A Multiple Intelligent Software Agent based Technique for Improving the Radar Detection of Low Observable Small Craft in Sea Clutter.

E. J. Hughes, M. Lewis
Cranfield University, RMCS
Shrivenham, Swindon, SN6 8LA
e.j.hughes@cranfield.ac.uk, m.b.lewis@cranfield.ac.uk,

Abstract

The current method of detection of a radar target is based on the setting of a threshold determined by the average of the background returns in the region of interest. Problems arise with this method when attempting to detect small targets in littoral waters since in designing the detector it is necessary to make assumptions concerning the statistical behaviour of the background clutter. Since only long term data is available and short term prediction is required there is an inevitable missed detection/false alarm problem.

The problems associated with detecting low observable targets using Track-before-Detect systems based on Hough transform or Dynamic Programming techniques are reviewed. An alternative self-adaptive spatio-temporal CFAR system and a multiple hypothesis tracker based on Multiple Intelligent Software Agents are described.

Introduction

Coastal seas are important because of their economic value. In Europe and the Asia-Pacific region, local seas are crucial to prosperity since the majority of trade reaches its destination through major shipping ports.

The fact that coastal seas are transportation routes means that they are threatened by illegal activities of many kinds. The "good order at sea" necessary for legitimate fishing, mineral extraction, sea passage and tourism is imperilled by terrorist groups, traffickers in drugs, arms and illegal immigrants attempting covert entry or other forms of hostile activity.

The use of the sea as a highway for commerce makes it a target for pirates, as has been shown in the East China Sea and the waters off Southeast Asia, East and West Africa, the Caribbean and elsewhere.

Ports and ships may be regarded by terrorists as lucrative targets. The seizure by Chechen rebels in 1996 of a Turkish ferry in the port of Trabazon carrying Russians is an example [1].

The cost of piracy on the high seas and sea robbery (when the crime is committed within territorial waters) cannot be precisely known. In 1996, one reliable source estimated the worldwide total to be at least US$300-450 million per year. A more recent estimate put the commercial losses from piracy at US$16 billion. [2]

It is believed that terrorist groups, such as al-Qaida-linked Jemaah Islamiyah, are studying maritime targets in Southeast Asia [3]. The prospect of a supertanker being attacked and set careering through crowded sea lanes or sailed to an oil-terminal and used to ignite it, is appalling, not least in environmental terms.

Small Targets in Littoral Waters

Particular threats that can be perceived are small craft, such as rubber boats, or men with MANPAD weapons on jet skis. Such threats are generally low observable radar targets for two reasons. Firstly the small physical size means that they do not produce strong reflections compared with the sea surface which they are close to and surrounded by. Secondly the sea itself is moving and produces varying reflections with random high amplitude peaks that may be mistaken for targets resulting in ‘false alarms’. All sea reflections are referred to collectively as ‘sea clutter’.

Target Detection

In commonly used methods of target detection, target returns that cross a detection threshold are taken as ‘potential targets’. A continuously updated table of confirmed and potential tracks is then used to classify the new target returns into valid detections for existing tracks, possible targets worth investigating or noise.

The majority of the returns that are processed by the radar originate from the either the sea surface or from land. The returns from the sea are not constant in amplitude and vary steadily. In addition the returns from different areas of the sea surface within the radar coverage will have different behaviours.

With low observable targets many returns will be below the detection threshold and there may be many missing detections along the track, resulting in targets being classified as noise if re-investigated, tracks never being initiated, tracks being deleted early or each track being maintained for an extended period.

In order to increase the probability of detection of weak targets, the detection threshold must be lowered with a consequent increase in the number of false alarms due to the spikes of the sea clutter crossing the threshold.

As the information from the received signal is limited, a false alarm must be treated as a true target, until it can be established as false. The increased false alarm rate causes problems with the association of returns with tracks and leads to an excessive number of false tracks being reported with the consequent risk of the tracking system becoming overwhelmed.

The Behaviour of Typical Clutter Characteristics

Traditional radar detection systems make a binary decision, based on a threshold derived from the clutter level in adjacent range cells, as to whether the return is from a target, or noise/clutter. The decision mechanism directly affects the probability of target detection and the probability of a false alarm. The discrimination of false alarms is ultimately performed in the tracking system, and therefore 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 practice, real clutter is spatially non-uniform, requiring the threshold to be adjusted locally to maintain a maximum probability of detection, whilst not exceeding the maximum tolerable probability of false alarm. CFAR systems attempt to address this problem. The premise 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, then a threshold level can be calculated that will achieve the maximum tolerable probability of false alarm. To estimate the low-order moments, samples are taken in range from around the return of interest.

The fundamental assumptions are that:

- the clutter is locally uniform, allowing statistics to be generated spatially;
- the statistics of the clutter are stationary allowing accurate estimates to be generated temporally;
- the shape of the clutter probability density function is known;
- a low number of samples (typically 30) will provide a sufficient estimate of the statistics;

Unfortunately these assumptions do not hold except for a limited range of scenarios. One scenario where none of the assumptions are likely to be valid is the littoral environment.

It has been found that to gather sufficient samples to obtain a reasonable estimate of the mean and standard deviation, the samples must be drawn from a spatio-temporal region. In order to make the samples as consistent as possible, the 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

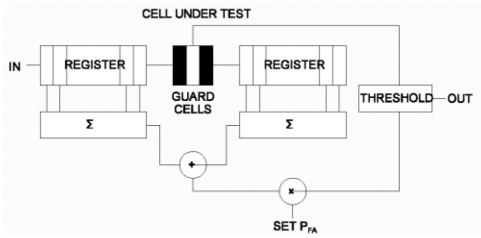

Figure 1 Constant False Alarm Rate by Cell Averaging

A key element in a modern radar system is the function that maintains the false alarm rate at a constant value. The false alarm rate is determined by the statistics of the clutter and the threshold level. In the case of noise the statistics are known and the threshold level can be set accordingly. If the noise power changes then the statistical parameters change and the threshold level should be changed (adaptive threshold).

There are a number of implementations of constant false alarm rate circuits. Each has its advantage and disadvantage. All 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 in the region of 1 - 2 dB.

In Figure 1 the signal or clutter seen in each range cell is stored in a series of registers. Every new pulse causes the previous set values to be overwritten. The great majority of cells will only contain clutter or noise. The threshold appropriate to any cell is determined by looking at neighbouring cells. The two cells immediately adjacent to the test cell are ignored because the target may be straddling the boundaries. The values in cells either side are added. This sum is *n* times the average, where *n* is the number of cells (10 - 20). The mean value can be scaled by a factor which determines the threshold level and hence the false alarm probability. The test cell is compared with this threshold and a yes/no decision made. Each range cell is tested in this way.

There are a number of disadvantages to this particular scheme. The statistical parameter is estimated from a relatively small number of samples and so will differ from the true population value. The estimate may be higher or lower that the true value leading to a range of possible threshold values. If there is any uncertainty in the threshold level then it must by set on the high side. The result is a slight reduction in the probability of detection that is equivalent to a small loss of signal to noise ratio, about 1 dB.

Using a large number of samples reduces the error in the measured parameter but increases the risk that some of the further 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 representative of the test cell.

A small number of cells will reduce the chance of non-representative values but gives a poorer measure of the parameter and a greater CFAR loss.

In some situation the clutter may change its characteristic in a short distance. For example there may be a band of rain with a well-defined edge or a cliff face dividing sea clutter from land clutter. As the test cell is moved down range the clutter boundary walks through the registers. One register will contain a clutter value that is representative of the value in the test cell and the other will have some non-representative clutter in it. If the test cell is in the low level clutter then the threshold will be set too high, the false alarm rate will be lower than it need be and the probability of detection will be reduced. If the test cell is in the high level clutter then the false alarm rate is too high, which cannot be allowed even though the probability of detection is slightly higher.

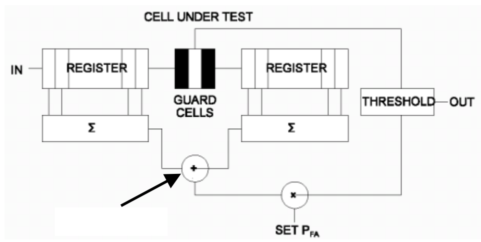

Figure 2 Modified Cell Averaging CFAR

The situation to avoid is increasing the false alarm rate. Instead of adding the contents of the two registers take the sum of each one separately and take the greater of the two to set the threshold (Figure 2). By doing this the excessive false alarm rate is avoided but at the expense of a slightly lower probability of detection, which is equivalent to a loss of signal to noise ratio.

Track-Before-Detect Techniques

Track before detect or 'pre-track' operation has been proposed in the past where either no threshold is applied, or a second low detection threshold is placed below the existing detection threshold to catch returns that did not quite cross the main threshold. These small amplitude detections are processed to see if they form tracks. The two main processing methods that have been proposed are Dynamic Programming [5] and Hough Transforms [5].

In radar, Dynamic Programming is an optimisation process that tries to identify the single most likely track through each cell. A number of algorithms exist and have been applied to such problems as human genome sequencing.

A Hough Transform treats the data from a number of radar scans as an image and the method looks for 'lines' within the data. The Hough Transform has its origin in particle physics where it has been used to locate particle trajectories.

Both methods are computationally intensive, the time to perform the calculation being proportional to N^3, additionally the Hough transform requires an extra step to re-associate returns with the set of possible tracks extracted from the transform. Fast approximate forms of the Radon, and the related Hough transform, also exist and require time proportional to $N^2 \log N$ [6].

A further problem is that both techniques are limited as to the geometry of the tracks that can be easily handled. Straight lines and circles are the geometries that these transform methods are best suited for.

Multiple Intelligent Software Agents

An intelligent agent is a form of software object that has the ability to store data internally; this data is known as the agent's state, and a set of methods, both public and private, that modify the agent's state dependent on the current input environment and the agent's current internal state. The software agent usually has the ability to affect its environment, thereby influencing its own future behaviour and the behaviour of other agents. The agents often have the ability to communicate directly with other agents in a system, enabling complex self-organising behaviour patterns to emerge. The use of cooperating agents leads to a highly parallel structure formed from simple elements. This allows the system to be flexible, expandable, robust and fast to process.

The Intelligent Agent Approach to Track Before Detect

The Track-Before-Detect problem has been investigated using multiple intelligent software agents, with the aim of producing novel alternative algorithms and has resulted in a Pre-Tracking system, rather than a Track-Before-Detect system that identifies tracks before applying thresholds. In the Pre-Tracking system, two thresholds are applied to the radar returns. The upper threshold is used to supply CFAR detections to the existing target tracking system (thus ensuring a level of performance no worse than conventional systems). The data crossing the lower threshold is used by the pre-tracker which is designed to tolerate a high level of false alarms.

The New MISA System

The key concept of the pre-track system is the exploitation of the spatio-temporal coherence of true target tracks, but with practical levels of processing. To achieve this, a self-adaptive spatio-temporal CFAR system is first used to identify 'interesting' radar returns. These 'interesting' returns are then passed to a pre-track system that attempts to associate the returns with previous returns according to a set of simple rules that define the likely feasible region that previous returns could lie in. The pre-track system does not make any explicit track predictions, unlike conventional multiple hypothesis trackers, but relies on associations between returns producing 'virtual' tracks within the data.

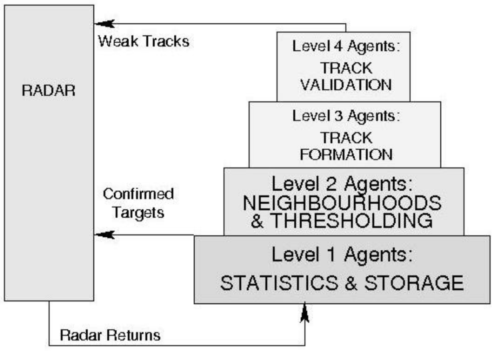

Figure 3
Functional Arrangement of MISA System

A system based on a hierarchical population of agents, each agent representing an individual radar cell that is allowed to self-organise into target tracks, has been constructed. Figure 3 shows a functional block diagram of the current MISA system. The radar system is shown on the left, feeding the radar returns into the lowest levels of the agent hierarchy. The radar returns at this point will have had all necessary processing applied prior to the application of a CFAR system and a threshold.

Level 1 and 2 agents form a Spatio-Temporal CFAR Subsystem whilst Levels 3 and 4 function as a multiple hypothesis track forming sub-system. The radar returns traverse the hierarchy, with high-confidence target detections being fed to the main radar tracker as track segments.

The Spatio-Temporal CFAR Subsystem, Agent Levels 1 and 2

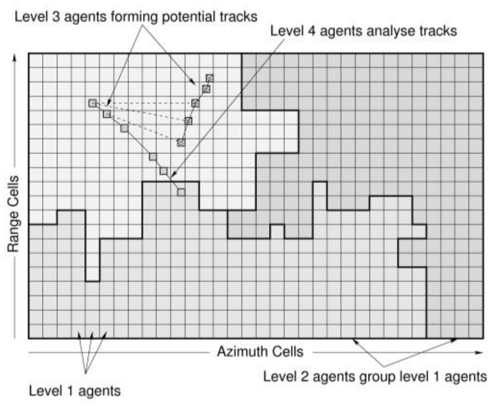

Figure 4. Cell-Level diagram of agent organisation

The basic functions of the Level 1 agents are to store a localised temporal history of the radar returns for their Level 1 range and azimuth cell, generate statistics of the stored data, and apply the two thresholds to classify a return as noise, a partial detection, or a full detection. The Level 1 agents of the hierarchy record the time and amplitude information of each return together with its detection classification, based on the two thresholds. Level 1 agents are organised into small clusters of similar cells having their own Level 2 agent as shown in Figure 4. All detections that cross the upper threshold are passed to the radar for processing as likely targets using the existing track algorithms. This guarantees that performance is no worse than conventional CFAR.

Level 2 Agents

Level 2 Agents are virtual agents formed by Level 1 agents communicating with near neighbours and linking to form clusters that have similar probability distributions, e.g., in a littoral environment, land clutter is likely to have a Rician like distribution, whilst sea clutter will more likely follow a log-normal, Weibull or *K*-distribution. The exact choice to determine 'similarity' is very dependent on how the threshold level is calculated.

The Level 2 agents adapt by exchanging Level 1 cells with other neighbouring Level 2 agents, as shown in Figure 4, in an attempt to form a cluster. Unlike the Level 1 agents the Level 2 agents do not have fixed spatial locations.

A small housekeeping structure is associated with each Level 2 agent. This monitors the statistics on the quantity and distribution of the detections and partial detections from the Level 1 agents it is responsible for, and also the statistics of Level 1 agents in the local vicinity (controlled by other Level 2 agents). These statistics, along with feedback from the main tracking system in the radar, are used to generate the upper threshold. Feedback from the Level 3 agents is used in conjunction with the statistics to set the lower threshold. These threshold levels for the Level 1 agents within the cluster are used for the classification of the radar returns. The distribution will affect the calculation of the positions of these thresholds relative to the mean, median and standard deviation etc., calculated by each of the Level 1 agents.

The Level 2 system creates dynamically re-configurable spatial awareness within the processing system, allowing better statistical estimates to be generated for the calculation of the threshold levels. This grouping allows spatial correlations of the underlying clutter to be exploited, as well as the temporal correlations held in the Level 1 agents.

Potential Track Formation, Level 3 Agents

Conceptually, as shown in Figure 4, Level 3 agents are formed with each being associated with a target return. When a Level 3 agent is created, it strives to form links with existing Level 3 agents that represent *virtual* tracks within the multi-agent system.

Using the 'An agent is a Detection' approach allows many track hypotheses to be formed for each return.

Agents that are marked as having the potential to be part of a track are then scanned to see if any previous links are recorded. If links exist they are checked to determine if the speed and direction changes are within the feasible region. The calculation of the feasible region for association of agents to allow links to be formed whilst keeping processing to an absolute minimum is one of the cornerstones of this research. Explicit forward prediction of likely positions is not used as the basis of the association error, only reverse checks on link and agent feasibility are performed.

If the new agent is within the feasible region, the importance of the link is calculated. This value can be used to prune the link set of the agent to reduce storage requirements.

Track Validation, Level 4 Agents

The primary function of a Level 4 agent is to assess the most likely path through a series of Level 3 agents and report the track to the main track database if it appears to be a true target. Level 4 agents are created when potential tracks are identified as a sequence of links formed between Level 3 agents. The Level 4 agent scans the track, looking for all the necessary correlations between stages that indicate a valid track is likely and eliminates unlikely tracks in the process. The Level 4 agent may also interrogate and analyse the target returns along the track in order to aid the track assessment by identifying possible missed detections. The Level 1 agent system is interrogated to see if a 'near miss' occurred when the data was thresholded. If a return is classified as belonging to a valid track at any time the Level 1 return may be promoted, the detection classification held in the temporal record being updated and the statistics describing the clutter updated accordingly. This process allows crisp tracks to be confirmed, some noise to be rejected, and areas of uncertainty to be identified.

As the number of agents reaches the upper limit of the processing capability, the life of the agents can be managed to allow a maximum population size to be maintained, whilst performance is allowed to degrade gracefully. This contrasts with conventional track formation where track overload can be catastrophic.

Once a track has been validated the track's elements are passed to the main radar tracker and the corresponding Level 3 agents notified that the track has been validated.

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.

In contrast the MISA process is, in effect, a simplified multiple hypothesis tracker, tightly coupled to a self-adaptive, context sensitive, spatio-temporal CFAR system. In environments with diverse clutter characteristics, the self-adaptive nature of the agent system self-organises using simple processing and by assuming that there will be too few data measurements to establish the clutter statistics accurately, a robust sub-optimal solution is formed.

The technique could be extended to IR and EO systems or used as a data fusion technique for multi-spectral sensing.

Acknowledgements

The prototype MISA pre-track system was developed as part of CRP programme contract no. RCSC/2/040.

In addition we wish to thank Dr., Sam Lycett, QinetiQ, Malvern and Dr., Andy Stove, Thales, Crawley for supplying radar return data that enabled the prototype MISA pre-track system to be evaluated.

References

1. "Turkey and the Chechens. End of a Love Affair?", L. Lingg, *News Insider,* 26/04/2005. http://www.newsinsider.org/editorials/turkey_chechnya.html

2. "The Growing Threat of Maritime Piracy", D. G. Wiencek, *China Brief,* Vol., 1, Iss., 1, 12/07/2001. http://www.jamestown.org/publications_details.php?volume_id=17&issue_id=629&article_id=4554

3. "Singapore PM says terrorists are studying maritime targets", *China Daily*, 04/06/2005, http://en.ce.cn/World/Asia-Pacific/200506/04/t20050604_3974156.shtml.

4. Arnold, J. Shaw, S. and Pasternack, H. 1993, *Efficient target tracking using dynamic programming*, IEEE Trans AES. Vol 29(1), 44--56

5. Storkey, A. J. Hambly, N. C. Williams, C. K. I. Mann, R. G, 2004, *Cleaning Sky Survey Databases using Hough Transform and Renewal String Approaches*, Monthly Notes of the Royal Astronomical Society Vol 347(1) 36—51

6. Brady, M. L. 1998, *A Fast Discrete Approximation Algorithm for the Radon Transform*, SIAM Jnl. on Computing, V27(1) 107—119