INTELLIGENT AGENTS

Multilevel agents could greatly improve the ability to track radar objects that prove difficult for conventional techniques

FOR RADAR SYSTEMS

by Evan Hughes and Michael Lewis

With difficult to observe targets and low-flying targets where multi-path effects can cause significant fading, many radar 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. 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. The consequent risk of this is the tracking system becoming overwhelmed.

Medium Pulse Repetition Frequency (PRF) radar systems allow all-round measurements of both the range and Doppler of targets in high-clutter environments to be made. Such radars use waveforms that are ambiguous in range, Doppler or both. Existing techniques that resolve these ambiguities require the number of detections input to the ambiguity resolution process to be kept to a small number, otherwise the number of false correlations or 'ghosts' becomes unworkably large.

PICKING TARGETS

Another significant issue that affects many look-down airborne radars is the difficulty in distinguishing between unwanted moving targets on the ground and targets of interest with low closing rates. Commonly, these unwanted targets are readily detectable, but must be excluded, by Doppler filtering for example, to keep the ambiguity resolution problem within bounds.

Track-before-detect or pre-track operation has been proposed 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 →

Improvements to radar track detection will reduce the false alarms that need to be monitored

detections are processed to see if they form tracks.

Our research uses a pre-track system that exploits spatio-temporal Doppler correlations to help reduce ghost targets, as well as reducing false alarms due to noise. Further, it makes use of intelligent agents. An intelligent agent is a form of software where its state depends on 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 co-operating agents leads to a highly parallel structure formed from simple elements. This allows the system to be flexible, expandable, robust and fast to process.

PROCESSING METHODS

The two main processing methods we have used are dynamic programming and Hough transforms. Dynamic Programming is an optimisation process that tries to identify the single most likely track through each cell. A Hough transform treats the data from a number of scans as an image and the method looks for 'lines' within the data. Both methods are computationally intensive with time complexity of $O(N^3)$. 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 have time complexity $O(N^2 \log N)$.

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 Constant False Alarm Rate (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.

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

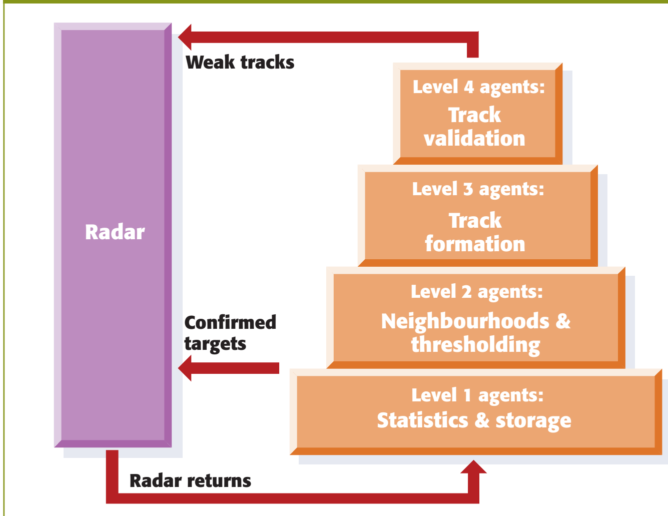

Fig 1: Functional arrangement of the MISA system

constructed. Fig 1 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 one and two agents form a spatio-temporal CFAR subsystem. Levels three and four function as a multiple hypothesis track-forming subsystem. The radar returns traverse the hierarchy, with high-confidence target detections being fed to the main radar tracker as track segments.

> "The capabilities of the tracker will determine the maximum false alarm rate that can be tolerated"

DECISION THRESHOLDS

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 just noise or 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. 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 and noise are spatially non-homogeneous, requiring the threshold to be adjusted to maintain a maximum probability of detection, without exceeding the maximum tolerable probability of false alarms. CFAR systems attempt to address this problem. The premise is that if the statistics of the noise or 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.

There are a number of assumptions that need to be made. The first is that the noise or clutter is locally homogeneous, allowing moments to be generated spatially. Second, the statistics of the noise/clutter are stationary, allowing accurate moments to be generated temporally. Third, the shape of the noise or clutter probability density function is known. And fourth, a low number of samples, typically 30, will provide a sufficient estimate of the moments;

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-Doppler region. In order to make the samples as homogeneous as possible, the region must be optimised to the current environment and, as 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.

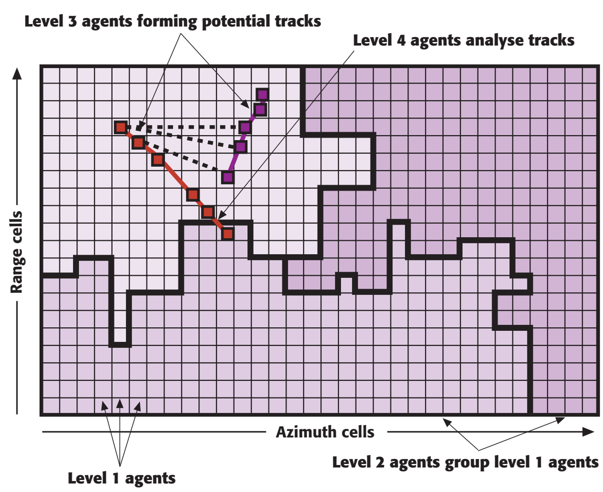

Fig 2: Agent organisation for one layer of Doppler cells

UNPICKING THE STRUCTURE

The invalidity of the above assumptions directly influenced the structure of the level one and two agent system.

The basic functions of the level one agents are to store a localised temporal history of the radar returns for their level-one range, azimuth and Doppler 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-one agents of the hierarchy record the time and amplitude information of each return together with its detection classification, based on the two thresholds. Level-one agents are organised into small clusters of similar cells having their own level-two agent as shown in fig 2. All detections that cross the upper threshold are passed to the radar for processing as likely targets using the existing →

> “Using the ‘agent is a detection’ approach allows many track hypotheses to be formed for each return”

track algorithms. This guarantees that the performance is no worse than conventional CFAR.

Level two agents are virtual agents formed by level-one agents communicating with near neighbours and linking to form clusters that have similar probability distributions. In a littoral environment, land clutter is likely to have a Rician-like distribution, whereas sea clutter will more likely follow a log-normal, Weibull or *K*-distribution. The exact choice to determine ‘similarity’ is highly dependent on how the threshold level is calculated.

The level-two agents adapt by exchanging level-one cells with other neighbouring level-two agents, in an attempt to form a cluster. Unlike the level-one agents the level-two agents do not have fixed spatial locations.

FEEDBACK STATISTICS

A small housekeeping structure is associated with each level-two agent. This monitors the statistics on the quantity and distribution of the detections and partial detections from the level-one agents it is responsible for, and also the statistics of level-one agents in the local vicinity, which are controlled by other level two 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-three agents is used in conjunction with the statistics to set the lower threshold. These threshold levels for the level-one 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 and so on, calculated by each of the level-one agents.

The level-two 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-one agents.

The output of the level-two agents is, in effect, a target and clutter map in range, azimuth and Doppler. In the case of medium PRF radar the targets within the map are ambiguous in range and Doppler. The ambiguities may be resolved by operating on N PRFs, typically eight, and requiring target data in a minimum number, M, typically three, in what is generally known as an M of N (three of eight) scheme. Individual level-one and two systems are used for each PRF. A single target will appear at a distinct but ambiguous range in M of N level-one and two maps. The MISA technique exploits spatio-temporal correlation between returns to extend the ambiguity resolution process of a traditional M of N detection scheme such as the coincidence algorithm. This efficient coupling of ambiguity resolution with target tracking function provides an enhanced capability to distinguish true from false targets, and allows a higher false alarm rate to be tolerated with the generation of fewer ghost targets.

Conceptually, level-three agents are each being associated with a target return. When a level-three agent is created, it strives to form links with existing level three agents that represent *virtual* tracks within the multi-agent system.

Using the ‘agent is a detection’ approach allows many track hypotheses to be formed for each return. A similar approach has been used in systems based on dynamic programming, but no threshold was applied and every possible return was considered, which required massive amounts of processing. The dynamic-programming approach has to rely on high-resolution Doppler information in order to reduce the number of possible linkages that can be formed. In the MISA approach, only coarse Doppler information is available, if any, and potentially very slow to stationary targets may also be considered. It is assumed that given the limited information available, many tracks could pass through each level-three agent.

Agents that are marked as having the potential to be part of a track are then scanned to see if any previous links are

Objects that are otherwise difficult to track could be identified using intelligent agents

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.

LINK IMPORTANCE

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.

The primary function of a level-four agent is to assess the most likely path through a series of level-three agents and report the track to the main track database if it appears to be a true target. Level-four agents are created when potential tracks are identified as a sequence of links formed between level three agents. The level-four 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-four 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-one 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-one 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.

> "The multi-level system has the potential to provide significant pre-track capability"

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, while 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 three agents notified that the track has been validated.

Results to date have demonstrated that, for many target scenarios, the ghost targets generated during the ambiguity resolution process are not correlated in a spatial-temporal-Doppler sense and are rejected easily by the MISA system.

There are indications, however, that some target types, such as radial targets in close formation, may lead to ghosts that appear to be correlated in space, time and Doppler for a significant period.

The levels of processing required to implement the level one and two systems has been investigated and the capabilities of the full self-adaptive spatio-temporal CFAR system demonstrated. Multi-agent code has been written which has allowed the full dynamic threshold control system to be integrated with the level three process and tested. The results, when compared against conventional methods including cell-averaging CFAR, indicate that the multi-level system has the potential to provide a very significant pre-track capability.

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 current CFAR techniques, by attempting to provide an optimal solution, can create a very fragile process.

In contrast, our 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. ■

Evan Hughes and Mike Lewis are researchers at the Cranfield University Royal Military College of Science