

Parametric Study of a Nonsuppressing CFAR Detector

J. D. WILSON

*Radar Analysis Branch
Radar Division*

February 25, 1983



NAVAL RESEARCH LABORATORY
Washington, D.C.

Approved for public release; distribution unlimited.

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER NRL Report 8671	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) PARAMETRIC STUDY OF A NONSUPPRESSING CFAR DETECTOR	5. TYPE OF REPORT & PERIOD COVERED Interim report on one phase of an NRL problem.	
	6. PERFORMING ORG. REPORT NUMBER	
7. AUTHOR(s) J. D. Wilson	8. CONTRACT OR GRANT NUMBER(s)	
9. PERFORMING ORGANIZATION NAME AND ADDRESS Naval Research Laboratory Washington, DC 20375	10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS 62712N; SF12141491 53-0609-00	
11. CONTROLLING OFFICE NAME AND ADDRESS Naval Sea Systems Command Washington, DC 20362	12. REPORT DATE February 25, 1983	
	13. NUMBER OF PAGES 15	
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)	15. SECURITY CLASS. (of this report) UNCLASSIFIED	
	15a. DECLASSIFICATION/DOWNGRADING SCHEDULE	
16. DISTRIBUTION STATEMENT (of this Report) Approved for public release; distribution unlimited.		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) Detection CFAR Nonsuppressing		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) This report describes a parametric analysis of a modified conventional constant false alarm rate (CFAR) detector. The modified CFAR detector is nonsuppressing in the sense that targets or other interfering signals in the reference cells of the detector will not <i>suppress</i> or prevent the detection of targets. The modification consists of dividing the set of reference cells used for determining the detection threshold in a conventional CFAR into <i>N</i> batches. Each batch is used to set a detection threshold and the test cell is compared to all the thresholds. If the test cell exceeds <i>M</i> or more of the <i>N</i> thresholds, a detection is declared. The values selected for the various parameters (Continued)		

DD FORM 1473
1 JAN 73

EDITION OF 1 NOV 65 IS OBSOLETE
S/N 0102-014-6601

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

20. ABSTRACT (Continued)

by the study were $M = 4$ and $N = 6$ batches of 4 samples each. With these values, the nonsuppressing CFAR required $\cong 1$ dB higher signal-to-noise ratio to achieve detection performance equivalent to the conventional CFAR in the absence of interfering targets. On the other hand, a large target in the reference cells of a conventional CFAR could completely suppress detection, while the same amount of interference causes only minor degradation in performance of the nonsuppressing CFAR detector.

CONTENTS

INTRODUCTION	1
DESCRIPTION OF CFAR DETECTORS	2
DESCRIPTION OF SIMULATION AND THRESHOLD DETERMINATION	4
DETECTION PERFORMANCE IN ABSENCE OF INTERFERENCE	7
DETECTION PERFORMANCE IN PRESENCE OF INTERFERENCE	9
SUMMARY	10
REFERENCES	11

PARAMETRIC STUDY OF A NONSUPPRESSING CFAR DETECTOR

INTRODUCTION

Historically, men were used to extract the presence and location of targets by observing radar displays. After the development of integrated circuits, the construction of devices to automatically detect and measure target locations became feasible, and automatic detectors were built to achieve fast reaction and massive data handling capability. A brief discussion of automatic processing is given in the "Radar Handbook" edited by Skolnik [1]. One of the practical detectors used in many applications is commonly termed a constant false alarm rate (CFAR) detector [2].

The CFAR detector basically operates as follows. For each transmitted pulse, the returned signal is sampled at the range resolution of the radar and each sample is compared to a threshold. This threshold is derived from the signal sampled over range intervals surrounding the test cell. These samples are called reference cells. Signals above the threshold are declared detections. Basically, the CFAR detector looks for a large amplitude signal existing over a short time interval relative to a surrounding background of much lower amplitude.

The performance of CFAR detectors has been evaluated both analytically and experimentally. Finn and Johnson [3] studied the performance of a simple CFAR detector for a single target in additive noise against a noise background as a function of the number of reference cells used. When the noise is identically Rayleigh-distributed across the reference cells and independent from reference cell to reference cell, the loss compared to optimum performance is on the order of less than 1 dB for the number of reference cells normally used. Experimentally many types of CFAR detectors have operated satisfactorily in a variety of environments using modest changes in thresholds according to the environment [4]. For example, detections are usually not made by the CFAR detector in radar signals containing heavy clutter because the statistic derived from the reference cells becomes large, which in turn creates a large threshold. However, the threshold must be increased slightly in these environments to maintain desired false-alarm rates. One reason why no single threshold yields a truly constant false-alarm rate is because the reference cells are often nonhomogeneous. Furthermore, the probability distributions may differ from the distribution used to set the thresholds. However, as previously noted, adjusting the threshold according to the environment usually will result in adequate performance.

An important case in which the performance of conventional CFAR detectors degrades is when either a target or interference spike is present in a reference cell and all other reference cells contain identically distributed independent noise. In this case the threshold derived from the reference cells is much too high; consequently, targets which would normally be detected are suppressed. To keep this from occurring, others [5,6] considered techniques of eliminating the reference cells containing the contaminating signal (target or interference spike) from the calculation of the threshold. Trunk, Cantrell, and Queen [5] considered removing all signals above a fixed threshold from being used in the threshold calculation. The detector used a ranking procedure so that the fixed threshold could be set and used. However, in many applications it is undesirable to lose the amplitude information through ranking and, consequently, this procedure cannot be used. Rickard and Dillard [6] describe a procedure for removing the largest signals in the reference cells from the calculation of the threshold. Good results were obtained. Unfortunately, the ranking circuits for locating the reference cells containing the largest amplitudes are cumbersome and somewhat complex making the implementation unattractive.

Another means of circumventing the losses which occur in CFAR detectors when another target appears in the reference cells is to use a log CFAR detector. The log CFAR detector takes the log of the time ordered amplitude cells $X(i)$ before the signal is passed through a conventional CFAR processor. Trunk [7] shows the detector works fairly well when several targets contaminate the reference cells. The log CFAR detector computes the threshold from the geometric mean rather than the arithmetic mean used in the conventional linear CFAR detector. Because the geometric mean is less sensitive than the arithmetic mean to a few large amplitudes, the detection losses are less in the log CFAR detector than in the conventional linear CFAR detector when the reference cells are contaminated. However, the log CFAR detector's threshold is very sensitive and is difficult to set to give good operation over a wide variety of signals.

The purpose of this report is to describe a CFAR detector which is easy to implement and which will detect targets with little loss in the presence of a contaminating signals in the reference cells. Its performance is evaluated in a conventional way through simulation. First, the thresholds for fixed false-alarm rates are set. The probability of detection is then evaluated in the absence and the presence of contaminating signals in the reference cells.

The basic idea of this nonsuppressing CFAR detector is to divide the set of reference cells into batches and estimate a threshold based on each batch. A contaminating signal in one of the reference cells will only contaminate the threshold based on the batch containing the extraneous signal. By only requiring the test cell to exceed M out of N thresholds, we may still detect (with some loss) in the presence of at least $N - M$ contaminating signals.

DESCRIPTION OF CFAR DETECTORS

The conventional and the nonsuppressing CFARs are discussed in this section. Figure 1 depicts the conventional CFAR. The signal is fed into a tapped delay line where the clock rate corresponds to the range resolution of the radar. As the signal propagates through the delay line, each sample or cell is tapped off to an appropriate destination. The first $N_R/2$ and the last $N_R/2$ cells (reference cells) are fed to an adder which forms the sum. The cells referred to as guard cells serve only to isolate the test cell from the reference cells. The number, G , of guard cells on each side of the test cell will be a function of the resolution of the radar and the expected size of the target, and are there to keep a target with appreciable return in more than one cell from being included in the reference cells and suppressing itself. The sum of the reference cells can be represented mathematically by

$$Y(i) = \sum_{j=1}^{N_R/2} [X(i - G - j) + X(i + G + j)], \quad (1)$$

where $X(K)$ refers to the sampled value of the K th range cell, and N_R and G are defined as above. For each i such that

$$N_R/2 + G < i \leq N_T - N_R/2 - G,$$

where N_T is the total number of range cells in a range sweep. There is a test cell, $X(i)$, which we wish to test for the presence of a target, and a value $Y(i)/N_R$, which is an estimate of the environment surrounding $X(i)$. Conventional CFAR operation can be expressed mathematically by declaring a detection in range cell i when

$$K'X(i) \geq Y(i). \quad (2)$$

K' is a multiplicative factor (which includes the number of reference cells, N_R) determined by either analysis, simulation, or experimentation to yield the desired probability of false alarm, P_{fa} .

Figure 2 depicts the nonsuppressing CFAR. Again the nonsuppressing CFAR employs a tapped delay line, but the individual samples in the reference cells feed into a number of adders, each of which forms a sum of a contiguous subset of the reference cells.

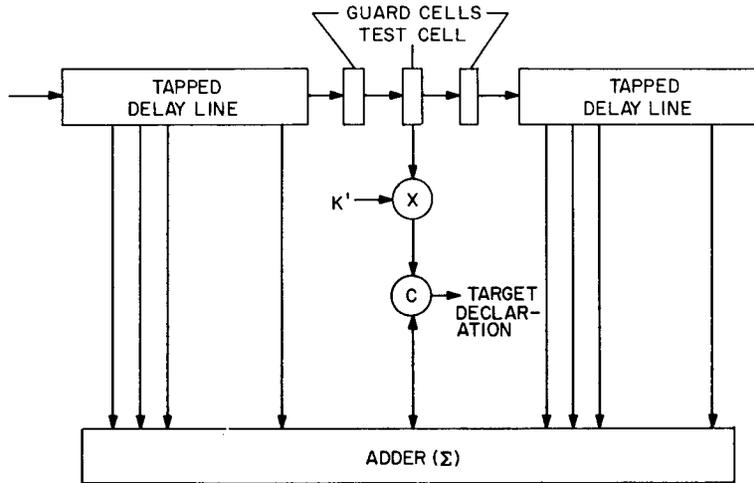


Fig. 1 - Conventional CFAR

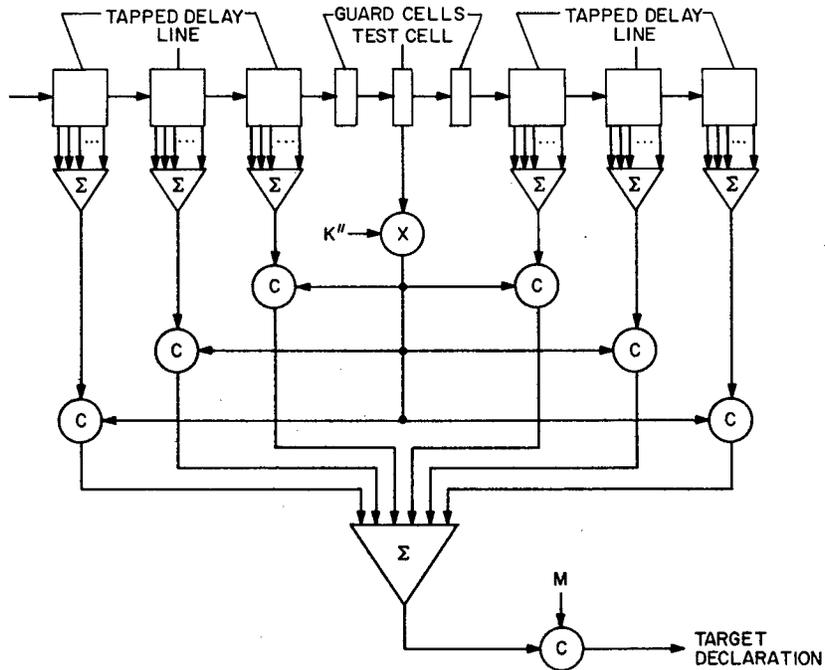


Fig. 2 - Nonsuppressing CFAR

For this detector, a detection in the i th range cell is declared if

$$\sum_{j=1}^{N_A} \delta_j(i) \geq M, \tag{3}$$

where M is an integer threshold and

$$\delta_j(i) = \begin{cases} 1 & \text{if } K''X(i) \geq Y_j(i) \\ 0 & \text{otherwise,} \end{cases} \tag{4}$$

with K'' being a multiplicative factor (including the number of cells added in each adder), $Y_j(i)$ is the sum in the j th adder and N_A is the number of adders.

In Fig. 1 and Fig. 2, the time-ordered sequence of range cell amplitudes is clocked through the delay line from left to right. The appropriate sums of the amplitude cells are performed each clock cycle. The test cell is multiplied by a factor K' or K'' to avoid a division of the sums. The remaining operations are simple comparators. In some applications it may be advantageous to implement the moving window integrator given by Eq. (1) in a recursive form as described in Ref. 5. The recursive implementation can yield a significant hardware savings when N_R is large.

DESCRIPTION OF SIMULATION AND THRESHOLD DETERMINATION

To perform an evaluation of the nonsuppressing CFAR detector by purely analytical means would be very difficult if not impossible. With the availability of a large digital computer, the numerical determination of probability of false alarm (P_{fa}) and probability of detection (P_d) of the CFAR is possible through simulation.

The signals generated for this simulation are linearly detected Gaussian noise for the interference environment and a constant level signal for the target. The random numbers generated by the computer are uniformly distributed between 0 and 1. The Gaussian distributed random variables G_1 and G_2 needed can be generated from these uniformly distributed random variables U_i by

$$G_1 = \sigma \sqrt{-2 \ln(U_1)} \sin(2\pi U_2), \quad (5)$$

and

$$G_2 = \sigma \sqrt{-2 \ln(U_1)} \cos(2\pi U_2),$$

where σ is the desired standard deviation of the Gaussian random signal. The linearly detected signal

$$R = \sqrt{G_1^2 + G_2^2} \quad (6)$$

is Rayleigh distributed with mean μ_R and standard deviation, σ_R given by

$$\mu_R = \sigma \sqrt{\frac{\pi}{2}} \quad (7)$$

and

$$\sigma_R = \sigma \sqrt{2 - \frac{\pi}{2}}, \quad (8)$$

where σ is the standard deviation of the underlying Gaussian process. Notice that the Rayleigh distributed signals may be obtained directly from the uniformly distributed signals U by

$$R = \sigma \sqrt{-2 \ln(U)}. \quad (9)$$

The noise samples in the reference cells were generated by Eq. (9) but when it was necessary to add targets to the noise then the noise components were generated by Eq. (5), the target added, and the envelope signal is generated by Eq. (6).

Since the CFAR we are evaluating required the test cell to exceed M out of N thresholds, each threshold based on a subset of reference cells, we found it convenient, especially for determining false alarm levels, to order the thresholds by value (e.g., $Y_1 \leq Y_2 \leq Y_3 \leq \dots \leq Y_N$) and check only the M^{th} value Y_M . The threshold T is derived from

$$T = K \hat{\mu}_R,$$

where K is a constant multiplier of an estimate of the mean of the Rayleigh distribution μ_R . The coefficient, K , is the multiplicative factor of Eqs. (2) and (4). The estimate $\hat{\mu}_R$ is made for each of the N subsets of reference cells. Conventionally, the multiplier K_1 is defined in terms of the underlying Gaussian process, that is,

$$T = K_1 \sigma.$$

Using Eq. (7) and the fact that $\hat{\mu}_R$ is an unbiased estimate of μ_R , we can define

$$T = K_1 \sqrt{\frac{2}{\pi}} \hat{\mu}_R, \quad (10)$$

where the threshold T is now defined in terms of the standard deviation of the underlying Gaussian process. We wish now to determine the value of K_1 which will yield a given probability of false alarm P_{fa} . We can do this by performing a large number of trials in each of which:

- The noise in each reference cell is simulated.
- The reference cells are summed in subsets and the subsets ranked.
- The noise in the test cell is simulated.
- An estimate of K_1 is made based on

$$K_1 = (S N_R) / (Y_M C), \quad (11)$$

where S is the value of the test cell, N_R is the number of reference cells in each subset, Y_M is the sum of the reference cells in the M th subset, and C is the conversion factor $\sqrt{2/\pi}$.

- The value of K_1 is placed in a histogram which is a sample density of the random value of the test cell normalized by an estimate of the standard deviation of the underlying Gaussian process.

In short, for each case, for the given S and Y_M , a false alarm would occur if the threshold were set by the value of K_1 calculated by Eq. (11). This is the relationship between threshold and false alarm for which we are looking. When the trials are completed, the sample distribution is obtained by integrating the sampled density from the tail. The value of the sample distribution at K_1 is the P_{fa} associated with the threshold T , given, by Eq. (10).

Since we are interested in small P_{fa} 's, a large number of trials would be necessary to have any confidence in our estimates of P_{fa} . To increase the number of trials which result in the occurrences we are measuring (in this case false alarms), we use importance sampling [8,9]. The basic principal of importance sampling is: if the event we wish to observe is unlikely with the true density function, obtain the samples from a modified density function with a higher probability of yielding values in the region of interest. Because we obtain more samples, each sample is weighted by

$$W = p(X_i) / P_m(X_i),$$

where $p(\cdot)$ is the true density function, $p_m(\cdot)$ is the modified density function, and X_i is the sample value. Since X_i is sampled from the density $p_m(\cdot)$, a typical value of X_i would yield values of $W \leq 1$. From Eq. (11), we see that we are examining a ratio of a test sample to a set of reference samples. If we select all samples from the modified density function, the output K_1 from Eq. (11) is unchanged and importance sampling will not help. However, following the example in Ref. 9, only the test cell is sampled from the modified density, the reference cells are obtained from the original density. The reasoning behind this is that the variance of the sum of the reference cells is much less than the variance of the test cell, and the behavior of the test cells sampled from the tails of the density have more impact on the outcome than the sum of the reference cells. This means that importance sampling is more effective with many reference cells; with fewer and fewer reference cells, more and more Monte Carlo trials must be run to accurately estimate the thresholds.

The thresholds themselves have no meaning except as intermediate results in generating P_d s for each detector. A set of typical thresholds is shown in Fig. 3. The notation to be used on the rest of the figures of the report follows. The symbols plotted are values obtained from simulation, the lines drawn are spline smoothed to the values. In the legend, the first label refers to the case where the thresholds are known. The second label refers to a conventional CFAR in which a moving window integrator (MWI) sums up 12 range cells from each side of the target and uses them to estimate one threshold. The third label refers to exceeding 4 thresholds of 6, each threshold estimated by 4 samples, to declare a detection, in this case, a false alarm.

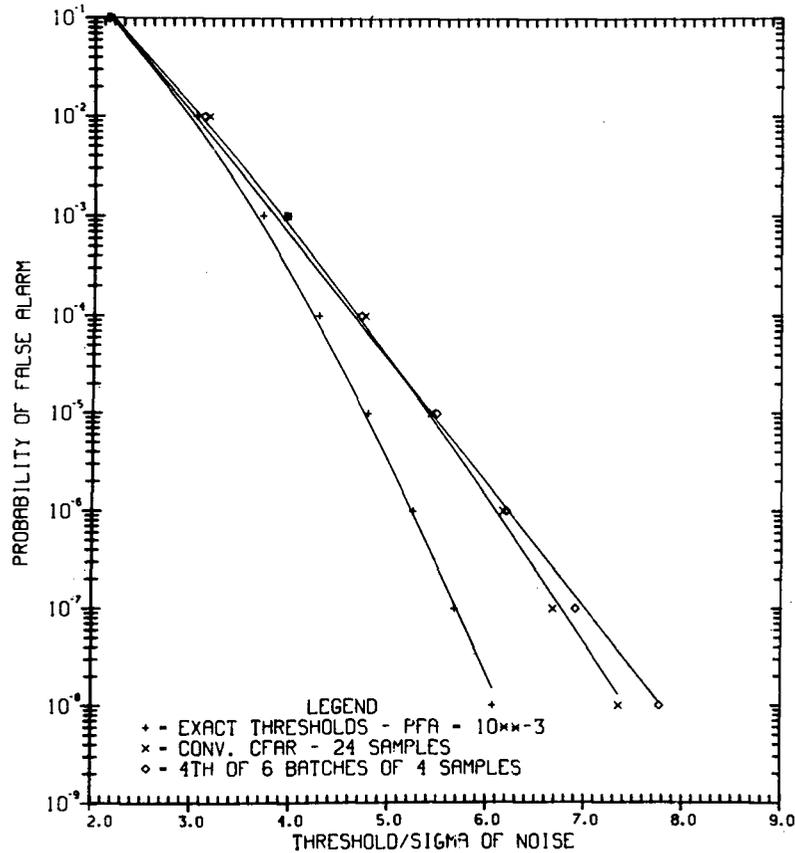


Fig. 3 — Typical examples of threshold values for various types of detectors

Once we have the thresholds for each detector, we may run a Monte Carlo to obtain the P_d . Since we are only interested in values of P_d between 0.1 and 0.99, a straightforward sampling scheme is sufficient. Importance sampling is not necessary. To estimate P_d , Eq. (11) is rewritten as

$$S \geq T = K_1 \times CY_M/N_R,$$

and declares a detection, where, as before, S is the test cell, N_R is the number of reference cells in each batch; Y_M is the sum of the reference cells in the M th batch; C is the conversion factor $\sqrt{2/\pi}$, and K_1 is the multiplier corresponding to a desired P_{fa} .

There are several parameters to be selected; the number of batches, the number of samples in each batch, and how many batches a test cell should exceed to be declared a detection. Each of these modifications defines a new detector requiring calculation of the threshold for the desired P_{fa} and P_d for various signal-to-noise ratios (S/N). Most of the P_d curves are for a P_{fa} of 10^{-3} because this threshold is much simpler to calculate than for the more common P_{fa} of 10^{-6} . Because we are making a

detection decision on each pulse, any further processing would lower the false-alarm rate considerably. For example, if there were two pulses on target and we require coincidence of detection on both pulses to declare a detection, the joint P_{fa} drops to 10^{-6} .

DETECTION PERFORMANCE IN THE ABSENCE OF INTERFERENCE

In determining appropriate parameters for the nonsuppressing CFAR, we rely heavily on the probability of detection curves. We include on each figure in this report which displays P_d curves, a reference curve of the P_d given the exact value of the noise level. This curve was obtained from the simulation by generating the test cell and comparing it to the known fixed threshold for Rayleigh noise.

There are three principal parameters to be determined

- the total number of reference samples
- the number of batches to divide the reference cells into
- the number of batches we must exceed to declare a detection.

The first parameter to examine is the total number of samples. Figure 4 is P_d for a conventional CFAR (one batch of reference cells), and Fig. 5 is the P_d for a nonsuppressing CFAR based on exceeding four of six thresholds. Note a gradual degradation of P_d as the number of reference cells decreases. The decline is fairly uniform until we get to the step between 24 samples and 12 samples, where the decline is much more noticeable. This strongly suggested 24 as the minimum number of reference samples. In the preceding, a tacit assumption was made that there would be six batches. This was investigated further in Fig. 6 where various combinations of batches and samples per batch (total ≈ 24) were evaluated. The only curves which can be identified easily on Fig. 6 are the P_d given exact thresholds and the P_d with conventional CFAR. The other P_d curves are essentially the same, given the resolution of the simulation.

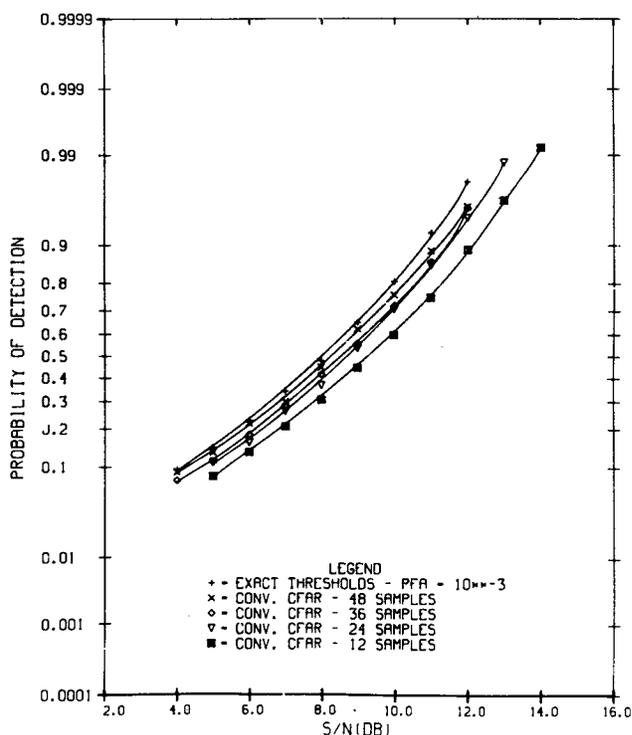


Fig. 4 — Detection performance of a conventional CFAR with various numbers of samples used to estimate the environment

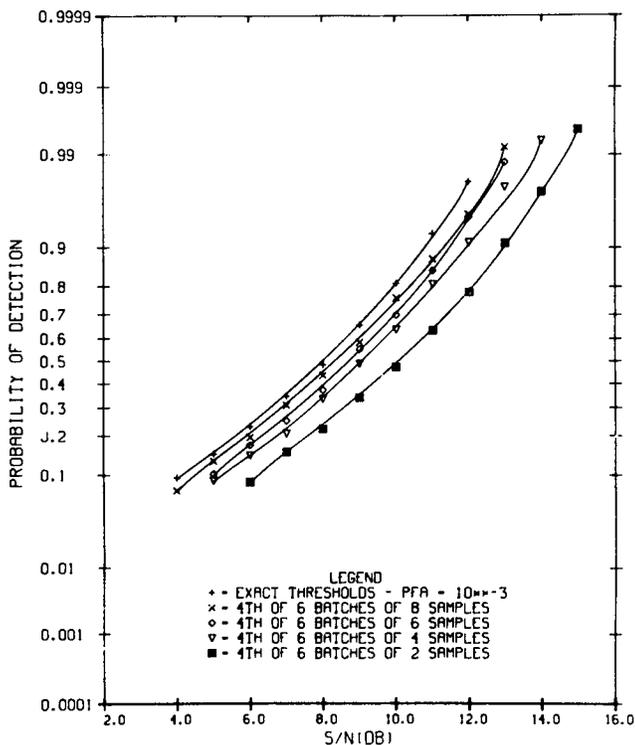


Fig. 5 — Detection performance of a nonsuppressing CFAR with various numbers of samples used to estimate the environment

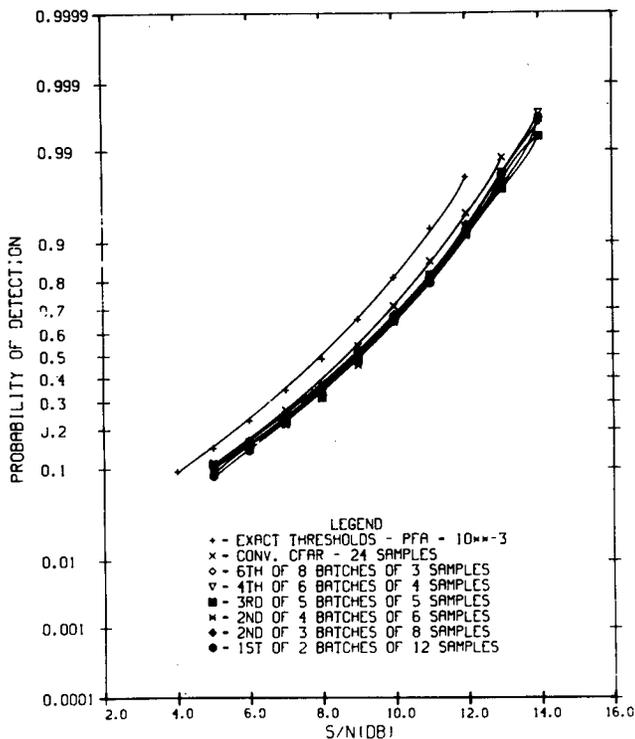


Fig. 6 — Detection performance of nonsuppressing CFARs with various values of M and N with approximately 24 samples used to estimate the environment

Tentatively setting the total number of thresholds at six, the effect of choosing different numbers of thresholds which the test cell must exceed to be called a detection is shown in Fig. 7. While there is a mild preference for using 4 out of 6 thresholds, the worse case of using 2 out of 6 thresholds only exacts a penalty of possibly 1 dB in additional S/N required for detection. A more serious objection to using only a few thresholds out of many to determine a detection is that the statistics of the environment are being estimated by nonrepresentative samples. The detection decision is based on samples from only a small portion of the range of possible functional values. Two random values could have very similar density functions over the small valued samples, but quite different probabilities for the extreme values. This could result in wrong P_{fa} if the wrong set of statistics is assumed. The closer the nonsuppressing CFAR resembles the conventional CFAR, the less chance of error.

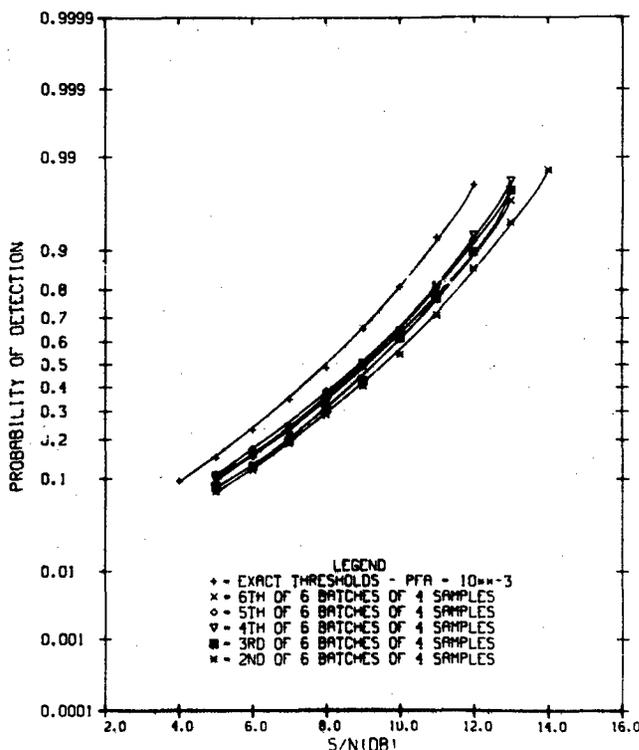


Fig. 7 — Detection performance of nonsuppressing CFARS with various values of M with $N = 6$ and 24 samples used to estimate the environment

We can now make some general observations. We would prefer the number of thresholds required for detection to be at least two less than the total number of thresholds. This allows at least one interfering target at the worst possible location, and two interfering targets if they are fortunately placed. We also wish to base the detection decision on as many samples as possible. With these remarks in mind, we selected three of the detection curves from Fig. 6 to display on Fig. 8. The performance of the 4 thresholds out of 6 detectors is marginally better than the 3 thresholds out of 5 detectors and uses fewer adders than the 6 thresholds out of 8 detectors; consequently, we have selected this detector for further evaluation.

DETECTION PERFORMANCE IN THE PRESENCE OF INTERFERENCE

To test for the effect of interference in the reference cells, we ran the P_d simulation with a fixed S/N of target and varied the size and number of interfering targets in the reference cells. We choose a S/N which yielded a P_d of 90% with a conventional CFAR. The candidate nonsuppressing CFAR

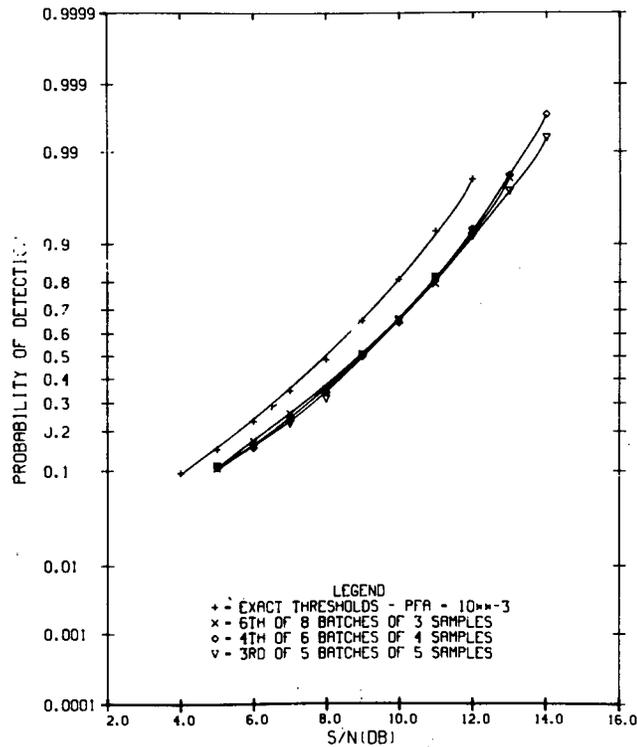


Fig. 8 — Detection performance of selected nonsuppressing CFARs

detector chosen to evaluate is one which calculates 6 thresholds and declares a detection when the test cell exceeds 4 of the thresholds. In Fig. 9, we see that with a zero interference level, the nonsuppressing CFAR detector performs slightly poorer than the conventional CFAR detector (note that the scale on the abscissa is not a dB scale). There are 4 cases considered for which the P_d is plotted against total interference. The first case is the conventional CFAR which is reduced to a 0.1 P_d by the time the interference is 3 times the size of the target. The second curve shows the effect of an interfering target in one batch of reference cells. The P_d decreases from 0.85 to a constant 0.8. This can be explained by the fact that as the interference gets larger, the threshold which is set by the interference becomes so large that the test cell can never exceed it, and the detector in the limit becomes a 4 out of 5 detectors. The same reasoning explains curve 3 with interference contaminating 2 thresholds, and in curve 4, with interference affecting 3 thresholds; we are trying to exceed 4 thresholds when it is only possible to exceed 3 when the interference is large.

Thus, we see that with only 1 or 2 thresholds contaminated by interfering targets, detection performance is still acceptable no matter what level of interference.

SUMMARY

We have seen that it is possible to implement a modification of a conventional CFAR with a minimum of additional hardware. It is relatively insensitive to interfering signals in the reference cells as long as the interference is limited in extent, such as another target. In fact, when declaring a detection on exceeding M thresholds out of N , the interference has little effect on detection as long as the interference only affects $N - M$ of the thresholds. The penalty we pay for using such a detector is typically less than 1 dB additional S/N required for detection at the 0.9 level. This performance is against Rayleigh noise. An important question left unanswered is the extent of robustness of this type of detector. How does it operate when the environment obeys statistical laws other than those assumed?

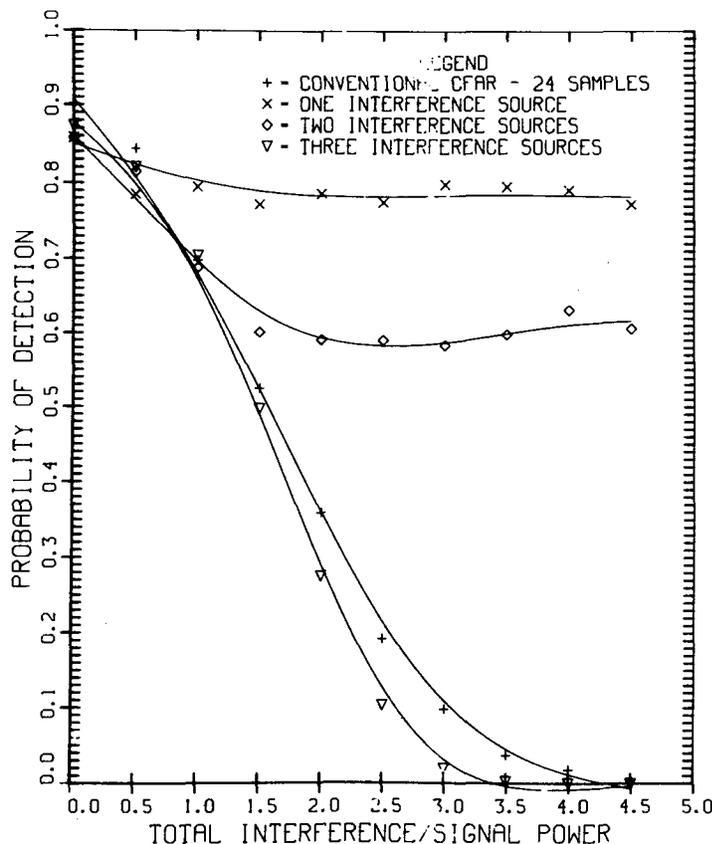


Fig. 9 — Effect on detection performance of a conventional CFAR and nonsuppressing CFAR of various numbers of interference sources with a signal level which yields 0.9 probability of detection for a conventional CFAR with no interference

REFERENCES

1. M.I. Skolnik, ed., *Radar Handbook*, McGraw-Hill Book Co., 1970.
2. D.K. Barton, ed., *International Cumulative Index on Radar Systems*, Institute of Electrical and Electronic Engineers, Inc. IEEE, New York, 1978.
3. H.M. Finn and R.S. Johnson, *Adaptive Detection Mode with Threshold Control as a Function of Spatially Sampled Clutter-Level Estimates*, RCA Rev. **29**, Sept. 1968, pp. 414-464.
4. P.G. Casner, Jr. and R.J. Prengaman, *Integration and Automation of Multiple Colocated Radars*, EASON-77 Record, IEEE Electronics and Aerospace Systems Convention, 1977.
5. G.V. Trunk, B.H. Cantrell, and F.D. Queen, *Modified Generalized Sign Test Processor for 2-D Radar*, IEEE Trans. Aerosp. Electron. Syst. **AES-10**, No. 5, Sept. 1974, pp. 574-582.
6. J.T. Rickard and G.M. Dillard, *Adaptive Detection Algorithms for Multiple-Target Situations*, IEEE Trans. Aerosp. Electron. Syst. **AES-13**, No. 4, July 1977, pp. 338-343.
7. G.V. Trunk, *Range Resolution of Targets Using Automatic Detectors*, IEEE Trans. Aerosp. Electron. Syst. **AES-14**, No. 5, Sept. 1978, pp. 750-755.

8. V.G. Hansen, *Importance Sampling in Computer Simulation of Signal Processors*, Computers and Electrical Eng., **1**, 1974, pp. 545-550.
9. R.L. Mitchell, *Importance Sampling Applied to Simulation of False Alarm Statistics*, IEEE Trans. Aerosp. Electron. Syst. **AES-17**, No. 1, Jan. 1981, pp. 15-24.