Dynamic Configuration of Time-Varying Waveforms for Agile Sensing and Tracking in Clutter

Sandeep P. Sira, *Student Member, IEEE*, Antonia Papandreou-Suppappola, *Senior Member, IEEE*, and Darryl Morrell, *Senior Member, IEEE*

***Abstract*—The advent of waveform-agile sensors has enabled the design of tracking systems where the transmitted waveform is changed on-the-fly in response to the tracker's requirements. This approach can provide performance improvements over individual optimization of the sensor waveform or the tracking algorithm. In this paper, we consider joint sensor configuration and tracking for the problem of tracking a single target in the presence of clutter using range and range-rate measurements obtained by waveform-agile, active sensors in a narrowband environment. We propose an algorithm to select and configure linear and nonlinear frequency-modulated waveforms to minimize the predicted mean square error (MSE) in the target state estimate; the MSE is predicted using the Cramér-Rao lower bound on the measurement error in conjunction with the unscented transform. We further extend our algorithm to match wideband environments, and we demonstrate the algorithm performance through a Monte Carlo simulation of a radar tracking example.**

***Index Terms*—Clutter, target-tracking, waveform-scheduling.**

I. Introduction

THE volume of data gathered by modern sensors often places an overwhelming demand on processing algorithms. The ability to intelligently direct such sensors to gather the most pertinent data can have significant impact on the performance of a system. In target tracking applications, for example, the advent of waveform-agile sensors, that can shape their transmitted waveforms on-the-fly, makes adaptive waveform configuration schemes possible. These schemes can change the transmitted waveform on a pulse-to-pulse basis to obtain the information that optimally improves the tracker's estimate of the target state. For example, these sensors can correspond to new radar systems currently under testing by the Navy with multiple transmit and receive elements capable of working independently. This system level optimization yields better tracking performance than optimizing the sensors and the tracker independently.

Manuscript received February 1, 2006; revised October 17, 2006. This work was supported in part by the Department of Defense, administered by the Air Force Office of Scientific Research by Grant No. AFOSR FA9550-05-1-0443, and by the DARPA, WAS program under NRL grant N00173-06-1-G006. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. J. Riba.

S. P. Sira and A. Papandreou-Suppappola are with the Department of Electrical Engineering, Arizona State University, Tempe, AZ 85287 USA (e-mail: ssira@asu.edu; papandreou@asu.edu).

D. Morrell is with the Department of Engineering, Arizona State University Polytechnic Campus, Mesa, AZ 85212 USA (e-mail: morrell@asu.edu).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TSP.2007.894418

The dynamic selection of waveforms for target tracking was first considered in [1], where the optimal waveform parameters were derived for tracking one-dimensional (1–D) target motion using a linear observations model with perfect detection in a clutter-free environment. The tracking was accomplished by a Kalman filter and waveforms with linear time-frequency characteristics were used. In this scenario, the problem of selecting the waveform to minimize the tracking error or the validation gate volume can be solved in closed form. This work was extended to include clutter and imperfect detection, but the linear observations model was still used [2].

Although waveform design optimization was investigated under different performance objectives, the applicability of the results to nonlinear target tracking applications is limited. For example, in [3] and [4], the tracking performance of different combinations of waveforms was compared using the expected value of the steady state estimation error. In this paper, the authors concluded that the upswept linear frequency-modulated chirp offers very good tracking performance. However, they did not consider dynamic adaptation of waveforms. Recently, the development of waveform libraries for target tracking was studied in [5]. Using an information theoretic criterion, the authors demonstrated that the maximum expected information about the target state could be obtained by a waveform whose duration was either its minimum or maximum allowable value. This finding, however, does not extend to other performance criteria and observations models, as we demonstrated in [6] and [7]. A combined optimization of the detection threshold and the transmitted waveform for target tracking was presented in [8], where a cost function based on the cumulative probability of track loss and the target state covariance was minimized. A strategy to select the optimal sequence of dwells of a phased array radar was developed in [9] by posing the problem as a partially observed Markov decision problem.

In the aforementioned works [1]–[5], [8], [9], little attention was paid to the application of the algorithms to tracking scenarios with nonlinear observations models. In addition, no systematic utilization of waveforms with varying time-frequency signatures was explored. In this paper, we present a configuration algorithm for waveform-agile sensors using generalized frequency-modulated (GFM) chirp signals in nonlinear scenarios. Specifically, we use a waveform library that is comprised of a number of chirps with different linear or nonlinear instantaneous frequencies. The proposed configuration algorithm simultaneously selects the phase function and configures the duration and frequency modulation (FM) rate of the transmitted waveforms. These waveforms are selected

and configured to minimize the predicted mean square error (MSE) in the tracker's estimate of the target state; the MSE is predicted using the Cramér-Rao lower bound (CRLB) on the measurement errors in conjunction with an unscented transform to linearize the observations model. We develop the configuration algorithms for both narrowband and wideband environments.

Our approach is generally applicable to waveform selection and configuration for multiple active sensors tracking a target. In our adopted scenario, two sensors generate range and Doppler measurements which are used by a tracker to estimate the Cartesian position and velocity coordinates of a single target moving in two dimensions. The received signal at the sensors includes reflections from clutter while additive noise causes imperfect detection of the target reflected signal. A particle filter with probabilistic data association is used to track the target. As we will demonstrate with numerical simulations, dynamic parameter selection for the transmitted waveforms improves tracking performance.

The paper is organized as follows. In Section II, we describe the GFM waveform structure and develop the narrowband model for the received signal. Section III describes the target dynamics, observations and clutter models, and develops the CRLB on the measurement noise covariance matrix. The target tracking algorithm using a particle filter with probabilistic data association is presented in Section IV while in Section V, we propose the waveform configuration algorithm. Section VI presents a simulation study while Section VII extends the algorithm to the wideband scenario.

II. Waveform Structure and Signal Model

A. Waveform Structure

We use a waveform library consisting of several types of GFM chirps. At each time epoch, each sensor transmits a pulse given by

$$s_T(t) = \sqrt{2}\mathrm{Re}\left[\sqrt{E_T}\tilde{s}(t)\exp(j2\pi f_c t)\right] \tag{1}$$

where f_c is the carrier frequency and E_T is the energy of the transmitted pulse. The complex envelope $\tilde{s}(t)$ is a GFM chirp that is defined as

$$\tilde{s}(t) = a(t)\exp\left(j2\pi b\xi\left(\frac{t}{t_r}\right)\right), \quad |t| \le \frac{T}{2} + t_f \tag{2}$$

where $a(t)$ is a trapezoidal envelope with rise/fall time $t_f \ll T/2$, b is a scalar FM rate parameter, $\xi(t/t_r)$ is the chirp phase function, and $t_r > 0$ is a reference time point. We can obtain different FM waveforms in our library with unique time-frequency signatures by choosing a phase function (and, thus, the waveform's instantaneous frequency $\zeta(t/t_r) = b\frac{d}{dt}\xi(t/t_r)$) from a set of possible phase functions. While the linear frequency-modulated chirp has been popular in radar and sonar, nonlinear frequency-modulated chirps can offer significant advantages [10], [11]. For example, waveforms with hyperbolic instantaneous frequency are Doppler-invariant and are similar to the signals used by bats and dolphins for echolocation [12]. As we will show, one advantage offered by nonlinear GFM chirps is minimal range-Doppler coupling which is an important feature in a tracking system.

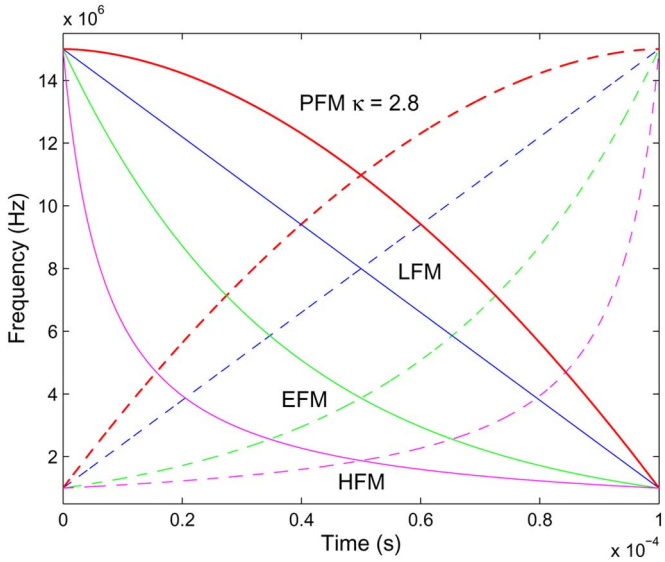

Fig. 1. Time-frequency plots of generalized FM chirp waveforms (downsweep—solid lines, upsweep—dotted lines) with duration 100 μs and frequency sweep 14 MHz.

TABLE I
FM Waveforms Used In The Configuration Scenarios

Waveform	Phase Function $\xi(t)$	Frequency sweep
LFM	$t/\gamma + (t+\lambda/2)^2/2$	$b\,\lambda$
PFM	$t/\gamma + (t+\lambda/2)^\kappa/\kappa$	$b\,\lambda^{\kappa-1}$
HFM	$\ln(t+\gamma+\lambda/2)$	$b\frac{\lambda}{\gamma(\gamma+\lambda)}$
EFM	$\exp(-(t+\lambda/2)/\gamma)$	$b\frac{1}{\gamma}(\mathrm{e}^{-\frac{\lambda}{\gamma}}-1)$

The waveforms considered in this paper are the linear FM (LFM), power FM (PFM), hyperbolic FM (HFM), and the exponential FM (EFM) chirps [10], [11]. We define the chirp duration as $\lambda = T + 2t_f$ and assume that $t_r = 1$ for the remainder of the paper. Note that the frequency sweep Δ_f can be specified by b, λ and $\xi(t)$. In order to completely specify a time-frequency signature for the waveform, we must also fix either the initial frequency $\zeta(-\lambda/2) = f_1$ or the final frequency $\zeta(\lambda/2) = f_2$. An additional parameter γ is, therefore, introduced in the phase functions of the considered waveforms as shown in Table I. For example, in the case of the LFM, $\zeta(-\lambda/2) = b/\gamma = f_1$ and $\zeta(\lambda/2) = f_1 + b\lambda = f_2$ can be used to uniquely determine b and γ once λ, Δ_f and either f_1 or f_2 are chosen. In this paper, we fix the larger of f_1 and f_2 and vary the remaining parameters to obtain various time-frequency characteristics. Fig. 1 shows the instantaneous frequencies of these waveforms for a duration of $\lambda = T + 2t_f = 100$ μs and a frequency sweep of $\Delta_f = 14$ MHz. Note that the PFM chirp is characterized by the power κ, and the LFM chirp is a special case of the PFM with $\kappa = 2$. We choose the time-varying amplitude $a(t)$ in (2) to have a trapezoidal structure as

$$a(t) = \begin{cases} \frac{\alpha}{t_f}\left(\frac{T}{2} + t_f + t\right), & \frac{-T}{2} - t_f \le t < \frac{-T}{2} \\ \alpha, & \frac{-T}{2} \le t < \frac{T}{2} \\ \frac{\alpha}{t_f}\left(\frac{T}{2} + t_f - t\right), & \frac{T}{2} \le t < \frac{T}{2} + t_f \end{cases} \tag{3}$$

where α is an amplitude chosen so that $\tilde{s}(t)$ in (2) has unit energy. A Gaussian envelope has often been used in the past [1]–[3], [5], [6], [8] due to the flexibility it affords in computing ambiguity functions and moments of the time and frequency distributions of the signal. However, the behavior of the phase function in the tails of the Gaussian envelope does not significantly affect the waveform's performance in range and Doppler estimation, since the tails contain only a small fraction of the signal energy. This can lead to biased conclusions on waveform performance when different phase functions are considered. Accordingly, a rectangular envelope would be most desirable, but the CRLB associated with such waveforms cannot be computed in closed form since the second moment of its spectrum is not finite [13]. Using a trapezoidal envelope as in (3) avoids the difficulties associated with the evaluation of the CRLB for waveforms with a rectangular envelope and yet approximates it well enough to provide a clearer comparison of phase function performance than using a Gaussian envelope [6].

B. Narrowband Signal Model

In this section, we develop the model for the received signal in a narrowband environment. When the transmitted waveform is reflected by the target, its velocity causes a Doppler scaling (compression or dilation) of the time of the complex envelope. The received signal is [14]

$$r(t) = s_R(t) + n(t) \tag{4}$$

where

$$s_R(t) = \sqrt{2}\mathrm{Re}\left[\sqrt{E_R}\tilde{s}\left(t - \tau - \frac{2\dot{r}}{c}t\right)e^{j2\pi f_c(t-\frac{2\dot{r}}{c}t)}\right] \tag{5}$$

and $n(t)$ is additive white Gaussian noise. In (5), $\tau = 2r/c$ is the delay of the received signal where r is the range of the target, c is the velocity of propagation, and E_R is the received signal energy. The radial velocity or range-rate of the target with respect to the sensor is $\dot{r}$, and we used the fact that $\dot{r} \ll c$. Under the assumption that the transmitter and receiver are at the same location, the received energy depends inversely on the fourth power of the target range [15]. We use this model for the SNR in our simulation study. Note that E_R will be a random variable due to the stochastic nature of the reflection and propagation processes [14]. For the purposes of this paper however, we treat E_R as nonrandom [1].

From (5), we note that the transmitted signal undergoes a Doppler scaling by the factor $1 - (2\dot{r}/c)$ as well as a shift in its carrier frequency. When the time-bandwidth product (TBP) of the waveform satisfies the narrowband condition

$$\mathrm{TBP} \ll \frac{c}{2\dot{r}} \tag{6}$$

the Doppler scaling can be approximated by a simple Doppler shift $\nu = -2f_c\dot{r}/c$ [14], and the received signal in (5) can be approximated by

$$s_R(t) \approx \sqrt{2}\mathrm{Re}\left[\sqrt{E_R}\tilde{s}(t-\tau)\exp\left(j2\pi(f_c t + \nu t)\right)\right]. \tag{7}$$

In other words, the approximation says that all frequencies in the signal are shifted equally. In the case of radar, the narrowband condition is easily met as the velocity of propagation is very large ($c \approx 3 \times 10^8$ m/s).

III. Target Dynamics and Measurement Models

We seek to estimate the motion of a target that is free to move in two dimensions. The target state is given by $\mathbf{X}_k = [x_k \quad y_k \quad \dot{x}_k \quad \dot{y}_k]^T$, where x_k and y_k correspond to the position, and $\dot{x}_k$ and $\dot{y}_k$ to the velocity, at time k in Cartesian coordinates. We use the superscript ^{T} to denote the transpose throughout the paper.

A. Target Dynamics

The target dynamics are modeled by a linear, constant velocity model given by

$$\mathbf{X}_k = F\mathbf{X}_{k-1} + \mathbf{W}_k. \tag{8}$$

The process noise is modeled by the uncorrelated Gaussian sequence $\mathbf{W}_k$. The constant matrix F and the process noise covariance Q are given by

$$F = \begin{bmatrix} 1 & 0 & \delta t & 0 \\ 0 & 1 & 0 & \delta t \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad Q = q\begin{bmatrix} \frac{\delta t^3}{3} & 0 & \frac{\delta t^2}{2} & 0 \\ 0 & \frac{\delta t^3}{3} & 0 & \frac{\delta t^2}{2} \\ \frac{\delta t^2}{2} & 0 & \delta t & 0 \\ 0 & \frac{\delta t^2}{2} & 0 & \delta t \end{bmatrix} \tag{9}$$

respectively, where δt is the sampling interval and q is a constant. The target motion is observed by two active sensors, Sensor A and Sensor B, which measure the range and range-rate of the target. The sensors are waveform-agile and can change their waveform from pulse-to-pulse.

B. Measurement Model

The received signal in (5) is processed by a bank of matched filters that yield maximum likelihood estimates of the delay and Doppler of the received signal. These, in turn, provide estimates of the range and range-rate of the target. With respect to sensor i, i = A or B, the nonlinear relation between the target state and its range and range-rate is given by $h^i(\mathbf{X}_k) = [r_k^i \quad \dot{r}_k^i]^T$, where

$$r_k^i = \sqrt{(x_k - x^i)^2 + (y_k - y^i)^2}$$
$$\dot{r}_k^i = \frac{(\dot{x}_k(x_k - x^i) + \dot{y}_k(y_k - y^i))}{r_k^i}$$

and sensor i is located at (x^i, y^i). The measurement originated from the target is given by

$$\mathbf{z}_{k,i} = h^i(\mathbf{X}_k) + \mathbf{V}_k^i \tag{10}$$

where the measurement error is modeled by the white Gaussian sequence $\mathbf{V}_k^i$. We assume that the measurement errors are uncorrelated with the process noise in (8).

C. Measurement Noise Covariance

We define the waveform parameter vector for sensor i at time k as $\boldsymbol{\theta}_k^i = [\xi_k^i \lambda_k^i b_k^i]^T$ where $\lambda_k^i = T + 2t_f$ is the signal duration. The impact of the waveform parameter vector on the tracker accuracy depends on the measurement errors in (10), which in turn depend on the resolution properties of the waveform. These properties are quantified by the covariance matrix $N(\boldsymbol{\theta}_k^i)$ of the process $\mathbf{V}_k^i$ in (10). Following [1], we approximate $N(\boldsymbol{\theta}_k^i)$ using the CRLB of the range and range-rate estimators for the waveform parameters. As we are configuring waveforms to be transmitted in a narrowband environment, we use the narrowband ambiguity function that is defined as [14]

$$AF_{\tilde{s}}(\tau,\nu) = \int_{-\infty}^{\infty} \tilde{s}\left(t+\frac{\tau}{2}\right)\tilde{s}^*\left(t-\frac{\tau}{2}\right)\exp(-j2\pi\nu t)dt. \tag{11}$$

The CRLB of the matched filter estimator can be obtained by inverting the Fisher information matrix which is the Hessian of the ambiguity function in (11), evaluated at the true target delay and Doppler [14]. Equivalently, the elements of the Fisher information matrix can also be determined directly from the complex envelope of the waveform. For the time-varying GFM chirps, we provide this computation and the resulting $N(\boldsymbol{\theta}_k^i)$ in Appendix .

It is important to note that the CRLB depends only on the properties of the ambiguity function at the origin. The location of the peak of the ambiguity function in the delay-Doppler plane is affected by its sidelobes which are not considered in the evaluation of the CRLB, thus limiting its effectiveness as a measure of the waveform's estimation performance. Another method of computing the measurement error covariance with explicit dependence on the sidelobes was proposed in [3]. This method is based on the notion of a resolution cell that encloses the ambiguity function contour at a given probability of detection with the true target location assumed to be distributed uniformly within the cell. However, the size of the resolution cell and the associated measurement error covariance increases with the probability of detection and thus the signal-to-noise ratio (SNR). It is therefore not feasible for use in an adaptive scheme. If high SNR is assumed, the sidelobes of the ambiguity function may be neglected and the estimator can be assumed to achieve the CRLB. We make this assumption and set $N(\boldsymbol{\theta}_k^i)$ equal to the CRLB, which is computed using the waveform in (2) with parameter $\boldsymbol{\theta}_k^i$.

D. Clutter Model

The measurement at time k consists of false alarms due to clutter as well as a single target return, if it is detected. For sensor i, the measurement is given by

$$\mathbf{Z}_k^i = \left[\mathbf{z}_{k,i}^1, \mathbf{z}_{k,i}^2, \ldots, \mathbf{z}_{k,i}^{m_k^i}\right] \tag{12}$$

where m_k^i is the number of measurements obtained at sensor i at time k and each $\mathbf{z}_{k,i}^m, m = 1,\ldots,m_k^i$, consists of a range and range-rate value. We also define $\mathbf{Z}_k = [\,\mathbf{Z}_k^A \quad \mathbf{Z}_k^B\,]$ as the measurement vector from both sensors at time k. The probability that the target is detected at time k by sensor i (or that $\mathbf{Z}_k^i$ in (12) includes a target-originated measurement) is given by $P_{d_k}^i$.

We assume that the number of false alarms follows a Poisson distribution with average ρV_k^i, where ρ is the density of the clutter and V_k^i is the validation gate volume. The validation gate is a region in the observation space centered on the next predicted observation. Only observations that fall within this region are validated as having been potentially originated by the target and are used in the tracking process [16]. The probability that m false alarms are obtained is

$$\mu(m) = \frac{\exp(-\rho V_k^i)(\rho V_k^i)^m}{m!}. \tag{13}$$

We also assume that the clutter is uniformly distributed in the volume V_k^i. The detector in each sensor compares the squared magnitude of the matched-filtered output to a threshold. Since the noise in the received signal in (2) is additive white Gaussian, the test statistic under the noise-only and target-present hypotheses follows an exponential distribution. Under these circumstances, it may be shown [3] that the probability of detection at time k can be modeled according to

$$P_{d_k}^i = P_f^{1/1+\eta_k^i} \tag{14}$$

where P_f is the desired probability of false alarm and η_k^i is the SNR at sensor i. Note that in (14) we have used the fact that the peak of the ambiguity function is used as an estimator of delay and Doppler shift.

IV. Target Tracking Using a Particle Filter

We seek to recursively estimate the probability distribution $p(\mathbf{X}_k|\mathbf{Z}_{1:k}, \boldsymbol{\theta}_{1:k})$ of the target state given the sequence of observations $\mathbf{Z}_{1:k}$ and waveform vectors $\boldsymbol{\theta}_{1:k}$ up to time k, where $\boldsymbol{\theta}_k = \left[\,\boldsymbol{\theta}_k^{A^T} \boldsymbol{\theta}_k^{B^T}\,\right]^T$ is a combined waveform parameter vector for both sensors at time k. The conditional mean of this density yields the state estimate. Since the observations are nonlinear functions of the target state, we use a particle filter as the tracker [17]. This filter propagates a set of N_s particles and corresponding weights, $\{\mathbf{X}_k^j, w_k^j\}, j = 1,\ldots,N_s$, in accordance with the kinematics model and the likelihood function. At each sampling instant, the particles are drawn from an importance density $q(\cdot)$, and the weights are updated as

$$w_k^j \propto w_{k-1}^j \frac{p\left(\mathbf{Z}_k|\mathbf{X}_k^j, \boldsymbol{\theta}_k\right) p\left(\mathbf{X}_k^j|\mathbf{X}_{k-1}^j\right)}{q\left(\mathbf{X}_k^j|\mathbf{X}_{k-1}^j, \mathbf{Z}_k, \boldsymbol{\theta}_k\right)}. \tag{15}$$

The kinematic prior $p(\mathbf{X}_k|\mathbf{X}_{k-1})$ is often used as the importance density, and as a result, the weights are proportional to the likelihood. For the observation model described in Section III, the measurements obtained by each sensor are independent and the likelihood function is

$$p\left(\mathbf{Z}_k|\mathbf{X}_k, \boldsymbol{\theta}_k\right) = \prod_{i=A,B} p\left(\mathbf{Z}_k^i|\mathbf{X}_k, \boldsymbol{\theta}_k^i\right). \tag{16}$$

We assume that the target generates at most one measurement at each sensor. Accordingly, if m_k^i measurements are received at sensor i at time k, they must be comprised of either one target-originated measurement together with $m_k^i - 1$ false alarms, or

m_k^i false alarms. In the former case, each measurement can be associated *a priori* with the target with equal probability. Given that the mth measurement is target-originated, the probability that it takes the value $\mathbf{z}_{k,i}^m$ can be evaluated from the measurement model in (10). The likelihood function for a sensor thus incorporates a contribution from each measurement, weighed by its measurement-to-target association probability. It can be shown that the likelihood function using probabilistic data association [16] is given by [18]

$$p(\mathbf{Z}_k^i|\mathbf{X}_k,\boldsymbol{\theta}_k^i) = (1-P_{d_k}^i)\mu(m_k^i)(V_k^i)^{-m_k^i} + P_{d_k}^i(V_k^i)^{-(m_k^i-1)}\mu(m_k^i-1)\frac{1}{m_k^i}\sum_{m=1}^{m_k^i} p(\mathbf{z}_{k,i}^m|\mathbf{X}_k,\boldsymbol{\theta}_k^i). \quad (17)$$

V. Dynamic Waveform Selection

The criterion we use for the dynamic waveform selection is the minimization of the tracking MSE. Note that we did not use the maximization of an information theoretic criterion that was recently applied to waveform selection [5].

The predicted squared error is a random variable, and we seek to minimize its expected value. This is because, when this selection is carried out, we do not have access to either the observation or the target state at the next sampling instant. We, therefore, attempt to minimize the cost function

$$J(\boldsymbol{\theta}_k) = E_{\mathbf{X}_k,\mathbf{Z}_k|\mathbf{Z}_{1:k-1}}\left((\mathbf{X}_k-\hat{\mathbf{X}}_k)^T\boldsymbol{\Lambda}(\mathbf{X}_k-\hat{\mathbf{X}}_k)\right) \quad (18)$$

over the space of allowable waveforms with parameter vectors $\boldsymbol{\theta}_k$. Here, $E(\cdot)$ is an expectation over predicted states and observations, $\boldsymbol{\Lambda}$ is a weighting matrix that ensures that the units of the cost function are consistent, and $\hat{\mathbf{X}}_k$ is the estimate of $\mathbf{X}_k$ given the sequence of observations $\mathbf{Z}_{1:k}$. Note that the cost in (18) results in a one-step ahead or myopic optimization. Although it is possible to formulate a nonmyopic cost function, the computational complexity associated with its minimization grows with the horizon of interest. The cost in (18) cannot be computed in closed form due to the nonlinear relationship between the target state and the measurement. We now present two methods of approximating it and obtaining the waveform that yields the lowest approximation cost.

A. Stochastic Approximation

The expectation in (18) can be expanded as

$$J_k(\boldsymbol{\theta}_k) = \int\int (\mathbf{X}_k-\hat{\mathbf{X}}_k)^T\boldsymbol{\Lambda}(\mathbf{X}_k-\hat{\mathbf{X}}_k) p(\mathbf{Z}_k|\mathbf{X}_k,\boldsymbol{\theta}_k)p(\mathbf{X}_k|\mathbf{Z}_{1:k-1},\boldsymbol{\theta}_{1:k-1})d\mathbf{X}_k d\mathbf{Z}_k. \quad (19)$$

It can be approximated using Monte Carlo (MC) methods for large N_x and N_z as

$$J_k(\boldsymbol{\theta}_k) \approx \frac{1}{N_x}\sum_{n=1}^{N_x}\frac{1}{N_z}\sum_{p=1}^{N_z}\left(\mathbf{X}_n-\hat{\mathbf{X}}_{n|\mathbf{Z}_p,\mathbf{Z}_{1:k-1}}\right)^T\boldsymbol{\Lambda}\left(\mathbf{X}_n-\hat{\mathbf{X}}_{n|\mathbf{Z}_p,\mathbf{Z}_{1:k-1}}\right) \quad (20)$$

where $\mathbf{X}_n, n = 1,\ldots,N_x$, are independent samples of predicted states drawn from the estimate of the density $p(\mathbf{X}_k|\mathbf{Z}_{1:k-1},\boldsymbol{\theta}_{1:k-1})$, which is obtained from the tracker, and $\mathbf{Z}_p, p = 1,\ldots,N_z$, are predicted observations drawn independently from the likelihood $p(\mathbf{Z}_k|\mathbf{X}_n,\boldsymbol{\theta}_k)$. The estimate of $\mathbf{X}_n$, given the observations $\mathbf{Z}_{1:k-1}$ and $\mathbf{Z}_p$, is denoted as $\hat{\mathbf{X}}_{n|\mathbf{Z}_p,\mathbf{Z}_{1:k-1}}$ and could be computed by a secondary particle filter.

The search for the waveform parameter that minimizes (20) can be accomplished by an iterative stochastic-steepest descent method. We note, however, that the gradient of (20) cannot be determined in closed form. As a result, we employ simultaneous perturbation stochastic approximation (SPSA) which is a finite differences approach for evaluating gradients of functions that are only observable through simulation in the presence of noise [19]. Using the approximation of (19) by (20) to evaluate the SPSA gradient, we can compute the estimate of the minimizing value of $\boldsymbol{\theta}_k$ [20].

This approach suffers from two drawbacks. First, the MC approximation in (20) is computationally intensive. This is due to the fact that each sample in the average requires an estimate of the predicted state $\hat{\mathbf{X}}_{n|Z_p,\mathbf{Z}_{1:k-1}}$ in (20), which is obtained using a particle filter. In addition, each iteration of the stochastic-steepest descent algorithm requires at least two evaluations of the cost function to calculate the gradient. Typically however, averaging is required to reduce the noise in the gradient estimate, and this further increases the computational burden. Secondly, a large number of iterations is required to achieve convergence due to the noise that accompanies each gradient estimate. The available range of the duration parameter is small, and this poses further challenges in achieving an accurate solution. We found the SPSA approach to be computationally intractable; thus, we propose a more computationally efficient method based on the unscented transform.

B. Unscented Transform Based Approximation

The Kalman filter covariance update equation provides a mechanism to recursively compute the covariance of the state estimate provided the observations and dynamics models are linear [21]. The observations model in (10) does not satisfy this requirement but it can be linearized and this approach can still be applied. The standard approach to this problem is the extended Kalman filter (EKF) which approximates the nonlinearity in (10) by a Taylor series expansion about a nominal target state and discards the higher order terms. An improvement on this method is the unscented transform [22]. The resulting unscented Kalman filter assumes that the density of the state given the observations is Gaussian and employs the unscented transform to compute its statistics under a nonlinear transformation. This approach has been shown to outperform the EKF [22]. Thus, we use the covariance update of the unscented Kalman filter to approximate the cost function as follows.

Let $P_{k-1|k-1}$ represent the covariance of the state estimate at time $k-1$. We wish to approximate the covariance $P_{k|k}(\boldsymbol{\theta}_k)$ that would be obtained if a waveform characterized by its parameter vector $\boldsymbol{\theta}_k$ was used to obtain a measurement at time k. First, the

dynamics model in (8) is used to obtain the predicted mean and covariance as

$$\hat{\mathbf{X}}_{k|k-1} = F\hat{\mathbf{X}}_{k-1|k-1} \text{ and}$$
$$P_{k|k-1} = FP_{k-1|k-1}F^T + Q \tag{21}$$

respectively. Next, we select $N+1$ sigma points $\boldsymbol{\mathcal{X}}_n$, $n = 0, 1, \ldots, N$, and corresponding weights $\mathcal{W}_n$ [22], to approximate the random variable $\hat{\mathbf{X}}_{k|k-1}$. For sensor i, a transformed set of sigma points $\boldsymbol{\mathcal{Z}}_n^i = h^i(\boldsymbol{\mathcal{X}}_n)$ is computed. Then, we calculate the covariances

$$P_{zz}^i = \sum_{n=0}^{N+1} \mathcal{W}_n \left(\boldsymbol{\mathcal{Z}}_n^i - \bar{\boldsymbol{\mathcal{Z}}}\right)\left(\boldsymbol{\mathcal{Z}}_n^i - \bar{\boldsymbol{\mathcal{Z}}}\right)^T$$
$$P_{xz}^i = \sum_{n=0}^{N+1} \mathcal{W}_n \left(\boldsymbol{\mathcal{X}}_n - \bar{\boldsymbol{\mathcal{X}}}\right)\left(\boldsymbol{\mathcal{Z}}_n^i - \bar{\boldsymbol{\mathcal{Z}}}\right)^T$$

where

$$\bar{\boldsymbol{\mathcal{Z}}} = \sum_{n=0}^{N+1} \mathcal{W}_n \boldsymbol{\mathcal{Z}}_n^i \text{ and } \bar{\boldsymbol{\mathcal{X}}} = \sum_{n=0}^{N+1} \mathcal{W}_n \boldsymbol{\mathcal{X}}_n.$$

Note that the expectation in (18) is also over all realizations of the clutter and thus over all possible measurement-to-target associations. The calculation of $P_{k|k}(\boldsymbol{\theta}_k)$ should reflect this fact. Intuitively, if there was no uncertainty in the origin of the measurements, all the information contained in the target-originated measurement could be unambiguously used to reduce the covariance of the state estimate. As is evident from (17), each measurement, whether target-generated or not, contributes to the update. The measurements that are not target-generated thus limit the reduction in the covariance during the update step. When probabilistic data association is used, it has been shown that the expected covariance of the state after an update with the measurement at sensor i is given by [23]

$$P_{k|k}^i\left(\boldsymbol{\theta}_k^i\right) = P_{k|k-1} - q_{2_k}^i P_{k|k}^c \tag{22}$$

where $P_{k|k}^c$ is the update due to the true measurement. Using the unscented transform, this update is given by

$$P_{k|k}^c = P_{xz}^i \left(P_{zz}^i + \left(N\left(\boldsymbol{\theta}_k^i\right)\right)^{-1} {P_{xz}^i}^T. \tag{23}$$

The scalar $q_{2_k}^i$ in (22) lies between 0 and 1 and is called the information reduction factor. It depends upon $P_{d_k}^i$, ρ and V_k^i, and serves to lower the reduction in the covariance that would have been obtained if there was no uncertainty in the origin of the measurements due to clutter. The computation of the information reduction factor involves a complicated integral which has to be evaluated by Monte Carlo methods. However, in this paper, we use some approximations that are available in [24].

Note that the information reduction factor is different for each sensor. This requires the update in (22) to be performed sequentially for each sensor. Accordingly, we first use the measurement of Sensor A to obtain $P_{k|k}^A(\boldsymbol{\theta}_k^A)$ using (22) and (23). A second update, using the measurements of Sensor B, is then carried out with $P_{k|k-1} = P_{k|k}^A(\boldsymbol{\theta}_k^A)$ in (22) to yield $P_{k|k}^B(\boldsymbol{\theta}_k^B)$. The second update requires another computation of all the covariance matrices in (23) using the unscented transform with $i = B$. Note that $P_{k|k}^B(\boldsymbol{\theta}_k^B)$ is the final predicted covariance $P_{k|k}(\boldsymbol{\theta}_k)$, as it includes information from both sensors A and B. The approximate cost $J(\boldsymbol{\theta}_k)$ in (18) is then computed as the trace of $\boldsymbol{\Lambda} P_{k|k}(\boldsymbol{\theta}_k)$.

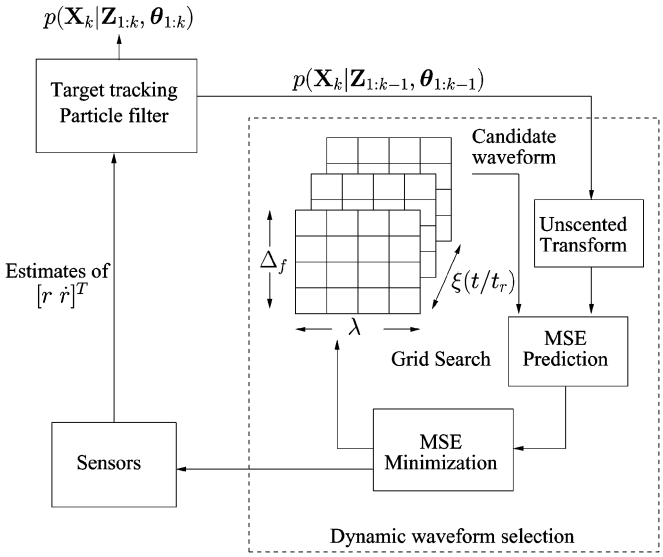

Fig. 2. Block diagram of the waveform selection algorithm.

C. Algorithm for Waveform Selection

The overall configuration and tracking algorithm is shown in Fig. 2. The dynamic waveform selection is performed by a search over the space of allowable waveforms for the candidate that results in the lowest cost, which is then chosen as the sensor configuration for the next sampling instant. While gradient-based methods could be used to optimize the duration and FM rate, that are continuous parameters, we chose to use a rectangular grid search approach over a finite search space since it is computationally cheaper and works well in practice. For each sensor, R grid points for λ and L grid points for Δ_f are evenly spaced over the intervals $[\lambda_{\min}, \lambda_{\max}]$ and $[0, \mathrm{BW}]$, respectively. Here, $\lambda_{\min}$ and $\lambda_{\max}$ are bounds that are determined by the constraints on the pulse duration, and BW is the maximum allowed frequency sweep. The values of the FM rates for each λ, corresponding to the phase function $\xi(t)$, are calculated for each frequency sweep Δ_f. We also consider the upsweep and downsweep of the frequency for each configuration as in Fig. 1, which leads to $2RL$ possible configurations per sensor for each phase function $\xi(t)$. The grid for a single phase function thus contains $(2\mathrm{RL})^2$ points.

The expected cost is computed at each grid point using the procedure described in Section V-B. The values of duration and frequency sweep that minimize the expected cost form the center of a new grid whose boundaries are taken as the immediate surrounding grid points. The expected cost is computed at each point on the new grid, and this procedure is repeated several times to find the configuration for each phase function that minimizes the expected cost. All combinations of these phase-specific configurations are now tested to determine the one that yields the lowest cost across all phase functions. A statement of this algorithm is provided in Table II. Note that this is a suboptimal version of a more exhaustive grid search that tests every possible combination of waveforms and their

TABLE II
ALGORITHM TO OBTAIN THE VALUE OF $\boldsymbol{\theta}_k$ THAT MINIMIZES THE PREDICTED COST $J(\boldsymbol{\theta}_k)$

- Compute $\hat{\mathbf{X}}_{k|k-1}$ and $P_{k|k-1}$ as in Equation (21).
- Compute $q_{2_k}^i$ and η_k^i, i = A,B as in Sections V-B and VI, respectively.
- For each waveform phase function ξ_s, $s = 1,\ldots,S$, set $\xi^A = \xi^B = \xi_s$ and obtain $(\boldsymbol{\theta}_s^A,\ \boldsymbol{\theta}_s^B)$ that minimizes the expected cost using the Grid Search algorithm.
- Find the waveform combination that minimizes $J(\boldsymbol{\theta})$:

$$\boldsymbol{\theta}^* = \arg\min_{s,s'} J(\boldsymbol{\theta}),$$

where $\boldsymbol{\theta} = [\boldsymbol{\theta}_s^{A^T}\ \boldsymbol{\theta}_{s'}^{B^T}]^T$, and $s, s' = 1,\ldots,S$.

Grid Search Algorithm

(i) Set $\lambda_{high}^A = \lambda_{high}^B = \lambda_{max}$, $\lambda_{low}^A = \lambda_{low}^B = \lambda_{min}$, $\Delta_{f_{high}}^A = \Delta_{f_{high}}^B$ = BW, and $\Delta_{f_{low}}^A = \Delta_{f_{low}}^B = 0$.
(ii) Set $\lambda^A = (\lambda_{high}^A + \lambda_{low}^A)/2$ and $\Delta_f^A = (\Delta_{f_{high}}^A + \Delta_{f_{low}}^A)/2$.
(iii) Repeat (iv)-(ix) below 5 times or until $J(\boldsymbol{\theta})$ stops decreasing.
(iv) Set i = A and i' = B.
(v) Use the unscented transform and compute P_{zz}^i, P_{xz}^i, as in Section V-B and $N(\boldsymbol{\theta}^i)$ as in Section III-C.
(vi) Compute $P_{k|k}^i(\boldsymbol{\theta}^i)$ as in (22), and $P_{zz}^{i'}, P_{xz}^{i'}$, as in Section V-B.
(vii) Perform a rectangular grid search over $\lambda^{i'} \in [\lambda_{low}^{i'}, \lambda_{high}^{i'}]$ and $\Delta_f^{i'} \in [\Delta_{f_{low}}^{i'}, \Delta_{f_{high}}^{i'}]$ to find $\boldsymbol{\theta}^{i'}$ that minimizes the expected cost $J(\boldsymbol{\theta})$, where $\boldsymbol{\theta} = [\boldsymbol{\theta}^{A^T}\ \boldsymbol{\theta}^{B^T}]^T$.
(viii) Set $\lambda_{low}^{i'}, \lambda_{high}^{i'}$ to the immediate grid points surrounding $\lambda^{i'}$, and similarly set $\Delta_{f_{low}}^{i'}, \Delta_{f_{high}}^{i'}$.
(ix) If i = A, exchange i and i'. Go to step (v).
(x) Return $\boldsymbol{\theta}^* = \boldsymbol{\theta}$.

(discretized) parameters for each sensor. Such a search, however, would entail a prohibitive computational expense.

D. Significance of the Weighting Matrix $\boldsymbol{\Lambda}$ in (18)

The weighting matrix $\boldsymbol{\Lambda}$ in (5) provides a mechanism to modify the effect of position and velocity variances on the waveform selection, as well as a mapping to cost function units. We have examined cases where the cost function depends on position variance only ($\boldsymbol{\Lambda} = \mathrm{diag}[1,1,0,0]$), and on velocity variance only ($\boldsymbol{\Lambda} = \mathrm{diag}[0,0,1,1]$). While the waveform selection in each case does not differ significantly from that obtained when $\boldsymbol{\Lambda} = \mathrm{diag}[1,1,1s^2,1s^2]$, the tracking performance degrades noticeably. This is due to the fact that the kinematics model in (8) uses the position *and* velocity estimates at time $k-1$ to estimate the position at time k, thus, introducing errors in position when the velocity is poorly estimated. This motivates the need for accurate velocity estimation. On the other hand, a poor position estimate at time $k-1$ would cause a poor estimate at time k irrespective of the accuracy of the velocity estimate. This motivates the need for position variance to influence the cost function, and thus the waveform selection as well.

VI. SIMULATIONS

Our simulation study is based on a radar tracking example in which the proposed dynamic waveform configuration is applied and successfully demonstrated. The GFM waveform library available to each sensor is shown in Table I and Fig. 1. The signal duration was constrained to lie in the range $\lambda_k^i \in [10\ \mu\mathrm{s}, 100\ \mu\mathrm{s}]$ with the rise/fall time of the envelope in (3) set to $t_f = 10$ ns. The maximum bandwidth allowed was 15 MHz. The sampling interval in (9) was $\delta t = 250$ ms resulting in 25 time steps, while the velocity of propagation was $c = 2.997925\times10^8$ m/s. The carrier frequency was $f_c = 10.4$ GHz. The target trajectory was generated with the initial condition $\mathbf{X}_0 = [0\,0\ 100\ 800]^T$ with $q = 0.1$ in (9). Sensors A and B were located at (0, -15,000) m and (15,631, 4,995) m, respectively. The SNR at each sensor was modeled according to $\eta_k^i = (r_0/r_k^i)^4$ where r_0 was the range at which a 0 dB SNR was obtained. In this example, r_0 was set to 50 km. The probability of false alarm was $P_f = 0.01$ and the validation gate was taken to be the 4-sigma region around the predicted observation [16]. The covariance of the initial estimate provided to the tracker was $P_0 = \mathrm{diag}[1000, 1000, 50, 50]$. The weighting matrix in (18) was set to $\boldsymbol{\Lambda} = \mathrm{diag}[1,1,1s^2,1s^2]$ so that the cost function in (18) had units of m^2. All results were averaged over 500 Monte Carlo simulations, conditioned on convergence. A track was classified as converged when no more than four continuous target-generated measurements fell outside the validation gate for either sensor [16].

A. Example 1—Waveform With Fixed Phase Function

In the first example, we consider the performance of the dynamic parameter configuration algorithm when the waveform library available to the sensors contains only the LFM chirp. The duration and FM rate (or equivalently the frequency sweep) of the transmitted LFM chirp are dynamically configured, and the performance is compared to that obtained when the parameters are fixed. For the fixed LFM chirps, we consider the shortest and longest allowable durations, and the frequency sweep is at its maximum. The averaged MSE for converged tracks, is shown in Fig. 3 for two clutter densities: $\rho = 0.0001$ and $\rho = 0.001$ false alarms per unit validation gate volume. As can be seen from the figure, the configuration algorithm provides improved performance. For example, when $k = 20$, the average error when using the configured waveform is about 10 dB less than that obtained by using the waveform with the largest time-bandwidth product, which is often a conventional choice in radar. It is interesting to note that bats exhibit dynamic adaptation of their transmitted sonar waveforms and use a variety of constant-frequency and FM chirps as they forage or hunt [25], [26]. Similar to our configuration algorithm, they do not necessarily use the waveform with the largest possible time-bandwidth product, opting instead for different pulse lengths, bandwidths and modulation, at different phases of the approach and capture of their prey.

Waveform adaptation also improves the number of tracks that converge. Specifically, when the fixed waveforms were used, an average of 25% and 38% of tracks did not converge in the case of the 10 μs pulse and the 100 μs pulse, respectively. When the waveform was dynamically configured, however, only 0.7% of the tracks did not converge. This can be attributed to the

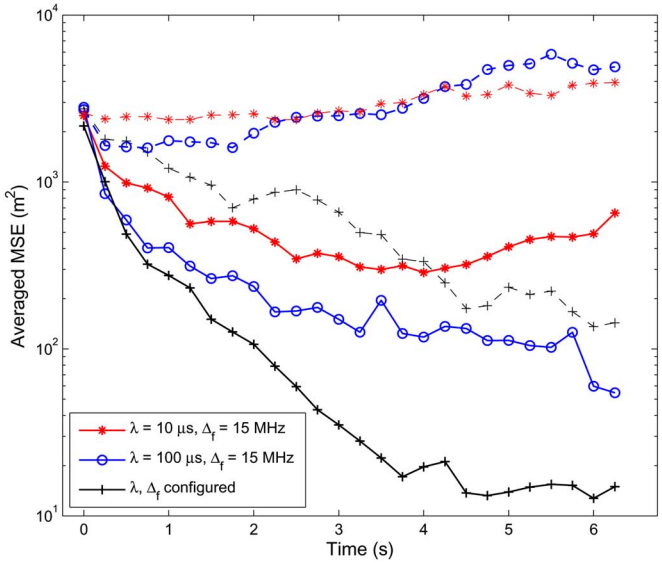

Fig. 3. Comparison of averaged MSE using fixed and configured LFM chirps only. Clutter density $\rho = 0.0001$ (solid lines) and $\rho = 0.001$ (dotted lines).

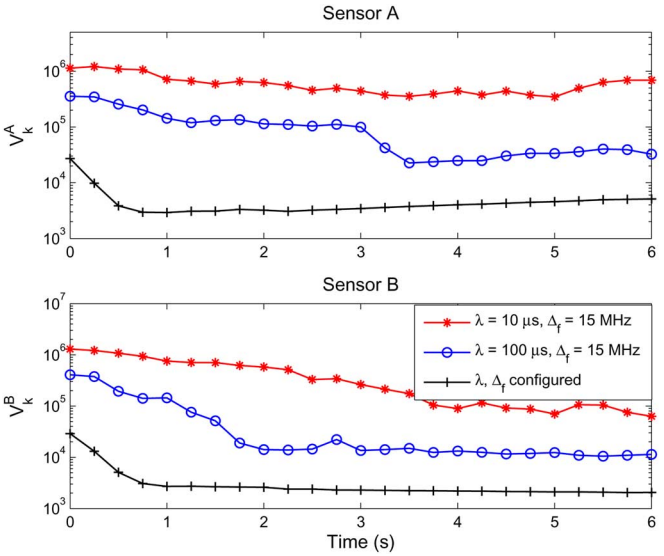

Fig. 4. Comparison of averaged validation gate volume using fixed and configured LFM chirps only with clutter density $\rho = 0.0001$.

number of false alarms that appear in the measurements at each sensor, which in turn is directly proportional to the validation gate volume. In Fig. 4, we compare the volume for the three LFM waveforms for $\rho = 0.0001$, and we find that the configured waveform results in a much lower volume for both sensors. The validation gate volume V_k^i for sensor i, $i = \mathrm{A}, \mathrm{B}$, is proportional to the determinant of the matrix $(P_{zz}^i + N(\boldsymbol{\theta}_k^i))$ in (23). As the search over the space of waveforms proceeds, it is apparent that this determinant will be minimized. In addition, the information reduction factor q_2^i in (22) increases as the validation gate volume decreases [23]. Thus, a waveform that results in a smaller validation gate also contributes more towards reducing the covariance of the state estimate.

B. Example 2—Waveforms With Agile Phase Function

In this simulation, in addition to dynamically configuring the duration and FM rate, we also select the phase function $\xi(t/t_r)$ in (2) of the generalized chirp waveforms with $t_r = 1$. The set of waveforms available to the sensors includes the PFM chirp with varying κ, the HFM and EFM chirps as described in Table I. For the simulations, we use values of κ of 2, 2.8, and 3. Recall that $\kappa = 2$ yields an LFM waveform. For this example, the clutter density was $\rho = 0.0001$ false alarms per unit validation gate volume. Fig. 5 shows the averaged MSE for converged tracks, that is obtained when the phase function is dynamically selected, in addition to the waveform duration and FM rate. In order to determine whether the agility in the phase function results in performance improvement, we also compare the performance obtained by configuring the duration and FM rate for each waveform, when it is the only one in the waveform library. Accordingly, we obtained the averaged MSE for each individual waveform as for the LFM chirp in Example 2 and these are shown in Fig. 5.

We find that there is some improvement in the tracking performance when the phase function is dynamically selected. For example, when $k > 9$, the waveform with agile phase function results in a 2–3 dB reduction in the MSE over those with fixed

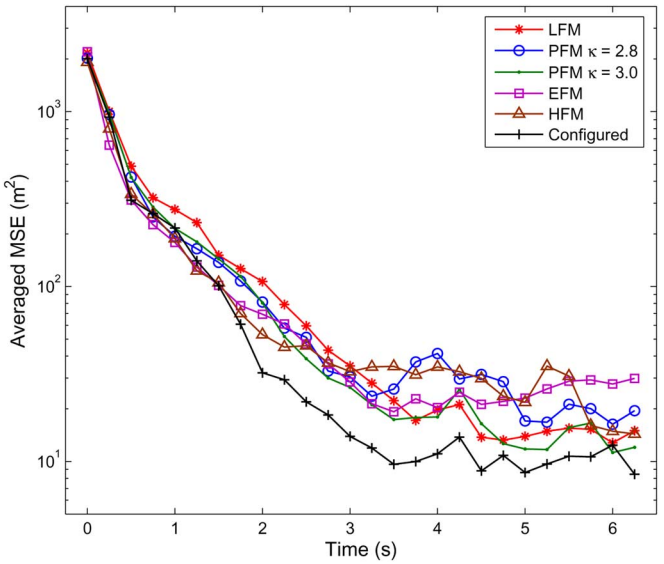

Fig. 5. Comparison of averaged MSE, conditioned on convergence, with and without phase function agility.

phase functions. A typical waveform selection for one simulation run is shown in Fig. 6. We note that the HFM and the EFM chirps are chosen at the first few instants while the PFM chirp with $\kappa = 3$ (parabolic instantaneous frequency) is used later. The LFM chirp and the PFM chirp with $\kappa = 2.8$ are not selected at all in any of the simulation runs.

In order to understand the waveform selection behavior, note that when the tracking starts, the range and range-rate of the target are poorly known and must be simultaneously estimated. Waveforms that offer little correlation between the estimation errors should accordingly be selected. Fig. 7 shows the correlation coefficient for the waveforms available to the tracker for a duration of 100 μs and a frequency sweep of 15 MHz. Note that the HFM and EFM chirps have smaller correlation coefficients than the PFM chirps. When the range and range-rate estimates of the targets improve, it is possible to utilize the correlation

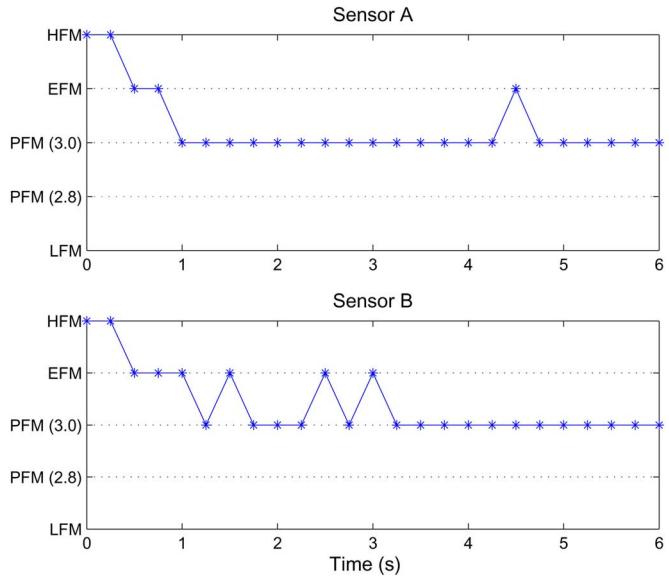

Fig. 6. Typical waveform selection when the phase function is configured together with the duration and FM rate of the waveform.

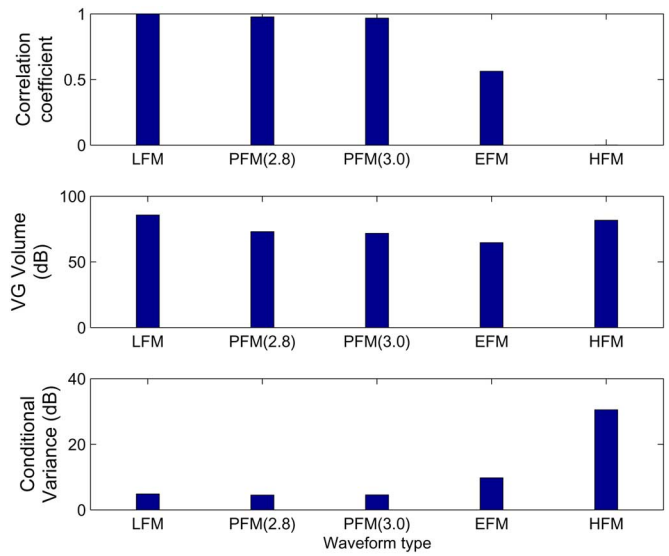

Fig. 7. Correlation coefficient (top), validation gate (VG) volume (middle), and conditional variance of range given range-rate at 10 dB SNR for the waveforms used in Section VI-B. The correlation coefficient for HFM is negligibly small.

between the estimation errors to improve the tracking performance. At this stage, waveforms that provide considerable correlation become more suitable. Thus, the conditional variance of the range given the range-rate becomes an important feature. As seen in Fig. 7, the PFM chirps offer lower values than the EFM and HFM chirps. However, minimization of the validation gate volume also occurs during the dynamic selection procedure. The volume for $P_{zz} = \mathrm{diag}[5, -50, -50, 1000]$ and an SNR of 20 dB is also shown in Fig. 7. Note that the PFM chirp with $\kappa = 3$ offers a lower volume than the LFM chirp and is therefore always selected during the later stages of the tracking sequence.

VII. Extension To Wideband Environments

When the narrowband condition imposed by (6) is not satisfied, the received signal model in (7) cannot be used since the scaling of the signal envelope is not negligible. This is especially true in sonar applications where $c \approx 1{,}500\ \mathrm{m/s}$, and with target velocities in the order of 10 m/s, we must have TBP $\ll 75$ to justify the narrowband assumption. As the TBP of sonar signals is generally much higher than 100, wideband processing is necessary.

In such situations, the received signal is given by

$$s_R(t) = \sqrt{2}\mathrm{Re}\left[\sqrt{E_R}\tilde{s}(\sigma t - \tau)\exp\left(j2\pi(f_c\sigma t)\right)\right] \qquad (24)$$

where $\sigma = (c - \dot{r})/(c + \dot{r}) > 0$ is the Doppler scale and $\tau = 2r/(c+\dot{r})$ is the delay. Since $\dot{r} \ll c$, we have $r \approx c\tau/2$ and $\dot{r} \approx c(1-\sigma)/2$. It can be shown that (7) is an approximation of (24) which holds when (6) is satisfied [14].

When the environment causes wideband changes in the signal, we must use the processing tool to match it. In this case, the wideband ambiguity (WAF) function provides a description of the correlation properties of the signal, and it is given by [27]

$$WAF_{\tilde{s}}(\tau, \sigma) = \sqrt{\sigma}\int_{-\infty}^{\infty}\tilde{s}(t)\tilde{s}^*(\sigma t + \tau)dt. \qquad (25)$$

As in the narrowband case, the CRLB can be evaluated by inverting the Hessian of the wideband ambiguity function and evaluating it at the true target delay and scale. Alternatively, the elements of the Fisher information matrix can be derived directly from the waveform as shown in Appendix . The wideband waveform selection algorithm has been applied to a sonar tracking example where perfect detection without clutter was assumed [7]. For the HFM chirp, the results indicated that the algorithm provides nearly 10 dB improvement in averaged MSE over fixed parameter waveforms.

VIII. Conclusion

In this paper, we presented a dynamic waveform selection and configuration algorithm for waveform-agile sensors, in which the waveforms are selected and configured to minimize a tracking error. In contrast with past research on this problem, we consider a nonlinear observations model to track a target moving in two dimensions as well as the use of waveforms with nonlinear time-frequency signatures. Our method of computing the predicted MSE is based on the unscented transform and the CRLB on the measurement error. We used a particle filter with probabilistic data association as the tracker; the particle filter provides robustness with respect to observation nonlinearities and sensor positioning.

The algorithm can be extended to multiple-step ahead or nonmyopic waveform scheduling by defining a cost function as a sum of the expected cost in (5) over a specified time horizon. Since two parameters, range and range-rate, are estimated at each instant, it is believed that scheduling two steps ahead may provide improved performance. The effect of the environment can also be incorporated by using matched generalized ambiguity functions [10] to obtain the CRLB on the measurement error for the given scenario. For example, dispersive environments may be considered in this framework.

One of the shortcomings of our approach is the requirement of high SNR which permits the sidelobes of the ambiguity function to be ignored. Such conditions cannot be guaranteed, and

a method of relating the signal to the measurement errors that does not rely on the CRLB would be a significant extension to our work. We also recognize that the false alarms due to clutter have been implicitly assumed to be uncorrelated at each sampling instant. This is an assumption often taken in waveform selection problems as well as in tracking [2], [16]. However, clutter cannot always be assumed uncorrelated. For example, in maritime applications, clutter was shown to exhibit spatial and temporal correlations [28]. Thus, further research is necessary to examine the waveform selection problem in this context.

APPENDIX

Let I denote the Fisher information matrix corresponding to the estimation of the delay and Doppler shift $[\tau\ \nu]^T$. For the narrowband case, the elements of I are given by [14]

$$\begin{aligned} I_{1,1} &= -\frac{\partial^2 AF_{\tilde{s}}(\tau,\nu)}{\partial\tau^2}\bigg|_{\substack{\tau=0\\ \nu=0}} \\ &= \int_{-\lambda/2}^{\lambda/2} (\dot{a}^2(t) + a^2(t)\Omega^2(t))dt \\ &\quad - \left[\int_{-\lambda/2}^{\lambda/2} a^2(t)\Omega(t)dt\right]^2 \\ I_{1,2} &= -\frac{\partial^2 AF_{\tilde{s}}(\tau,\nu)}{\partial\tau\partial\nu}\bigg|_{\substack{\tau=0\\ \nu=0}} = \int_{-\lambda/2}^{\lambda/2} ta^2(t)\Omega^2(t)dt \\ I_{2,2} &= -\frac{\partial^2 AF_{\tilde{s}}(\tau,\nu)}{\partial\nu^2}\bigg|_{\substack{\tau=0\\ \nu=0}} = \int_{-\lambda/2}^{\lambda/2} t^2a^2(t)dt \end{aligned}$$

where $\Omega(t) = 2\pi(b\frac{d}{dt}\xi(t/t_r) + f_c)$ and $I_{2,1} = I_{1,2}$. The CRLB for the measurement of $[\,r\ \dot{r}\,]$ at sensor i is then given by

$$N(\boldsymbol{\theta}^i) = \frac{1}{\eta}\Gamma I^{-1}\Gamma \tag{26}$$

where $\Gamma = \mathrm{diag}[c/2, c/(2f_c)]$ and η is the SNR.

For the wideband case, $I_{1,1}$ is identical to that for the narrowband case even though the wideband ambiguity function is used to compute it. The other elements of I are [29]

$$\begin{aligned} I_{1,2} &= -\frac{\partial^2 WAF_{\tilde{s}}(\tau,\sigma)}{\partial\tau\partial\sigma}\bigg|_{\substack{\tau=0\\ \sigma=1}} \\ &= \int_{-\lambda/2}^{\lambda/2} t(\dot{a}^2(t) + a^2(t)\Omega^2(t))dt \\ &\quad - \int_{-\lambda/2}^{\lambda/2} a^2(t)\Omega(t)dt \cdot \int_{-\lambda/2}^{\lambda/2} ta^2(t)\Omega(t)dt \\ I_{2,2} &= -\frac{\partial^2 WAF_{\tilde{s}}(\tau,\sigma)}{\partial\sigma^2}\bigg|_{\substack{\tau=0\\ \sigma=1}} \\ &= \int_{-\lambda/2}^{\lambda/2} t^2(\dot{a}^2(t) + a^2(t)\Omega^2(t))dt \\ &\quad - \left[\int_{-\lambda/2}^{\lambda/2} ta^2(t)\Omega(t)dt\right]^2 - \frac{1}{4}. \end{aligned}$$

In this case also, the CRLB is given by (26) but with $\Gamma = \mathrm{diag}[c/2, c/2]$.

REFERENCES

[1] D. J. Kershaw and R. J. Evans, "Optimal waveform selection for tracking systems," *IEEE Trans. Inf. Theory*, vol. 40, pp. 1536–1550, Sep. 1994.

[2] ——, "Waveform selective probabilistic data association," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 33, pp. 1180–1188, Oct. 1997.

[3] C. Rago, P. Willett, and Y. Bar-Shalom, "Detection-tracking performance with combined waveforms," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 34, pp. 612–624, Apr. 1998.

[4] R. Niu, P. Willett, and Y. Bar-Shalom, "Tracking considerations in selection of radar waveform for range and range-rate measurements," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 38, pp. 467–487, Apr. 2002.

[5] S. D. Howard, S. Suvorova, and W. Moran, "Waveform libraries for radar tracking applications," presented at the Int. Conf. Waveform Diversity and Design, Edinburgh, U.K., Nov. 2004.

[6] S. P. Sira, A. Papandreou-Suppappola, and D. Morrell, "Time-varying waveform selection and configuration for agile sensors in tracking applications," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, Mar. 2005, vol. 5, pp. 881–884.

[7] S. P. Sira, A. Papandreou-Suppappola, and D. Morrell, "Waveform scheduling in wideband environments," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, Toulouse, France, May 2006, vol. 1, pp. I-697–I-700.

[8] S. M. Hong, R. J. Evans, and H. S. Shin, "Optimization of waveform and detection threshold for range and range-rate tracking in clutter," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 41, no. 1, pp. 17–33, Jan. 2005.

[9] B. F. L. Scala, W. Moran, and R. J. Evans, "Optimal adaptive waveform selection for target detection," in *Proc. Int. Conf. Radar*, 2003, pp. 492–496.

[10] A. Papandreou-Suppappola, R. L. Murray, B. G. Iem, and G. F. Boudreaux-Bartels, "Group delay shift covariant quadratic time-frequency representations," *IEEE Trans. Signal Process.*, vol. 49, no. 11, pp. 2549–2564, Nov. 2001.

[11] A. Papandreou-Suppappola, "Time-varying processing: Tutorial on principles and practice," in *Applications in Time-Frequency Signal Processing*, A. Papandreou-Suppappola, Ed. Boca Raton, FL: CRC, 2002, pp. 1–84.

[12] R. A. Altes and E. L. Titlebaum, "Bat signals as optimally Doppler tolerant waveforms," *J. Acoust. Soc. Amer.*, vol. 48, pp. 1014–1020, Oct. 1970.

[13] C. E. Cook and M. Bernfeld, *Radar Signals: An Introduction to Theory and Application*. Boston, MA: Artech House, 1993.

[14] H. L. Van Trees, *Detection Estimation and Modulation Theory, Part III*. New York: Wiley, 1971.

[15] M. I. Skolnik, *Radar Handbook*, M. I. Skolnik, Ed., 2nd ed ed. New York: McGraw-Hill, 1990.

[16] Y. Bar-Shalom and T. E. Fortmann, *Tracking and Data Association*. Boston, MA: Academic, 1988.

[17] M. S. Arulampalam, S. Maskell, N. Gordon, and T. Clapp, "A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking," *IEEE Trans. Signal Process.*, vol. 50, no. 2, pp. 174–188, Feb. 2002.

[18] R. Niu, P. Willett, and Y. Bar-Shalom, "Matrix CRLB scaling due to measurements of uncertain origin," *IEEE Trans. Signal Process.*, vol. 49, no. 7, pp. 1325–1335, Jul. 2001.

[19] J. C. Spall, "Multivariate stochastic approximation using a simultaneous perturbation gradient approximation," *IEEE Trans. Autom. Control*, vol. 37, pp. 332–341, Mar. 1992.

[20] S. P. Sira, D. Morrell, and A. Papandreou-Suppappola, "Waveform design and scheduling for agile sensors for target tracking," in *Proc. Asilomar Conf. Signals, Syst. Comput.*, Nov. 2004, vol. 1, pp. 820–824.

[21] B. D. Anderson and J. B. Moore, *Optimal Filtering*. Englewood Cliffs, NJ: Prentice-Hall, 1979.

[22] S. Julier and J. Uhlmann, "A new extension of the Kalman filter to nonlinear systems," presented at the Int. Symp. Aerosp./Defense Sens., Simulation and Controls, Orlando, FL, 1997.

[23] T. Fortmann, Y. Bar-Shalom, M. Scheffe, and S. Gelfand, "Detection thresholds for tracking in clutter—A connection between estimation and signal processing," *IEEE Trans. Autom. Control*, vol. 30, pp. 221–229, Mar. 1985.

[24] D. J. Kershaw and R. J. Evans, "A contribution to performance prediction for probabilistic data association tracking filters," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 32, pp. 1143–1147, Jul. 1996.
[25] J. A. Simmons and R. Stein, "Acoustic imaging in bat sonar: Echolocation signals and the evolution of echolocation," *J. Comparative Psychol.*, vol. 135, no. 1, pp. 61–84, Mar. 1980.
[26] J. A. Simmons, P. A. Saillant, and S. P. Dear, "Through a bat's ear," *IEEE Spectrum*, vol. 29, pp. 46–48, Mar. 1992.
[27] D. A. Swick, "A review of wideband ambiguity functions," Naval Res. Lab., Washington, DC, Tech. Rep. 6994, 1969.
[28] K. Ward, C. Baker, and S. Watts, "Maritime surveillance radar part 1: Radar scattering from the ocean surface," *Proc. Inst. Elect. Eng.*, vol. 137, no. Pt. F, 2, pp. 51–62, Apr. 1990.
[29] D. W. Ricker, *Echo Signal Processing*. Boston, MA: Kluwer Academic, 2003.

Sandeep P. Sira (S05) received the M.Tech. degree in electrical engineering from the Indian Institute of Technology, Kanpur, in 1999.

He is a Graduate Student and Research Assistant with the Department of Electrical Engineering, Arizona State University, Tempe, where he is pursuing the Ph.D. degree. His research interests include waveform-agile sensing, target tracking, and detection and estimation theory. He was a commissioned officer in the Corps of Signals, Indian Army, from 1988 to 2003.

Antonia Papandreou-Suppappola (S'87–M'91–SM'03) received the Ph.D. degree in electrical engineering in 1995 from the University of Rhode Island.

She held a Navy-supported research faculty position with the University of Rhode Island. She is currently an Associate Professor at Arizona State University, Tempe. Her research interests are in the areas of integrated sensing and processing, waveform agile sensing, and time-frequency signal processing. She has published more than 100 refereed journal articles, book chapters, and conference papers.

Dr. Papandreou-Suppappola is the recipient of the NSF CAREER Award in 2002. She is currently serving as an Associate Editor for the IEEE TRANSACTIONS ON SIGNAL PROCESSING and Treasurer of the Conference Board of the IEEE Signal Processing Society.

Darryl Morrell (SM'04) received the Ph.D. degree in electrical engineering in 1988 from Brigham Young University, Salt Lake City, UT.

He is currently an Associate Professor with the Department of Engineering, Arizona State University, Tempe, (Polytechnic campus) where he is participating in the design and implementation of a multidisciplinary undergraduate engineering program using innovative, research-based pedagogical, and curricular approaches. His research interests include stochastic decision theory applied to sensor scheduling and information fusion. He has received funding from the Army Research Office, the Air Force Office of Scientific Research, and DARPA to investigate different aspects of Bayesian decision theory, with applications to target tracking, target identification, and sensor configuration and scheduling problems in the context of complex sensor systems and sensor networks. His publications include more than 50 refereed journal articles, book chapters, and conference papers.