

PREDICTING A HYBRID SYSTEM*

J. E. BOYD[†], D. D. SWORDER[‡], AND R. G. HUTCHINS[§]

Abstract. Predicting the future state of a hybrid system is made exceptionally difficult by uncertainty in the current operating mode of the plant and the possibility of a plant modal transition during the interval over which a prediction is to be made. Unjustified approximations and errors in the system model typically create large errors in the prediction and, for some applications even worse, yield an unrealistic estimate of prediction error covariance. The effect can be an unreasonably optimistic appraisal of the accuracy of a prediction. In this paper, we document the progress to date of an investigation into hybrid prediction using the recently developed Gaussian Wavelet Estimator. We show that the path-length 2 version of this sophisticated estimator can provide consistent predictions and error covariance estimates when applied to an example ship defense problem.

1. Introduction. Hybrid systems have a characteristic state space decomposition connecting a finite dimensional Euclidean space (representing, for example, position and velocity), and a set with S elements (representing the operating regimes). A time-continuous process $\{x_t\}$ in the former, called the *base state*, is modulated by a *modal-state* process, which can be specified by an indicator process $\{\phi_t\}$ from the discrete set¹. Hybrid models are useful for representing multimode systems for purposes of estimation and control [11]. The synthesis problems that arise are difficult to solve, however, because of the nonlinearity in the equations of base-state evolution.

For the case in which a linear-Gauss-Markov (LGM) model is used for base-state evolution under each modal hypothesis, and the measurements are linear (or linearizable), considerable effort has been devoted to developing efficient algorithms for state estimation. Many are derived by first introducing a time-discrete approximation to the time-continuous system dynamics. It is well known that the conditional mean of the base state is then given by a weighted sum of the outputs of a family of Kalman filters, each filter tuned to a specific modal history.

While this representation of the base-state estimate is simply derived, the result is not useful in most applications. As time increases, the number of possible modal paths grows geometrically and the implementation of the optimal estimator soon becomes impossible. For example, in an encounter involving multiple aircraft tracked with multiple sensors and having multiple motion modes, the number of possible hypotheses delineating the temporal evolution of motion/observation path grows rapidly. At some

*Received on February 25, 2002; accepted for publication on November 19, 2002. Please send correspondence to J. E. Boyd.

[†]Cubic Defense Systems, San Diego, CA. E-mail: john.boyd@cubic.com

[‡]Department of ECE, University of California, San Diego, CA. E-mail: sworder@ece.ucsd.edu

[§]Department of ECE, Naval Postgraduate School, Monterey, CA. E-mail: hutchins@nps.navy.mil

¹Representing a Markov process by a sequence of indicator vectors has proven useful to many researchers. Although he does not claim to have invented it, Robert Elliott certainly contributed to the popularity of the technique in his seminal paper on filtering for Markov processes [3].

point—often very early—a complexity constraint must be imposed to bound the size of the algorithm. If a Gaussian sum is used to approximate the conditional distribution of the hybrid state, complexity can be controlled by limiting the number of terms in the sum.

Gaussian sums have proved their worth in various applications and are frequently used [1, 11]. Mixture reduction, reducing the number of terms in the Gaussian sum by mixing, is required to avoid hypothesis bloating (see [10, Section 3]). Mixture reduction can be based simply on the number of terms retained or on the depth of the path used as a conditioning index: a path-length-one (PL1) algorithm would use S filters, each associated with the current hypothesized mode; a PL2 algorithm would use S^2 filters, each associated with the current mode and its progenitor, and so on.

Methods of improving hybrid estimators have included both more sophisticated algorithms (software) and more sophisticated sensors (hardware). When the latter include direct modal measurements, various approaches to sensor fusion have been proposed. High-accuracy measurement of the mode is attractive because it obviates the need for a high-order filter bank.

We consider the problem of predicting a hybrid system. Predicting the future state of a dynamical system is problematic in any case because of the always imprecise knowledge of the current state and the effect of noise applied to the plant during the prediction interval. Additional complications make predicting a hybrid system even more difficult: besides uncertainty in the current plant base state, imperfect knowledge of the current modal state contributes uncertainty to selecting the proper modal dynamics to be applied to estimate the evolution of the base state. Moreover, the possibility of a modal transition during the prediction interval further increases the uncertainty in applying mode-dependent dynamics to the most recent base-state estimate. These factors obviously increase the uncertainty of a prediction, and for many applications, a realistic estimate of this uncertainty, the prediction error covariance, is required. The obvious sensitivity of prediction error and prediction error covariance to incorrect modal estimates (and to poor modal error covariance) suggest that a successful predictor must be based on a filter with good modal estimation performance, including realistic assessment of modal error covariance. This is in contrast with multiple model filters whose kinematic estimates may be satisfactory despite lower quality modal identification [8]. For this reason, the present study focused on applying the GWE to the prediction problem.

In this paper, a variable path-length sensor fusion algorithm, the Gaussian Wavelet Estimator (GWE), is presented. The GWE can be used with or without a direct modal sensor—the latter having cost advantage but performance disadvantage. Base-state prediction using the GWE is next considered and applied to a sample problem: defending a ship against an incoming missile. Performance of the PL2 predictor is evaluated with and without a modal sensor, and possible methods to improve the

predictor are suggested.

2. Multiple Model Estimators. Suppose the plant model of the system is defined on the probability space $(\Omega, \mathcal{F}, \mathcal{P})$ and the time interval $[0, \mathbf{T}]$. There is a right continuous filtration $\{\mathcal{F}_t; 0 \leq t \leq \mathbf{T}\}$ and right continuous, \mathcal{F}_t -adapted random processes: $\{\phi_t\}$, a piecewise constant modal process, and $\{w_t\}$, a Brownian motion. The modal state is a unit vector in \mathbf{R}^S with state space $\{\mathbf{e}_1, \dots, \mathbf{e}_S\}$. The coordinate in ϕ_t with value one marks the current mode of operation².

Both time-continuous and time-discrete models have been used in applications, the latter formed by sampling the relevant variables of the former at rate $1/T$. The coefficients in the linear, time-discrete hybrid model are modulated by the modal-state process. The latter are typically represented with a Markov process with transition matrix Π [11]:

$$(1) \quad x[k] = \sum_i (A_i x[k-1] + C_i w[k]) \phi_i[k-1],$$

$$(2) \quad \phi[k] = \Pi \phi[k-1] + \xi[k],$$

where $\{w[k]\}$ is an $\mathcal{F}[k]$ -unit Gaussian white sequence, and $\{\xi[k]\}$ is an $\mathcal{F}[k]$ -martingale increment sequence. The vector $x[k]$ is the base state. Label the local plant-noise covariance $R_x(i) > 0$.

The sensors give a linear measurement of the base-state vector and, if appropriate, a direct measurement of the mode:

$$(3) \quad y[k] = Hx[k] + n[k]$$

$$(4) \quad z[k] = \mathbf{D}\phi[k] + \eta[k]$$

where $\{n[k]\}$ is an $\mathcal{F}[k]$ -Gaussian white sequence with positive covariance R_x . The discernibility matrix, \mathbf{D} , measures the fidelity of the modal measurement, with $\mathbf{D}_{i,j} = \mathcal{P}\{z[k] = \mathbf{e}_i | \phi[k] = \mathbf{e}_j\}$, and $\{\eta[k]\}$ is a martingale increment. The measurement rate is determined by the sensors and their utilization policy and may differ from the basic clock rate—usually it is much slower. The sensor gains can be made time dependent so that at times of no measurement the observation is uninformative.

²For notational convenience in what follows, \mathbf{S} will designate an integer index set $\{1, \dots, S\}$. The vector \mathbf{e}_i is the i -th canonical unit vector in a space whose dimension is obvious from the context. Where no confusion will arise, a subscript may identify time, the component of a vector, or the element of an indexed family, with the meaning determined by context. Similarly a superscript may denote a power operator or an element of an indexed family. If a process $\{y_t\}$ is sampled every T seconds, the discrete sequence so generated is written $\{y[k]\}$, where the index denotes sample number rather than time. Conditional expectation is denoted with a circumflex, with the relevant σ -field apparent from context. A Gaussian random variable with mean \hat{x}_t and covariance P_{xx} is indicated by $x \sim \mathbf{N}(\hat{x}_t, P_{xx})$, with the same symbol used for the density function itself where no confusion will arise. If A is a positive symmetric matrix and x a conformal vector, $x'Ax$ is denoted $\|x\|_A^2$.

The initial plant state values are assumed to be independent with probability distributions $x[0] \sim \mathbf{N}(\hat{x}[0], P_{xx}[0])$ and $\phi[0] \sim \hat{\phi}[0]$. Denote the filtration generated by $\{y[k]\}$ by $\mathcal{Y}[k]$ and that by both measurements by $\mathcal{G}[k]$. The basic filtering problem is that of estimating $x[k]$ and $\phi[k]$ on the basis of $\mathcal{G}[k]$ (or $\mathcal{Y}[k]$ if only one sensor is available). Prediction means estimating $x[k+n]$ and $\phi[k+n]$ on the basis of $\mathcal{G}[k]$ (respectively, $\mathcal{Y}[k]$).

2.1. The Gaussian Wavelet Estimator. The system model is a hybrid in which the base state is represented in the conventional manner, and the mode is represented by a random regime process. Unfortunately, the equation for the true mean-square estimate of the comprehensive state is not implementable. For this reason, alternatives have been proposed that include moment approximations (the PME in [11]) and PL1 approximations (the MMAE when the mode is constant [17], and the IMM when the mode is a nonconstant Markov process [9]). The GWE can employ an arbitrary path length (for PL1 see [7] and for PL2 without modal measurements see [15]). The PL2 approximation to the conditional distribution of the state is given by the Gaussian sum

$$(5) \quad q[k] = \sum_{i,j,l=1}^S \alpha_{ijl}[k] \mathbf{N}(m_{jl}[k], P_{jl}[k]).$$

To provide a consistent notation, modal sequence identifiers will be written as subscripts p, i, j, l : p is the current mode, i is the immediate predecessor, j precedes by two time steps, and l by three steps. In (5), the conditional density at the beginning of the k th time interval would be a Gaussian sum. The mean and variance of a term in the sum $\mathbf{N}(m_{jl}[k], P_{jl}[k])$ would be associated with the two-term modal sequence jl preceding the new interval. The weight assigned a particular Gaussian element $\alpha_{ijl}[k]$ depends on the present (mode i) as well as the modes leading up to the present (jl). (The IMM and MMAE use the simpler form: $q[k] = \sum_{i=1}^S \alpha_i[k] \mathbf{N}(m_j[k], P_j[k])$, but the interpretation is the same.)

Equation (5) and the distributions that follow will be written in either their normalized ($\sum_{i,j,l} \alpha_{ijl}[k] = 1$) or, equivalently, in their unnormalized ($\sum_{i,j,l} \alpha_{ijl}[k] \neq 1$) form, usually without comment. The coefficients $\{\alpha_{ijl}\}$ are viewed as the conditional probabilities of the associated modal sequences, and the Gaussian primitives are viewed as the conditional densities of the base state given the modal path. If at time $t = kT$ the $\mathcal{G}[k]$ -density of the hybrid state is approximated by (5), the fundamental mapping is that from $\{\alpha_{ijl}[k], m_{jl}[k], P_{jl}[k], y[k+1], z[k+1]\}$ to $\{\alpha_{ijl}[k+1], m_{jl}[k+1], P_{jl}[k+1]\}$. The time $k+1$ filter variables must be determined from those at time k together with the modal- and base-state observations of the system at time $k+1$.

The dual sensor PL2-GWE can be written:

GWE-recurrence

Extrapolate: $i, j, l \in \mathbf{S}$

$$(6) \quad m_{ijl}^-[k+1] = A_i m_{jl}[k]$$

$$(7) \quad P_{ijl}^-[k+1] = A_i P_{jl}[k] A_i' + R_x(i)$$

Update: $i, j, l \in \mathbf{S}$

$$(8) \quad m_{ijl}^+[k+1] = m_{ijl}^-[k+1] + K_{ijl}[k+1](y[k+1] - H m_{ijl}^-[k+1])$$

$$(9) \quad P_{ijl}^+[k+1] = (I - K_{ijl}[k+1]H)P_{ijl}^-[k+1]$$

where $R_{ijl}[k+1] = R_x + H'P_{ijl}^-[k+1]H$ is the local output covariance and $K_{ijl}[k+1] = P_{ijl}^-[k+1]H'R_{ijl}[k+1]$ is the local ‘‘Kalman gain.’’ Denote $L_{ijl} = |D_{yy}^{ijl}[k+1]|^{\frac{1}{2}} \exp \frac{1}{2} \Delta \|m_{ijl}^+[k+1]\|_{D_{ijl}^+[k+1]}^2$, where information matrices D are defined to be the inverse of the corresponding covariance matrix, P , and, in particular, $D_{yy} = P_{yy}^{-1}$, the inverse of the base-state measurement covariance. The weighting coefficients in the GWE satisfy:

GWE modal recurrence

$\mathcal{Y}[k]$ Update: for $i, j, l \in \mathbf{S}$

$$(10) \quad \alpha_{ijl}^-[k+1] = L_{ijl} \alpha_{ijl}[k+1]$$

$\mathcal{G}[k]$ Update and Mode Transition: for $p, i, j, l \in \mathbf{S}$

$$(11) \quad \alpha_{pijl}[k+1] = z[k+1]' \mathbf{D}_{.p} \Pi_{pi} \alpha_{ijl}^-[k+1].$$

From this, various moments of interest can be computed; e.g.,

$$(12) \quad \hat{\phi}_p[k+1] = \sum_{i,j,l} \alpha_{pijl}[k+1] / \sum_{p,i,j,l} \alpha_{pijl}[k+1]$$

and

$$(13) \quad \hat{x}[k+1] = \sum_{i,j,l} \alpha_{ijl}^-[k+1] m_{ijl}^+[k+1] / \sum_{i,j,l} \alpha_{ijl}^-[k+1].$$

Note here that p indexes the plant mode at the ‘‘next time,’’ $k+1$. So in (11), for example, the post-transition weights $\alpha_{pijl}[k+1]$ form an augmented set relative to their predecessors, $\alpha_{ijl}^-[k+1]$, as each of the preceding weights represents a candidate initial condition that must be extrapolated assuming each possible post-transition mode.

Before the iteration is complete, mixture reduction is needed in the GWE. In [10], densities are merged on the basis of their size and distance from the ‘‘principle

components.” In the GWE, mixture reduction is achieved by retaining a single element from each path of length two preceding the current modal state. This is done using the conventional Gaussian sum merging formula [4]. As a notational convenience, denote the normalized family of α_{pijl} (or similar families) by $\bar{\alpha}_{pijl}$. Then the mixture reduction formula is:

$$(14) \quad \alpha_{pij}[k+1] = \sum_l \bar{\alpha}_{pijl}[k+1]; \quad m_{ij}[k+1] = \sum_l m_{ijl}^+[k+1] \bar{\alpha}_{ijl}^-$$

$$(15) \quad P_{ij}[k+1] = \sum_l (P_{ijl}^+[k+1] + (m_{ijl}^+[k+1] - m_{ij}[k+1])(\cdot)') \bar{\alpha}_{ijl}^-.$$

The recurrence of the GWE-estimator is complete.

Shown in Figure 1 is a diagram of the processing steps of the GWE. Beginning with the current (time step = k) best estimates of the plant, the GWE extrapolates the base state forward to time $k+1$, applying each possible plant mode to each possible initial condition, using estimated modal probabilities as appropriate. The resulting filter variables represent the best estimate of the plant at time $k+1$ prior to any new information (observations), and they are used to form plant outputs which represent the predicted base and modal states at time $k+1$. These computations can be performed at time k , so they do, in fact, represent a prediction for the plant outputs. Next, the base-state measurements for time $k+1$ are incorporated to improve the state estimates and reduce error covariance. In this formulation of the filter, the adjustment of modal state probabilities for a possible mode transition event and incorporation of any modal observation are combined in a single processing step. The Gaussian sums are then merged, or mixed, to reduce their number by the factor S by which they were increased in the first step. This is done by “summing out” the distinct terms resulting from the least recent time index, l . This final step thus advances the modal path fragment window underlying the Gaussian sums by trimming off the oldest coordinate.

2.2. Comparison to other Hybrid Estimators . The IMM is a PL1-algorithm that is particularly effective in hybrid tracking applications because it uses a sophisticated merging step to achieve performance approaching PL2. The conventional IMM is a $\mathcal{Y}[k]$ -algorithm, but fusion of the modal measurement was accomplished in [5], and a $\mathcal{G}[k]$ -IMM was presented there. We denote this imager-enhanced IMM by IMM-IE.

The GWE has much in common with the IMM and the MMAE; e.g., a bank of parallel Kalman filters, but the hypothesis merging steps are considerably different. The latter two maintain S active hypotheses and no pruning is required. But the GWE, even if PL1, requires pruning to avoid growth in complexity. For the present study, the PL2 version of the GWE was used with a dual sensor architecture.

The IMM-IE has a form similar to that of the GWE, but it begins with a less

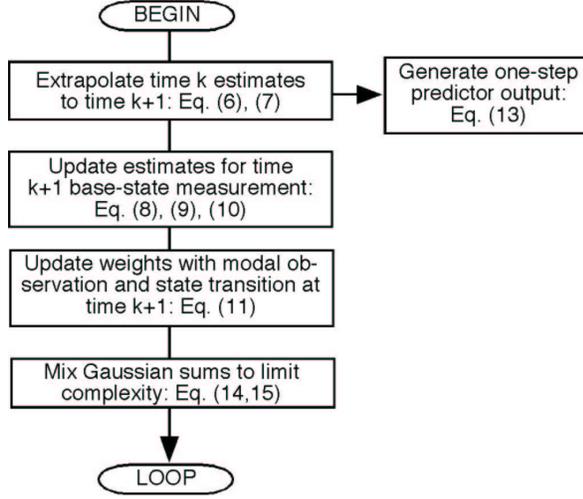


FIG. 1. Processing flow for the GWE.

differentiated Gaussian sum. As was the case in the GWE, extrapolation-update in the IMM-IE utilizes a set of S Kalman filters. The filters are initialized with $\{m_l[k], P_l[k]; l \in \mathbf{S}\}$, and after an observation, each generates a residual, $r_l[k+1] = y[k+1] - Hm_l[k+1]$ with pseudo-covariance $P_{yy}^l[k+1]$. The IMM-IE employs the residuals to update the modal estimate [5], as does the MMAE. Let $M_l = |D_{yy}^l[k+1]|^{1/2} \exp -\frac{1}{2} \|r_l[k+1]\|_{D_{yy}^l[k+1]}^2$. Then

IMM-IE modal recurrence

$\mathcal{Y}[k]$ Update: for $l \in \mathbf{S}$

$$(16) \quad \alpha_l^- [k+1] = M_l \alpha_l [k+1]$$

$\mathcal{G}[k]$ Update and Mode Transition for $l \in \mathbf{S}$

$$(17) \quad \alpha_l [k+1] = z[k+1]' \mathbf{D}_{.l} \sum_{p=1}^S \Pi_{lp} \alpha_p [k+1].$$

The IMM-IE does not require mixture reduction to maintain the correct number of hypotheses. Merging the moments after a measurement is, however, required for good performance, and the required formulas are given in the references.

The IMM-IE and the GWE require S Kalman filters, but the GWE extrapolates more initial conditions. If the performance of the two algorithms is equivalent, the reduced complexity of the IMM-IE makes it the appropriate choice. The fundamental difference between the algorithms arises from the way they utilize the $\{y[k]\}$ observation to improve modal estimation. Look first at the IMM-IE. Each filter in the bank

generates a residual. The bigger this residual is, normalized by the associated information matrix $D_{yy}^l[k+1]$, the smaller is the factor that multiplies the mode probability. This is plausible because an incorrect filter will have a large, and probably biased, residual for the most part. By reducing the influence of the filters that generate a sequence of large residuals, the associated modal hypothesis is appropriately reduced. A similar use of the residuals is found in the MMAE.

The GWE adjusts the modal probabilities on a different basis. It too uses the residuals generated by the filter bank, though there are more of them because there are several initial conditions per filter. The GWE looks not at the residual itself, but instead at the impact the residual has on the conditional mean state. And not the mean itself nor the increment in the mean, but rather the GWE emphasizes hypotheses that maximize the change in the normalized magnitude of m_{ijl} . This is harder to rationalize than is residual minimization; the role of the change in size of the conditional mean vector has received little attention in the literature. However, the improvement this somewhat unintuitive weighting makes in estimation accuracy is clear [12, 13].

3. Predicting a Hybrid System. The problem of predicting a system subject to changes in operating mode raises questions capable of inspiring philosophical debate. How is the possibility of one or more changes in plant mode during the prediction interval to be handled? If such changes are not possible (or even highly unlikely), the most current aggregate base-state estimate and error covariance might be extrapolated forward in time using the plant dynamics specified by the most probable mode. Given the likelihood, however, that the modal probability will be distributed over several modes, this is seen to be a rather crude approach. Certainly a more accurate prediction will be obtained if all modal components of the density estimate (means and covariances) are extrapolated using the path-appropriate modal dynamics and new combining weights are computed. But how are these modal path probabilities to be estimated? Given the Markov assumption of mode transitions, modal probabilities decay toward their steady state when extrapolated in the absence of observation. Is this built-in effect sufficient? What if the prediction interval is so long more than one transition is likely?

A further subtlety arises from the time-discrete approximation to the underlying time-continuous base state, and this matter is seen to affect the filter estimates as well as prediction. Suppose, to be specific, that the system is observed only every 10 seconds. A substantial interval elapses between observations. On what clock interval should a GWE simulation be run? Typically, discretization of a time-continuous system is done in such a way that the time-discrete interval is “small” relative to plant dynamics. Following this guideline, early simulations of the GWE were conducted with a simulation time step much smaller than the nominal observation interval [12], typically 0.1 sec. Presumably an advantage to a PL2 estimator accrues in part from

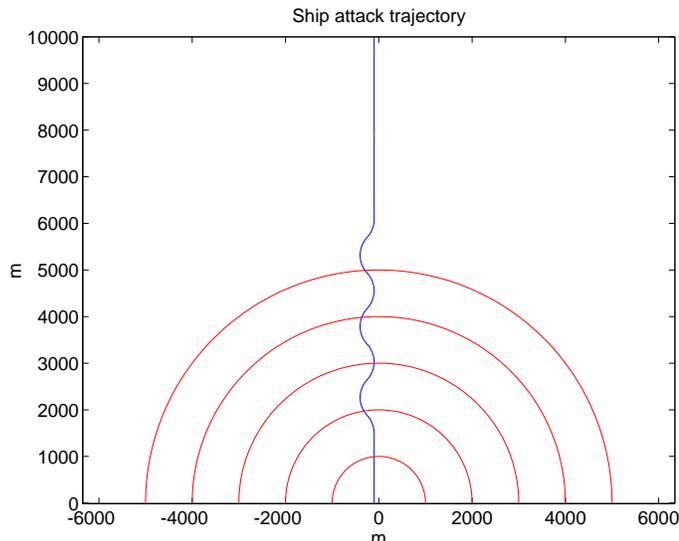


FIG. 2. *Missile trajectory for the example ship defense problem. The concentric rings represent kilometers from the ship. The trajectory shown represents the final 37.5 sec of missile flight.*

the “memory” the estimator has of the already-observed performance of the plant. Using a time-step of 0.1 sec., however, limits estimator “memory” to the preceding 0.2 sec. Depending on the plant dynamics, this may induce only a negligible advantage. If, on the other hand, the filter is run on the observation clock (10 seconds for this hypothetical case), its memory is extended to 20 seconds. Surely this should improve filter (and therefore, prediction) accuracy.

Some of these ambiguities clearly exist only because of plant modeling and filtering approximations. Given a tractable, time-continuous plant model and an optimal, conditional-mean estimator, the prediction should be the conditional mean of plant state at time $t + \tau$ given observations to time t . Tuning an approximate estimator to best predict an approximation to a time-continuous plant may require experimentation and judgement.

4. Example: a Ship Defense Problem. To illustrate and study predicting a hybrid system, we sought a problem involving a multimodal plant and having a natural penalty for inaccurate covariance estimation. A ship defense problem has these characteristics. Suppose a surface skimming missile approaches a ship at about Mach 1 with a constant altitude, say 5 m. It is desired to shoot the incoming missile with a machine gun having substantially higher-speed projectiles. We assume the missile performs evasive maneuvers as it approaches the ship: a series of random jinks having, however, known turn rates and turn-duration statistics. A representative trajectory is shown in Figure 2.

Such a missile defense system actually exists. The Phalanx Close-in Weapon Sys-

tem (CIWS) has been installed on all US Navy combatants [2]. According to the 1989 edition of *Jane's Weapon Systems* [6], it consists of a self-contained, closed-loop tracking system and a six-barrel Gatling gun, having a firing rate of 4500 rds/min, muzzle velocity of 1035 m/sec, and magazine capacity of 900 rounds. Some versions of the gun have a dispersion pattern created by barrel clamps. Interestingly, magazine size is deliberately limited to prevent overheating the gun barrels. Further capabilities of the defensive system, including maximum range, are not stated, but these details more than suffice as inspiration for our illustrative problem. For our problem, defending the ship requires tracking the missile, computing the time of flight for machine gun rounds to the approximate position of the missile, and predicting the position of the missile and $1\text{-}\sigma$ (or perhaps $2\text{-}\sigma$ or $3\text{-}\sigma$) prediction error covariance at the time gun rounds arrive. Atmospheric slowing of the projectiles and second order corrections to the actual time and position of engagement³ were neglected for this study.

Presumably, the smaller the prediction error covariance can be made, the higher the probability of kill will be, as the density of gun rounds can be increased if the missile position at time of engagement is better localized. We therefore require that the prediction error be small and that its covariance be realistic.

4.1. System (missile) model. We modeled the missile as subsonic, having a constant speed of 280 m/sec. The maneuver model for the missile is shown in Table 1. The table indicates that turns last about $1/.7 = 1.4$ s, while coasts last $1/.3 = 3.3$ s on average. Given the missile is turning left or right, it is 6 times more likely to transition to the other turn than it is to begin coasting. This characterizes the jinking behavior shown in Figure 2.

Because the missile maneuver model is Markov, it is memoryless. While the predictor may have high confidence of the current mode, it does not “know” how long this mode has obtained, and it is unable to predict with certainty when the next transition may occur, nor what that transition will be. These facts make the prediction problem more difficult. Alternative formulations of the maneuver model (see discussion below, Section 5) may diminish this greatest difficulty of the prediction problem.

4.2. Measurement suite. For the simulation, we used the simplest radar model, assuming Gaussian position errors of 20 m in the X (east-west) direction and 10 m in Y . Radar measurement frequency was 1 Hz. The imager, when used, was of high quality but low frequency. The simulated imager gave an indication of maneuver mode (left turn, coast, or right turn) once per second. The indication was correct with probability .900; if turning, the indication was no turn with probability .099 and the opposite turn with probability .001. If coasting, the incorrect reports

³Because the missile moves toward the ship during the flight of the defensive projectiles, the actual time of engagement will be earlier than we compute.

TABLE 1

Markov missile maneuver model. The model is specified by transition rates from row state to column state. Negative (diagonal) entries are the negative reciprocal of the mean time spent in the state.

	Left	Coast	Right
Left	-0.7	0.1	0.6
Coast	0.15	-0.3	0.15
Right	0.6	0.1	-0.7

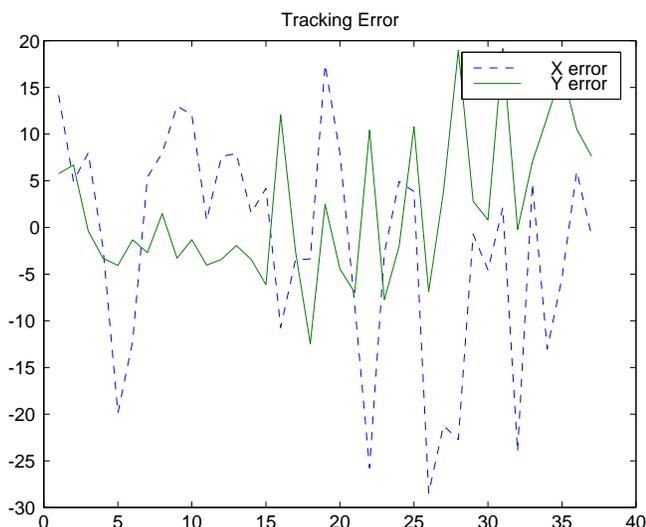


FIG. 3. Filter tracking error using the radar and imager. For the sample path shown, the root-mean-square value was about 16 m.

were balanced, with .050 probability to each of the turn modes.

4.3. Tracking performance. Figure 3 illustrates the accuracy with which the target was tracked for one sample path when both radar and imager were used. The relatively low radar sample rate (1 Hz) relative to the maneuver duration (about 3 sec) prevents much filtering of radar noise during the turns, and the peak X errors are seen to be similar to the radar model. For the run shown, the time root-mean-square averages of tracking error were 14.4 m in X and 6.5 m in Y . The time evolution of filter error covariance for the same sample path is shown in Figure 4. The filter has some trouble tracking the missile through the jinks, and this is correctly reflected in the error covariance. The rms time averages of these functions were 14.9 m and 7.4 m, confirming that the filter error covariance reasonably matches the filter tracking error shown in Figure 3.

We also ran the filter with no imager. Tracking performance was comparable to that with the good imager. An example run showed 19.8 m rms tracking error in X

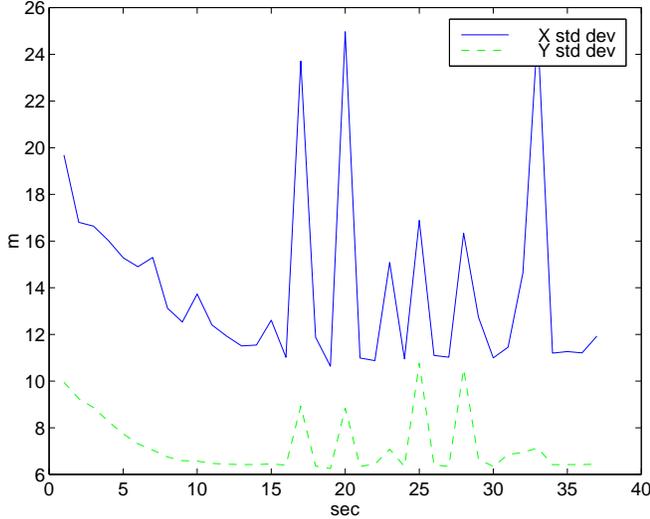


FIG. 4. Square root of filter tracking error covariance using radar and imager. The effect of the maneuvers is clearly evident for $t \in [16, 34]$.

and 8.3 m in Y . Error covariance time rms averages were 16.9 m and 8.1 m, just noticeably larger than those for the filter with imager. As expected, the imager made only a small improvement in tracking accuracy, and the error covariance averages again matched the filter error satisfactorily.

4.4. Prediction performance. All the prediction results presented in this paper were generated with the PL2 GWE using 1-second forward extrapolations and mode probability mixing. We first ran the predictor assuming a 1030 m/sec projectile. This is only about three times the missile velocity, so the missile travels a significant distance during the bullet time-of-flight. For example, at the maximum range shown in Figure 2, the prediction time is nearly 10 sec. Because the jinks, once they begin, last only about 3 seconds each, it is clear that the maneuvers made by the missile during flight of the bullets are crucially important to prediction accuracy. Figure 5 is a representation of the predictor performance with no imager. The balloons represent the covariance ellipses calculated by the predictor. They are centered on the predicted point of the encounter, and a line segment indicates the displacement from the actual missile position at that time.

It is clear in Figure 5 that at long ranges, this ship defense system is seriously challenged. The $1\text{-}\sigma$ ellipses are more than 1 km across, and in many cases, they do not overlap the missile trajectory. On the other hand, comparison of the ellipse size and the corresponding prediction error shows that the two quantities are reasonably related. If gunfire strategy were to uniformly fill an ellipse (or better, perhaps, a $2\text{-}\sigma$ or $3\text{-}\sigma$ version) with bullets at the predicted encounter time, such a strategy might

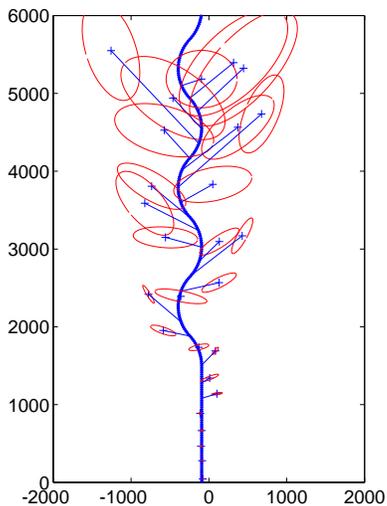


FIG. 5. *One-sigma covariance ellipses for target position prediction. Line segments from ellipse centers represent prediction error.*

indeed be effective. Unfortunately, the firing rates required to implement this strategy are probably not realizable: it would require a million bullets to fill a square kilometer with projectiles spaced 1 m apart.

Figure 6 presents the prediction accuracy averaged for 20 simulation trials. Encounters are planned from about 15 seconds into the missile visible flight, with the first predictions being over 7 seconds in length. The decreasing bias in Y error is caused by the reduced advance that would be caused by a turn in either direction. As distance-to-go diminishes, this effect is reduced. The strong oscillations in X -prediction error result from the rapid maneuver tempo: more than one maneuver mode change occurs during the prediction interval.

Adding an imager to the sensor suite slightly improves tracking performance, but more importantly in many problems, greatly improves mode identification. Figure 7 shows the result of the improved sensor suite in this application. While the prediction error is somewhat reduced relative to Figure 6, the results are not dramatic because the long prediction intervals allow the unpredictable missile maneuvers during projectile time-of-flight to dominate the prediction error. Better modal identification is not the key to improved performance of this system.

If, as is clear, the long time intervals and rapid pace of modal changes are the problems to be overcome, what further can be done? A higher observation rate might be considered. This will improve tracking accuracy for nearly any filter (see the benchmark tracking problems at the 1994 American Control Conference, e.g., [14]), and modal identification as well for most hybrid filters. But it will not shorten the time interval over which the predictions must be made, and that is the central issue

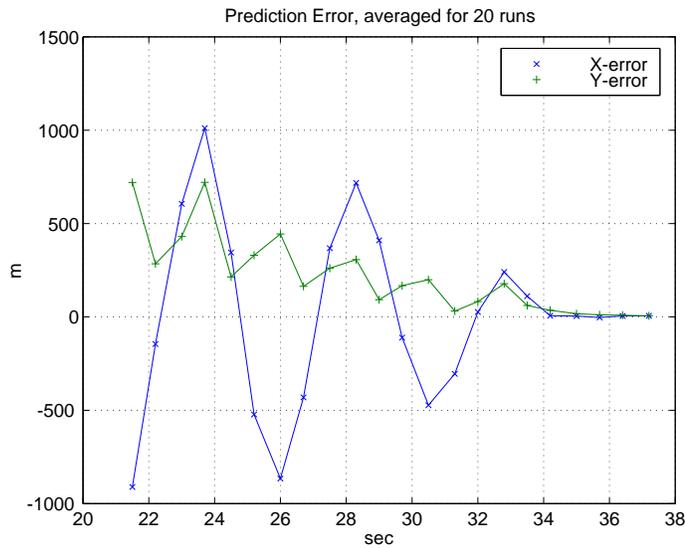


FIG. 6. Filter prediction error, 20-sample average. The time axis indicates the time of defensive weapon encounter with incoming missile. If the missile is not destroyed, impact with own ship is at 37.4 s. Projectile speed is 1030 m/sec; no imager is used in the tracker.

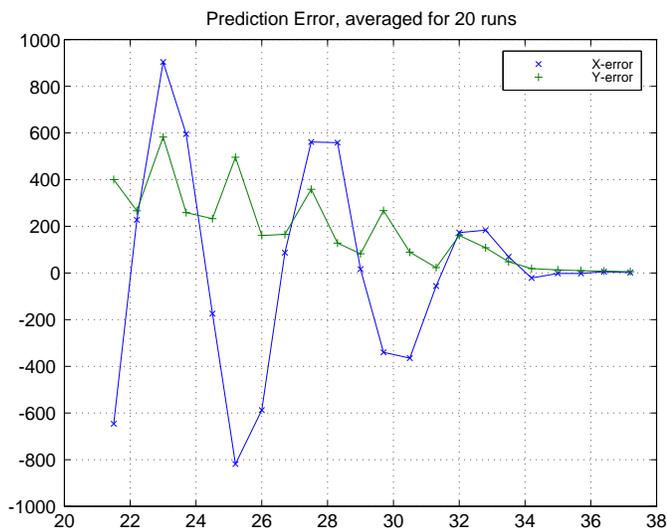


FIG. 7. Filter prediction error (m) using radar and the good imager, 20-sample average. The time axis indicates the time (s) of countermeasure encounter with incoming missile. If the missile is not destroyed, impact with own ship is at 37.4 s.

here.

Prediction error is nearly zero over the last 3 seconds of flight, so defensive weapon system use could be delayed until that time. But the missile is less than a kilometer from the ship when the turns end and the prediction error recedes. If the gun has

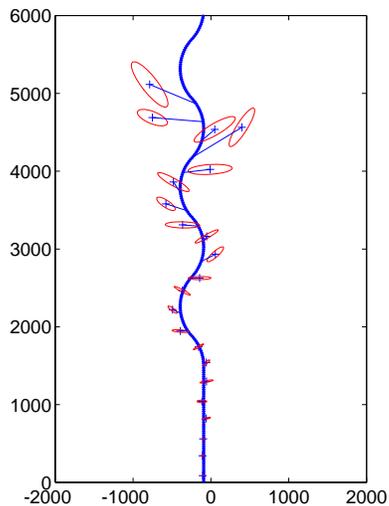


FIG. 8. *One-sigma covariance ellipses for target position prediction assuming the defensive weapon has 2000 m/sec projectiles. Sensor for this case was radar only.*

sufficient effect to cover a 200 meter ellipse, this range can be extended to perhaps 5 seconds out, or 1.4 km. This may still be too close for comfort.

As a means of reducing the prediction time required, we postulated a higher speed projectile: 2000 m/sec. Shown in Figure 8 are the prediction error covariance ellipses for this case. Comparing to Figure 5 shows a dramatic reduction in size of the ellipses and prediction error vectors. The prediction errors, plotted in Figure 9, are under 200 m over 9 seconds out, that is out to 2.6 km. And they are under 300 m as far out as 3.4 km.

5. Conclusions and future work. We found that a hybrid system can indeed be predicted and that the PL2 GWE gives results that could be useful in some circumstances. Prediction error and error covariance were found to be consistent with each other in a variety of model parameter and observation suite settings. The major difficulty of hybrid prediction appears when prediction intervals are long compared to expected modal durations.

Although Markov modal transitions are usually assumed, this is not always the case. In [16], for example, modal transitions are modeled with a renewal process. The resulting “memory” this refinement gives the filter can lead to improved performance in both filtering and predictions. We plan to look further into the example ship defense problem with the more sophisticated renewal modal transitions in the near future.

Several assumptions made on the example problem prevent the results from being conclusively applied to a real system. The crude radar model we used differs in important ways from a real radar. Hypersonic and transonic projectiles do not enjoy constant speed during defensive weapon fly-out. Second order timing issues should be

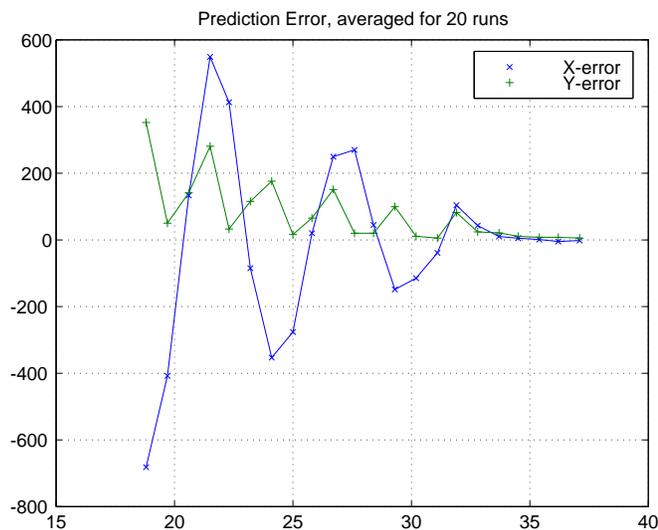


FIG. 9. Filter prediction error, 20-sample average, for the 2000 m/sec projectiles. Sensor suite is radar only.

modeled, and the firing rate and magazine capacity of the defensive weapon should also be considered. Effectively leading the target (in the duck hunting sense) will remain nearly impossible if the target can make one or more maneuver changes during projectile fly-out. Clearly, directed energy weapons would have a great advantage in this application because the speed of the energy almost eliminates the need for prediction.

REFERENCES

- [1] B. D. O ANDERSON AND J. B. MOORE, *Optimal Filtering*. Prentice-Hall, Inc., 1979.
- [2] RAYTHEON CO, *Phalanx close-in weapon system (CIWS)*, <http://www.raytheon.com/es/esproducts/dssphlx/dssphlx.htm>, 2001.
- [3] R. J. ELLIOTT, *Exact adaptive filters for markov chains observed in gaussian noise*, *Automatica*, 30:9(1994).
- [4] R. J. ELLIOTT, L. AGGOUN, AND J. B. MOORE, *Hidden Markov Models: Estimation and Control*, Springer-Verlag, 1995.
- [5] J. S. EVANS AND R. J. EVANS, *Image-enhanced multiple model tracking*, *Automatica*, 35(199), pp. 1769–1786.
- [6] *Jane's Weapon Systems*. Jane's Information Group, Coulsdon Surrey, UK, 1989.
- [7] C. T. LEONDES, D. D. SWORDER, AND J. E. BOYD, *Multiple model methods in path following*, *Journal of Mathematical Analysis and Applications*, 251(2000), pp. 609–623.
- [8] X. R. LI AND Y. BAR-SHALOM, *Design of an interacting multiple model algorithm for air traffic control tracking*, *IEEE Trans. on Control Systems Technology*, 1:3(1993), pp. 186–194.
- [9] E. MAZOR, A. AVERBUCH, Y. BAR-SHALOM, AND J. DAYAN, *Interacting multiple model methods in target tracking: a survey*, *IEEE Trans. on Aerospace and Electronic Systems*, 34:1(1998), pp. 103–123.
- [10] L. Y. PAO, *Multisensor multitarget mixture reduction algorithms for tracking*, *Journal of*

- Guidance Control and Dynamics, 17:6(1994), pp. 1205–1211.
- [11] D. D. SWORDER AND J. E. BOYD, *Estimation Problems in Hybrid Systems*, Cambridge University Press, Cambridge, UK, 1999.
 - [12] D. D. SWORDER AND J. E. BOYD, *A new merging formula for multiple model trackers*, Proc. of the SPIE: Signal and Data Processing of Small Targets, 4048(2000), pp. 498–509.
 - [13] D. D. SWORDER AND J. E. BOYD, *Sensor fusion in estimation algorithms*, Journal of the Franklin Institute, 1:1(2003), pp. 1–10.
 - [14] D. D. SWORDER, J. E. BOYD, F. DUFOUR, AND P. BERTRAND, *Dual-path algorithms for the benchmark tracking problem*, Proc. of the American Control Conference, pp. 2088–2092, June 1994.
 - [15] D. D. SWORDER, J. E. BOYD, R. J. ELLIOTT, AND R. G. HUTCHINS, *Data fusion using multiple models*, Proceedings of the Asilomar Conf. on Circuits, Systems, and Computers, pp. 1749–1753, Nov 2000.
 - [16] D. D. SWORDER, M. KENT, R. VOJAK, AND R. G. HUTCHINS, *Renewal models for maneuvering targets*, IEEE Trans. on Aerospace and Electronic Systems, 31:1(1995), pp. 138–150.
 - [17] J. R. VASQUEZ AND P. S. MAYBECK, *Density algorithm based moving-bank MMAE*, Proc. of the Conf. on Decision and Control, pp. 4117–4122, Dec. 1999.

