

Evaluation of a Second Order Adaptive CFAR in the Littoral Environment

E. J. Hughes, M. Lewis
Cranfield University at the Defence Academy of the United Kingdom
Shrivenham, Swindon, SN6 8LA

Abstract

This paper examines seven common CFAR techniques and describes a spatially adaptive CFAR that captures the best performance in respect of P_D and P_{fa} in the presence of long tail and target corrupted clutter.

Introduction

Detection of small targets in littoral environments forms a challenge for modern radar systems. Often high target densities are present, and the scene is complex with the local clutter statistics varying both spatially and temporally.

After describing the problems encountered when attempting to detect targets in littoral waters using common CFAR (Constant False Alarm Rate) adaptive detection threshold techniques a deeper examination of Detection and CFAR techniques follows together with a discussion of the assumptions on which the design of CFAR systems is based and how these are flawed and a spatially adaptive CFAR window is suggested.

An assessment of the performance of seven CFARs is described with the aim of answering the question 'How do the different forms of CFAR cope with either inappropriate threshold multipliers or the typical compound distributions that are observed from non-homogenous clutter regions?'

The conclusion is that the ideal CFAR for use with large window sizes is either a 2nd order ordered statistic method, for when more target corruption is expected than true long-tailed clutter; or a Fisher-Tippett 2nd order CFAR for when long-tailed statistics may be present but with mild target corruption. The Fisher-Tippett CFAR is recommended on the

basis of simple processing as the best compromise solution.

Having established the form of CFAR to be employed and the number of cells required in the window a self-organising software intelligent agent based spatially adaptive CFAR method is described. The CFAR has a lowered threshold in order to detect smaller targets allowing the probability of detection to be improved, whilst providing effective control of the false alarm rate.

The Problem of Detecting Small Targets in Littoral Waters

In commonly used methods of target detection, target returns that cross a detection threshold derived from the clutter level in range cells adjacent to the cell under test are taken as 'potential targets'. The decision mechanism directly affects the probability of target detection and the probability of a false alarm.

Discrimination against false alarms is ultimately performed in the tracking system, and the capabilities of the tracker will determine the maximum false alarm rate that can be tolerated, and therefore the minimum value for the decision threshold. In littoral environments the target density can be very high and the environment is unknown and temporally unstable. In order to calculate a threshold location, the underlying clutter probability density function shape must either be known, or it must be possible to estimate the shape of the distribution from the radar

alarms being generated. The result is a slight reduction in the probability of detection equivalent to a small loss of signal to noise ratio.

Using a large number of samples reduces the error in the measured parameter but increases the risk that some of the farther samples are not representative of the clutter in the region of the test cell. It also increases the chance that the cells may contain targets, which again are not observations. To detect very small targets, the threshold must be lowered resulting in an increased false alarm rate. If the threshold is reduced too far, the number of false alarms increases dramatically. It is therefore not a trivial matter to depress the detection threshold and simultaneously control the false alarm rate in a difficult littoral environment.

Detection and CFAR

The detection of targets relies on being able to separate target returns from background signal. In practice, a cell will contain either a background signal (receiver noise + clutter echo), or target plus background. The premise underlying classical CFAR processing is that if the statistics of the noise/clutter are known and a good estimate of the low-order moments (or central moments) is generated from the measured data (typically 30 samples) a threshold level can be calculated that will achieve the maximum probability of detection together with the maximum tolerable probability of false alarm.

In reality, both the statistical distribution of the target and the statistical distribution of the background are unknown, and in a non-stationary littoral environment, it may also not be possible to determine the instantaneous probability distributions of either.

All this requires that the threshold be adjusted locally to maintain a maximum probability of detection, whilst not exceeding the maximum tolerable probability of false alarm. The best that can be achieved in reality is to set a threshold level that provides the highest number of total returns that the tracker can

process and separate into target tracks and clutter. Due to the complexity and non-stationary statistics of the environment, the probability of detection that results will be the best that is achievable for the target at the instant of observation.

It has been found that in order to gather sufficient samples to obtain a reasonable estimate of the mean and standard deviation of the clutter statistics, the samples must be drawn from a spatio-temporal region. In order to make the samples as consistent as possible, the topology of each local region must be optimised to the current environment and since this is unknown and dynamic, the region must be adaptive. As the statistics are non-stationary, only a limited time history may be used. Although sources of thermal noise are likely to be independent, clutter samples tend to be highly correlated. Thus the number of truly independent samples is reduced, again leading to poor estimates of the statistics.

Constant False Alarm Rate (CFAR) Techniques

All CFARs maintain a constant false alarm rate at the expense of detection probability and so introduce a processing loss. This loss depends upon the particular implementation and is often in the region of 1 - 2 dB and is related to the number of samples that are gathered in order to determine the statistical properties of the local background.

The optimal CFAR process for thermal noise is the Cell Averaging CFAR, in particular such techniques perform optimally in uniform Rayleigh clutter [1]. Although the detection threshold is adaptive there are a number of disadvantages to this scheme. The statistical parameter is estimated from a relatively small number of samples and so is likely to differ from the true population value. The uncertainty in the threshold level results in it being set high side to prevent an excessive number of false representative of the clutter.

On the other hand a small number of cells will reduce the chance of non-representative values being present but gives a poorer

measure of the parameter and a greater CFAR loss.

Very many variants of the cell averaging CFAR have been developed with one of the most successful in practice being the Order Statistic CFAR (OSCFAR). Here a non-parametric ranking process is used to determine the threshold level. The OSCFAR performs robustly in clutter regions, but the threshold is often set pessimistically in regions of pure noise. The OSCFAR is related strongly to the process of median filtering.

Anastassopoulos and Lampropoulos [1] discuss a number of other CFAR schemes that attempt to overcome the problems of clutter edges, multiple targets and non-Rayleigh clutter. Although a solution is presented the authors admit that the CFAR proposed requires an execution time that is nearly a factor of three greater than an order statistic CFAR of the same length.

Second-Order CFAR

If higher order statistics can be estimated accurately, any variations in both the spread and the shape of the clutter distribution could be accounted for and be used to set a threshold dependent not only on the mean of the clutter level but also the standard deviation. The disadvantage would appear to be the need for an even larger number of samples than is required for simple CA-CFAR.

In order to answer the question ‘*how many samples are needed in practice?*’ seven first and 2nd order CFAR systems were considered. The processes are summarised in equations 2 to 7, where T_1 is the threshold for the first order CFAR, T_2 is the 2nd order threshold level, v_i are the background voltage magnitude samples from the local region and k_1 and k_2 are the multiplication factors required to give the desired false alarm probability for the 1st and 2nd order CFARs respectively.

- Linear CA: $v > k_1 \times \mathbf{mean}(\mathbf{v})$

In CA-CFAR the threshold is set at the mean of the levels either side of the cell being examined multiplied by a factor.

$$T_1 = k_1 \frac{1}{N} \sum_{i=1}^N v_i \quad (1)$$

- 2nd Order Linear CA:

$$v > k_2 \times \mathbf{std}(\mathbf{v}) + \mathbf{mean}(\mathbf{v})$$

The simplest 2nd order CFAR is an extension of linear cell-averaging [2]. The concept is instead of setting the threshold as a factor times the mean background voltage, the threshold is set to the mean background voltage plus a factor times the standard deviation.

$$T_2 = \frac{1}{N} \sum_{i=1}^N v_i + k_2 \sqrt{\frac{1}{N-1} \sum_{i=1}^N v_i^2 - \left(\frac{1}{N} \sum_{i=1}^N v_i \right)^2} \quad (2)$$

- OS: $v > k_1 \times \mathbf{median}(\mathbf{v})$

- 2nd OS:

$$v > k_2 \times \mathbf{range}(25-75)(\mathbf{v}) + \mathbf{median}(\mathbf{v})$$

Order Statistic methods are described by equations 3 and 4 where V is the vector of all local target voltage samples, P_{25} , P_{50} and P_{75} are the 25th percentile, 50th percentile (median) and 75th percentile respectively.

$$T_1 = k_1 P_{75}(V) \quad (3)$$

$$T_2 = P_{50}(V) + k_2 (P_{75}(V) - P_{25}(V)) \quad (4)$$

- 2nd Order Fisher-Tippet CA:

$$\log(v) > k_2 \times \mathbf{std}(\log(\mathbf{v})) + \mathbf{mean}(\log(\mathbf{v}))$$

Given that sea and land clutter can often have a long-tailed probability distribution, it makes sense to investigate approaches that estimate the shape (i.e. length of tail) directly. If the clutter has a log-normal probability density function, then after passing the received signal through a logarithmic amplifier, the probability density function will be normally distributed, with the mean and standard deviation of the probability density function describing its shape uniquely.

If the clutter follows a Weibull distribution, then after a logarithmic amplifier, the statistics will follow a Fisher-Tippet distribution. Equation 5 describes a 2nd order CFAR process based on processing the logarithm of voltage.

$$T_2 = \frac{1}{N} \sum_{i=1}^N \log(v_i) + k_2 \sqrt{\frac{1}{N-1} \sum_{i=1}^N \log(v_i)^2 - \left(\frac{1}{N} \sum_{i=1}^N \log(v_i) \right)^2} \quad (5)$$

In order to compare the performance of a range of CFAR systems, square-law cell averaging, and root-law cell averaging have also been considered and are described by equations 6 and 7.

- Square CA: $v^2 > k_1 \times \mathbf{mean}(v^2)$

$$T_1 = k_1 \frac{1}{N} \sum_{i=1}^N v_i^2 \quad (6)$$

- Root CA: $\sqrt{v} > k_1 \times \mathbf{mean}(\sqrt{v})$

$$T_1 = k_1 \frac{1}{N} \sum_{i=1}^N \sqrt{v_i} \quad (7)$$

Comparative CFAR Performance Assessment

A simple experiment has been devised to establish the relative performance of the different CFAR methods as the number of available samples (window size) is varied under *realistic* operating conditions.

The hypothesis is that as the window size increases, the aggregate performance of all of the CFAR methods will saturate, with no one method being superior. It is however anticipated that the false alarm performance of some methods may be more attractive than others.

Unknown clutter experiment

The purpose was to investigate how the different forms of CFAR cope with either inappropriate threshold multipliers or the typical compound distributions that are observed from non-homogenous clutter regions. As the exact structure of the compound probability density functions are

not known, the number of assumptions used in the analysis has been minimised by using a Monte-Carlo process where few *a-priori* assumptions are needed.

The experiments were conducted based on an unknown clutter distribution. Rayleigh noise was generated and a multiplier, k , was calculated to provide a given $P_{f.a.}$. K -distributed samples were drawn and a threshold calculated based on the CFAR method being evaluated and the previously calculated multiplier k .

A K -distributed clutter sample and a ‘target plus K -distributed clutter sample’ were tested against the threshold, allowing actual P_D and P_{fa} to be established (from 10,000 trials).

The experiments were repeated for a range of window sizes

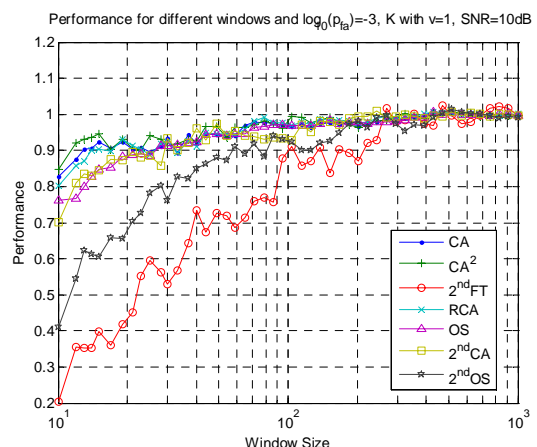


Figure 1. Plot of relative CFAR performance with respect to window size, for the case of long-tailed K-distributed clutter.

Figure 1 shows that as the window size increases the performance of the various CFARs becomes comparable thus supporting the hypothesis. The conclusion from Figure 1 is that for an ~300 sample window the performance of all the CFARs under consideration would be similar.

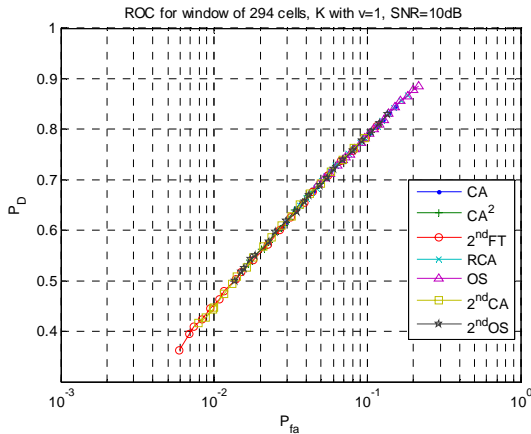


Figure 2. Receiver Operating Curve for 294 cell window and long-tailed clutter

Figure 2 shows that the Receiver Operating Curve for a 294 point window indicates that all of the 7 methods have comparable performance, as they all lie on the same ROC trajectory. It is clear however, that the methods do not behave in the same manner as the cell averaging approaches appear to be producing large numbers of false alarms due to a lower than expected threshold (and therefore a corresponding increase in the probability of detection).

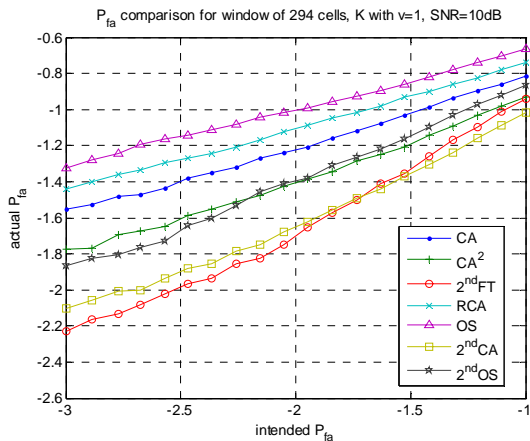


Figure 3 Plot of actual false alarm versus the intended false alarm probability for long-tailed K-distributed clutter.

This behaviour is confirmed in Figure 3 where the actual P_{fa} is plotted with respect to the intended P_{fa} . The figure shows that all of the CFAR methods produce more false alarm returns than intended, however some are much more extreme than others. The 2nd order CFAR methods outperform the 1st order methods, with the 1st order methods all consistently setting the threshold much too

low and therefore generating very large numbers of false-alarms. The experiment demonstrates that in unknown long-tailed clutter, if 2nd order statistics can be estimated, then a more appropriate threshold choice can be made.

It is clear that with large window sizes (>200 cells) the performance difference between the different CFARs is minimal, i.e. when plotted, the CFAR ROC curves all lie on the same ROC trajectory. Thus a method is now required that can gather groups of 100+ homogenous cells from the scene of interest. The second experiment, looking at target corruption of these regions, will determine how accurately the regions need to be defined.

Target corruption experiment

In a real scenario, the CFAR threshold may be set based on an *a-priori* assumption of the noise statistics (e.g. Rayleigh), yet the true clutter may be very long-tailed, or the local cells corrupted by targets or other environmental artefacts. The effect of corruption is that the probability density function describing the background becomes a compound distribution, comprised of clutter-like and target-like samples. If a few target returns are included in the background sample, the tail of the distribution is ‘stretched’, however the distribution is now multi-modal (i.e. multiple peaks of high density), rather than long-tailed.

The experiments were modified to produce a compound background distribution. A percentage of the samples had contributions from a Swerling 2 target added and the threshold location was calculated from the modified sample set, based on the CFAR method being evaluated and the calculated multiplier k .

A Rayleigh distributed ‘noise’ sample and a ‘target plus noise sample’ were tested against the threshold, allowing actual P_D and P_{fa} to be established (from 10,000 trials);

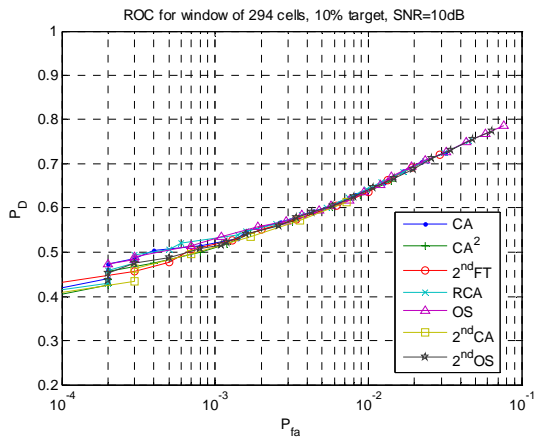


Figure 4 Receiver Operating Curve for 294 cell window and a compound distribution, where 10% of the cells have been corrupted by target-like returns

The ROC curve for a 294 point window shown in Figure 4 indicates that all 7 methods have comparable performance, i.e. no one method outperforms the other methods significantly as all the methods lie on essentially the same ROC curve trajectory.

Again it is clear, that the methods do not behave in the same manner. Figure 5 shows the P_{fa} performance of all 7 methods. All of the methods provide fewer (better) false alarms than the design intention, with an associated reduction in the P_D . Since the false alarms do not explode out of control, unlike in the true long-tailed distribution; although undesired, all of the methods could tolerate some degree of target corruption. Interestingly, there is no clear distinction between the performance of first and second order CFAR methods, rather the order-statistic methods (best being 1st order, followed by the 2nd order) are the best performing, then root and linear cell averaging, and then the Fisher-Tippett CFAR etc. The performance with changing window size is very similar to the long-tailed clutter experiment with 200 to 300 cell window sizes providing good performance for all of the CFAR methods.

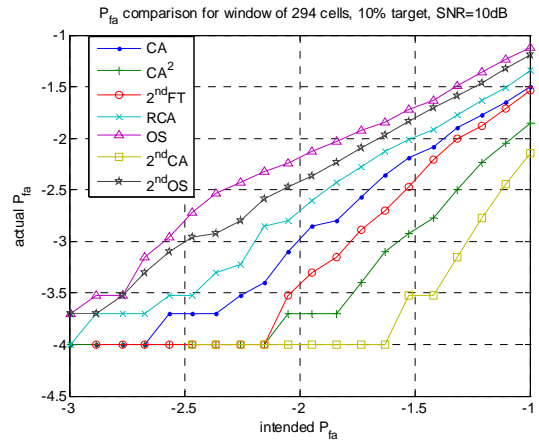


Figure 5 Graph of actual P_{fa} versus intended P_{fa} for a 294 cell window and 10% corruption by target-like returns.

Experiment summary

Method trade-off front based on P_{fa} performance rank, SNR=10db, 294 cell window

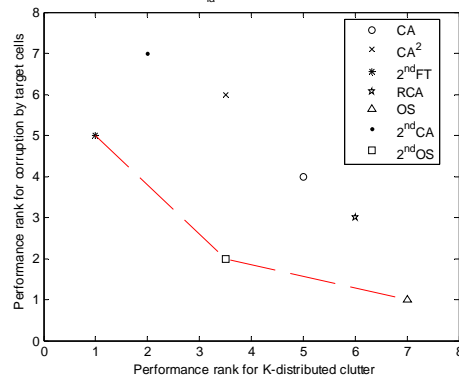


Figure 6. Trade Off Surface for various CFAR Surfaces

A trade-off comparison can be made using the P_{fa} performance for the larger range-azimuth window since the ROC performance can be considered 'equivalent',

The ranks of the different methods for the two different noise distributions have been plotted in Figure 6 (rank=1 is best, 7 is worst).

The trade-off surface reveals:

- for clutter-dominated scenarios (e.g. spatio-temporal CFAR), the 2nd order Fischer-Tippett approach is recommended.
- For noise-alone operation where target corruption may be an issue, 1st order order-statistic CFAR is best
- For mixed operation where both target corruption and long-tailed clutter may be present, 2nd order order-statistic CFAR is the best compromise.

- A 200+ cell window is best when 2nd order CFAR is being applied.

The two experiments have demonstrated 3 key observations in the comparison of 1st and 2nd order CFAR methods.

The first observation is that with window sizes greater than approximately 100 cells, the performance of all of the CFAR methods converge to lie on the same ROC trajectory, indicating that no one method could be considered superior when both P_D and P_{fa} performance are considered simultaneously.

The second observation is that in the presence of long-tailed clutter, with an unknown distribution, the 2nd order CFAR methods provide the most reliable threshold locations (when moderate to large window sizes are considered). Many of the 1st order methods place too low a threshold and the number of false alarms may be much higher than anticipated.

The 3rd observation is that when the gathered cells are partially corrupted by target returns, resulting in a compound distribution, the order statistic methods provide the best performance, however all of the methods set the threshold high, reducing the P_D , but not incurring an increase in the number of false alarms that are observed.

Choice of CFAR

The ideal CFAR for use with large window sizes is therefore either a 2nd order ordered statistic method, for when more target corruption is expected than true long-tailed clutter; or a Fisher-Tippett 2nd order CFAR for when long-tailed statistics may be present but with mild target corruption. In practice however, the Fisher-Tippett CFAR is far simpler to process and has been selected as the best compromise solution.

As previously observed the homogeneity and stationarity of the clutter in the littoral environment is poor. If a large number of spatial samples is gathered, implying that the statistics are gathered over a wide tempo-spatial area, the region around the cell-under-test must be as clear as possible of artefacts

such as buoys, harbour walls, cliffs etc. in order to maximise the performance of the Fisher-Tippett CFAR. The problem now becomes how to gather an effective 200+ cells whilst minimising corruption by target returns and also attempting to collect from regions where the statistics are as homogenous as possible.

To overcome these problems a novel self-organising system based on the use of multiple intelligent software agents (MISA) has been developed. The agent system reacts to features in the environment according to simple rules and modifies the areas over which the statistics gathering processes are performed accordingly such that the spatio-temporal data gathering is more effective and therefore better suited to the 2nd order CFAR process.

The key design philosophy has been to recognise that as the statistics of the scene are changing too rapidly to allow calculation to sufficient accuracy for idealised detection algorithms; any processing that is applied can only ever be sub-optimal.

The Multiple Intelligent Software Agent Spatio-Temporal CFAR Subsystem

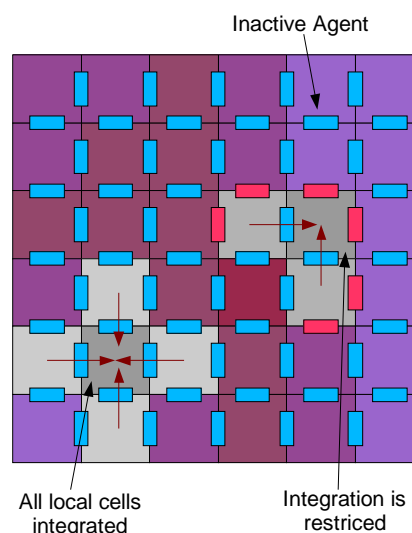


Figure 7 Layout of cells and agents

The CFAR cells are arranged as elements of a range-azimuth map. The map contains two types of agents: range-azimuth cells containing two identical IIR filters that perform temporal integration of successive

target returns and the square of the returns and ‘bridging agents’ located between pairs of range-azimuth cells that attempt to assess whether the pair of cells have similar statistics or not. This process is illustrated in Figure 7. The ‘bridging’ agents use the temporal statistics gathered by the range-azimuth cells in order to inform their decision, and if they decide that the pair of cells have statistically different characteristics, they prevent the spatial data gathering across the boundary of the statistical agents. The main effect is that land masses and other anomalous areas of clutter are formed into distinct areas. This ability to restrict spatial data gathering allows the statistics to better represent distinct areas of homogenous clutter. Local threshold calculations may be made based upon data which is comprised primarily of one probability density function, rather than being a compound distribution. In addition the problems observed with traditional CFAR where targets are difficult to detect in the vicinity of clutter boundaries have been largely eliminated.

Temporal Processing

Temporal processing is performed by two identical IIR filters that perform temporal integration of successive target returns, and the square of the returns. The IIR filter that calculates the mean is described by the following recurrence relationship

$$T_{\mu}(R, \theta, t) = \frac{0.9T_{\mu}(R, \theta, t-1) + I(R, \theta, t)}{1 + 0.9} \quad (8)$$

Where $T_{\mu}(R, \theta, t)$ is the temporal mean at each range, azimuth and time, $I(R, \theta, t)$ is the new raw input data. The filters produce the sum of exponentially decaying contributions from previous radar returns. The factor of 0.9 in the numerator and denominator provides the time constant of the filter. Approximately 10 scans of data influence the mean value that is calculated. In very non-stationary clutter, it is important to gather data only from a temporal history that can be assumed stationary. In some measured data that has been analysed, a maximum of 10 scans could be observed

before the stationarity of the data was questionable.

A similar filter, $T_{\sigma}(R, \theta, t)$, that sums the squares of the input voltages is also applied with $I(R, \theta, t)$ replaced by its square. Thus the variance (and therefore standard deviation) may be approximated as $T_{\sigma}(R, \theta, t) - T_{\mu}(R, \theta, t)^2$.

This simple approach to generating the mean and standard deviation of spatio-temporal regions allows the Fisher-Tippett 2nd order CFAR process to be applied very easily. In order for the Fisher-Tippett CFAR to be used, the logarithm of the radar return amplitude is passed to the statistic gathering process.

Spatial Processing

For the spatial processing, each range-azimuth cell has 4 intelligent agents, the bridging or B agents, shared with its neighbours around its borders, as shown in Figure 5. The B agents prevent the spatial integration from being disturbed by fixed targets and other boundaries. Each B agent monitors the $T_{\mu}(R, \theta, t)$ and $T_{\sigma}(R, \theta, t)$ values of the cells on either side of it, and if either $T_{\mu}(R, \theta, t)$ or $T_{\sigma}(R, \theta, t)$ are consistently different, it switches to a blocking state and prevents spatial integration occurring across the boundary. Each B agent maintains μ and σ values, the μ value, B_{μ} , being:

$$B_{\mu}(R+, \theta, t) = 0.9B_{\mu}(R+, \theta, t-1) + \text{sgn} \left(\begin{array}{c} T_{\mu}(R, \theta, t) \\ -T_{\mu}(R+1, \theta, t) \end{array} \right) \quad (9)$$

Where the notation $B(R+, \theta, t)$ denotes the agent that lies between cells (R, θ) and $(R+1, \theta)$ etc. The agent $B(R, \theta+, t)$ is the equivalent in the orthogonal grid direction. The use of the signum function rather than the raw difference results in an indication of the median rate of dissimilarity rather than the mean of the difference between the agents. Again, the factor of 0.9 used in the equation has been set to provide reasonable results on data where stationarity may be approximated over 10 scans of data.

The decision as to whether the agent should block is generated by identifying the B agents

which separate cells having the greatest dissimilarity (one agent for T_μ data and one for T_σ) Thus the agents with the largest difference between means, and the largest difference between the squared returns are identified. The magnitudes of these two values are then used to set a threshold to determine the bridging agent's activity. The agent will record $B(R, \theta+, t)=0$ if either the $|B_\mu|$ or $|B_\sigma|$ is greater than 70% of $|B_{\mu|\sigma}^{(\max)}|$.

Expressed in formal logic the truth value for the blocking action, for a single azimuth B agent is

$$B(R, \theta+, t) \Leftrightarrow \left(\begin{array}{l} B_\mu(R, \theta+, t) \\ \neg \left(> 0.7 |B_\mu^{(\max)}| \mid B_\sigma(R, \theta+, t) > 0.7 |B_\sigma^{(\max)}| \right) \end{array} \right) \quad (10)$$

Where TRUE and FALSE correspond to 1 and 0 respectively

This empirical rule provides a trade-off between adequate detection of clutter edges and false activations in the scene of the respective maximal values.

The spatial integration of the means is then described by:

$$S_\mu(R, \theta, t) = \frac{\left(\begin{array}{l} 0.9S_\mu(R, \theta, t-1) + 0.7T_\mu(R, \theta, t) \\ + \sum_4 (S_\mu(R \pm 1, \theta \pm 1, t-1) B(R \pm, \theta \pm, t)) \end{array} \right)}{\left(0.9 + \sum_4 B(R \pm, \theta \pm, t) + 0.7 \right)} \quad (11)$$

The integration of the squared returns is performed in a similar manner.

In the absence of blocking agent action statistics are gathered over an approximately 5×5 cell area. The result is that the statistics are based on ~ 250 samples allowing a reasonable estimate of the 2nd order statistics to be made.

CFAR Threshold

A threshold is calculated based on the S results and the Fisher-Tippet CFAR equation is used to threshold the input data in I . To

prevent moving targets from disrupting the mean and standard deviations, target detections, based on the thresholds calculated in the previous scan, are censored from the integration process. The censoring process simply prevents T level updates for any cells in which detections have been made. The censoring process is monitored to ensure that the censoring is only applicable to transient (i.e. moving targets) and that fixed targets are captured by the bridging agents.

The controlled spatial integration allows more samples to be gathered and more stable and accurate estimates of mean and variance to be obtained with edges in the scene preserved as sharp discontinuities. This process allows accurate thresholds to be determined to within a few cells of features within the environment. The new adaptive spatio-temporal CFAR is essentially an edge-preserving low-pass filter in a similar vein to median filters and Beltrami Flow [3]. However, median filtering and Beltrami flow derive the clutter-edge information independently for each scan and therefore rely on the noise levels to be much less than the intensity difference of the edges to be preserved. In contrast, with the new agent system, the clutter 'edges' are determined temporally over a sequence of scans and can therefore tolerate significantly more noise and are more stable between scans.

Conclusion

Many existing CFAR approaches will produce very good results if the clutter statistics are known exactly, but can perform badly if there is even a small error in the estimated parameters. The result is that by attempting to provide an optimal solution a very fragile process is created.

The self-adaptive spatio-temporal CFAR is proving effective at gathering large numbers of statistically homogeneous data samples from complex and difficult environments. The ability to gather large sample sizes means that robust estimates of threshold locations can be generated through 2nd order CFAR processing, reducing fluctuations in false

alarm rates and allowing depressed thresholds to be used in combination with a pre-track system. Even though the approach is essentially cell-averaging CFAR, the performance is proving to be extremely reliable in complex environments and processing losses are small as accurate threshold locations can be calculated.

Acknowledgements

The work reported in this paper was funded by the Electro-Magnetic Remote Sensing (EMRS) Defence Technology Centre, established by the UK Ministry of Defence and run by a consortium SELEX Galileo,

Thales UK, Roke Manor Research and Filtronic.

References

1. Anastassopoulos, V. Lampropoulos, G. 1992, *A New and Robust CFAR Detection Algorithm*, IEEE Trans AES. Vol 28(2), 420 -- 427
2. M. Skolnik, *Radar Handbook*, 2nd Edition, McGraw Hill, 1990
3. A. Spira, R. Kimmel, and N. Sochen, *Efficient Beltrami Flow using a Short Time Kernel*, Scale Space 2003, Scotland, UK, June 2003. LNCS vol. 2695, pp. 511-522