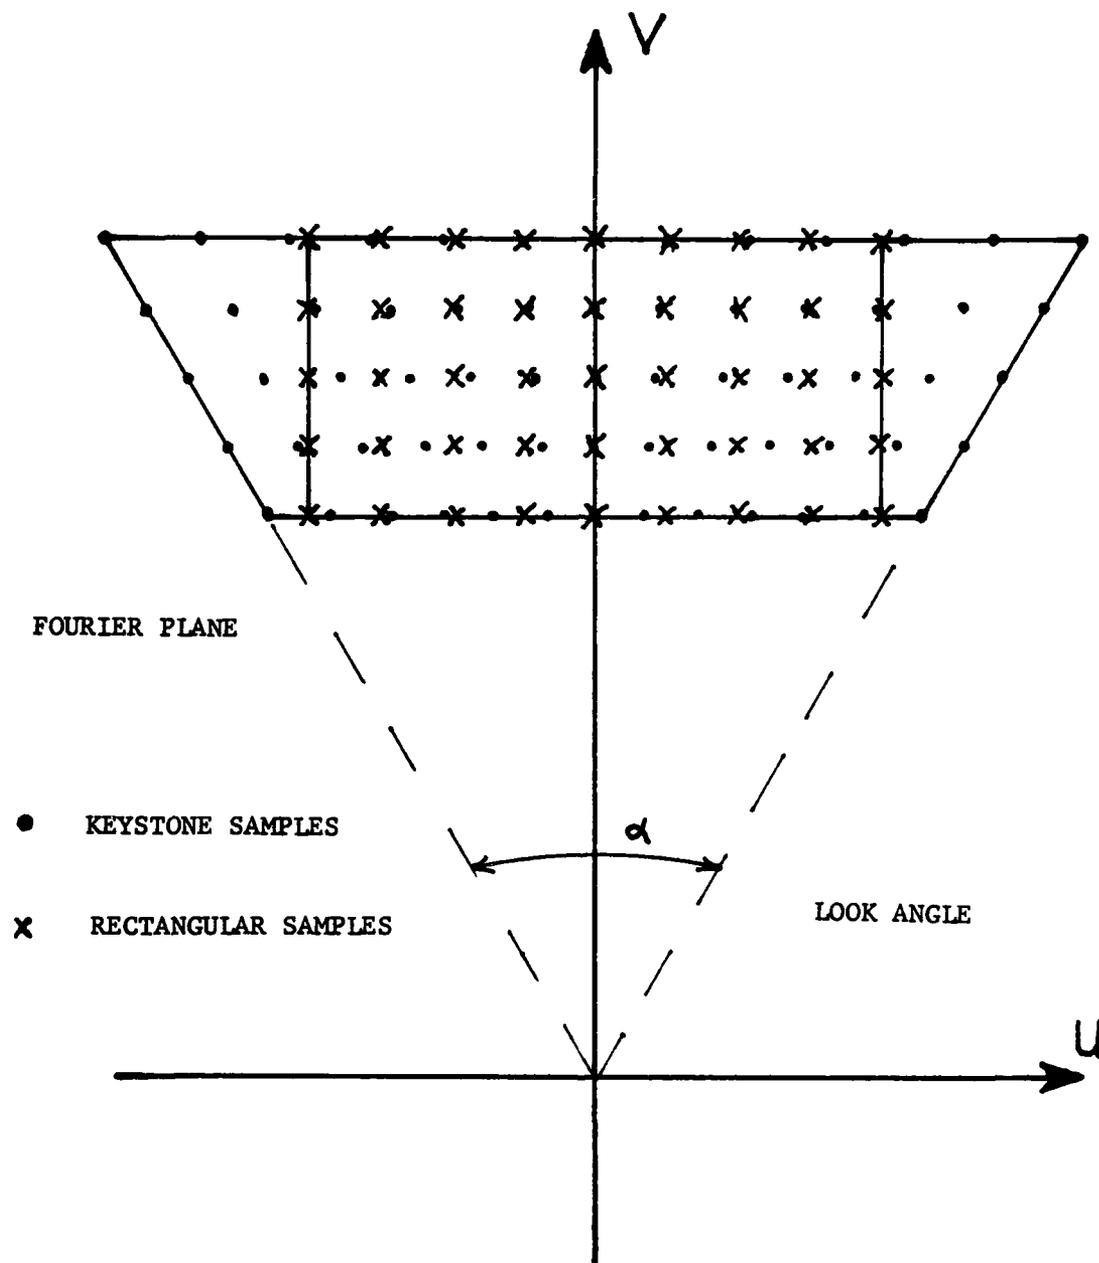


Figure 4.3 Keystone raster and interpolation to rectangular raster.

is that the azimuthal data are usually highly oversampled, typically by a factor of 5 or 10 to 1 [50]. This oversampling is present independent of whether the data are on a polar or keystone raster and is in addition to the wider azimuth bandwidth. The oversampling is illustrated in Figure 4.4 a) and 4.4 b).

Most SAR processors perform a presumming or decimation operation on the azimuthal data, during real time collection, in order to reduce the bulk storage requirements to more manageable levels [50]. In either case, only the resulting data are available for later use and the large azimuth to range bandwidth ratio remains. The presummer usually takes the form of a simple adder (averager). In other cases, the data values are scaled by a window, Hamming for instance, prior to the presum operation in order to achieve some desired results. The computations required to perform this windowing often restrict its use. What is proposed is an adaptive presummer that sums groups of values (along the range bin in question) around the rectangular grid points to which a nearest-neighbor interpolation step is going to be performed. This is illustrated in Figure 4.4 a). Since the oversampling is usually quite high, an original keystone sample will be available quite near a rectangular grid point. In the adaptive presummer case, the maximum position error between rectangular and keystone samples is one-half the keystone sampling interval. Employing the normal method of summing each group of Q points as they arrive will generate a maximum position error of one-half the rectangular sampling interval. The obvious result of adaptive presumming is a reduction position error by a factor of Q . This technique can also be employed in the polar-to-rectangular case

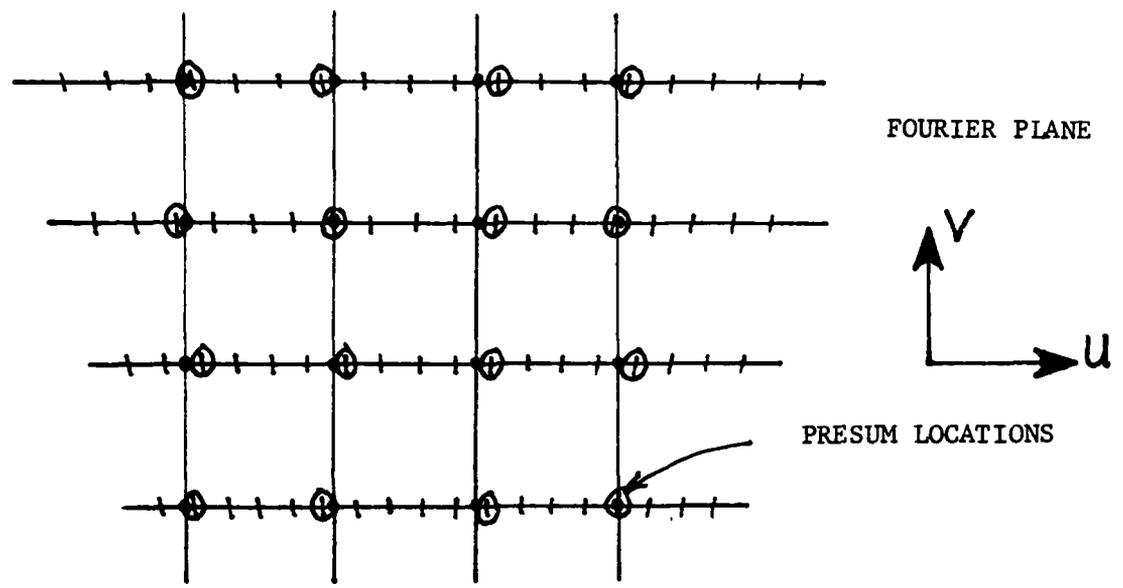


Figure 4.4 a) Oversampling and adaptive presumming in keystone-to-rectangular nearest-neighbor interpolation.

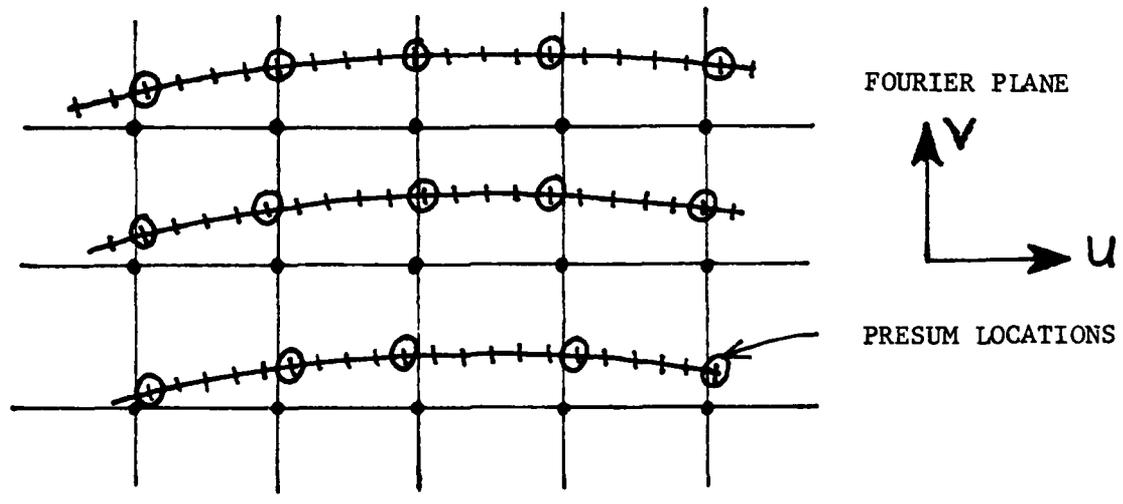


Figure 4.4 b) Oversampling and adaptive presumming in polar-to-rectangular nearest-neighbor interpolation.

where the adaptive presumming is performed along the circular arcs in a manner to generate samples nearest the rectangular grid points. See Figure 4.4 b). Again, if the included angle (look angle) is small, then the error due to nearest-neighbor interpolation may be negligible.

4.2 Jitter Analysis

In this section, a model is derived to characterize the nearest-neighbor interpolation scheme. This error will be referred to as jitter noise in the sense that the interpolated value can be considered to be the exact value with some noise signal modifying the value. As Schwartz pointed out in [53], nearest-neighbor interpolation is fairly insensitive to noise. Therefore, this technique should be fairly robust with respect to noisy data. The purpose of this analysis is to characterize the jitter induced noise and derive a model that accurately reflects the effects of nearest-neighbor interpolation in the keystone case.

Since the jittering is occurring in only one dimension, the following analysis will be for only one dimension, i.e., for a specific range bin. Denote the jittered samples of the RT as $R^J(\omega_k)$ where ω_k is discrete frequency. The jittered samples can be considered to be the correct rectangular samples, $R(\omega_k)$, modified by some additional signal. This modifying signal is assumed to be due to the small error in position between the rectangular and keystone samples. Using Δ to represent the error in position or jitter distance, then,

$$R^J(\omega_k) \approx R(\omega_k) + R'(\omega_k)\Delta\omega_k, \quad (4.2)$$

where $R'(\omega_k)$ is some measure of the slope over the jitter distance. What (4.2) says is that the jittered signal, or noise, is approximately equal to some measure of the slope (to be jittered) times the jitter distance.

One of the purposes of this analysis is to derive a signal independent model for characterizing some aspects of the jitter noise. In the case of trying to derive a signal independent relationship, the best that could be hoped for is some sort of bound on the error. Bounding $R'(\omega_k)$ by the maximum value of the slope, a bound on the maximum jitter noise amplitude, $N^J(\omega_k)$, is

$$N^J(\omega_k) = \max[R'(\omega_k)]\max[\Delta\omega_k]. \quad (4.3)$$

Substituting (4.3) into (4.2),

$$R^J(\omega_k) \approx R(\omega) + \max[R'(\omega_k)]\max[\Delta\omega_k], \quad (4.4)$$

where $\max[\Delta\omega_k]$ represents the largest possible jitter distance, one-half the distance between keystone samples. A better model that may not strictly bound the jitter error, but will be more representative of the signal value, is obtained by substituting the average jitter distance for the maximum jitter distance. In this case, (4.3) becomes

$$N^J(\omega_k) = \max[R'(\omega_k)]E\{\Delta\omega_k\}. \quad (4.5)$$

Employing the above analysis, the derived model will be composed of two factors. The first factor is the expected jitter (interpolation) distance as a function of the oversampling rate and other terms. Once this value is calculated, if a bound can be placed on some measure of the rate of change in the sampled Fourier domain, then an equation can

be derived that models the variation of this measure of the jittered signal. First, a model for the expected jitter distance is derived.

Referring to Figure 4.3, let M represent the total number of keystone samples on each horizontal line. This value of M is equal to QN , i.e., the product of the number of rectangular samples times the oversampling rate, Q . The term ξr denotes the sampling interval in the vertical direction and α is the look angle. With these definitions, the separation between keystone samples is

$$S = \frac{2n\xi r \tan(\alpha/2)}{M}, \quad (4.6)$$

where n represents the range bin (a horizontal line of data). Assuming that the distance between rectangular grid points is equal, then any given rectangular point is uniformly distributed in the interval between two keystone points. With this assumption, the average jitter distance J_d as a function of range bin, n , is

$$J_d = \frac{n\xi r \tan(\alpha/2)}{2M}. \quad (4.7)$$

This equation assumes a constant azimuth sampling interval for a given range bin. This assumption is not strictly correct. Equation (4.7) is derived from the fact that the jitter distance is one-half the separation of the keystone samples, and on the average, the misalignment will be the keystone separation divided by four. Equation 4.7 can easily take into account oversampling that may be present in a specific system.

In reference to (4.2), two different measurable quantities are apparent. Either the magnitude or the phase could be employed as a measure of the signal variation. The measure proposed in this chapter

is the phase of the complex signals. The reasons for using the phase as a measure follow. One interpretation of the source for the high resolution that SAR systems apparently provide is that the phase is the single most important factor. In work indirectly related to SAR, Oppenheim and Lim [54] have demonstrated the importance of phase in reconstructing images. In another interpretation, the phase component in SAR data can be considered to represent the relative times at which reflected radar pulses are received. These relative receive times represent the spatial distance of the targets to the transmitter. By correlating the reflections over the look angle, the specific targets are resolved; the relative signal magnitudes are not critical to the resolution. Employing the above reasoning, phase can be seen to be the critical factor in SAR imagery. Mathematically, this can be presented as a stationary phase approximation [26]. If a complex function has a fast varying phase and a slowly varying magnitude, then the integral (Fourier transform) of this function is primarily due to the effects of the phase. In SAR systems, the phase of the complex data is in fact varying rapidly with reference to the magnitude; consequently, a stationary phase approximation would indicate that the phase of this signal is the important quantity for image generation. For these reasons, a measure of phase more accurately reflects the effects of jitter than a measure of magnitude.

In order to derive bounds on the rate of change in the phase of SAR data, the effects of sampling on the complex data have to be understood first. To simplify matters, sampling theory will be discussed with reference to a one-dimensional signal; the results clearly apply to the

2D case. Assume that the desired resolution in the spatial domain is δ (meters). In order to resolve this cell in the spatial domain, the Fourier domain must have a spectral bandwidth of at least $4\pi/\delta$ (radians/meters). For a map size of X (meters), sampling theory states that the Fourier domain must be sampled at a minimum rate of $\pi/2X$. These two relationships imply that $4X/\delta$ complex Fourier domain samples are required. Since in the examples to be considered here the FFT length will be the dominating factor (fixed), the above relationship, i.e.,

$$N \geq 4X/\delta \quad (4.8)$$

establishes either the resolution or the image size given the other. In practice, the factor limiting the image size is often the beamwidth of the antenna used to illuminate the scene. Using the above relationships, a bound can be placed on the maximum rate of change (the slope) of the phase. Consider a target that is at the extreme edge of the scene. This target can be considered to be an object at the center of the scene that has been phase shifted (by multiplication with the complex exponential $e^{j\pi X/2}$) to the edge of the scene. The phase of this complex exponential generates the greatest slope in the phase of the complex data. Ignoring the effects of the complex reflectivity of the target, uncompensated motion of the aircraft and other factors that influence the recorded data, and thus the phase, then the phase will change by a factor of π over a period corresponding to the highest frequency of the sampled data. Since δ and X are constrained by N and sampling theory, it is clear that the highest frequency, represented by X , is sufficiently sampled. Therefore, at the minimum sampling rate

specified by (4.8), the phase is sampled twice over the interval in which it changes by π . With the sampling rate of the Fourier domain set at $2\pi/N\delta$, the minimum period in the Fourier domain is $4\pi/N\delta$. The maximum average slope of the phase is then

$$\phi_{\max} = \left(\frac{\pi}{1} \right) / \left(\frac{2\pi}{N\delta} \right) = \frac{N\delta}{2}. \quad (4.9)$$

Using the equation for mean jitter distance, (4.7), a relationship for the mean phase jitter can be obtained. Since the jitter distance is a function of the range bin, denoted by n in equation (4.7), assume that the calculated value for the jitter at W_0 (represented by some value of n , say n_1) is a fair value for the mean phase jitter over all the range bins. Then

$$E\{\phi_j\} = \frac{N\delta}{2} \left[\frac{W_0 \tan(\alpha/2)}{(QN-1)} \right] \approx \frac{\delta W_0 \tan(\alpha/2)}{2Q} \quad (4.10)$$

where

$$\frac{4\pi f_0}{c} = W_0 = \frac{n_1 4\pi}{\delta}. \quad (4.11)$$

In the simulations employed later in this chapter, the complex array is assumed to be square, not rectangular. As a consequence, the range bandwidth determines the azimuth bandwidth. Further, the actual look angle of perhaps 10 degrees is dominated by the modified (imposed) azimuth bandwidth which corresponds to a look angle of approximately one degree. For this case, the jitter distance can be calculated to be

$$J_d = \left[\frac{W_0}{W_{\min}} \right] \left[\frac{2\pi}{\delta(QN-1)} \right] \quad (4.12)$$

where

$$\frac{W_0}{W_{\min}} = \frac{2\delta f_0}{2\delta f_0 - C} \quad (4.13)$$

Thus the maximum phase jitter error is approximately

$$E\{\phi_j\} = \left[\frac{2\delta f_0}{2\delta f_0 - C} \right] \left[\frac{2\pi}{\delta(QN-1)} \right] \frac{N\delta}{2} \approx \frac{2\pi\delta f_0}{Q(2\delta f_0 - C)} \quad (4.14)$$

If a value of $\tan(\alpha/2)$ corresponding to the modelled azimuth bandwidth is employed in (4.10), then equation (4.14) is obtained.

Two modifications need to be made to equations (4.10) and (4.14). In the analysis it was assumed that the final square rectangular raster, after interpolation, barely met the necessary sampling requirements. However, in many cases the final Fourier data will be oversampled. The original analysis showed that as a consequence of barely meeting the sampling requirement, the phase changed by π over the sampling interval. If, in fact the final Fourier raster does oversample the data by a factor of K , then the phase change will be π/K between samples. The second factor that needs to be considered is that the above analysis considered the largest phase shift to be due to the largest spatial offset. In the assumed square region under consideration, this phase is a factor of 1.414 larger (because of the diagonal distance) than originally discussed. These effects modify (4.10) and (4.14) to:

$$E\{\phi_j\} = 0.707\delta W_0 \frac{\tan(\alpha/2)}{KQ} \quad (4.10a)$$

and

$$E\{\phi_j\} = \frac{2.828\delta\pi f_0}{KQ(2\delta f_0 - C)} \quad (4.14a)$$

In the next section, the validity of the jitter distance model will be

verified and relationships between the theoretical jitter and measured phase jitter will be examined and discussed. Further, a simulation is provided to demonstrate the affects of nearest-neighbor interpolation.

4.3 Experimental Results

Three experiments are discussed in this section. The first is a program that generates a keystone raster and a rectangular raster and calculates the average jitter distance introduced by nearest-neighbor interpolation. Results of this experiment will be presented first to verify (4.7). In the second experiment, data are generated in a keystone format and nearest-neighbor interpolation is used to transform this data set to a rectangular raster. The jittered keystone values, now in a rectangular raster, are compared to the theoretical values of the 2D FT obtained by sampling on a rectangular raster. Average phase jitter is calculated and compared to the theoretical values. Further, the image generated by the jittered keystone values is compared to the image obtained from the correct rectangularly sampled data. The expected result of higher oversampling rates generating better images will be verified. The third experiment attempts to generate the jittered FD data by employing the theoretical model for jitter, equation (4.2). Correct data are corrupted by the use of noise as predicted in equation (4.2). Images generated by this experiment are compared to those generated from the real jittered data. The purpose is to further verify the validity of the jitter model for nearest-neighbor interpolation. A result of these experiments is to demonstrate the

applicability of nearest-neighbor interpolation when there exists significant oversampling and adaptive presumming is employed.

Table 4.1 contains some results for the theoretical jitter distance mean, equation (4.7) and the actual measured mean found by computer simulation. It can be seen that the model provides an excellent approximation to the measured values. Further, the model appears to provide a very good estimate of the jitter error for nearly any combination of parameters. Other experiments have shown that even at large look angles the model and measured values differ by at most approximately one percent.

Table 4.1 Measured and Theoretical Jitter Distance.

Look Angle	Eq. (4.3)	Simulation	Eq. (4.3)	Simulation
	Q=1	Q=1	Q=7	Q=7
20°	0.1106	0.1106	0.0158	0.0157
30°	0.1681	0.1676	0.0240	0.0238
40°	0.2284	0.2276	0.0326	0.0323
50°	0.2926	0.2916	0.0418	0.0413
60°	0.3623	0.3610	0.0518	0.0512

As a preliminary step to discussing the results of the SAR simulation, some of the implementation details need to be described. The model employed in this work starts after the mixing and demodulation stage and just prior to sampling. Starting at this point allows a flexible sampling scheme, keystone or rectangular for instance. Since the two data channels available after the demodulation are assumed to represent the real and imaginary components of the 2D FT of the scene, a

method is required for generating these data. Two significantly different methods were considered. In the first technique, the 2D DFT (computed with FFT's) is calculated for some arbitrary image and samples of this data set are employed in the simulation. There are two major problems with employing this model. Both of these problems arise from the desired ability to obtain samples over the conic region for arbitrary frequency offsets. In order to simulate relatively large frequency offsets, the originally calculated DFT must be orders of magnitude more dense than the sampled raster. This is because the DFT is constrained to lie in a 0 to 2π region and in order to accommodate the sampling requirements around the offset mapped (by the sampling rate) into 0 to 2π , the data will have to be very dense. This high density implies that very large FFT's will have to be used in the calculation of the data set in order to accurately support a sampling of this data set over a small subset of the 0 to 2π region. A further problem is that using FFT's implies that samples are available only on a rectilinear raster. Therefore, if samples are required on a keystone or polar grid, interpolation will be necessary. Since the objective is to obtain these original sample values with as little error as possible, a very accurate interpolator would be required.

In the second method, an image is constructed from a set of rectangles. These rectangles can be of arbitrary size and can be positioned anywhere in the image scene. Since the 2D FT of a rectangle can be calculated exactly (the product of two sinc waveforms) and by introducing a complex phase, e^{jx} , the exact 2D FT of an arbitrarily located rectangle can be expressed in a closed-form expression. By

superposition, any combination of shifted rectangles can be used to construct the image. The final result is a sum of sinc waveforms, each multiplied by a complex phase, representing the relative spatial offsets. Because this model is continuous, it is not biased towards any specific coordinate structure. Given an arbitrary position in the w_1, w_2 plane, the exact complex value of the 2D FT of the image can be calculated for that point. Thus by designing an appropriate method for generating w_1, w_2 values, exact samples of the 2D FT of the image can be obtained on any desired raster. This is the model employed in the following simulations. This model also accurately reflects the fact that prior to sampling, the two demodulated data channels contain continuous data. The major disadvantages of this technique are: 1) The class of images is restricted to those that can be constructed from rectangles. Admittedly, any image could be represented with arbitrary accuracy in this manner, but the large number of independent sinc waveforms needed to represent this image may be prohibitive and 2) The computations required to evaluate this sum of phase shifted sinc waveforms for each sample point can become quite large.

In order to calculate the phase jitter, sample values of the 2D FT of an image are obtained on a keystone format (to some specified oversampling rate) and sample values of the 2D FT of the image are also calculated for the rectangular raster. Then after finding the keystone sample that is the nearest-neighbor to a specific rectangular point, the phases of each are calculated and the difference is stored. This is performed for each point in the rectangular grid and the average value of all the stored phase errors is calculated. This value is displayed

in Table 4.2 for various keystone oversampling rates and as a function of different required resolutions and, because the FFT length is not altered, different rectangular sampling rates. Also, in Table 4.2 are the theoretical phase jitter values calculated from equation (4.10a). The theoretical values are for the center range bin only, not the average over all range bins. The error introduced by this assumption is quite small. In Figure 4.5 is the image employed in this analysis.

Inspection of these results indicates that the theoretical value tends to bound, but not strictly, the average phase jitter error. Reasons why this bound is not strict follow. Since the jitter distance model is quite accurate, the error must be related to the bounds on the phase slope. Although as pointed out earlier, using the average jitter distance is also a cause of error, at least in terms of strictly bounding the error. The primary source of error in the phase slope model is the assumption that the maximum phase slope can be approximated by $N\delta/2$. To obtain this result it is assumed that the phase changes linearly over the sampling interval. Clearly, this will not in general be true. The true maximum phase slope will be modulated by a function of the data. As an example, in the image used in this example, if the rectangle at $(x,y) = (13,13)$ is removed, then the measured phase jitter for $\delta=10$ meters and a keystone oversampling rate of 5 is 0.474. A small change in the image induces a large change in the measured statistics of the image. Since the model doesn't account for specific features in the image, the predicted phase jitter error is constant while the measured values change, sometimes quite substantially, as noted above. However, equation (4.10a) does satisfy the objective of

Table 4.2 Measured and Theoretical Phase Jitter Error,

Resolution	Rectangular Rate, K	Keystone Rate, Q	Theory	Measured
5.0	1.0	3.0	1.527	0.763
5.0	1.0	5.0	0.916	0.758
5.0	1.0	7.0	0.654	0.756
5.0	1.0	9.0	0.509	0.754
10.0	2.0	3.0	0.752	0.408
10.0	2.0	5.0	0.451	0.401
10.0	2.0	7.0	0.322	0.403
10.0	2.0	9.0	0.251	0.403
20.0	4.0	3.0	0.373	0.128
20.0	4.0	5.0	0.224	0.118
20.0	4.0	7.0	0.160	0.115
20.0	4.0	9.0	0.124	0.115

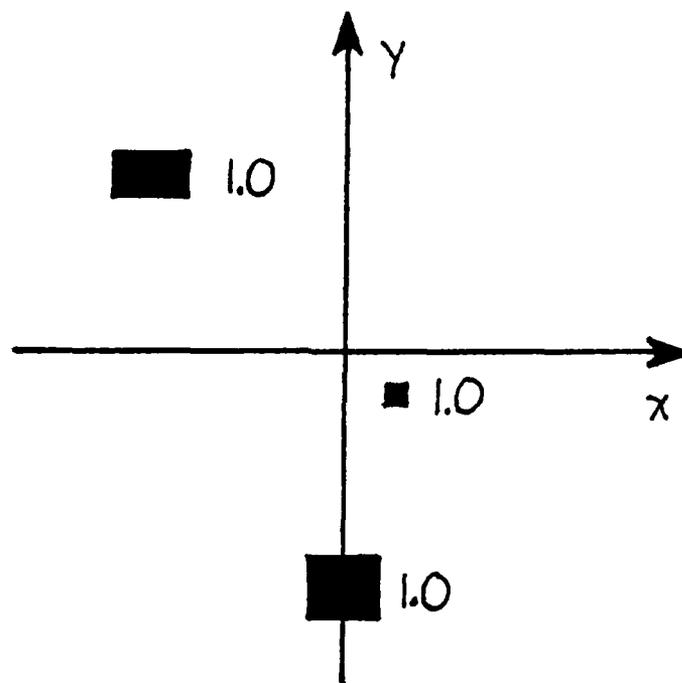


Figure 4.5 a) Ideal scene #1.

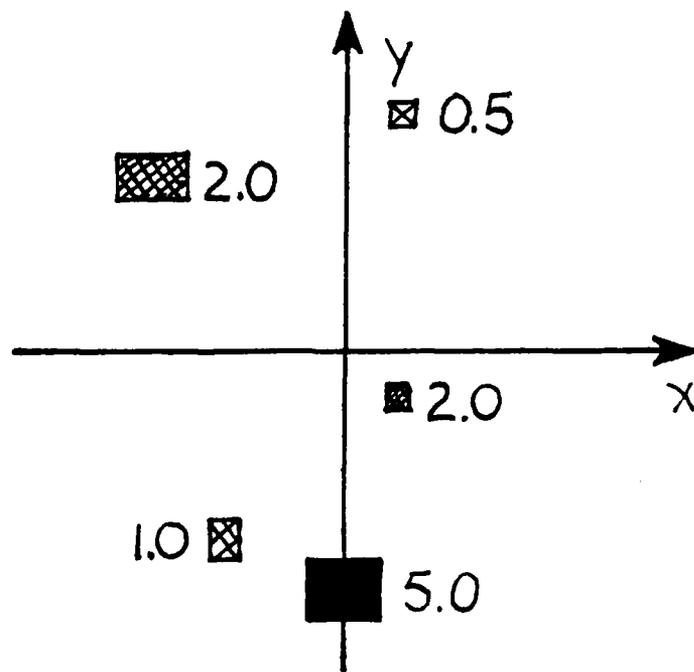


Figure 4.5 b) Ideal scene #2.

obtaining an equation that approximates the phase jitter.

In Figure 4.6a is the actual image generated by sampling the 2D FT of scene #1 (Figure 4.5a) on a rectangular raster. The desired resolution is 10 meters and the FFT length is 64 for an effective rectangular oversampling rate of approximately 2. In Figures 4.6b - 4.6d are the results of sampling on a keystone raster with varying oversampling rates. The expected result is that higher oversampling rates should produce higher quality images, i.e., less phase jitter noise. These results appear to support this contention. Experiments with scene #2 using this model have generated similar results. These are illustrated in Figures 4.7a - 4.7d.

The third experiment really consists of two parts, in which two different noise models are used. In the first part, a noise model corresponding to (4.5) is used. This corresponds to a signal independent bound on the slope of the FD data. Since the bound is independent of the signal, this model will generate uncorrelated noise. From (4.7), the expected value of the jitter distance is J_d . Substituting J_d and the constant K representing the maximum slope of the FD data, (4.2) can be rewritten as

$$R^J(\omega_k) = R(\omega_k) + KJ_d. \quad (4.15)$$

Breaking (4.15) down into the real and imaginary components, (dropping the ω_k dependence)

$$[s_r^J + js_i^J] = [s_r + js_i] + [KJ_d + jKJ_d], \quad (4.16)$$

the subscripts r and i denote the real and imaginary components of the complex value. Equation (4.16) also represents the method of



Figure 4.6 a) Reconstruction of #1
from rectangularly sampled data.

Figure 4.6 b) Reconstruction of #1
from keystone sampled data, $Q=1$.

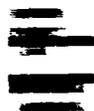


Figure 4.6 c) Reconstruction of #1
from keystone sampled data, $Q=3$.

Figure 4.6 d) Reconstruction of #1
from keystone sampled data, $Q=5$.

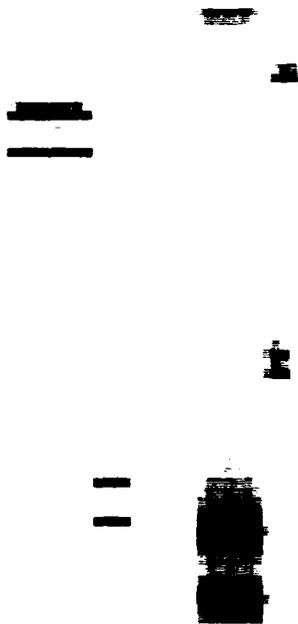


Figure 4.7 a) Reconstruction of #2 from rectangularly sampled data.

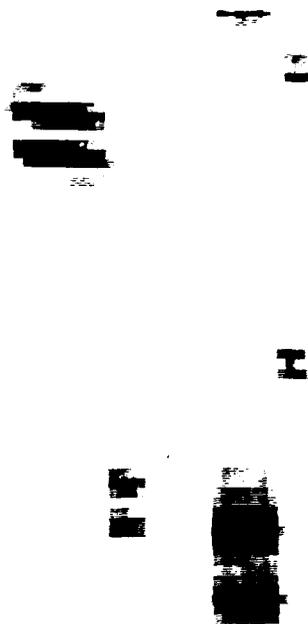


Figure 4.7 b) Reconstruction of #2 from keystone sampled data, $Q=1$.

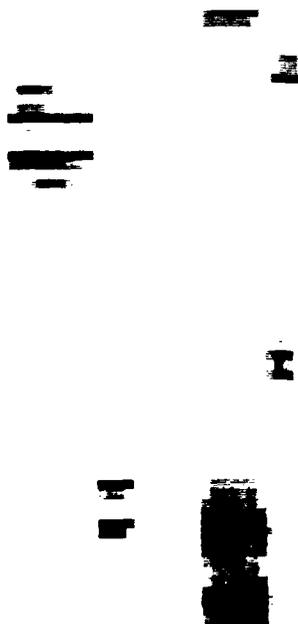


Figure 4.7 c) Reconstruction of #2 from keystone sampled data, $Q=3$.

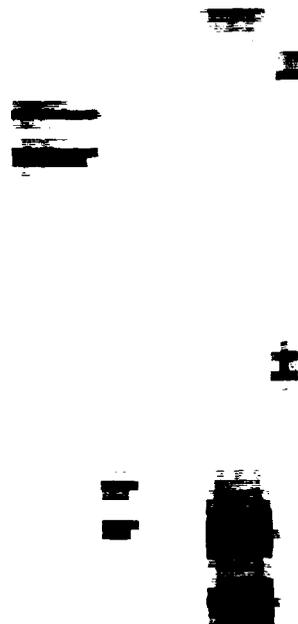


Figure 4.7 d) Reconstruction of #2 from keystone sampled data, $Q=5$.

implementation. Each real and imaginary sample of the rectangularly sampled FD is modified by the addition of a random number, KJ_d , where J_d is independent for the real and imaginary components. J_d is independent and uniformly distributed over one-half the keystone sampling interval. The expected result of this experiment is that the final image will appear to be that obtained from the unjittered data with noise added to the image after the 2D FT. This is due to the linearity of the FT and the fact that the noise is uncorrelated to the signal. The results illustrated in Figures 4.8b and 4.8d are clearly quite different from the results in Figures 4.8a and 4.8c.

The major reason for this discrepancy is that the noise model employed in this experiment is uncorrelated to the signal. If the signal independent model is used, then a constant amount of noise is the assumed jitter error, even if there is no signal. Clearly, when the signal is jittered, the amplitude of the noise is related to the signal. In equation (4.2), this relationship is the derivative of the signal. The model in (4.16) is a realization of the signal independent bound employed in experiment two. In order for that model to be independent of the signal, an uncorrelated model had to be assumed. This is the source of using the maximum possible phase shift as an approximation to the greatest slope. An improvement to this model would include some measure of the signal. Although this technique is impossible for general analysis, it can be employed in simulations to test the validity of the jitter model.

An improvement over the signal independent approximation of (4.3) would be to use, as the measure of the slope of the signal, a

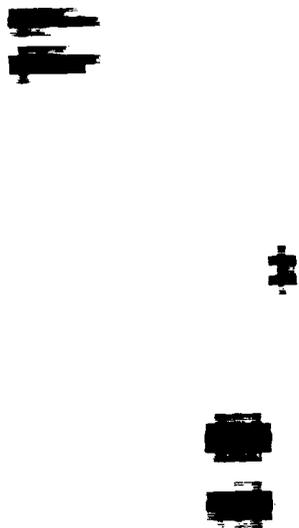


Figure 4.8 a) Reconstruction of #1
from keystone sampled data, $Q=5$.



Figure 4.8 b) Reconstruction of #1
from rectangularly sampled data
corrupted by additive noise.



Figure 4.8 c) Reconstruction of #2
from keystone sampled data, $Q=5$.

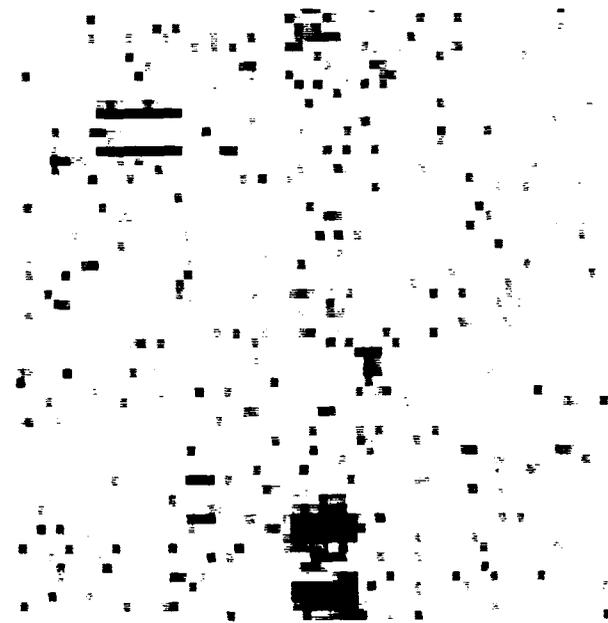


Figure 4.8 d) Reconstruction of #2
from rectangularly sampled data
corrupted with additive noise.

differential. With this model, (4.5) becomes

$$N^J(\omega_k) = J_d [R(\omega_k) - R(\omega_{k-1})] \quad (4.17)$$

which is certainly a better approximation to the instantaneous slope than a bound on the maximum slope is. Note that J_d has been substituted in for the expected value of the jitter distance. This model can also be expanded in terms of its components, as in (4.16), and this represents the method of implementation. As before, J_d is independent of the signal. The two values used to modify the real and imaginary components are also independent of each other and J_d is uniformly distributed over one-half the keystone sampling interval.

The results of this experiment are illustrated in Figures 4.9b and 4.9d. Comparing these results to those in Figures 4.9a and 4.9c, very good agreement is achieved in both examples. These pictures are quite comparable, even to the degree of smearing present in the targets. A conclusion is that the correlated noise model, (4.17), quite accurately reflects the affects of jittering the data.

4.4 Conclusion

It was shown that the phase jitter model provides some measure of the actual phase error introduced by nearest-neighbor interpolation. That the model employing a bound on the maximum slope is scene-independent should allow its use as a general design tool for digital SAR systems. Some of the features of this phase model follow. First, the use of phase error as a measure of data quality in SAR has been shown to reflect some aspects of image quality fairly well. Second,

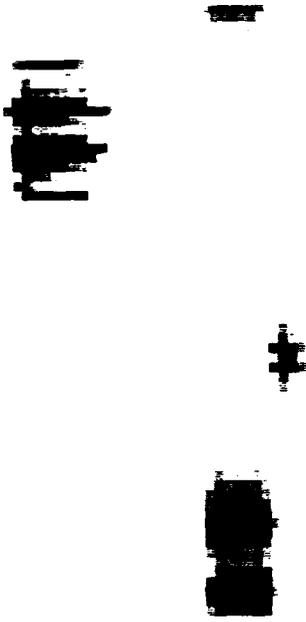


Figure 4.9 a) Reconstruction of #1
from keystone sampled data, $Q=5$.

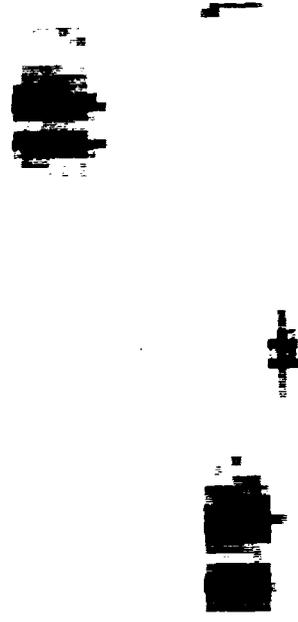


Figure 4.9 b) Reconstruction of #1
from rectangularly sampled data
corrupted by multiplicative noise.

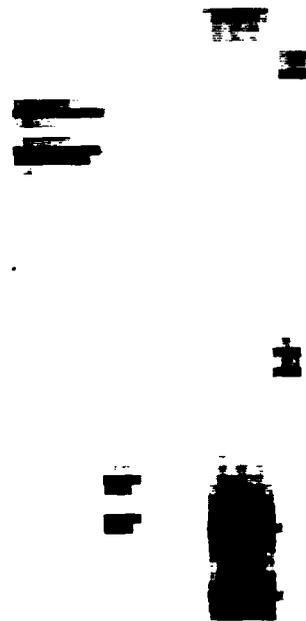


Figure 4.9 c) Reconstruction of #2
from keystone sampled data, $Q=5$.



Figure 4.9 d) Reconstruction of #2
from rectangularly sampled data
corrupted by multiplicative noise.

equations (4.10a) and (4.14a) do provide qualitative relationships between various parameters and the expected phase error. The last two experiments provided further verification of the reliability of these models. The important result is that a correlated noise model quite accurately reflects the affects of jittering the data. While the correlated model was shown to be quite good and the independent model appeared to be rather poor, it should be noted that these results reflect a rather fine level of detail in the image. In a real SAR system, the average target will not be as relatively large as in these simulations. Further, sampling requirements will likely reduce the level of jitter noise below that simulated in these experiments. In this case, the uncorrelated, signal independent noise model may be satisfactory for modelling the affects of jittering.

As a result of this work, it has become apparent that flexible sampling schemes can greatly reduce the computational requirements of a digital SAR system. In cases where flexible sampling is not available, attention should be paid to the selection of data formats. By the intelligent choice of rasters, the computational burden, normally imposed by coordinate transformations, can be greatly reduced. An area requiring further work is the choice of optimal sampling strategies and data formats to minimize both the interpolation error and computational requirements.

5. CONCLUSION

This thesis considered three distinct but related topics. A purpose of this conclusion is to draw this work together in a more cohesive manner. In Chapter two, some results concerning the effects of discretization on Gerchberg's and Papoulis' algorithms were studied. It was identified that a major problem in any implementation of these techniques is the finite number of samples of data, or filter coefficients, that can be stored and manipulated. An obvious result of this finite implementation is the inability to perform ideal filtering. Performing this non-ideal processing introduces error into the resultant signal. Because these extrapolation techniques are recursive, error produced in one pass is modified, or compounded, in the next pass. Further, even in the continuous case, these algorithms obtain the MNLS solution only after an infinite number of iterations. The net effects of employing finite records and performing only a finite number of iterations are to produce convergence behaviour as illustrated in Figure 2.5.

The emphasis of Chapter two was to characterize the behaviour of Gerchberg's and Papoulis' algorithms as a function of the record lengths employed in the realization. Two key results were obtained from this analysis. First, both algorithms can be characterized as contraction mappings for any finite record length. The implication is that for any initial guess, the algorithm will converge. Second, for most realistic applications of these techniques, the optimal solution is obtained after a relatively small number of iterations in comparison to the number

required to approach the fixed point. From a practical standpoint, this implies that the algorithm should not be allowed to converge to the fixed point, but rather, be terminated after a much smaller number of iterations. Precise means of determining when this optimal solution is obtained and the quality of the extrapolation have not been addressed and offers one area of further study.

A motivation for studying these techniques was the desire to employ these and similar methods in the missing cone problem. As pointed out in Chapter three, other researchers had observed convergence behaviour similar to that in the 1D case. Observations on early implementations of PSA also generated similar results and motivated the detailed study of Gerchberg's and Papoulis' algorithms. The extension of the 1D results to the 2D case supplies a heuristic explanation for the convergence behaviour observed in both the algorithms proposed in this work, as well as in other techniques. Further, knowledge of the convergence behaviour aided the selection of the optimal solutions.

Results published in Chapter three illustrate the degree of recovery possible by employing the PSA and AIM algorithms. As some of the examples indicate, usable reconstructions can be obtained with as little as 35% of the data. An important feature of these techniques is that they require operations on the order of the same number of operations required for the initial reconstruction. Also of significance is that while techniques proposed by other researchers require $O(N^6)$ operations, PSA and AIM need only approximately $O(N^3)$ operations and produce superior results.

During the study of Gerchberg's and Papoulis' algorithms, other researchers published results rediscovering some older extrapolation methods. Some of these methods, particularly those proposed by Jain and Ranganath, provide a means of obtaining the exact MNL solution as opposed to the approximation generated by Gerchberg's or Papoulis' algorithm. Incorporation of Jain's methods into the PSA and AIM algorithms is one possible extension of this work. It is conjectured that these improved algorithms would produce results superior to those published here.

One possible extension of the PSA and AIM algorithms is to situations where the missing cone consists of several unconnected regions, i.e., several missing cones. The algorithms as described in Chapter three could very readily handle this data format. The only modifications required would be a more flexible indexing scheme for the substitution of known projection (or Fourier domain) data. Another extension of PSA and AIM is to the data geometry present in the SAR case. In fact, the limited data situation illustrated in Figure 4.2 was the original motivation for studying tomographic reconstruction schemes, and consequently, the missing cone problem. At present, other researchers are applying tomographic concepts to the SAR case. Considering the far superior performance of the convolutional-backprojection reconstruction technique over the direct Fourier method, it seems reasonable to apply CBP techniques to the SAR case. An obvious extension to employing CBP in SAR is to apply extrapolation algorithms similar to PSA and AIM as a means of further improving the resolution in SAR systems.

REFERENCES

1. R.W. Gerchberg, "Super-resolution through error energy reduction," Optic Acta, vol. 2, no. 9, pp. 709-720, 1974.
2. G.T. Herman, Image Reconstruction from Projections, The Fundamentals of Computerized Tomography. New York: Academic Press, 1980.
3. R.O. Harger, Synthetic Aperture Radar Systems. New York: Academic Press, 1970.
4. R.M. Mersereau and A.V. Oppenheim, "Digital reconstruction of multidimensional signals from their projections," Proc. IEEE, vol. 62, pp. 1319-1338, Oct. 1974.
5. T.S. Huang, J.L.C. Sanz, H. Fan, J. Shafii and B. Tsai, "Numerical comparison of several algorithms for band-limited extrapolation," Submitted to Applied Optics, 1982.
6. N. Levinson and R.M. Redheffer, Complex Variables. New York: Holden-Day, chapter 3, 1970.
7. A. Papoulis, "A new algorithm in spectral analysis and band-limited extrapolation," IEEE Trans. CAS, vol. CAS-2, no. 9, pp. 735-742, Sept. 1975.
8. J.L.C. Sanz and T.S. Huang, "On the Gerchberg-Papoulis algorithm," Submitted to IEEE Trans. CAS, 1983.
9. L. Landweber, "An iteration formula for Fredholm equations of the first kind," Amer. J. Math, 73, pp. 615-624, 1951.
10. M.S. Sabri and W. Steenaart, "An approach to band-limited signal extrapolation: The extrapolation matrix," IEEE Trans. CAS, vol. CAS-25, no. 2, pp. 74-78, Feb. 1978.
11. J.A. Cadzow, "An extrapolation procedure for band-limited signals," IEEE Trans. ASSP, vol. ASSP-27, no. 1, pp.4-12, Feb. 1979.

12. M.S. Sabri and W. Steenaart, "Comments on 'An extrapolation procedure for band-limited signals' ", IEEE Trans. ASSP, vol. ASSP-28, no. 2, p. 254, April 1980.
13. J.A. Cadzow, "Observations on the extrapolation of band-limited signal problem," IEEE Trans. ASSP, vol. ASSP-29, no. 6, pp. 1208-1209, Dec. 1981.
14. M.S. Sabri and W. Steenaart, "Rebuttal to 'Observations on the extrapolation of a band-limited signal problem' ", IEEE Trans. ASSP, vol. ASSP-29, no. 6, p. 1209, Dec. 1981.
15. A.K. Jain and S. Ranganath, "Extrapolation algorithms for discrete signals with applications in spectral estimation," IEEE Trans. ASSP, vol. ASSP-29, no. 4, pp. 830-845, Aug. 1981.
16. C.R. Rao and S.K. Mitra, Generalized Inverse of Matrices and Applications. New York: Wiley, 1971.
17. T.S. Huang and J.L.C. Sanz, "Four models for the band-limited extrapolation problem," Proc. of the Topical Meeting on Signal Recovery and Synthesis with Incomplete Information and Partial Constraints, Incline Village, Nev., Jan. 12-14, 1983.
18. J.L.C. Sanz and T.S. Huang, "A unified Hilbert space approach to iterative linear signal restoration," Submitted to the Journal of the Optical Society of America, 1983.
19. M. Vidyasagar, Nonlinear Systems Analysis. New York: Prentice-Hall, chapter 3, 1978.
20. J.M. Ortega and W.C. Rheinolt, Iterative Solution of Non-linear Equations in Several Variables. New York: Academic Press, 1970.
21. R. Schaefer, R. Mersereau and M. Richards, "Constrained iterative restoration algorithms," Proc. IEEE, vol. 69, pp. 432-450, April 1981.
22. J.L.C. Sanz and T.S. Huang, "Discrete and continuous band-limited signal extrapolation," Submitted to IEEE Trans. ASSP, 1982.

23. R.E. Crochiere and L.R. Rabiner, "Optimum FIR digital filter implementations for decimation, interpolation and narrow-band filtering," IEEE Trans. ASSP, vol. ASSP-23, no. 5, Oct. 1975.
24. S.W. Roland, "Computer implementation of image reconstruction formulas," Topics in Applied Physics, Image Reconstruction from Projections, Implementation and Applications. New York: Springer-Verlag, Editor: G.T. Herman, pp. 9-79, 1979.
25. J. Radon, "(On the determination of functions from their integrals along certain manifolds)," Sachsische Akademie Der Wissenschaften (Leipzig) — Mathematisch — Naturwissenschaftliche Klasse. Sitzungserichte, vol. 69, pp. 262-277, 1917 (German).
26. L.A. Shepp and B.F. Logan, "The Fourier reconstruction of a head section," IEEE Trans. Nucl. Sci., vol. NS-21, pp. 21-42, 1974.
27. H. Stark, "Direct Fourier reconstruction in computer tomography," IEEE Trans. ASSP, vol. ASSP-29, no. 2, pp. 237-244, April 1981.
28. A. Papoulis, Systems and Transforms with Applications in Optics. New York: McGraw-Hill, chapter 5, 1968.
29. D.C. Munson, Jr. and W.E. Higgins, "Faster algorithms for computer-aided tomography and spotlight mode synthetic aperture radar," Proc. 16'th Asilomar Conf. on Circuits, Systems and Computers, 1982.
30. G.N. Ramachandran and A.V. Lakshminarayanan, "Three-dimensional reconstruction from radiographs and electron micrographs: Application of convolutions instead of Fourier transforms," Proc. National Academy of Sciences, USA, vol. 68, no. 9, pp. 2236-2240, Sept. 1971.
31. M. Schwartz and L. Shaw, Signal Processing: Discrete Spectral Analysis, Detection and Estimation. New York: McGraw-Hill, chapter 6, 1975.
32. J.P. Burg, "Maximum entropy spectral analysis," Proc. 37'th Meeting of the Society of Exploratory Geophysicists, Oklahoma City, OK., 1967.
33. A. Papoulis, "Maximum entropy and spectral estimation: A review," IEEE Trans. ASSP, vol. ASSP-27, no. 6, pp. 1176-1187, Dec. 1981.

34. S.J. Wernecke and L.R. D'Addario, "Maximum entropy image reconstruction," IEEE Trans. on Computers, vol. C-26, pp. 351-364, April 1977.
35. N. Baba, K. Murata, K. Okada and Y. Fujimoto, "Image reconstruction from electron micrographs by maximum entropy methods," Optik 58, no. 4, pp. 233-239, 1981.
36. N. Baba, K. Murata, K. Okada and Y. Fujimoto, "Three-dimensional image reconstruction from electron micrographs and spatial filtering for the image," Optik 54, no. 2, pp.97-105, 1979.
37. R. Gordon, R. Bender and G.T. Herman, "Algebraic reconstruction techniques (ART) for three-dimensional electron microscopy and x-ray photography," Journal Theoretical Biology, #29, pp. 471-481, 1970.
38. K.M. Hanson and G.W. Wecksung, "Bayesian approach to limited-angle CT reconstructions," Proc. of the Topical Meeting on Signal Recovery and Synthesis with Incomplete Information and Partial Constraints, Incline Village, Nev., Jan 12-14, 1983.
39. L. Garnero and J. Brunol, "Tomographic imaging with a limited view angle using an expansion on a set of eigenfunctions adapted to space-limited objects," Proc. of the Topical Meeting on Signal Recovery and Synthesis with Incomplete Information and Partial Constraints, Incline Village, Nev., Jan 12-14, 1983.
40. A. Lent and H. Tuy, "An iterative method for the extrapolation of band-limited functions," Journal of Mathematics Analysis and Applications, vol. 83, no. 2, pp. 554-565, 1981.
41. S. Renjen and T.S. Huang, "Experimental results on the missing cone problem in computer-aided tomography," IEEE Symposium on Computer Applications in Medicine, Washington D.C., Nov. 2-4, 1981.
42. H.K. Tuy, "An algorithm for incomplete range of view reconstruction," Proc. of the Topical Meeting on Signal Recovery and Synthesis with Incomplete Information and Partial Constraints, Incline Village, Nev., Jan. 12-14, 1983.
43. J.R. Fienup, "Reconstruction of an object from the modulus of its Fourier transform," Optics Letters, vol. 3, no. 1, pp. 27-29, July 1978.

44. D.A Hayner, S. Renjen, T.S. Huang and W.K. Jenkins, "Algorithms and experimental results on image reconstruction from limited data," Proc. IEEE Int. Conf. on ASSP, Paris, May 3-5, 1982.
45. P.A. Rattey and A.G. Lindgren, "Sampling the 2-D Radon transform," IEEE Trans. ASSP, vol. ASSP-29, no. 5, Oct. 1981.
46. R.W. Schafer and L.R. Rabiner, "A digital signal processing approach to interpolation," Proc. of the IEEE, vol. 61, pp. 692-702, June 1973.
47. L.R. Rabiner and B. Gold, Theory and Application of Digital Signal Processing. New York: Prentice-Hall, chapter 3, 1975.
48. W.G. Weis and W.K. Jenkins, "Theory and simulation of digital FFT processors for spotlight mode synthetic aperture radar," Lockheed Missles and Space Corporation, Report #D566409, Dec. 1977.
49. J.J. Kovaly, Synthetic Aperture Radar. Dedham: Artech House, 1976.
50. J.C. Kirk, Jr., "A discussion of digital processing in synthetic aperture radar," IEEE Trans. Aerosp. Electron. Syst., vol. AES-11, pp. 326-337, May 1975.
51. J.L. Auterman et. al., Unclassified excerpts from "Airborne spotlight radar system," Radar and Optics Division, Environmental Research Institute of Michigan, AFAL-TR-75-105, Dec. 1975.
52. K. Tomiyasu, "Tutorial review of synthetic aperture radar (SAR) with applications to imaging the ocean surface," Proc. of the IEEE, vol. 66, no. 5, May 1978.
53. D.A. Schwartz, "Analysis and experimental investigation of three synthetic aperture radar formats," Coordinated Science Laboratory, University of Illinois, Urbana, Ill. Technical Report #T-94, March 1980 (M.S. Thesis).
54. A.V. Oppenheim and J.S. Lim, "The importance of phase in signals," Proc. of the IEEE, vol. 69, no. 5, May 1981.

VITA

David Alan Hayner was born on April 30, 1957 in Ancon, Canal Zone, Panama. He obtained his Bachelor of Science in Computer Engineering from the University of Illinois in May, 1979. In January of 1979 he entered the Graduate college of the University of Illinois and received his Master of Science in Electrical Engineering in October of 1980. During the summer of 1978 Mr. Hayner was employed by International Business Machines Corporation, Rochester, Minnesota and during the summer of 1979 he worked for Lawrence Livermore National Laboratories, Livermore, California. From January 1979 until May 1982 he was employed by the Department of Electrical Engineering at the University of Illinois as a graduate Teaching Assistant. For his efforts in teaching he was awarded the Harold L. Olesen Award for Excellence in Undergraduate Teaching by a Graduate Student in 1980.

Since June 1981, Mr. Hayner has worked as a Research Assistant in the Coordinated Science Laboratory. During his graduate studies he has published one technical report, two conference papers and a book chapter. He presently has accepted employment with Gould Laboratories, Rolling Meadows, Illinois. Mr. Hayner is a member of the Institute of Electrical and Electronics Engineers (IEEE).

FILMED

FILMED

8