

Optimal Adaptive Waveform Selection for Target Detection

B. F. La Scala
Dept. of Electrical
and Electronic Engineering
University of Melbourne
Melbourne, Australia
E-mail: bfls@ieee.org

W. Moran
Dept. of Electrical
and Electronic Engineering
University of Melbourne
Melbourne, Australia
E-mail: b.moran@ee.mu.oz.au

R. J. Evans
Dept. of Electrical
and Electronic Engineering
University of Melbourne
Melbourne, Australia
E-mail: r.evans@ee.mu.oz.au

Abstract—Modern phased array radars are able to adaptively modify their performance to the environment. To make full use of this capability, scheduling algorithms need to be designed. This paper poses the problem of adaptive waveform scheduling for detecting new targets in the context of finite horizon stochastic dynamic programming. The result is a scheduling algorithm that minimises the time taken to detect new targets, detecting these targets in accordance with importance, while minimising the use of radar resources.

I. INTRODUCTION

Modern phased array radars, with flexible waveform generation and beam steering capability, are able to adaptively modify their performance to suit a variety of environments. This power has not yet been fully exploited, in part because of the lack of suitable scheduling algorithms. This paper describes an optimal waveform selection algorithm for the detection of new targets. It does not examine the problem of maintaining tracks on established targets. In the rest of this paper we will use the term “target detection” to refer to the identification of new targets, rather than the detection on subsequent scans of targets already under track.

Phased array radars can direct their beam in any direction without inertia. Thus the radar can switch between the tasks of tracking existing targets and acquiring new targets essentially instantaneously. Such a radar thus achieves the multi-mission capability of target acquisition and target tracking, unlike typical mechanically scanned radars. This flexibility also allows the system designer to consider the task of searching for new targets separately from, and independently of, that of updating established tracks, even if the same radar is performing both tasks.

In general, a waveform can be tailored to achieve good Doppler or good range resolution, but not both simultaneously. This is a problem in heavy clutter environments, typified by an airborne radar seeking to detect slow moving ground targets or by a littoral radar attempting to detect submarine periscopes in the presence of sea clutter. In both cases the part of the return

ascribable to the clutter can be orders of magnitude larger than that from the target or targets of interest. Waveforms, to a greater or lesser extent, smear the clutter into the target region, thus reducing detectability. Once a track has been established on a target, the appropriate choice of waveform can be based on the track state estimates. This problem has been examined in works such as [3], [17], [9], [14], [8].

The problem we are considering in this paper is the detection of new targets, so the results in the works cited above are not directly applicable. The efficient search for new targets has been examined in [19] and [18] and an overview of these results, and related work, can be found in [2, Ch 14]. While these works provide guidelines for parameter selection in a variety of cases, they do not pose the problem adaptively. Clutter mapping and optimal scheduling would permit the tailoring of waveforms and beam shapes to best match the working environment of the radar. These techniques, alone or in combination, offer the possibility of adaptive adjustment of the sensor modes to optimise performance. Because of the high data rates, manual optimisation of the performance of a modern radar by pulse tailoring is not possible. There is significant potential for improvement in new target detection if adaptive waveform selection is considered part of the detection process.

The simplest schemes for adaptive waveform management ascribe a cost function to the clutter/target environment for each individual pulse and select the waveform that optimises the cost function on a pulse by pulse basis. While such a “greedy” scheme would radically improve performance over conventional fixed waveform radars, more can be gained by scheduling waveforms over a number of pulses, so as to optimise the sum of the costs over these pulses.

In this paper we pose the adaptive waveform scheduling problem for new target detection as a stochastic dynamic programming problem of the type known as a *partially observed Markov decision problem*. Solutions of this type of problem are optimal control policies that maximise an objective function. In this framework, the adaptive waveform selection problem for target detection becomes the selection of which sequence of waveforms to use to maximise the overall rewards of target detection. These rewards could take into account factors such

This work was supported in part by the Defense Advanced Research Projects Agency of the US Department of Defense and was monitored by the Office of Naval Research under Contract No. N00014-02-1-0802. The United States government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright notation hereon.

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 Optimal Adaptive Waveform Selection for Target Detection				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) Dept. of Electrical and Electronic Engineering University of Melbourne Melbourne, 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 UU	18. NUMBER OF PAGES 5	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

as timeliness of detection or target importance and/or threat. In addition, the location of targets already under track can also be incorporated in the reward structure, as well as a clutter map if known, to prevent radar resources being wasted on known scatterers. In the most general case, the overall reward for target detection could be a combination of all of these types of factors.

The result is an adaptive waveform selection algorithm that minimises the time taken to detect new targets, detecting these targets in accordance with importance, while minimising the use of radar resources.

II. PROBLEM OUTLINE

Methods for waveform selection to improve both detection and tracking performance have been examined in [15], [14] and [9]. In these works, the problem was one of both acquiring and tracking a single target in clutter. In this paper we will only consider the problem of detecting new targets, rather than tracking established targets. Once a target is detected and confirmed, it is handed over to the track update process and is no longer a concern of the target detection process.

We pose the problem of optimal adaptive waveform selection for target detection as a finite horizon stochastic control problem. While we only consider detection performance here, this approach can be extended to consider both detection and tracking. The problem of optimal beam scheduling to maximise tracking performance using this type of approach was examined in [10].

The problem posed in [10] used an infinite horizon as they were concerned with tracking performance. Since we are only considering the detection problem it is appropriate to use a short, finite horizon. Unlike [10], the short time horizon allows us to assume that the scene does not change during this interval, i.e. the targets do not move appreciably. This assumption is reasonable as the number of dwells used to confirm a track on a new target is typically very small, see for example [18] and [6].

In addition, we will consider the detection problem in each radar beam to be independent of other beams. That is, a target detection in one beam does not provide any information on the likelihood of a detection in a neighbouring beam. This implies that the beams are spaced with minimal overlap. While this is not always true, it is a reasonable assumption and provides a useful starting point for developing this adaptive waveform selection method. Therefore, in the remainder of this paper we will consider the detection problem in a particular beam in isolation.

The format of the remainder of this paper is as follows. The next section sets up the problem of adaptive waveform selection for target detection as a stochastic control problem. Section IV shows how the effect of the choice of waveform on the probability of detection is incorporated into the model. Section V discusses a number of choices for the objective function that is to be maximised, while an optimal solution method is outlined in Section VI.

III. STOCHASTIC CONTROL PROBLEM

We divide the area covered by a particular radar beam into a grid in range-Doppler space, with the cells in range indexed by $\tau = 1, \dots, N$ and those in Doppler indexed by $\nu = 1, \dots, M$. We make no assumptions about the number of targets that may be present, thus the number of possible scenes or hypotheses about the radar scene is 2^{NM} . Let the space of hypotheses be denoted by $\mathcal{H}$. In a very simple example, the range space could be divided into two cells (i.e. a target is either near or far) and Doppler into three cells (i.e. a target is either receding, stationary or approaching) then the set of hypotheses, $\mathcal{H}$, has 2^6 elements corresponding to all the possibilities ranging from no targets present to 6 targets with one in each cell. Note, we assume that the resolution of the radar and the size of the cells are such that at most one scatterer can be distinguished in each cell.

The state of our model is then $x(k) = i$ where $i \in \mathcal{H}$, i.e. it is one of the possible scenes, and is fixed over the time interval of interest.

The radar provides noisy measurements $y(k) \in \mathcal{H}$ of the true scene, $x \in \mathcal{H}$. The probability of receiving a particular measurement $y(k) = j$ will depend on both the true, underlying scene and on the choice of waveform used to generate the measurement.

Let $u(k)$ be the control variable that indicates which waveform is chosen at time k to generate measurement $y(k+1)$, where $u(k) \in \mathcal{U}$. Then $\mathbf{B}(u(k)) = (b_{ji}(u(k)))_{i,j \in \mathcal{H}}$ is the measurement probability matrix where

$$b_{ji}(u(k)) = \Pr(y(k+1) = j | x = i, u(k)).$$

In other words, $b_{ji}(u)$ is the probability of a detection in all the cells considered to have a scatterer present under hypothesis j given that the true scene is given by hypothesis i and is observed with waveform u .

Define $\pi = \{u(0), u(1), \dots, u(T)\}$ where $T+1$ is the maximum number of dwells that can be used to detect and confirm targets for a given beam. Then π is a sequence of waveforms that could be used for that detection process. Let

$$V(x) = E\left[\sum_{k=0}^T c(x, u(k))\right]$$

where $c(x, u(k))$ is the reward earned when the scene x is observed using waveform $u(k)$. This cost function will typically express the capacity of the waveform to discriminate potential targets in the particular clutter environment expressed by state x . Then the aim of our problem is to find the sequence π^* that satisfies

$$V^*(x) = \max_{\pi} E\left[\sum_{k=0}^T c(x, u(k))\right]. \quad (1)$$

Now, our original aim was to design an optimal waveform selection algorithm that can adapt to the actual state of the radar scene. However, knowledge of the actual state is not available. Instead, we only have access to noisy measurements

of the scene. To handle this, let $Y^k = (y(1), y(2), \dots, y(k))$ and $U^k = (u(0), u(1), \dots, u(k))$ and then define

$$p_i(k) \triangleq \Pr(x = i | Y^k, U^{k-1})$$

that is, the vector $\mathbf{p}(k)$ is the conditional density of the state given the measurements and the controls. Using Bayes' rule and the Law of Total Probability the following recursion can be derived for $\mathbf{p}(k+1)$ [7], [13]

$$\begin{aligned} p_j(k+1) &= \Pr(x = j | Y^{k+1}, U^k) \\ &= \frac{b_{(y(k+1), j)}(u(k)) p_j(k)}{\sum_{i \in \mathcal{H}} b_{y(k+1), i}(u(k)) p_i(k)} \end{aligned}$$

Let $\bar{\mathbf{B}}(j, u)$ be the matrix with the vector $(b_{j,i}(u))_{i \in \mathcal{H}}$ down the diagonal and zeros elsewhere, then in matrix notation

$$\mathbf{p}(k+1) = \frac{\bar{\mathbf{B}}(y(k+1), u(k)) \mathbf{p}(k)}{\mathbf{1}' \bar{\mathbf{B}}(y(k+1), u(k)) \mathbf{p}(k)} \quad (2)$$

where $\mathbf{1}$ is a column vector of ones. Note, that the denominator of (2) is

$$\mathbf{1}' \bar{\mathbf{B}}(y(k+1), u(k)) \mathbf{p}(k) = \Pr(y(k+1) | Y^k, U^k)$$

which is not a function of $\mathbf{p}(k)$.

The conditional state density vector $\mathbf{p}(k)$ is also known as the *information state* and it is a sufficient statistic for the true state x . That is, our original stochastic control problem in terms of x can be rewritten in terms of $\mathbf{p}$ [1]. In other words, the optimal control policy π^* that is the solution of (1) is also the solution of

$$V^*(\mathbf{p}(0)) = \max_{\pi} E \left[\sum_{k=0}^T c(\mathbf{p}(k), u(k)) \right] \quad (3)$$

where $\mathbf{p}(0)$ is the *a priori* probability density of the scene.

IV. CALCULATING THE MEASUREMENT PROBABILITIES

The key feature of this model for waveform selection is the manner in which the measurement probabilities vary with the choice of waveform. Recall that the measurement probability $b_{ji}(u(k))$ is

$$b_{ji}(u(k)) = \Pr(y(k+1) = j | x = i, u(k))$$

which is the probability of obtaining the detection pattern described by hypothesis j when the true scene is i and is measured by waveform $u(k)$.

To illustrate how the measurement probabilities will vary with waveform choice consider the simple case when there is only a single cell in range and three in Doppler space. These three correspond to receding targets, (near) stationary targets and approaching targets. The set of possible scenes $\mathcal{H}$ is then

- $i = 1$ no targets present
- $i = 2$ a single, receding target
- $i = 3$ a single, stationary target
- $i = 4$ a single, approaching target
- $i = 5$ a receding target and a stationary target
- $i = 6$ a receding target and an approaching target
- $i = 7$ a stationary target and an approaching target
- $i = 8$ a receding, a stationary and an approaching target

Suppose the true scene is $i = 2$, i.e. there is a single, receding target. If the chosen waveform $u(k)$ has relatively good Doppler resolution then

$$b_{22}(u(k)) = \Pr(y(k+1) = 2 | x = 2, u(k))$$

will be high. However, the hypotheses that correspond to scenes in which other targets are present, as well as a receding target, will also be moderately high as noise and clutter may produce detections in the other cells also. That is, $b_{j2}(u(k))$ for $j = 5, 6, 8$ will be significant. The remaining hypotheses, $b_{j2}(u(k))$ for $j = 1, 3, 4, 7$ all do not contain a receding target so they will be the least likely as they require that the true target is not detected and that clutter or noise produces a false return in another cell.

On the other hand, suppose the chosen waveform has very poor Doppler resolution then we might find that

$$b_{j2} = \frac{1}{8}$$

for all j , i.e. the waveform provides no useful information about the scene.

V. OBJECTIVE FUNCTION

We will consider two basic forms for the objective function at time k , $c(\mathbf{p}(k), u(k))$, in the adaptive radar scheduling problem. The first, more simple function is

$$c(\mathbf{p}(k), u(k)) = \sum_{i \in \mathcal{H}} p_i(k) \log(p_i(k)) \quad (4)$$

where $p_i(k) \log(p_i(k))$ is set to 0 when $p_i(k)$ is sufficiently small. This function is at a maximum when the entropy is minimised. Therefore, maximising this objective function will produce a sequence of waveforms that determines the scene as accurately as possible over the allowed number of dwells. The inclusion of a discounting factor, γ^k , $0 < \gamma < 1$, in the value function (3) would modify the solution so that the most accurate estimate of the scene was obtained as quickly as possible.

The second form of the objective function allows the inclusion of information such as a known clutter map and allows possible targets to be classified according to their importance. In this case

$$c(\mathbf{p}(k), u(k)) = \sum_{i \in \mathcal{H}} p_i(k) \log(p_i(k)) \sum_{\tau, \nu} c_{\tau, \nu} \delta_{\tau, \nu}(i) \quad (5)$$

where $\delta_{\tau, \nu}(i) = 1$ if cell (τ, ν) contains a scatterer under hypothesis i and $c_{\tau, \nu}$ is a weight on the importance of detecting a scatterer in that cell. For cells that correspond to areas of known clutter or "uninteresting" targets such as receding targets that are far from the radar, this weight would be small. For "interesting" targets such as those at close range and/or approaching the radar this weighting factor would be large. The use of the second form of the objective function in (3) yields the optimal sequence of waveforms to detect the targets as accurately as possible with the more important targets detected more quickly and accurately than those of low importance. A discount factor can be included if it is also desired that the detection be performed as quickly as possible.

VI. SOLUTION METHOD

The stochastic control problem given by equations (3) and (2) belongs to the class of problems known as Partially Observed Markov Decision Problems (POMDP). An overview of these types of problems and methods for solving them can be found in [13] and [12]. There are a number of algorithms for finding both optimal and near-optimal solutions to these types of problems over a finite horizon. A survey of optimal algorithms for finite horizon problems can be found in [4].

The various algorithms make use of the fact that over a finite horizon the objective function of a POMDP is piecewise linear and convex. Thus the objective function for a given horizon, can be represented by a set of vectors. The various solution methods provide different ways of finding this set of vectors. As the size of the adaptive waveform selection problem is large, we will use the highly efficient refinement of the Witness algorithm [11] known as Incremental Pruning [5].

The dynamic programming algorithm [1] shows that the solution of the problem (3) can be found by proceeding backwards with the recursion

$$\begin{aligned} V_T(\mathbf{p}(T)) &= c(\mathbf{p}(T)) \\ V_k(\mathbf{p}(k)) &= \min_{u(k) \in \mathcal{U}} E[c(\mathbf{p}(k), u(k)) + \\ &\quad V_{k+1}(\mathbf{p}_{y(k+1)}^{u(k)}(k+1))], k = T-1, \dots, 0 \end{aligned}$$

where $\mathbf{p}_y^u$ is the solution of (2) when u is used at dwell $k+1$, generating observation y . The terminal cost $c(\mathbf{p}(T))$, does not depend on u since no further scans are made. The key result here is that it is possible to define a new objective function $\bar{V}$ in terms of a given objective function V . That is, we can write

$$\bar{V}(\mathbf{p}) = \max_{u \in \mathcal{U}} \left(c(\mathbf{p}, u) + \sum_{y \in \mathcal{H}} Pr(y|\mathbf{p}, u) V(\mathbf{p}_y^u) \right). \quad (6)$$

In [5] it is shown that this can be broken into a series of simpler combinations of other objective functions

$$\begin{aligned} \bar{V}(\mathbf{p}) &= \max_{u \in \mathcal{U}} V^u(\mathbf{p}) \\ V^u(\mathbf{p}) &= \sum_y V_y^u(\mathbf{p}) \\ V_y^u(\mathbf{p}) &= \frac{1}{|\mathcal{H}|} c(\mathbf{p}, u) + Pr(y|\mathbf{p}, u) V(\mathbf{p}_y^u). \end{aligned}$$

Over a finite horizon, the objective function V can be expressed as $V(\mathbf{p}) = \max_{\alpha \in \mathcal{S}} \mathbf{p}'\alpha$ for some finite set of vectors $\mathcal{S}$. This means that it is possible to write

$$\begin{aligned} V_y^u(\mathbf{p}) &= \max_{\alpha \in \mathcal{S}_y^u} \mathbf{p}'\alpha \\ V^u(\mathbf{p}) &= \max_{\alpha \in \mathcal{S}^u} \mathbf{p}'\alpha \\ \bar{V}(\mathbf{p}) &= \max_{\alpha \in \bar{\mathcal{S}}} \mathbf{p}'\alpha \end{aligned}$$

for some finite sets of vectors $\mathcal{S}_y^u$, $\mathcal{S}^u$ and $\bar{\mathcal{S}}$ for all $u \in \mathcal{U}$ and $y \in \mathcal{H}$. These sets have a unique representation of minimum size and [5] provides an efficient method for generating these sets given $\mathcal{S}$.

VII. EXAMPLE

In order to calculate $b_{ji}(u(k))$ we need to make some assumptions about the likelihood of a detection in a cell when there are scatterers in nearby cells. A common assumption in tracking in clutter problems is that a target located in a given cell will not effect a detection in any other. While this is an overly idealised assumption, we will use it here as it allows ready comparison with other work in this area.

A. Information State Recursion

Under the independence assumption above, the probability of target existence in a cell now only depends on detections in the cell and not on measurements in neighbouring cells. Thus the original information state recursion (2) (which is a vector of length 2^{NM}) reduces to NM independent scalar equations. Define

$$p_{\tau,\nu}(k) \triangleq Pr(\epsilon_{\tau,\nu}|Y^k, U^{k-1})$$

where $\epsilon_{\tau,\nu}$ is the event that there is a target in cell (τ, ν) . In other words, $p_{\tau,\nu}(k)$ is the probability of the existence of a target in cell (τ, ν) given all the information available up to time k . Let $y_{\tau,\nu}(k+1)$ be the measurement in the cell (τ, ν) at time $k+1$ (i.e. either a detection, $d_{\tau,\nu}$ or no detection, $\bar{d}_{\tau,\nu}$) then the recursion for $p_{\tau,\nu}$ is

$$p_{\tau,\nu}(k+1) = \frac{1}{\Delta} Pr(y_{\tau,\nu}(k+1)|\epsilon_{\tau,\nu}, u(k)) p_{\tau,\nu}(k)$$

where

$$\begin{aligned} \Delta &= Pr(d_{\tau,\nu}|\epsilon_{\tau,\nu}, u(k)) p_{\tau,\nu}(k) + \\ &\quad Pr(\bar{d}_{\tau,\nu}|\bar{\epsilon}_{\tau,\nu}, u(k)) (1 - p_{\tau,\nu}(k)) \end{aligned}$$

and $\bar{\epsilon}_{\tau,\nu}$ is the complementary event that there is no target in cell (τ, ν) .

B. Measurement Probabilities

Let the probability of a detection in cell (τ, ν) $Pr(d_{\tau,\nu}|\epsilon_{\tau,\nu}, u) = P_d^{\tau,\nu}(u)$ and the probability of a false alarm $Pr(\bar{d}_{\tau,\nu}|\bar{\epsilon}_{\tau,\nu}, u) = P_f^{\tau,\nu}(u)$. To calculate $P_d^{\tau,\nu}(u)$ and $P_f^{\tau,\nu}(u)$ we use the receiver model described in [15]. In this model, all targets have a Swerling 1 distribution and the noise is additive, white and Gaussian with known power. The extension to other models is straightforward.

Under the model of [15], when there is no target present, the output of the matched filter receiver is a complex Gaussian random variable with zero mean and variance given by

$$\sigma_0^2 = 2N_0\xi$$

where N_0 is the known, ambient noise power and ξ is the energy of the transmitted pulse. For convenience, we will assume ξ is the same for all possible waveforms as was done in both [15] and [9], although this is not required by our model.

The matched filter output in a cell centred on (τ_0, ν_0) when the target return has an actual time delay of τ and Doppler shift of ν is still zero mean and Gaussian, however the variance is given by

$$\sigma_1^2 = 2N_0\xi + 2\sigma_A^2\xi^2\mathcal{A}(\tau_0 - \tau, \nu_0 - \nu)$$

where σ_A^2 is the variance of the amplitude of the target return and $\mathcal{A}$ is the ambiguity function. The ambiguity function specifies the output of the matched filter in the absence of noise. It is given by the equation [16]

$$\mathcal{A}(\tau, \nu) = \frac{1}{(\int |s(\lambda)|^2 d\lambda)^2} \left| \int s(\lambda) s^*(\lambda - \tau) e^{2\pi j \nu \lambda} d\lambda \right|^2$$

where $s(t)$ is the transmitted baseband signal.

Recall, the magnitude square of a complex Gaussian random variable $y \sim \mathcal{N}(0, \sigma^2)$ is exponentially distributed, with the density

$$x = y^2 \sim \frac{1}{2\sigma^2} e^{-x/2\sigma^2}$$

so, if there is no target in the cell centred on (τ, ν) under hypothesis i and the detection threshold is D then

$$\begin{aligned} P_f^{\tau, \nu}(u) &= \int_D \frac{1}{2\sigma_0^2} \exp\left(\frac{-x}{2\sigma_0^2}\right) dx \\ &= \exp\left(\frac{-D}{2\sigma_0^2}\right) \\ &= P_f \end{aligned}$$

for all u since the energy of the transmitted pulse is assumed to be the same in all cases.

In the case when a target is present in a cell, assuming its actual location in the cell has a uniform distribution

$$\begin{aligned} P_d^{\tau, \nu}(u) &= \frac{1}{|A|} \int_{(\tau_a, \nu_a \in A)} \int_D \frac{1}{2\sigma_1^2} \exp\left(\frac{-x}{2\sigma_1^2}\right) dx d\tau_a d\nu_a \\ &= \frac{1}{|A|} \int_{(\tau_a, \nu_a \in A)} \exp\left(\frac{-D}{2\sigma_1^2}\right) d\tau_a d\nu_a \end{aligned}$$

where A is the resolution cell centred on (τ, ν) with volume $|A|$.

C. Objective Function

Under the independence assumption, it can be shown that the first form of the objective function

$$c(\mathbf{p}(k), u(k)) = \sum_{i \in \mathcal{H}} p_i(k) \log(p_i(k))$$

reduces to

$$\begin{aligned} c(\mathbf{p}(k), u(k)) &= \sum_{\tau, \nu} \{ p_{\tau, \nu}(k) \log(p_{\tau, \nu}(k)) + \\ &\quad (1 - p_{\tau, \nu}(k)) \log(1 - p_{\tau, \nu}(k)) \}. \end{aligned}$$

The second form of the objective function discussed in Section V becomes

$$\begin{aligned} c(\mathbf{p}(k), u(k)) &= \sum_{\tau, \nu} \{ c_{\tau, \nu} p_{\tau, \nu}(k) \log(p_{\tau, \nu}(k)) + \\ &\quad (1 - p_{\tau, \nu}(k)) \log(1 - p_{\tau, \nu}(k)) \} \end{aligned}$$

where $c_{\tau, \nu}$ is a weighting factor that reflects the importance of detecting a scatter in cell (τ, ν) .

VIII. CONCLUSION

By posing the target detection problem as a stochastic dynamic programming problem we are able to produce schemes for optimal waveform selection over a finite horizon. We are also able to develop a flexible framework that can be extended in a number of ways. These include changes to the way the detection probabilities are calculated to remove the idealised assumption that nearby scatterers do not interfere with one another. This framework can also be extended to consider tracking as well as detection performance.

REFERENCES

- [1] D. Bertsekas. *Dynamic Programming and Optimal Control*, volume 1. Athena Scientific, Belmont, MA, USA, 2nd edition, 2001.
- [2] S. Blackman and R. Popoli. *Design and Analysis of Modern Tracking Systems*. Artech House, Boston, USA, 1999.
- [3] W. D. Blair. Toward the integration of tracking and signal processing for phased array radar. In *Proc. of SPIE conference on Signal and Data Processing of Small Targets*, volume 2235, pages 303–316, Orlando, Florida, USA, 1994.
- [4] A. R. Cassandra. Optimal policies for solving partially observed Markov decision processes. Technical Report CS-94-14, Brown University, Providence, RI, USA, 1994.
- [5] A. R. Cassandra, M. L. Littman, and N. L. Zhang. Incremental pruning: A simple, fast, exact method for partially observable Markov decision processes. In *Proc. of 13th Annual Conference on Uncertainty in Artificial Intelligence (UAI-97)*, Providence, Rhode Island, USA, 1997.
- [6] R. A. Dana and D. Moraitis. Probability of detecting a Swerling I target on two correlated observations. *IEEE Trans. on Aerospace and Electronic Systems*, 17(5):727–730, 1981.
- [7] R. J. Elliott, L. Aggoun, and J. B. Moore. *Hidden Markov Models: Estimation and Control*. Springer-Verlag, New York, USA, 1995.
- [8] S.-M. Jong and Y.-H. Jung. Optimal scheduling of track updates in phased array radars. *IEEE Trans. on Aerospace and Electronic Systems*, 34(3):1016–1022, 1998.
- [9] D. J. Kershaw and R. J. Evans. Waveform selective probabilistic data association. *IEEE Trans. on Aerospace and Electronic Systems*, 33(4):1180–1188, 1997.
- [10] V. Krishnamurthy and R. J. Evans. Hidden Markov model multiarm bandits: A methodology for beam scheduling in multitarget tracking. *IEEE Trans. on Signal Processing*, 49(12):2893–2908, 2001.
- [11] M. L. Littman. The Witness algorithm for solving partially observed Markov decision processes. Technical Report CS-94-40, Brown University, Providence, RI, USA, 1994.
- [12] W. S. Lovejoy. A survey of algorithmic methods for partially observed markov decision processes. *Annals of Operations Research*, 28:47–66, 1991.
- [13] G. E. Monahan. A survey of partially observable markov decision processes: Theory, models and algorithms. *Management Science*, 28(3):1–16, 1982.
- [14] R. Niu, P. Willett, and Y. Bar-Shalom. Tracking considerations in selection of radar waveform for range and range-rate measurements. *IEEE Trans. on Aerospace and Electronic Systems*, 38(2):467–487, 2002.
- [15] C. Rago, P. Willett, and Y. Bar-Shalom. Detection-tracking performance with combined waveforms. *IEEE Trans. on Aerospace and Electronic Systems*, 34(2):612–624, 1998.
- [16] M. I. Skolnik. *Introduction to Radar Systems*. McGraw-Hill, 3rd edition, 2001.
- [17] D. Strömberg. Scheduling of track updates in phased array radars. In *Proc. of 1996 National Radar Conference*, pages 214–219, Ann Arbor, Michigan, USA, 1996.
- [18] G. V. Trunk, J. D. Wilson, and P. K. Hughes, II. Phased array parameter optimization for low-altitude targets. In *Proc. of IEEE Radar Conference*, pages 196–200, Washington, D.C., USA, 1995.
- [19] G. van Keuk and S. S. Blackman. On phased array radar tracking and parameter control. *IEEE Trans. on Aerospace and Electronic Systems*, 29(1):186–194, 1993.