Session 21

21.1

Constant False Alarm Rate (CFAR) Detection

21.2

- When we considered the problem of optimal signal detection, we assumed we knew the distribution of the noise under the null hypothesis.

- We not only assumed we knew the form of the distribution (e.g., Gaussian or Rayleigh), but we assumed we had perfect knowledge of the parameters of the distribution.

    - What if we are wrong about the parameters (or the distribution for that matter?)

    - Incorrect knowledge of the noise distribution results in errors in the likelihood ratio threshold.

    - This can *drastically* effect actual detection.

21.3

Assume that we must decide between two simple hypotheses

$$H_0 : \quad X \sim \exp(\mu_0),$$

$$H_1 : \quad X \sim \exp(\mu_1).$$

(assume $\mu_1 > \mu_0$.)

The resulting test will be a threshold test of the form

$$\phi(X) = \begin{cases} 1, & \text{for } X \geq x_0; \\ 0, & \text{for } X < x_0. \end{cases}$$

The threshold x_0 that yields a probability of false alarm α is

$$x_0 = -\mu_0 \ln \alpha.$$

If we have an error in μ_0, we will have a significantly different false alarm probability:

$$\alpha = (\alpha_{\text{nom}})^{\mu_{\text{nom}}/\mu_0}.$$

21.4

The effects of inaccurate noise estimates

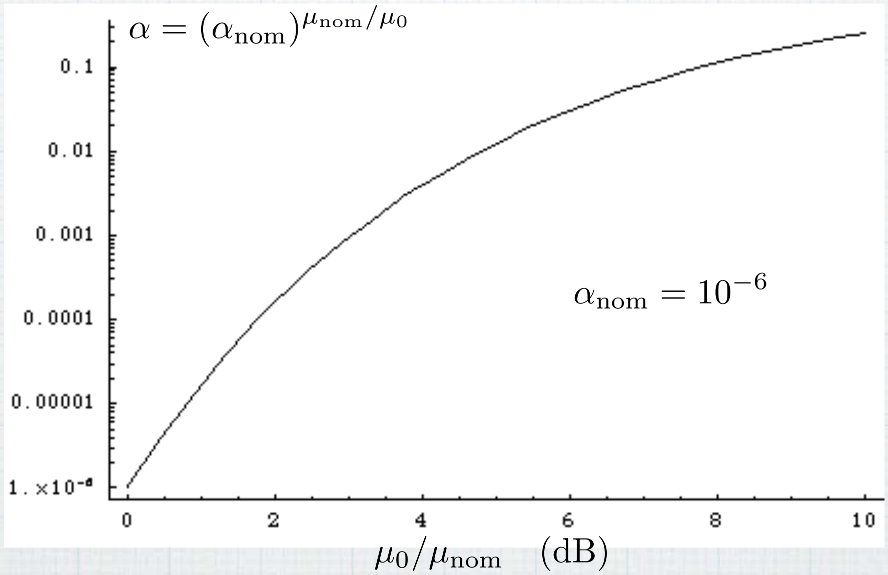

21.5

The exponential Detection Problem Revisited

Assume that we must decide between two simple hypotheses

$$H_0 : \quad Y \sim \exp(\mu_0),$$

$$H_1 : \quad Y \sim \exp(\mu_1).$$

Now if we think of

$$\mu_1 = \mu_0 + \mu_s,$$

where

$$\mu_s = \text{signal component of } \mu_1,$$

then if we define the *signal-to-noise ratio* as

$$S = \mu_s/\mu_0,$$

we can rewrite μ_1 as

$$\mu_1 = \mu_0(1+S),$$

and our simple hypotheses can be rewritten as

$$H_0 : \quad Y \sim \exp(\mu_0) \qquad \text{versus} \qquad H_1 : \quad Y \sim \exp(\mu_0(1+S)).$$

The most powerful test of size α is given by

$$\phi(Y) = \begin{cases} 1, & \text{for } Y > Y_0, \\ 0, & \text{for } Y \leq Y_0, \end{cases}$$

where the threshold Y_0 is given by

$$Y_0 = -\mu_0 \ln \alpha.$$

The power of the test is given by

$$\beta = P(Y > Y_0 | H_1) = \cdots = \alpha^{1/(1+S)}.$$

n.b. The threshold Y_0 is a function of μ_0 and the probability of false alarm α.

If we don't know the value of μ_0, we cannot set the threshold Y_0 that will yield our size α test. How should we proceed?

In principle, μ_0 could take on a broad range of positive values.

We could view H_0 as a the composite hypothesis that $\mu_0 \in (0, \infty)$.

We could then use a generalized likelihood ratio test to solve the problem.

Under hypothesis H_0, this would correspond to finding the maximum likelihood estimate $\hat{\mu}_0$ and using it in place of μ_0. But for one sample measurement, this does not yield a good estimate.

However, if we had N i.i.d. measurements $X_1, \ldots, X_N$ of the noise, we could use the maximum likelihood (and minimum variance unbiased) estimate

$$\hat{\mu}_0 = \frac{1}{N} \sum_{i=1}^{N} X_i$$

in place of μ_0.

- In a “typical” radar scenario, targets are sparsely located against a background of noise and clutter.
- There tends to be regions of local statistical homogeneity in this noise/clutter background because the physical environment giving rise to it often has homogeneous statistics.
- However, there can be significant changes in the local scattering characteristics as you move through the scattering environment.
- There can be sharp boundaries between scattering regions.

- This suggests that one approach to estimating the background noise power for target detection in a particular resolution cell is to average the measured noise power in surrounding resolution cells.
- This is an example of a class of detection techniques called *Constant False Alarm Rate* (CFAR) techniques.
- We will see where the term *Constant False Alarm Rate* comes from, but more important than the constant false alarm rate is a robustness to changes in the average noise power.

A Generic CFAR Processor

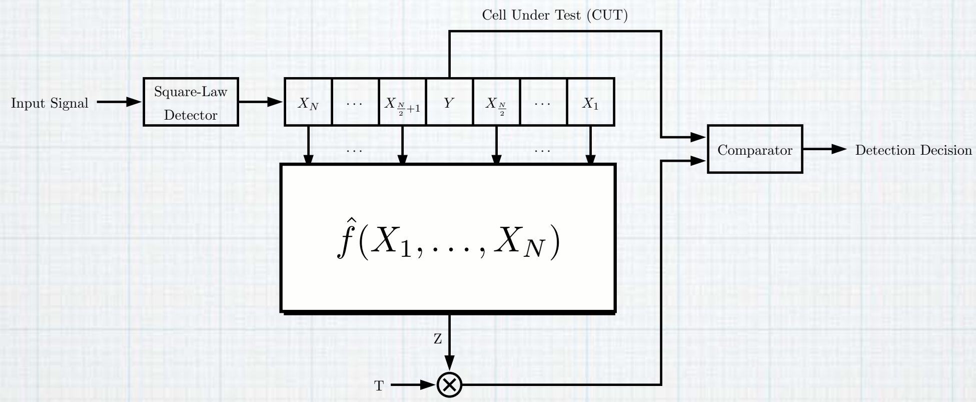

The resolution *cell under test* (CUT) with measurement Y is tested for the presence of the target using a threshold computed using neighboring resolution cell measurements $X_1, \ldots, X_N$.

The statistic $Z = \hat{f}(X_1, \ldots, X_N)$ is an estimate of the noise power.

The threshold scaling factor T sets the threshold level by scaling the statistic the statistic Z. This works because the threshold is the product of a constant and the average noise power.

A More Specific Class of CFAR Processors

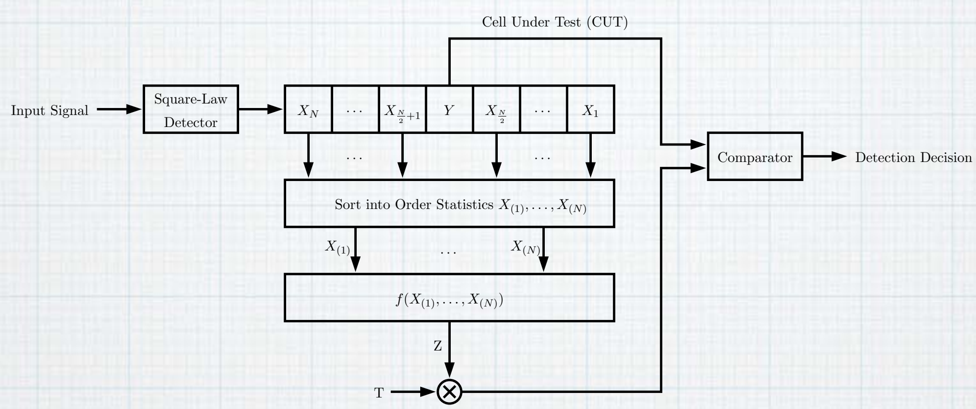

This processor can compute

- Mean
- Median
- Arbitrary Order Statistics
- Linear Combination of Order Statistics.

Cell-Averaging CFAR (CA-CFAR)

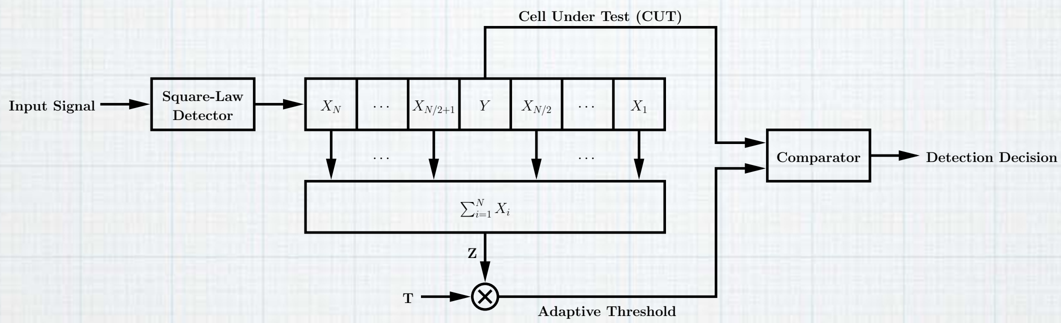

In CA-CFAR, we have that the statistic Z/N is just the sample mean.

It can be shown that Z/N is the maximum-likelihood estimate of μ_0. (It is also the *minimum variance unbiased estimate* (MVUE) of μ_0 and an *efficient estimate*—satisfying the Cramer-Rao lower bound.)

If we assume that $X_1, \ldots, X_N$ are i.i.d exponential with mean μ (drop subscript for simplicity) we have

$$f_{X_i}(x) = \frac{1}{\mu} e^{-x/\mu} \cdot 1_{[0,\infty)}(x).$$

The moment generating function of each X_i is

$$\Phi_{X_i}(s) = \left(\frac{1}{1 - \mu s} \right).$$

The moment generating function of Z is

$$\Phi_Z(s) = \left(\frac{1}{1 - \mu s} \right)^N.$$

Because the test is a threshold test comparing the CUT Y to the threshold TZ, the probability of false alarm is

$$\begin{aligned}\alpha &= \mathrm{E}_Z\left[P[Y > TZ|H_0]\right] \\ &= \mathrm{E}_Z\left[\int_{TZ}^{\infty} \frac{1}{\mu} e^{-y/\mu} dy\right] \\ &= \mathrm{E}_Z\left[\exp(-TZ/\mu)\right] \\ &= \int_{-infty}^{\infty} e^{-\frac{TZ}{\mu}} f_Z(z) dz \\ &= \Phi_Z\left(-\frac{T}{\mu}\right),\end{aligned}$$

where $\mathrm{E}_Z[\cdot]$ denotes expectation w.r.t. Z.

Substituting this into the expression for $\Phi_Z(s)$, the false alarm probability is

$$\alpha = (1+T)^{-N}.$$

Note that the false alarm rate is not a function of the mean noise power μ. Hence the term *constant false-alarm rate*.

The threshold scaling factor yielding a size α test is

$$T = (\alpha)^{-1/N} - 1.$$

Similarly, the detection probability can be calculated under the alternative hypothesis H_1 and given by

$$\begin{aligned}\beta &= \mathrm{E}_Z\left[P[Y > TZ|H_1]\right] \\ &= \mathrm{E}_Z\left[\int_{TZ}^{\infty} \frac{1}{\mu(1+S)} e^{-y/\mu(1+S)} dy\right] \\ &= \mathrm{E}_Z\left[\exp(-TZ/\mu)\right] \\ &= \Phi_Z\left(-\frac{T}{\mu(1+S)}\right) \\ &= \left[1 + \frac{T}{(1+S)}\right]^{-N} \\ &= \left(\frac{1+S}{1+T+S}\right)^{N}.\end{aligned}$$

Combining these results, we find that

$$\beta = \left(\frac{1+S}{\alpha^{-1/N}+S}\right)^N .$$

In the limit, as $N \to \infty$, we have

$$\begin{aligned}
\lim_{N\to\infty} \alpha &= \lim_{N\to\infty} (1+\epsilon/N)^{-N} \\
&= \exp\{-\epsilon\} \\
\lim_{N\to\infty} \beta &= \lim_{N\to\infty} (1+\epsilon/N(1+S))^{-N} \\
&= \exp\{-\epsilon/(1+S)\} \\
\Rightarrow \quad \beta &\to \alpha^{1/(1+S)}, \text{ as } N \to \infty
\end{aligned}$$

CA-CFAR Detection Performance

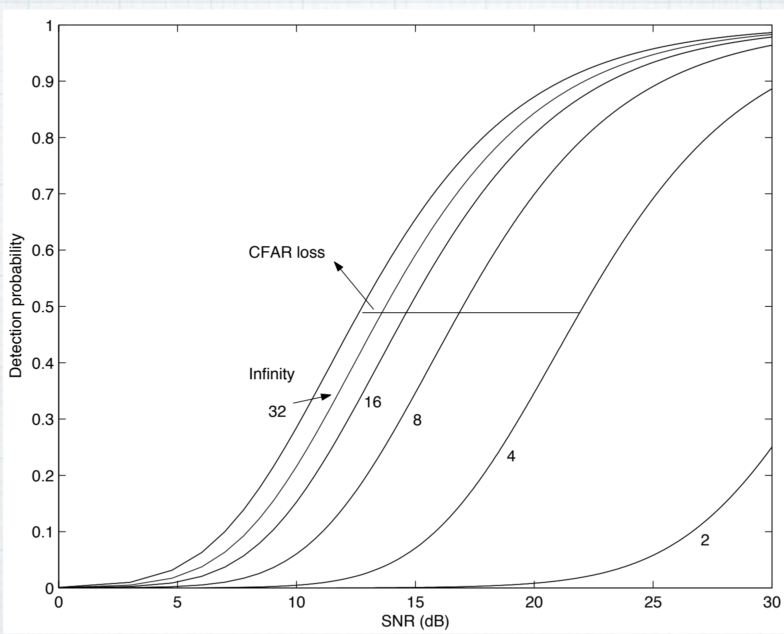

21.18

CA-CFAR P_d versus N and SNR (dB) for a desired $P_{fa} = 1 \times 10^{-6}$

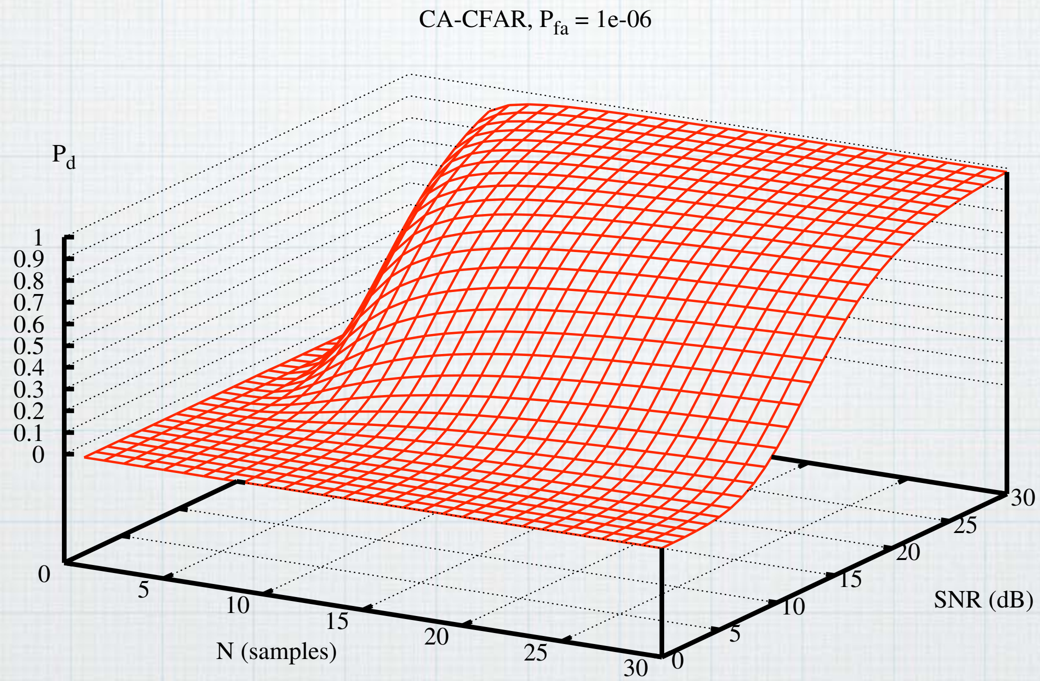

Figure from: Michael F. Rimbert, *Constant False Alarm Rate Detection Techniques Based on Empirical Distribution Function Statistics*, Ph.D Thesis, School of Electrical and Computer Engineering, Purdue University, August 2005.

21.19

CA-CFAR Threshold Map, $N = 8$, $P_{FA} = 0.5$

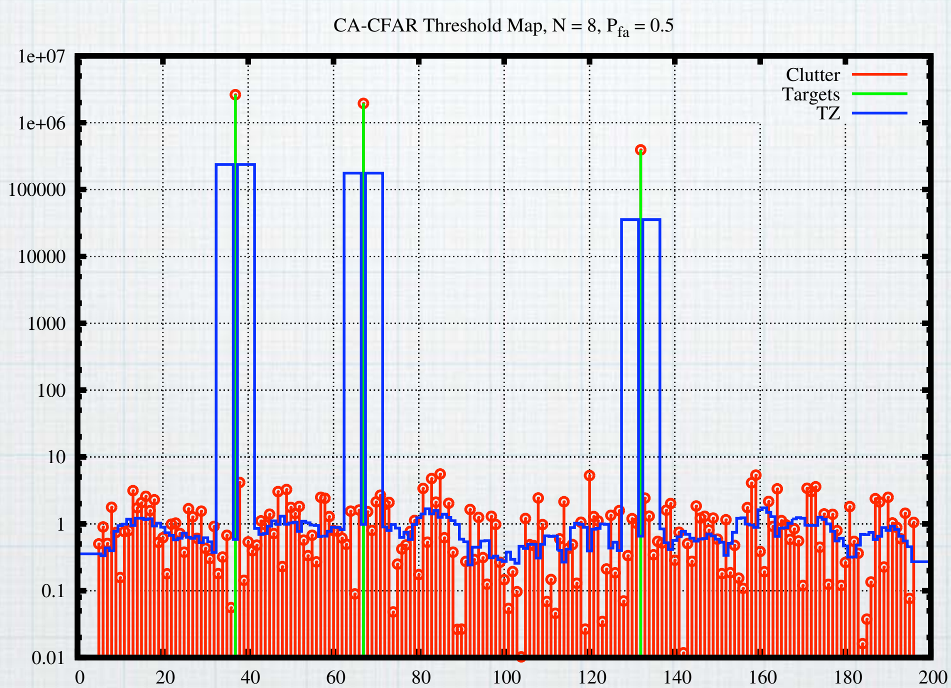

Figure from: Michael F. Rimbert, *Constant False Alarm Rate Detection Techniques Based on Empirical Distribution Function Statistics*, Ph.D Thesis, School of Electrical and Computer Engineering, Purdue University, August 2005.

CA-CFAR Threshold Map, $N = 8$, $P_{FA} = 1 \times 10^{-6}$

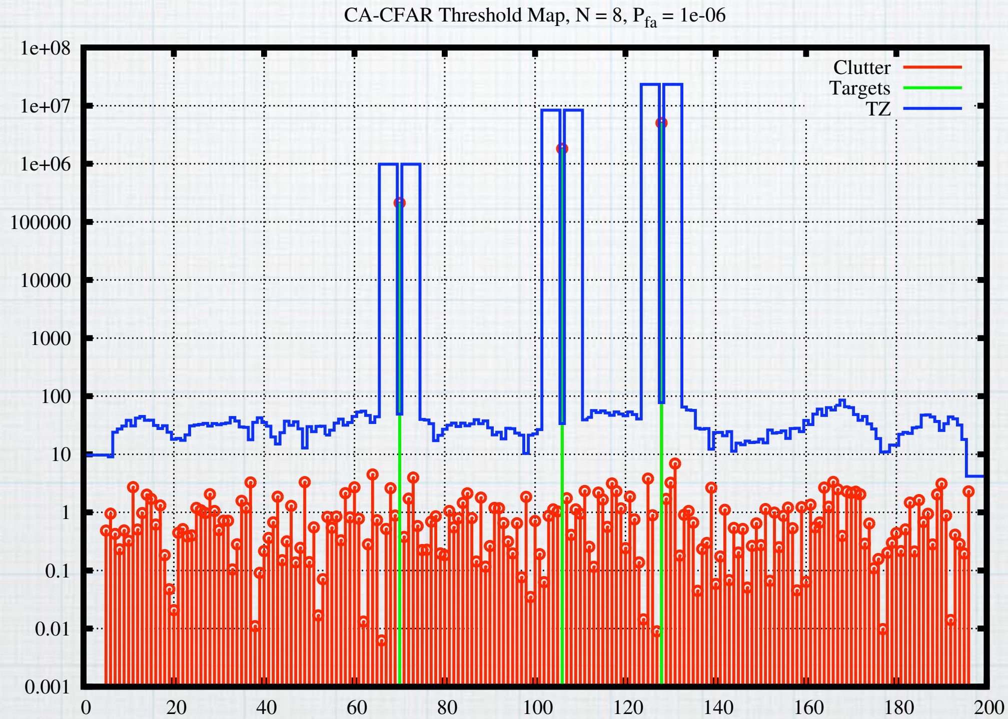

Figure from: Michael F. Rimbert, *Constant False Alarm Rate Detection Techniques Based on Empirical Distribution Function Statistics*, Ph.D Thesis, School of Electrical and Computer Engineering, Purdue University, August 2005.

CA-CFAR Limitations

- The performance of CA-CFAR suffers when statistical homogeneity of the reference window samples is violated. This commonly occurs when:

  1. Reference window contains interfering targets

  2. Reference window contains “clutter edges”—boundaries between regions with differing scattering characteristics.

CA-CFAR – Interfering Targets

21.22

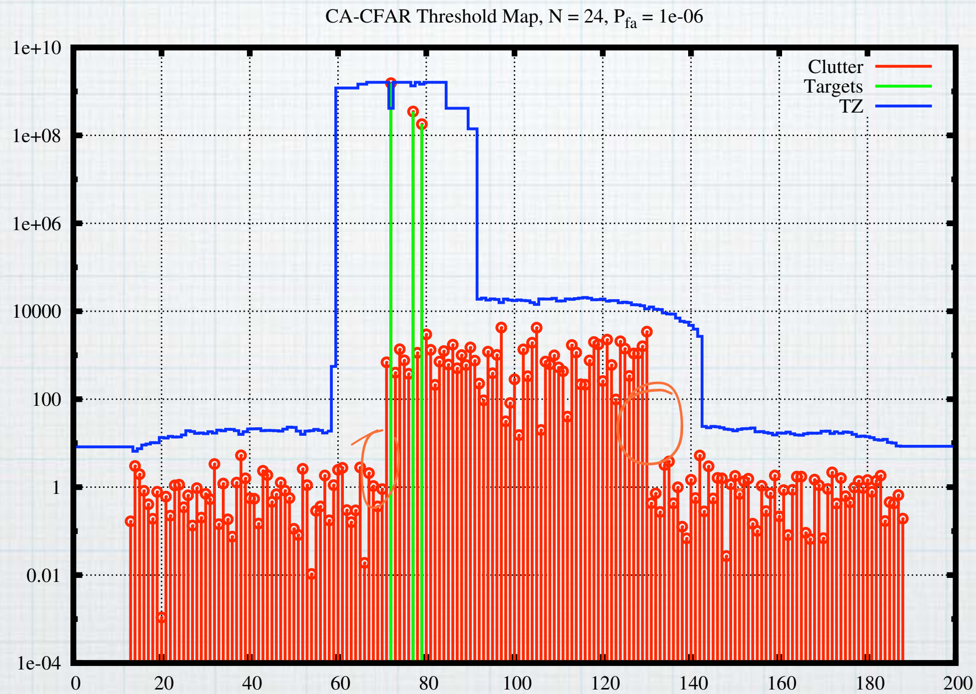

Figure from: Michael F. Rimbert, *Constant False Alarm Rate Detection Techniques Based on Empirical Distribution Function Statistics*, Ph.D Thesis, School of Electrical and Computer Engineering, Purdue University, August 2005.

CA-CFAR – Clutter Edges

21.23

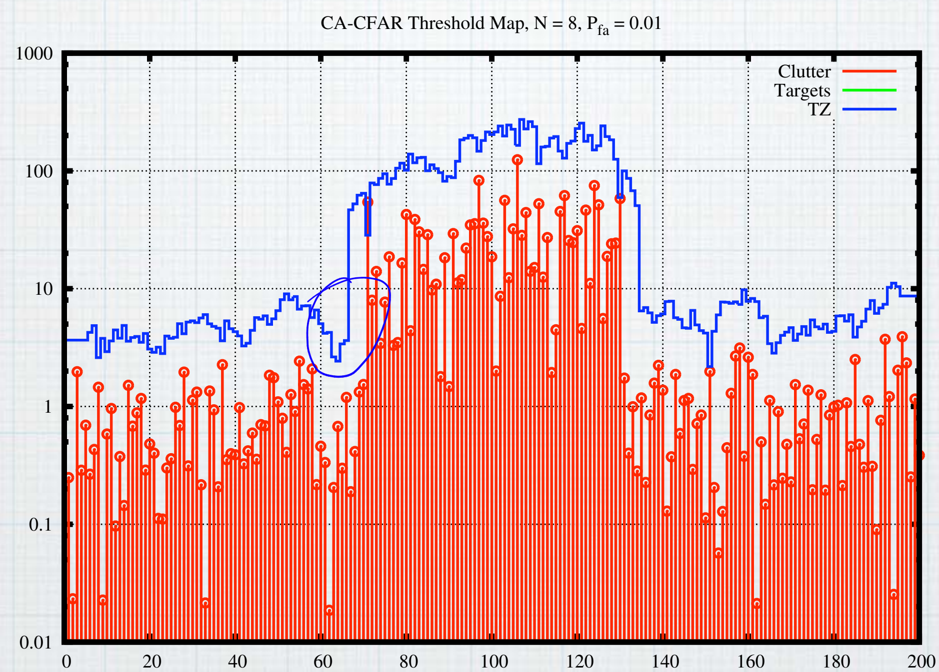

Figure from: Michael F. Rimbert, *Constant False Alarm Rate Detection Techniques Based on Empirical Distribution Function Statistics*, Ph.D Thesis, School of Electrical and Computer Engineering, Purdue University, August 2005.

21.24

GO-CFAR and SO-CFAR Processors

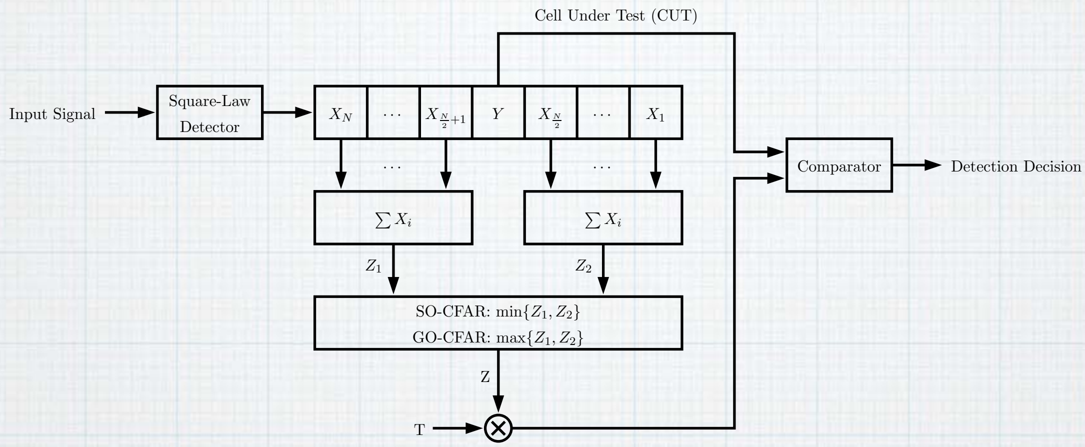

- *Greatest of CFAR* (GO-FAR) reduces excessive false alarms near clutter edges, but poor detection performance in the presence of interfering targets.
- *Smallest of CFAR* (SO-FAR) has improved performance in the presence of interfering targets, but high false alarm rate near clutter edges.

21.25

CFAR Processor Performance Comparison

Problem	**Processor**			
	CA-CFAR	**GO-CFAR**	**SO-CFAR**	**OS-CFAR**
Clutter Edges	Poor	Good	Poor	Good
Interfering Targets	Poor	Poor	Good	Good

Figure from: Michael F. Rimbert, *Constant False Alarm Rate Detection Techniques Based on Empirical Distribution Function Statistics*, Ph.D Thesis, School of Electrical and Computer Engineering, Purdue University, August 2005.

Order-Statistic CFAR (OS-CFAR)

Herman Rohling, "Radar CFAR Thresholding in Clutter and Multiple Target Situations," *IEEE Transactions on Aerospace and Electronic Systems,* vol. 19, pp. 608–621, 1983.

- In OS-CFAR, the average noise power in a region is estimated using an *order statistic*, or ranked sample of the noise power samples in the reference window.
- For example, we might use the *sample median* instead of the *sample mean* to estimate the average noise power.
- While an order statistic estimate is not the maximum likelihood estimate if the samples are independent and statistically homogeneous (i.i.d.), order statistics (e.g., the sample median) are much more robust to deviations from this ideal.