WAVEFORM DESIGN AND SCHEDULING FOR AGILE SENSORS FOR TARGET TRACKING

Sandeep P. Sira, Darryl Morrell and Antonia Papandreou-Suppappola *

Department of Electrical Engineering, Arizona State University
ssira@asu.edu, morrell@asu.edu, papandreou@asu.edu

ABSTRACT

Waveform agile sensors are capable of modifying transmitted waveforms to adapt to a particular scenario of interest. In this paper, we investigate the problem of optimizing the pulse length and frequency sweep rate of transmitted sonar waveforms to track a target moving in two dimensions using two sensors. We use Simultaneous Perturbation Stochastic Approximation (SPSA) in conjunction with a particle filter tracker to optimize the waveform parameters to minimize the expected squared tracking error at each time (greedy minimization). Simulation results show that waveforms with optimized parameters provide better performance than those that use fixed parameters.

1. INTRODUCTION

Modern sensor systems consist of a network of functionally compatible components that can be independently designed and optimized. In active target tracking systems, for example, the sensor and tracking sub-systems have typically been treated as separate entities [1]. However, significant performance improvement can be obtained by an integrated optimization of the system as a whole, wherein agile sensors dynamically design the transmitted waveform so as to obtain information that contributes best to the overall system objective. For target tracking systems, this involves optimal waveform design and scheduling at the sensor front-end, in response to the tracker's current requirements, thus leading to better tracking performance.

In [2, 3], the optimal selection and design of waveforms to track an underwater target moving in a one-dimensional, clutter-free environment was considered. The target dynamics and observation models were linear which allowed the use of the Kalman filter tracker. Closed form solutions for the optimal waveform and its parameters were derived to minimize the mean square tracking error. Most realistic target tracking scenarios, however, involve non-linearities and these solutions cannot be extended to more complex target motion.

*This work was supported by the DARPA Integrated Sensing and Processing program through a contract with Raytheon Missile Systems.

Adaptive waveform design and scheduling involves a search over a space of allowable waveforms and their parameters to choose the waveform that minimizes (or maximizes) a given criterion. In most cases, the performance metric cannot be expressed as a function of the design parameter; it is only observable as a realization of a stochastic process. The optimal waveform selection problem does not admit an exact solution and stochastic optimization techniques [4] are required.

In this paper, we consider tracking a target moving in two-dimensions (2-D) by two sensors that transmit linear frequency modulated (LFM) chirp signals. We present a waveform scheduling algorithm that selects the duration and chirp rate of the next transmitted waveform for each sensor to minimize the predicted mean square tracking error. The waveform selection is performed using a stochastic steepest descent algorithm that uses SPSA [5] to estimate the gradient of a cost function. The target tracking and cost function evaluation are carried out using particle filters [6].

2. SYSTEM MODEL

2.1. Target dynamics model

Let $\mathbf{X}_k = [x_k \;\; y_k \;\; \dot{x}_k \;\; \dot{y}_k]^T$ represent the state of a target at time k, where x_k and y_k are the x and y position coordinates respectively, and $\dot{x}_k$ and $\dot{y}_k$ are the respective velocities. The target dynamics model consists of a linear, constant velocity model:

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

Here, $\mathbf{W}_k$ are independent, zero-mean Gaussian samples with covariance matrix Q that model the acceleration of the target. F and Q are given by

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

where δt is the sampling interval and q is a constant.

2.2. Sensor waveforms and observations model

The target is free to move in 2-D and is tracked by two fixed sensors, A and B. Each sensor transmits an LFM waveform with a complex Gaussian envelope given by

$$s(t) = \left(1/(\pi\lambda^2)\right)^{\frac{1}{4}} \mathrm{e}^{-\frac{(t/t_r)^2}{2\lambda^2}} \mathrm{e}^{j2\pi b(t/t_r)^2}. \qquad (3)$$

Here, λ parameterizes the duration of the Gaussian envelope, b is the chirp rate, and $t_r = 1$ is a reference time. For sensor i, $i =$ A, B, $\boldsymbol{\theta}_k^i = [\lambda_k^i, b_k^i]^T$ represents the LFM waveform parameter vector at time k.

The transmitted waveforms are reflected by the target, and sensor i measures the time delay τ^i and Doppler shift ν^i of the received signal. The range and range-rate of the target are given by $r^i = c\tau^i/2$ and $\dot{r}^i = c\nu^i/(2\omega_c)$, where c is the velocity of propagation of the waveform and ω_c is the carrier frequency. With respect to sensor i located at (x^i, y^i), $[r_k^i \dot{r}_k^i]^T = h(\mathbf{X}_k)$ relates the observation to the target state where

$$\begin{aligned} r_k^i &= \sqrt{(x_k - x^i)^2 + (y_k - y^i)^2} \\ \dot{r}_k^i &= (\dot{x}_k(x_k - x^i) + \dot{y}_k(y_k - y^i))/r^i. \end{aligned}$$

The observation $\mathbf{Z}_k$ is thus a nonlinear function of the target state and is modeled as

$$\mathbf{Z}_k = h(\mathbf{X}_k) + \mathbf{V}_k \,, \qquad (4)$$

where the measurement errors are modeled by $\mathbf{V}_k$, a zero-mean, Gaussian noise process with covariance matrix $N(\boldsymbol{\theta}_k)$. Here, $\boldsymbol{\theta}_k = [\boldsymbol{\theta}_k^{A^T} \boldsymbol{\theta}_k^{B^T}]^T$ is a combined waveform parameter vector for both sensors at time k. The measurement error thus depends explicitly upon the transmitted waveform through $N(\boldsymbol{\theta}_k)$. Due to the nonlinearity in the observation model (4), we use a particle filter to recursively estimate the target state.

2.3. Sensor characterization

In this section, we derive a relationship between the LFM waveform parameters $\boldsymbol{\theta}_k$ and the observation noise covariance $N(\boldsymbol{\theta}_k)$. We use the Cramer-Rao lower bound (CRLB) on the variance of the range and range-rate errors to approximate the noise covariance. We assume that the observation noise at each of the sensors is independent. Accordingly, $N(\boldsymbol{\theta}_k)$ will be a block diagonal matrix of the respective noise covariance matrix for each sensor.

The ambiguity function (AF) of a signal $s(t)$ is given by

$$AF_s(\tau, \nu) = \int_t s\left(t + \frac{\tau}{2}\right) s^*\left(t - \frac{\tau}{2}\right) e^{-j2\pi\nu t} dt,$$

where τ and ν can correspond to the errors in the estimates of the delay and Doppler shift of the pulse when reflected off a target. The negative of the second derivatives of the AF, evaluated at $\tau = 0$ and $\nu = 0$, yield the elements of the Fisher information matrix [7]. Denoting the signal-to-noise ratio (SNR) as η, the Fisher information matrix for the LFM waveform in (3) is

$$I = (2\pi)^2 \eta \begin{bmatrix} \frac{1}{2(2\pi)^2\lambda^2} + 2b^2\lambda^2 & 2b\lambda^2 \\ 2b\lambda^2 & \frac{\lambda^2}{2} \end{bmatrix}. \qquad (5)$$

The CRLB on the variance of the error in the estimate of $[\tau, \nu]^T$ is given by I^{-1}. In a matched-filter receiver, the maximum likelihood estimates are jointly asymptotically Gaussian with covariance matrix I^{-1} [7]. We restrict our attention to cases where the SNR is high and there is no clutter. Hence, the side lobes of the AF may be neglected, and I^{-1} becomes a suitable characterization of the optimal receiver.

Since $r = c\tau/2$ and $\dot{r} = c\nu/(2\omega_c)$, the CRLB on the error variance of the estimate of $[r, \dot{r}]^T$ is given by $\Gamma I^{-1} \Gamma^T$ where $\Gamma = \mathrm{diag}(c/2, c/(2\omega_c))$. I^{-1} depends explicitly on the waveform parameters (see (5)). The measurement error covariance at the ith sensor is $N(\boldsymbol{\theta}_k^i) = \Gamma(I_k^i)^{-1}\Gamma^T$ and $N(\boldsymbol{\theta}_k) = \mathrm{diag}(N(\boldsymbol{\theta}_k^A), N(\boldsymbol{\theta}_k^B))$.

3. STOCHASTIC OPTIMIZATION

3.1. Performance criterion

We seek to configure the sensors with the waveform parameter $\boldsymbol{\theta}_k$ that minimizes the predicted mean square tracking error, given the sequence of observations up to time $k-1$. We define the performance metric as

$$J(\boldsymbol{\theta}_k) = E_{\mathbf{X}_k, \mathbf{Z}_k | \mathbf{Z}_{1:k-1}} \left\{ (\mathbf{X}_k - \hat{\mathbf{X}}_k)^T (\mathbf{X}_k - \hat{\mathbf{X}}_k) \right\}, \quad (6)$$

where $E\{\cdot\}$ is the expectation over $\mathbf{X}_k$ and $\mathbf{Z}_k$, and $\hat{\mathbf{X}}_k$ is the estimate of $\mathbf{X}_k$ given the sequence of observations from 1 to k that is obtained by the particle filter. The value of $\boldsymbol{\theta}_k$ that minimizes $J(\boldsymbol{\theta}_k)$ in (6) is the optimal waveform parameter.

3.2. Stochastic steepest descent

In order to obtain an estimate of the optimal waveform parameter, $\hat{\boldsymbol{\theta}}_k$, we use an iterative stochastic steepest descent algorithm that updates $\hat{\boldsymbol{\theta}}_k$ according to

$$\hat{\boldsymbol{\theta}}_k^{l+1} = \pi_\Theta\left(\hat{\boldsymbol{\theta}}_k^l - a_l \widehat{\nabla J}_l(\hat{\boldsymbol{\theta}}_k^l)\right), \quad l = 0, \ldots, L-1, \quad (7)$$

where $\widehat{\nabla J}_l(\hat{\boldsymbol{\theta}}_k^l)$ is the estimate of the gradient of the cost function at parameter $\hat{\boldsymbol{\theta}}_k^l$ and $\pi_\Theta(\cdot)$ is a projection operator that constrains $\hat{\boldsymbol{\theta}}_k^{l+1}$ to lie within the set of allowed waveform parameters. The gradient is approximated using SPSA [5], which requires estimates of the cost function as discussed next.

3.3. Estimation of the cost function

The cost function is the expected squared tracking error in (6) which can be expanded as

$$J(\boldsymbol{\theta}_k) = \int_{\mathbf{X}_k} \int_{\mathbf{Z}_k} (\mathbf{X}_k - \hat{\mathbf{X}}_k)^T (\mathbf{X}_k - \hat{\mathbf{X}}_k)\, p(\mathbf{Z}_k|\mathbf{X}_k, \boldsymbol{\theta}_k) \\ p(\mathbf{X}_k|\mathbf{Z}_{1:k-1}, \boldsymbol{\theta}_{1:k-1})\, d\mathbf{X}_k\, d\mathbf{Z}_k.$$

This evaluation must be carried out *before* the waveform at time k is transmitted. The expectation in (6) is thus over predicted states and observations. Accordingly, the squared error is a random variable, and its expectation may be computed by Monte Carlo simulations. The expectation is approximately given by

$$\hat{J}(\boldsymbol{\theta}_k) = \frac{1}{N}\sum_{n=1}^{N}\frac{1}{M}\sum_{m=1}^{M}\left(\mathbf{X}_k^n - \hat{\mathbf{X}}_{k|\mathbf{Z}_k^m, \mathbf{Z}_{1:k-1}}\right)^T \\ \cdot \left(\mathbf{X}_k^n - \hat{\mathbf{X}}_{k|\mathbf{Z}_k^m, \mathbf{Z}_{1:k-1}}\right), \quad (8)$$

where $\mathbf{X}_k^n$ are independent samples drawn from the density $p(\mathbf{X}_k|\mathbf{Z}_{1:k-1}, \boldsymbol{\theta}_{1:k-1})$ and $\mathbf{Z}_k^m$ are drawn independently from $p(\mathbf{Z}_k|\mathbf{X}_k^n, \boldsymbol{\theta}_k)$. The estimate of $\mathbf{X}_k^n$, given the observations $\mathbf{Z}_{1:k-1}$ and $\mathbf{Z}_k^m$ is $\hat{\mathbf{X}}_{k|\mathbf{Z}_k^m, \mathbf{Z}_{1:k-1}}$, and it is computed by a secondary particle filter. As N, M in (8) approach infinity, $\hat{J}(\boldsymbol{\theta}_k)$ approaches $J(\boldsymbol{\theta}_k)$ in (6).

3.4. Gradient approximation using SPSA

When the gradient of a cost function cannot be calculated in closed form, a finite differences approach may be used to approximate the gradient [8]. This involves determining the rate of change of the cost function over perturbations in the parameters. As the number of parameters increases, the method becomes computationally expensive. The SPSA method [5] follows the same approach but simultaneously perturbs all elements of the waveform parameter, thus reducing computational complexity.

Let $\boldsymbol{\Delta}_l = \{\Delta_{l1}, \ldots, \Delta_{l4}\}$ be a vector of symmetric Bernoulli distributed independent random variables, where l is the iteration index in (7). Also, let

$$\hat{J}_l^{(+)} = \hat{J}_l(\hat{\boldsymbol{\theta}}_k^l + c_l\boldsymbol{\Delta}_l), \quad \hat{J}_l^{(-)} = \hat{J}_l(\hat{\boldsymbol{\theta}}_k^l - c_l\boldsymbol{\Delta}_l) \quad (9)$$

represent noisy measurements of the cost function, evaluated as in (8) using parameters obtained by simultaneously perturbing $\hat{\boldsymbol{\theta}}_k^l$. The scalar sequence c_l satisfies $c_l > 0$ and $\sum_{l=1}^{\infty}(a_l/c_l)^2 < \infty$. The estimate of the gradient at the lth iteration is then given by

$$\widehat{\nabla J}_l(\hat{\boldsymbol{\theta}}_k^l) = \left[\frac{\hat{J}_l^{(+)} - \hat{J}_l^{(-)}}{2c_l\Delta_{l1}} \cdots \frac{\hat{J}_l^{(+)} - \hat{J}_l^{(-)}}{2c_l\Delta_{l4}}\right]^T. \quad (10)$$

A block diagram of the stochastic optimization procedure is shown in Figure 1. The target tracker maintains an estimate of the posterior density $p(\mathbf{X}_{k-1}|\mathbf{Z}_{1:k-1}, \boldsymbol{\theta}_{1:k-1})$ of the state given the observations by a collection of N_s particles $\mathbf{X}_{k-1}^j$ and weights $w_{k-1}^j, j = 1, \ldots, N_s$. At each time index, future states are predicted and then estimated by a secondary particle filter using observations that depend upon the current estimate of $\boldsymbol{\theta}_k$. The squared error in the estimate is used to compute the cost function which in turn is used to approximate the gradient using SPSA. Finally, the waveform parameter vector is updated according to (7).

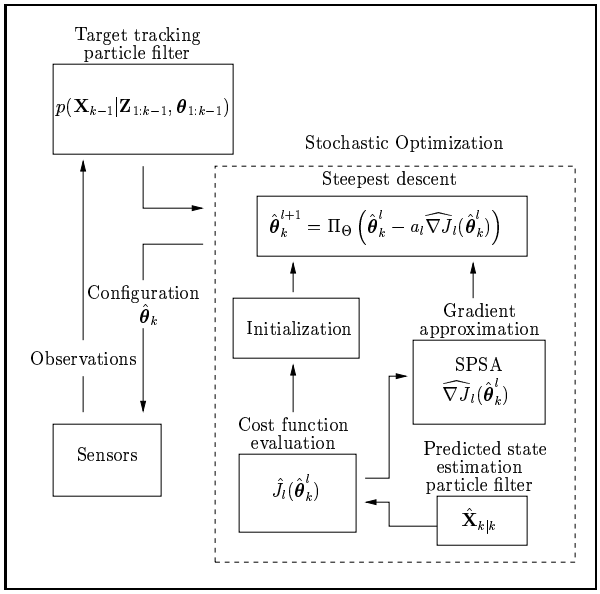

Fig. 1. Block diagram of the stochastic optimization procedure outlined in Section 3.

The components of the gradient corresponding to the duration and the chirp rate have significantly different magnitudes. Better convergence of the gradient descent algorithm is obtained by using different values of a_l in (7) and c_l in (9) for each component of the gradient [9].

In order to ensure that the steepest descent algorithm in (7) converges to the global minimum (over the constrained set of parameters) of the error performance surface, it is essential to choose the initial waveform parameter $\hat{\boldsymbol{\theta}}_k^0$ carefully. This may be achieved by evaluating the cost function at various test parameter values and initializing (7) with the one that yields the lowest cost. Furthermore, the noise inherent in the cost function evaluation in (8) and the gradient estimation in (10) may cause the number of iterations, L in (7), to become prohibitively large before convergence is achieved. This is particularly true if one or more of the waveform parameters is constrained to a range of values that is very small in comparison to the variations in the cost function evaluations. It is therefore advantageous to terminate the iteration at some practical value and compare the final and initial costs. In case the steepest descent algorithm has not resulted in a cost reduction, the initial value of the parameter may be chosen as $\boldsymbol{\theta}_k$.

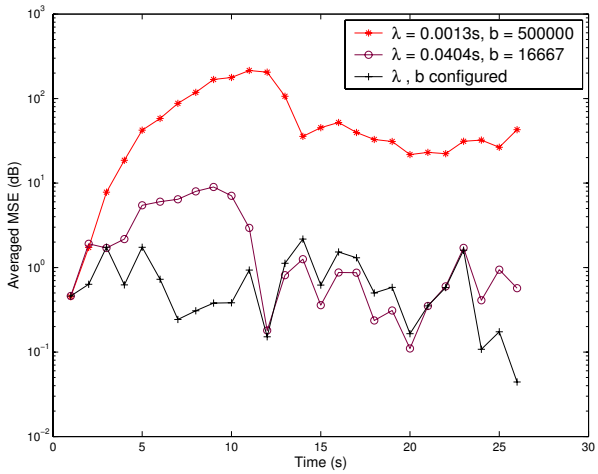

Fig. 2. Averaged mean square error (MSE) conditioned on convergence with $q = 0.1$.

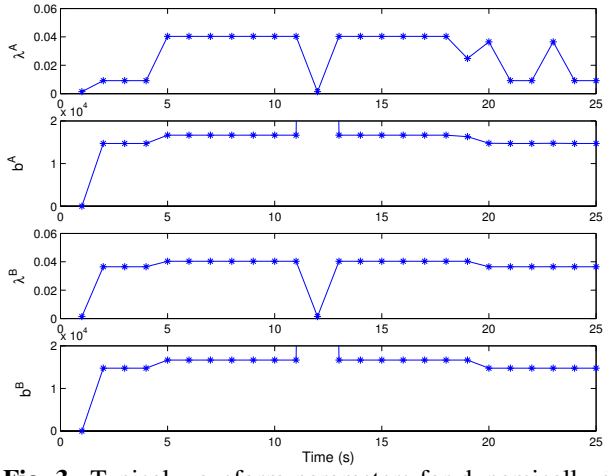

Fig. 3. Typical waveform parameters for dynamically designed waveforms for $q = 0.1$. The chirp rate for both sensors at time step 12 is 500000.

4. SIMULATION RESULTS

The simulation setup consists of a single underwater target that is free to move in 2-D. The carrier frequency is ω_c = 25 kHz and the velocity of sound in water is 1500 m/s. The effective pulse length, T_s, is chosen to be the time interval over which the signal amplitude is greater than 0.1% of its maximum value. This in turn determines the value of $\lambda = T_s/\alpha$, where $\alpha = 7.4338$. The pulse length is constrained to lie in the range [0.01,0.3] s while the bandwidth is limited to $B = 5$ kHz. We assume perfect target detection and zero probability of false alarm. To initialize $\hat{\boldsymbol{\theta}}_k^0$ in (7), we form all waveform combinations for which each sensor can use the minimum or maximum value of λ and $b \in \{-b_{max}, 0, b_{max}\}$. Here $b_{max} = B/(\lambda\alpha)$ is the maximum allowable chirp rate at a given value of λ. The expected cost is evaluated at each combination, and the combination that results in the lowest cost is chosen as $\hat{\boldsymbol{\theta}}_k^0$. We use $L = 50$ iterations and $N = M = 10$ averages for the cost function evaluation in (8).

The target moves along a parabolic trajectory generated according to the relation $x_k = 4y_k^2/2500$. The velocity of the target is obtained as the rate of change of position with $\delta t = 2$ s. A small additive noise is added to generate the final trajectory. Sensors A and B are located at (230,0) m and (-230,0) m. As the target dynamics model in (1) does not include acceleration, we test the tracking performance using different values of the process parameter q in (2). As q increases, the target dynamics model permits larger accelerations thus improving the filter's ability to track the target. However, it also increases the tracking error.

In the first simulation, $q = 0.1$. Figure 2 compares the average squared tracking error of converged runs in a 500 run simulation using sensors with fixed and dynamic waveforms. The convergence or divergence in each run was verified by examining the estimated position. When the sensors were dynamically configured as in Section 3, the filter never diverged. The percentage of runs that converged when fixed waveforms were used is shown in Table 1. We note that the fixed waveform filters do not always successfully track the target. The waveform parameters selected by the algorithm in a typical simulation are shown in Figure 3.

In the second simulation, we set $q = 1$. There is a significant improvement in the ability of the fixed waveform filters to track the target as shown in Table 1 since the number of converged runs is much larger when $q = 1$ than when $q = 0.1$. However, there is also an increase in the average tracking error for the filters, conditioned on convergence, as demonstrated in Figure 4. It is apparent that the dynamically configured sensors yield the best performance. The typical waveform parameters selected for this example are shown in Figure 5.

5. DISCUSSION

The waveform selection represents an adaptation to the uncertainty in the present state estimate, and it is influenced by the instantaneous sensor-target geometry. In Figure 3, during the initial and terminal parts of the trajectory, the sensors select short duration waveforms, favoring position estimation. Note that errors in position contribute more to the total tracking error than do errors in velocity. When $q = 0.1$, the target dynamics model provides more reliable information

Duration λ	Chirp rate b	$q = 0.1$	$q = 1$
0.0013 s	500000	14%	68%
0.0404 s	16667	22.4%	93%

Table 1. Percentage of converged runs for fixed waveforms.

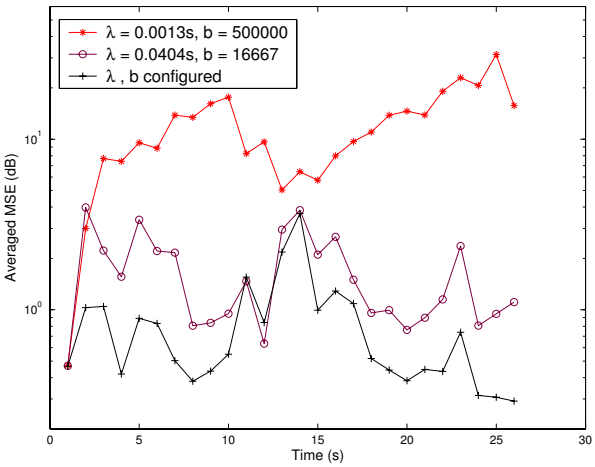

Fig. 4. Averaged mean square error conditioned on convergence with $q = 1$.

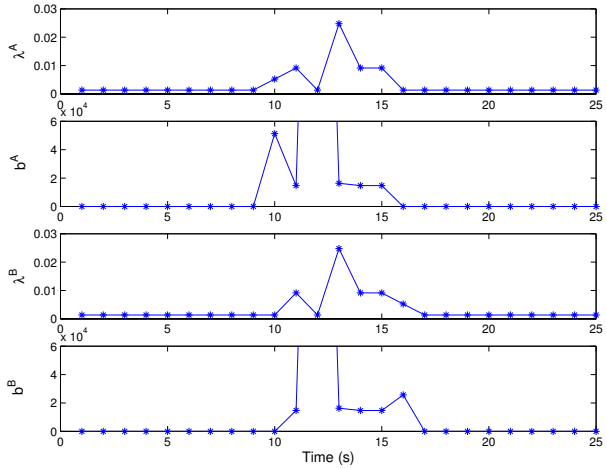

Fig. 5. Typical waveform parameters for dynamically designed waveforms for $q = 1$. The chirp rate for both sensors at time step 12 is approximately 500000.

to the tracking filter. When the range and range-rate estimation errors are correlated, waveforms with the longest possible duration and the highest allowable chirp rate simultaneously result in the best estimates of both parameters. This minimizes the variance of the range estimation errors conditioned on the range-rate which can be shown to be $\sigma^2_{r|\dot{r}} = c^2/(8\eta b^2\lambda^2)$. This characterizes the selection of the waveforms until the target becomes collinear with the sensors. The resulting system geometry causes the observation errors to maximally degrade the state estimate. The sensors respond by selecting the shortest pulse at this instant so as to accurately estimate position.

When $q = 1$ as in Figure 5, the filter places little trust on the motion model and relies more on the observation. The waveform selected has the shortest duration favoring position estimation. The corresponding chirp rate selected is zero; thus the transmitted waveform is a Gaussian waveform instead of an LFM chirp. This is due to the fact that the correlation between the estimation errors, that may be exploited to reduce the estimate error variance, is small.

6. CONCLUSION

The problem of optimally selecting waveforms to minimize the mean square tracking error does not admit an exact solution due to the nonlinearities present in most systems. In this paper, we have presented one possible approach to the problem of scheduling sensors using stochastic optimization to predict the tracking error at the next time step and a steepest descent algorithm to search for the configuration that minimizes this error. We applied the method to the tracking of an underwater target using two sensors and found that the resulting filter is more robust and yields better performance than when waveforms with fixed parameters are used. Due to the iterative nature of the stochastic optimization, the proposed technique can be computationally costly. We have recently developed a method, based on the unscented transform, to provide a real time implementation of the waveform selection algorithm [10].

7. REFERENCES

[1] Y Bar-Shalom and T. E. Fortmann, *Tracking and Data Association*, Academic Press, Boston, 1988.

[2] D. J. Kershaw and R. J. Evans, "Optimal waveform selection for tracking systems," *IEEE Transactions on Information Theory*, vol. 40, pp. 1536–1550, September 1994.

[3] D. J. Kershaw and R. J. Evans, "Waveform selective probabilistic data association," *IEEE Trans. on Aerospace and Electronic Systems*, vol. 33, pp. 1180–1188, October 1997.

[4] S. Andradottir, "A review of simulation optimization techniques," *Proc. of the Winter Simulation Conf.*, pp. 151–158, 1998.

[5] J. C. Spall, "Multivariate stochastic approximation using a simultaneous perturbation gradient approximation," *IEEE Trans. on Aut. Control*, vol. 37, pp. 332–341, March 1992.

[6] M. S. Arulampalam, S. Maskell, N. Gordon, and T. Clapp, "A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking," *IEEE Transactions on Signal Processing*, vol. 50, pp. 174–188, February 2002.

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

[8] H. Robbins and S. Monro, "A stochastic approximation method," *An. of Math. Stat.*, vol. 22, pp. 400–407, Sept. 1951.

[9] J. C. Spall, "Implementation of the simultaneous perturbation method for stochastic optimization," *IEEE Trans. on Aerospace and Electronic Systems*, vol. 34, no. 3, pp. 817–823, July 1998.

[10] S. P. Sira, A. Papandreou-Suppappola, and D. Morrell, "Time-varying waveform selection and configuration for agile sensors for target tracking applications," *ICASSP*, March 2005.