

Detection of Thrust Vector Control Nozzle Failures Using Multiple Model Adaptive Estimation

Peter D. Hanlon
Air Force Research Laboratory/MNAL
Eglin AFB, FL 32542

Johnny H. Evers
Air Force Research Laboratory/MNGN
Eglin AFB, FL 32542

I. Abstract

We investigated the feasibility of detecting the failure of a thrust vector control nozzle, like those proposed on advanced concepts of air-to-air missiles (Figure 1), using a modified Multiple Model Adaptive Estimation (MMAE) algorithm with a Neyman-Pearson based Hypothesis Testing Algorithm (NPHTA). These nozzles are currently being researched as a means to allow stable flight control at high angles of attack. It is likely that the missile would tumble if a nozzle failed during a high angle of attack maneuver, which could produce catastrophic results if the nozzle failure occurred as the missile is separating from the aircraft. Our goal was to determine if the NPHTA version of the MMAE could detect a nozzle failure during the separation flyout maneuver, prior to any high angle of attack maneuvers. The specific failure modes that we examined were single failures of a nozzle valve that was stuck either open or closed. We found that under ideal conditions the NPHTA MMAE could detect a nozzle valve failure within one clock cycle (0.02 seconds) during this initial separation flyout and thus could test the nozzle valves prior to high angle of attack maneuvers.




Figure 1. Air Superiority Missile System (ASMS).

II. Introduction

The modified MMAE, diagrammatically shown in Figure 2, uses a Kalman filter to produce a residual (the difference between the filter estimate of the measurements and the actual measurements) which the NPHTA tests, using several different models of the various failure hypotheses, to determine the hypothesis that most closely matches the true failure state. Previous research [1, 2, 3] has shown that the NPHTA requires only a *single* Kalman filter residual, provides a systematic methodology for the MMAE designer to obtain the desired failure detection performance, often detects failures faster than the standard MMAE, and is not subject to variations in failure detection performance due to failure injection timing. For this application we modeled a valve that is stuck closed as differences between the control input matrices (yielding hypotheses in terms of ΔB), and modeled a valve that is stuck open as differences in the control input vectors (yielding hypotheses in terms of $\Delta u(t_i)$). These known differences are used to compute the probability distribution of a Kalman filter residual for each of the hypotheses, which are then used, along with the desired probability of detection (P_D) and probability of false alarm (P_{FA}), to conduct a series of binary hypothesis tests (one hypothesis versus another single hypothesis) to determine

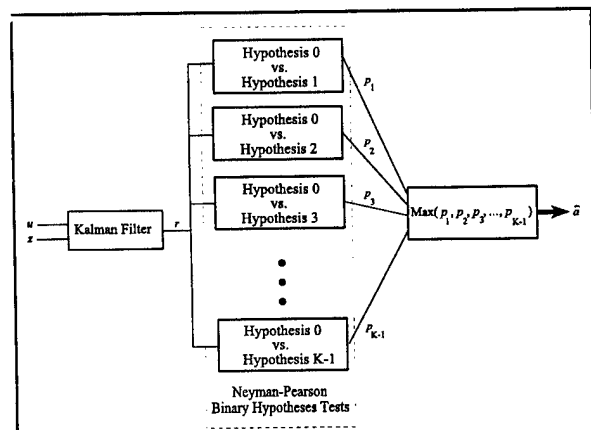


Figure 2. Multiple Model Adaptive Estimation Algorithm using the Neyman-Pearson Hypothesis Testing Algorithm

1998 1127 008

the most likely hypothesis. Initial results indicate that the NPHTA MMAE detects these nozzle failures within 0.02 seconds after the failure is injected.

III. Theory Development

3.1 Truth Model Development

We started our investigation by developing a model, subsequently called the truth model, that gives us the best estimates of the missile dynamics under the various hypothesized nozzle failures. We started with an existing linearized model of the form:

$$\mathbf{x}_T(t_i) = \Phi_T \mathbf{x}_T(t_{i-1}) + \mathbf{B}_T \mathbf{u}_T(t_{i-1}) \quad (1)$$

where \mathbf{x}_T is the truth model state vector

Φ_T is the truth model state transition matrix

\mathbf{B}_T is the truth model control input matrix

\mathbf{u}_T is the truth model system input vector.

Initially the truth model consisted of nine states (forward velocity, angle of attack, pitch rate, pitch angle, sideslip angle, roll rate, roll angle, yaw rate, and yaw angle) and three inputs (pseudo-elevator, pseudo-aileron, and pseudo-rudder). Given the initial rocket chamber pressure and the thrust vector nozzle throat size, we computed the force that could be applied to the missile by the thrust vector nozzles during the initial separation flyout, and then, given the distance from the missile center of gravity and the nozzle throat, we computed the moment induced by the thrust vector nozzles. The angular rates due to these moments were computed using the given roll, pitch, and yaw angular momenta. The model in Eq (1) was augmented by adding the induced angular rates of the thrust vector nozzles, thus giving us a nine input model where the thrust vector nozzle inputs are either a one (for when the nozzle is fired and induces the additional angular rate) or zero (for when the nozzle is off).

The model in Eq (1) was further modified to include wind gusts (using the Dryden wind gust model for moderate chop at medium altitude) and to model the noise corrupted measurements that would be available to the MMAE. We assumed that an inexpensive rate gyro is used to measure the missile angular rates. Therefore, we developed a linearized model of the form:

$$\begin{aligned} \mathbf{x}_T(t_i) &= \Phi_T \mathbf{x}_T(t_{i-1}) + \mathbf{B}_T \mathbf{u}_T(t_{i-1}) + \mathbf{G}_T \mathbf{w}_T(t_{i-1}) \\ \mathbf{z}_T(t_i) &= \mathbf{H}_T \mathbf{x}_T(t_i) + \mathbf{v}_T(t_i) \end{aligned} \quad (2)$$

where \mathbf{G}_T is the truth model noise input matrix

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

$$E\{\mathbf{w}_T(t_i) \mathbf{w}_T^T(t_j)\} = \begin{cases} \mathbf{Q}_T, & t_i = t_j \\ 0, & t_i \neq t_j \end{cases} \quad (3)$$

\mathbf{z}_T is the truth model measurement vector

\mathbf{H}_T is the truth model output matrix

\mathbf{v}_T is an additive white Gaussian measurement noise input that is used in the truth model, assumed to be independent of \mathbf{w}_T , and zero-mean with

$$E\{\mathbf{v}_T(t_i) \mathbf{v}_T^T(t_j)\} = \begin{cases} \mathbf{R}_T, & t_i = t_j \\ 0, & t_i \neq t_j \end{cases} \quad (4)$$

We chose to simulate the TVC nozzle failure modes using the system input vector, \mathbf{u} . The nominal system input is shown in the first column of Figure 3, which shows that all of the nozzles are pulsed for 0.06 seconds starting 0.1 seconds into the simulation. If there are no nozzle failures, the net result of this input is to cancel the effects of each individual nozzle on the missile dynamics. We simulated a TVC nozzle stuck in the closed position by setting the corresponding element of the system input vector to zero, which means the TVC nozzle is closed, throughout the simulation. A failure of this type for nozzle 1 is shown in the second column of Figure 3. We assume that the most likely time that a TVC valve would stick open is at the tail end of the nominally commanded pulse input. Therefore we simulated a TVC nozzle stuck in the open position by setting the corresponding element of the system input vector to one, which means the TVC nozzle is open, after the initial pulse is commanded. This type of failure for nozzle 1 is shown in the third column of Figure 3.




Figure 3. System Inputs for various failure models.

Classical MMAE algorithms have used a bank of Kalman filters, each based on a different failure model, to generate a series of residuals that are used by a single Bayesian hypothesis testing algorithm. Previous research [1, 2, 3] has shown that the Kalman filter residuals within a classical MMAE structure in which hypotheses differ in the \mathbf{B} , \mathbf{H} , and Φ matrices, are normally distributed random vectors with the same known (and precomputable) steady state covariances, but different means. The

residual covariances are the same because the Gaussian noise inputs, w_T and v_T , are not a function of B , H , and Φ matrices which means that they are assumed to be the same for each of the hypotheses. This development can easily be extended to include hypotheses that differ in the system input vector. Therefore, the Kalman filter residuals will have Gaussian distributions that differ only in the mean, and are dependent on the hypotheses of the Kalman filter internal models.

3.2 Kalman Filter Equations

We assume a time invariant Kalman filter model with stationary noises (and eventually a steady state constant-gain Kalman filter for implementation) associated with a particular hypothesized failure status, which will be denoted with the subscript k . Thus we have:

$$\begin{aligned} x_k(t) &= \Phi_k x_k(t_{i-1}) + B_k u_k(t_{i-1}) + G_k w_k(t_{i-1}) \\ z_k(t) &= H_k x_k(t) + v_k(t) \end{aligned} \quad (5)$$

where x_k is the Kalman filter model state vector

Φ_k is the Kalman filter model state transition matrix

B_k is the Kalman filter model control input matrix

u_k is the Kalman filter model system input vector

G_k is the Kalman filter model noise input matrix

w_k is an additive white Gaussian discrete-time dynamics noise input used in the Kalman filter model, with zero mean and

$$E\{w_k(t) w_k^T(t_j)\} = \begin{cases} Q_k, & t_i = t_j \\ 0, & t_i \neq t_j \end{cases} \quad (6)$$

z_k is the Kalman filter model measurement vector

H_k is the Kalman filter model output matrix

v_k is an additive white Gaussian measurement noise input that is used in the Kalman filter model, assumed to be independent of w_k , and zero-mean with

$$E\{v_k(t) v_k^T(t_j)\} = \begin{cases} R_k, & t_i = t_j \\ 0, & t_i \neq t_j \end{cases} \quad (7)$$

The k^{th} Kalman filter algorithm uses the k^{th} design model to define time propagation and measurement update equations of the state estimates and the state estimate covariance matrix. The state estimate propagation equation based on the k^{th} Kalman filter model is:

$$\begin{aligned} \hat{x}_k(t_i^+) &= \Phi_k \hat{x}_k(t_{i-1}^+) + B_k u(t_{i-1}) \\ \hat{z}_k(t_i^+) &= H_k \hat{x}_k(t_i^+) \end{aligned} \quad (8)$$

where \hat{x}_k is the Kalman filter state estimate vector

$\hat{z}_k(t_i^+)$ 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.

Also, the state estimate covariance matrix propagation equation is:

$$P_k(t_i) = \Phi_k P_k(t_{i-1}^+) \Phi_k^T + G_k Q_k G_k^T \quad (9)$$

The state estimates are updated using:

$$\hat{x}_k(t_i) = \hat{x}_k(t_i^+) + K_k(t_i) r_k(t_i) \quad (10)$$

where the Kalman filter gain is:

$$K_k(t_i) = P_k(t_i) H_k^T A_k(t_i)^{-1} \quad (11)$$

and the filter-computed residual covariance matrix A_k is:

$$A_k(t_i) = H_k P_k(t_i) H_k^T + R_k \quad (12)$$

The filter residual vector, shown in Eq (10), is defined as:

$$r_k(t_i) \triangleq z(t_i) - H_k \hat{x}_k(t_i^+) = z_r(t_i) - H_k \hat{x}_k(t_i^+) \quad (13)$$

which is simply the difference between the measurements (z) and the Kalman filter estimates, based on its model, of those measurements before they are taken [$H_k \hat{x}_k(t_i^+)$]. Finally, the filter state estimate covariance matrix is updated using:

$$P_k(t_i) = P_k(t_i^+) - K_k(t_i) H_k P_k(t_i^+) \quad (14)$$

The steady state Kalman filter can be precomputed by iterating Eqs (9), (11), (12), and (14) until steady state is reached. With this steady state implementation, the state covariance matrix, the steady state Kalman filter gain, and the steady state Kalman filter residual covariance matrices are assumed to be constant and therefore do not need to be computed in real time. The steady state Kalman filter equations become:

$$\hat{x}_k(t_i) = \Phi_k \hat{x}_k(t_{i-1}^+) + B_k u(t_{i-1}) \quad (15)$$

for propagating the state estimates and

$$\hat{x}_k(t_i) = \hat{x}_k(t_i^+) + K_k r_k(t_i) \quad (16)$$

for updating the state estimates.

3.3 Neyman-Pearson Based Hypothesis Testing

Previous research [1, 2, 3] has developed a methodology to compute the mean and covariance of the Kalman filter residual given that the Kalman filter model differs from the true in the state transition matrix, control input matrix, and output matrix. The various hypotheses that were explored for this research are defined by ΔB_k , where

$$\Delta B_k \triangleq B_T - B_k \rightarrow B_k = B_T - \Delta B_k \quad (17)$$

We denote a set of these parameters as θ and the parameter space of all possible parameter variations as Θ . A particular hypothesis is constructed by defining a subset of the parameter space:

$$h_k: \theta \in \Theta_k \quad (18)$$

We also assume that the various hypotheses form a disjoint covering of the parameter space, so that

$$\Theta = \Theta_0 \cup \Theta_1 \cup \dots \cup \Theta_{N-1} \quad (19)$$

If we test h_0 versus h_1 versus ... h_{N-1} , then we have a N-ary hypothesis test. The primary hypothesis, which is the assumed true hypothesis at the time of the test, will be denoted h_0 . The other hypotheses will be called the alternate hypotheses. Thus, a binary hypothesis test is a test of the primary hypothesis, h_0 , versus the alternative hypothesis, h_1 . N-ary hypothesis testing is an extension of binary hypothesis testing by simply making N-1 binary hypothesis tests between the primary hypothesis, h_0 , and the N-1 alternate hypotheses, h_1 through h_{N-1} , to obtain the desired N-ary hypothesis test. This is sufficient because the covering in Eq (19) is assumed to be disjoint. If each subspace Θ_k contains a single element (i.e. represents a single flight control failure status), then the hypothesis h_k is called a simple hypothesis. We will assume that the hypotheses we are researching are simple hypotheses (each subspace contains only a single TVC nozzle failure status hypothesis).




Figure 4. Neyman-Pearson Hypothesis Test.

The basic concept behind the Neyman-Pearson Hypothesis test is graphically represented in Figure 4. We start by assuming the one hypothesis, called the primary hypothesis denoted as hypothesis zero is Figure 4, is currently correct. With this hypothesis (hypothesis zero) the Kalman filter residual would have the distribution (assumed Gaussian) shown on the left of Figure 4, i.e. with a mean of $mean_0$. Under another possible hypothesis, called the alternate hypothesis and denoted hypothesis one in Figure 4, the residual has a different distribution which is shown on the right of Figure 4, i.e. same covariance but with a mean of $mean_1$. Given the desired probability of false alarm, we can compute a threshold, Γ . If the sample of the residual is above Γ , then we would conclude that hypothesis one is most likely correct.

Likewise, if the sample was below Γ we would conclude that hypothesis zero was most likely correct.

There are two types of errors that can occur when testing hypotheses. Type I errors, called false alarms, occur when the alternate hypothesis is chosen when the true hypothesis is the primary hypothesis. The probability that this type of error can occur is called either the probability of false alarm, denoted P_{FA} , or the size of the test, denoted α . Type II errors, called missed detection, occur when the primary hypothesis is chosen when the true hypothesis is the alternate hypothesis. The probability that this type of error can occur is called the probability of a miss. This probability is $1-P_D$, where P_D is the probability of detection, also called the power of the test.

Note that, these definitions for a false alarm and a missed detection are used for Neyman-Pearson hypothesis testing [4]. Researchers in the failure identification field usually define a false alarm as a declaration of a failure, when in fact one had not occurred, and a missed detection as a declaration of no-failure, when in fact a failure had occurred. These definitions agree with the Neyman-Pearson definitions above, if the current hypothesis is the no-failure hypothesis. Disagreement occurs when a failure hypothesis is in force. For example, if the assumed primary hypothesis is a failed elevator, and the hypothesis test declares a no-failure when in fact the elevator failure remains true, the Neyman-Pearson definition would define this a false alarm but the usual terminology would define this a missed detection. We choose to use the Neyman-Pearson definitions because they are defined, without ambiguity, by the design parameters P_{FA} and the P_D .

A binary test, like the one just described, of $h_0: \theta \in \Theta_0$ (the parameter set is in the parameter subspace Θ_0) versus $h_1: \theta \in \Theta_1$ of a residual would take the form

$$\phi\{r(t_i)\} = \begin{cases} 1 & \sim h_1, \quad r \in R \\ 0 & \sim h_0, \quad r \in A. \end{cases} \quad (20)$$

This is read [4] "the test function $\phi(r)$ equals 1, and hypothesis h_1 is accepted, if the residual lies in the h_0 rejection region R. The test function equals zero, and hypothesis h_0 is accepted, if the residual lies in the h_0 acceptance region A."

The Neyman-Pearson Lemma [4: 107] tells us that the test of the form

$$\phi\{r\} = \begin{cases} 1, & \text{if } f_{\theta_1}(r) > \Gamma f_{\theta_0}(r) \\ \zeta, & \text{if } f_{\theta_1}(r) = \Gamma f_{\theta_0}(r) \\ 0, & \text{if } f_{\theta_1}(r) < \Gamma f_{\theta_0}(r) \end{cases} \quad (21)$$

for some test threshold $\Gamma \geq 0$ and some $0 \leq \zeta \leq 1$, is the most powerful test (the test with the largest P_D) of size $\alpha > 0$ (for a given P_{FA}) for testing $h_0: \theta \in \Theta_0$ versus

$h_1: \theta = \theta_1$. This test can be rewritten by defining the likelihood ratio as:

$$l(\mathbf{r}) = \frac{f_{\theta_1}(\mathbf{r})}{f_{\theta_0}(\mathbf{r})} \quad (22)$$

so that the most powerful test of size α for testing $h_0: \theta = \theta_0$ versus $h_1: \theta = \theta_1$ is a likelihood ratio test of the form

$$\phi\{\mathbf{r}\} = \begin{cases} 1, & \text{if } l(\mathbf{r}) > \Gamma \\ \zeta, & \text{if } l(\mathbf{r}) = \Gamma \\ 0, & \text{if } l(\mathbf{r}) < \Gamma \end{cases} \quad (23)$$

If $l(\mathbf{r}) = \Gamma$ with probability zero, then $\zeta = 0$ and the threshold Γ , for a given size (P_{FA}) and power (P_D), can be computed from:

$$P_{FA} = P_{\theta_0}[l(\mathbf{r}) > \Gamma] = \int_{\Gamma}^{\infty} f_{\theta_0}(l) dl \quad P_D = \int_{\Gamma}^{\infty} f_{\theta_1}(l) dl \quad (24)$$

where $f_{\theta_0}(l)$ is the density function of $l(\mathbf{r})$ under h_0 .

3.4 Single Residual Equivalence

We will now show that the likelihood ratio in Eq (22), which operates on a single residual, is equivalent to a likelihood ratio that operates on two different residuals. This will allow us to use the structure in Figure 1, where a *single* Kalman filter residual is used with multiple hypothesis tests, in place of the classical MMAE structure. First, we assume that the density functions used in the likelihood ratio are conditioned on the measurement history $\mathbf{Z}(t_{i-1})$. We will compare two residuals from any two different Kalman filters, \mathbf{r}_j and \mathbf{r}_k , generated at time t_i . The Kalman filter models used in these filters use two different hypotheses that differ from the true system model only in the control input matrix, allowing us to describe these residuals in terms of $\Delta\mathbf{B}_k$ and $\Delta\mathbf{B}_j$. We compare the two likelihood ratios to check for equality and for simplicity of notation, we are suppressing the time argument, t_i , of the residuals:

$$\frac{f_{\theta_1}(\mathbf{r}_j)}{f_{\theta_0}(\mathbf{r}_k)} \stackrel{?}{=} \frac{f_{\theta_1}(\mathbf{r}_j)}{f_{\theta_0}(\mathbf{r}_k)}$$

$$\frac{\beta_1 \exp\left\{-\frac{1}{2}[\mathbf{r}_j - E_{Z(t_{i-1})}\{\mathbf{r}_j\}]_{h_1}^T \mathbf{A}_1^{-1} [\mathbf{r}_j - E_{Z(t_{i-1})}\{\mathbf{r}_j\}]_{h_1}\right\}}{\beta_0 \exp\left\{-\frac{1}{2}[\mathbf{r}_k - E_{Z(t_{i-1})}\{\mathbf{r}_k\}]_{h_0}^T \mathbf{A}_0^{-1} [\mathbf{r}_k - E_{Z(t_{i-1})}\{\mathbf{r}_k\}]_{h_0}\right\}} \stackrel{?}{=} \frac{\beta_1 \exp\left\{-\frac{1}{2}[\mathbf{r}_j - E_{Z(t_{i-1})}\{\mathbf{r}_j\}]_{h_1}^T \mathbf{A}_1^{-1} [\mathbf{r}_j - E_{Z(t_{i-1})}\{\mathbf{r}_j\}]_{h_1}\right\}}{\beta_0 \exp\left\{-\frac{1}{2}[\mathbf{r}_j - E_{Z(t_{i-1})}\{\mathbf{r}_j\}]_{h_0}^T \mathbf{A}_0^{-1} [\mathbf{r}_j - E_{Z(t_{i-1})}\{\mathbf{r}_j\}]_{h_0}\right\}} \quad (25)$$

The primary hypothesis, h_0 , and the alternative hypothesis, h_1 , are not necessarily the same hypotheses that were used for the Kalman filter models of the filters that generated the residuals.

Examining Eq (25), we note that the numerators and leading coefficients are equal. Thus the equality will hold if and only if:

$$\mathbf{r}_k - E_{Z(t_{i-1})}\{\mathbf{r}_k\}|_{h_0} \stackrel{?}{=} \mathbf{r}_j - E_{Z(t_{i-1})}\{\mathbf{r}_j\}|_{h_0} \quad (26)$$

We did not restrict the hypothesis h_1 , so if the equality of Eq (26) holds, we can use this development for any alternative hypothesis. We assume that the hypothesis, h_0 , can consist of any of the modeling differences that we are investigating. First, we define the error between the k^{th} Kalman filter state estimate and the true system model states as:

$$\epsilon_k(t_i) \triangleq \mathbf{x}_T(t_i) - \hat{\mathbf{x}}_k(t_i) \quad (27)$$

Now we expand one of the terms in Eq (13), using the definitions of Eq (5) for the true system model, and Eqs (8), (17), and (27) to get:

$$\begin{aligned} \mathbf{r}_k(t_i) \triangleq \mathbf{z}(t_i) - \mathbf{H}_k \hat{\mathbf{x}}_k(t_i) &= [\mathbf{H}_T \mathbf{x}_T(t_i) + \mathbf{v}_T(t_i)] - \mathbf{H}_k \hat{\mathbf{x}}_k(t_i) \\ &= \mathbf{H}_T [\Phi_T \mathbf{x}_T(t_{i-1}) + \mathbf{B}_T \mathbf{u}_T(t_{i-1}) + \mathbf{G}_T \mathbf{w}_d(t_{i-1})] + \mathbf{v}_T(t_i) \\ &\quad - \mathbf{H}_k [\Phi_k \hat{\mathbf{x}}_k(t_{i-1}) + \mathbf{B}_k \mathbf{u}_k(t_{i-1})] \\ &= \mathbf{H}_T \Phi_T \mathbf{x}_T(t_{i-1}) - \mathbf{H}_k \Phi_k \hat{\mathbf{x}}_k(t_{i-1}) + \mathbf{H}_T \mathbf{G}_T \mathbf{w}_d(t_{i-1}) + \mathbf{v}_T(t_i) \\ &\quad + \mathbf{H}_T \mathbf{B}_T \mathbf{u}_T(t_{i-1}) - \mathbf{H}_k \mathbf{B}_k \mathbf{u}_k(t_{i-1}) \\ &= \mathbf{H}_T \Phi_T \mathbf{x}_T(t_{i-1}) - \mathbf{H}_T \Phi_T \hat{\mathbf{x}}_k(t_{i-1}) + \mathbf{H}_T \mathbf{G}_T \mathbf{w}_d(t_{i-1}) + \mathbf{v}_T(t_i) \\ &\quad + [\mathbf{H}_T \mathbf{B}_T - \mathbf{H}_T (\mathbf{B}_T - \Delta\mathbf{B}_k)] \mathbf{u}(t_{i-1}) \\ &= \mathbf{H}_T \Phi_T [\mathbf{x}_T(t_{i-1}) - \hat{\mathbf{x}}_k(t_{i-1})] + \mathbf{H}_T \mathbf{G}_T \mathbf{w}_d(t_{i-1}) + \mathbf{v}_T(t_i) \\ &\quad + \mathbf{H}_T \Delta\mathbf{B}_k \mathbf{u}(t_{i-1}) \\ &= \mathbf{H}_T \Phi_T \epsilon_k(t_{i-1}) + \mathbf{H}_T \Delta\mathbf{B}_k \mathbf{u}(t_{i-1}) + \mathbf{H}_T \mathbf{G}_T \mathbf{w}_d(t_{i-1}) + \mathbf{v}_T(t_i) \end{aligned} \quad (28)$$

We take the conditional expectation of this to find the mean of the residual:

$$\begin{aligned} E_{Z(t_{i-1})}\{\mathbf{r}_k(t_i)\} &= \mathbf{H}_T \Phi_T E_{Z(t_{i-1})}\{\epsilon_k(t_{i-1})\} + \mathbf{H}_T \Delta\mathbf{B}_k \mathbf{u}(t_{i-1}) \\ &\quad + \mathbf{H}_T \mathbf{G}_T E_{Z(t_{i-1})}\{\mathbf{w}_d(t_{i-1})\} + E_{Z(t_{i-1})}\{\mathbf{v}_T(t_i)\} \\ E_{Z(t_{i-1})}\{\mathbf{r}_k(t_i)\} &= \mathbf{H}_T \Phi_T E_{Z(t_{i-1})}\{\epsilon_k(t_{i-1})\} + \mathbf{H}_T \Delta\mathbf{B}_k \mathbf{u}(t_{i-1}) \end{aligned} \quad (29)$$

We substitute Eq (28) and Eq (29) into Eq (26) and cancel terms to get:

$$\begin{aligned} \mathbf{H}_T \Phi_T [\epsilon_k(t_{i-1}) - E_{Z(t_{i-1})}\{\epsilon_k(t_{i-1})\}] + \mathbf{H}_T \mathbf{G}_T \mathbf{w}_d(t_{i-1}) + \mathbf{v}_T(t_i) \\ \stackrel{?}{=} \mathbf{H}_T \Phi_T [\epsilon_j(t_{i-1}) - E_{Z(t_{i-1})}\{\epsilon_j(t_{i-1})\}] + \mathbf{H}_T \mathbf{G}_T \mathbf{w}_d(t_{i-1}) + \mathbf{v}_T(t_i) \\ - \epsilon_k(t_{i-1}) - E_{Z(t_{i-1})}\{\epsilon_k(t_{i-1})\} \stackrel{?}{=} \epsilon_j(t_{i-1}) - E_{Z(t_{i-1})}\{\epsilon_j(t_{i-1})\} \end{aligned} \quad (30)$$

Using the definition in Eq (27) we get:

$$\begin{aligned} [\mathbf{x}_T(t_{i-1}) - \hat{\mathbf{x}}_k(t_{i-1})] - E_{Z(t_{i-1})}\{\mathbf{x}_T(t_{i-1}) - \hat{\mathbf{x}}_k(t_{i-1})\} \\ \stackrel{?}{=} [\mathbf{x}_T(t_{i-1}) - \hat{\mathbf{x}}_j(t_{i-1})] - E_{Z(t_{i-1})}\{\mathbf{x}_T(t_{i-1}) - \hat{\mathbf{x}}_j(t_{i-1})\} \\ \mathbf{x}_T(t_{i-1}) - \hat{\mathbf{x}}_k(t_{i-1}) - E_{Z(t_{i-1})}\{\mathbf{x}_T(t_{i-1})\} + \hat{\mathbf{x}}_k(t_{i-1}) \\ \stackrel{?}{=} \mathbf{x}_T(t_{i-1}) - \hat{\mathbf{x}}_j(t_{i-1}) - E_{Z(t_{i-1})}\{\mathbf{x}_T(t_{i-1})\} + \hat{\mathbf{x}}_j(t_{i-1}) \\ \mathbf{x}_T(t_{i-1}) - E_{Z(t_{i-1})}\{\mathbf{x}_T(t_{i-1})\} \stackrel{?}{=} \mathbf{x}_T(t_{i-1}) - E_{Z(t_{i-1})}\{\mathbf{x}_T(t_{i-1})\} \end{aligned} \quad (31)$$

Therefore

$$\frac{f_{\delta_1}(r_j)}{f_{\delta_0}(r_k)} = \frac{f_{\delta_1}(r_j)}{f_{\delta_0}(r_j)} \quad (32)$$

Thus we can perform the hypothesis test using a single Kalman filter residual, testing it for multiple hypotheses, and this is equivalent to performing the hypothesis test using multiple residuals, each one tested for a single hypothesis.

3.5 Binary Hypothesis Testing

Now we can apply the Neyman-Pearson hypothesis test to the residual of a Kalman filter. We need an estimate of the residual covariance for a particular hypothesis, which we can precompute, as described in Section 3.1, by using a Kalman filter that uses the hypothesized system model. We will use the steady state Kalman filter estimate of the residual covariance matrix, thus the residual covariance will be considered constant. Note that the estimate for the residual covariance is not a function of the system input matrix, \mathbf{B} . Thus, if we only have modeling differences in the system input matrix, so only $\Delta\mathbf{B}$ is nonzero, then the residual covariance matrix will be identical for each of the hypotheses. We will develop the Neyman-Pearson based hypothesis testing algorithm for the case of $\Delta\mathbf{B} \neq \mathbf{0}$.

Under these assumptions, we have a Neyman-Pearson detector for "common covariances, uncommon means" as developed by Scharf [4: 111]. We will develop the Neyman-Pearson detector in the same manner as Scharf does, but we will extend it to a single-time sample hypothesis test using a multidimensional random variable (a single time sample of the residual).

We start with a single time sample of a Kalman filter residual and design a binary test using the Neyman-Pearson Lemma. We denote the mean of the residual given the past measurement history and a particular hypothesis as

$$\begin{aligned} E_{z(t-1)}\{r(t)\} |_{h_0} &= m_0(t) \\ E_{z(t-1)}\{r(t)\} |_{h_1} &= m_1(t). \end{aligned} \quad (33)$$

Using Eq (22), the likelihood ratio is a ratio of two normal density functions, which becomes the ratio of two exponential terms when the leading β terms are canceled. Thus the likelihood ratio becomes:

$$\begin{aligned} l(r(t)) &= \exp \left\{ -\frac{1}{2} [r(t) - m_1(t)]^T A_k^{-1} [r(t) - m_1(t)] \right. \\ &\quad \left. + \frac{1}{2} [r(t) - m_0(t)]^T A_k^{-1} [r(t) - m_0(t)] \right\} \\ &= \exp \left\{ \frac{1}{2} [r(t)^T A_k^{-1} m_1(t) + m_1(t)^T A_k^{-1} r(t) - m_1(t)^T A_k^{-1} m_1(t) \right. \\ &\quad \left. - r(t)^T A_k^{-1} m_0(t) - m_0(t)^T A_k^{-1} r(t) + m_0(t)^T A_k^{-1} m_0(t) \right\}. \end{aligned} \quad (34)$$

Note that each term within the exponential is a scalar and

we can use the symmetric property of the covariance matrix A_k to get:

$$\begin{aligned} (A_k^{-1})^T &= A_k^{-1} \\ r(t)^T A_k^{-1} m_1(t) &= (r(t)^T A_k^{-1} m_1(t))^T \\ &= m_1(t)^T (A_k^{-1})^T r(t) = m_1(t)^T A_k^{-1} r(t) \\ r(t)^T A_k^{-1} m_0(t) &= (r(t)^T A_k^{-1} m_0(t))^T \\ &= m_0(t)^T (A_k^{-1})^T r(t) = m_0(t)^T A_k^{-1} r(t). \end{aligned} \quad (35)$$

Also, we can use the natural log of the likelihood ratio in place of the likelihood ratio for decision-making because the natural logarithm is a monotonic function. Therefore we get the log likelihood ratio:

$$\begin{aligned} L(r(t)) &= \ln \{l(r(t))\} \\ &= m_1(t)^T A_k^{-1} r(t) - m_0(t)^T A_k^{-1} r(t) \\ &\quad - \frac{1}{2} (m_1(t)^T A_k^{-1} m_1(t) - m_0(t)^T A_k^{-1} m_0(t)) \\ &= (m_1(t) - m_0(t))^T A_k^{-1} r(t) \\ &\quad - \frac{1}{2} (m_1(t) - m_0(t))^T A_k^{-1} (m_1(t) + m_0(t)) \\ &= s(t) - b(t) \end{aligned} \quad (36)$$

where

$$\begin{aligned} s(t) &= (m_1(t) - m_0(t))^T A_k^{-1} r(t) \\ b(t) &= \frac{1}{2} (m_1(t) - m_0(t))^T A_k^{-1} (m_1(t) + m_0(t)). \end{aligned} \quad (37)$$

We call $s(t)$ the signal and $b(t)$ the bias of the log likelihood function.

Rewriting the Neyman-Pearson Lemma to use the log likelihood ratio in Eq (36), we get:

$$\phi\{r(t)\} = \begin{cases} 1, & L(r(t)) > \eta = \ln(\Gamma) \\ 0, & L(r(t)) \leq \eta. \end{cases} \quad (38)$$

Now we only need to compute the threshold η in order to define the Neyman-Pearson hypothesis test algorithm completely. We first need to find the probability density function for the test statistic $L(r(t))$. We note from Eq (36) that L is a linear function of r , therefore L is also normally distributed and we need only compute the mean and covariance to define the probability density function. We'll start by computing the mean under both the primary and alternate hypotheses.

First we define:

$$\begin{aligned} d(t) &= m_1(t) - m_0(t) \\ -L(r(t)) &= d(t)^T A_k^{-1} r(t) \\ &\quad - \frac{1}{2} d(t)^T A_k^{-1} (m_1(t) + m_0(t)) \\ &= s(t) - b(t). \end{aligned} \quad (39)$$

Now we compute the mean of the test statistic, given that the primary hypothesis is in force:

$$\begin{aligned} E_{z(t-1)}\{L(r(t))\} |_{h_0} &= d(t)^T A_k^{-1} E_{z(t-1)}\{r(t)\} |_{h_0} \\ &\quad - \frac{1}{2} d(t)^T A_k^{-1} (m_1(t) + m_0(t)) \\ &= d(t)^T A_k^{-1} m_0(t) - \frac{1}{2} d(t)^T A_k^{-1} m_1(t) - \frac{1}{2} d(t)^T A_k^{-1} m_0(t) \end{aligned}$$

$$\begin{aligned}
E_{z(t_i)}\{L(\mathbf{r}(t_i))\}|_{h_0} &= \frac{1}{2} \mathbf{d}(t_i)^T \mathbf{A}_k^{-1} (\mathbf{m}_0(t_i) - \mathbf{m}_1(t_i)) \\
&= -\frac{1}{2} \mathbf{d}(t_i)^T \mathbf{A}_k^{-1} \mathbf{d}(t_i) = -\frac{1}{2} D(t_i)
\end{aligned} \quad (40)$$

where $D(t_i)$ will be called the hypothesis discrimination measure.

Likewise, we compute the mean of the test statistic given that the alternative hypothesis is in force:

$$\begin{aligned}
E_{z(t_i)}\{L(\mathbf{r}(t_i))\}|_{h_1} &= \mathbf{d}(t_i)^T \mathbf{A}_k^{-1} E_{z(t_i)}\{\mathbf{r}(t_i)\}|_{h_1} \\
&\quad - \frac{1}{2} \mathbf{d}(t_i)^T \mathbf{A}_k^{-1} (\mathbf{m}_1(t_i) + \mathbf{m}_0(t_i)) \\
&= \mathbf{d}(t_i)^T \mathbf{A}_k^{-1} \mathbf{m}_1(t_i) - \frac{1}{2} \mathbf{d}(t_i)^T \mathbf{A}_k^{-1} \mathbf{m}_1(t_i) - \frac{1}{2} \mathbf{d}(t_i)^T \mathbf{A}_k^{-1} \mathbf{m}_0(t_i) \\
&= \frac{1}{2} \mathbf{d}(t_i)^T \mathbf{A}_k^{-1} (\mathbf{m}_1(t_i) - \mathbf{m}_0(t_i)) \\
&= \frac{1}{2} \mathbf{d}(t_i)^T \mathbf{A}_k^{-1} \mathbf{d}(t_i) = \frac{1}{2} D(t_i).
\end{aligned} \quad (41)$$

Now to compute the conditional variance of the test statistic under either hypothesis:

$$\begin{aligned}
\text{var}_{z(t_i)}\{L(\mathbf{r}(t_i))\} &= E_{z(t_i)}\{L(\mathbf{r}(t_i))^2\} - [E_{z(t_i)}\{L(\mathbf{r}(t_i))\}]^2 \\
&= E_{z(t_i)}\{[s(t_i) - \mathbf{b}(t_i)]^2\} - [E_{z(t_i)}\{s(t_i) - \mathbf{b}(t_i)\}]^2 \\
&= E_{z(t_i)}\{s(t_i)^2 - 2s(t_i)\mathbf{b}(t_i) + \mathbf{b}(t_i)^2\} - [E_{z(t_i)}\{s(t_i)\} - \mathbf{b}(t_i)]^2 \\
&= E_{z(t_i)}\{s(t_i)^2\} - 2E_{z(t_i)}\{s(t_i)\}\mathbf{b}(t_i) + \mathbf{b}(t_i)^2 \\
&\quad - [E_{z(t_i)}\{s(t_i)\}^2 - 2E_{z(t_i)}\{s(t_i)\}\mathbf{b}(t_i) + \mathbf{b}(t_i)^2] \\
&= E_{z(t_i)}\{s(t_i)^2\} - E_{z(t_i)}\{s(t_i)\}^2 = \text{var}_{z(t_i)}\{s(t_i)\}.
\end{aligned} \quad (42)$$

Now use Eq (37) to get:

$$\begin{aligned}
\text{var}_{z(t_i)}\{L(\mathbf{r}(t_i))\} &= E_{z(t_i)}\{s(t_i)^2\} - E_{z(t_i)}\{s(t_i)\}^2 \\
&= E_{z(t_i)}\{[\mathbf{d}(t_i)^T \mathbf{A}_k^{-1} \mathbf{r}(t_i)]^2\} - E_{z(t_i)}\{[\mathbf{d}(t_i)^T \mathbf{A}_k^{-1} \mathbf{r}(t_i)]\}^2 \\
&= E_{z(t_i)}\{[\mathbf{d}(t_i)^T \mathbf{A}_k^{-1} \mathbf{r}(t_i)] [\mathbf{d}(t_i)^T \mathbf{A}_k^{-1} \mathbf{r}(t_i)]^T\} \\
&\quad - E_{z(t_i)}\{[\mathbf{d}(t_i)^T \mathbf{A}_k^{-1} \mathbf{r}(t_i)]\} E_{z(t_i)}\{[\mathbf{d}(t_i)^T \mathbf{A}_k^{-1} \mathbf{r}(t_i)]^T\} \\
&= \mathbf{d}(t_i)^T \mathbf{A}_k^{-1} E_{z(t_i)}\{\mathbf{r}(t_i) \mathbf{r}(t_i)^T\} \mathbf{A}_k^{-1} \mathbf{d}(t_i) \\
&\quad - \mathbf{d}(t_i)^T \mathbf{A}_k^{-1} E_{z(t_i)}\{\mathbf{r}(t_i)\} E_{z(t_i)}\{\mathbf{r}(t_i)^T\} \mathbf{A}_k^{-1} \mathbf{d}(t_i) \\
&= \mathbf{d}(t_i)^T \mathbf{A}_k^{-1} [E_{z(t_i)}\{\mathbf{r}(t_i) \mathbf{r}(t_i)^T\} - E_{z(t_i)}\{\mathbf{r}(t_i)\} E_{z(t_i)}\{\mathbf{r}(t_i)^T\}] \mathbf{A}_k^{-1} \mathbf{d}(t_i) \\
&= \mathbf{d}(t_i)^T \mathbf{A}_k^{-1} \mathbf{A}_k \mathbf{A}_k^{-1} \mathbf{d}(t_i) \\
&= \mathbf{d}(t_i)^T \mathbf{A}_k^{-1} \mathbf{d}(t_i) = D(t_i).
\end{aligned} \quad (43)$$

To summarize, in the case of a TVC nozzle failure we know that the test statistic is normally distributed as follows:

$$\begin{aligned}
h_0 : L(\mathbf{r}(t_i)) &\sim N\left[-\frac{D(t_i)}{2}, D(t_i)\right] \\
h_1 : L(\mathbf{r}(t_i)) &\sim N\left[\frac{D(t_i)}{2}, D(t_i)\right].
\end{aligned} \quad (44)$$

We substitute this into the known functional form for a normal density, which can then be used to compute the threshold η for a given P_{FA} . Using the definition given in Eq (24), we get:

$$\begin{aligned}
\alpha = P_{FA} &= \int_{\eta}^{\infty} \frac{1}{(2\pi D(t_i))^{1/2}} \exp\left\{-\frac{1}{2} \frac{\left(x + \frac{D(t_i)}{2}\right)^2}{D(t_i)}\right\} dx \\
&= \int_{\left[\frac{\eta + \frac{D(t_i)}{2}}{D(t_i)^{1/2}}\right]}^{\infty} \frac{1}{(2\pi)^{1/2}} \exp\left\{-\frac{y^2}{2}\right\} dy \\
&= 1 - \gamma\left(\frac{\eta + \frac{D(t_i)}{2}}{D(t_i)^{1/2}}\right) = 1 - \gamma(g)
\end{aligned} \quad (45)$$

where

$$\begin{aligned}
g &= \frac{\eta + \frac{D(t_i)}{2}}{D(t_i)^{1/2}} \quad \gamma(g) = \int_{-g}^g \frac{1}{(2\pi)^{1/2}} \exp\left\{-\frac{x^2}{2}\right\} dx \\
\gamma(-g) &= 1 - \gamma(g).
\end{aligned} \quad (46)$$

Similarly, we can compute the P_D :

$$\begin{aligned}
P_D &= \int_{\eta}^{\infty} \frac{1}{(2\pi D(t_i))^{1/2}} \exp\left\{-\frac{1}{2} \frac{\left(x - \frac{D(t_i)}{2}\right)^2}{D(t_i)}\right\} dx \\
&= \int_{\left[\frac{\eta - \frac{D(t_i)}{2}}{D(t_i)^{1/2}}\right]}^{\infty} \frac{1}{(2\pi)^{1/2}} \exp\left\{-\frac{y^2}{2}\right\} dy \\
&= 1 - \gamma\left(\frac{\eta - \frac{D(t_i)}{2}}{D(t_i)^{1/2}}\right) = 1 - \gamma(g - D(t_i)^{1/2}).
\end{aligned} \quad (47)$$

For most applications, the designer would set P_{FA} and P_D to achieve the desired hypothesis testing performance, and that would dictate a unique test threshold, Γ , and a desired hypothesis discrimination measure, $D(t_i)$, needed to achieve the P_{FA} and P_D . However, this does not give us a unique $\mathbf{d}(t_i)$ needed to achieve the desired discrimination measure, even for a given \mathbf{A}_k^{-1} . Let's assume that one of the hypotheses, say h_0 , is correct, so that $\mathbf{m}_0(t_i) = \mathbf{0}$. This gives us $\mathbf{d}(t_i) = \mathbf{m}_1(t_i)$, which we can use, along with the precomputable \mathbf{A}_k^{-1} , to get a class of residual means that would achieve the desired discrimination measure.

3.6 N-ary Hypothesis Testing

The algorithm developed in the previous section performs a single binary hypothesis test, which can be extended to any N -ary test by using the algorithm against several alternative hypotheses. We simply define the mean of the residual given a specific alternative hypothesis as:

$$E_{z(t_i)}\{\mathbf{r}(t_i)\}|_{h_k} = \mathbf{m}_k(t_i) \quad (48)$$

where k denotes the specific alternative hypothesis, and define $L_k(\mathbf{r}(t_i))$ as the test statistic and $D_k(t_i)$ as the discrimination measure for the k^{th} hypothesis. The development for each of the alternative hypotheses is the same as in the previous section, except that we replace the

subscript 1 with k so that we are testing hypothesis 0 against hypothesis k .

All of the alternative hypothesis tests are run in parallel, so all of the test statistics and discrimination measures are computed in parallel (not sequentially) with those of the other hypotheses. Once the discrimination measures are computed for each of the alternative hypotheses, we collect the hypotheses whose discrimination measure exceeds the desired discrimination measure (which is based on P_{FA} and P_D). We compare the computed test statistics for each hypothesis against the test threshold and keep the hypotheses whose test statistic exceed the test threshold. We then choose the most likely hypothesis by choosing the hypothesis with the largest test statistic.

IV. Simulation Results

4.1 Computation of Kalman Filter Residual

We computed the residual mean and covariance using the development in previous research [1, 3] and compared it to the actual residual for various failure hypotheses. A representative sample of these comparisons are shown in Figures 5, 6, and 7 where the actual residual is plotted against the computed mean of the residual along with dashed lines that show one standard deviation above and below the computed mean. Thus, we would expect the 68% of the actual residual samples would fall within the dashed lines. We show the case where there is no failure, Figure 5, the case where nozzle 1 is stuck closed, Figure 6, and the case where nozzle 1 is stuck open, Figure 7. Note that in all these cases the actual residual matches the computed mean and covariance, thus verifying our development of the expected residual mean and covariance.

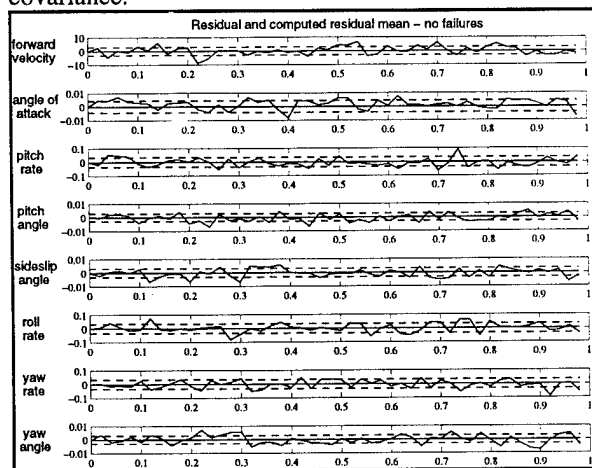


Figure 5. Computed and Actual Kalman Filter Residual for the case of no TVC Nozzle Failure.

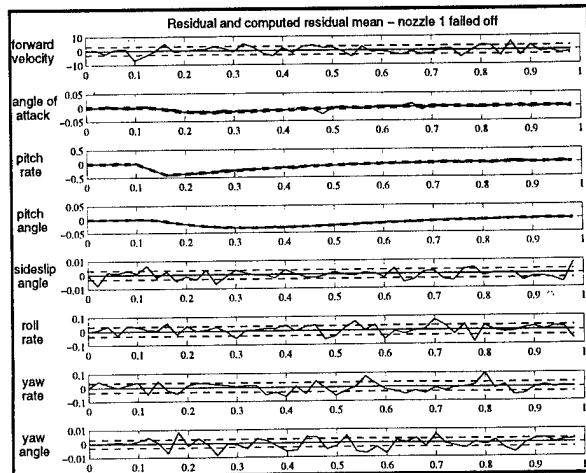


Figure 6. Computed and Actual Kalman Filter Residual for the case of a TVC Nozzle Failed to Off.




Figure 7. Computed and Actual Kalman Filter Residual for the case of a TVC Nozzle Failed to On.

4.2 Failure Identification Performance

Using the development in Section 3, we implemented a Neyman-Pearson Hypothesis Testing Algorithm (NPHTA), shown in Figure 2. The output of this algorithm is a declaration of the failure status, not a set of conditional probabilities. We chose to compare the declared failure status to the true failure status to evaluate the performance of this algorithm. In Figure 8 we present the failure identification performance of this structure, by plotting the agreement (denoted by 1) and disagreement (denoted by 0) of the declared failure status with the true failure status. The nozzle "off" failure were induced at 0.1 seconds into the simulation of the nozzle "on" failures occur at 0.16 seconds. These results were obtained using the full 9-state Kalman filter models of the nozzle "on"




Figure 8. Failure Identification Performance of the Neyman-Pearson Hypothesis Test MMAE Using 9-State Kalman Filter Models.

failures occur at 0.16 seconds. These results were obtained using the full 9-state Kalman filter models.

The design parameters for the NPHTA are the probability of detection and probability of false alarm, which were 0.999 and 0.01 respectively, for this research. Performance sensitivity to the choice of these two parameters could be investigated for a given application, but these are chosen as reasonable, representative values.

We experimented with reducing the computational cost of implementing the NPHTA MMAE by using 3-state Kalman filter models instead of the full 9-state models. We used only the angular rate states (pitch rate, roll rate, and yaw rate), thus the complete missile dynamics were not accurately represented. The results, Figure 9, were quite good. The NPHTA MMAE identified the "off" nozzle failures quickly, but eventually the inaccurate modeling caused the failure state to be misidentified at about 0.35 seconds into the simulation. Also, some of the stuck "on" nozzle failures show some brief difficulty in sorting out which failure status is correct, but eventually correctly identify the failure status. These results were fairly good because




Figure 9. Failure Identification Performance of the Neyman-Pearson Hypothesis Test MMAE Using 3-State Kalman Filter Models.

the nozzle failures produced such dramatic changes in the missile flight path that the differences between the various failure hypotheses were quite dramatic, therefore, the NPHTA could distinguish between these various hypotheses.

V. Conclusions

A Neyman-Pearson Hypothesis Testing Algorithm (NPHTA) has been developed in Section 3. We extended the standard development for a scalar process, to a multidimensional process. We also showed that, by using the NPHTA, we could perform multiple hypothesis tests on a single Kalman filter residual, instead of having to form a residual (and separate Kalman filter) for each hypothesis. Thus, we would use a Kalman filter "bank" of a single Kalman filter, followed by the NPHTA. This structure was implemented and tested, and the results were presented in Section 4.

The NPHTA MMAE structure provides a more rigorous framework for attaining the desired failure identification performance. The MMAE designer specifies the desired P_D and P_{FA} , which gives the required discrimination measure. With this measure

specified, the MMAE designer would iterate on the design of the input dither until the required discrimination measure is attained within the desired identification time. This alleviates the need to perform extensive Monte Carlo testing that is required to estimate the P_D and P_{FA} of the standard MMAE structure. Therefore, we have shown the initial feasibility of using the NPHTA MMAE concept to detect thrust vector nozzle valve failures during the initial separation flyout of an advance air superiority missile.

VI. References

1. Hanlon, P. D. *Practical Implementation of Multiple Model Adaptive Estimation Using Neyman-Pearson Based Hypothesis Testing and Spectral Estimation Tools*. PhD dissertation, AFIT/DS/ENG/96-07. School of Engineering, Air Force Institute of Technology (AU), Wright-Patterson AFB OH, September 1996.
2. Hanlon, P. D., and P. S. Maybeck. "Interrelationship of Single-Filter and Multiple-Model Adaptive Algorithms," to be published in *IEEE Transactions of Aerospace and Electronic Systems*.
3. Hanlon, P. D., and P. S. Maybeck. "Characterization of Kalman Filter Residual in the Presence on Mismodeling," submitted to *IEEE Transactions of Aerospace and Electronic Systems*, October 1997, to be altered before final publication.
4. Scharf, L. L. *Statistical Signal Processing*. Reading, Massachusetts: Addison-Wesley, 1991.