

Multiple-Model Adaptive Estimation Using a Residual Correlation Kalman Filter Bank

PETER D. HANLON

PETER S. MAYBECK, Fellow, IEEE
Air Force Institute of Technology

We propose a modified multiple model adaptive estimation (MMAE) algorithm that uses the time correlation of the Kalman filter residuals, in place of their scaled magnitude, to assign conditional probabilities for each of the modeled hypotheses. This modified algorithm, denoted the residual correlation Kalman filter bank (RCKFB), uses the magnitude of an estimate of the correlation of the residual with a slightly modified version of the usual MMAE hypothesis testing algorithm to assign the conditional probabilities to the various hypotheses that are modeled in the Kalman filter bank within the MMAE. This concept is used to detect flight control actuator failures, where the existence of a single frequency sinusoid (which is highly time correlated) in the residual of an elemental filter within an MMAE is indicative of that filter having the wrong actuator failure status hypothesis. This technique results in a delay in detecting the flight control actuator failure because several samples of the residual must be collected before the residual correlation can be estimated. However, it allows a significant reduction of the amplitude of the required system inputs for exciting the various system modes to enhance identifiability, to the point where they may possibly be subliminal, so as not to be objectionable to the pilot and passengers.

Manuscript received December 24, 1997; released for publication November 10, 1999.

IEEE Log No. T-AES/36/2/05216.

Refereeing of this manuscript was handled by S. Shrier.

Authors' address: Dept. of Electrical Engineering and Computer Science, Air Force Institute of Technology, 2950 P St., Bldg. 640, Wright-Patterson Air Force Base, OH 45433-7765, E-mail: (Peter.Maybeck@afit.af.mil).

U.S. Government work not supported by U.S. copyright.

0018-9251/00/\$10.00 2000 IEEE

I. INTRODUCTION

A multiple model adaptive estimator (MMAE) [1, 2, 6, 8, 9, 11] consists of a bank of parallel Kalman filters, each with a different model, and a hypothesis testing algorithm as shown in Fig. 1. Each of the internal models of the Kalman filters can be represented by a discrete value of a parameter vector (\mathbf{a}_k ; $k = 1, 2, \dots, K$). The Kalman filters are provided a measurement vector (\mathbf{z}) and the input vector (\mathbf{u}), and produce a state estimate ($\hat{\mathbf{x}}_k$) and a residual (\mathbf{r}_k). The hypothesis testing algorithm uses the residuals to compute conditional probabilities (p_k) of the various hypotheses that are modeled in the Kalman filters, conditioned on the history of measurements received up to that time, and to compute an estimate of the true parameter vector ($\hat{\mathbf{a}}$). The conventional MMAE computes conditional probabilities (p_k) in a manner that exploits three of four characteristics of Kalman filter residuals that are based on a correctly modeled hypothesis—that they should be Gaussian, zero-mean, and of computable covariance—but does not exploit the fact that they should also be white. The algorithm developed herein addresses this directly, yielding a complement to the conventional MMAE.

One application of MMAE is flight control sensor/actuator failure detection and identification, where each Kalman filter has a different failure status model (\mathbf{a}_k) that it uses to form the state estimate ($\hat{\mathbf{x}}_k$) and the residual (\mathbf{r}_k). The hypothesis testing algorithm assigns conditional probabilities (p_k) to each of the hypotheses that were used to form the Kalman filter models. These conditional probabilities indicate the relative correctness of the various filter models, and can be used to select the best estimate of the true system failure status, weight the individual state estimates appropriately, and form a probability-weighted average state estimate ($\hat{\mathbf{x}}_{\text{MMAE}}$).

A primary objection to implementing an MMAE-based (or other) failure detection algorithm is the need to dither the system constantly to enhance failure identifiability. The MMAE compares the magnitudes of the residuals (appropriately scaled to account for various uncertainties and noises) from the various filters and chooses the hypothesis that corresponds to the residual that has a history of having smallest (scaled) magnitude. Large residuals must be produced by the filters with models that are incorrect to be able to discount these incorrect hypotheses. The residual is the difference between the measurement of the system output and the filter's prediction of what that measurement should be, based on the filter-assumed system model. Therefore, to produce the needed large residuals in the incorrect filters, we need to produce a history of sufficiently large system outputs, so we need to dither the system constantly and thereby

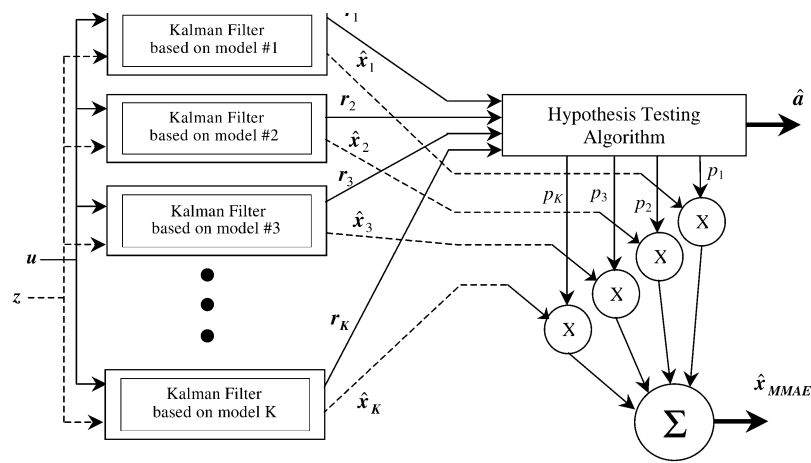


Fig. 1. MMAE algorithm.

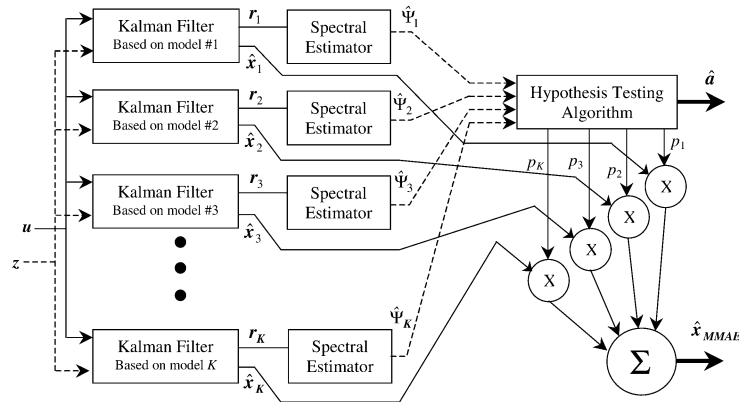


Fig. 2. MMAE algorithm using RCKFB.

excite the system states. For flight control failure applications, we would need to move the aircraft continually in all axes to produce the desired failure detection performance, to which most pilots (not to mention passengers) would strenuously object.

We propose constructing, as shown in Fig. 2, a Kalman filter bank, denoted the residual correlation Kalman filter bank (RCKFB), with outputs that are estimates of the power spectral density of each of the residuals from the Kalman filters in the bank. Previous research has shown that, if the Kalman filter model is correct, the residual is a white sequence with zero mean, and it is no longer zero-mean if the model is incorrect [4, 5]. This causes a change in the residual correlation, but not the covariance of the residual. If there are model differences in the control input matrix, the change in mean of the residual is a summation of input terms [4, 5]. If the dither input is a sinusoid, then these terms produce a residual sinusoidal component at the same frequency as the input. Thus, the appearance of a purposeful dither in the residual of the filter is indicative of an incorrect failure status hypothesis in that filter model, since any dithers in actual sensor outputs

(\mathbf{z}) would be compensated by the ($\mathbf{H}\hat{\mathbf{x}}(t_i^-)$) term in the residual, ($\mathbf{z} - \mathbf{H}\hat{\mathbf{x}}(t_i^-)$), if the assumed model were correct. This effect is clearly evident in Fig. 3 where the residual vector from the Kalman filter based on a model that assumes a fully functional aircraft, shows the presence of the elevator dither input, shown in Fig. 4, when an elevator failure occurs (note particularly the pitch and pitch rate residuals).

Since we *know* the frequency of the input, we can use the spectral content of the residual *at this particular frequency* to indicate the presence of mismodeling. Fig. 5 is the Fourier transform of 200 data points of the residual for the pitch rate residual element shown in Fig. 3 (solid line), along with the same residual element from the same Kalman filter when there is no failure (dotted line). The “spike” in the solid line occurs at the elevator dither input frequency. Note that, at this particular frequency, the spectral content of the residual with the mismodeling is significantly greater than the spectral content for the correctly modeled residual. This figure shows that the spectral content of the residual clearly indicates the presence or absence of the mismodeling.




Fig. 3. Time history of 9-dimensional residual vector from Kalman filter with fully functional aircraft model in presence of elevator failure occurring at the initial time (angles measured in rad, rates in rad/s, time in s).

II. THEORY DEVELOPMENT

A. Multiple Model Adaptive Estimation Equations

1) *Basic Kalman Filter Equations:* We assume a steady state Kalman filter model (and eventually a steady state constant-gain Kalman filter for implementation) associated with a particular hypothesized failure status, which is denoted with the subscript \mathbf{k} . Thus we have

$$\begin{aligned} \mathbf{x}_k(t_i) &= \Phi_k \mathbf{x}_k(t_{i-1}) + \mathbf{B}_k \mathbf{u}(t_{i-1}) + \mathbf{G}_k \mathbf{w}_k(t_{i-1}) \\ \mathbf{z}_k(t_i) &= \mathbf{H}_k \mathbf{x}_k(t_i) + \mathbf{v}_k(t_i) \end{aligned} \quad (1)$$

where

\mathbf{x}_k is the Kalman filter model state vector,
 Φ_k is the Kalman filter model state transition matrix,
 \mathbf{B}_k is the Kalman filter model control input matrix,
 \mathbf{u} is the system input vector,
 \mathbf{G}_k is the Kalman filter model noise input matrix,

\mathbf{w}_k is an additive white discrete-time dynamics noise input used in the Kalman filter model, with zero-mean and

$$E\{\mathbf{w}_k(t_i)\mathbf{w}_k^T(t_j)\} = \begin{cases} \mathbf{Q}_k, & t_i = t_j \\ \mathbf{0}, & t_i \neq t_j \end{cases} \quad (2)$$

\mathbf{z}_k is the Kalman filter model measurement vector,
 \mathbf{H}_k is the Kalman filter model output matrix,
 \mathbf{v}_k is an additive white measurement noise input that is used in the Kalman filter model.

This noise input is assumed to be independent of \mathbf{w}_k , and zero-mean with

$$E\{\mathbf{v}_k(t_i)\mathbf{v}_k^T(t_j)\} = \begin{cases} \mathbf{R}_k, & t_i = t_j \\ \mathbf{0}, & t_i \neq t_j \end{cases} \quad (3)$$

Note that the Kalman filter model and the truth model are both linear models, but the dimensionality of these two models may not necessarily be the same. In most cases, the Kalman filter model is a reduced order




Fig. 4. Time history of dither control inputs (in rad; time in s).




Fig. 5. Fourier transform of pitch rate residual from fully functional Kalman filter with mismatching (solid line) and no mismatching (dotted line).

version of the truth model (the Kalman filter model states are often a subset of the truth model states).

The Kalman filter algorithm uses this model to define time propagation and measurement update equations of the Kalman filter state estimates and the Kalman filter state estimate covariance matrix. The Kalman filter state estimate propagation equation based on the Kalman filter model is

$$\begin{aligned}\hat{\mathbf{x}}_{\mathbf{k}}(t_i^-) &= \Phi_{\mathbf{k}} \hat{\mathbf{x}}_{\mathbf{k}}(t_{i-1}^+) + \mathbf{B}_{\mathbf{k}} \mathbf{u}(t_{i-1}) \\ \hat{\mathbf{z}}_{\mathbf{k}}(t_i^-) &= \mathbf{H}_{\mathbf{k}} \hat{\mathbf{x}}_{\mathbf{k}}(t_i^-)\end{aligned}\quad (4)$$

where

$\hat{\mathbf{x}}_{\mathbf{k}}$ is the Kalman filter state estimate vector,
 $\hat{\mathbf{z}}_{\mathbf{k}}(t_i^-)$ is the Kalman filter estimate of the measurement vector before it becomes available,
 t_i^- is the time just before the measurement update at the i th time sample, and
 t_{i-1}^+ is the time just after the measurement update at the $(i-1)$ time sample,

and the state estimate covariance matrix propagation equation:

$$\mathbf{P}_{\mathbf{k}}(t_i^-) = \Phi_{\mathbf{k}} \mathbf{P}_{\mathbf{k}}(t_{i-1}^+) \Phi_{\mathbf{k}}^T + \mathbf{G}_{\mathbf{k}} \mathbf{Q}_{\mathbf{k}} \mathbf{G}_{\mathbf{k}}^T. \quad (5)$$

The Kalman filter state estimates are updated using:

$$\hat{\mathbf{x}}_{\mathbf{k}}(t_i^+) = \hat{\mathbf{x}}_{\mathbf{k}}(t_i^-) + \mathbf{K}_{\mathbf{k}}(t_i) \mathbf{r}_{\mathbf{k}}(t_i) \quad (6)$$

where the Kalman filter gain is

$$\mathbf{K}_{\mathbf{k}}(t_i) = \mathbf{P}_{\mathbf{k}}(t_i^-) \mathbf{H}_{\mathbf{k}}^T \mathbf{A}_{\mathbf{k}}(t_i)^{-1} \quad (7)$$

and the Kalman filter-computed residual covariance matrix $\mathbf{A}_{\mathbf{k}}$ is

$$\mathbf{A}_{\mathbf{k}}(t_i) = \mathbf{H}_{\mathbf{k}} \mathbf{P}_{\mathbf{k}}(t_i^-) \mathbf{H}_{\mathbf{k}}^T + \mathbf{R}_{\mathbf{k}}. \quad (8)$$

The Kalman filter residual vector, shown in (6), is defined as

$$\mathbf{r}_{\mathbf{k}}(t_i) = \mathbf{z}(t_i) - \mathbf{H}_{\mathbf{k}} \hat{\mathbf{x}}_{\mathbf{k}}(t_i^-) = \mathbf{z}_{\mathbf{T}}(t_i) - \mathbf{H}_{\mathbf{k}} \hat{\mathbf{x}}_{\mathbf{k}}(t_i^-) \quad (9)$$

which is simply the difference between the measurements (\mathbf{z}) and the Kalman filter estimates, based on its model, of those measurements before they are taken ($\mathbf{H}_{\mathbf{k}} \hat{\mathbf{x}}_{\mathbf{k}}(t_i^-)$). Finally, the Kalman filter state estimate covariance matrix is updated using

$$\mathbf{P}_{\mathbf{k}}(t_i^+) = \mathbf{P}_{\mathbf{k}}(t_i^-) - \mathbf{K}_{\mathbf{k}}(t_i) \mathbf{H}_{\mathbf{k}} \mathbf{P}_{\mathbf{k}}(t_i^-). \quad (10)$$

The steady state values of the Kalman filter estimate of the state covariance matrix can be precomputed by iterating (5), (7), (8), and (10) until steady state of the covariance and gain matrices is reached. Once this value for the state covariance matrix is found, the steady state Kalman filter gain $\mathbf{K}_{\mathbf{k}}$ and the steady state Kalman filter residual covariance matrix $\mathbf{A}_{\mathbf{k}}$ are computed using (7) and (8). With this steady state implementation, the state covariance matrix, the steady state Kalman filter gain, and the steady state Kalman filter residual covariance matrices

therefore do not need to be computed in real time. The steady state Kalman filter equations become

$$\hat{\mathbf{x}}_{\mathbf{k}}(t_i^-) = \Phi_{\mathbf{k}} \hat{\mathbf{x}}_{\mathbf{k}}(t_{i-1}^+) + \mathbf{B}_{\mathbf{k}} \mathbf{u}(t_{i-1}) \quad (11)$$

for propagating the state estimates and

$$\hat{\mathbf{x}}_{\mathbf{k}}(t_i^+) = \hat{\mathbf{x}}_{\mathbf{k}}(t_i^-) + \mathbf{K}_{\mathbf{k}} \mathbf{r}_{\mathbf{k}}(t_i) \quad (12)$$

for updating the state estimates..

2) *Hypothesis Testing Algorithm*: The hypothesis testing algorithm (HTA) simultaneously tests the residuals of the Kalman filter bank under multiple hypotheses. Previous research [12] has shown that, if the Kalman filter model matches the true system model, the residual has a mean of zero and a precomputable covariance matrix $\mathbf{A}_{\mathbf{k}}$. Therefore, this Kalman filter residual is a white Gaussian sequence of mean zero and covariance

$$\mathbf{A}_{\mathbf{k}} = \mathbf{H}_{\mathbf{k}} \mathbf{P}_{\mathbf{k}}^- \mathbf{H}_{\mathbf{k}}^T + \mathbf{R}_{\mathbf{k}}. \quad (13)$$

Therefore we get that the conditional density function of the measurement (\mathbf{z}) at t_i for the \mathbf{k} th Kalman filter, conditioned on the measurement history ($\mathbf{Z}(t_{i-1}) = [\mathbf{z}^T(t_1) \cdots \mathbf{z}^T(t_{i-1})]^T$), is

$$f_{\mathbf{z}(t_i)|\mathbf{h}, \mathbf{Z}(t_{i-1})}(\mathbf{z}_i | \mathbf{h}_{\mathbf{k}}, \mathbf{Z}_{i-1}) = \beta_{\mathbf{k}} \exp\{\cdot\}$$

where

$$\begin{aligned}\beta_{\mathbf{k}} &= \frac{1}{(2\pi)^{m/2} |\mathbf{A}_{\mathbf{k}}|^{1/2}} \quad \text{and} \\ \{\cdot\} &= \{-\frac{1}{2} \mathbf{r}_{\mathbf{k}}^T(t_i) \mathbf{A}_{\mathbf{k}}^{-1} \mathbf{r}_{\mathbf{k}}(t_i)\}.\end{aligned}\quad (14)$$

The scalar likelihood quotient is defined as

$$q_{\mathbf{k}}(t_i) = \mathbf{r}_{\mathbf{k}}^T(t_i) \mathbf{A}_{\mathbf{k}}^{-1} \mathbf{r}_{\mathbf{k}}(t_i). \quad (15)$$

We define the conditional probability for a particular hypothesis as

$$p_{\mathbf{k}}(t_i) = \Pr\{\mathbf{h} = \mathbf{h}_{\mathbf{k}} | \mathbf{Z}(t_i) = \mathbf{Z}_i\}. \quad (16)$$

We can compute the conditional probability for a particular hypothesis by [10, 13]:

$$p_{\mathbf{k}}(t_i) = \frac{f_{\mathbf{z}(t_i)|\mathbf{h}, \mathbf{Z}(t_{i-1})}(\mathbf{z}_i | \mathbf{h}_{\mathbf{k}}, \mathbf{Z}_{i-1}) \cdot p_{\mathbf{k}}(t_{i-1})}{\sum_{j=1}^K f_{\mathbf{z}(t_i)|\mathbf{h}, \mathbf{Z}(t_{i-1})}(\mathbf{z}_i | \mathbf{h}_{\mathbf{j}}, \mathbf{Z}_{i-1}) \cdot p_{\mathbf{j}}(t_{i-1})}. \quad (17)$$

In this equation we use the prior conditional probabilities, $p_{\mathbf{k}}(t_{i-1})$, to weigh the conditional densities of the current measurements, assuming each hypothesis, and then normalize it over the complete set of such numerator terms. For failure identification applications, where we usually want to choose the most likely hypothesis out of the set of possible hypotheses, we can choose the hypothesis with the largest conditional probability. These conditional probabilities can also be used to weight and blend the various hypotheses, depending on the particular application.

In practice, these conditional probabilities will fluctuate rapidly from one time sample to the next. If the hypothesis testing algorithm is choosing the hypothesis with the highest conditional probability, these fluctuations could cause momentary incorrect hypothesis declarations. To alleviate this phenomenon, the conditional probability calculations are modified and compared with a decision threshold. The decision threshold is defined such that a hypothesis is chosen only if its conditional probability is greater than the threshold. The threshold is set to avoid momentary incorrect hypothesis declarations while hopefully providing adequate hypothesis testing performance. The performance is usually measured by the time it takes to detect and declare the correct flight control failure for the failure identification application.

B. Residual Correlation Kalman Filter Bank

1) *Periodogram Equations:* Kay develops several spectral estimation techniques; we have chosen to use the periodogram [7] for this research since its characteristics are well researched and it is used as a conceptual basis for many other spectral estimation techniques. One version of the periodogram utilizes the fast Fourier transform, which could greatly aid in implementing this technique because this transform is available on commercial chip sets. Other spectral estimation techniques still need to be researched to determine which one or ones produce the desired performance. The periodogram is based on estimating the autocorrelation of the residual and then taking the discrete Fourier transform of the autocorrelation to produce an estimate of the power spectral density.

We must first make some assumptions to be able to estimate the autocorrelation of the residual. We assume that the residual sequence is a series of samples of a stationary process, so that the probability distributions do not change with time. We also assume that the residual is ergodic in the autocorrelation function, which implies that the expected value of the time-averaged autocorrelation function is the same irrespective of the length of time averaging.

For a multidimensional sequence, the ergodic estimate of the autocorrelation is

$$\hat{\mathbf{A}}_{\mathbf{k}}(p) = \frac{1}{N} \sum_{n=0}^{N-1-|p|} \mathbf{r}_{\mathbf{k}}(t_{i-n}) \mathbf{r}_{\mathbf{k}}^T(t_{i-n-|p|}). \quad (18)$$

The periodogram is the discrete Fourier transform of this sequence, thus:

$$\begin{aligned} \hat{\Psi}_{\mathbf{k}}(f; t_i) &= \sum_{p=-(N-1)}^{N-1} \hat{\mathbf{A}}_{\mathbf{k}}(p) \exp(-j2\pi f p) \\ &= \sum_{p=-(N-1)}^{N-1} \frac{1}{N} \sum_{n=0}^{N-1-|p|} \mathbf{r}_{\mathbf{k}}(t_{i-n}) \mathbf{r}_{\mathbf{k}}^T(t_{i-n-|p|}) \\ &\quad \times \exp(-j2\pi f p) \end{aligned} \quad (19)$$

where N is the number of data samples that are collected over time.

Note that by dividing by N instead of $(N-1)$ or $(N-|p|-1)$ in (18), we are using a biased estimate of the autocorrelation. We need the $1/N$ factor to make this a viable estimate of the power spectral density.

If the residual sequence is a scalar sequence, then (19) can be shown to be equivalent to

$$\hat{\Psi}_{\mathbf{k}}(f; t_i) = \frac{1}{N} \left| \sum_{n=0}^{N-1} \mathbf{r}_{\mathbf{k}}(t_{i-n}) \exp(-j2\pi f p) \right|^2 \quad (20)$$

which is simply the squared absolute value of the N -point Fourier transform of the residual sequence. This is much easier to implement and executes much faster than (19), since it exploits the fast Fourier transform routines. If the residual sequence is multidimensional, we can approximate (19) as a diagonal matrix with components given by (20) if the cross-correlation terms between elements of the residual vector are negligible. For the specific flight control application that we are studying, most of these cross terms are negligible. Since we are attempting to find practical implementations of the MMAE algorithm, and (19) is computationally intensive, we chose to implement (20) for this research. For the remainder of this paper, $\hat{\Psi}_{\mathbf{k}}(f; t_i)$ will denote an m -vector of scalar components given by (20) where m is the number of elements of the residual vector.

2) *Modified Hypothesis Testing Algorithm:* We propose altering the Kalman filter bank by estimating the spectral content of each of the residuals using (20). This structure is shown in Fig. 2, where the HTA is a slightly modified version of the HTA in Fig. 1. Kay [7] shows that this can be interpreted as filtering the residual with a bandpass filter centered at f and a 3 dB bandwidth of $1/N$, sampling the output, and computing the squared magnitude. The $1/N$ factor is needed to make the estimate a power spectral density. Clearly, as more data samples are used (N increases) the filter bandwidth narrows and more of the out-of-bandwidth noise is rejected. To use this estimate of the power spectral density in the hypothesis testing algorithms that are commonly used in MMAE, namely (9) and (13)–(17), we consider $\hat{\Psi}_{\mathbf{k}}(f; t_i)$ to play the role of $\mathbf{z}(t_i)$, and then must subtract its conditional mean to form the analog of $\mathbf{r}_{\mathbf{k}}(t_i)$ in (9), compute its covariance (the analog of (13)), and justify the Gaussian form of (14).

The conditional mean of $\hat{\Psi}_{\mathbf{k}}(f; t_i)$, denoted here as $\bar{\hat{\Psi}}_{\mathbf{k}}(f; t_i)$, can be computed analytically as in [4], but this is computationally burdensome. Alternatively, each of its scalar components can be approximated with the value corresponding to the average “floor” value obvious in Fig. 5 (or its point value at f) computed *a priori* in performance analyses assuming no mismodeling or computed adaptively in real time.

The covariance of $\Psi_{\mathbf{k}}(f; t_i)$ can similarly be evaluated analytically [4] at significant computational expense, or approximated with sample statistics; unlike $\mathbf{A}_{\mathbf{k}}$ in (14), it may be time varying rather than constant.

To justify the Gaussianity of $\hat{\Psi}_{\mathbf{k}}(f; t_i)$, note that in (19) and (20) we are multiplying two Gaussian random variables together, which will yield a generalized chi-squared distributed random variable, and then summing several of the chi-squared random variables together (we used $N = 100$ for this research). Since the samples of the residual are independent (if the \mathbf{k} th hypothesis is assumed correct) and identically distributed ($\mathbf{A}_{\mathbf{k}}$ given by (13) is assumed constant), we can use the Lindeburg–Levy Theorem [18] to show that the distribution of $\hat{\Psi}$ converges in distribution to a normal distribution. We generated histograms of $\hat{\Psi}$ using 120 data points and observed that, for the case where the residual has a non-zero mean (when a failure occurs), this approximation works well. However, when the residual is zero-mean (when the Kalman filter model is based on the correct hypothesis), then $\hat{\Psi}$ appears to be more chi-squared distributed. Further research needs to be accomplished to characterize the distribution of $\hat{\Psi}$ properly for a small number of data samples.

In this particular application, we can evaluate $\hat{\Psi}_{\mathbf{k}}(f; t_i)$ at $f = f_0$, where f_0 is the *known* frequency of the input, which gives us the modified version of (14):

$$f_{\psi(f_0; t_i) | \mathbf{h}, \mathbf{Z}(t_{i-1})}(\psi_i | \mathbf{h}_{\mathbf{k}}, \mathbf{Z}_{i-1}) = \beta_{\mathbf{k}\psi} \exp\{\cdot\} \quad (21)$$

where

$$\beta_{\mathbf{k}\psi} = \frac{1}{(2\pi)^{m/2} |\mathbf{A}_{\mathbf{k}\psi}|^{1/2}},$$

$$\{\cdot\} = \left\{ -\frac{1}{2} \mathbf{r}_{\mathbf{k}\psi}^T(f_0; t_i) \mathbf{A}_{\mathbf{k}\psi}^{-1} \mathbf{r}_{\mathbf{k}\psi}(f_0; t_i) \right\}$$

$$\text{and } \mathbf{r}_{\mathbf{k}\psi}(f_0; t_i) = \hat{\Psi}_{\mathbf{k}}(f_0; t_i) - \overline{\hat{\Psi}_{\mathbf{k}}}(f_0; t_i)$$

and (15) becomes

$$q_{\psi}(t_i) = \mathbf{r}_{\mathbf{k}\psi}^T(f_0; t_i) \mathbf{A}_{\mathbf{k}\psi}^{-1} \mathbf{r}_{\mathbf{k}\psi}(f_0; t_i) \quad (22)$$

with the power spectral density estimate covariance matrix denoted $\mathbf{A}_{\mathbf{k}\psi}$.

We observed that, for the dither input levels that we used for this research, the mean for the no-failure hypothesis is very small when compared with the mean of the spectral density estimate for an actuator failure hypothesis. This is clearly seen in Fig. 5, where the no-failure case has a mean of about 1, while the failed case has a mean of about 40. For *this* application, we found that the change in mean effects of $\hat{\Psi}$ was much greater than the effect of using an incorrect distribution. This was determined by first modeling the power spectral density estimate as a zero-mean process, the $\hat{\Psi}_{\mathbf{k}}(f; t_i)$ in (21) is set to a zero vector, and then modeling it as a non-zero mean

process whose mean, $\Psi_{\mathbf{k}}(f; t_i)$ in (21), is estimated using sample statistics. In both cases the power spectral density estimate covariance matrix, denoted $\mathbf{A}_{\mathbf{k}\psi}$, was assumed constant and also computed using 400 samples of the fully functional residual power spectral density estimate. Many software packages, such as MATLAB, have an intrinsic sample covariance computation available, which makes this method quite easy to implement. We found the failure detection performance, to be described in Section IIIB, for each of these cases was indistinguishable, thus demonstrating that the change in mean effects of $\hat{\Psi}$ was much greater than the effect of using an incorrect distribution.

At this point, the development of the HTA for the RCKFB, (21) and (22), parallels the development of the HTA for the standard Kalman filter bank (SKFB), (13)–(17). Other research [14] has described several modifications to the HTA, such as β stripping, lower bounding of the conditional probabilities, filter tuning, smoothing of the conditional probabilities, exponential penalty increase, and propagating several time samples before updating. The same modifications used for the HTA were applied to the RCKFB, however, further research is needed to define the specific effects of each of the modifications on the performance of the MMAE using the RCKFB.

III. SIMULATION RESULTS

We present a comparison of the failure identification performance of two MMAE structures, one that uses the SKFB, shown in Fig. 1, and the other using the RCKFB, shown in Fig. 2, that was developed in Section IIB. The failure identification performance of the SKFB is presented in Section IIIA and the RCKFB performance in Section IIIB. Section IIIC is a comparison of the failure identification performances for these two structures.

Previous research [4, 5] has found that the mean of the residual will have elements of the input for the case of an actuator failure, and for a sensor failure the residual will have elements of the state estimates. If we use a sinusoid for the input, then the sinusoid input elements cause the residual in a filter with the incorrect hypothesis to be a sinusoid. This sinusoid appears as a spike in the power spectral density at the frequency of the input, which was shown in Fig. 5. Therefore, since we can stipulate the frequency of the system input (since we can generate the dither input ourselves), we can use the spectral content at that frequency to indicate the presence of an actuator failure, which was the basis for developing the RCKFB. However, we cannot readily stipulate the frequency of the state estimates, unless extensive simulation is done to find the particular input that causes sinusoidal state estimates. Since the state estimates are not necessarily sinusoidal, we might




Fig. 6. Overall failure identification of SKFB: conditional probability for each hypothesized failure, vs. time (s).

need to estimate the power spectral density across the entire spectrum to try to find indications of the presence of a sensor failure. Therefore, we believe that this technique would not work well for sensor failures unless the appropriate input is used that causes the states estimates to become strongly correlated in a readily distinguishable manner. For this reason, we implemented the RCKFB structure using only actuator failure models and restricted the SKFB structure to actuator failure hypotheses only, to make a fair comparison of failure identification algorithms. The usual implementation of these structures includes sensor failure hypotheses [3, 14–17].

Previous research [14] found input levels, shown in Fig. 4, that worked well for flight control failure identification. Such a nonsubliminal input dither is very objectionable, so we investigated the failure identification performance of the SKFB and RCKFB for several different input levels. We used several multipliers, ranging from 0.01 to 1.0, to generate these various input levels.

A. SKFB Failure Identification Performance

The failure identification performance of the SKFB structure for this particular application was

previously researched [14] and some enhancement through the modifications listed in Section IIB2 was accomplished. We present a summary of these results to facilitate performance comparisons with the RCKFB structure.

We found, in the previous research [14], that the results from all of the failure cases can be summarized by observing only the conditional probability of the particular failure that is actually occurring, because the fully functional hypothesis mirrors this probability and all other hypothesized failures are near zero throughout the simulation. The overall performance of the SKFB failure identification performance can be summarized by showing the conditional probabilities for each of the failure hypotheses on a series of plots, shown in Fig. 6, where each plot shows the SKFB failure identification performance for a particular failure occurring at 1.0 s into the simulation.

Detection of sensor failures occurs quite rapidly with the exception of the forward velocity sensor. This was due to the large measurement noise of that particular sensor when compared with the small change in forward velocity of the aircraft during the simulation time [14]. This caused the conditional

probabilities for this hypothesis to grow very slowly, as reflected in Fig. 6.

At this point we restricted our research to actuator failures due to reasons that were mentioned earlier. This slightly changed the actuator plots of Fig. 6 by removing the small “ripples” that are evident in the aileron and rudder plots, which was due to some bleeding of the conditional probabilities into the sensor failure hypotheses. However, removing the sensor failure hypotheses did not perceptibly change the point where the conditional probabilities started to grow. Thus the conditional probabilities shown in the top part of Fig. 6 are a good representation of the actuator failure detection performance for the SKFB MMAE structure. One simple method of declaring a failure is to choose a probability threshold that is used as a declaration trigger. For example, if the threshold was set to a probability of 0.5 and a left elevator failure occurred, when the left elevator conditional probability exceeds 0.5, then a failed left elevator would be declared. Fig. 6 shows that this would occur at about 1.5 s into the simulation, or about 0.5 s after the failure actually occurred. Using this threshold, we averaged the actuator failure performance of this MMAE structure over 5 Monte Carlo runs, and then averaged the left and right actuator failure performance results to get an average over 10 test samples. We averaged the left and right actuator failure performances because we assume that they are independent. This assumption is based on the orthogonal system input (Fig. 4), where the left actuator input is almost exactly out-of-phase with the right input (a slight offset from exactly out-of-phase was needed to allow some excitation of the sensors).

Several simulations were accomplished using different control input strengths which were generated by multiplying the original control input, shown in Fig. 4, by various multipliers ranging from 1.0 to 0.01. This maintained the frequency and phase of the original control input, but reduced the amplitude to the point where it could possibly be considered subliminal. The results of these simulations are graphically shown in Figs. 7, 8, and 9, for failures of the elevators, ailerons, and rudders, respectively. Note that the vertical axis is a logarithmic (base 10) scale, and that the SKFB results are given by the locus demarked with x's, on each of these figures.

The identification time is the amount of time from the occurrence of the failure to when the conditional probability crosses the probability threshold at 0.5. At very small inputs, the conditional probability grew so slowly, that the randomness of the probabilities would cause it to cross the threshold a few times. For these cases, the identification time is the amount of time from the occurrence of the failure to when the conditional probability crosses the probability threshold for the last time and stays there.




Fig. 7. Comparative failure identification performance for elevator failure.




Fig. 8. Comparative failure identification performance for aileron failure.




Fig. 9. Comparative failure identification performance for rudder failure.




Fig. 10. Left elevator Kalman filter residual for left elevator failure at 1 s (angles in rad, rates in rad/s; versus time in s).

Note the dashed line portion of the plots in Figs. 7, 8, and 9. These indicate where false alarms were observed, with an increase in the number of false alarms as the input was decreased. The dashes are to show where the performance may be deemed unacceptable due to the increase in the false alarm rate.

These figures show that failure performance drops off dramatically as the input magnitude decreases. This is particularly evident for the elevator and aileron failures (Figs. 7 and 8), where identification times are beyond the termination of the simulation (7 s after the failure) for input multipliers of less than 0.025. These results make good sense, because stronger inputs would cause larger residuals from the incorrect Kalman filters when the failure occurs. Likewise, smaller inputs will produce relatively smaller residuals when the failure occurs, so the conditional probability of the failure hypothesis requires more time to grow. In particular, note that the rudder failure identification

(Fig. 9) is still viable at an input multiplier level of 0.025, but drops off quickly after that.

Note in Fig. 9 that, as the rudder input is increased, the failure identification time actually increases, which does not correspond to the anticipated results of better identification performance for stronger inputs. Other research [14] found that this algorithm had trouble identifying rudder failures, possibly because of cross-coupling between axes causing misidentification of a sensor failure. This is not the case here, since, at this point, we have limited the SKFB to actuator failure hypotheses to make a fair comparison with the RCKFB structure.

There are two reasons that the failure identification time increases with increasing rudder inputs, one is the relatively stronger rudder input (Fig. 4), compared with the elevator and aileron inputs, and the other is the timing of the failure. Fig. 3 shows that a mismodeled actuator causes the input of that particular actuator to be reflected in the residuals. A relatively




Fig. 11. Actual failure identification performance of RCKFB: conditional probabilities for each hypothesized failure vs. time (s).

stronger rudder input, compared with the elevator and aileron inputs, would cause much larger residuals in the residuals of the Kalman filter that mismodeled the rudder actuators. These larger residuals take longer to decay once the rudder actuators are appropriately modeled. Likewise, the timing of the failure also influences the failure identification performance.

Fig. 10 shows that a rudder failure at 1 s into the simulation occurs when the pitch rate and pitch angle residuals are close to their peak. This requires more time for the residual to decay than would have been required if the failure occurred when the residuals were near a zero-crossing. If the Kalman filter residual with the correct model takes longer to decay, the failure identification time will correspondingly increase.

B. RCKFB Failure Identification Performance

We implemented the RCKFB with a modified HTA shown in Fig. 2. We used the form of the RCKFB, shown in (20), that exploits the fast Fourier transform, but ignores the cross-coupling terms of the residual. These cross-coupling terms are not small in many cases. For example, the residual for a left elevator failure (Fig. 3) shows that both the pitch rate and pitch angle strongly reflect the elevator input. Note that the pitch rate is clearly a sinusoid, and would be strongly correlated in time with itself. The pitch angle is also a sinusoid at the same frequency, so the cross-correlation between these two sinusoids would be quite strong. By ignoring these cases of strong cross-correlations, we are implementing an algorithm that does not exploit any of the information available from these cross terms. We considered the computational cost of exploiting these cross terms to be unacceptable, so we implemented the algorithm

using only the autocorrelation of the individual elements of the residual.

A summary of the failure detection performance, following the format of Fig. 6, is presented in Fig. 11. As previously noted, sensor failure detection and identification was not investigated.

In Section II we noted that Kay [7] shows that this algorithm essentially filters the residual with a bandpass filter, centered at the frequency of interest and with a bandwidth of $1/N$. If the bandwidth of this filter is small, the other input frequencies will not have spectral content in the bandwidth of this bandpass filter. Thus, we need to collect enough data samples so that the bandwidth of the filter will not have any spectral content from the other inputs. The more data samples that we need to collect, the longer the delay will be from the time when the sinusoid is present (when the failure occurs) to when we can compute the spectral estimate of this filter and detect the sinusoid. Thus, there is a tradeoff between the number of data samples needed to distinguish between the input sinusoids and the failure detection time. This would dictate a wide spacing of the input sinusoids, to allow the bandpass filters to have larger bandwidths, so fewer data samples are needed. Unfortunately, most mechanical system (like aircraft) have very narrow low pass system response bandwidths (the bandwidth was about 2.2 Hz for this particular aircraft). The system response to the input must be seen in the residual for any failure identification scheme to work, thus we need to specify input frequencies that are within the system response bandwidth. This requires that the inputs be located close together in the frequency spectrum, which requires narrow bandpass filters for this structure, which in turn requires more data samples to be collected. Our inputs were at 0.5, 1, and 2 Hz, which, when normalized by dividing

by the sampling frequency, become 0.01, 0.02, and 0.04, respectively. The two closest frequencies were 0.5 Hz away from each other, or 0.01 in normalized frequency. The bandpass filter would need to have a bandwidth of less than 1 Hz (0.5 Hz on each side of the center frequency), which is 0.02 in normalized frequency, to be able to distinguish these input frequencies from each other. This requires 100 data samples ($2/N = 0.02 \Rightarrow N = 100$) for the required bandpass filter bandwidth, which is the value of N chosen for this research. This value caused a delay in the detection of the failure because at least half of the 100 data samples (which translates to 1 s) had to show strong correlation at the input sinusoid frequency for the spectral estimator even to start to identify the presence of the sinusoid. Fig. 11, when compared with Fig. 6, shows that this delay caused this algorithm to take about 1 s longer to identify the actuator failures.

The aircraft model that we used for this simulation assumed complete axis decoupling. If this assumption is correct, we can increase the bandwidth of the bandpass filters, thus decreasing the number of required data samples. The elevator input was at a frequency of 1 Hz, while the rudder input was at 0.5 Hz and the aileron input was at 2 Hz. Since the pitch axis would not affect the roll and yaw axes, the elevator input would not appear in the spectra for the longitudinal elements of the residual (roll and yaw). Thus the elevator input frequency did not need to be considered when we were designing the bandwidth of the bandpass filter. Thus, the two closest frequencies were actually 1.5 Hz apart, which means that we could have used only 34 data samples to distinguish between the rudder and aileron input frequencies. Time constraints prevented testing the failure identification performance using other values of N .

We tested the failure performance of this structure at the various input strengths that were used for the SKFB testing. These results are graphically presented in Figs. 7, 8, and 9 by the loci demarked with o's. Note that the results are an average of only two Monte Carlo runs with both left and right actuators, for a total of four test samples. We observed a large drop in failure identification time for the strong rudder inputs. The reason for this phenomenon is illustrated in Fig. 10. Note, when the failure occurs at 1 s, the pitch rate and pitch angle elements of the residual are at their peak and take about 0.5 s to die out. During this decay, these waveforms appear to be part of the sinusoid that was clearly present before the failure occurred. The sinusoid in the fully functional Kalman filter residual (Fig. 3) is also delayed by about 0.5 s. Thus, the left elevator failure Kalman filter residual will show a slowly decaying spectral content at the elevator input frequency, and the fully functional Kalman filter residual will show a slowly building spectral content at that frequency. Thus, the HTA

will not identify the failure until the decay is nearly complete.

C. Comparative Failure Identification Performance

The simulation results for both the SKFB and the RCKFB are graphically shown in Figs. 7, 8, and 9 to facilitate a comparison of the failure identification performance of the two structures. Note that, for a strong input level, the RCKFB structure took longer to identify the flight control actuator failures because 100 data samples had to be collected and the spectral estimate computed, which delayed the failure identification. Collecting a smaller number of data samples at these high input levels might produce faster failure identification times.

As the input decreases, the failure identification performance generally decreases for the SKFB structure, and seems to remain fairly constant for the RCKFB structure. The exception to the general rule is the decrease in rudder failure identification performance as the input strength increases for both the SKFB and RCKFB structures. This was caused by both the relative strength of the rudder input, which caused a larger residual amplitude, and the timing of the failure, which occurred just as the residual peaked, so the residual took longer to decay.

Both structures have a minimum input level below which their performance diminishes rapidly. As the input strength is decreased, the approach of this minimum input level for good failure identification performance is signaled by an increase in the number of false alarms. This suggests that the false alarm rate could be monitored (by using statistical testing for estimating the false alarm rate) and this could be used as feedback to set the input level appropriately. Note also, that the minimum failure time identification for the SKFB, particularly for a rudder failure, could be obtained by setting the input level just above the point where the false alarm rate begins to increase and the failure identification performance starts to deteriorate.

The important feature to notice in Figs. 7, 8, and 9 is that the RCKFB provides the best failure identification performance at low input levels. In particular, only the RCKFB structure still provided failure identification performance for input levels below 1/20 of the original input used for this research.

V. CONCLUSIONS

We developed a new kind of multiple model filter bank structure in Section IIB, in which an estimation of the spectral content of the residuals is used by the hypothesis testing algorithm. The actuator failure identification performance of this RCKFB was tested and the results were presented in Section IIIB, and a comparison of this performance against the standard MMAE structure was made in Section IIIC. When

an actuator failure occurs, elements of the control input are added to the residual of a filter with an incorrect hypothesis. If the control input elements are sinusoids, then the residual includes a sinusoid, with the same frequency as the control input elements. Therefore, we can use the spectral content of the residual, at the input frequencies, to indicated the presence of an actuator failure. We found that, when the control input is fairly strong, this technique takes longer than the standard MMAE to make the correct failure identification, due to the time used to collect a sufficient number of samples to generate an adequate spectral estimate. However, this technique provides failure identification performance at small input levels where the standard MMAE no longer functions. This technique looks quite promising as a method to provide good failure identification with subliminal input dithers, which could be used in conjunction with the standard MMAE structure to provide failure identification performance at low input levels and to corroborate the failure identification of the SKFB structure at higher input levels.

These results strongly suggest a combined SKFB and RCKFB MMAE. The same filters in the SKFB MMAE structure, Fig. 1, would be used as the filters in the RCKFB MMAE structure in Fig. 2. Thus a single bank of Kalman filters would be implemented with the residuals used as a direct input to the standard HTA and to the spectral estimators with a modified HTA. The outputs of these two structures would be combined so that at strong input levels the fast failure identification of the SKFB structure would be exploited, and at low input levels the complementary RCKFB performance would be exploited. This suggests that the SKFB MMAE structure would provide good failure identification during maneuvers and the RCKFB MMAE structure could corroborate the failure detection of the SKFB and provide good flight control actuator failure identification, using a subliminal dither, during nonmaneuvering flight profiles. Further research needs to accomplished to determine how to combine these structures appropriately.

REFERENCES

- [1] Chang, C. B., and Athans, M. (1977)
Hypothesis testing and state estimation for discrete systems with finite-valued switching parameters. Publication ESL-P-758, MIT Electronic Systems Laboratory, Lexington, MA, June 1977.
- [2] Chang, C. B. (1978)
State estimation for discrete systems with switching parameters.
IEEE Transactions on Aerospace and Electronic Systems, **AES-14**, 4 (May 1978), 418–425.
- [3] Eide, P. K., and Maybeck, P. S. (1996)
An MMAE failure detection system for the F-16.
IEEE Transactions on Aerospace and Electronic Systems, **32**, 3 (July 1996), 1125–1136.

- [4] Hanlon, P. D. (1996)
Practical implementation of multiple model adaptive estimation using Neyman–Pearson based hypothesis testing and spectral estimation tools.
Ph.D. dissertation, AFIT/DS/ENG/96-07, School of Engineering, Air Force Institute of Technology, Wright-Patterson AFB, OH, Sept. 1996.
- [5] Hanlon, P. D., and Maybeck, P. S. (2000)
Characterization of Kalman filter residuals in the presence of mismodeling.
IEEE Transactions on Aerospace and Electronic Systems, **36**, 1 (2000).
- [6] Hawkes, R. M., and Moore, J. B. (1976)
Performance bounds for adaptive estimation.
Proceedings of the IEEE, **64** (Aug. 1976), 1143–1150.
- [7] Kay, S. M. (1988)
Modern Spectral Estimation.
Englewood Cliffs, NJ: Prentice Hall, 1988.
- [8] Lainiotis, D. G. (1976)
Partitioning: A unifying framework for adaptive systems, I: Estimation.
Proceedings of the IEEE, **64** (Aug. 1976), 1126–1142.
- [9] Lashlee, R. W., Jr., and Maybeck, P. S. (1988)
Spacestructure control using moving bank multiple model adaptive estimation.
In *Proceedings of the IEEE Conference on Decision and Control*, San Antonio, TX, Dec. 1988.
- [10] Li, X. R., and Bar-Shalom, Y. (1996)
Multiple-model estimation with variable structure.
IEEE Transactions on Automatic Control, **41**, 4 (Apr. 1996), 482–439.
- [11] Magill, D. T. (1965)
Optimal adaptive estimation of sample stochastic processes.
IEEE Transactions on Automatic Control, **AC-10**, 4 (Oct. 1965), 434–439.
- [12] Maybeck, P. S. (1979)
Stochastic Models, Estimation, and Control, Vol. 1.
New York: Academic Press, 1979. Republished Arlington, VA: Navtech, 1994.
- [13] Maybeck, P. S. (1979)
Stochastic Models, Estimation, and Control, Vol. 2.
New York: Academic Press, 1979. Republished Arlington, VA: Navtech, 1994.
- [14] Maybeck, P. S., and Hanlon, P. D. (1995)
Performance enhancement of a multiple model adaptive estimator.
IEEE Transactions on Aerospace and Electronic Systems, **31**, 4 (Oct. 1995), 1240–1253.
- [15] Maybeck, P. S., and Pogoda, D. L. (1989)
Multiple model adaptive controller for the STOL F-15 with sensor/actuator failures.
In *Proceedings of the IEEE Conference on Decision and Control*, Tampa, FL, Dec. 1989, 1566–1572.
- [16] Maybeck, P. S., and Stevens, R. D. (1991)
Reconfigurable flight control via multiple model adaptive control methods.
IEEE Transactions on Aerospace and Electronic Systems, **27**, 3 (May 1991), 470–480.
- [17] Menke, T. E., and Maybeck, P. S. (1995)
Sensor/actuator failure detection in the VISTA F-16 by multiple model adaptive estimation.
IEEE Transactions on Aerospace and Electronic Systems, **31**, 4 (Oct. 1995), 1218–1229.
- [18] Scharf, L. L. (1991)
Statistical Signal Processing.
Reading, MA: Addison-Wesley, 1991.

Peter D. Hanlon was born on 1 June 1956 in Woburn, MA. He attended Burlington County College and Rutgers University, Cook College, where he studied meteorology. He received the B.S. in engineering in 1985 from the University of Central Florida, Orlando, and the M.S. and Ph.D. in electrical engineering in 1992 and 1996, respectively, from the Air Force Institute of Technology.



He enlisted in the Air Force in 1979, trained as a voice processing specialist, and was assigned to Athens, Greece in 1981 where he flew as a crew member on RC-135s. He was accepted into the Airman's Education and Commissioning Program (AECPP) in 1983 and was assigned to the University of Central Florida, Orlando, FL until his graduation in 1985. After his commissioning, he was assigned to the Geophysics Laboratory at Hanscom Air Force Base, MA, where he served as either program manager or project engineer for several space shuttle, sounding rocket, and high altitude balloon research projects until 1989, when he was assigned to the staff of the Aeronautical Systems Division, Wright-Patterson Air Force Base. He attended the Air Force Institute of Technology from 1991–1996. After serving as an engineer at the Wright Laboratory, Eglin Air Force Base, Florida, he joined the faculty of the Department of Electrical Engineering and Computer Engineering at the United States Military Academy at West Point, where he is currently an Assistant Professor.

Peter S. Maybeck (M'74—SM'84—F'87) was born in New York, NY on February 9, 1947. He received the B.S. and Ph.D. degrees in aeronautical and astronautical engineering from M.I.T., Cambridge, in 1968 and 1972, respectively.

In 1968, he was employed by the Apollo Digital Autopilot Group of The C. S. Draper Laboratory, Cambridge, MA. From 1972 to 1973, he served as a military control engineer for the Air Force Flight Dynamics Laboratory and then joined the faculty of the Air Force Institute of Technology in June of 1973. He is currently Professor of Electrical Engineering, responsible for the graduate sequence in estimation and stochastic control and for individual advanced digital filtering and control courses. Current research interests concentrate on using optimal estimation techniques for guidance systems, tracking, adaptive systems, and failure detection purposes.

Dr. Maybeck is the author of numerous papers on applied optimal filtering as well as the book, *Stochastic Models, Estimation and Control* (Academic Press, Vol. 1–1979, Vols. 2 and 3–1982; republished by Navtech in 1994). He is a member of Tau Beta Pi, Sigma Gamma Tau, Eta Kappa Nu, and Sigma Xi. He was recipient of the DeFlorez Award (ingenuity and competence of research), the James Means Prize (excellence in systems engineering) and the Hertz Foundation Fellowship at M.I.T. in 1968. In all years from 1975 to 1998, he received commendation as outstanding Professor of Electrical Engineering at A.F.I.T. In December of 1978, he received an award from the Affiliate Societies Council of Dayton as one of the twelve outstanding scientists in the Dayton, OH area. In March of 1980, he was presented with the Eta Kappa Nu Association's C. Holmes MacDonald Award, designating him as the outstanding electrical engineering professor in the United States under the age of 35 (he had placed second in this national competition for 1977 as well). In 1985, he received the Frederick Emmons Terman Award, the highest national award to a Professor of Electrical Engineering given by the American Society of Engineering Education. He is a member of the A.I.A.A., and he is the current I.E.E.E. Dayton Section Student Activities chairman and a member of the I.E.E.E. Executive Committee of Dayton, and he previously served as Chairman of the local Automatic Control Group.

