

Chapter 10

APPLICATION OF SENSOR SCHEDULING CONCEPTS TO RADAR

William Moran

University of Melbourne, Melbourne, Australia

Sofia Suvorova

University of Melbourne, Melbourne, Australia

Stephen Howard

Defence Science and Technology Organisation, Edinburgh, Australia

1. Introduction

In this chapter, we spend some time illustrating the ideas on sensor scheduling in a specific context: that of a radar system. A typical pulse radar system operates by illuminating a scene with a short pulse of electromagnetic energy at radio frequency (RF) and collecting the energy reflected from the scene. Newer radar systems have the capability to modify the method of illumination in several ways. Parameters such as time between pulses or pulse repetition interval (PRI), carrier frequency, the shape of the transmitted waveform, the shape and direction of the beam, etc., which were typically hard-wired in previous generations of radars, may now be regarded as controllable during system operation. All of these parameters admit the possibility of modification based on knowledge of the scene, so as best to extract the required information. We remark that scheduling “on the transmit side” is quite different from the possibility of changing processing after receipt of the return. If different methods of processing are contemplated for the radar return, then in principle these can be done

in parallel, so that in some sense no choice between the two approaches is required. On the transmit side, in particular, parameter selections may have to be made in an especially timely way to keep pace with the system's operating tempo.

Ideally we would like, on the basis of what we now know about the scene, including "targets of interest" and uninteresting reflections, to schedule the transmission parameters over the long-term to maximize the operationally important information returned to the operator on time scales required for operational purposes. Such a goal would demand scheduling algorithms capable of undertaking non-myopic scheduling on a pulse-by-pulse basis. As we have seen in earlier Chapters, computational considerations make such a goal infeasible at the moment. Simplifications that permit suboptimal solutions are required. Ideas associated with these are discussed in the following sections; but first we describe the basic mathematics of radar.

2. Basic Radar

For our purposes, a radar system is comprised of a transmitter and a receiver of electromagnetic radiation at a common frequency and with a common clock. The clock maintains accurate times for transmission and reception of signals, and importantly it keeps a record of phase of the carrier.

2.1 Range-Doppler Ambiguity

The transmitter is assumed to emit a signal $\mathbf{w}_{\text{Tx}}(t)$. This signal is a rapidly oscillating sinusoid (the "carrier") on which is imposed a relatively slowly varying modulation $\mathbf{s}_{\text{Tx}}(t)$ (the "waveform"):

$$\mathbf{w}_{\text{Tx}}(t) = \mathbf{s}_{\text{Tx}}(t) \sin(2\pi f_c t), \quad (10.1)$$

where f_c is the carrier frequency.

For the moment we are not concerned with the direction of transmission or polarization of the transmit beam. While the total amount of energy depends on a number of factors, some of it hits each reflecting element in the scene and a small proportion of this energy is reflected back to the antenna of the receiver. To put energy issues in perspective, the total energy returning from an element of the scene falls off as the inverse fourth power of the distance of the element from the radar. This means that the amount of energy entering a radar from a small target at distance is often tiny even relative to the thermal noise in the receiver. Significant processing is needed to extract the signal from the noise in such circumstances.

The return to the receiver is generally assumed to be a linear superposition of reflections from the different elements of the scene. As a result, in the current analysis of the response of the radar to the scene, it is enough to consider the response to a single reflector, which we shall refer to as the *target*.

We assume that the receiver is collocated with the transmitter, so that the waveform returns to the receiver a time delay of $\tau = 2R/c$ later, where R is the distance to the target and c is the speed of light. The return signal is $\mathbf{w}_{\text{Tx}}(t - \tau)$ multiplied by an attenuation and phase shift factor due to the distance involved, the nature of the target, and atmospheric conditions. At this stage, we ignore this factor. On return to the receiver, the carrier is stripped off (demodulated).

There are two channels in the receiver; the in-phase (I) channel corresponding to demodulation by $\sin(2\pi f_c t)$, and the quadrature (Q) channel corresponding to demodulation by $\cos(2\pi f_c t)$. This duplication of effort is required to sense the phase (or at least the phase-differences) in the return. After demodulation, the signal in the I channel of the receiver is $s_{\text{Tx}}(t - \tau) \cos(2\pi f_c \tau)$ and in the Q channel is $-s_{\text{Tx}}(t - \tau) \sin(2\pi f_c \tau)$. We can now regard this as a complex waveform

$$s_{\text{Tx}}(t - \tau) e^{-2\pi i f_c \tau}. \quad (10.2)$$

In this way, we can think even of the transmitted signal $\mathbf{w}_{\text{Tx}}(t)$ as having values in the complex domain. The real and complex parts of the signal are imposed on the I and Q channels respectively to form the transmit waveform, so that again the complex return is

$$s_{\text{Tx}}(t - \tau) e^{-2\pi i f_c \tau}. \quad (10.3)$$

At this point, the signal is filtered against another signal $s_{\text{Rx}}(t)$. From a mathematical perspective, this produces the filtered response¹

$$\int_{\mathbb{R}} s_{\text{Tx}}(t - \tau) \overline{s_{\text{Rx}}(t - x)} dt. \quad (10.4)$$

So far we have assumed that the target is stationary. When the target moves the return is modified by the Doppler effect. If this is done correctly, it results in a “time dilation” of the return. For our purposes, where the signals have small bandwidth relative to the carrier frequency, and the pulses are relatively short, the narrowband approximation allows this time dilation to be closely approximated by a frequency shift, so that if the target has a radial velocity v ,

¹In this chapter, overbar is used to denote complex conjugation.

the return becomes

$$\mathbf{s}_{\text{Txc}}(t - \tau)e^{2\pi i f_c(1-2v/c)(t-\tau)}.$$

This is demodulated and filtered as in equation (10.4) to obtain

$$\begin{aligned} A_{\mathbf{s}_{\text{Txc}}, \mathbf{s}_{\text{Rx}}}(\tau, f_d) &= \int_{\mathbb{R}} e^{-2\pi i f_d(t-\tau)} \mathbf{s}_{\text{Txc}}(t - \tau) \overline{\mathbf{s}_{\text{Rx}}(t - x)} dt \\ &= \int_{\mathbb{R}} e^{-2\pi i f_d t} \mathbf{s}_{\text{Txc}}(t) \overline{\mathbf{s}_{\text{Rx}}(t - (x - \tau))} dt. \end{aligned} \quad (10.5)$$

This function of time delay τ and Doppler f_d is called the *ambiguity function* of the transmit and receive signals. To be more precise, this is referred to as the *cross ambiguity function* since the filtering kernel used by the receiver is, in principle, different from the transmitted signal. When the filtering kernel and transmit signal are the same, this filtering process is called *matched filtering* and its result the *auto-ambiguity function*. Propagation through the transmission medium and electronic processing in the receiver add noise to the signal and the effect of the matched filter is to maximize the signal power relative to the noise power in the processed signal.

No discussion of the ambiguity function is complete without a mention of Moyal's Identity. This states that if f and g are finite energy signals (i.e., in $L^2(\mathbb{R})$) then $A_{f,g}$ is in $L^2(\mathbb{R}^2)$ and for f_1, g_1 also in $L^2(\mathbb{R})$,

$$\langle A_{f_1, g_1}, A_{f, g} \rangle = \langle f_1, f \rangle \overline{\langle g_1, g \rangle}. \quad (10.6)$$

In particular this says that if consider the auto-ambiguities of f and f_1 , we obtain

$$\langle A_{f_1, f_1}, A_{f, f} \rangle = |\langle f_1, f \rangle|^2. \quad (10.7)$$

Finally, letting $f = f_1$ we find

$$\|A_{f, f}\|^2 = \|f\|^4. \quad (10.8)$$

It is clear also that $A_{f, f}(0, 0) = \|f\|^2$ and that this is the maximum value of the ambiguity (by the Cauchy-Schwarz Inequality). This implies that the ambiguity function must spread out; it cannot be concentrated close to the origin. Thus, there is always a non-trivial trade-off between Doppler and time delay (range) measurements. The choice of waveforms determines where this trade-off is pitched. At one extreme, for example, long sinusoidal waveforms will give good Doppler measurements but bad range measurements.

We return to a general range-Doppler scene; again for the moment no azimuth and, in this treatment will not discuss polarization. Such a scene is a

linear combination $\sum_k \sigma(\tau_k, \phi_k) \delta(t - \tau_k, f - \phi_k)$ of reflectors (scatterers) at delays $\tau_k = 2r_k/c$ and Doppler shifts $\phi_k = 2f_c v_k/c$, where the range and radial velocity of the k^{th} target are r_k and v_k respectively. The processed return is

$$\sum_k \alpha_k A_{\mathbf{s}_{\text{Txc}}, \mathbf{s}_{\text{Rx}}}(t - \tau_k, f - \phi_k). \quad (10.9)$$

We end this section with a few remarks on the use of waveforms in practice. First, as we have already remarked, there are good reasons to choose $\mathbf{s}_{\text{Rx}} = \mathbf{s}_{\text{Txc}}$. While this is by no means always the case, there have to be other good reasons (in terms of achievable side-lobes in the ambiguity, for example) not to do this. For the remainder of this section we make that assumption and restrict our attention to matched filtering and the auto-ambiguity function.

In view of the nature of the ambiguity function, it may seem that an optimal choice of waveform to achieve very accurate range resolution is a very short pulse; as sharp as the available bandwidth will allow. While this solution is sometimes adopted, it has disadvantages, the most important of which is that such a pulse carries very little energy. In view of the $1/r^4$ decay in return energy from a scatterer at distance r from the radar, it is important to transmit as much energy as other considerations will allow. Accordingly a “pulse compression” approach is adopted. Waveforms \mathbf{s}_{Txc} are chosen that are relatively long pulses (though, as we shall see in section 2.3, not too long) in time but are such that their auto-correlations (i.e., their ambiguity at zero Doppler) has a sharp peak at zero time delay and has small side-lobes away from this peak. In other words, the ambiguity function at zero Doppler is close to a “thumbtack” in the range direction.

We remark that for the kind of radar system we are interested in here (so-called “pulse-Doppler radars”), the pulse is short enough that Doppler has an insignificant effect on it. Effectively then the ambiguity is constant in the Doppler direction. One might ask why we have spent so much time on the ambiguity if it is so simple in this situation. The first reason is that this is totally general to any kind of radar and exhibits the limits of radar as a sensor. The second reason is that, while the effects of Doppler are relatively trivial in this context, the study of the ambiguity function quickly shows why this is the case. The final reason is that when we do Doppler processing in Section 2.3 it will be important to see that theory in the context of radar ambiguity.

Probably the most commonly used (complex) waveform in radar is a linear frequency modulated (LFM) or “chirp” pulse:

$$\mathbf{s}_{\text{Txc}}(t) = \exp 2\pi i \gamma t^2. \quad (10.10)$$

The “chirp rate” γ determines the effectiveness of the pulse. We illustrate the ambiguity with this example. The key feature is that the ambiguity has a ridge

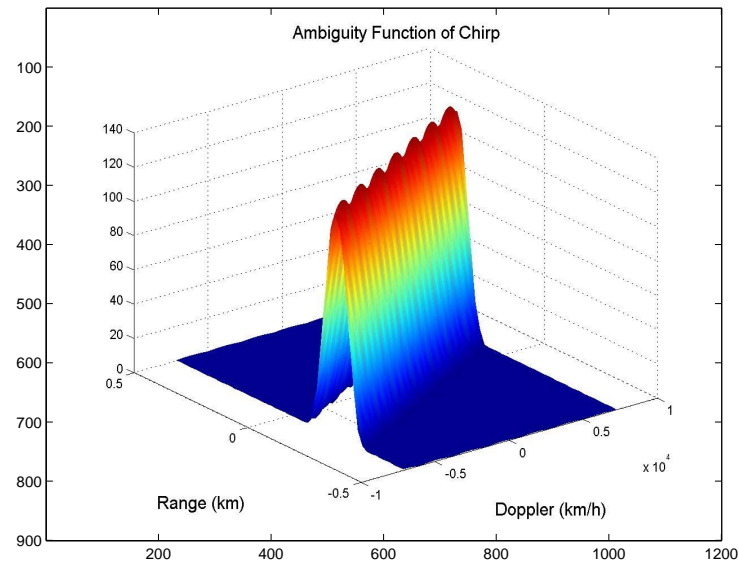


Figure 10.1. Absolute value of ambiguity of a linear frequency modulated pulse.

at an angle to the range axis and decays rapidly away from this ridge. Along the ridge is a slight ripple. The angle made by this ridge is a function of γ ; in fact, in the correct units, the tangent of the angle is γ .

2.2 Elevation and Azimuth

In addition to the measurements of range and, potentially, Doppler discussed in the previous section, it is usually important to have a good estimate of the azimuth of the target. In some applications the elevation is of interest too. The focus of this discussion will be on azimuth, while we remark that the elevation can be treated in essentially the same way. Many radars have a dish antenna of some kind that has a fixed beam-pattern, but it is in the philosophy of this book to consider systems in which the radar has some flexibility in the shape of its beam, as well as the ability to rapidly switch the direction of the beam. This is the case for systems using phased-array antennas and we discuss ambiguity for these systems.

Typically we are interested in scenes consisting of complex (that is, range and phase) information at each azimuth angle, range bin and Doppler bin.² Thus we shall regard a *scene* as a complex function $\sigma(r, v, \alpha)$ of range r , radial velocity v and azimuth α . More often, it will be described by a collection of scatterers at points (r_k, v_k, α_k) where the complex magnitude of a scatterer is $\sigma(r_k, v_k, \alpha_k)$. It will also be convenient, with some abuse of notation, to think of the scene in terms of delay $\tau_k = 2r_k/c$ and Doppler shift $\phi_k = 2f_c v_k/c$. The phase of the complex number represents the phase shift on the return induced by the carrier and the scatterer.

The array is assumed to have M antenna elements forming a horizontal linear array, which are separated by a uniform distance δ . The transmitter and receiver share the same antenna elements. Minor variants of this analysis will accommodate alternative schemes. The m^{th} antenna element transmits a relatively slowly varying signal $\mathbf{s}_{\text{Tx}}(t)$ modulated onto a rapidly varying carrier to produce the waveform $\mathbf{s}_{\text{Tx}}(t)e^{2\pi i f_c t}$. For this analysis, it is important that a time delay corresponding to the distance across the array face makes an insignificant difference to the signal $\mathbf{s}_{\text{Tx}}(t)$; this is the role of the “slowly varying” assumption. With this approximation, the k^{th} scatterer receives the waveform

$$\mathbf{s}_{\text{Tx}}(t - r_k)e^{2\pi i f_c m \delta \sin \alpha_k} e^{2\pi i f_c (t - r_k)}, \quad (10.11)$$

from the m^{th} antenna element.

We may assign a weight a_m (which can be complex) to each antenna element. This weight includes both a phase-shift and attenuation. Assuming omni-directional antenna elements, the k^{th} scatterer sees the total transmission from all channels (omitting the terms associated with carrier frequency) as

$$\begin{aligned} \sum_{m=1}^M a_m \mathbf{s}_{\text{Tx}}(t - r_k) e^{2\pi i f_c m \delta \sin \alpha_k} &= \mathbf{s}_{\text{Tx}}(t - r_k) \sum_{m=1}^M a_m e^{2\pi i f_c m \delta \sin \alpha_k} \\ &= \mathbf{s}_{\text{Tx}}(t - 2r_k) \phi_{\text{tr}}(\alpha_k), \end{aligned} \quad (10.12)$$

where $\phi_{\text{tr}}(\alpha) = \sum_{m=1}^M a_m e^{2\pi i f_c m \delta \sin \alpha}$ denotes the transmit pattern of the antenna array as a function of azimuth. If the elements are not omni-directional, the transmit pattern will need to be modified to accommodate the element pattern. We remark that the phase and attenuation serve different purposes in the design of the transmit pattern. The phase is used to point the beam in a

²Processed returns are typically quantized into discrete azimuth, range and Doppler values called *bins*.

given direction. Thus, to produce a beam with a main lobe at an angle α from boresight, a_m should equal $|a_m| = \exp 2\pi i \theta_m$ where $\theta_m = -m \sin \alpha$. The attenuation is used to shape the beam. Typically, this is manifested as a trade-off between sharpness of the main beam on the one hand, and height and spread of side-lobes on the other.

The return in channel m' due to the k^{th} scatterer will be, assuming that the scatterers are far enough away from the radar that the return can be regarded as a plane wave,³

$$\sigma(r_k, v_k, \alpha_k) \mathbf{s}_{\text{Txc}}(t - 2r_k) e^{4\pi i t f_c v_k / c} \phi_{\text{tr}}(\alpha_k) e^{-2\pi i f_c m' \delta \sin \alpha_k}, \quad (10.13)$$

where the $e^{4\pi i t f_c v_k / c}$ term in the product represents the effect of Doppler in shifting the frequency of the carrier. The data received then will be the sum over all scatterers in the scene

$$\sum_k \sigma(r_k, v_k, \alpha_k) \mathbf{s}_{\text{Txc}}(t - 2r_k) \phi_{\text{tr}}(\alpha_k) e^{4\pi i t f_c v_k / c} e^{-2\pi i f_c m' \delta \sin \alpha_k}. \quad (10.14)$$

In order to extract azimuthal information, the signal is Fourier transformed in the m' variable. It may also be “tapered” (or “windowed”) by additional weights $w_{m'}$. Thus we obtain

$$\begin{aligned} & \sum_m w_m e^{2\pi i f_c m \delta \sin \alpha} \sum_k \sigma(r_k, v_k, \alpha_k) \mathbf{s}_{\text{Txc}}(t - 2r_k) \phi_{\text{tr}}(\alpha_k) e^{4\pi i t f_c v_k / c} e^{-2\pi i f_c m \delta \sin \alpha_k} \\ &= \sum_k \sigma(r_k, v_k, \alpha_k) \mathbf{s}_{\text{Txc}}(t - 2r_k) \phi_{\text{tr}}(\alpha_k) e^{4\pi i t f_c v_k / c} \sum_m w_m e^{2\pi i f_c m \delta (\sin \alpha - \sin \alpha_k)} \end{aligned} \quad (10.15)$$

If we write $\theta(\alpha, \alpha_k) = \sum_m w_m e^{2\pi i f_c m \delta (\sin \alpha - \sin \alpha_k)}$, we obtain the processed return

$$\sigma(r_k, v_k, \alpha_k) \mathbf{s}_{\text{Txc}}(t - 2r_k) \phi_{\text{tr}}(\alpha_k) e^{4\pi i t f_c v_k / c} \theta(\alpha, \alpha_k). \quad (10.16)$$

Finally this is filtered as in Section 2.1 to obtain

$$\sum_k \sigma(r_k, v_k, \alpha_k) A_{\mathbf{s}_{\text{Txc}}, \mathbf{s}_{\text{Rx}}}(u - \tau_k, f - \phi_k) \phi_{\text{tr}}(\alpha_k) \theta(\alpha, \alpha_k). \quad (10.17)$$

The important conclusion to draw from this analysis is that the range-Doppler-azimuth ambiguity splits into a product of terms corresponding to azimuthal ambiguity $\theta(\alpha, \alpha_k) \phi_{\text{tr}}(\alpha_k)$ on the one hand and range-Doppler ambiguity $A(u - \tau_k, f - \phi_k)$ on the other.

³This is the so-called “far-field” assumption.

Range-Doppler ambiguity has been discussed at length in Section 2.1. The azimuthal ambiguity is

$$\begin{aligned}\theta(\alpha, \alpha_k)\phi_{\text{tr}}(\alpha_k) &= \sum_m w_m e^{2\pi i f_c m \delta (\sin \alpha - \sin \alpha_k)} \sum_{m'=1}^M a_{m'} e^{2\pi i f_c m' \delta \sin \alpha_k} \\ &= \sum_{m, m'=1}^M a_{m'} w_m \exp\left(2\pi i f_c \delta (m \sin \alpha + (m' - m) \sin \alpha_k)\right).\end{aligned}\quad (10.18)$$

Typically if we are looking in the direction α , we would want to point the transmit beam in that direction. This corresponds to a choice $a_{m'} = r_{m'} e^{-2\pi i f_c m' \delta \sin \alpha}$ where $r_{m'} > 0$ for all m' . This gives an azimuthal ambiguity of

$$\sum_{m, m'=1}^M r_{m'} w_m \exp\left(2\pi i f_c \delta ((m - m')(\sin \alpha - \sin \alpha_k))\right) = \Gamma(\sin \alpha - \sin \alpha_k).\quad (10.19)$$

A choice of $w_m = r_m$ so that the transmit and receive beam-patterns are the same will give

$$\Gamma(s) = \left| \sum_{m=1}^M w_m \exp 2\pi i \delta f_c m s \right|^2.\quad (10.20)$$

The choice $w_m = 1$ for all m produces

$$\Gamma(s) = \left| \sum_{m=1}^M \exp 2\pi i \delta f_c m s \right|^2 = \left(\frac{\sin \pi \delta f_c (M + 1) s}{\sin \pi \delta f_c s} \right),\quad (10.21)$$

which is a slowly decaying *sinc* response. Alternative choices of attenuation produce reduced sidelobes but at the expense of wider central beam. Roughly speaking, the flatter the coefficients w_m the more pronounced the sidelobes are. If the coefficients are tapered toward the ends of the antenna, then typically the sidelobes will be lower but the main lobe is broader. Thus a more sharply focused beam results in greater sidelobes. There is a considerable literature on the choice of these “spatial filters” to shape the beam. We refer the reader to the books mentioned at the beginning of this chapter for more details.

2.3 Doppler Processing

As we have seen in considering the ambiguity function, an individual waveform is capable of detecting the movement of a target toward or away from the radar. However, in practice this is rarely the way the detection of range-rate is achieved. If only range-rate (and not range) is required, as is the case, for example, with radars for the detection of highway speeding infringements, then a

standard solution is to transmit a continuous sinusoid. The shift in frequency is detected by a filter that eliminates the zero-Doppler component usually dominating the return and resulting from the direct return of the transmitted signal to the receiver from the transmitter and fixed elements in the scene.

When knowledge of range is required along with range-rate, the choice of waveforms represents a compromise to accommodate various conflicting requirements. On the one hand they should be long to produce “energy on target” and to allow Doppler to have a measurable effect. On the other hand the longer the pulse (in a *monostatic* mode; i.e., when the transmitter and receiver are collocated) the longer the “blind range.” Since in this scheme the receiver can only receive the return when the transmitter is not illuminating the scene, there is a distance of half the length of the pulse (in meters) that cannot be viewed by the radar. In any case, within a pulse there should be some modulation to produce good pulse compression.

Since an acceptable blind range conflicts with the pulse length required to produce measurable intra-pulse Doppler, radar engineers have had to resort to a trick that overcomes both the Doppler and blind range constraints. The method of “Doppler processing” involves the transmission of a short waveform over multiple pulses, with a significant “listening period” between the pulses. While on each pulse the Doppler has negligible effect, it does have a measurable effect over the sequence of pulses. It is important for this process that the oscillator used to time the operation of the system maintain coherence over the extent of these multiple pulses.

We assume that N_{dp} pulses $s_{Tx_c}(t - n\Delta)$ are transmitted at a uniform time separation of Δ in such a way as to retain coherence; i.e., the phase information from the carrier is retained.

We have learned from the preceding two sections that

- 1 We can consider a single target;
- 2 We can separate the effects of range-Doppler on the one hand and azimuth (and elevation if required) on the other.

Accordingly we treat the case of a single scatterer at delay τ_1 and Doppler frequency ϕ_1 and ignore the azimuthal dependence. As before, this exposition can take place “at baseband;” i.e., without inclusion of the carrier term $\exp 2\pi i f_c t$. The return from the n^{th} pulse is then

$$s_{Tx_c}(t - \tau_1 - n\Delta)e^{2\pi i \phi_1 t}, \quad (10.22)$$

with an appropriate multiplicative factor to account for the decay due to range and the reflectivity of the scatterer. The intra-pulse time variable t is known as “fast time” and the pulse count n as “slow time.”

It is assumed that the inter-pulse time Δ is long enough that most of the return from a given transmit pulse arrives back within the subsequent listening period. Of course this cannot be guaranteed, especially if there are large distant scatterers. The problem of “range-aliasing” is a significant one in radar. We shall return to it briefly later.

We collect the N_{dp} returns and stack them alongside each other so that ranges match. The effect is to appear to shift the pulses in time by an appropriate multiple of Δ . Thus we obtain

$$s'_{\text{rec},n}(t, m) = p(t - \tau_1)e^{2\pi i\phi_1(t+n\Delta)}. \quad (10.23)$$

Now we take a Fourier transform in the n (slow time) direction to obtain

$$\begin{aligned} S_{\text{rec}}(t, \nu) &= \sum_{n=0}^{N_{\text{dp}}-1} e^{-2\pi i n \frac{\nu}{N_{\text{dp}}-1}} e^{2\pi i\phi_1(t-n\Delta)} \mathbf{s}_{\text{Txc}}(t - \tau_1)\theta \\ &= \mathbf{s}_{\text{Txc}}(t - \tau_1)e^{2\pi i\phi_1 t} \sum_{n=0}^{N_{\text{dp}}-1} e^{2\pi i n(\phi_1\Delta - \frac{\nu}{N_{\text{dp}}})} \\ &= \mathbf{s}_{\text{Txc}}(t - \tau)e^{t\pi i\phi_1 t} \psi(\nu, \phi_1), \end{aligned} \quad (10.24)$$

where $\psi(\nu, \phi) = \sum_{n=0}^{N_{\text{dp}}-1} e^{2\pi i n(\phi\Delta - \frac{\nu}{N_{\text{dp}}})}$. Note that this is (up to a phase factor) just a “periodic sinc” function with a peak when $\nu = 2N f_c \frac{v_k}{c} \Delta$ and ν 's which differ from this by multiples of N_{dp} .

Finally we perform a matched filter against the original waveform $\mathbf{s}_{\text{Rx}}(t)$ to obtain

$$\begin{aligned} \int_{\mathbb{R}} \mathbf{s}_{\text{Txc}}(t - \tau_1)e^{2\pi i\phi_1 t} \mathbf{s}_{\text{Rx}}(t - u) dt \psi(\nu, \phi_1) \\ = e^{2\pi i u \phi_1} A(u - \tau_1, \phi_1) \psi(\nu, \phi_1), \end{aligned} \quad (10.25)$$

where

$$A_{\mathbf{s}_{\text{Txc}}, \mathbf{s}_{\text{Rx}}}(u, d) = \int_{\mathbb{R}} \mathbf{s}_{\text{Rx}}(t - u)e^{2\pi i d t} \mathbf{s}_{\text{Txc}}(t) dt, \quad (10.26)$$

is the ambiguity function as described in Section 2.1.

Observe now that with the assumption of small Dopplers and short pulses, the ambiguity $A(u - \tau_1, \phi_1)$ is approximately $A(u - \tau_1, 0)$. This results in a separation of the Doppler and range ambiguities in the same way as we split off the azimuthal ambiguity. The Doppler ambiguity happens in slow time, the range ambiguity in fast time.

It might be asked now what happened to Moyal's identity and the constraints it imposes on the range and Doppler ambiguities. Moyal's identity still applies, of course. We observe that the Doppler has been sampled at the rate given by Δ . This means that there is again an aliasing problem for targets of which the range rate gives a frequency in excess of the Nyquist frequency for this sampling rate. The smaller Δ is, the less the effect of Doppler aliasing, but the greater the effect of range aliasing. The total ambiguity of the pulse train in this case has peaks at points of the form $(n\Delta, \frac{m}{\Delta})$ in the range-Doppler plane.

There is a trade-off then between Doppler aliasing and range aliasing. This multi-pulse concept can be fitted into the ambiguity picture given in Section 2.1 and the trade-off between range and Doppler aliasing is just a consequence of Moyal's Identity. From the scheduling viewpoint, the time Δ and the number of pulses N_{dp} are both available for modification in real time.

2.4 Effects of Ambiguity

Given a scene, the radar return is produced by "convolving" it with the *generalized ambiguity function*.

$$G(t; \nu, d; \alpha, \beta) = A(t, 2f_c \frac{\nu}{c}) \phi_{tr}(\beta) \theta(\alpha, \beta) \psi(\nu, v_k), \quad (10.27)$$

where, by convolution, we mean here that the processed return is just

$$S_{pr}(u, \nu, \alpha) = \sum_k \sigma(r_k, v_k, \alpha_k) G(u - 2r_k; \nu, v_k; \alpha, \alpha_k). \quad (10.28)$$

In assessing the effect of a given radar mode, it is enough then to store for each waveform, taper and transmit pattern the generalized ambiguity, including (where appropriate) Doppler processing. Moreover, when we wish to test a waveform against a predicted scene, so as to choose one that is optimal in some quantifiable sense, it is this object that we use.

In fact the stored information can be somewhat reduced by first taking out the transmit pattern $\phi_{tr}(\beta)$, since it is scene independent. Then we may regard θ as a function of $\sin(\alpha)$; i.e., we can define $\theta_r(y) = \sum_m w_m e^{2\pi i m \delta y}$, so that

$$\theta(\alpha, \beta) = \theta_r(\sin(\alpha) - \sin(\beta)). \quad (10.29)$$

Thus we may make the generalized ambiguity

$$G_r(u, \nu, v, y) = A(u, 2f_c \frac{v}{c}) \theta_r(y) \psi(\nu, v), \quad (10.30)$$

and then the convolution equation (10.28) becomes

$$S_{\text{pr}}(t, \nu, \alpha) = \phi_{\text{tr}}(\beta) \sum_k \sigma(r_k, v_k, \alpha_k) G_r(t - r_k, \nu, v_k, \sin(\alpha) - \sin(\alpha_k)). \quad (10.31)$$

3. Measurement in Radar

Here we discuss the issue of measurement. As we have seen in Section 2.4, the radar measurement process amounts to a smearing of the scene by the ambiguity function. In order to incorporate these ideas into tracking and sensor scheduling problems, we need to find a simplification of the measurement process. We restrict attention for the moment to a radar that is undertaking only range and Doppler measurements, as in Section 2.1.

Our treatment here follows very closely that of van Trees in [235, Ch. 10]. As shown there, the log-likelihood associated with the measurement of range and Doppler of a single point target at a delay τ_a and Doppler ω_a resulting from the use of a waveform $\mathbf{w}(t)$ and a matched-filter is (leaving aside constants associated with the radar parameters that are unchanged throughout our calculations),

$$\Lambda(\tau, \omega) = |\mathbf{b}|^2 |A_{\mathbf{w}}(\tau - \tau_a, \omega - \omega_a)|^2 + 2\Re\left\{ \mathbf{b} A_{\mathbf{w}}(\tau - \tau_a, \omega - \omega_a) \mathbf{n}^*(\tau, \omega) \right\} + |\mathbf{n}(\tau, \omega)|^2, \quad (10.32)$$

where \mathbf{b} is a complex Gaussian random variable representing the reflectivity of the target, and $\mathbf{n}(\tau, \omega)$ is the integral of the receiver noise $\mathbf{N}(t)$ emanating from the filtering and Doppler processing; i.e.,

$$\mathbf{n}(\tau, \omega) = \int_{-\infty}^{\infty} \mathbf{N}(t) \mathbf{w}^*(t - \tau) e^{-i\omega t} dt. \quad (10.33)$$

The Fisher information matrix \mathbf{J} associated with this measurement is

$$\mathbf{J} = \begin{pmatrix} J_{\tau\tau} & J_{\tau\omega} \\ J_{\omega\tau} & J_{\omega\omega} \end{pmatrix} = -\mathbb{E} \left(\begin{bmatrix} \frac{\partial^2 \Lambda(\tau, \omega)}{\partial \tau^2} & \frac{\partial^2 \Lambda(\tau, \omega)}{\partial \tau \partial \omega} \\ \frac{\partial^2 \Lambda(\tau, \omega)}{\partial \omega \partial \tau} & \frac{\partial^2 \Lambda(\tau, \omega)}{\partial \omega^2} \end{bmatrix} \right), \quad (10.34)$$

where the expectation is taken over the independent random processes, \mathbf{b} and $\mathbf{N}(t)$. In our situation, it can be assumed that $J_{\tau\omega} = J_{\omega\tau}$. Computations yield, again within a constant factor,

$$\begin{aligned}
 J_{\tau\tau} &= \int_{-\infty}^{\infty} \omega^2 |W(\omega)|^2 \frac{d\omega}{2\pi} - \left(\int_{-\infty}^{\infty} \omega |W(\omega)|^2 \frac{d\omega}{2\pi} \right)^2 \\
 J_{\tau\omega} &= \int_{-\infty}^{\infty} t \mathbf{w}(t) \mathbf{w}'(t) dt - \int_{-\infty}^{\infty} \omega |W(\omega)|^2 \frac{d\omega}{2\pi} \int_{-\infty}^{\infty} t |\mathbf{w}(t)|^2 dt \\
 J_{\omega\omega} &= \int_{-\infty}^{\infty} t^2 |\mathbf{w}(t)| dt - \left(\int_{-\infty}^{\infty} t |\mathbf{w}(t)|^2 dt \right)^2.
 \end{aligned} \tag{10.35}$$

For any given waveform this can be calculated. Of course the Fisher information matrix determines the Cramér-Rao lower bound for an estimator of the variables τ, ω . In [234], van Trees argues that in many circumstances it provides a reasonable approximation to the inverse of the covariance matrix of the measurement process. As is easily seen, \mathbf{J}^{-1} is the inverse of the covariance matrix of a Gaussian that is the best approximation to $|A(\tau, \omega)|^2$ at its peak at $(0, 0)$.

Accordingly \mathbf{J} has been used as the basis for assessing the performance of a waveform in measuring the time delay and Doppler of a target for the purposes of scheduling. The sensor is characterized by a measurement noise covariance matrix

$$\mathbf{R} = \mathbf{T} \mathbf{J}^{-1} \mathbf{T}, \tag{10.36}$$

where \mathbf{T} is the transformation matrix between the time delay and Doppler measured by the receiver and the target range and velocity.

In particular it forms the basis of the seminal work of Kershaw and Evans on this topic. We describe this work next.

4. Basic Scheduling of Waveforms in Target Tracking

The aim of this section is to describe the work of Kershaw and Evans [131, 130] for scheduling of waveforms. This work has been at the basis of much subsequent work in the area of waveform scheduling.

4.1 Measurement Validation

Our assumption is that the radar produces measurements of the form

$$\mathbf{y}_r^k = \mathbf{H}\mathbf{x}_k + \omega_r^k, \quad (10.37)$$

where $r = 1, \dots, m_k$ are the detections at each time instance k , \mathbf{H} is measurement matrix and the noise ω_r^k is Gaussian with zero mean and covariance \mathbf{R}_k as discussed in Section 3, identical for all $r = 1, \dots, m_k$.

At time k the system will have produced, based on the previous measurements, an estimated position $\mathbf{x}_{k|k-1}$ of the target and an innovation covariance matrix \mathbf{S}_k . The details about how this is done are given in the next section. Based on these estimates we establish a *validation gate*. This will be the ellipsoid

$$\{\mathbf{y} : (\mathbf{y} - \mathbf{H}\mathbf{x}_{k|k-1})^\top (\mathbf{S}_k)^{-1} (\mathbf{y} - \mathbf{H}\mathbf{x}_{k|k-1}) < g^2\}, \quad (10.38)$$

centered on the estimated measurement $\mathbf{H}\mathbf{x}_{k|k-1}$ of the target at time k , where g is a threshold specified in advance. We shall need the volume of this ellipsoid for later calculations. In the range-Doppler case, it is

$$V_k = \pi g^2 \det(\mathbf{S}_k)^{1/2}. \quad (10.39)$$

Only measurements \mathbf{y}_r^k lying in this ellipsoid will be considered; the rest are discarded. This will facilitate computation and in any case eliminates outliers that might corrupt the tracking process.

The measurements that result then are:

- 1 A finite number of points \mathbf{y}_r^k .
- 2 For each of these points a covariance matrix \mathbf{R}_k .

4.2 IPDA Tracker

For convenience, we first describe the PDA tracker. Then we add the features specific for the IPDA tracker. We refer to the paper of Musicki, Evans and Stankovic [177] for a more detailed description of this.

For the moment we fix a waveform \mathbf{w} used to make the measurements. The target is assumed to follow a standard Gauss-Markov model. Target position $\mathbf{x}_k = (x_k, \dot{x}_k, \ddot{x}_k)$ in range-Doppler-acceleration space at time k moves with essentially zero acceleration in range according to the dynamics

$$\mathbf{x}_{k+1} = \mathbf{F}\mathbf{x}_k + \nu_k \quad (10.40)$$

where ν_k is Gaussian with zero mean and covariance matrix \mathbf{Q} (both \mathbf{Q} and \mathbf{F} are independent of k). The matrix \mathbf{F} is of the form

$$\mathbf{F} = \begin{pmatrix} 1 & \Delta & \frac{\Delta^2}{2} \\ 0 & 1 & \Delta \\ 0 & 0 & 1 \end{pmatrix}, \quad (10.41)$$

where Δ is the time step between the epochs.

Each measurement has associated with it an *innovation*

$$\mathbf{v}_r^k = \mathbf{y}_r^k - \mathbf{H}\mathbf{x}_{k|k-1}, \quad (10.42)$$

where $\mathbf{x}_{k|k-1}$ is the estimate of \mathbf{x}_k given the previous measurements, \mathbf{y}_r^j for $j < k$ and all r .

The covariance matrix of this statistic is the *innovation covariance matrix*

$$\mathbf{S}_k = \mathbf{H}\mathbf{P}_{k|k-1}\mathbf{H}^\top + \mathbf{R}_k. \quad (10.43)$$

Note that this is independent of the particular measurement \mathbf{y}_r^k and only depends on the waveform used at time k through the error covariance matrix \mathbf{R}_k and prior data. Here $\mathbf{P}_{k|k-1}$ is the error covariance associated with the state estimator $\mathbf{x}_{k|k-1}$ and is calculated using the standard Riccati equations

$$\mathbf{x}_{k|k-1} = \mathbf{F}\mathbf{x}_{k-1|k-1}, \quad \mathbf{P}_{k|k-1} = \mathbf{F}\mathbf{P}_{k-1|k-1}\mathbf{F}^\top + \mathbf{Q}. \quad (10.44)$$

We write P_D for the target detection probability and P_G for the probability that the target, if detected, is in the validation gate. Note that P_G is just the total Gaussian probability inside the validation gate and can be assumed to be unity in practice, since P_G is greater than 0.99 when $g > \sqrt{2} + 2$.

We assume that clutter is uniformly distributed with density ρ and write

$$b_k = 2\pi\rho\sqrt{\det(\mathbf{S}_k)}\frac{(1 - P_D P_G)}{P_D} = 2\frac{\rho V_k (1 - P_D P_G)}{g^2 P_D} \quad (10.45)$$

and

$$e_r^k = \exp\left(-\mathbf{v}_r^k(\mathbf{S}_k)^{-1}\mathbf{v}_r^k\right). \quad (10.46)$$

Then the probabilities that none of the measurements resulted from the target b_0^k and that the r^{th} measurement is the correct one are, respectively,

$$b_0^k = \frac{b_k}{b_k + \sum_r e_r^k}, \quad b_r^k = \frac{e_r^k}{b_k + \sum_r e_r^k}. \quad (10.47)$$

Note that (of course) $\sum_{r=0}^{m_k} b_r^k = 1$.

We now (using the standard PDA methodology — see [83] for the most lucid version of this) replace the innovation \mathbf{v}_r^k by the combined innovation

$$\mathbf{v}_k = \sum_{r=1}^{m_k} b_r^k \mathbf{v}_r^k. \quad (10.48)$$

Next we use the Kalman update to obtain an estimate of the state vector $\mathbf{x}_{k|k}$:

$$\mathbf{x}_{k|k} = \mathbf{x}_{k|k-1} + \sum_{r=1}^{m_k} b_r^k \mathbf{K}_k \mathbf{v}_r^k, \quad \text{where } \mathbf{K}_k = \mathbf{P}_{k|k-1} \mathbf{H}^\top (\mathbf{S}_k)^{-1}. \quad (10.49)$$

The update of the error covariance matrix is

$$\mathbf{P}_{k|k} = \mathbf{P}_{k|k-1} - \sum_r b_r^k \mathbf{K}_k \mathbf{S}_k (\mathbf{K}_k)^\top + \mathbf{P}_k, \quad (10.50)$$

where

$$\mathbf{P}_k = \sum_{r=1}^{m_k} b_r^k \left(\mathbf{K}_k \mathbf{v}_r^k (\mathbf{K}_k \mathbf{v}_r^k)^\top - \mathbf{K}_k \mathbf{v}_k (\mathbf{K}_k \mathbf{v}_k)^\top \right). \quad (10.51)$$

A standard approximation to it is used to obtain an estimate for the error covariance. This technique of Fortmann et al. in [83], replaces the random elements in (10.50) by their expectations. In fact, we replace \mathbf{P}_k and b_0^k by

$$\bar{\mathbf{P}}_k = \mathbb{E}[\mathbf{P}_k | Y^{k-1}], \quad \bar{b}_0^k = \mathbb{E}[b_0^k | Y^{k-1}], \quad (10.52)$$

where Y^{k-1} represents the measurement history at time $k-1$.

We refrain from repeating the calculations of [83] here. The results are that, if we write

$$\begin{aligned} q_1 &= \frac{P_D}{2} \int_0^g r^3 e^{-r^2/2} dr \\ I_2(m, b, g) &= \frac{P_D}{2} \sum_{m=1}^{\infty} \frac{e^{-\rho V_k} (\rho V_k)^{m-1}}{(m-1)!} \left(\frac{2}{g^2} \right)^{m-1} \\ &\quad \times \int_0^g \dots \int_0^g \frac{e^{-u_1^2} u_1^2}{b + \sum_{r=1}^m e^{-u_r^2/2}} du_1 \dots du_r \end{aligned} \quad (10.53)$$

where b_k is defined in equation (10.45), V_k in equation (10.39) and g is the “radius” of the validation gate, then

$$\bar{\mathbf{P}}_{k|k}(\theta_k) \approx \mathbf{P}_{k|k-1} - (P_D P_G - q_1 + q_2) \sum_r b_r^k \mathbf{K}_k \mathbf{S}_k(\mathbf{K}_k)^\top. \quad (10.54)$$

Note that Fortmann et al. introduce further approximations in [83]. Specifically, they note that for typical values of g , $P_G \approx 1$ and $q_1 \approx P_D$. Thus equation (10.54) simplifies to

$$\bar{\mathbf{P}}_{k|k} \approx \mathbf{P}_{k|k-1} - q_2 \sum_r b_r^k \mathbf{K}_k \mathbf{S}_k(\mathbf{K}_k)^\top. \quad (10.55)$$

Calculation of $I_2(m, b, g)$ in equation (10.53) requires a numerical integration scheme for a range of values of m , b and g . Kershaw and Evans [131] use an approximation to q_2 in the two-dimensional case:

$$q_2 \approx \frac{0.997 P_D}{1 + 0.37 P_D^{-1.57} \rho V_k}. \quad (10.56)$$

The IPDA tracker introduces the notion of *track existence* into the PDA tracker. We write χ_k to be the event of track existence at time k . This is assumed to behave as a Markov chain in the sense that

$$\mathbf{P}(\chi_k | Y^{k-1}) = p_1 \mathbf{P}(\chi_{k-1} | Y^{k-1}) + p_2 (1 - \mathbf{P}(\chi_{k-1} | Y^{k-1})), \quad (10.57)$$

where p_1 and p_2 are between 0 and 1. The choice of these is to some extent arbitrary, though simulations (R. Evans, private communication) indicate that the performance of the system is not very sensitive to the choice. On the other hand, if we assume that p_1 and p_2 are equal, then $\mathbf{P}(\chi_{k+1} | Y^k) = p_1$ and is independent of the previous history. A more appropriate choice then would have $p_1 > p_2$. There is clearly room for more work here. We write $f(\mathbf{y} | Y^{k-1})$ for the measurement Gaussian density corresponding to the mean $\mathbf{x}_{k|k-1}$ and covariance matrix $\mathbf{P}_{k|k-1}$, so that this represents the density of the predicted next measurement. Define

$$P_G = \int_{V_k} f(\mathbf{y} | Y^{k-1}) d\mathbf{y}. \quad (10.58)$$

This is the probability that the next measurement will be inside the validation gate. Let

$$p(\mathbf{y} | Y^{k-1}) = \frac{1}{P_G} f(\mathbf{y} | Y^{k-1}) \quad (\mathbf{y} \in V_k). \quad (10.59)$$

This is the conditional probability density conditioned on the measurement falling in the gate.

Now let

$$\delta_k = \begin{cases} P_G P_D & \text{if } m_k = 0; \\ P_G P_D \left(1 - \sum_{r=1}^{m_k} \frac{p(\mathbf{y}_r^k | Y^{k-1})}{\rho_k(\mathbf{y}_r^k)} \right) & \text{if } m_k > 0; \end{cases} \quad (10.60)$$

where P_D is the *a priori* probability of detection of a target.

We update the probability of track existence by

$$\mathbf{P}(\chi_k | Y^k) = \frac{1 - \delta_k}{1 - \delta_k \mathbf{P}(\chi_k | Y^{k-1})} \mathbf{P}(\chi_k | Y^{k-1}). \quad (10.61)$$

The conditional probabilities that measurement r at time k originated from the potential target are now given by

$$b_r^k = \begin{cases} \frac{1 - P_D P_G}{1 - \delta_k} & \text{if } r = 0; \\ \frac{P_D P_G \frac{p(\mathbf{y}_r^k | Y^{k-1})}{\rho_k(\mathbf{y}_r^k)}}{1 - \delta_k} & \text{if } r > 0. \end{cases} \quad (10.62)$$

These are then used, as explained in the preceding part of this section, to calculate the state estimate $\mathbf{x}_{k|k}$ and the error covariance $\mathbf{P}_{k|k}$ (see equations (10.49–10.52)). The Markov chain property (10.57) is used to update the track existence probability. It is this value which is thresholded to provide a true or false track detection test.

5. Measures of Effectiveness for Waveforms

To optimize the choice of waveform at each epoch it is necessary to calculate a cost function for all available waveforms in the library. This cost function should be a function of:

- 1 the predicted clutter distribution at that epoch based on a clutter mapper as described in, for example, [177];
- 2 the estimated position of the potential target at that epoch, based on the IPDA tracker.

This section discusses two potential waveform evaluators: the single noise covariance (SNC) matrix and the integrated clutter measure (ICM).

5.1 Single Noise Covariance Model

The approximate equation (10.54) for the error covariance gives a measure of the effectiveness of each waveform. The parameter θ is used to range over a collection of waveforms, and so the error covariances $\mathbf{R}_k = \mathbf{R}_k(\theta)$ depend on this parameter θ , as do all objects calculated in terms of them. As already indicated in Section 4.2, the covariance of the innovation \mathbf{v}_r^k equation (10.43) and the gain matrix \mathbf{K}_k (see equation (10.49)) are, respectively:

$$\mathbf{S}_k(\theta) = \mathbf{H}\mathbf{P}_{k|k-1}\mathbf{H}^\top + \mathbf{R}_k(\theta), \quad \mathbf{K}_k(\theta) = \mathbf{P}_{k|k-1}\mathbf{H}^\top\mathbf{S}_k(\theta)^{-1}. \quad (10.63)$$

Substituting these in the approximate equation for the gain matrix, we obtain

$$\bar{\mathbf{P}}_{k|k}(\theta) = \mathbf{P}_{k|k-1} - q_2\mathbf{P}_{k|k-1}\mathbf{H}^\top \left(\mathbf{H}\mathbf{P}_{k|k-1}\mathbf{H}^\top + \sum_r b_r^k \mathbf{R}_k(\theta) \right)^{-1} \mathbf{H}\mathbf{P}_{k|k-1}^\top. \quad (10.64)$$

A suitable choice of waveform will be made to minimize the estimated track error $\bar{\mathbf{P}}_{k|k}(\theta)$, or rather its determinant (or trace as it was originally used in [131, 130]), over all choices of θ :

$$\theta_k^* = \underset{\theta}{\operatorname{argmin}} \det(\bar{\mathbf{P}}_{k|k}(\theta)). \quad (10.65)$$

This choice of waveform is used for the next measurement.

5.2 Integrated Clutter Measure

In this section, we assume that an estimate is available of the clutter distribution in the range-Doppler plane. A simple method for doing this is given in [177]. Write $\gamma_k(t, f)$ for the power estimate of clutter at range (time) t and Doppler (frequency) f at discrete time k . We assume too, as in the work of Kershaw and Evans, that an estimate of the target state is available.

Our aim, for any potential waveform w with ambiguity A_w , is to calculate

$$F_k(w) = \left| \iint_{V_k} \left(\iint_{\mathbb{R}^2} A_w(t-t', f-f') \gamma_k(t', f') dt' df' \right) dt df \right|. \quad (10.66)$$

This is the integrated “spill-over” into the validation window of the sidelobes of the waveform due to the clutter. It makes sense to minimize this quantity over all waveforms. That is, the waveform optimizer we envisage will choose the optimal waveform at the k^{th} time instance by

$$w_k^* = \underset{w}{\operatorname{argmin}} F_k(w). \quad (10.67)$$

We remark that the implementation of this measure requires significant approximation that we shall discuss further.

Before going into the details of these approximations, we note that, because phase information from complex scatterers is usually uninformative, this waveform evaluator is replaced by

$$F_k(w) = \iint_{V_k} \left(\iint_{\mathbb{R}^2} |A_w(t - t', f - f')| |\gamma_k(t', f')| dt' df' \right) dt df. \quad (10.68)$$

This means that no account is taken of cancellation effects over multiple sidelobes, but these are so phase dependent as to be unlikely to exist robustly. By making this change from equation (10.66) to equation (10.68), we reduce the possibility of large waveform sidelobes (with the center of the ambiguity on the target) being in areas of high clutter.

5.3 Approximation of ICM

Assume that the clutter and target dynamics and that the measurements are modeled by linear equation, perturbed by Gaussian noise

$$\begin{aligned} \mathbf{x}_k &= \mathbf{F}\mathbf{x}_{k-1} + \nu_k; \\ \mathbf{y}_k &= \mathbf{H}\mathbf{x}_k + \omega_k. \end{aligned} \quad (10.69)$$

The noise ν_k in the dynamical model and the measurement noise ω_k are assumed to be stationary, and mutually independent, zero-mean with covariance, in the case of ν_k , \mathbf{Q} . The measurements \mathbf{y}_k are used to estimate the target and clutter probability densities recursively. At time k these densities are described by the mean $\mathbf{x}_{k|k-1}$ and covariance $\mathbf{P}_{k|k-1}$. This, in particular allows us to calculate the validation window V_k and the clutter power $|\gamma_k(t', f')|$ for Equation (10.68).

To make Equation (10.68) quickly calculable, we use the following approximations. First, the tracking is assumed to be performed by a PDA filter based tracker. Thus the estimate of target state is given by a single Gaussian distribution or by the mixture of Gaussians, all described by individual means and covariances. The validation window V_k , usually thought of as an ellipsoid, will also be approximated by a single Gaussian. Also the clutter power $|\gamma_k(t', f')|$ is given by its estimate at time k by discrete values for $\gamma = \{\gamma_1, \gamma_2, \dots, \gamma_m\}$, $t' = \{t'_1, t'_2, \dots, t'_m\}$ and $f' = \{f'_1, f'_2, \dots, f'_m\}$, where each γ_i is the power of the clutter in the i^{th} range-Doppler cell given by (t_i, f_i) .

The second approximation is the approximation of the absolute value of the ambiguity function A_w as a Gaussian mixture

$$|A_w(\mathbf{y})| = \sum_{i=-n}^n \alpha_i \mathcal{N}(\mathbf{y}; \delta_j, \mathbf{R}_j). \quad (10.70)$$

This can be obtained by a least squares fitting algorithm and is performed off-line once for each waveform. The finer the resolution of the fit, the more realistic the algorithm is. In the context of tracking this approximation is interpreted as an approximation of the measurement noise ω distribution by any of the Gaussians $\mathcal{N}(\mathbf{y}; \delta_j, \mathbf{R}_j)$ with mean δ_j and covariance \mathbf{R}_j . Note, that δ_0 is centered in the main lobe of the ambiguity function and is equal to zero and \mathbf{R}_0 is given by covariance \mathbf{R} as discussed in Section 3. Each δ_j for $j \in \{-n, \dots, -1, 1, \dots, n\}$ is centered in the j^{th} sidelobe of the ambiguity function. Note, that $2n+1$ is the number of “significant” sidelobes and is determined by a fitting algorithm. We assume that the measurements originated by the sidelobes can be detected with probability relative to α_j and to the power of the clutter scatterer γ , from which the radar measurement is originated. Thus, we postulate that at time k there could be up to $2n+1$ measurements originated from a single scatterer.

The target measurement probability density, used for calculating the validation window, is given by the mean $\mathbf{H}\mathbf{x}_{k|k-1}$ and covariance $\mathbf{S} = \mathbf{H}\mathbf{P}_{k|k-1}\mathbf{H}^T + \mathbf{R}_0$. This too is an approximation, justified since in a typical situation target power is much smaller than that of clutter and its detection in the sidelobes can be neglected.

The validation gate is given by an ellipsoid, for which probability that the measurement is in the gate is equal some P_G , usually greater than 0.99.

The probability of target detection in the main lobe is affected by the amount of “spill-over” from the clutter into the validation window; the less spill-over, the easier it is to detect the target.

With these approximations, Equation (10.68) is simplified as follows:

$$\begin{aligned} F_k(w) &= \iint_{\mathbb{R}^2} \chi(V_k) \left(\iint_{\mathbb{R}^2} |A_w(t-t', f-f')| |\gamma_k(t', f')| dt' df' \right) dt df \\ &\approx \iint_{\mathbb{R}^2} \mathcal{N}(\mathbf{y}; \hat{\mathbf{y}}, \mathbf{S}) \sum_{j=1}^m \gamma_j \sum_{i=-n}^n \alpha_i \mathcal{N}(\mathbf{y}; (t_j, f_j)^T + \delta_i, \mathbf{R}_i) d\mathbf{y} \\ &= \sum_{j=1}^m \sum_{i=-n}^n \gamma_j \alpha_i \mathcal{N}(\hat{\mathbf{y}}; (t_j, f_j)^T + \delta_i, \mathbf{S} + \mathbf{R}_i). \end{aligned} \quad (10.71)$$

In the second line, the characteristic function $\chi(V_k)$ is replaced by a Gaussian function. In the last line the formula for the Gaussian product is used. Without a loss of generality we can assume that the target power is 1, then γ_j in Equation (10.71) can be seen as the clutter-to-target power ratio and $F_k(w)^{-1}$ represents SNR in the validation gate. While this is not rigorously justified, it appears to be a reasonable assumption that works well in simulations. Assuming that the detection was obtained by a likelihood ratio test realized via, for example, a Neyman-Pearson type of detector, we write well-known formula for probability of target detection P_d as

$$P_d = 1 - \Phi(\Phi^{-1}(1 - P_f) - F_k(w)^{-1}), \quad (10.72)$$

where Φ is a standard normal probability distribution function, and P_f is desired probability of false alarm.

5.4 Simulation Results

We illustrate the ideas described in Section 5.3 with a simple example. Given a scenario with a constant velocity small target trajectory and random, large, slow moving clutter, two waveforms are scheduled using ICM and compared to the results with those for the same waveforms used in a non-scheduled way. The clutter-to-target power ratio γ is similar for all clutter and about 60 dB. The two waveforms are up-sweep and down-sweep linear frequency modulated (LFM) waveforms, their ambiguity function is approximated by a Gaussian mixture with five Gaussian terms as shown in Figure 10.2. The tracking is performed using an IPDA filter [178], as described earlier. Target measurement using i^{th} waveform is present with probability of detection given in Equation (10.72) with $P_f = 10^{-5}$.

Clutter measurements were generated by sidelobes of the ambiguity function with probability $\alpha_i, i = -n, \dots, -1, 1, \dots, n$ and with probability $\alpha_0 = 1$ by main lobe. For a given scenario, 1000 runs were performed and the results averaged over them. The cost as calculated in Equation (10.71). Position RMS error and probability of target existence are represented in Figure 10.3. We observe that, with scheduling, probability of target existence is maintained close to 1, in contrast to the two non-scheduled cases. A similar conclusion holds for RMS error: with scheduling the track error is smaller than without scheduling.

6. Scheduling of Beam Steering and Waveforms

In this section a scene consists of a collection of scatterers varying over range, Doppler and azimuth, that have been spread (in range-Doppler) by the




Figure 10.2. Approximation of ambiguity function by Gaussian mixture for up and down sweep waveforms.




Figure 10.3. Position root-mean-square error and probability of track existence.

waveform and (in azimuth) by the beam-shape. These ideas are described earlier. We investigate a radar system capable of rapid beam steering and of waveform switching. The transmit waveform is chosen from a small library of such. The operational requirement of the radar is to track a number of maneuvering targets while performing surveillance for new potential targets. Tracking is accomplished by means of an LMIPDA (Linear Multi-target Integrated Probabilistic Data Association) tracker as described in [176], which is a variant of the IPDA tracker described in Section 4 designed to handle many targets efficiently. Interacting multiple models (IMM) is used to model maneuvering targets in the tracker. LMIPDA provides a probability of track existence, permitting a “track-before-detect” technique to be adopted. “False alarm” tracks are maintained until the probability of track existence falls below a threshold. In derivations of formulae in this section the SNC case (described in Section 5.1) is assumed. LMIPDA-IMM is a recursive algorithm combining multi-target data association (LMIPDA) with maneuvering target state estimation implemented using IMM.

The aim is to maintain the tracks of the existing targets to within a specified accuracy as determined by the absolute value of the track error covariance matrix. However, this has to be done within the time available given that a full scan has to be performed within a prescribed interval. In this section we give an algorithm for scheduling revisits to measure the targets while maintaining surveillance.

6.1 Tracking of Multiple Maneuvering Targets

We assume that a radar system tracks T targets where T is a random variable $0 \leq T \leq T_0$ and the t^{th} target is in state \mathbf{x}_k^t at epoch k . In addition the radar undertakes surveillance to discover new targets. This surveillance is assumed to require a certain length of time, say τ_{scan} within every interval of length τ_{total} . The remainder of the time is spent measuring targets being tracked. We aim to schedule revisit times to targets within these constraints.

For IMM-based tracking algorithms, for example, the target trajectory is approximated by an average over a finite number of given dynamic models. In this context, we assume that the dynamical models are independent of the target and associated to each is a corresponding state propagation matrix \mathbf{F}_m ($m = 1, 2, \dots, M$). The recursion for state transition is

$$\mathbf{x}^t(k) = \mathbf{F}_m(k)\mathbf{x}^t(k-1) + \nu_m^t(k), \quad (10.73)$$

where the index m is a possible value of a random variable $M(k)$, the *dynamical model* which takes any discrete value $[1, 2, \dots, M]$. Process noise

$\nu_1^t(k), \dots, \nu_M^t(k)$ is Gaussian, depends on both target and dynamical model and is independent between different values of each of these indices. The covariance matrix of $\nu_m^t(k)$ is denoted by $\mathbf{Q}_m^t(k)$.

In the tracker, the dynamical model of the t^{th} target $M^t(k)$ is assumed to evolve as a Markov Chain with given transition probabilities, denoted by

$$\pi_{m,\ell}^t = P\{M^t(k) = m | M^t(k-1) = \ell\}; j, \ell \in [1, \dots, M]. \quad (10.74)$$

It is assumed that N different *measurement modes* are available for each target, each given by a measurement matrix \mathbf{H}_n^t $n = 1, 2, \dots, N$:

$$\mathbf{y}^t(k) = \mathbf{H}_n^t(k)\mathbf{x}^t(k) + \omega_n^t(k) \quad (10.75)$$

where here $\mathbf{y}^t(k)$ is the measurement to be obtained from the t^{th} target at time k , $\omega_n^t(k)$ is the measurement noise, and $n = n(k)$ is a control variable for the measurement mode. The variable $\tilde{t} = \tilde{t}(k)$ represents the choice of target to which the beam is steered at the k^{th} epoch. The measurement noise $\omega_1^t(k), \dots, \omega_N^t(k)$ are zero mean white and uncorrelated Gaussian noise sequences with the covariance matrix of $\omega_n^t(k)$ denoted by $\mathbf{R}_n^t(k)$.

The waveforms impinge on the measurement process through the covariance matrix of the noise $\omega_n^t(k)$ as described in Section 3.

The choice of measurement is made using the control variable $n(k)$. In this example two choices are made at each epoch, the target to be measured and the waveform used. More than one target may be in the beam and then measurements of each target will be updated using the LMIPDA-IMM algorithm.

6.2 Scheduling

As already stated, at each epoch a target track and a beam direction have to be selected. The scheduler has a list $\Delta = \{\delta_1, \delta_2, \dots, \delta_K\}$ of “revisit intervals”. Each of the numbers δ_k is a number of epochs representing the possible times between measurements of any of the existing targets. It is assumed for the purposes of scheduling and tracking that during any of these revisit intervals the target dynamics do not change.

In order to determine which target to measure and which waveform to use, for each existing target and each waveform the track error covariance $P_{k-1|k-1}^t$ is propagated forward using the Kalman update equations and assuming each of the different potential revisit intervals in the list Δ in the dynamics. In the absence of measurements the best we can do is to use the current knowledge to predict forward and update the covariance matrix, dynamic model probability

density and probability of track existence. The algorithms is now modified as follows. *Forward prediction* is performed separately for each dynamical model. Because the dynamics of the target depends on the revisit time $\delta \in \Delta$ this calculations are performed for each revisit time. *Covariance update* is normally done with the data, but since we are interested in choosing the best sensor mode at this stage, the following calculations are required. If the target does not exists there will be no measurements originating from the target and the error covariance matrix is equal to the prior covariance matrix, if the target exists, is detected, and the measurement is received then the error covariance matrix is updated using the Kalman update equations. The expected covariance update is calculated using Bayes rule, namely,

$$\mathbf{P}_{k|k}(j, \delta) = (I - \psi_{k|k-1} P_D P_G \mathbf{K}(\phi, \delta) \mathbf{H}) \mathbf{P}_{k|k-1}(j, \delta), \quad (10.76)$$

where $\psi_{k|k-1}$ is the *a priori* probability of track existence, $P_D P_G$ is the probability that target is detected and its measurement is validated. $\mathbf{K}(\phi, \delta)$ is a Kalman gain calculated for each sensor mode; that is, for the waveform ϕ and revisit time δ . Both ϕ and δ take discrete values from waveform library and revisit time set Δ .

$$\mathbf{K}(\phi, \delta) = \mathbf{P}_{k|k-1}(j, \delta) \mathbf{H} \mathbf{S}^{-1}(\phi), \quad (10.77)$$

where \mathbf{S} is innovation covariance matrix, calculated as usual

$$\mathbf{S}(\phi) = \mathbf{H} \mathbf{P}_{k|k-1} \mathbf{H}^T + \mathbf{R}_\phi.$$

The above calculations are made for all combinations of revisit times in Δ and waveforms in the library. In considering a non-myopic approach, the number of combinations grows exponentially in the number of steps ahead, and soon becomes impractical for implementation. Having obtained the error covariance matrix for all possible combinations of sensor modes, the optimal sensor mode (waveform) is then chosen for each target to be the one which gives the longest re-visit time, while constraining the absolute value of the determinant of the error covariance matrix to be smaller than the prescribed upper limit K or the smallest value of the determinant, to ensure that a solution to the following equation exist. In mathematical language the objective is to obtain

$$\phi, \delta = \arg \max \Delta, \text{ subject to } |\mathbf{P}_{k|k}(\phi, \delta)| \leq \max\{K, \min_{\phi, \delta} |\mathbf{P}_{k|k}(\phi, \delta)|\}. \quad (10.78)$$

Scheduling is then done to permit a full scan over the prescribed scan period while also satisfying the constraints imposed by the revisit times obtained by the sensor scheduler. Once a target is measured, its revisit time is re-calculated.

The solution to the above equation is not necessary unique for ϕ . In cases when there is more than one ϕ , the waveform that gives the smallest determinant of the error covariance matrix is chosen.

For the case of N -step-ahead scheduling, the revisit times and waveforms are calculated while the target states are propagated forward over N measurements. The cost function is the absolute value of the determinant of the track error covariance after the N^{th} measurement. Only the first of these measurements is done before the revisit calculation is done again for that target, so the second may never be implemented.

6.3 Simulation results

Here we demonstrate the effects, in simulation, of scheduling as described in Sections 4, 5 and 6. Specifically, we compare random choice of waveform with a scanning beam against one-step-ahead and two-step-ahead beam and waveform scheduling. All three simulations were performed 100 times on the same scenario. In the first case, measurements were taken at each scan with no further measurements beyond the scan measurements permitted. In these experiments we used a small waveform library consisting of three waveforms: an up-sweep chirp, a down-sweep chirp and an unmodulated pulse. In the unscheduled case, waveforms were chosen randomly from this library. The simulated scene corresponded to a surveillance area of 15 km by 15 km. The scene contained two maneuvering land targets in stationary land clutter, which had small random Doppler to simulate movement of vegetation in wind. While the level of fidelity of the clutter is low, it is sufficient to demonstrate the principles of scheduling. The number of clutter measurements at each epoch was generated by samples from a Poisson distribution with mean ~ 5 per scan per square kilometer. Target measurements were produced with probability of detection 0.9. The target state x^t consisted of target range, target range rate and target azimuth. The targets were performing the following maneuvers: constant velocity, constant acceleration, constant deceleration and coordinated turns with constant angular velocity. In the scheduling cases, surveillance time used approximately 80% of each scan period; the remaining 20% was allocated, as described above, to maintaining tracks of existing targets.

The outcome of these experiments suggests that in the presence of clutter tracking performance can be improved with scheduling and even more with multiple-step-ahead scheduling as opposed to one-step-ahead. The results are represented in Figure 10.4.

It should be observed that in Figure 10.4 the RMS error was considerably worse, especially during the early part of the simulation, for the unscheduled

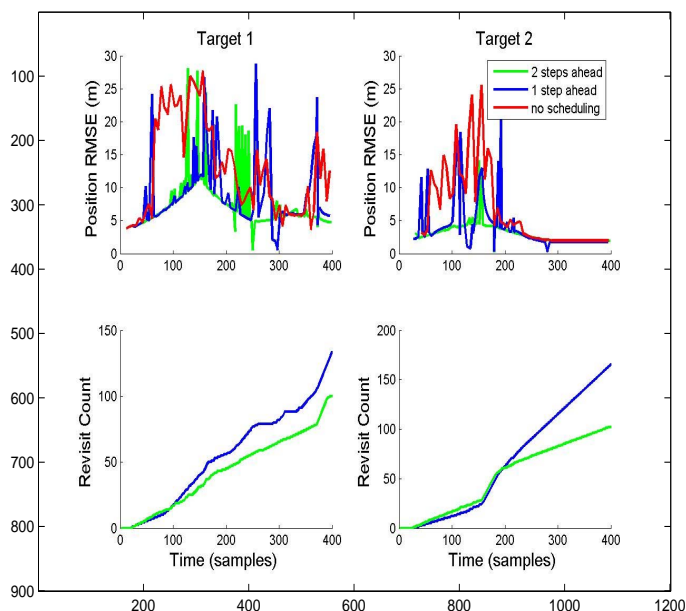


Figure 10.4. Root-mean-square error (RMSE) and revisit count for one-step-ahead versus two-step-ahead beam and waveform scheduling.

case. In fact the RMS error in the unscheduled case is larger immediately after significant maneuvers as can be expected. Of course, in this case the revisit time is fixed and is not plotted in the second subplot. One observes that, for the two-step-ahead case, tracking accuracy is improved (top plots) slightly over the one-step-ahead case, but with a significant reduction in revisits necessary to maintain those tracks.

7. Waveform Libraries

We have discussed how to estimate the effectiveness of individual waveforms in Section 5. We envisage that the choice of waveforms will be from a relatively small library. Others have considered the design of waveforms by the choice of parameters in a generic waveform type. For example, see [215], one could choose waveforms of the form

$$w_{\alpha,\beta}(t) = e^{2\pi i \alpha t^\beta}. \quad (10.79)$$

This has the form of a *power frequency modulated* waveform. Choice of different values of α and β provides waveforms with different range and Doppler properties. This is effectively an infinite library. Since our interest is mainly

in understanding the concepts of sensor management, we shall restrict attention to the situation where the number of waveforms is finite. In the context of radar, where the time between pulses is very small, computational issues are crucial. Scheduling problems are notoriously computationally intensive, and so it makes sense to reduce that problem by minimizing the number of waveforms from which to make the choice. Here we discuss how to choose libraries, bearing in mind computational constraints.

In order to choose between different radar modes, or more specifically waveforms, at a given epoch, a method is needed to measure the effectiveness of such a mode in a given context. There are three issues associated with this *cost function* assignment that need to be addressed. The first is the apparently obvious statement that optimization of the cost should, in principle, produce a desired operational outcome. However, it is often the case that, for reasons of feasibility, cost functions represent only part of the total operational requirement. For example, in the context of tracking a target such as a missile, it is fairly normal to measure the cost in terms of an estimate of the track accuracy such as the estimated mean-square error of the estimated target position. On the other hand, from an operational perspective, at least for some of the engagement between the missile and its target, it may be much more important that the track be maintained; i.e., that the missile be not lost by the radar. While it would appear that track error is in some sense related to track maintenance, the optimal choice of radar mode to reduce track error may be different from that required to maintain track. For real radar systems, we propose that cost functions be based on the precise operational needs at the time with the objective of maximizing the probability of achieving the operational goals.

The second issue is that the cost function is a function both of the library of radar modes and of the environment. The information available about the environment is available in two forms: information that is acquired through exogenous sources (perhaps prior to the current collection) and information that has been acquired in the current sequence of measurements. The former are extremely important, especially in the context of radars observing an urban situation where the effects on the return due to buildings, fences, and other aspects of the built environment can be significant. We shall, nevertheless, ignore them in the current discussion. We focus only on the information that is acquired during the deployment of the radar system. Of course if computational constraints were not at issue, we would base cost functions on all the data that has been acquired so far. But this is infeasible, except in rare cases. Some processing has to be done, itself limited by computational considerations, to extract key features that can be fed to the cost function. In much of the literature on scheduling of waveforms for tracking, this “extracted information” is just the track error covariance.

The third and final issue we wish to address, would seem to contradict the first. There we have espoused the idea that cost functions should be operationally significant. However, this means that theoretical developments will be fragmented according to the cost function used. We propose therefore, for the purposes of developing a theory of waveform and other mode libraries, a generic cost function that has some relevance to operational costs.

One approach to this, at least in the context of tracking, is to calculate the *expected information* to be obtained from a measurement with a given waveform, given the current state of knowledge of the scene. In this context, it is expressed by knowledge of the target parameters – typically velocity and position. This is defined as the mutual information between the target variable (range and Doppler) and the processed (e.g., matched filtered) radar return resulting from the use of the waveform.

The *utility* of a waveform library is defined by averaging over a distribution on the possible state covariance matrices, the maximum of this expected information over all waveforms. A subset of a library is *irrelevant* if it does not contribute to the utility in the sense that the library has the same utility with or without that subset. It is possible to use this concept to obtain parsimonious libraries involving LFM waveforms and more general “chirped” collections of waveforms.

We point out that it is usually difficult to find an explicit and usable expression for this distribution of state covariance matrices in practice. Nonetheless the very nature of this formalism permits us to make some general statements about waveform libraries.

We use the model described in Section 3. In the context of our discussion in this section, we represent the measurement obtained using the waveform ϕ as a Gaussian measurement with covariance \mathbf{R}_ϕ . The current state of the system is represented by the state covariance matrix \mathbf{P} . Of course, the estimated position and velocity of the target is also important for the tracking function of the radar. But in this context they play no role in the choice of waveforms. In a clutter-rich (and varying) scenario, the estimate of the target parameters will clearly play a more important role. The *expected information* obtained from a measurement with such a waveform, given the current state of knowledge of the target, is

$$I(X; Y) = \log \det(\mathbf{I} + \mathbf{R}_\phi^{-1} \mathbf{P}). \quad (10.80)$$

This is the mutual information between the target variables (range and Doppler) X and the processed radar return Y resulting from the use of the waveform ϕ . \mathbf{I} is the identity matrix. We use this expected information as the *measure of ef-*

fectiveness of the waveform ϕ in this context. The more information we extract from the situation the better.

We assume a knowledge of the possible state covariances \mathbf{P} generated by the tracking system. This knowledge is statistical and is represented by a probability distribution $F(\mathbf{P})$ over the space of all positive definite matrices. This distribution will be a function of the previous choices of waveform since their Fisher matrices play a part in the calculation of this distribution, so that there is an inherent circularity in the definition. Indeed every update of the tracker essentially “convolves” the distribution $F_n(\mathbf{P})$ of the covariances at time n (updated by the dynamics) with the system noise and the measurement noise. If the initial distribution $F_0(\mathbf{P})$ is sufficiently “flat” (invariant in some sense), it may be that such “convolutions” will have no effect. For the purposes of this discussion, we assume that $F(\mathbf{P})$ is a fixed distribution.

We define the *utility* of a waveform library $\mathcal{L} \subset L^2(\mathbb{R})$, with respect to a distribution F , to be

$$G_F(\mathcal{L}) = \int_{\mathbf{P} > 0} \max_{\phi \in \mathcal{L}} \log \det(\mathbf{I} + \mathbf{R}_\phi^{-1} \mathbf{P}) dF(\mathbf{P}). \quad (10.81)$$

Thus we have assumed that the optimal waveform is chosen in accordance with the measure of effectiveness defined in equation (10.80) and have averaged this over all possible current states, as represented by the covariance matrices \mathbf{P} and in accordance with their distribution $F(\mathbf{P})$.

We consider two libraries \mathcal{L} and \mathcal{L}' to be *weakly equivalent*, with respect to the distribution F , if $G_F(\mathcal{L}) = G_F(\mathcal{L}')$, and *strongly equivalent* if $G_F(\mathcal{L}) = G_F(\mathcal{L}')$ for all F .

We call the subset $\mathcal{S} \subset \mathcal{L}$ *weakly irrelevant* with respect to the distribution F if $G_F(\mathcal{L} \setminus \mathcal{S}) = G_F(\mathcal{L})$, and *strongly irrelevant* if $G_F(\mathcal{L} \setminus \mathcal{S}) = G_F(\mathcal{L})$ for all F . Here, \setminus denotes the usual set difference. In what follows we will work in receiver coordinates, i.e., treat \mathbf{T} above as \mathbf{I} . This amounts to a change in parameterization of the positive definite matrices in the integral in (10.81). Again, an invariance assumption (under the dynamics of the Gauss-Markov system) on the distribution $F(\mathbf{P})$ will make this choice of coordinates irrelevant.

While this theory continues to be explored, we discuss here just one result. As we said in section 2.1, LFM “chirps” are important for radar applications.

7.1 LFM Waveform Library

We investigate an LFM (“chirp”) waveform library. In this case the library consists of

$$\mathcal{L}_{\text{chirp}} = \{\phi_0(t) \exp(i\lambda t^2/2) \mid \lambda_{\min} \leq \lambda \leq \lambda_{\max}\} \quad (10.82)$$

where $\phi_0(t)$ is an unmodulated pulse, λ_{\min} and λ_{\max} are the minimum and maximum chirp rates supported by the radar. For this library the corresponding measurement covariance matrices are of the form [181, 131]

$$\mathbf{R}_\phi = S(\lambda)\mathbf{R}_{\phi_0}S(\lambda)^\top, \quad (10.83)$$

where

$$S(\lambda) = \begin{pmatrix} 1 & 0 \\ \lambda & 1 \end{pmatrix}. \quad (10.84)$$

We have been able to prove (see [226]) that this library is strongly equivalent to

$$\mathcal{L}'_{\text{chirp}} = \{\phi_0(t) \exp(i\lambda_{\min}t^2/2), \phi_0(t) \exp(i\lambda_{\max}t^2/2)\}. \quad (10.85)$$

That is, we do just as well if we keep only the LFMs with the minimum and maximum rates. In the case of the LFM waveform library, the error covariance matrices (despite a popular assumption [184]) are not rotations of each other, but the results of shearing transformations. For \mathbf{R}_{ϕ_0} a diagonal matrix with ρ_1, ρ_2 on the diagonal, direct computation yields the following expression for the mutual information $I(X; Y)$:

$$I(X; Y) = \frac{P_{11}}{\rho_2} \lambda^2 - 2 \frac{P_{12}}{\rho_2} \lambda + \frac{|\mathbf{P}|}{|\mathbf{R}_0|} + 1 + \frac{P_{11}}{\rho_1} + \frac{P_{22}}{\rho_2}. \quad (10.86)$$

This is a quadratic in λ with positive second derivative since \mathbf{P} and \mathbf{R}_{ϕ_0} are both positive definite, and therefore achieves its maximum at the end points; i.e., at maximum or minimum allowed sweep rate. The optimal sweep rate is chosen to be

$$\lambda_\phi = \begin{cases} \lambda_{\max}, & \text{if } \lambda_{\min} + \lambda_{\max} > \frac{P_{12}}{P_{11}} \\ \lambda_{\min}, & \text{otherwise.} \end{cases} \quad (10.87)$$

7.2 LFM-Rotation Library

For the library under consideration here, we start with an unmodulated waveform ϕ_0 and allow both the “chirping” transformations (10.82) and the fractional Fourier transformations, i.e.,

$$\mathcal{L}_{\text{FrFT}} = \{\exp(i\theta(\mathbf{t}^2 + \mathbf{f}^2)/2)\phi_0 \mid \theta \in \Theta\}, \quad (10.88)$$

That is, we consider all transformations of the following form.

$$\mathcal{L}_{\text{FrFT}} = \{\exp(i\theta(\mathbf{t}^2 + \mathbf{f}^2)/2) \exp(i\lambda\mathbf{t}^2/2)\phi_0 \mid \lambda_{\min} \leq \lambda \leq \lambda_{\max}, \theta \in \Theta\} \quad (10.89)$$

where the set Θ is chosen so as not to violate the bandwidth constraints of the radar, and \mathbf{f} is the operator on $L^2(\mathbb{R})$ defined by

$$\mathbf{f}\phi(t) = i\frac{d}{dt}\phi(t), \quad (10.90)$$

Note that \mathbf{f} and \mathbf{t} commute up to an extra additive term (the ‘‘canonical commutation relations’’). To be precise,

$$[\mathbf{t}, \mathbf{f}] = \mathbf{t}\mathbf{f} - \mathbf{f}\mathbf{t} = -i\mathbf{I}. \quad (10.91)$$

For this library the corresponding measurement covariance matrices are given by (10.83) with

$$S(\theta, \lambda) = \begin{pmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{pmatrix} \begin{pmatrix} 1 & 0 \\ \lambda & 1 \end{pmatrix}. \quad (10.92)$$

In the case of a finite number of waveforms in the library, we observe that the utility of the rotation library improves with the number of waveforms in the library. There exists a unique θ that maximizes the mutual information $I(X; Y)$ and, similar to the pure chirp library case, maximum and minimum allowed chirp rate λ

$$\mathcal{L}'_{\text{FrFT-chirp}} = \{\exp(i\theta(\mathbf{t}^2 + \mathbf{f}^2)/2) \exp(i\lambda\mathbf{t}^2/2)\phi_0 \mid \lambda \in \{\lambda_{\min}, \lambda_{\max}\}, \theta \in \Theta\} \quad (10.93)$$

is strongly equivalent to $\mathcal{L}_{\text{FrFT-chirp}}$.

8. Conclusion

Our aim in this chapter has been to illustrate some of the ideas of sensor scheduling in the context of radar. This has been a highly subjective view of the issues and ideas, based to a large extent on our own ideas and work on this subject. To understand the sensor scheduling problem for radar, and elsewhere, it is important to be aware of the various functionalities of the sensors and their capabilities. To do this for radar, we have described the way in which a pulse-Doppler radar works, at least roughly. This has given some indication of what level of flexibility is available to the designers of such a system, and what aspects might be dynamically and adaptively tunable. We remark that, in reality, radar engineers have to work at a much finer level of detail. Issues

associated with the limitations of the electronics play a significant role in the design of a radar, and robustness is a key objective.

We have discussed some relatively simple techniques in radar scheduling. Our focus has been to show, in a few areas, that scheduling is useful. It will improve detection and tracking performance and reduce revisit time to tracked targets while maximizing surveillance time. Some recent work by Sira et al. [214], has shown, for example, that scheduling will significantly improve detection of small targets in sea clutter. Nehorai [242] has also shown similar results.

An important issue that has been largely neglected in the sensor scheduling community is that of library design. In whatever sensor scheduling context, the library of available sensor modes will be a key design component in the system. What principles should guide the system designer in this context? We have touched on this subject albeit briefly and only in the context of radar waveforms in Section 7. A closely related aspect of the design of a scheduling system, which has also not received much attention in the literature, is that of the design of measures of effectiveness for the various sensor modes, and in the radar context for waveforms in particular. Our belief is that ultimately these measures will be based on operational criteria, but that these criteria will drive “local” more easily computable cost functions that are more closely tied to the system. We have initiated a discussion of measures of effectiveness within the context of waveforms for radar in Section 5.

This is an embryonic subject and much more needs to be done. In a sense, all of the work done so far really addresses idealizations of the true problems. More work is needed, in particular, on the problems associated with implementation. The short time between pulses in a radar system make it very difficult to schedule on a pulse-to-pulse basis. The calculations associated with, in particular, clutter mapping which involve large amounts of data, are probably too computationally intensive to be done on this time scale. Compressive sensing ideas [71] may play a role in reducing this complexity.

Work is needed too on the potential gains of non-myopic scheduling in this context. A real system will be required to work across multiple time scales, integrating multiple tracking functions, and scheduling over many pulses. Scheduling will be used both to optimize allocation of the resource between different operational functions, as we have discussed in Section 6, and to improve operational performance by, for example, waveform and beam-shape scheduling. For military functions, issues such as jamming are important, and scheduling is destined to play a role here too. Ultimately, we anticipate that, against a sophisticated adversary, game-theoretic techniques will be used to drive radar scheduling to mitigate jamming.

Radar scheduling is replete with problems, some of which we have touched on, and many of which we have not, and very few solutions. It represents an exciting and challenging area of research that will drive much of the research activity in radar signal processing over the next few years.

References

- [1] R. Agrawal, M. V. Hegde, and D. Teneketzis. Asymptotically efficient adaptive allocation rules for the multiarmed bandit problem with switching cost. *IEEE Transactions on Automatic Control*, 33:899–906, 1988.
- [2] R. Agrawal, M. V. Hegde, and D. Teneketzis. Multi-armed bandits with multiple plays and switching cost. *Stochastics and Stochastic Reports*, 29:437–459, 1990.
- [3] R. Agrawal and D. Teneketzis. Certainty equivalence control with forcing: revisited. *Systems and Control Letters*, 13:405–412, 1989.
- [4] R. Agrawal, D. Teneketzis, and V. Anantharam. Asymptotically efficient adaptive allocation schemes for controlled Markov chains: finite parameter space. *IEEE Transactions on Automatic Control*, 34:1249–1259, 1989.
- [5] R. Agrawal, D. Teneketzis, and V. Anantharam. Asymptotically efficient adaptive control schemes for controlled I.I.D. processes: finite parameter space. *IEEE Transactions on Automatic Control*, 34:258–267, 1989.
- [6] S.-I. Amari. *Methods of Information Geometry*. American Mathematical Society - Oxford University Press, Providence, RI, 2000.
- [7] V. Anantharam, P. Varaiya, and J. Walrand. Asymptotically efficient allocation rules for the multiarmed bandit problem with multiple plays — part I: I.I.D. rewards. *IEEE Transactions on Automatic Control*, 32:968–976, 1987.
- [8] V. Anantharam, P. Varaiya, and J. Walrand. Asymptotically efficient allocation rules for the multiarmed bandit problem with multiple plays — part II: Markovian rewards. *IEEE Transactions on Automatic Control*, 32:977–982, 1987.

- [9] P. S. Ansell, K. D. Glazebrook, J. Niño-Mora, and M. O’Keefe. Whittle’s index policy for a multi-class queueing system with convex holding costs. *Mathematical Methods of Operations Research*, 57:21–39, 2003.
- [10] M. S. Arulampalam, S. Maskell, N. Gordon, and T. Clapp. A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking. *IEEE Transactions on Signal Processing*, 50:174–188, 2002.
- [11] M. Asawa and D. Teneketzis. Multi-armed bandits with switching penalties. *IEEE Transactions on Automatic Control*, 41:328–348, 1996.
- [12] J. Banks and R. Sundaram. Switching costs and the Gittins index. *Econometrica*, 62:687–694, 1994.
- [13] Y. Bar-Shalom. *Multitarget Multisensor Tracking: Advanced Applications*. Artech House, Boston, MA, 1990.
- [14] Y. Bar-Shalom and W. D. Blair. *Multitarget-Multisensor Tracking: Applications and Advances, Volume III*. Artech House, Boston, MA, 2000.
- [15] A. R. Barron. Complexity regularization with application to artificial neural networks. In *Nonparametric Functional Estimation and Related Topics*, pages 561–576. Kluwer Academic Publishers, 1991.
- [16] A. G. Barto, W. Powell, and J. Si, editors. *Learning and Approximate Dynamic Programming*. IEEE Press, New York, NY, 2004.
- [17] M. Beckmann. *Dynamic Programming of Economic Decisions*. Springer-Verlag, New York, NY, 1968.
- [18] R. Bellman. On the theory of dynamic programming. *Proceedings of the National Academy of Sciences*, 38:716–719, 1952.
- [19] R. Bellman. A problem in the sequential design of experiments. *Sankhya*, 16:221–229, 1956.
- [20] R. Bellman. *Adaptive Control Processes: a Guided Tour*. Princeton University Press, Princeton, NJ, 1961.
- [21] R. Bellman and S. Dreyfus. *Applied Dynamic Programming*. Princeton University Press, Princeton, NJ, 1962.
- [22] D. A. Berry and B. Fristedt. *Bandit problems: sequential allocation of experiments*. Chapman and Hall, 1985.
- [23] D. P. Bertsekas. *Dynamic Programming and Optimal Control*, volume 1. Athena Scientific, 1995.

- [24] D. P. Bertsekas. *Dynamic Programming and Optimal Control*, volume 2. Athena Scientific, 1995.
- [25] D. P. Bertsekas. *Dynamic Programming and Optimal Control, Vols. I-II*. Athena Scientific, Belmont, MA, 3rd edition, 2005.
- [26] D. P. Bertsekas and D. A. Castañón. Rollout algorithms for stochastic scheduling. *Heuristics*, 5:89–108, 1999.
- [27] D. P. Bertsekas and S. E. Shreve. *Stochastic Optimal Control: The Discrete Time Case*, volume 1. Academic Press, 1978.
- [28] D. P. Bertsekas and J. N. Tsitsiklis. *Neuro-Dynamic Programming*. Athena Scientific, Belmont, MA, 1996.
- [29] D. Bertsimas and J. Niño-Mora. Conservation laws, extended polymatroids and multiarmed bandit problems; a polyhedral approach to indexable systems. *Mathematics of Operations Research*, 21:257–306, 1996.
- [30] D. Bertsimas and J. Niño-Mora. Restless bandits, linear programming relaxations, and a primal-dual index heuristic. *Operations Research*, 48:80–90, 2000.
- [31] D. Bertsimas, I. C. Paschalidis, and J. N. Tsitsiklis. Branching bandits and Klimov’s problem: achievable region and side constraints. *IEEE Transactions on Automatic Control*, 40:2063–2075, 1995.
- [32] P. Billingsley. *Probability and Measure*. John Wiley and Sons, New York, NY, 1995.
- [33] S. S. Blackman. *Multiple-Target Tracking with Radar Applications*. Artech House, Boston, MA, 1986.
- [34] D. Blackwell. Discrete dynamic programming. *Annals of Mathematical Statistics*, 33:719–726, 1962.
- [35] D. Blackwell. Discounted dynamic programming. *Annals of Mathematical Statistics*, 36:226–235, 1965.
- [36] W. D. Blair and M. Brandt-Pearce. Unresolved Rayleigh target detection using monopulse measurements. *IEEE Transactions on Aerospace and Electronic Systems*, 34:543–552, 1998.
- [37] G. Blanchard and D. Geman. Hierarchical testing designs for pattern recognition. *Annals of Statistics*, 33(3):1155–1202, 2005.
- [38] D. Blatt and A. O. Hero. From weighted classification to policy search. In *Neural Information Processing Symposium*, volume 18, pages 139–146, 2005.

- [39] D. Blatt and A. O. Hero. Optimal sensor scheduling via classification reduction of policy search (CROPS). In *International Conference on Automated Planning and Scheduling*, 2006.
- [40] H. A. P. Blom and E. A. Bloem. Joint IMM-PDA particle filter. In *International Conference on Information Fusion*, 2003.
- [41] A. G. B. S. J. Bradtke and S. P. Singh. Learning to act using real-time dynamic programming. *Artificial Intelligence*, 72:81–138, 1995.
- [42] L. Breiman, J. Friedman, R. Olshen, and C. J. Stone. *Classification and Regression Trees*. Wadsworth, Belmont, CA, 1983.
- [43] M. V. Burnashev and K. S. Zigangirov. An interval estimation problem for controlled observations. *Problems in Information Transmission*, 10:223–231, 1974. Translated from *Problemy Peredachi Informatsii*, 10(3):51–61, July-September, 1974.
- [44] L. Carin, H. Yu, Y. Dalichaouch, A. R. Perry, P. V. Czipott, and C. E. Baum. On the wideband EMI response of a rotationally symmetric permeable and conducting target. *IEEE Transactions on Geoscience and Remote Sensing*, 39:1206–1213, June 2001.
- [45] A. R. Cassandra. *Exact and Approximate Algorithms for Partially Observable Markov Decision Processes*. PhD thesis, Department of Computer Science, Brown University, 1998.
- [46] A. R. Cassandra, M. L. Littman, and L. P. Kaelbling. Incremental pruning: A simple, fast, exact method for partially observable Markov decision processes. In *Uncertainty in Artificial Intelligence*, 1997.
- [47] D. A. Castañón. Approximate dynamic programming for sensor management. In *IEEE Conference on Decision and Control*, pages 1202–1207. IEEE, 1997.
- [48] D. A. Castañón. A lower bound on adaptive sensor management performance for classification. In *IEEE Conference on Decision and Control*. IEEE, 2005.
- [49] D. A. Castañón and J. M. Wohletz. Model predictive control for dynamic unreliable resource allocation. In *IEEE Conference on Decision and Control*, volume 4, pages 3754–3759. IEEE, 2002.
- [50] R. Castro, R. Willett, and R. Nowak. Coarse-to-fine manifold learning. In *IEEE International Conference on Acoustics, Speech and Signal Processing*, May, Montreal, Canada, 2004.

- [51] R. Castro, R. Willett, and R. Nowak. Faster rates in regression via active learning. In *Neural Information Processing Systems*, 2005.
- [52] R. Castro, R. Willett, and R. Nowak. Faster rates in regression via active learning. Technical report, University of Wisconsin, Madison, October 2005. ECE-05-3 Technical Report.
- [53] E. Çinlar. *Introduction to Stochastic Processes*. Prentice-Hall, Englewood Cliffs, NJ, 1975.
- [54] H. S. Chang, R. L. Givan, and E. K. P. Chong. Parallel rollout for online solution of partially observable Markov decision processes. *Discrete Event Dynamic Systems*, 14:309–341, 2004.
- [55] H. Chernoff. Sequential design of experiments. *Annals of Mathematical Statistics*, 30:755–770, 1959.
- [56] H. Chernoff. *Sequential Analysis and Optimal Design*. SIAM, 1972.
- [57] A. Chhetri, D. Morrell, and A. Papandreou-Suppappola. Efficient search strategies for non-myopic sensor scheduling in target tracking. In *Asilomar Conference on Signals, Systems, and Computers*, 2004.
- [58] E. K. P. Chong, R. L. Givan, and H. S. Chang. A framework for simulation-based network control via hindsight optimization. In *IEEE Conference on Decision and Control*, pages 1433–1438, 2000.
- [59] Y. S. Chow, H. Robins, and D. Siegmund. *Great Expectations: The theory of Optimal Stopping*. Houghton Mifflin Company, Boston, MA, 1971.
- [60] D. Cochran. Waveform-agile sensing: opportunities and challenges. In *IEEE International Conference on Acoustics, Speech, and Signal Processing*, pages 877–880, Philadelphia, PA, 2005.
- [61] D. Cochran, D. Sinno, and A. Clausen. Source detection and localization using a multi-mode detector: a Bayesian approach. In *IEEE International Conference on Acoustics, Speech, and Signal Processing*, pages 1173–1176, Phoenix, AZ, 1999.
- [62] D. A. Cohn, Z. Ghahramani, and M. I. Jordan. Active learning with statistical models. *Advances in Neural Information Processing Systems*, 7:705–712, 1995.
- [63] D. A. Cohn, Z. Ghahramani, and M. I. Jordan. Active learning with statistical models. *Journal of Artificial Intelligence Research*, pages 129–145, 1996.

- [64] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. John Wiley and Sons, New York, NY, 1991.
- [65] N. Cristianini and J. Shawe-Taylor. *Support Vector Machines and Other Kernel Based Learning Methods*. Cambridge University Press, Cambridge, UK, 2000.
- [66] I. Csiszár. Information-type measures of divergence of probability distributions and indirect observations. *Studia Sci. Math. Hung.*, 2:299–318, 1967.
- [67] I. Csiszár and J. Körner. *Information Theory: Coding Theorems for Discrete Memoryless Systems*. Academic Press, Orlando FL, 1981.
- [68] M. H. DeGroot. *Optimal Statistical Decisions*. McGraw Hill, 1970.
- [69] C. Dellacherie and P. A. Meyer. *Probabilities and Potential B: Theory of Martingales*. North-Holland, Amsterdam, 1982.
- [70] E. V. Denardo. *Dynamic Programming Models and Applications*. Prentice-Hall, Englewood Cliffs, NJ, 1982.
- [71] D. Donoho. Compressed sensing. *IEEE Trans. on Information Theory*, 52(4):1289–1306, April 2006.
- [72] A. Doucet. On sequential Monte Carlo methods for Bayesian filtering. Uk. tech. rep., Dept. Eng. Univ. Cambridge, 1998.
- [73] A. Doucet, N. de Freitas, and N. Gordon. *Sequential Monte Carlo Methods in Practice*. Springer Publishing, New York, NY, 2001.
- [74] A. Doucet, B.-N. Vo, C. Andrieu, and M. Davy. Particle filtering for multi-target tracking and sensor management. In *International Conference on Information Fusion*, 2002.
- [75] N. Ehsan and M. Liu. Optimal bandwidth allocation in a delay channel. submitted to *JSAC*.
- [76] N. Ehsan and M. Liu. Optimal channel allocation for uplink transmission in satellite communications. submitted to *IEEE Transactions on Vehicular Technology*.
- [77] N. Ehsan and M. Liu. Server allocation with delayed state observation: sufficient conditions for the optimality an index policy. submitted to *PEIS*.
- [78] N. Ehsan and M. Liu. On the optimal index policy for bandwidth allocation with delayed state observation and differentiated services. In *IEEE*

- Annual Conference on Computer Communications*, volume 3, pages 1974–1983, Hong Kong, April 2004.
- [79] N. Ehsan and M. Liu. Properties of optimal resource sharing in delay channels. In *IEEE Conference on Decision and Control*, volume 3, pages 3277–3282, Paradise Island, Bahamas, 2004.
- [80] N. El Karoui and I. Karatzas. Dynamic allocation problems in continuous time. *Annals of Applied Probability*, 4(2):255–286, 1994.
- [81] V. V. Federov. *Theory of optimal experiments*. Academic Press, Orlando, 1972.
- [82] R. A. Fisher. *The design of experiments*. Oliver and Boyd, Edinburgh, 1935.
- [83] T. E. Fortmann, Y. Bar-Shalom, M. Scheffé, and S. Gelfand. Detection thresholds for tracking in clutter — A connection between estimation and signal processing. *IEEE Transactions on Automatic Control*, 30(3):221–229, March 1985.
- [84] Y. Freund, H. S. Seung, E. Shamir, and N. Tishby. Selective sampling using the query by committee algorithm. *Machine Learning*, 28(2-3):133–168, August 1997.
- [85] E. Frostig and G. Weiss. Four proofs of Gittins’ multi-armed bandit theorem. Technical report, The University of Haifa, Mount Carmel, 31905, Israel, November 1999.
- [86] N. Geng, C. E. Baum, and L. Carin. On the low-frequency natural response of conducting and permeable targets. *IEEE Transactions on Geoscience and Remote Sensing*, 37:347–359, January 1999.
- [87] J. C. Gittins. Bandit processes and dynamic allocation indices. *Journal of the Royal Statistical Society: Series B (Methodological)*, 41(2):148–177, 1979.
- [88] J. C. Gittins. *Multi-Armed Bandit Allocation Indices*. John Wiley and Sons, New York, NY, 1989.
- [89] J. C. Gittins and D. M. Jones. A dynamic allocation index for sequential design of experiments. *Progress in Statistics, Euro. Meet. Statist.*, 1:241–266, 1972.
- [90] K. D. Glazebrook, J. Niño Mora, and P. S. Ansell. Index policies for a class of discounted restless bandits. *Advances in Applied Probability*, 34(4):754–774, 2002.

- [91] K. D. Glazebrook and D. Ruiz-Hernandez. A restless bandit approach to stochastic scheduling problems with switching costs. Preprint, March 2005.
- [92] G. Golubev and B. Levit. Sequential recovery of analytic periodic edges in the binary image models. *Mathematical Methods of Statistics*, 12:95–115, 2003.
- [93] N. J. Gordon, D. J. Salmond, and A. F. M. Smith. A novel approach to non-linear and non-Gaussian Bayesian state estimation. *IEE Proceedings on Radar and Signal Processing*, 140:107–113, 1993.
- [94] E. Gottlieb and R. Harrigan. The umbra simulation framework. Sand2001-1533 (unlimited release), Sandia National Laboratory, 2001.
- [95] C. H. Gowda and R. Viswanatha. Performance of distributed CFAR test under various clutter amplitudes. *IEEE Transactions on Aerospace and Electronic Systems*, 35:1410–1419, 1999.
- [96] R. M. Gray. Vector quantization. *IEEE ASSP Magazine*, pages 4–29, Apr. 1984.
- [97] J. A. Gubner. *Probability and Random Processes for Electrical and Computer Engineers*. Cambridge University Press, New York, NY, 2006.
- [98] P. Hall and I. Molchanov. Sequential methods for design-adaptive estimation of discontinuities in regression curves and surfaces. *Annals of Statistics*, 31(3):921–941, 2003.
- [99] P. Hanselman, C. Lawrence, E. Fortunato, B. Tenney, and E. Blasch. Dynamic tactical targeting. In *Conference on Battlefield Digitization and Network-Centric Systems IV*, volume SPIE 5441, pages 36–47, 2004.
- [100] J. P. Hardwick and Q. F. Stout. Flexible algorithms for creating and analyzing adaptive sampling procedures. In N. Flournoy, W. F. Rosenberger, and W. K. Wong, editors, *New Developments and Applications in Experimental Design*, volume 34 of *Lecture Notes - Monograph Series*, pages 91–105. Institute of Mathematical Statistics, 1998.
- [101] T. Hastie, R. Tibshirani, and J. H. Friedman. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer Series in Statistics, Basel, CH, 2001.
- [102] J. Havrda and F. Chárvat. Quantification method of classification processes. *Kiberbetika Cisló*, 1(3):30–34, 1967.

- [103] Y. He and E. K. P. Chong. Sensor scheduling for target tracking in sensor networks. In *IEEE Conference on Decision and Control*, pages 743–748, 2004.
- [104] Y. He and E. K. P. Chong. Sensor scheduling for target tracking: A Monte Carlo sampling approach. *Digital Signal Processing*, 16(5):533–545, September 2006.
- [105] M. L. Hernandez, T. Kirubarajan, and Y. Bar-Shalom. Multisensor resource deployment using posterior Cramér-Rao bounds. *IEEE Transactions on Aerospace and Electronic Systems*, 40(2):399–416, April 2004.
- [106] A. O. Hero, B. Ma, O. Michel, and J. Gorman. Applications of entropic spanning graphs. *IEEE Signal Processing Magazine*, 19(2):85–95, 2002.
- [107] A. O. Hero, B. Ma, O. Michel, and J. D. Gorman. Alpha divergence for classification, indexing and retrieval. Technical Report Technical Report 328, Comm. and Sig. Proc. Lab. (CSPL), Dept. EECS, The University of Michigan, 2001.
- [108] K. J. Hintz. A measure of the information gain attributable to cueing. *IEEE Transactions on Systems, Man and Cybernetics*, 21(2):237–244, 1991.
- [109] K. J. Hintz and E. S. McVey. Multi-process constrained estimation. *IEEE Transactions on Systems, Man and Cybernetics*, 21(1):434–442, January/February 1991.
- [110] M. Horstein. Sequential decoding using noiseless feedback. *IEEE Transactions on Information Theory*, 9(3):136–143, 1963.
- [111] R. Howard. *Dynamic Programming and Markov Processes*. John Wiley and Sons, New York, NY, 1960.
- [112] C. Hue, J.-P. Le Cadre, and P. Perez. Sequential Monte Carlo methods for multiple target tracking and data fusion. *IEEE Transactions on Signal Processing*, 50:309–325, 2002.
- [113] C. Hue, J.-P. Le Cadre, and P. Perez. Tracking multiple objects with particle filtering. *IEEE Transactions on Aerospace and Electronic Systems*, 38:791–812, 2002.
- [114] M. Isard and J. MacCormick. BraMBLe: A Bayesian multiple-blob tracker. In *International Conference on Computer Vision*, 2001.

- [115] T. Ishikida. *Informational Aspects of Decentralized Resource Allocation*. PhD thesis, University of California, Berkeley, 1992.
- [116] T. Ishikida and P. Varaiya. Multi-armed bandit problem revisited. *Journal of Optimization Theory and Applications*, 83:113–154, 1994.
- [117] J. Jacod and P. Protter. *Probability Essentials*. Springer-Verlag, 2003.
- [118] A. H. Jazwinski. *Stochastic Processes and Filtering Theory*. Academic Press, New York, NY, 1970.
- [119] S. Ji, R. Parr, and L. Carin. Non-myopic multi-aspect sensing with partially observable Markov decision processes. *IEEE Transactions on Signal Processing*, 55(6):2720–2730, 2007.
- [120] S. Julier and J. Uhlmann. Unscented filtering and non-linear estimation. *Proceedings of the IEEE*, 92:401–422, 2004.
- [121] L. P. Kaelbling, M. L. Littman, and A. R. Cassandra. Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101:99–134, 1998.
- [122] R. Karlsson and F. Gustafsson. Monte Carlo data association for multiple target tracking. In *IEE Workshop on Target Tracking: Algorithms and Applications*, 2001.
- [123] H. Kaspi and A. Mandelbaum. Multi-armed bandits in discrete and continuous time. *Annals of Applied Probability*, 8:1270–1290, 1998.
- [124] K. Kastella. Discrimination gain for sensor management in multitarget detection and tracking. In *IEEE-SMC and IMACS Multiconference*, volume 1, pages 167–172, 1996.
- [125] K. Kastella. Discrimination gain to optimize classification. *IEEE Transactions on Systems, Man and Cybernetics—Part A: Systems and Humans*, 27(1), January 1997.
- [126] M. N. Katehakis and U. G. Rothblum. Finite state multi-armed bandit problems: Sensitive-discount, average-reward and average-overtaking optimality. *Annals of Applied Probability*, 6:1024–1034, 1996.
- [127] M. N. Katehakis and A. F. Veinott, Jr. The multi-armed bandit problem: Decomposition and computation. *Mathematics of Operations Research*, 12:262–268, 1987.
- [128] M. J. Kearns, Y. Mansour, and A. Y. Ng. A sparse sampling algorithm for near-optimal planning in large Markov decision processes. In *International Joint Conference on Artificial Intelligence*, pages 1324–1331, 1999.

- [129] F. P. Kelly. Multi-armed bandits with discount factor near one: The Bernoulli case. *Annals of Statistics*, 9:987–1001, 1981.
- [130] D. J. Kershaw and R. J. Evans. Optimal waveform selection for tracking systems. *IEEE Transactions on Information Theory*, 40(5):1536–50, September 1994.
- [131] D. J. Kershaw and R. J. Evans. Waveform selective probabilistic data association. *IEEE Transactions on Aerospace and Electronic Systems*, 33(4):1180–88, October 1997.
- [132] G. P. Klimov. Time sharing service systems I. *Theory of Probability and its Applications (in Russian: Teoriya Veroyatnostei i ee Primeneniya)*, 19:532–551, 1974.
- [133] G. P. Klimov. Time sharing service systems II. *Theory of Probability and its Applications (in Russian: Teoriya Veroyatnostei i ee Primeneniya)*, 23:314–321, 1978.
- [134] E. D. Kolaczyk and R. D. Nowak. Multiscale likelihood analysis and complexity penalized estimation. *Annals of Statistics*, 32(2):500–527, 2004.
- [135] A. Korostelev and J.-C. Kim. Rates of convergence for the sup-norm risk in image models under sequential designs. *Statistics and Probability Letters*, 46:391–399, 2000.
- [136] A. P. Korostelev. On minimax rates of convergence in image models under sequential design. *Statistics and Probability Letters*, 43:369–375, 1999.
- [137] A. P. Korostelev and A. B. Tsybakov. *Minimax Theory of Image Reconstruction*. Springer Lecture Notes in Statistics, 1993.
- [138] C. Kreucher, D. Blatt, A. Hero, and K. Kastella. Adaptive multi-modality sensor scheduling for detection and tracking of smart targets. *Digital Signal Processing*, 16(5):546–567, 2005.
- [139] C. Kreucher, A. Hero, K. Kastella, and D. Chang. Efficient methods of non-myopic sensor management for multitarget tracking. In *IEEE Conference on Decision and Control*, 2004.
- [140] C. Kreucher, A. O. Hero, and K. Kastella. Multiple model particle filtering for multi-target tracking. In *Workshop on Adaptive Sensor Array Processing*, 2004.

- [141] C. Kreucher, K. Kastella, and A. Hero. Multi-target sensor management using alpha divergence measures. In *International Conference on Information Processing in Sensor Networks*, 2003.
- [142] C. M. Kreucher, A. O. Hero, and K. Kastella. A comparison of task driven and information driven sensor management for target tracking. In *IEEE Conference on Decision and Control*, 2005.
- [143] C. M. Kreucher, A. O. Hero, K. D. Kastella, and M. R. Morelande. An information-based approach to sensor management in large dynamic networks. *Proceedings of the IEEE*, 95(5):978–999, May 2007.
- [144] C. M. Kreucher, K. Kastella, and A. O. Hero. Information based sensor management for multitarget tracking. In *SPIE Conference on Signal and Data Processing of Small Targets*, 2003.
- [145] C. M. Kreucher, K. Kastella, and A. O. Hero. Multitarget tracking using the joint multitarget probability density. *IEEE Transactions on Aerospace and Electronic Systems*, 39(4):1396–1414, 2005.
- [146] C. M. Kreucher, K. Kastella, and A. O. Hero. Sensor management using an active sensing approach. *Signal Processing*, 85(3):607–624, 2005.
- [147] V. Krishnamurthy. Algorithms for optimal scheduling and management of hidden Markov model sensors. *IEEE Transactions on Signal Processing*, 50(6):1382–1397, 2002.
- [148] V. Krishnamurthy and R. J. Evans. Hidden Markov model multiarmed bandits: A methodology for beam scheduling in multitarget tracking. *IEEE Transactions on Signal Processing*, 49(12):2893–2908, 2001.
- [149] V. Krishnamurthy and R. J. Evans. Correction to hidden Markov model multi-arm bandits: A methodology for beam scheduling in multi-target tracking. *IEEE Transactions on Signal Processing*, 51(6):1662–1663, 2003.
- [150] A. Krogh and J. Vedelsby. Neural network ensembles, cross validation, and active learning. *Advances in Neural Information Processing Systems*, 7:231–238, 1995.
- [151] W. S. Kuklinski. Adaptive sensor tasking and control. In *MITRE 2005 Technology Symposium*. MITRE Corporation, 2005.
- [152] S. Kullback. *Information Theory and Statistics*. Dover, 1978.
- [153] P. R. Kumar and P. Varaiya. *Stochastic Systems: Estimation, Identification, and Adaptive Control*. Prentice Hall, 1986.

- [154] H. Kushner. *Introduction to Stochastic Control*. Holt, Rinehart and Winston, New York, NY, 1971.
- [155] B. F. La Scala, B. Moran, and R. Evans. Optimal scheduling for target detection with agile beam radars. In *NATO SET-059 Symposium on Target Tracking and Sensor Data Fusion for Military Observation Systems*, 2003.
- [156] T. Lai and H. Robbins. Asymptotically efficient adaptive allocation rules. *Advances in Applied Mathematics*, 6:4–22, 1985.
- [157] R. E. Larson and J. L. Casti. *Principles of Dynamic Programming, Parts 1-2*. Marcel Dekker, New York, NY, 1982.
- [158] X. Liao, H. Li, and B. Krishnapuram. An m -ary KMP classifier for multi-aspect target classification. In *IEEE International Conference on Acoustics, Speech, and Signal Processing*, volume 2, pages 61–64, 2004.
- [159] M. L. Littman. The witness algorithm: Solving partially observable Markov decision processes. Technical Report CS-94-40, Brown University, 1994.
- [160] J. Liu and R. Chen. Sequential Monte Carlo methods for dynamic systems. *Journal of the American Statistical Association*, 1998.
- [161] C. Lott and D. Teneketzis. On the optimality of an index rule in multi-channel allocation for single-hop mobile networks with multiple service classes. *Probability in the Engineering and Informational Sciences*, 14:259–297, 2000.
- [162] W. S. Lovejoy. A survey of algorithmic methods for partially observed Markov decision processes. *Annals of Operations Research*, 28(1):47–65, 1991.
- [163] D. MacKay. Information-based objective functions for active data selection. *Neural Computation*, 4:590–604, 1992.
- [164] D. MacKay. *Information Theory, Inference and Learning Algorithms*. Cambridge University Press, 2004.
- [165] R. Mahler. Global optimal sensor allocation. In *National Symposium on Sensor Fusion*, volume 1, pages 167–172, 1996.
- [166] A. Mandelbaum. Discrete multiarmed bandits and multiparameter processes. *Probability Theory and Related Fields*, 71:129–147, 1986.

- [167] A. Mandelbaum. Continuous multi-armed bandits and multiparameter processes. *Annals of Probability*, 15:1527–1556, 1987.
- [168] S. Maskell, M. Rollason, N. Gordon, and D. Salmond. Efficient particle filtering for multiple target tracking with application to tracking in structured images. In *SPIE Conference on Signal and Data Processing of Small Targets*, 2002.
- [169] M. McClure and L. Carin. Matched pursuits with a wave-based dictionary. *IEEE Transactions on Signal Processing*, 45:2912–2927, December 1997.
- [170] S. P. Meyn and R. L. Tweedie. *Markov Chains and Stochastic Stability*. Springer-Verlag, London, 1993.
- [171] J. Mickova. Stochastic scheduling with multi-armed bandits. Master's thesis, University of Melbourne, Australia, 2000.
- [172] M. I. Miller, A. Srivastava, and U. Grenander. Conditional mean estimation via jump-diffusion processes in multiple target tracking/recognition. *IEEE Transactions on Signal Processing*, 43:2678–2690, 1995.
- [173] G. E. Monahan. A survey of partially observable Markov decision processes: Theory, models and algorithms. *Management Science*, 28(1):1–16, 1982.
- [174] M. Morelande, C. M. Kreucher, and K. Kastella. A Bayesian approach to multiple target detection and tracking. *IEEE Transactions on Signal Processing*, 55(5):1589–1604, 2007.
- [175] S. Musick and R. Malhotra. Chasing the elusive sensor manager. In *IEEE National Aerospace and Electronics Conference*, volume 1, pages 606–613, 1994.
- [176] D. Mušicki, S. Challa, and S. Suvorova. Multi target tracking of ground targets in clutter with LMIPDA-IMM. In *International Conference on Information Fusion*, Stockholm, Sweden, July 2004.
- [177] D. Mušicki and R. Evans. Clutter map information for data association and track initialization. *IEEE Transactions on Aerospace and Electronic Systems*, 40(4):387–398, April 2001.
- [178] D. Mušicki, R. Evans, and S. Stankovic. Integrated probabilistic data association. *IEEE Transactions on Automatic Control*, 39(6):1237–1240, June 1994.

- [179] P. Nash. *Optimal Allocation of Resources Between Research Projects*. PhD thesis, Cambridge University, 1973.
- [180] A. Nedic and M. K. Schneider. Index rule-based management of a sensor for searching, tracking, and identifying. In *Tri-Service Radar Symposium*, Boulder Colorado, June 2003.
- [181] A. Nehorai and A. Dogandžić. Cramér-Rao bounds for estimating range, velocity and direction with an active array. *IEEE Transactions on Signal Processing*, 49(6):1122–1137, June 2001.
- [182] J. Niño-Mora. Restless bandits, partial conservation laws, and indexability. *Advances in Applied Probability*, 33:76–98, 2001.
- [183] J. Niño-Mora. Dynamic allocation indices for restless projects and queuing admission control: a polyhedral approach. *Mathematical Programming, Series A*, 93:361–413, 2002.
- [184] R. Niu, P. Willett, and Y. Bar-Shalom. From the waveform through the resolution cell to the tracker. In *IEEE Aerospace Conference*, March 1999.
- [185] R. Nowak, U. Mitra, and R. Willett. Estimating inhomogeneous fields using wireless sensor networks. *IEEE Journal on Selected Areas in Communications*, 22(6):999–1006, 2004.
- [186] M. Orton and W. Fitzgerald. A Bayesian approach to tracking multiple targets using sensor arrays and particle filters. *IEEE Transactions on Signal Processing*, 50:216–223, 2002.
- [187] D. G. Pandelis and D. Teneketzis. On the optimality of the Gittins index rule in multi-armed bandits with multiple plays. *Mathematical Methods of Operations Research*, 50:449–461, 1999.
- [188] J. Pineau, G. Gordon, and S. Thrun. Point-based value iteration: An anytime algorithm for POMDPs. In *International Joint Conference on Artificial Intelligence*, August 2003.
- [189] M. K. Pitt and N. Shephard. Filtering via simulation: Auxiliary particle filters. *Journal of the American Statistical Association*, 94:590–599, 1999.
- [190] F. Pukelsheim. *Optimal Design of Experiments*. John Wiley and Sons, New York, NY, 1993.
- [191] M. L. Puterman, editor. *Dynamic Programming and its Applications*. Academic Press, New York, NY, 1978.

- [192] M. L. Puterman. *Markov Decision Problems: Discrete Stochastic Dynamic Programming*. John Wiley and Sons, New York, NY, 1994.
- [193] R. Raich, J. Costa, and A. O. Hero. On dimensionality reduction for classification and its application. In *IEEE International Conference on Acoustics, Speech, and Signal Processing*, Toulouse, May 2006.
- [194] A. Rényi. On measures of entropy and information. In *Berkeley Symposium on Mathematics, Statistics and Probability*, volume 1, pages 547–561, 1961.
- [195] R. Rifkin and A. Klautau. In defense of one-vs-all classification. *Journal of Machine Learning Research*, 5:101–141, January 2004.
- [196] S. M. Ross. *Applied Probability Models with Optimization Applications*. Dover Publications, New York, NY, 1970.
- [197] S. M. Ross. *Introduction to Stochastic Dynamic Programming*. Academic Press, New York, NY, 1983.
- [198] N. Roy, G. Gordon, and S. Thrun. Finding approximate POMDP solutions through belief compression. *Journal of Artificial Intelligence Research*, 23:1–40, 2005.
- [199] P. Runkle, P. Bharadwaj, and L. Carin. Hidden Markov model multi-aspect target classification. *IEEE Transactions on Signal Processing*, 47:2035–2040, July 1999.
- [200] P. Runkle, L. Carin, L. Couchman, T. Yoder, and J. Bucaro. Multi-aspect identification of submerged elastic targets via wave-based matching pursuits and hidden Markov models. *J. Acoustical Soc. Am.*, 106:605–616, August 1999.
- [201] J. Rust. Chapter 14: Numerical dynamic programming in economics. In H. Amman, D. Kendrick, and J. Rust, editors, *Handbook of Computational Economics*. Elsevier, North Holland, 1996.
- [202] J. Rust. Using randomization to break the curse of dimensionality. *Econometrica*, 65:487–516, 1997.
- [203] W. Schmaedeke and K. Kastella. Event-averaged maximum likelihood estimation and information-based sensor management. *Proceedings of SPIE*, 2232:91–96, June 1994.
- [204] M. K. Schneider, G. L. Mealy, and F. M. Pait. Closing the loop in sensor fusion systems: Stochastic dynamic programming approaches. In *American Control Conference*, 2004.

- [205] D. Schulz, D. Fox, and J. Hightower. People tracking with anonymous and ID-sensors using Rao-Blackwellised particle filter. In *International Joint Conference on Artificial Intelligence*, 2003.
- [206] N. Secomandi. A rollout policy for the vehicle routing problem with stochastic demands. *Operations Research*, 49:796–802, 2001.
- [207] C. E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27:379–423, 1948.
- [208] A. N. Shireyaev. *Optimal Stopping Rules*. Springer-Verlag, 1978.
- [209] A. N. Shireyaev. *Probability*. Springer-Verlag, 1995.
- [210] A. Singh, R. Nowak, and P. Ramanathan. Active learning for adaptive mobile sensing networks. In *International Conference on Information Processing in Sensor Networks*, Nashville, TN, April 2006.
- [211] D. Sinno. *Attentive Management of Configurable Sensor Systems*. PhD thesis, Arizona State University, 2000.
- [212] D. Sinno and D. Cochran. Dynamic estimation with selectable linear measurements. In *IEEE International Conference on Acoustics, Speech, and Signal Processing*, pages 2193–2196, Seattle, WA, 1998.
- [213] D. Sinno, D. Cochran, and D. Morrell. Multi-mode detection with Markov target motion. In *International Conference on Information Fusion*, volume WeD1, pages 26–31, Paris, France, 2000.
- [214] S. P. Sira, D. Cochran, A. Papandreou-Suppappola, D. Morrell, W. Moran, S. Howard, and R. Calderbank. Adaptive waveform design for improved detection of low RCS targets in heavy sea clutter. *IEEE Journal on Selected Areas in Signal Processing*, 1(1):55–66, June 2007.
- [215] S. P. Sira, A. Papandreou-Suppappola, and D. Morrell. Time-varying waveform selection and configuration for agile sensors in tracking applications. In *IEEE International Conference on Acoustics, Speech, and Signal Processing*, volume 5, pages 881–884, March 2005.
- [216] R. D. Smallwood and E. J. Sondik. The optimal control of partially observable Markov processes over a finite horizon. *Operations Research*, 21:1071–1088, 1973.
- [217] E. J. Sondik. *The Optimal Control of Partially Observable Markov Processes*. PhD thesis, Stanford University, 1971.

- [218] E. J. Sondik. The optimal control of partially observable Markov processes over the infinite horizon: Discounted costs. *Operations Research*, 26(2):282–304, 1978.
- [219] N. O. Song and D. Teneketzis. Discrete search with multiple sensors. *Mathematical Methods of Operations Research*, 60:1–14, 2004.
- [220] Statlog. Landsat MSS data.
- [221] L. D. Stone, C. A. Barlow, and T. L. Corwin. *Bayesian Multiple Target Tracking*. Artech House, Boston, MA, 1999.
- [222] M. Stone. Cross-validatory choice and assessment of statistical predictions. *Journal of the Royal Statistical Society, Series B*, 36:111–147, 1974.
- [223] C. Striebel. Sufficient statistics in the optimum control of stochastic systems. *Journal of Mathematical Analysis and Applications*, 12:576–592, 1965.
- [224] K. Sung and P. Niyogi. Active learning for function approximation. *Proc. Advances in Neural Information Processing Systems*, 7, 1995.
- [225] R. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, Cambridge, MA, 1998.
- [226] S. Suvorova, S. D. Howard, W. Moran, and R. J. Evans. Waveform libraries for radar tracking applications: Maneuvering targets. In *Defence Applications of Signal Processing*, 2004.
- [227] I. J. Taneja. New developments in generalized information measures. *Advances in Imaging and Electron Physics*, 91:37–135, 1995.
- [228] G. Tesauro. Temporal difference learning and TD-gammon. *Communications of the ACM*, 38(3), March 1995.
- [229] S. Tong and D. Koller. Support vector machine active learning with applications to text classification. *International Conference on Machine Learning*, pages 999–1006, 2000.
- [230] J. N. Tsitsiklis. A lemma on the multiarmed bandit problem. *IEEE Transactions on Automatic Control*, 31:576–577, 1986.
- [231] B. E. Tullsson. Monopulse tracking of Rayleigh targets: A simple approach. *IEEE Transactions on Aerospace and Electronic Systems*, 27:520–531, 1991.

- [232] M. Van Oyen, D. Pandelis, and D. Teneketzis. Optimality of index policies for stochastic scheduling with switching penalties. *Journal of Applied Probability*, 29:957–966, 1992.
- [233] M. P. Van Oyen and D. Teneketzis. Optimal stochastic scheduling of forest networks with switching penalties. *Advances in Applied Probability*, 26:474–479, 1994.
- [234] H. L. van Trees. *Detection, Estimation, and Modulation Theory: Part I*. John Wiley and Sons, New York, NY, 1968.
- [235] H. L. van Trees. *Detection, Estimation and Modulation Theory, Part III*. John Wiley and Sons, New York, NY, 1971.
- [236] V. N. Vapnik. *Statistical Learning Theory*. John Wiley and Sons, New York, NY, 1998.
- [237] V. N. Vapnik. An overview of statistical learning theory. *IEEE Transactions on Neural Networks*, 10(5):988–999, 1999.
- [238] P. P. Varaiya, J. C. Walrand, and C. Buyukkoc. Extensions of the multiarmed bandit problem: The discounted case. *IEEE Transactions on Automatic Control*, 30:426–439, 1985.
- [239] M. Veth, J. Busque, D. Heesch, T. Burgess, F. Douglas, and B. Kish. Affordable moving surface target engagement. In *IEEE Aerospace Conference*, volume 5, pages 2545–2551, 2002.
- [240] P. Vincent and Y. Bengio. Kernel matching pursuit. *Machine Learning*, 48:165–187, 2002.
- [241] A. Wald. *Sequential Analysis*. John Wiley and Sons, New York, NY, 1947.
- [242] J. Wang, A. Dogandžić, and A. Nehorai. Maximum likelihood estimation of compound-Gaussian clutter and target parameters. *IEEE Transactions on Signal Processing*, 54:3884–3898, October 2006.
- [243] R. B. Washburn, M. K. Schneider, and J. J. Fox. Stochastic dynamic programming based approaches to sensor resource management. In *International Conference on Information Fusion*, volume 1, pages 608–615, 2002.
- [244] R. R. Weber. On Gittins index for multiarmed bandits. *Annals of Probability*, 2:1024–1033, 1992.
- [245] R. R. Weber and G. Weiss. On an index policy for restless bandits. *Journal of Applied Probability*, 27:637–648, 1990.

- [246] C. C. White III. Partially observed Markov decision processes: A survey. *Annals of Operations Research*, 32, 1991.
- [247] P. Whittle. Multi-armed bandits and Gittins index. *Journal of the Royal Statistical Society: Series B (Methodological)*, 42:143–149, 1980.
- [248] P. Whittle. Arm-acquiring bandits. *Annals of Probability*, 9:284–292, 1981.
- [249] P. Whittle. *Optimization Over Time: Dynamic Programming and Stochastic Control*. John Wiley and Sons, New York, NY, 1983.
- [250] P. Whittle. Restless bandits: Activity allocation in a changing world. *Journal of Applied Probability*, 25A:287–298, 1988.
- [251] P. Whittle. Tax problems in the undiscounted case. *Journal of Applied Probability*, 42(3):754–765, 2005.
- [252] R. Willett, A. Martin, and R. Nowak. Backcasting: Adaptive sampling for sensor networks. In *Information Processing in Sensor Networks*, 26-27 April, Berkeley, CA, USA, 2004.
- [253] I. J. Won, D. A. Keiswetter, and D. R. Hanson. GEM-3: A monostatic broadband electromagnetic induction sensor. *J. Environ. Eng. Geophys.*, 2:53–64, March 1997.
- [254] G. Wu, E. K. P. Chong, and R. L. Givan. Burst-level congestion control using hindsight optimization. *IEEE Transactions on Automatic Control*, 47:979–991, 2002.
- [255] R. W. Yeung. *A First Course in Information Theory*. Springer, 2002.
- [256] H. Yu and D. P. Bertsekas. Discretized approximations for pomdp with average cost. In *Conference on Uncertainty in Artificial Intelligence*, pages 619–627, 2004.
- [257] Y. Zhang, L. M. Collins, H. Yu, C. E. Baum, and L. Carin. Sensing of unexploded ordnance with magnetometer and induction data: Theory and signal processing. *IEEE Transactions on Geoscience and Remote Sensing*, 41:1005–1015, May 2003.
- [258] Y. Zhang, X. Liao, and L. Carin. Detection of buried targets via active selection of labeled data: application to sensing subsurface uxo. *IEEE Transactions on Geoscience and Remote Sensing*, 42(11):2535–2543, 2004.

- [259] F. Zhao, J. Shin, and J. Reich. Information-driven dynamic sensor collaboration. *IEEE Signal Processing Magazine*, pages 61–72, March 2002.

Index

- Acoustic underwater sensing, 27
- Action, 98
- Action selector, 102, 103
- Action-sequence approximations, 109
- Active learning rate, 193
- Adaptive partition, 77
- Adaptive sampling, 177, 181
- ADP, 26
- Airborne laser scanning sensor, 178
- Airborne sensor, 97
- alpha divergence, 37
- alpha entropy, 36
- Ambiguity function, 243
- Ambiguity function, 224, 242
- Approximate Dynamic Programming, 26
- Approximate dynamic programming, 108
- Array, 97
- Average reward MDP, 13
- Azimuthal ambiguity, 228

- Base policy, 111
- Battery status, 102
- Bayes rule, 23, 49, 103, 183, 247
- Bayes update, 23
- Bayes' rule, 277
- Bayesian CRB, 40
- Bayesian filtering, 64
- Beam scheduling, 245
- Beamshaping, 226
- Belief state, 41, 98, 276
- Belief-state approximation, 114
- Belief-state feedback, 98
- Belief-state simplification, 114
- Belief-state space, 98, 101
- Bellman equation, 12
- Bellman's principle, 99
- Blind range, 230
- Bound, 105, 109
- Boundary fragment class, 192
- Box-counting dimension, 190
- Brownian motion, 65
- Burnashev-Zigangirov algorithm, 183
- BZ algorithm, 183

- Carrier, 222
- Cartesian product, 102
- CFAR, 71
- Chapman-Kolmogorov equation, 275
- Chernoff bound, 39, 187
- Chernoff exponent, 39
- Chirp waveform, 225
- Chirp waveform library, 253
- Classification, 203
- Classification reduction of optimal policy search, 43
- Closed loop, 102
- Combined innovation, 237
- Communication resource, 101
- Completely observable rollout, 114
- Complexity regularized estimator, 192
- Conditional entropy, 270
- Constant false alarm rate, 71
- Continuation set, 131
- Control architecture, 111
- Control, receding horizon, 100
- Controller, 102
- Coupled partitions, 75
- Covariance update, 247
- CRB, 211
- CROPS, 43
- Cross ambiguity function, 224
- Cross-cue, 260
- Cusp-free boundary set, 196

- D-optimal experimental design, 210
- DAI, 124
- DARPA, 261
- Data fusion, 262
- Discount factor, 276
- Discounted reward MDP, 13
- Divergence, 96
 - Alpha-divergence, 96, 107
 - Renyi-divergence, 96, 107
- Domain knowledge, 106, 118
- Dominating control law, 136
- Dominating machine, 136

306 FOUNDATIONS AND APPLICATIONS OF SENSOR MANAGEMENT

- Doppler aliasing, 232
- Doppler processing, 230, 232, 233
- DTT, 261
- DTT:TEST, 265
- Dynamic allocation index, 124
- Dynamic programming, 122

- Electrically scanned array, 97
- Electromagnetic induction sensor, 213
- EMI, 213
- Empirical loss, 201
- Expected information gain, 42
- Expected loss, 201
- Expected value-to-go, 106

- f-divergence, 37
- Feature, 108
- Feedback, 98
- Filter, 102
 - Auto-ambiguity function, 224
 - Bayesian filter, 64
 - Cross-ambiguity function, 224
 - Extended Kalman filter, 114
 - IPDA tracker, 235
 - Kalman filter, 103, 114
 - Matched filter, 224, 231
 - Measurement filter, 102
 - Multi-target particle filter, 73
 - Non-linear filter, 65
 - Particle filter, 71, 72, 103, 111, 114, 118
 - PDA tracker, 235, 241
 - SIR particle filter, 71
 - Unscented Kalman filter, 114
- Fisher information, 40, 234
- Fisher information matrix, 233
- Fokker-Planck equation, 65
- Foresight optimization, 109
- Forward induction, 129
- Forward prediction, 247
- FPE, 65
- Function approximator, 107

- Gauss-Markov model, 235
- Gaussian, 103, 114
- Gittins index, 124, 125, 129, 131–133, 138, 141, 146, 280
- GMTI, 60, 260
- Gridlock, 261
- Ground surveillance, 97

- Heterogeneous sensors, 96
- Heuristics, 106
- Hidden Markov model, 276
- Hindsight optimization, 109
- Hitting time, 280
- Homogeneous Markov process, 274
- Horizon, 96, 99
- Hyperspectral imaging, 54, 55, 258

- ID, 258
- IMM, 245
- Independent partitions, 75
- Index-type allocation policies, 122
- Indexability, 143
- Information divergence, 37
- Information gain, 42, 96, 107, 115
- Information gain sandwich bound, 45
- Information state, 41, 98, 276
- Innovation, 236
- Innovation covariance matrix, 236
- IPDA tracker, 235
- ISAR, 60
- ISR, 60, 257
- Itô equation, 65

- JMPD, 59
- Joint multi-target probability density, 59
- Joint probabilistic data association, 61
- JPDA, 61
- JPG, 214
- JSTARS, 48, 257

- Kalman filter, 103, 114, 278
- Kalman gain, 247
- Kalman update, 247
- Kalman update equations, 237, 246
- Kernel matching pursuits, 215
- Keyhole spacecraft, 258
- Kinematic prior, 64
- KL divergence, 271
- Klimov's problem, 147
- KMP, 215
- KP, 64

- Landmine sensing, 220
- Landsat radar, 54
- LFM, 225, 243, 252, 253
- Likelihood ratio test, 39
- LMIPDA, 245
- Lookahead, 100

- MAB, 121
- Magnetometer sensor, 213
- MAP detector, 39
- Marginalized information gain, 46
- Markov chain, 238, 239, 246, 273
- Markov decision process, 10, 13
- Markov process, 10
- Markov property, 273
- Markov reward processes, 275
- Matched filter, 231
- MDP, 10, 13
- Measure of effectiveness, 252
- Measurement, 102
- Measurement filter, 102

- MHT, 61
- MI, 42
- MIG, 46
- Mobile sensor, 96
- Model sensitivity, 51
- Monostatic radar, 230
- Monte Carlo, 103, 104
- Moving target, 96
- Moyal's Identity, 232
- Moyal's identity, 224
- MTI radar, 48
- Multi-armed bandit, 121, 122
 - Arm-acquiring bandit, 137
 - Bandit process, 123
 - Classical multi-armed bandit, 123
 - Multiple Plays, 140
 - Restless bandits, 142
 - Superprocess, 134
 - Switching Penalties, 138
- Multi-target tracking, 48
- Multiple hypothesis tracking, 61
- Mutual information, 42, 251, 254
- Myopic policy, 125

- Narrow band approximation, 223
- National asset, 258
- Neurodynamic programming, 108
- Neyman-Pearson detector, 243

- Objective function, 99
- Obscuration, 97
- Observable, 102
- Observation, 102
- Observation law, 103, 276
- Observation space, 276
- OID, 72
- Optimal design of experiments, 210
- Optimal policy, 98

- Parallel rollout, 111
- Parametric approximation, 107
- Partial conservation laws, 144
- Partially observable Markov decision process, 19
- Partially observable Markov processes, 276
- Partially Observed Markov Decision Problems, 19
- Particle filter, 103, 111, 114, 118
- Passive learning rate, 193
- PDA tracker, 235, 237
- Phased array antennas, 226
- PMHT, 61
- Policy, 11, 99, 124
 - Adaptive policy, 30
 - Admissible policy, MDP, 11
 - Admissible policy, POMDP, 20
 - Base policy, 110, 111, 114
 - Completely observable rollout, 114
 - CROPS, 43
 - EVTG approximation, 117
 - Forward induction policy, 126
 - Index-type policy, 124
 - Information gain, 47
 - Information gain policy, 117
 - Iteration operator, 16
 - Markov policy, 11, 13
 - MDP policy, 12
 - Multistep lookahead policy, 125
 - Myopic, 95
 - Myopic policy, 41, 117, 125
 - Non-myopic, 95
 - Optimal policy, 12
 - Parallel rollout, 111
 - Policy iteration, 19, 110
 - POMDP policy, 20
 - Random policy, 117
 - Rollout, 110
 - Search, 42
 - Single stage policy, 13
 - Stationary, 100
 - Stationary policy, 13
- Policy improvement, 110
- Policy iteration, 19
- Policy, optimal, 98
- POMDP, 19, 41, 95, 278
- POMDP approximation, 95
- popup threat, 259
- Predator, 262
- PRI, 221
- Principal components, 208
- Principle of Optimality, 12
- Probabilistic bisection algorithm, 182
- Probabilistic multiple hypothesis tracker, 61
- Probability of decision error, 39
- Probability space, 278
- Proxy for performance, 45
- Pulse compression, 225
- Pulse-Doppler radar, 225

- Q-function, 103
- Q-learning, 108
- Q-value, 100
- Q-value approximation, 98, 104

- Rényi divergence, 37
- Rényi entropy, 36
- Radar
 - FOPEN, 258
 - Hyperspectral imaging, 258
 - Laser, 258
 - Pulse-Doppler radar, 225
- Radar system, 221, 257
- Radar:Beam scheduling, 245
- Range aliasing, 232
- Range ambiguity, 228

308 *FOUNDATIONS AND APPLICATIONS OF SENSOR MANAGEMENT*

- Ranking, 101, 104
- RDP, 191
- Receding horizon, 100
- Receiver operating characteristic, 215
- Recursive dyadic partitions, 191
- Reduction to classification, 43, 115
- Regret, 106
- Reinforcement learning, 108
- Relaxation, 105
- Resource management, 95
- Revisit time, 245–249
- Reward, 99, 275
- Reward surrogation, 115
- Riccati equations, 236
- ROC, 215
- Rollout, 110

- Sampling importance resampling, 71
- Scanned array, 97
- Scheduling policy, 124
- SDP, 124
- Sensor motion, 96
- Sensor scheduling, 8, 31, 34, 53, 60, 96, 110, 154, 157, 163, 221, 232, 234, 259
- Sensor trajectory, 97
- Sensor usage cost, 113
- Sequential data selection, 210
- Shannon entropy, 270
- Shannon entropy policy, 42
- Shannon mutual information, 272
- Shannon, Claude, 35
- Sigma-field, 278
- SIR, 71
- SNCR, 60
- State, 98
- State space, 98
- State-transition law, 274
- State-transition matrix, 274
- State-transition probability, 274
- Stationary, 100
- Stationary MDP, 13
- Stationary policy, 100
- Stochastic matrix, 274
- Stopping set, 131

- Stopping time, 126, 279
- Sufficient statistic, 22, 131
- Surrogate reward, 115
- Surveillance, 97
- Switching index, 139, 140

- T-step-look-ahead policy, 125
- Tactical asset, 258
- Target identification, 46
- Target motion, 96
- Target tracking, 46, 96, 102, 106, 110
- Tax problem, 147
- Terrain classification, 54
- Terrain elevation, 97
- Theater asset, 258
- Time-homogeneous Markov process, 274
- Topographical map, 97
- Total reward MDP, 13
- Track existence, 239
- Track existence, 238, 247
- Track life, 264
- Tracking, 96, 102, 106, 110, 163
- Tracking error, 113
- Training, 107
- Transition law, 103, 274
- Transition matrix, 274
- Transition probability, 274
- Twenty questions game, 178

- U-2, 259
- UAV, 258
- UGS, 258
- Uncertainty reduction measures, 41
- Unobservable states, 102
- UXO, 203, 213

- Validation gate, 235, 238, 242
- Value iteration, 18
- Value-to-go, 106

- Waveform libraries, 249
- Waveform library utility, 251
- Waveform scheduling, 234
- Waveform selection, 55