Control of Observation Systems with Application to Radar

Alan C. O'Connor
School of Engineering and Applied Sciences, Harvard University, Cambridge, MA 02138
E-mail: aoconnor@eecs.harvard.edu

Abstract: We formulate the problem of managing sensors as one of optimal control of the resulting state estimator covariance matrix. Originally posed in the early days of optimal control, this problem has received little attention, despite the proliferation of active sensing systems and sensor networks and the resulting need to target observations to obtain useful state estimates while minimizing energy expenditure. We consider the problem where the parameters of an observation system, to be selected from an given set, are controlled in-the-loop. We give a necessary condition for optimality and a numerical scheme for finding these solutions. Finally, using a convenient parametrization of the ambiguity function for a certain class of pulsed radar systems, we apply this theory to optimal tracking of a target with known dynamics.

Key Words: Sensor Management, Observation Control, Radar

1 INTRODUCTION

In this paper, we consider the problem of observation control for linear dynamical systems. Often, optimal control problems require the synthesis of a control law, with the measurement process, by which data for state estimation are obtained, taken as a given. However, in certain situations, it is advantageous to decide online which observations to take. As an example, we consider the problem of selecting the optimal radar pulses for tracking a target. In this and similar cases, the measurement process should be treated not as a given, but should be adjusted in an inner control loop.

1.1 Applications of Observation Control

The possibility of controlling the observation process can arise in several ways. In some situations, the sensors have adjustable parameters. This is the case, for example, in some radar systems, where the amplitude and spectral characteristics of the emitted waveform can be adjusted to improve target tracking. The question of waveform amplitude control for radar was considered in [2]. In the analogous biological system, bats are known to adjust echolocation pulses in the course of target pursuit.

In other applications, the quality of data available for state estimation depends on the placement of fixed sensors or the trajectories followed by mobile sensors. Selection and optimal placement of sensors is important for industrial process monitoring. The project described in [11] deployed a sensor network composed of underwater gliders to obtain data about ocean currents. Fulfillment of the science objectives of the project depended on optimization of the trajectories of the sensor platforms.

Finally, in large sensor networks it may be necessary for a central controller to select a subset of the sensor nodes to query. Such networks might be involved in tracking of chemical plumes for public safety, in monitoring air pollutants for emissions control, or for tracking military targets.

This material is based upon work supported by, or in part by, the U.S. Army Research Laboratory and the U.S. Army Research Office under contract/grant number W911NF-07-1-0376.

1.2 Previous Work

[1] gives a very clear formulation of the observation control problem. That paper, as well as [7] and [14], all consider selecting the best observation functional from a finite set. The author of [5] addresses the problem of choosing the observation functional from an arbitrary compact set, but crucially, he assumes that the dynamics of the system under observation are entirely deterministic and known precisely, thus achieving a considerable simplification. That paper also addresses the question of the optimal timing of a discrete observation process. There is also an extensive literature on the optimal placement of sensors for distributed parameter systems (DPS). See [10] for a survey of that field.

1.3 Plan of the Paper

In this paper, we restrict ourselves to systems described by linear stochastic differential equations and to observations that are linear functionals of the state plus Gaussian noise. We show that the formulation of the objectives of sensor management as an optimization problem permit the application of tools from control theory. To this end, we demonstrate that a linear function of the state estimator error variance is a natural metric for comparing the performance of observation policies. We present necessary criteria for optimality of policies and describe a gradient-descent method for solving the resulting boundary-value problem. Finally, we demonstrate the application of these ideas to a problem of radar waveform selection for optimal target tracking.

2 EVALUATION OF SENSING POLICIES

Many different measures of the effectiveness of sensor management policies have been proposed. [12] selects which sensors in a network to query by comparing the expected information gain over a finite time horizon. In [9], the evaluation is made based on the Rényi divergence of probability distributions representing multiple targets.

For linear dynamical systems perturbed by Gaussian noise, the most typical objective has been minimizing the trace of the state-estimate error covariance Σ as in [4, 7]. However, as noted in [8], the trace of the covariance can not be used to judge when the variance of particular states necessary for making a decision are small. That author proposes a number of other measures that are better in this respect, including the logarithm of the determinant of the covariance, which he interprets as entropic information.

The author of [2] formulates the observation control problem as the minimization of the variance of an arbitrary linear function of the state vector x. For example, $w = Gx$ so that $\mathcal{E}(\hat{w}-w)(\hat{w}-w)' = \mathcal{E}G(\hat{x}-x)(\hat{x}-x)'G' = \operatorname{tr} G'G\Sigma$. This approach has particular relevance when the ultimate goal is to make decisions based on the state estimate. We motivate our choice of loss function for evaluate sensor management policies by the following connection with stochastic optimal control.

2.1 Separation Theorem for the LQG Problem

Setting up the optimal control for a linear-quadratic-Gaussian problem gives us a reasonable choice for how to penalize the state-estimator variance. Consider a dynamical system with state vector $x(t)$, control $u(t)$, and disturbance $w(t)$ described by the stochastic differential equation

$$dx = A(t)x\,dt + B(t)(u\,dt + dw). \tag{1}$$

Let observations of the state be generated according to

$$dy = C(t)x\,dt + d\nu. \tag{2}$$

We take $w(t)$ and $\nu(t)$ to be Brownian motion processes (see [3] for a nice introduction to Itô calculus). Suppose the objective is to select $u(x,t)$ and $C(x,t)$ to minimize the cost

$$\eta = \mathcal{E}\int_{t_0}^{t_f} x'(t)Lx(t)\;dt + x'(t_f)L_f x(t_f) \tag{3}$$

Theorem 1 (Separation Theorem)**.** *The optimal control is $u^*(t) = -B'(t)K(t)\hat{x}(t)$, and the optimal loss η^* is:*

$$\begin{aligned}\eta^* &= \mathcal{E}x'(0)K(0)x(0) \qquad (4)\\ &\quad + \mathcal{E}\int_{t_0}^{t_f}\big((x-\hat{x})'KBB'K(x-\hat{x}) + \operatorname{tr} KBB'\big)\;dt,\end{aligned}$$

where $\hat{x}(t)$ is the Kalman filter estimate of $x(t)$, and $K(t)$ is a solution to the Riccati equation

$$\dot{K} = -KA - A'K + KBB'K - L, \qquad K(t_f) = L_f. \tag{5}$$

The first term of (4) captures the cost due to the initial state of the system. The second term captures the additional cost due to using the state-estimate in the feedback loop in place of the true state. The last term captures the loss resulting from process noise.

Plugging in the optimal feedback control $u = -B'Kx(t)$ and using the fact that the optimal state-estimator is unbiased ($\mathcal{E}(x-\hat{x}) = 0$), we obtain the modified system

$$\begin{aligned}\frac{d}{dt}\hat{x} = \frac{d}{dt}\mathcal{E}x &= \underbrace{(A(t) - B(t)B'(t)K(t))}\hat{x}(t) \qquad (6)\\ &= \tilde{A}(t)\hat{x}(t), \qquad (7)\end{aligned}$$

and an equation for the estimator error variance

$$\begin{aligned}\dot{\Sigma} &= \frac{d}{dt}\mathcal{E}(x-\hat{x})(x-\hat{x})' \qquad (8)\\ &= \tilde{A}(t)\Sigma + \Sigma\tilde{A}'(t) + BB' - \Sigma C'(t)R^{-1}C(t)\Sigma. \qquad (9)\end{aligned}$$

Corollary 1. *For a given optimal control problem (1,2,3), the loss is minimized when we chose the observation functional $C(t)$ to minimize the state-estimator covariance $\Sigma(t)$ weighted by $Q(t) = K(t)BB'K(t)$, where K is the solution to the Riccati equation (5).*

Proof. The optimal policy for the control of the stochastic system with imperfect observation is to feedback the (Kalman filtered) state estimate with a gain given by the solution of the Riccati equation (5). The details of the observation functional do not appear in (5), and so the only part of the optimal cost η^* which can be minimized with respect to the choice of $C(t)$ is the second term of (4). That term may be rewritten

$$\mathcal{E}\int_{t_0}^{t_f}(x-\hat{x})'KBB'K(x-\hat{x})\,dt = \int_{t_0}^{t_f}\operatorname{tr} KBB'K\Sigma\,dt. \tag{10}$$

Therefore, to minimize the cost η, the observation functional $C(t)$ should be chosen to minimize

$$\eta_c = \int_{t_0}^{t_f}\operatorname{tr} KBB'K\Sigma\;dt. \tag{11}$$

□

3 OBSERVATION CONTROL

Consider the non-linear differential equation

$$\begin{aligned}\dot{\Sigma}(t) &= A\Sigma + \Sigma A' + BB' - \Sigma C'(t)R^{-1}C(t)\Sigma, \qquad (12)\\ \Sigma(0) &= \Sigma_0 \qquad (13)\end{aligned}$$

and think of the observation functional $C'(t)R^{-1}C(t)$ as a control, which we may use to steer the estimator variance in a desired direction. The observation control problem can take the form of either of the following:

1. Given constraints on the observation effort over an interval of time, minimize a functional $\eta = \int_{t_0}^{t_f}\operatorname{tr} Q\Sigma\,dt + \operatorname{tr} Q_f\Sigma(t_f)$ of the state-estimator covariance.
2. Given initial and final values for the state-estimator covariance, find the observation policy that achieves the desired transfer while minimizing the observation effort.

In each of these cases, we may consider sub-cases where $C(t)$ is restricted to be in some finite set, in some compact set $\mathcal{C}$, or even $C(t) = \alpha C,\ \alpha \in \mathbf{R}^+$ for a fixed C, that is where we control only the amplitude of the observation functional. The first two sub-cases correspond roughly to two of the problems posed by Chernous'ko in [5]. The amplitude-only observation control is considered in [2]. In this paper, we will consider only the first problem, and in particular the case where the instantaneous observation "effort" is limited by a constraint of the form $\|C(t)\| \leq k$.

3.1 Application of the Maximum Principle

An observation strategy must consider not only an instantaneous reduction in variance but also the effect of the dynamics on future time evolution. The Pontryagin maximum principle gives us a principled way to trade off between present and future costs. Let $A(t) \in \mathbf{R}^{n\times n}$ and $B(t) \in \mathbf{R}^{n\times m}$ be matrices describing the dynamics of $x(t)$. Let the observation functional $C(t)$ be from some specified set $\mathcal{C} \subseteq \mathbf{R}^{m\times n}$. Assume $C(t)$ is bounded and measurable. Then the variance $\Sigma(t) = \mathcal{E}(x-\hat{x})(x-\hat{x})'$ of the estimator for x evolves according to Equation 12. The problem is to find $C(t) \in \mathcal{C}$ to minimize

$$\eta = \int_{t_0}^{t_f} \operatorname{tr} Q\Sigma(t)\ dt + \operatorname{tr} Q_f\Sigma(t_f), \qquad (14)$$

where Q and Q_f are positive semi-definite matrices. The control Hamiltonian is then

$$\mathcal{H} = \operatorname{tr} \Lambda'\big(A\Sigma + \Sigma A' + BB' - \Sigma C'(t)C(t)\Sigma\big) + Q\Sigma, \qquad (15)$$

The differential equation for the costate is given by

$$\dot{\Lambda} = -\frac{\partial \mathcal{H}'}{\partial \Sigma}, \qquad (16)$$

$$= -(A - \Sigma C'C)'\Lambda - \Lambda(A - \Sigma C'C) - Q. \qquad (17)$$

$$\Lambda(t_f) = Q_f. \qquad (18)$$

The costate matrix is symmetric positive semi-definite.

3.2 Necessary Condition for Optimality

Denote the rows of $C(t)$ by $\{c_1, c_2, \ldots, c_m\}$. Then a straightforward application of the Pontryagin maximum principle requires the optimal $C(t)$ satisfy:

$$C^*(t) = \arg\max_C \operatorname{tr} \Lambda\Sigma C'(t)C(t)\Sigma \qquad (19)$$

$$= \arg\max_C \operatorname{tr} C(t)\Sigma\Lambda\Sigma C'(t) \qquad (20)$$

$$= \arg\max_{c_1,\ldots,c_m} c_1\Sigma\Lambda\Sigma c_1' + \cdots + c_m\Sigma\Lambda\Sigma c_m' \qquad (21)$$

If $C(t)$ has a single row, this reduces to

$$C^*(t) = \arg\max_c c\Sigma\Lambda\Sigma c' \qquad (22)$$

If the set $\mathcal{C}$ is defined parametrically by $\mathcal{C} = \{c \in (\mathbf{R}^n)^* : \varphi(c) = 0\}$, for $\varphi(\cdot)$ a differentiable function, the Lagrange multiplier method leads to the requirement that at c^*, the normal vector to $\mathcal{C}$ should be parallel to $\Sigma\Lambda\Sigma c'$. More precisely,

$$\Sigma\Lambda\Sigma c' = \lambda\ \nabla\varphi\big|_c. \qquad (23)$$

If $\mathcal{C} = S^{n-1}$, the normal vector is c' and the condition is in the form of an eigenvalue problem: $\Sigma\Lambda\Sigma c' = \lambda c'$. A look at (22) shows that we should let c' be the eigenvector associated with the largest eigenvalue.

4 NUMERICAL TECHNIQUE

We now present a gradient-descent based method for finding a policy satisfying the necessary condition for optimality. First we note that in the problem formulated above, we have an initial value for $\Sigma(t)$ and a final value for $\Lambda(t_f) = Q_f$. If we were given $\Lambda(t_0)$, we could integrate forward in time, at each point selecting the observation control that minimized $\hat{\mathcal{H}}$. Thus, rather than trying to minimize the cost η directly, we propose a scheme in which we apply corrections to $\Lambda^k(t_0)$ in order to minimize $\|\Lambda^k(t_f) - Q_f\|^2$. This approach has two features to recommend it. First, the differential equation for $\Lambda(t)$ is linear, so the matrix derivative $\frac{\partial\|\Lambda^k(t_f)-Q_f\|^2}{\partial\Lambda^k(t_0)}$ is relatively easy to compute, and second, for problems with a unique trajectory satisfying the maximum principle, we have the following bound connecting minimization of this final error in the costate and minimization of the cost:

Theorem 2 (Error Bounds). *Let $\{C^k(t), \Sigma^k(t), \Lambda^k(t)\}$ be a triplet satisfying the differential equations 12, 17, and the necessary condition for optimality 21, but with an incorrect final value for the costate: $\Lambda^k(t_f) \neq Q_f$. Let $\eta_k = \operatorname{tr}\big(\Sigma^k(t_f)\Lambda^k(t_f)\big)$. Then η^*, the optimal cost for the problem 14, is bounded as follows:*

$$\eta_k - \|\Sigma^+(t_f)\|\cdot\|E\| \ \leq \eta^* \leq\ \eta_k + \|\Sigma^k(t_f)\|\cdot\|E\| \qquad (24)$$

where $E = \big(Q_f - \Lambda^k(t_f)\big)$ and $\|\cdot\|$ indicates the Frobenius norm of a matrix. $\Sigma^+(t)$ is the solution of $\dot{\Sigma}^+ = A\Sigma^+ + \Sigma^+A' + BB'$, $\Sigma^+(0) = \Sigma_0$ which satisfies $\Sigma^+(t) \geq \Sigma^*(t)$.

Proof. First, to obtain the upper bound, note that $\eta^* \leq \operatorname{tr}\big(\Sigma^k(t_f)Q_f\big) = \operatorname{tr}\big(\Sigma^k(t_f)(\Lambda^k(t_f) + E)\big)$. Then using the triangle inequality, $\operatorname{tr}\big(\Sigma^k(t_f)Q_f\big) \leq \eta_k + \|\Sigma^k(t_f)\|\|E\|$. To obtain the lower bound, we interchange the roles of the desired penalty Q_f and the actual one $\Lambda^k(t_f)$, and repeat the previous argument. The only problem is that we don't know $\Sigma^*(t_f)$, but in any case, its norm is less than $\Sigma^+(t_f)$ as defined above. □

Thus, by minimizing the final error in the costate, we are guaranteed to approach the optimal cost η^*, and the bounds give us the penalty we incur for early termination of a gradient descent algorithm.

5 RADAR WAVEFORM SELECTION

We now consider the application of these ideas to parameter selection in a pulsed Doppler radar system. The envelope of each outgoing pulse is given by $u(t)$, and we assume for concreteness that it is a Gaussian pulse with linear FM chirp,

$$u(t) = \left(\frac{2k^2}{\pi}\right)^{1/4} e^{-(k^2-jb)t^2}. \qquad (25)$$

The center frequency of the emitted pulse is f_0, and N_0 is the power spectral density of the Gaussian noise present in the return. We will assume that the pulse repetition frequency, that is the frequency at which pulses are emitted, reflections received, and the two signals correlated, is large enough with respect to the time constants of the target motion so that a continuous-time approximation to this system makes sense. The optimization problem is to choose the parameters $\{k(t), b(t)\}$, which give the pulse durations and FM rates for a sequence of emitted pulses so as to minimize the tracking variance of a target with known dynamics.

We momentarily change coordinates from range (r) and radial velocity (v_r) to delay ($\tau = 2r/c$) and Doppler shift ($\phi = 2f_0v_r/c$), which are more natural in the subsequent calculations. The variances of the delay and Doppler shift estimates are typically bounded in terms of the quantities α and β. For our waveform these are

$$\begin{aligned}\alpha^2 &= \frac{(2\pi)^2}{2E}\int t^2 u(t)u^*(t)dt \\ &= \frac{(2\pi)^2}{2E}\int t^2\left(\frac{2k^2}{\pi}\right)^{1/2}\exp\left(-2k^2t^2\right)\, dt \\ &= \frac{\pi^2}{2Ek^2}. \end{aligned} \tag{26}$$

Define $(\Delta f)^2 = \frac{4(k^4+b^2)}{\pi^2k^2}$. Then β^2 is given by

$$\begin{aligned}\beta^2 &= \frac{1}{2E}\int u'(t)(u^*)'(t)dt = \frac{(2\pi)^2}{2E}\int f^2\left|U(f)\right|^2\, df \\ &= \frac{(2\pi)^2}{2E}\int \frac{f^2}{\sqrt{2\pi(\Delta f)^2/16}}\exp\left(\frac{-16f^2}{2(\Delta f)^2}\right)\, df \\ &= \frac{(2\pi)^2}{2E}\frac{(\Delta f)^2}{16} = \frac{k^4+b^2}{2Ek^2}. \end{aligned} \tag{27}$$

The off-diagonal terms of the Fisher information matrix $\mathcal{I}$ are:

$$\begin{aligned}&\frac{2\pi}{N_0}\Im\int tu\frac{\partial u^*}{\partial t}dt \\ &\quad= \frac{2\pi}{N_0}\Im\int tu(t)\left(-2jbtu^*(t)\right)dt \\ &\quad= -2b\left(\frac{2E}{2\pi N_0}\alpha^2\right) = \frac{-b\alpha^2}{\pi}\left(\frac{2E}{N_0}\right). \end{aligned} \tag{28}$$

Finally, for the choice of amplitude above, the energy of the signal is $E = \frac{1}{2}\int u(t)u^*(t)\,dt = \frac{1}{2}$. So, the Fisher information matrix is given by

$$\mathcal{I} = \frac{1}{N_0}\begin{pmatrix}\beta^2 & \frac{-b\alpha^2}{\pi} \\ \frac{-b\alpha^2}{\pi} & \alpha^2\end{pmatrix}. \tag{29}$$

The Cramér-Rao bound states $\mathrm{V}ar\begin{pmatrix}\tau \\ \phi\end{pmatrix} \geq \mathcal{I}^{-1}$, and making the change back to range and radial velocity coordinates, we obtain

$$\mathrm{V}ar\begin{pmatrix}r \\ v_r\end{pmatrix} \geq R := \frac{\frac{1}{4}c^2N_0}{\alpha^2\beta^2 - \frac{b^2\alpha^4}{\pi^2}}\begin{pmatrix}\alpha^2 & \frac{b\alpha^2}{\pi f_0} \\ \frac{b\alpha^2}{\pi f_0} & \frac{\beta^2}{f_0^2}\end{pmatrix}. \tag{30}$$

Finally, R^{-1} is

$$R^{-1} = \frac{4f_0^2}{c^2N_0}\begin{pmatrix}\frac{k^4+b^2}{f_0^2k^2} & \frac{-b\pi}{f_0k^2} \\ \frac{-b\pi}{f_0k^2} & \frac{\pi^2}{k^2}\end{pmatrix} \tag{31}$$

Now R^{-1} is a symmetric matrix with determinant $\kappa^2 := \frac{4\pi^2}{c^2N_0}$, which for a fixed signal-to-noise ratio, is independent of the parameters, and gives the area of the uncertainty ellipse. We can parametrize R^{-1} by

$$R^{-1} = \kappa\exp(\Omega\theta)\begin{pmatrix}\zeta & 0 \\ 0 & \frac{1}{\zeta}\end{pmatrix}\exp(\Omega'\theta), \tag{32}$$

where $\Omega = \left(\begin{smallmatrix}0 & 1 \\ -1 & 0\end{smallmatrix}\right)$ is the infinitesimal generator of rotations, $\tan\theta$ is

$$-\frac{b^2+k^4-\pi^2f_0^2+\sqrt{\left(b^2+k^4+\pi^2f_0^2\right)^2-4\pi^2f_0^2k^4}}{2\pi bf_0}, \tag{33}$$

and ζ is

$$\frac{b^2+k^4+\pi^2f_0^2+\sqrt{\left(b^2+k^4+\pi^2f_0^2\right)^2-4\pi^2f_0^2k^4}}{2\pi f_0k^2}. \tag{34}$$

Surprisingly, if we restrict ourselves to $b > 0, k > 0$, the parametrization is injective, with inverses given by

$$b = \pi f_0\tan\theta\left(\frac{\zeta^2-1}{\zeta^2\tan^2\theta+1}\right) \tag{35}$$

and

$$k^2 = \pi f_0\zeta\left(\frac{\tan^2\theta+1}{\zeta^2\tan^2\theta+1}\right), \tag{36}$$

so that, given any symmetric matrix with $\det R^{-1} = \kappa^2$, describing a desired allocation of the measurement uncertainty, we can synthesize a corresponding FM-chirped Gaussian pulse. To our knowledge, this is the first time these formulas have appeared, although they can be derived from the more general result on transformations of the ambiguity function in [13] by evaluating the resulting integrals.

5.1 Derivation of the Optimal Policy

Let the target dynamics and observations obey the stochastic differential equations

$$d\mathbf{x} = A\mathbf{x}dt + Bd\mathbf{w} \tag{37}$$

$$d\mathbf{y} = \mathbf{x}dt + Dd\nu. \tag{38}$$

The Cramér-Rao bound on the variance is typically very tight, so here we will assume that $\mathcal{E}(D(t)d\nu)(D(t)d\nu)' = R = \frac{1}{\kappa}V(t)\operatorname{diag}(\zeta^{-1},\zeta)V'(t)dt$, with $V(t)$ orthogonal. The problem we wish to solve is:

$$\begin{aligned}\text{minimize}\quad & \eta = \operatorname{tr} Q_f\Sigma(t_f) \\ \text{w.r.t.}\quad & \theta(t)\in[0,2\pi], \zeta(t)\in[\zeta_{min},\zeta_{max}]\end{aligned} \tag{39}$$

where $\Sigma(t) = \mathcal{E}(x-\hat{x})(x-\hat{x})'$ obeys

$$\dot{\Sigma} = A\Sigma + \Sigma A' + BB' - \Sigma R^{-1}\Sigma, \quad \Sigma(0) = \Sigma_0. \tag{40}$$

The maximum principle indicates that we should maximize $\hat{\mathcal{H}} = \operatorname{tr}\Lambda\Sigma R^{-1}\Sigma$, where if we substitute our parametrization $R^{-1} = \kappa V\operatorname{diag}(\zeta,\zeta^{-1})V'$, we have

$$\hat{\mathcal{H}} = \kappa\left(\zeta v_1'\Sigma\Lambda\Sigma v_1 + \frac{1}{\zeta}v_2'\Sigma\Lambda\Sigma v_2\right). \tag{41}$$

If the bounds $\{\zeta_{min}, \zeta_{max}\}$ are independent of θ, this is extremized by picking $v_1(t), v_2(t)$ to be the two eigenvectors of $\Sigma\Lambda\Sigma$ with associated eigenvalues $\lambda_1(t) \geq \lambda_2(t) > 0$. The only extremum of (41) with respect to ζ occurs $0 = \left(\lambda_1 - \lambda_2/\zeta^2\right)$, but this is a minimum. Thus it remains only to check the endpoints. If $\zeta_{min}\zeta_{max} > 1$ and we assign $\lambda_1 = \lambda_{max}, \lambda_2 = \lambda_{min}$, where $\{\lambda_{min}, \lambda_{max}\}$ are the eigenvalues of $\Sigma\Lambda\Sigma$, then we pick $\zeta(t) = \zeta_{max}$. Conversely, if $\zeta_{min}\zeta_{max} < 1$ and we assign $\lambda_1 = \lambda_{min}, \lambda_2 = \lambda_{max}$, we maximize $\hat{\mathcal{H}}$ by picking $\zeta(t) = \zeta_{min}$.

5.2 Geometric Interpretation

We wish to demonstrate why an extreme value for the eccentricity of the uncertainty ellipse is always optimal. Suppose we have an uncertainty ellipse for the current state estimate. This ellipse represents a level set of the probability density over the state space. If Σ is the current value of the state-estimator covariance, then $\log \rho(x) \propto (\hat{x} - x)'\Sigma^{-1}(\hat{x} - x)$. If this is multiplied by a second Gaussian distribution with covariance R, representing the uncertainty of a new observation to be incorporated into the state estimate, the product density will have covariance $(\Sigma^{-1} + R^{-1})^{-1}$. Suppose for simplicity that both of these are diagonal: $\Sigma = \operatorname{diag}(s_1, s_2)$ and $R = \operatorname{diag}(r_1, r_2)$. Then if

$$\rho_1(\hat{x} - x) = \rho_1(z) \propto \exp\left(-\frac{1}{2}\left(\frac{z_1^2}{s_1^2} + \frac{z_2^2}{s_2^2}\right)\right), \quad (42)$$

$$\rho_2(x - dy) = \rho_2(z) \propto \exp\left(-\frac{1}{2}\left(\frac{z_1^2}{r_1^2} + \frac{z_2^2}{r_2^2}\right)\right), \quad (43)$$

the product $\rho_1\rho_2$ is proportional to

$$\exp\left(-\frac{1}{2}\left(z_1^2\left(\frac{1}{s_1^2} + \frac{1}{r_1^2}\right) + z_2^2\left(\frac{1}{s_2^2} + \frac{1}{r_2^2}\right)\right)\right). \quad (44)$$

The area of the uncertainty ellipse of the product $\rho_1\rho_2$ is

$$A = \pi c\sqrt{\det \Sigma_3} = \pi c\sqrt{\frac{s_1^2 r_1^2}{s_1^2 + r_1^2}}\sqrt{\frac{s_2^2 r_2^2}{s_2^2 + r_2^2}} \quad (45)$$

$$= \frac{\pi c s_1 s_2 r_1 r_2}{\sqrt{(s_1^2 + r_1^2)(s_2^2 + r_2^2)}}. \quad (46)$$

Since the numerator of the last expression is out of our control, as the product $r_1 r_2 = \det R = 1$ is constant, the area A is minimized by maximizing the denominator

$$\arg\max_{r_1, r_2} \sqrt{(s_1^2 + r_1^2)(s_2^2 + r_2^2)} = \arg\max_{r_1, r_2} \left(s_1^2 r_2^2 + s_2^2 r_1^2\right).$$

So if $s_1 > s_2$, as in Figure 1, we should let $r_2 > r_1$ and pick the extreme values: $r_2 = r_{max}, r_1 = 1/r_{max}$.

Now, we make the connection between minimizing the area of the uncertainty ellipse and maximization of the control Hamiltonian: if we let $v_i = \sqrt{\Lambda'}w_i$, and $P = \sqrt{\Lambda}\Sigma\sqrt{\Lambda'}$, which we can do since Λ is always symmetric positive-definite, then this is

$$\hat{\mathcal{H}} = \zeta\|Pw_1\|^2 + \frac{1}{\zeta}\|Pw_2\|^2 \quad (47)$$

We have made a coordinate change $z \to \sqrt{\Lambda}z$ that transforms the density $\rho_1(z)$ so that radius squared has units of cost (i.e. the same units as η). In the transformed space, $\hat{\mathcal{H}}$ is maximized, and the area of the uncertainty ellipse and hence cost are minimized, by picking w_1 the eigenvector of P with largest eigenvalue. The eigenvectors of P are the same as those of P^2, and in the original coordinates $\sqrt{\Lambda^{-1}}P^2\sqrt{\Lambda^{-1}} = \Sigma\Lambda\Sigma$. This is all shown in Figure 1.

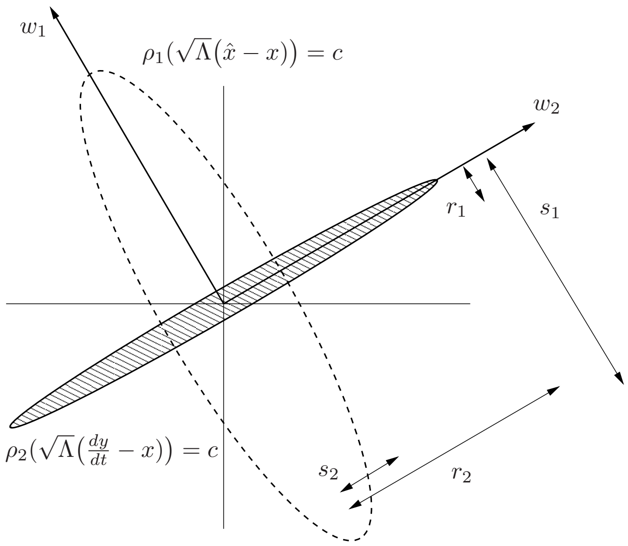

Figure 1: Diagram showing a level set of $\rho_1\left(\sqrt{\Lambda}(\hat{x} - x)\right)$ (dashed line), and the uncertainty region associated with the optimal observation policy $\rho_2\left(\sqrt{\Lambda}(\frac{dy}{dt} - x)\right)$ (shaded).

6 SIMULATION RESULTS

We compare the performance of the best constant policy to two different approximations to the optimal time-varying policy for tracking a periodically oscillating target. The first time-varying policy is a greedy approximation obtained by selecting $\{\theta(t), \zeta(t)\}$ at each time instant neglecting the effect of future observations. In particular, we used $\Phi_{-A'}(t, t_f)Q_f\Phi_{-A}(t, t_f)$ in place of the costate $\Lambda(t)$ when extremizing the control Hamiltonian. The second time-varying policy was obtained by applying the numerical scheme described in Section 4 until the error in the final costate was reduced to $\|E\| = 2.3 \times 10^{-5}$. We let $A = \left(\begin{smallmatrix} 0 & 1 \\ -\pi^2 & 0 \end{smallmatrix}\right)$. There is assumed to be dynamical noise added to the velocity only, so $B = \left(\begin{smallmatrix} 0 \\ \sqrt{10} \end{smallmatrix}\right)$. The initial variance is $\Sigma_0 = \left(\begin{smallmatrix} 10 & 0 \\ 0 & 2 \end{smallmatrix}\right)$, and the penalty functional $Q_f = \left(\begin{smallmatrix} 10 & 0 \\ 0 & 1 \end{smallmatrix}\right)$ weighs position errors at the final time more heavily than velocity errors. $t_0 = 0, t_f = 1$ and the integration time-step is fixed at 5×10^{-4}. The extremal values of the shape parameter are $\zeta_{min} = 1, \zeta_{max} = 5$. Finally, $\kappa = 1.5$. For all three policies, $\zeta(t) \equiv \zeta_{max}$. The three policies for choosing $R^{-1}(t)$ as parametrized by $\theta(t)$ are shown in Figure 2 and the resulting trajectories for the estimator variance are shown in Figure 3. The final values for the covariance and the corresponding scores η are shown in Table 1. The lower bound given in Theorem 2 implies the optimal cost is greater than 2.178. For this example, the optimized policy for shaping the radar pulses achieved a cost η less than a third of that obtained by the best constant policy.

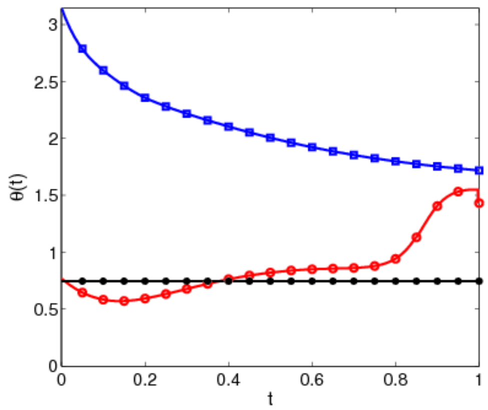

Figure 2: Three policies for radar pulse selection: best constant θ, greedy (blue squares), and optimized (red circles).

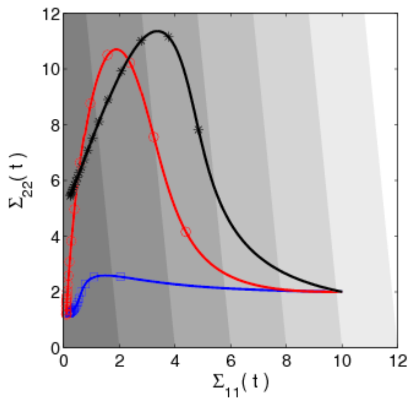

Figure 3: Trajectories of the variance equation for the policies shown in Figure 2: constant-θ (black stars), greedy (blue squares), and optimized (red circles). The shading shows the contours of the cost function $\eta = \operatorname{tr} Q_f \Sigma(t_f)$.

Table 1: Results from the Three Observation Policies

Policy	$\Sigma(t_f)$	η
Constant $\theta = 0.746$	$\begin{pmatrix} 0.240 & -0.104 \\ -0.104 & 5.436 \end{pmatrix}$	7.831
Greedy	$\begin{pmatrix} 0.193 & -0.056 \\ -0.056 & 1.230 \end{pmatrix}$	3.158
Optimized	$\begin{pmatrix} 0.100 & 0.010 \\ 0.010 & 1.179 \end{pmatrix}$	2.179
Optimal		≥ 2.178

7 CONCLUSIONS

The efficient management of sensors is an important task for the system designer in a number of application areas. Whether in environmental monitoring, battlefield awareness, target tracking, or a host of other applications, limited resources available for sensing must be allocated so as to maximize the usefulness of the resulting data for decision making. This report argued, through a connection to optimal stochastic control, that penalizing linear functions of the error covariance is a natural way to measure the effectiveness of a sensor management policy.

Posed in this way, the sensor management problem is one amenable to solution by the tools of optimal control. In particular, we used the Pontryagin maximum principle to derive a necessary condition for optimality, and described a method for finding policies satisfying this condition.

Finally, a problem in radar pulse selection was analyzed. The parametrization of the uncertainties associated with linear-FM chirped Gaussian pulses reduced this to a optimal control problem of the type considered here. Some remarks on the geometric interpretation of the necessary condition for optimality were given, along with the results of a small simulation comparing a few different waveform selection policies.

8 ACKNOWLEDGMENTS

We would like to thank N. Khaneja for first suggesting the radar problem as an application of observation control, P. Owrutsky for help in the derivation of the inverse formulas (35-36), and the reviewer for several helpful suggestions.

REFERENCES

[1] M. Athans, *On the determination of optimal costly measurement strategies for linear stochastic systems*, Automatica **8** (1972), 397–412.

[2] M. Athans and F. Schweppe, *Optimal waveform design via control theoretic concepts*, Information and Control **10** (1967), 335–377.

[3] R.S. Bucy and P.D. Joseph, *Filtering for stochastic processes with applications to guidance*, Interscience Publications, 1968.

[4] W.H. Chen and J.H. Seinfeld, *Optimal location of process measurements*, Intl. J. Control **21** (1975), no. 6, 1003–1014.

[5] F.L. Chernous'ko, *On optimization of the tracking process*, J. Applied Mathematics and Mechanics, translated from Prikl. Mat. Mekh. **33** (1969), no. 1, 101–111.

[6] C.E. Cook and M. Bernfeld, *Radar signals*, Academic Press, 1967.

[7] F.N. Grigor'yev and N.A. Kuznetsov, *Control of the observation process in continuous systems*, Problems of Control and Information Theory **6** (1977), no. 3, 181–201.

[8] B. Grocholsky, *Information-theoretic control of multiple sensor platforms*, Ph.D. thesis, University of Sydney, Australia, 2002.

[9] C.M. Kreucher, A.O. Hero III, K.D. Kastella, and M. R. Morelande, *An information based approach to sensor management in large dynamic networks*, Proc. IEEE **95** (2007), no. 5, 978–999.

[10] C.S. Kubrusly and H. Malebranche, *Sensors and controllers location in distributed systems - a survey*, Automatica **21** (1985), no. 2, 117–128.

[11] N.E. Leonard, D.A. Paley, F. Lekien, R. Sepulchre, D.M. Fratantoni, and R.E. Davis, *Collective motion, sensor networks, and ocean sampling*, Proc. IEEE **95** (2007), no. 1, 48–74.

[12] J. Liu, D. Petrovic, and F. Zhao, *Multi-step information-directed sensor querying in distributed sensor networks*, Proc. ICASSP, 2003, pp. V–145–148.

[13] F.B. Reis, *A linear transformation of the ambiguity function plane (corresp.)*, IRE Trans. on Information Theory **8** (1962), no. 1, 59.

[14] A.I. Yashin, *On the choice of optimal observation process*, Engineering Cybernetics, translated from Tekh. Kibern. (1969), no. 2, 105–108.