Sensor/Actuator Failure Detection in the Vista F-16 by Multiple Model Adaptive Estimation

TIMOTHY E. MENKE

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

Multiple model adaptive estimation (MMAE) is applied to the Variable In-flight Stability Test Aircraft (VISTA) F-16 flight control system at a low dynamic pressure flight condition (0.4 M at 20000 ft). Single actuator and sensor failures are addressed first, followed by dual actuator and sensor failures. The system is evaluated for complete or "hard" failures, partial or "soft" failures, and combinations of hard and soft actuator and sensor failures. Residual monitoring is discussed for single and dual failure scenarios. Performance is enhanced by the application of a modified Bayesian form of MMAE, scalar residual monitoring to reduce ambiguities, automatic dithering where advantageous, and purposeful commands.

Manuscript received September 18, 1992; revised April 12, 1994.

IEEE Log No. T-AES/31/4/14135.

Authors' address: Department of Electrical and Computer Engineering, Air Force Institute of Technology/ENG, Wright-Patterson AFB, OH 45433.

0018-9251/95/$10.00 © 1995 IEEE

I. INTRODUCTION

For many applications, it is highly desirable to develop an aircraft flight control system with reconfigurable capabilities: able to detect and isolate failures of sensors and/or actuators and then to employ a control algorithm that has been specifically designed for the current failure mode status. One means of accomplishing this, in a manner that is ideally suited to distributive computation, is multiple model adaptive estimation (MMAE) [1–4] and control (MMAC) [5–7].

Assume that the aircraft system is adequately represented by a linear perturbation stochastic state model, with a (failure status) uncertain parameter vector affecting the matrices defining the structure of the model or depicting the statistics of the noises entering it. Further assume that the parameters can take on only discrete values: either this is reasonable physically (as for many failure detection formulations), or representative discrete values are chosen throughout the continuous range of possible values. Then a Kalman filter is designed for each choice of parameter value, resulting in a bank of K separate "elemental" filters. Based upon the observed characteristics of the residuals in these K filters, the conditional probabilities of each discrete parameter value being "correct", given the measurement history to that time, are evaluated iteratively. In MMAC configurations, a separate set of controller gains is associated with each elemental filter in the bank. The control value of each elemental controller is weighted by its corresponding probability, and the adaptive control is produced as the probability weighted average of the elemental controller outputs. As one alternative (using maximum a posteriori, or MAP, rather than minimum mean square error, or MMSE, criteria for optimality), the control value from the single elemental controller associated with the highest conditional probability can be selected as the output of the adaptive controller.

Previous efforts investigated the application of a MMAC algorithm to a short takeoff and landing (STOL) F-15 [8, 9]. The system was modeled with four elemental controllers designed for a healthy aircraft, failed pitch rate sensor, failed stabilator, or failed "pseudosurface"—a combination of canards, ailerons, and trailing edge flaps. Conclusions from this study indicated that the elemental filters must be carefully tuned to avoid masking of "good" versus "bad" models. This observation is not compatible with loop transmission recovery (LTR) tuning techniques. Other research efforts demonstrated the effectiveness of the MMAC algorithm using seven elemental controllers designed for a healthy aircraft, one of three actuator failures, or one of three sensor failures [10, 11]. The study included effects of single and double failures, and partial failures as well as hard failures. It also demonstrated the effectiveness of alternate techniques

to resolve ambiguities using modified computational techniques and scalar residual monitoring.

This research extends beyond this earlier work in three important aspects. First, both longitudinal and lateral-directional channels are considered, rather than just considering longitudinal dynamics. Second, out of necessity with the greater number of sensors and actuators and cross coupling of channels, disambiguation and failure declarations received attention through purposeful pilot commands and automatic control dithering as well as scalar residual monitoring. Finally, it was decided not to design "just a paper-study MMAC" composed of simplistic elemental controllers, but to take an existing full-scale flight control system (that of the VISTA F-16 (variable in-flight stability test aircraft)) and to embed an MMAE algorithm into its front end to provide failure detection capacity. This is intended to provide a more realistic evaluation of performance in an actual implementation, at the expense of not providing individual elemental controllers designed explicitly for each failure status condition. As a result, the emphasis is concentrated on *failure identification* performance versus *reconfigurable control* characteristics. If the failure mode is correctly identified, the algorithm is considered to have operated successfully, even if (due to lack of subsequent controller reconfiguration) the resulting closed-loop system is unstable and this could be addressed via MMAC versus MMAE-based control.

II. MMAC AND MMAE-BASED CONTROL

Let a denote the vector of uncertain parameters in a given linear stochastic model for a dynamic system, in this case depicting the failure status of sensors and actuators of the aircraft. These parameters can affect the matrices defining the structure of the model or depicting the statistics of the noises entering it. In order to make simultaneous estimation of states and parameters tractable, it is assumed that a can take on only one of K discrete representative values. If we define the hypothesis conditional probability $p_k(t_i)$ as the probability that a assumes the value a_k (for $k = 1,2,\ldots,K$), conditioned on the observed measurement history to time t_i:

$$p_k(t_i) = \Pr[a = a_k \mid Z(t_i) = Z_i] \tag{1}$$

then it can be shown [1–4] that $p_k(t_i)$ can be evaluated recursively for all k via the iteration:

$$p_k(t_i) = \frac{f_{z(t_i)|a,Z(t_{i-1})}(z_i \mid a_k, Z_{i-1}) p_k(t_{i-1})}{\sum_{j=1}^{K} f_{z(t_i)|a,Z(t_{i-1})}(z_i \mid a_j, Z_{i-1}) p_j(t_{i-1})} \tag{2}$$

in terms of the previous values of $p_1(t_{i-1}),\ldots,p_K(t_{i-1})$, and conditional probability densities for the current measurement $z(t_i)$ to be defined explicitly in (12).

Notionally, the measurement history random vector $Z(t_i)$ is made up of partitions $z(t_1),\ldots,z(t_i)$ that are the measurements available at the sample times $t_1,\ldots,t_i$; similarly, the realization Z_i of the measurement history vector has partitions $z_1,\ldots,z_i$. Furthermore, the Bayesian MMAC output is the probability weighted average [5–7]:

$$u_{\mathrm{MMAC}}(t_i) = \sum_{k=1}^{K} u_k[\hat{x}_k(t_i^*), t_i] p_k(t_i). \tag{3}$$

Here $u_k[x(t_i),t_i]$ is a deterministic optimal full-state feedback control law based on the assumption that the parameter vector equals a_k, and $x_k(t_i^+)$ is the state estimate generated by a Kalman filter similarly based on the assumption that $a = a_k$. If the parameter were in fact equal to a_k, then certainty equivalence [5] would allow the LQG (linear system, quadratic cost, Gaussian noise) optimal stochastic control to be generated as one of the $u_k[x_k(t_i^+),t_i)]$ terms in the summation of (3).

More explicitly, let the model corresponding to a_k be described by an "equivalent" discrete-time model [4, 5, 11] for a continuous-time system with sampled data measurements:

$$x_k(t_{i+1}) = \Phi_k(t_{i+1},t_i)x_k(t_i) + B_k(t_i)u(t_i) + G_k(t_i)w_k(t_i) \tag{4}$$

$$z(t_i) = H_k(t_i)x_k(t_i) + v_k(t_i) \tag{5}$$

where x_k is the state, u is a control input, w_k is discrete-time zero-mean white Gaussian dynamics noise of covariance $Q_k(t_i)$ at each t_i, z is the measurement vector, and v_k is discrete-time zero-mean white Gaussian measurement noise of covariance $R_k(t_i)$ at t_i, assumed independent of w_k; the initial state $x(t_0)$ is modeled as Gaussian, with mean x_{k0} and covariance P_{k0} and is assumed independent of w_k and v_k. Based on this model, the Kalman filter [11] is specified by the measurement update:

$$A_k(t_i) = H_k(t_i)P_k(t_i^-)H_k^T(t_i) + R_k(t_i) \tag{6}$$

$$K_k(t_i) = P_k(t_i^-)H_k^T(t_i)A_k^{-1}(t_i) \tag{7}$$

$$\hat{x}(t_i^+) = \hat{x}_k(t_i^-) + K_k(t_i)[z_i - H_k(t_i)\hat{x}_k(t_i^-)] \tag{8}$$

$$P_k(t_i^+) = P_k(t_i^-) - K_k(t_i)H_k(t_i)P_k(t_i^-) \tag{9}$$

and the propagation relation:

$$\hat{x}_k(t_{i+1}^-) = \Phi_k(t_{i+1},t_i)\hat{x}_k(t_i^+) + B_k(t_i)u(t_i) \tag{10}$$

$$P_k(t_{i+1}^-) = \Phi_k(t_{i+1},t_i)P_k(t_i^+)\Phi_k^T(t_{i+1},t_i) + G_k(t_i)Q_k(t_i)G_k^T(t_i). \tag{11}$$

The MMAE algorithm is composed of a bank of K separate Kalman filters, each based on a particular value $a_1,\ldots,a_K$ of the parameter vector, as depicted in Fig. 1. Instead of generating a control vector u_k, the MMAE generates a probabilistically weighted

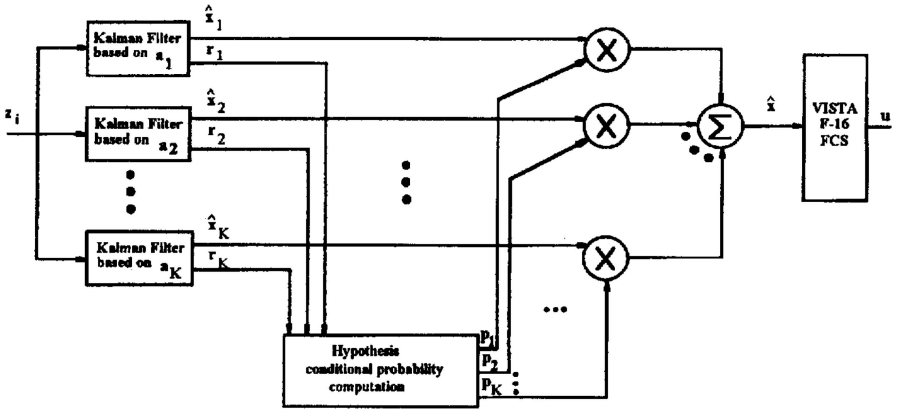

Fig. 1. MMAE algorithm.

state estimate vector, $x_{\mathrm{MMAE}}(t_i)$. In an MMAE-based controller, this state estimate is used by the flight control system to generate the control vector for the aircraft. When the measurement z_i becomes available at time t_i, the residual vector r_k is generated in each of the K filters according to the bracketed term in (8), and used to compute $p_1(t_i),\ldots,p_K(t_i)$ via (2). Each numerator density function in (2) is given by the Gaussian form:

$$f_{z(t_i)|a,Z(t_{i-1})}(z_i \mid a_k, Z_{i-1}) = \frac{1}{(2\pi)^{(m/2)}|A_k(t_i)|^{1/2}} \exp[\cdot] \tag{12}$$

$$[\cdot] = [-\tfrac{1}{2} r_k^T(t_i) A_k^{-1}(t_i) r_k(t_i)] \tag{13}$$

where m is the measurement dimension and $A_k(t_i)$ is calculated in the kth Kalman filter as in (6). The denominator in (2) is simply the sum of all the computed numerator terms and thus is the scale factor required to ensure that the $p_k(t_i)$s sum to one.

One expects the residuals of the Kalman filter based upon the "best" model to have a mean squared value most in consonance with its own computed $A_k(t_i)$ consistently over time, while "mismatched" filters will have larger residuals than anticipated through $A_k(t_i)$. Therefore, (2), (3), and (6)–(12) will most heavily weight the filter based upon the most correct assumed parameter value. However, the performance of the algorithm depends on there being significant differences in the characteristics of residuals in "correct" versus mismatched filters. Each filter should be tuned for best performance when the "true" values of the uncertain parameters are identical to its assumed value for these parameters. One should specifically avoid the "conservative" philosophy of adding considerable dynamics pseudonoise, often used to guard against divergence, since this tends to mask the differences between good and bad models. If, as a result of such tuning, one of the filters should diverge (which is clearly indicated by large magnitude residuals or term computed in (13)), it can be restarted with the current state estimate from the MMAE as computed from the nondivergent filters.

III. MODELING EFFORT

Aerodynamic Model: A six-degree-of-freedom nonlinear aerodynamic model provided data to generate a linearized perturbation model utilized in this study. The data base resides within the Flight Dynamics Laboratory at Wright-Patterson AFB, OH. The linearized model includes increments for pitch attitude, pitch rate, angle of attack, velocity, roll angle, sideslip angle, roll rate, and yaw rate. Normal and lateral accelerations are computed. Control effects are given by left and right stabilators, left and right flaperons, rudder and leading edge flaps. The model is developed with constant thrust.

Flight Control System: The flight control system (FCS) model is a Fortran representation of the VISTA F-16 FCS [13, 14]. The model realistically depicts the true system by including longitudinal, lateral, and directional channels. Each channel provides command force gradients, command limiting, signal magnitude and rate limiting (accomplished in controller software), gain scheduling, biases, filtering characteristics, and true surface position and rate limiting. Sensor measurements are corrected for position error where applicable. The flight control system requires seven sensor inputs for proper performance: velocity, angle of attack, pitch rate, normal acceleration, roll rate, yaw rate, and lateral acceleration.

The development of a detailed model allows for a realistic evaluation of the MMAE algorithm. The FCS and linearized aerodynamic models were validated separately and as a system using a six-degrees-of-freedom nonlinear simulation. Results indicated excellent correlation, provided that the constraints of the linear aerodynamic perturbation model were not violdated. Given the short convergence times typical for a fault detection and isolation algorithm, this is not a restrictive constraint.

IV. ALGORITHM IMPLEMENTATION

Hypothesized Failures: For the single failure scenarios, the parameter space, denoted by the vector quantity a, was discretized into twelve hypothesized hard failures: left stabilator, right stabilator, left flaperon, right flaperon, rudder, velocity sensor, angle of attack sensor, pitch rate sensor, normal acceleration sensor, roll rate sensor, yaw rate sensor, and the lateral acceleration sensor. Additionally, the no-failure aircraft condition was included to provide an initial system configuration prior to failure transition. Total or "hard" actuator failures are modeled by zeroing out the appropriate columns of the control input matrix B and hard sensor failures are modeled by zeroing out the corresponding rows of the measurement matrix H.

Only hard failures and no-failure conditions are used for the design of elemental filters for the MMAE, but "soft" failures are also simulated during testing (with the desire that the MMAE would form a probability-weighted average of estimates based on the associated hard failure and no-failure assumptions). Partial or soft actuator failures are modeled by multiplying the appropriate column of the control input matrix B of (4) by a factor of effectiveness. Soft sensor failures are modeled by increasing the variance in the R matrix.

Dual failure scenarios are similar to the single failure scenarios until the first failure is identified. Upon identification (the criterion used was that the probability in the elemental filter must be above 0.9 for 10 sample periods), a new bank of filters is loaded. Each filter in the new bank assumes the previously identified single hard failure and any other hypothesized second hard failure, including no other failure. See Fig. 2. A no-failure filter is also included to allow the system to back out of the decision tree structure in the event that a misidentification occurs.

Bayesian Form: The final probability-weighted average of the state estimates, computed as shown in Fig. 1, is produced by a Bayesian form of the MMAE algorithm. Practical implementation reqires a lower bound when computing the probabilities according to (2). The addition of a lower bound precludes the algorithm from assigning any single $p_k(t_i)$ a value of zero, which would prevent it from being considered in future probability computations. From the iterative nature of (2), if $p_k(t_{i-1})$ were assigned a value of zero for one of the filters, subsequent probability calculations for that filter would also assign a probability of zero (i.e., $p_k(t_i) = 0$). The addition of a lower bound provides another favorable characteristic. The number of iterations required to increase a very small, but non-zero, p_k is directly proportional to the magnitude of the p_k. By providing a lower bound, we allow p_k values, previously not important to the combined state estimate, to increase in a timely manner if the system failure status changes. This lower bounding is a simpler means of allowing for time-varying hypothesis probabilities than explicitly modeling hypothesis transitions, as by Markov process models [4].

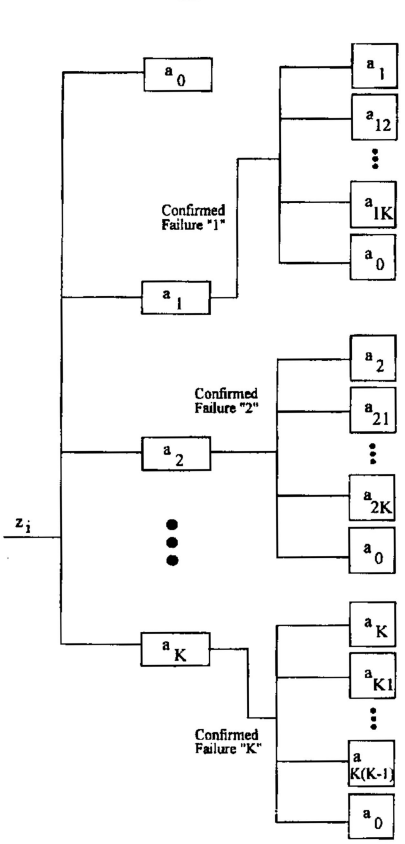

Fig. 2. Hierarchical structure for multiple failures.

Beta Dominance: As discussed earlier in Section II, the probabilities, $p_k(t_i)$, are calculated according to (2). Earlier efforts [2, 4, 6, 10] noted that the leading coefficient preceding the exponential term in (12) does not provide any useful information in the identification of the failure. As discussed in Section II, the likelihood quotient,

$$L_k(t_i) = r_k^T(t_i) A_k^{-1} r_k(t_i) \tag{14}$$

compares the residual with the internally computed residual covariance of the hypothesized filter. Filters with residuals that have mean square values most in consonance with their internally computed covariance are assigned the higher probabilities by the MMAE algorithm. However, if the likelihood quotients were nearly identical in magnitude for all k, the probability computations would be based upon the magnitude of the determinants of the $A_k(t_i)$ matrices, resulting in an incorrect assignment of the probabilities. This effect is known as "beta dominance". Because sensor failures, as simulated by zeroing out a row of H, yield smaller $A_k(t_i)$ values, "beta dominance" produces a tendency to generate false alarms about sensor failures.

Previous efforts removed the term preceding the exponential in (12). Since the denominator of (2) contains the summation of all numerator terms, excluding the terms preceding the exponentials in the calculation of the probabilities does not alter the fact that the computed probabilities sum to one.

Scalar Residual Monitoring: Incorrect or ambiguous failure identification may be resolved through the use of scalar residual monitoring. Eqs. (2), (12), and (13) demonstrate the relationship of the probability calculations, the probability density function, and the likelihood quotient. These three equations demonstrate the dependency of the probability calculations on the magnitude of the likelihood quotient, (13). The likelihood quotient is merely the sum of scalar terms relating the product of any two scalar components of the residual vector and the internally computed covariance for those two components. If a sensor failure occurs, the single scalar term associated solely with that sensor should have a residual value whose magnitude is much larger than the associated variance in all of the elemental filters except for the filter designed to "look" for that sensor failure. Scalar residual monitoring can be used as an additional vote when attempting to reduce or eliminate failure identification ambiguities.

Purposeful Commands: Failure detection and isolation using the MMAE algorithm requires a stimulus to disturb the system from a quiescent state. The performance of the MMAE algorithm depends upon the magnitude of the residuals within incorrect filters having large residual values. Residuals are the difference between measurements and filter predictions of those measurements. Incorrect filters will provide poor estimates relative to the filter based on the "true" system status. Small deviations from a quiescent state will be virtually indistinguishable from system noise, providing poor failure detection and identification. Having justified the need for stimuli to "shake up" the system, rationale was developed to select stimuli, control deflections, and improve performance. Previous efforts selected a pitch down maneuver to aid in the identification process for the longitudinal axis of an aircraft, with generally favorable results [8–11]. However, fundamental differences exist between earlier research and this effort. Earlier efforts concentrated on applying the MMAC algorithm, evaluating its performance, and designing algorithms to maintain stability and control in the longitudinal axis. A longitudinal pitch down maneuver was sufficient to provide enough system excitation for good performance. A three-axis sophisticated control system requires excitation in multiple axes to provide adequate residual growth in filters whose hypothesis does not reflect the true system failure status. The purposeful commands used in this effort were longitudinal stick pulls, lateral stick pulses, and varying amounts of rudder application. Ordinary aircraft maneuvering would probably be more than sufficient to provide adequate excitation and good performance; straight-and-level flight would be more challenging (though less flight critical) for a failure detection system.

Autonomous Dithering: Autonomous dithering enhances failure detection and identification by providing sufficient excitation in benign nonmaneuvering flight conditions or as a pilot-selectable option. A number of dither signals were evaluated, including square waves, triangle waves, combinations of these forms, and sine waves. Pulse trains using a square wave form produced good performance with one drawback, failures are not detected until the application of the pulse. Additionally, pilots may find the application of a dither signal of sufficient strength to provide good failure detection and isolation objectionable, unless they were able to turn such dithering on or off themselves. Sufficient data was not available to relate pilot comments and normal and lateral accelerations in this application, so dithers were developed to be as subliminal as possible while yielding desired identifiability of failures.

V. PERFORMANCE

The application of the MMAE algorithms to the VISTA F-16 aircraft in a low dynamic pressure case provided an interesting test for this technique. The flight condition, 0.4 Mach at an altitude of 20000 ft, demonstrated algorithm performance in a low dynamic pressure scenario. Earlier efforts studied the VISTA F-16 at a higher dynamic pressure and emphasized different failure scenarios and characteristics [15]. The goal was to evaluate the ability of the MMAE algorithm to detect and isolate failures within the flight control system and not to evaluate the ability of the controller to maintain control of the vehicle after the identification of the failure. An added benefit of using the VISTA F-16 flight control system was the absence of any single-failure induced loss of control. Dual failure elemental filters were designed for every two-failure actuator and sensor combination. The figures presented in this section are single data runs as opposed to Monte Carlo runs averaged over a number of runs, in order to exhibit real-time signal characteristics (Monte Carlo runs *were* used to corroborate performance attributes over multiple experimental trials).

For single-failure plots, the following elemental filter abbreviations are used: FF, fully functional; A1, left stabilator; A2, right stabilator; A3, left flaperon; A4, right flaperon; A5, rudder; S1, velocity sensor; S2, angle of attack sensor; S3, pitch rate sensor; S4, normal acceleration sensor; S5, roll rate sensor; S6, yaw rate sensor; and S7, lateral acceleration sensor. Dual-failure plot nomenclature is described subsequently.

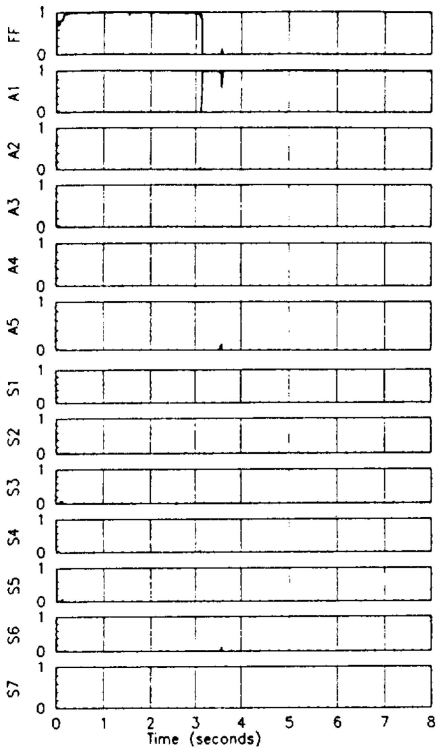

Fig. 3. Probabilities for left stabilator failure using purposeful roll command.

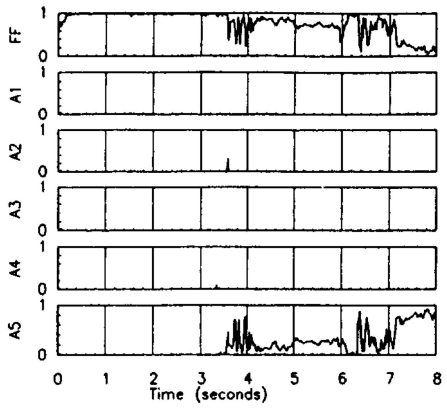

Fig. 4. Probabilities for 50% reduced effectiveness rudder failure using subliminal dither pulse.

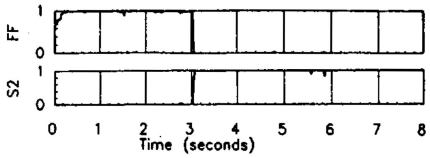

Fig. 5. Probabilities for angle of attack sensor failure using subliminal dither pulse.

Single Hard Actuator Failures: Fig. 3 demonstrates a left stabilator failure induced at 3.0 s. Prior to the insertion of the failure in the truth model, essentially all of the probability is properly contained within the FF elemental filter. At 3.0 s, the failure is inserted and the left stabilator elemental filter (A1) collects the probability. The time delay between failure insertion and failure detection is approximately 0.2 s. In this scenario, a purposeful roll command is used to excite the roll channel (a roll command aids the algorithm in not only identifying the stabilator failure, but which stabilator has failed). Small dither signals of magnitude less than 0.2 *g*s are used to excite the pitch and yaw channels to enhance failure detection. The roll step command is inserted into the flight control system at 3.0 s and has a magnitude of 13.5 lateral stick lbs adn is held for 1.55 s. The side stick application times were selected to coincide with the failure insertion at 3.0 s (to excite the system and generate residual growth) and to demonstrate the detection capability with the command in and with the command released. All of the single hard actuator failures resulted in good performance (quick detection, proper identification, failure declaration remains locked).

Single Soft Actuator Failures: Two types of excitation signals are utilized throughout the research effort: purposeful commands and autonomous dither signals. Autonomous dither signals provide failure detection capability in benign (i.e., nonmaneuvering) flight conditions. Fig. 4 demonstrates a 50% reduced effectiveness rudder failure using a subliminal dither pulse. In this scenario, three sinusoidal dither signals are applied to the pitch, roll, and yaw channels, of amplitudes 12.5 lbs, 11.0 lbs, and 30.0 lbs, respectively. Prior to the failure insertion at 3.0 s, the probability is contained within the FF filter. After the failure insertion at 3.0 s, the probability is shared between the FF elemental filter and the complete rudder failure (A5) elemental filter. This correlates very well with the anticipation that the probability should be equally shared between the two filters. In this straight-and-level flight condition, the subliminal dither pulse is sufficient to identify a partial rudder failure and properly blend the outputs of the two filters. This scenario highlights the strengths of the Bayesian approach. For scenarios with a reduction in actuator effectiveness below 50%, the failure is typically not detected. For scenarios with a reduction in actuator effectiveness greater than 75%, the algorithm behaves the same as a hard actuator failure. These results are reported for a straight-and-level flight condition; maneuvering flight may produce better detection capability since more of the actuator effectiveness is demanded for good flight performance.

Single Hard Sensor Failures: Fig. 5 demonstrates an angle of attack sensor failure using a subliminal dither signal. The failure is induced at 3.0 s, and it is identified by the angle of attack (S2) elemental filter at 3.1 s: within 0.1 s. All of the single hard sensor failures are identified quickly and accurately.

Single Soft Sensor Failures: Single soft sensor failures are modeled by increasing the variances in the R matrix. Scenarios with variances increased

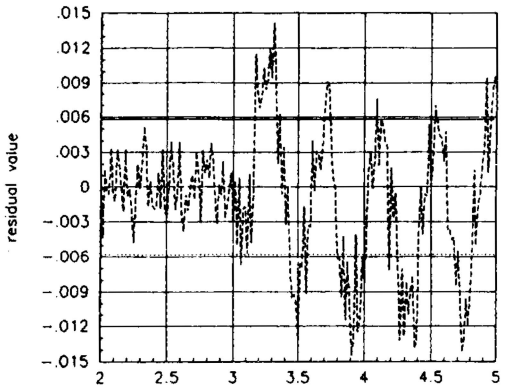

Fig. 6. Single scalar velocity residual for FF filter given left stabilator failure.

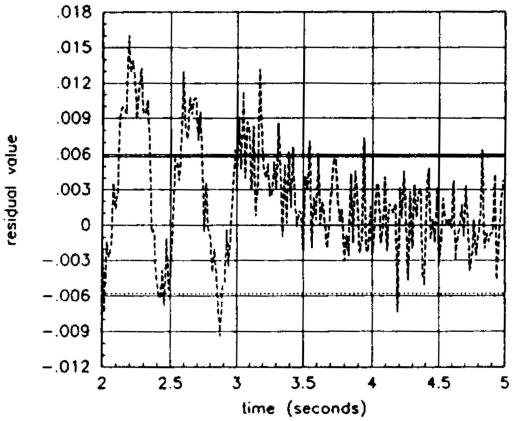

Fig. 7. Single scalar velocity residual for left stabilator filter given left stabilator failure.

by a factor of 2 to 10 have been evaluated, and the performance is generally poor. In some cases, the failure is detected after 3 or more seconds. When inspecting the scenarios, a spiking phenomena is observed in the probability plots. The spiking occurs in the correct filter in virtually every scenario, but the total probability is not sufficient for detection by the algorithm. More work must be done in this area of failure detection, including allowing sensor bias shifts, which should be easier to detect than increased noise variance.

Residual Monitoring for the Single Failure Scenarios: As explained in Section IV, residual monitoring provides additional voting when confronted with ambiguities. Research efforts demonstrated good residual monitoring performance when used with a sinusoidal dither signal. Figs. 6 and 7 present the single scalar velocity residuals for the FF and failed-left-stabilator elemental filters, respectively. In this scenario, a left stabilator failure is inserted into the simulation at 3.0 s. From Fig. 6, it can be seen that the FF elemental filter has the correct hypothesis until 3.0 s. The velocity residual appears white and is within the $+/-2.0\sigma$ bounds as computed by the elemental filter. After 3.0 s, the velocity residual for this elemental filter violates the $+/-2.0\sigma$ bounds and has a frequency that matches the sinusoidal dither frequency. In contrast, Fig. 7 presents the velocity residual for the left stabilator elemental filter. Prior to 3.0 s, the velocity residual violates the $+/-2.0\sigma$ bounds and has a frequency that matches the sinusoidal dither frequency. At 3.0 s, the characteristics of the velocity residual noticeably change. Almost immediately, the residual appears white. By 3.3 s the residual lies within the $+/-2.0\sigma$ bounds and remains within the bounds for the duration of the simulation. For the scenarios addressed in this effort, the velocity, normal acceleration, lateral acceleration residuals provide the most meaningful indications of failure transitions.

Dual Hard Failures: Table I presents the nomenclature for interpreting the dual failure plots.

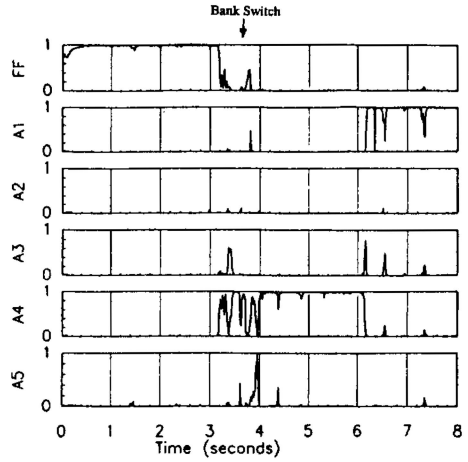

Fig. 8. Probabilities for right flaperon failure followed by left stabilator failure using sinusoidal dither signal.

The filter abbreviations are identical in meaning to the single failure elemental filter prior to the identification of the first failure. After the identification of the first failure, a new bank of filters is brought on line, as shown in Fig. 2. At this point the designations change, and each elemental filter symbol represents the first failure and a second hypothesized failure as indicated by the symbol. The FF designation remains the same throughout all of the scenarios (single and double failures). As an example, after a left stabilator failure has been identified and a new bank of filters if brought on line, the elemental filter designations change meaning. In this example, the left stabilator (A1) elemental filter remains the same. The designator A2 is interpreted as the dual stabilator elemental failure filter, the A3 designator is interpreted as the left stabilator and left flaperon failure filter, and so on.

Fig. 8 represents a right flaperon failure followed by a left stabilator failure. The right flaperon failure is inserted into the simulation at 3.0 s followed by the left stabilator failure at 6.0 s. The bank switching

TABLE I
Multiple Failure Nomenclature After Bank Switch

Designator	Description
FF	*Fully Functional Filter*
A1	*First Failure + Left Stabilator Failure*
A2	*First Failure + Right Stabilator Failure*
A3	*First Failure + Left Flaperon Failure*
A4	*First Failure + Right Flaperon Failure*
A5	*First Failure + Rudder Failure*
S1	*First Failure + Velocity Sensor Failure*
S2	*First Failure + Angle of Attack Sensor Failure*
S3	*First Failure + Pitch Rate Sensor Failure*
S4	*First Failure + Normal Acceleration Sensor Failure*
S5	*First Failure + Roll Rate Sensor Failure*
S6	*First Failure + Yaw Rate Sensor Failure*
S7	*First Failure + Lateral Acceleration Sensor Failure*

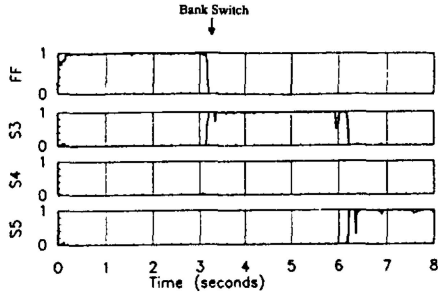

Fig. 9. Probabilities for pitch rate sensor failure followed by roll rate sensor failure using sinusoidal dither signal.

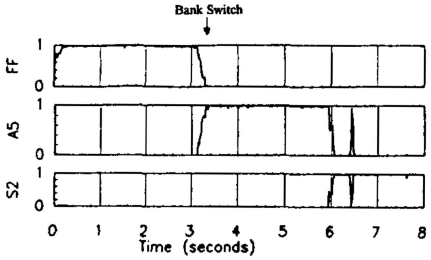

Fig. 10. Probabilities for rudder failure followed by angle of attack sensor failure using sinusoidal dither pulse.

criteria is satisfied at 3.6 s, after which the elemental filter designators assume their new definitions (as just discussed). Prior to 3.0 s, the FF elemental filter contains the probability. The failure is inserted into the simulation at 3.0 s and detected by the right flaperon (A4) elemental failure at 3.2 s. Between 3.2 and 3.6 s, the majority of the probability is shared between the left flaperon (A3), the right flaperon (A4), and the rudder (A5) elemental filters. At 3.6 s, the bank is switched. The new bank assumes a right flaperon failure and any other second hypothesized failure. Between 3.6 and 4.0 s, the probability is shared between the right flaperon (A4) elemental filter, the FF elemental filter, the right flaperon and rudder (A5) elemental filter, and the right flaperon and angle of attack sensor (S2) failure filter. Between 4.0 and 6.0 s, the probability is contained in the right flaperon (A4) failure filter. The second failure is inserted at 6.0 s into the truth model. Except for a small spike in the right flaperon and right stabilator (A3) filter, the probability collects in the right flaperon and left stabilator (A1) failure filter within 0.2 s. In this scenario (straight and level flight), the sinusoidal dither signal applies the excitation signal. By increasing the signal strength, the failure identification performance can be enhanced; however, the dither signal can no longer be labeled subliminal. Maneuvering flight usually produces improved performance over benign flight conditions.

Fig. 9 presents a dual sensor failure. A pitch rate sensor failure is induced at 3.0 s. The failure is identified at 3.2 s. At this point in the simulation, the bank switching criterion is met and a new bank of filters is brought on line. The new bank of filters assumes a pitch rate failure and any other second hypothesized failure. The roll rate sensor failure is induced at 6.0 s. With the exception of a temporary dropout, the elemental filter assuming both a pitch rate sensor failure and a roll rate sensor failure contains all of the probability after 6.2 s. Dual sensor failures generally provide good results.

Fig. 10 demonstrates a mixed actuator and sensor failure. In this scenario, a rudder failure is induced at 3.0 s. The rudder failure is identified at 3.3 s and a bank switch occurs before 3.4 s. The new bank of filters assumes a rudder failure and any other second

TABLE II
Multiple Hard Failure Summary Matrix

	LS	RF	RD	VS	AOA	PR	NZ	RR	YR	NY
LS		Good	Good	Good	Good	Good	Good	Good	Good	Good
RS	Fair	Good	Good	Poor	Good	Good	Good	Good	Good	Good
LF	Poor	Good	Poor	Fair	Good	Good	Good	Good	Good	Good
RF	Poor		Poor	Fair	Good	Good	Good	Good	Good	Good
RD	ND	Good		Fair	Good	Good	Good	Good	Good	Good
VS	Good	Good	Good		Good	Good	Good	Good	Good	Good
AOA	ND	Good	Good	Good		Good	Good	Good	Good	Good
PR	ND	Good	Poor	Fair	Good		Good	Good	Good	Good
NZ	Good	Good	Good	Good	Good	Good		Good	Good	Good
RR	ND	ND	Poor	Fair	Good	Good	Good		Good	Good
YR	ND	ND	ND	Poor	Good	Good	Good	Good		Poor
NY	Fair	Good	Good	Fair	Good	Good	Good	Good	Good	

hypothesized failure. At 6.0 s, an angle of attack sensor failure is induced. With the exception of a single spike, all of the probability is contained within the rudder-failure- and angle-of-attack-sensor-failure elemental filter by 6.1 s.

Multiple Hard Failure Results: Table II summarizes the multiple hard failure results. The column header lists the left stabilator (LS), the right flaperon (RF), the rudder (RD), the velocity sensor (VS), the angle of attack sensor (AOA), the pitch rate sensor (PR), the normal acceleration sensor (NZ), the roll rate sensor (RR), the yaw rate sensor (YR), and the lateral acceleration sensor (NY). The right stabilator and left flaperon are not included since no additional insight is gained by evaluating these failures. The first column lists all of the second failures. Any dual failure scenario can be evaluated by finding the intersection of the appropriate column (first failure) and row (second failure). The terminology within the table is in some sense subjective and is defined as follows: ND (no detection—the signal was not detected in the appropriate elemental filter); Poor (some probability spiking occurred in the appropriate failed filter); Fair (probability lock—the probability in the appropriate filter was at 0.988 (max) for some appreciable period of time); Good (probability lock and hold—the probability in the appropriate filter was at 0.988 for some appreciable period of time and was consistently 0.988 through the end of the simulation run at 8.0 s). The ND and Poor results are shaded to accentuate the problem areas. Particularly evident from this shading is that there is much less difficulty if the first failure is a sensor rather than an actuator. In many cases, this can be attributed to the loss of actuator effectiveness when employing a dither excitation signal. The four outlined cells (LSRS, VSLS, RFRD, RFPR) are examples of initial poor performance improved by changing the dither signal strength and/or modifying the dither signal form. For these four cases, the preliminary performance with the original dither signal was either ND or Poor. By modifying the dither signal once the first failure is detected, the performance is significantly enhanced. In many cases, the dither signal can still be considered subliminal. The increased dither signal strength compensates for the loss of actuator authority, resulting in a match of the original unfailed aircraft responses to the dither. Due to time constraints, only four of the shaded areas were retested with an improved dither signal; however, the four cases were chosen to cover a representative set of failures. A dual actuator case demonstrates the most difficult failure detection point for the MMAE algorithm. Physically it would be quite easy to detect dual failed actuators through purposeful commands. Results indicate that the remainder of the shaded regions could be improved by increasing the dither signal strength. The results are good for all dual failure combinations in which the first failure is a sensor, except for three velocity sensor cases adn one lateral acceleration case. The velocity sensor failures which involve a stabilator can be attributed to a reduction in the excitation dither signal magnitude brought about by the loss of actuator authority. This is easily corrected by increasing the dither signal strength to prefailure excitation levels, as demonstrated in the Good rating in the VSLS block of Table II. The other two exceptions involve a second failure of the yaw rate sensor. A known bias exists in the yaw rate sensor residuals. The solution to this problem would be to reinitialize or restart the filter periodically to maintain a very small bias below a predetermined threshold. The threshold would be established empirically based upon the failure identification performance. Overall, the results for the multiple hard failure scenarios are good.

Dual Failure Residual Monitoring: By using sinusoidal dither signals, dual failure residual monitoring can provide additional voting when ambiguities arise. Results demonstrate [14] clear failure indications for the first and second failures. Typically the most interesting residuals are provided by the FF filter, the first failure elemental filter, and the second dual failure elemental filter. Each of these three filters provide a vote in the determination of the true system status. For the dual failure scenarios, the most useful residuals are the velocity, pitch rate, normal acceleration, and the lateral acceleration. Other residuals provide less obvious indications of the failure status.

Simultaneous Failures: Failure scenarios included dual failures separated by 3.0, 0.5, and finally 0.1 s (simultaneous). For the first scenario tested, namely a right flaperon failure followed by a left stabilator failure, the dual failures are identified and rated as Good by the subjective rating system if the separation is 3.0 or 0.5 s. For the simultaneous failure (0.1 s separation), the left stabilator failure is detected first and the bank switch occurs. Some spiking occurs in

each of the flaperons but the second failure is not detected. This result is consistent with Table II. For the first two relatively widely separated failures, the results agree with the RFLS block in the table (Good). The last simultaneous case resulted in a LSRF scenario and the results again agree with Table II (Poor). Increasing the dither signal strength would compensate for the degradation in actuator effectiveness and significantly improve the rating. The other three failure scenarios that were investigated, a normal acceleration sensor failure and a pitch rate sensor failure, a roll rate sensor failure followed by a lateral acceleration sensor failure, and a roll rate sensor failure followed by a rudder failure are rated Good for the widely separated (3.0, 0.5 s) and simultaneous (0.1 s) failures.

Overall Performance Trends. Single Failures: Single hard actuator and sensor failures are easily identified. Results indicate only one "forced" misidentification in over 2000 runs. The application of a rudder kick and hold when attempting to identify a rudder failure resulted in a declaration of a yaw rate failure. The result was not surprising since a rudder failure and a yaw rate failure appear similar in the state variables. When the rudder failed, the lack of a rudder input produced very little yaw rate, resulting in a situation in which the algorithm could not distinguish between a rudder failure or a yaw rate sensor failure. By applying differential stabilator and flaperon inputs, a yaw rate could be produced to resolve the ambiguity. Single soft actuator failures resulted in good performance when the actuator effectiveness is reduced below 50%. The failure results for a reduction below 75% effectiveness match the hard failure results. For scenarios with actuator effectiveness between 50 and 75%, a sharing of the probability is observed between the two filters with the most correct hypothesis as anticipated. Single soft sensor failures, modeled by increasing a measurement noise variance by a factor of 2 to 10, demonstrate poor results. Some spiking phenomena is present in the probability for the elemental filter with the correct hypothesis, but the overall probability levels are not sufficient to ensure good performance. A larger variance might produce better results. Results have demonstrated the algorithm is sensitive to bias shifts, and a soft sensor failure modeled by a large enough bias shift would be detected. Residual monitoring is very useful, when combined with a sinusoidal dither signal or purposeful command, to provide an additional vote when attempting to resolve ambiguities.

Overall Performance Trends. Multiple Failure: Multiple hard failure results indicate good performance when the first failure is a sensor. For an actuator first failure, the results are typically degraded. By rerunning a subset of the ND or Poor cases and increasing the dither signal strength, the results for the first actuator failure subset are significantly improved. Overall, all of the test runs support the claim that the MMAE algorithm can accurately identify dual hard failures. Residual monitoring is useful in dual failure scenarios when the system is excited using a sinusoidal dither signal. Residual monitoring, using a sinusoidal dither signal, provides clear indications of failures through the "whiteness", frequency, and magnitude (in relation to the filter computed 2σ bounds) of the residual. An implementable algorithm might remove any residual bias and count the zero crossings to determine the "whiteness" of the signal. Simultaneous failures provide results consistent with Table II.

While many dither wave forms were considered and tested, the sine wave dither form yielded the best performance. The signal is continuous, arguably subliminal, and demonstrates good performance. It is extremely useful in resolving ambiguities. The continuous nature of the wave (versus intermittent pulse trains) provides constant failure detection coverage and can easily be modulated in amplitude or frequency to provide the best detection signal for a given failure scenario. A continuously alternating frequency may provide even better algorithm performance at all flight conditions. Amplitude modulation accounts for varying dynamic pressures, atmospheric disturbances, or aircraft status (increased dither signal strengths being required if an actuator failure or partial failure has been declared). In this effort, wideband dither signals were not investigated, but sinusoidal dither forms did provide better overall performance than pulse train forms.

VI. SUMMARY

A MMAE algorithm has been evaluated for single and multiple actuator and sensor failures illustrating the performance of the algorithm when applied to a VISTA F-16 FCS using a linearized aerodynamic model. A modified Bayesian approach allows for a blending of state estimates and provides lower bounds on computed probabilities to enhance algorithm convergence properties. Compensation for "beta dominance" enhances algorithm performance by not allowing the term preceding the exponential in (12) to enter into the calculations. This term biases the calculation of the probabilities toward the elemental filter with an $A_k(t_i)$ matrix having the smallest determinant thus increasing sensor failure false alarms. Scalar residual monitoring effectively aids in resolving ambiguities.

For single failures, the algorithm demonstrates good convergence characteristics during purposeful commands and dither signals. Optimizing the dither to improve algorithm performance is effective. However, large dither signals cannot be considered subliminal and may be considered objectionable by a pilot; allowing him to turn the dither on and off may be more useful practically.

For multiple failures, the algorithm demonstrates good convergence for hard actuator and sensor failures. For actuator first failures, an increased dither signal strength is required to compensate for the loss of actuator effectiveness in exciting the system states. Actuator soft failures provide good results when the actuator effectiveness is reduced below 50%. Sensor soft failures (modeled by moderately increased measurement noise) are not well detected, but do not impact control performance much either. Simultaneous failures produce results consistent with identical failures separated in time by three seconds. Residual monitoring is useful in multiple failure scenarios when coupled with a sinusoidal dither signal.

REFERENCES

[1] Magill, D. T. (1965)
Optimal adaptive estimation of sampled stochastic processes.
IEEE Transactions on Automatic Control, **AC-10**, 5 (Oct. 1965), 434–439.

[2] Athans, M., and Chang, C. B. (1976)
Adaptive estimation and parameter identification using multiple model estimation algorithm.
Technical note 1976-28, ESD-TR-76-184, Lincoln Laboratory, Lexington, MA, June 1976.

[3] Maybeck, P. S. (1979)
Stochastic Models, Estimation and Control, Vol. 1.
New York: Academic Press, 1979.

[4] Maybeck, P. S. (1982)
Stochastic Models, Estimation and Control, Vol. 2.
New York: Academic Press, 1982.

[5] Maybeck, P. S. (1982)
Stochastic Models, Estimation and Control, Vol. 3.
New York: Academic Press, 1982.

[6] Athans, M., et al. (1977)
The stochastic control of the F-8C aircraft using a multiple model adaptive control (MMAC) method—Part 1: Equilibrium flight.
IEEE Transactions on Automatic Control, **AC-22**, 5 (Oct. 1977), 768–780.

[7] Greene, C. S., and Willsky, A. S. (1980)
An analysis of the multiple model adaptive control algorithm.
In *Proceedings of the IEEE Conference on Decision and Control*, Albuquerque, NM, Dec. 1980, 1142–1145.

[8] Pogoda, D. L. (1988)
Multiple model adaptive controller for the STOL F-15 with sensor/actuator failures.
M.S. thesis, Air Force Institute of Technology, Wright-Patterson AFB, OH, Dec. 1988.

[9] 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.

[10] Stevens, R. D. (1989)
Characterization of a reconfigurable multiple model adaptive controller using a STOL F-15 model.
M.S.E.E. thesis, Air Force Institute of Technology, Wright-Patterson AFB, OH, Dec. 1989.

[11] 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.

[12] Stratton, G. L. (1991)
Actuator and sensor failure detection using a multiple model adaptive technique for the VISTA/F-16.
M.S.E.E. thesis, Air Force Institute of Technology, Wright-Patterson AFB, OH, Dec. 1991.

[13] Menke, T. E., and Maybeck, P. S. (1992)
Multiple model adaptive estimation applied to the VISTA F-16 flight control system with actuator and sensor failures.
In *Proceedings of the National Aerospace and Electronics Conference*, Dayton, OH, May 1992.

[14] Menke, T. E. (1992)
Multiple model adaptive estimation applied to the VISTA F-16 with actuator and sensor failures.
M.S.A.E. thesis, Air Force Institute of Technology, Wright-Patterson AFB, OH, June 1992.

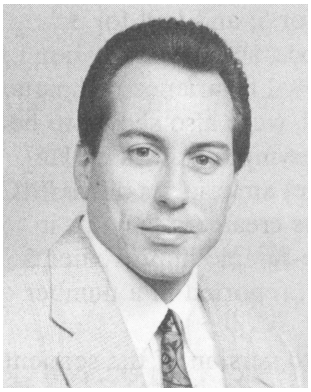

Timothy E. Menke was born in Minneapolis, MN on April 20, 1963. He received his B.S. degree in aerospace engineering from the University of Kansas, Lawrence, in May of 1985 and his M.S. degree in aerospace engineering in 1992 from the School of Engineering, Air Force Institute of Technology, Wright-Patterson AFB.

After attending Officer Training School and obtaining his commission in the United States Air Force, he was assigned to Wright-Patterson Air Force Base in the Flight Control Engineering Division.

Mr. Menke is a member of Sigma Gamma Tau.

Peter S. Maybeck (S'70—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 a Professor of Electrical Engineering. Current research interests concentrate on using optimal estimation techniques for guidance systems, tracking, adaptive systems and failure detection purposes.

Dr. Maybeck is 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). 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 1993, 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 Fellow of the IEEE and a member of the A.I.A.A., and he is the current IEEE Dayton Section Student Activities Chairman and a member of the IEEE Executive Committee of Dayton, and he previously served as Chairman of the local Automatic Control Group.