ADAPTIVE LATTICE FILTERING

FOR RADAR APPLICATIONS

by

© Carey James Gibson, B.Sc., M.Eng.

A Thesis

Submitted to the School of Graduate Studies

in Partial Fulfillment of the Requirements

for the Degree

Doctor of Philosophy

McMaster University

January, 1982

ADAPTIVE LATTICE FILTERING

FOR RADAR APPLICATIONS

DOCTOR OF PHILOSOPHY (1982)
(Electrical Engineering)

McMASTER UNIVERSITY
Hamilton, Ontario

TITLE: Adaptive Lattice Filtering for Radar Applications

AUTHOR: Carey James Gibson, B.Sc. (Dalhousie University)
M.Eng. (McMaster University)

SUPERVISOR: Dr. S. Haykin

NUMBER OF PAGES: xv, 197

ABSTRACT

This thesis examines the lattice-structure prediction-error filter, and its application to air-traffic-control radar for the detection of targets (such as aircraft) obscured by clutter (unwanted reflections from the ground or weather systems). The digitally implemented lattice-structure filter adapts to and eliminates the clutter spectrum, producing an output only when a target causes a change in the input signal. Conventional MTI filters do not perform this detection as reliably.

Adaptation to an input signal results from the recursive calculation of the lattice-structure filter's reflection coefficients. Six algorithms for this calculation were examined and compared using simulated radar data. A number of adaptive methods for continuously implementing these algorithms were also analysed. These included the standard gradient and least-squares methods, and two new methods developed in this thesis, the simple gradient and adaptive gradient methods. The harmonic-mean algorithm and the standard and simple gradient methods were selected as most appropriate for this application.

The adaptive learning characteristics (both stationary and non-stationary) of these lattice methods were studied theoretically and experimentally, and quantitative relationships were developed describing their behaviour. The performance of the lattice-structure as a radar clutter filter was examined in terms of improvement factor, receiver-operator-characteristic, and sub-clutter visibility. Both simulated and

actual radar data were used. The actual radar data included signals from aircraft, bird flocks, ground clutter, and several types of weather clutter. The performance of the lattice-structure filter with this data was found to be more consistent and consistently better than the conventional MTI filter.

ACKNOWLEDGEMENTS

I would like to express my sincere thanks to Dr. Simon Haykin, for his supervision, help, and encouragement during this research. I also gratefully acknowledge the contributions of: Stan Kesler and Brian Currie, for their helpful assistance; Peter Metford and others with the Communications Research Laboratory, for many stimulating discussions; the National Science and Engineering Council of Canada, for their support; and Linda Hunter, for typing this thesis.

However, my greatest appreciation is reserved for my parents, whose support and love were very important to me during this period of my life.

TABLE OF CONTENTS

CHAPTER 3 (cont'd.)

LIST OF ILLUSTRATIONS

LIST OF TABLES

LIST OF PRINCIPAL SYMBOLS

$a_{m,i}$	PEF coefficient for stage i of filter with order m
$b_m(n)$	Backward prediction-error output for filter order m at time n
$f_m(n)$	Forward prediction-error output
m	Filter order
n	Data sample number (time index for discrete time series)
$v_m(n)$, $y_m(n)$	Recursive summations used by Method II lattice filter
$x(n)$	Input data for filter
$E[\ ]$	Expectation operator
$\|\ \|$	Absolute value operator
$*$	Complex conjugation operator
f	Frequency (normalized to the sampling frequency)
$\rho_m(n)$	Reflection coefficient of lattice filter
ω	Adaptive weighting factor (simple gradient method)
μ	Adaptive weighting factor (standard gradient method)
$\alpha_m(n)$	Adaptive step size for simple gradient method

Abbreviations:

CNR	Clutter-to-noise ratio
CRL	Communications Research Laboratory (McMaster)
I and Q	Inphase and quadrature

Abbreviations (cont'd.)

MTD	Moving target detector
MTI	Moving target indicator
P_D	Probability of detection
P_{FA}	Probability of false alarm
PEF	Prediction-error filter
PRF	Pulse repetition frequency
ROC	Receiver-operator-characteristics
SCR	Signal-to-clutter ratio
SCV	Sub-clutter visibility
TDL	Tapped delay line
PPI	Plan-position indicator
CFB	Canadian Forces Base

CHAPTER 1

INTRODUCTION

1.1 The Detection Problem with Scanning Radar

This thesis is concerned mainly with the design and analysis of a specific type of adaptive digital filter having a lattice structure, and with the analysis of its performance for radar enhancement using both simulated and actual radar data. Therefore, before discussing the filter structure in detail, it would be useful to describe the problem to which it is directed.

The radar signals in question come from a scanning radar used for air-traffic control in an airport environment. This radar uses a very directional antenna which repeatedly transmits a short, monotonic pulse of radio energy and then listens for echoes from objects (such as aircraft) in the environment. (The time delay between pulse and echo indicates the object's range.) While this is happening, the antenna is slowly scanning around in a clockwise direction to cover the entire area in azimuth.

This leads to an important concept, the radar resolution cell. This cell is a small segment in range and azimuth of the radar environment, which may be thought of as the space effectively occupied by a point object, as seen by the radar. Alternatively, the cell may be defined as the closest spacing objects may have to be resolved separately

by the radar. In range, the cell dimension is determined by the pulse length and is typically about 0.1 km. In azimuth, the dimension is determined by the half-power beamwidth of the antenna pattern, generally less than two degrees.

In typical radars, because of this antenna pattern, between ten and twenty pulses may pass through the cell before the antenna has scanned past the cell. An object in the cell may return an echo from each of these pulses, forming a time series of samples representative of the object. If these pulse echoes are detected coherently by the radar receiver, then the time series contains both amplitude and phase information; the amplitude being related to the size and reflectivity of the object (radar cross-section), and the phase being related to the change in range of the object (Doppler shift). This time series of samples, taken at a single range and analysed either in blocks or continuously, forms the basis for many radar signal processing systems, including the lattice-structure filters studied in this thesis.

The need for some type of processing of the radar signal is dictated by the fact that the desired echoes from a target (such as an aircraft) are not the only signals detected. The additional, unwanted signals, known as "clutter", come from a number of sources within the radar environment and can easily obscure the unprocessed target signal.

The most common source of clutter is reflections from the ground, or from objects at ground level (trees, buildings, etc.). This clutter is especially strong near the radar site, but may occur at any range from high objects or hilly terrain. Another major source of clutter is

weather disturbances. Rain, snow, and other storm conditions can result in strong reflections from large areas of the sky, obscuring any targets in the area. A different sort of weather-related clutter results from atmospheric inversion; areas of denser air which reflect or refract the radar pulse. Echoes created this way are often moving and may resemble targets or cover significant areas. Even small objects such as birds can be detected (especially in flocks) and may be mistaken for aircraft.

1.2 Traditional Versus Adaptive Processing

Processing of the radar signal, for the purposes of this thesis, will be limited to the function of filtering out a clutter signal, which might be obscuring a target signal, while leaving the target signal (if present) in detectable form. (Other processing functions may be applied to the same signals for other purposes.)

The traditional processing technique commonly in use today is the moving-target indicator (MTI) filter which assumes that the target of interest is moving (radially to the radar) and that the clutter is not. It is essentially a high-pass filter applied to detect the Doppler signal in the time series of samples from a constant range. This is usually implemented using a fixed tapped-delay-line (TDL) canceller, comparing two or more successive pulses for any changes.

The assumption that the target is moving and the clutter is not results in two problems. First, strong targets may be lost due to a low radial velocity. This includes slow-flying bird flocks, which are

a definite threat to air safety near airports. Second, storms of even moderate intensity are often moving fast enough to break through the MTI filter and thus obscure much of the radar screen. Yet, it is during such storms that reliable radar service is most needed for air safety. What is needed is a technique not requiring the MTI assumption.

One such technique is to use an adaptive filtering system, which assumes only that the target and clutter have different radial velocities. This presents a new problem, namely, how to tell what is a target and what is clutter. To solve this, we may use other characteristic differences between the two signals. Clutter returns (weather, ground, etc.) normally have consistent statistical characteristics over a large area (and thus a large number of samples in the time series). Target returns, on the other hand, have different spectral statistics, normally covering a very small area limited by the resolution cell of the radar.

A number of different schemes have been proposed to take advantage of these characteristic differences. One method of doing this involves classifying the clutter on the basis of spectral information from conventional analysis, and then choosing a fixed filter having appropriate clutter-rejection characteristics (based on the classification) [2, 30]. This method has been shown to be an improvement over the fixed MTI filter. However, the results are not as good as they might be, largely due to the limited response of the classification, end-bias effects, and other limitations. Another scheme showing some success uses a bank of filters (or a fast Fourier transformer [FFT]) centered on different frequency bands covering the Doppler spectrum.

A thresholding system is used to reduce the outputs of filters responding over a large area (the clutter bands) [42]. This scheme, however, is limited to a small number (e.g., 8) of filters having correspondingly wide responses, and targets falling within these bands are also reduced.

Another approach is to use a filter which adapts itself to remove the clutter signal. (This thesis falls into this category of approach). A filter which performs this function is known as a prediction-error filter (PEF), where the output of the filter is the error in the prediction process; that is, any signal components which were not predicted, so as to be filtered out. In the field of signal processing, it is often desirable to make use of a filter like this, which adapts itself to the input signal in such a way that the error output is minimized (e.g., to eliminate noise, interference, echoes, or other unwanted signals). This is one aspect of linear prediction [1], the basic assumption of which is that the signal in question can be modelled as a linear combination of previous inputs and/or outputs of the filter.

One way in which the PEF may be used to detect a target in clutter takes advantage of the different area coverage of the two types of signals. By using a large number of previous samples of the time series and their statistics, the value of the next sample in the time series can be predicted. If the predicted value is then subtracted from the actual value, a signal component with similar statistics (e.g., the continuing clutter signal) will be subtracted (filtered out) while signal components with different statistics (e.g., targets) will not.

Another way in which the target may be detected is quite similar to this, except that the filter is allowed to adapt quickly, so that it adapts to targets as well as clutter. Then, there is a significant output from the filter only during a sudden change in the spectral characteristics of the input signal. This occurs during the transition from no-target to target conditions, whether or not clutter is present.

In previous attempts to use the PEF for this type of application, a tapped-delay-line structure has been used [43]. This structure, in different forms, has been studied for many years and its characteristics are generally well known [3]. However, some of these characteristics make this particular structure a bad choice for this type of processing. One of the most important of these is the convergence properties. As the clutter and target signals have specific lengths of time during which they are present in the time series, it is desirable that the adaptive filter converge to a new signal in a set period of time (for instance, so that the clutter may be adapted to, while the target is not). Other negative characteristics include tendencies to instability, poor resolution, large levels of quantization noise with higher filter orders, and the inability to alter the filter order during operation. Most of these characteristics are related to the fact that this structure is adapted by minimizing a single, global error criterion [4].

1.3 Lattice-Structure Filtering

The negative characteristics of the tapped-delay-line structure have led to the examination of a more recently developed structure for adaptive prediction-error filtering, the lattice structure. This structure was originally proposed by Burg [5, 6], for use in spectral estimation, and independently derived by Itakura and Saito [7]. Unlike the tapped-delay-line structure, the operating criterion for the lattice-structure PEF minimizes the error output of each filter stage independently. Due mainly to this difference, the lattice structure escapes those negative characteristics peculiar to the tapped-delay-line structure.

The convergence properties of the adaptive lattice-structure PEF are well behaved; they are independent of the signal statistics and may be controlled over a wide range by a simple adaptive constant. In addition, the lattice coefficients (called reflection coefficients) may be chosen to guarantee stability and the filter order may be adjusted at will (i.e., intermediate filter-order outputs are always available). Also, the lattice-structure PEF generally shows higher resolution than its tapped-delay-line counterpart in reflecting the signal spectrum, and shows high insensitivity to quantization noise, regardless of filter order [8]. These factors suggest the lattice structure would be a good choice for application to radar, and this suggestion is affirmed by the results presented later in this thesis.

In addition to filtering the signal, the lattice structure PEF also provides spectral information about the signal, and thus could be used as the basis for a radar Doppler processor [31]. Alternatively,

this information may be used for classification of the clutter signal, which was the first radar application for which this structure was suggested [14]. The lattice-structure PEF has also found application in the fields of geophysics [9], speech analysis [4, 10, 11], channel equalization [12], and noise cancelling [13].

The basic development of the lattice-structure PEF is presented in Chapter 2. This reviews the construction of the lattice structure, starting from the basic tapped-delay-line structure, and presents some important properties of the lattice structure. Also analysed are a number of basic algorithms which have been proposed for minimizing the prediction-error of the filter. One of these, the harmonic-mean algorithm, is chosen as the best available for this application.

Chapter 3 deals with methods by which the algorithm used may be implemented as a continuously adaptive filter. Several methods are developed and these are compared with those developed by other researchers. The learning characteristics of these methods are studied in Chapter 4, and the formulae developed to describe their primary convergence properties are confirmed by experimental simulations.

Application of the lattice structure to radar signals begins in Chapter 5 with the analysis of simulated signals and presentation of the performance measures and standards used in the analysis of radar data processing. Analysis of actual radar data and discussion of the same are given in Chapters 6 and 7, and the results and recommendations for further work are summarized in Chapter 8.

CHAPTER 2

THE LATTICE-STRUCTURE PREDICTION-ERROR FILTER

2.1 The Prediction-Based Lattice Filter

The prediction-error filter (PEF) may be defined as a structure which combines successive samples of the input signal multiplied by coefficients, so that the output (prediction error) of the filter is minimized. There are two kinds of tapped-delay-line PEF, depending on the form of prediction error utilized. Based on a given sequence of input samples, a forward PEF is designed to minimize the mean-square value of the forward prediction error, defined as the difference between the predicted value of the input one step into the future and its actual value. On the other hand, a backward PEF is designed to minimize the mean-square value of the backward prediction error, defined as the difference between the predicted value of the input one step into the past and its actual value.

In this thesis, the input is denoted by $x(n)$, the forward prediction error by $f_m(n)$, and the backward prediction error by $b_m(n)$, where the subscript m denotes the filter order.

Consider the tapped-delay-line filter of Fig. 2.1, which is a forward PEF operating on an input $x(n)$ to produce an output $f_m(n)$. By definition, the first coefficient $a_{m,0}$ is unity. The remaining PEF coefficients $a_{m,i}$ ($i = 1, 2, \ldots, m$) are adjusted to minimize the mean-

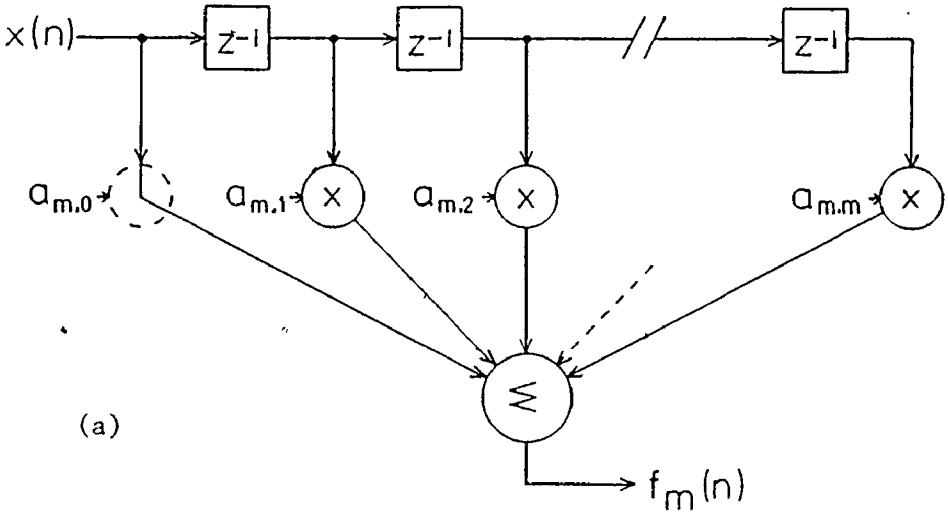

(a)

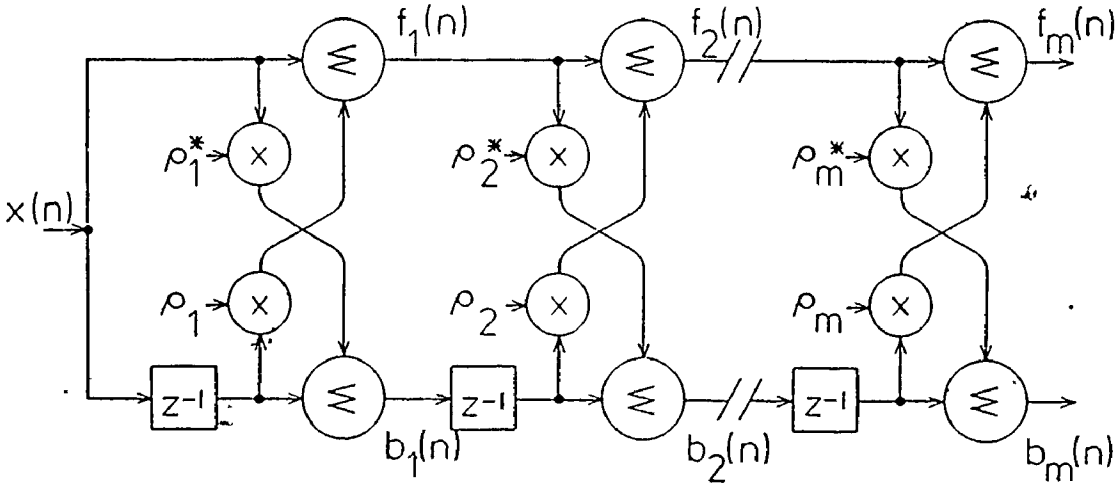

(b)

Figure 2.1: Prediction-Error Filter Structures
(a) Tapped-Delay-Line (Forward), (b) Lattice

square value of the output $f_m(n)$ by predicting the value of the input $x(n)$. The corresponding backward PEF has the same form as Fig. 2.1(a) except that the PEF coefficients $a_{m,i}$ $(i = 0, 1, \ldots, m)$ are complex-conjugated and reversed in order. With an input $x(n)$ applied to the backward PEF, the resulting output is $b_m(n)$. A forward PEF and a backward PEF may be combined in the form of a lattice structure [14], as shown in Fig. 2.1(b). The number of stages (filter orders) contained in the forward, backward, and lattice-structure PEF's are the same. The forward and backward outputs of stage $m+1$ of the lattice are given by, respectively,

$$f_{m+1}(n) = f_m(n) + \rho_{m+1}(n)\, b_m(n-1) \tag{2.1}$$

and

$$b_{m+1}(n) = b_m(n-1) + \rho^*_{m+1}(n)\, f_m(n), \tag{2.2}$$

for $n \geq 1$ and $0 \leq m \leq M-1$ where M is the filter order, $\rho_m(n)$ are the reflection coefficients (also known as partial correlation or PARCOR coefficients) and $\rho^*_m(n)$ are their complex conjugates. (The name "reflection coefficient" comes from comparing the coefficient's effect on the reduction of prediction error power in each lattice stage with the transmission of power through a terminated, two-port network [15].) The outputs of each stage of the lattice are the same as the final outputs of forward and backward PEF's of corresponding orders.

There are a number of definite advantages in using the lattice-structure PEF. One of the most important of these is the fact that the backward prediction errors are orthogonal to each other [15]. (They are

also orthogonal to the forward prediction errors, other than of the same order.) This allows for the decoupling of successive stages. Thus, the reflection coefficient and prediction errors can be calculated separately for each stage of the filter, without reference to any other stages. This independence results in the desirable convergence and stability properties of the lattice structure, as noted in the previous chapter. In addition, as with other prediction-error filters, the lattice structure is minimum phase (minimum delay) [15].

Let us now consider the relationship of the reflection coefficients to the PEF coefficients of the tapped-delay-line filter shown in Fig. 2.1. The PEF coefficients normally obey the constraints:

$$a_{m,i} = 1 \text{ for } i = 0,$$

$$a_{m,i} = 0 \text{ for } i > m \text{ or } i < 0, \tag{2.3}$$

and $$-1 \leq a_{m,m} \leq 1 \ .$$

The lattice filter in Fig. 2.1 can be replaced by an equivalent tapped-delay-line PEF of the same order to give an identical filtering action. The basic relationship between the two types of filters is that the reflection coefficient $\rho_i(n)$ equals the final coefficient $a_{i,i}$ of an i^{th} order tapped-delay-line PEF for $1 \leq i \leq m$. The PEF coefficients are then calculated from the Levinson recursion [15]:

$$a_{m,i} = a_{m-1,i} + a_{m,m} \cdot a^*_{m-1,\,m-i} \tag{2.4}$$

by starting with m = 2 and working up to the order of the filter. This relationship can also be developed empirically by equating the outputs of the two filters and working recursively through the equations

$$f_m(n) = f_{m-1}(n) + a_{m,m} \cdot b_{m-1}(n-1)$$

$$= \sum_{i=0}^{m} a_{m,i} \cdot x(n-i) \qquad (2.5)$$

and

$$b_m(n) = b_{m-1}(n-1) + a^*_{m,m} \cdot f_{m-1}(n)$$

$$= \sum_{i=0}^{m} a^*_{m,m-i} \cdot x(n-i) \qquad (2.6)$$

The reflection coefficients are also directly related to the autocorrelation function of the input time series. The autocorrelation values may be uniquely determined from the reflection coefficients, provided the zero-lag autocorrelation value is known [15].

2.2 Some Observations on Spectral Characteristics

Having computed the PEF coefficients, the spectral estimate of the signal can then be computed using the maximum-entropy method developed by Burg [6]. The idea of this method is to choose the spectrum (in the form of a non-negative function of frequency) which corresponds to the most random or the most unpredictable time series whose autocorrelation function agrees with a set of known values. The method

derives its name from the fact that this condition corresponds to the concept of maximum entropy as used in information theory. According to this method, the power spectrum of the signal is given as:

$$P(f) = \frac{p_m/W}{\left|1 + \sum_{i=1}^{m} a_{m,i} \exp\left[-j2\pi f i(\Delta t)\right]\right|^2} \qquad (2.7)$$

over the frequency range $-W \leq f \leq W$ where $W = 1/[2\cdot(\Delta t)]$ and Δt is the sampling period. The p_m is equal to the output error power, that is $p_m = E[\,|f_m(n)|^2]$.

Building on the above theoretical development, a number of practical observations can be made about the application of the lattice filter to signals for which some characteristics of the spectrum are already known:

(1) As the transfer function of a lattice filter is made up only of a number of zeros, the filter responds best to a signal having a spectrum fitting the all-pole or autoregressive model. An example of this is shown in Fig. 2.2, which shows the frequency response of a lattice filter adapted to a single frequency equal to 0.5. We see that changing the filter order from 1 to 5 has a negligible effect on the frequency response of the filter, confirming the fact that with the input spectrum containing a single pole, a single-order filter is sufficient. Zeros in the spectrum may be represented in an all-pole model by approximating the spectrum with a large number of poles [16]. To do this accurately would require a very

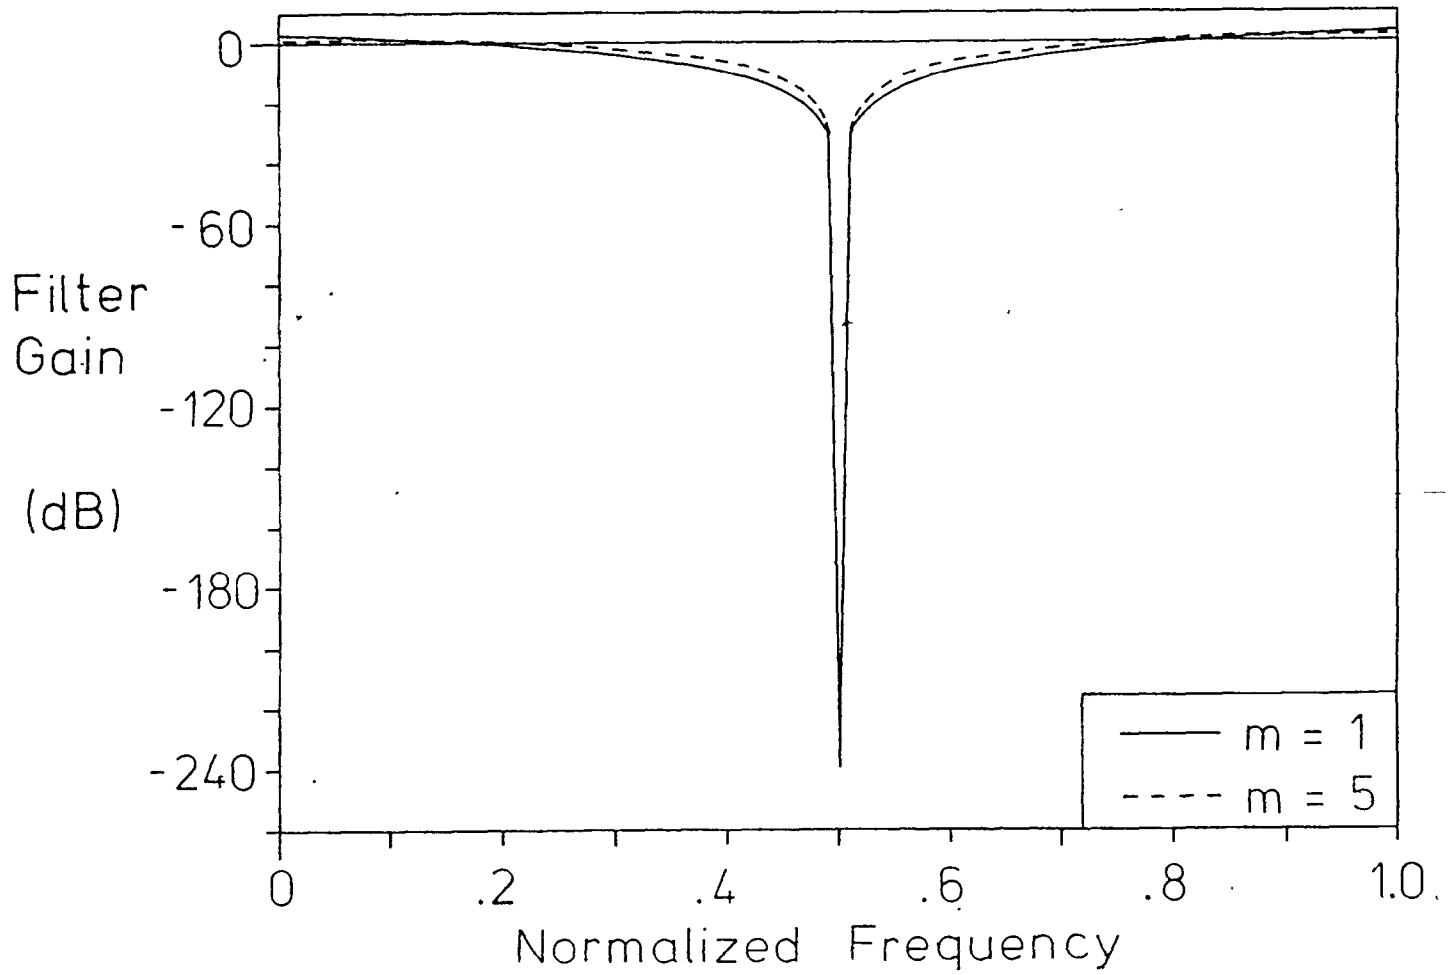

Figure 2.2: Frequency Response of a Lattice-Structure PEF Adapted to a Single Frequency Signal

large filter order. However, there is often some physical justification for only approximately locating zeros (with a lower filter order) while accurately locating poles in the spectrum. The justification is that most sensing devices (e.g., seismometers, microphones, and even ears) only approximate the troughs in a spectrum while accurately representing the peaks [16]. Thus, as the data collected often does not accurately represent the spectral zeros, there is no point in having a filter of greater accuracy with respect to the zeros. (A modified form of the lattice structure has been developed [17, 18] which responds well to all zero spectra.) The use of an autoregressive (AR) model to represent the clutter signal we wish to eliminate is also justified. Indeed, some of the earliest work on this question suggests that a Gaussian spectrum is good approximation for clutter [32]. More recent work has resulted in a second-order AR model for both ground and weather clutter [40]. This model fits the clutter's power spectral density P(f) with the formula

$$P(f) = 1/(1 + |f/f_c|^n) \quad , \qquad (2.8)$$

where f_c is the cutoff frequency. The constant n, which determines the high frequency rolloff rate, has been experimentally observed at values ranging from 2 to 4 (depending on the radar frequency and the clutter type) [41], corresponding to a filter having one or two poles. Power spectra recorded for bird clutter [33] also appear to fit this model. However, a comprehensive study of available models has led one researcher to conclude that there is "no

unique representation for the power spectral density of the clutter" [30]. The matter is further complicated by the assumption in such models of a stationary signal, which is not the case for scanning radar. Also, the spectrum of a small clutter sample area may be quite untypical of the whole clutter population. This is not a great problem, however, as the clutter spectrum can be represented as a fairly simple envelope function containing the actual complex, finely detailed structure of the spectrum. This fine detail is mainly related to the short-term variations and disturbances of the clutter, while the envelope is derived from longer term parameters. Thus, it is only necessary for the filter to adapt to the optimum clutter spectrum envelope in order to do its job. As this envelope can almost always be drawn as a broad spectral peak (exceptions apparently being sea-clutter and some man-made interference), the envelope can be represented by an autoregressive model.

(2) The number of filter stages needed to filter a given signal properly will be equal to or greater than the number of peaks in the spectrum of the signal. For sharp spectral peaks an equal (or slightly greater) number of filter stages would be sufficient. However, for broad peaks (e.g., Gaussian) more than one filter stage is needed for each peak. A broad peak is treated, in effect, as the super-positioning of a number of narrower peaks. Generally, the number of stages needed for each such peak would be quite small (typically, two or three). This property will be discussed further in Chapter 5.

(3) Burg [6] has shown that an important property of the sequence of reflection coefficients is that it converges to zero if the infinite-order output error power p_∞ does not vanish. This means that $\rho_m(n)$ goes to zero as m goes to infinity. This limit can be useful in determining whether much higher order coefficients need be calculated for unknown complex spectra. The sequence terminates with ρ_m if ρ_{m+1} equals zero, resulting in a maximum-entropy spectrum specified by m coefficients. An exception is $\rho_m = 1$, in which case the sequence terminates and the spectrum consists only of a pure set of m delta functions. (Note: ρ_m may often approach very close to one without equalling one).

(4) Finally, as can reasonably be expected, the variance of successive reflection coefficients becomes larger as the coefficient value becomes smaller [9]. Therefore, a sensible place to end the series would be a filter order at which this variance is on the verge of becoming a significant factor in obscuring the true reflection coefficient value. (Note: Several consecutive orders should be tested above the order in question to ensure it does not represent a local abnormality). Typically, this point occurs with m having an approximate value of 0.1 n.

Now, let us consider the manner in which these reflection coefficients are calculated.

2.3 Reflection Coefficient Algorithms

The adaptive nature of the PEF is a direct consequence of choosing

the filter coefficients to minimize a measure of the prediction error of the filter. The measure almost universally used is the least-squares error, which originated with Gauss in the early 1800's [1]. To satisfy this measure, the mean (or, more generally, the expected value) of the squared prediction error is minimized by adaptively controlling the coefficients. (This can be done by setting the partial derivative of the mean squared-error with respect to each coefficient equal to zero, and solving for that coefficient). In the case of the lattice-structure PEF, this is effected locally for each filter stage, with each reflection coefficient chosen to minimize the prediction error of that stage. However, as there are two prediction errors ($f_m(n)$ and $b_m(n)$, see Fig. 2.1) for each lattice filter stage, there are a number of different algorithms for minimizing the expectations of the squared prediction errors, depending on how the expectations of these values are combined. Six algorithms, having been suggested from various sources, are presented below.

To simplify the presentation of these algorithms, let us define the expectations of the squared value of the forward and backward prediction errors, respectively, as

$$F_m(n) = E[|f_m(n)|^2], \tag{2.9}$$

and

$$B_m(n) = E[|b_m(n)|^2]. \tag{2.10}$$

Another factor which occurs in all the algorithms is the expectation of the negative cross-power of the forward and backward prediction errors,

given by

$$C_m(n) = -E[f_m(n) \cdot b_m^*(n-1)]. \qquad (2.11)$$

(It should be noted that, although these algorithms were originally developed using only real-valued variables, they are broadened here to deal also with complex values [38]. The asterisk denotes complex conjugation).

(1) One of the simplest algorithms is the forward algorithm, which simply minimizes the forward mean square error, $F_m(n)$. This results in reflection coefficients with values of

$$\rho_m^F(n) = C_{m-1}(n) / B_{m-1}(n-1) \qquad (2.12)$$

(where the superscript F denotes the forward algorithm). This may seem an obvious choice, as it is generally only the forward prediction errors that are of interest in a particular application. However, it should be remembered that the backward prediction errors are still used for the calculation of higher filter orders. This algorithm has found favour with a number of researchers (e.g., [19]) who consider its results to be better than any other algorithm for steady-state signals. Actually, this algorithm gives generally better results only for signals whose amplitude is increasing with time, a fact which does not appear to have been realized by anyone else. (The typical "steady-state" signal is really a special case of this, with the increase all occurring at time zero). Thus, this algorithm would not be a good choice for most non-stationary signals. (An example of this behaviour will be given in Chapter 5).

(2) A similar algorithm is the backward algorithm, which minimizes the backward mean square error $B_m(n)$. For this algorithm, the reflection coefficients take the values

$$\rho_m^B(n) = C_{m-1}(n) / F_{m-1}(n) \tag{2.13}$$

(where the superscript B denotes the backward algorithm). This algorithm is much like the forward algorithm, except that it performs better for signals with decreasing amplitudes, and it does not seem to have found any favor.

(3) A method using both forward and backward algorithms (as presented above) was introduced by Griffiths [20]. This forward and backward algorithm (denoted F & B) uses different values of the reflection coefficient to calculate the forward and backward prediction errors of each lattice stage. (Thus, for this algorithm only, the forward and backward reflection coefficients are not complex conjugates of each other, as shown in Fig. 2.1(b)). The reflection coefficient used to give the forward prediction error is calculated using the forward algorithm, and the backward reflection coefficient is calculated using the backward algorithm.

All three of the above algorithms have a common theoretical problem. To ensure the stability of the filter, it is necessary that

$$|\rho_m(n)| \leq 1, \tag{2.14}$$

a condition that these three algorithms cannot generally meet. In practice, this does not appear to be a problem (as will be seen in Chapter 5). However, the following three algorithms all meet the condition in eq. (2.14).

(4) Makhoul [4] has shown that if the magnitude of either $\rho_m^F(n)$ or $\rho_m^B(n)$ is greater than one, the magnitude of the other is necessarily less than one. Based on this, he has suggested the minimum algorithm, which chooses for each stage either $\rho_m^F(n)$ or $\rho_m^B(n)$, depending on which is lesser. This is given by

$$\rho_m^M(n) = \frac{C_{m-1}(n)}{\max\left[F_{m-1}(n),\ B_{m-1}(n-1)\right]} \tag{2.15}$$

(where M denotes the minimum algorithm and max $[\cdot]$ takes the maximum of its arguments).

(5) Next, there is the geometric-mean algorithm, used first by Itakura and Saito [7], which uses the geometric mean of the values computed by using the forward and backward algorithms, as shown by

$$\rho_m^G(n) = \frac{C_{m-1}(n)}{\left[F_{m-1}(n) \cdot B_{m-1}(n-1)\right]^{\frac{1}{2}}} \tag{2.16}$$

(where G denotes the geometric-mean algorithm). This algorithm does not directly minimize any error criterion.

(6) Perhaps the commonest algorithm is the harmonic-mean algorithm developed by Burg [9]. This algorithm minimizes the sum of $F_m(n)$ and

$B_m(n)$, which is a reasonable error criterion for both stationary and non-stationary signals. This results in a reflection coefficient which is the harmonic mean of $\rho_m^F(n)$ and $\rho_m^B(n)$, as given by

$$\rho_m^H(n) = \frac{2\, C_{m-1}(n)}{F_{m-1}(n) + B_{m-1}(n-1)} \tag{2.17}$$

(where H denotes the harmonic-mean algorithm).

The last three algorithms above were shown by Makhoul [4] to be special cases of a more general formula,

$$\rho_m^r(n) = S \cdot \left[\tfrac{1}{2}\left(|\rho_m^F(n)|^r + |\rho_m^B(n)|^r\right)\right]^{1/r}, \tag{2.18}$$

where S is the sign of $C_{m-1}(n)$ and $r \leq 0$ (to guarantee stability). Specifically, $r \to 0$ results in $\rho_m^G(n)$, $r \to -\infty$ results in $\rho_m^M(n)$, and $r = -1$ results in $\rho_m^H(n)$. It is worth noting that, for a true steady-state signal, all six algorithms would converge to the same value.

To choose the algorithm which best suited our application, these algorithms were compared experimentally using simulated radar data. The results of these simulations are presented in Chapter 5. Based on those results, the harmonic-mean algorithm was chosen as the most promising, and in the following chapters (except Chapter 5) this algorithm alone is used to calculate the reflection coefficients $\rho_m(n)$. The implementation of this algorithm, as a practical adaptive system, is the subject of the next chapter.

CHAPTER 3

ADAPTIVE LATTICE IMPLEMENTATION METHODS

3.1 Introduction

The equations in section 2.3 calculate the lattice reflection coefficients in terms of theoretical "expected" values of the prediction error statistics. In the case of block processing (for a stationary signal), the transition from a theoretical model to the practical world merely involves replacing the expectations of the various terms with the averages (summations over n) of those terms for all available data samples.

For a continuously adaptive system, however, the transition is more complex. The reflection coefficients are now computed recursively (i.e., updated with each new sample rather than being totally recomputed). Starting with an initial (arbitrary) value, the reflection coefficient is updated by correction terms, which are calculated to provide the greatest reduction of prediction errors per correction unit. This is known as a steepest descent, or gradient approach [21]. Also, it is generally desirable to apply a weighting factor to the prediction error terms used in these calculations. This allows the filter to essentially "forget" the effects of samples beyond a certain distance or time (which may not represent current signal statistics), including any start-up

effects. It should be noted that weighting (or windowing) the prediction errors is not in any way equivalent to weighting the input samples, which would result in a stationary signal of reduced resolution.

3.2 Simple Gradient Method

Our first method of recursively calculating the reflection coefficients may be called the simple gradient method, due to the use of an approximation which results in a simpler implementation than the more common standard gradient method. The simple gradient method starts by considering the new coefficient as being the sum of the old coefficient and a correction term. This correction term is just the difference between the new and old values of the coefficient, as given by

$$\rho_{m+1}(n) - \rho_{m+1}(n-1) = \frac{-2 \sum_{i=1}^{n} [f_m(i) \cdot b_m^*(i-1)]}{\sum_{i=1}^{n} [|f_m(i)|^2 + |b_m(i-1)|^2]} + \frac{2 \sum_{i=1}^{n-1} [f_m(i) \cdot b_m^*(i-1)]}{\sum_{i=1}^{n-1} [|f_m(i)|^2 + |b_m(i-1)|^2]} \tag{3.1}$$

(Note the difference between n and n-1 as the limits on the summations). This equation can be rewritten as the sum of the old coefficient and a new update term which contains only information from the present time interval (i.e., the inputs to that filter stage), both multiplied by a third term. This results in the equation

$$\rho_{m+1}(n) - \rho_{m+1}(n-1) = \left[\frac{-2\, f_m(n) \cdot b_m^*(n-1)}{|f_m(n)|^2 + |b_m(n-1)|^2} - \rho_{m+1}(n-1)\right]$$

$$\cdot \quad \frac{|f_m(n)|^2 + |b_m(n-1)|^2}{\sum_{i=1}^{n} [|f_m(i)|^2 + |b_m(i-1)|^2]} \quad . \tag{3.2}$$

Rearranging equation (3.2) gives:

$$\rho_{m+1}(n) = [1-\gamma_m(n)]\, \rho_{m+1}(n-1) - \frac{2\gamma_m(n) \cdot f_m(n) \cdot b_m^*(n-1)}{|f_m(n)|^2 + |b_m(n-1)|^2} , \tag{3.3}$$

where

$$\gamma_m(n) = \frac{|f_m(n)|^2 + |b_m(n-1)|^2}{\sum_{i=1}^{n} [|f_m(i)|^2 + |b_m(i-1)|^2]} \quad . \tag{3.4}$$

It can be seen that for the steady state (constant power) case, $\gamma_m(n) \approx 1/n$ where n is the number of data samples processed. If, however, $\gamma_m(n) = \gamma$ is held constant in the calculation, then it may be replaced by using the weighting factor ω as defined in the formula

$$\omega = 1 - \gamma = 1 - 1/n'$$

$$\approx e^{-1/n'} \quad (\text{for } n' >> 0) , \tag{3.5}$$

where n' is the theoretical data adaptive length of the filtering action. (For $n' \geq 10$, the exponential form of (3.5) is less than 0.5 percent from the actual value). The resulting constant ω has a value in the range

$0 \leq \omega \leq 1$, with smaller values giving quicker adaptation. Rewriting equation (3.3) with ω gives:

$$\rho_{m+1}(n) = \omega \cdot \rho_{m+1}(n-1) + \alpha_m(n) \cdot f_m(n) \cdot b_m^*(n-1) , \qquad (3.6)$$

where $\rho_{m+1}(0) = 0$ (for the normal case) and the adaptive step size $\alpha_m(n)$ is given as

$$\alpha_m(n) = -2(1-\omega)/[\,|f_m(n)|^2 + |b_m(n-1)|^2\,] \quad . \qquad (3.7)$$

The recursive relationship in equation (3.6) can also be written as the sum:

$$\rho_{m+1}(n) = \sum_{i=1}^{n} [\omega^{(n-i)} \cdot \alpha_m(i) \cdot f_m(i) \cdot b_m^*(i-1)] \quad . \qquad (3.8)$$

An implicit condition on this recursive relationship is that the power of the prediction error $f_m(n)$ or $b_m(n)$ is not a time-varying function. This condition, as will be seen in Chapter 4, results in a somewhat poorer response for non-stationary signals, when compared with the standard gradient method. However, this shortcoming is partially offset by the desirable convergence properties, and the reduced storage and computation time requirements (less than the standard gradient method by 20% and 30%, respectively) of the simple gradient method.

The simple gradient method was first presented in 1979 by the author of this thesis [23], and the derivation and performance of the method were detailed in a paper [24] which followed. The computer implementation of this method can be found in Appendix A. The equations can

be easily modified to apply to any of the algorithms presented in section 2.3.

An interesting variation on the above method is to adjust the data adaptive length n' (and therefore ω) dynamically within the computation, so as to be proportional to the rate of change of the signal statistics. One way of doing this, which shall be called the adaptive gradient method, is to replace $\gamma_m(n)$ in equation (3.3) with the factor $(|\rho_m(n-1)-\rho_m'(n)|)/2$, where $\rho_m'(n)$ is the instantaneous estimate of $\rho_m(n)$, as given by

$$\rho'_{m+1}(n) = \frac{-2\ f_m(n)\ b_m(n-1)}{|f_m(n)|^2 + |b_m(n-1)|^2} \quad . \tag{3.9}$$

This substitution results in the reflection coefficients being calculated recursively by

$$\rho_m(n) = \rho_m(n-1) - \frac{(|\rho_m(n-1) - \rho_m'(n)|)\ (\rho_m(n-1) - \rho_m'(n))}{2} \quad . \tag{3.10}$$

(Some results from the adaptive gradient method are presented in Chapter 6).

3.3 Standard Gradient Method

We will call the second method of recursively calculating the reflection coefficients the standard gradient method because it is the method that is generally arrived at by applying gradient techniques to the lattice equations [10, 11, 12]. This is a more rigorous approach, which avoids the constant power assumption of the simple gradient method, at the expense of increased complexity of operation and performance.

The standard gradient method consists of replacing the expectations of the various terms, used for the reflection coefficient calculation, with corresponding averages in which an exponential weighting has been added to the individual values. This weighting allows the characteristics of older data samples to be progressively forgotten, giving greater importance to current data. The reflection coefficients thus calculated are given by

$$\rho_{m+1}(n) = \frac{-2 \sum_{i=1}^{n} [\mu^{(n-i)} \cdot f_m(i) \cdot b_m^*(i-1)]}{\sum_{i=1}^{n} \mu^{(n-i)} \cdot [|f_m(i)|^2 + |b_m(i-1)|^2]} , \qquad (3.11)$$

where μ is a weighting constant, normally in the range $0 < \mu \leq 1$. This is equivalent to weighting the forward and delayed backward prediction errors by $\mu^{(n-i)/2}$. (As noted in the previous section, this weighting has no effect on the stationarity of the input). Small values of μ result in quick adaptation to new signal characteristics, while values of μ near one result in slower convergence and the reduction of noise (by integration).

The exponential form of the weighting function is chosen because it can be applied as part of a simple recursive implementation. Thus, with each new sample, a new value for $\rho_{m+1}(n)$ is calculated from

$$\rho_{m+1}(n) = \frac{v_{m+1}(n)}{y_{m+1}(n)} , \qquad (3.12)$$

where

$$v_{m+1}(n) = \mu \cdot v_{m+1}(n-1) - 2\, f_m(n) \cdot b_m^*(n-1) \tag{3.13}$$

and

$$y_{m+1}(n) = \mu \cdot y_{m+1}(n-1) + |f_m(n)|^2 + |b_m(n-1)|^2 \quad . \tag{3.14}$$

(In the normal case, the initial conditions are $v_{m+1}(0) = y_{m+1}(0) = 0$). This method may be easily modified to apply to any of the algorithms presented in section 2.3. The computer implementation of this method is presented in Appendix A.

A number of researchers (e.g.: [10, 20, 22]) have derived this gradient equation in a different form, generalized here to apply to complex data as:

$$\rho_{m+1}(n+1) = \rho_{m+1}(n) - \frac{f_m(n)\, b_{m+1}^*(n) + f_{m+1}(n)\, b_m^*(n-1)}{y_{m+1}(n)} \quad . \tag{3.15}$$

By using equations (2.1) and (2.2) to substitute for the values of $f_{m+1}(n)$ and $b_{m+1}(n)$ in equation (3.15), this equation takes a form similar to that given by equation (3.12).

Another derivation for the standard gradient method results from the application of Kalman filtering theory [21, 25] to the lattice structure. Because each stage of the lattice is optimized independently, the Kalman filtering equations are applied to a single stage. The basis for this is to rewrite equation (2.1) and the complex conjugate of equation (2.2) in the form of the Kalman measurement equation for that stage,

$$\begin{bmatrix} f_{m-1}(n) \\ b^*_{m-1}(n-1) \end{bmatrix} = \begin{bmatrix} b_{m-1}(n-1) \\ f^*_{m-1}(n) \end{bmatrix} \cdot \begin{bmatrix} -\rho_m(n) \end{bmatrix} + \begin{bmatrix} f_m(n) \\ b^*_m(n) \end{bmatrix} \quad . \qquad (3.16)$$

When the Kalman equations are used to minimize the measurement noise vector $[f_m(n), b^*_m(n)]$, the resulting solution is the harmonic mean algorithm using the standard gradient method (eq. (3.11) to (3.15)) with $\mu = 1$ (the Kalman approach does not allow for an adaptive constant). If only $f_m(n)$ is used for the measurement noise, the solution becomes the forward algorithm, and likewise using $b_m(n)$ results in the backward algorithm. The application of Kalman filtering theory to the lattice structure is developed in detail in Appendix B.

3.4 Comparisons with Other Methods

At this point, it would be useful to compare the above two methods of implementing the lattice structure adaptively, with algorithms proposed by other researchers. There are two methods in particular which deserve attention: the modified least-mean square (LMS) approach of Griffiths [13], and a least-squares method developed by Morf, et. al. [26, 27].

Griffiths [13] has proposed the use of an algorithm, which is generalized to apply to complex data, as shown by

$$\rho_m(n+1) = \rho_m(n) - \frac{\alpha}{\sigma^2_m(n)} \left[f_m(n) \cdot b^*_{m-1}(n-1) + f_{m-1}(n) \cdot b^*_m(n) \right] , \qquad (3.17)$$

where α is an adaptation parameter, and $\sigma^2_m(n)$ is the power estimate at the m^{th} stage given by

$$\sigma_m^2(n) = \beta \cdot \sigma_m^2(n-1) + (1-\beta)\left[|f_{m-1}(n)|^2 + |b_{m-1}(n-1)|^2\right] \quad . \tag{3.18}$$

Using equations (2.1) and (2.2) to replace $f_m(n)$ and $b_m(n)$ with $f_{m-1}(n)$ and $b_{m-1}(n-1)$ gives:

$$\rho_m(n) = \left(1 - \frac{\alpha \cdot \left[|f_{m-1}(n)|^2 + |b_{m-1}(n-1)|^2\right]}{\sigma_m^2(n)}\right)\rho_m(n-1)$$

$$- \frac{2\alpha}{\sigma_m^2(n)}\left[f_{m-1}(n) \cdot b_m^*(n-1)\right] \quad . \tag{3.19}$$

The novel aspect of Griffiths' algorithm is the introduction of the second adaptation parameter β. Two values of β are of special interest, namely, $\beta = 0$ and $\beta = 1 - \alpha$. Setting $\beta = 0$ results in equation (3.19) taking on the same form as equation (3.6) (where $\omega = 1 - \alpha$) which is the simple gradient as described previously. On the other hand, setting $\beta = 1 - \alpha$ results in equation (3.19) being identical with the standard gradient method (equations (3.12-14) , where $\mu = (1 - \alpha)/\alpha$). Other values of β may be used in the range $0 \leq \beta \leq 1 - \alpha$ ($\beta > 1 - \alpha$ is unstable). However, there seems little justification in doing so, as $\beta = 1 - \alpha$ has the strongest theoretical basis, and there are no computational improvements until $\beta = 0$.

The method recently developed by Morf and Lee [26, 27] is generally known as the "least-squares" method (although all the methods presented in this thesis minimize a measure of the least-squares error). A detailed derivation of this method is presented in [28]. Generalizing

Morf's method to complex data and summing the recursions of Morf's equations results in the forward reflection coefficient (used to calculate the forward prediction error) having the value

$$\rho^{F}_{m+1}(n) = \frac{-\sum_{i=1}^{n} \left[\frac{\mu^{n-i}}{1-\gamma_m(i-1)} \; f_m(i)\; b^*_m(i-1) \right]}{\sum_{i=1}^{n} \left[\frac{\mu^{n-i}}{1-\gamma_m(i-1)} \; |b_m(i-1)|^2 \right]} , \tag{3.20}$$

and the backward coefficient having the value

$$\rho^{B}_{m+1}(n) = \frac{-\sum_{i=1}^{n} \left[\frac{\mu^{n-i}}{1-\gamma_m(i-1)} \; f_m(i)\; b^*_m(i-1) \right]}{\sum_{i=1}^{n} \left[\frac{\mu^{n-i}}{1-\gamma_m(i-1)} \; |f_m(i)|^2 \right]} , \tag{3.21}$$

where

$$\gamma_{m+1}(n) = \gamma_m(n) + \frac{|b_m(n)|^2}{\sum_{i=1}^{n} \left[\frac{\mu^{n-i}}{1-\gamma_m(i)} \; |b_m(i)|^2 \right]} \tag{3.22}$$

and $\gamma_o(n) = 0$. (In one version [26] of this equation, the summation in (3.22) is multiplied by $1-\mu$). This method, as presented by Morf, et. al., is an implementation of the forward and backward algorithm for reflection coefficient calculation (see section 2.3). However, only one value of $\gamma_m(n)$ is used for both forward and backward coefficients at each order, that value being in this case calculated from only the backward prediction errors. It is suggested in some papers (e.g., [27]) that the equation (3.22) could also be written as

$$\gamma_{m+1}(n) = \gamma_m(n-1) + \frac{|f_m(n)|^2}{\sum_{i=1}^{n} \left[\frac{\mu^{n-i}}{1-\gamma_m(i-1)} \; |f_m(i)|^2 \right]} \qquad (3.23)$$

However, this equivalence would hold only for a truly stationary input (for large n). Also, in implementing the forward and backward algorithm, this method fails to ensure the stability of the resulting reflection coefficients.

Morf's "least-squares" method can be easily applied to implement any of the reflection coefficient algorithms in section 2.3. For instance, the harmonic-mean algorithm, implemented by Morf's method, becomes

$$\rho_{m+1}(n) = \frac{-2 \sum_{i=1}^{n} \left[\frac{\mu^{n-i}}{1-\gamma_m(i-1)} \; f_m(i)\, b_m^*(i-1) \right]}{\sum_{i=1}^{n} \left[\frac{\mu^{n-i}}{1-\gamma_m(i-1)} \; \{|f_m(i)|^2 + |b_m(i-1)|^2\} \right]} , \qquad (3.24)$$

where $\gamma_m(n)$ is now given by

$$\gamma_{m+1}(n) = \gamma_m(n-1) + \frac{|f_m(n)|^2 + |b_m(n-1)|^2}{\sum_{i=1}^{n} \left[\frac{\mu^{n-i}}{1-\gamma_m(i-1)} \; \{|f_m(i)|^2 + |b_m(i-1)|^2\} \right]} . \qquad (3.25)$$

Comparing equation (3.24) with the corresponding equation (3.11) for the standard gradient method of implementation, it can be seen that the only difference is the introduction in Morf's method of an additional weighting term, $[1-\gamma_m(i-1)]^{-1}$. In fact, as $\gamma_o(n) = 0$, the calculation for the first reflection coefficient $\rho_1(n)$ is identical. (Note also the similarity of

$\gamma_m(n)$, as given by equation (3.25) to the corresponding value in equation (3.11), used for the derivation of the simple gradient method).

The term $\gamma_m(n)$ deserves careful attention. For a true steady-state signal ($|x(n)|^2 = E[|x(n)|^2]$), we can show that

$$\gamma_m(n) \to 1 - \mu^m \quad \text{as } n \to \infty \quad . \tag{3.26}$$

For lower values of n, the value of $\gamma_m(n)$ would be larger than that given in equation (3.26), starting at n = 1 with $\gamma_m(1) = 1$ and quickly converging towards the value given above. As n increases, the $\gamma_m(n)$ terms in equations (3.20), (3.21), and (3.24) become constant and cancel out, so that Morf's method approaches the standard gradient method, becoming identical as $n \to \infty$. Thus, the only behavioural difference between the two methods for steady-state signals would be a somewhat quicker convergence by Morf's method during the initial "start-up" period of the filter, and this only if the reflection coefficients were given initial, non-zero estimates to start from.

Let us now examine the behaviour of this method for signals with non-stationary or changing statistics. (We will assume that the filter has run long enough to "forget" the initial start-up period mentioned above). Deviating from the derivation of equation (3.26), we can make a good approximation of $\gamma_1(n)$ as

$$\gamma_1(n) \simeq \frac{(1-\mu)\ |x(n)|^2}{E[|x(n)|^2]} \quad , \tag{3.27}$$

with similar approximations for higher orders. (Note: $E[|x(n)|^2]$ is

calculated such that $\gamma_1(n)$ is always ≤ 1). Morf and Lee [26] state that $\gamma_m(n)$ is a log-likelihood variable related to the input, in that it is a measure of the deviation of successive samples from a Gaussian distribution, with $\gamma_m(n)$ tending to one for non-Gaussian components. However, it is clear from equation (3.27) that $\gamma_1(n)$ is actually responding to deviations in the power of the signal from its expected value. In fact, if the change in signal statistics includes an increase in signal power, the result will be an increase in the value of $\gamma_1(n)$, (and quicker convergence to the new statistics). If the change includes a decrease in power, $\gamma_1(n)$ will decrease in value (resulting in an undesirable, slower-than-normal convergence). Similar conditions apply to higher filter orders. $\gamma_m(n)$ is also quite sensitive to noise. Based on this examination, the validity of the statement of Morf and Lee given above is questionable.

A better indicator of signal non-stationarity, based on the forward and backward algorithm, was suggested by Griffiths [20]. The forward and backward reflection coefficients for each state have the same optimum values for a stationary input, moving apart as the input becomes less stationary. Griffiths' indicator of signal change is formed of the inner product between a vector with components $\rho_m^F(n)$ and a vector with components $\rho_m^B(n)$.

Satorius and Pack [8] have compared Morf's method to the standard gradient method and reported the convergence of Morf's method to be slightly better. However, they used the harmonic-mean algorithm for the standard gradient method (as compared to the forward and backward algorithm

for Morf's method), and preloaded the lattice using Morf's method while only giving the gradient lattice an initial estimate [29]. As will be seen in the following chapters, the forward and backward algorithm in general converges better than other algorithms (although having some undesirable characteristics), and preloading the lattice greatly reduces the convergence time, thus rendering the comparison of Satorius and Pack next to meaningless. A more recent comparison by Honig and Messerschmitt [44] finds no discernable differences in performance between the two methods, for stationary signals.

In fact, there appear to be no advantages to Morf's method significant enough to warrant the increased complexity involved. (For cases where quicker adaptation to signal changes is required, the adaptive gradient method discussed in section 3.2 would probably be a better choice). Finally, Morf and Lee [26] describe their method as an "exact solution" to the least-squares problem while the gradient methods are "approximations". Considering the above analysis, this statement is unjustifiable.

CHAPTER 4

LEARNING CHARACTERISTICS OF THE ADAPTIVE LATTICE FILTER

4.1 First-Order Convergence Properties for the Simple Gradient Method

In Chapter 3, two major recursive methods were developed for continuously updating the calculation of the reflection coefficients of the adaptive lattice-structure PEF. These methods allow the lattice to adapt to a signal by progressively learning its statistical properties in a controlled manner. In order to use these methods properly, it is necessary to understand these learning characteristics. In this section, we shall begin by examining the primary characteristics of the simple gradient method, as applied to the harmonic-mean lattice algorithm.

An important characteristic of the adaptive filter is the rate at which the reflection coefficients converge to their optimum values for given (stationary) input signal statistics. This rate of convergence is controlled by the adaptive weighting factor (ω). A small value of this factor results in a quickly converging filter which is sensitive to momentary fluctuations in signal statistics (noise), while a large value results in slower convergence and integrates out the noise. A quantitative evaluation of this relationship for the simple gradient method is developed below, as applied to the first stage of the filter. Higher filter orders have similar convergence properties, but they must first decouple from lower orders of the filter. This decoupling action will

be discussed in section 4.3.

The instantaneous estimate of the first reflection coefficient at time n can be defined as

$$\rho_1'(n) = \frac{-2 \cdot f_0(n) \cdot b_0^*(n-1)}{|f_0(n)|^2 + |b_0(n-1)|^2} \quad . \tag{4.1}$$

Combining equations (3.6), (3.7), and (4.1) results in the recursive relationship:

$$\rho_1(n) = \omega\rho_1(n-1) + (1-\omega)\ \rho_1'(n) \quad . \tag{4.2}$$

For a truly stationary process beginning at time n = 0, the instantaneous estimation of equation (4.1) for n > 1 will, in fact, be equal to the optimum (asymptotic) value of the reflection coefficient, $\hat{\rho}_1$. Using this fact and given the initial value of the reflection coefficient $\rho_1(0)$, (for example, the filter's start-up values, or the value for a previous time series to which the filter has adapted), the filter's convergence equation can be computed by repeated application of the recursion equation (4.2) as

$$\begin{aligned} \rho_1(n) &= \omega^n\ \rho_1(0) + (1-\omega)\ \omega^{n-1}\ \rho_1' + (1-\omega) \sum_{i=0}^{n-2} [\omega^i\ \hat{\rho}_1] \\ &= \omega^n \cdot \rho_1(0) + (1-\omega)\ \omega^{n-1}\ \rho_1'(1) + (1-\omega^{n-1})\ \hat{\rho}_1 \quad . \end{aligned} \tag{4.3}$$

From this, the fractional error in the reflection coefficient at time n can be computed as

$$\varepsilon_1(n) = \frac{\hat{\rho}_1 - \rho_1(n)}{\hat{\rho}_1} = \omega^{n-1} \left[1 - \frac{\omega\, \rho_1(0) + (1-\omega)\, \rho_1'(1)}{\hat{\rho}_1} \right] \qquad (4.4)$$

The factor $\rho_1'(1)$ need not be known for most practical applications of this filter. Indeed, for the initial start-up case where $\rho_1(0) = x(0) = 0$, we have $\rho_1'(1) = 0$, resulting in the simplified versions

$$\rho_1(n) = (1-\omega^{n-1}) \cdot \hat{\rho}_1 \qquad (4.5)$$

and

$$\varepsilon_1(n) = \omega^{n-1} \qquad (4.6)$$

for equation (4.3) and (4.4), respectively. Such a case is plotted in Fig. 4.1 (along with the standard gradient method), where the filter is adapting to an ideal signal consisting of a single frequency. The exponential nature of this curve is easily observed.

For the transition case where $\rho_1(0)$ is known but not equal to zero, given values of ω approaching unity (which is the common case), and therefore $\rho_1'(1) \approx \rho_1(0)$, equations (4.3) and (4.4) can be simplified as follows, respectively:

$$\rho_1(n) \approx (1-\omega^{n-1})\, \hat{\rho}_1 + \omega^{n-1}\, \rho_1(0) \qquad (4.7)$$

and

$$\varepsilon_1(n) \approx \omega^{n-1} \cdot (1-\rho_1(0)/\hat{\rho}_1). \qquad (4.8)$$

This measure of convergence error can also be written in terms of

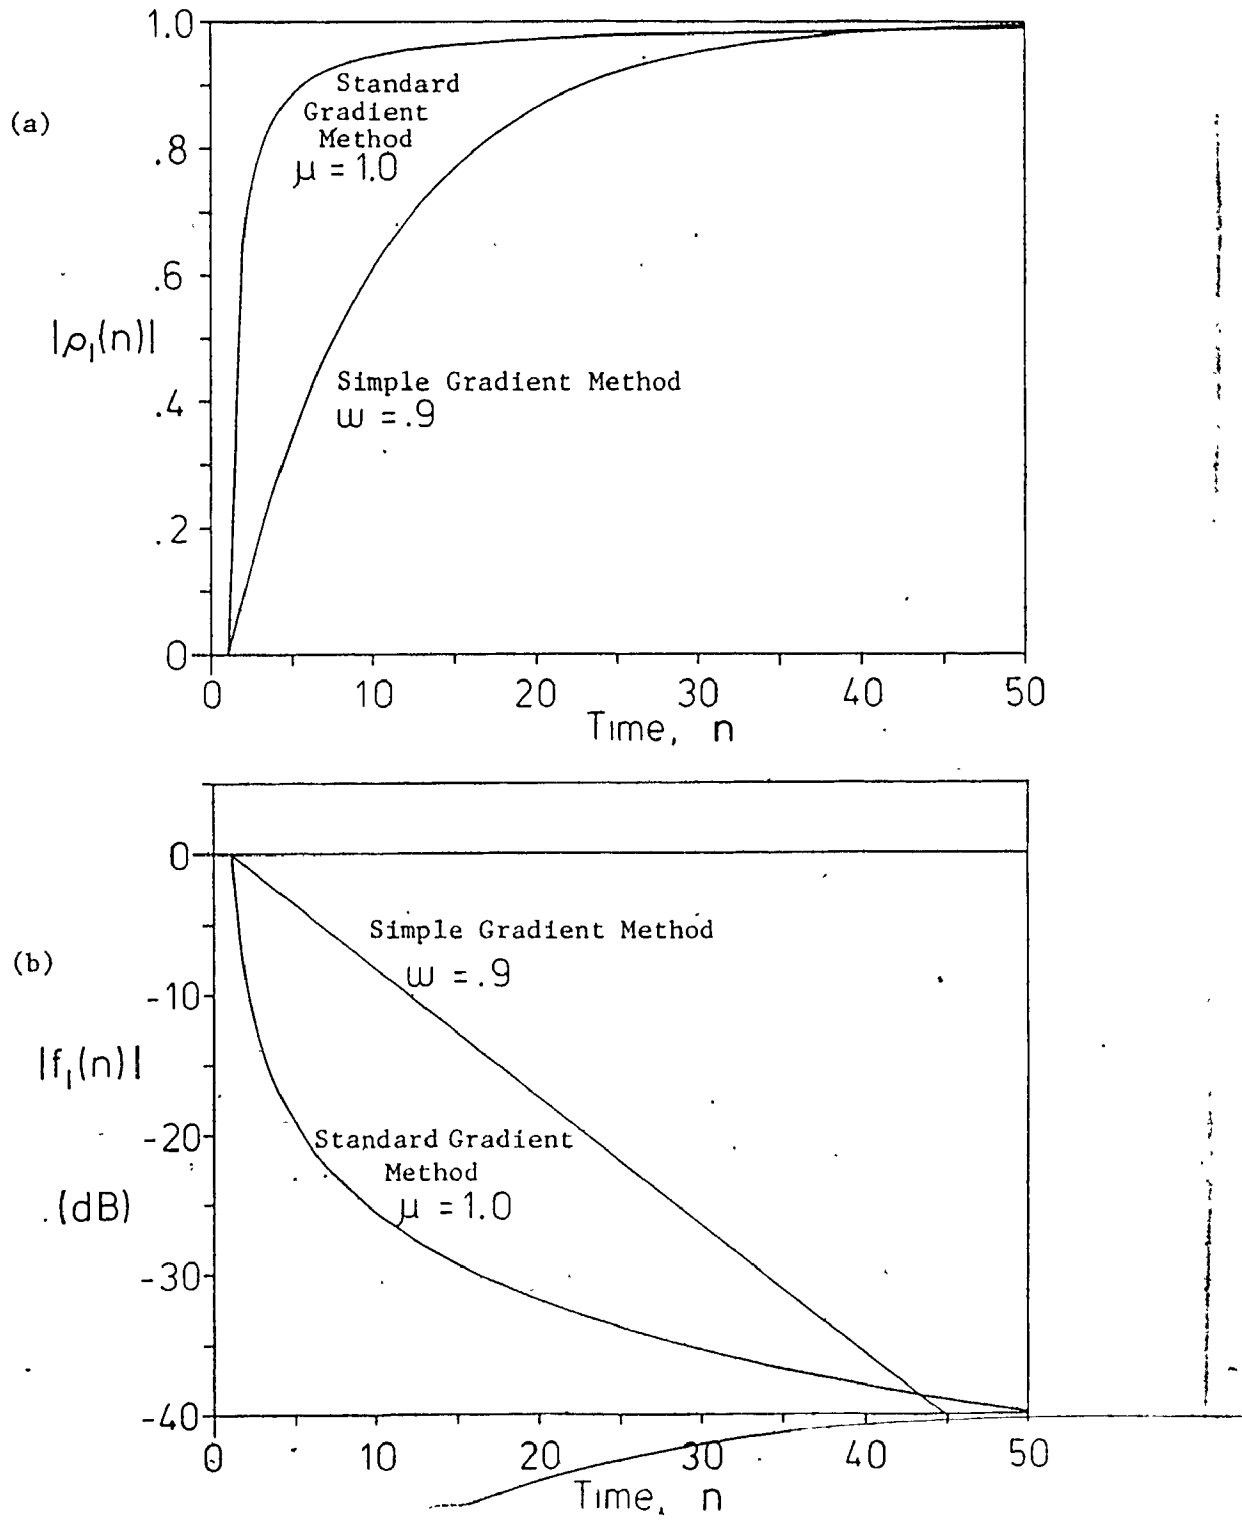

Figure 4.1: Convergence Behaviour of First Filter Order (with ideal signal)
(a) Reflection Coefficient, (b) Output

the ratio of the data length actually processed to the theoretical data adaptive length n' by applying equation (3.5). Thus, for example, equation (4.1) becomes

$$\varepsilon_1(n) = e^{-(n-1)/n'}; \text{ for } n > 0 \text{ and, } n' >> 0 . \qquad (4.9)$$

Although these convergence formulae are developed only for the reflection coefficients, the component of the signal to which the filter stage is responding will show the same convergence (in reverse) at the output. This can be seen in Fig. 4.1(b), where we have plotted the output of the first stage of the filter (corresponding to the reflection coefficient in Fig. 4.1(a)) adapting to an ideal signal consisting of a single frequency. The output converges exponentially (log-scale) from the initial input value to a value approaching zero. (The minimum value is set by the noise level in the signal).

4.2 First-Order Convergence Properties for the Standard Gradient Method

In this section, we shall examine the convergence properties of the standard gradient method governing the adaptation of the first stage of the filter to a stationary input. (This development parallels that of the simple gradient method in the previous section). The rate of convergence is controlled by an adaptive weighting factor (μ). Small values of μ give quicker convergence, and larger values result in less sensitivity to noise statistics.

As with the development of the simple gradient method, instantaneous estimates can be made for the numerator and demoninator terms used in the

calculation of the first reflection coefficient by the standard gradient method. For a truly stationary process, these estimates $v_1'(n)$ and $y_1'(n)$, as defined by

$$v_1'(n) = -2f_0(n) \cdot b_0^*(n-1) \tag{4.10}$$

and

$$y_1'(n) = |f_0(n)|^2 + |b_0(n-1)|^2, \tag{4.11}$$

are equal to the optimum values $\hat{v}_1$ and $\hat{y}_1$ (for $n > 1$), such that $\hat{v}_1/\hat{y}_1 = \hat{\rho}_1$. Combining equations (3.13) with (4.10) and (3.14) with (4.11) gives the recursion relationships:

$$v_1(n) = \mu v_1(n-1) + v_1'(n) \tag{4.12}$$

and

$$y_1(n) = \mu y_1(n-1) + y_1'(n) \tag{4.13}$$

Repeated applications of these recursions result in the following formula for the reflection coefficient:

$$\rho_1(n) = \frac{v_1(n)}{y_1(n)} = \frac{\sum_{i=0}^{n-2} [\mu^i \cdot \hat{v}_1] + \mu^{n-1} \cdot v_1'(1) + \mu^n \cdot v_1(0)}{\sum_{i=0}^{n-2} [\mu^i \cdot \hat{y}_1] + \mu^{n-1} \cdot y_1'(1) + \mu^n \cdot y_1(0)} \tag{4.14}$$

or, for $\mu \neq 1$,

$$\rho_1(n) = \frac{\frac{1-\mu^{n-1}}{1-\mu} \cdot \hat{v}_1 + \mu^{n-1} \cdot v_1'(1) + \mu^n \cdot v_1(0)}{\frac{1-\mu^{n-1}}{1-\mu} \cdot \hat{y}_1 + \mu^{n-1} \cdot y_1'(1) + \mu^n \cdot y_1(0)} \tag{4.15}$$

These equations are difficult to simplify significantly, except for the initial start-up case where $v_1(0) = y_1(0) = x(0) = 0$. Then, $b_0(0) = 0$, and therefore $v_1'(1) = 0$, and $y_1'(1) = |f_0(1)|^2$. In a stationary environment, the forward and backward prediction-error powers are equal, so that $y_1'(1) = \hat{y}_1/2$, simplifying equation (4.15) to

$$\rho_1(n) = \frac{\frac{1-\mu^{n-1}}{1-\mu} \cdot \hat{v}_1}{\left(\frac{1-\mu^{n-1}}{1-\mu} + \frac{\mu^{n-1}}{2}\right) \cdot \hat{y}_1} = \frac{2 \cdot (1-\mu^{n-1})}{(2-\mu^{n-1}-\mu^n)} \cdot \hat{\rho}_1 \qquad (4.16)$$

From this, the fractional error in the reflection coefficient at time n can be computed as

$$\varepsilon_1(n) = \frac{\hat{\rho}_1 - \rho_1(n)}{\hat{\rho}_1} = \frac{\mu^{n-1} - \mu^n}{2 - \mu^{n-1} - \mu^n} \qquad (4.17)$$

Another special case of interest is when $\mu = 1$, for which equation (4.14) simplifies to

$$\rho_1(n) = \frac{(n-1) \cdot \hat{v}_1 + v_1'(1) + v_1(0)}{(n-1) \cdot \hat{y}_1 + y_1'(1) + y_1(0)} \qquad (4.18)$$

For the initial start-up case as described above, this simplifies further to

$$\rho_1(n) = \frac{(n-1)\ \hat{v}_1}{(n-1/2)\ \hat{y}_1} = \frac{2n-2}{2n-1}\ \hat{\rho}_1 \qquad (4.19)$$

The corresponding value of the fractional error in $\rho_1(n)$ then becomes

$$\varepsilon_1(n) = \frac{\hat{\rho}_1 - \rho_1(n)}{\hat{\rho}_1} = \frac{1}{2n-1} \quad . \tag{4.20}$$

As with the simple gradient method, these convergence rates can also be applied to the relevant signal component at the filter stage output. The reflection coefficients and filter output for the case above are plotted in Fig. 4.1. In comparing the response to an ideal signal (consisting of a single frequency) shown in these figures for the two methods here described, two observations are worth noting. First, the standard gradient method does not exhibit the simple exponential behaviour observed in the simple gradient method. Second, although initially the standard gradient method converges more quickly than the simple gradient method, later on (at time n = 43 in this example), the simple gradient method catches up with and begins to converge faster than the standard gradient method.

Finally, though no theoretical justification can be given at this time, it would be reasonable to assume that equations (4.17) and (4.20) could be multiplied by the factor $(1 + \rho_1(0)/\hat{\rho})$ for the more general case when $\rho_1(0) \neq 0$. This assumption is based on the convergence results for the simple gradient method.

4.3 Decoupling and Higher-Order Convergence

Determination of the convergence behaviour of higher-order stages of the filter is a more complex procedure than that detailed above for the first stage. This is because each successive stage must "decouple" from

the preceding stage before it can begin to converge to its optimum value. Both the time taken to decouple and the starting values of the reflection coefficients are very much dependent on the input signal, as well as the action of all preceding stages.

The decoupling action proceeds as follows: Each stage of the filter attempts to minimize its error output by adapting to the most prominent spectral component of the input signal for that stage (the output of the preceding stage). Thus, when a new signal is presented to the filter, the first stage starts to adapt to the most prominent spectral component of that signal. At the same time, because the first stage has not yet started to filter it out, the second stage (and higher orders) also is attempting to adapt to the same component of the signal. However, as that component is continually changing due to the improving action of the first stage, proper adaptation by the second stage is impaired. This continues until a point at which the first stage has adapted sufficiently to that signal component so that it is no longer the most prominent spectral component in the output of the first stage. As further adaptation by this stage has little overall importance with respect to the signal as perceived by the following stage, the first stage can be said to be "adapted" at that point (although further adaptation takes place). The second stage then "decouples" from the first (i.e., its action becomes independent of it) and begins adapting to the next most prominent component remaining in the signal after filtering by the first stage. Higher order stages of the filter decouple in the same fashion.

A good example of this decoupling action is shown in Fig. 4.2,

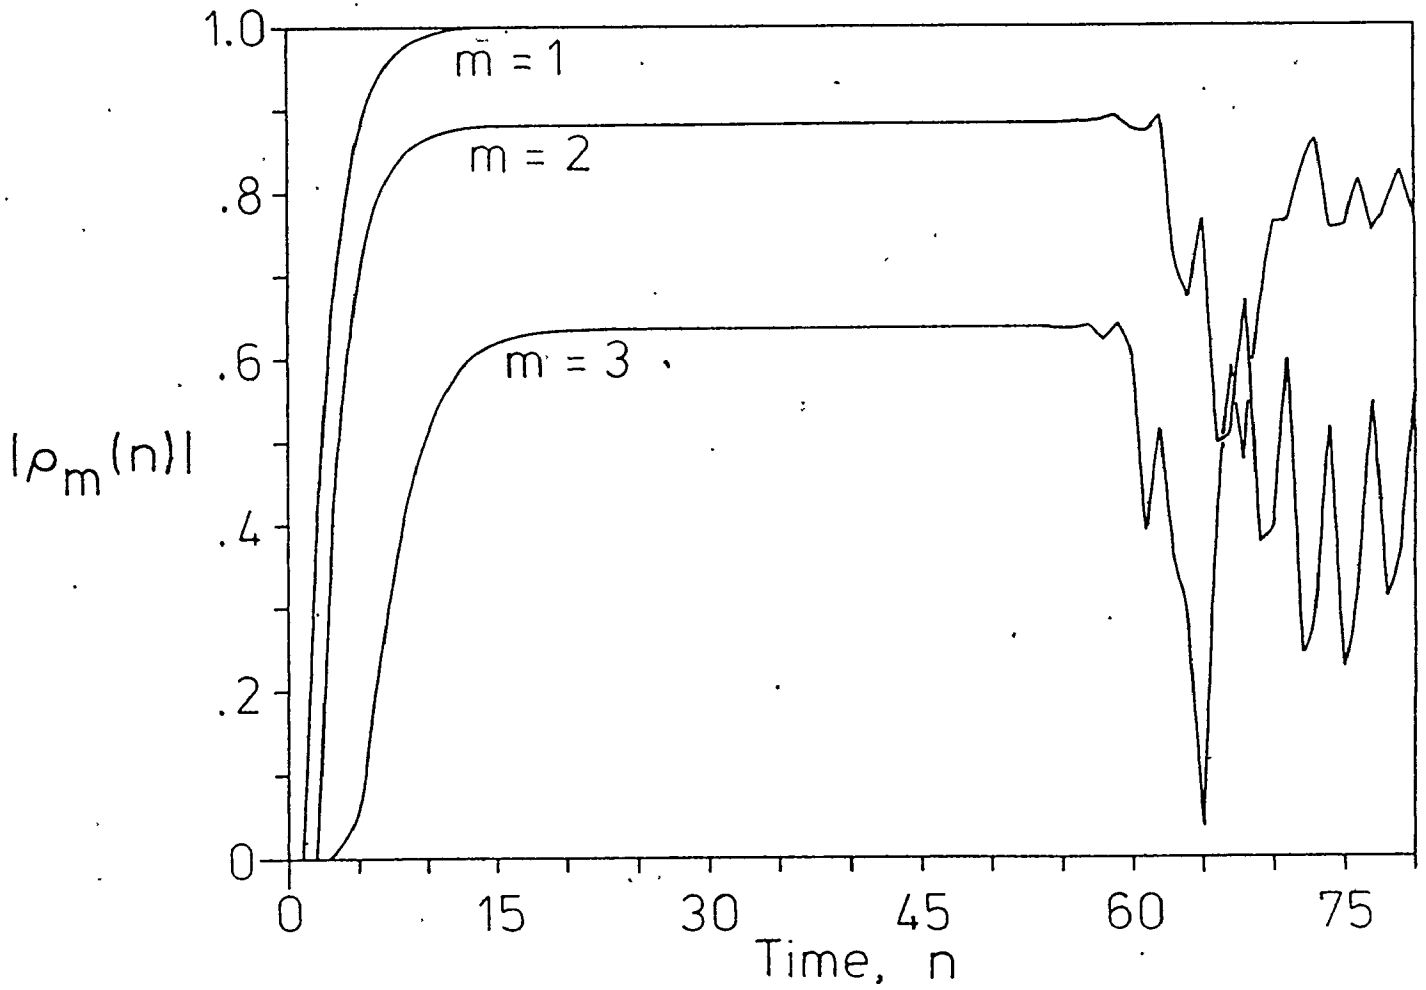

— Figure 4.2: Decoupling of Higher Filter Orders for a Signal Consisting of a Single Frequency (Simple Gradient Method, $\omega = .6$)

where the filter (simple gradient method) is adapting to an ideal signal consisting of only one frequency. As can be seen, all three filter orders (as well as higher orders) initially attempt to adapt to the same signal. This goes on until a point (at time $n \approx 60$) where the first coefficient has reduced the single frequency to a level below the quantization noise of the system, at which point the higher order stages become erratic as they respond to that noise. (At $n = 60$, the fractional error in ρ_1 is $\varepsilon_1 = 6.3 \times 10^{-14} = 2^{-44}$. As the computer used had a nominal accuracy of 48 bits, an accuracy of 44 bits is very reasonable for the trigometric functions involved).

After decoupling, the convergence of the higher-order stages follows approximately the same formulae as the convergence of the first stage (for example, n-1 for the first stage in equations (4.5) to (4.9) is replaced by n-2 for the second stage, and so on). However, a complicating factor in this convergence is that the filtering action of the first stage can, and often will, affect to some degree signal components other than the one to which it has adapted. This results in residual coupling after the major decoupling has taken place, which will impair to some extent the convergence behaviour of the higher stages resulting in increased convergence time [20]. Under such conditions, the filter stages may never become completely orthogonal (decoupled) - although in a practical sense they may be considered as such.

Obviously from the above discussion, besides being dependent on the filter's adaptive weighting constant, the time taken to decouple the second stage from the first is very much dependent on the characteristics

of the input signal. This stage decoupling time (τ) for a specific signal can be defined for the simple gradient method as

$$\tau = an' \quad , \tag{4.21}$$

where a is a dimensionless decoupling constant unique to the specific type of signal in question. Using this definition, the adaptive weighting factor of the filter (equation (3.5)) can be rewritten as

$$\omega = e^{-a/\tau} \quad . \tag{4.22}$$

(There is no similar development for the standard gradient method).

The above relationship of ω and τ has a very useful application in selective filtering. In such cases, it may be desired to filter out interference in the form of a continuous signal component while not filtering out information in the form of a temporary signal component, or to identify a signal of known duration by its reflection coefficients while providing optimum rejection of short-term statistical deviations (noise). This generally can be done by setting the relaxation time subject to the constraints

$$t_a << \tau << t_b \quad , \tag{4.23}$$

where t_a is the maximum duration of the signal not to be adapted to and t_b is the minimum duration of the signal to be adapted to. (τ may be given any value in this range. However, as the filter order increases, values closer to t_a are to be preferred to give higher filter orders a chance to adapt).

Knowing the value of the decoupling constant 'a' for a given

signal type (determined experimentally), the filter can be set to accept or reject signals of any desired duration by the use of equation (4.22) with the desired value of τ. (This subject will be dealt with further in the following chapters).

It may seem at first that the total relaxation time of the filter, not the stage relaxation time, should be used in equation (4.23), but two considerations argue against this for many types of signals such as intermittent information in interference (for example, in a radar environment the signal of interest may consist of a target embedded in a background of clutter). First, the short-term signal is often responsible for the most prominent signal spectrum component while it is present. Therefore, the first stage is the most important consideration as it tries to adapt to the short-term signal and $t_a << \tau$ is desirable. Second, many commonly encountered continuous signals, especially of the interference type, can be modelled as a coloured-noise process whose spectrum has essentially a single peak with a Gaussian distribution. In such cases, it is predominantly the first stage which filters the signal out while successive stages deal with the remaining residuals, as was discussed in section 2.2. Thus, for $\tau < t_b$, that signal will essentially be filtered out, although not perfectly.

However, there will be exceptions to this rule where it may be desirable to use a smaller value t_b/m' in place of t_b in equation (4.23) to allow the higher order stages of the filter to adapt to the signal as well. The constant m' may assume any of a number of values, generally ranging from 1 (as above) to m (the filter order) or even above for

cases where residual coupling increases the higher-order convergence and decoupling time. It should be noted, however, that the lower the value of m', the better the rejection of short-term signal variations as factors in the calculation of the reflection coefficients.

As noted previously, τ may have any value in the range covered by equation (4.23). Two such values are the arithmetic and geometric means of t_a and t_b/m', given respectively by

$$\tau_a = (t_a + t_b/m')/2 \tag{4.24}$$

and

$$\tau_g = \left[\frac{t_a \cdot t_b}{m'}\right]^{\frac{1}{2}} . \tag{4.25}$$

One or other of these values should generally prove to be a good compromise in the choice of τ. The arithmetic mean has proved closer to the experimentally optimum value in our studies.

Before we leave this discussion of stationary learning characteristics to examine some non-stationary characteristics of the lattice structure filter, there are several other factors affecting convergence which are worth mentioning. One area of interest is that of reducing the convergence time of the filter by pre-setting the initial values of the reflection coefficients (for example, to their first instantaneous estimates), or by pre-loading the filter with data before allowing adaptations to begin. This is especially advantageous when the filter is set for a long convergence time. This procedure assumes knowledge beforehand as to the starting

point of the signal which is to be adapted to, an assumption which may not be valid for many circumstances. (A comparison of the filter behaviour when pre-loaded may be found in Chapter 5). Other factors which may affect the convergence behaviour of these filters include low signal-to-noise ratios, gradient noise, and quantization noise. These factors deserve further attention but were not examined for this thesis.

4.4 Non-Stationary Learning Characteristics

We have thus far discussed the response of the filter to a stationary signal; however, for a non-stationary signal, the calculations become much more complex. Obviously, the reflection coefficients could never hope to achieve their optimum values. How close they come to these values becomes a trade-off between the error introduced by the filter lagging behind in the adaptive process (i.e., giving too much weight to older, out-of-date samples), and the error introduced by noise (i.e., giving too much weight to the latest samples and their statistical variation). As the error due to lag in the adaptive process is greater for high values of the adaptive weighting factor and the error due to noise is greater for lower values, there is an optimum value of this weighting factor for which the total error is minimized. This value would depend both on the rate of change of the signal statistics and on the signal's relative noise-power in a complex manner, and where these are not known beforehand, it becomes a purely experimental determination. Because it does not have the implicit assumption of a stationary signal, the standard gradient method of determining the reflection coefficients detailed above would be expected

to perform better for non-stationary signals.

A signal may be described as non-stationary when its statistical characteristics vary with time. One way of simulating such a signal is to take a single frequency (complex sinewave) carrier, and modulate it with another sinewave. While in a strict sense, this noise-free modulated sinewave cannot be called non-stationary, on a sample-to-sample basis its statistics (frequency content and/or complex amplitude) are changing with time. Linear prediction-error filters are unable to predict this modulation (such filters can predict, at best, the instantaneous carrier signal, tracking the modulation). As a result, the adaptive filter responds to the modulated signal in the same way as it would to a true non-stationary signal (the filter cannot tell the difference), and the use of a modulated sinewave as the experimental representation of a non-stationary signal is justified for this case. In fact, such a signal is ideal for experimental investigations because it has few parameters to be varied.

In this section, we shall use double-side-band suppressed carrier (DSB-SC) and standard amplitude modulation (AM), and frequency modulation (FM) techniques to study the effects of these parameters on the lattice filter performance. (Much of the work presented in this section is also being published separately by the author [49]). The DSB-SC modulated signal will be examined first because it is the simplest of these signals, having only two frequency components in its spectrum.

The convergence behaviour for such a signal is plotted in Fig. 4.3 (for the reflection coefficients and the filter outputs, respectively). This indicates that, as expected, the standard gradient method gives much

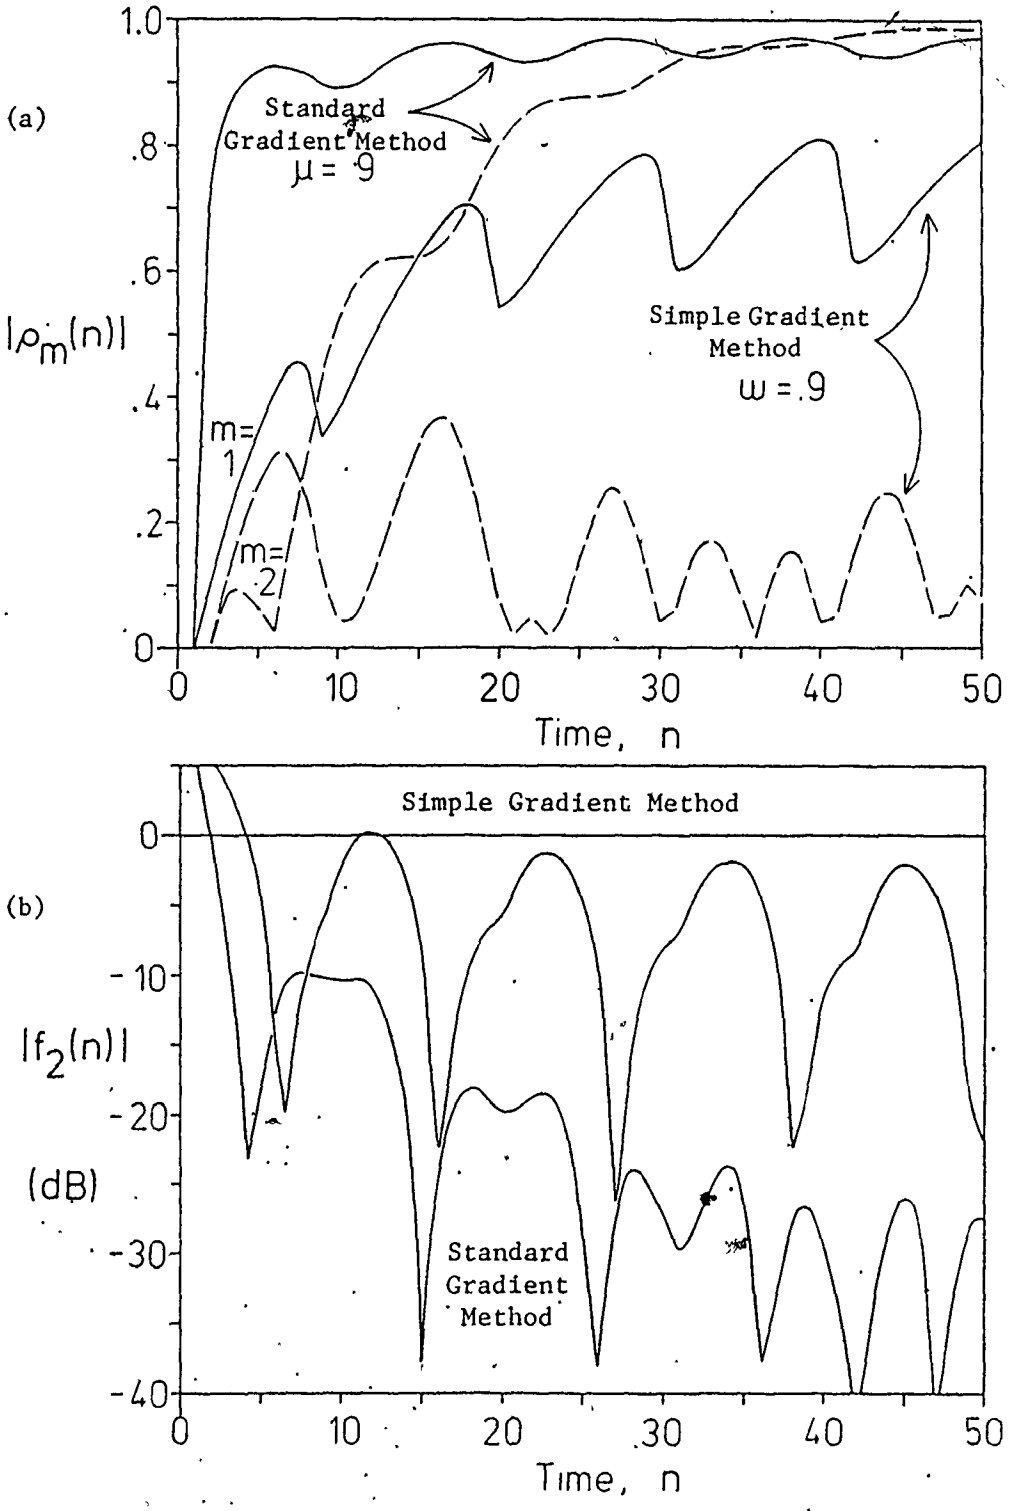

Figure 4.3: Convergence Behaviour for a Non-Stationary Signal
(a) Reflection Coefficients, (b) Filter Outputs

better results for this type of signal, giving about 32 dB of attenuation as opposed to the simple gradient method, which reduces the signal by only about 7 dB over the period observed. (The larger attenuation dips in the response result from the large variations in the input level due to the beating of the two frequencies). The plots shown are for a filter of order 2; for filter order 1, the two methods give very similar results, with much less attenuation. For orders above 2, little improvement is seen. The initial convergence follows the formulae given in the previous sections, but the convergence soon reaches an azimotopic value and goes no further.

Let us now examine how this azimotopic value of the filter error outputs (or equivalently, the signal attenuation) varies as we vary the parameters of the filter and of the signal. The filter parameters to be varied consist of the filter order and the adaptive constant (μ). The only important parameter of the signal is the modulating frequency f_m, which determines the rate of change (or degree of non-stationarity) of the signal. Our investigations have shown that the results are independent of the carrier frequency (f_m is normalized to the sampling frequency). The adapted filter outputs as a function of f_m are plotted in Fig. 4.4, for different values of μ. As can be seen, for the first order filter output, the results are independent of the adaptive constant. Only the standard gradient method is presented in these plots, as the results with the simple gradient method are similar, although not as smooth. For the second-order output, while the simple gradient method shows little improvement over the first order, the standard gradient

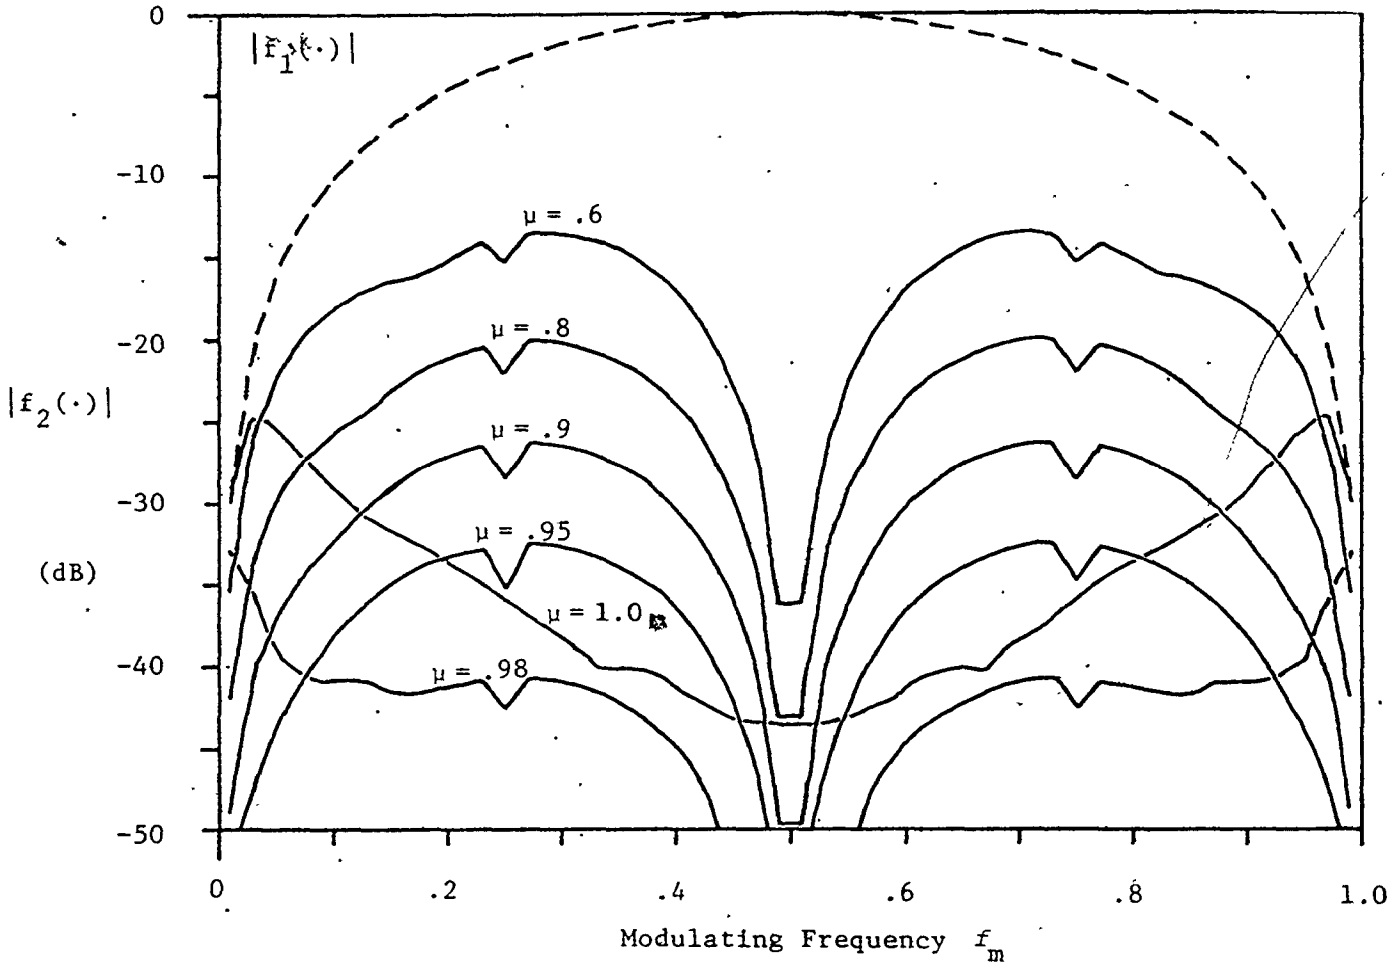

Figure 4.4: Filter Output Versus Modulating Frequency for DSB-SC Modulation (Dotted line is first order output for all μ up to .98)

method shows a significant improvement in reducing the signal level. This reduction is very dependent on the adaptive constant, with the optimum value of μ being in the range of 0.98. There is no significant improvement for the third-order output.

One step up in complexity from DSB-SC is the standard AM signal, where a third frequency component (the carrier) and another parameter (percentage modulation) have been added. The adapted filter-error outputs for modulations of 100%, 33%, 14%, and 7% are plotted in Fig. 4.5, (a) to (c), for the filter orders 1 to 3, respectively. (The dotted lines in these plots are the next lower order outputs using DSB-SC, included for comparison). These AM curves are very similar to the DSB-SC curves, except that: (a) a third filter order is needed to handle the additional frequency component, and (b) the modulation level is a new factor; as it decreases, so does the error output.

Figure 4.5(d) replots the filter outputs with 33% modulation for different values of the adaptive constant μ. Comparison of this figure with Fig. 4.4 shows basically similar performance to that obtained using DSB-SC. In this case (for AM), the first two filter orders are independent of μ (and of the choice of adaptive method), while the third-order output shows a strong dependence on μ, with the optimum value now being around 0.95. (There is no significant improvement from using filter orders above 3 with the standard gradient method). As before, this method out-performs the simple gradient method.

Finally, let us examine the results using frequency modulation to create the signal. Again, the first 3 filter orders are plotted in Fig. 4.6, with the percent modulation parameter now being replaced by

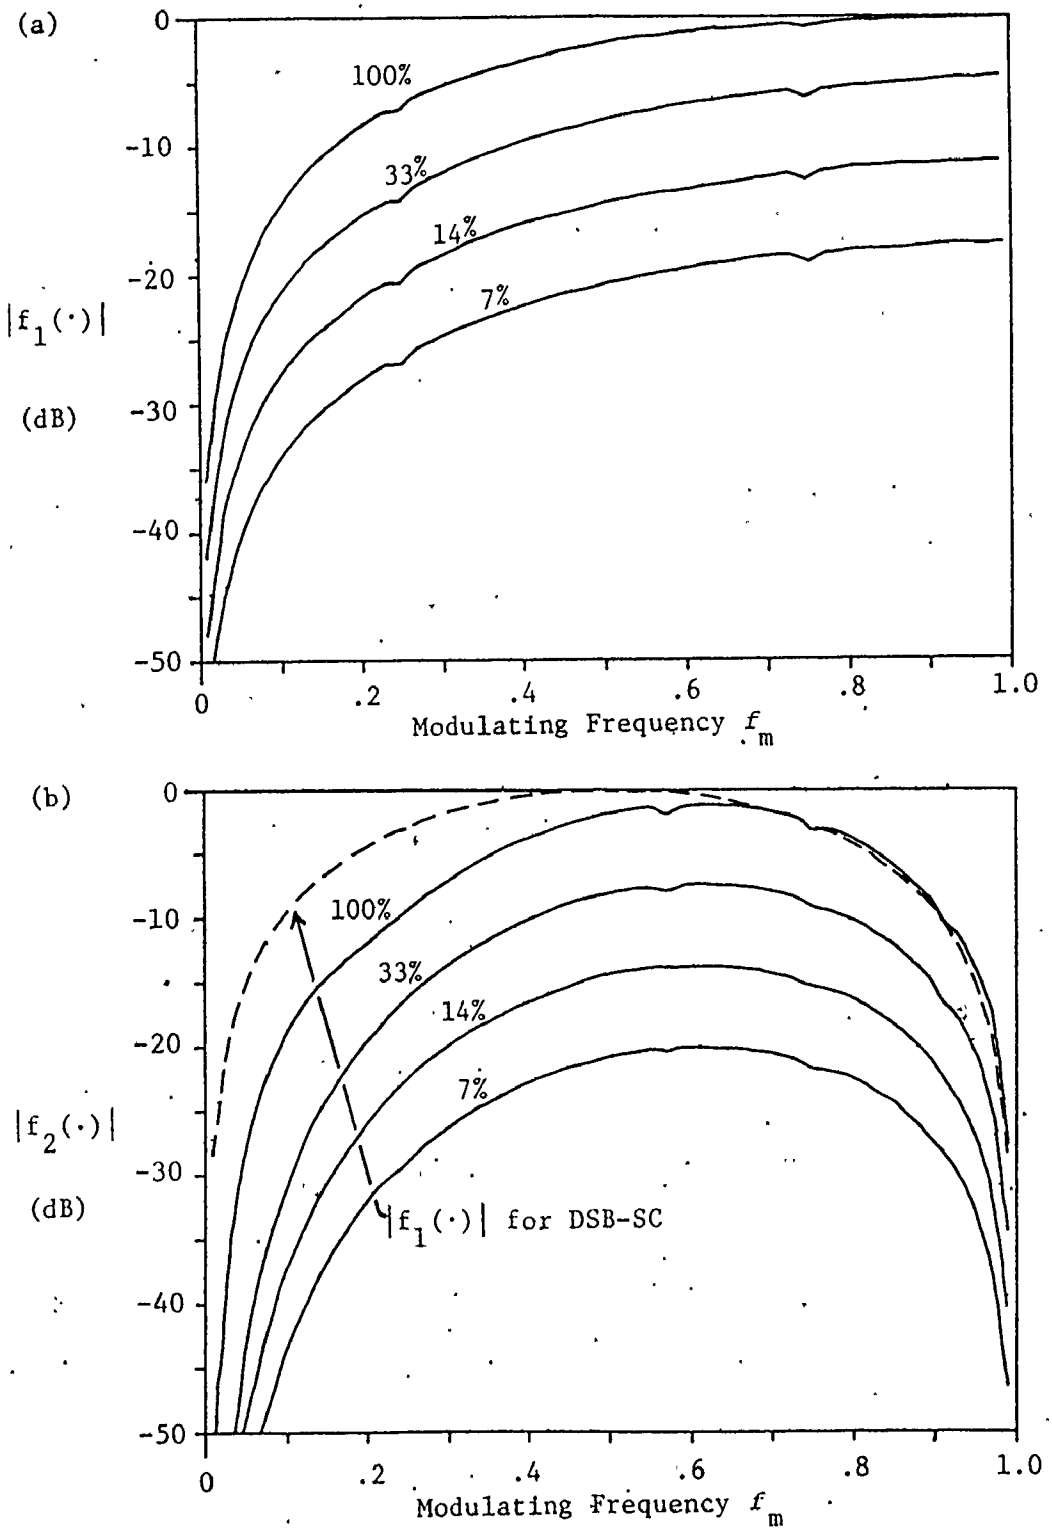

Figure 4.5: Filter Output Versus Modulating Frequency for AM
(a) m = 1 (all μ up to 0.98), (b) m = 2 (all μ up to 0.98)
(c) m = 3 (μ = 0.9) , (d) m = 3, 33% AM

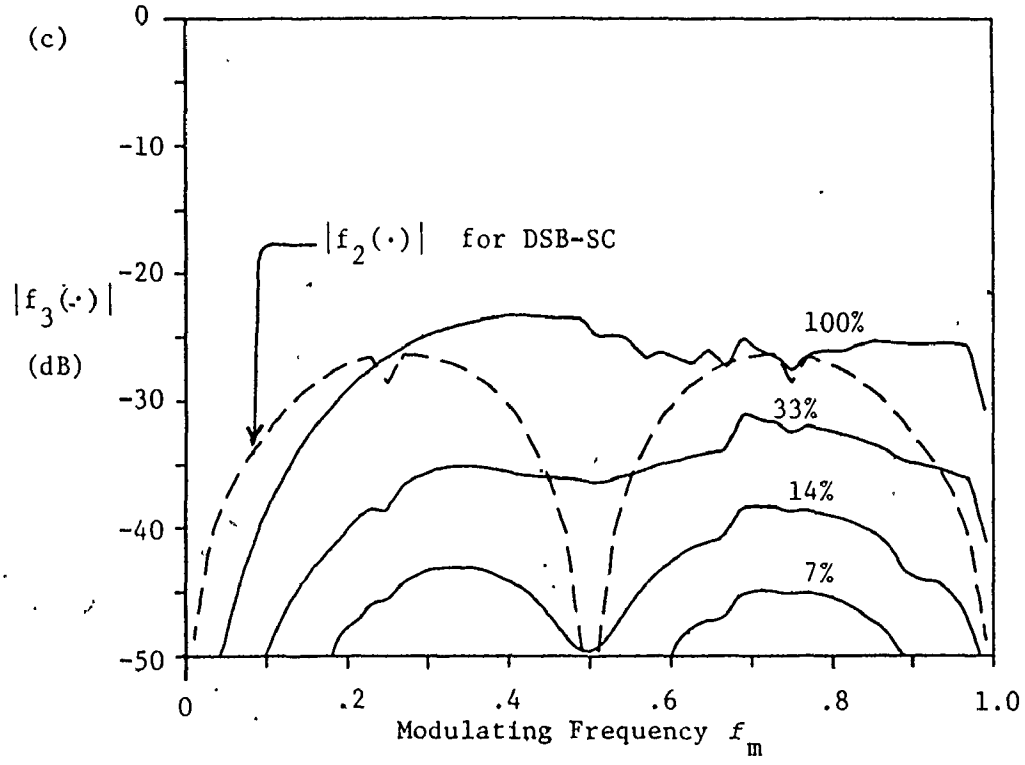
(c)
0
-10
-20
-30
-40
-50
$|f_3(\cdot)|$
(dB)
$|f_2(\cdot)|$ for DSB-SC
100%
33%
14%
7%
0
.2
.4
.6
.8
1.0
Modulating Frequency f_m

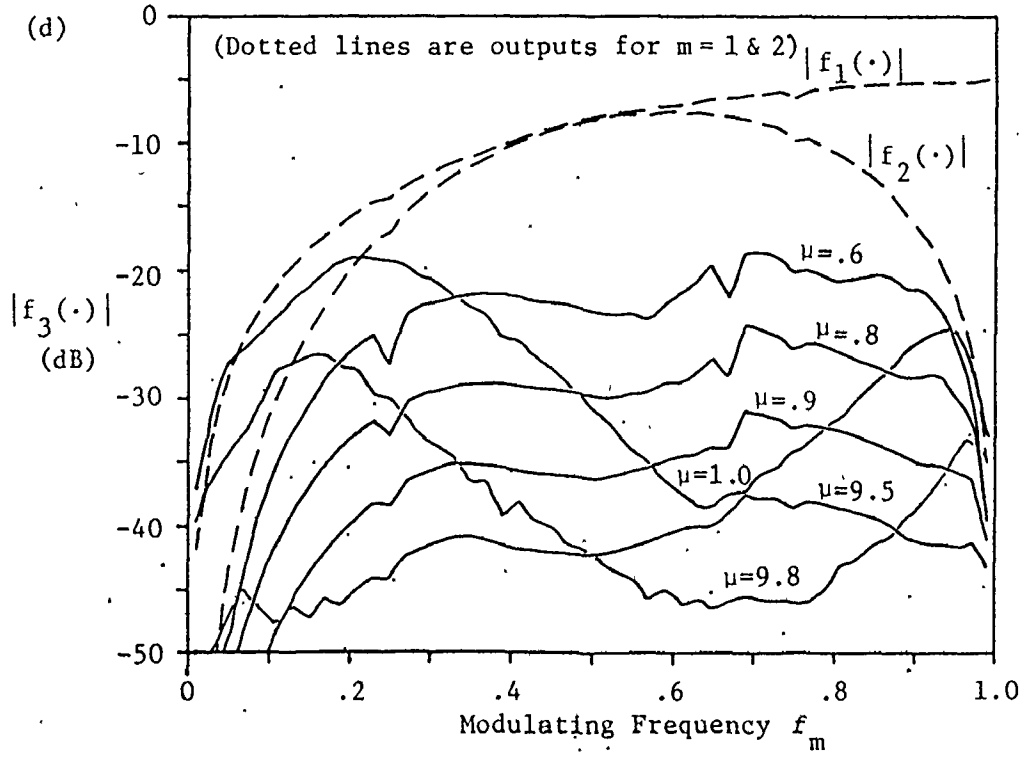
(d)
(Dotted lines are outputs for m = 1 & 2)
0
-10
-20
-30
-40
-50
$|f_3(\cdot)|$
(dB)
$|f_1(\cdot)|$
$|f_2(\cdot)|$
μ=.6
μ=.8
μ=.9
μ=1.0
μ=9.5
μ=9.8
0
.2
.4
.6
.8
1.0
Modulating Frequency f_m

(Figure 4.5)

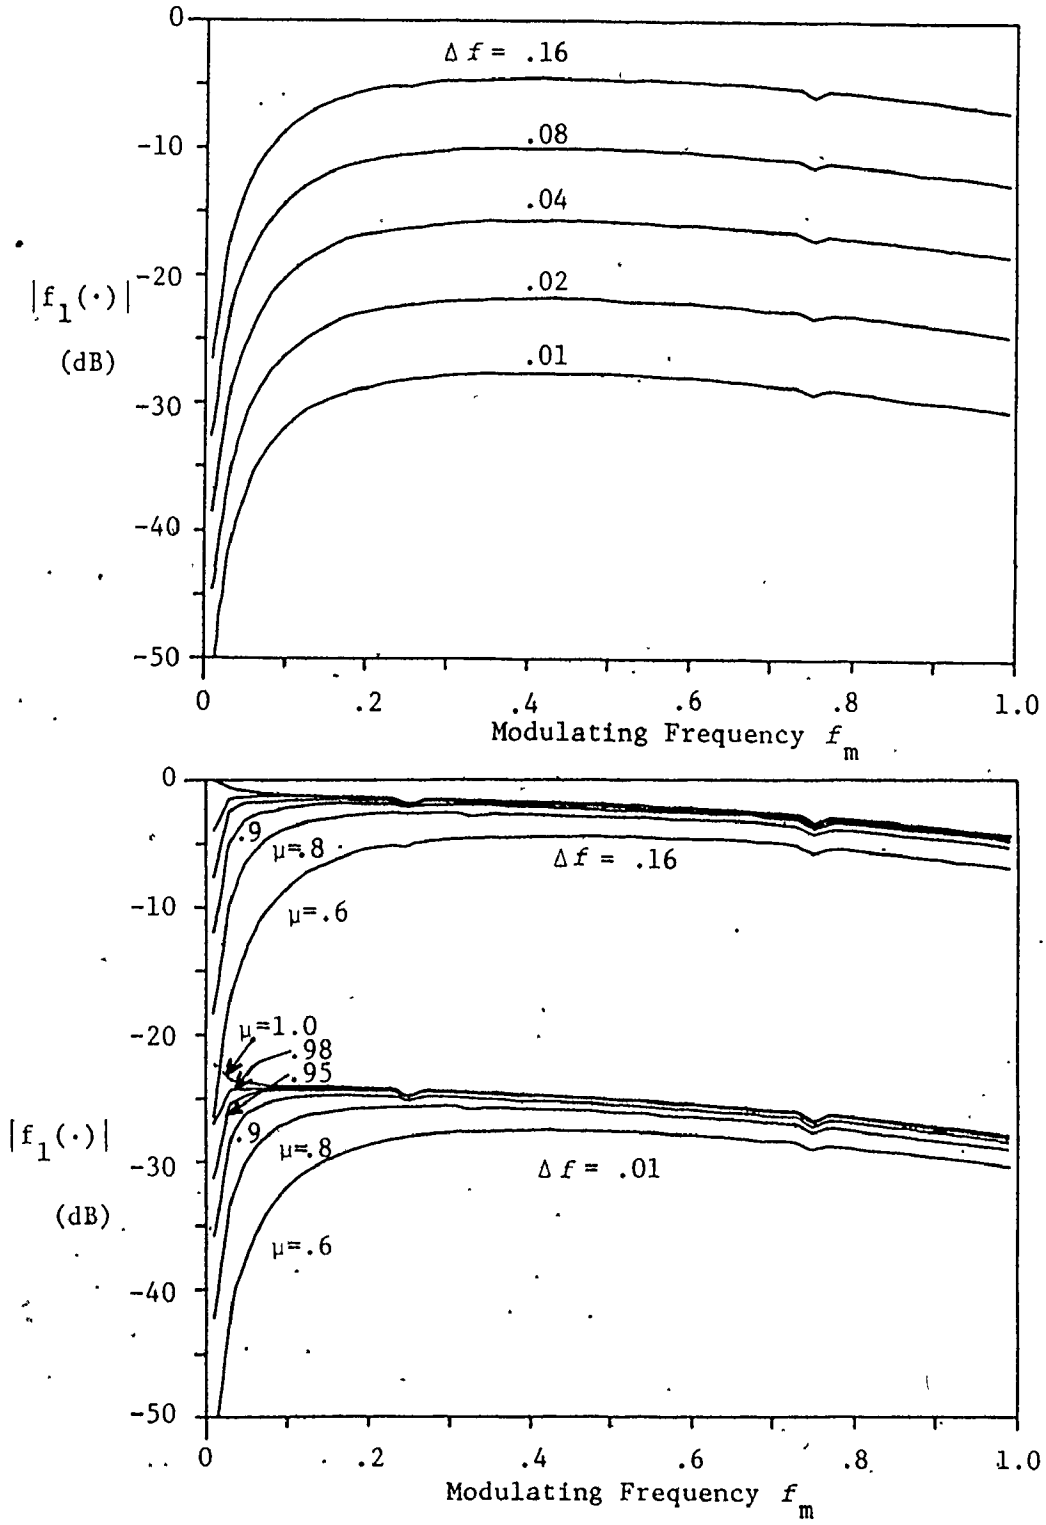

Figure 4.6: Filter Output Versus Modulating Frequency for FM
(a) m = 1 (μ = 0.6), (b) m = 1 (Δf = 0.16, 0.01)
(c) m = 2 (μ = 0.6), (d) m = 3 (μ = 0.6)

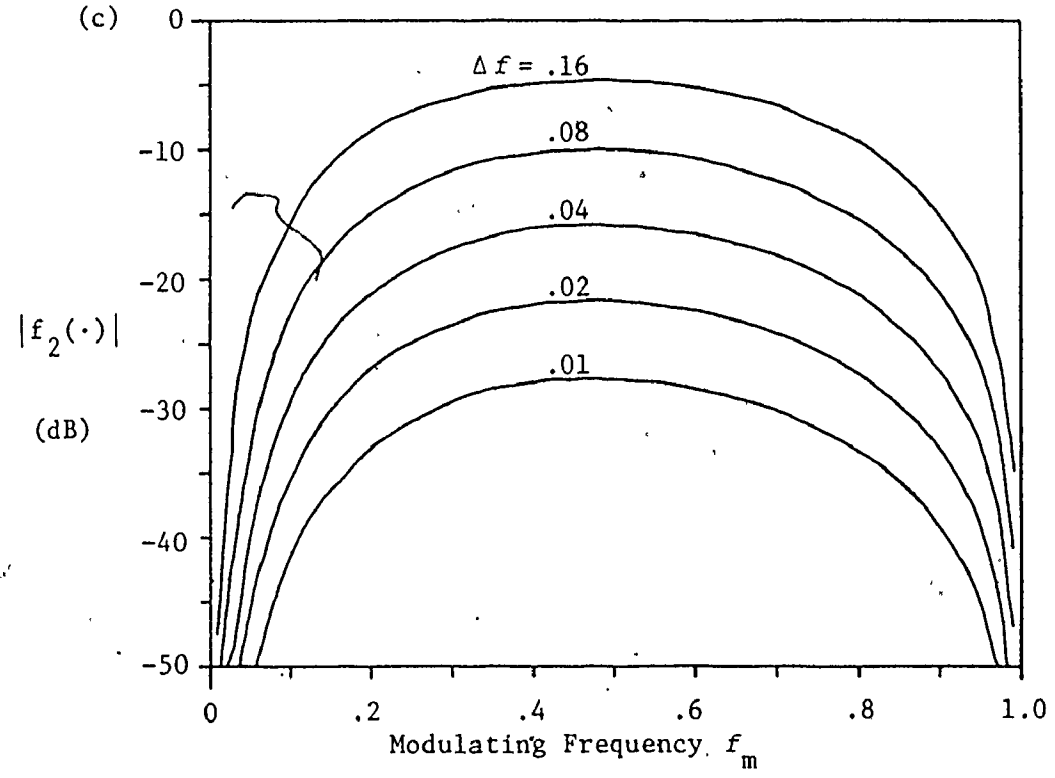
(c)
0
-10
-20
-30
-40
-50
$|f_2(\cdot)|$
(dB)
Δf = .16
.08
.04
.02
.01
0
.2
.4
.6
.8
1.0
Modulating Frequency f_m

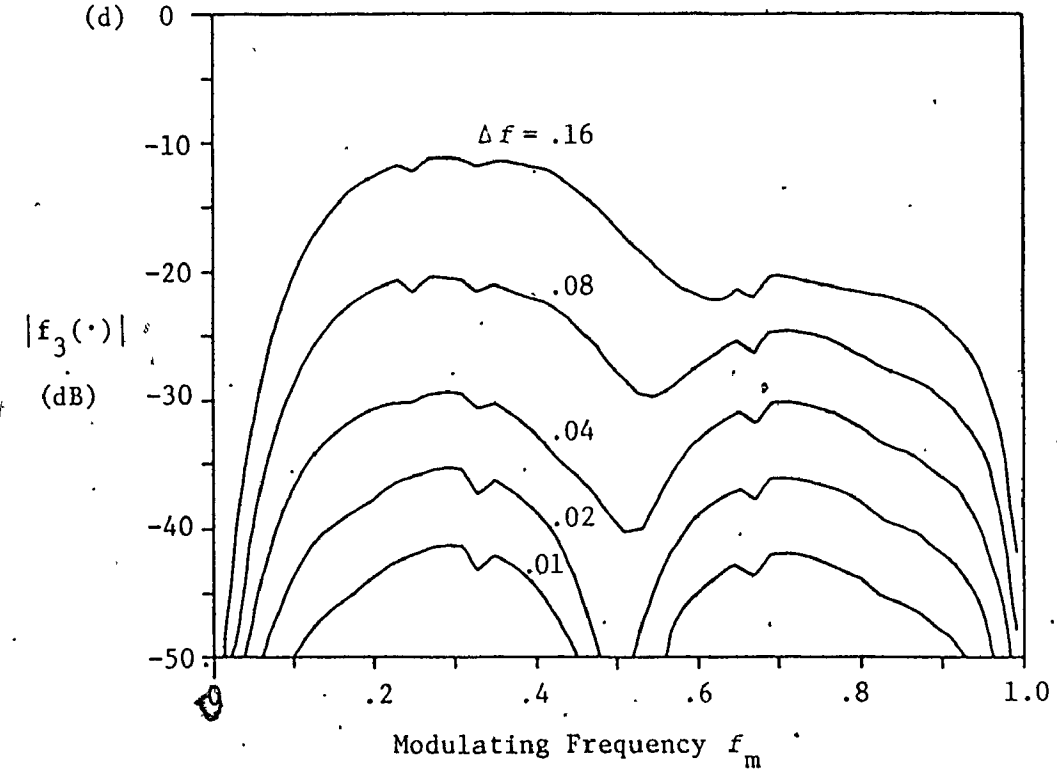
(d)
0
-10
-20
-30
-40
-50
$|f_3(\cdot)|$
(dB)
Δf = .16
.08
.04
.02
.01
.2
.4
.6
.8
1.0
Modulating Frequency f_m

(Figure 4.6)

the parameter of frequency deviation. For FM signals, there is a small but noticeable dependence on μ for the first two filter orders, as can be seen in Fig. 4.6(b) for the first order. Comparing Fig. 4.6 with the corresponding Fig. 4.5 for AM signals, it is apparent that the behaviours are quite similar. (This similarity grows more pronounced if the FM plots are redone with the modulation index β held constant instead of the frequency deviation Δf). As with the previous signals, the error output for the third filter order shows a strong dependence on the adaptive constant, with optimum values of μ ranging from 0.6 for $\Delta f = 0.16$, to 0.9 for $\Delta f = 0.01$. Also, there is again no significant improvement for filter orders above three.

A theoretical explanation of the non-stationary behaviour described here is, at the present time, not possible. However, certain trends are visible:

(1) There is a definite limit on the number of filter stages needed, beyond which there is little improvement in performance (3 stages for standard AM and FM, 2 stages for DSB-SC).

(2) The highest of these stages has a performance which is strongly dependent on μ, the optimum value of which appears to be a function of the modulation level and type. The standard gradient method performs much better than the simple gradient method at this stage.

(3) The lower stages of the filter show a performance which is largely independent of adaptive constant or method chosen, as well as taking a very basic shape when plotted.

(4) The response of the filter to the varying signal is basically the same whether it is the amplitude or the frequency of the signal which is being varied.

(5) Finally, as could be expected, the filter performance improved as the modulation level decreased and as the rate of variation (f_m) approached zero.

The non-stationary behaviour of the lattice-structure PEF is a subject of which little is understood (in a quantitative sense). Particularly needed are explanations of the dependence of this behaviour on the adaptive constant and the modulating frequency. Such research goes beyond the scope of this thesis, but definitely deserves future attention.

CHAPTER 5

ANALYSIS USING SIMULATED RADAR DATA

5.1 Introduction and Performance Measures

In the previous chapters, some of the properties of the lattice-structure PEF were analysed using sinusoidal waveforms as the filter input. Such simple input signals are very useful for examining basic characteristics of the filter, but they are not adequate for dealing with the response of the filter to a complex signal, such as radar clutter and targets. In Chapters 6 and 7, actual radar data will be examined, allowing the proper evaluation of the filter under some specific operating conditions. However, as only a limited number of discrete cases could be covered this way, there is also a need for artificially synthesizing radar signals representing the general case, which may be varied to determine how the filter responds with different signal parameters. Such a signal is achieved through computer simulation, and will be used in this chapter for qualitative comparisons of the different lattice algorithms and judging the effects of noise and target velocity (and other parameters) on the response of the filter.

The simulated radar data used for this analysis consisted of signals representative of weather clutter and target signals. The weather clutter signal was generated by an algorithm developed by Hawkes [30],

based on a generalized form of the autocorrelation function of radar clutter, starting from a random distribution. Data sets representing widely different clutter conditions were generated by varying the clutter algorithm's parameters and were used to test the response of the different adaptive lattice filtering set-ups. Most of the performance data presented in this chapter resulted from the use of a data set representative of an "average" rainstorm which was uniform and non-moving over the region of interest. These results are consistent with those from the other data sets. (Another example, a data set representing a very heavy rainstorm with some velocity, is presented at the end of this chapter). The target signal was simulated by the product of a complex sinusoid representing the Doppler modulation related to target radial velocity, and a Gaussian envelope approximating the horizontal beam-pattern of the radar antenna.

At this point, it is necessary to define and describe some of the performance measures used to analyse radar clutter filters. Perhaps the most important [34, 35] of these is the Improvement Factor, defined as: "The signal-to-clutter ratio at the output of the system divided by the signal-to-clutter ratio at the input of the system, averaged uniformly over all target radial velocities of interest" [36]. With the simulated data, it is easy enough to average over all possible target radial velocities (at the end of this chapter, some examples are plotted of the improvement factor as a function of velocity). However, for the real data in Chapters 6 and 7, only a limited number of radial velocities were available, so that averaging takes place using only those

available targets.

Another useful radar performance measure consists of the receiver-operator-characteristic (ROC) curves, which plot the probability of detection of a radar target versus the probability of false alarms under various circumstances. When determined experimentally, these probabilities are simply the numbers of detections and false alarms recorded (peaks exceeding a set signal threshold level). This measure implies some additional processing beyond the clutter filter, which in the cases presented here consists of a simple amplitude detector and threshold circuit. Another important performance measure, the sub-clutter-visibility (SCV), will be presented in Chapter 7.

5.2 Comparisons in Adapted Response

Our first application for the simulated radar clutter arises in the selection of the lattice algorithm to be used for subsequent research. As will be seen, there is no "right" choice for all possible signals, as different algorithms may perform better or worse for different types of signals. This is why it is important to use a representative radar signal, such as the average rainstorm clutter mentioned above, to compare these algorithms. (Most of the results presented in this section were previously published by the author [37]).

Figure 5.1 plots the clutter attenuation (reciprocal of filter gain) of the adapted (fifth order) lattice filter versus adaptive constant for four of the reflection coefficient algorithms presented in section 2.3. These are the harmonic-mean (H), the minimum (M), the

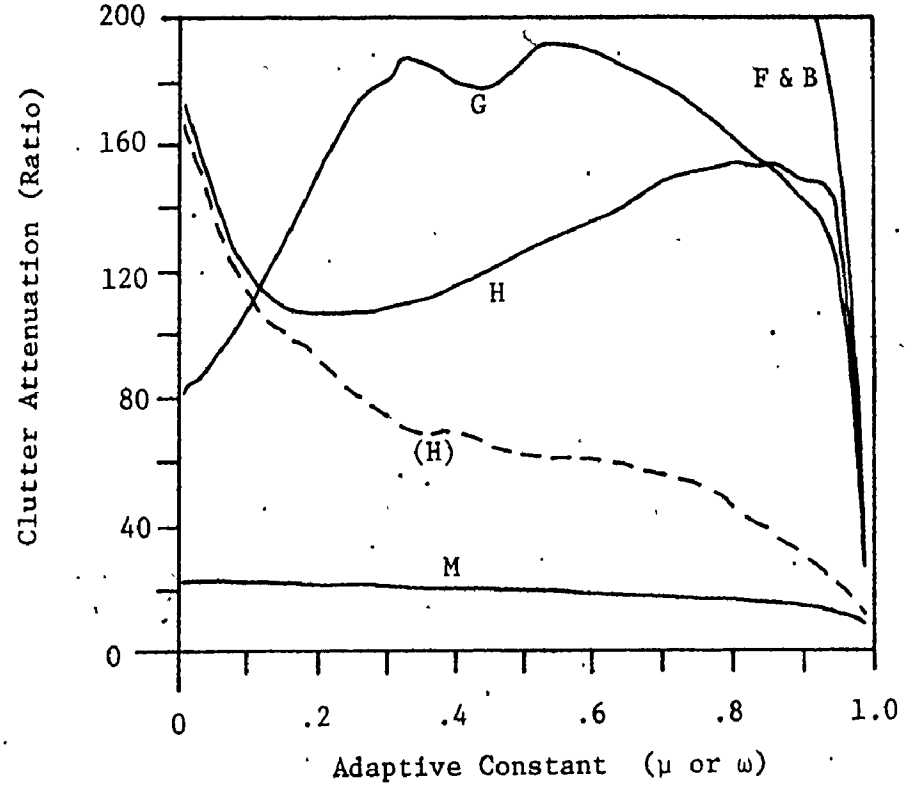

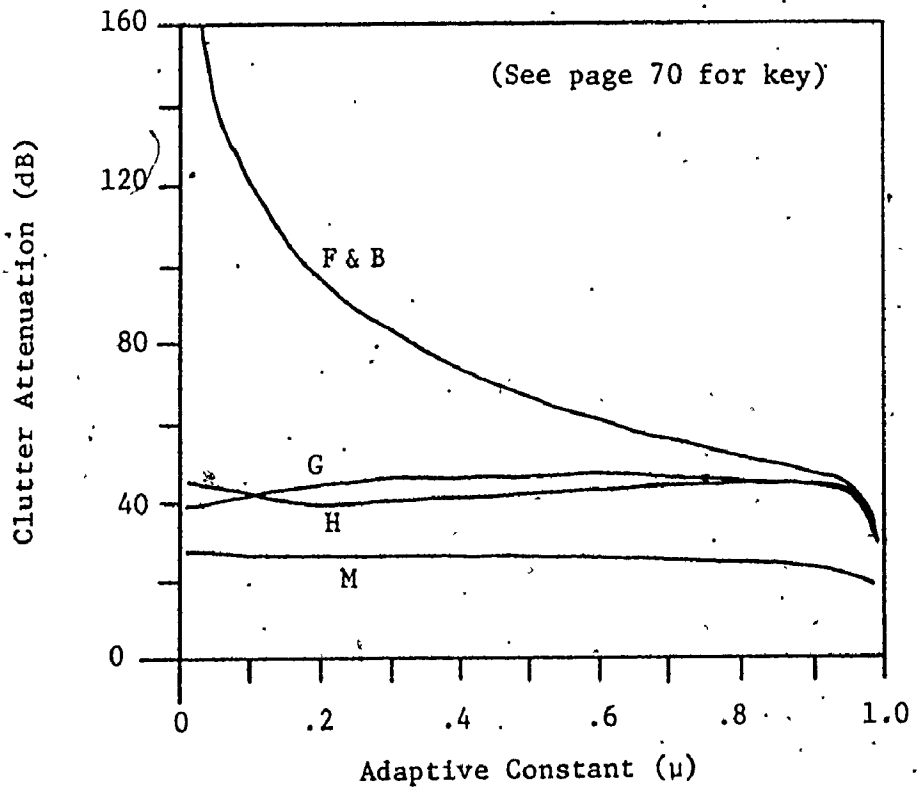

Figure 5.1: Attenuation of Clutter Signal Using Different Lattice Algorithms (m = 5)

geometric-mean (G), and the forward-and-backward (F & B) algorithms, all of which attempt to minimize both forward and backward prediction errors. The two other algorithms (forward (F), and backward (B)) are compared to these using a different approach in the next section. Only the first three of these algorithms guarantee the stability of the filter structure. Also represented in this figure are the two adaptive implementation methods (as developed in Chapter 3), plotted for the harmonic-mean algorithm only (this relationship is representative of all the algorithms):

The sudden drop in attenuation as the adaptive constant goes above $\sim$0.97 in Fig. 5.1 is a result of the limited length of the data sample ($n = 256$), which does not give enough time for the filters to adapt properly with higher values of the constant. The relative performance of the two methods of adaptive implementation is basically as might be expected, with the simple gradient method response equal to that of the standard gradient method at $\omega = \mu = 0$, and getting progressively worse as the adaptive constant approaches 1.0 (due to the inherent assumption of signal stationarity for this method).

It may seem from Fig. 5.1 that the F & B algorithm has a large advantage over the other algorithms for this radar clutter, an advantage which increases as the adaptive constant approaches zero. This is probably due to the reflection coefficients being able to assume independent values for the forward and backward branches, thus providing an additional degree of freedom to deal with short-term statistical variations in the signal. It has been noted that the differences between the forward

and backward coefficients are a good indication of the degree of these variations [20].

The seeming advantage of the F & B algorithm described above disappears, however, when we examine the improvement factor performance measure, plotted in Fig. 5.2 for the same conditions as Fig. 5.1. The same freedom which allows the F & B algorithm to do so well in attenuating the clutter signal also appears to make it difficult for this algorithm to distinguish between targets and clutter. Although this algorithm does not perform well in this regard, there is no sign of instability in the filter, although this algorithm is the only one of the four not guaranteeing stability.

Another consistently poor performer is the minimum algorithm. The problem with this algorithm probably results from the recursive implementation. In a block-processing implementation, the forward or backward decision of this algorithm would be made only once for each filter stage, while in the recursive implementation, each stage could be switching back and forth in response to new data.

A notable factor in Fig. 5.2 (and also in 5.1) is that the curves are not smooth, but show considerable variations in response as the adaptive constant changes. This can be attributed to the fact that as the constant is increased, the algorithms are more strongly influenced by data samples further in the past. The geometric-mean algorithm seems to be particularly sensitive to this influence, suggesting that its response to non-stationary signals may be somewhat erratic.

The harmonic-mean algorithm, on the other hand, shows little of this

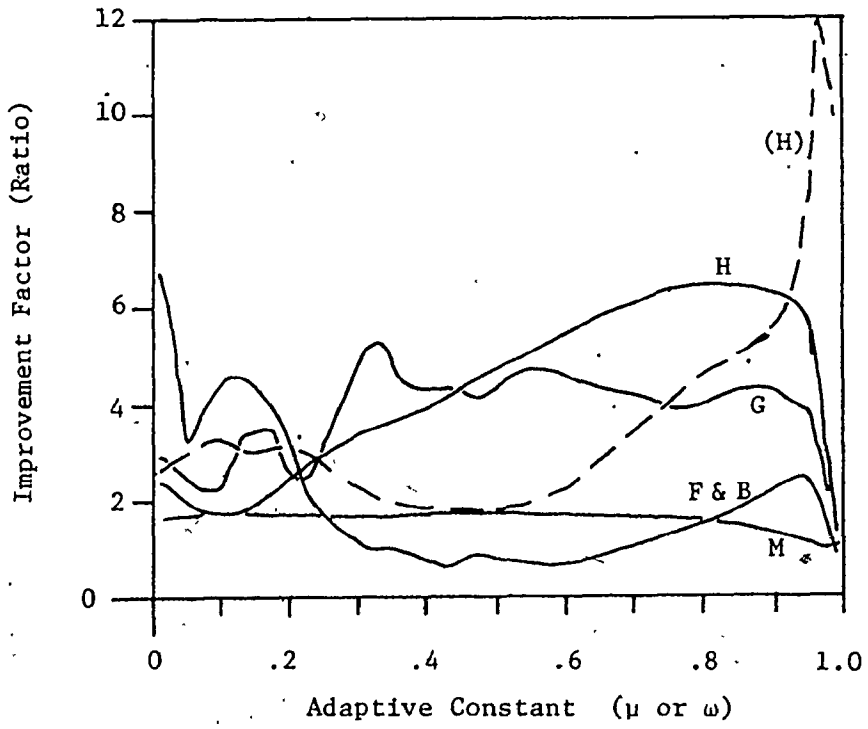

Figure 5.2: Comparison of Radar Improvement Factor (I) Using Different Lattice Algorithms (m = 5)

Key to Figures 5.1 - 5.3:

F & B - Forward and Backward Algorithm
M - Minimum Algorithm
G - Geometric-Mean Algorithm
H - Harmonic-Mean Algorithm
} Standard Gradient Method

(H) - Harmonic-Mean Algorithm - Simple Gradient Method

influence. Based on this fact, and on the fact that this algorithm apparently gives the highest improvement of the four for this method, the harmonic-mean algorithm was selected as the best for this application. Unless otherwise noted, this algorithm will be the one used throughout the remainder of this thesis. (It should be remembered that, although this algorithm may be best for the radar filtering problem, it is not necessarily so for other types of signals). The filter order used ($m = 5$) did not appear to be a major factor in these measurements. The filter responded well at $m = 3$ and above for this clutter signal. As the filter order increased, there was a gradual decline in improvement factor (especially for the F & B algorithm).

In comparing the two methods of recursive implementation, the standard gradient method does better for most measurements, as expected. The peak enhancement found with the simple gradient method near $\omega = 1.0$ is less expected. In this range, neither method has time to adapt to the target signal. The standard gradient method, however, would start to adapt to the target, while the simple gradient method may be unable to respond appropriately to the target signal due to its assumption of signal invariance, an assumption which increases as ω increases. Although both methods are plotted only for the harmonic-mean algorithm, their relative performance is representative of all the algorithms.

Before leaving this section on comparisons in adapted response, let us examine the filter response in the frequency domain. Figure 5.3 compares the effects of algorithm, method and filter order on lattice filters adapted to the same average clutter signal as before. The actual clutter spectrum is Gaussian in shape, centered at zero frequency and

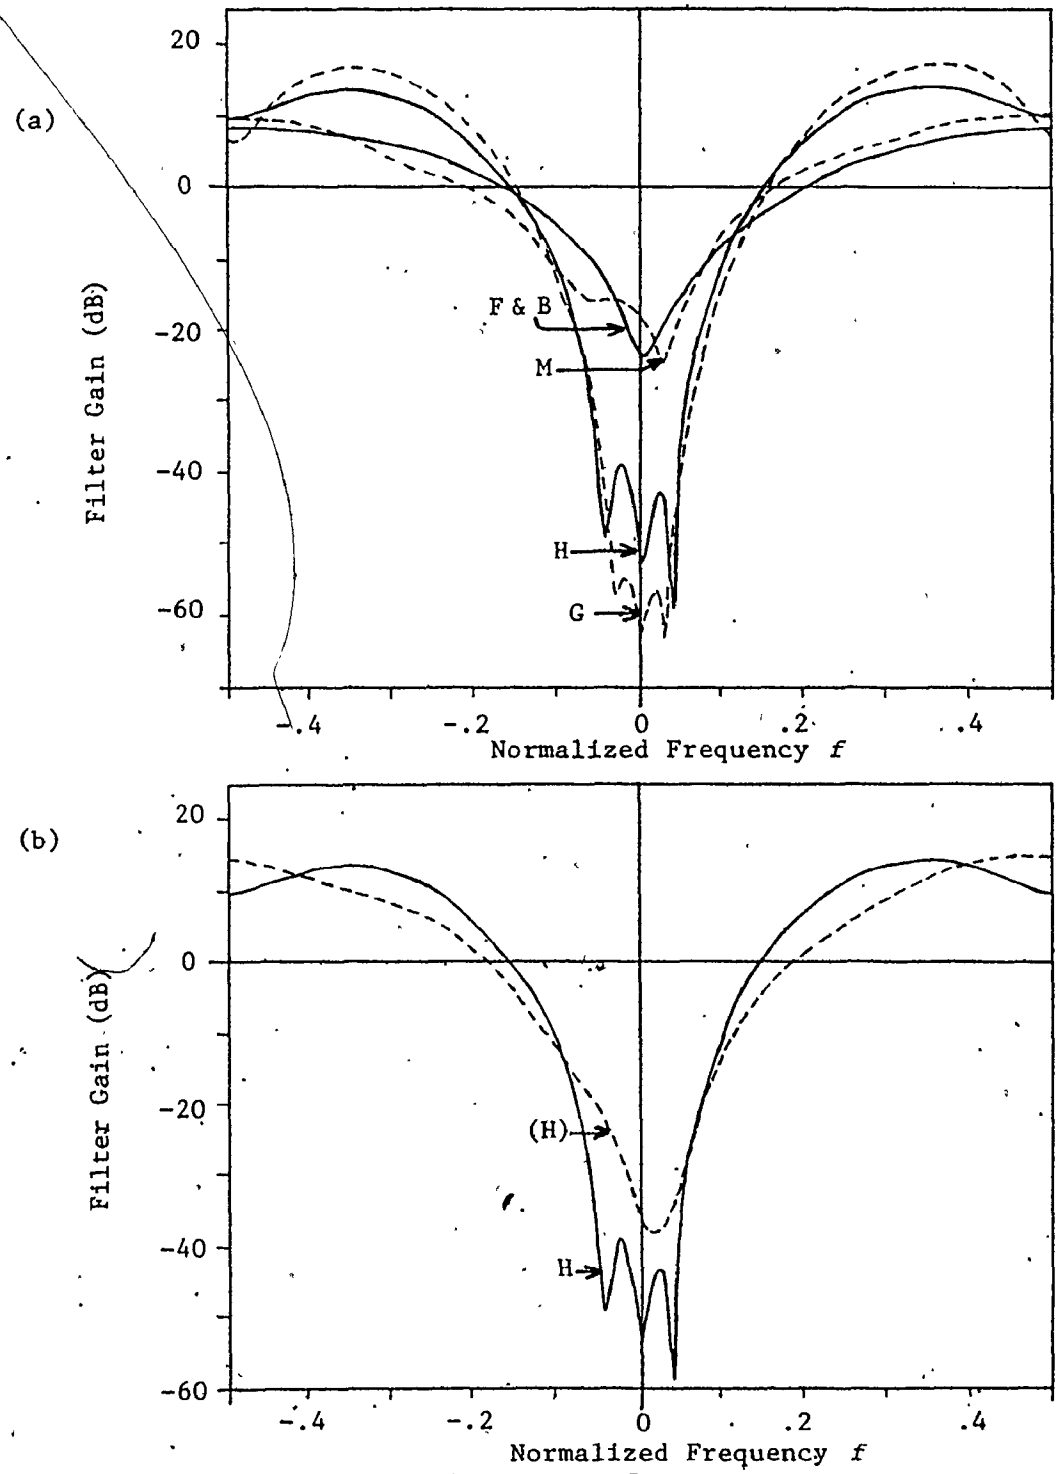

Figure 5.3: Frequency Response of Lattice Filter Adapted to Radar Clutter Data ($\mu = \omega = 0.9$):
(a) Different Algorithms (Standard Gradient Method)
(b) Different Methods (Harmonic-Mean Algorithm, m=5)
(c) Different Filter Orders (Harmonic-Mean Algorithm, Standard Gradient Method)

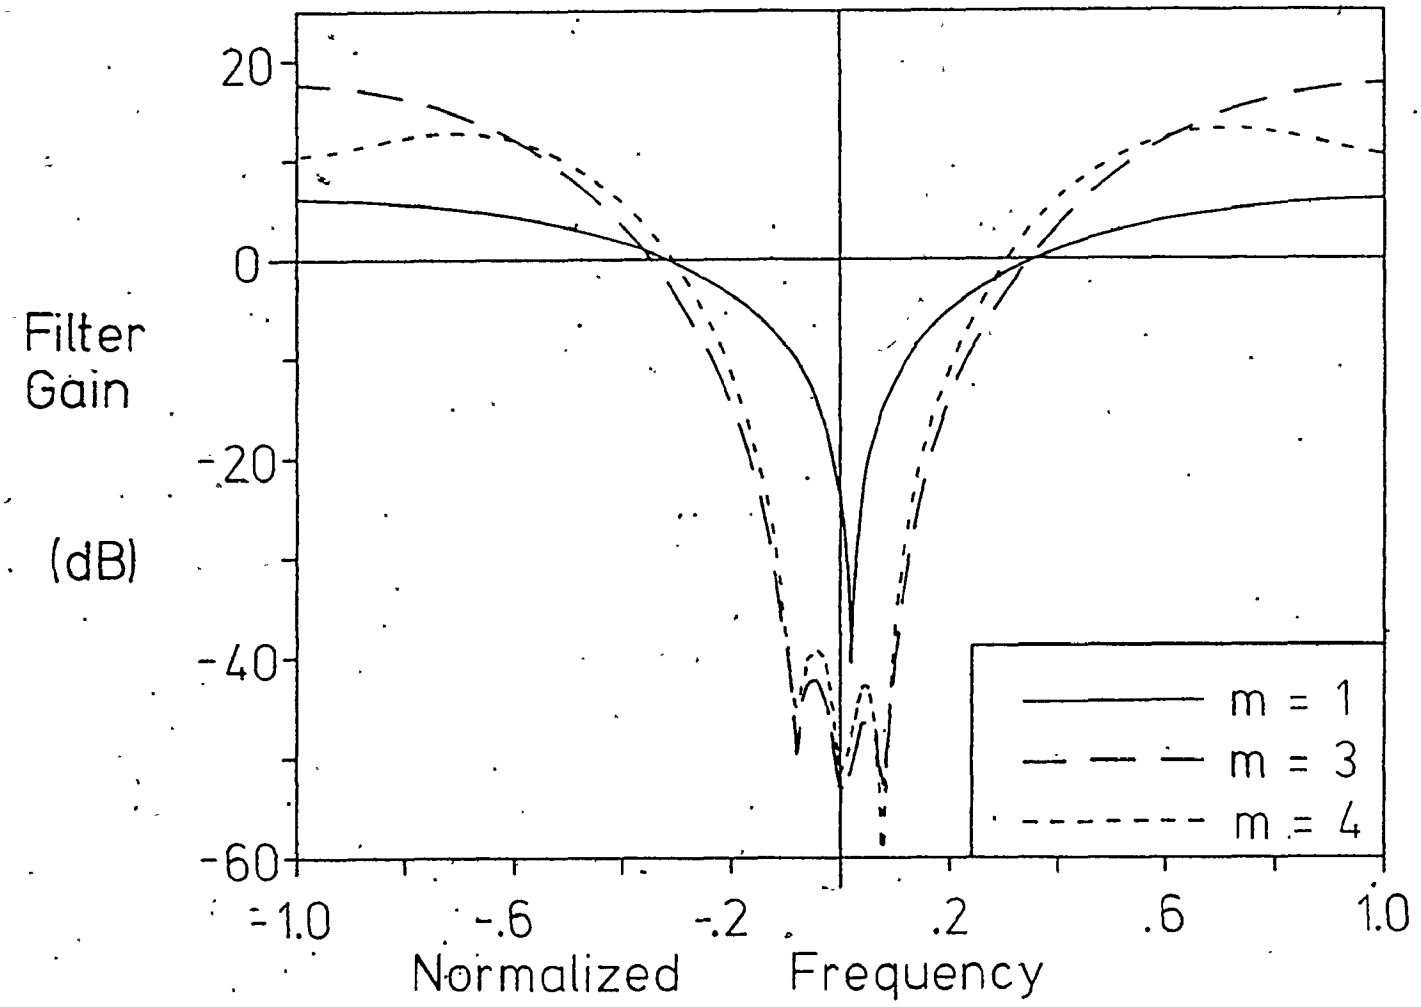
20
0
-20
-40
-60
Filter
Gain
(dB)
-1.0
-.6
-.2
.2
.6
1.0
Normalized Frequency
m = 1
m = 3
m = 4

Figure 5.3(c)

having a dynamic range of about 60 dB. The filters approximate the spectrum as a series of zeros, filtering out narrow bands close together. (This can be clearly seen in the frequency responses of the harmonic and geometric-mean algorithms). As can be seen in Fig. 5.3(c), the first three orders of the lattice filter are used to approximate the inverse of this Gaussian spectrum. The fourth stage of the filter provides some attenuation to flatten the artificial gain introduced by the first three orders, away from the main peak. Fifth and higher-order stages have little effect on the spectrum. More complex clutter signals require higher filter orders. These spectra need not be symmetrical around zero frequency.

5.3 Some Comparisons in Convergence Properties

Let us now turn our attention to the convergence properties of the lattice filters in response to the simulated radar clutter; first to confirm that these properties agree with the theoretical development presented in Chapter 4, and second, as a basis for further comparison.

Figures 5.4 and 5.5 plot the convergence behaviour of the lattice filter reflection coefficients using the simple and standard gradient methods respectively, in response to the same simulated radar clutter signal as was used earlier in this chapter. It can be seen that the convergence behaviour of the first reflection coefficient for each method (the top plots in Figures 5.4 and 5.5) does indeed follow the theoretical relationships presented earlier, with irregularities

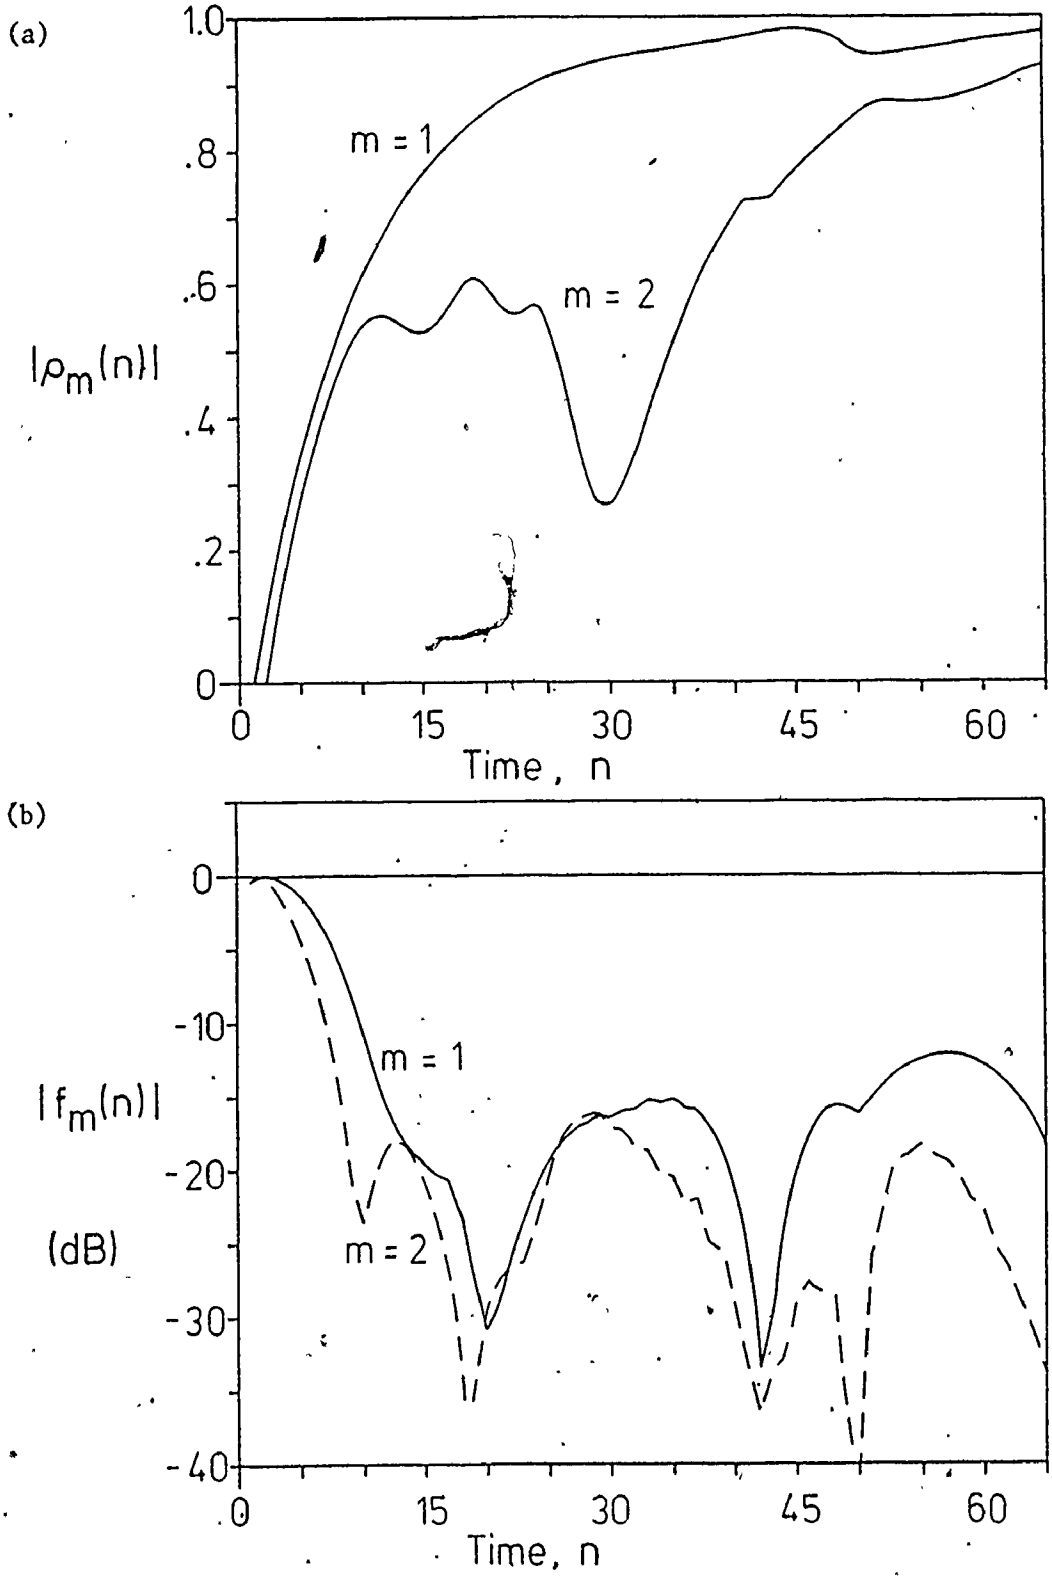

Figure 5.4: Convergence Behaviour for Radar Clutter Using the Simple Gradient Method (ω = 0.9) (a) Reflection coefficients, (b) Output

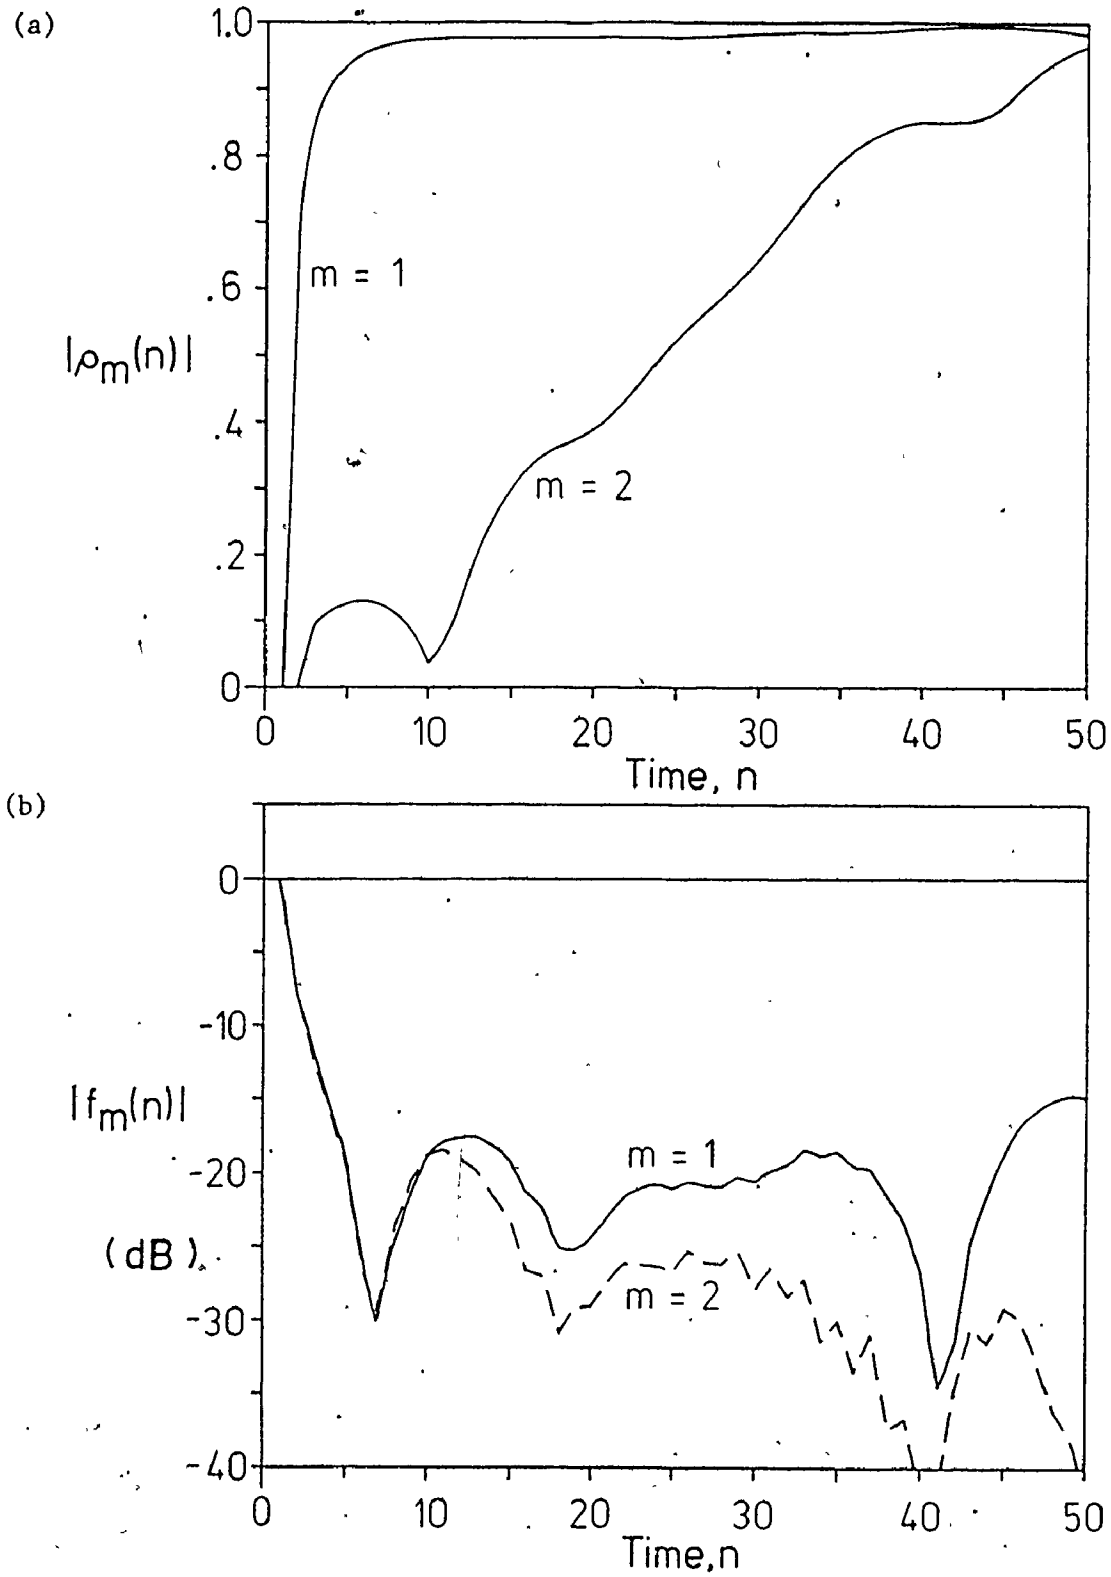

Figure 5.5: Convergence Behaviour for Radar Clutter Using the Standard Gradient Method ($\mu = 0.9$) (a) Reflection Coefficients (b) Output

caused by statistical fluctuations in the data. Similarly, after decoupling, the second order tends to exhibit the same behaviour, although the irregularities are more pronounced, especially with the standard gradient method. Also, the filter outputs, as seen in the lower plots, tend to exhibit the expected behaviour at first, but quickly reach a point beyond which the average attenuation of the clutter is near constant. (The variations beyond this point are due to fluctuations in the data).

The decoupling point of the second order of the filter is quite important to our study. For the reflection coefficients calculated by the simple gradient method with the clutter data, decoupling takes place at about $n = 25$ $(= \tau)$ for $\omega = 0.9$. This corresponds to a fractional error in the first reflection coefficient given by $\varepsilon_1(\tau) = 0.08$ which gives a decoupling constant for this particular signal type as $a = 2.64$. (See sections 4.1 and 4.3 for definitions). Having thus determined the decoupling constant, a value for the adaptive constant ω can be chosen to discriminate between target and clutter. Assuming typical values of target duration as 20 samples and clutter duration as 100 samples or more, a good value of ω to discriminate the two is approximately 0.95 (for a second-order filter), assuming the clutter is completely stationary (which is not generally true).

The decoupling point using the standard gradient method of reflection coefficient calculation occurs at $n = 8$ for $\mu = 0.9$ (Fig. 11). This, in turn, corresponds to an error in the first coefficient of $\varepsilon_1(\tau) = 0.044$, approximately half the error level at which the simple

gradient method coefficients decoupled. There is no clear explanation as to why the standard gradient coefficient must be twice as accurate as the simple gradient coefficient before decoupling occurs. As with the simple gradient method, the adaptive weighting factor μ of the standard gradient method may be adjusted to provide the optimum decoupling time for a given signal, although no corresponding decoupling constant exists.

In comparing Figures 5.4 and 5.5, it may appear that the filter using the standard gradient method calculation results in greater attenuation of the data. However, this is largely due to the fact that with this method the filter has an effectively smaller adaptive constant, allowing faster convergence to statistical fluctuations in the data. In practice, the filter using the standard gradient method is slightly better than that using the simple gradient method, for this type of data. As noted before, the first method is computationally more efficient, as well as being more easily used.

The convergence behaviour of the lattice-structure PEF may also be used to compare the relative merits of different algorithms for the reflection coefficient calculation. This is especially true in the case of the forward and the backward algorithms, as the convergence behaviours demonstrate an interesting property of these algorithms. Figure 5.6 compares the behaviour of these two algorithms (denoted F and B) with that of the harmonic-mean algorithm (denoted H), which is initially the harmonic mean of the two. It can be seen in Fig. 5.6(a) that initially the forward algorithm does better and the backward algorithm does worse than the median value of the harmonic mean. This is due

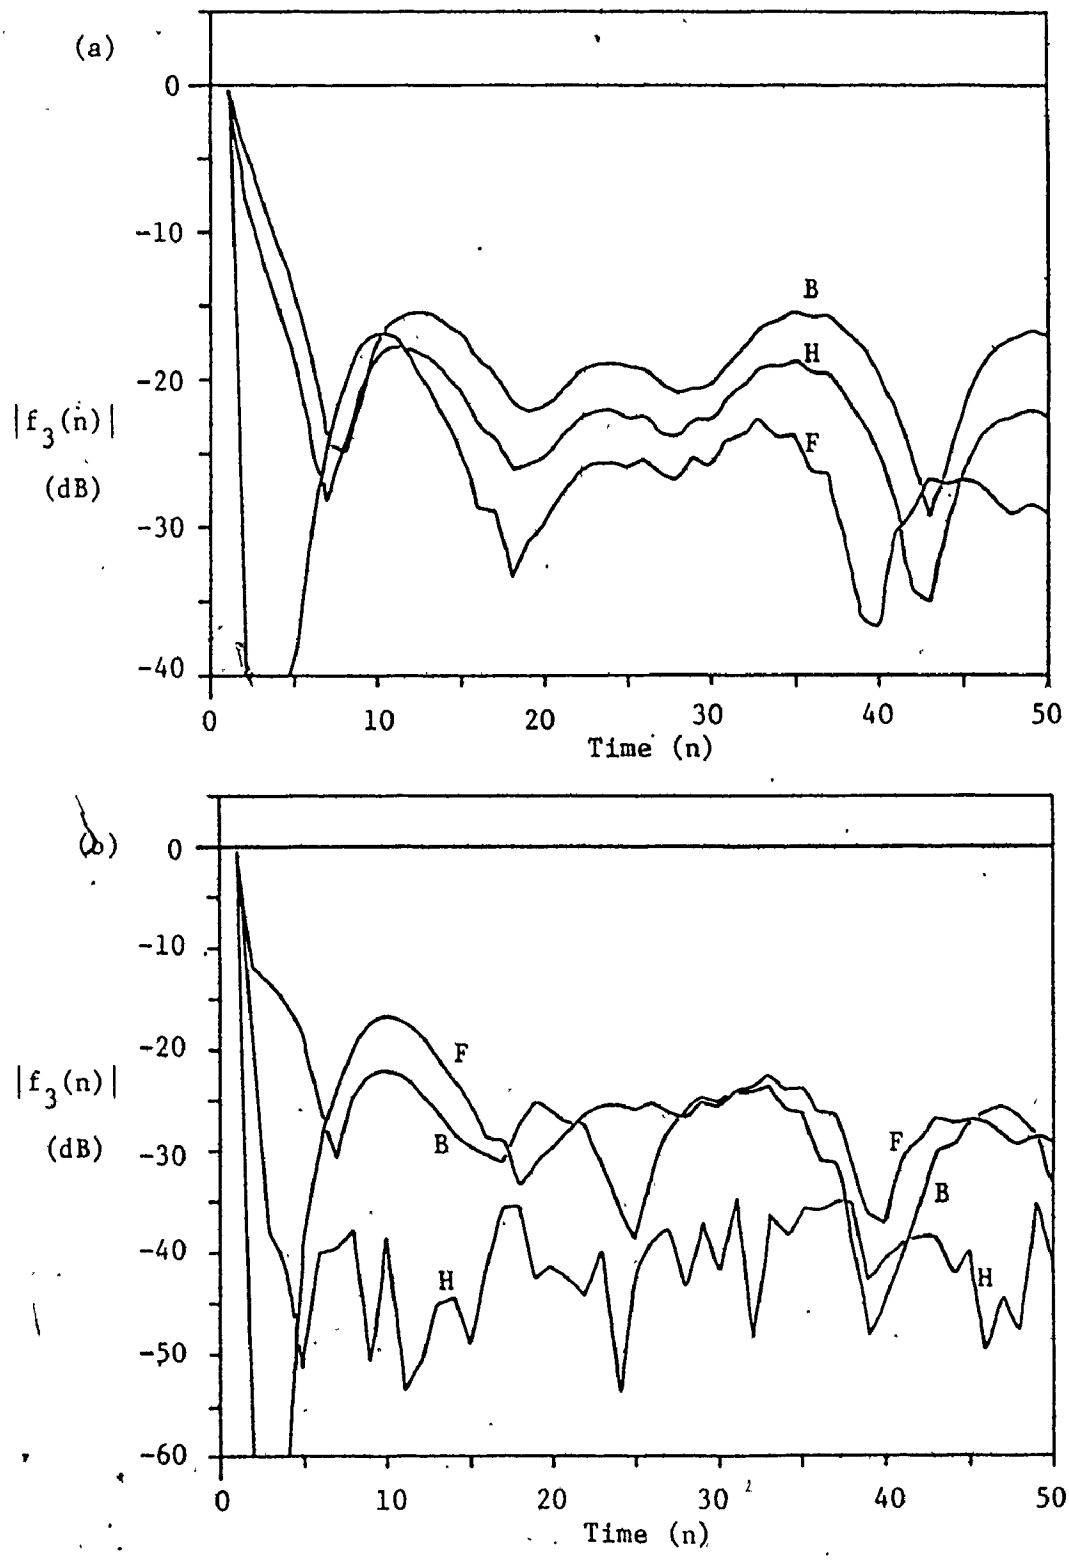

Figure 5.6: Convergence Behaviour for Radar Clutter Comparing F, B, and H Algorithms (Standard Gradient Method, $\mu = 1.0$, $m = 3$)

(a) Normal Initial Conditions, (b) Pre-Loaded Initial Conditions

to the fact that the forward algorithm favours a time series growing in amplitude with time, while the backward algorithm favours a time series decreasing with time. Because of the initial conditions normally used, the time series in question effectively has a large growth at time n = 0 (from zero to its average value), which accounts for the observed behaviour. However, when the filter has adapted long enough to forget the initial conditions, or when it has been pre-loaded with data, as in Fig. 5.6(b), the true advantage of the harmonic-mean algorithm for time-varying data can be seen.

Finally, bearing in mind the relationship of the standard gradient method with the Kalman filtering theory (see section 3.3 and Appendix 2), it is informative to compare the convergence behaviour of the lattice to that of a Kalman tapped-delay-line filter (TDL). This is done in Fig. 5.7. In this case, the TDL filter initially converges more quickly than the lattice, but when both have been allowed to adapt, or are pre-loaded, the lattice provides better clutter attenuation. Also, with the lattice, we have the option of adjusting the adaptive constant to further improve its performance and control its adaptive rate. (There is little significant difference in performance between the forward Kalman TDL, as presented here, or the backward TDL, or a Kalman filter constrained to minimize both forward and backward TDL prediction errors).

5.4 Other Factors Affecting Performance

In this section, we shall briefly examine the effects of noise, target velocity, and changes in clutter parameters on the lattice filter

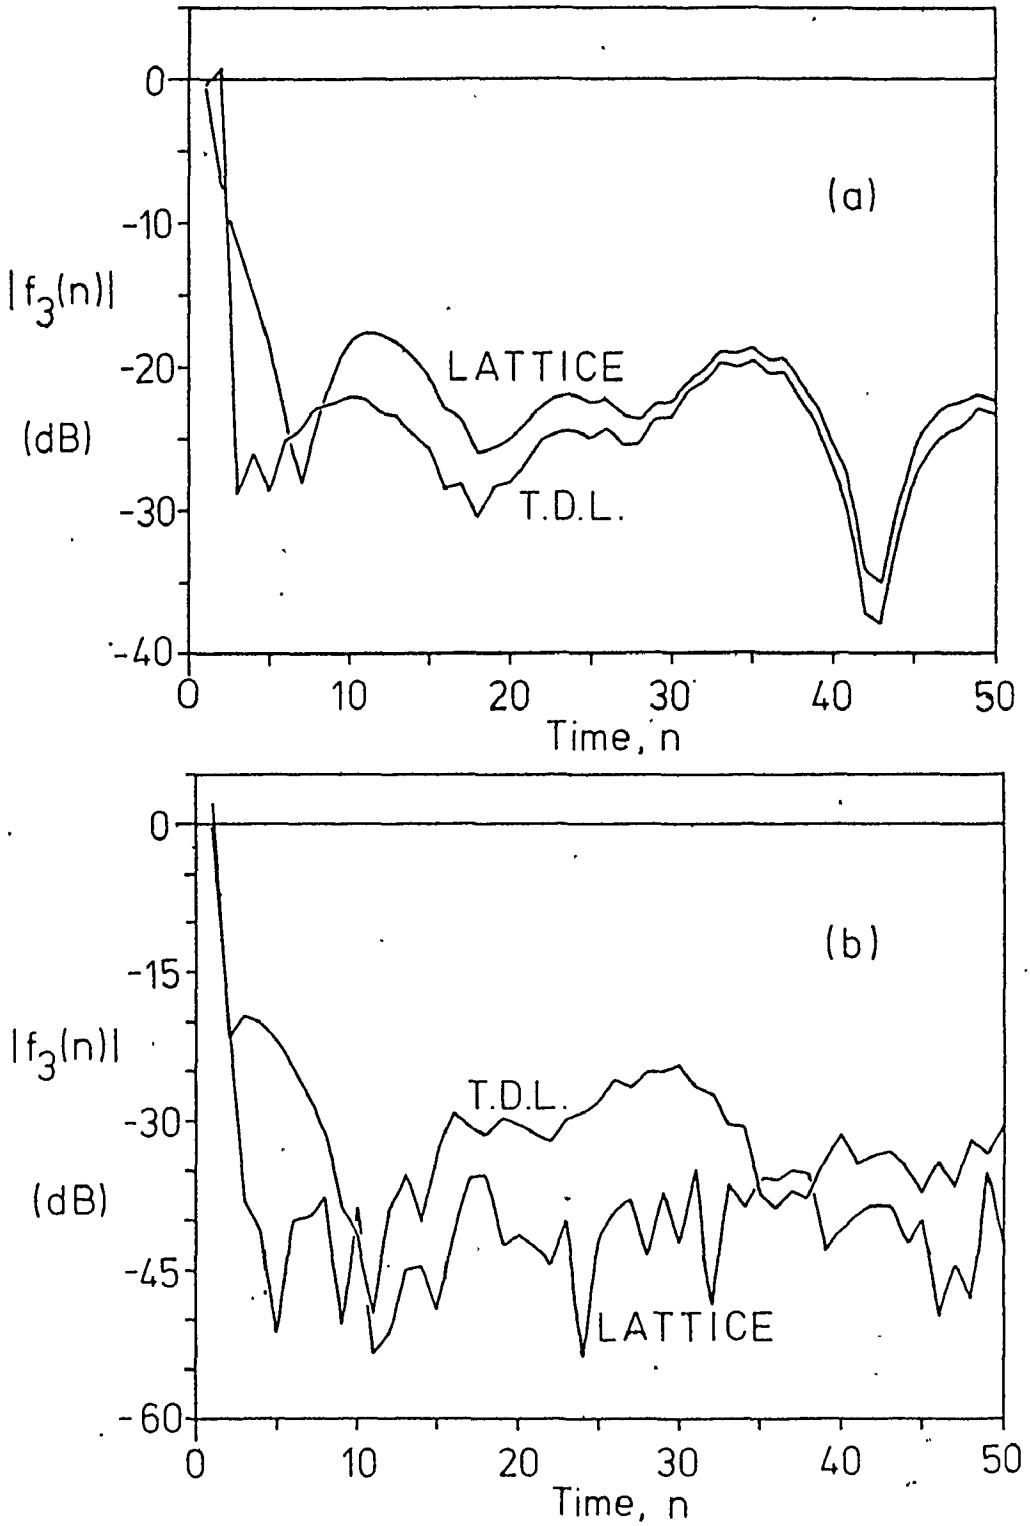

Figure 5.7: Comparison of Convergence Behaviour for Kalman Forward TDL Filter and Lattice Filter (Standard Gradient Method, $\mu = 1.0$)

(a) Normal Initial Conditions, (b) Preloaded Initial Conditions

performance [23]. (For this examination, we shall use the simple gradient method of lattice implementation with $\omega = 0.9$ and $m = 2$). Figure 5.8 plots the improvement factor as a function of target velocity, and part of the frequency response of the lattice filter. The signal is the same simulated average rain clutter data that was used previously, with varying amounts of white Gaussian noise added, as indicated.

A number of interesting points arise from Fig. 5.8. First, although the filter performance is progressively degraded by the addition of noise, this degraded performance is not ill-behaved. In fact, for high white noise conditions, an improvement factor near unity and a basically flat frequency response are to be expected, as the filter is basically inactive. Thus, the filter is behaving properly in the presence of noise. Second, at no point does the improvement factor drop significantly below unity, which is important because otherwise targets might be lost due to the action of the filter (this happens with conventional MTI filtering). Third, the improvement factor curves roughly reproduce the frequency response curves of the filter, as might be expected. The adaptive constant has some effect on the shape of these curves, especially those of the improvement factor. Reducing the constant results in a lower overall improvement curve, as the filter partially adapts to the targets. Increasing the constant above $\omega = 0.9$ tends to raise the low point of the improvement curve.

The observations given above also apply when the clutter conditions are altered significantly. Figure 5.9 plots the same parameters as Fig. 5.8, this time using a simulated clutter signal corresponding to a very severe storm, moving through the radar area with a significant

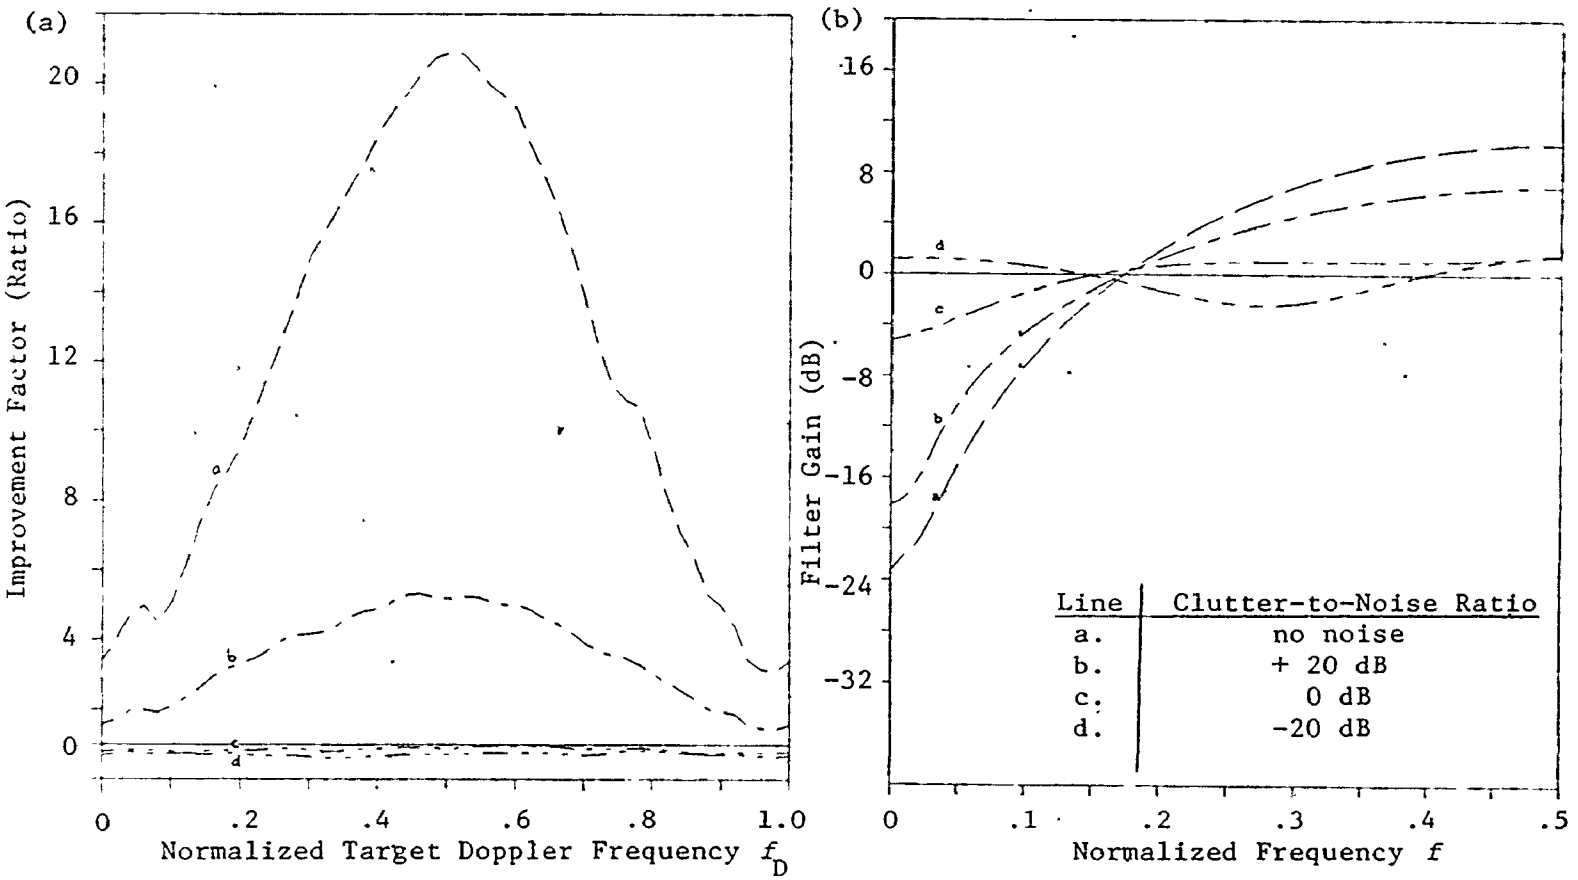

Figure 5.8: Effect of Noise on Lattice Filter Response for Average Rain Clutter
(a) Improvement Factor, (b) Part of Frequency Response

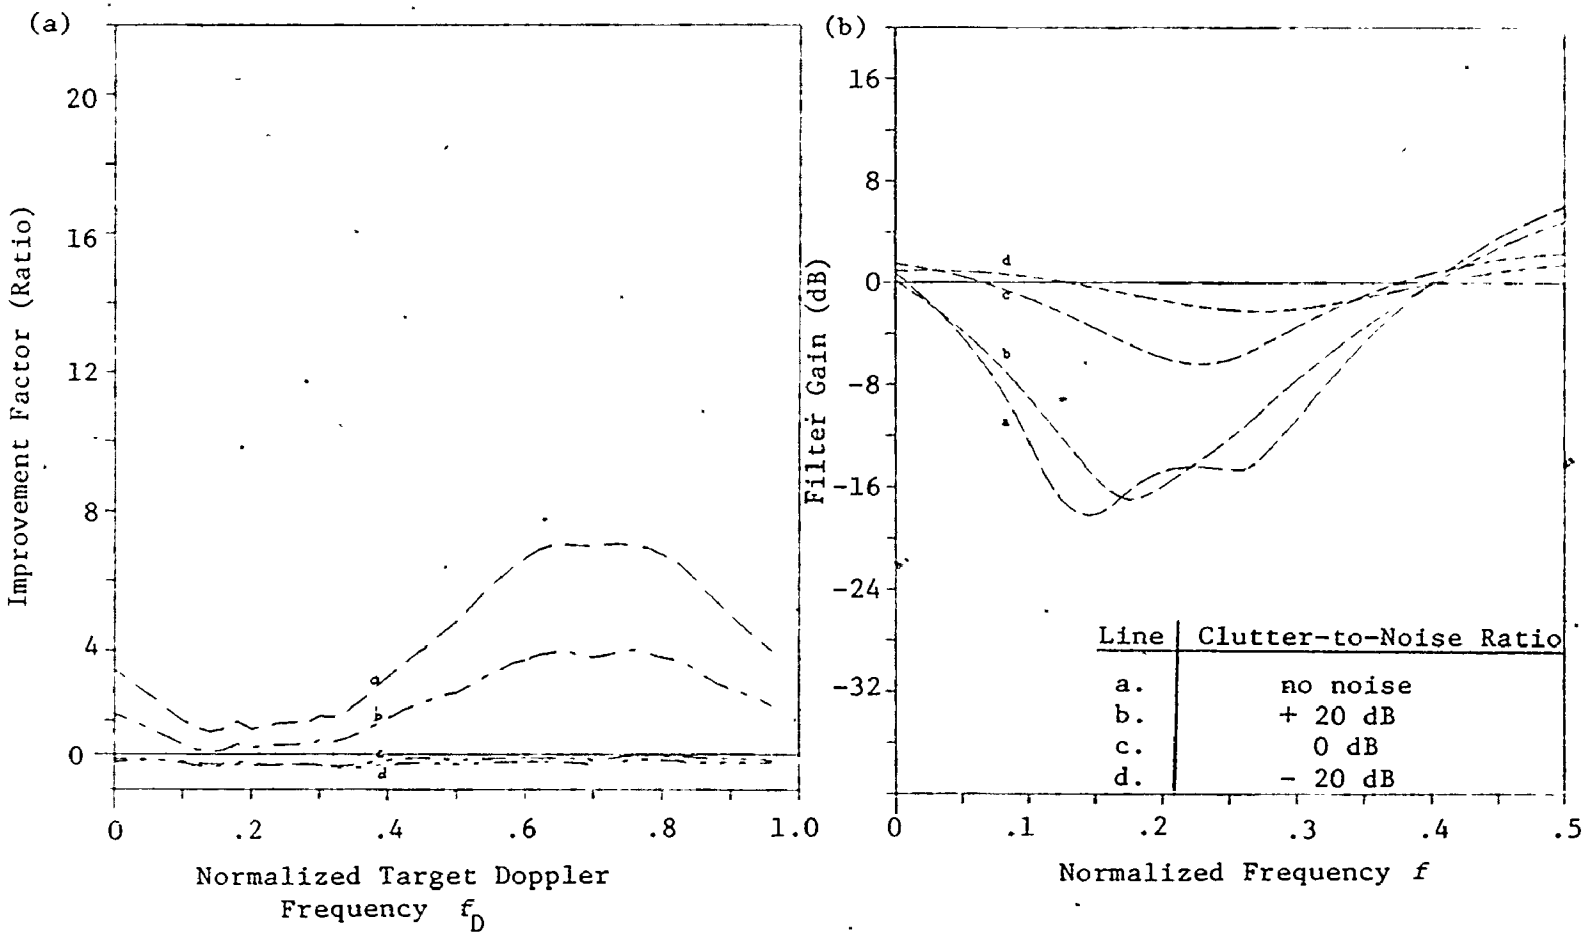

Figure 5.9: Effect of Noise on Lattice Filter Response for Severe Storm Clutter
(a) Improvement Factor, (b) Part of Frequency Response

radial velocity. Two additional observations may be made about this figure. First, the improvement factors are reduced from those in the previous figure, as might be expected due to the fluctuating nature of the severe storm clutter data. Second, the response curves are no longer centered at zero frequency, due to the radial velocity of the storm. This is unlike most conventional clutter filters, which are generally centered at zero frequency. Thus, using the lattice filter would improve the chance of detecting a target with zero radial velocity obscured by a moving weather disturbance. Overall, the adaptive lattice filter performed reasonably well with the strong clutter, especially considering that the conventional MTI filter would actually degrade the radar performance under these conditions.

Finally, a limited examination was made of the probability of false alarm versus probability of detection statistics for the above cases. The results were much the same for both types of clutter, with the response growing worse as the target signal became weaker, and as the noise became stronger. However, such performance measures become much more meaningful when applied to actual radar data, as will be done in the next chapters. (Some preliminary results, based on actual radar data, have been published by the author [50].)

CHAPTER 6

ANALYSIS OF ACTUAL RADAR DATA

6.1 Introduction

In the previous chapter, simulated radar data was used to analyse some properties of the lattice-structure PEF. In Chapters 6 and 7, actual radar data will be used as the basis of our study. It will be seen that the simulated and actual radar clutter have significantly different characteristics, necessitating a corresponding adjustment of the lattice PEF parameters. In particular, the actual clutter is quite non-stationary and more random than the simulated data. This chapter examines examples of the various types of actual radar clutter and target data which were recorded, and the resulting responses of standard MTI and lattice structure filters. Chapter 7 presents performance-measure statistics based on this data.

The data used for this analysis was collected over the period of October 2 - 15, 1978, at the Bagotville Canadian Forces Base in Quebec, It was recorded from a Texas Instruments ASR-8 surveillance radar (air-traffic control) on a wideband video recorder for later analysis on the HP computer of the Communications Research Laboratory at McMaster University. Details of the radar and recording system are provided in Appendix C. The data was selectively sampled and transferred into the computer using a control and interface unit designed as a part of this thesis,

and described in Appendix D.

Computational requirements limited the amount of the recorded data which could be fully analysed. Selected range rings of data were analysed, including several examples of every type of clutter present in the data. In all cases, the performances of the filters were very similar for all the samples of a particular clutter type, even though these samples were from different ranges and recorded at different times. In this presentation, generally only results from one or two examples of each clutter type will be shown. The clutter analysed this way includes rain, snow, ice pellets (freezing rain), inversion, and ground clutter. In addition, all aircraft and bird flocks which could be found in the data and properly identified were sampled.

The results presented in this chapter consist mainly of plots of the actual radar signal and the associated filter output signals. These plots are presented in the form of signal amplitude versus sample number (corresponding to azimuth) for a small section of a range ring having typical (or special) characteristics. Accompanying some of these plots are periodograms of the signals. There are generally four lines on the plots, these being the input signal and the outputs of a multi-pulse canceller MTI filter (see Appendix A for details) and two lattice-structure filters, one using the simple gradient method and the other using the standard gradient method for adaptive implementation. Unless otherwise noted, a filter order of five was used for the comparison of the three filter types.

The radar installation from which the data was collected made use of a staggered pulse-repetition frequency (PRF), with four sequenced

periods (as discussed in Appendix C). This can present a problem to the lattice-structure filter, which assumes a constant time period between samples. The solution used in this thesis was the parallel operation of four separate lattice filters, each taking only every fourth sample in series. (Actually, the same filter was used with four separate memories). This solution was compared with the use of one lattice filter on all the data, and was found to be slightly better. When plotted, the outputs of the four lattice filters were averaged together for simplicity, as were the input and MTI output signals for the same four sample pulses.

Because of the parallel filtering structure, and the non-stationary nature of the data, it could be expected that the optimum values of the adaptive constants used for the lattice-structure filters would be lower than those of the stationary simulated data of the previous chapter. In fact, as will be shown in section 7.1, the best performance is achieved with $\omega = 0.0$ for the simple gradient method and $\mu = 0.1$ for the standard gradient method. At these values, the filters are able to partially adapt to the targets as well as the clutter, and target detection is based on the response to the transition region between clutter and target. In this region, the signal statistics belong to neither clutter nor target, but are changing between the two. These values of ω (0.0) and μ (0.1) are used throughout this chapter, unless otherwise noted.

6.2 Target Data

6.2.1 Aircraft

In the course of examining the recorded data, six aircraft targets which could be identified were found. Of these, one was later discarded as unreliable, leaving the five described in Table 6.1. (These five targets may seem to be a small data base from which to form any conclusions or derive any statistics. However, they are representative of the mid-size aircraft found at many airports, and more importantly, cover the range of radial velocities, including zero, typical of such aircraft. As such, they do form an acceptable data base). The original identification of these aircraft took place at the airport where the radar data was being recorded. Figure 6.1 shows typical primary radar PPI (plan-position-indicator) displays available from the radar for air-traffic control. Figure 6.1(a) is a "normal video" display of ground clutter surrounding the radar location, while Fig. 6.1 (b) shows the "MTI video" display of an aircraft shortly after take-off. (The radar processor which created these displays is described in Appendix C).

TABLE 6.1: Aircraft Targets

Aircraft Record #	Description	Range (nautical miles)	Predominant Movement
1	Boeing 737 jet (in light snow)	4½	Radially (inward)
2	Boeing 737 jet	22	Radially (inward)
3	Viscount propjet	7	Tangentially
4	Viscount propjet	12	Radially (outward)
5	Hellcopter (military)	7	Radially (inward)

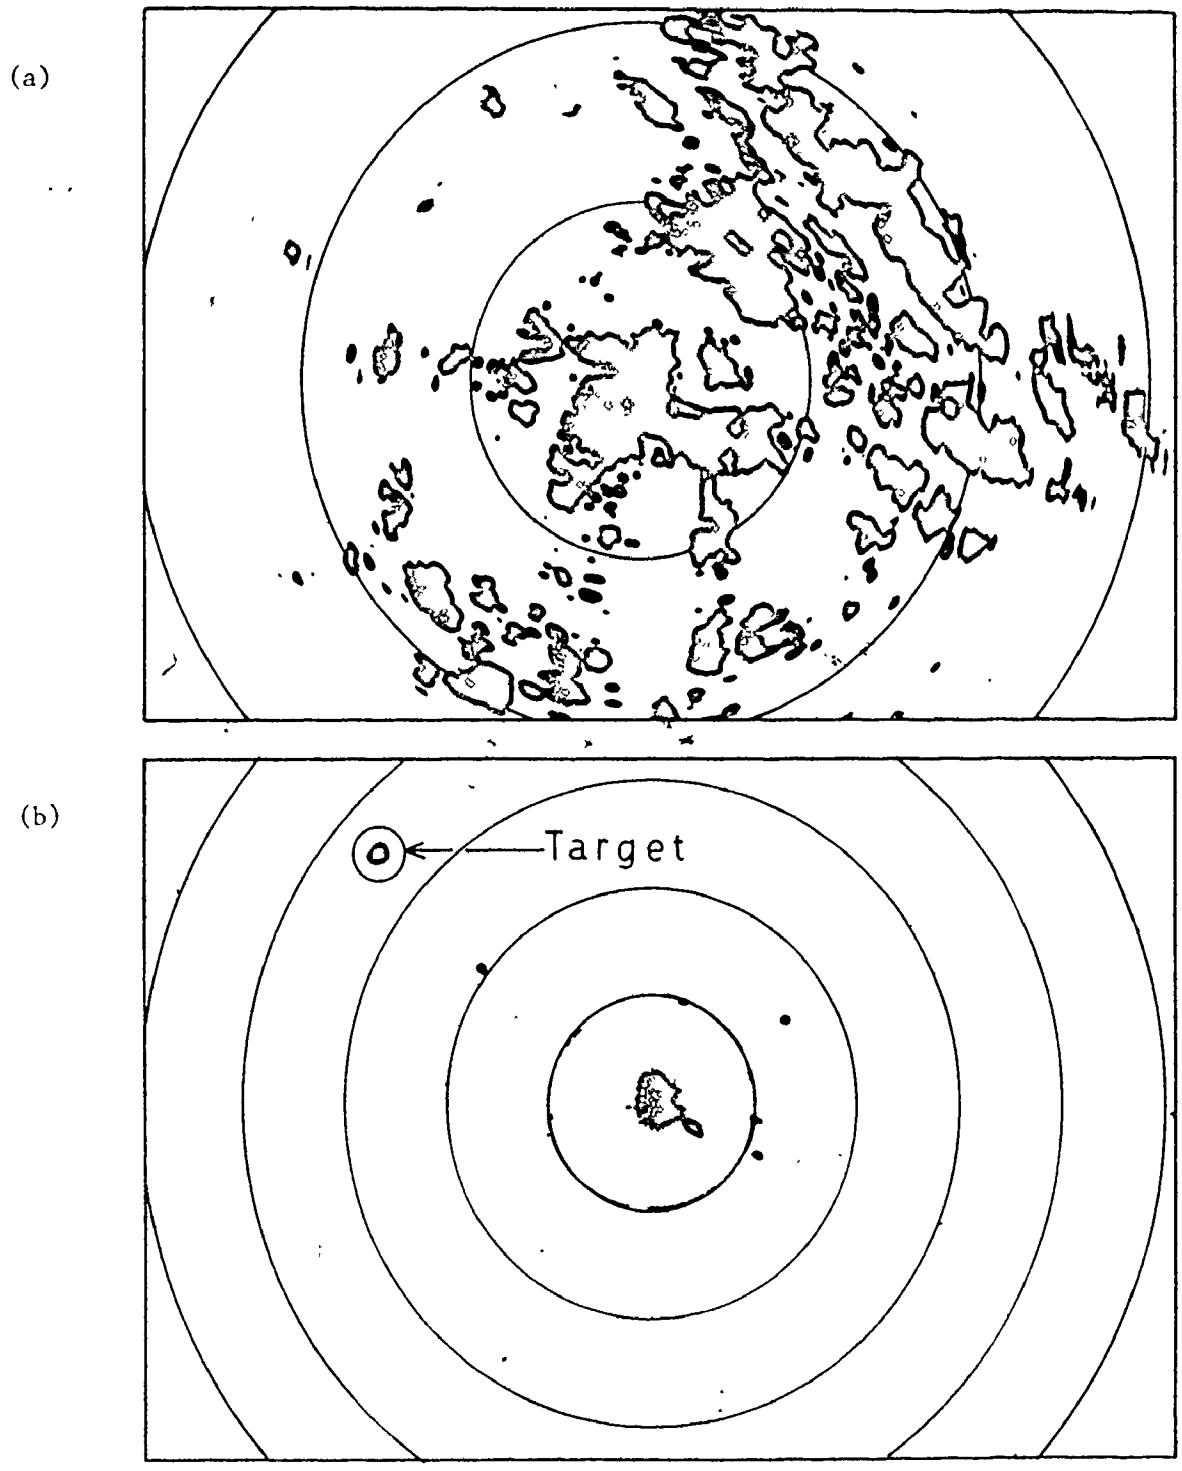

Figure 6.1: Operating PPI Displays (Negatives of Actual Photographs)
(a) Enhanced "Normal Video" (no targets visible), 10 nautical mile range rings
(b) "MTI video" (aircraft at 3½ nautical miles in NW), 1 nautical mile range rings

After the recordings were taken back to the laboratory, they were played back into another PPI, and a suitable range ring containing the target was selected and transferred into the computer for analysis. An example of this is given in Fig. 6.2, for the same aircraft as above, taken a few minutes later at a greater range (aircraft record #4). This figure shows the laboratory PPI display (with the selected range ring), and the corresponding amplitude display from the computer. Besides the target, there are peaks in the amplitude display from ground clutter and a noise background (presumably from the recorder).

(In comparing Figures 6.1(a) and 6.2(a), it may appear that Fig. 6.2(a) is more finely detailed. In fact, this is because much of the data displayed in Fig. 6.1(a) was lost to the PPI in Fig. 6.2(a), because the laboratory system was geared for accurate data transferral, rather than maximum visibility. The laboratory PPI displayed only the positive half of the in-phase channel of data, with three out of four scans eliminated and without the limiting/integration enhancement of the radar PPI display).

Let us now turn our attention briefly to the response of the lattice and MTI filters to the aircraft target signals. (Aircraft record #5 will be used as an example). Figure 6.3 shows the amplitude response of the input and outputs using the standard MTI and lattice-structure filters (using the two methods of adaptation) for this target. A very noticeable factor in Fig. 6.3(a) is that the target peaks in the filter outputs are much lower than the peak in the input. This is typical of the lattice-filter outputs, but the MTI filter output is generally closer to the input peak level for many targets. The lower level here

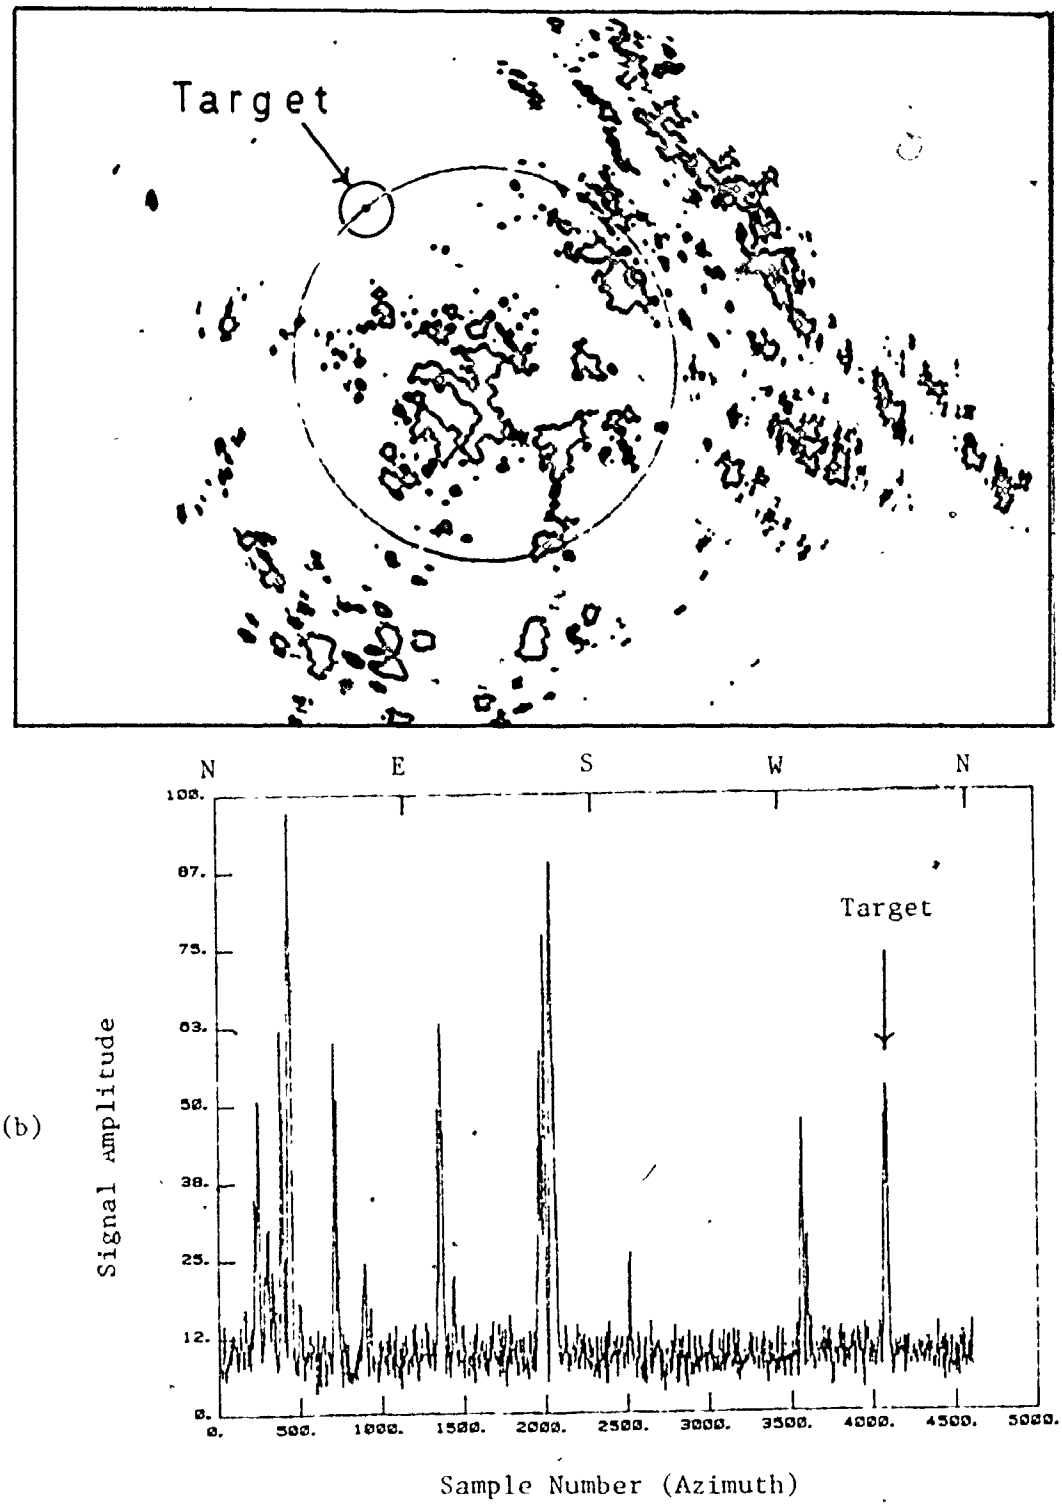

(b)

Figure 6.2: Data for Aircraft Record #4

(a) Laboratory PPI Display (aircraft at 12 nautical miles in NW, crossing ring)
(b) Signal Amplitude Plot of Range Ring shown above

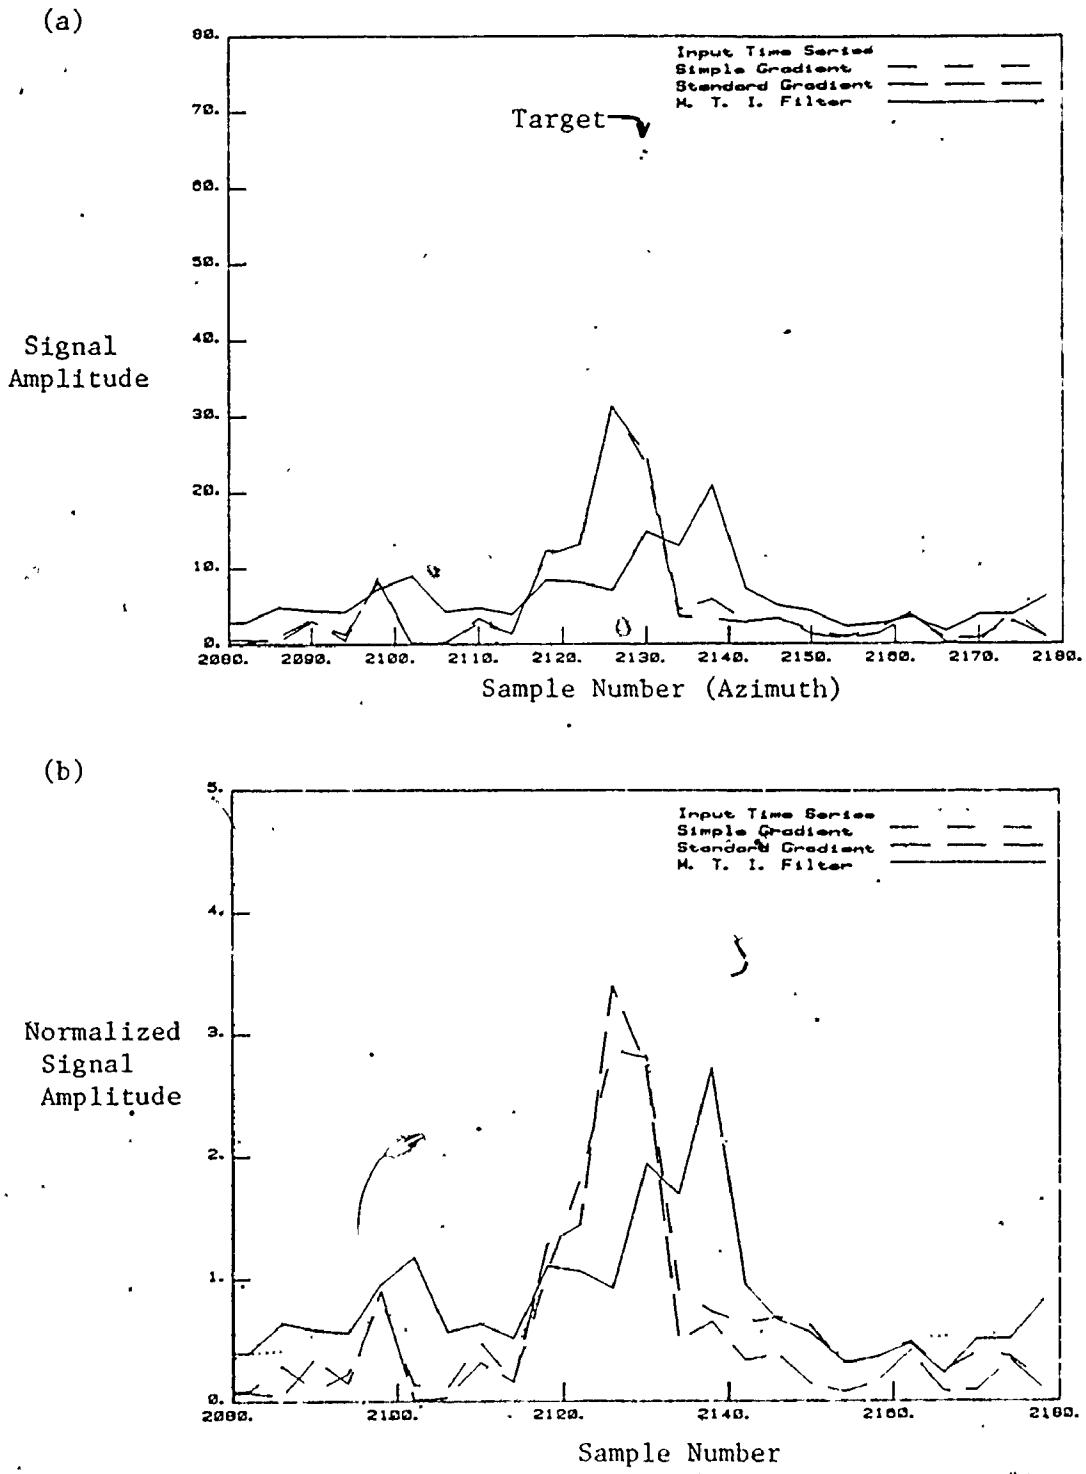

Figure 6.3: Comparisons of Filter Outputs for Aircraft Record #5
(a) Absolute Amplitude
(b) Amplitude Normalized Over 100 Samples

is due to the target's relatively low radial velocity. (The lattice filter's output peak level is basically insensitive to target velocity). A less noticeable but equally important factor is that the background noise and clutter are also reduced in these outputs, so that when the plotted signals are normalized (as in Fig. 6.3(b)), the peak values assume their proper prominence. In the normalized plots of all five aircraft targets, the lattice-structure filters produced higher peak responses than the MTI filters, and the lattice response was approximately the same for all five aircraft, although the input signal peak structures varied significantly. (Plots of aircraft data in the presence of clutter will be presented later in this chapter).

The response of the filters to the above aircraft signal in the frequency domain is also of interest. Figure 6.4(a) plots the periodograms, calculated as the squared magnitude of the discrete Fourier transform of the target area (the central 40 data samples of Fig. 6.3(a)), covering the frequency range of the aircraft's Doppler signal. This can be compared with the corresponding periodograms of a faster moving aircraft (aircraft record #1) shown in Fig. 6.4(b). The most noticeable difference between the two plots is in the response of the MTI filter, which for the faster aircraft responds strongly at the Doppler frequency, and for the slower aircraft misses this frequency completely. (It does, in this case, respond to some higher frequency harmonics from the slower aircraft, which are not shown here). The lattice-structure filters, on the other hand, show basically the same responses for both target signals. (The relatively low level of these responses is not a problem; it is largely due to the

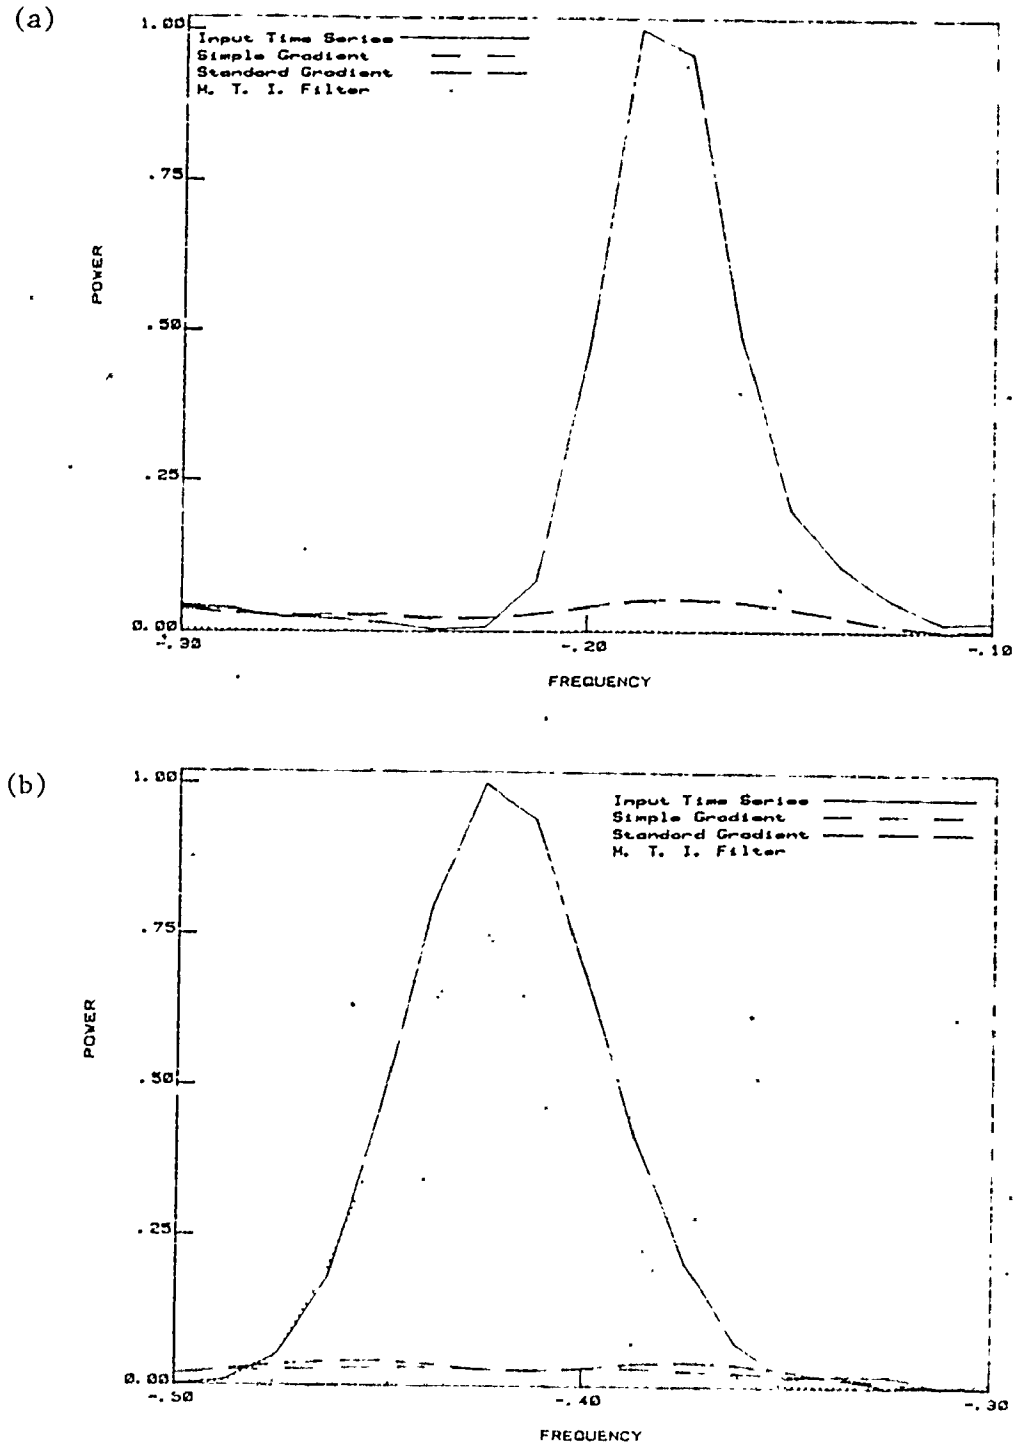

Figure 6.4: Periodograms of Aircraft Target Data
(a) Aircraft Recording #5
(b) Aircraft Recording #1

lack of normalization of the output data, as well as the partial whitening of the output spectrum by the filter).

Finally, Fig. 6.5 compares the output of a lattice-structure filter using the adaptive gradient method with the output using the simple gradient method ($\omega = 0.5$) for the same aircraft as above. As can be seen, the outputs are very similar. Other comparisons have shown that, in terms of both amplitude response and improvement factors, the adaptive gradient method performs almost identically with the simple gradient method set at $\omega = 0.5$, for this data.

6.2.2 Bird Flocks

The inclusion of bird flock-data as a type of target signal may at first seem contradictory, as radar returns from birds are generally considered to be a form of clutter. But there are several good arguments for considering bird flocks as targets. The first of these is the nature of the data itself. As will be seen, the data from bird flocks is very similar to that from slow-moving aircraft. Because the spatial distribution of birds within the flock is typically close to the limits of resolution of the radar cell, the flock often appears to the radar to be essentially a point-source reflector. Also, the velocity spread of birds within the flock is generally less than or of the same order as the velocity spread of an aircraft target, due to antenna scanning modulation. (It should be noted that this thesis deals with birds in flocks only. Birds moving independently over a wide area could be expected to have completely different characteristics).

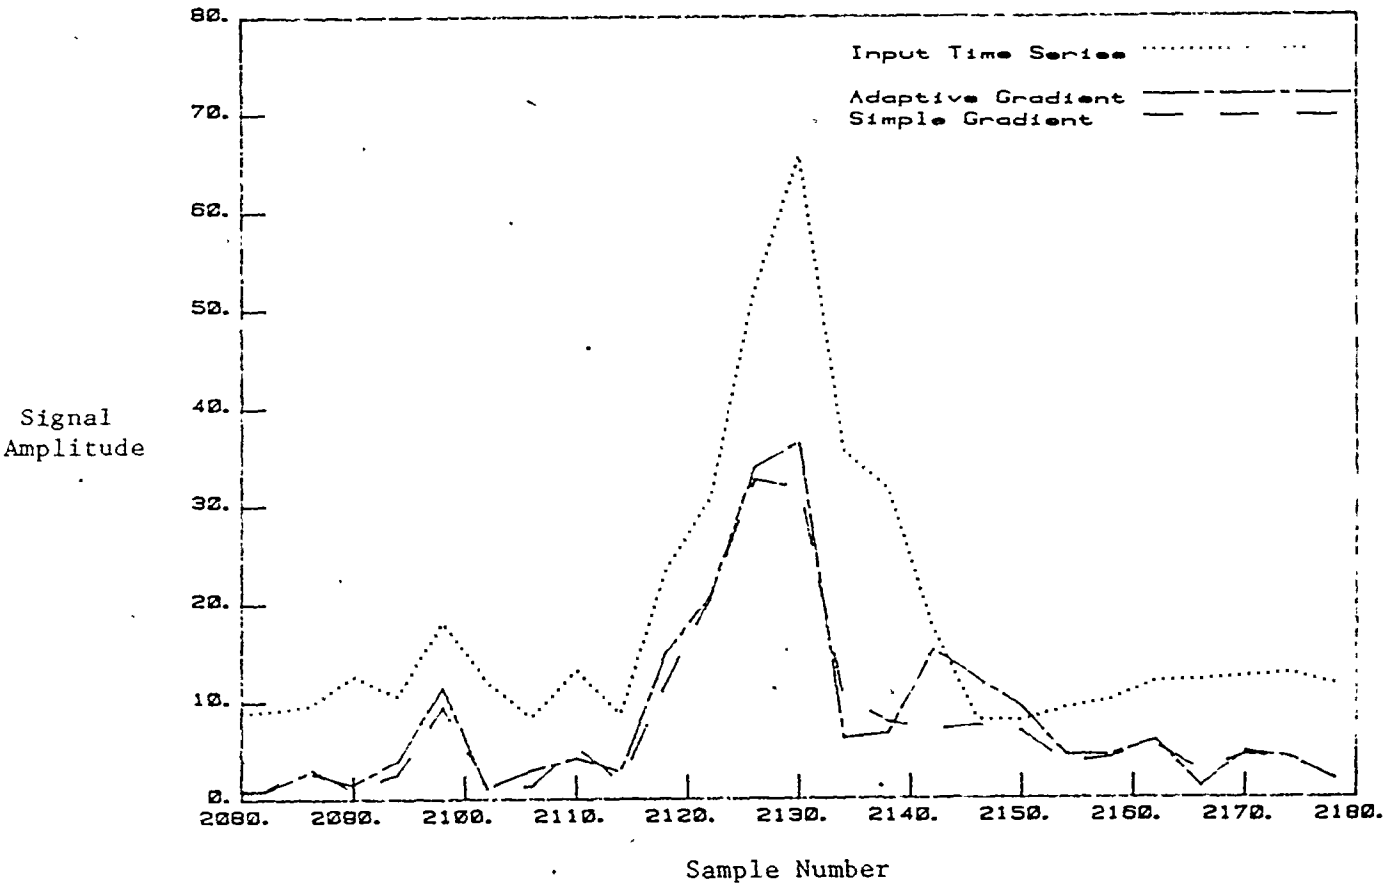

Figure 6.5: Comparison of Adaptive and Simple Gradient Methods of Lattice Filtering for Aircraft Record #5

The second major argument for considering bird flocks as targets is the consideration of air safety. Birds pose a real collision threat to aircraft, especially when the aircraft are landing and taking off. One New York airport reported forty million dollars worth of aircraft damage from bird collisions in the year 1975 alone, and in Canada, twelve Starfighter jets have been lost due to bird strikes over an eight-year period. Thus, it seems very desirable to treat bird-flock returns as targets to be avoided, rather than as clutter to be ignored and eliminated.

The performance of the standard MTI filters presently used by most radars can best be described as hit-and-miss for this sort of data. Faster flocks are detected as targets, while returns from slower flocks are eliminated as clutter. The lattice-structure filters presented here, however, detect the flocks as targets with a high degree of reliability. In practice, this detection of bird flocks as targets could lead to a confusion between the flocks and slow aircraft. This is not a serious problem, however, as the basic safety procedure for both types of targets is the same; other aircraft must avoid this area. On the other hand, to mistake bird flocks for another type of clutter would be extremely dangerous. The consistent detection of bird flocks as targets could make air travel much safer.

The time period over which this data was recorded at Bagotville was chosen to coincide with the annual southward migration of snow geese through the area. Although the expected large-scale migration did not occur in the vicinity of the radar location, a number of smaller flocks

did pass through the area and were recorded. A complicating factor in dealing with the bird-flock data was the desire to have visual confirmation of the flocks recorded, which effectively limited the range to a very few miles around the radar. At these ranges, there was also a large contribution from ground clutter, making it difficult to obtain a confirmed recording of birds and nothing else. However, it was possible to isolate a few examples of recorded bird-flock clutter signals, which are described in Table 6.2. (Although these flocks resembled snow geese in size and behaviour, it was generally impossible to make a positive identification at the ranges involved).

Figure 6.6 presents examples of the normalized response of lattice and MTI filters to two of the bird-flock records, typical of the variety present in this data. The top plot belongs to a large, fast-moving flock, and both data and filter responses are very similar to those recorded for aircraft targets. The bottom plot belongs to a much slower flock. (There is a significant background clutter level in the lower plot). As can be seen, the MTI filter detects the one flock well and completely misses the other flock, while the lattice filters detect both these flocks (and the

TABLE 6.2: Bird-Flock Recordings

Flock Record #	Description	Range (nautical miles)	Predominant Movement
1	Large flock in formation	2¼	Radially (inward)
2	50 birds in "V" formation	5	Radially (outward)
3	(Same)	6	Radially (outward)
4	Uncertain due to range	11	Tangentially

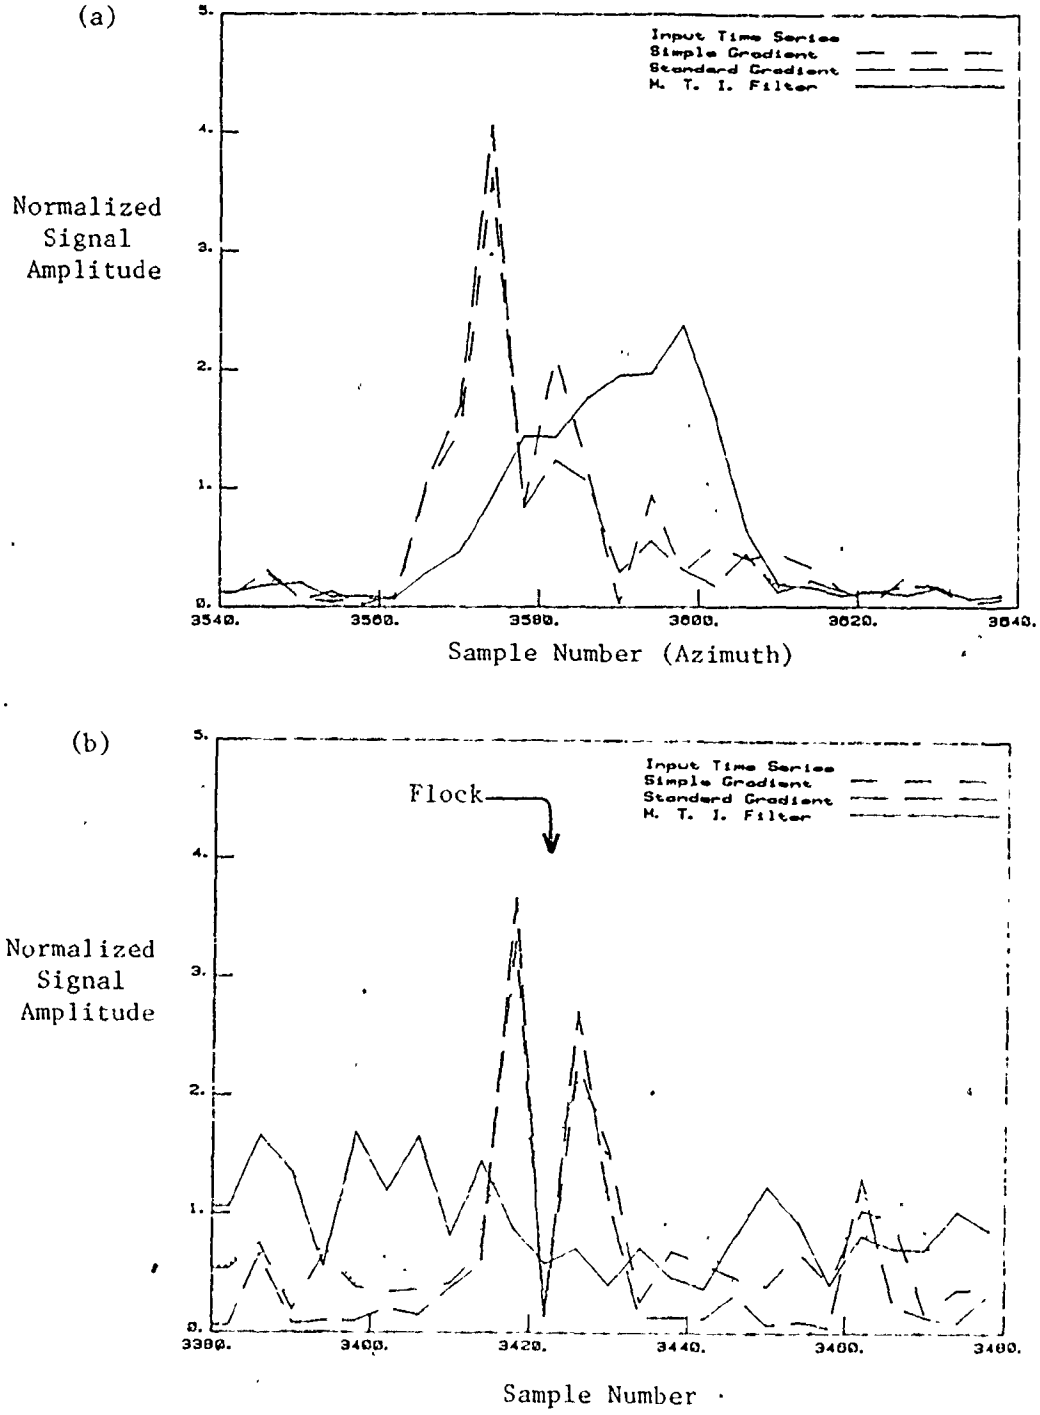

Figure 6.6: Comprisons of Filter Outputs for Bird Flocks
(a) Flock Record #1
(b) Flock Record #2

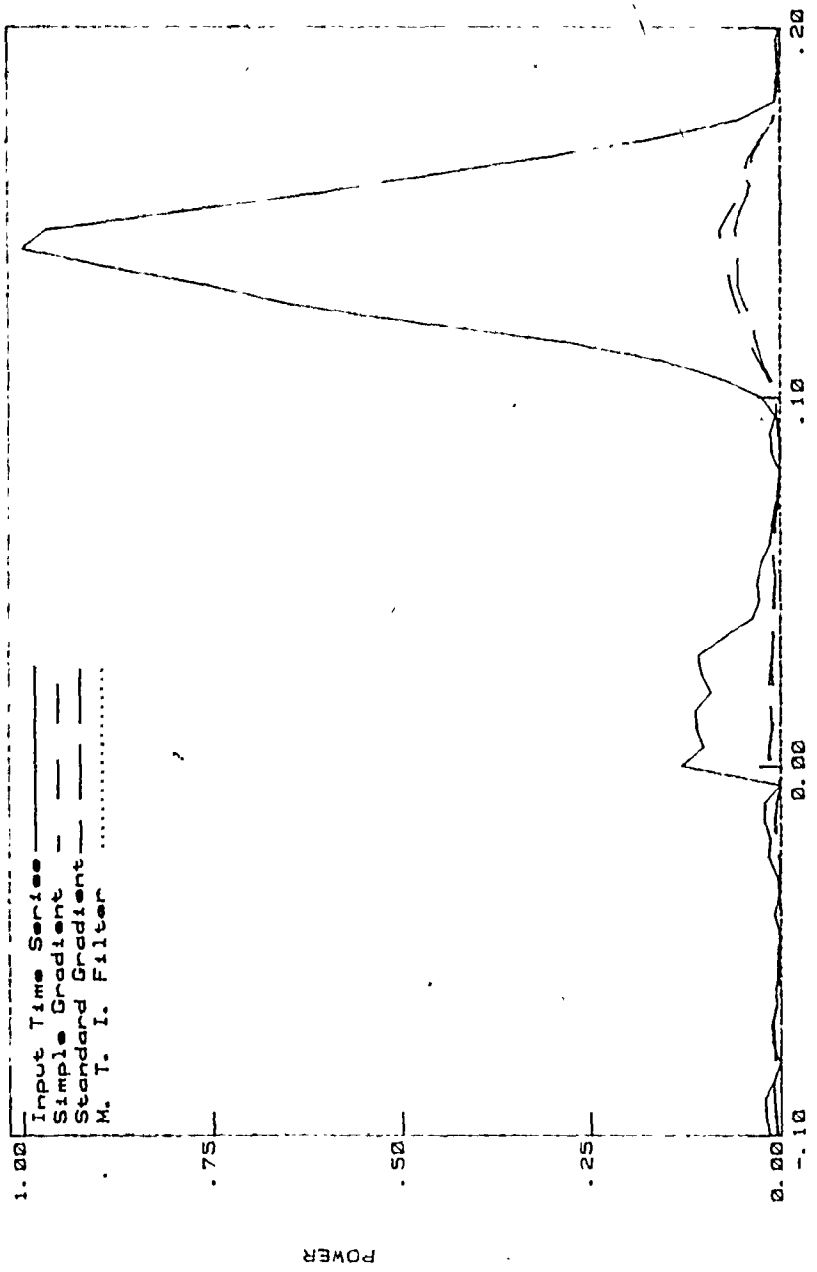

Figure 6.7: Periodogram of Bird-Flock Data (Flock Record #4)

others) equally well. Figure 6.7 shows a typical response in the spectral domain, again for a slow-moving flock. (Note the contribution of ground clutter at zero frequency).

6.3 Clutter Data

6.3.1 Rain Clutter

During fine weather conditions, when only ground clutter is present (for example, see Fig. 6.1), the performance of conventional MTI filters is generally acceptable (except for the loss of tangentially moving targets). With the addition of any weather clutter, however, the situation changes greatly. Consider the case of a "light" rainstorm (as defined by the weather office), carried by 14 knot winds from the east. The operating PPI displays during such a rainstorm are shown in Fig. 6.8. Even though the radar was using circular polarization at the antenna, which reduces the returns from rain considerably, significant areas of the screen are obscured by the rain, for MTI as well as normal displays. This is especially true in the enhancement mode, as seen in frames (c) and (d) of Fig. 6.8 which also show the effect of switching from high to low antenna beam at about 15 nautical miles. ("Enhancement" was the normal operating mode at the Bagotville installation. As described in Appendix C, it includes limiting and integration of the signal). Obviously, aircraft in the region obscured by the rainstorm could not be detected.

In order to remove the rainstorm returns from the MTI display, many radars (including the ASR-8) use what is called constant-false-alarm-rate (CFAR) processing. This processing tests each data point

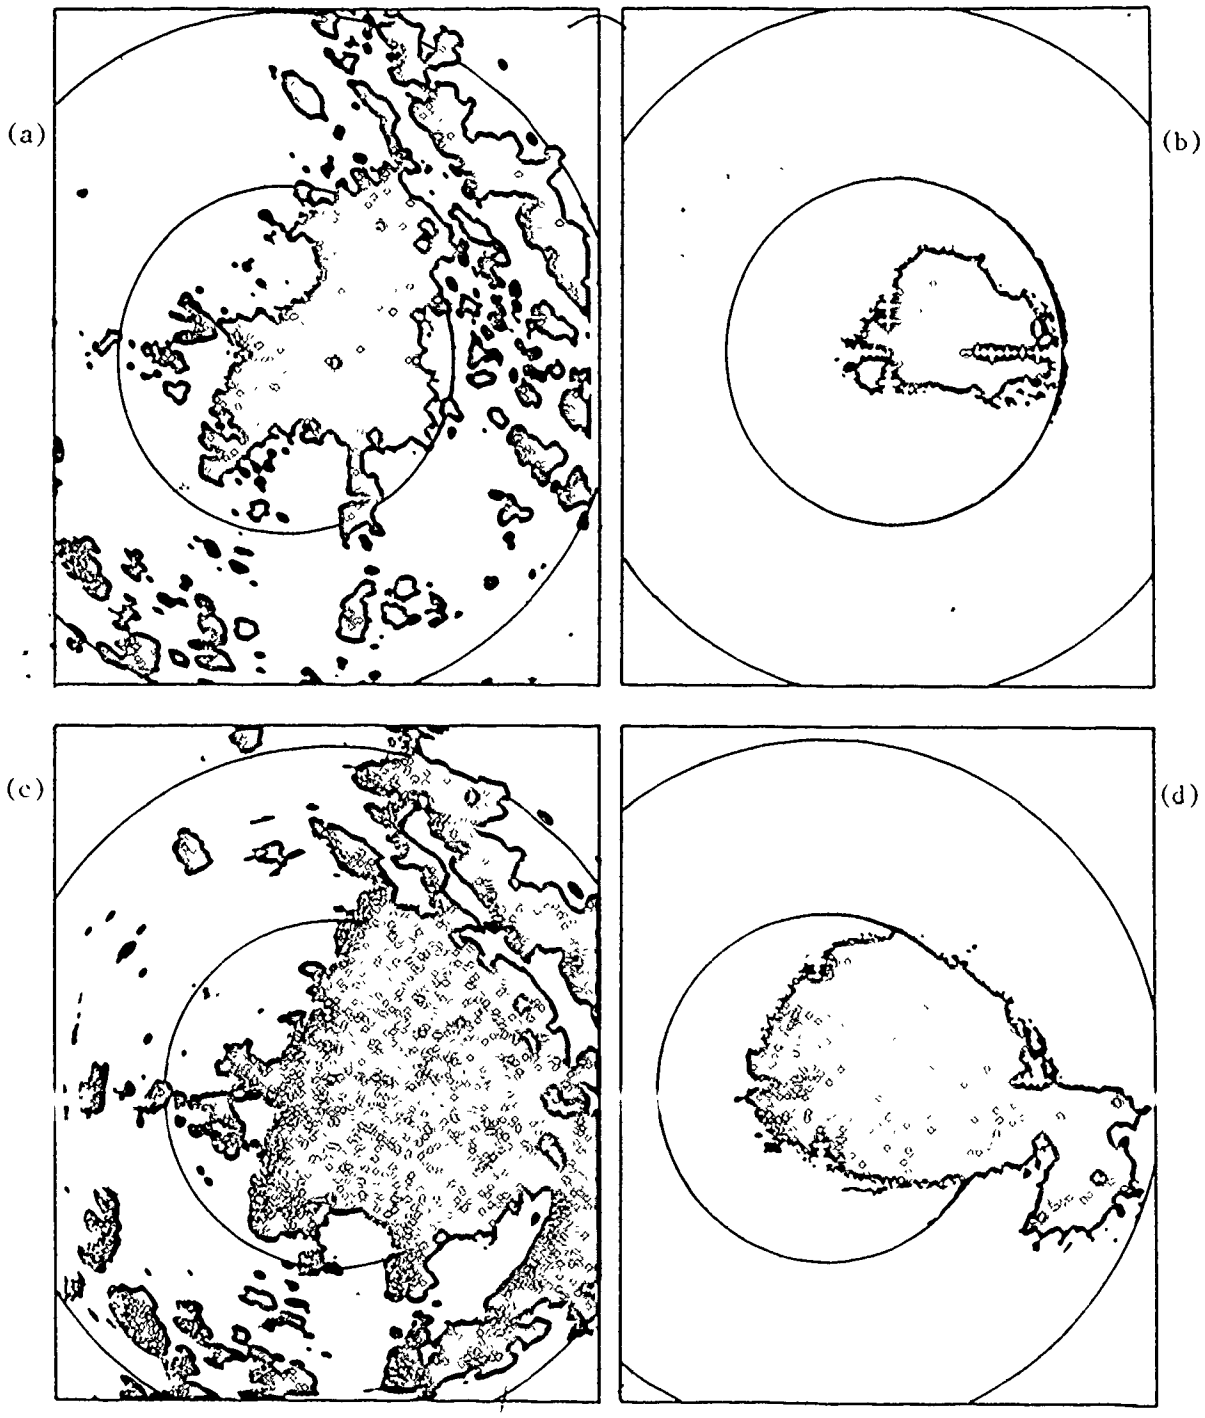

Figure 6.8: Rain Clutter PPI Displays (10 Nautical Mile Rings)

(a) "Normal" (no enhancement), (b) "MTI" (no enhancement)
(c) "Normal" (enhanced), (d) "MTI" (enhanced)

from the MTI filter against a threshold of the average value from surrounding points. Thus, detected returns covering a large area (such as weather clutter) are reduced to the level of system noise, while isolated returns (targets) are largely unaffected (in theory). An example of the application of CFAR processing, to the MTI display in Fig. 6.8(d), is shown in Fig. 6.9. The lower picture of this figure includes a weather outline map. (Details of this processing are presented in Appendix C). The CFAR display is not much of an improvement, however, as target returns occurring in regions of strong clutter (such as within the central outline of Fig. 6.9(b)) are also eliminated from the display. Thus, the conventional MTI filter's inability to distinguish between moving clutter and targets becomes a real problem. This is the problem for which the adaptive lattice-structure filter is intended to be a solution.

Figure 6.10 presents the data for a range ring sampled from the above storm. Comparing the amplitude plots of this figure and Fig. 6.2, it can be seen that the rain clutter is a large part of the signal. The MTI and lattice filter outputs are compared for sectors of this ring in the plots of Fig. 6.11. Plots (a) and (b) compare these outputs for sectors where the rain is moving tangentially (no radial velocity) and radially. All the filters perform approximately equally for the low-velocity clutter. For the higher-velocity clutter, on the other hand, the lattice-structure filters still do approximately as well as before, while the MTI filter does much worse.

The effects of varying the filter orders can be seen by comparing plots (b) and (c), for the same sector of data. In plot (b), the filter

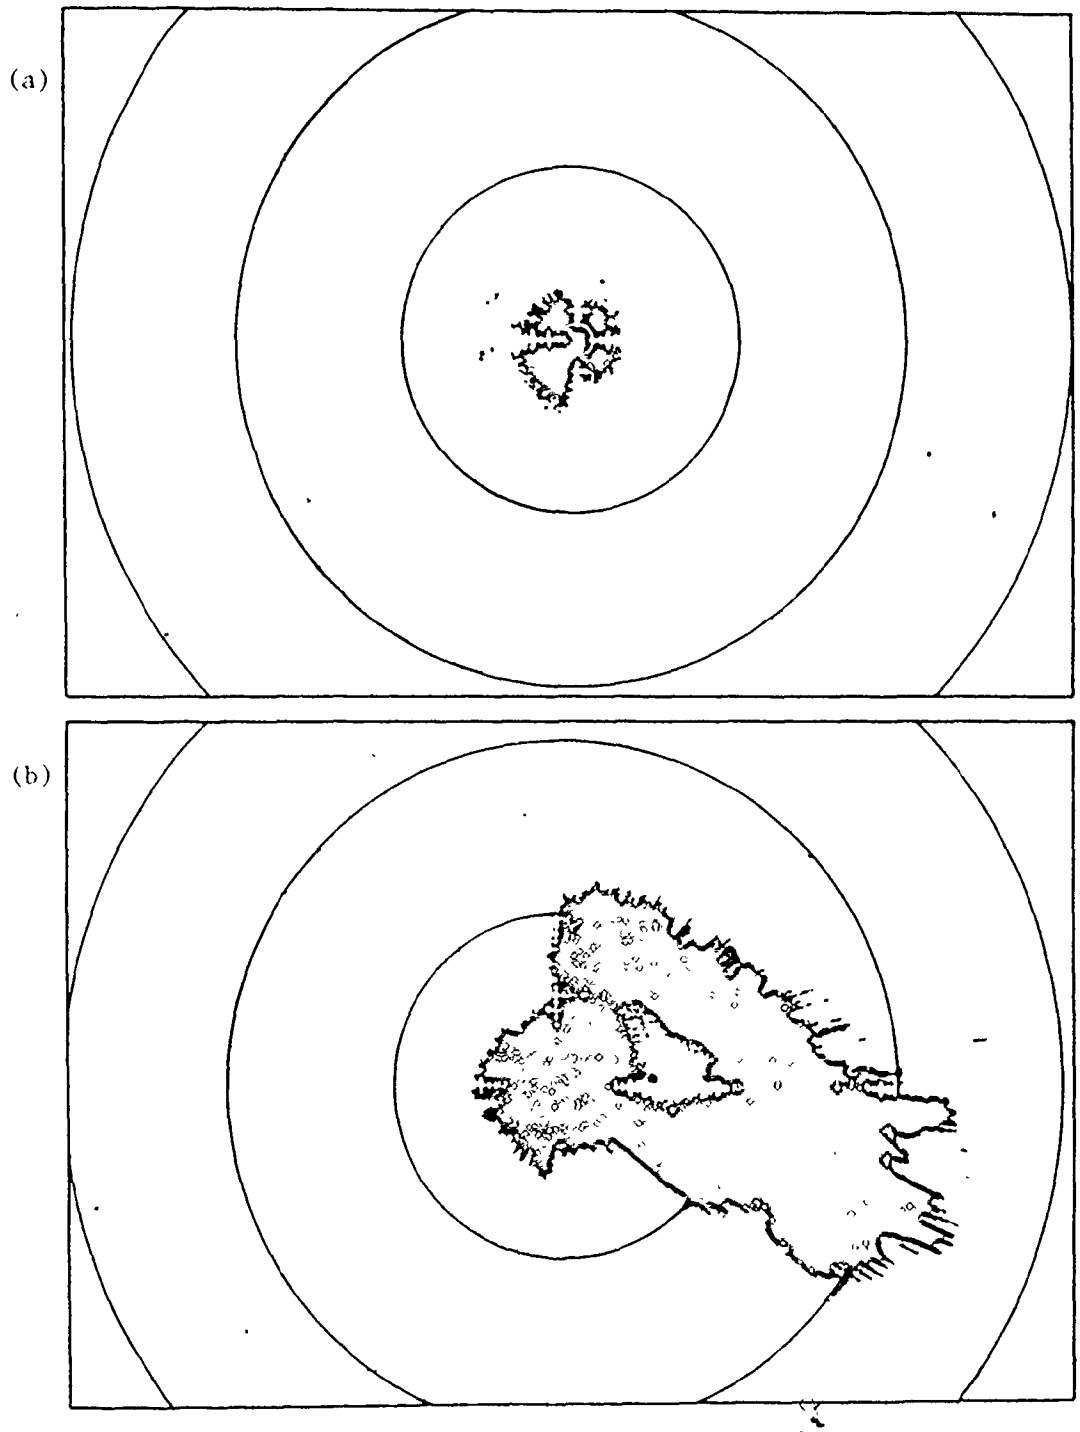

Figure 6.9: Rain Clutter PPI Displays Using CFAR (10 nautical mile rings)

(a) CFAR alone
(b) CFAR plus Weather Map

(a)

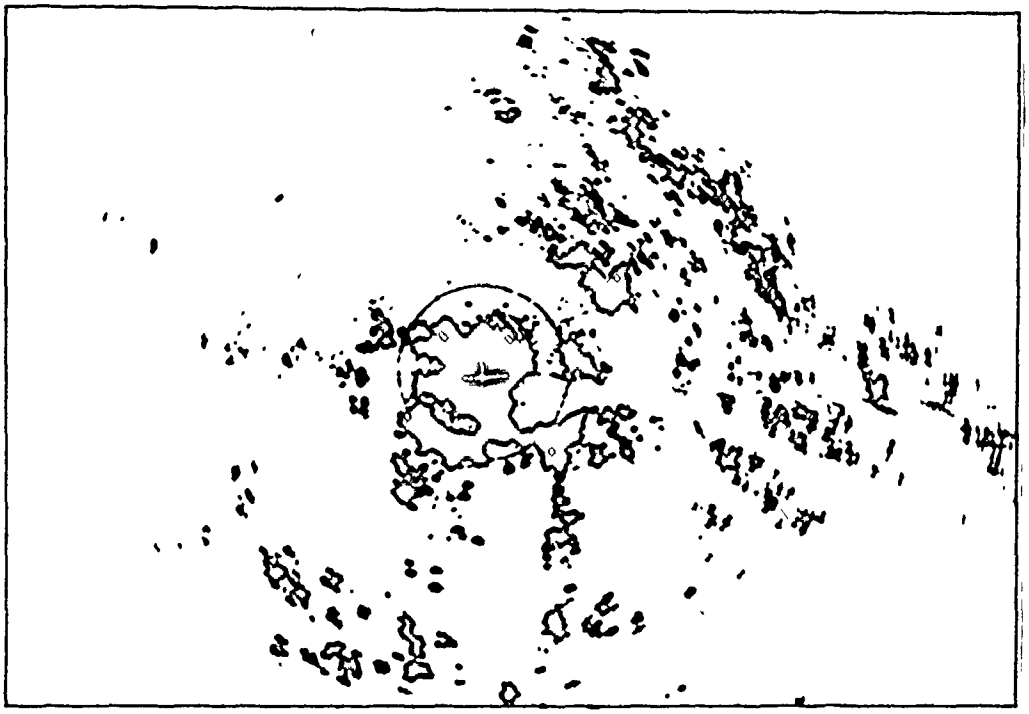

(b)

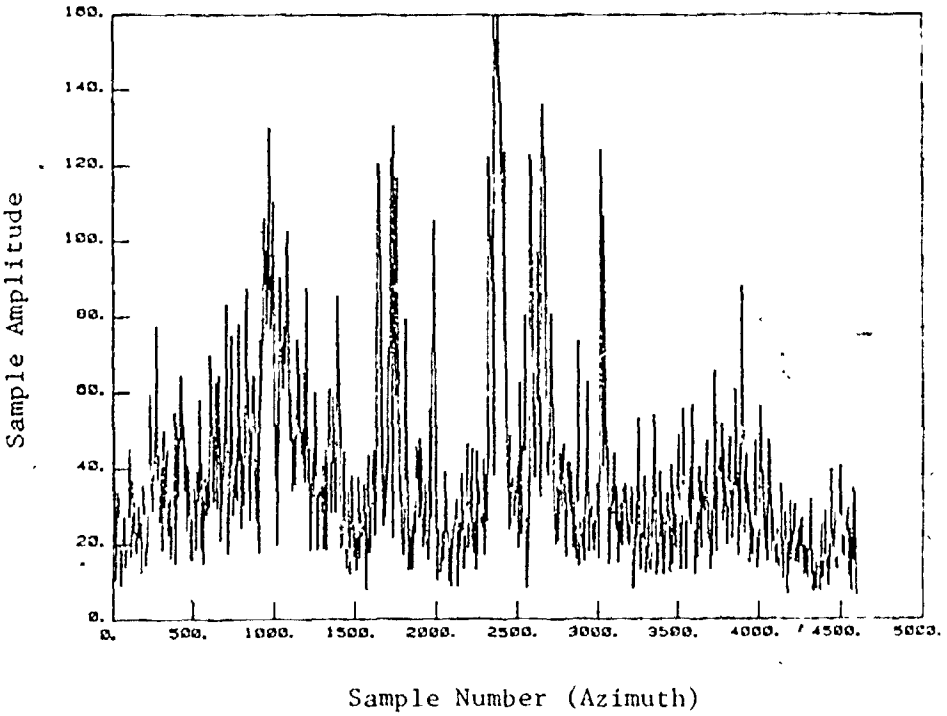

Figure 6.10: Data for Rain Clutter (range ring at 5 nautical miles)

(a) Laboratory PPI Display
(b) Amplitude Plot of Range Ring shown above

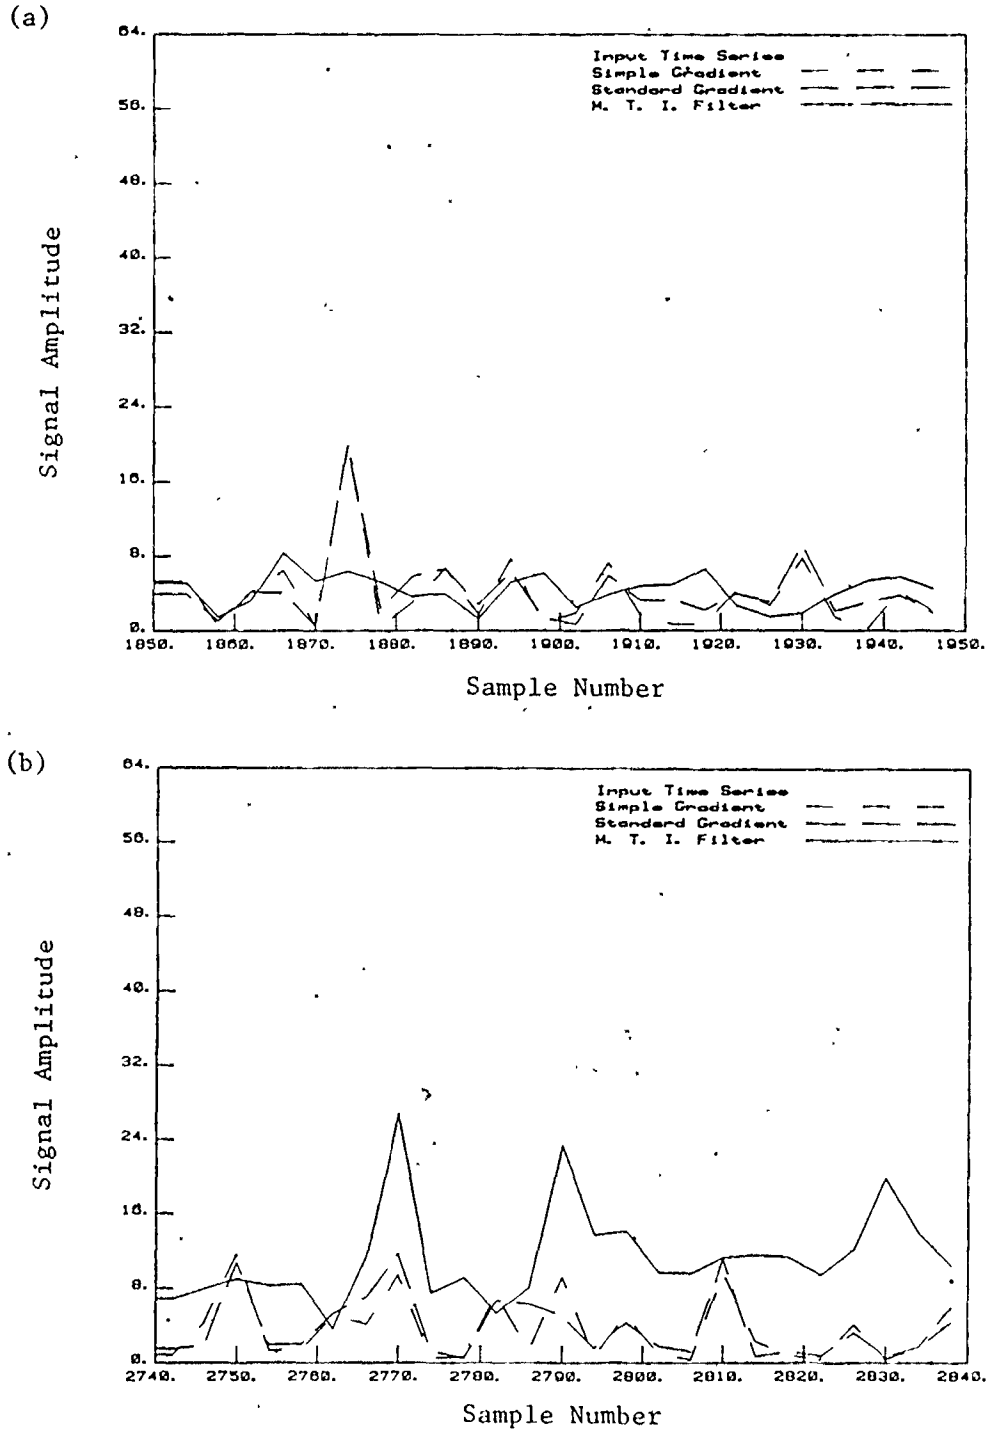

Figure 6.11: Comparisons of Filter Outputs for Rain Clutter

(a) Sector with Rain Velocity Tangential
(b) Sector with Rain Velocity Radial

(c)

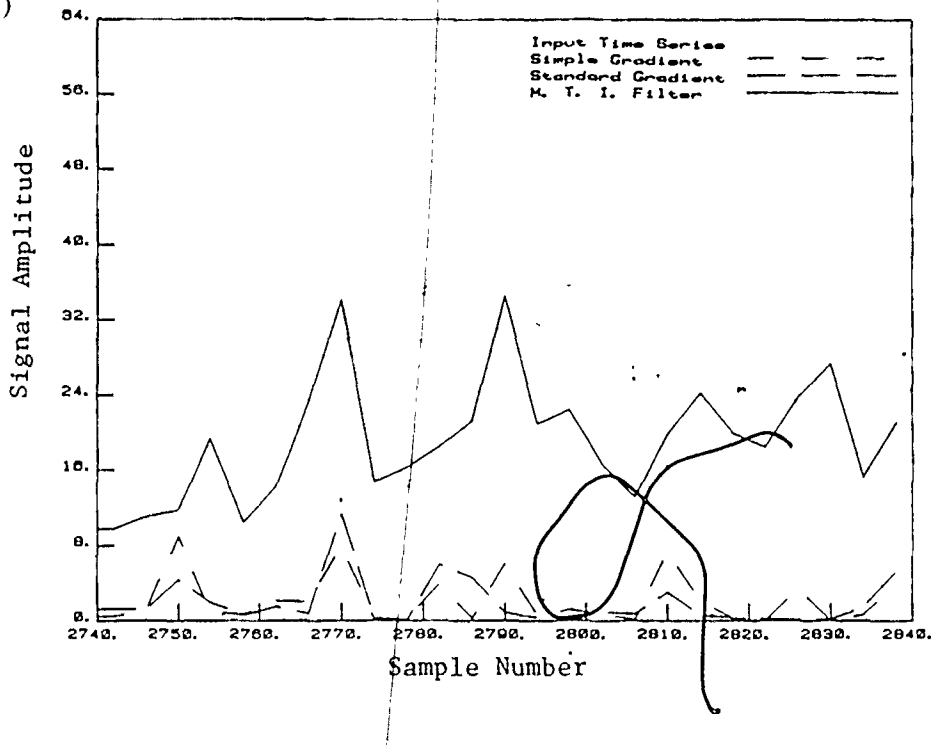

(d)

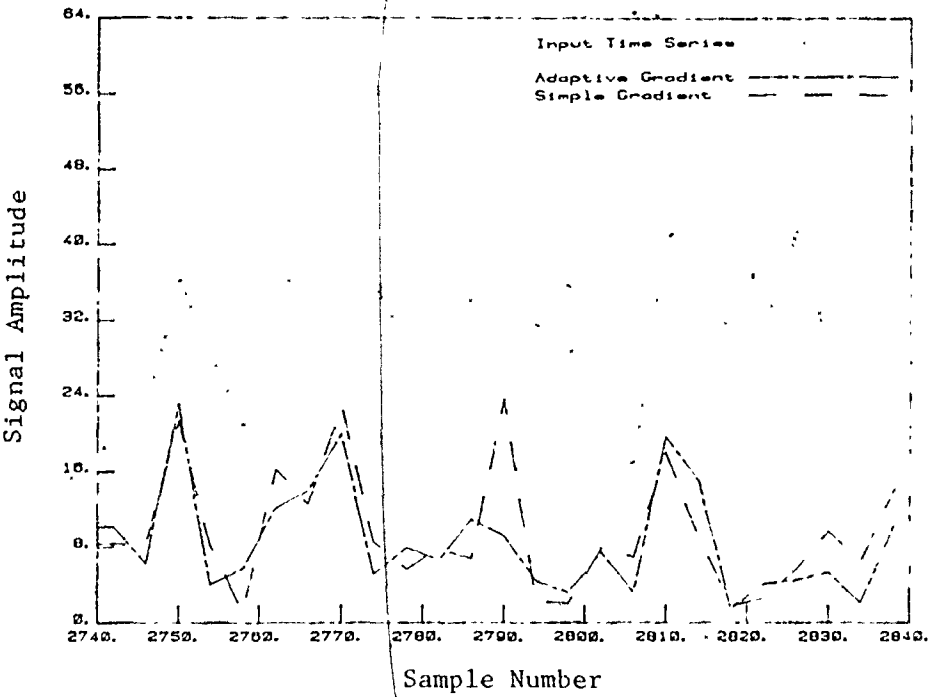

(Figure 6.11 cont'd.)

(c) Same as in (b) for Different Filter Orders (m = 2 for MTI, m = 9 for Lattice Filters)

(d) Same as in (B) for Adaptive Gradient vs. Simple Gradient Method (ω = 0.5)

order is five for all filters. In plot (c), the MTL filter order has been decreased to two (identical to that of the radar installation), with a marked worsening in performance. The lattice filters in plot (c) had a filter order of nine (the highest tested), and show a slight improvement in performance. For all types of clutter, the lattice filters showed this same gradual improvement with filter order, with performance remaining fairly good for filter orders down to about two. The MTI filter, on the other hand, showed a significant improvement with higher filter orders for rain clutter only. For other clutter types, there was little significant difference between filter orders of two and five. (More details are presented in Chapter 7). Finally, plot (d) of Fig. 6.11 compares the performance of the simple gradient (with $\omega = 0.5$) and adaptive gradient implementations of the lattice filter, for the same sector. As noted in section 6.2.1, the performance of the two is very similar.

Periodograms of the filter outputs for two sectors of the rain-clutter range-ring are presented in Fig. 6.12. The sector in the upper plot has a low average Doppler frequency, while the other is much higher. As can be seen, the lattice filters have basically the same performance for both, reducing the higher peaks in the spectrum by 20 to 30 dB, thus whitening the spectrum. The MTI filter, on the other hand, does as well as the lattice filters for the low Doppler frequency case (better near zero), but for the higher Doppler frequencies does very poorly, as can be seen in the lower plot.

Of great interest is the performance of these filters when a target is present in addition to the clutter. This can be tested by adding a known

(a)

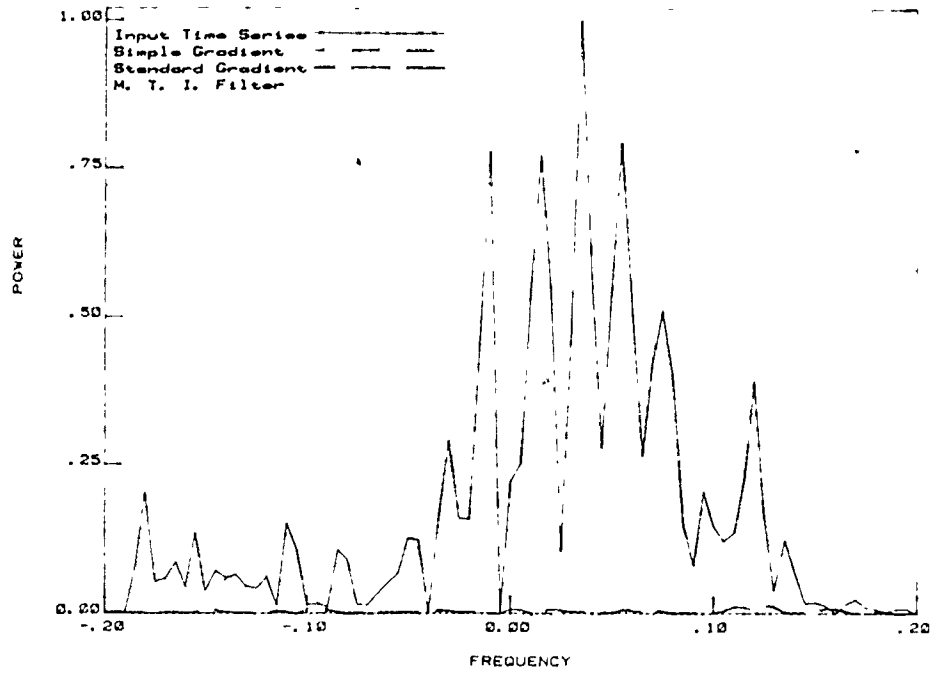

(b)

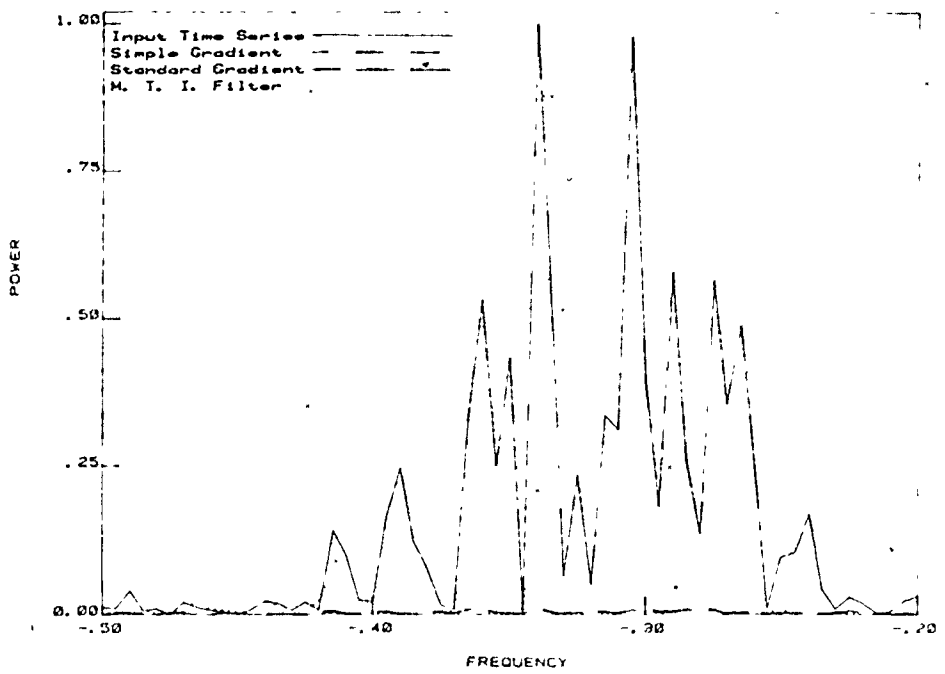

Figure 6.12: Periodograms for Rain Clutter (100 Samples)

(a) Sector with Rain Velocity Radial
(b) Sector with Rain Velocity Tangential

target to the clutter record, and then comparing the resultant filter outputs. An example of this is presented in Fig. 6.13. (The target, #5 in Table 6.1, is added to rain clutter of the same sector as in Fig. 6.11, b-d). In this example, the target is badly obscured by the clutter, and the MTI filter fails to sort the mixture out. The lattice filters, however, provide a reasonably good detection.

6.3.2 Other Precipitation Clutter

Rain was not the only precipitation clutter seen during the time at Bagotville. Periods of snow and of ice pellets (freezing rain) were also recorded, and an example of each is given here. The example of snow clutter comes from a moderate snow storm, during which there were light surface winds from the south-east. Figure 6.14 is a picture of the MTI display, taken during this storm. From this picture, it can be surmised that the winds at higher altitudes were probably of higher velocity and from a more southerly direction (common weather conditions), resulting in the clutter "breaking through" the MTI filter. The example of ice-pellet clutter comes from a period of light ice-pellet precipitation, with light winds from the north-west. In both these cases, as with the rain clutter, significant portions of the radar PPI screen were obscured by the clutter, creating the problems discussed in the previous section.

Comparisons of the MTI and lattice-structure filter outputs are presented in Figures 6.15 and 6.16 for sectors of ice-pellet clutter and snow clutter, respectively, taken from the storms described above. In both figures, the upper plots come from sectors of clutter having a large

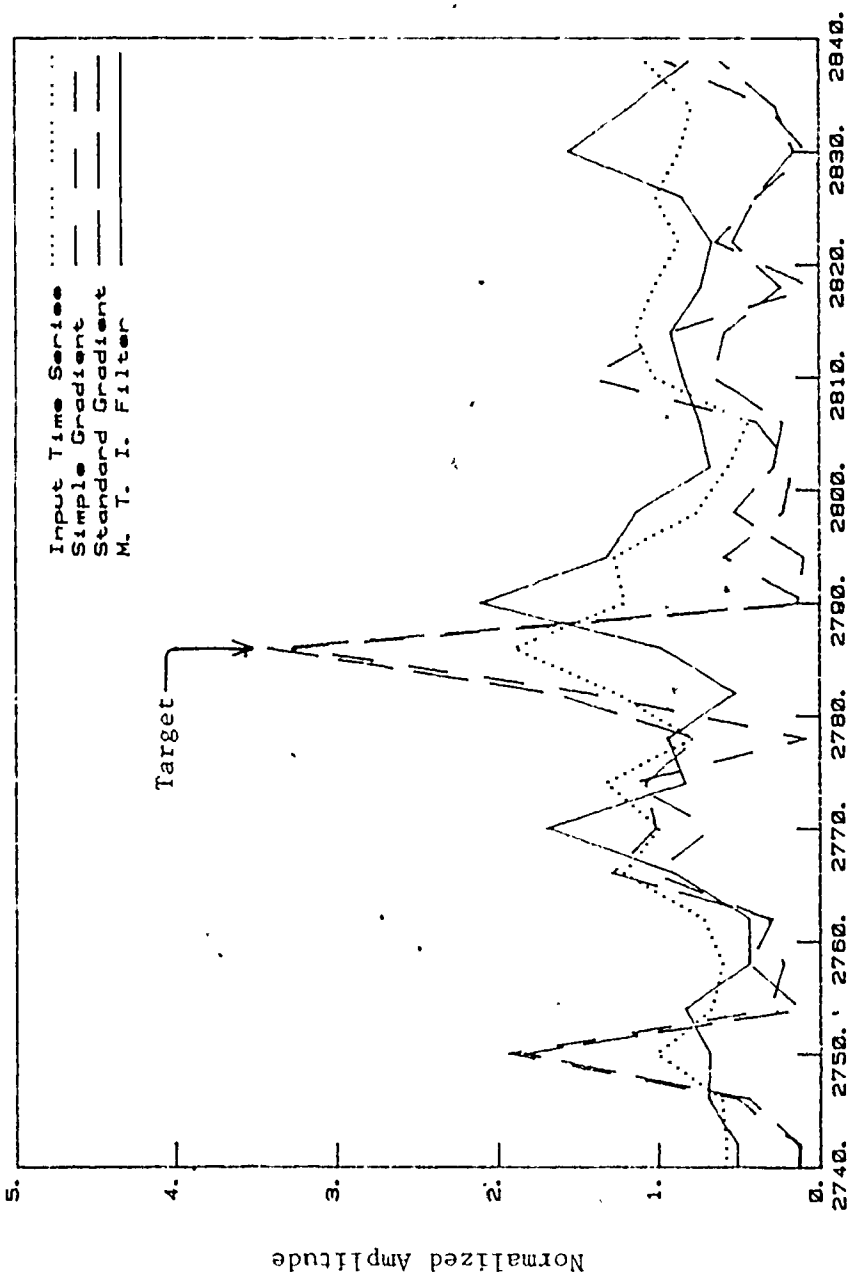

Figure 6.13: Comparison of Normalized Filter Outputs for Rain Clutter with Target Added (#5 in Table 6.1)

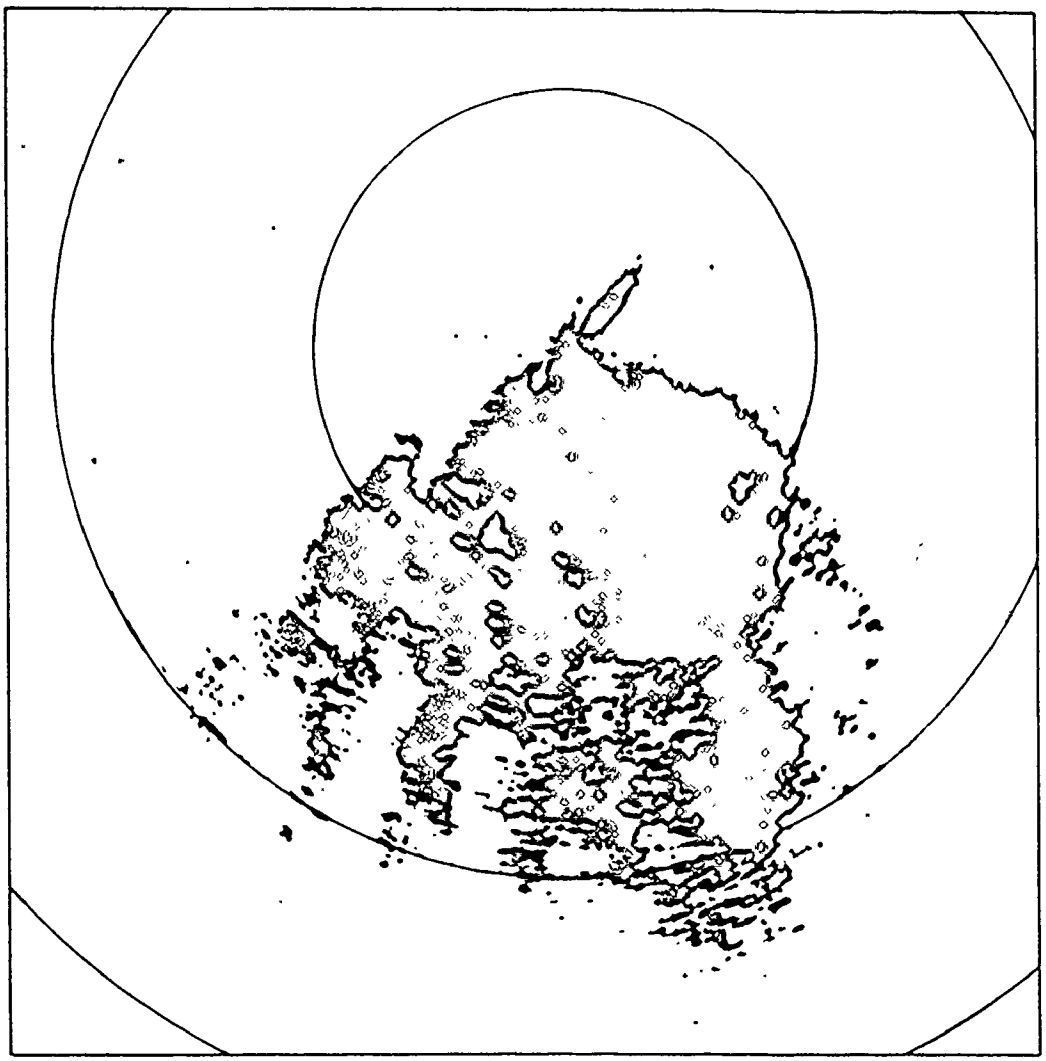

Figure 6.14: MTI Display of Snow Clutter on Operating PPI
(5 nautical mile range rings)

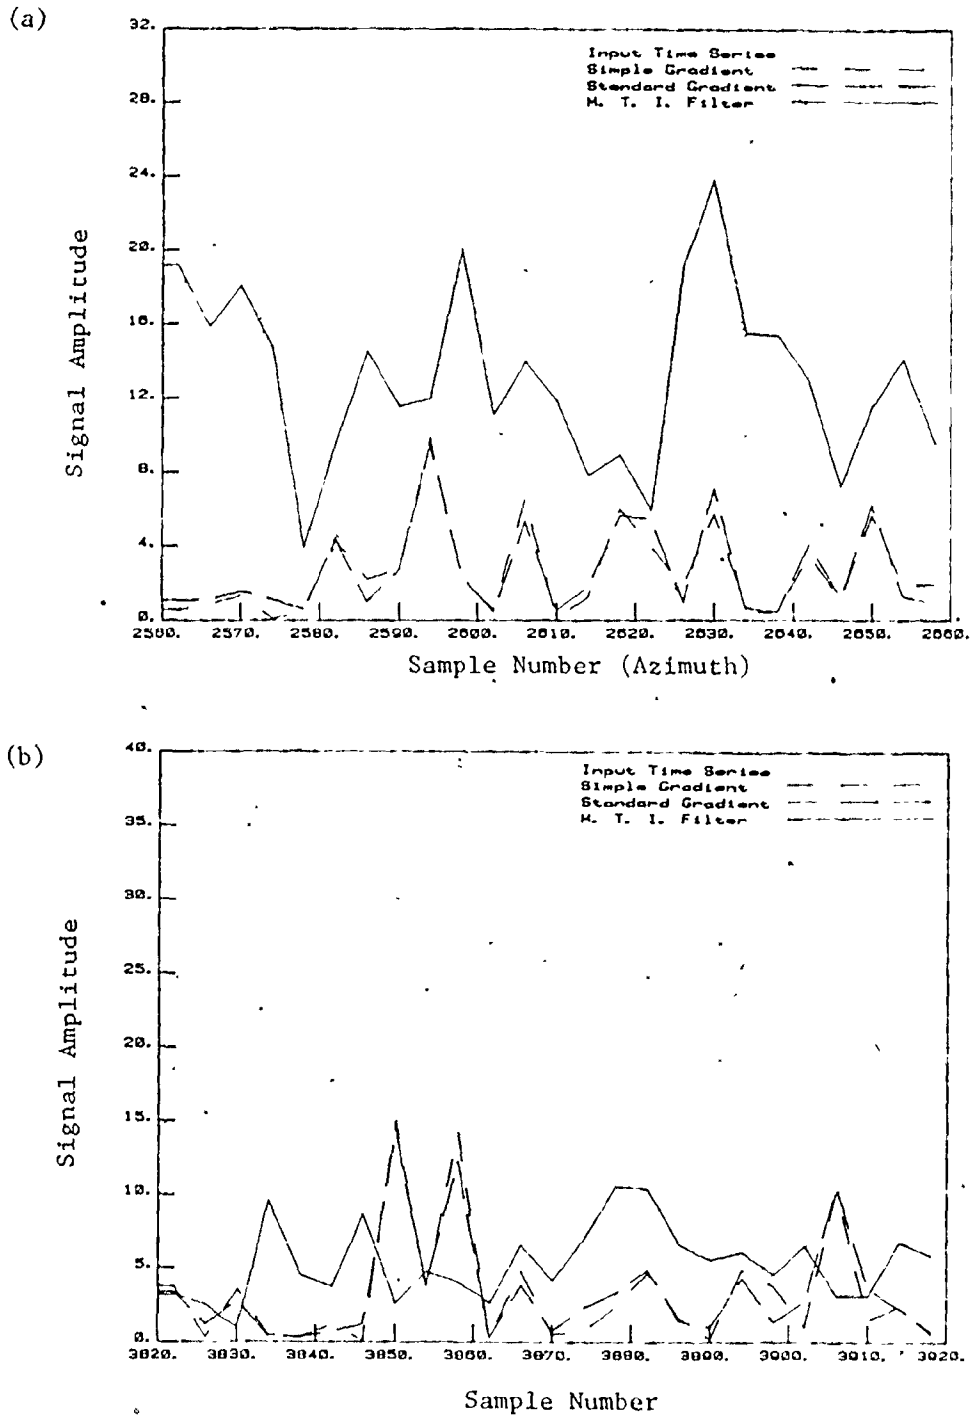

Figure 6.15: Comparisons of Filter Outputs for Ice-Pellet Clutter (a) Moving Radially, (b) Moving Tangentially

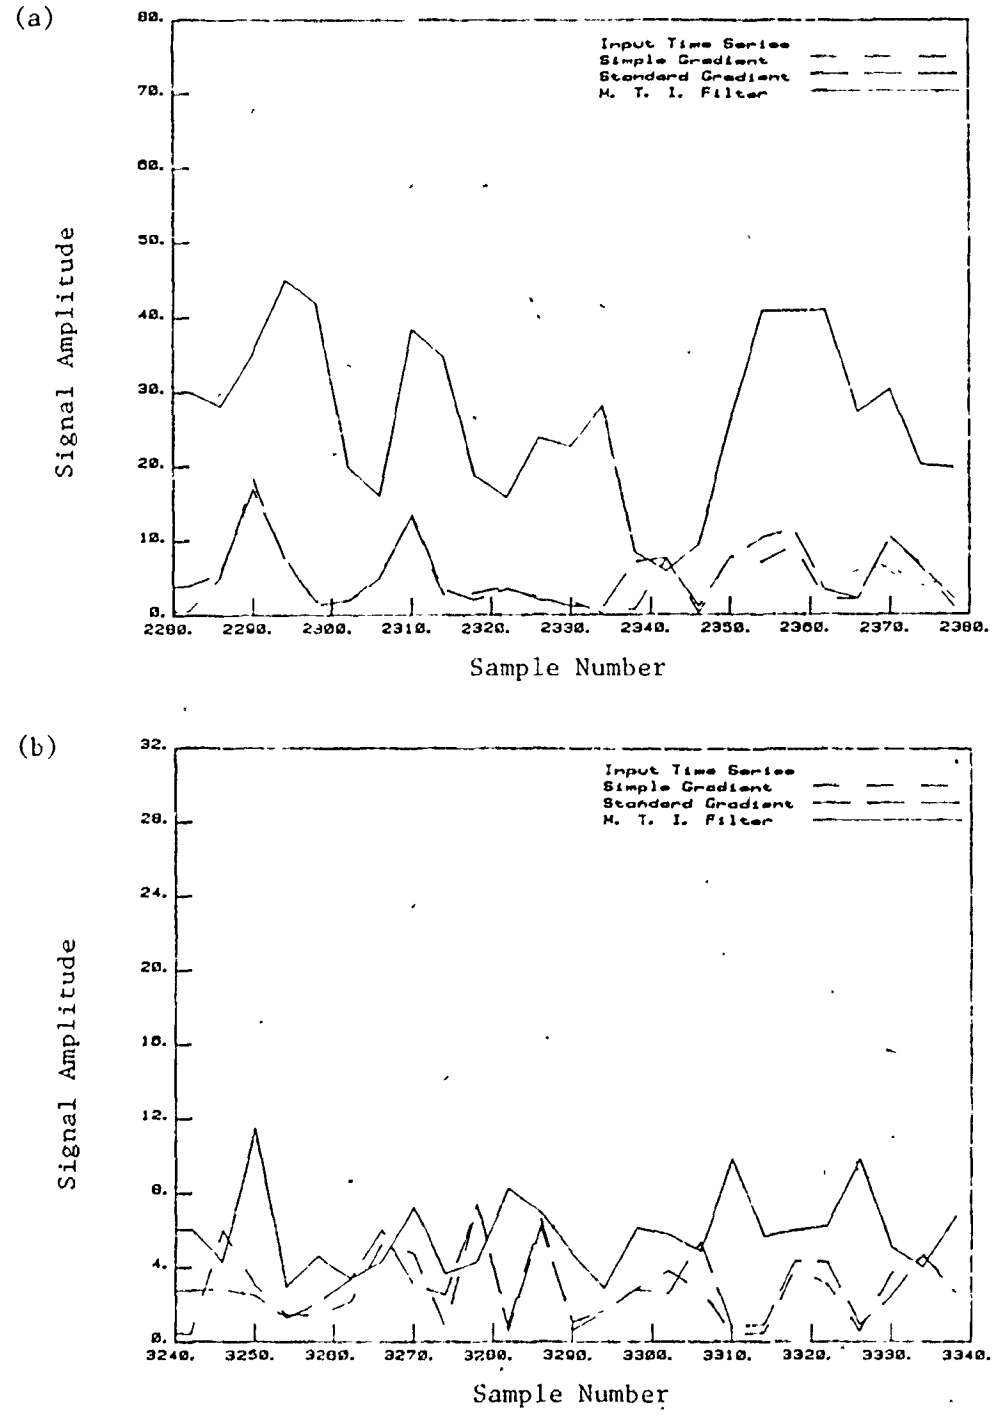

Figure 6.16: Comparisons of Filter Outputs for Snow Clutter
(a) Moving Radially, (b) Moving Tangentially

radial velocity, while the lower plots come from sectors where the movement is basically tangential. The ice-pellet clutter data was from a ring sampled at a range of 10 nautical miles, while the snow clutter was sampled at a range of 6½ nautical miles. As can be seen, the response in these figures is similar to that for rain clutter (as shown in Fig. 6.11). The lattice-structure filters performed consistently well, while the MTI filter's performance ranged from almost as good as the lattice filters' performance (for low radial velocity), to a very poor response.

Figure 6.17 plots the periodograms of the filter outputs (as shown in Fig. 6.15) for ice-pellet clutter. Again, the results are similar to the peridograms for rain clutter (in Fig. 6.12), except that the MTI response for radially moving ice-pellet clutter is much worse. This is because the Doppler velocity of this clutter happens to be centered at exactly the frequency where the MTI filter has its maximum gain, (i.e., 0.5, normalized to the sampling frequency. (The lower plot also includes a peak resulting from ground clutter). Figure 6.18 plots the periodograms of the filter outputs (as shown in Fig. 6.16(a)) for a sector of snow clutter. The slight increase in randomness of this plot, compared to previous plots, may have resulted from the lower density of the snow-flakes, allowing greater wind-related turbulence. The response curve of the MTI filter can be viewed in this plot, having unity gain at a frequency of 0.5, and dropping to zero gain at zero frequency. As with rain clutter, the lattice filters provided between 20 and 30 dB of attenuation of the periodogram peaks for both ice-pellet

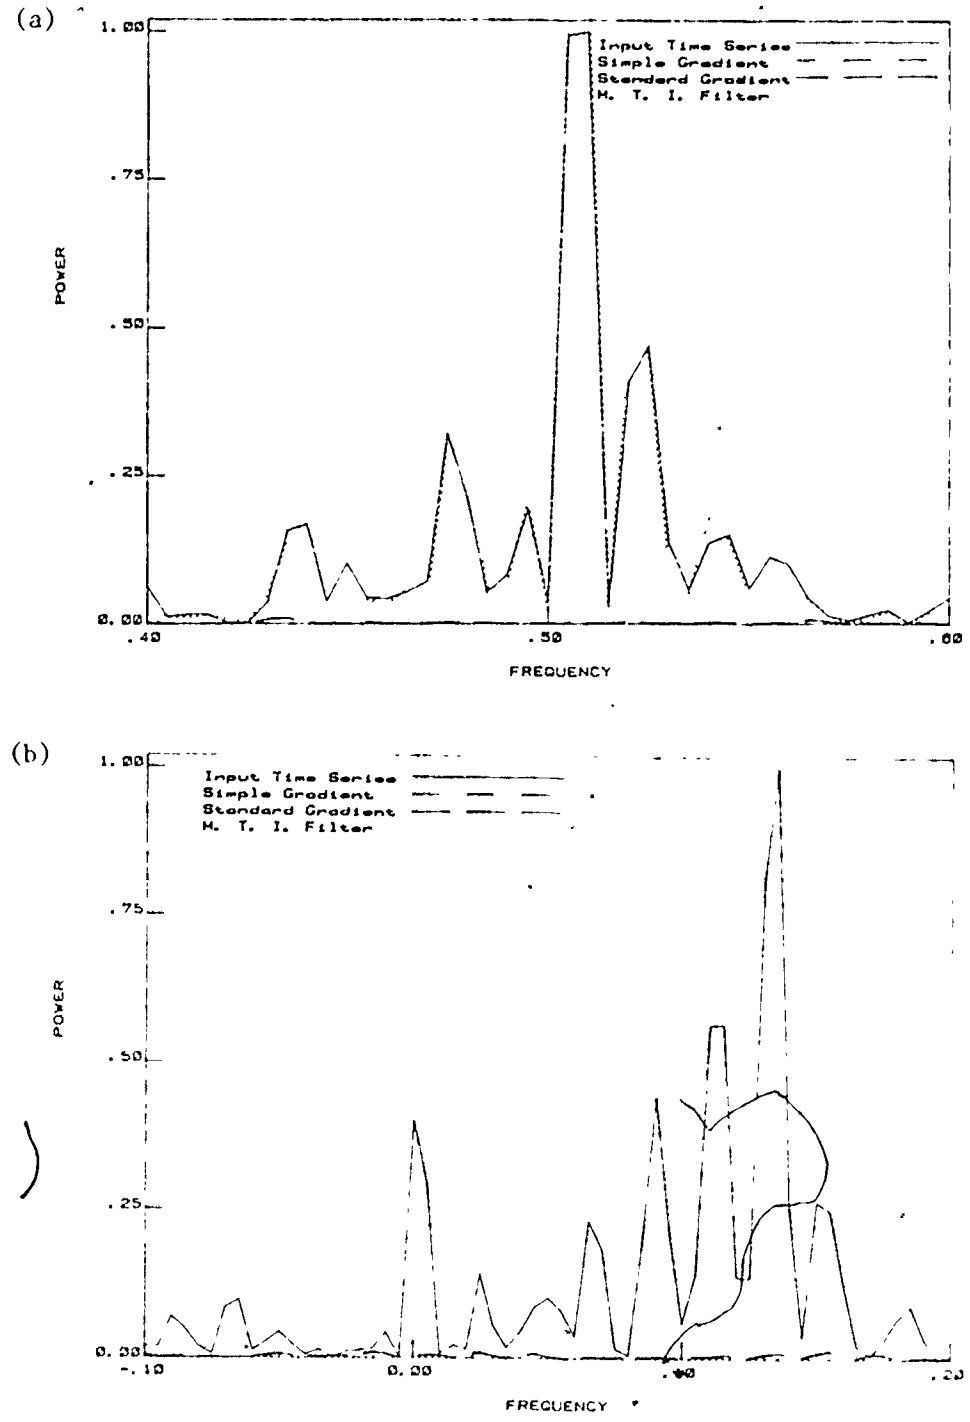

Figure 6.17: Periodograms for Ice-Pellet Clutter
(a) Moving Radially, (b) Moving Tangentially

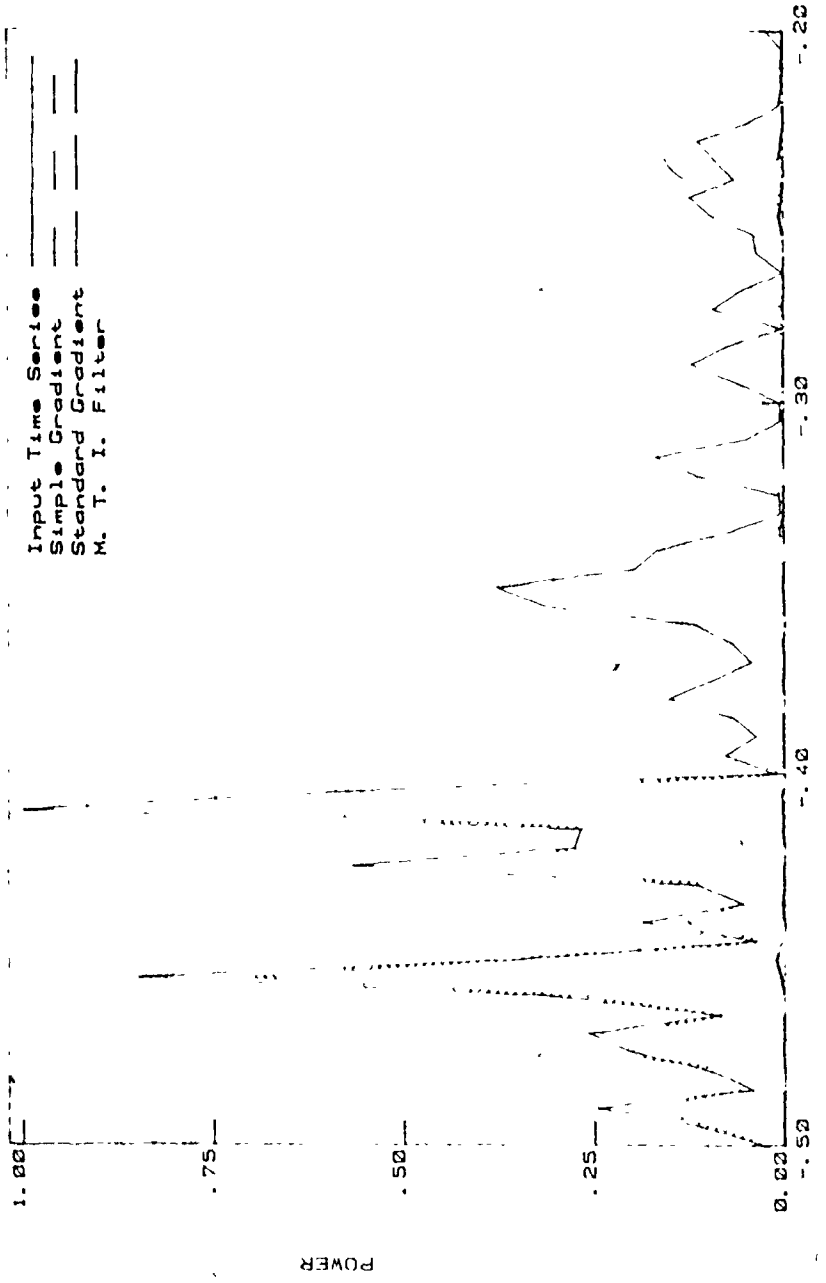

Figure 6.18: Periodograms for Snow Clutter

and snow clutter, significantly whitening the clutter spectrums.

Finally, the responses of the filters were compared, with known target signals added to the clutter input. The example shown in Fig. 6.19 is for the addition of an aircraft target (#1 in Table 6.1) to the ice-pellet clutter from Fig. 6.15(a). As can be seen, the MTI filter fails to detect this target, while the lattice-structure filters do so quite well. In summary, then, the effects of snow and ice-pellet clutter on the performance of lattice-structure and MTI filters are very similar to the effects of rain clutter, as presented in section 6.3.1.

6.3.3 Ground and Inversion Clutter

This section deals with two types of clutter not caused by precipitation, namely ground clutter and inversion clutter (thought to be ground echoes refracted or reflected by layers of denser air near the ground). This may seem to be a strange grouping, but as will be seen, the characteristics of these two types of clutter are quite similar.

Ground clutter is always present in a radar environment. A typical set of operating displays with only ground clutter (and a target) present were shown in Fig. 6.1. The lower display in the figure showed how the MTI filter almost completely eliminated the ground clutter, as it was designed to do. This figure may be compared with Fig. 6.20, containing the corresponding displays with inversion clutter present. This data was recorded on a clear, cold evening with moderate (13 knot) winds from the south-east. The MTI display shows two "blobs" of scattered points, aligned with the wind. (In other displays of this type of clutter, the points have short tails, all aligned radially to the display center).

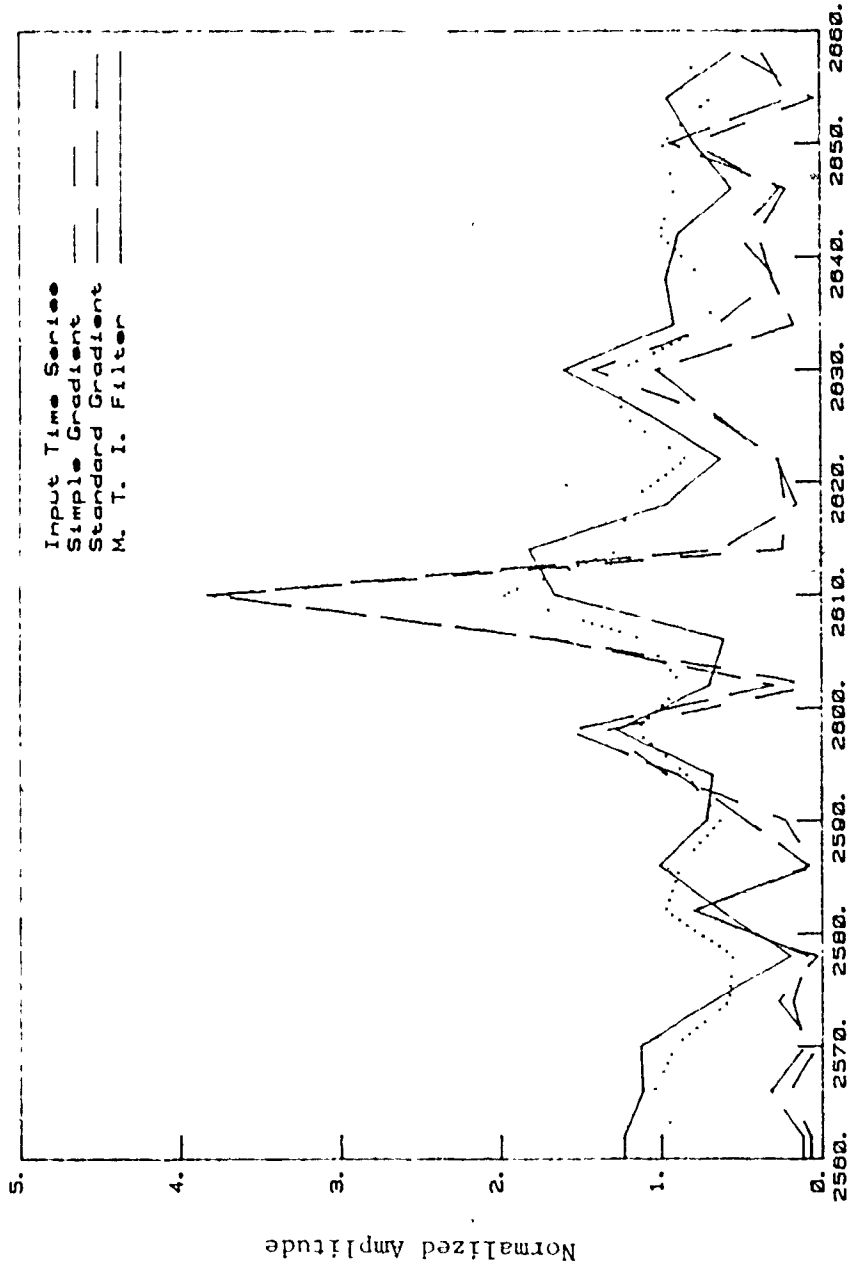

Figure 6.19: Comparison of Normalized Filter Outputs for Ice-Pellet Clutter with Target Added (#1 in Table 6.1)

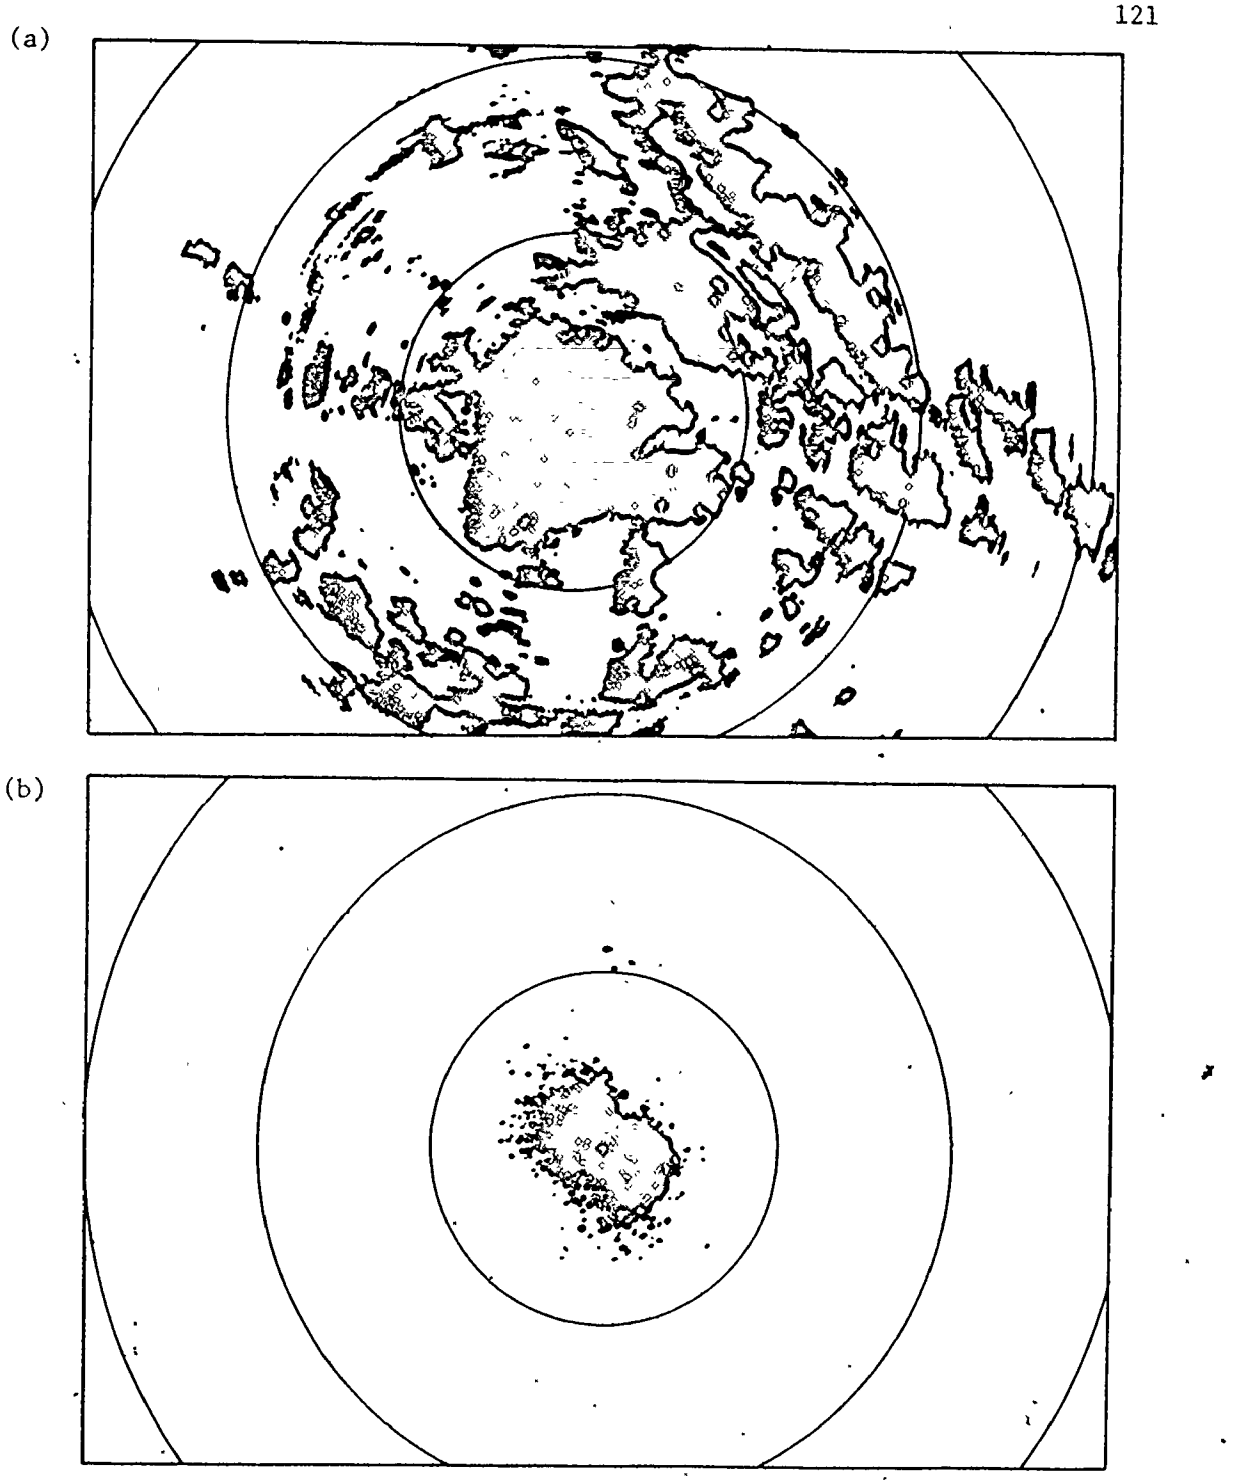

Figure 6.20: Inversion Clutter PPI Displays (10 nautical mile range rings)
(a) "Normal" Video, (b) "MTI" Video

Figure 6.21 shows the data for a range ring sampled from this clutter. This may be compared with Fig. 6.2 for ground clutter. (Note: Figure 6.2 is for a range ring at 12 nautical miles. The ground clutter range ring analysed in this section was sampled at 6 nautical miles, but is similar in appearance.) The high, random level of the inversion clutter is apparent.

The identification of atmospheric inversion as the source of clutter attributed to it here is not certain. This sort of clutter is often called "angel clutter", and some reports suggest that it is caused mainly by birds [42]. This seems unlikely for three reasons. They are: (a) No birds were seen in the area at the times when the clutter was present; (b) The clutter movement was aligned with the wind direction (and therefore presumably related); and, (c) The clutter always occurred during the same weather conditions. These were late afternoons or evenings with clear skies when the temperature was dropping rapidly. These are conditions which often lead to the formation of an inversion layer near to the ground [48]. This comes about when the ground cools faster than the air mass above it, resulting in a fairly stable layer of denser air. Movement of the layer or its boundaries could impart a Doppler shift to signals refracted or reflected by it, resulting in MTI detection. However, whether or not this is the true cause of this type of clutter is not very important to this study.

Figure 6.22 compares the MTI and lattice-structure filter outputs for sectors of the ground clutter and inversion clutter range-rings, respectively. Both filter types do very well for both types of clutter, although the response to inversion clutter is slightly noisier than the

(a)

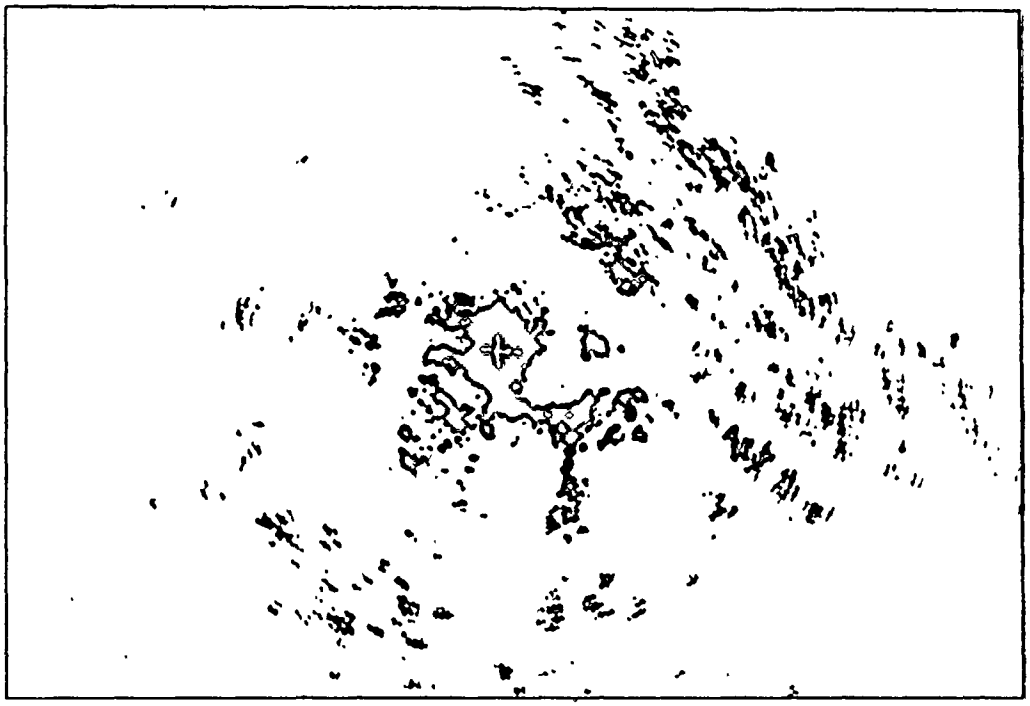

(b)

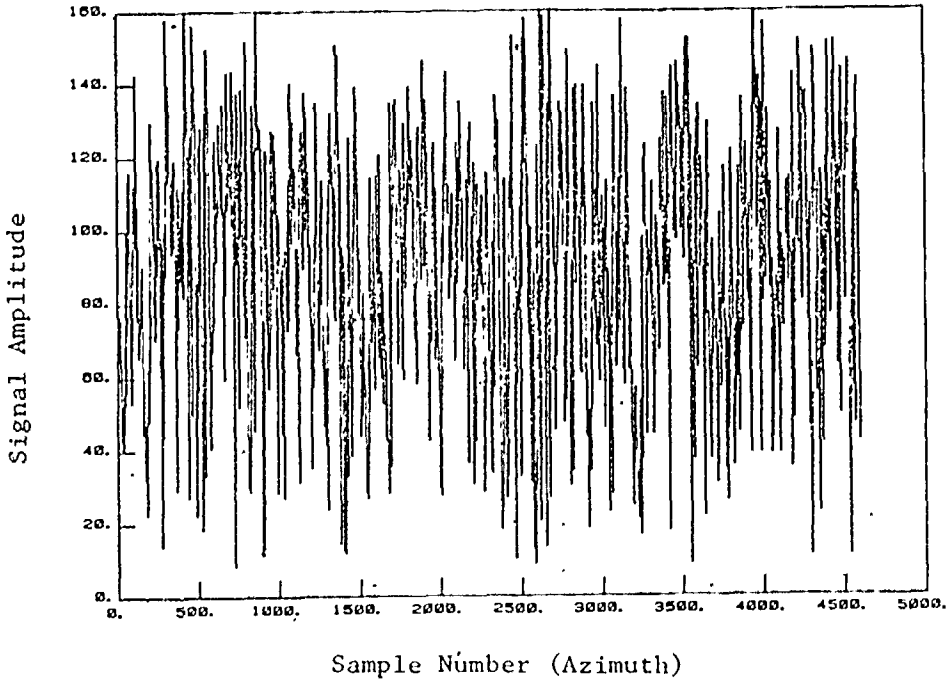

Figure 6.21: Data for Inversion Clutter (range ring at 1.5 nautical miles)
(a) Laboratory PPI Display, (b) Amplitude Plot of above Range Ring

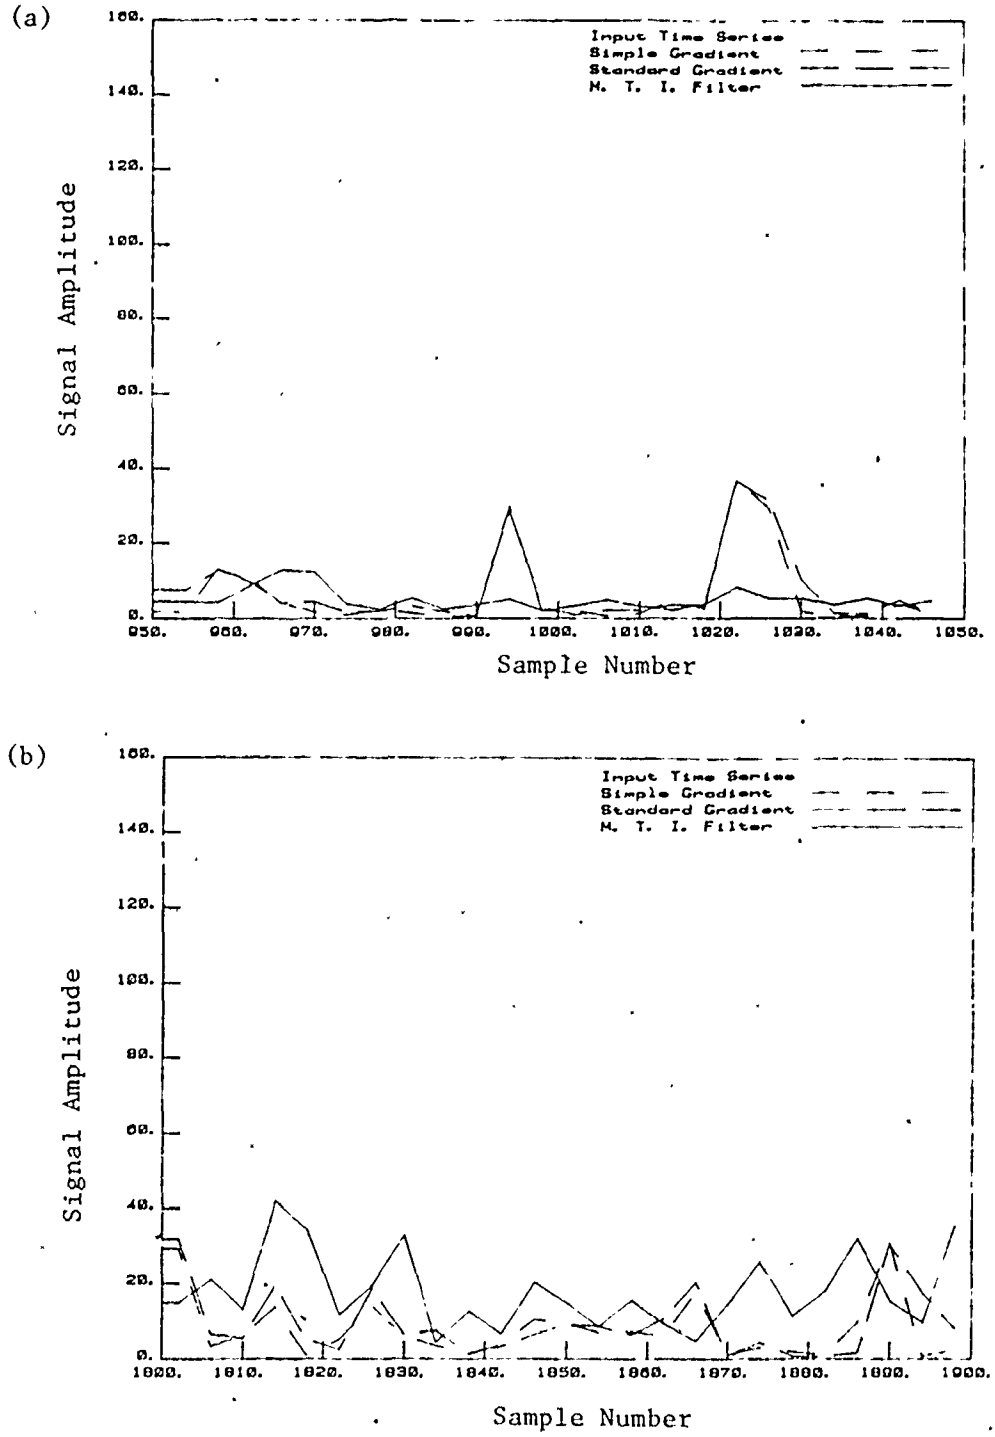

Figure 6.22: Comparisons of Filter Outputs for Ground and Inversion Clutter

(a) Ground Clutter, (b) Inversion Clutter

response to ground clutter. The small peaks in the lattice filter outputs for ground clutter are due to sudden changes in the clutter. The MTI response is slightly worse than that of the lattice structures for inversion clutter, but not significantly so.

Periodograms of the responses for these same sectors of ground and inversion clutter are plotted in Fig. 6.23. Although the filter output responses were plotted in this figure, they are not readily visible, as they are too close to the zero axis. In comparing the two plots, it is notable that: (a) the central peak for inversion clutter is offset from zero slightly, and (b) the inversion plot contains a number of lower-level, spurious peaks away from the central peak. Larger spurious peaks probably account for the points detected by the MTI filter in the clutter area (such as in Fig. 6.20). In general, the lattice filters provided about 30 dB of attenuation of the higher peaks for both ground and inversion clutter.

Finally, Fig. 6.24 compares the filter outputs when a known target signal (#1 in Table 6.1) is added to the same ground clutter record. It can be seen that both types of filter do a good job of detecting targets obscured by this type of clutter. How good a job they do will be the subject of the next chapter.

(a)

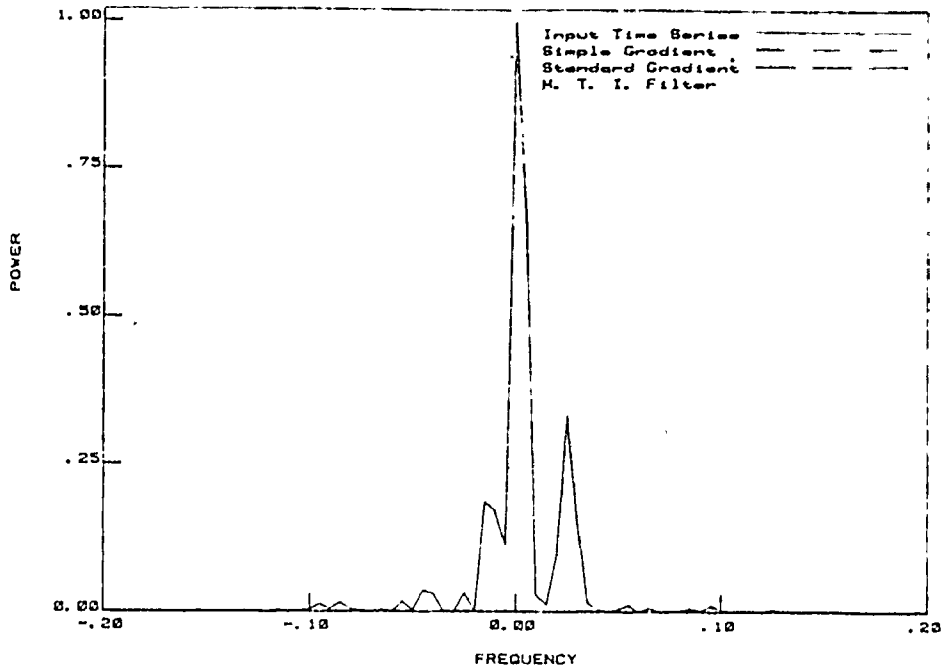

(b)

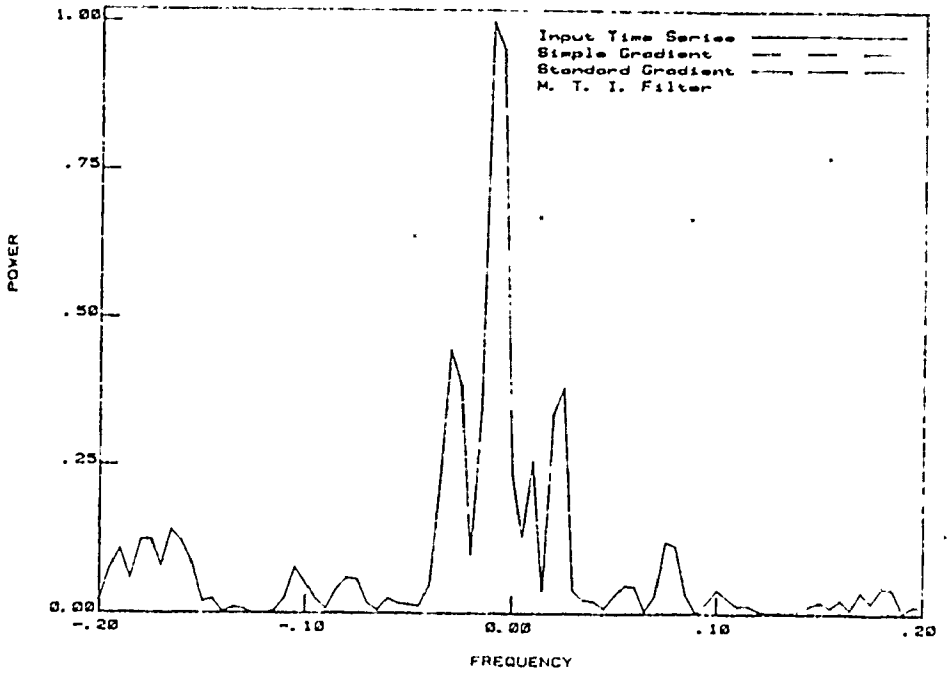

Figure 6.23: Periodograms for Ground and Inversion Clutter
(a) Ground Clutter, (b) Inversion Clutter

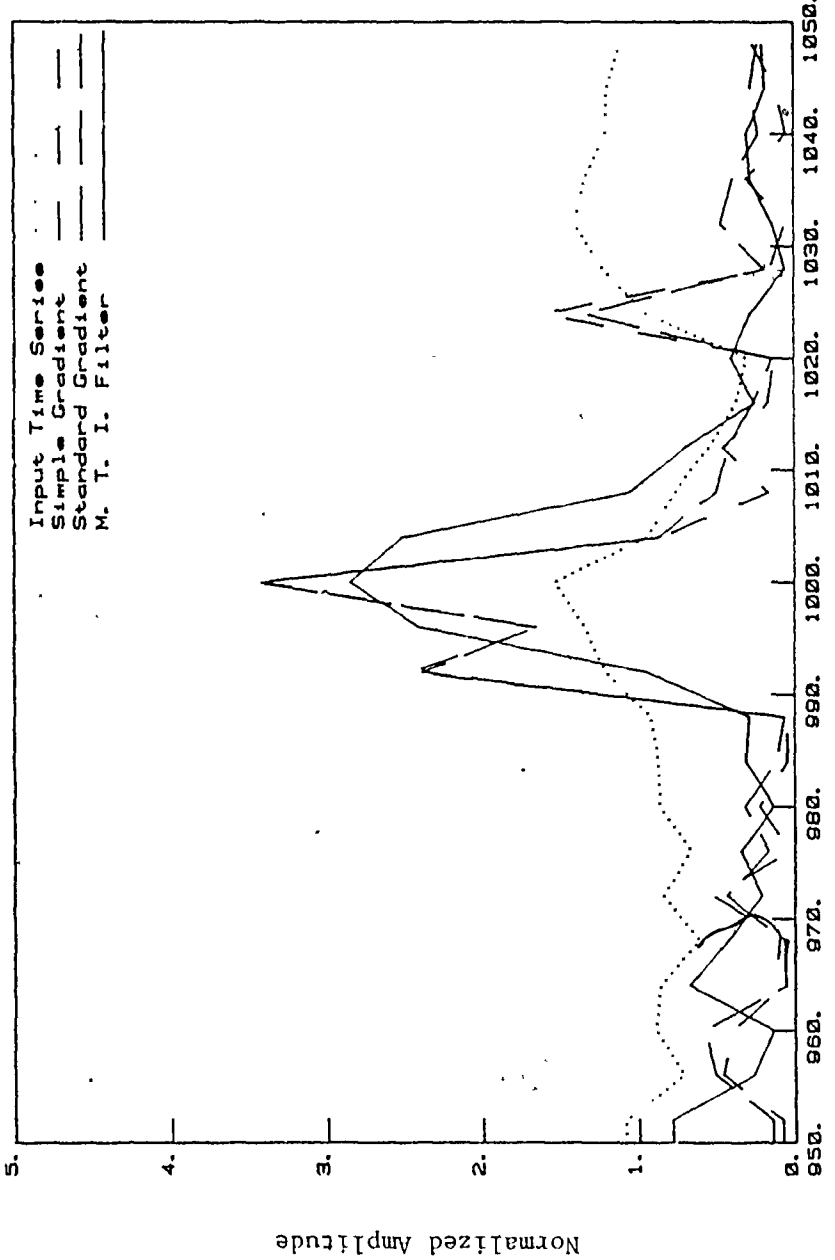

Figure 6.24: Comparison of Normalized Filter Outputs for Ground Clutter with Target Added (#1 in Table 6.1)

CHAPTER 7

PERFORMANCE MEASURES FOR ACTUAL RADAR DATA

7.1 Improvement Factors

The improvement factor of a radar clutter filter (as discussed in section 5.1) is defined as: "The signal-to-clutter ratio at the output of the system divided by the signal-to-clutter ratio at the input of the system, averaged uniformly over all target radial velocities of interest" [36].

This can be written as

$$I = \frac{(T_{out}/C_{out})}{(T_{in}/C_{in})}, \qquad (7.1)$$

where T and C are the target and clutter powers. In applying this performance measure to the actual radar data collected as part of this thesis, the improvement factors (I) were calculated from the ratios of the target gains in the filter (T_{out}/T_{in}) to the clutter gains in the same filter (C_{out}/C_{in}).

The target amplitudes used in computing these gains were measured by first taking the amplitudes of individual data samples, and averaging over four samples. Then, the largest of the averaged values in the target area was selected, for both the input and the output of the filter. This averaging is short enough to ensure that the antenna pattern does not become a factor, but long enough to reduce the effect of noise and elim-

inate interference from the staggered pulse-repetition-frequency of the radar. (By taking the largest averaged value, the target is measured at the peak of the antenna pattern).

The above process was applied to the five aircraft signals described in section 6.2, and the five resulting gains were averaged together to form a single target gain for each filtering set-up. This corresponds to the "... averaged uniformly over all target velocities of interest" part of the definition quoted above. Five targets may seem to be a small number to form a "uniform" average, but these targets do represent a fair cross-section of the possible velocities (including zero). Besides, these were the only such signals readily available, and thus became the targets "of interest".

The clutter gains were measured by averaging the signal amplitudes over the entire range ring, and taking the ratio of the average output power to the average input power. These gains were measured for more than one range-ring of each type of clutter, with these rings generally recorded at different ranges and times, to ensure uniformity of the results.

The improvement factors plotted in this chapter were calculated using the average of clutter gains from several such range-rings of each clutter type; however, improvement factors were also calculated using the gains of the individual rings. In general, the measured improvement factors for the single clutter rings and for the average for a clutter type were within 1/2 dB, except where noted. The exceptions involved only the MTI filters and apparently related to the overall velocity of the specific clutter record. As the above variance was much less than that between different clutter types, this was taken as evidence that the selected range-

rings were representative of the clutter present.

The measured improvement factors, as a function of filter order, are plotted in Figures 7.1 to 7.5 for rain, snow, ice-pellet, ground, and inversion clutter, respectively. Each of these plots has five separate curves; one for the MTI filter, and two each for the simple gradient and standard gradient implementations of the lattice filter. The lattice-structure filters had adaptive constant values of 0.2 and 0.6. (These values were chosen to be representative of the useful range of these constants.

The set of curves for the lattice filters presents a very similar pattern in each of these five figures, with the only significant difference being in the overall level of the set of curves. Even this was fairly constant, with rain, snow, and ground clutter all resulting in approximately the same improvement factor levels (for each lattice filter). The improvement factors were about 3 dB higher for inversion clutter and 2 dB lower for ice-pellet clutter. A common factor in all the lattice filter curves is the increase in improvement factor with increasing filter order. The rate of this increase slows down at higher filter orders, suggesting that there is an upper limit to the improvement. (Lattice filter orders above nine were not tested due to a desire for simplicity in the final system, and because little additional improvement was expected). As expected, the standard gradient method of adaptively implementing the lattice filter performed better than the simple gradient method.

Another factor of interest in the lattice filter response is the effect of the adaptive constant. As can be seen in Figures 7.1 to 7.5, lower values of adaptive constant resulted in better performance. Figure

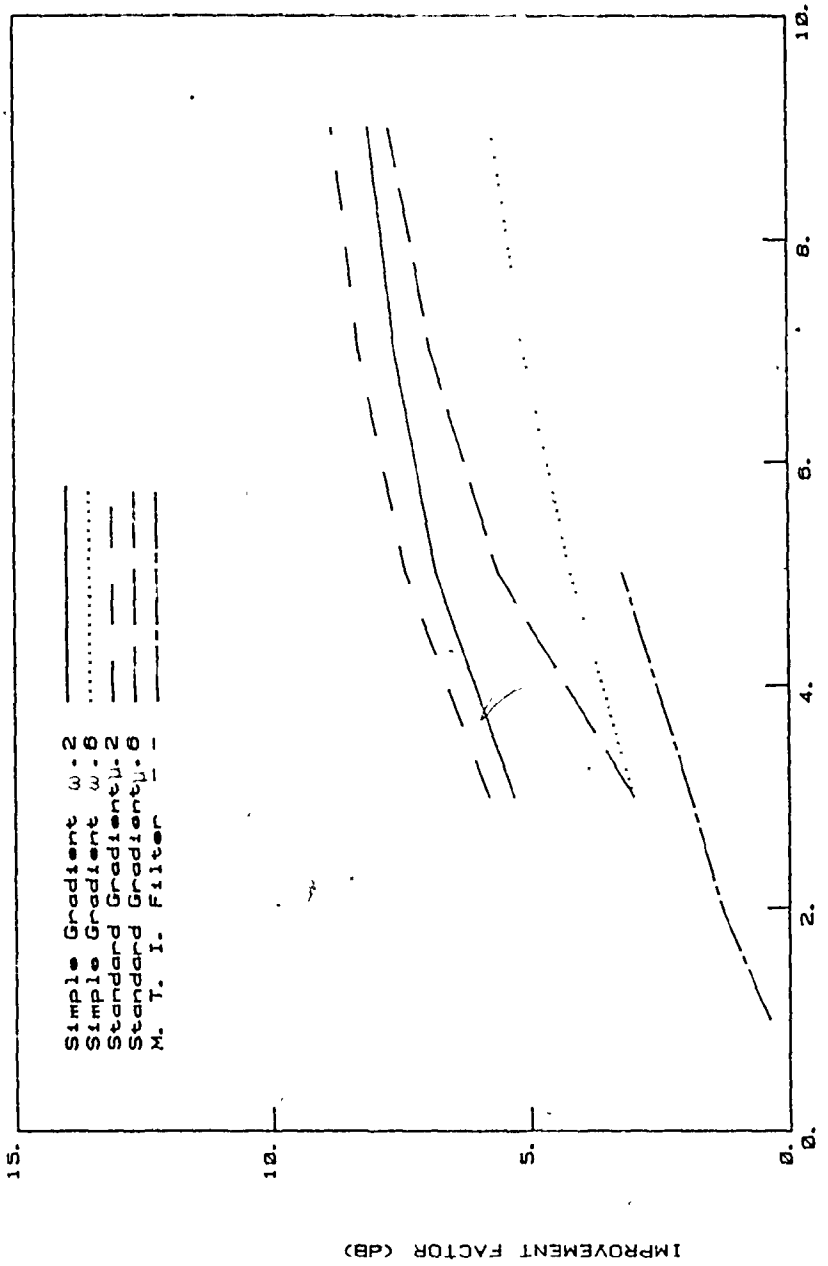

Figure 7.1: Improvement Factors for Rain Clutter

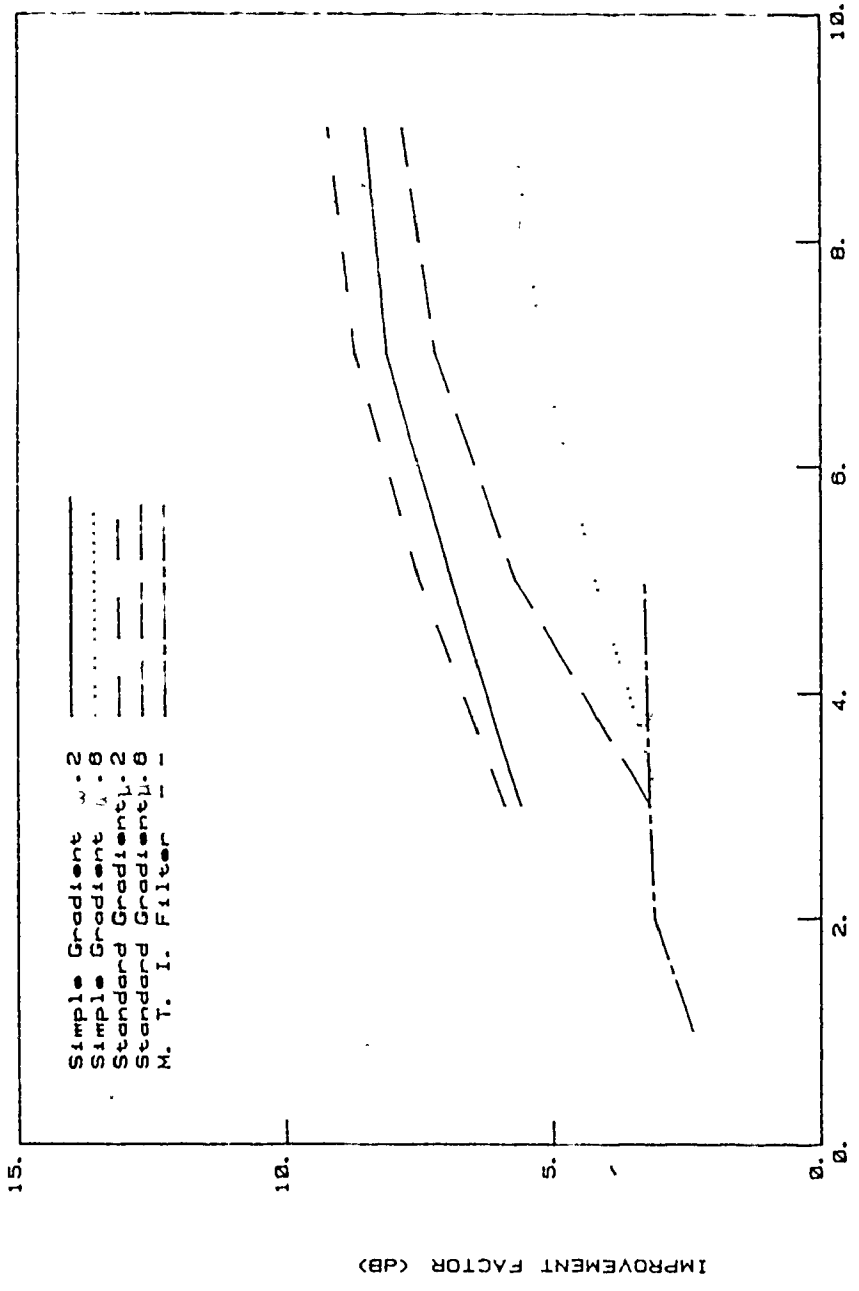

Figure 7.2: Improvement Factors for Snow Clutter

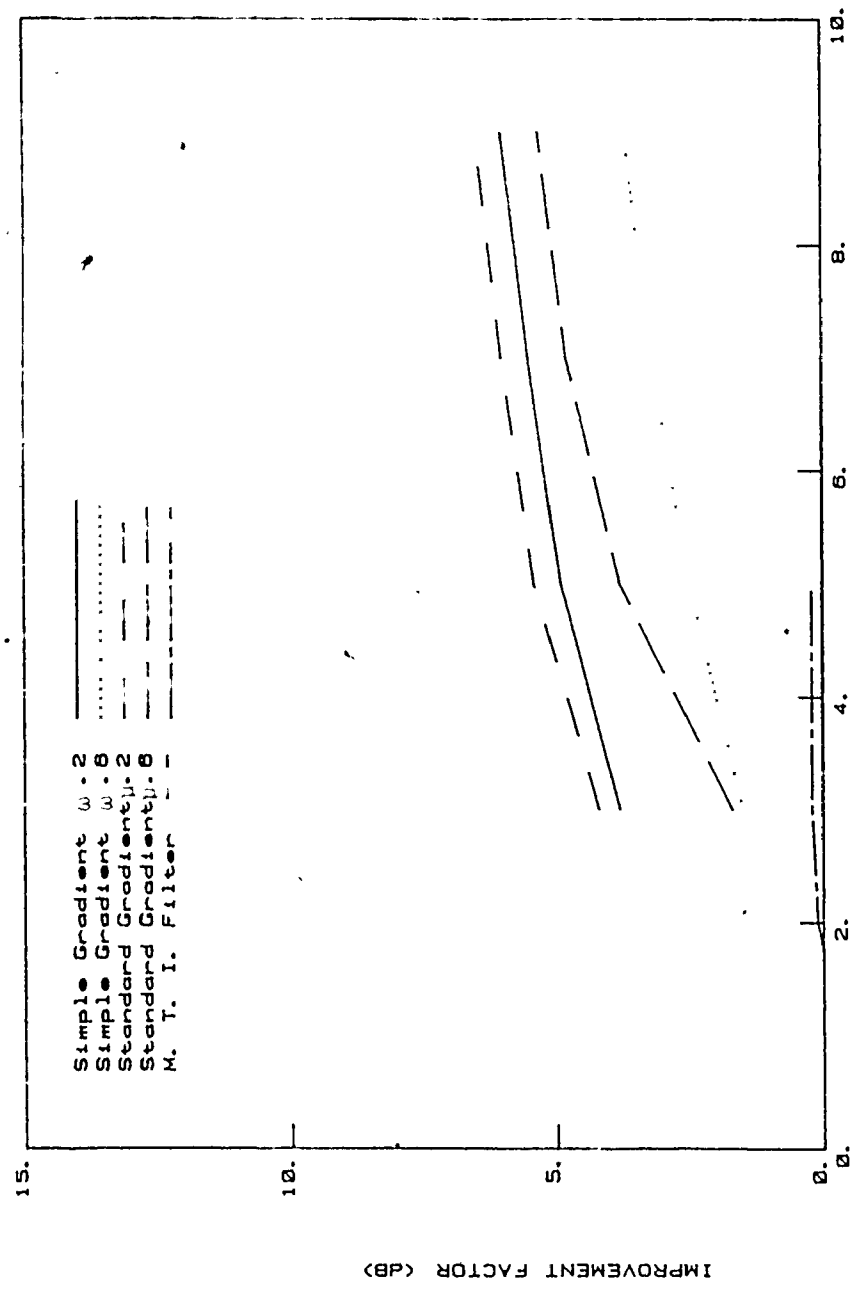

Figure 7.3: Improvement Factors for Ice-Pellet Clutter

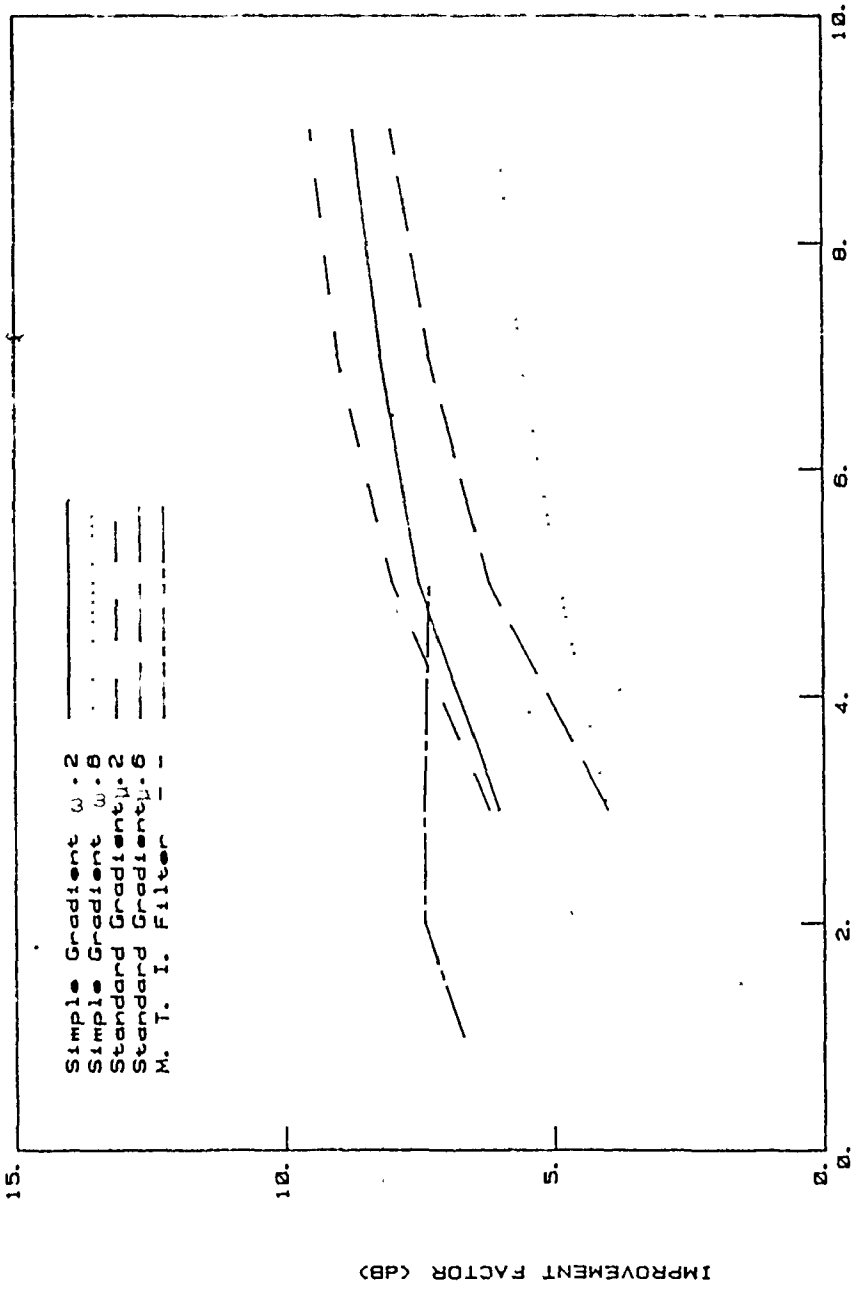

Figure 7.4: Improvement Factors for Ground Clutter

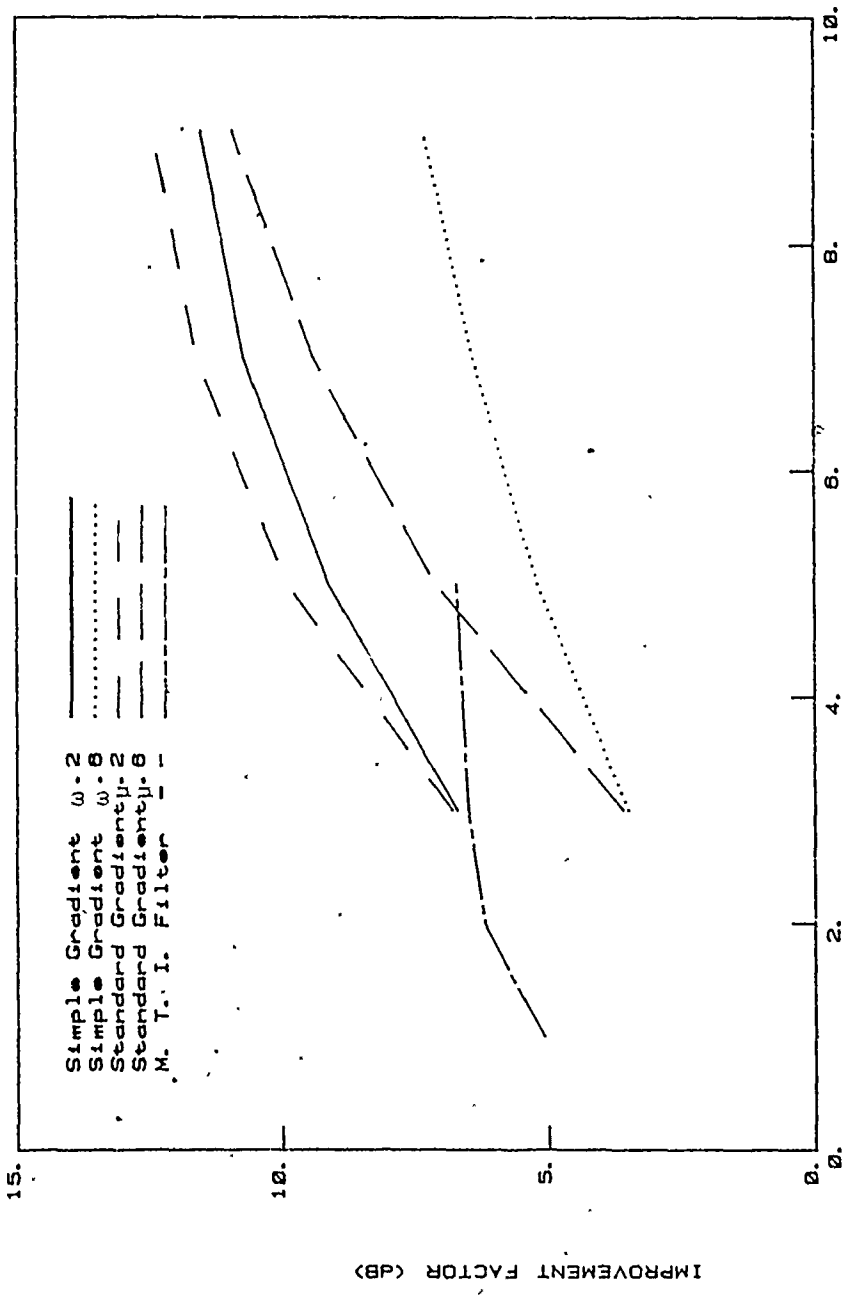

Figure 7.5: Improvement Factors for Inversion Clutter

7.6 plots improvement factors versus adaptive constants for both implementation methods (filter order = 9), and for three of the clutter types. (The other clutter types give very similar performance). The best values of adaptive constant for the standard gradient method occur in the range from 0 to 0.2, with an optimum value of $\mu = 0.1$. For the simple gradient method, on the other hand, the optimum performance occurs only at $\omega = 0$. (Note: At $\omega = 0 = \mu$, both methods simplify to a common set of equations).

The curves in Figures 7.1 to 7.5 for the MTI filter, unlike those for the lattice filters, show considerable variation with the type of clutter. The MTI filter performed best for ground clutter and worst for ice-pellet clutter. In addition, there was some variation in performance within specific clutter types, probably due to the velocity of the clutter. In particular, the single range-rings presented in section 6.3 for snow and inversion clutter both resulted in MTI performance about 2 dB below that shown in Figures 7.2 and 7.5 for an average of several rings. Some tests were also made using bird targets (from section 6.2.2) instead of the aircraft. This resulted in a further 3 dB reduction in MTI performance, with little effect on the lattice filters. The effect of MTI filter order also varied, depending on the type of clutter. For rain clutter, the improvement factor increased with filter order, up to $m = 5$. For the other clutter types, however, there was little increase beyond $m = 2$. (In fact, for some range-rings of clutter, there was a decrease in level with increasing filter order). This tends to validate the common choice of $m = 2$ for conventional radar MTI filters.

The relative performance of the lattice and MTI filters can be assessed from Figures 7.1 to 7.5. (In this instance, we will compare the

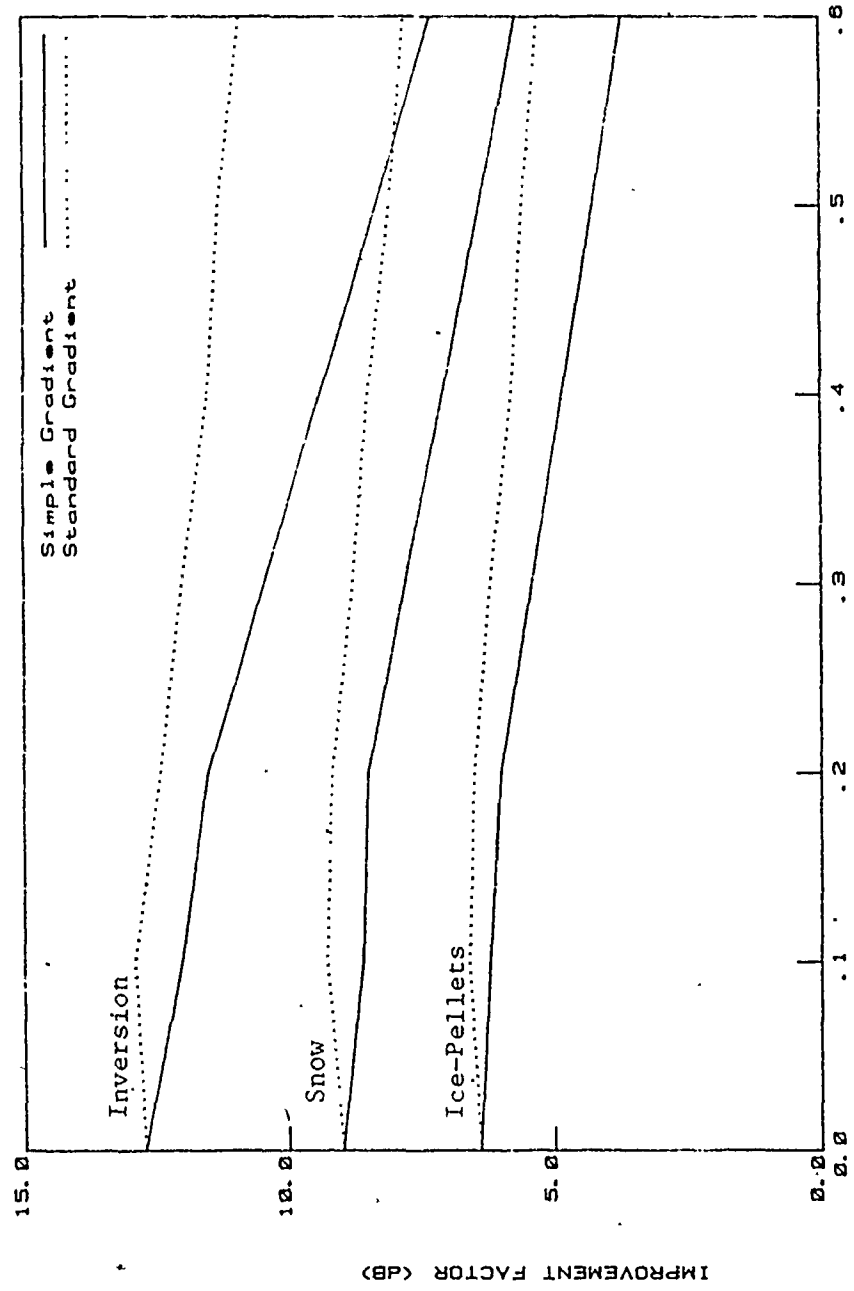

Figure 7.6: Improvement Factors Versus Adaptive Constants (m = 9)

fifth orders of the MTI and standard gradient method lattice filters, with $\mu = 0.2$). The most important observations are that the lattice filter always performed as well as, or better than, the MTI filter, and that the lattice filter performed consistently well. The improvement factors of the lattice and MTI filters were approximately the same for ground clutter, while for inversion clutter the lattice was 3 to 5 dB better, and for the precipitation clutter the lattice was 4 to 6 dB better.

This does not tell the whole story, however. When the whole range-ring is analysed as a unit, the performance is averaged over the whole spread of Doppler velocities present in the clutter. Also, some ground clutter is present in every range-ring. If only small sections of a range-ring are considered (as was done in Chapter 6), the performance ranges from nearly identical (for well-behaved clutter moving tangentially), to improvement factors with as much as 20 dB difference, in favour of the lattice. Figure 7.7 shows a histogram of the improvement factors for 21 such sections, selected from all types of clutter (mainly at points of highest and lowest clutter velocities). Although not too significant in a statistical sense, this figure does show the wider spread of improvement factor values for the MTI filter compared to the lattice filter. Also, the MTI filter improvement factors are often negative (degrading performance), unlike the lattice filter improvement factors, which never go negative.

It would be desirable to compare these results to the results of researchers using other new techniques on actual radar data. However, this is difficult for a number of reasons. The most important of these is that the numerical results presented here are significantly degraded

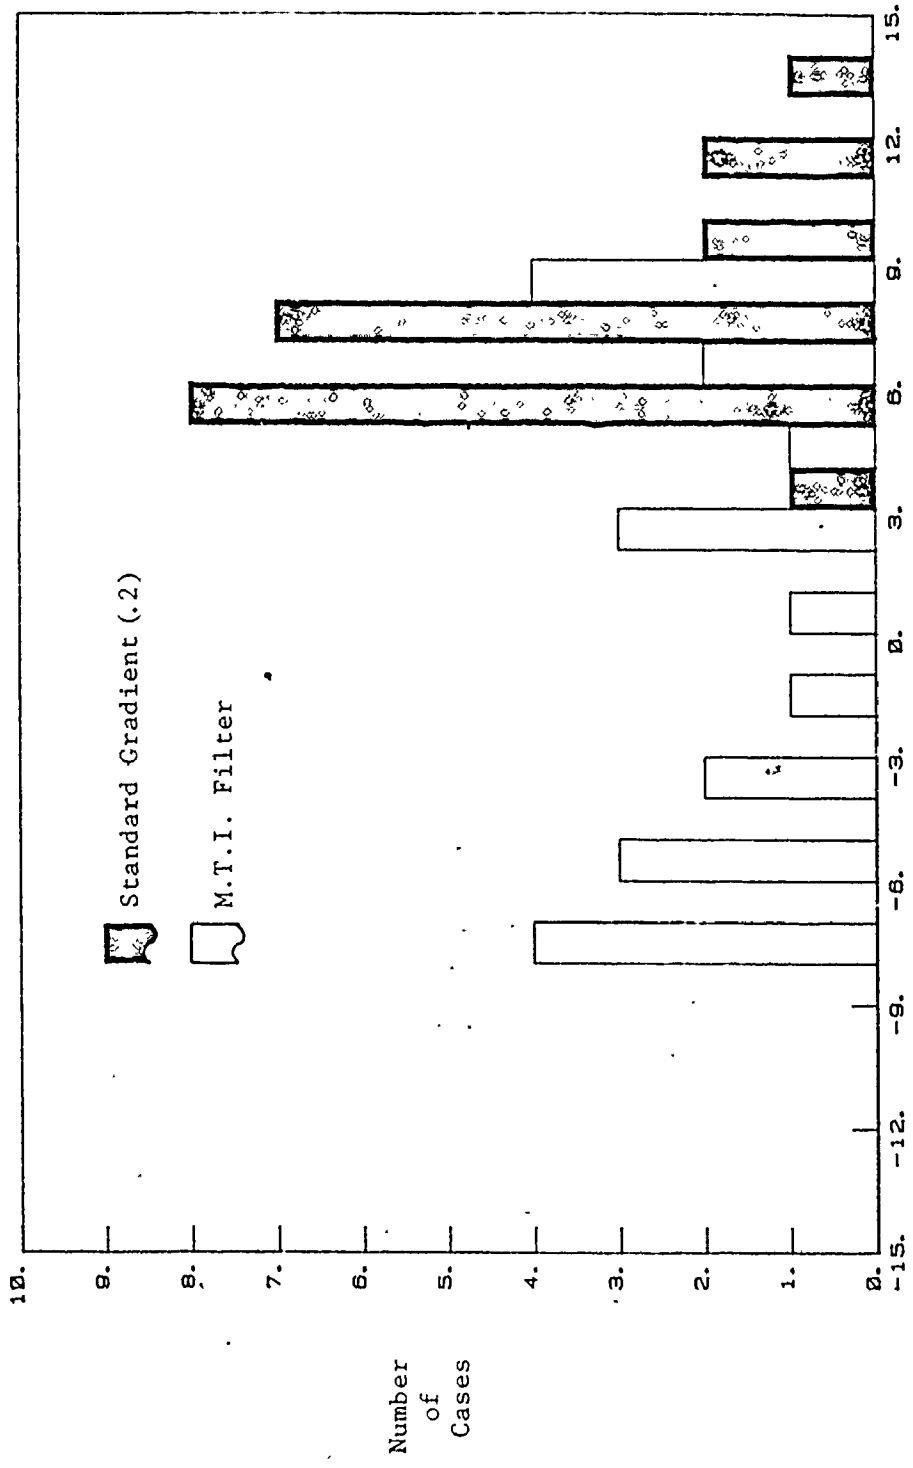

Figure 7.7: Histogram of Sectional Improvement Factors (m = 5)

from their potential values by the limited dynamic range of the data used. This was a result of the signal-to-noise ratio of the video recorder (see Appendix C), which limited the dynamic range to a maximum of about 6 bits. Somewhere between 12 and 16 bits would have been desirable.

Another problem is that very little has been published analysing new techniques using actual radar data, and what has been published often does not include a comparative, quantitative analysis. One adaptive approach, which has received attention recently, is the Moving Target Detector (MTD) system [42]. This complex system uses an adaptive ground clutter map to detect slow targets and an MTI filter followed by an FFT to detect faster targets in weather clutter. The FFT outputs (usually 8 points) are stored in bins which are tested against a threshold based on the response of the surrounding area for that bin. Detection depends on weather and target returns being found in separate bins. Additional processing of the FFT bin and ground-clutter map outputs determines the detection. It should be noted that this system uses several adaptive techniques together, while the lattice-structure filter (as described in this thesis) is only a single technique. In fact, the lattice filter could be used as one component of a similar system, as will be discussed in Chapter 8.

Some results are available using actual radar data to compare the MTD system against MTI filtering alone [47]. However, direct comparison of the MTD system with lattice-structure filtering would have to take account of the fundamental differences between the adaptive clutter cancellation of the lattice and the clutter classification and non-parametric

detection of the MTD. To make a valid comparison, the lattice filter should be part of a system similar to the MTD, with Maximum Entropy spectral estimation (calculated from the lattice reflection coefficients) taking the place of the FFT stage. Simulation studies (using block-processing for the lattice) have shown that this system is significantly better than the MTD with ground clutter, and only slightly inferior with weather clutter [31].

In conclusion, the results presented here should not be considered as an absolute measure of the performance of the lattice-structure filter for this application. Rather, this presentation should be considered as a "proof-of-concept" for this application of the structure. What has been shown here (and in the following sections of this chapter) is that the lattice filter reliably reduces clutter signals (compared to target signals) and consistently out-performs the MTI filter, and that it does so using a simple structure and under adverse conditions.

7.2 Receiver-Operator-Characteristics

Improvement factors are not the only performance measure applied to radar systems. Often, the performance of a system is evaluated on the basis of the probabilities of false alarms and of detections, for given signal-to-interference ratios of a particular system. Such values form the receiver-operator-characteristic (ROC) of the system.

Unlike the improvement factor, the ROC applies to the whole radar system, from antenna to detection and display, and not just to a particular section of the system. For our use, we will hold all these other factors constant, except for the clutter filtering system (lattice or

MTI), to allow a comparative analysis of the response. Also unlike the improvement factor, the ROC approach to the data is non-linear (non-parametric). The data is tested against a threshold, and if it exceeds this, it is judged to be a target. If this is correct, it becomes a "detection"; if incorrect, it is a "false alarm". The proportionate numbers of these two cases at a particular threshold become the basic ROC statistic. Obviously, the desire is to maximize the number of detections while minimizing the number of false alarms, or to reduce the signal-to-interference ratio required to achieve specified limits on these values.

The probability of false alarm (P_{FA}) and probability of detection (P_D) values may be determined, for specified conditions, either theoretically or experimentally (as is done in this thesis). Theoretically, P_{FA} and P_D are defined only given a specified value of the ratio of signal to receiver noise. When receiver noise is replaced by a correlated interference source (clutter), as is the case here, these values are not well defined, and depend largely on the clutter probability density. (Such factors as range, weather, wind, temperature, etc., all have effects). The P_D and P_{FA} values presented here were measured for targets combined with clutter at specified signal-to-clutter ratios (SCR). This is unlike the improvement factor results, where clutter and target gains were evaluated separately. No averaging was applied to the signals before threshold measurements.

The probabilities of detection used in this analysis were calculated by repeatedly adding a target to different locations in the clutter signals (using the five aircraft targets described in section 6.2.1) and testing

the resulting signals against 25 levels of threshold. The same thresholds were used with the clutter signals alone to calculate the probabilities of false alarms. This process was repeated for each SCR.

(The SCR values were measured using the peak target signal and average clutter signal amplitudes. The peak target levels could possibly have a significant noise component. Expected levels of these peak signals, determined by fitting the target returns to the approximate antenna pattern, were found to be an average of 3.4 dB below the measured peak signals for the five aircraft targets. The measured peak levels are generally considered the more appropriate choice of value, and were used here. However, adjustment for the expected values could be achieved if desired by subtracting 3.4 dB from the SCR values presented in this section).

ROC curves resulting from analysis of the actual radar data are presented in Figures 7.8 and 7.9 for rain and ground clutter, respectively. These figures plot the log of P_{FA} versus the SCR (in dB), with the probability of detection set at 0.5. This is a reasonable value for P_D and was chosen because measurement at this value is more accurate than that at any other value. Four curves are shown in each plot. These are the responses of MTI filtering, lattice filtering using the simple gradient method ($\omega = 0$), lattice filtering using the standard gradient method ($\mu = 0.1$), and the unfiltered data, for reference purposes. (These values of ω and μ were determined to be optimum in section 7.1). All filters were fifth order. The ROC curves were calculated using 20 range-rings of each clutter type (10^5 data points). Thus, the points plotted

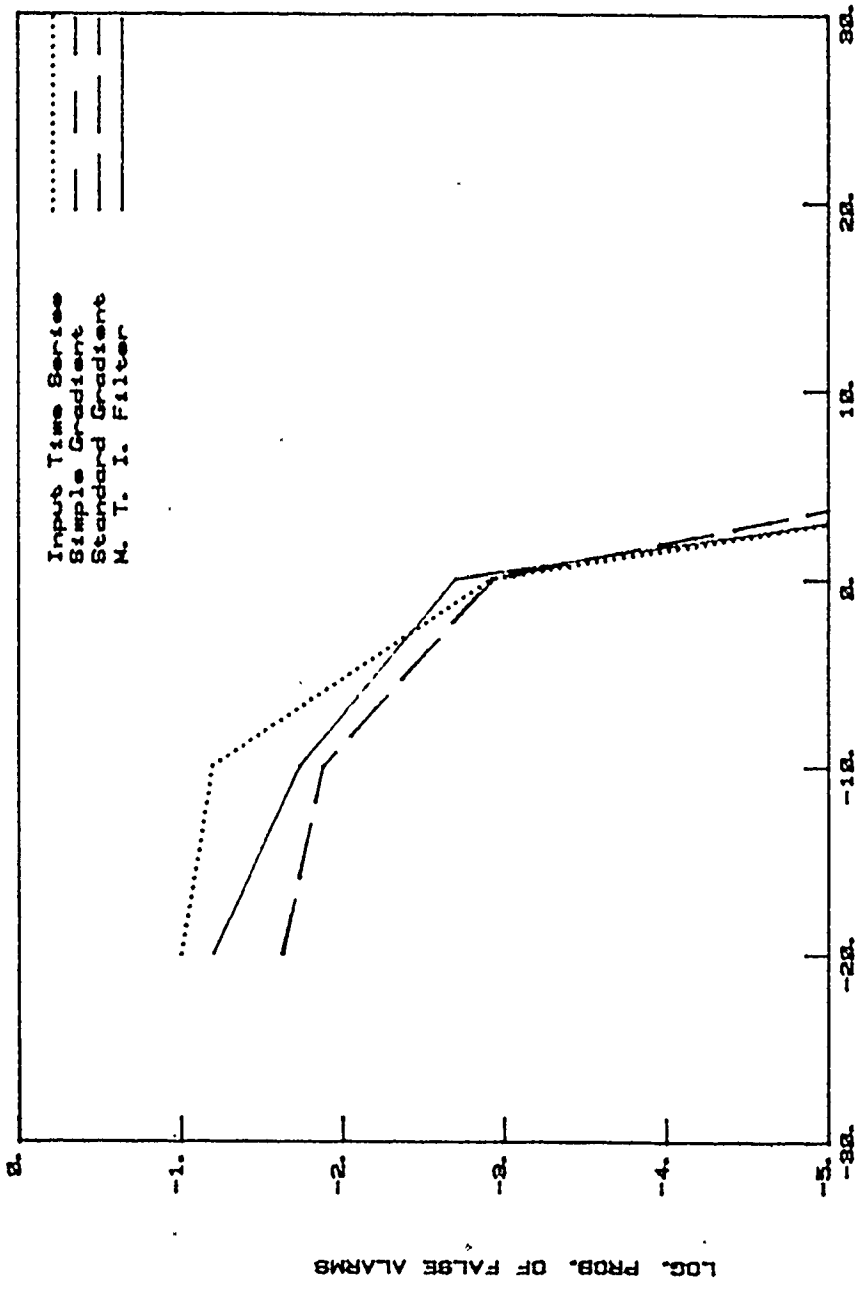

Figure 7.8: ROC Curves for Rain Clutter (m = 5)

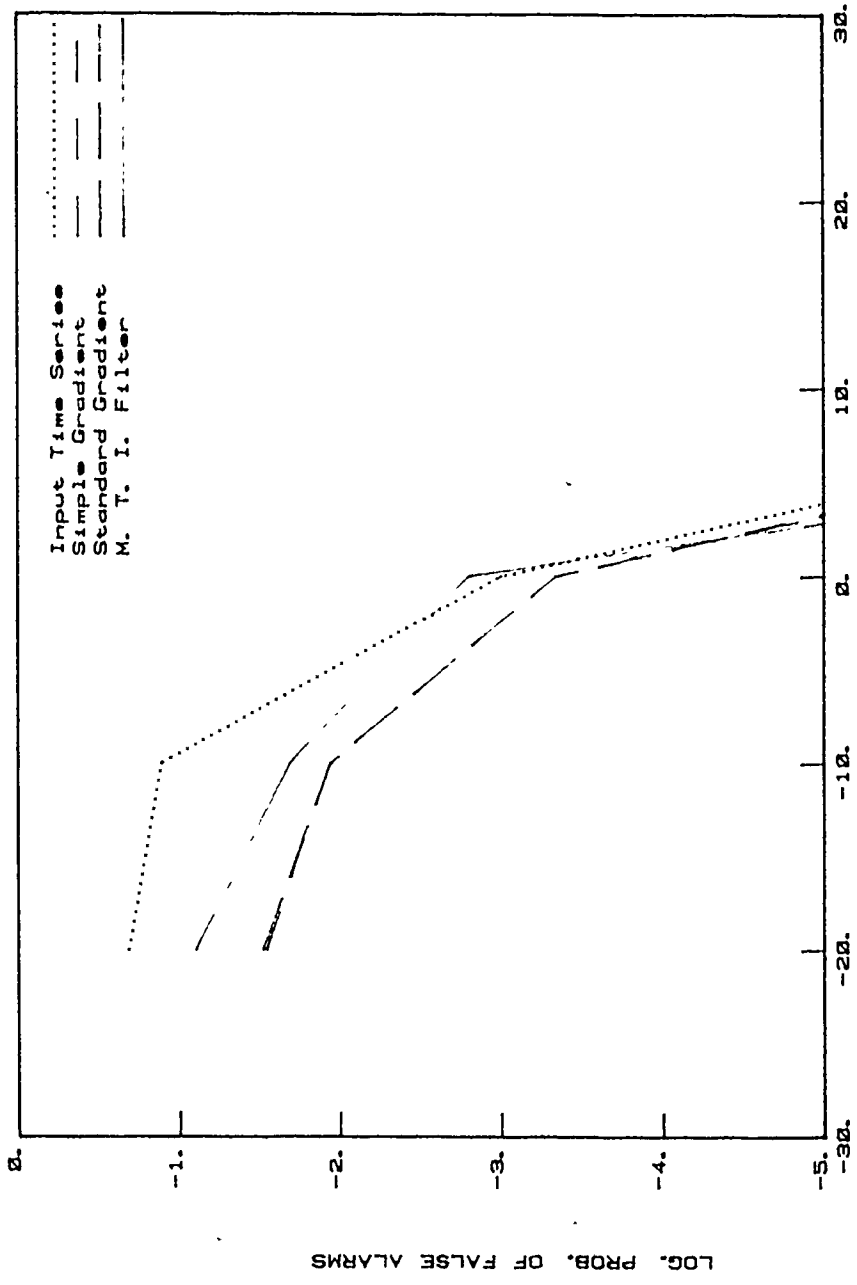

Figure 7.9: ROC Curves for Ground Clutter (m = 5)

for log (P_{FA}) = -5 do not have the same accuracy as the rest of these plots. This data was sampled from the rain and ground clutter conditions described in section 6.3.

These figures indicate that for SCR values above approximately 5 dB, reliable detection with low false alarm rates occurs regardless of any processing used. As the SCR drops to negative values, however, the probability of false alarm increases. In this region, both MTI and lattice filtering provide some reduction of false alarms, with the lattice filtering doing significantly better than the MTI. The performance of the two methods of lattice filtering was approximately equal. From additional data not presented here, it is clear that the lattice filters not only gave lower P_{FA} values than the MTI filter, but also were more consistent in their performance. For SCR values below -20 dB, the measured P_{FA} values were close to the values at -20 dB.

Only the two clutter types presented in Figures 7.8 and 7.9 were analysed to the level of accuracy shown in these figures, due to the great amount of computing required to produce statistical results of this type. Each of these figures represents about 60 hours of central-processor time on the CRL's HP-1000 computer. However, less accurate results were obtained with all five types of clutter described in section 6.3 (using only one range-ring of each type), for a large variety of filter orders and adaptive constants. (This required over 300 more computer hours.) These results were consistent with those presented above, and were used in the calculation of sub-clutter visibilities in the next section.

7.3 Sub-Clutter Visibility

Sub-clutter visibility (SCV) is defined as "The ratio by which the target echo power may be weaker than the coincident clutter echo power and still be detected with specified detection and false-alarm probabilities. Target and clutter powers are measured on a single pulse return and all target radial velocities are assumed equally likely." [36]. For this thesis, we shall soften the definition slightly by assuming that the radial velocities of the five aircraft targets used (see section 6.2.1) are a good enough approximation of "...equally likely". P_D and P_{FA} are here specified as 0.5 and 0.01, respectively; and SCV values are derived from the corresponding SCR values measured as part of the ROC analysis for section 7.2.

The sub-clutter visibility, as a function of filter order, is presented in Figures 7.10 to 7.14 for rain, snow, ice-pellet, ground, and inversion clutter, respectively. Higher SCV values indicate better performance. There are four curves in each figure, one for MTI filtering, one for lattice filtering using the simple gradient method ($\omega = 0.2$), and two for lattice filtering using the standard gradient method (μ = 0.2 & 0.6). This presentation is similar to the improvement factor presentation in section 7.1. The SCV performance of the lattice filters as a function of adaptive constant was also similar to that for improvement factors (in Fig. 7.6). The SCV for the simple gradient method peaked at $\omega = 0$ for all clutter types, while for the standard gradient method, the peak values were found with ω in the range from 0 to 0.1.

As before, these figures show the lattice filters out-performing

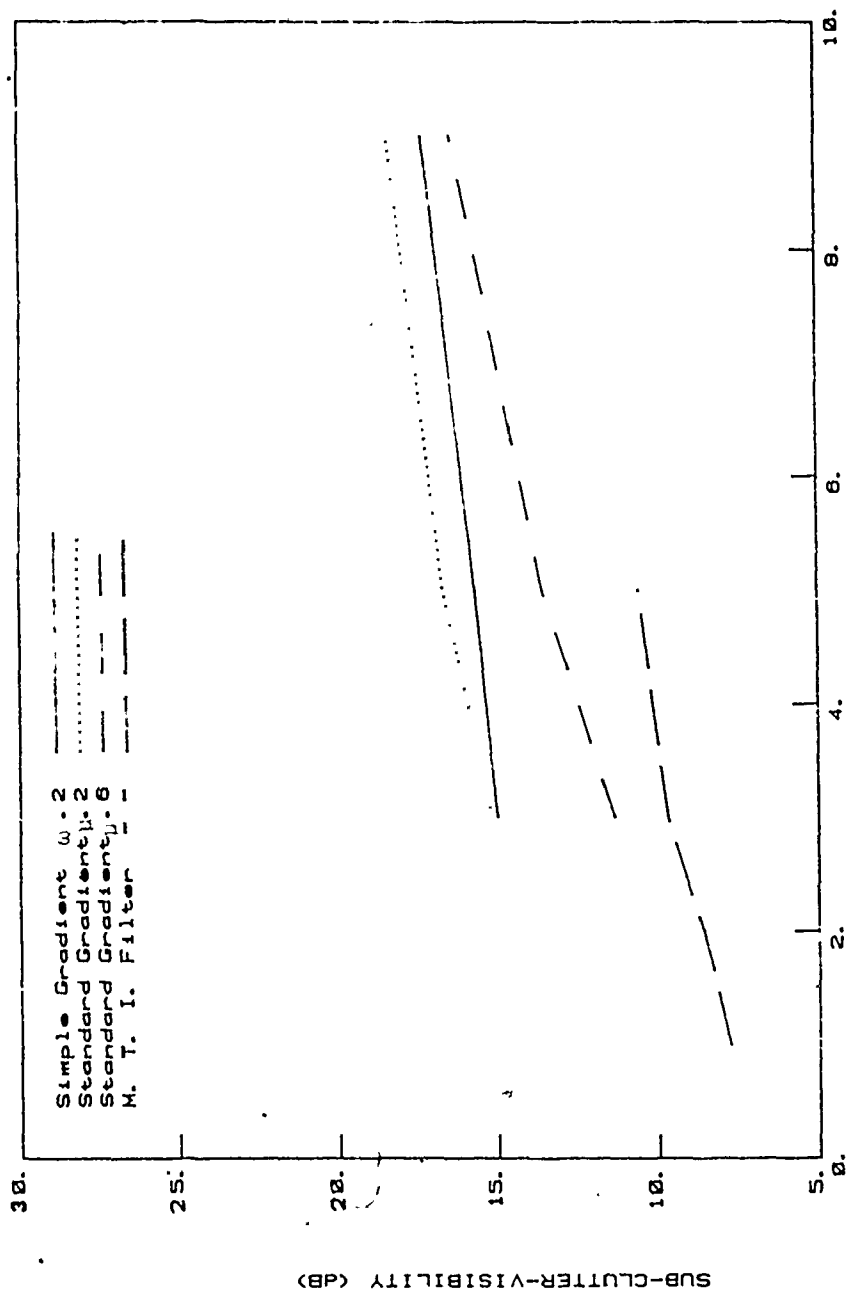

Figure 7.10: Sub-Clutter Visibility for Rain Clutter

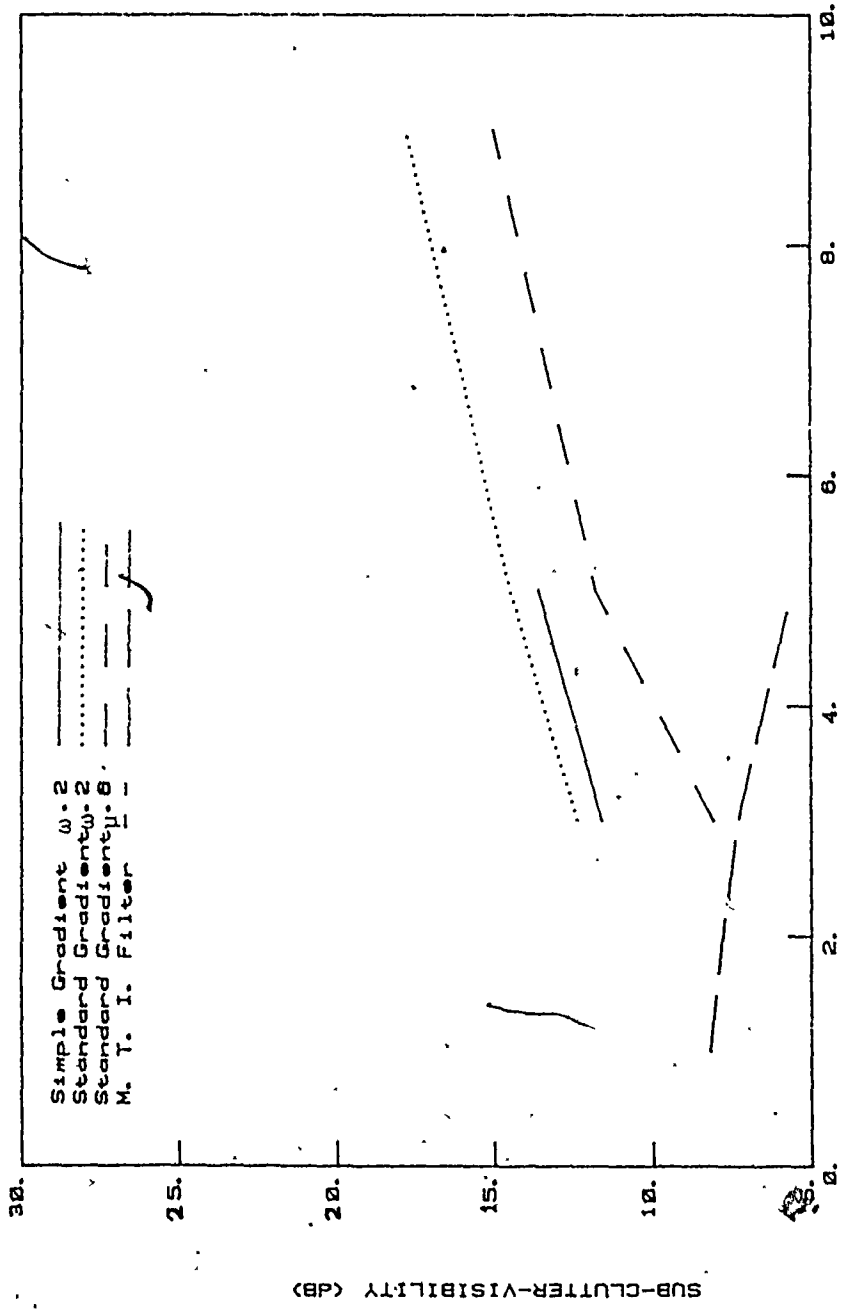

Figure 7.11: Sub-Clutter Visibility for Snow Clutter

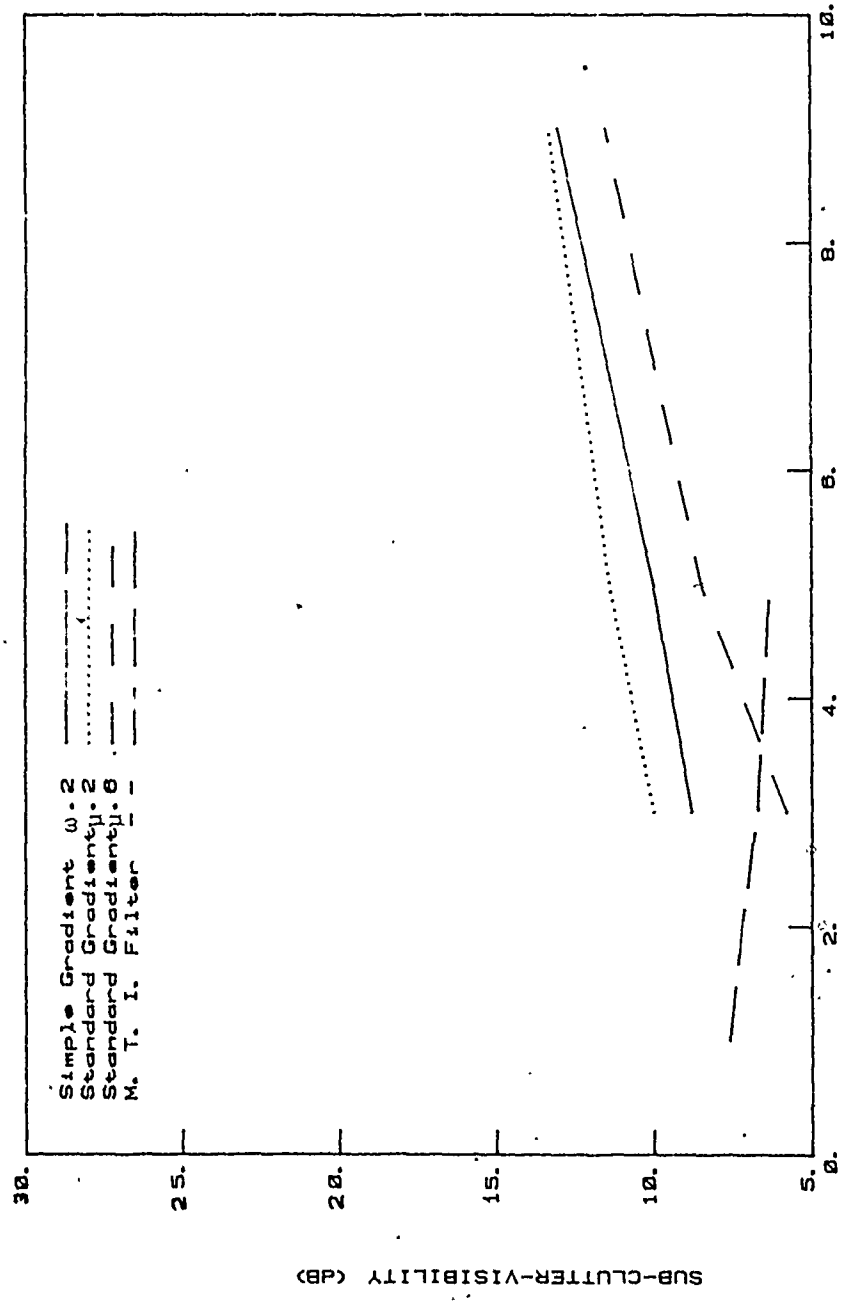

Figure 7.12: Sub-Clutter Visibility for Ice Pellet Clutter

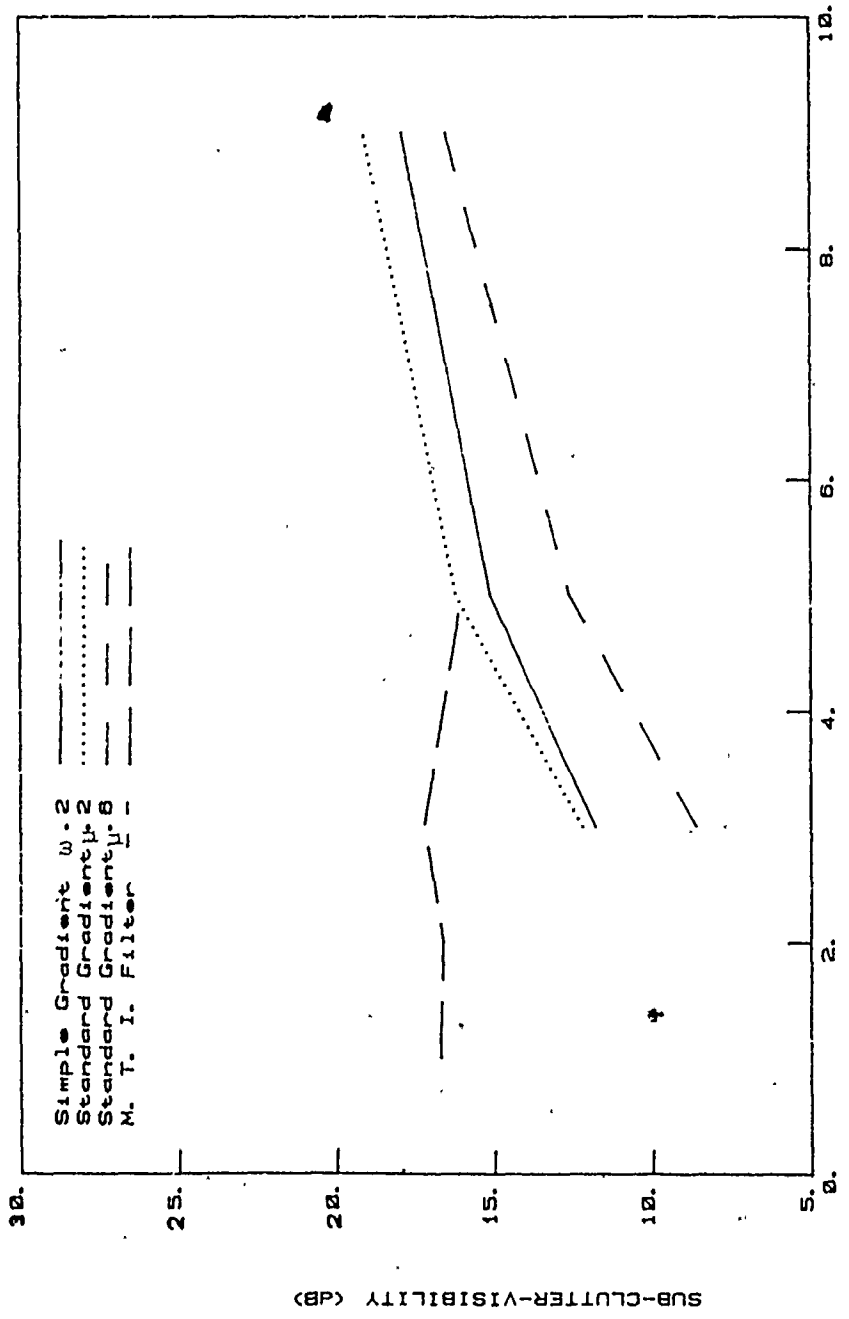

Figure 7.13: Sub-Clutter Visibility for Ground Clutter

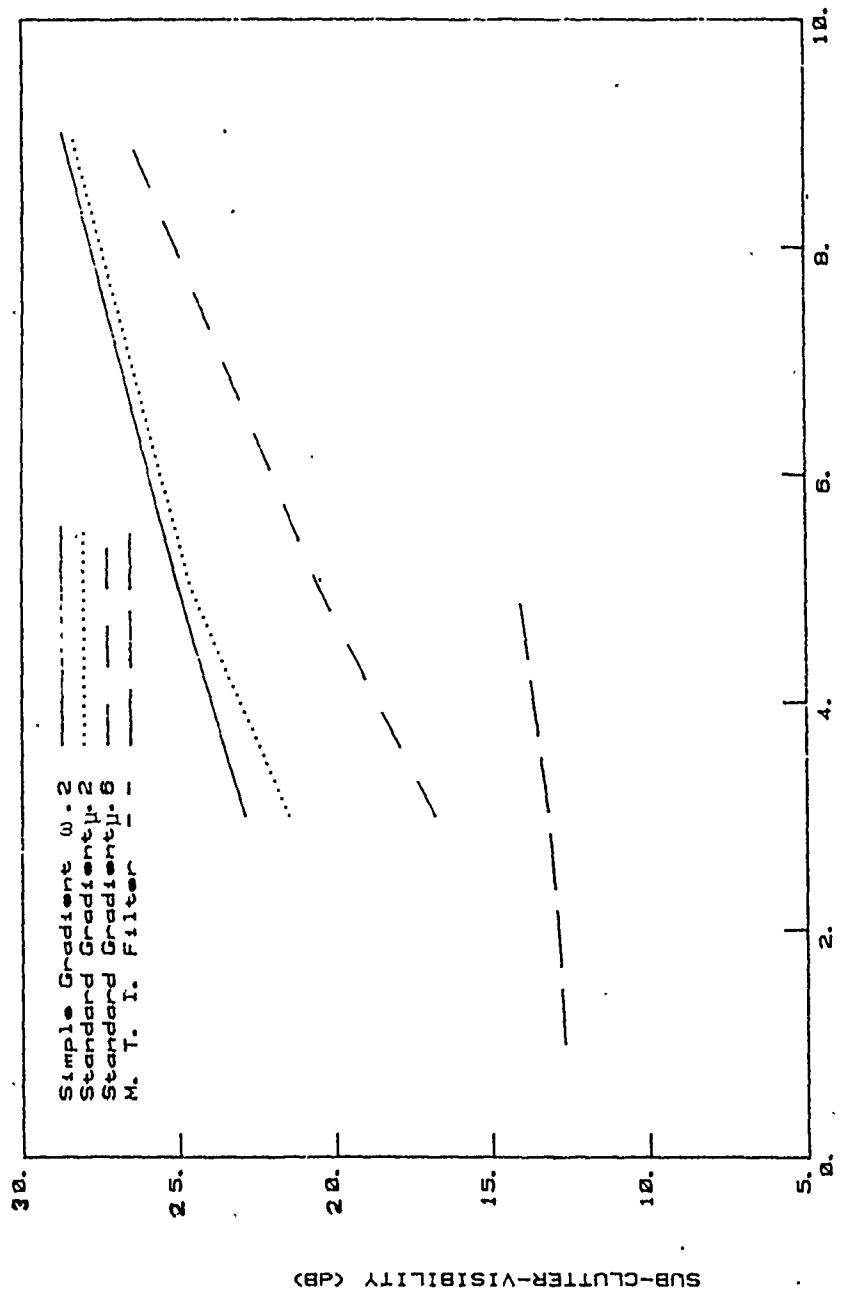

Figure 7.14: Sub-Clutter Visibility for Inversion Clutter

or equalling the MTI filter. Looking at the SCV performance, given a filter order of five for example, the best lattice filter response equalled the MTI filter response for ground clutter, and exceeded it by 5 dB for ice-pellet clutter, 6 dB for rain clutter, 9 dB for snow clutter, and 11 dB for inversion clutter. The lattice filters showed a steady increase in performance with filter order, while the performance of the MTI filter was relatively independent of filter order. The standard gradient method ($\mu = 0.2$) gave the best performance of the lattice filters in these examples. For inversion clutter, the lattice filters did especially well on the absolute scale.

In summary, the SCV results presented here, and the other results presented in this chapter amply demonstrate that the application of the lattice-structure-filter in radar systems, replacing the conventional processing, is very desirable.

CHAPTER 8

CONCLUSION

8.1 Summary of Major Findings

The contributions of this thesis fall into two main categories. These are: (a) additions to the body of theory on lattice-structure prediction-error filtering (particularly in adaptive implementation and convergence properties), and (b) demonstration of the usefulness of this technique for radar signal processing (the elimination of clutter).

After studying the basic lattice structure, a comparative analysis was made of the different algorithms (i.e. minimization criteria) used for the calculation of the lattice reflection coefficients. Theoretical considerations favoured the harmonic-mean algorithm, and this was confirmed by simulation studies using computer-generated radar data.

To implement the lattice structure in a continuously adaptive form, a method of recursively updating the reflection coefficients is needed (generally, a gradient approach to the minimization criteria). Several existing methods were examined, the best suited being the standard gradient method. Two new methods were developed as part of this research, a simple gradient method and an adaptive gradient method. The simple gradient method performed almost as well as the standard gradient, but required less computer time and memory. The adaptive

gradient method adjusts its adaptive rate according to the degree of change in the signal statistics. As this was not appropriate for the particular application under study, only the standard and simple gradient methods were used extensively in this thesis.

A major contribution was made to the understanding of the learning properties of lattice filters implemented by these methods. Equations describing the stationary convergence properties of the lattice structure (using both standard and simple gradient methods) were developed theoretically and confirmed experimentally. The decoupling action of higher lattice stages for convergence was also examined in detail. In addition, a brief study was made of the non-stationary learning characteristics of the lattice.

All of this led up to the application of the lattice-structure PEF to the problem of eliminating radar clutter, in order to detect radar targets. This was examined first using simulated radar data, and then with actual radar data. The filters' output signals and performance measures used to determine the filters' effectiveness were examined in detail. Performance measures examined included improvement factors, receiver-operator-characteristics, and sub-clutter visibility. These results were compared with the results using a conventional MTI filter. In all cases, the lattice filters performed consistently well, and consistently better than the MTI filter, which at times performed very poorly. This led to the conclusion that the lattice-structure PEF is well suited to application as a clutter elimination filter, and would make a good replacement for the systems presently in use.

The need to eliminate clutter signals from the radar return is a

very real concern. In conventional radar systems, targets (such as aircraft) can often be obscured by clutter, or lost due to the processing used to remove clutter. This is a definite threat to air safety. There is a strong theoretical basis for using the lattice-structure PEF in this situation. All that can generally be assumed about the targets and clutter is that they have different Doppler spectra, and that targets are concentrated spatially while clutter tends to be diffuse.

Conventional MTI filters often fail completely because they require the target and clutter Doppler signals to fall into specific, pre-defined ranges. The performance of Doppler radar processors using banks of filters or FFTs is degraded by the spectral inefficiency of these techniques in responding to the Doppler spectra of the signals. To give the best performance in separating targets from clutter, the clutter filter response must be a good estimate of the actual clutter spectrum. An adaptive filter is the only reasonably simple solution for all types of clutter. Of the adaptive (prediction-error) filter structures, only the lattice structure has the properties needed to differentiate the targets and clutter on the basis of spatial distribution (so as to eliminate only the clutter). Given this theoretical basis, and the experimental results confirming it, the arguments for using the lattice structure for radar clutter filtering are very strong.

8.2 Recommendations for Future Research

Arising out of this research, there are a number of areas which could be suggested for future study. These can be roughly divided into

two groups, those dealing with the theory and operation of the lattice structure, and those dealing with the application of this structure to radar signal processing.

Some of the suggested research areas falling into the first group include:

(1) Lattice algorithms - With some justification, the harmonic-mean algorithm received the most attention in this thesis. However, some of the other algorithms (especially the forward and backward algorithm) deserve further attention, including simulation studies and statistical analysis.

(2) Adaptive gradient methods - Although not appropriate for this research, these methods could be very useful in other applications. In particular, the adaptive gradient method developed in section 3.2 should be pursued and compared with the "least-squares" method of Morf and Lee [26,27]. (This least-squares method was judged in section 3.4 to have serious theoretical flaws).

(3) Decoupling and convergence properties of higher-order lattice stages - Obtaining a satisfactory theoretical solution for this problem may prove very difficult, if not impossible. However, it is an area which should be examined.

(4) Non-stationary learning characteristics - This thesis barely scratched the surface of this important field; much more could be done.

Falling into the second grouping, concerning the application of the lattice to radar clutter, are the following suggestions:

(5) Further analysis of existing data - This could include analysis of the actual radar data using other lattice algorithms and methods, as well as direct comparisons to other adaptive clutter filtering systems presently being researched.

(6) Classification of radar clutter - Preliminary results from this research indicate that the lattice filter using block processing provides much better spectral estimates (used for classification) than the continuous implementation. However, there may exist some way of manipulating the reflection coefficients of the continuous implementation, so as to take advantage of the best properties of both the block and continuous implementations. In either case, the classification could be based directly on the reflection coefficients, rather than the spectral estimate.

(7) Real-time operation - This is perhaps the most important recommendation of those presented here. If the lattice filter is ever to be applied operationally to radar signal processing (and this would be very desirable), a necessary first step is a lattice radar-processor operating in real time. This would provide a demonstration of the system's usefulness, a large volume of statistical results (necessary for measuring the system's true performance), and a tool for further system refinement and development. The following paragraphs examine some of the practical considerations for a real-time processing system of this type.

The system would start with the digitized I and Q signals from the radar receiver. (See Appendix C).. The analog-to-digital conversion should have a dynamic rangeof 12 to 16 bits and a sampling rate in excess of 2 MHz. These signals would then be fed into the lattice-structure filter system. There would be, in effect, one lattice filter for each

range bin in the data, with each filter operating independently on its own range ring. In practice, there would probably be only a single lattice filter structure, for which the memory (reflection coefficients and backward filter outputs) would be switched between range bins. The memory switching could be either by sequential addressing, or by memory rotation (as is used for the MTI filtering of the radar processor described in Appendix C). The filtering could be shared among a small number of filter structures if a single structure would not cycle fast enough.

The filter structures might be implemented either as hard-wired digital circuitry or programmed in a dedicated computer, depending on economic and processing speed considerations. An array processor would be well suited for this application, when provided with direct input capabilities. The most economical lattice filter to implement in this system would use the simple gradient method with $\omega = 0$. (Constraining ω to be zero reduces the memory requirements by 50% over the simple gradient method, and 60% over the standard gradient method, and reduces computation time by eliminating multiplications of ω). However, for research purposes more flexibility of the filter would be useful. The filter order could be set as high as the processing equipment would allow, with an order of five or above desirable.

One environmental factor, which would degrade the lattice filter performance for separating targets from clutter, is points of sharply defined ground clutter (such as towers and some buildings). These points would look like zero-velocity targets to the lattice filter. This problem could be eliminated by prefacing the lattice filter with an adaptive ground-

clutter map. This map would store continuously the non-moving I and Q signals from multiple rotations of the radar antenna. The appropriate map values would then be subtracted from the I and Q signals before filtering. (This process requires phase coherence between pulses).

The lattice-radar processing system described above could be used in a variety of operational environments. Some of the most interesting of these are the automated target tracking setups to be used for air-traffic control. To operate properly, these tracking systems require a greater level of clutter rejection (especially weather clutter) than is available with conventional systems. They also require that tangentially moving targets not be lost for long periods. Figure 8.1 shows a block diagram of a lattice-radar processor, as it would be applied for a tracking system. The lattice processor would fulfil the same role as the MTD system [42] in automated tracking setups, and could potentially perform better. However, whatever the radar application, the lattice filter can offer significant improvements in clutter rejection and target enhancement.

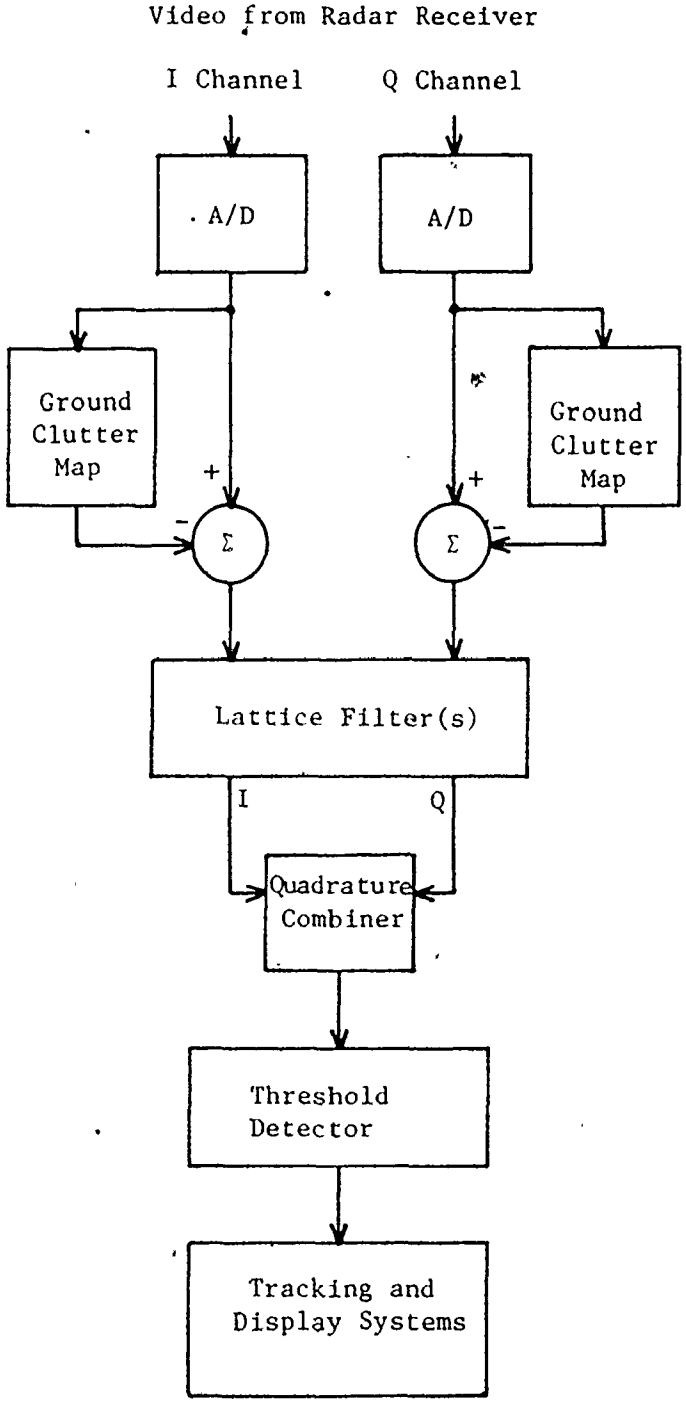

Figure 8.1: Lattice Radar Processor

REFERENCES

[1] J. Makhoul, "Linear Prediction: A Tutorial Review", Proc. IEEE, vol. 63, pp. 561-580, (1975).

[2] S. Haykin and C. Hawkes, "Adaptive Digital Filtering for Coherent MTI Radar", Information Sciences 11, pp. 335-359, (1976).

[3] B. Widrow, J. McCool, M. Larimore, and C. Johnson Jr., "Stationary and Nonstationary Learning Characteristics of the LMS Adaptive Filter", Proc. IEEE, vol. 64, pp. 1151-1162, (1976).

[4] J. Makhoul, "Stable and Efficient Lattice Methods for Linear Prediction", IEEE Trans. on Acoustics, Speech, and Signal Processing, vol. ASSP-25, pp. 423-427, (1977).

[5] J. Burg, "Maximum Entropy Spectral Analysis", presented to 37th Mtg. of Society of Exploration Geophysicists, Oklahoma City, (1967).

[6] J. Burg, "Maximum Entropy Spectral Analysis", Ph.D. Thesis, Stanford University, (1975).

[7] F. Itakura and S. Saito, "Digital Filtering Techniques for Speech Analysis and Synthesis", presented at the 7th Int. Cong. Acoustics, Budapest, paper 25-C-1, (1971).

[8] E. Satorius and J. Pack, "Application of Least Squares Lattice Algorithms to Adaptive Equalization", IEEE Trans. on Communications, vol. COM-29, pp. 136-142, (1981).

[9] D. Riley and J. Burg, "Time and Space Adaptive Deconvolution Filters", presented to 42nd Mtg. of Society of Exploration Geophysicists, (1972).

[10] T. Carter, "Study of an Adaptive Lattice Structure for Linear Prediction Analysis of Speech", Proc. 1978 IEEE Int. Conf. on ASSP, vol. 78CH1285-6, pp. 27-30, (1978).

[11] R. Viswanathan and J. Makhoul, "Sequential Lattice Methods for Stable Linear Prediction", Proc. EASCON-76 (IEEE), pp. 155A-H, (1976).

[12] J. Makhoul, "A Class of All-Zero Lattice Digital Filters: Properties and Applications", IEEE Trans. on ASSP, vol. ASSP-26, pp. 304-314, (1978).

[13] L. Griffiths, "An Adaptive Lattice Structure for Noise-Cancelling Applications", Proc. 1978 IEEE Int. Conf. on ASSP, vol. 78CH1285-6, pp. 87-90, (1978).

[14] S. Kesler, "Nonlinear Spectral Analysis of Radar Clutter", Ph.D. Thesis, McMaster University, (1978).

[15] S. Haykin (editor), "Nonlinear Methods of Spectral Analysis", (Springer-Verlag, 1979).

[16] M. Alam, and A. Sage, "Sequential Estimation and Identification of Reflection Coefficients by Minimax Entropy Inverse Filtering", Comput. and Elec. Engng., vol. 2, pp. 315-338, (1975).

[17] M. Atashroo, "Autocorrelation Prediction", Proc. 1977 IEEE Int. Conf. on ASSP, vol. 77CH1197-3, pp. 5-9, (1977).

[18] D. Parikh, N. Ahmed, and S. Stearns, "An Adaptive Lattice Algorithm for Recursive Filters", IEEE Trans. on ASSP, vol. ASSP-28, pp. 110-111, (1980).

[19] M. Honig and D. Messerschmitt, "Convergence Properties of the Adaptive Digital Lattice Filter", Proc. 1980 IEEE Int. Conf. on ASSP, vol. CH1559-4, pp. 984-988, (1980).

[20] L. Griffiths, "A Continuously-Adaptive Filter Implemented as a Lattice Structure", Proc. 1977 IEEE Int. Conf. on ASSP, vol. 77CH1197-3, pp. 683-686, (1977).

[21] Private communication from Dr. S. Haykin, "Adaptive Systems", EE776 course notes, McMaster University, (1980).

[22] J. Makhoul and R. Viswanathan, "Adaptive Lattice Methods for Linear Prediction", Proc. 1978 IEEE Int. Conf. on ASSP, vol. 78CH1285-6, pp. 83-86, (1978).

[23] C. Gibson, S. Haykin, and S. Kesler, "Maximum Entropy (Adaptive) Filtering Applied to Radar Clutter", Proc. 1979 IEEE Int. Conf. on ASSP, vol. 79CH1379-7, pp. 166-169, (1979).

[24] C. Gibson and S. Haykin, "Learning Characteristics of Adaptive Lattice Filtering Algorithms", IEEE Trans. on ASSP, vol. ASSP-28, pp. 681-691, (1980).

[25] R. Kalman, "A New Approach to Linear Filtering and Prediction Problems", Trans. ASME, Journal of Basic Engineering, vol. 82D, pp. 34-45, (1960).

[26] M. Morf and D. Lee, "Recursive Least Squares Ladder Forms for Fast Parameter Tracking", Proc. IEEE 1978 Conf. on Decision and Control, vol. CH1392-0, pp. 1362-1367, (1978).

[27] D. Lee and M. Morf, "A Novel Innovations Based Time-Domain Pitch Detector", Proc. 1980 IEEE Int. Conf. on ASSP, vol. CH1559-4, pp. 40-44, (1980).

[28] J. Pack and E. Satorius, "Least Squares, Adaptive Lattice Algorithms", Naval Ocean Systems Center, Technical Report 423, San Diego, Calif., (1979).

[29] E. Satorius and S. Alexander, "Channel Equalization Using Adaptive Lattice Algorithms", IEEE Trans. on Communications, vol. COM-27, pp. 899-905, (1979).

[30] C. Hawkes, "Adaptive Digital Filtering for Coherent MTI Radar", Ph.D. Thesis, McMaster University, (1975).

[31] S. Haykin and H. Chan, "Computer Simulation Study of a Radar Doppler Processor Using the MEM", Proc. IEE, vol. 127, pt. F, #6, pp. 464-470, (1980).

[32] E. Barlow, "Doppler Radar", Proc. IRE, vol. 37, pp. 340-355, (1949).

[33] W. Flock, "Radar Signature Studies Applicable to Bird-Aircraft Collision Avoidance", Final Report for period June 18, 1970 to Nov. 15, 1973; Air Force Weapons Laboratory, Air Force System Command, Kirtland Air Force Base, NM.

[34] H. Ward, "Clutter Filter Performance Measures", Proc. 1980, IEEE Int. Radar Conf., vol. CH1493-6, pp. 231-239, (1980).

[35] F. Nathanson, "Further Thoughts on Clutter Filter Performance Measures", Proc. 1980 IEEE Int. Radar Conf., vol. CH1493-6, pp. 240-242, (1980).

[36] IEEE Standard Dictionary of Electrical and Electronics terms, 2nd Ed., J. Wiley and Sons, (1977).

[37] C. Gibson and S. Haykin, "A Comparison of Algorithms for the Calculation of Adaptive Lattice Filters", Proc. 1980 IEEE Int. Conf. on ASSP, vol. CH1559-4, pp. 978-983, (1980).

[38] S. Haykin and S. Kesler, "The Complex Form of the Maximum Entropy Method for Spectral Estimation", Proc. IEEE, vol. 64, pp. 822-823, (1976).

[39] "A Brief Description of the ASR-8 Airport Surveillance Radar", Publ. No. SP05A-EG76, Texas Instruments Inc., (1976).

[40] W. Fischbein, S. Graveline, and O. Rittenbach, "Clutter Attenuation Analysis", Tech. Report ECOM 2808, (1967).

[41] N. Currie, F. Dyer, and R. Hayes, "Radar Land Clutter Measurements at Frequencies of 9.5, 16, 35, and 95 GHz", Tech. Report #3, Project A-1485, Georgia Inst. of Tech., Engineering Experiment Station, (1975).

[42] L. Cartledge and R. O'Donnell, "Description and Performance Evaluation of the Moving Target Detector", Project Report ATC-69, Lincoln Laboratory, Mass. Inst. of Tech., (1977).

[43] W. Burdic, "Detection of Narrowband Signals Using Time-Domain Adaptive Filters", IEEE Trans. on Aerospace and Electronic Systems, vol. AES-14, pp. 578-591, (1978).

[44] M. Honig and D. Messerschmitt, "Comparison of LS and LMS Lattice Predictor Algorithms Using Two Performance Criteria", submitted to IEEE Trans. on ASSP, (1981).

[45] S. Tretter, "Intro. to Discrete-Time Signal Processing"; John Wiley and Sons, (1976).

[46] D. Godard, "Channel Equalization Using a Kalman Filter for Fast Data Transmission", IBM J. of Research and Development, pp. 267-273 (1974).

[47] T. Irabu, et. al., "On the Performance of a 2-Dimensional Clutter Rejection System", Proc. 1980 IEEE Int. Radar Conf., pp. 311-316, (1980).

[48] Environment Canada, "Weather Ways", Supply and Services Canada, (1978).

[49] C. Gibson and S. Haykin, "Non-Stationary Learning Characteristics of Adaptive Lattice Filters", accepted for publication in Proc. 1982 IEEE Int. Conf. on ASSP.

[50] C. Gibson and S. Haykin, "Performance Studies of Adaptive Lattice Prediction-Error Filters for Target Detection in a Radar Environment Using Real Data", Proc. 1981 IEEE Int. Conf. on ASSP, vol. 81CH1610-5, pp. 1054-1057, (1981).

APPENDIX A

COMPUTER IMPLEMENTATION OF ADAPTIVE LATTICE FILTERING METHODS

This appendix describes the computer routines used to implement the various methods of lattice-structure adaptive filtering developed in this thesis. Figure A.1 is the flow diagram of these routines. Figures A.2, A.3, and A.4 contain the Fortran listings of routines using the simple gradient, adaptive gradient, and standard gradient approaches, respectively. (As presented, these routines all implement Burg's harmonic-mean algorithm; however, they are easily modified to implement the other algorithms detailed in section 2.3). Finally, the routine used for MTI filtering is presented in Fig. A.5. This routine implements a normalized multi-stage delay-line canceller, as shown in Fig. A.5(b). These four routines are designed to be interchangeable within a program.

The lattice filtering routines are called once for each data sample. When a new data sample is received, it is passed through the filter, the filter is updated, and the outputs and new filter reflection coefficient values are returned to the calling program. The input sample is sent in the first positions of the A array, which on its return contains the filter outputs (the forward prediction errors $f_o(n)$ to $f_m(n)$). The B array contains the backward prediction errors, and the C array contains the reflection coefficients. (For the standard gradient method routine, the

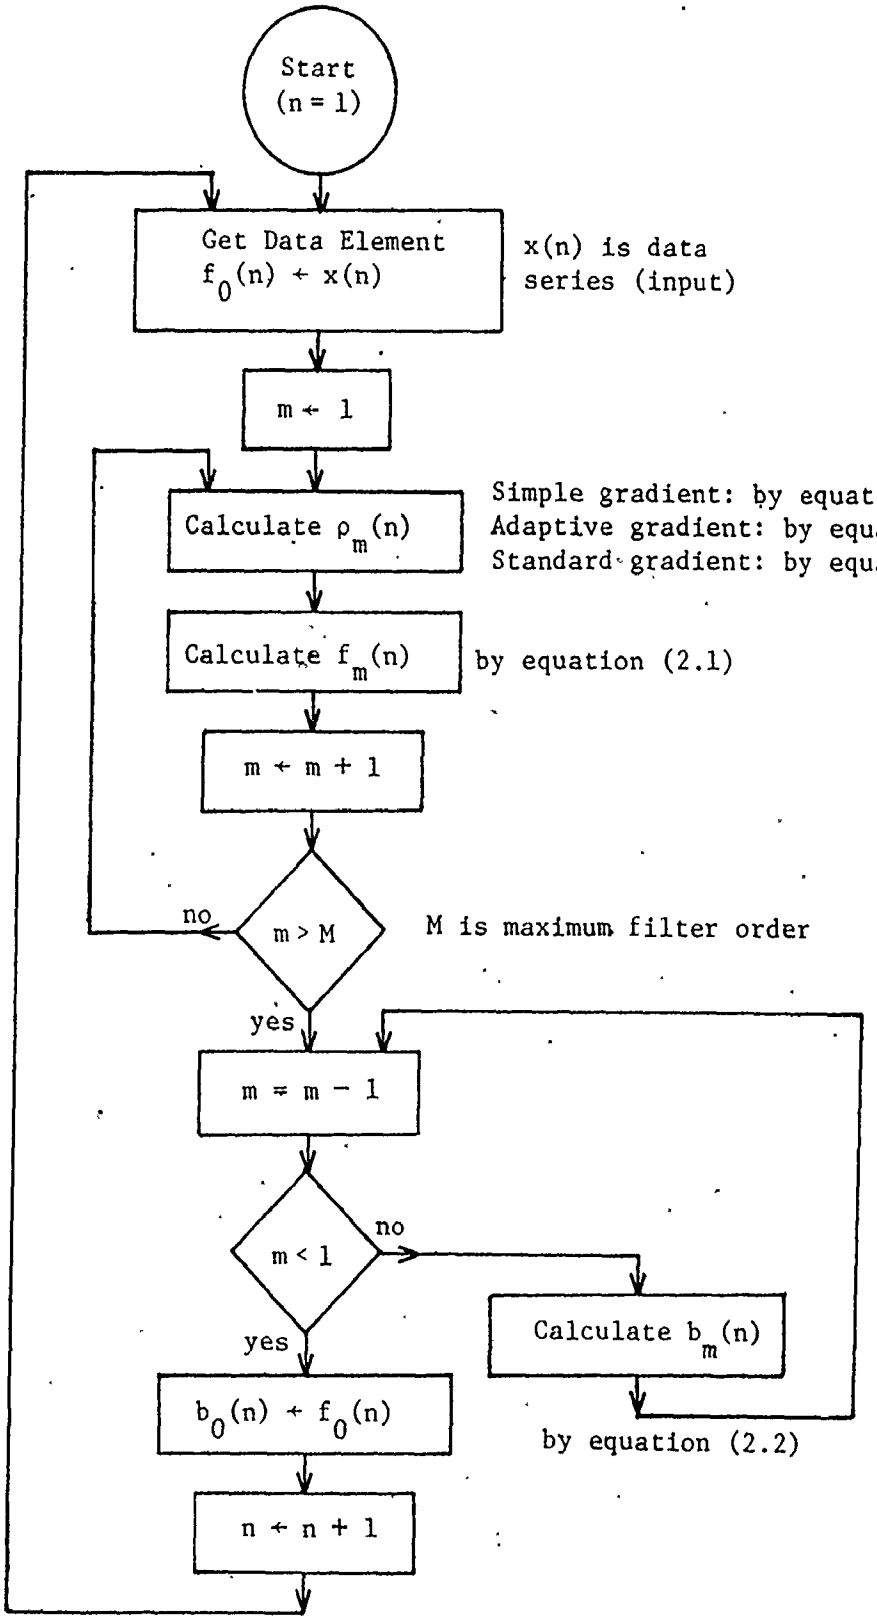

Figure A.1: Flow Diagram of Lattice Filtering Routines

```
      SUBROUTINE ADMEF (A,B,C,N,XA)
C
C  ADMEF USES BURG LATTICE TO ADAPTIVELY FILTER
C   COMPLEX DATA SERIES.
C  ADMEF HANDLES ONE ELEMENT OF SERIES AT A TIME.
C  ( FORWARD FILTER OUTPUTS ARE IN A,
C   BACKWARD FILTER OUTPUTS ARE IN B,
C   FILTER COEFFICIENTS ARE IN C ).
C  WRITTEN BY C. GIBSON, JUNE 1978.
C
C  A(1),A(2) ARE COMPLEX INPUT DATA FOR ONE ELEMENT OF SERIES.
C  A(2M+1),A(2M+2) ARE COMPLEX OUTPUT OF FILTER ORDER M (M=1 TO N).
C  B(2N) IS ARRAY OF BACKWARD FILTER OUTPUTS.
C  C(2N) IS ARRAY OF FILTER COEFFICIENTS.
C  N IS MAXIMUM FILTER ORDER.
C  XA IS EXPONENTIAL MEMORY FACTOR ( FRACTION OF OLD FILTER
C   COEFF. USED FOR NEW FILTER COEFF. ).
C
C  B AND C ARRAYS ARE UPDATED BY ADMEF AND MUST BE
C   INITIALLIZED (=0) BEFORE FIRST CALL.
C  IN ALL ARRAYS, ODD SUBSCRIPTS ARE REAL PART, EVEN
C   SUBSCRIPTS ARE IMAGINARY PART.
C
      REAL A(1), B(1), C(1)
      M = 2*N - 1
      XB = 2.*XA - 2.
      DO 10 I=1,M,2
      J = I + 1
      X = A(I)*A(I) + A(J)*A(J) + B(I)*B(I) + B(J)*B(J)
      IF(X.EQ.0.) X = 1.
      C(I) = XA*C(I) + XB*(A(I)*B(I) + A(J)*B(J))/X
      C(J) = XA*C(J) + XB*(A(J)*B(I) - A(I)*B(J))/X
      A(I+2) = A(I) + B(I)*C(I) - B(J)*C(J)
      A(J+2) = A(J) + B(J)*C(I) + B(I)*C(J)
   10 CONTINUE
      IF(N.LT.2) GO TO 30
      DO 20 K=3,M,2
      I = M - K + 1
      J = I + 1
      B(I+2) = B(I) + A(I)*C(I) + A(J)*C(J)
      B(J+2) = B(J) + A(J)*C(I) - A(I)*C(J)
   20 CONTINUE
   30 B(1) = A(1)
      B(2) = A(2)
      RETURN
      END
```

Figure A.2: Routine Implementing Simple Gradient Method

```
      SUBROUTINE AUTL1 (A,B,C,N)
C
C  AUTL1 USES BURG LATTICE TO ADAPTIVELY FILTER
C   COMPLEX DATA SERIES.
C  AUTL1 HANDLES ONE ELEMENT OF SERIES AT A TIME,
C  ( FORWARD FILTER OUTPUTS ARE IN A;
C   BACKWARD FILTER OUTPUTS ARE IN B,
C   FILTER COEFFICIENTS ARE IN C ).
C  WRITTEN BY C. GIBSON, 1981.
C
C  A(1),A(2) ARE COMPLEX INPUT DATA FOR ONE ELEMENT OF SERIES.
C  A(2M+1),A(2M+2) ARE COMPLEX OUTPUT OF FILTER ORDER M (M=1 TO N).
C  B(2N) IS ARRAY OF BACKWARD FILTER OUTPUTS.
C  C(2N) IS ARRAY OF FILTER COEFFICIENTS.
C  N IS MAXIMUM FILTER ORDER.
C
C  B AND C ARRAYS ARE UPDATED BY ADMEF AND MUST BE
C   INITIALLIZED (=0) BEFORE FIRST CALL.
C  IN ALL ARRAYS, ODD SUBSCRIPTS ARE REAL PART, EVEN
C   SUBSCRIPTS ARE IMAGINARY PART.
C
      REAL A(1), B(1), C(1)
      M = 2*N - 1
      DO 10 I=1,M,2
      J = I + 1
      X = (A(I)*A(I) + A(J)*A(J) + B(I)*B(I) + B(J)*B(J))/2.
      IF(X.EQ.0.) X = 1.
      Y = C(I) + (A(I)*B(I) + A(J)*B(J))/X
      Z = C(J) + (A(J)*B(I) - A(I)*B(J))/X
      D = SQRT( Y*Y + Z*Z )/2.
      C(I) = C(I) - D*Y
      C(J) = C(J) - D*Z
      A(I+2) = A(I) + B(I)*C(I) - B(J)*C(J)
      A(J+2) = A(J) + B(J)*C(I) + B(I)*C(J)
   10 CONTINUE
      IF(N.LT.2) GO TO 30
      DO 20 K=3,M,2
      I = M - K + 1
      J = I + 1
      B(I+2) = B(I) + A(I)*C(I) + A(J)*C(J)
      B(J+2) = B(J) + A(J)*C(I) - A(I)*C(J)
   20 CONTINUE
   30 B(1) = A(1)
      B(2) = A(2)
      RETURN
      END
```

Figure A.3: Routine Implementing Adaptive Gradient Method

```
      SUBROUTINE LATT (A,B,C,D,N,XA)
C
C  LATT USES BURG LATTICE METHOD TO ADAPTIVELY
C   FILTER COMPLEX DATA SERIES.
C  LATT HANDLES ONE ELEMENT OF SERIES AT A TIME.
C  ( FORWARD FILTER OUTPUTS ARE IN A,
C   BACKWARD FILTER ELEMENTS ARE IN B,
C   FILTER COEFFICIENTS ARE IN C AND D ).
C  WRITTEN BY C. GIBSON, MARCH 1979.

C
C  A(1),A(2) ARE COMPLEX INPUT DATA FOR ONE ELEMENT OF SERIES.
C  A(2M+1),A(2M+2) ARE COMPLEX OUTPUT OF FILTER ORDER M (M=1 TO N).
C  B(2N) IS ARRAY OF BACKWARD FILTER ELEMENTS.
C  C(2N) IS ARRAY OF FILTER COEFFICIENT NUMERATORS.
C  D(N) IS ARRAY OF FILTER COEFFICIENT DENOMINATORS.
C  N IS MAXIMUM FILTER ORDER.
C  XA IS EXPONENTIAL MEMORY FACTOR ( FRACTION OF OLD FILTER
C   COEFF. USED FOR NEW FILTER COEFF. ).
C
C  B, C, AND D ARRAYS ARE PREVIOUSLY SET AND UPDATED BY LATT.
C  THEY MUST BE INITIALLIZED BEFORE FIRST CALL.
C  IN A,B,C ARRAYS ODD SUBSCRIPTS ARE REAL PART, EVEN
C   SUBSCRIPTS ARE IMAGINARY PART.
C
      REAL A(1),B(1),C(1),D(1)
      M = 2*N - 1
      K = 0
      DO 10 I=1,M,2
      J = I + 1
      K = K + 1
      D(K) = XA*D(K) + A(I)*A(I) + A(J)*A(J) + B(I)*B(I) + B(J)*B(J)
      IF(D(K).EQ.0.) D(K) = 1.
      C(I) = XA*C(I) - 2.*(A(I)*B(I) + A(J)*B(J))
      C(J) = XA*C(J) - 2.*(A(J)*B(I) - A(I)*B(J))
      A(I+2) = A(I) + (B(I)*C(I) - B(J)*C(J))/D(K)
      A(J+2) = A(J) + (B(J)*C(I) + B(I)*C(J))/D(K)
   10 CONTINUE
      IF(N.LT.2) GO TO 30
      DO 20 L=3,M,2
      I = M - L + 1
      J = I + 1
      K = K - 1
      B(I+2) = B(I) + (A(I)*C(I) + A(J)*C(J))/D(K)
      B(J+2) = B(J) + (A(J)*C(I) - A(I)*C(J))/D(K)
   20 CONTINUE
   30 B(1) = A(1)
      B(2) = A(2)
      RETURN
      END
```

Figure A.4: Routine Implementing Standard Gradient Method

```
      SUBROUTINE MTI (A,N,M)
C
C  COMPUTES NORMALLIZED M.T.I. DIGITAL FILTER.
C  WRITTEN BY C. GIBSON, JAN. 1981.
C  COMPLEX INPUT IN A(1),A(2).
C  COMPLEX OUTPUT IN A(2N+1),A(2N+2).
C  N IS FILTER ORDER (MIN.=1, MAX.=10).
C  M=0 FOR NO NORMALLIZATION.
C
      REAL A(1),B(10),C(10)
      DATA B,C/20*0./
      I = 1
      B(I) = B(I) + A(1)
      C(I) = C(I) + A(2)
      X = 2.
   10 I = I + 1
      IF(I.GT.N) GO TO 20
      B(I) = B(I) + B(I-1)
      C(I) = C(I) + C(I-1)
      X = X + X
      GO TO 10
   20 J = N + N + 2
      A(J-1) = B(N)
      A(J) = C(N)
      IF(M.NE.1) GO TO 30
      A(J-1) = A(J-1)/X
      A(J) = A(J)/X
   30 I = I - 1
      IF(I.EQ.1) GO TO 40
      B(I) = -B(I-1)
      C(I) = -C(I-1)
      GO TO 30
   40 B(I) = -A(1)
      C(I) = -A(2)
      RETURN
      END
```

(a)

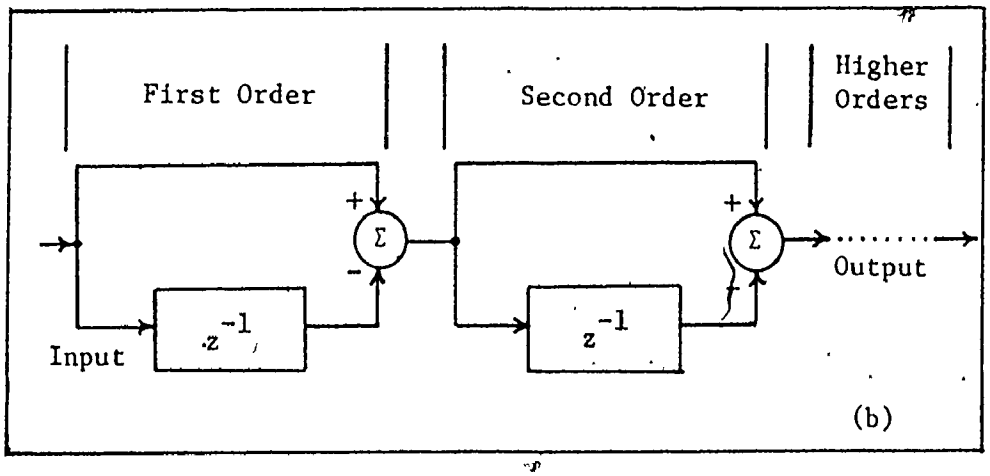

Figure A.5: Routine Implementing MTI Filter

(a) Printout, (b) Block Diagram

C array contains only the numerators of the reflection coefficients, and the D array contains the denominators). The other two arguments, N and XA, are the filter order (m) and the adaptive contstant (μ or ω), respectively.

These routines perform complex functions on complex values; however, they do so using real arithmetic and variables. This is done using two real variable locations in each linear array for each complex value. The two variables are located adjacently, with the real part first, followed by the imaginary part. This arrangement is similar to that used in most computers for storing complex arrays. Thus, for use in most programs using complex variables, only an EQUIVALENCE statement would normally be needed to relate the A, B, and C arrays to their complex counterparts. (Note: The D array contains only real values).

The B and C arrays must be dimensioned to 2 x N for the complex values $b_1(n)$ to $b_m(n)$ and $\rho_1(n)$ to $\rho_m(n)$, and the A array to 2 x N + 2 for the complex values $f_o(n)$ to $f_m(n)$ ($f_o(n) = b_o(n)$). This dimensioning takes place in the calling program, allowing the subroutine to take any filter order. There is no storage of these values within the lattice routines. The MTI routine, on the other hand, does store its intermediate values internally, thus limiting its filter order. So as to be interchangeable with the other routines, the MTI output is placed in the locations of the A array corresponding to the selected filter order.

APPENDIX B

KALMAN FILTERING THEORY APPLIED TO THE LATTICE STRUCTURE

B.1 Tapped-Delay-Line Prediction-Error Filter

Let us first review the application of Kalman filtering to the tapped-delay-line (TDL) structure. Figure 2.1(a) shows the basic structure of a tapped-delay-line prediction-error filter (PEF) of order m, operated in the forward direction. (The operator z is the unit-delay operator). The structure of Fig. 2.1(a) implements the equation

$$f_m(n) = d(n) - \underset{\sim}{u}^T(n)\ \underset{\sim}{a}(n)\ , \tag{B.1}$$

where $f_m(n)$ is the forward prediction error measured at the filter output, and d(n) is the desired response (namely, the input u(n), which is labelled x(n) in Fig. 2.1) for the filter to predict. The superscript T denotes transposition. The u(n) and a(n) are vectors of previous input samples and filter tap-weights, respectively, as shown by

$$\underset{\sim}{u}(n) = \begin{bmatrix} u(n-1) \\ u(n-2) \\ \vdots \\ u(n-m) \end{bmatrix} \quad , \quad \underset{\sim}{a}(n) = \begin{bmatrix} a_1(n) \\ a_2(n) \\ \vdots \\ a_m(n) \end{bmatrix} \tag{B.2}$$

We assume that both u(n) and a(n) are complex-valued.

In an adaptive prediction filter, the tap-weight vector $\underset{\sim}{a}(n)$ is continuously updated to minimize the error output $f_m(n)$ by best predicting the input data stream. The Kalman filtering implementation of this adaptive process is accomplished through the series of identifications given in Table B.1. Using the formulas for the complex form of Kalman filtering theory [25, 45], the PEF can then be computed, for the steady-state case, through the recursive application of the following procedure (the asterisk denotes complex conjugation):

(1) Compute the error output $f_m(n)$ using eq. (B.1) ;

(2) Compute the Kalman gain $\underset{\sim}{G}(n)$:

$$\underset{\sim}{G}(n) = \underset{\sim}{P}(n)\ \underset{\sim}{u}^*(n)\ [\underset{\sim}{u}^T(n)\ \underset{\sim}{P}(n)\ \underset{\sim}{u}^*(n) + \varepsilon_{min}]^{-1} , \qquad (B.3)$$

where ε_{min} is the minimum value of the mean-square error and $\underset{\sim}{P}(n)$ is the error-covariance matrix;

(3) Compute the new estimate of $\underset{\sim}{a}(n)$:

$$\underset{\sim}{a}(n+1) = \underset{\sim}{a}(n) + \underset{\sim}{G}(n)\ f_m(n) ; \qquad (B.4)$$

(4) Compute the new estimate of the error covariance matrix $\underset{\sim}{P}(n)$:

$$\underset{\sim}{P}(n+1) = \underset{\sim}{P}(n) - \underset{\sim}{G}(n)\ \underset{\sim}{u}^T(n)\ \underset{\sim}{P}(n) ; \qquad (B.5)$$

(5) Increase n by one, and repeat from step (1).

The initial values of $\underset{\sim}{a}(0)$ and $\underset{\sim}{P}(0)$ are normally set to zero except for the diagonal elements of $\underset{\sim}{P}(0)$ which are set to an estimate of the average value of the optimum tap-weight magnitudes (e.g.: 0.75). The minimum value of mean-square error ε_{min} is normally chosen between 0.01

Table B.1: Kalman-TDL Identifications

Kalman Filter		Tapped Delay Line PEF	
State vector	$\underset{\sim}{x}(n)$	$\underset{\sim}{a}(n)$	
State transition matrix	$\underset{\sim}{\phi}(n+1, n)$	$\underset{\sim}{I}$	(identity matrix)
State noise vector	$\underset{\sim}{v}(n)$	$\underset{\sim}{0}$	(null vector)
Observed data vector	$\underset{\sim}{y}(n)$	$d(n)$	$(=u(n))$
Measurement matrix	$\underset{\sim}{C}(n)$	$\underset{\sim}{u}^T(n)$	
Measurement noise vector	$\underset{\sim}{w}(n)$	$f_m(n)$	

Table B.2: Kalman-Lattice Identifications

Kalman Filter	Lattice Filter, to Minimize -		
	$f_m(n)$	$b_m(n)$	both
State vector $\underset{\sim}{x}(n)$	$-\rho_m(n)$	$-\rho_m(n)$	$-\rho_m(n)$
State transition matrix $\underset{\sim}{\phi}(n+1, n)$	$\underset{\sim}{I}$	$\underset{\sim}{I}$	$\underset{\sim}{I}$
State noise vector $\underset{\sim}{v}(n)$	$\underset{\sim}{0}$	$\underset{\sim}{0}$	$\underset{\sim}{0}$
Observed data vector $\underset{\sim}{y}(n)$	$f_{m-1}(n)$	$b^*_{m-1}(n-1)$	$\begin{bmatrix} f_{m-1}(n) \\ b^*_{m-1}(n-1) \end{bmatrix}$
Measurement matrix $\underset{\sim}{C}(n)$	$b_{m-1}(n-1)$	$f^*_{m-1}(n)$	$\begin{bmatrix} b_{m-1}(n-1) \\ f^*_{m-1}(n) \end{bmatrix}$
Measurement noise vector $\underset{\sim}{w}(n)$	$f_m(n)$	$b^*_m(n)$	$\begin{bmatrix} f_m(n) \\ b^*_m(n) \end{bmatrix}$

and 0.001. These initial values have little effect beyond the first few iterations of the filter [46].

B.2 Lattice-Structure Prediction-Error Filter

An intrinsic property of the lattice structure PEF is that the output error signal of each successive stage is orthogonal to every other stage [4, 15]. This allows the adaptation of each stage to be optimized independently of each other, rather than the global optimization applied to the tapped-delay-line PEF. Thus, in applying the Kalman filtering theory to the lattice structure, only a single stage of the filter need be considered.

The adaptive operation of the lattice structure consists of continuously updating the reflection coefficients $\rho_m(n)$ to minimize the prediction errors $f_m(n)$ and $b_m(n)$ in some form. The Kalman filtering implementation of this process depends on the identifications given in Table B.2. (Note that in deriving these identifications, the complex conjugate of equation (2.2) was used).

Using ϵ_{min} as the minimum mean-square error, for the steady-state case, the correlation matrix $\underset{\sim}{R}(n)$ of the measurement noise vector takes the value:

$$\underset{\sim}{R}(n) = \begin{bmatrix} \epsilon_{min} & 0 \\ 0 & \epsilon_{min} \end{bmatrix} \tag{B.6}$$

The Kalman gain can then be computed as [25, 45]

$$\underset{\sim}{G}(n) = \underset{\sim}{P}(n)\ \underset{\sim}{C}^H(n)\ [\underset{\sim}{C}(n)\ \underset{\sim}{P}(n)\ \underset{\sim}{C}^H(n) + \underset{\sim}{R}(n)]^{-1}\ , \tag{B.7}$$

where the superscript H denotes Hermitian transposition, and the error covariance P(n) is a scalar. Define

$$\alpha_m(n) = \frac{\varepsilon_{min}}{P(n)} \qquad (B.8)$$

Then, substituting (B.6) and the definition of the measurement matrix $\underset{\sim}{C}$ which minimizes both $f_m(n)$ and $b_m(n)$ into eq. (B.7), we get the Kalman gain

$$\underset{\sim}{G}(n) = \frac{[b^*_{m-1}(n-1),\ f_{m-1}(n)]}{\alpha_m(n) + |f_{m-1}(n)^2| + |b_{m-1}(n-1)|^2} \qquad (B.9)$$

Next, the state vector is given by [25, 45]

$$\underset{\sim}{x}(n+1) = \underset{\sim}{x}(n) + \underset{\sim}{G}(n)\ [\underset{\sim}{y}(n) - \underset{\sim}{C}(n)\ x(n)] \qquad (B.10)$$

Substituting (B.9) and the definitions of the state vector x(n), the observed data vector $\underset{\sim}{y}(n)$ and the measurement matrix $\underset{\sim}{C}$ (see Table B.2) in eq. (B.10), we get

$$\rho_m(n+1) = \frac{\alpha_m(n)\ \rho_m(n) - 2\ f_{m-1}(n)\ b^*_{m-1}(n-1)}{\alpha_m(n) + |f_{m-1}(n)|^2 + |b_{m-1}(n-1)|^2} \qquad (B.11)$$

The new estimate of the error covariance P(n) is calculated as

$$P(n+1) = P(n) - \underset{\sim}{G}(n)\ \underset{\sim}{C}(n)\ P(n)$$

$$= P(n)\ \frac{\alpha_m(n)}{\alpha_m(n) + |f_{m-1}(n)|^2 + |b_{m-1}(n-1)|^2} \qquad (B.12)$$

Using this result, the new value of $\alpha_m(n)$ becomes:

$$\alpha_m(n+1) = \varepsilon_{min} P^{-1}(n+1)$$

$$= \alpha_m(n) + |f_{m-1}(n)|^2 + |b_{m-1}(n-1)|^2 \qquad \text{(B.13)}$$

The recursive procedure for calculating the filter reflection coefficients thus becomes the calculations of eq. (B.11) and (B.13) in sequence for each filter order m, incrementing the time index n, and repeating the cycle. The initial values of $P_m(0)$ and ε_{min} (and thus $\alpha_m(0)$) are normally zero.

This procedure is mathematically identical to the standard gradient method implementation, as formulated in section 3.3, but with $\mu = 1$. As a result, these calculations apply only for the steady-state (stationary) case. They can, however, be extended to the non-stationary case by the use of the adaptive constant μ. As a consequence of choosing the Kalman-lattice identifications to minimize both forward and backward prediction errors, the harmonic-mean algorithm has been implemented. If either of the other two identifications of Table B.2 were used, either the forward algorithm or the backward algorithm would be implemented in the same way.

APPENDIX C

RADAR INSTALLATION AND RECORDING SYSTEM

C.1 Radar Transmitter and Antenna

This appendix describes the more important features of the ASR-8K Air-Surveillance Radar installed at the Bagotville CFB, Quebec. Also described is the video recorder used to record data from this radar. A block diagram of the total system is shown in Fig. C.1. Details of the radar installation come from the system manual [39].

Table C.1 presents the pertinent details for the radar transmitter. In fact, two separate transmitters were available (Channel A and Channel B), using the same antenna and processing. The radar could be operated with A, B, or both transmitters operating (diversity mode). For this research, only the Channel A transmitter was in use. (The diversity mode interlaces the returns from both transmitters).

An important feature of this radar is the staggered PRF, where the pulse-to-pulse period takes one of four values, in sequence. This is done to eliminate blind speeds (using conventional MTI filtering) to velocities in excess of 2150 knots (~Mach 3). With a constant PRF, these blind speeds would occur at multiples of 109 knots.

The radar antenna is described in Table C.2. This antenna has two feed horns resulting in two beams, each having approximately the same pattern. One of these beams, known as the active beam, is used for

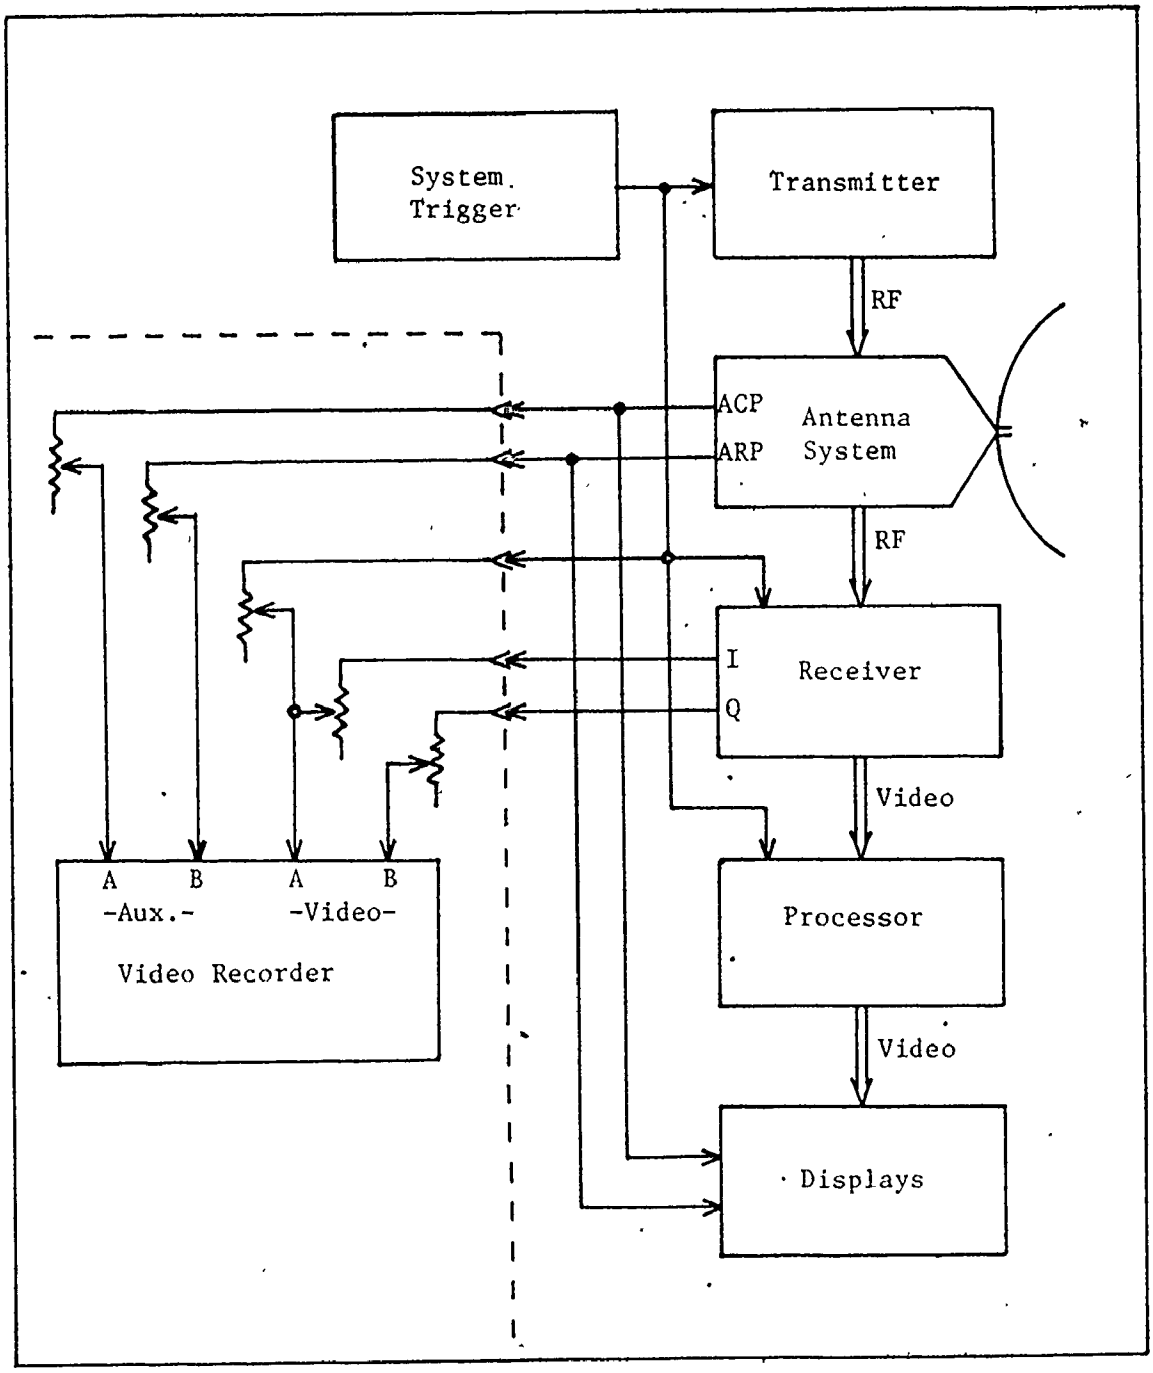

Figure C.1: Radar System Block Diagram

both transmitting and receiving. The other, passive beam is for receiving only, and has its pattern elevated 3.5^{o} with respect to the active beam, so as to reduce ground clutter returns. The radar receiver switches between the two beams by way of a programmed range-azimuth gate (RAG). The passive (high) beam is used close to the radar site, while the active (low) beam is used further out.

C.2 Radar Receiver

Figure C.2 is a block diagram of the radar receiver. Of particular note are the PIN diode RF attenuators in the input paths from the antenna horns. These allow the operator to remotely control both receiver sensivitity (one of six preset levels), and STC (sensitivity-time-constant, one of three preset curves plus off) for each of the antenna signals.

Table C.1: Transmitter Characteristics

Type - S-band Klystron (coherent, pulse-to-pulse)

Frequency - 2786 MHz (Channel A), 2714 MHz (Channel B)

Power - 1.08 MW (Channel A, VSWR = 1.13:1), 1.10 MW (Channel B)

Pulse width - 0.6 μsec.

Resulting range resolution - ~90 meters

Pulse repetition frequency (PRF) - Average of 1041 Hz (staggered)

Interpulse periods - 871, 961, 830, and 1177 μsec.

(960.8 μsec. average)

Resulting range limit (including allowance for recycling) - 61 naut. mi.

Table C.2: Antenna Installation Characteristics

Scanning antenna (clockwise in horizontal plane)

Rate of rotation - 12.5 rpm

Gain - Active beam: 33.5 dB

Passive beam: 32.5 dB

Vertical pattern - Cosecant over 30°

Horizontal beamwidth - 1.35° (-3 dB)

Pulses per beamwidth - 19

Polarization - Circular (horizontal also available)

Additional outputs -

ARP (Azimuth reference pulse): one per rotation, when antenna points north

ACP (Azimuth change pulse): 4096 per rotation, giving positional accuracy < 0.1°

RAG beam switching (in receiver) -

Azimuth	Passive Beam Range	Active Beam Range
90°-345°	0-15.1 naut. mi.	> 15.1 naut. mi.
345°-90°	0-24.8 naut. mi.	> 24.8 naut. mi.

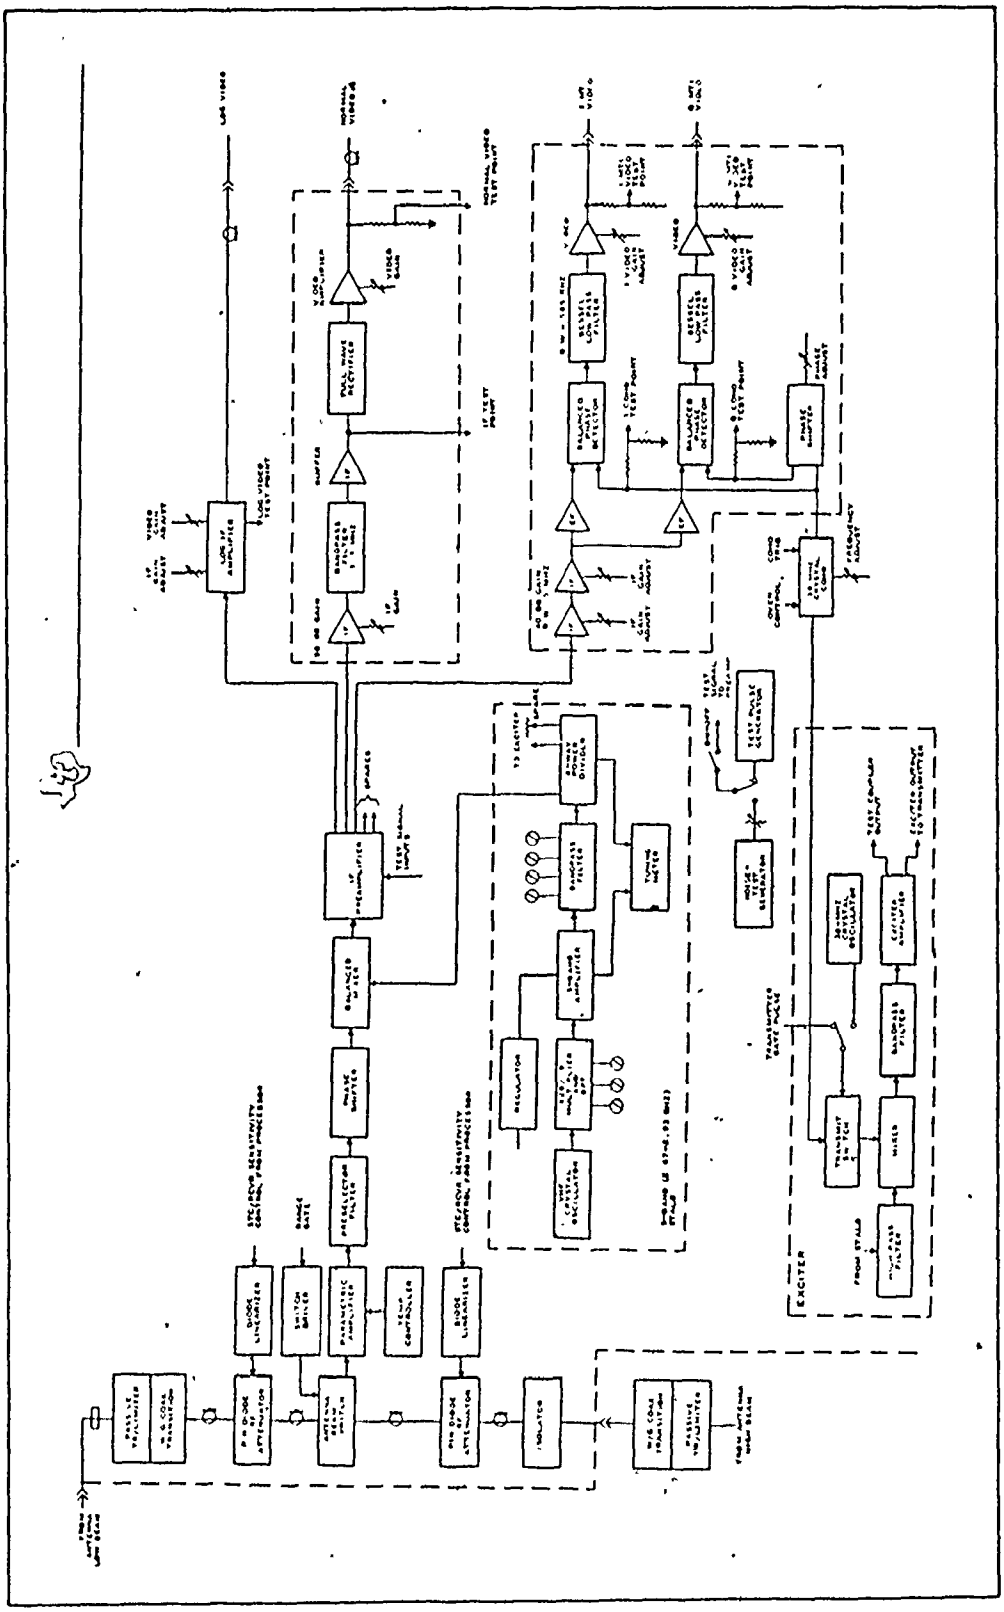

Figure C.2: Receiver Block Diagram

The STC attempts to correct for the decrease of signal amplitude with range by applying an exponential gain function to the return from each pulse. Most of the recordings used in this research were made with the STC off; however, some were made using the STC-1 response, which is shown in Fig. C.3.

After passing through the initial receiver stages, the IF signal is sent to three different detector blocks in the receiver. A non-coherent detector produces a normal video signal, which will be used in the normal display (mainly ground clutter). A logarithmic detector (similar to the normal detector) produces a video signal used for weather mapping. Two coherent detectors, one shifted from the other by 90^{o}, produce coherent inphase (I) and quadrature (Q) video signals, which are used by the processor to generate the MTI video display. These bipolar I and Q signals (un-processed) were recorded on the wide-band video-recorder and returned to the laboratory for analysis. (The normal video signal can be recreated from these I and Q signals).

C.3 Radar Processor

A block diagram of the processor is shown in Fig. C.4. Although the processor outputs were not used in this research (being very non-linear), a description of the features available is useful for comparison purposes. The four receiver signals described above are here processed into two signals, a normal video and an MTI video, which can then be displayed on the air-controller's PPI screen. The signals are processed digitally, with a sampling rate of 2.14 MHz.

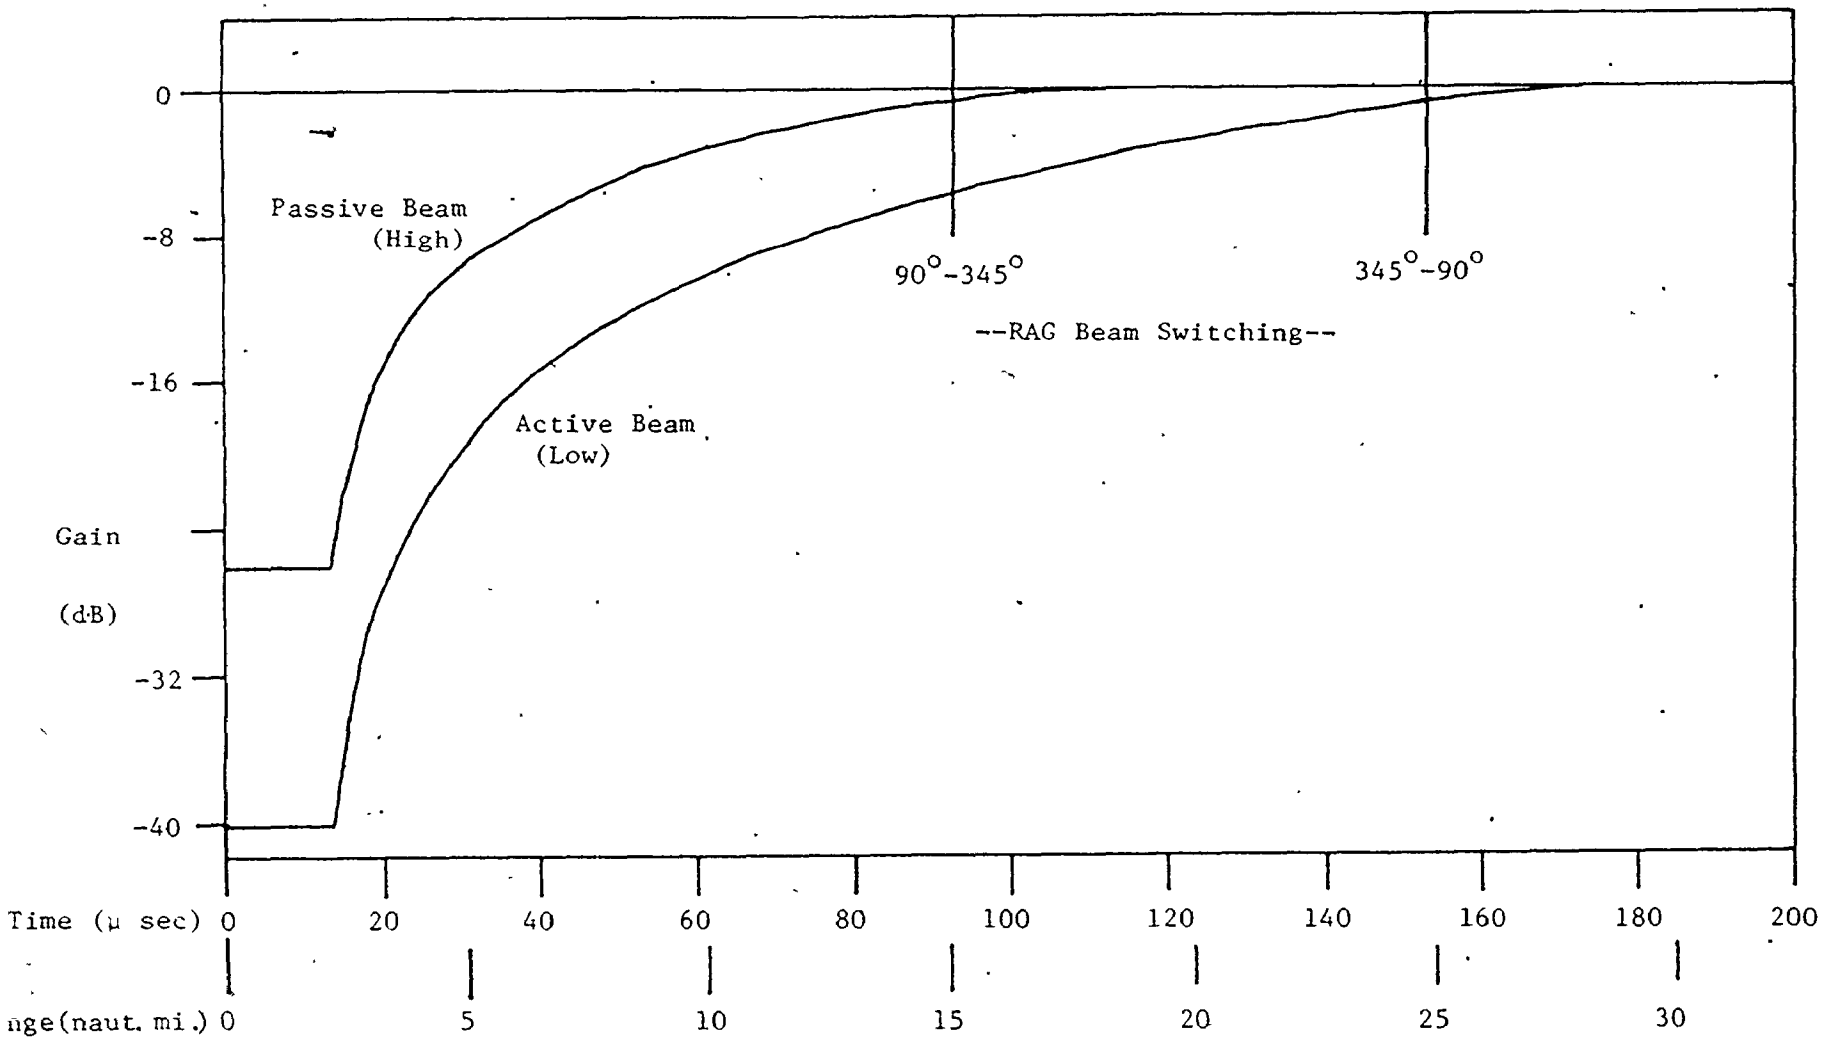

Figure C.3: STC-1 Response

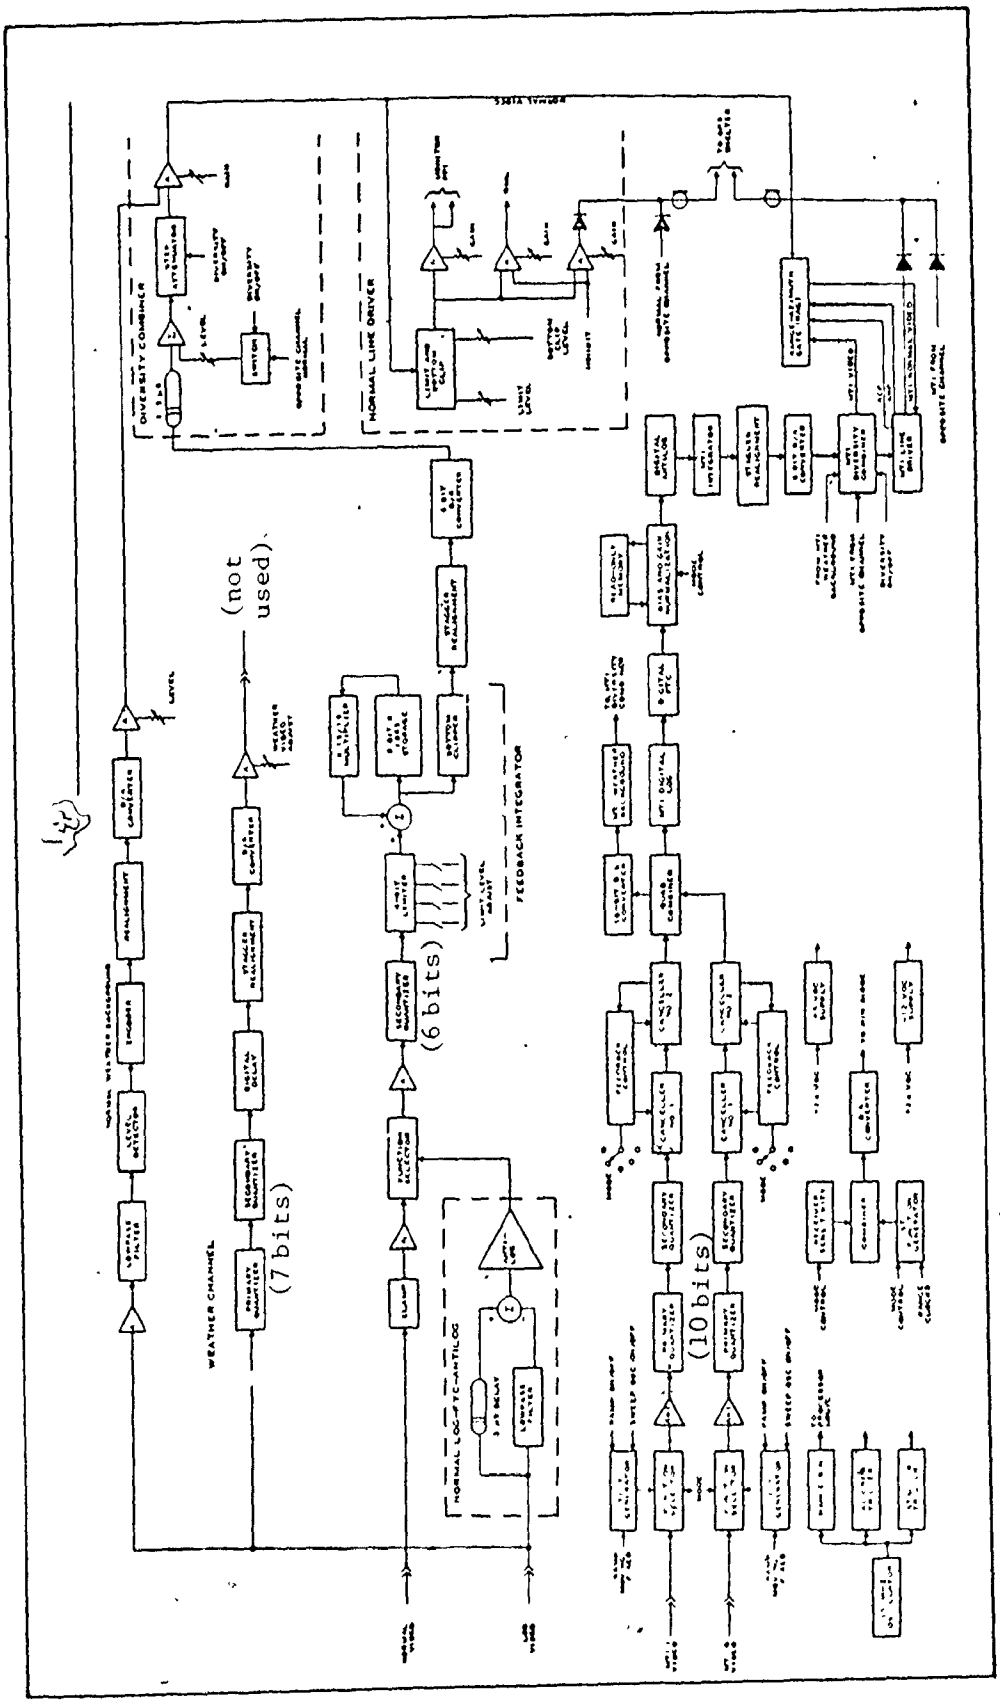

Figure C.4: Radar Processor Block Diagram

The I and Q signals go through separate MTI filters before being combined as one signal. The ASR-8 allows several modes of MTI filtering; in this installation, the mode selected was to cascade two single delay (two-pulse) cancellers with no feedback to form dual (3 pulse) cancellers (see Fig. A.5) having a typical sine-squared response. Each canceller stage has storage for 1665 range bins. After being combined, the MTI signal may go through any of three additional, operator-selected processing steps. These are: CFAR (constant false alarm rate), enhancement, and the addition of a weather outline. After these steps, the MTI signal goes through a range-azimuth gate which switches to the normal channel beyond 30-45 nautical miles, so as not to lose weak, outlying targets. (The theoretical system improvement factor for the MTI channel is 34 dB).

The normal channel also has the three operator-selected steps, which are basically identical to those of the MTI channel. The CFAR (also known as log-FTC-antilog) takes the logarithm of each range bin, subtracts the average of the log of the previous five range bins, and takes the anti-log to reduce signals covering large areas to the level of receiver noise. The enhancement mode uses a feedback integrator to smooth out the responses and reduce noise. The weather background uses a level detector and encoder to produce an overlay of weather systems. There are several weather modes available on the ASR-8; the one in use consisted of a multi-level outline with blanking, which is diagrammed in Fig. C.5.

C.4 Video Recorder

The radar data and reference signals were recorded on an RCA ADVISER-series wideband video recorder/reproducer. (ADVISER stands for Airborne Dual-channel Variable Input, Severe Environment Recorder/Reproducer). This tape deck allows about 45 minutes of recording time using 10½ inch reels of 2-inch wide tape, running at 15 inches per second. The unit has two wide-band video channels, which were used to record the I and Q video signals, and two 15 kHz bandwidth auxiliary channels, which were used to record the ACP and ARP signals. (The system trigger was mixed into the I video signal before recording).

Each video channel on the recorder had a frequency response of 10 Hz - 6 MHz ± 3 dB and a maximum SNR of 36 dB. This frequency range was more than adequate for these signals. The SNR, however, limited the effective dynamic range of the signals to 6 bits, when digitized for analysis. As the radar's MTI processor worked with a 10-bit dynamic range, it is to be expected that there would be some degradation of performance as a result (up to 18 dB). However, this did not greatly affect the validity of the data analysis in this thesis, which was mainly comparative in nature.

APPENDIX D

LABORATORY DATA TRANSFER SYSTEM

This appendix describes the experimental system which was used for laboratory analysis of the actual radar data. The data was recorded as described in Appendix C. The system described here recovers the data from the video recorder, displays it, samples selected range rings from the data, and transfers these samples into the computer for analysis. Figure D.1 is a block diagram of this system. The video signals played back by the video recorder (described in Appendix C) are split into two routes, one leading through wideband amplifiers to the PPI display, the other leading through low-pass filters (1 MHz) for data sampling. These low-pass filters feed an oscilloscope and two analog-to-digital (A/D) converters, which are controlled by and feed into the Interface and Control Unit (ICU). The I video signal is also fed into the System Trigger Detector, which strips the trigger from the video (previously combined as one signal during recording) for use by the ICU and PPI. (The trigger is shaped by a pulse generator for the PPI). The auxiliary channels from the recorder, containing the ACP and ARP radar signals, are fed into the Synchro Generator, which uses them to form the three-phase servo-motor signals required by the PPI for scanning. (In the PPI's normal environment, these signals would be supplied by a servo-generator attached to the antenna). The Synchro Generator

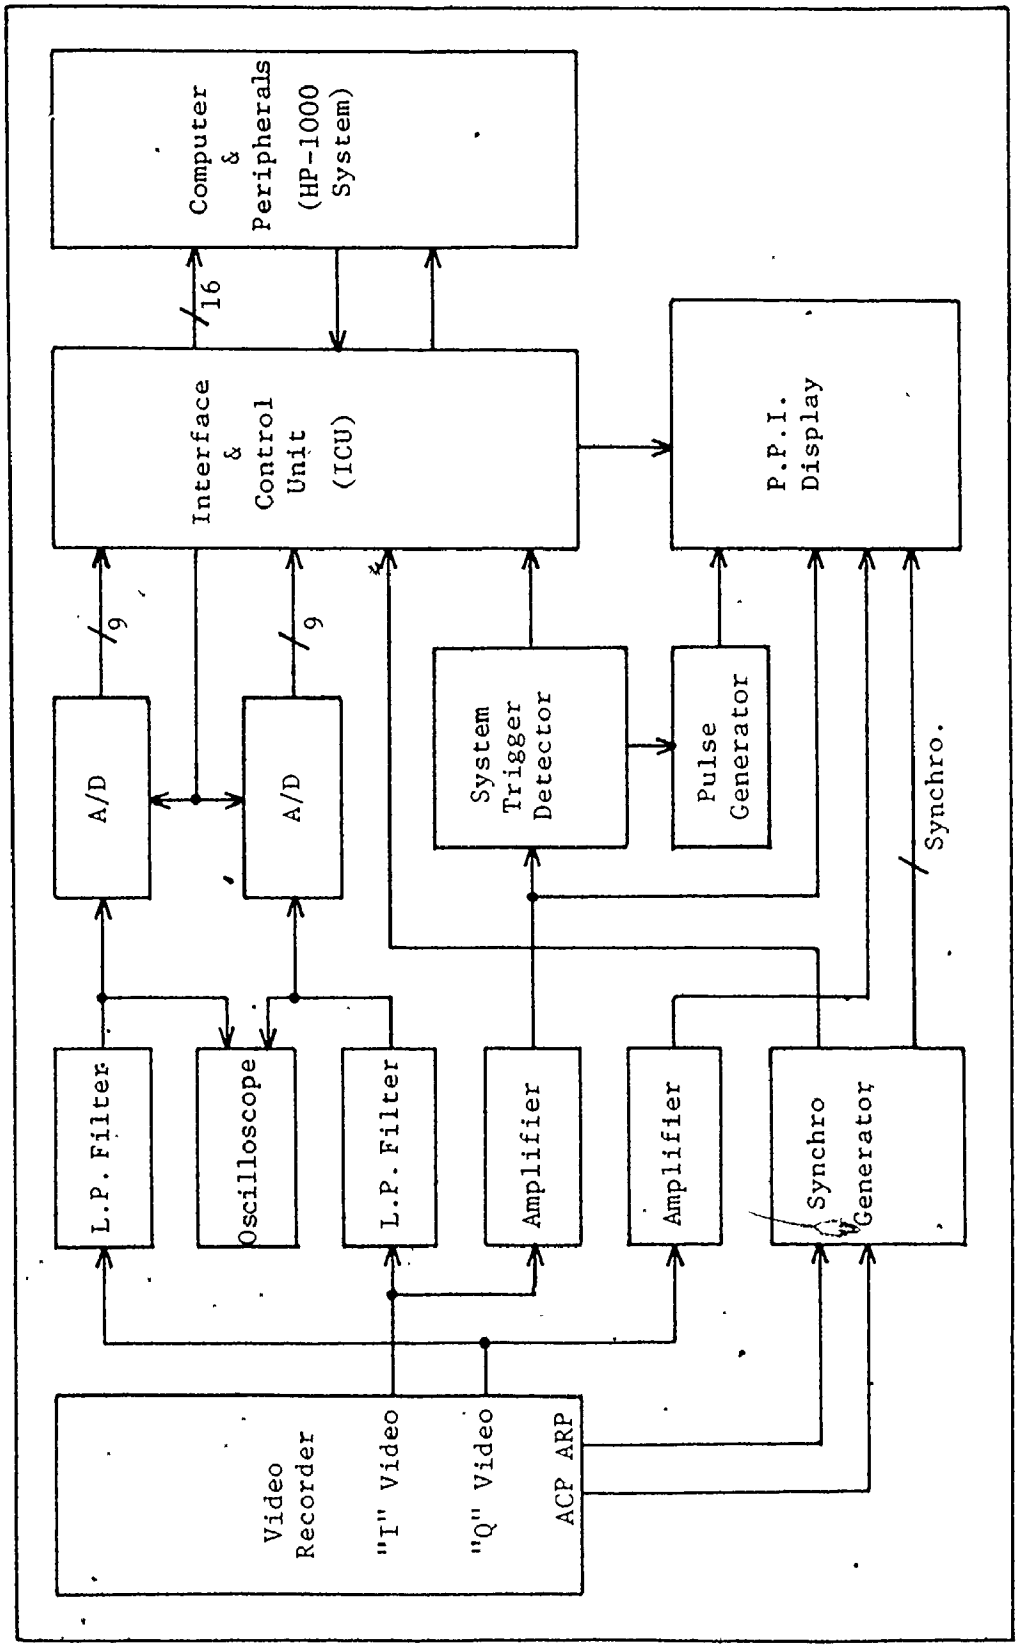

Figure D.1: Laboratory System Block Diagram

also supplies a shaped ARP pulse to the ICU.

The Synchro Generator unit was designed by Brian Currie, of the Communications Research Laboratory (CRL) at McMaster University, for this sort of application. The System Trigger Detector and the ICU were designed and built by the author of this thesis, specifically for the research herein. The A/D units, which feature 8 bits output and a conversion time under 200 nsec., were assembled by CRL staff using commercially available units. All the other devices shown are standard, general-purpose laboratory instruments, with the exception of the PPI, which was a military system on loan to the CRL.

Figure D.2 shows the oscilloscope traces of the I and Q video channels. The system trigger is the negative-going pulse at the beginning of the I channel return data. This data is preceded by a pulse burst (actually part of the waveform for the preceding return), which is used to reset the radar system. This burst of pulses, and strong negative peaks in the data both make the detection of the system trigger difficult at times. This is solved by the System Trigger Detector circuit diagrammed in Fig. D.3. This circuit works by detecting the pulse burst and locking out for a period to cover the burst, and then detecting the next negative signal (the system trigger). After this, the circuit locks out any further detections for the remainder of the return period.

Finally, the ICU is presented in Fig. D.4 and Table D.1. In addition to what is shown, this unit also contains line drivers for the 16 bits of data from the A/D's, drivers for the LED indicators corresponding to the numbered boxes in the figure, and a power supply. These

(a)

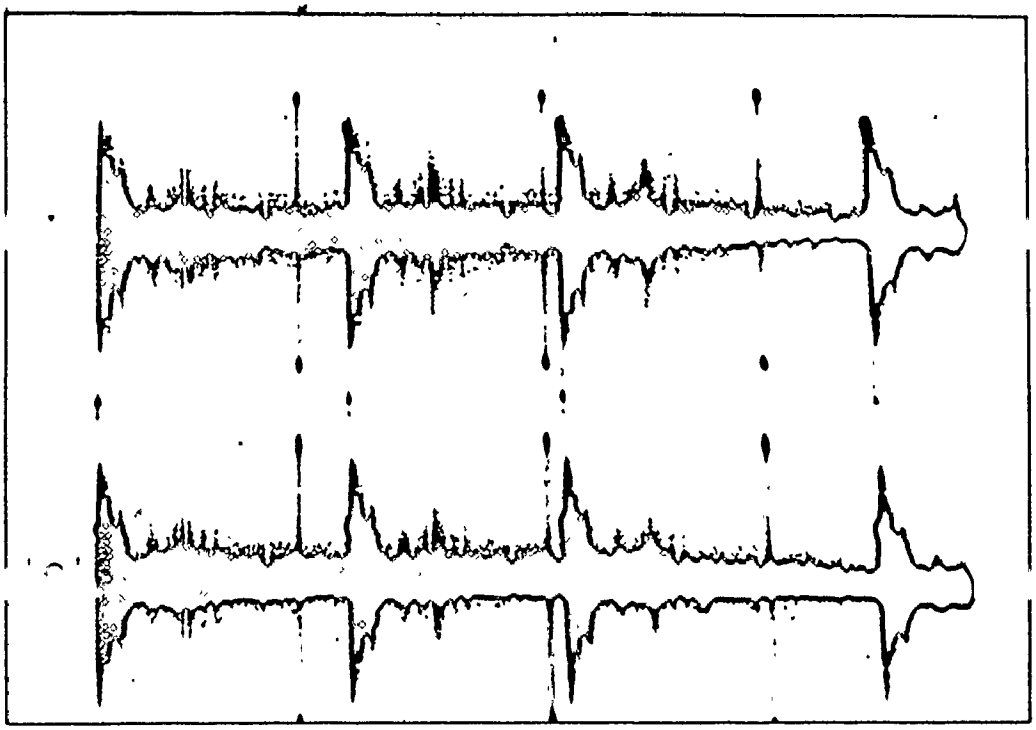

(b)

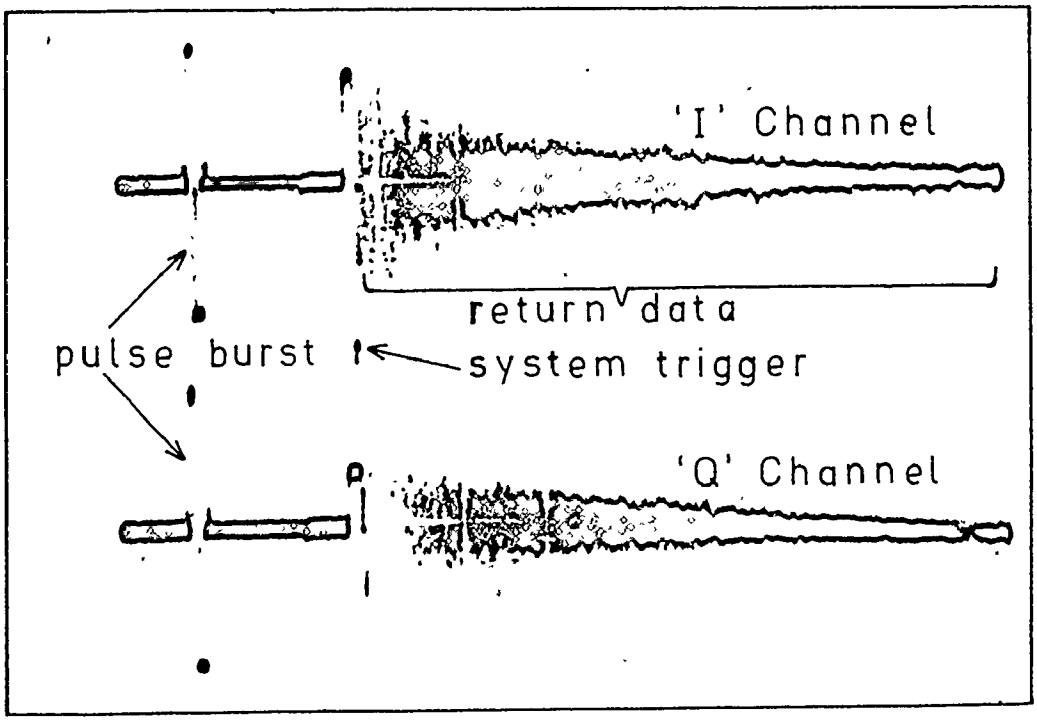

Figure D.2: Oscilloscope Displays of I & Q Channels

(a) 4 Pulses (showing staggered PRF)
(b) Single Pulse

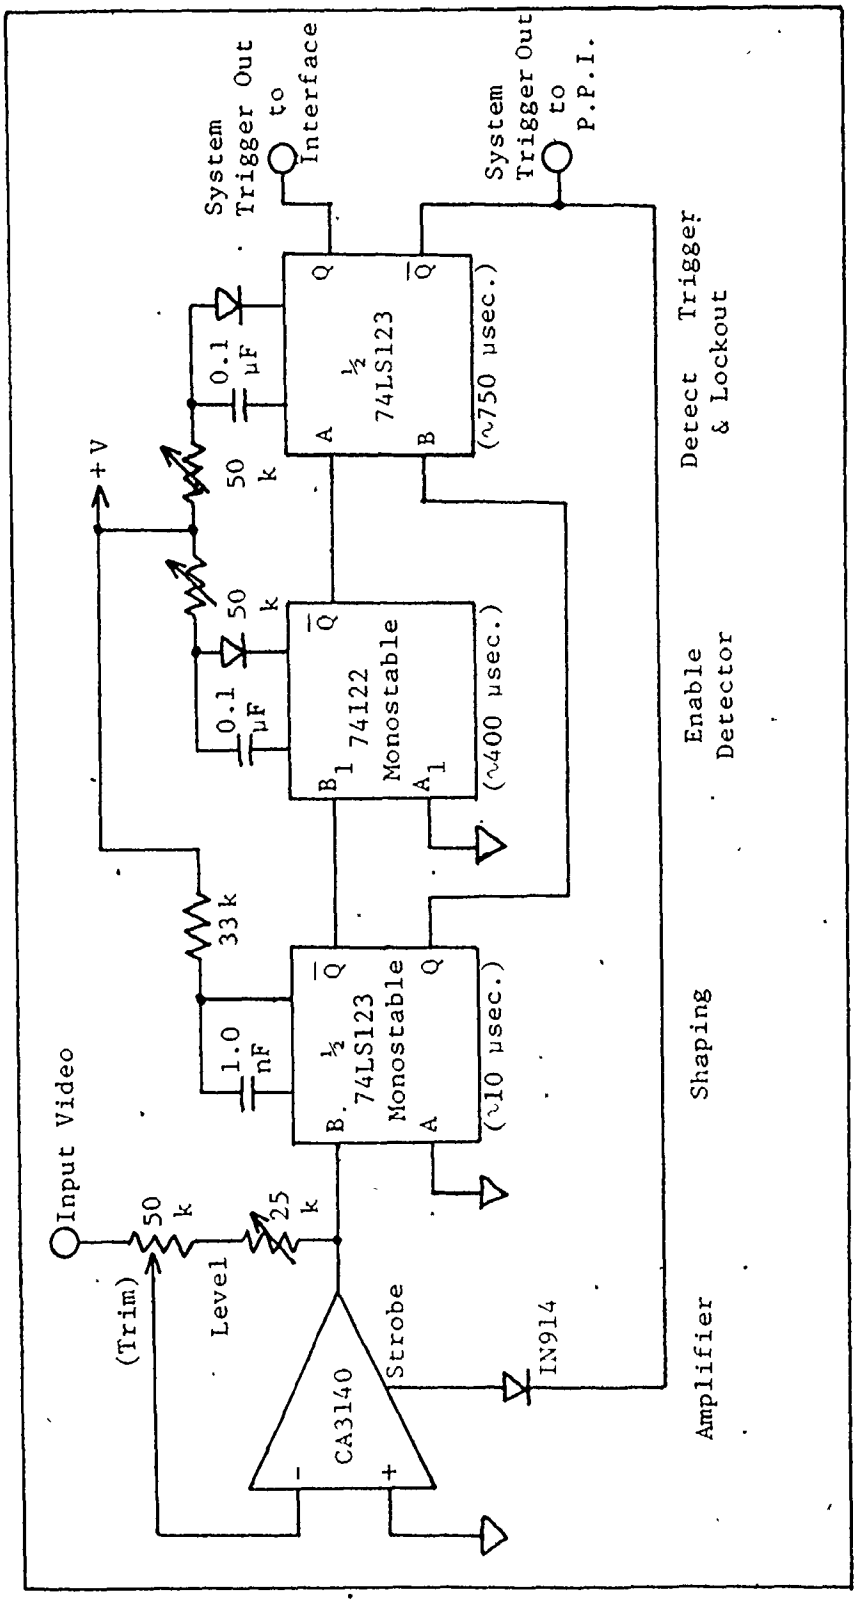

Figure D.3: System Trigger Detector

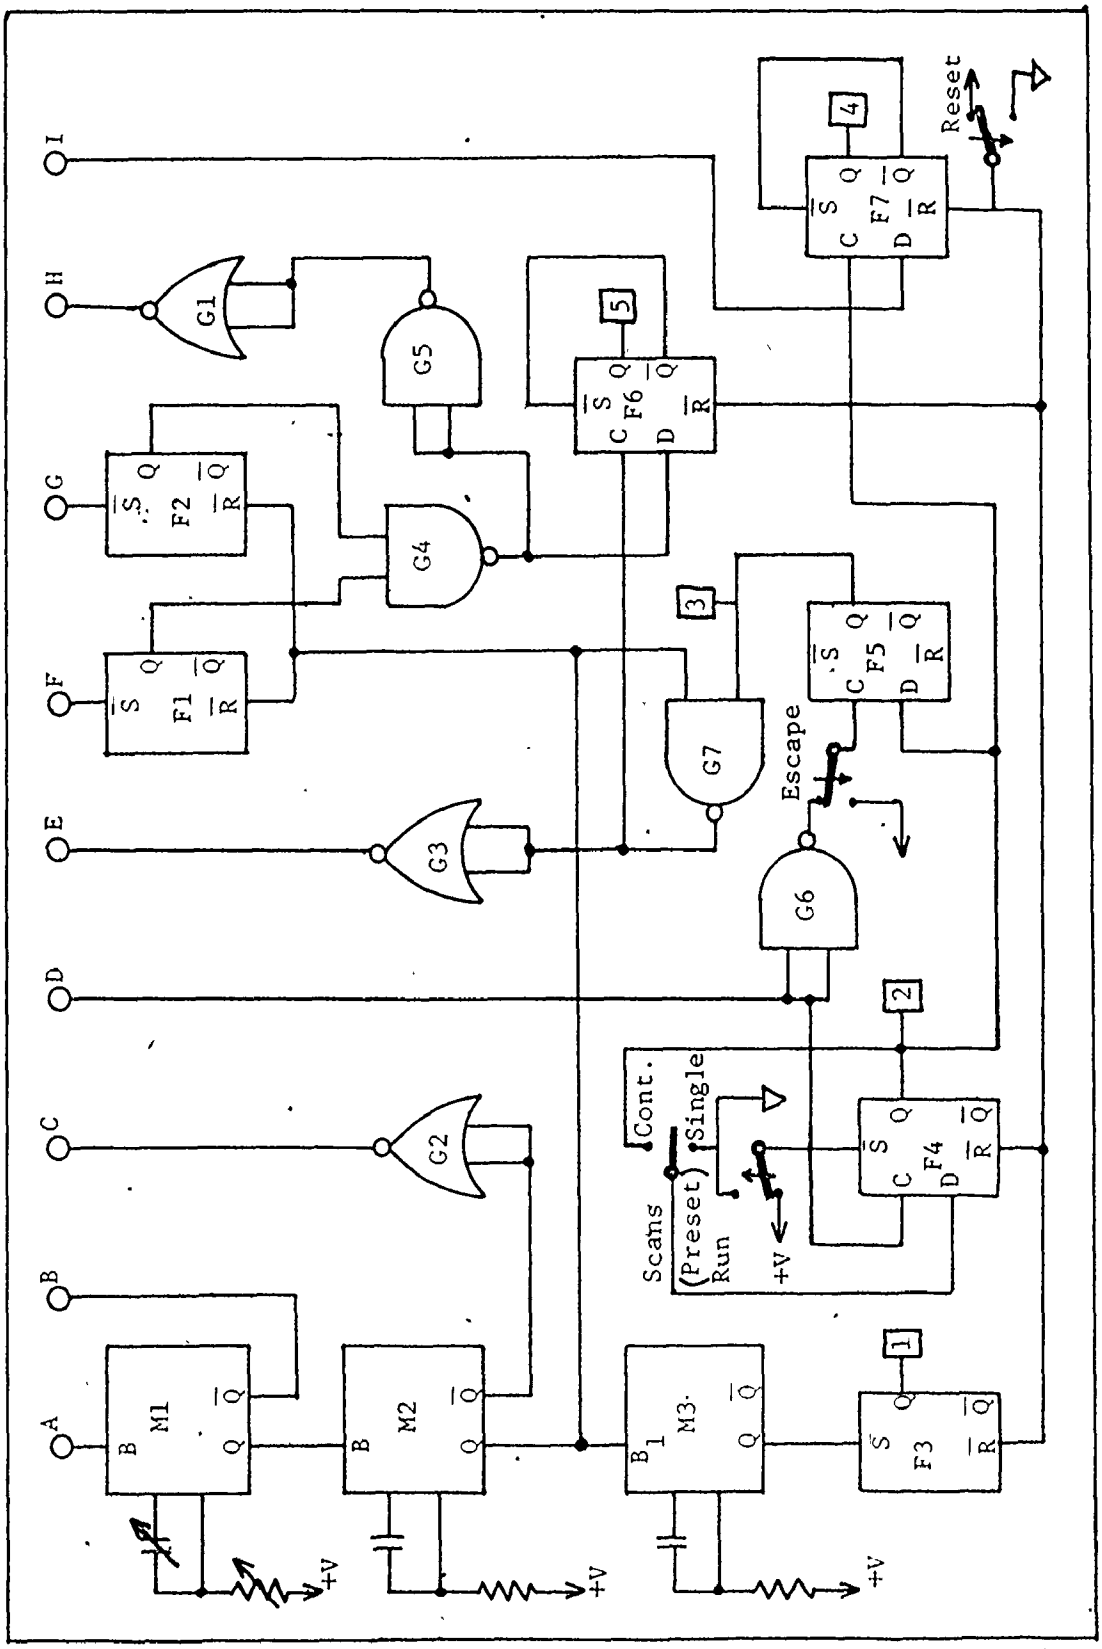

Figure D.4: Interface and Control Unit (ICU)

Table D.1: Interface and Control Unit Key

Label	Meaning
A	System trigger input (from system trigger detector)
B	Range delay output (for time measurement)
C	Range mark output to PPI
D	ARP input (from synchro generator)
E	"Encode" output to A/D's
F	"Ready" input from A/D #1
G	"Ready" input from A/D #2
H	Flag output to computer (data ready)
I	Flag input from computer (computer ready)
1	System trigger fault (L.E.D.)
2	Preset to start transfer (L.E.D.)
3	Running transfer (L.E.D.)
4	Computer fault (L.E.D.)
5	A/D fault (L.E.D.)

line and LED drivers, and the line driving gates G1 to G3 in Fig. D.4, are all high level, 50 ohm output NOR gates designed for line driving. The ICU serves two main functions. These are the control of the range-ring selection and sampling, and the indication of system faults. The system trigger is fed into variable monostable M1, which delays for a period equal to the desired range, and then triggers monostable M2 to give a short pulse. This pulse is sent to the PPI to form the range-ring display, to reset the "A/D ready" flip-flops F1 and F2, and to the A/D encode lines if a "run" signal is present. If M2 does not fire within a prescribed time, M3 and F3 indicate a fault.

The north mark (ARP) is fed into a circuit formed by F4, F5, and G6, which starts the data transfer at the next north mark in the data (when preset) and ends the transfer on the following north mark (unless set to run continuously). When both A/D's have completed their conversions, a signal is sent to the computer (by circuit F1, F2, G1, G4 and G5) saying that the data is ready to be picked up. If an A/D fails to signal ready in a reasonable time, or if the computer is not ready when the data is, a fault is indicated (by flip-flops F6 and F7, respectively). The data is transferred into the computer as one 16-bit integer word, which is later broken down into I and Q parts.