
CHAPTER 8

AUTOMATIC DETECTION, TRACKING, AND SENSOR INTEGRATION

G. V. Trunk
Naval Research Laboratory

8.1 INTRODUCTION

Since the invention of radar, radar operators have detected and tracked targets by using visual inputs from a variety of displays. Although operators can perform these tasks very accurately, they are easily saturated and quickly become fatigued. Various studies have shown that operators can manually track only a few targets. To correct this situation, automatic detection and tracking (ADT) systems were attached to many radars. As digital processing increases in speed and hardware decreases in cost and size, ADT systems will become associated with almost all but the simplest radars.

In this chapter, automatic detection, automatic tracking, and sensor integration systems for air surveillance radar will be discussed. Included in this discussion are various noncoherent integrators that provide target enhancement, thresholding techniques for reducing false alarms and target suppression, and algorithms for estimating target position and resolving targets. Then, an overview of the entire tracking system is given, followed by a discussion of its various components such as tracking filter, maneuver-following logic, track initiation, and correlation logic. Next, multiscan approaches to automatic tracking such as maximum likelihood are discussed. Finally, the chapter concludes with a discussion of sensor integration and radar netting, including both colocated and multisite systems.

8.2 AUTOMATIC DETECTION

The statistical framework necessary for the development of automatic detection was applied to radar in the 1940s by Marcum,¹ and later Swerling² extended the work to fluctuating targets. They investigated many of the statistical problems

associated with the noncoherent detection of targets in Rayleigh noise. (NOTE: If the quadrature components are gaussian-distributed, the envelope is Rayleigh-distributed and the power is exponentially distributed.) Marcum's most important result was the generation of curves of probability of detection (P_D) versus signal-to-noise ratio (S/N) for a detector which sums N envelope-detected samples (either linear or square-law) under the assumption of equal signal amplitudes. However, in a search radar, as the beam sweeps over the target, the returned signal amplitude is modulated by the antenna pattern. Many authors investigated various detectors (weightings), comparing detection performance and angular estimation results with optimal values; and many of these results are presented later in this section.

In the original work on detectors, the environment was assumed known and homogeneous, so that fixed thresholds could be used. However, a realistic environment (e.g., containing land, sea, and rain) will cause an exorbitant number of false alarms for a fixed-threshold system that does not utilize excellent coherent processing. Three main approaches, adaptive thresholding, nonparametric detectors, and clutter maps, have been used to solve the false-alarm problem. Both adaptive thresholding and nonparametric detectors are based on the assumption that homogeneity exists in a small region about the range cell that is being tested. The adaptive thresholding method assumes that the noise density is known except for a few unknown parameters (e.g., the mean and the variance). The surrounding reference cells are then used to estimate the unknown parameters, and a threshold based on the estimated density is obtained. Nonparametric detectors obtain a constant false-alarm rate (CFAR) by ranking the test samples (ordering the samples from smallest to largest), usually with the reference cells. Under the hypothesis that all the samples (test and reference) are independent samples from an unknown density function, the test sample has a uniform density function, and, consequently, a threshold which yields CFAR can be set. Clutter maps store an average background level for each range-azimuth cell. A target is then declared in a range-azimuth cell if the new value exceeds the average background level by a specified amount.

Optimal Detector. The radar detection problem is a binary hypothesis-testing problem in which H_0 denotes the hypothesis that no target is present and H_1 is the hypothesis that the target is present. While several criteria (i.e., definitions of optimality) can be used to solve this problem, the most appropriate for radar is the Neyman-Pearson.³ This criterion maximizes the probability of detection P_D for a given probability of false alarm P_{fa} by comparing the likelihood ratio L [defined by Eq. (8.1)] to an appropriate threshold T which determines the P_{fa} . A target is declared present if

$$L(x_1, \dots, x_n) \frac{p(x_1, \dots, x_n|H_1)}{p(x_1, \dots, x_n|H_0)} \geq T \quad (8.1)$$

where $p(x_1, \dots, x_n|H_1)$ and $p(x_1, \dots, x_n|H_0)$ are the joint probability density functions of the n samples x_i under the conditions of target presence and target absence, respectively. For a linear envelope detector the samples have a Rayleigh density under H_0 and a rician density under H_1 , and the likelihood ratio detector reduces to

$$\prod_{i=1}^n I_0\left(\frac{A_i x_i}{\sigma^2}\right) \geq T \quad (8.2)$$

where I_0 is the Bessel function of zero order, σ^2 is the noise power, and A_i is the target amplitude of the i th pulse and is proportional to the antenna power pattern. For small signals ($A_i < \sigma$), the detector reduces to the square-law detector

$$\sum_{i=1}^n A_i^2 x_i^2 \geq T \quad (8.3)$$

and for large signals ($A_i \gg \sigma$), it reduces to the linear detector

$$\sum_{i=1}^n A_i x_i > T \quad (8.4)$$

For constant signal amplitude (i.e., $A_i = A$) these detectors were first studied by Marcum¹ and were studied in succeeding years by numerous other people. Detection curves for both linear and square-law detectors are given in Chap. 2. The most important facts concerning these detectors are the following:

- The detection performances of the linear and square-law detectors are similar, differing only by less than 0.2 dB over wide ranges of P_D , P_{fa} , and n .
- Since the signal return of a scanning radar is modulated by the antenna pattern, to maximize the S/N when integrating a large number of pulses with no weighting (i.e., $A_i = 1$) only 0.84 of the pulses between the half-power points should be integrated, and the antenna beam-shape factor (ABSF) is 1.6 dB.⁴ The ABSF is the number by which the midbeam S/N must be reduced so that the detection curves generated for equal signal amplitudes can be used for the scanning radar.
- The collapsing loss for the linear detector can be several decibels greater than the loss for a square-law detector⁵ (see Fig. 8.1). The collapsing loss is the additional signal required to maintain the same P_D and P_{fa} when unwanted noise samples along with the desired signal-plus-noise samples are integrated. The number of signal samples integrated is N , the number of extraneous noise samples integrated is M , and the collapsing ratio $\rho = (N + M)/N$.
- Most automatic detectors are required not only to detect targets but to make angular estimates of the azimuth position of the target. Swerling⁶ calculated the standard deviation of the optimal estimate by using the Cramer-Rao lower bound. The results are shown in Fig. 8.2, where a normalized standard deviation is plotted against the midbeam S/N . This result holds for a moderate or large number of pulses integrated, and the optimal estimate involves finding the location where the correlation of the returned signal and the derivative of the antenna pattern is zero. Although this estimate is rarely implemented, its performance is approached by simple estimates.

Practical Detectors. Many different detectors (often called *integrators*) are used to accumulate the radar returns as a radar sweeps by a target. A few of the most common detectors⁷ are shown in Fig. 8.3. Though they are shown in the figure as being constructed with shift registers, they would normally be implemented with random-access memory. The input to these detectors can be linear, square-law, or log video. Since linear is probably the most commonly used, the advantages and disadvantages of the various detectors will be stated for this video.

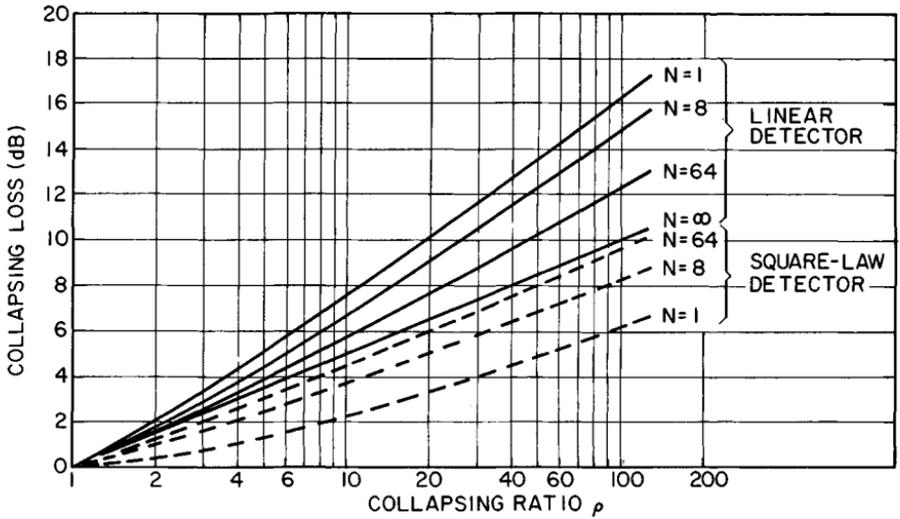


FIG. 8.1 Collapsing loss versus collapsing ratio for a probability of false alarm of 10^{-6} and a probability of detection of 0.5. (Copyright 1972, IEEE; from Ref. 5.)

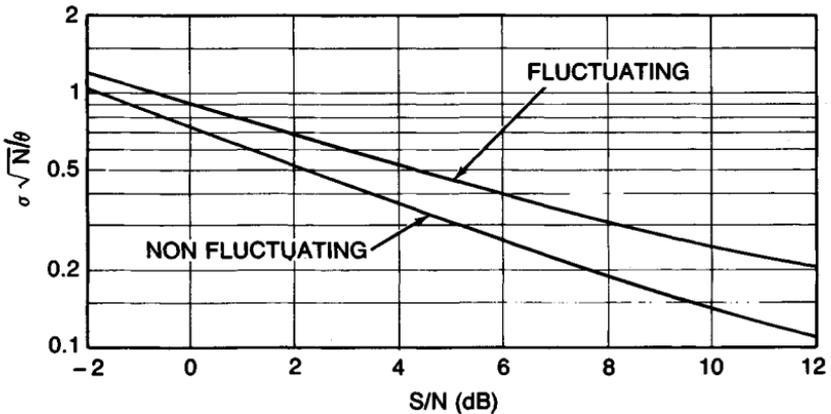


FIG. 8.2 Comparison of angular estimates with the Cramer-Rao lower bound. σ is the standard deviation of the estimation error, and N is the number of pulses within the 3-dB beamwidth, which is v . The S/N is the value at the center of the beam. (Copyright 1956, IEEE; after Ref. 6.)

Moving Window. The moving window in Fig. 8.3a performs a running sum of n pulses in each range cell;

$$S_i = S_{i-1} + x_i - x_{i-n} \quad (8.5)$$

where S_i is the sum at the i th pulse of the last n pulses and x_i is the i th pulse. The performance⁸ of this detector for $n \approx 10$ is only 0.5 dB worse than the optimal detector given by Eq. (8.3). The detection performance can be obtained by using

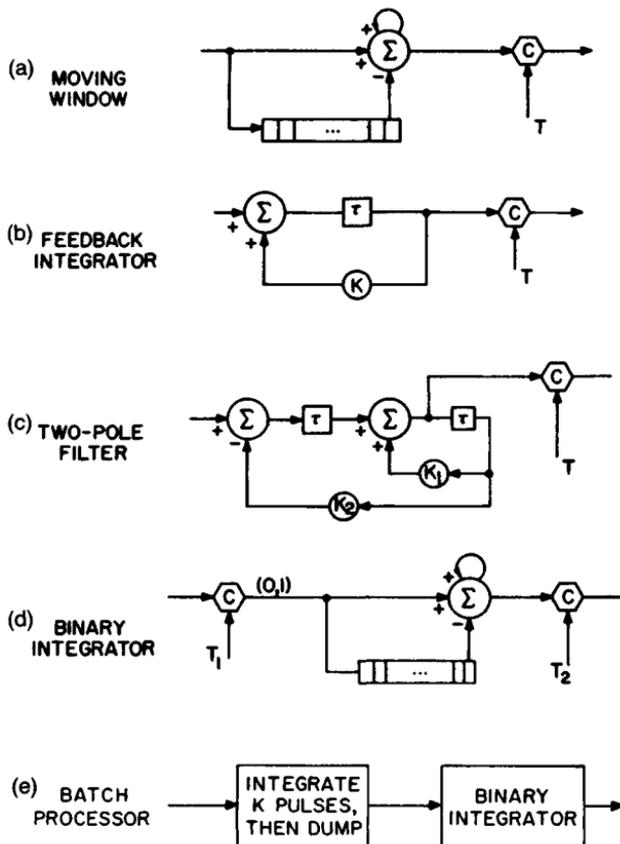


FIG. 8.3 Block diagrams of various detectors. The letter C indicates a comparison, τ is a delay, and loops indicate feedback. (From Ref. 7.)

an ABSF of 1.6 dB and the detection curves in Chap. 2. The angular estimate that is obtained by either taking the maximum value of the running sum or taking the midpoint between the first and last crossings of the detection threshold has a bias of $n/2$ pulses, which is easily corrected. The standard deviation of the estimation error of both estimators is about 20 percent higher than the optimal estimate specified by Cramer-Rao bound. A disadvantage of this detector is that it is susceptible to interference; that is, one large sample from interference can cause a detection. This problem can be minimized by using limiting. A minor disadvantage is that the last n pulses for each range cell must be saved, resulting in a large storage requirement when a large number of pulses are integrated. However, because of the availability of large memories of reduced size and cost, this is a minor problem.

The detection performance discussed previously is based on the assumption that the target is centered in the moving window. In the real situation the radar scans over the target, and decisions which are highly correlated are made at every pulse. Hansen⁹ analyzed this situation for $N = 2, 4, 8,$ and 16 pulses and

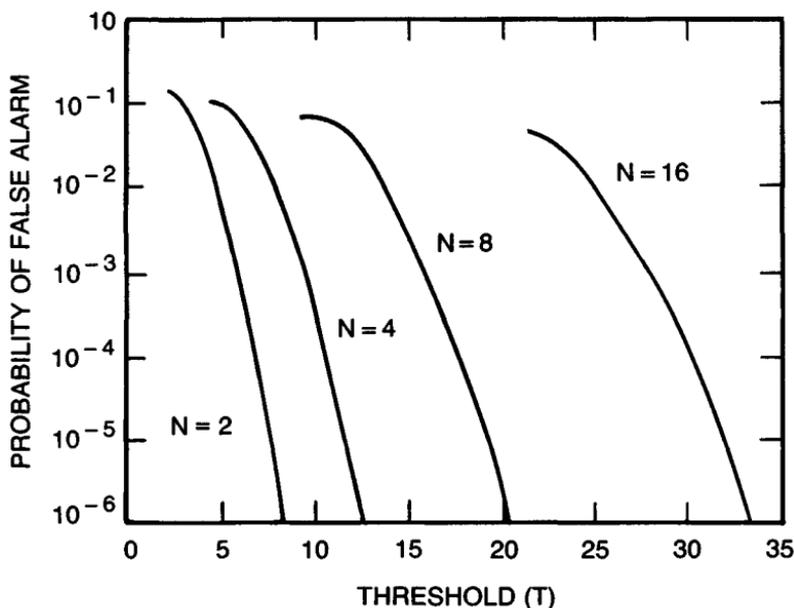


FIG. 8.4 Single-sweep false-alarm probability P_{fa} versus threshold for moving window. The noise is Rayleigh-distributed with $\sigma = 1$. (Copyright 1970, IEEE; after Ref. 9.)

calculated the detection thresholds shown in Fig. 8.4, the detection performance shown in Fig. 8.5, and the angular accuracy shown in Fig. 8.6. Comparing Hansen's scanning calculation with the single-point calculation, one concludes that 1 dB of improvement is obtained by making a decision at every pulse. The angular error of the beam-splitting procedure is about 20 percent greater than the optimal estimate. For large signal-to-noise ratios, the accuracy (rms error) of the beam-splitting and maximum-return procedures will be limited by the pulse

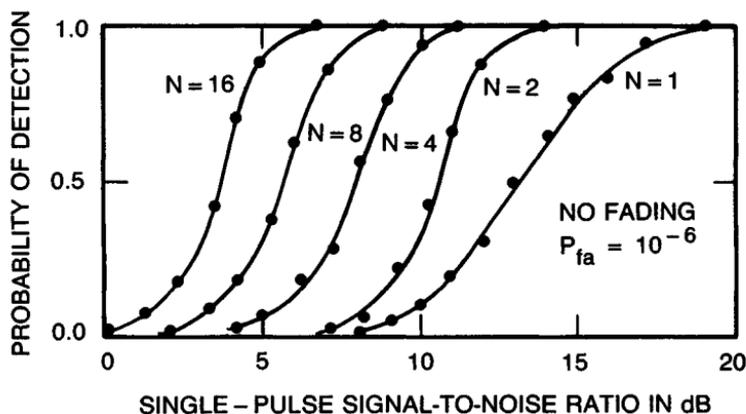


FIG. 8.5 Detection performance of the analog moving-window detector for the no-fading case. (Copyright 1970, IEEE; after Ref. 9.)

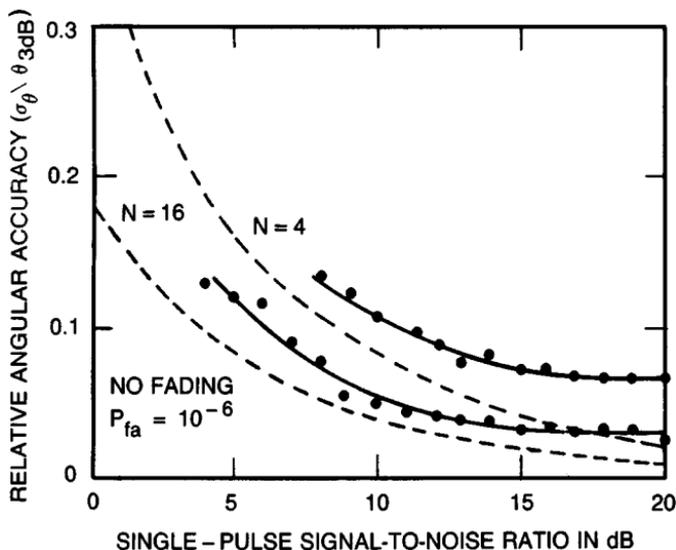


FIG. 8.6 Angular accuracy obtained with beam-splitting estimation procedure for the no-fading case. Broken-line curves are lower bounds derived by Swerling,⁶ and points shown are simulation results. (Copyright 1970, IEEE; after Ref. 9.)

spacing¹⁰ and will approach

$$\sigma(\hat{\theta}) = \Delta\theta/\sqrt{12} \quad (8.6)$$

where $\Delta\theta$ is the angular rotation between transmitted pulses. Consequently, if the number of pulses per beamwidth is small, the angular accuracy will be poor. For instance, if pulses are separated by 0.5 beamwidth, $\sigma(\hat{\theta})$ is bounded by 0.14 beamwidth. However, improved accuracy can be obtained by using the amplitudes of the radar returns. An accurate estimate of the target angle is given by

$$\hat{\theta} = \theta_1 + \frac{\Delta\theta}{2} + \frac{1}{2a\Delta\theta} \ln(A_2/A_1) \quad (8.7)$$

where

$$a = 1.386/(\text{beamwidth})^2 \quad (8.8)$$

and A_1 and A_2 are the two largest amplitudes of the returned samples and occur at angles θ_1 and $\theta_2 = \theta_1 + \Delta\theta$ respectively. Since the estimate should lie between θ_1 and θ_2 and Eq. (8.7) will not always yield such an estimate, $\hat{\theta}$ should be set equal to θ_1 if $\hat{\theta} < \theta_1$ and $\hat{\theta}$ should be set equal to θ_2 if $\hat{\theta} > \theta_2$. The accuracy of this estimator is given in Fig. 8.7 for the case of $n = 2$ pulses per beamwidth. This estimation procedure can also be used to estimate the elevation angle of a target in multibeam systems where θ_1 and θ_2 are the elevation-pointing angles of adjacent beams.

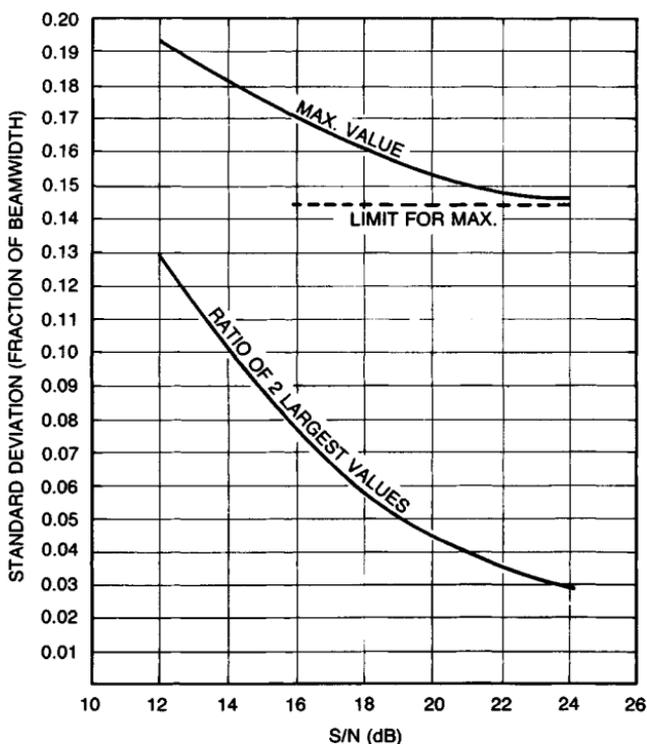


FIG. 8.7 Angular accuracy using two-pulse estimates.

Feedback Integrator. The amount of storage required can be reduced significantly by using a feedback integrator shown in Fig. 8.3b:

$$S_i = KS_{k-1} + x_i \quad (8.9)$$

For a feedback value of K , the effective number of pulses integrated M is $M = 1/(1 - K)$, and for optimal (maximum P_D) performance $M = 0.63 N$, where N is the number of pulses between the 3-dB antenna beamwidth.¹¹ The detection performance is given by the detection curves for M pulses with ABSF = 1.6 dB. Although the feedback integrator applies an exponential weighting into the past, its detection performance is only 1 dB less than that of the optimal integrator.⁸ Unfortunately, difficulties are encountered when using the feedback integrator to estimate the azimuth position.¹¹ The threshold-crossing procedure yields estimates only 20 percent greater than the lower bound, but the bias is a function of S/N and must be estimated. On the other hand, the maximum value, though it has a constant bias, has estimates that are 100 percent greater than the lower bound. Furthermore, the exponential weighting function essentially destroys the radar antenna sidelobes. Because of these problems, the feedback integrator has limited utility.

Two-Pole Filter. The two-pole filter in Fig. 8.3c requires the storage of an intermediate calculation in addition to the integrated output and is described mathematically by

$$y_i = x_i - k_2 z_{i-1} \quad (8.10)$$

and

$$z_i = y_{i-1} + k_1 z_{i-1} \quad (8.11)$$

where x_i is the input, y_i is the intermediate calculation, z_i is the output, and k_1 and k_2 are the two feedback values. The values^{12,13} which maximize P_D are given by

$$k_1 = 2 \exp(-\xi \omega_d \tau / \sqrt{1 - \xi^2}) \cos(\omega_d \tau) \quad (8.12)$$

and

$$k_2 = \exp(-2\xi \omega_d \tau / \sqrt{1 - \xi^2}) \quad (8.13)$$

where $\xi = 0.63$, $N\omega_d \tau = 2.2$, and N is the number of pulses between the 3-dB points of the antenna. With this rather simple device a weighting pattern similar to the antenna pattern can be obtained. The detection performance is within 0.15 dB of the optimal detector, and its angular estimates are about 20 percent greater than the Cramer-Rao lower bound. If the desired number of pulses integrated is changed (e.g., because of a change in the antenna rotation rate of the radar), it is only necessary to change the feedback values k_1 and k_2 . The problems with this detector are that (1) it has rather high detector sidelobes, 15 to 20 dB, and (2) it is extremely sensitive to interference (i.e., the filter has a high gain resulting in a large output for a single sample that has a high value).

Binary Integrator. The binary integrator is also known as the dual-threshold detector, M-out-of-N detector, or rank detector (see "Nonparametric Detectors" later in this section), and numerous individuals have studied it.¹⁴⁻¹⁸ As shown in Fig. 8.3d, the input samples are quantized to 0 or 1, depending on whether or not they are less than a threshold T_1 . The last N zeros and ones are summed and compared with a second threshold $T_2 = M$. For large N , the detection performance of this detector is approximately 2 dB less than the moving-window integrator because of the hard limiting of the data, and the angular estimation error is about 25 percent greater than the Cramer-Rao lower bound. Schwartz¹⁶ showed that within 0.2 dB the optimal value of M for maximum P_D is given by

$$M = 1.5\sqrt{N} \quad (8.14)$$

when $10^{-10} < P_{fa} < 10^{-5}$ and $0.5 < P_D < 0.9$. The optimal value of P_n , the probability of exceeding T_1 when only noise is present, was calculated by Dillard¹⁸ and is shown in Fig. 8.8. The corresponding threshold T_1 is

$$T_1 = \sigma(-2 \ln P_n)^{1/2} \quad (8.15)$$

A comparison of the optimal (best value of M) binary integrator with various other procedures is given in Figs. 8.9 and 8.10 for $P_D = 0.5$ and 0.9, respectively.

The binary integrator is used in many radars because (1) it is easily implemented, (2) it ignores interference spikes which cause trouble with integrators that directly use signal amplitude, and (3) it works extremely well when the noise has a non-Rayleigh density.¹⁹ For $N = 3$, comparison of the optimal binary integrator (3 out of 3), another binary integration (2 out of 3), and the moving-window detector in log-normal interference (an example of a non-Rayleigh density) is shown in Fig. 8.11. The optimal binary integrator is much better than the moving-

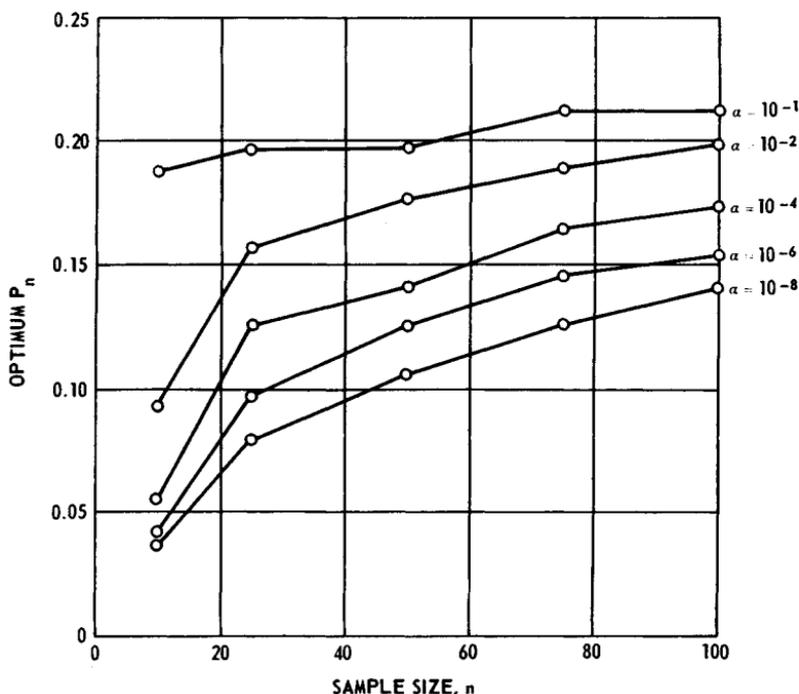


FIG. 8.8 Optimum values of P_n as a function of the sample size n and the probability of false alarm α ; Ricean distribution with $S/N = 0$ dB per pulse. (Copyright 1967, IEEE; from Ref. 18.)

window integrator. The optimal values for log-normal interference were calculated by Schleher¹⁹ and are $M = 3, 8,$ and 25 and $N = 3, 10,$ and 30 , respectively.

A modified version of binary integration is sometimes used when there is a large number of pulses. It also has flexibility to integrate a different number of pulses. The modified binary moving window (MBMW) differs from the ordinary binary moving window (OBMW) by the introduction of a third threshold. When the second threshold is reached, one counts the number of consecutive pulses for which the second threshold is exceeded. When this number equals the third threshold, a target is declared. The performance of the MBMW and a comparison with the OBMW were given in Ref. 20. The major conclusion to be drawn is that the larger the value of N , the larger the difference in performance between the MBMW and OBMW detectors. For instance, with respect to the OBMW, the MBMW incurs losses of 0.15, 0.53, 1.80, and 2.45 dB for $N = 8, 16, 24,$ and 32 pulses, respectively.

Batch Processor. The batch processor (Fig. 8.3e) is very useful when a large number of pulses are in the 3-dB beamwidth. If KN pulses are in the 3-dB beamwidth, K pulses are summed (batched) and either a 0 or a 1 is declared, depending on whether or not the batch is less than a threshold T_1 . The last N zeros and ones are summed and compared with a second threshold M . An alternative version of this detector is to put the batches through a moving-window detector.

The batch processor, like the binary integrator, is easily implemented, ignores interference spikes, and works extremely well when the noise has a non-Rayleigh

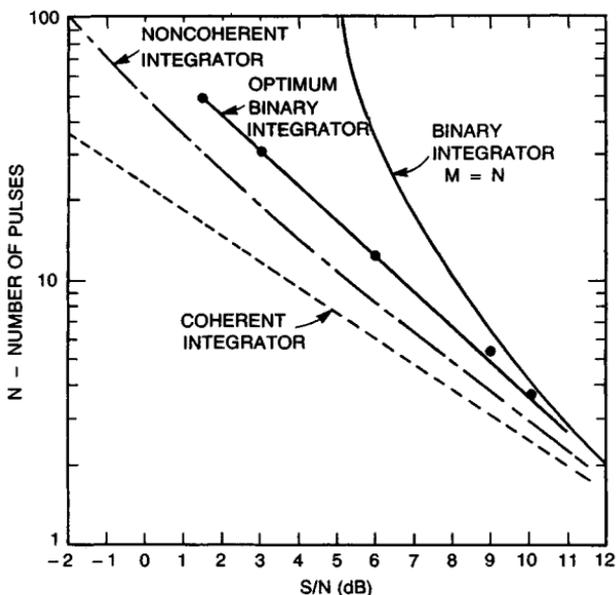


FIG. 8.9 Comparison of binary integrator (M out of N) with other integration methods ($P_{fa} = 10^{-10}$; $P_D = 0.5$). (Copyright 1956, IEEE; after Ref. 16.)

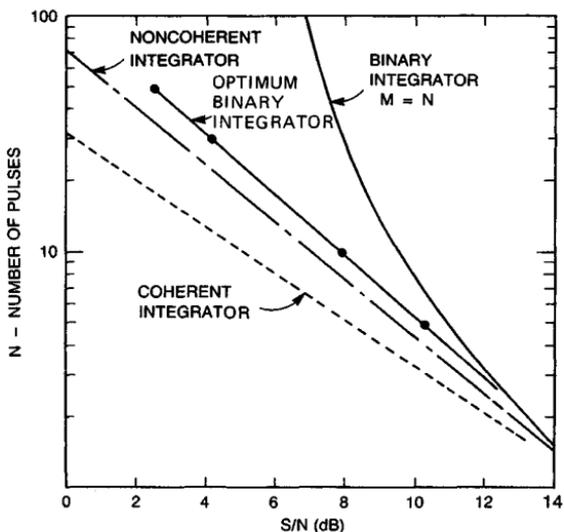


FIG. 8.10 Comparison of binary integrator (M out of N) with other integration methods ($P_{fa} = 10^{-10}$; $P_D = 0.90$). (Copyright 1956, IEEE; after Ref. 16.)

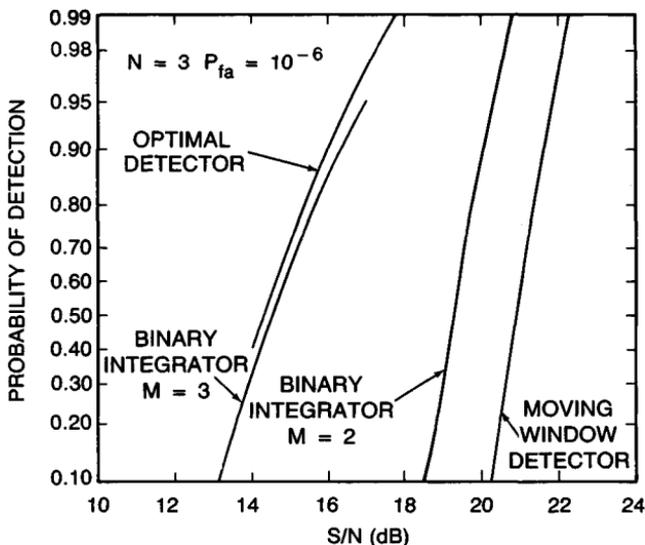


FIG. 8.11 Comparison of various detectors in log-normal ($\sigma = 6$ dB) interference ($N = 3$; $P_{fa} = 10^{-6}$). (Copyright 1975, IEEE; after Ref. 19.)

density. Furthermore, the batch processor requires less storage, detects better, and estimates angles more accurately than the binary integrator. For instance, if there were 80 pulses on target, one could batch 16 pulses, quantize this result to a 0 or a 1, and declare a target with a 3-out-of-5 (or 2-out-of-5) binary integrator. With an 8-bit analog-to-digital converter, the storage requirement per range cell is 17 bits (12 bits for the batch and 5 for the binary integrator) for the batch processor as opposed to 80 bits for the binary integrator and 640 bits for the moving window. The detection performance of the batch processor for a large number of pulses integrated is approximately 0.5 dB worse than the moving window. The batch processor has been successfully implemented by the Applied Physics Laboratory²¹ of Johns Hopkins University. To obtain an accurate azimuth estimate $\hat{\theta}$, approximately 20 percent greater than the lower bound,

$$\hat{\theta} = \frac{\sum B_i \theta_i}{\sum B_i} \quad (8.16)$$

is used, where B_i is the batch amplitude and θ_i is the azimuth angle corresponding to the center of the batch.

False-Alarm Control. In the presence of clutter, if fixed thresholds are used with the previously discussed integrators, an enormous number of detections will occur and will saturate and disrupt the tracking computer associated with the radar system. Four important facts should be noted:

- A tracking system should be associated with the automatic detection system (the only exception is when one displays multiple scans of detections).
- The P_{fa} of the detector should be as high as possible without saturating the tracking computer.

- Random false alarms and unwanted targets (e.g., stationary targets) are not a problem if they are removed by the tracking computer.
- Scan-to-scan processing can be used to remove stationary point clutter or moving-target indication (MTI) clutter residues.

One can limit the number of false alarms with a fixed-threshold system by setting a very high threshold. Unfortunately, this would reduce target sensitivity in regions of low noise (clutter) return. Three main approaches—adaptive threshold, nonparametric detectors, and clutter maps—have been used to reduce the false-alarm problem. Adaptive thresholding and nonparametric detectors assume that the samples in the range cells surrounding the test cell (called *reference cells*) are independent and identically distributed. Furthermore, it is usually assumed that the time samples are independent. Both kinds of detectors test whether the test cell has a return sufficiently larger than the reference cells. Clutter maps allow variation in space, but the clutter must be stationary over several (typically 5 to 10) scans. Clutter maps store an average background level for each range-azimuth cell. A target is then declared in a range-azimuth cell if the new value exceeds the average background level by a specified amount.

Adaptive Thresholding. The basic assumption of the adaptive thresholding technique is that the probability density of the noise is known except for a few unknown parameters. The surrounding reference cells are then used to estimate the unknown parameters, and a threshold based on the estimated parameters is obtained. The simplest adaptive detector, shown in Fig. 8.12, is the cell-averaging CFAR (constant false-alarm rate) investigated by Finn and Johnson.²² If the noise has a Rayleigh density, $p(x) = x \exp(-x^2/2\sigma^2)/\sigma^2$, only the parameter σ (σ^2 is the noise power) needs to be estimated, and the threshold is of the form $T = K\sum x_i = Kn\sqrt{\pi/2}\hat{\sigma}$, where $\hat{\sigma}$ is the estimate of σ . However, since T is set by an estimate $\hat{\sigma}$, it has some error and must be slightly larger than the threshold that one would use if σ were known exactly a priori. The raised threshold causes a loss in target sensitivity and is referred to as a CFAR loss. This loss has been

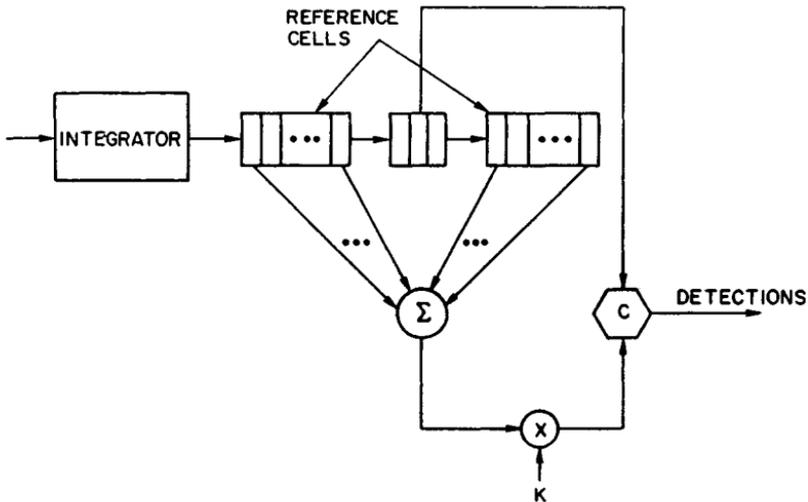


FIG. 8.12 Cell-averaging CFAR. The letter C indicates a comparison. (From Ref. 7.)

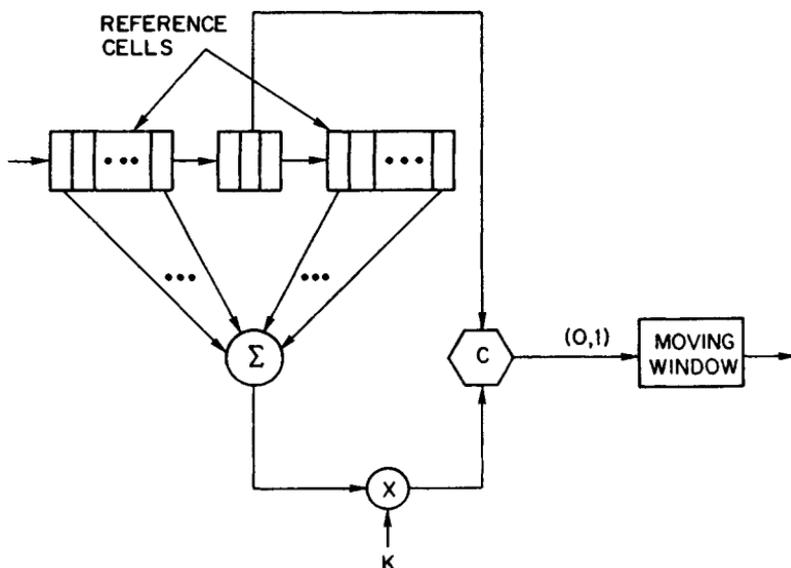
TABLE 8.1 CFAR Loss for $P_{fa} = 10^{-6}$ and $P_D = 0.9^*$

Number of pulses integrated	Loss for various numbers of reference cells, dB					
	1	2	3	5	10	∞
1	...	15.3	7.7	3.5	0	
3	...	7.8	5.1	3.1	1.4	0
10	6.3	3.3	2.2	1.3	0.7	0
30	3.6	2.0	1.4	1.0	0.5	0
100	2.4	1.4	1.0	0.6	0.3	0

*After Ref. 23.

calculated²³ and is summarized in Table 8.1. As can be seen, for a small number of reference cells the loss is large because of the poor estimate of σ . Consequently, one would prefer to use a large number of reference cells. However, if one does this, the homogeneity assumption (i.e., all the reference cells are statistically similar) might be violated. A good rule of thumb is to use enough reference cells so that the CFAR loss is below 1 dB and at the same time not let the reference cells extend beyond 1 nmi on either side of the test cell. For a particular radar this might not be feasible.

If there is uncertainty about whether or not the noise is Rayleigh-distributed, it is better to threshold individual pulses and use a binary integrator as shown in Fig. 8.13. This detector is tolerant of variations in the noise density because by setting K to yield a 1 with probability 0.1, a $P_{fa} \approx 10^{-6}$ can be obtained by using a 7-out-of-9 detector. While noise may be non-Rayleigh, it will probably be very Rayleigh-like out to the tenth percentile. Furthermore, one can use feedback based on several

**FIG. 8.13** Implementation of a binary integrator. The letter C indicates a comparison. (From Ref. 7.)

scans of data to control K in order to maintain a desired P_{fa} either on a scan or a sector basis. This demonstrates a general rule: to maintain a low P_{fa} in various environments, adaptive thresholding should be placed in front of the integrator.

If the noise power varies from pulse to pulse (as it would in jamming when frequency agility is employed), one must CFAR each pulse and then integrate. While the binary integrator performs this type of CFAR action, analysis^{24,25} has shown that the ratio detector in Fig. 8.14 is a better detector. The ratio detector sums signal-to-noise ratios and is specified by

$$\sum_{i=1}^n \frac{x_i^2(j)}{\frac{1}{2m} \sum_{k=1}^m [x_i^2(j+1+k) + x_i^2(j-1-k)]} \quad (8.17)$$

where $x_i(j)$ is the i th envelope-detected pulse in the j th range cell and $2m$ is the number of reference cells. The denominator is the maximum-likelihood estimate of σ_i^2 , the noise power per pulse. It will detect targets even though only a few returned pulses have a high signal-to-noise ratio. Unfortunately, this will also cause the ratio detector to declare false alarms in the presence of narrow-pulse interference. To reduce the number of false alarms when narrow-pulse interference is present, the individual power ratios can be soft-limited²⁵ to a small enough value so that interference will cause only a few false alarms. A comparison of the ratio detector with other commonly used detectors is shown in Figs. 8.15 and 8.16 for nonfluctuating and fluctuating targets. A typical performance in sidelobe jamming when the jamming level varies by 20 dB per pulse is shown in Fig. 8.17. By employing a second test to identify the presence of narrow-pulse interference, a detection performance approximately halfway between the limiting and nonlimiting ratio detectors can be obtained.

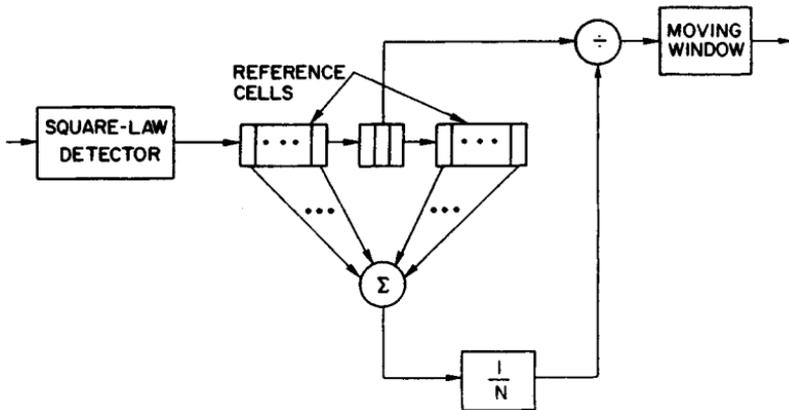


FIG. 8.14 Ratio detector. (From Ref. 7.)

If the noise samples are dependent in time or have a non-Rayleigh density such as the chi-square density or log-normal density, it is necessary to estimate two parameters and the adaptive detector is more complicated. Usually several pulses are integrated so that one can assume the integrated output has a gaussian probability density. Then the two parameters that must be estimated are the

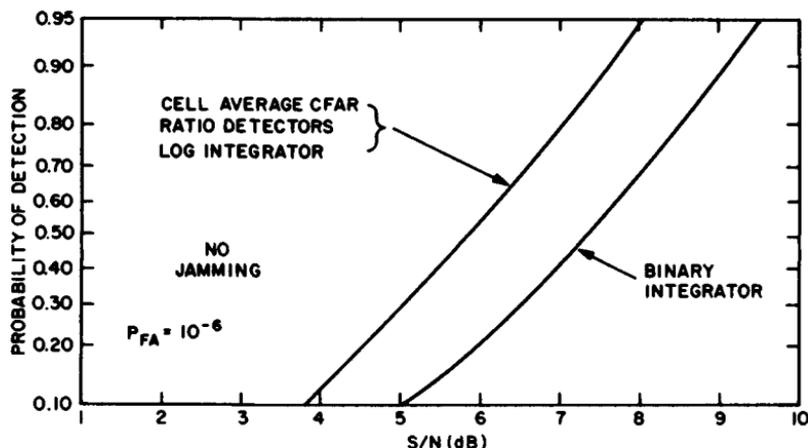


FIG. 8.15 Curves of probability of detection versus signal-to-noise ratio per pulse for the cell-averaging CFAR, ratio detectors, log integrator, and binary integrator: nonfluctuating target, $N = 6$, and probability of false alarm = 10^{-6} . (From Ref. 25.)

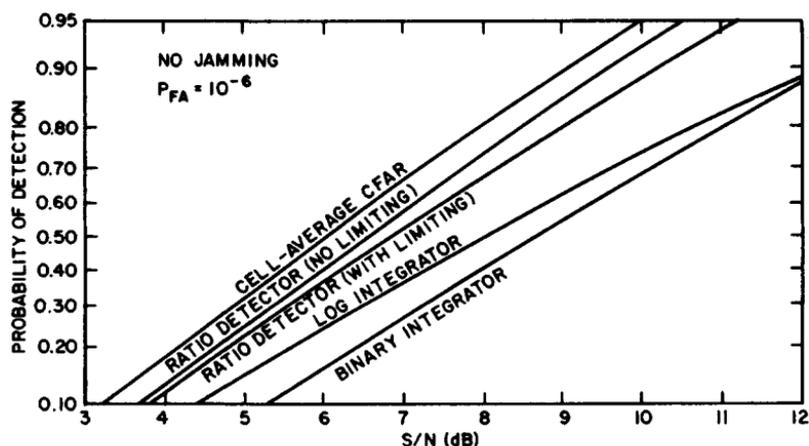


FIG. 8.16 Curves of probability of detection versus signal-to-noise ratio for the cell-averaging CFAR, ratio detectors, log integrator, and binary integrator: Rayleigh, pulse-to-pulse fluctuating target, $N = 6$, and probability of false alarm = 10^{-6} . (From Ref. 25.)

mean and the variance, and a threshold of the form $T = \hat{\mu} + K\hat{\sigma}$ is used. Though the mean is easily obtained in hardware, the usual estimate of the standard deviation

$$\hat{\sigma} = \left[\frac{1}{N} \sum (x_i - \bar{x})^2 \right]^{1/2} \quad (8.18)$$

where

$$x = \frac{1}{N} \sum x_i \quad (8.19)$$

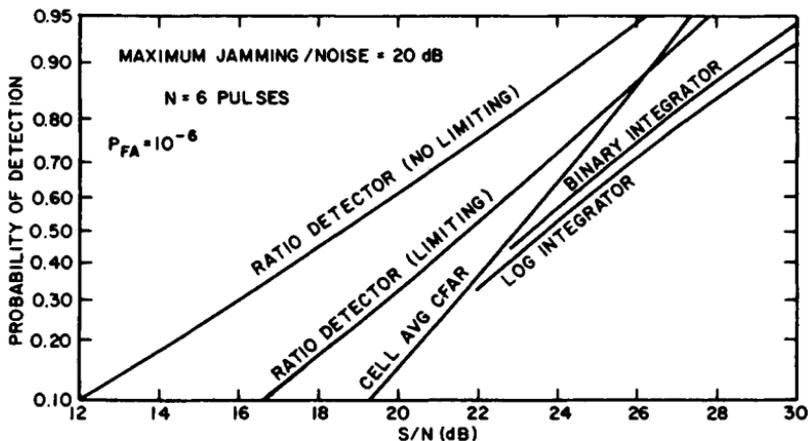


FIG. 8.17 Curves of probability of detection versus signal-to-noise ratio for the cell-averaging CFAR, ratio detectors, log integrator, and binary integrator: Rayleigh, pulse-to-pulse fluctuations, probability of false alarm = 10^{-6} , and maximum jamming-to-noise ratio = 20 dB. (From Ref. 25.)

is more difficult to implement. Consequently, the mean deviate defined by

$$\sigma = A \sum |x_i - \bar{x}| \quad (8.20)$$

is sometimes used because of its ease of implementation. Nothing can be done to the binary integrator to yield a low P_{fa} if the noise samples are correlated. Thus, it should not be used in this situation. However, if the correlation time is less than a batching interval, the batch processor will yield a low P_{fa} without modifications.

Target Suppression. Target suppression is the loss in detectability caused by other targets or clutter residues in the reference cells. Basically, there are two approaches to solving this problem: (1) remove large returns from the calculation of the threshold,²⁶⁻²⁸ or (2) diminish the effects of large returns by either limiting or using log video. The technique that should be used is a function of the particular radar system and its environment.

Rickard and Dillard²⁷ proposed a class of detectors D_K , where the K largest samples are censored (removed) from the reference cells. A comparison of D_0 (no censoring) with D_1 and D_2 for a Swerling 2 target and a single square-law detected pulse is shown in Fig. 8.18, where N is the number of reference cells, β is the ratio of the power of the interfering target to the target in the test cell, and the bracketed pair (m, n) indicates the Swerling models of the target and the interfering target, respectively. As shown in Fig. 8.18, when one has an interfering target, the P_D does not approach 1 as S/N increases. Another approach²⁶ which censors samples in the reference cell if they exceed a threshold is briefly discussed in the subsection "Nonparametric Detectors."

Finn²⁸ investigated the problem of the reference cells spanning two continuous different "noise" fields (e.g., thermal noise, sea clutter, land clutter, etc.). On the basis of the samples, he estimated the statistical parameters of the two noise fields and the separation point between them. Then, only those reference cells which are in the noise field containing the test cell are used to calculate the adaptive threshold.

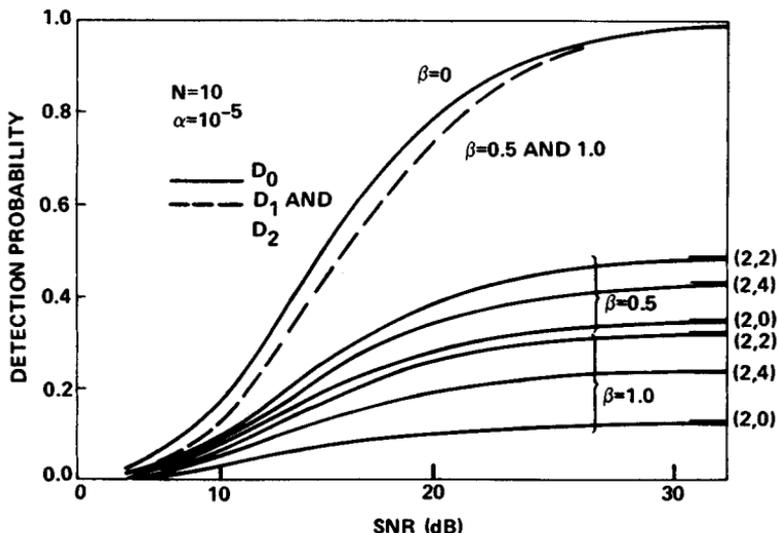


FIG. 8.18 Detection probability versus SNR for Swerling Case 2 primary target. (Copyright 1977, IEEE; from Ref. 27.)

An alternative approach for interfering targets is to use log video. By taking the log, large samples in the reference cells will have little effect on the threshold. The loss associated with using log video is 0.5 dB for 10 pulses integrated and 1.0 dB for 100 pulses integrated.²⁹ An implementation of the log CFAR³⁰ is shown in Fig. 8.19. In many systems the antilog shown in Fig. 8.19 is not taken. To maintain the same CFAR loss as for linear video, the number of reference cell M_{\log} for the log CFAR should equal

$$M_{\log} = 1.65 M_{\text{lin}} - 0.65 \quad (8.21)$$

where M_{lin} is the number of reference cells for linear video. The effect of target suppression with log video is discussed later in this section (Table 8.2).

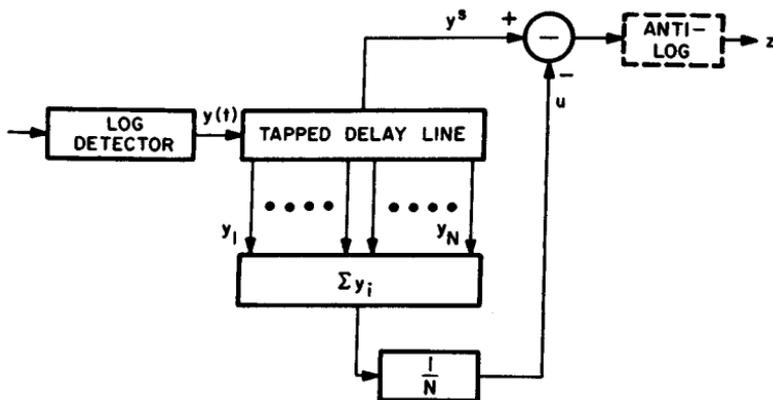


FIG. 8.19 Block diagram of cell-averaging log-CFAR receiver. (Copyright 1972, IEEE; from Ref. 30.)

Nonparametric Detectors. Usually nonparametric detectors obtain CFAR by ranking the test sample with the reference cells.^{31,32} Ranking means that one orders the samples from the smallest to the largest and replaces the smallest with rank 0, the next smallest with rank 1, ..., and the largest with rank $n-1$. Under the hypothesis that all the samples are independent samples from an unknown density function, the test sample has equal probability of taking on any of the n values. For instance, referring to the ranker in Fig. 8.20, the test cell is compared with 15 of its neighbors. Since in the set of 16 samples, the test sample has equal probability of being the smallest sample (or equivalently any other rank), the probability that the test sample takes on values 0, 1, ..., 15 is 1:16. A simple rank detector is constructed by comparing the rank with a threshold K and generating a 1 if the rank is larger, a 0 otherwise. The 0s and 1s are summed in a moving window. This detector incurs a CFAR loss of about 2 dB but achieves a fixed P_{fa} for any unknown noise density as long as the time samples are independent. This detector was incorporated into the ARTS-3A postprocessor used in conjunction with the Federal Aviation Administration airport surveillance radar (ASR). The major shortcoming of this detector is that it is fairly susceptible to target suppression (e.g., if a large target is in the reference cells, the test cell cannot receive the highest ranks).

If the time samples are correlated, the rank detector will not yield CFAR. A modified rank detector, called the modified generalized sign test (MGST),²⁶ maintains a low P_{fa} and is shown in Fig. 8.21. This detector can be divided into three parts: a ranker, an integrator (in this case a two-pole filter), and a threshold (decision process). A target is declared when the integrated output exceeds two thresholds. The first threshold is fixed (equals $\mu + T_1/K$ in Fig. 8.21) and yields $P_{fa} = 10^{-6}$ when the reference cells are independent and identically distributed. The second threshold is

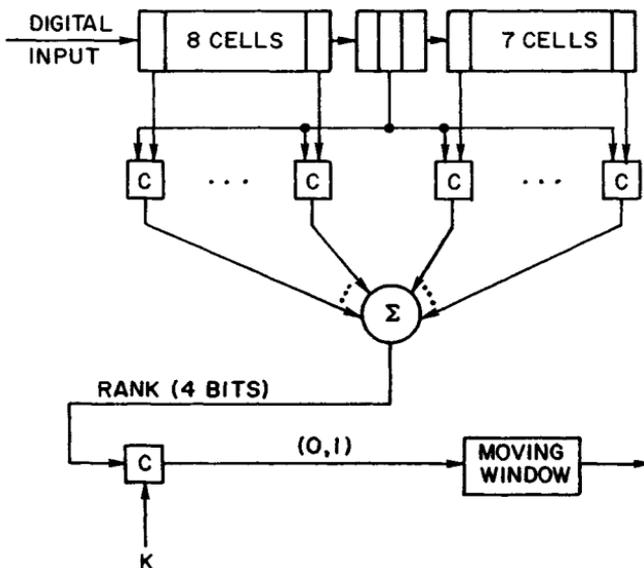


FIG. 8.20 Rank detector: output of a comparator C is either a zero or a one. (From Ref. 7.)

adaptive and maintains a low P_{fa} when the reference samples are correlated. The device estimates the standard deviation of the correlated samples with the mean deviate estimator, where extraneous targets in the reference cells have been excluded from the estimate by use of a preliminary threshold T_2 .

The rank and MGST detectors are basically two-sample detectors. They decide that a target is present if the ranks of the test cell are significantly greater than the ranks of the reference cells. Target suppression occurs at all interfaces (e.g., land, sea) where the homogeneity assumption is violated. However, some tests exist, such as the Spearman Rho and Kendall Tau tests,³³ that depend only on the test cell. These tests use the fact that as the antenna beam sweeps by a point target, the signal return increases and then decreases. Thus, for the test cell the ranks should follow a pattern, first increasing and then decreasing. Although these detectors do not require reference cells and hence have the useful property of not requiring homogeneity, they are not generally used because of the large CFAR losses that occurs for moderate sample sizes. For instance, the CFAR losses are approximately 10 dB for 16 pulses on target and 6 dB for 32 pulses on target.³³

The basic disadvantages of all nonparametric detectors are that (1) they have relatively large CFAR losses, (2) they have problems with correlated samples, and (3) one loses amplitude information, which can be a very important discriminant between target and clutter.³⁴ For example, a large return ($\sigma \approx 1000 \text{ m}^2$) in a clutter area is probably just clutter breakthrough. See "Contact Entry Logic" in Sec. 8.3.

Clutter Mapping. A clutter map uses adaptive thresholding where the threshold is calculated from the return in the test cell on previous scans rather than from the surrounding reference cells on the same scan. This technique has

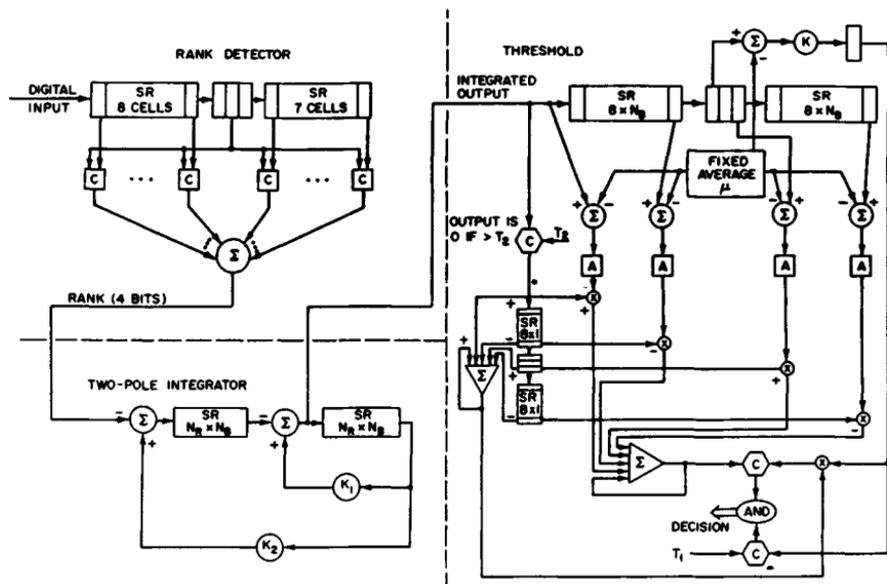


FIG. 8.21 Modified generalized sign test processor. (Copyright 1974, IEEE; from Ref. 26.)

the advantage that for essentially stationary environments (e.g., land-based radar against ground clutter), the radar has interclutter visibility—it can see between large clutter returns. Lincoln Laboratory³⁵ in its moving-target detector (MTD) used a clutter map for the zero-doppler filter very effectively. The decision threshold T for the i th cell is

$$T = A S_{i-1} \quad (8.22)$$

where

$$S_i = K S_{i-1} + X_i \quad (8.23)$$

S_i is the average background level, X_i is the return in the i th cell, K is the feedback value which determines the map time constant, and A is the constant which determines the false-alarm rate. In the MTD used for ASR application K is 7:8, which effectively averages the last eight scans. The main utility of clutter maps is with fixed-frequency land-based radars. While clutter maps can be used with frequency-agile radars and on moving platforms, they are not nearly as effective in these environments.

Target Resolution. In automatic detection systems, a single large target will probably be detected many times, e.g., in adjacent range cells, azimuth beams, and elevation beams. Therefore, automatic detection systems have algorithms for merging the individual detections into a single centroided detection. Most algorithms have been designed so that they will rarely split a single target into two targets. This procedure results in poor range resolution capability. A merging algorithm³⁶ often used is the adjacent-detection merging algorithm, which decides whether a new detection is adjacent to any of the previously determined sets of adjacent detections. If the new detection is adjacent to any detection in the set of adjacent detections, it is added to the set. Two detections are adjacent if two of their three parameters (range, azimuth, and elevation) are the same and the other parameter differs by the resolution element: range cell ΔR , azimuth beamwidth θ , or elevation beamwidth γ .

A simulation³⁶ was run to compare the resolving capability of three common detection procedures: linear detector with $T = \hat{\mu} + A\hat{\sigma}$, linear detector with $T = B\hat{\mu}$, and log detector with $T = C + \hat{\mu}$. The constants A , B , and C are used to obtain the same P_{fa} for all detectors. The estimates $\hat{\mu}$ and $\hat{\sigma}$ of μ and σ were obtained from either (1) all the reference cells or (2) the leading or lagging half of the reference cells, choosing the half with the lower mean value. The simulation involved two targets separated by 1.5, 2.0, 2.5, or 3.0 range cells and a third target 7.0 range cells from the first target. When the two closely spaced targets were well separated, either 2.5 or 3.0 range cells apart, the probability of detecting both targets (P_{D2}) was < 0.05 for the linear detector with $T = \hat{\mu} + A\hat{\sigma}$; $0.15 < P_{D2} < 0.75$ for the linear detector with $T = B\hat{\mu}$; and $P_{D2} > 0.9$ for the log detector. A second simulation, involving only two targets, investigated the effect of target suppression on log video, and the results are summarized in Table 8.2. One notes an improved performance for small S/N (10 to 13 dB) when one calculates the threshold using only the half of the reference cells with the lower mean value. The resolution capability of the log detector which uses only the half of the reference cells with the lower mean is shown in Fig. 8.22. The probability of resolving two equal-amplitude targets does not rise above 0.9 until they are separated in range by 2.5 pulse widths.

By assuming that the target is small with respect to the pulse width and that the pulse shape is known, the resolution capability can be improved by fitting

TABLE 8.2 Probability of Detecting with Log Video Two Targets Separated by 1.5, 2.0, 2.5, or 3.0 Range Cells*

Thresholding technique	Target separation	S/N of target no. 2				
		10	13	20	30	40
All reference cells	1.5	0.0	0.04	0.0	0.00	0.00
	2.0	0.0	0.22	0.54	0.14	0.10
	2.5	0.04	0.24	0.94	0.62	0.32
	3.0	0.0	0.24	0.88	0.92	0.76
Reference cells with minimum mean value	1.5	0.0	0.0	0.00	0.0	0.02
	2.0	0.10	0.32	0.44	0.12	0.04
	2.5	0.18	0.58	0.98	0.46	0.28
	3.0	0.22	0.66	0.98	0.82	0.74

* S/N of target 1 is 20 dB. S/N of target 2 is 10, 13, 20, 30, or 40 dB. After Ref. 36.

the known pulse shape to the received data and comparing the residue square error with a threshold.³⁷ If only one target is present, the residue should be only noise and hence should be small. If two or more targets are present, the residue will contain signal from the remaining targets and should be large. The results of resolving two targets with $S/N = 20$ dB are shown in Fig. 8.23. These targets can be resolved at a resolution probability of 0.9 at separations

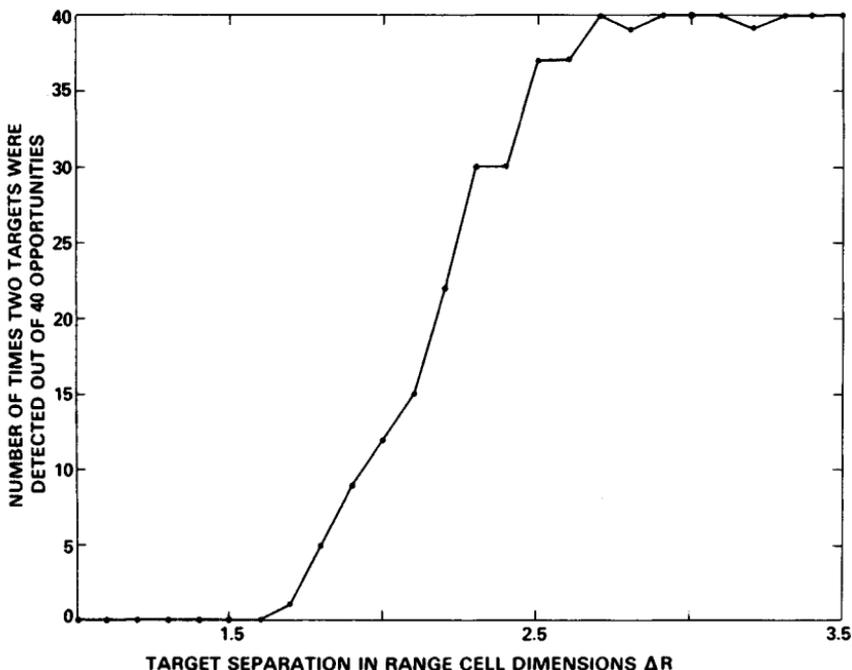


FIG. 8.22 Resolution capability of a log detector which uses the half of the reference cells with the lower mean. (Copyright 1978, IEEE; from Ref. 36.)

varying between one-fourth and three-fourths of a pulse width, depending on the relative phase difference between the two targets. Furthermore, this result can be improved further by processing multiple pulses.

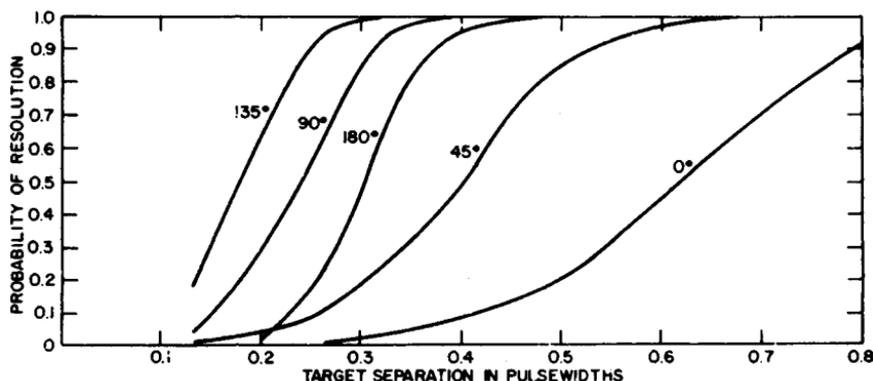


FIG. 8.23 Probability of resolution as a function of range separation: sampling rate $\Delta R = 1.5$ samples per pulse width; target strengths—nonfluctuating, $A_1 = A_2 = 20$ dB; phase differences = 0° , 45° , 90° , 135° , and 180° . (Copyright 1984, IEEE; from Ref. 37.)

Detection Summary. When only 2 to 4 samples (pulses) are available, a binary integrator should be used to avoid false alarms due to interference. When a moderate number of pulses (5 to 16) are available, a binary integrator, a rank detector, or a moving-window integrator should be used. If the number of pulses is large (greater than 20), a batch processor or a two-pole filter should be used. If the samples are independent, a one-parameter (mean) threshold can be used. If the samples are dependent, one can either use a two-parameter (mean and variance) threshold or adapt a one-parameter threshold on a sector basis. These rules should serve only as a general guideline. It is *highly recommended* that before a detector is chosen the radar video from the environment of interest be collected and analyzed and that various detection processes be simulated on a computer and tested against the recorded data.

8.3 AUTOMATIC TRACKING

Track-while-scan (TWS) systems are tracking systems for surveillance radars whose nominal scan time (revisit time) is from 4 to 12 s for aircraft targets. If the probability of detection (P_D) per scan is high, if accurate target location measurements are made, if the target density is low, and if there are only a few false alarms, the design of the correlation logic (i.e., associating detections with tracks) and tracking filter (i.e., filter for smoothing and predicting track positions) is straightforward. However, in a realistic radar environment these assumptions are seldom valid, and the design of the automatic tracking system is complicated. In actual situations one encounters target fades (changes in signal strength due to multipath propagation, blind speeds, and atmospheric conditions), false alarms (due to noise, clutter, interference, and jamming), and poor radar parameter estimates (due to noise, unstabilized antennas, unresolved targets, target splits,