SENSOR MANAGEMENT FOR RADAR: A TUTORIAL

Bill Moran
The University of Melbourne, Parkville, Vic 3010, Australia
and
Prometheus Inc. 21 Arnold Ave, Newport, RI 02840, USA

Sofia Suvorova
The University of Melbourne, Parkville, Vic 3010, Australia

Stephen Howard
Electronic Warfare and Radar Division
Defence Science and Technology Organisation
Edinburgh, SA5100,Australia

Abstract In this chapter we describe some of the ideas being pursued in sensor scheduling as they apply to radar. A modern phased-array pulse-Doppler radar has several different parameters available for scheduling: waveform, beam-shape, beam direction, pulse repetition interval, etc. Choice of different values for these parameters provides different transmit modes for the radar and these modes in turn provide a variety of "blurrings" of the image of the scene. The application of ideas in scheduling to the different possible modes of the transmit phase of such a radar, has been shown in simulation to improve many aspects of the performance in tracking and detection of targets. We give a quick introduction to the ideas of radar followed by a discussion of some of the theoretical ideas involved, and with results of some simulations. We end with a discussion of the theoretical problem of scheduling the measurements and tracking of a one-dimensional system.

Keywords:
Radar; sensor scheduling; waveform; beam-shape; control; detection; tracking; revisit time; myopic; non-myopic.

1. Introduction

A radar system is a way of viewing a scene using electromagnetic radiation at wavelengths that can be processed using electronic equipment. Since ambient radiation at these wavelengths tends to be low in power, typically radars provide the illumination as well as the viewing system. The control of the source of radiation leads to major advantages, as well as some disadvantages. The most important disadvantage is that the amount of illumination is limited. Most radar systems in use are *monostatic*; that is, their illumination source and receiver are collocated. This has the advantage of shared electronics and antennas. Much effort is currently going into *multistatic* radar systems, but in this chapter we will focus only on monostatic radars. For such radars the energy returning to the receiver from a scatterer is inversely proportional to the 4th power of the distance. This means that, to achieve significant range, radars have to rely on a mix of high transmission power, clever ideas in the use of waveforms, sophisticated antenna design to focus the energy, and high performance signal processing.

Our aim in this chapter is to describe ideas being explored for the control of radar systems. Since we are not assuming any expertise in radar, we begin with a short description of the ideas of radar theory. This description focuses on the most commonly used form of the technology, namely a pulse-Doppler radar system. After that we dicuss some of the basic ideas in sensor management and then give results of simulations that show the kind of improvement that the use of sensor scheduling might produce. We have focused on work we have been associated with, and have omitted much excellent work of other workers in this burgeoning subject. Finally we discuss a theoretical problem in sensor management.

2. Radar Fundamentals

In this first section we discuss the basic ideas in a pulse-Doppler radar system. Our treatment is brief and focuses on the underlying theory rather than on the important issues of implementation.

2.1 Ambiguity and Radar

Illumination of the scene is provided by a signal that is emitted from the radar system. This signal is usually a waveform that is relatively slowly varying superimposed on a rapidly oscillating sinusoidal carrier. Thus it can be expressed as

$$\mathbf{s}(t) = \mathbf{w}(t).\cos(2\pi f_c t). \tag{1}$$

Here $\mathbf{w}(t)$ is the slowly varying waveform, and f_c is the carrier frequency. It is important to make the rather obvious observation at this stage that all signals transmitted and received are real-valued. However, it is possible to represent complex waveforms in such a way that they can be transmitted. Thus for a complex waveform $\mathbf{w}(t)$ we transmit the signal

$$\mathbf{s}(t) = (\Re\mathbf{w}(t)).\cos(2\pi f_c t) - (\Im\mathbf{w}(t)).\sin(2\pi f_c t). \tag{2}$$

On return, the "in-phase" or I component can be separated from the "quadrature" or Q component by demodulation against $\cos(2\pi(f_c t))$ and $\sin(2\pi f_c t))$ respectively. Much of the theory of radar processing takes place in the complex domain. It is convenient, and a powerful theoretical device, to replace the signal (2) by its complex version:

$$\mathbf{s}_c(t) = \mathbf{w}(t)).\exp(2\pi i f_c t), \tag{3}$$

so that $\mathbf{s}(t) = \Re\big(\mathbf{s}_c(t)\big)$. The carrier is often in the range 1–30GHz. The waveform will typically occupy a bandwidth that is less than 1/10 of that.

The superposition principle allows us to assume just a single scatterer in the view of the radar. The transmitted signal hits this scatterer whose distance (we measure distance and time in the same units) from the (collocated) transmitter and receiver is r. Assume that the scatterer is stationary. The return signal will be a delayed version of the original, delayed by the total round trip time from the radar to the scatterer. Specifically the signal voltage at the antenna of the receiver is

$$\mathbf{s}_u(t) = A\mathbf{s}(t - 2r) \tag{4}$$

where A represents the overall attenuation and includes a phase change (so is complex) due to reflection.

In the receiver some noise is added ("receiver noise"), arising from thermal activity generated within the components of the receiver. For distant scatterers the return signal is often so weak that this thermal noise can become a significant issue. We write

$$\mathbf{s}_r(t) = \mathbf{s}_u(t) + N(t),$$

where $N(t)$ is a white Gaussian process, for the signal after the initial stages of the receiver. Thermal noise is to a good approximation white and Gaussian.

Now we consider the possibility that the target is moving relative to the radar. The scattered waveform is modified by the Doppler effect. If this is done correctly it results in a "time dilation" of the return signal, so that, if the target has a radial velocity v, the return signal $\mathbf{s}_u(t)$ becomes

$$\mathbf{s}_u(t) = A\mathbf{s}(\alpha t - 2r),$$

where

$$\alpha = \frac{(1 - \frac{v}{c})}{(1 + \frac{v}{c})}.$$

When v is much smaller than c this is approximated by $\alpha = (1 - 2v/c)$. A further approximation is possible if, as is usually the case, the signal is "narrow band"; that is, if its (Fourier) spectrum is essentially in a range $(f_c - \delta, f_c + \delta)$ and its reflection in the origin, where δ is small compared to f_c. For most radar applications, this is a reasonable assumption since the signal modulating the carrier will have relatively low bandwidth. In this case, the return signal is approximated by shifting the frequency of the return from a stationary target at the same range by $f_d = (2v/c)f_c$, the so-called "Doppler frequency". This is best written in terms of the complex signal

$$\mathbf{s}_u(t) = \Re\big(\mathbf{w}(t - \frac{2R}{c}).e^{2\pi i f_c(1-2v/c)(t-\frac{2R}{c})}\big) \tag{5}$$

This equation is the standard one used in most radar calculations.

When the return is received, it is demodulated to strip off the carrier frequency. Typically, the return is "mixed with", that is multiplied by, $\cos 2\pi f t$ and then low-pass filtered to eliminate the high frequency component of the mixed signal. This is the demodulation phase refered to earlier.

In the complex domain, the demodulated signal is as described in (5). The signal is then filtered against another chosen signal $\mathbf{v}(t)$, often $\mathbf{v}$ is chosen to be the same as $\mathbf{w}$ (*match-filtering*); that is, it is correlated with that signal, resulting in

$$A\mathbf{w}, \mathbf{v}(x, f) = \int_{\mathbf{R}} \mathbf{v}(t)^* \mathbf{w}(t - x) e^{2\pi i f t}\, dt, \tag{6}$$

after a slight change of variable.

A general scene may be regarded as a function of range and Doppler, corresponding to a "reflectivity" assignment $\rho(t, f)$ to each value of range and Doppler. We include in this description of the scene the attenuation due to range of the scatterer. The superposition principle says that the resulting return is a convolution in range and Dopper of the scene with the ambiguity:

$$R(\tau, f) = \iint_{\mathbf{R}^2} \rho(\tau', f') A_{\mathbf{w},\mathbf{v}}(\tau - \tau', f - f')\, d\tau' df' \tag{7}$$

By varying the waveform, we are able to vary the shape of the ambiguity and thereby the kind of blurring that the radar process does to the scene. Evidently it would be best if there were no blurring, that is, if the ambiguity were a "thumbtack" with a spike at the origin and zero elsewhere. Unfortunately, there is a fundamental limitation that prevents this. It is known in various forms, in particular, as (one form of) the Heisenberg Uncertainty Principle, and as Moyal's Identity. In the latter formulation, it is expressed as follows:

$$||A_{\mathbf{w},\mathbf{v}}||_{L^2(\mathbf{R}^2)} = ||\mathbf{w}||_{L^2(\mathbf{R})}.||\mathbf{v}||_{L^2(\mathbf{R})} \tag{8}$$

It states that the L^2 norm of the ambiguity function as a function on $\mathbf{R}^2$ is the product of the L^2 norms of the transmit signal and the filtering signal as functions on $\mathbf{R}$. Since signals have finite energy, the ambiguity must be an L^2 function, and have a lower bound on its L^2 norm. Accordingly a "thumbtack" is impossible. The range-Doppler must be "blurred" by the imaging process in radar.

2.2 Beam-forming

In addition to finding range and Doppler, a radar usually needs to estimate the direction of a target. This is done by pointing the illumination in particular directions and "filtering" the return according to which direction it comes from.

The classical way to form a beam in radar is to use a paraboloidal dish. The beam is pointed in a given direction by mechanically steering the dish. Both the transmit and return beams are "spatially filtered" by the dish. Returns from particular directions are emphasized and those from other directions are attenuated. More and more this approach is being replaced by an electronically steered array antenna. Typically, this is comprised of a multiplicity of small antenna elements to which the transmit signal is fed. By varying the phase of the signal across the array it is possible to steer the direction of the beam, and by varying the voltage applied to each element it is possible to reshape the beam. The direction and the shape of the transmit beam can be varied rapidly. This is particularly important in a situation where the radar is performing multiple functions such as tracking several targets while detecting new targets. As a receive antenna, such a system can simultaneously steer many beams by means of the processing of the returns at each antenna element.

In neither the mechanical nor the electronic approaches is the beam perfectly sharp. This is inevitable since the aperture of the system is finite in extent. In the case of the electronic array, this problem is compounded by the fact that the array has discrete elements, rather than a continuum. However, in the latter case it is controllable. As a result of this imperfection, again the scene is "blurred"; in this case the directions of the scatterers are averaged over the response of the antenna. In the case of an electronic array, it is possible to change the "blurring" as well as beam-direction quickly. Thus in a phased-array system there is scope for the control of the illumination.

2.3 Doppler Processing and Pulse Compression

One way of copying with the ambiguity trade-off problem forced by Moyal's Identity (8) is to use a technique called *Doppler processing.* There are several issues associated with the accurate measurement of range and Doppler:

- A short pulse gives more accurate range measurement;
- A longer pulse has more energy in it, and the more energy used in illumination the more will be scattered back;
- The effect of the Doppler of typical targets on short pulses is essentially trivial.

An imperfect solution to the problems arising from the contradictory (to Moyal's Identity) requirements of good range and Doppler measurement is adopted by a *pulse-Doppler radar.* The solution involves the following mechanisms:

DP-1) Pulses of a length short enough to incur relatively little Doppler effect but long enough to individually give relatively high energy on target are chosen;

DP-2) These pulses are chosen in such a way that their auto-correlations are close to a spike with small side-lobes;

DP-3) A number of such pulses are transmitted with long gaps between them to give time for the Doppler to have effect across the whole sequence of pulses.

The effect of DP-2) is to produce a virtual pulse whose length is the width of the central lobe. Of course, this is never completely perfect since it does have side-lobes, but waveforms have been described for which the performance in this respect is excellent. DP-3) means that the Doppler frequency shift is being sampled at a discrete set of time points. If the sampling rate is faster than the Nyquist of the Doppler frequency shift, then the Doppler can be unambiguously extracted.

One might ask why Moyal's Identity does not cause problems here. Of course it does. Whatever the sampling rate, there are Doppler frequencies that are ambiguous and correspond to side-lobes in the overall ambiguity of the series of pulses. It is important to choose the sampling rate to be high enough that this does not happen for targets of interest. On the other hand, if the sampling rate is high then returns of earlier pulses from distant targets can appear after later pulses have been transmitted. This *range-aliasing* also corresponds to side-lobes in the overall ambiguity. Thus Doppler processing also suffers the same problems as a single waveform. However, it provides a mechanism for control of the position of the side-lobes to best fit the context. Moreover, it is possible to view the sampling rate, as well as the number of pulses used in this processing, as control parameters in scheduling a sensor.

3. Sensor Management — Overview

Conventional radars typically employ the same waveform and beampattern over many pulses. The received signal can be, and often is, processed in several ways to extract different kinds of information, or in response to knowledge gained from the environment, but on the transmit side, the mode of operation of the radar system is essentially static. In these systems it may be possible to modify the waveform used offline but not during the processing period. Recent advances in hardware have made the possibility of changing transmit modes, and indeed most parameters quickly; if not between pulses then at least on a scale of a few tens of pulses. Moreover, as in the case of the receive-side adaptivity, these modifications can take into account the knowledge of the environment gained about the scene.

The key features of a managed sensor system are that it senses the environment and chooses an appropriate waveform, beam-pattern, pulse repetition interval (PRI), etc (collectively called the *sensor mode*) to best extract the required information. Any such system must have, at least, the following components in addition to the basic sensor and ancillary components:

SM-1) A method of estimating the current (that is at the time of transmission of next pulse) state of the environment. This is done on the basis of prior measurements together with some model of the dynamics of the environment. It may be important to estimate not only the scatterers of interest (*targets*) but also those that are not of interest (*clutter*), since knowledge of the latter may be useful for determination of an optimal radar mode.

SM-2) A measure of effectiveness of each potential sensor mode. This should be a function of both the mode (as defined above) and of the environment, or at least the estimate of it mentioned in

SM-1). Most importantly, it should be based on the operational problem at hand.

SM-3) A library of modes from which the optimal mode is chosen. This might be just a finite library, but also might be an infinite parameterized family of, say, waveforms.

SM-4) A method for finding the optimal choice of mode over one or more epochs, based on the measure of effectiveness.

We note that, at its simplest, the optimization will be on an epoch by epoch basis (the so-called "greedy" or "myopic" approach). In this case, the mode is chosen just to optimize for the next epoch and defer consideration of future behavior. A more sophisticated system would look several epochs ahead in applying the measure of effectiveness, though it would also update the scheduling policy on an epoch by epoch basis. Such an approach is, *a priori*, very computer intensive, and much work is needed to develop shortcuts to calculation of the optimal policy. Sometimes it may be appropriate to choose to measure the effectiveness of a policy only at the last epoch of application of that policy.

It should be noted that this regime allows the possibility that the sensor is spread over several platforms and/or is comprised of several physically different sensors within each platform. It can encompass trajectory control for platforms and even control of data rates in connecting platforms to each other and to a central node. In each case the system can be viewed as consisting of many real or virtual sensors, where a virtual sensor can be a particular mode of a sensor, a position of a platform, a particular bit of a measurement made by a sensor, etc. Thus the sensor management problem may seen in all of these cases as one of choosing to switch between many different sensors, where the choice is made on the basis on knowledge of the environment. This view is schematically represented in Figure 1.

The ultimate goal of research in this area is to "close the loop" in radar signal processing by producing algorithms for scheduling of beam-directions, beam-shapes, waveforms and other radar modalities so as to optimally extract information from the environment (targets and clutter). Several sub-objectives contribute to this. As we have already said, in order to choose the best modality for a given radar environment, an estimate of that environment needs to be available at the time of making the selection, a method of assessing the effectiveness of a given modality in a given environment is required, as well as an optimal scheduling algorithm to make the selection of an optimal modality for each of a number of future epochs. Because of space constraints, we limit our discussion to

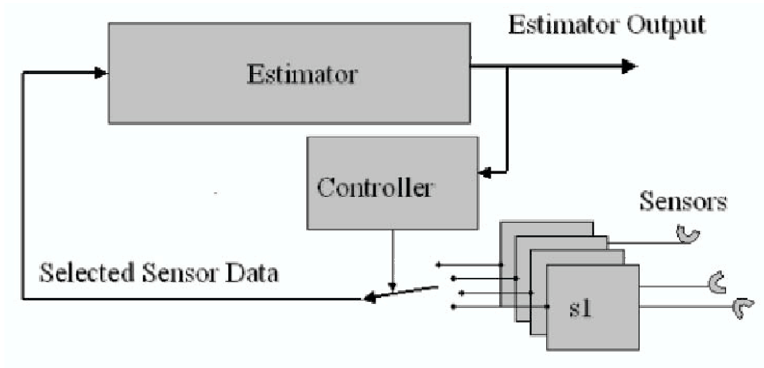

Figure 1. Schematic of Sensor Management

simulations for just one- and two-step ahead scheduling. Before proceeding to the simulation work, we discuss the theory of waveform libraries. The choice of the library of modes between which the sensor can switch is, of course, an important consideration in the development of scheduled radar systems.

4. Theory of Waveform Libraries

With the advent of radars capable of waveform agility, the design of optimal waveform libraries comes into question. The purpose of this section is to consider the design of such waveform libraries for radar tracking applications, from an information theoretic point of view. We note that waveform libraries will depend in general on the specific applications in which the systems are to be used. Airborne radars will require different libraries from ship-borne ones. Radars used in a tracking mode will require different optimal libraries than radars in a surveillance mode.

The idea of selecting waveforms adaptively based on tracking considerations was introduced in the papers of Kershaw and Evans [3, 4]. There they used a cost function based on the predicted track error covariance matrix.

In designing or improving a waveform library certain questions arise. Firstly it is important to establish the measure of effectiveness (MoE) for individual waveforms (cost function) and then to extend this to an MoE for the library. If a particular set of waveforms is added, will this improve the library in these terms and, on the other hand, how much will removing some waveforms reduce the utility of the library? It is the purpose of this chapter to develop an information theoretic framework for addressing such questions, at least from the target tracking point of view and to look at its application to specific waveform collections.

We use the basic sensor model proposed in [4]. While this has limitations, it is simple and therefore useful as a starting point for discussion of the problem. In this model, the sensor is characterized by a measurement noise covariance matrix which is waveform dependent

$$\mathbf{R}_\phi = \mathbf{T}^T \mathbf{J}_\phi^{-1} \mathbf{T}, \tag{9}$$

where $\mathbf{J}_\phi$ is the Fisher information matrix corresponding to the measurement using waveform $\phi \in L^2(\mathbf{R})$, and $\mathbf{T}$ is the transformation matrix between the time delay and Doppler measured by the receiver and the target range and velocity. The Fisher information matrix is given by an expression involving the normalized second order time and frequency moments of the waveform ϕ. It is also expressible in terms of the Hessian of the squared absolute value of the ambiguity function of the waveform at the origin of the range-Doppler plane. This calculation is done in [6].

It should be pointed out that the use of the Fisher matrix here is an approximation. It really corresponds to the Cramér-Rao lower bound on the estimator for the target from this measurement. It can be shown that the estimator here is *asymptotically efficient* (see[2], pp. 38–39) in that the covariance matrix approaches the Cramér-Rao lower bound over a large number of measurements (*loc. cit.*).

We note that the Hessian equivalence means that the Fisher matrix expresses purely local information about the ambiguity function at its peak. It says nothing about the structure of the ambiguity away from that peak. This local nature of the Fisher matrix is of some concern when considering its use in expressing a measure of effectiveness for a waveform. It can be argued, however, that this is a reasonable approach for tracking (where the return is "gated" in the vicinity of the predicted target position and Doppler) and in relatively low clutter situations. In a detection problem in a highly cluttered environment, the side-lobes will play a significant role and alternative measures of effectiveness ought to be considered.

In the context of our discussion in this chapter, we represent the measurement obtained using the waveform ϕ as a Gaussian measurement with covariance $\mathbf{R}_\phi$. 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

$$I(X;Y) = \log\det(\mathbf{I} + \mathbf{R}_\phi^{-1}\mathbf{P}). \tag{10}$$

This is the mutual information between the target variable (range and Doppler) X and the processed (with a matched filter) radar return Y, resulting from the use of the waveform ϕ. $\mathbf{I}$ is the identity matrix. We use this expected information as the MoE 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 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.

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

$$G_F(\mathcal{L}) = \int_{\mathbf{P}>0} \max_{\phi\in\mathcal{L}} \log\det(\mathbf{I} + \mathbf{R}_\phi^{-1}\mathbf{P})\, dF(\mathbf{P}). \tag{11}$$

Thus we have assumed that the optimal waveform is chosen in accordance with the MoE defined in equation (10) 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.

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 (11).

Having defined the utility of a waveform library we go on to investigate the utilities of a few libraries. Specifically, we consider libraries generated from a fixed waveform ϕ_0, usually an unmodulated pulse of some fixed duration, by *symplectic transformations*. Such transformations form a group of unitary transformations on $L^2(\mathbf{R})$ and include linear frequency modulation as well as the Fractional Fourier transform (FrFT) in a sense that we shall make clear.

Under such transformations $\phi = \mathbf{U}\phi_0$, the ambiguity function of the waveform ϕ_0, is modified according to the following equation.

$$|A_\phi(\mathbf{x})| = |A_{\phi_0}(\mathbf{S}^{-1}\mathbf{x})| \tag{12}$$

where $\mathbf{x} = (t, f)^T$ and $\det(\mathbf{S}) = 1$, and ϕ_0 ranges over all members of $L^2(\mathbf{R})$. Indeed, a reasonable definition of *symplectic transformation* in this context is any unitary operator on $L^2(\mathbf{R})$ that transforms the ambiguity function according to equation (12). There is a technical problem here that requires resolution. A waveform is *not* determined by the absolute value of its ambiguity. Thus there may be more than one transformation $\mathbf{S}$ under which equation (12) is valid. It turns out that in this case the the transformation is unique.

It is not hard to see that such transformations form a group. Suppose that U_1 and U_2 are symplectic in this sense and S_1 and S_2 correspond to them. Then

$$|A_{U_1 U_2 \phi_0}(\mathbf{x})| = |A_{U_2 \phi_0}(\mathbf{S}_1^{-1}\mathbf{x})| = |A_{\phi_0}(\mathbf{S}_2^{-1}\mathbf{S}_1^{-1}\mathbf{x})| = |A_{\phi_0}((\mathbf{S}_1\mathbf{S}_2)^{-1}\mathbf{x})|. \tag{13}$$

Furthermore, under symplectic transformations, it is relatively easy to show, using the Hessian formula for calculating the Fisher information matrix, that the measurement covariance matrix transforms as

$$\mathbf{R}_{U\phi_0} = \mathbf{S}^T \mathbf{R}_{\phi_0} \mathbf{S} \tag{14}$$

when S is associated with U.

An LFM (“chirp”) waveform library consists of

$$\mathcal{L}_{\text{chirp}} = \{\exp(i\lambda \mathbf{t}^2/2)\phi_0 \mid \lambda_{\min} \le \lambda \le \lambda_{\max}\} \tag{15}$$

where ϕ_0 is an unmodulated pulse, $\lambda_{\min}$ and λ_{max} are the minimum and maximum chirp rates supported by the radar, and $\mathbf{t}$ is the (unbounded) operator on $L^2(\mathbf{R})$ defined by

$$\mathbf{t}\phi(t) = t\phi(t). \tag{16}$$

It follows that

$$(\exp(i\lambda \mathbf{t}^2/2)\phi)(t) = \exp(i\lambda t^2/2)\phi(t). \tag{17}$$

For this library the corresponding measurement covariance matrices are given by (14) with

$$\mathbf{S}(\lambda) = \begin{pmatrix} 1 & 0 \\ \lambda & 1 \end{pmatrix}. \tag{18}$$

It is relatively easy to see that

$$\mathcal{L}'_{\text{chirp}} = \{\exp(i\lambda_{\min}\mathbf{t}^2/2)\phi_0, \exp(i\lambda_{\max}\mathbf{t}^2/2)\phi_0\} \tag{19}$$

is strongly equivalent to $\mathcal{L}_{\text{chirp}}$. That is, we do just as well if we keep only the LFMs with the minimum and maximum rates. In range-Doppler coordinates, the error covariance matrix for each LFM can be represented by

$$R(\lambda) = \mathbf{S}(\lambda)^T R_0 \mathbf{S} S(\lambda), \tag{20}$$

where R_0 is a diagonal matrix with ρ_1, ρ_2 on the diagonal; that is, a covariance matrix for the rectangular pulse [1, 4]. Direct computations give the following expression for the mutual information $I(X;Y)$:

$$I(X;Y) = 4\frac{P_{11}}{\rho_2}\frac{\lambda^2}{4} - 4\frac{P_{12}}{\rho_2}\frac{\lambda}{2} + \frac{|P|}{|R|} + 1 + \frac{P_{11}}{\rho_1} + \frac{P_{22}}{\rho_2}. \tag{21}$$

This is a quadratic in λ with positive second derivative since P and R are both positive definite, and therefore achieves its maximum at the end points, i.e at maximum or minimum allowed sweep rate.

Another way to create a waveform library is to take an ambiguity and rotate it. In this case, the new waveform is a fractional Fourier transform of the old one.

$$\mathcal{L}_{\mathrm{FrFT}} = \{\exp(i\theta(\mathbf{t}^2 + \mathbf{f}^2)/2)\phi_0 \mid \theta \in \Theta\}, \tag{22}$$

where the set $\Theta \subset [0, 2\pi]$ can be chosen so as not to violate the bandwidth constraints of the radar, and $\mathbf{f}$ is the operator on $L^2(\mathbf{R})$ defined by

$$\mathbf{f}\phi(t) = i\phi'(t). \tag{23}$$

For this library the corresponding transformation in range-Doppler space is given by the rotation

$$\mathbf{R}(\theta) = \begin{pmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{pmatrix}. \tag{24}$$

It is possible to consider combinations of the rotation and chirping transformations applied to an unmodulated waveform ϕ_0; that is, we consider all transformations of the following form:

$$\mathcal{L}_{\mathrm{FrFT}} = \{\exp(i\theta(\mathbf{t}^2 + \mathbf{f}^2)/2)\exp(i\lambda\mathbf{t}^2/2)\phi_0 \mid \lambda_{\min} \leq \lambda \leq \lambda_{\max}, \theta \in \Theta\} \tag{25}$$

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

$$\mathbf{f}\phi(t) = i\phi'(t), \tag{26}$$

where $\cdot'$ denotes differentiation in time. Note that $\mathbf{f}$ and $\mathbf{t}$ commute up to an extra additive term (the "canonical commutation relations"). To be precise,

$$[\mathbf{t}, \mathbf{f}] = \mathbf{tf} - \mathbf{ft} = -i\mathbf{I}. \tag{27}$$

For this library the corresponding measurement covariance matrices are given by (14) with

$$\mathbf{S}(\theta, \lambda) = \begin{pmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{pmatrix} \begin{pmatrix} 1 & 0 \\ \lambda & 1 \end{pmatrix}. \tag{28}$$

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. We can show that there exists a unique θ which maximizes the mutual information $I(X;Y)$ and, in a similar fashion to the pure chirp library case,

$$\mathcal{L}'_{\text{FrFT-chirp}} = \{\exp(i\theta(\mathbf{t}^2 + \mathbf{f}^2)/2)\exp(i\lambda\mathbf{t}^2/2)\phi_0 \mid \lambda \in \{\lambda_{\min}, \lambda_{\max}\}, \theta \in \Theta\} \quad (29)$$

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

5. Sensor scheduling simulations and result

Here we discuss simulations for sensor scheduling problems over up to two epochs into the future. The difficulties here reside in the design of the cost function and tracking of the scene. Our aim here is to show that sensor scheduling does, at least in simulation, achieve performance improvement.

Several aspects are common to all of the simulations described here. The scenarios all involve multiple maneuvering and crossing targets in simulated clutter. The simulated targets move according to an interacting multiple models (IMM) method; that is, at each epoch one of a finite number of dynamical models is chosen. The choice changes from epoch to epoch according to a Markov chain. Each of the dynamical models is linear. Process noise is, in each case, white and independent from epoch to epoch. Measurement is made using a waveform from a small finite library of waveforms, that we specify in each case.

A brief description of the tracking and waveform scheduling aspects of the scheme is as follows:

Tracking Since we are tracking multiple maneuvering targets, we use an iterated multiple modes (IMM) based tracker. This assumes that each target assumes at each epoch one of a finite number of dynamical models, such as “constant velocity”, “constant linear acceleration”, “fast left turn”, etc, and implements a filter for each such dynamical model. As is normal in IMM the dynamical model is assumed to evolve by means of a Markov chain. We remark that the models and transition matrices are not identical with those used in constructing the scene. All noise on the processes is assumed Gaussian and independent between epochs. Multiple targets and clutter are addressed by an integrated probabilistic data association tracker, specifically the LMIPDA-IMM algorithm described in [5]. This is a recursive algorithm combining a multi-target data association algorithm (LMIPDA) with manoeuvring target state estimation implemented using IMM. Each track carries along with it a "probability of track existence" which is updated at each epoch along with the track. In addition the probability of each dynamical model is updated from the measurements.

Waveform Scheduling The choice of measurement is made using the control variable $n(k)$. In fact two choices are made at each epoch, the target to be measured and the waveform used. The waveforms impinge on the measurement process through the covariance matrix of the noise $\omega_n^t(k)$. In this model, the sensor is characterised by a measurement noise covariance matrix which is waveform dependent

$$R_\phi = T^T J_\phi^{-1} T, \tag{30}$$

where J_ϕ is the Fisher information matrix corresponding to the measurement using waveform ϕ and T is the transformation matrix between the time delay and Doppler measured by the receiver and the target range and velocity. It is assumed that N different *measurement modes* are available for each target, each given by a measurement matrix H_n^t $n = 1, 2 \ldots, N$.

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. In the absence of measurements, as will be the case in the study of revisit times, the best we can do is to use current knowledge to predict forward and update the covariance matrix, dynamic model pdf and probability of track existence. The tracking and scheduling algorithms now becomes as follows:

- *IMM mixing* as in [5] is conducted as usual;

- *Forward prediction* is then performed separately for each dynamical model.

- *Covariance update:* this 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 *a priori* covariance matrix, if the target exists, is detected, and the measurement is received then the error covariance matrix is updated using the Kalman equation.

- The dynamic model and track existence pdfs are updated. If the target does not exist it produces no measurement; if it does and is detected the expected measurement pdf,dynamical model and track existence pdfs are using the LMIPDA-IMM filter.

- The next step is to combine the estimates for all dynamics models $j = 1, \ldots, M$ into one, using the standard "IMM combination" formulae [5]. We refer the interested reader to this paper for details.

5.1 One- and Two-Step Ahead Scheduling

Our first aim is to do a simple comparison of one-step and two-step scheduling of waveforms and other radar parameters. The modes of the radar system (beam-direction and waveform) are chosen for the next one or two PRIs based on the predicted scene over that time. We note that in the two-step case the choice of radar mode is updated on a PRI by PRI basis. A comparison between one and two-step ahead scheduling is an important issue, since if it is shown that the improvement achieved by two step ahead optimal scheduling over just one-step ahead scheduling is slight, it is reasonable to guess that one-step ahead scheduling is for practical purposes optimal. Since multi-step scheduling is inherently much more computationally intensive, it is best avoided if it results in only a marginal improvement. We emphasize that, of course, results of this kind are very likely to be scenario dependent unless there is some inherently mathematical reason why optimal multi-step ahead scheduling is achievable by a myopic approach. That would appear unlikely. We emphasize too that this work has been done on a simulator. The structure of the scene is highly artificial and the clutter models very simplistic.

We have compared one-step and two-step ahead scheduling using two performance measures. The first is the root mean square error of the track estimation; this is a fairly obvious measure of the performance of the tracker. The second measure was the number of track updates. Since the sensor is managed in such a way that track updating is done only when the predicted track error exceeds a threshold, this also gives a measure of how far the estimation process is diverging from the actual target state.

We refrain here from giving detailed descriptions of the experiments. Their outcome suggests that, in the presence of clutter, the tracking performance can be improved with multiple step ahead scheduling as opposed to one step ahead. The results are represented in Figures 2 and 3. One observes that for two steps ahead the tracking accuracy is improved, albeit slightly, while the number of times the track had to be

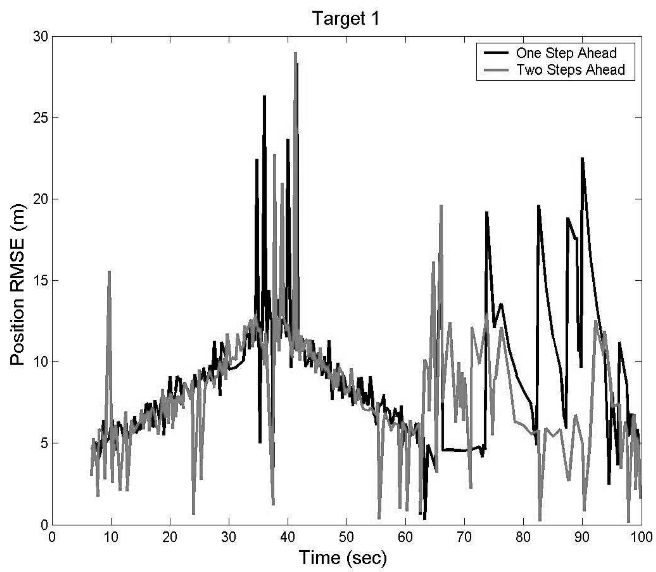

Figure 2. Root Mean Square Error (RMSE)

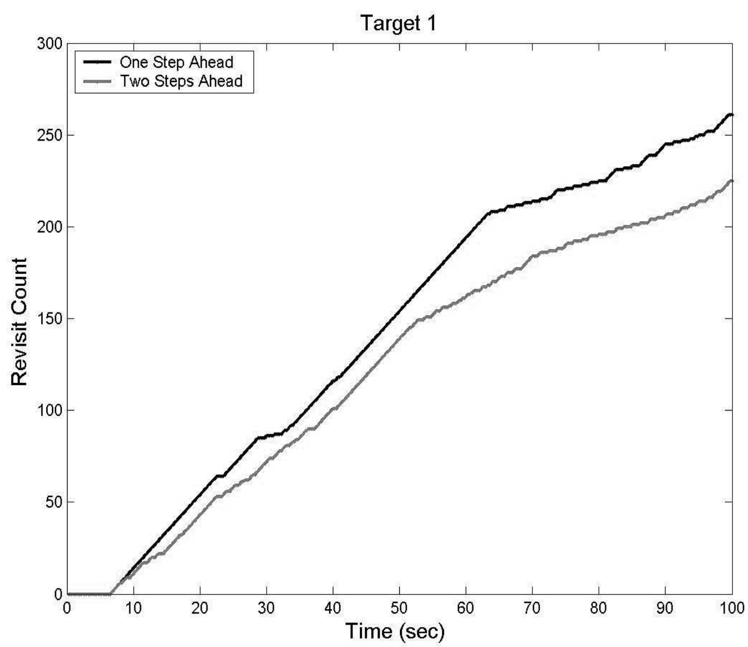

Figure 3. The number of track updates

updated is reduced. In both cases the improvement is not large, and is worse immediately after the aircraft has maneuvered. Once the aircraft has settled back into a linear model again the two-step ahead scheduler does better.

5.2 Scheduling of Waveform Libraries

The next series of experiments is focused on how the choice of waveform libraries affect the problem of tracking of maneuvering targets.

As in the previous experiments, at each epoch we would like to select a waveform (or really the error covariance matrix associated with a measurement using this waveform) so that the measurement will minimize the uncertainty of the dynamic model of the target. We study two possible measures: entropy of the *a posteriori* pdf of the models and *mutual information* between the dynamic model pdf and measurement history. Both of these involve making modifications to the LMIPDA-IMM approach that are described in [5]. Since we want to minimize the entropy *before* taking the measurement, we need to consider the *expected* value of the cost. To do this we replace the measurement $\mathbf{z}$ in the IMM equations by its expected value. In the case of the second measure, for a model we have

$$I(\Gamma; Z) = -\sum_{\gamma=1}^{M} P\{\gamma\} \log P\{\gamma\} + \int P\{\mathbf{z}\} \sum_{\gamma=1}^{M} P\{\gamma|\mathbf{z}\} \log P\{\gamma|\mathbf{z}\} d\mathbf{z}, \tag{31}$$

where $P\{\gamma\}$ is the *a priori* probability of the model $\gamma \in \Gamma$, and $\mathbf{z}$ is the measurement.

Simulations were performed for both cost functions. Target trajectories in range and Doppler were randomly created. The maneuvers for the trajectories were generated using a given transition probability matrix. We identified four maneuvers: 0 acceleration; 10m/s^2 acceleration; 50m/s^2 acceleration; -10m/s^2 acceleration.

In the experiments we considered rotation-LFM waveform libraries with 1 waveform (max upsweep chirp), 2 waveforms (max upsweep and max downsweep chirps), and 6 waveforms (maximum upsweep, maximum downsweep chirps and 2 rotations 0.2π and 0.4π as defined in equation (22) to the left for the maximum upsweep and maximum downsweep chirps).

The results are presented in Figures 4, 5, 6, and 7. Clearly, for either cost function, waveform scheduling using the six-waveform library outperforms waveform scheduling using the two-waveform library, which in turn outperforms no scheduling (one waveform) in both estimation accuracy (Figures 4 and 6) and correct identification of target maneuver (Figures 5 and 7).

5.3 Re-visit Time Scheduling

Finally in this section on simulations, we briefly describe a project that includes many of the ideas we have presented already. The crucial problem is to use scheduling to reduce the amount of time spent on tracking known targets while retaining a given level of track accuracy. By doing this we permit the sensor to spend more time in surveillance for new targets.

We postulate a radar system tracking T targets where T is a random variable $0 \leq T \leq T_0$ and the tth target is in state $x^t(k)$ 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 T_{scan} within every interval of length T_{total}. The remainder of the time is spent measuring targets being tracked. We aim to schedule revisit times to targets within these constraints.

At each epoch a target track and a beam direction have to be selected. The scheduler has a list $\Delta = \{\delta_1, \delta_2, \ldots, \delta_K\}$ of "revisit intervals". Each of the numbers δ_k is a number of epochs representing the possible times

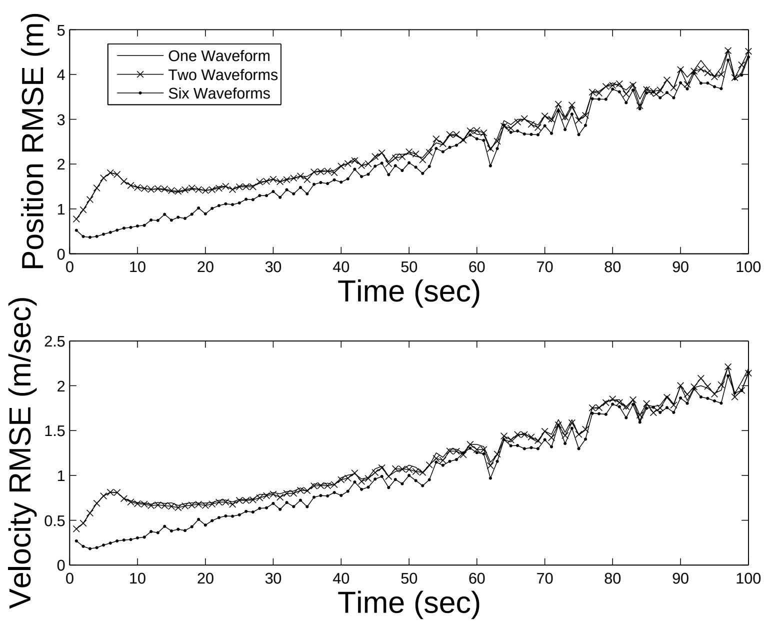

Figure 4. Root Mean Square Error for Entropy Cost

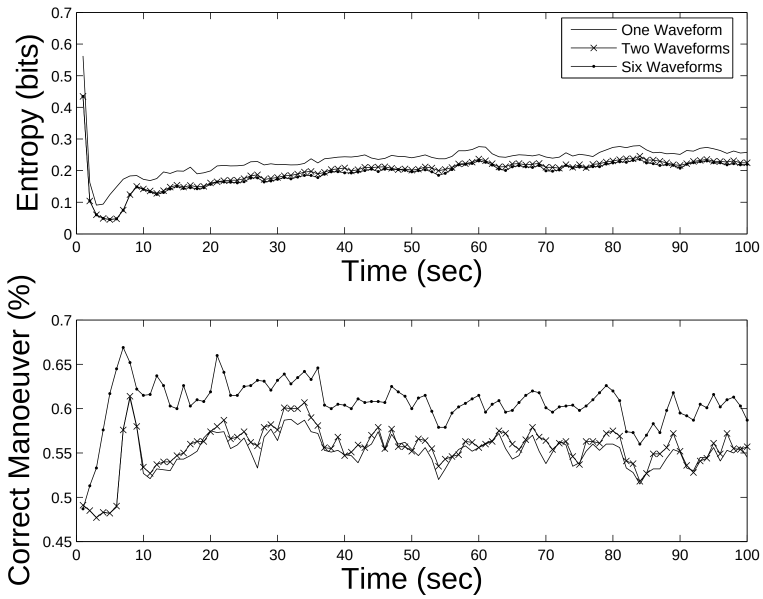

Figure 5. Cost Function and Correct Maneuver Identification for Entropy Cost

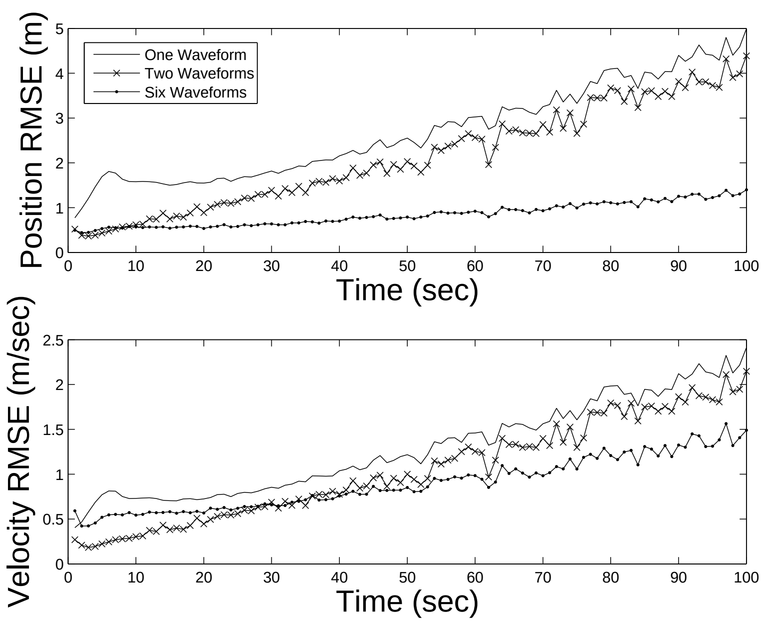

Figure 6. Root Mean Square Error for Mutual Information Cost

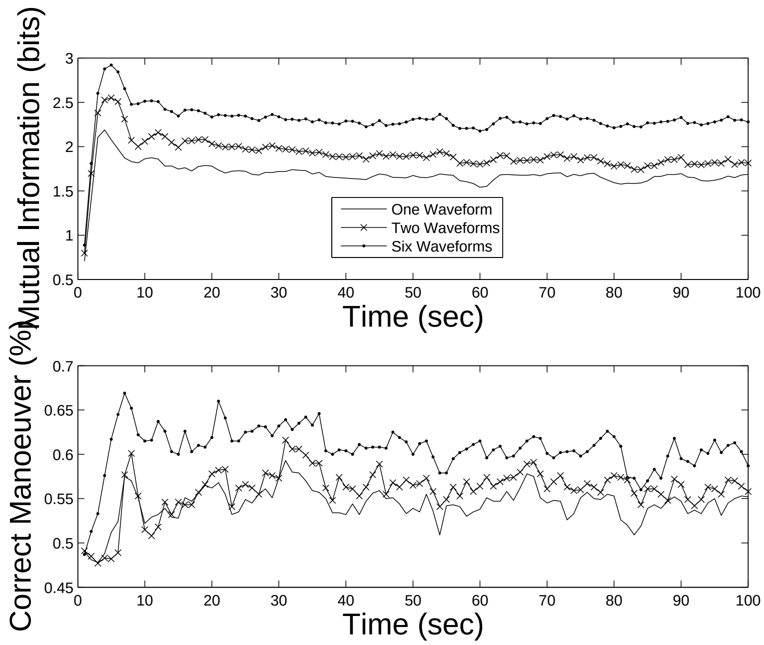

Figure 7. Cost Function and Correct Maneuver Identification for Mutual Information Cost

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, though the simulator permits target maneuvers on an epoch by epoch basis.

The LMIPDA-IMM calculations are performed for all combinations of revisit times in Δ and waveforms in the library. Evidently then 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. In other words, our objective is

$$\phi, \delta = \arg\max \Delta, \text{ subject to } |\det(P_{k|k})| \leq K. \quad (32)$$

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.

We note that for many manoeuvring targets there may be no solution to the scheduling problem that satisfies the constraints. However, we have not been able to simulate a situation in which this happens.

We have, on the other hand done simple simulations for the case of one-step ahead and two-step ahead scheduling. In the latter case, the revisit times and waveforms are calculated while the target states are propagated forward over two measurements, with the cost function being the absolute value of the determinant of the track error covariance after the second measurement. Only the first of these measurements is done before the revisit calculation is done again for that target, so that the second may never be implemented.

Simulations were performed to compare the effects of no scheduling with random choice of waveform against one-step and two-step ahead beam and waveform scheduling as described in the last section. 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. The waveforms were chosen at random from the three waveforms in the library. The simulated scene corresponded to a surveillance area of 15km by 15km contained two maneuvering land targets in stationary land clutter which had small random Doppler to simulate movement of vegetation in wind. The number of clutter measurements at each epoch was generated by samples from a Poisson distribution with mean ~ 5 per scan per sq.km. 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 these experiments we used the waveform library consisting of three waveforms: an up-sweep chirp, a down-sweep chirp and an unmodulated pulse. In the scheduling cases, surveillance time used approximately 80 percent of each scan period, the remaining 20% being allocated as described above to the maintenance of tracks of existing targets.

The outcome of 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 8. It should be observed in Figure 8 that RMS error was considerably worse especially during the early part of the simulation for the unscheduled case. In fact the RMS error in the unscheduled case is 5larger immediately after significant manoeuvres 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

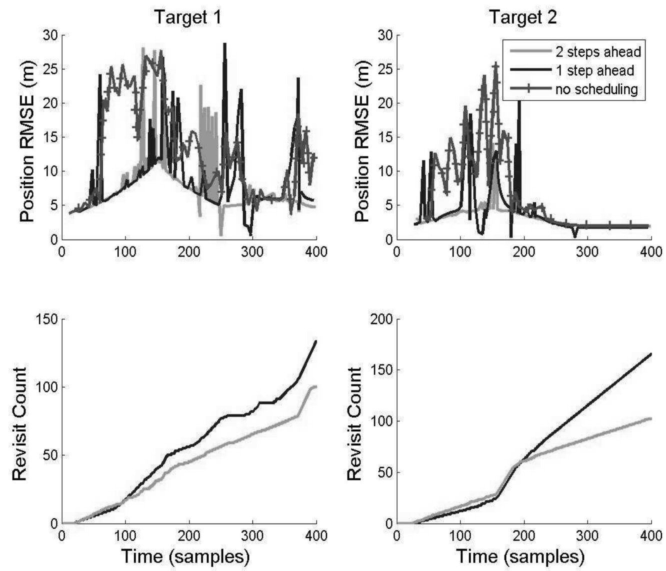

Figure 8. Root Mean Square Error (RMSE) and Revisit Count for one vs. two step ahead beam and waveform scheduling

ahead case, tracking accuracy is improved (top plots) slightly over the one-step ahead case but with a significant reduction in revisit times to maintain those tracks.

References

[1] A. Nehorai 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.

[2] Steven M. Kay. *Fundamentals of Statistical Signal Processing*. Prentice-Hall, 1993.

[3] D.J. Kershaw and R.J. Evans. Optimal waveform selection for tracking systems. *IEEE Transactions on Information Theory*, 40(5):1536–50, September 1994.

[4] D.J. Kershaw and R.J. Evans. Waveform selective probabilistic data association. *IEEE Aerospace and Electronic Systems*, 33(4):1180–88, October 1997.

[5] Darko Mušicki, Subhash Challa, and Sofia Suvorova. Multi target tracking of ground targets in clutter with LMIPDA-IMM. In *7th International Conference on Information Fusion*, Stockholm, Sweden, July 2004.

[6] H.L. van Trees. *Detection, Estimation and Modulation Theory, Part III.* Wiley, New York, 1971.