

On the Use of Entropy for Optimal Radar Resource Management and Control

Paul E. Berry and David A. B. Fogg

Defence Science and Technology Organisation, PO Box 1500, Edinburgh, SA 5111 Australia
paul.berry@dsto.defence.gov.au

Abstract—The use of information entropy as a quantitative measure of uncertainty for radar resource management and control objectives is explored and applied to issues of current interest. It is seen to be rigorous and objective, and therefore potentially superior to heuristic, rule-based approaches for problems which can be formulated in probabilistic terms. It is particularly appropriate for sensor systems in general which have as their objective the acquisition of information, but which are dominated by uncertainty and subject to time and resource constraints. Examples of the application of this control and management methodology are given to radar problems on widely differing time-scales: the scheduling of track updates in a beam-agile radar (ms), and the tasking of a constellation of SAR surveillance satellites for maritime search and tracking (hours).

Index Terms—Radar, Radar Detection, Adaptive Control, Resource Management, Tracking, Synthetic Aperture Radar.

I. INTRODUCTION

IN modern radars there is considerable scope for selecting, and hence optimizing, control parameters and scheduling control actions for the purpose of enhancing overall performance over a range of time-scales. Any attempt to optimize control should be performed adaptively, that is in response to changed external circumstances, and within the context of an overall objective which, in the case of radars, is the acquisition of information. However, information derived from radar detection, estimation and data processing algorithms is fundamentally uncertain due to noise and clutter affecting detectability and generating false alarms, unpredictable target motion between detections, and ambiguities in the associations between detections and established tracks. If the degree of uncertainty in information can be quantified then constrained radar resources can be managed so as to minimise it over time. Some previous approaches [7] have formulated the radar control problem in terms of a cost function which represents radar time with constraints on the errors (i.e. uncertainty) in the information; the present authors take the view that as available radar time is fixed it should more properly be represented as a constraint and the objective should be to minimise the uncertainty.

For estimation, prediction and control problems which can be formulated probabilistically, that is in terms of a time-varying probability distribution over a state space, a suitable

measure of uncertainty is information entropy. In fact it can be shown [1] that information entropy is a *unique* measure of uncertainty (except for an arbitrary multiplicative constant) which satisfies a set of postulated axioms which such a measure should be reasonably expected to satisfy, in particular that it increases with increasing uncertainty and is a maximum when all states are equally probable. The use of entropy in sensing in the context of robotic navigation is discussed in [11].

Entropy can be used to measure the uncertainty in the current state of a system but, regarded as a random variable, can be more usefully used to predict the uncertainty in a future state of a system. Control parameters may then be selected or optimised and actions scheduled so to minimise the expected future uncertainty. This enables a methodology to be established in which (i) the estimate for the current state of the system is updated as measurements are made (which in the case of a radar generally means the numbers, locations and speeds of targets), (ii) a future state of the system is predicted (statistically) for a particular choice of control parameters from an allowable set (which may comprise a mix of discrete- and continuous-valued interdependent variables), (iii) the expected (i.e. average) uncertainty of the future state is quantified (using entropy), and (iv) the control parameter values selected which minimise the expected uncertainty (subject to time and resource constraints).

Probabilistically, this is the best that can be achieved, assuming that relevant probabilities are known for the performance of the sensors and the evolution of the system between observations. Two applications of the methodology are presented in this paper: the optimal scheduling of track updates in a multifunction phased array radar (PAR), and the optimal tasking of a constellation of synthetic aperture radar (SAR) satellites for tracking maritime targets.

II. THEORETICAL PRELIMINARIES

A random variable X (scalar or vector) is defined on a space Ω whose states may be discrete or continuous, bounded or unbounded. A probability density function $p(X)$ is defined over Ω such that $\int_{\Omega} p(X)dX = 1$ and with the

Report Documentation Page

Form Approved
OMB No. 0704-0188

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE 14 APR 2005	2. REPORT TYPE N/A	3. DATES COVERED -	
4. TITLE AND SUBTITLE On the Use of Entropy for Optimal Radar Resource Management and Control		5a. CONTRACT NUMBER	
		5b. GRANT NUMBER	
		5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)		5d. PROJECT NUMBER	
		5e. TASK NUMBER	
		5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Defence Science and Technology Organisation, PO Box 1500, Edinburgh, SA 5111 Australia		8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSOR/MONITOR'S ACRONYM(S)	
		11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release, distribution unlimited			
13. SUPPLEMENTARY NOTES See also ADM001798, Proceedings of the International Conference on Radar (RADAR 2003) Held in Adelaide, Australia on 3-5 September 2003.			
14. ABSTRACT			
15. SUBJECT TERMS			
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified	UU
			18. NUMBER OF PAGES 6
			19a. NAME OF RESPONSIBLE PERSON

entropy of X defined as

$$H(X) = - \int_{\Omega} p(X) \log p(X) dX \quad (1)$$

In practice, X will represent the state of a target or set of targets in terms of their locations and velocities and at any time will not be known precisely due to sensor imperfections, target mobility and ambiguous discrimination between targets. It is therefore appropriate to represent X , or an estimate of it, in terms of its probability density function, $p(X)$, which will evolve over time as targets move between observations and increase uncertainty in their locations, and can be updated by measurements from sensors as they arise using Bayes' Rule:

$$p(X | Y) = \frac{L(Y | X)p(X)}{\Pr\{Y\}} \quad (2)$$

where $L(Y | X) = \Pr\{Y | X\}$ is the likelihood function for the sensor or set of sensors, $p(X)$ is the *a priori* distribution, $p(X | Y)$ is the *a posteriori* distribution following the observation Y , and the probability of observing Y is $\Pr\{Y\} = \int_{\Omega} L(Y | X)p(X)dX$. In practice, Bayesian estimation may be implemented efficiently as a Kalman filter if the appropriate conditions apply (linearity and Gaussian-distributed errors). Note that discrete state space versions of these expressions are possible and will be used.

A particularly useful result is that if the random variable $X = (x, y)^T$ defined over $\Omega = \Omega_x \times \Omega_y$, with $x \in \Omega_x$ and $y \in \Omega_y$ independent random variables, then $p(x, y) = p(x)p(y)$ which leads to $H(x, y) = H(x) + H(y)$. By induction, if $X = (x_1, x_2, \dots, x_n)^T$ with x_1, \dots, x_n mutually independent, then

$$H(x_1, x_2, \dots, x_n) = \sum_{i=1}^n H(x_i) \quad (3)$$

III. MULTIFUNCTION PAR CONTROL APPLICATION

The objective is to continuously track a number of independent airborne targets using a single multifunction phased-array radar by observing them at intermittent times so as to determine their locations and update their tracks. The update rate for each target should ideally be as frequent as possible in order to maximise the probability of the target being within the beam, and the dwell time (time on target) as long as possible in order to maximise the SNR, and hence probability of detection, whilst keeping the false alarm rate low. However, radar time generally needs to be shared between a number target tracking jobs, as well as the surveillance and weapons guidance functions. If the update rate or dwell time are too low then a target may not be detected resulting in additional looks being scheduled with high priority in order to revisit it. The decision-making

required to schedule the track updates therefore needs to make complex trade-offs over time to ensure that the radar's resources are used efficiently, and that as the radar becomes overloaded, its performance degrades gracefully rather than catastrophically ([7]-[10]).

Suppose that there are N targets to be tracked and each target has a dynamic equation of the form

$$\mathbf{x}(k+1) = \mathbf{F}(k)\mathbf{x}(k) + \mathbf{v}(k) \quad (4a)$$

and a measurement equation

$$\mathbf{z}(k) = \mathbf{H}(k)\mathbf{x}(k) + \mathbf{w}(k) \quad (4b)$$

where $\mathbf{v}(k)$ and $\mathbf{w}(k)$ are sequences of zero-mean, white Gaussian noise processes, as normally specified for Kalman filter trackers [4]. Then $\mathbf{x}(k)$ for times t_k , $k = 0, 1, \dots$, is a multivariate Gaussian distribution which may be estimated in terms of its mean $\hat{\mathbf{x}}(k)$ and covariance matrix $\mathbf{P}(k)$. The estimation is performed in two stages in a Kalman filter: the predicted state $\hat{\mathbf{x}}(k+1|k)$ and $\mathbf{P}(k+1|k)$ based upon the dynamic equation between times t_k and t_{k+1} , and the updated state $\hat{\mathbf{x}}(k+1|k+1) \rightarrow \hat{\mathbf{x}}(k+1)$ and $\mathbf{P}(k+1|k+1) \rightarrow \mathbf{P}(k+1)$ based upon a set of measurements $\mathbf{z}(k+1)$ at time t_{k+1} .

The state vector $\hat{\mathbf{x}}$ will generally represent the target location and velocity in three dimensions, and \mathbf{P} the corresponding covariances, but for the purpose of beam scheduling to maintain tracks' we are only interested in the elements of the covariance matrix representing the error in target azimuth and elevation. Let this be the matrix $\mathbf{Q}_i(t)$, for the i^{th} target at time t , then the entropy representing the uncertainty in its position is (Theorem 9.4.1, [2])

$$h_i(t) = \frac{1}{2} \log \{4\pi^2 e^2 | \mathbf{Q}_i(t) |\} \quad (5)$$

Using the result of Section II, the entropy associated with the joint system of N independent targets at time t is

$$H(t) = \sum_{i=1}^N h_i(t) = \frac{1}{2} \sum_{i=1}^N \log \{4\pi^2 e^2 | \mathbf{Q}_i(t) |\} \quad (6)$$

This expression provides the means for objectively quantifying the overall uncertainty associated with a set of independent targets and balancing the resources allocated to them by scheduling their individual updates so as to minimise it. Alternatively, the optimal control problem could be formulated so as to specify an acceptable level of uncertainty as a constraint and minimise the resources necessary to maintain that level. This might be an appropriate formulation for a set of high priority targets which must be tracked at all costs, with remaining radar resources applied to low priority targets and other functions.

This formulation is fairly general and could be applied to complex situations. However, in the following it is applied to a simplified and analytically tractable situation so as to demonstrate that it yields intuitively reasonable results. For the special case when the target azimuth and elevation

coordinates are independent, the Kalman filter equations decouple and each coordinate can be predicted and updated separately. If, for the i^{th} target, x_i and y_i are the azimuthal and elevation coordinates and $\sigma_{i,x}$ and $\sigma_{i,y}$ are their variances then $|\mathbf{Q}_i| = \sigma_{i,x}\sigma_{i,y}$ and

$$\log |\mathbf{Q}_i| = \log \sigma_{i,x} + \log \sigma_{i,y} \quad (7)$$

Hence at time t_n

$$H(t_n) = \log 2\pi e + \frac{1}{2} \sum_{i=1}^N \log \sigma_{i,x}^n + \log \sigma_{i,y}^n \quad (8)$$

The decoupled Kalman filter for the dynamic equation applied to the i^{th} target enables the variance to be predicted at time τ later from the current covariances, thus:

$$\sigma_{i,x}^2(k+1|k) = \sigma_{i,x}^2(k|k) + 2b_{i,x}\tau + d_{i,x}\tau^2$$

$$\sigma_{i,y}^2(k+1|k) = \sigma_{i,y}^2(k|k) + 2b_{i,y}\tau + d_{i,y}\tau^2$$

where $b_{i,x}$ and $b_{i,y}$ are the position/velocity covariances for the x_i and y_i variables for the i^{th} target, and $d_{i,x}$ and $d_{i,y}$ are the velocity variances.

Suppose that at this time it is proposed to select a target and point the beam towards its predicted position, then the consequence is that either the target is detected or it is not. In each case we can compute what the entropy will be and from our knowledge of the detection probability compute the expected entropy. For targets which are not selected the expected entropy will be the same as the predicted entropy. We can then decide which of the N targets should be observed so as to minimise the overall entropy.

If the target is detected then the decoupled Kalman filter measurement equations will update the covariances thus:

$$\frac{1}{\sigma_{i,x}^2(k+1|k+1)} = \frac{1}{\sigma_{i,x}^2(k+1|k)} + \frac{1}{\bar{\sigma}_{i,x}^2}$$

where $\bar{\sigma}_{i,x}$ is the variance of the measurement error in the x_i coordinate. Therefore the entropy associated with the i^{th} target after detection, will be

$$\frac{1}{2} \log \sigma_{i,x}(k+1|k+1) + \frac{1}{2} \log \sigma_{i,y}(k+1|k+1)$$

The probability of this occurring is

$$\Pr\{\text{target detected}\} = \Pr\{\text{target detected AND target in beam}\}$$

$$= \Pr\{\text{target detected} | \text{target in beam}\} \Pr\{\text{target in beam}\}$$

$$= P_D P_B$$

where $P_D = (P_{FA})^{1/(1+SNR)}$ for a Swerling type I fluctuating target and P_B is estimated by integrating the target probability distribution over the 3dB pencil beam width.

Note that the Kalman filter does not take account of non-detections. If a target is assumed to exist, a region observed and the target not detected then this is information which could be used by a Bayes filter to update the target's locational probability distribution. However, the updated distribution would no longer be Gaussian so a Kalman filter

would not be applicable. Instead, in conventional tracking, the estimate is treated as though no observation occurred. If no detection occurs then the entropy stays as

$$\frac{1}{2} \log \sigma_{i,x}(k+1|k) + \frac{1}{2} \log \sigma_{i,y}(k+1|k)$$

with probability $1 - P_D P_B$. So the expected entropy resulting from observing the i^{th} target is

$$\frac{1}{2} P_D P_B \log \sigma_{i,x}(k+1|k+1) + \frac{1}{2} (1 - P_D P_B) \log \sigma_{i,x}(k+1|k) \\ + \frac{1}{2} P_D P_B \log \sigma_{i,y}(k+1|k+1) + \frac{1}{2} (1 - P_D P_B) \log \sigma_{i,y}(k+1|k)$$

$$= \frac{1}{4} P_D P_B \log \left(\frac{\bar{\sigma}_{i,x}^2}{\sigma_{i,x}^2(k+1|k) + \bar{\sigma}_{i,x}^2} \right) + \frac{1}{2} \log \sigma_{i,x}$$

$$+ \frac{1}{4} P_D P_B \log \left(\frac{\bar{\sigma}_{i,y}^2}{\sigma_{i,y}^2(k+1|k) + \bar{\sigma}_{i,y}^2} \right) + \frac{1}{2} \log \sigma_{i,y}$$

and the total expected entropy *change* is simply

$$= \frac{1}{4} P_D P_B \log \left(\frac{\bar{\sigma}_{i,x}^2}{\sigma_{i,x}^2(k+1|k) + \bar{\sigma}_{i,x}^2} \right)$$

$$+ \frac{1}{4} P_D P_B \log \left(\frac{\bar{\sigma}_{i,y}^2}{\sigma_{i,y}^2(k+1|k) + \bar{\sigma}_{i,y}^2} \right)$$

which should be made as negative as possible.

Selecting the target which creates the biggest change in total entropy is then simply a matter of finding for which value of i the following expression is the least:

$$\left(\frac{\bar{\sigma}_{i,x}^2}{\sigma_{i,x}^2(k+1|k) + \bar{\sigma}_{i,x}^2} \right)^{P_D P_B} \left(\frac{\bar{\sigma}_{i,y}^2}{\sigma_{i,y}^2(k+1|k) + \bar{\sigma}_{i,y}^2} \right)^{P_D P_B}$$

This will tend to select those targets which have the largest predicted error covariances, the smallest measurement errors and the largest probabilities of being detected which conforms with intuition.

The function to be optimised in optimisation theory is usually referred to as an objective function but in operations research and engineering it is often referred to as a utility function which is to be maximised or a cost function which is to be minimised. The problem of optimal scheduling of track updates has been posed in both forms. In [8] and [9] it is formulated as a nonlinear optimal control problem with the objective of minimising the radar energy needed to update N tracks, and in [10] a cost function is minimised which is the expectation of the sum of the discounted costs for each target over an infinite time horizon. Blackman in [3] postulates the existence of a utility function $U(x, y)$ for two target attributes x and y (such as errors in range and angle) which is to be maximised and assumes that their independence leads to a linear sum $U(x, y) = aU_1(x) + bU_2(y)$ (a *separable* objective function) but states that the choice of a utility function is a subjective matter. This paper has demonstrated that a utility function can be defined objectively and uniquely, and that separability arises naturally from (3). Separability is a desirable attribute of a mathematical programming problem.

Having defined an objective function which quantifies uncertainty, it is necessary to decide for which future point in time, or over what time period, its expectation should be minimised, that is the choice of time horizon. The objective function proposed in this paper in terms of entropy would be used to determine the sequence of optimal track updates which reduces the overall uncertainty at some future point in time. The length of time should be sufficiently long that each possible track update is capable of being considered, but not so long that the prediction is unrealistic. Following the next track update, the sequence will need to be revised if there is a failure to detect the target. This will occur from time to time for each target.

The complexity of an agile beam radar is such that the preceding analysis has had necessarily to be simplified so as to render it analytically tractable for the purpose of demonstrating the control methodology, but it could be extended without difficulty to other adaptive radar control issues. For instance it could be used in conjunction with sophisticated trackers such as interactive multiple model (IMM) trackers and represent optimal searching to reacquire targets, and it could incorporate the target dwell time as a control parameter, as well as the choice of waveform and detection thresholds in determining the optimal schedule. However, the precise formulation of the optimal scheduling problem would be highly dependent on the particular system characteristics.

IV. SAR SATELLITE CONSTELLATION TASKING APPLICATION

A fixed constellation of satellites with on-board SAR sensors is used to search for, estimate the location of, and track surface maritime targets. When within range of a ground station, they transmit sensed data for processing and are tasked for future observations. It is possible to control the look-angle of the SAR and its resolution but the appropriate choice of controls for each fly-pass depends upon the information sought. For example, the strategy adopted for an area or barrier search will differ from that for locating a previously detected target or tracking a number of targets. Since space-based surveillance resources are limited it is important that they be applied effectively. If there is sufficient time to process data and task sensors between fly-passes, then it is only necessary to plan ahead to the next fly-pass of a satellite which occurs at a predetermined time.

Bayesian estimation (see [5]) is used in preference to Kalman filtering for wide area surveillance applications because it can deal with search as well as tracking, and the likelihood function $L(\cdot|\cdot)$ (see (2)) can represent non-linearities in the measurement model such as discrete ambiguities in location (e.g. due to multi-path effects), and target/environmental constraints (e.g. embodies knowledge of land masses, and targets masked by cloud cover affecting electro-optic sensors).

The region \mathfrak{R} of interest is discretised into numbered cells, each cell of which has sufficiently small area that the

probability of it being occupied by two or more independent targets is negligible. For simplicity the cells may be defined by the imposition of a rectangular grid upon \mathfrak{R} . Without loss of generality, let there be N cells where the probability of a cell being occupied by a target at a time t is $p_i(t)$, $i = 1, \dots, N$. Let the surveillance satellites have access to the region \mathfrak{R} at discrete instants of time $t_1, t_2, \dots, t_n, \dots$, which are known in advance and will not, in general, be equally spaced, then we write $p_i^n \equiv p_i(t_n)$. These times are referred to as epochs and the area accessed by a satellite at the n^{th} epoch is denoted $A_n \subset \mathfrak{R}$. Note that p_i^n may be interpreted to mean either the probability of cell i containing a target at epoch n (as opposed to it not containing a target), or alternatively the probability that a single target in the region \mathfrak{R} occupies cell i (as opposed to any other cell). These represent alternative formulations of the problem and which of them is relevant will depend upon the information sought, as will be discussed. More general formulations are possible for multiple targets.

The SAR sensors' performances are modelled in terms of their probabilities of detection and false alarm. Depending upon the resolution of the sensor and the intervening environment, a sensor's probability of detection for a cell is p_d (if a target exists) and false alarm probability is p_{fa} . Note that false alarm probability per cell has to be calculated from the false alarm distribution P_{fa} per unit area (i.e., a Poisson distribution in 2-D) and the area of the cell i.e., $p_{fa} = P_{fa}\Delta A$ where ΔA is the area of a cell. For simplicity it is assumed that all cells are equal in area and that the false alarm distribution is uniform, hence p_{fa} is spatially constant for a given choice of resolution. The updating of the probability distribution is based upon the observations using Bayes' rule (2) which uses a likelihood function appropriate to the way in which the state space is structured.

A target motion update is computed to predict target motions between epochs n and $n+1$. This is usually based upon a Gauss-Markov target motion model which is equivalent to the dynamic equation (4a) but other models are possible which may be more appropriate to maritime targets (see [6]).

A. Search

The surveillance information objective is to determine the existence of targets within a region \mathfrak{R} or detect targets entering or departing a region. There is no requirement to track them and therefore no need to maintain accurate estimates of their locations or be able to associate new detections with previous detections. A cell-based formulation is therefore chosen in which cells are assumed to be independent and no unique association is made between

measurements and targets. This is adequate for determining the existence of targets within \mathfrak{R} in the search phase of a surveillance operation. The probability p_i^n is therefore to be interpreted as the probability of a target occupying cell i at epoch n . In other words each cell is an entire state space and the probability of cell i not containing a target is $1 - p_i^n$. All cells are independent which means that only cells which are observed (whether or not a detection occurs) are updated using Bayes' rule. This takes the form

$$p_i^{n+1}(1) = \frac{p_d \hat{p}_i^{n+1}}{p_d \hat{p}_i^{n+1} + p_{fa} (1 - \hat{p}_i^{n+1})}$$

$$p_i^{n+1}(0) = \frac{(1 - p_d) \hat{p}_i^{n+1}}{(1 - p_d) \hat{p}_i^{n+1} + (1 - p_{fa})(1 - \hat{p}_i^{n+1})}$$

for a detection and a non-detection event respectively in the observed cell i at the $(n+1)^{\text{th}}$ epoch where \hat{p}_i^{n+1} is the *a priori* probability. By virtue of (3), the independence of the cells implies that the entropy of the joint system, representing uncertainty in the existence of multiple targets within \mathfrak{R} , is simply the sum of the individual cell entropies

$$H^n = \sum_{i=1}^N h_i^n = - \sum_{i=1}^N p_i^n \log p_i^n + (1 - p_i^n) \log(1 - p_i^n)$$

The choice of control action α which maximises the expected change in entropy at the next fly pass of a SAR satellite at epoch $(n+1)$ is

$$\max_{\alpha} \sum_{i \in A(\alpha)} \{\hat{h}_i^{n+1} - \bar{h}_i^{n+1}\}$$

where α determines both the size of the swathe, and hence number of cells in the set $A(\alpha)$ which are observed, as well as the sensor resolution through p_d and p_{fa} . The expected entropy for a single cell i is (see [6])

$$\begin{aligned} \bar{h}_i^{n+1} &= \hat{h}_i^{n+1} - \hat{p}_i^{n+1} [p_d \log p_d + (1 - p_d) \log(1 - p_d)] \\ &\quad - (1 - \hat{p}_i^{n+1}) [p_{fa} \log p_{fa} + (1 - p_{fa}) \log(1 - p_{fa})] \\ &\quad + [p_d \hat{p}_i^{n+1} + p_{fa} (1 - \hat{p}_i^{n+1})] \log [p_d \hat{p}_i^{n+1} + p_{fa} (1 - \hat{p}_i^{n+1})] \\ &\quad + [1 - p_d \hat{p}_i^{n+1} - p_{fa} (1 - \hat{p}_i^{n+1})] \times \\ &\quad \log [1 - p_d \hat{p}_i^{n+1} - p_{fa} (1 - \hat{p}_i^{n+1})] \end{aligned}$$

B. Track

Under the assumption that there exists a single target of interest which can be discriminated from all other targets, all observations can be associated with it for the purpose of tracking. A Bayesian estimator can be derived which simultaneously updates position and velocity, which is a generalisation of the dynamic equation (4a). The likelihood function for the single target tracking problem is more complex than for the search problem.

Suppose the swathe $A \subset \mathfrak{R}$ is swept, comprising the set of M cells which are observed and which result in the set of detections $\{j_1, j_2, \dots, j_k\}$. The likelihood function used in Bayes' rule has different forms depending upon whether cell i corresponds to any of the M cells of swathe A , or not in A (L_{III}), and on whether, if it is in A , it is one of the $\{j_1, j_2, \dots, j_k\}$ (L_{II}), or not (L_I). Note that all of $\{j_1, j_2, \dots, j_k\}$ must be in A (because a sensor can only obtain detections where it observes).

Case I corresponds to the situation in which the target is contained in A but does not belong to the set of detections. So the likelihood function is the probability that the target is not detected but that there are k false alarms in A (and hence $M-k-1$ non-false alarms).

$$\begin{aligned} \text{Case I: } i \in A \text{ but } i \notin \{j_1, j_2, \dots, j_k\} \text{ for } k = 0, \dots, M-1 \\ L_I(j_1, j_2, \dots, j_k \mid i \in A \cap i \notin \{j_1, j_2, \dots, j_k\}) \\ = (1 - p_d) p_{fa}^k (1 - p_{fa})^{M-k-1} \end{aligned}$$

Case II corresponds to the situation in which the target is contained in A but *does* belong to the set of detections. So the likelihood function is the probability that the target *is* detected and that there are $k-1$ false alarms in A (and hence $M-k$ non-false alarms).

$$\begin{aligned} \text{Case II: } i \in A \text{ and } i \in \{j_1, j_2, \dots, j_k\} \text{ for } k = 1, \dots, M \\ L_{II}(j_1, j_2, \dots, j_k \mid i \in A \cap i \in \{j_1, j_2, \dots, j_k\}) \\ = p_d p_{fa}^{k-1} (1 - p_{fa})^{M-k} \end{aligned}$$

Case III corresponds to the situation in which the target is not in A , so all detections must be false alarms and the question of whether the actual target is detected or not does not arise. Hence the likelihood function is the probability of k false alarms and $M-k$ non-false alarms.

Case III: $i \notin A$ for $k = 0, \dots, M$

$$L_{III}(j_1, j_2, \dots, j_k \mid i \notin A) = p_{fa}^k (1 - p_{fa})^{M-k}$$

Note that all of the above assumes that the probabilities of detection p_d and false alarm p_{fa} (per cell) are spatially uniform in A and are known functions of the sensor resolution, which, together with the look angle, determine the size of A . The probabilities can in principle be made spatially and temporally non-uniform so as to represent environmental effects.

A general expression may be derived in the single target case (see [6]) for the expectation of the entropy over all possible sets of observations $\underline{l} = \{j_1, j_2, \dots, j_k\} \in \Lambda$ at the $(n+1)^{\text{th}}$ epoch, where Λ is the set of all possible sets of observations:

$$\begin{aligned} \bar{h}^{n+1} &= \hat{h}^{n+1} \\ &\quad - \sum_i \hat{p}_i^{n+1} \sum_{\underline{l} \in \Lambda} L(\underline{l} \mid i) \log L(\underline{l} \mid i) + \sum_{\underline{l} \in \Lambda} P^{n+1}(\underline{l}) \log P^{n+1}(\underline{l}) \end{aligned}$$

where $L(\underline{l}|i)$ is the sensor's likelihood function defined as the probability of obtaining measurement set $\underline{l} = \{j_1, j_2, \dots, j_k\}$ given that the target is in cell i , \hat{p}^{n+1} is the *a priori* probability that the target is in cell i (following target motion update), and $P^{n+1}(\underline{l}) = \sum_{i \in \mathcal{R}} L(\underline{l}|i) \hat{p}_i^{n+1}$ is the

probability of obtaining the measurement set \underline{l} . The multiple detections arise from false alarms with at most one detection corresponding to the actual target.

The expected entropy expression has the following physical interpretation: it represents the change in entropy from the *a priori* entropy with the second term on the r.h.s. representing an increase in entropy due to sensor imperfections and the third term representing a decrease but dependent on the sensor coverage. Therefore the second and third terms explicitly represent the trade-off between coverage and resolution.

V. IMPLEMENTATION ISSUES

As additional variables are incorporated into stochastic problems the joint state spaces increase combinatorially so methods need to be explored to control the size of the state space such as exploiting independence between variables. The problems should be formulated as generally as possible in the first instance and then rational approximations applied. This enables the magnitude of the approximation error to be gauged. In the case of the single target tracking likelihood function the generally small magnitude of p_{fa} will tend to limit the number of multiple (false) detections so an approximation is possible to the number of states which need to be accommodated.

Bayesian estimation has tended to be avoided for real-time applications in the past because of its computational overheads but recent developments in particle filtering have rendered the approach more feasible. Its ability to represent nonlinearities in the dynamic and measurement models makes it more robust than Kalman filtering for tracking and data fusion applications.

Entropy computations tend also to be computationally burdensome due to the calculation of logarithmic terms over all possible states. The expected entropy is even more demanding because the expectation is taken over all possible measurements. However expressions for the expected entropy have been presented in terms of the expected *change* in entropy following an observation. These expressions are capable of rational approximation within the context of specific problems of interest and reduce to separable problems in certain instances. However, the question of interest is not whether one is able to obtain a truly optimal solution, but whether one is able to obtain the best solution possible within the available time and computational resources. This may require the use of efficient but sub-optimal heuristic search methods such as constraint programming.

A technique which is able to exploit interdependencies between random variables so as to reduce the complexity of the joint state space of all possible relevant variables is Bayesian Belief Networks. This involves the exploration of hypotheses, for example an assumption about the total number of targets in a region, and updates the sub-space probabilities subject to that assumption. The probability of that hypothesis also needs to be monitored, or perhaps the entropy over the probabilities for all competing hypotheses, so as to switch beliefs as necessary.

The problems presented have been stated fairly generally without much discussion of constraints and optimisation techniques because the precise formulation of an optimisation or scheduling problem depends crucially on the nature of the variables (discrete or continuous, bounded or unbounded) and the constraints (linear or nonlinear, equality, or inequality constraints).

VI. CONCLUSION

Entropy has been shown to be a useful and practical way of measuring uncertainty in information derived from radar observations on disparate timescales. It can be used to monitor the quality of information already obtained, to predict the quality of future information resulting from a proposed control action or parameter value, or, most usefully, to select the optimal control action and parameter value which results in the minimal expected uncertainty and hence makes best use of the radar resources. Depending upon the problem of interest, the entropy objective function can be coupled with resource constraints to formulate a constrained optimisation problem over time if necessary.

REFERENCES

- [1] E. A. Bender, "Mathematical Methods in Artificial Intelligence". Los Alamitos, CA: IEEE Computer Society Press, 1996, ch. 12.
- [2] T. M. Cover and J. A. Thomas, "Elements of Information Theory". New York: Wiley, 1991.
- [3] Y. Bar-Shalom, "Multitarget-Multisensor Tracking: Applications and Advances Volume II". Boston: Artech House, 1992, ch.8.
- [4] Y. Bar-Shalom and X-R. Li, "Estimation and Tracking: Principles, Techniques and Software". Boston: Artech House, 1993.
- [5] L. D. Stone, C. A. Barlow and T. L. Corwin, "Bayesian Multiple Target Tracking". Boston: Artech House, 1999.
- [6] P. E. Berry, D. A. B. Fogg and C. Pontecorvo, "Optimal search, location and tracking of surface maritime targets by a constellation of surveillance satellite", DSTO Research Report., to be published.
- [7] S. P. Noyes, "Calculation of the next time for track update in the MESAR phased array radar", *IEE Colloquium on Target Tracking and Data Fusion*, UK, June 1998.
- [8] S-M. Hong and Y-H Jung, "Optimal scheduling of track updates in phased array radars", *IEEE Transactions on Aerospace and Electronic Systems*, vol. 34, pp. 1016-1022, July 1998.
- [9] Y-H Jung and S-M. Hong, "Optimal scheduling of detection and tracking parameters for a phased array radar", pp. 1111-1116, *SICE '99*, July, 1999
- [10] V. Krishnamurthy and R. J. Evans, "Hidden Markov Model multiarm bandits: a methodology for beam scheduling in multitarget tracking", *IEEE Transactions on Signal Processing*, vol. 49, no. 12, pp. 2893-2908, December 2001.
- [11] J. Manyika and H. Durrant-Whyte, *Data Fusion and Sensor Management: A Decentralized Information-Theoretic Approach*. New York: Ellis Horwood, 1993, p. 42.