Modelling Information and Decision-Making under Uncertainty for Integrated Surveillance Operations

P E Berry, D L Hall, D A B Fogg and V Fok*
Surveillance Systems Division
Defence Science & Technology Organisation
PO Box 1500
Salisbury SA 5108
Australia
+61 (0) 88259 5899
paul.berry@dsto.defence.gov.au

Abstract

In order to obtain a realistic assessment of the capability of military surveillance assets it is necessary to consider the complete range of contributing factors not just the performance of sensors in isolation. If simulation is to be the means for carrying out the assessment, the environment and all relevant physical entities, stationary and mobile, can readily be incorporated. As well as platforms and sensor models, representations of the processes involved in sensor tasking, information collation, fusion and dissemination also need to be incorporated. Of particular importance is the fact that surveillance information is inherently uncertain and incomplete due to imperfect sensors, inadequate coverage, environmental effects, conflicting data and degradation of information over time. Also of importance is the modelling of decision-making based upon this uncertain information. In order to establish a basis for constructing and using Measures of Merit (MoM) in surveillance systems assessment studies, the paper begins by reviewing progress in the area of incorporating into the assessment human requirements associated with conducting surveillance and extracting information from surveillance sensors and, where advantageous, combining it with information from other sources.

The paper describes progress within Surveillance Systems Division in modelling processes for surveillance operations and formulating appropriate Measures of Merit for their assessment. The purpose of this work is to compare the effectiveness of different options for surveillance and strategies for integrating surveillance sensors through sensor and information management. Specific issues of interest to surveillance are addressed as an initial application of the model and to demonstrate benefits in the approach. These issues are representative of the information required at different levels in the military structure ranging from the sensor operator to the surveillance mission commander.

The approach adopted explicitly represents information and its associated uncertainty using a probabilistic formulation. Evidence generated by surveillance sensors and other information sources over time is accommodated by updating information using Bayesian inference. Measures of Merit appropriate to the mission of interest are derived from the uncertain surveillance information: the value of information is measured using Shannon's entropy, whereas operational

* This work was supported by the Australian Department of Defence.

outcomes are expressed as expectations of random variables. The potentially large state space for the surveillance system is rendered manageable by exploiting the conditional interdependencies of variables to construct a Bayesian network.

Entropy based techniques enable weighted contributions from disparate systems to be combined to form a single measure which varies with time indicating the improvement or degradation of knowledge relevant to the user requirement.

The models of decision-makers are implemented as software agents conforming to a BDI (beliefs, desires and intentions) architecture and forming a C2 hierarchy. Measures of Merit are required not only to quantify the effectiveness of the operation as a whole for the purpose of assessment, but as indicators of how well the various goals and sub-goals set for the agents have been achieved, and to determine whether in order to achieve an overall goal, a sub-goal needs to be changed. Criteria for performing actions and for changing goals in accordance with the prior plan are therefore MoM-based. As well as providing MoM for analyses and simulations, these techniques can be used as the basis for command support tools.

1. **Introduction**

This paper describes a model for simulating surveillance operations involving a diverse range of assets for the purpose of assessing effectiveness and facilitating the exploration of issues of surveillance system architecture and operational doctrine. It is intended that such a model should provide a testbed for examining proposed concepts and comparing competing options aimed at enhancing surveillance operational effectiveness in the Australian theatre of operations, and thereby contribute to systems studies involving consideration of capability.

Since the purpose of the model is to predict surveillance operational effectiveness it must address the contribution of the surveillance system to the overall outcome of a military operation. Therefore it needs not only to incorporate representations of the individual sensors, their platforms and the opposing targets, but also the way the information emanating from those sensors is managed and manipulated, and how the sensors are coordinated and controlled. For, fundamentally, it is through the process of generating and utilising information that surveillance effectiveness can be understood and explored. Consequently, information, or its effect, needs to be represented in some way, not to mention the dynamics of the military operation, as the outcome of a complex operation is critically dependant upon the timing of significant events.

Therefore, as with most military simulations, the model is required to represent the locations and dynamics (or at least kinematics) of physical entities, and aspects of the environment which they occupy, particularly as it affects sensor performance. However, the need to represent information renders the approach adopted in this paper to be somewhat different to the norm. The model under consideration is primarily intended as a constructive simulation model, meaning that there is no interaction with any human player serving the role of, say, a commander, pilot or sensor operator as would often be the case in a wargame or training simulator. All entities therefore need to be preprogrammed to behave according to an understood or agreed set of rules. These can simply be responses to physical events or prevailing conditions and require the programmer to

ensure that the entities will not experience conditions inappropriate to their internal rules. When entities not in physical contact with each other interact they are effectively communicating and responding to information received. The representation of the *effect* of information specific to an entity on its behaviour is often incorporated by the inclusion of entity scripts, as with Virtual Prototype's Stage and MODSAF. There may additionally be issues to do with the time taken to receive and respond to messages containing information but this can usually be represented by incorporating delays.

If one considers surveillance as a continuous estimation process in which measurements from a range of sensors are processed, combined and analysed over a period of time, with a surveillance 'picture' emerging gradually rather than instantaneously, never achieving 'ground truth' because of the dynamics of the situation and the scarcity of resources, and the operation of the sensors themselves potentially influenced by the information collected, then it becomes clear that a simple deterministic rule-based approach to modelling surveillance is inadequate. For when is a condition considered to be fully satisfied or an event to have unequivocally occurred to warrant a prescribed response when information is always incomplete and can even be contradictory?

It follows that in addition to the physical domain (targets, platforms, sensors and environment) there needs to be an explicit representation of the 'informational' domain. Whereas in the deterministic approach information can be represented simply by the assignment of values to variables, surveillance information is characterised by its inherent uncertainty, the consequence of which being that a given variable can assume one of a range of possible values, each with a probability, the sum of which is unity. We need to be able to quantify the information uncertainty so that the system can seek to minimise it within the context of the surveillance mission, and can take account of the effect of new evidence as it becomes available.

At this point it is worth cataloguing the sources of uncertainty. The sensors themselves have performance limitations relating to power, range, noise and signal return strength (dependant upon target cross-section), the result of which being that at any instant a given target may not be detected or that a false alarm occurs. There will be environmental effects which determine not only whether a target is detected (eg. clutter, cloud cover, ionospheric state) but also the errors associated with the measurement such as location estimate (eg. over-the-horizon radar (OTHR) coordinate registration error). Because of resource limitations a sensor or collection of sensors may have inadequate coverage in relation to the dynamics of the situation being monitored. Where there is overlap between sensors' ranges there may be ambiguity with regard to whether detections relate to the same or different targets, and if they are the same disgreement in the estimates of speed or location. There may be competing hypotheses for a track formed from associations between detections from a single sensor, or for a track formed from the fusion of tracks from several sensors. The uncertainty associated with information will increase over time unless evidence is continually received because of the dynamics of the situation being monitored and the unpredictability of the target behaviours.

What is required is a formalism which is capable of taking account of all of these effects in a consistent manner and which can be practically implemented within a simulation model. If this can be done and uncertainty quantified then surveillance effectiveness can be gauged by the level

of uncertainty associated with the relevant information provided to the overall mission commander. However, this is only part of the story, for information collected can be used to enact responses which influence the way the operation evolves, both in terms of controlling the surveillance assets effectively and in terms of the force response which aims to achieve the desired outcome (mission objective). In fact, the representation of this 'feedback loop' is essential, for without it there would be just a one-way upward flow of information, the accumulation of which would be simply a measure of the performance of the surveillance system. Furthermore there would be no way of enhancing the effectiveness of the system as a whole except through the manner in which the resulting information is processed to form the surveillance picture.

Therefore the formalism must also account for actions taken in response to information received given that information is inherently uncertain. In effect we are modelling decision-making under uncertainty but we assume that the possible decisions have already been prescribed, through the compilation of a plan, so that the purpose of the decision model is simply to determine the timing of a decision point and select an appropriate response at that point. There are three types of action taken in response to information: actions which simply affect the way information is processed such as assuming the truth of a particular hypothesis which causes its consequences to be further explored, actions which determine the way in which sensors and their platforms are employed and affect the information subsequently collected, and force response actions which contribute to the outcome of the overall mission which is the ultimate means by which surveillance effectiveness can be gauged.

The proposed model accounts for uncertainty through the construction of an appropriate state space from the random variables of interest and over which is defined a joint probability distribution which is updated in response to new sensor measurements using Bayesian estimation and also incorporates the effects of information degradation over time by representing stochastically the unpredictable behaviours of targets. The size of the state space constructed over all possible random variables would normally be unwieldy so the concept of Bayesian networks is used to decompose the problem. Decisions are modelled using statistical decision theory.

The modelling formalism based upon Bayesian inference and statistical decision theory is implemented within a modular model architecture. This is done so that the scenario, within the context of which the surveillance system is assessed, can be modified or replaced without impacting upon the physical entity models or the surveillance system component models. Thus the intention has been to embody as much as possible of the scenario-specific algorithms and data in software modules which are essentially plans. The intention of the paper is to demonstrate the scalability of both the modelling methodology and the model architecture rather than to develop a detailed and precise model of any specific surveillance architecture or surveillance operation. The scalability of the approach ensures that a model can be developed to any required level of fidelity (ie. faithfulness to reality). For this to be possible, all stages in the surveillance process need to be represented, from the specific sensors, platform, targets and environment occupying the physical domain through to the military operation supported by the surveillance activities, and only within the context of which effectiveness can be gauged. By adopting a

globally consistent modelling approach to each stage of the surveillance process, scalability is ensured.

The modelling methodology relates to how the surveillance aspects of an overall operation are modelled, in terms of both the management of the information arising from the sensors, and the control and coordination of the sensors and their platforms, in pursuance of the surveillance mission objective as determined by an operational commander requiring surveillance information. The result of the modelling procedure is a detailed specification for how information is represented and processed and how decisions based upon it are made, in terms of data and algorithms. The scalability of the methodology facilitates the inclusion of more detailed and accurate representations of stages in the surveillance process without necessitating a redesign of the entire model.

The model architecture is the detailed specification of the software needed to implement the algorithmic model. By adopting an *agent-oriented* architectural design in which closely related data and algorithms are encapsulated (as in object-oriented design) and, furthermore, are associated with goals which are interrelated through a goal hierarchy, a modular design results which is consistent with the modelling methodology. In addition to scalability, another requirement of the model architecture is to be able to deal with all of the potential interactions that could conceivably occur in a complex, dynamic system. These are difficult to predict so we have chosen to implement the architecture within a flexible and visual modelling environment provided by MATLAB's 'STATEFLOW' product in order to accomodate any future requirement for model architecture reconfiguration.

The specific operation modelled for the purpose of demonstration has been chosen to address a cross-section of the surveillance modelling issues at all stages in the surveillance process. The 'surveillance process' referred to is the general sequence of activities leading to the satisfaction of the surveillance information requirement, namely: search, detect, locate, track, estimate speed, classify, identify.

2. Scenario incorporating surveillance

The modelling formalism and model architecture are explained within the context of a specific surveillance operation in which Orange (opposing) forces are following a prescribed course of action and Blue (friendly) forces are responding in accordance with a plan which is driving the response of the surveillance system to physical events. The Blue force plan is specific to the scenario and Blue force mission, of course, but as the driver of the surveillance system to be modelled it is necessary to appreciate the overall context of its employment.

We have chosen to base our model of surveillance operations on a simple test scenario and incorporate appropriate surveillance concepts into that scenario designed to study future technology to allow the exercising of air dominance and offensive strikes. The test scenario has been developed and analysed only to determine which analysis techniques will be appropriate for the fully detailed and rigorous study. This scenario and analyses of the scenario have been previously reported on by various groups working in DSTO [Marlow *et al*., 2000; Hall and Berry,

2000]. Physical entities such as Orange and Blue aircraft and surveillance platforms and sensors are defined in the scenario and events associated with entities satisfying certain conditions are generated during the simulation of this scenario. These events (such as detections of enemy entities and attacks on the Blue garrison) are the basis of the information used to measure the entropy of the surveillance system and effectiveness of the decisions and surveillance plans generated during the simulation. The scenario is summarised here

An island off the coast of Australia (Airamak) is under threat of occupation by an adversary (Orangeland) which is demonstrating an intent to attack. Blue forces on Airamak are to garrison the island against an attack but have been assessed as incapable of protecting the assets (port, runway, fuel) on the island. Protective forces (airforce, GBAD) travel to Airamak but the island is attacked prior to their arrival. The Blue garrison and protective forces will be reinforced by a Blue naval force. This naval force will be attacked by the enemy as it transits to Airamak. Once it arrives the resupply of Airamak begins while the defence of Airamak against the Orange force continues. We are able to vary the numbers of entities in the scenario and the times between events occurring. This scenario has 3 main phases – a preparation phase, a reinforcement phase and a re-supply phase. During the preparation phase the Blue force assembles the naval force in preparation for the reinforcement phase, and the Blue garrison comes under attack by the Orange force. We are only concerned in this report with the preparation phase.

While Airamak is under threat of invasion basic wide area surveillance (WAS) will be performed of the area around the island to enable situation awareness and threat assessment. The surveillance assets available are satellite sensors for ocean surveillance and OTHR for air surveillance. If the perceived level of threat to the island changes significantly (caused by information from the WAS or intelligence) then the surveillance requirements could possibly change such that focal area surveillance (FAS) of specific areas could be necessary. Assets available for this part of the scenario include an AEW&C aircraft and retasking of the other WAS assets. Being able to measure the effectiveness of tasking and retasking these assets in an integrated manner is, of course, part of the impetus for this study. The effectiveness of the surveillance parameters determined by entropy based algorithms (such as optimum satellite sensor look angle) is also of prime importance to the study. Information from FAS (such as identification of a squadron of enemy fighters en-route to Airamak) or other events during the simulation (such as an actual attack on Airamak by the Orange fighters) will cause the surveillance requirements of the Blue force to alter again. The physical models of the entities, the simulation, the entropy and probabilistic based algorithms, and information processing modules to determine decisions and measure the performance of the surveillance system have all been implemented in the MATLAB technical computing environment.

3. **Formalism for the information model**

The approach taken to the problem of modelling processing of information from sensors and the controlling of sensors in response to that information is to decompose the entire process into a set of modules which are organised hierarchically. This hierarchical organisation fits well with the hierarchical approach to specifying a plan, which can be considered to comprise a goal hierarchy. This enables the decision modelling modules, which are designed to achieve goal satisfaction, to

be separate from the inference modules which simply process information from sensors and generate measures which are used by the decision modules. This model architecture simplifies the interfaces between the modules so that they can be flexibly reconfigured for different scenarios.

In general, a goal at a particular level is invoked by a goal at the next higher level when an invocation condition is satisfied. The goal will be maintained until either it has been satisfied (as determined by a goal satisfaction condition) or has been terminated by the next higher level because it has been deemed to be no longer relevant. While the goal is maintained it executes its plan, which involves it invoking sub-goals when appropriate conditions are satisfied.

3.1 Goal Hierarchy

The purpose of a surveillance system, as with any asset supporting a military operation, is to achieve an objective set by a commander. In the case of surveillance the objective may be stated in terms of information that is sought regarding targets in a particular area and during a particular timespan, which may conform to a particular type or be behaving in a particular way. We adopt a goal-based approach to specifying a military objective in which the objective is associated with a measure, which can be determined and assesses the degree to which it has been achieved, and with a plan which determines how it is to be achieved, as well as a condition which determines when it has been achieved. The measure associated with a surveillance system goal could be the quality of information delivered by a certain time, or the time taken to deliver information of a specified quality.

However, information measures really just assess the performance of the surveillance system. If the intention is to assess effectiveness then what needs to be determined is the difference the system makes to the overall outcome of the military operation, not just the surveillance-specific component. Therefore a higher level goal is needed stating the overriding military objective, associated with a Measure of Outcome (MoO) which determines how well the objective has been met. This enable the sensitivity of the outcome to be explored as a function of the surveillance system's characteristics. If resource utilisation is of interest then a more complex Measure of Effectiveness (MoE) can be specified which factors the utility of the outcome against the cost of the resources used. This could be useful for comparing force capability options.

We see then that the surveillance system goal is just a sub-goal of the overall operational goal. It is, in fact, possible to specify a goal hierarchy with a corresponding measure, plan and termination condition for each goal. For convenience we refer to the highest operational level as Level 4. At this level the plan is designed to achieve, or optimise, the overall mission objective measured in terms of some physical outcome such as damage sustained, number of casualties or area of territory occupied. The Level 4 plan itself consists of sub-goals corresponding to plans which are invoked when appropriate conditions are satisfied. These conditions relate to information about the threat posed by observed targets. Therefore Level 4 also needs to perform a threat assessment based upon information arising from the surveillance system as well as other possible sources such as intelligence reports. We see that at Level 4, as with all other levels, there are two distinct processes: executing a plan by deciding actions in response to information about

events or conditions, and generating the required information upon which the decisions are based.

Level 3 corresponds to the surveillance system and incorporates situational awareness, that is the compilation of all information relevant to the mission at hand as specified in the Level 3 goal. When invoked it executes its plan by tasking sub-components of the surveillance system to undertake either wide area surveillance or focal area surveillance. It is concerned with where the system should direct its attention and what specific target information is sought in response to changing requirements from Level 4 as the threat escalates.

Level 2 is referred here to as the multisensor integration system and is responsible for directly tasking the sensors and their platforms. It is concerned with surveillance integration issues such as sensor cueing, coordination and information fusion. Two surveillance integration issues have been selected for representation at this level, namely the process of first searching for and then locating maritime targets using satellite-based sensors (the task requires a different strategy for each stage of the process), and the association of airborne target tracks arising from wide area and focal area surveillance sensors and the consequent fusion of information arising from each.

Level 1 is concerned with the direct control of the sensors and their platforms and the goals at this level are essentially control objectives.

As mentioned above, each level incorporates two components: one executing a plan and the other processing information. The information processing at each level is based upon Bayesian inference. It essentially updates probability distributions for parameters of interest to the plan execution module arising from new sensor data. The plan execution module has to determine when and how to respond to changes in that information so as to achieve its goal. Whereas all possible responses are prescribed in the plan, the precise timing and satisfaction of conditions for their invocation need to be ascertained. This employs elements of statistical decision theory. These are discussed in the next sections.

3.1.1 Goals chosen for scenario

The Level 4 goal is to minimise damage sustained by Airamak facilities. This requires an attrition module to represent the attack on Airamak facilities by Orange force aircraft, the outcome depending upon whether the island's air defences have been sufficiently forewarned by the surveillance system of an impending attack.

Level 3 has two goals, each relating to a requirement for surveillance information requested from Level 4 and depending upon the perception of the threat posed by Orangeland. One is searching for, estimating the location of and tracking all airborne and maritime targets in the region of Airamak. The other is to determine if any of them have departed from Orangeland and are on course to Airamak and seek out further information about them which could be used to assess Orange force intentions. The first goal is invoked for different regions, the area of operations around Airamak or the vicinity of Orangeland.

There are three Level 2 goals, namely wide area surveillance of maritime targets using satellite surveillance sensors, wide area surveillance of airborne targets using over-the-horizon-radar, and focal area surveillance of airborne targets using a combination of AEW&C and ground based air defence radar.

The Level 1 goals are specific to sensors and their platforms. The satellite-borne sensor has two alternate goals, one which searches a region with no *a priori* knowledge in order to determine the existence of maritime targets within it, another which determines the locations of known targets within a region, each having a strategy which is optimised to satisfy the goal in question. The over-the-horizon radar has a goal which is to determine the existence and speed of, and track, all airborne targets within a given region. The focal area surveillance assets (AEW&C and ground based air defence radar) have a goal which is for them jointly to detect, locate, track and identify targets already deemed to be of interest and which have been cued by the OTHR.

3.2 Inference model

The inference model updates information relevant to each level in the goal hierarchy as a result of measurements arising from surveillance sensors. At each level information is represented as a probability distribution over an appropriately defined state space. These distributions need to be initialised, need to be updated using Bayes' rule as measurement data arrives, and need to evolve over time in a manner which represents the degradation of information when no data arrives. The inference model also uses this information to generate entropy-based measures required by the decision model. The conditional dependence of states at one level on the states immediately below is such that the inference model is, in fact, a Bayesian network [Jenson, 1996]. This means that as evidence arrives from the sensors at Level 1, state probability updates are propagated upwards through the levels.

3.2.1 States chosen for scenario

State spaces need to be carefully designed so that they are mutually exclusive and, within the context of the chosen scenario, exhaustive due to the probabilistic formulation of the problem.

Level 4 has only three discrete states (threat states) which represent Blue force's perceptions regarding Orange force's intentions resulting from a threat assessment. State 1 represents an initial heightened state of alert resulting from the deteriorated political situation but with no expectation of Orangeland having hostile intentions; state 2 represents an expectation of Orangeland intending to provoke or cause political embarrassment through military action but with no intention of causing material damage; and state 3 represents belief that an attack is in preparation or in progress with a definite intention to cause material damage to Airamak facilities. The random variable over the states is indicated as $\boldsymbol{T}$ in Figure 1.

Level 3 (situational awareness) has states representing the number, locations, speeds, tracks and identities of all airborne and maritime targets within the region as point estimates with associated error distributions. This surveillance 'picture' is compiled from information at level 2 from specific sensors by combining, extracting or fusing data. Its validity is dependant upon

hypotheses relating to the correctness of associating tracks (in the case of airborne target surveillance) and of extracting point estimates for locations (in the case of maritime target surveillance) from a continuous distribution. Point estimates for tracks and locations for airborne targets arising from the over-the-horizon radar are inherently uncertain due to the nature of the sensor. In the case of the satellite-based sensor, the uncertainty for point estimates of maritime target locations arising from the continuous p.d.f. originates not from sensor resolution limitations (it is assumed to be ideal) but from target motion during the satellite revisit times. The probabilities for these hypotheses are monitored as they influence inferences drawn about the state of threat at Level 4. The random variables for the target tracks and identities are shown as ***t*** and ***ID*** respectively in Figure 1. The assumptions regarding the targets' origin, destination and number are shown as ***O***, ***D*** and ***N*** respectively.

Level 2 (multisensor integration) directly tasks the sensors to acquire the information which is compiled at Level 3. The sensors may not be capable of generating the required information directly or individually, but through a process involving the integration of sensors and their information such as cueing of one sensor by another, and association of information obtained by different sensors. The means by which sensors can be integrated to provide greater effectiveness are assumed to be built into the plans. At this level the states monitored are, for wide area surveillance, the estimated number and locations of all maritime targets, recorded not as point estimates, but as a probability density function over the region of interest, and the estimated number and tracks of detected airborne targets recorded as point estimates of location and velocity with associated errors. In Figure 1, subscripts *m* and *a* refer to maritime and airborne targets respectively.

Since Level 1 is concerned with generating measurements, estimation is at the level of determining whether a detection has occurred or not and, if so, what the associated error is. This process is inherent to the sensor models used and so is not explicitly represented in the surveillance system model. In fact the over-the-horizon radar model used generates tracks directly from sequences of detections. For present purposes we choose to assume that these tracks are correct in the sense of being associated with true targets, but that not all targets need necessarily be tracked and that errors are associated with both the tracks and the location estimates. In future we could explicitly represent the tracking process in a manner consistent with the overall formalism using Bayesian estimation if the need arises. In Figure 1, ***y*** refers to location measurements whereas ***x*** refers to location estimates.

3.2.2 Inference algorithms

If there exists a prior probability distribution *p(x)* for a random variable *x* whose value is to be inferred from a measurement *y* then Bayes' rule states that in the light of the new evidence *y* the posterior distribution can be obtained thus

$$\Pr\{x \mid y\} = \frac{L(y \mid x)\, p(x)}{p(y)}$$

where *L(y|x)* is the likelihood function, the probability for measurement *y* being obtained for a given system state *x*. The distributions over the states at each level are inferred from the

distributions over the states at the level below using appropriately derived likelihood distributions.

We begin by describing the estimation models for the sensors. It is important to clearly state the assumptions behind these models as this affects the fidelity of the overall surveillance system model.

The satellite-borne sensor is assumed be steerable with a look angle determined in advance. Within its swathe it detects maritime targets with a specified probability of detection but also receives false detections. The region of interest is represented as a collection of rectangular cells, the size chosen such that at most only one detection can occur within it at any time. There can be any number of targets within the region of interest. Each time the satellite sweeps across part of the region it downloads data from which information about target detections is extracted. Between satellite revisits targets are assumed to move according to a Gauss-Markov target motion model based upon historical data, for example relating to shipping lanes. A joint probability density function is defined for the number of targets and their locations and is globally updated on the basis of new detection events.

The likelihood function for this estimation process is stated but not derived. The cells composing the region are labelled from *1* to *M* with *k* targets under the satellite swathe. The set of targets is labelled $T = \{ j_1 \ldots j_k \}$. If *m* detections occur with the set of detections labelled $D = \{ i_1 \ldots i_m \}$ then the likelihood of obtaining the *m* detections *D* given that there are *k* targets *T* is

$$L(D \mid T\} = p_d^{m'} (1 - p_d)^{k-m'} P_{fa}^{m-m'} (1 - P_{fa})^{M-(k+m-m')}$$

where p_d is the probability of a target being successfully detected and P_{fa} is the probability of a false alarm occuring in a cell. This likelihood function enables the joint probability distribution for the number of targets and their locations to be updated in response to detection events given that some detections are false alarms and some targets are undetected.

The over-the-horizon radar model generates tracks from detections. It represents the probabilistic detection of targets from a signal-to-noise ratio derived from environmental and target characteristics, and forms these into confirmed tracks using the latest 5 out of 8 detections. It does not represent any false detections nor does it represent multiple track hypotheses. Although a sophisticated sensor/environment model, its representation of tracking is relatively crude and capable of improvement. However it is considered suitable for present purposes with the emphasis on modelling the entire surveillance system to a reasonable level of fidelity. We assume that all tracks generated correspond to actual targets but that not all targets need necessarily be successfully tracked or even detected. Also a track may be lost and later regained. For simplicity, a track is represented as a point location estimate and velocity estimate, with associated errors. Within the context of the scenario, of interest is the origination of the tracks, ‘specifically did a target originate from Orangeland?’

We also need to estimate the total number of airborne targets given the performance of the sensor and the fact the tracks can be lost and regained. We estimate the probability distribution for the total number of targets at any time given the number of targets being tracked and events

corresponding to the acquisition of a new track or the loss of an existing track in an increment of time Δt. If the rate of generation of new tracks from an untracked target is assumed Poissonian with rate λ and the rate of loss of tracks from a tracked target is also Poissonian with rate μ then the likelihood functions are

$$L(\text{new track} \mid \mathrm{N_T} \,\&\, n_t) = (1 - n_t \mu \Delta t)(N_T - n_t)\lambda \Delta t$$
$$L(\text{track lost} \mid \mathrm{N_T} \,\&\, n_t) = n_t \mu \Delta t[1 - (N_T - n_t)\lambda \Delta t]$$
$$L(\text{no change} \mid \mathrm{N_T} \,\&\, n_t) = [1 - n_t \mu \Delta t][1 - (N_T - n_t)\lambda \Delta t]$$

to $O(\Delta t)$ where N_T is the unknown number of targets and n_t is the current number of tracked targets. These likelihood functions enable the probability distribution $\Pr\{N_T \mid n_t\}$ to be updated as tracks are gained or lost. This assumes that Δt is sufficiently small that at most one track is gained or lost within the time increment.

The focal area surveillance (FAS) assets within the scenario comprise an AEW&C aircraft and the Airamak ground based air defence radar. We model these in combination by assuming that AEW&C target information can be relayed to the GBAD radar, effectively extending its range. We also assume that all targets within range are successfully detected, tracked and identified. We assume that the focal area surveillance assets are cued by the OTHR. The issue of interest for the scenario modelled is the association of tracks arising from the OTHR with focal area surveillance tracks, for this enables information about the origin of targets to be combined with information about their identity, enabling their intentions to be deduced. Under these assumptions it is fairly easy to assign probabilities to associations between the relatively uncertain and incomplete set of OTHR tracks, and the more accurate and complete set of FAS tracks and consequently fuse the information associated with each. For those targets beyond range of the FAS assets, their numbers, tracks and identities will require estimation.

We estimate the association between OTHR and FAS tracks as follows. Suppose target track data is represented as $\boldsymbol{t}$ and incorporates estimates of location and velocity in two dimensions. Let two targets have absolute tracks A and B denoted by $t_A^{(f)}$ and $t_B^{(f)}$ according to the FAS, and estimated tracks C and D denoted by $t_C^{(w)}$ and $t_D^{(w)}$ according to the WAS (wide area surveillance, ie. OTHR) with associated distributions $P_C(t^{(w)})$ and $P_D(t^{(w)})$ then the problem is to associate the estimated tracks from the WAS to the absolute tracks from the FAS and to determine probabilities for those associations. If the WAS does not succeed in tracking some or all of the targets then uniform distributions can be assumed for their locations and speeds since they are known to exist. Then the probability of associating tracks A with C and B with D are:

$$\frac{P_C(t_A^{(f)})P_D(t_B^{(f)})}{P_C(t_A^{(f)})P_D(t_B^{(f)}) + P_C(t_B^{(f)})P_D(t_A^{(f)})}$$

which may be used to assign probabilities to deductions regarding the origination and identity of targets.

Maritime target data arising from the satellite-borne sensor in its search mode is used to update the joint p.d.f. for the number of targets and their locations. This attempts to minimise the entropy for the target number determining an optimal search strategy. When this entropy is below

a critical threshold the target number is assumed known and locations estimated conditioned upon this number. In its locate mode of operation the satellite seeks to minimise the conditional entropy for the target locations conditioned upon the assumed number. The strategy adopted is to look only where targets are believed to be rather than where there are none. When this conditional entropy is below a threshold then point estimates for target locations can be extracted. Errors associated with these point estimates are due to target motion between satellite revisits rather than sensor resolution error. Tracks can be formed from associations between point estimates.

Situational awareness at Level 3 estimates the entire surveillance picture formed from the compilation of all detection, location, track and identity data for airborne and maritime targets. It also filters information by focussing attention on targets originating from Orangeland and potentially on course for Airamak. The existence of targets whose data is consistent with this behaviour constitutes a hypothesis whose probability is determined.

Threat assessment at Level 4 estimates the threat state (ie. Blue force's perception of Orange's intentions) as a probability over the three states chosen for the scenario. The random variables whose probabilities are updated as a result of Bayesian estimation may be organised into a Bayesian network as shown in Figure 1. The labels on the links, L1, L2, L3 and L4, refer to the levels in the architectural hierarchy at which the processing is effected. State probability distributions are also updated due to the unpredictable motion of targets using appropriate target motion models. Threat states also potentially evolve over time due to changes in Orange force intentions. This can be represented as a discrete-time Markov chain with specified transition probabilities.

The probabilities for the threat states depends upon the arrival of a HUMINT report about aircraft having taken off from an Orangland airbase which has the effect of increasing belief that Orange forces will attempt to penetrate Airamak defences for purposes unknown. A likelihood function is required which expresses the probability of such a report being generated while in each of the threat states. The specification of such probabilities is subjective but provide a way of representing the utilisation of surveillance information within the context of an operation in a manner consistent with the formalism applied to the rest of the model. Then

$$\Pr\{T \mid \text{report received}\} = \frac{L(\text{report received} \mid T)\Pr\{T\}}{\Pr\{\text{report received}\}}$$

The threat state probabilities also depend upon the situational awareness information from Level 3. The likelihood function required in this case is

$$L(n \geq 1 \,\&\, ID = \text{fighter} \,\&\, O = \text{Orangeland} \,\&\, D = \text{Airamak} \mid T) =$$

$$\Pr\{\text{at least one fighter originating from Orangeland and destined towards Airamak} \mid \text{threat state is } T\}$$

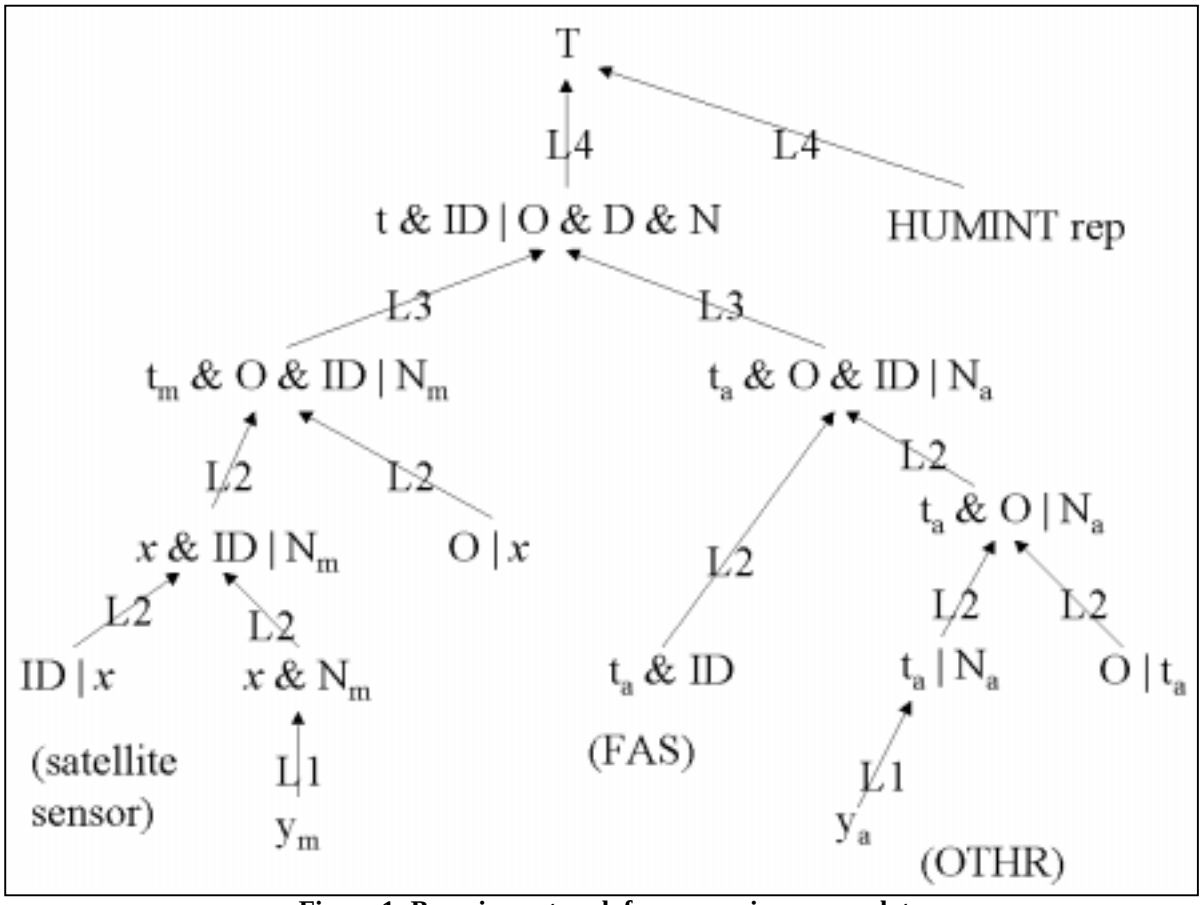

Figure 1: Bayesian network for processing sensor data

3.3 Decision model

The decision model for each level in the goal hierarchy determines responses to information. As mentioned previously the responses can be of three types: information processing, sensor tasking and force response. There need to be criteria for determining when responses are required.

In the case of information processing the decision taken would be to enhance information relating to a particular hypothesis, for example using a target motion model appropriate to the belief that a target is of a particular type, rather than using a generic target motion model. The reason why a decision needs to be taken at all is because of limited computing resources. With unlimited resources all possible hypotheses can be explored concurrently. We choose to represent this type of criterion in terms of the information entropy of the hypothesis concerned, where entropy is a measure of the degree of confidence in the information regarding the truth of the competing hypotheses. If, for the sake of simplicity, we assume that there are unlimited computing resources available in the system being modelled, then the criteria for believing particular hypotheses are arbitrary as all can be believed concurrently.

When a decision is taken to change the tasking of sensors and platforms then the objective must be to enhance the information collected. This information must be relevant to the goal and its

level in the goal hierarchy. This requires a measure of the utility of information. A decision point occurs because the utility of information obtained by changing the tasking is greater than that of not changing it, and the decision taken should be that which maximises the expectation of the utility of the information sought. Again we choose entropy as a measure of the utility of information for deciding when and how to change the tasking of sensors.

A decision at Level 4 must contribute to the achievement of the overall operational objective. Since the primary purpose of the model which is the subject of this paper is to model surveillance within an operational context, then the modelling of decision-making at this level needs only to be sufficient to reflect the variability of the outcome arising from the effectiveness of the surveillance system. Decision modelling at this level is constrained by operational doctrine and effectively just provides a 'test harness' for stimulating and measuring the response of the surveillance system. It would be possible to relate expected outcomes (or their measures) to current threat states and possible actions in a probabilistic sense through a utility function, and then select an action which optimises the expected outcome over a given time horizon. This would be a more sophisticated approach to the modelling of Level 4 decisions but would be consistent with the formalism adopted for Levels 1, 2 and 3 using information utility functions based upon entropy.

4. **Model architecture**

The choice of software system for implementing the model was influenced strongly by the following requirements:

- software modules need to be readily reusable and reconfigurable for different scenarios,
- rapid prototyping capability with visual interfaces for ease of development,
- ability to incorporate legacy sensor models,
- the need to link a continuous-time and space physical simulation with a discrete-event information-driven simulation,
- ability to specify system concurrency desirable due to concurrency being inherent to goal-based approach

4.1 Physical model

As previously noted, we have implemented the test scenario with surveillance operations by using the suite of MATLAB software available to model the entities and their movement throughout the scenario. We simulated the scenario by interfacing the physical models with each other via SIMULINK, a simulation package which is also part of the MATLAB environment. The information processing and information based decisions, (i.e. the information model) is implemented in STATEFLOW, another package available with MATLAB.

The physical entities are defined by their positions within the simulation at each time-step (latitudes and longitudes) and their velocity. The propagation of the entities is calculated accurately and efficiently by making use of the navigational functions available in the Mapping (MATLAB) Toolbox. We have also modelled the propagation of the satellites in the scenario accurately by numerically solving Kepler's equation – the fundamental equation in orbital and

celestial mechanics [Vinti, 1998] and modelled the access which a sensor may have to a target by taking into account the 3-dimensional geometry of the system. We plan to improve the satellite modelling in the future by interfacing the MATLAB architecture with the commercial-off-the-shelf (COTS) product Satellite Tool-Kit. The over the horizon radar (OTHR) has been modelled by utilising the database results from a FORTRAN model of OTHR previously developed at DSTO and interfacing this with the MATLAB architecture. In the future we hope to interface the FORTRAN program directly to the MATLAB environment. Constraints on resources and time has precluded any attempts to do this until recently. Interfacing the modelling environment with existing models has always been of great importance to our choice of simulation architecture and MATLAB satisfies this requirement. It has allowed us to quickly model these entities and effects and interface them with the SIMULINK toolbox to produce a simulation of the scenario. SIMULINK is a package designed to assist modellers in interfacing software models to produce discrete or continuous temporal based simulations. This has in turn enabled us to concentrate on the more important aspects of the study, namely the algorithms and framework with which to handle the information and decisions important throughout the simulation (i.e. the discrete event simulation). Figure 2 shows the top level of the simulation program. Each model is defined as a separate system and input and output data flow between models via data links defined graphically in the program as a type of signal connection. In this way multiple copies of models can be instantiated in the simulation. Data (information) and events are passed from the simulation to the STATEFLOW model of information handling by the same connections. This is all observable graphically as an animation of the entities' movements and any detections which may be generated by sensors.

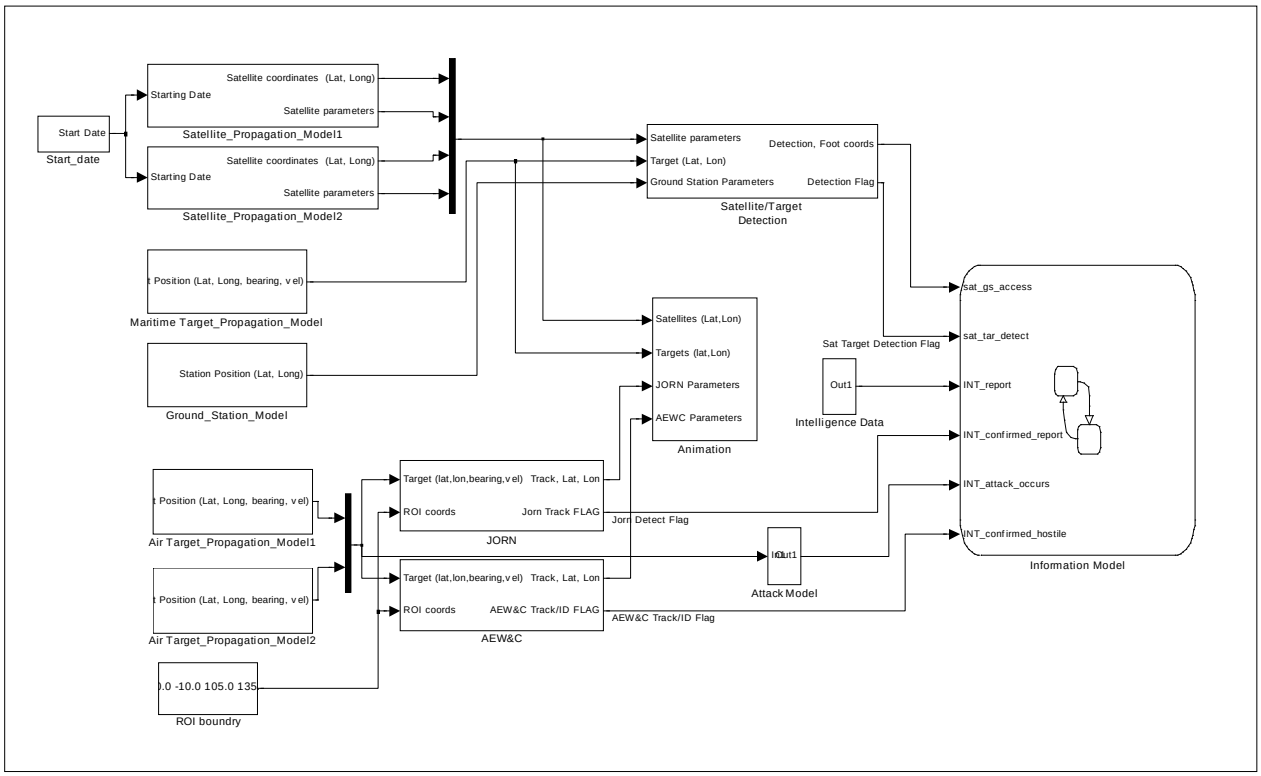

Figure 2: Physical simulation model specified in SIMULINK

We have also developed a stochastic attrition based model of the attack which could take place once enemy fighters arrive at Airamak. This model has been described [Hall and Berry, 2000] and shown to be effective in producing the probability distributions of entities (fuel tanks, aircraft, GBAD, runway) surviving an attack. These results can be used as measures of effectiveness of some of the decision aid algorithms in the information model. We present the attrition model in Figure 3.

Each block represents the events which occur during the preparation phase and subsequent layers have higher fidelity models of the events. The blocks are modelled by the following logic. The Force Rate of Effort F_{ROE} is interpreted as the rate of attack for a single enemy aircraft in units of attack per unit time (or rate of firing a surface to air missile per unit time, G_{ROE}, in the case of GBAD). Then the probability of an attack occurring in $(t, t+\Delta t)$ is given by

$$P_{attack} = N_{aircraft}(t) F_{ROE} \Delta t$$

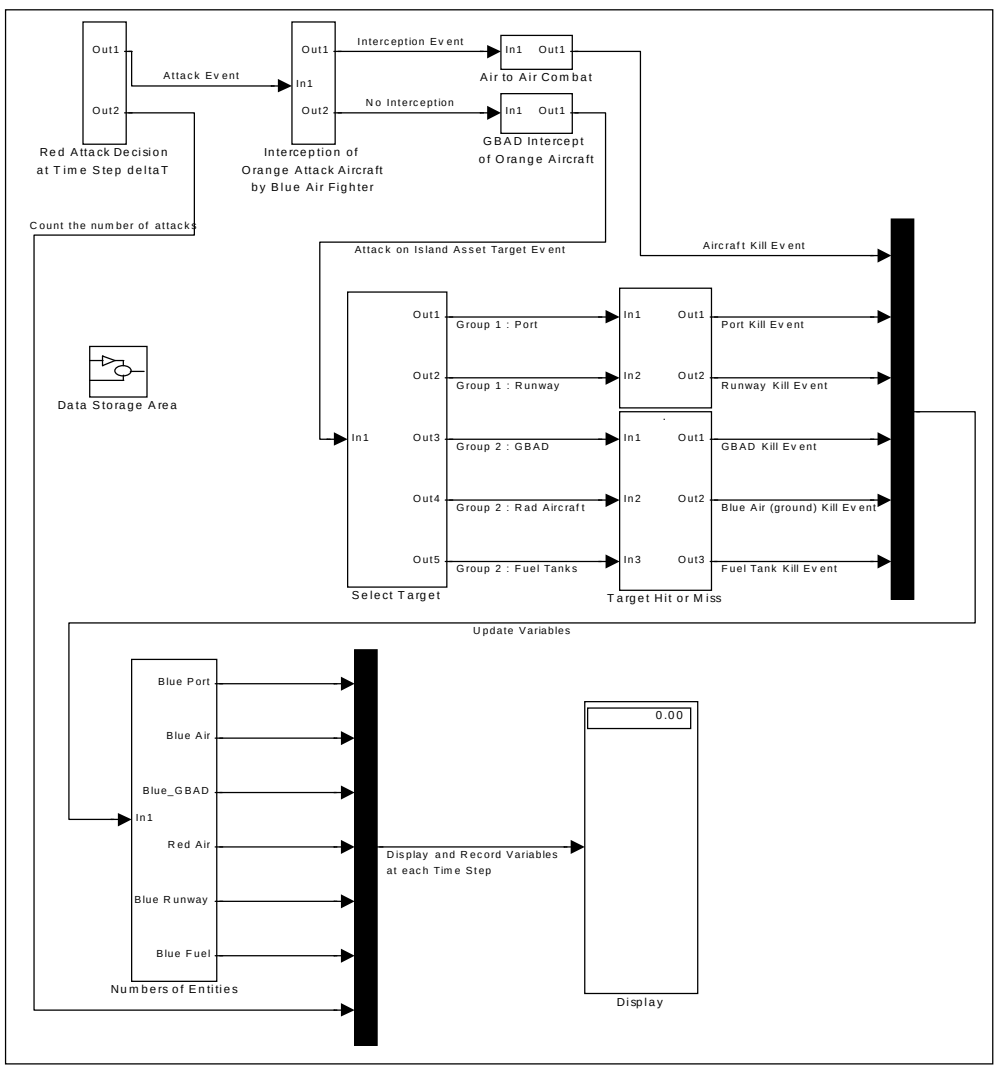

Figure 3: Attrition model used in SIMULINK

During a time step Δt if an Orange aircraft is attacking, the Blue fighter force may intercept an attack. Intercepts can only take place once the fighters are operational. We can assume that if an intercept is successful (p_{int}^{blue}) then the probability of the Orange aircraft being shot down is p_{kill}^{orange}. If an intercept results in a successful shooting down then $N_{aircraft}^{orange}(t)$, the number of Orange aircraft, is reduced by 1 at the end of the time step. There is a probability (p_{kill}^{blue}) that the Orange aircraft destroys the Blue fighter during air-to-air combat. The number of Blue aircraft will be reduced by 1 at the end of the time step if this is case. If an intercept is not successful the Orange aircraft will continue on its mission to attack a ground target k with kill probability $P_k(i)$. The probabilities used in the model are presented in Table 1. Note that in the top level of the simulation only 2 Orange aircraft have been defined. Each of these represent a squadron of 5 enemy aircraft as indicated in Table 1.

Table 1.

Event	Probability
p_{kill}^{blue}	0.5
p_{kill}^{orange}	0.4
$P_k(i)$ (GBAD)	Prob of destroying 1 installation = 0.5
$P_k(i)$ (Blue aircraft on ground)	Prob of destroying 1 aircraft = 0.5
$P_k(i)$ (fuel tank)	Prob of destroying 1 tank = 0.5
$P_k(i)$ (runway)	Prob of destroying 50% of the runway = 0.1
$P_k(i)$ (port)	Prob of destroying 20% of the port area = 0.15
p_{int}^{blue}	$0.56\, N_{aircraft}^{blue}(t) / N_{aircraft}^{blue}(t=0)$
Variable	**Value**
F_{ROE}	1.75 sorties per day per 10 aircraft (0.007 attacks per hour per craft)
G_{ROE}	17 missiles fired per day per 6 installations (0.12 missiles fired per hour per GBAD)
$N_{aircraft}^{blue}(0)$	12
$N_{aircraft}^{orange}(0)$	10 (2 squadrons of 5 aircraft)
$N_{GBAD}^{blue}(0)$	6
$N_{fuel}^{blue}(0)$	3
$N_{runway}^{blue}(0)$	1
$N_{port}^{blue}(0)$	1

4.2 Surveillance information model

The goal hierarchy of section 3.1 corresponds to an agent-based software architecture which has previously been implemented in software systems such as dMARS [d'Inverno *et al*, 1998]. We have chosen to implement it in MATLAB's STATEFLOW as this has been designed to interface closely with SIMULINK which was used for the physical model. STATEFLOW is a graphical design and development tool based on MATLAB, and is specifically designed to visually model and simulate complex discrete-event systems based on finite state machine theory. The modelling capabilities of STATEFLOW allow one to:

- graphically model systems using both State transition diagram notation and Flow diagram notation all on the same STATEFLOW chart
- create charts within charts to provide hierachical organisation for the diagrams.
- graphically define parallelism, junctions and history
- graphically represent common code structures such as "for loops" and "if-then-else" statements via flow diagrams
- utilise any existing functionality from MATLAB, SIMULINK or other toolboxes
- readily have multiple instantiations of the one STATEFLOW chart for use in the same scenario
- incorporate direct event-broadcasting capability
- incorporate temporal conditions for event scheduling purposes

The features of STATEFLOW during execution allow one to:

- integrate event information-based simulation with SIMULINK's discrete-time, continuous-time or hybrid environment
- visualise the state transitions of the model during run-time

The debugging features of STATEFLOW provides users with:

- syntax checking for the model prior to execution
- run-time checking for transition conflicts, cyclic problems, state consistency, and data range violations
- run-time debugging that supports break points, model stepping, data browsing and analysis of code coverages

5. ***Conclusions***

A formalism for representing information uncertainty in models of surveillance operations has been presented, and a corresponding model architecture proposed and implemented for simulating them which takes account of the need to represent both the continuous-time and -space physical domain, and the discrete-event informational domain. Furthermore the model architecture can be readily reconfigured for different operational situations and surveillance architectures, and can easily incorporate existing models of sensors, environment, signal processing, tracking and data fusion algorithms if required in order to enhance its fidelity. Both the formalism and model architecture are scalable in the sense of being capable of being extended to represent and implement increasingly complex systems and operations.

6. ***References***

[Jenson, 1996] Finn V. Jenson. *An Introduction to Bayesian Networks.* University College London Press, London, 1996.

[Marlow *et. al*., 2000] David Marlow. *Quantitative analysis of large-scale scenarios in support of the AIR6000 capability definition process.* Proceedings of the Defence Operations Analysis Symposium, Melbourne, 2000.

[Vinti, 1998] John P. Vinti. *Orbital and Celestial Mechanics.* Published by the American Institute of Aeronautics and Astronautics, Inc., Virginia, 1998.

[Hall and Berry, 2000] Damian L. Hall and Paul E. Berry. *Merits of Monte-Carlo Simulation for the AIR 6000 Test Scenario.* Divisional Discussion Paper, Surveillance Systems Division, DSTO, Salisbury, 2000.

[d'Inverno *et al*, 1998] d'Inverno M, Kinny D, Luck M, Wooldridge M. *A formal specification of dMARS*. Intelligent Agents IV. Agent Theories, Architectures, and Languages. 4th International Workshop, ATAL'97 Proceedings. Springer-Verlag, Berlin. 1998.