ROBERT E. THURBER

ADVANCED SIGNAL PROCESSING TECHNIQUES FOR THE DETECTION OF SURFACE TARGETS

Recent advances in the technology of signal processing components and subsystems have made possible the implementation of complex signal processing tasks in a simple, cost-effective manner. This has led to novel techniques for processing surveillance radar data in real time to extract weak targets in clutter and interference environments.

INTRODUCTION

The primary functions of surveillance radars are to recognize the presence of targets of interest and to report their positions. In the simplest form, this is accomplished by transmitting pulsed electromagnetic energy in a given direction and listening for an echo that is reflected from a target. The time delay to the received echo determines the distance to the target, and the direction of transmission is a measurement of its bearing. A volume of space can be surveyed by sequentially transmitting and listening in different directions until all possible target locations are covered.

A common type of surveillance radar has an antenna that produces a pattern that is narrow in azimuth (bearing) and continuously rotates through 360° at a constant rate of 1 to 10 seconds. The antenna pattern (beamwidth) in elevation determines the height of targets that can be observed, and the listening time between the short transmissions determines the maximum range at which targets can be observed. This radar, referred to as a two-dimensional (2D) range and bearing radar, is used to search for everything from air and surface targets to weather systems from land, sea, and airborne platforms.

One of the earliest ways to process data from 2D radars was to convert the received energy to a voltage that modulated the intensity of a cathode ray tube. The distance from the center of the tube represented the range to a target, and the bearing was determined by the position around the tube face. The presence of a target and its position in range and bearing were determined by a person observing the display.

It is relatively easy for an operator to distinguish between a return from an air or surface contact and background thermal noise. This method of signal processing works well in the absence of interfering signals and when only a small number of true targets are being processed. Unfortunately, this is not generally the case, and over the past 40 years much work has been done to enhance the returns from desired targets, automate the process of detecting and locating targets, and reduce the probability of making false target declarations.

This work has led to systems that perform almost flawlessly in their roles of air defense and air traffic control, but they are expensive to produce and maintain. There are, however, many types of surveillance radar systems not related to air safety or defense that rely solely on a display and an operator for signal processing. They still suffer from limitations imposed by the environment. A very important class of radar in this category is sea surface surveillance systems. They perform such important functions as shipboard surface surveillance, coastal surveillance, and ship collision avoidance.

Sea surface surveillance radars are designed to operate reliably for long periods of time, but their performance is limited by the environment. Signals returned from the sea tend to mask signals from true targets; an operator must distinguish between clutter and targets, relying only on the signal's intensity and character as presented on the display. Because more and more demands are being placed on sea surface surveillance radars to perform in such vital roles as ship collision avoidance, there is a need to improve their target detection capability in the presence of clutter and to automate the detection process.

APL sponsored a project to apply the latest advances in signal processing components and subsystems to the problem of producing a simple, low-cost radar signal processor that could significantly enhance the performance of surveillance radars. By using high-speed digital conversion and processing circuits and a large quantity of solid-state memory, several novel techniques for processing surveillance radar data were realized. They enhanced the performance of sea surface surveillance radars and are proving to have wider applications in processing radar signals and other high-speed data in real time.

Our work resulted in a project for the Pacific Missile Test Center to instrument their sea surface surveillance radars with signal processors and to integrate the data from noncollocated radars to form a data base of composite surface target tracks. Two missile and gun test firing ranges are being instrumented. One is located off the coast of California north of Los Angeles, and the other is off the island of Kauai in Hawaii. These surface surveillance radars are a primary means of determining if the test ranges are clear of surface craft so that ordnance firing can take place.

SIGNAL PROCESSING

The processing of signals from surveillance radars can take many forms, but there are only three principal objectives. The first is to be able to distinguish the signal of a true target from noise, interference, electronic transmissions, or returns from nontargets. This objective is usually accomplished by filtering the radar data with a device that is matched to those characteristics of the desired target that differ from the characteristics of the competing interference. It has the effect of enhancing the target signals as opposed to the background interference.

The second objective deals with the process of declaring the presence of a target. When there are only a few targets in a reasonably clutter-free environment, the process can be carried out by an operator viewing a display scope; however, to obtain reliable and repeatable detection of targets in dense or cluttered environments, the process must be automated.

In order to present a reliable data base to the user of the surveillance radar data, there should be a minimum of false target declarations and missed targets. Thus, the third objective of surveillance radar signal processing is to reduce the probability of false targets being reported while endeavoring to maintain maximum target detectability.

TARGET SIGNAL ENHANCEMENT

The target signals at the output of a radar receiver are optimized by matching the receiver filter characteristics to the transmission waveform. For instance, if the transmission is a continuous train of rectangular pulses of RF energy, the receiver bandpass would be the Fourier transform of that waveform, or $(\sin x)/x$. This maximizes the signal versus background noise on a per transmission basis but does not enhance the desired target signal over returns from clutter such as the sea or weather.

To enhance the target signal further, the data can be observed for several consecutive transmissions and matched to a filter that looks for some true target characteristics, such as a phase change caused by the Doppler shift of a moving target. The more sophisticated air surveillance radars use a system in which the amplitude and phase of the data for every possible target location for several radar transmissions are stored and filtered to extract only those signals that exhibit a phase change representative of a moving target. This is commonly referred to as a moving target indicator. Implementing the system requires a very stable analog delay line filter or a high-speed digital filter and a radar with a very stable RF system so that the data phase change over several transmissions can be measured very accurately (to 1° or less).

In general, this type of system would not be cost effective in a sea surface surveillance radar because of the expense of building a high stability radar and the difficulty in measuring the Doppler shift on very slowly moving targets. However, with the advent of high-speed large-scale digital integrated circuits and low-cost high-speed solid-state memory, a digital filter that operates on amplitude-only radar data is very attractive. To exploit the possible inclusion of these advanced circuit components, several methods of enhancing the target signal-to-clutter ratio using only received signal amplitude were investigated.

One way that has been considered to improve the signal-to-clutter ratio using only radar amplitude data relies on the moving target Doppler effect. If a moving target is in the same location (same radar resolution cell) as clutter, the target voltage vectorially adds to the clutter voltage. This causes the data from that cell to be amplitude modulated at the Doppler frequency of the target-minus-clutter velocity. Therefore, the amplitude data can be filtered to enhance targets that are moving if they are in the same location as clutter returns.

This process is ideally suited to sea surface surveillance radars because it does not require phase data, which are generally unobtainable from these simple radars, and because the returns from the sea almost always guarantee the presence of clutter. In general, however, sea surface surveillance radars are high resolution (they transmit a short pulse with a narrow beamwidth) in order to reduce the area of the sea illuminated per pulse and hence the amplitude of the sea clutter return. This means that the number of radar cells of data that must be stored in a filter is large, typically about 16,000 per transmission.

With today's technology, a digital transversal filter that stores four transmissions of data and operates at 10 megahertz is reasonably inexpensive and would easily handle the requirements of a high-resolution surveillance radar. The configuration chosen for the processor is shown in Fig. 1. The unit accepts logarithmically scaled analog data from the radar and uniformly quantizes them to form digital words that are stored in memory for three radar transmission periods. Therefore, at any given time, there are for every radar resolution cell four data points that can be weighted and summed to form the filtered output. Weighting is done by adding the logarithm of the multiplicative constant to the quantized logarithmic data; the resultant is linearized before the summing operation. The selection of the weights determines the response of the filter to the target data and the improvement in the target-to-clutter ratio. This is discussed more fully in the next section.

This method of filtering the data does nothing to reduce random interference signals originating from

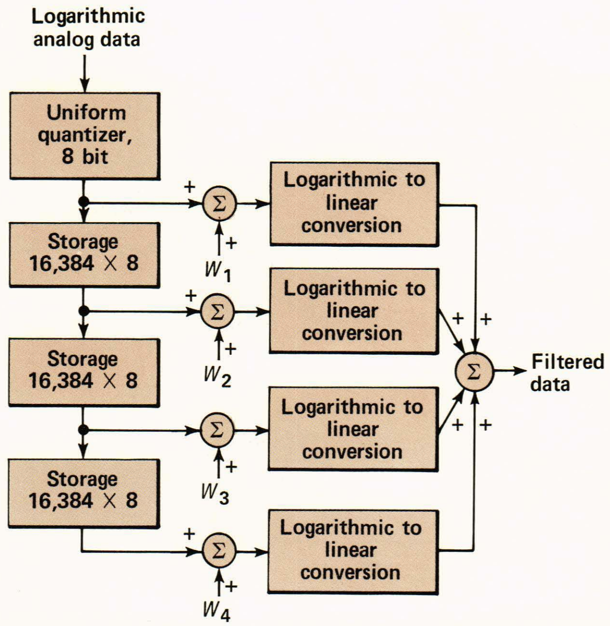

Figure 1 – Architecture of the digital transversal filter used to filter radar data at a 10 megahertz rate. The filter features logarithmic data input followed by uniform quantization. The filter weights are programmable so that the filter response can be matched precisely to the environment.

other radars and from electronic transmitting equipment that are constant sources of false target indications in most radar systems. In fact, signals that occur in a given radar cell for only one transmission interval can corrupt an otherwise valid response from a sequence of transmissions. Because of the flexibility in selecting weights for the transversal filter, the problem can be eliminated by detecting the interference and changing the filter weights to compensate. This is discussed further in the section on Suppressing Interference by Reprogramming the Moving Target Indicator Filter.

OPTIMIZED MOVING TARGET INDICATOR FOR THE DETECTION OF SLOW-MOVING TARGETS

When trying to enhance the target signal by observing the phase change from transmission to transmission, one can unambiguously determine a target's radial velocity up to the Nyquist rate by the following relation:

$$f_d = \frac{2V_r f_0}{c}, \tag{1}$$

where f_d is the Doppler frequency shift, f_0 is the transmitted frequency, V_r is the radial velocity of the target, and c is the velocity of propagation. In the case of slow-moving surface targets, the range of velocity that is of interest is less than the maximum unambiguous value. It is therefore desirable to tailor the filter weights to maximize the response over this range of velocity.

As the antenna scans across a target, the target return is amplitude modulated according to the antenna beamshape. This, in general, reduces the effectiveness of a moving target indicator filter; the weighting functions are therefore selected to minimize the effect. When the response for slow targets in clutter was maximized, it was found that the filter response resulting from antenna modulation became predominant. It was observed experimentally that the effect enhances the signal-to-sea-clutter ratio for slow-moving and stationary targets in sea clutter by as much as a factor of 3. The weights tend to differentiate a point target in distributed clutter, causing the target to stand out.

This result was not initially expected. Further analysis of the phenomenon resulted in the determination of weighting functions that optimized the enhancement of slow-moving targets in sea clutter on the basis of both Doppler and antenna modulation. The optimum weights turned out to be a function of target range. Although a compromise set of weights allows operation over all ranges with minimum performance degradation, the flexibility of the transversal filter implemented in the surface processor permits the separate selection of weights in four different range intervals. Nearly optimum operation is obtained by this configuration.

To optimize the performance, an interference improvement factor (IIF) is derived for a generalized square law detector. This improvement factor, which can be written as a ratio of two quadratic forms, incorporates the effects of antenna beamshape, interference composed of clutter plus noise, and the target Doppler frequency. It can be maximized to improve target detectability by eigenvalue/eigenvector techniques to yield the optimal weights, W_1, W_2, W_3, W_4, to be applied by the digital transversal filter.

It can be shown that the interference improvement factor is given by

$$\mathrm{IIF} = \frac{1}{\tfrac{1}{2}\sigma^2} \cdot \frac{\sum_{n=1}^{N}\sum_{m=1}^{N} W_n W_m \left[\left(\frac{a_n{}^2}{2\sigma^2}\right)\left(\frac{a_m{}^2}{2\sigma^2}\right)\right.}{\sum_{n=1}^{N}\sum_{m=1}^{N} W_n W_m (1+\rho_{nm}{}^2)} + \frac{\left.\frac{1}{2\sigma^2}(a_n{}^2 + a_m{}^2) + \frac{1}{\sigma^2} a_n a_m \rho_{nm} \cos(\phi_n - \phi_m)\right]}{\sum_{n=1}^{N}\sum_{m=1}^{N} W_n W_m (1+\rho_{nm}{}^2)}, \tag{2}$$

where

$$\rho_{nm} = \{R\rho_c[(n-m)T] + \delta_{nm}\}/(1+R),$$
$$R = \sigma_c^2/\sigma_o^2 = \text{clutter-to-noise power ratio},$$
$$\phi_n - \phi_m = W_d(n-m)T,$$

$$a_n = \text{signal amplitude} = \exp\left[-2.772588\left(\frac{\ddot{\theta} t_0}{\Theta_{3dB}} + [n - \tfrac{1}{2}(N+1)]\frac{\ddot{\theta} T}{\Theta_{3dB}}\right)^2\right] : n = 1, 2, \ldots, N,$$

$\ddot{\theta}$ = antenna rotation rate, degrees per second,
Θ_{3dB} = 3-decibel one-way antenna beamwidth,
t_0 = time offset of center of N pulse batch from beam peak,
$\rho_\vartheta(\tau) = \exp(-2\pi^2 \sigma_f^2 \tau^2)$,
σ_f = standard deviation of clutter spectrum, in hertz,
T = interpulse period of radar,
W_d = Doppler frequency in radians,
N = Number of pulses per batch,
$\frac{1}{2}\sigma^2$ = peak per pulse signal/interference ratio.

A Gaussian antenna beamshape and a Gaussian shaped clutter spectrum have been assumed in the above expressions.

The interference improvement factor can be expressed as a ratio of two quadratic forms,

$$IIF = \frac{\mathbf{W}^{\mathrm{T}} A \mathbf{W}}{\mathbf{W}^{\mathrm{T}} C \mathbf{W}} , \tag{3}$$

where A and C are real symmetric matrices and C is positive definite. $\mathbf{W}$ is a column vector corresponding to the weights W_1 . . . W_4 while $\mathbf{W}^{\mathrm{T}}$ is its transpose, a row vector. The interference improvement factor can be maximized by solving the following eigenvector equation:

$$C^{-1} A \hat{\mathbf{W}} = \lambda_1 \hat{\mathbf{W}}. \tag{4}$$

λ_1 is the maximum eigenvalue of the matrix $C^{-1}A$, and $\hat{\mathbf{W}}$ is its corresponding eigenvector. $\hat{\mathbf{W}}$ specifies the optimal weights; λ_1 is also the maximum value of the improvement factor.

Plot (a) in Fig. 2, the interference improvement factor versus radial velocity, has been optimized for target velocities of 0 to 10 knots at a range of 20 nautical miles. The radar parameters used are for the surface surveillance radars at the Point Mugu, Calif., Pacific Missile Test Center Range. They are S-band radars with a 0.1 microsecond pulse width at 550 pulses per second and an azimuth beamwidth of 0.9°. They are located 1500 feet above sea level. For comparison, (b) is the same plot using a binomial weighting function, as is normally the case for an air surveillance moving target indicator system.

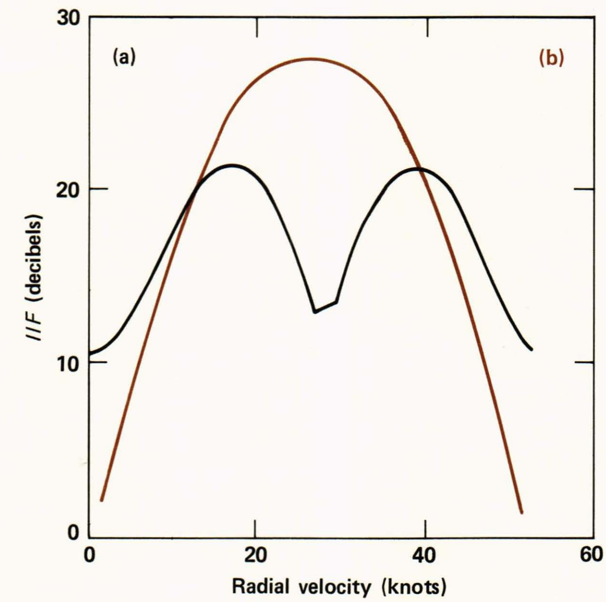

Figure 2 – Plots of the interference improvement factor versus target radial velocity for two sets of filter weights. Plot (a) is for a set of filter weights optimized for target velocities of 0 to 10 knots at 20 nautical miles. Plot (b) is for binomial weights generally used for an air surveillance moving target indicator system.

SUPPRESSING INTERFERENCE BY REPROGRAMMING THE MOVING TARGET INDICATOR FILTER

In most environments, RF signals originating from other equipment appear in the received data. In the case of a radar, the interfering signals most often come from other radars operating at or near the same transmission frequency. This causes pulses that can be falsely detected as a target. However, those pulses almost never appear at the same location (range cell) on consecutive transmissions because they are not synchronized to the radar transmissions. The occurrence of interference can be detected by comparing the data amplitude in a given range cell with the amplitude in the same range cell on previous radar transmissions. In the case of the transversal filter, there are always four consecutive transmissions of data available for comparison, and, if the two largest values in a set of four for a given cell differ significantly, it is very likely that the largest data point is an interference signal.

If allowed to remain in the data stream, this value would corrupt the filter output for the four transmissions where it is present in the stored data. In the past, systems have handled the problem by censoring the range cell at the filter output for the time the interference is present in the filter or by removing the interference from the data in the filter and substituting a value in its place based on true data in the vicinity of the interference. Censoring the range cell causes loss of data and can cause loss in performance in a heavy interference environment. Replacing the interference with "reasonable" data requires both calculation and insertion of the substitute data.

The design of the filter developed for the surface radar processor permits a somewhat different approach to solving the interference problem. The filter weights are additive constants that are selected from a read-only memory and summed with the logarithmic radar data. Eight different sets of weights are stored in memory; they can be selected on a range-cell-to-range-cell basis. When interference is detected in a range cell, the weight set that eliminates the interference and produces an optimum response for the remaining data points is selected. In this manner, all interference is eliminated from the filtered output, and the valid data are always filtered in an optimum fashion.

The maximization procedure discussed in the previous section can also be applied with reduced dimension matrices if there are missing pulses in the n pulse batch owing to interference or jamming. If the ith pulse is determined to be an interference pulse by comparison with the amplitudes of other pulses in the batch, then, by striking out the ith column element of $\mathbf{W}^T$, the ith row and column of matrices A and C, and the ith row element of $\mathbf{W}$, the optimization can still be accomplished. For instance, if the three-pulse case is reduced to a two-pulse case when the middle pulse is from interference, then

$$\mathbf{W}^T A \mathbf{W} = (W_1 \, W_2 \, W_3) \begin{pmatrix} A_{11} & A_{12} & A_{13} \\ A_{21} & A_{22} & A_{23} \\ A_{31} & A_{32} & A_{33} \end{pmatrix} \begin{pmatrix} W_1 \\ W_2 \\ W_3 \end{pmatrix} \quad (5)$$

becomes

$$(W_1 \, W_3) \begin{pmatrix} A_{11} & A_{13} \\ A_{31} & A_{33} \end{pmatrix} \begin{pmatrix} W_1 \\ W_3 \end{pmatrix} \quad (6)$$

Dimensions are similarly reduced for the quadratic form $\mathbf{W}^T C \mathbf{W}$. When Eq. 6 is used, we obtain the best weights for the two surviving pulses. Thus, it is not necessary to throw out the results for the whole batch.

Figure 3 illustrates the weights and how they are used for a normal target response and for one with interference. As the antenna scans across a target, the target-received amplitude is modulated, as shown in the top left illustration. The optimal weights determined for this case using the Pacific Missile Test Center's surface surveillance radar parameters are -0.98, $+1$, $+1$, -0.98; they are shown in the middle left illustration. The result of multiplying the target response by these weights is shown in the bottom left illustration; the filter output is the sum of the four weighted values shown.

The illustrations on the right are for the case where the third pulse is contaminated by interference. The optimal weights (-0.66, $+1$, 0, $+0.31$) and the weighted outputs are shown.

CONSTANT FALSE ALARM RATE PROCESSING

The automation of the detection process can take many forms, from simply setting a fixed threshold on the input signals to determining an adaptive detection threshold by measuring the mean and variance of signals returned from the local environment. In the ideal automatic target detector system, a constant false alarm rate (CFAR) is maintained, regardless of the environment. This, in turn, establishes the signal/background-noise ratio required for target detection in the given environment.

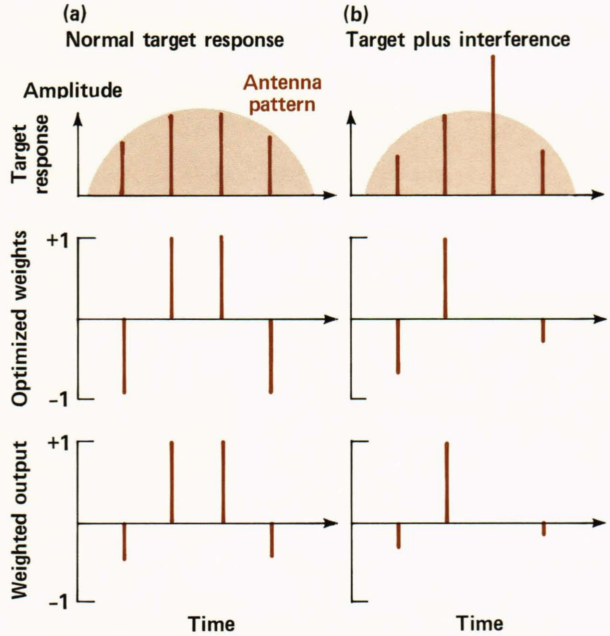

Figure 3 – An illustration of the effects of the transversal filter weights for two cases. For the normal target response, the antenna-modulated target is weighted by the values shown; the weighted results are illustrated in (a). If an interfering pulse is detected, as in (b), the weights are changed to cancel that interfering pulse and reoptimize the filter response.

The basic function of the CFAR processor is to determine an adaptive detection threshold against which the incoming signals are compared. To maintain a uniform CFAR, the processor must adapt to the environment as a function of time and spatial distribution. The latest advances in medium- and large-scale high-speed integrated circuits make possible a real-time digital implementation of the required algorithms for threshold determination.

The CFAR processor that was implemented is shown in Fig. 4. To set a threshold against which all incoming signals are compared in order to determine the presence of a target, the processor samples the local environment and attempts to determine the statistical distribution of the data. This is done in two stages.

The first stage is referred to as a moving window adaptive video processor. As data are obtained from the receiver, the data around the cell that is being processed are averaged to determine the mean value in the immediate vicinity. The whole process moves in time as data are received, hence the term "moving window processor." The mean value is subtracted from each cell before it is sent for further processing to produce a data set with zero mean. If all the data

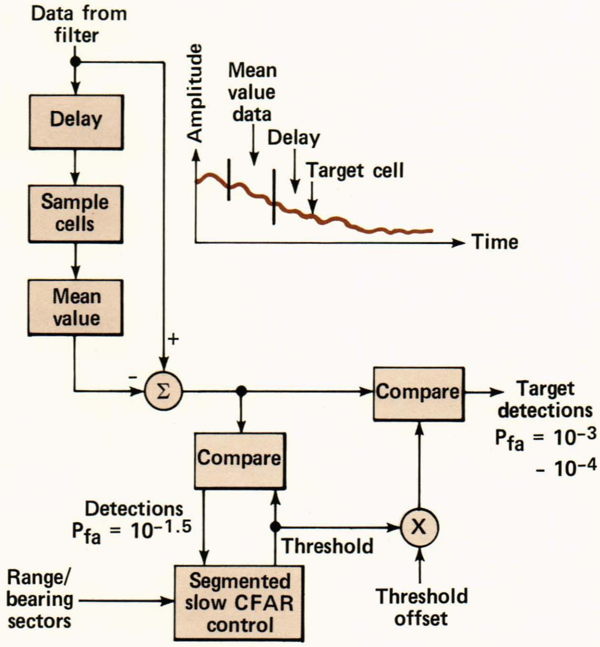

Figure 4 – The configuration of the constant false alarm rate processing circuits. The statistical mean of the input data is calculated and removed from the data in a moving window fashion. A threshold is then set by closed-loop control that is segmented to produce separate thresholds for 4096 range/bearing sectors.

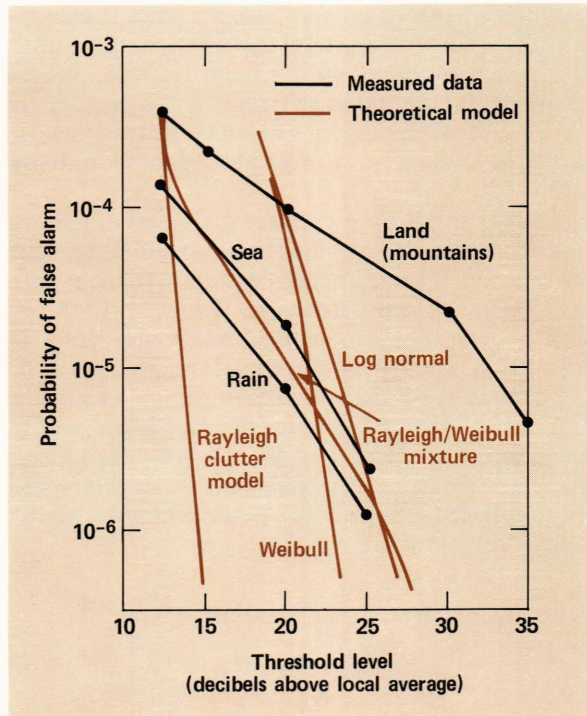

Figure 5 – Measured and theoretical probability distributions for an S-band radar.

have the same statistical distribution but vary in intensity, a simple threshold on this zero mean data set would produce a uniform CFAR over the entire surveillance area. This is not generally the case because the location of the clutter data source changes within the search area, and different types of clutter data usually exhibit different statistical characteristics.

The second stage of the CFAR processor is designed to minimize the effect of varying clutter distributions. The zero mean data are thresholded where the threshold value is determined in closed-loop fashion to keep a CFAR of 1 false alarm in 32 opportunities. A separate threshold is calculated and updated in area segments. In the processor developed for surface surveillance radars, there is a total of 4096 segments covering the search areas, each containing an independent threshold based on the closed-loop calculation. The exact size of the segments is set to provide the best match for the spatial distribution of the sea clutter at the radar site.

The false alarm rate produced in the closed-loop system is set high in order to control its value precisely. A false alarm rate of 1 in 32 is too large to be sent for further processing, so each threshold is multiplied by a constant and compared to the zero mean data to produce the final detection output at a lower false alarm rate.

RETROSPECTIVE DATA PROCESSING

Because of the statistical nature of many types of clutter, a CFAR detection device must set a high detection threshold to maintain a reasonably low false alarm rate. A rate of 1 in 10^6 to 10^8 opportunities is typically required at the input to an automatic tracking system. By setting the threshold to obtain this rate, detections of small- and medium-sized targets can be missed.

Most theoretical studies of radar processors and most radar range equation predictions are based on Gaussian noise and clutter assumptions that are seldom valid in practice. As an example, Fig. 5 illustrates measured cumulative probability distributions of the envelope of clutter returns from an S-band pencil-beam 3D radar in three environments (sea clutter, land clutter, and rain clutter). Also shown is the cumulative Rayleigh probability distribution that would describe the clutter envelope if the clutter or noise was in fact Gaussian. Three conclusions from this and similar clutter studies are that

1. Clutter statistics are not Gaussian;
2. No single non-Gaussian model accurately describes the variety of possible clutter types and situations;
3. The detection threshold required to maintain a low CFAR (e.g., 10^{-6} to 10^{-8}) can be 10 to 20 decibels above that required with Gaussian noise (thus leading to a 10 to 20 decibel detection loss).

A retrospective processor uses all contacts from several past radar scans, examining all possible target trajectories formed from stored contacts for each input detection. The retrospective processing architecture enables the processor to sort through thousands of contacts efficiently, eliminating many false alarms and retaining those contacts that describe reasonable trajectories. This approach overcomes the problem by allowing the contact detection process to operate at a much higher false alarm rate (e.g., 10^{-3}), one at which the detection loss for non-Gaussian clutter is greatly reduced and the differences between clutter types are smaller.

This process is ideally suited to sea surface surveillance radars. The range resolution of these radars is exceptionally fine, permitting the use of narrow trajectory filters in the range dimension. A filter width on the order of 7 knots is feasible and gives very good performance on contact data in the sea clutter region when operated with a false alarm rate of 1 in 1000, or about 2000+ false contacts per scan. A detailed description of the retrospective concept as applied to sea surface surveillance data is given in the next section.

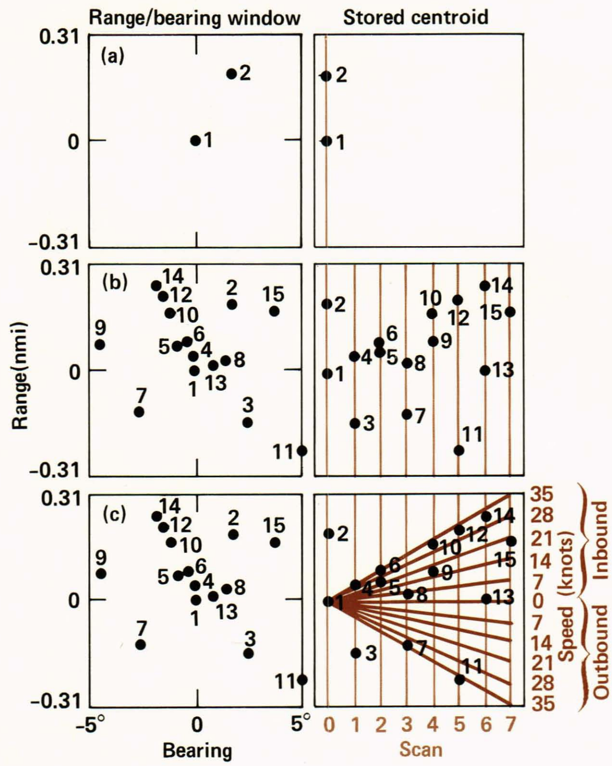

Figure 6 – The retrospective process: (a) a single scan of data, (b) eight scans of data, and (c) eight scans of data with trajectory filters applied.

RETROSPECTIVE PROCESSING CONCEPT

The major elements of the retrospective processing concept are (a) retrospective correlation using a bearing/range link structure, (b) updates of a velocity profile mask, and (c) an output decision based on the updates of the mask.

Linked list data structures are crucial to an efficient search of large collections of data. The bearing/range-ordered linked list structure provides automatically the first dimension of a correlation process and allows a logical ordering of processes to follow a scanning radar. The major advantages of bearing/range-linked correlation with multiscan data are the removal of old data and restoration of the link. Scan-to-scan linking is always newest to oldest, so that the memory for the oldest scan can be reset or overwritten without disturbing the link for active scans.

To determine whether or not a target is present on the basis of M scans of data, patterns of detection are evaluated for their reasonableness under the assumption that the data represent true target returns. A velocity profile is constructed in the range-only dimension with data accumulated over bearing zones representative of the bearing resolution of the radar. Each profile consists of M bits, indicating whether a detection occurred in a cell corresponding to that profile in each of the M scans preceding the contact of interest.

The velocity profile updates produce a hit/miss pattern for each velocity profile, resulting in a multiscan hit pattern rather than a simple count of detections associated with a profile. This allows a quality to be assigned to each pattern of hits and misses that is compared to a threshold. Only if the quality in any profile for a given contact exceeds the threshold is that contact output by the process.

As an example, the various steps are shown in Fig. 6 for the special case of a processor for surface target detections with a seven-scan correlation and 10 7-knot velocity profiles. This filter examines every contact and compares it with contacts from the previous seven scans if they fall within a certain range and bearing window about the contact. The contact of interest is output only if it forms a probable track with contacts from the previous seven scans. The range and bearing windows are set to examine all prior contacts that could result from a target moving at up to 35 knots in any direction.

Figure 6a shows a window with a single scan of data. Contact 1 is the contact to be examined, and contact 2 is another that falls in the range/bearing window. At the right of the range/bearing window, all contacts that appear in the window on this scan are plotted as a function of range.

Figure 6b is the same range/bearing window about contact 1, showing all the contacts from the previous seven scans. The false alarm rate used in this plot is 10^{-4}. It is clear from this eight-scan history that contact 1 forms a target track with contacts 4, 6, 10, 12, and 14 from previous scans.

Here again, to the right of the range/bearing window, the contacts within the window on each scan are plotted versus range. The function of the retrospective data processor is to look at the contacts from the previous seven scans and see if they form a probable track with the contact of interest. To do this, the stored contacts are examined for patterns of detection data, and the reasonableness of the patterns is evaluated under the assumption that the data represent a true target return.

Figure 6c shows how the detection patterns are determined. A set of velocity profiles is constructed in the range-only dimension for contact data within the range/bearing window. The profiles are set up in 7-knot increments for target range rates of 0 to 35 knots inbound and outbound, making a total of 10 profiles. The detection pattern is determined for each profile. There are 127 possible detection patterns, each representing a certain probability that the contact is real. That probability is computed assuming Poisson clutter spatial distributions. A threshold set on this probability determines whether a contact is to be output.

EXPERIMENTAL RESULTS

The primary objective of the project was to demonstrate the signal processing concepts of noncoherent moving target indicators, constant false alarm rate processing, and retrospective data processing using actual radar data in an operational environment. To accomplish this, a signal processing unit was designed and built that contained a high-speed video digitizer and a digital transversal filter, as well as a two-stage CFAR processor. The processor was interfaced with a digital magnetic tape system so that data about potential targets could be recorded in the field for offline processing.

Figure 7 shows the preprocessor and the associated data recording equipment that were used in initial experiments. The preprocessor consists of an amplitude processor in which input video is digitized and filtered, a processing section where filtered data are detected with a well-regulated false alarm rate in all environments, and data formatting circuits for interfacing with the digital magnetic tape system.

Digital recordings of the contacts are made and returned to the laboratory for reduction. The centroids of potential targets form the input to the retrospective processor. The initial retrospective processor was developed in software on a general-purpose computer, which permitted us to perfect algorithms without building hardware. The first version of the retrospective algorithm was used extensively to filter data collected in several sea surface surveillance environments. Once the basic technique was proven, hardware and high-speed microprocessor-based software designs to operate in real time were developed.

Figures 8 and 9 illustrate qualitatively the operation of the retrospective processor in a sea clutter environment. An AN/FPS-114 radar mounted on Laguna Peak on the California coast was instrumented as shown in Fig. 10. The S-band radar performs sea surveillance for range-safety purposes. It has a high spatial resolution antenna (0.1 microsecond pulse width by 0.9° azimuth beamwidth) and a 4 second rotation period, but no coherent Doppler processing. It is approximately 1500 feet above sea level. Figure 8 is a plan position indicator (PPI) display of a raw video recording from a single scan of the radar. After onsite analog-to-digital conversion, the video signals pass through a digital-range-averaging logarithmic

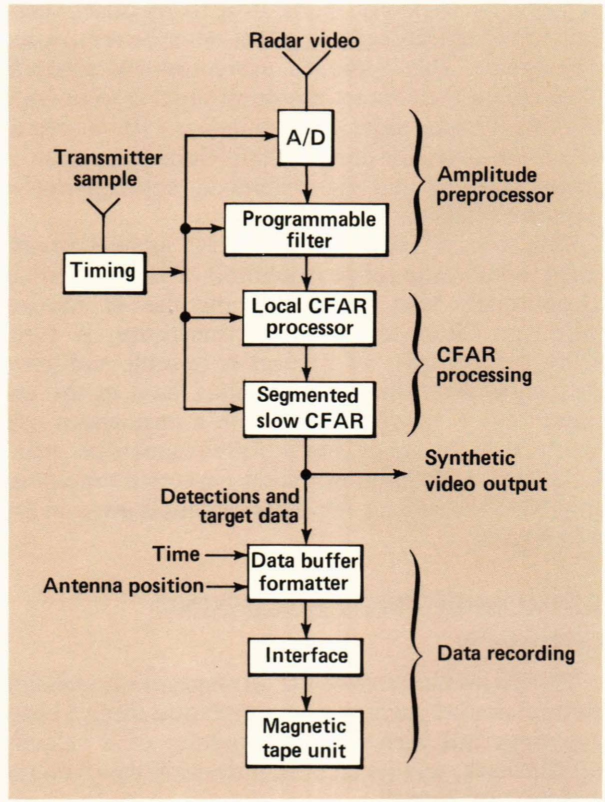

Figure 7 – The signal processor and data recording configuration used in the initial experiments on a sea surface surveillance radar. The high-speed data filtering and CFAR processing were part of the processor taken into the field. Potential target detections were recorded for retrospective processing in the laboratory.

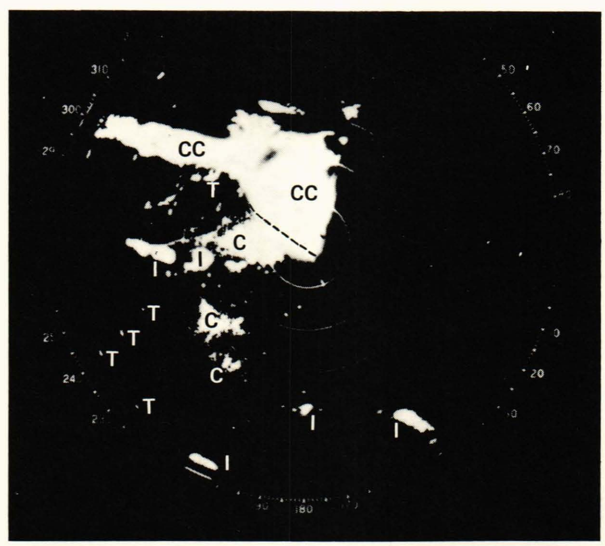

Figure 8 – Plan position display of raw radar video from the sea surveillance experiment. The display covers 70 nautical miles showing the Channel Islands (I) and the coast of California (CC). Close-in sea clutter (C) and some long-range targets (T) are visible.

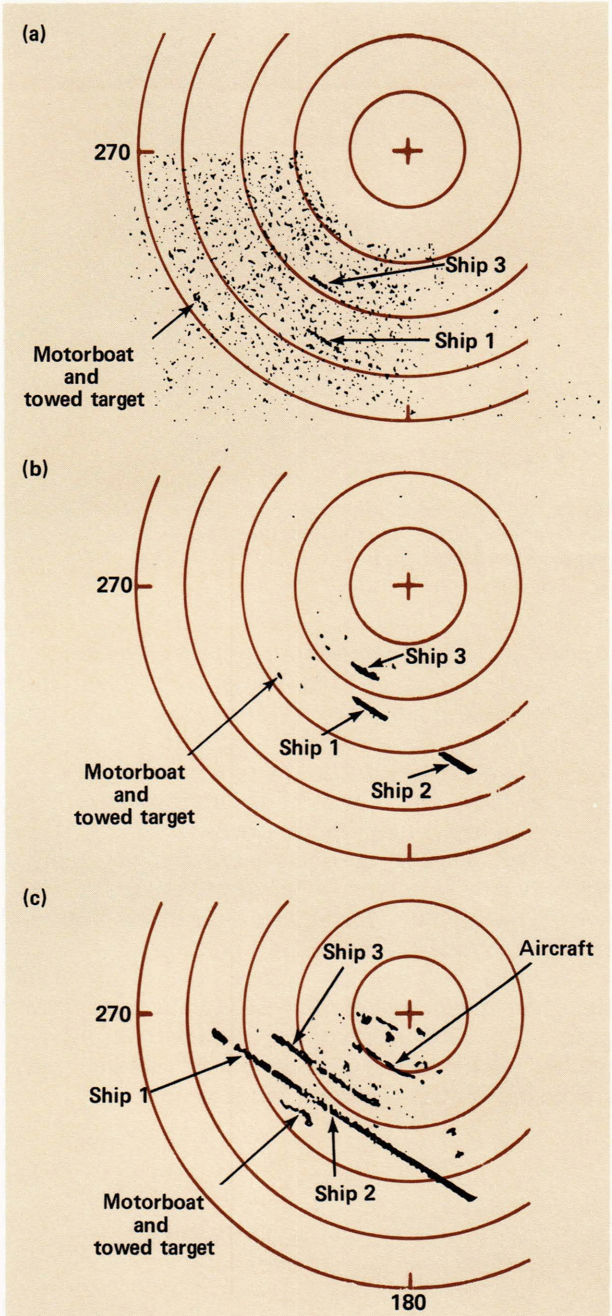

Figure 9 – The operation of the retrospective processor in a sea-clutter environment: (a) contact data after CFAR processing (100 scans, 2 nmi range rings, 9 nmi total range); (b) contact data after retrospective processing (100 scans, 3 nmi range rings); (c) contact data after retrospective processing (1000 scans, 3 nmi range rings). The long target tracks are commercial vessels in the Santa Barbara Channel and the close-in tracks are fishing and sport craft. Only data in the area of interest were recorded and presented.

CFAR device to produce the contacts shown in Fig. 9a. A high false alarm rate (approximately 10^{-3}) is maintained in the device, resulting in approximately 2000 false alarms per scan. False alarms from sea clutter are indistinguishable from small target returns when viewed on a single-scan basis.

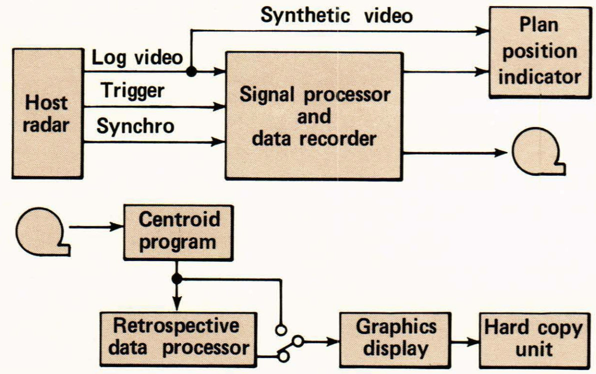

Figure 10 – Data collection and analysis system used in signal processing experiments.

Figure 9b illustrates two effects of passing 100 scans of radar contacts through the retrospective processor. The false alarm rate has been reduced by at least four orders of magnitude, and the ships and boats in the channel are clearly visible. The large reduction in false alarm rate that the retrospective processor produces is further illustrated by Fig. 9c, which shows the effects of passing 1000 scans (about 1 hour) of radar contacts through the processor. Only a few actual false alarms are visible. This ability to reduce false alarm rates greatly by exploiting the spatial and temporal correlation properties of sea clutter allows much greater target sensitivity to be achieved.

AUTOMATED RANGE SURVEILLANCE SYSTEM

The successful demonstration of these signal processing concepts using actual radar data resulted in a project for the Pacific Missile Test Center to instrument its sea surface surveillance radars with signal processors and to integrate the data from noncollocated radars to form a composite surface target track data base. This system is known as the Automated Range Surveillance System. The radars are located at two missile and gun test firing ranges, one off the coast north of Los Angeles and one off the coast of the island of Kauai in Hawaii. Each range is instrumented with three noncollocated sea surface surveillance radars for locating surface craft. These radar data are used primarily to locate all surface craft in the test area so that an evaluation of a safe ordnance firing condition can be made.

Prior to the project for automating these radars, all detection and entering of target tracks were performed by operators viewing raw radar data on a display scope. The Automated Range Surveillance System is designed to enhance target detectability, automate the detection process, integrate the target data from all surface surveillance radars, and form target tracks from this composite data base. It provides reliable target detection in all environments and produces a data base that accurately locates the present position of all surface targets and predicts future target positions by means of the track data. The end result is a reliable, accurate method for ensuring the safety of test firing operations.

The locations of the three radars used at the California range are shown in Fig. 11, the optical radar horizon being indicated for each radar. They are all identical S-band radars (3000 megahertz transmitting frequency) with a 0.1 microsecond transmission pulse width at a rate of 1000 per second and an antenna beamwidth of 0.9°.

Figure 12 is a block diagram of the complete Automated Range Surveillance System for the Point Mugu installation. The system consists of three identical radar signal processing units, one for each radar; data communication equipment for transmitting target data to the range control center; and integration, tracking, and display equipment at the range control center.

The radar signal processing unit accepts the raw analog radar data and digitizes the data at a 10 megahertz rate. The digitized data are filtered in two digital transversal filters operating in parallel. One filter is programmed to provide the target enhancement features on the basis of the amplitude-only moving target indicator algorithms, and the second is programmed to integrate the data to provide maximum target detectability in a thermal-noise-limited environment. The outputs of the filters go to the CFAR processor to provide target detections, which are combined at this point and sent to a microprocessor-based digital processor for further processing.

The detection data are provided at a high false alarm rate to ensure good target detectability in all environments. This is typically set at a false alarm probability of 10^{-3} to 10^{-4}, giving 5000 to 10,000 false alarms every radar scan (4 seconds). After the data are input to the microprocessor, they are centroided to yield only one data point per potential target. For the radars at Point Mugu, this reduces the 5000 to 10,000 contacts to 1000 to 2000 potential target centroids. The centroids are sent to another microprocessor that performs the retrospective data filter algorithm. The output of the retrospective data filter comprises only those potential targets that passed the trajectory correlation criteria. Essentially, these are true targets, and their number ranges from 8 to 12 up to 200, depending on the radar location and the number of pleasure craft in the area.

The data required to transfer the information about this small number of targets can be transmitted from a standard 9600-baud telephone transmission system. Figure 13 is a photograph of a radar signal processing unit for the Point Mugu system.

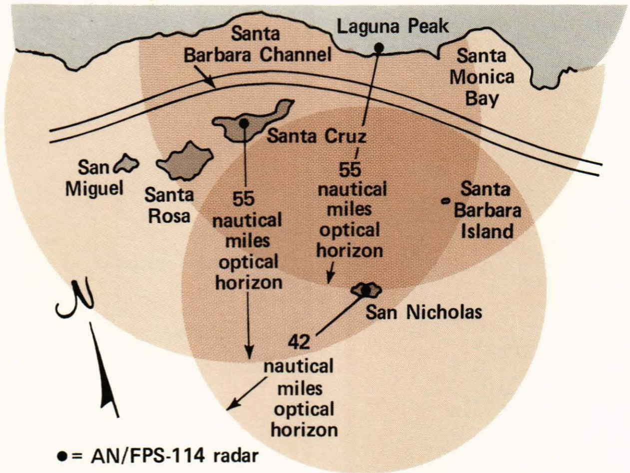

Figure 11 – The locations of the AN/FPS-114 surface surveillance radars at the Pacific Missile Test Center Range, Point Mugu, Calif.

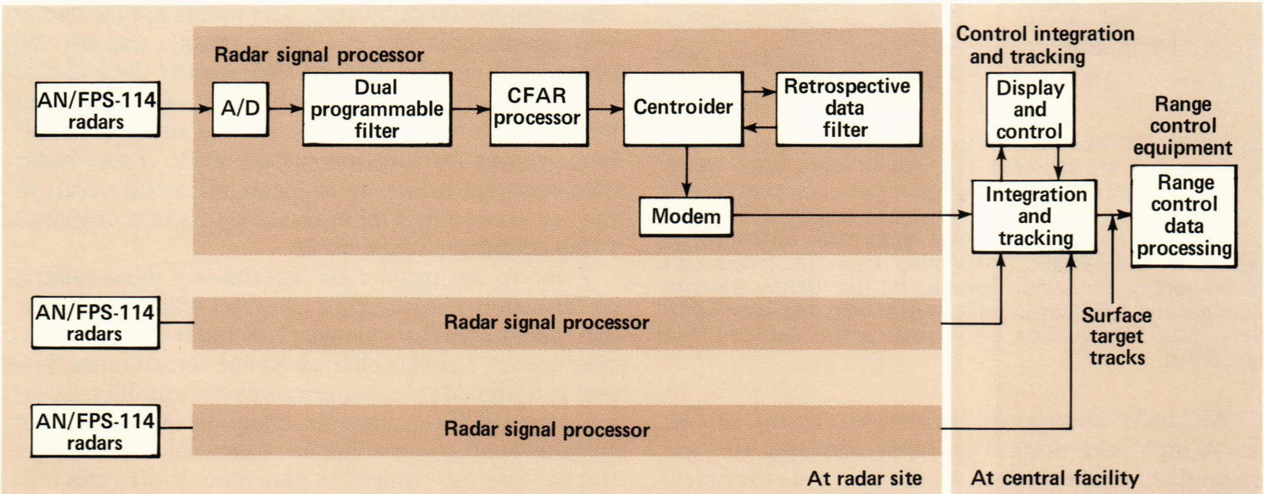

Figure 12 – The Automated Range Surveillance System that is being developed at APL for the Pacific Missile Test Center. This system is an integrated automatic detection and tracking system that is based on the signal processing technique developed for the detection of surface targets.

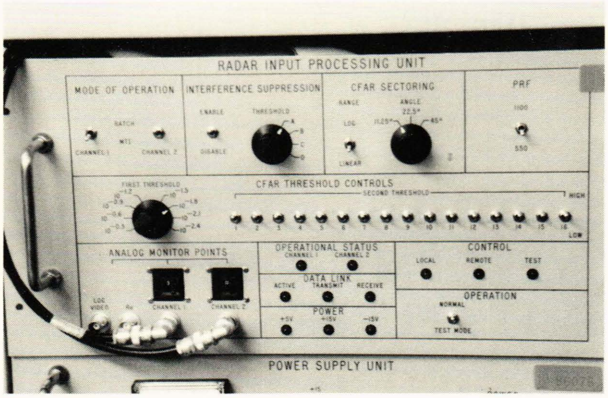

Figure 13 – A radar signal processor that is attached to the AN/FPS-114 surface surveillance radars. These units perform automatic target detection and report retrospective filtered data to the tracking and display units.

The target data from all three radars are combined in a microprocessor at the Range Operations Center at Point Mugu Naval Air Station. The composite data base is used to form target tracks, which are then passed to the operations computer that is part of the Range Operations Center. The target track data are also used to form display data that are formatted as a track history and interfaced with a color display console, which is part of the Automated Range Surveillance System hardware, and also with the standard Navy display consoles used in the Operations Center.

The first phase of the project was to design and install one radar signal processing unit at a radar site, transmit the data to the Range Operations Center, and generate the display data for the consoles. This task was completed during the first quarter of 1983. Production versions of the radar signal processing units are to be completed by early 1984, along with the software required for integration and tracking.

The system for the test range in Hawaii is conceptually the same as that at Point Mugu, with minor hardware and software changes to accommodate different radar types and locations. That system is also scheduled for 1984 delivery.

FUTURE APPLICATIONS

Although the signal processing concepts described were applied primarily to surveillance radars and, in particular, sea surface surveillance, they have a much broader potential application. The retrospective data filter, for instance, can be applied to any data base that is sampled periodically and in which the true signals obey some trajectory pattern not characteristic of the noise data. This has been tried on several different data bases from 2D and 3D air surveillance radars and infrared surveillance devices.

The concept of adapting the data filter in the presence of interfering signals can be applied to any surveillance radar that uses a digital transversal filter to enhance target signals. At present, most radars using this type of processing throw away all the data present in the filter when interference is detected in a range cell. Under conditions of heavy interference, this can cause an overall degradation in performance.

CONCLUSIONS

Given the state of the art in signal processing today, the development of a cost-effective digital signal processing system for sea surface surveillance radars is a reality. The ready availability of low-cost high-speed digital circuitry, high-density solid-state memory, and microprocessors makes possible the implementation of novel techniques for target signal enhancement and false alarm reduction that were heretofore unattainable.

Several unique techniques have been perfected and demonstrated on data from sea surface surveillance radars. A much broader application of some of these processing concepts is envisioned for the future.

BIBLIOGRAPHY

F. R. Castella, "An Optimized Semi-Coherent MTI for Slow Target Detection," International Conf. on Digital Signal Processing (Sep 1980).

F. R. Castella, "Optimization of a Generalized Square Law Detector," *IEE Proc.* **130** P&F, No. 5 (1983).

R. J. Prengaman, R. E. Thurber, and W. G. Bath, "A Retrospective Detection Algorithm for Extraction of Weak Targets in Clutter and Interference Environments," *Proc. International Conf. Radar 82* (Oct 1982).

R. J. Prengaman, R. E. Thurber, and W. G. Bath, "A Retrospective Detection Algorithm for Extraction of Weak Targets in Clutter and Interference Enviroments," *Proc. EASCON '83* (Sep 1983).

R. J. Prengaman, R. E. Thurber, and W. G. Bath, "Retrospective Processing – A New Technique for the Radar Detection of Surface Targets," *Applied Physics Laboratory Developments in Science and Technology, Fiscal Year 1982,* JHU/APL DST-10 (1983).

R. E. Thurber, R. Rzemien, S. M. Myers, G. W. Riffle, and R. J. Clevering, "Surface Radar Processor," *Applied Physics Laboratory Developments in Science and Technology, Fiscal Year 1980,* JHU/APL DST-8.

ACKNOWLEDGMENTS – Many people contributed to the success of this work. B. D. Dobbins provided the opportunity to pursue our ideas under IR&D funding and encouraged us to promote our techniques to potential users. R. J. Prengaman provided the concepts and guided the project. F. R. Castella and W. G. Bath provided the analytic background for the semicoherent moving target indicator features. Dr. Bath and L. A. Biddison provided the analytic background for the retrospective data processor. Ms. Biddison also developed the detailed algorithms used in the microprocessor-based real-time processor. G. W. Riffle, W. L. Hom, and G. A. Miller designed the hardware for the Automated Range Surveillance System, and M. R. Newcomer and Ms. Biddison developed the microprocessor software. C. P. Richards provided the detailed mechanical and electrical design and assembly of the Automated Range Surveillance System.