JOURNAL OF GUIDANCE, CONTROL, AND DYNAMICS
Vol. 22, No. 6, November–December 1999

Robust Failure Detection for Reentry Vehicle Attitude Control Systems

Ramses M. Agustin*
Massachusetts Institute of Technology, Cambridge, Massachusetts 02139
and
Rami S. Mangoubi,† Roger M. Hain,‡ and Neil J. Adams§
Charles Stark Draper Laboratory, Cambridge, Massachusetts 02139

This paper presents a robust failure detection methodology for the attitude control system of reusable launch vehicles (RLVs). In particular, we consider the problem of estimating the thrust from multiple jets firing from an RLV reaction control system (RCS), as well as the related problem of distinguishing between failures in the RCS and the aerosurfaces. For accurately known vehicle and sensor models, the Kalman filter provides the optimal estimate for the jet thrust, in the least-squares sense. During reentry, however, plant model uncertainties are a major problem for such a filter as the vehicle's aerodynamics vary widely with rapidly changing Mach number, making gain scheduling impractical. Consequently, the Kalman filter's performance degrades. Even if the Mach number were accurately known, rapid gain scheduling may not be desirable or even possible, because of the large data storage requirements it entails. Transient, robust H_∞ or game-theoretic filters are proposed for next-generation RLV, and a prototype design is demonstrated for the Space Shuttle Orbiter's attitude determination system. Simulation results demonstrate that the robust filters can be insensitive to plant model uncertainties over a much wider range of Mach numbers than a traditional Kalman filter, while remaining sensitive to failures in the aerosurfaces and the RCS jets.

I. Introduction

DURING reentry, reusable launch vehicles (RLVs) typically rely on aerosurfaces and a reaction control system (RCS) for attitude control. For example, the Space Shuttle Oribiter, which is used in this paper as a platform for designing a prototype robust filter for next-generation RLV, has aerosurfaces that consist of the rudder, elevons, speedbrake, and body flap, whereas the RCS consists of bipropellant jets that are fired in the appropriate direction to provide desired thrust and augment the aerosurfaces. During on-orbit operation, only the RCS is used. Shuttle reentry begins with the deorbit burn where the Orbiter is oriented in a tail-first position and jets are fired to slow the vehicle and allow capture by the Earth's gravity and atmosphere. The RCS jets reorient the Orbiter to a high angle of attack, nose-first position. During the first part of reentry, the Orbiter is oriented to a 40-deg angle of attack that is maintained until the vehicle descends and decelerates to below Mach 10. After passing through Mach 10, the Orbiter gradually reduces its angle of attack to 10 deg. The RCS are the sole attitude effectors until the atmospheric dynamic pressure is large enough for the aerosurfaces to become useful. At some altitude inside the transition region, the control surfaces are activated, and attitude control is provided by both the RCS and the aerosurfaces. During the late stages of reentry, control is achieved using the aerosurfaces alone.

Our objective is to detect and isolate failures in the attitude control system. During reentry, the challenge is to distinguish between failures in the aerosurfaces (more specifically, the ailerons, as rudders are not used in the early stages of reentry), and the RCS jets. In addition, we also would like to have an accurate estimate of the jet thrust both in the presence and in the absence of aileron failures. Correct isolation of a failure in the attitude control system is essential to obtaining accurate thrust estimates, which are used for monitoring the health of the RCS. For the purposes of this paper, we choose the reentry flight phase because it is the most difficult stage for the attitude control system's failure detection and isolation. Specifically, during reentry, the vehicle undergoes rapid changes in aerodynamic flight properties as the Mach number decreases, and precise knowledge of these flight properties is not available. The combination of rapid change and model uncertainty also makes it difficult to rely on a single, accurate Orbiter model and a Kalman filter design based on that model, or on multiple models and gain scheduling. Moreover, the uncertainty also makes it difficult to distinguish between failures in the RCS and in the elevon control surfaces. Even in the absence of any failure, the Kalman filter's jet thrust estimation performance degrades considerably in the presence of model uncertainty.

Even if the Mach number were accurately known, rapid gain scheduling may not be desirable or even possible because of the large data storage requirements it entails. In general, however, modeling errors caused by inaccurate knowledge of the Mach number and other factors do exist. As a result, the Kalman filter will not yield accurate thrust estimates. A desirable solution, therefore, is a filter architecture that can rapidly distinguish between anomalies in the RCS and in the elevons. The filter architecture must also provide jet thrust estimates even if a jet misfires, or an aerosurface fails.

The most often used technique to add some measure of robustness to an estimator is to increase the design process noise covariance, i.e., to overdesign the Kalman filter. However, there are limitations to this technique. In particular, it is shown in Refs. 1 and 2 that robust game theoretic or H_∞ filters handle model uncertainties better than robustified Kalman filters. The design of these filters is based on the small gain theorem.[2–8] Such filters are robust to a general class of noise and plant model uncertainties. The steady-state robust filter derived in Ref. 2 has been used in Ref. 9 to estimate the Space Shuttle Orbiter's RCS jet thrust. The results in Ref. 9 demonstrate that steady-state robust filters can yield improved thrust estimates for a wide range of Mach numbers, although the filter response is unacceptably slow.

Received 22 January 1998; revision received 28 May 1999; accepted for publication 4 June 1999.

*Graduate Student, Department of Mechanical Engineering, 77 Massachusetts Avenue; also Draper Lab Fellow, Control and Dynamical Systems Division, C. S. Draper Laboratory, 555 Tech Square, Cambridge, MA 02139.

†Technical Staff, Control and Dynamical Systems Division, 555 Tech Square.

‡Technical Staff, Control and Dynamical Systems Division, 555 Tech Square; currently Technical Staff, Harvard Smithsonian Center for Astrophysics, Cambridge, MA, 02138.

§Associate Director, Control and Dynamical Systems Division, 555 Tech Square.

The effect of model uncertainty on the performance of state estimators and failure detection algorithms is addressed in Refs. 2 and 10–13, among others. In Ref. 14 a geometric interpretation of the concept of analytical redundancy leads to a procedure involving singular-value decompositions for determining redundancy relations that are maximally insensitive to model uncertainties. An alternative approach is found in Ref. 15, where they assume that model errors may be deduced from the uncertainties of a set of underlying parameters. The partial derivatives of the residuals with respect to these parameters are then computed, and the residual generator with lowest partial sensitivity is selected. In Refs. 16–18, a bound on the effect of model uncertainties on the residual is estimated. This bound is then used to set the threshold accordingly. Robust detection methods based on the unknown input observer include those of Refs. 12, 19, and 20. Roughly speaking, an unknown input observer is a filter whose output is zero in the absence of failure regardless of the uncertainty and disturbance. The authors in Refs. 12 and 20 contributed significantly to the problem of robust detection using different techniques.

Previous work using H_∞ techniques include the work of Ref. 21, where a steady-state frequency domain-based filter design is used to attenuate the effect of disturbances. In Ref. 22 a steady-state H_∞/μ robust filter for failure detecion is introduced. This filter requires the solution of two Riccati equations and, as such, is robust to both disturbances and model uncertainty. As mentioned earlier, the use of transient filters for this application is motivated by the work in Ref. 9. In this paper we build upon the work in Ref. 22 by designing a filter architecture based on the transient, discrete-time, robust game theoretic or H_∞ filters derived in Refs. 2, 6, and 7 that can distinguish between failures in the RCS and the aerosurfaces. The failure detection architecture presented consists of two robust filters. One filter is tuned to estimate robustly the jet thrust even when jets misfire, whereas the other is tuned to detect failures in the aerosurfaces. Moreover, if a failure in the aerosurfaces is detected, then it is possible to reconfigure the jet thrust estimator to obtain an accurate estimate of the jet thrust. A one-filter architecture is also possible, but the two-filter architecture produced quicker detection.

In the next section we formulate the problem and demonstrate the deleterious effect of model uncertainty on the performance of the Kalman filter, motivating the use of robust filtering. In Sec. III we present a brief formulation of the robust filtering problem. The robust filter's equations, as well as the failure detection and isolation (FDI) architecture, are also presented. Results are presented in Sec. IV and conclusions in Sec. V.

II. Problem Description

The RCS consists of 44 bipropellant jets that, together with the aerosurfaces, provide attitude control and limited three-axis maneuvering capability. Thirty-eight of the jets are primary jets, each capable of providing 870 lb of thrust in vacuum. The jets are the sole attitude effectors for the initial part of reentry because there is insufficient dynamic pressure for the control surfaces to be effective in the thin atmosphere. As the vehicle loses altitude, falling into the increasingly denser atmosphere, dynamic pressure increases. The aerosurfaces are then activated and augment the jets until there is sufficient dynamic pressure to control the attitude with the aerosurfaces alone.

We consider the Space Shuttle Orbiter's lateral dynamics for bank and sideslip. A linear model of the rigid-body reentry rotational dynamics is used as derived by Zacharias.[23] Although originally given in continuous time, the dynamics are discretized for use with the Orbiter's digital computing system. The state-space representation for the Orbiter, linearized at values of the state and control for an operating point, is given by

$$x_{k+1} = A_k x_k + B_k u_k + G_k w_k + \mathcal{T}\theta_k$$

$$y_k = C_k x_k + D_k u_k + E_k v_k \tag{1}$$

where x_k is the state of the system at time k, u_k is the aerosurfaces control input, θ_k are the jet inputs, w_k is the process noise, y_k is the measurement (angular body rates) vector, and v_k is the sensor noise. In our case there are four plant states. The first two states are the vehicle attitudes: bank angle and sideslip angle, and the next two states are the corresponding rates. The measurements consist of vehicle body roll and yaw rates. These measurements are combinations of the bank and sideslip rate states, and the linearized equations at each flight regime relating these quantities are contained in the C_k matrix. As the Orbiter descends during reentry, the A_k, B_k, G_k, C_k, D_k, and E_k matrices change in accordance with changes in Mach number and angle of attack. In addition, a new state is defined for the jet thrust estimate. A simple, scalar, high-bandwidth Gauss–Markov model is used to represent a multiple-jet firing in a particular direction, i.e.,

$$\theta_{k+1} = a_\theta \theta_k + g_\theta v_k \tag{2}$$

The parameters a_θ and g_θ determine the bandwidth and amplitude of the model. For simplicity, we consider jets firing in only a single direction. Multiple Gauss–Markov models can be added to estimate thrust in multiple directions. The preceding model is augmented to the plant model of Eq. (1) to design the Kalman and robust filters. The interaction between the jets and the states is also modeled by augmenting the matrices A_k in the same equation. Consequently, the actual filter design contains five states. Only a single additional state is necessary to model multiple jets firing in the same direction. Multiple jets create scalar increases in thrust, easily modeled by a first-order Gauss–Markov process. This state, after being initialized to the number of jets commanded to fire, estimates the total thrust in a specific direction. Note that this additional state is not able to identify which jet has failed, but allows the comparison of the total number of jets firing in a specific direction with the total number commanded to fire, enabling a failed jet to be detected.

Our objective is to design a filter based on a time-invariant plant model, obtained by linearizing around a certain operating point, that performs well over a flight envelope, or a range of operating points. This makes gain scheduling at each time step unnecessary and provides robustness to uncertainties. As mentioned earlier, the filter is to robustly detect and isolate failures in the lateral dynamics attitude control system, i.e., to determine whether a failure is in the elevons or in the RCS. An additional goal is to obtain a robust estimate of the jet thrust even if a jet misfires or a failure in the aerosurfaces occurs. Correct isolation between jet and aerosurface failure is essential for accurate jet thrust estimation because without proper failure isolation irregularities in thrust due to aerosurface failure could be interpreted as off-nominal jet performance.

To motivate the need for robust filtering, a Kalman filter is designed based on the linear, time-invariant state-space model for the nominal condition of $M = 7.5$ and an angle of attack of $\alpha = 35$ deg. This filter is then tested on both the design plant and on a perturbed plant model based on a Mach number of $M = 8.8$ and $\alpha = 38$ deg. A step-on-step-off command with two jets firing in the same direction is used in this simple simulation. Aerosurfaces are used for trim adjustment in response to fluctuations in atmospheric conditions as modeled by the process noise.

Simulation results for the Kalman filter are shown in Fig. 1. The solid line represents the commanded jet thrust magnitude. The dashed lines represent a 15% error margin for the thrust estimate. The jet thrust estimates given by the Kalman filter are represented by the dash-dotted line for the design plant and a dotted line for the perturbed plant.

The estimates for the nominal plant and for the perturbed plant reveal that the Kalman filter works well for the correct plant model, but provides a severely degraded estimate when the model is perturbed. Note that the error in the degraded estimate for the perturbed plant is large enough to lead to the conclusion that more jets are fired than commanded. As mentioned in the introduction, it is desirable to avoid, or at least minimize, gain scheduling. Gain scheduling of a Kalman filter would be necessary because of the presence of model uncertainty and the rapid change in aerodynamic conditions during reentry, which, in the midreentry region of concern, move from Mach 10 at about 170,000-ft altitude to less than Mach 2 at about

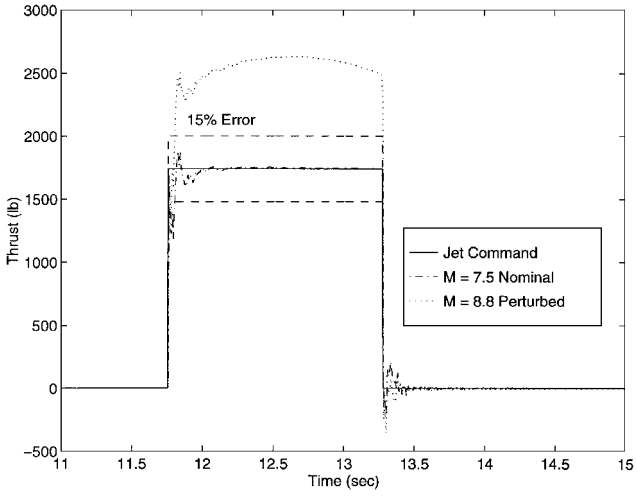

Fig. 1 Kalman filter's jet thrust estimates for a nominal and perturbed plant.

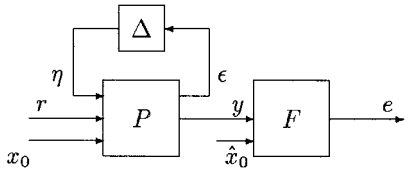

Fig. 2 General representation of the robust estimation problem.

75,000-ft altitude in approximately 20 min. New linear environment models that would need to accompany a set of gain-scheduled traditional Kalman filters would be necessary at an inordinately large number of linearization points throughout reentry because of the highly nonlinear nature of the flight regime during the reentry flight phase. The effectiveness of the Kalman filter would be reduced because of the time lost to multiple reinitializations of the filter and increased overhead in switching linearization points. In comparison, the time lost to the slight overhead of robust filters with a much less restrictive number of gain-scheduled linearization points is minimal.

In the next section we give a general formulation of the robust estimation problem and present the robust filter equations. These equations are derived in Ref. 2. The section also describes the filter architecture used for our robust failure detection and thrust estimation problem.

III. Robust Estimation

Figure 2 is a general input/output representation of a nominal plant $\boldsymbol{P}$ with modeling uncertainties Δ and an estimator $\boldsymbol{F}$. The vector r represents the combined process and measurement noise, x_0 the initial state vector of the nominal plant, $\hat{x}_0$ the initial state estimate, y the measurement vector, and e the estimation error. A known control input signal u may easily be accommodated in this formulation as part of the exogenous signal r. See Ref. 2 for details. The signals ϵ and η represent the interaction between the nominal plant and the perturbation. In what follows $\|r\|$ represents the 2-norm of the signal r over the time interval of interest. For instance, in discrete-time $\|r\| \equiv (\Sigma r_i' r_i)^{1/2}$ is the ℓ_2 norm of r. The norms of the other signals, namely η, ϵ, y, and e, are similarly defined. Weights can be added to these models, or alternatively the effect of the weights can be incorporated by varying the plant model parameters, to be described shortly. Finally, $\|x_0\|^2_{\bar{X}_0}$ and $\|x_0 - \hat{x}_0\|^2_{\bar{P}_0^{-1}}$ represent, respectively, the weighted Euclidean norms of the initial condition x_0 and the initial estimation error $x_0 - \hat{x}_0$.

The robust estimator seeks to bound the induced ℓ_2 norm of the operator from the input disturbances r and initial estimation error $(x_0 - \hat{x}_0)$ to the estimation error e, provided the model perturbations Δ have bounded induced ℓ_2 norm. Mathematically, this translates to the following performance criterion:

$$\sup_{(r,\, x_0 - \hat{x}_0) \neq 0} J_1 \equiv \frac{\|e\|^2}{\|r\|^2 + \|x_0 - \hat{x}_0\|^2_{\bar{P}_0^{-1}}} < 1$$

$$\forall \Delta \ni \|\Delta\|^2 \equiv \sup_{\epsilon \neq 0} \frac{\|\eta\|^2}{\|\epsilon\|^2} < 1 \tag{3}$$

The approach used to achieve the performance goal of Eq. (3) consists of treating η and x_0 as additional inputs to P and treating ϵ as an additional plant output. The revised objective is then to bound the induced norm of the mapping from the augmented input to the augmented output by 1. Thus, a new performance criterion is defined as

$$\sup_{(r,\eta,\, x_0 - \hat{x}_0) \neq 0} J_2 < 1 \tag{4}$$

where

$$J_2 \equiv \frac{\|e\|^2 + \|\epsilon\|^2}{\|r\|^2 + \|\eta\|^2 + \|x_0\|^2_{\bar{X}_0} + \|x_0 - \hat{x}_0\|^2_{\bar{P}_0^{-1}}}$$

Equation (4) is a robust performance or small gain condition. It is easy to show that satisfying this criterion guarantees the original performance condition of Eq. (3). It is in turn possible to achieve the condition of Eq. (4) by solving a minmax estimation problem whose objective function is given by

$$\min_{\hat{x}} \max_{r,\eta,\, x_0} J_3 \tag{5}$$

where

$$J_3 \equiv \|e\|^2 + \|\epsilon\|^2 - \gamma^2 \left(\|r\|^2 + \|\eta\|^2 + \|x_0 - \hat{x}_0\|^2_{\bar{P}_0^{-1}} \right)$$

subject to the plant's dynamic constraints. Specifically, if a solution to the preceding minmax problem is possible for a value of $\gamma \leq 1$, then $J_3 - \gamma^2 \|x_0\|^2_{\bar{X}_0} < 0$, which is equivalent to satisfying Eq. (4). The original criterion of Eq. (3) is then satisfied as well.

The preceding performance criterion is different from that of the Kalman filter's. Specifically, the robust filter attempts to minimize the ratio of the signal 2-norm of the augmented output to the signal 2-norm of the augmented input, while the Kalman filter minimizes the 2-norm of the estimation error squared. Furthermore, because of the relationship of the game-theoretic optimization to risk sensitivity, the difference between the Kalman filter and the robust filter has a stochastic interpretation as well (see Refs. 2 and 7 for details).

A. Robust Discrete Filter Equations

With $r_k = [w_k' v_k' \nu_k' u_k']'$ and $d_k = [r_k' \eta_k']'$, we write a time-varying nominal plant state-space representation as

$$x_{k+1} = A_k x_k + B_k d_k \tag{6a}$$

$$\epsilon_k = S_k x_k + T_k d_k \tag{6b}$$

$$e_k = M_k (x_k - \hat{x}_k) \tag{6c}$$

$$y_k = C_k x_k + D_k d_k \tag{6d}$$

The matrices S_k and T_k represent how the state and disturbance interact with the perturbation Δ and are a function of the uncertainties in A_k, B_k, C_k, and D_k (see for instance Ref. 2, pp. 61, 62). The error e_k is defined by a weighting matrix M_k that may be used to emphasize certain states over other states. We note that the state x_k in Eqs. (6a–6d) would include plant states, as well as any additional states such as the thrust θ_k in the preceding section, and possible shaping filter states for modeling disturbance processes.

The solution for the robust estimation problem of Eq. (5), as derived in Ref. 2, is then given as

$$\hat{x}_{k+1} = (\bar{A}_k - K_k \bar{C}_k)\hat{x}_k + K_k y_k \tag{7a}$$

$$K_k = \left(\bar{B}_k \bar{D}_k' + \bar{A}_k H_k^{-1} \bar{C}_k'\right) \left(\bar{D}_k \bar{D}_k' + \bar{C}_k H_k^{-1} \bar{C}_k'\right)^{-1} \tag{7b}$$

where

$$\bar{A}_k = A_k + \gamma^{-2} B_k Z_k^{-1} F_k' \tag{8}$$

$$\bar{B}_k = B_k Z_k^{-\frac{1}{2}} \tag{9}$$

$$\bar{C}_k = C_k + \gamma^{-2} D_k Z_k^{-1} F_k' \tag{10}$$

$$\bar{D}_k = D_k Z_k^{-\frac{1}{2}} \tag{11}$$

$$F_k = S_k' T_k + A_k' X_{k+1} B_k \tag{12}$$

$$H_k = P_k^{-1} - \gamma^{-2} M_k' M_k \tag{13}$$

$$Z_k = I - \gamma^{-2}(T_k' T_k + B_k' X_{k+1} B_k) \tag{14}$$

and the matrices X_k and P_k are, respectively, positive definite solutions to the following two Riccati equations:

$$X_k = A_k' X_{k+1} A_k + S_k' S_k + \gamma^{-2} F_k Z_k^{-1} F_k' \qquad X_N = 0 \tag{15}$$

$$P_{k+1} = (\bar{A}_k - K_k \bar{C}_k) H_k^{-1} (\bar{A}_k - K_k \bar{C}_k)'$$

$$+ (\bar{B}_k - K_k \bar{D}_k)(\bar{B}_k - K_k \bar{D}_k)' \qquad P_0 = \check{P}_0 \tag{16}$$

In the preceding equations the assumption is made that Z_k and H_k are positive definite and that $X_0 < \gamma^2 \bar{X}_0$.

The solution to this problem is an extension of both the H_∞ optimal estimator and the Kalman filter for nominal systems. If there are no model perturbations, then $S_k = T_k = 0$ in Eq. (6b), so that the Riccati equation (15) for X_k is superfluous, i.e., $X_k = 0$. In that case $\bar{X}_0 = 0$ by assumption. The estimator is reduced to solving one Riccati equation based on the nominal plant dynamics. The Riccati equation (16) and the gain (7b) are then the same as those of the Kalman filter except that the term $H_k = (P_k^{-1} - \gamma^{-2} M_k' M_k)$ replaces P_k^{-1}. In the case of the Kalman filter, P_k^{-1} is a measure of the information available at time k prior to taking a measurement. Subtracting $\gamma^{-2} M_k' M_k$ from P_k^{-1} means that we believe we have less information available. This filter is the standard game theoretic filter. Note that the smaller the choice of γ, the more conservative we are.

With $\Delta \equiv 0$, γ is no longer constrained to be less than unity. The minimum value of γ for which $H_k > 0, \forall k$ gives the solution to the optimal H_∞ problem. On the other hand, if $\gamma \to \infty$, then $H_k^{-1} \to P_k$ in both the Riccati equation and the optimal gain. In that case using any M_k, no identity becomes superfluous, and one recovers the Kalman filter. When $\Delta = \phi$, the minimax or game theoretic estimator can, therefore, be viewed as an extension of the Kalman filter and the H_∞ optimal estimator, where decreasing the design parameter γ trades off-nominal performance in the minimum variance sense to provide robustness to disturbance modeling error.

B. Robust FDI Filter Architecture

Figure 3 shows the filter architecture used for attitude control system fault detection. The output from the plant goes into two robust filters. One of the strengths of robust filter design is that various ways of representing the uncertainty and weighting matrices are possible. This flexibility can be used to tune in each of the two filters to a different set of states. This can be done, for example, by

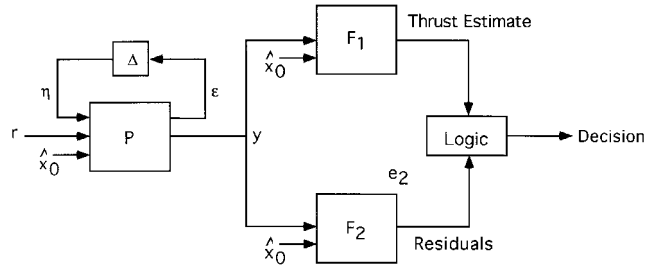

Fig. 3 Filter architecture.

varying the values of the weighting matrix M_k in Eq. (6c). Other techniques are also possible, such as those involving the matrices S_k and T_k in the preceding section. For more details, see for instance Ref. 5.

In the architecture of Fig. 3, one filter concentrates on the RCS, whereas the other concentrates on the aerosurfaces. Filter F_1 provides accurate thrust estimates. The filter is designed so as to robustly minimize the sum of squared error $\sum(\theta_k - \hat{\theta}_k)^2$ (Eq. 2). Specifically, the filter parameters are selected so that the induced 2-norm of the mapping between the disturbances as well as the possible aerosurface failures (modeled as potential inputs) on the one hand, and the estimation error on the other, is kept to a minimum over the entire range of plants under consideration, i.e., Mach 5.5 to 8.8. Over the same range, the mapping between a jet thrust input and the thrust estimate is kept as close to unity as possible.

The second robust filter F_2 is designed to detect failures in the aerosurfaces only and is therefore insensitive to perturbations in the jet thrust input. This filter is designed so as to minimize the output residual square error $\sum(y_k - \hat{y}_k)^2$. In this design parameters are chosen so that the induced 2-norm of the mapping from disturbances and thrust input to output residual error is as small as possible over the entire range of plants. On the other hand, the same mapping is highly responsive to aerosurface failures, modeled as inputs.

The details of the two filters' design can be seen in Ref. 25. An alternative design, where an aileron failure state is introduced using a Gauss–Markov model, is also possible (see Ref. 2). In that case, the objective would be to minimize that state's estimation error.

There is no systematic design procedure that can easily lead to the desired robust filter design. Any robust filter design is an iterative procedure that requires the analysis of different mappings.

The detection logic works as follows. Once an aileron failure is detected by the second filter, the first filter can be easily modified to provide accurate thrust estimates. This is done by adjusting the parameters so that the induced 2-norm of the mapping between the newly discovered failure and the thrust estimation error is kept as low as possible. If a jet failure occurs, filter F_1 can provide an accurate thrust estimate, thus allowing for jet failure detection.

Because we do not consider simultaneous failures for the purposes of this paper, using these two filters in parallel allows for the proper isolation of both kinds of failures in the Space Shuttle Orbiter's attitude control system, aerosurfaces and jets, and permits accurate thrust estimates even in the presence of aerosurface failures. Finally, a one-filter architecture is possible, but the detection response is slower.

A primary strength of robust filters, lessened performance sensitivity to specific off-nominal conditions, also contributes to slower estimation time response when compared with a traditional Kalman filter designed around the appropriate plant dynamics. Compare, for example, settling time after detecting jet thrust in a traditional Kalman filter design, Fig. 1, and a robust estimator, Fig. 4 in the next section. One method used in this robust FDI filter architecture to speed up the robust filter's response time is to reset periodically the Riccati matrix P to larger values. This causes the robust estimator to converge more quickly to a perturbed estimate than it otherwise would with mature, small values in the Riccati matrix. Because an off-nominal condition is by its nature a transient phenomenon, past information contained in the Riccati matrix is not useful, so little is lost by resetting the matrix in this manner. Naturally, a careful trade-off must be made between quicker response times to perturbations

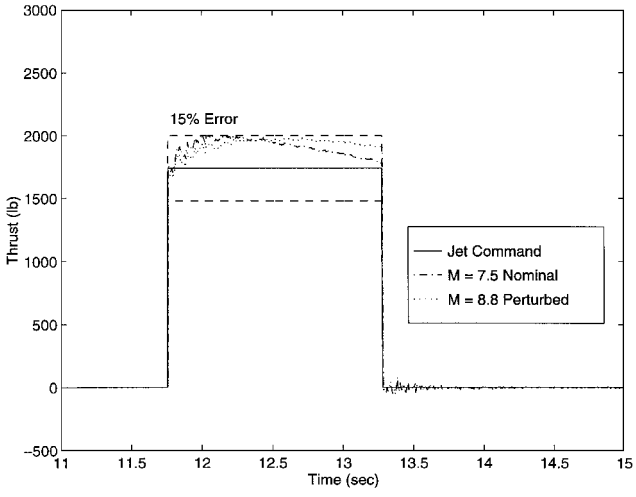

Fig. 4 Robust filter's F_1 jet thrust estimates for a nominal and perturbed plant.

with more frequent Riccati matrix resets and decreased estimation accuracy while the reset filter converges to a new solution.

In the next section we will demonstrate the use of this resetting technique for the robust filter and compare results with the same technique applied to the traditional Kalman filter. In cases where we have applied this technique, we will refer to the filter as a transient filter because resetting the Riccati matrix will cause the gains to be time-varying. Where we also examine the performance of the robust and Kalman filters without this technique, we will refer to the filter as a steady-state filter because after initialization the gains will have converged to a steady-state value and will effectively no longer be time-varying.

IV. Results

The robust FDI architecture just described is tested using the same simulation used to present the Kalman filter results in Sec. II. Nominal and perturbed systems were defined based on different flight regimes. Specifically, parametric perturbations are derived from two sets of linear time-invariant system matrices for different flight operating conditions (see Ref. 2, pp. 60–62), namely, $M = 7.5$, $\alpha = 35$ deg and $M = 8.8$, $\alpha = 38$ deg, where M is the Mach number and α is the angle of attack. Certain nonlinear effects such as jet self-impingement, variation in atmospheric pressure, and other transient effects are ignored. The operating points $M = 7.5$ and 8.8 were chosen because of the high degree of jet activity present between these two Mach numbers during reentry. Whereas good performance with the robust filters will be demonstrated at these two points, it is anticipated that some gain-scheduling of robust estimators may still be necessary over the entire reentry flight regime. Nonetheless, gain-scheduling requirements for robust estimators will be significantly less severe than gain-scheduling requirements for traditional Kalman filters.

Referring back to Eqs. (1), the process noise input matrix G_k is taken as $\frac{1}{20} \times$ (diagonal elements of B) to indicate a $\pm 5\%$ uncertainty in aerosurface deflections. A sampling interval of $T = 0.005$ s is chosen to achieve the minimum average estimation error in shortest convergence time. Whereas this time interval is shorter than the current Orbiter sample interval, it is realistic for the next generation RLV for which this FDI architecture is designed.

A. Jet Thrust Estimation

We will first discuss results for jet thrust estimation, which is given by filter F_1 in the architecture of Fig. 3. Simulation results comparing the performance of the Kalman and robust filters are shown in Fig. 1 (Sec. II) and Fig. 4. The filters are restarted at the start and end of the jet firing sequence, i.e., whenever the number of jets commanded changes. These plots show the effectiveness of the robust filter's insensitivity to model uncertainty. The dashed lines indicate the 15% error margin, and it is clear that the robust filter estimates are within that margin for both nominal (Mach 7.5) and perturbed (Mach 8.8) plants, whereas the Kalman filter estimates lie outside the error margin for the perturbed plant. Note also that the performance of the robust filter in Fig. 4 for either the nominal or perturbed plant compares well with that of the nominal, or optimal, Kalman filter shown in Fig. 1.

Comparable performance of both filters also indicates that the filters have similar bandwidths. Although the robust filter does, in fact, have a smaller bandwidth than the Kalman filter, this is not always the case. Generally speaking, the robust filter is designed to have a low gain in the region where uncertainty in the model dynamics has the greatest negative impact (see Ref. 22 for an example). Where it is possible to use this reduced gain approach in robust estimator design, the quick response time of the Kalman filter is preserved as much as possible. The ability to design the robust filter in this manner for this RLV application is aided by the accommodation of some gain scheduling of filters over the widely varying dynamics of the various reentry regions.

In this simulation jets were fired during three intervals over a 15-s simulation period. Figure 5 shows the robust filter's transient performance during the first 2 s of that simulation period. One of four jets commanded fails to fire at 0.35 s and remains off until 0.74 s. The spikes in the plots of Fig. 5 are caused by the filter's restarting. Overall, these results demonstrate clearly that while the transient robust filter is insensitive to model perturbation, it is highly sensitive to unexpected jet malfunction. Note also that the performance of the robust filter is maintained not only at the two flight regimes that where used in the filters design, Mach 7.5 and 8.8, but also at an intermediate regime that was not explicity used in the design of the filter. This demonstrates that the robust filter is usable over a range of flight regimes and is not just specifically tuned for favorable performance at only two points in the regime.

Results comparing the Kalman filter with the robust filter over the entire 15-s simulation period are presented in Table 1. Two different simulation runs are used: one where jets fire as commanded and the other with jet failures. The firing patterns commanded are the same

Table 1 Sum of squared error (10^7 lb^2) in jet estimate

Simulation run	Kalman filter	Robust filter
Unfailed jet (steady-state)		
Nominal plant ($M = 7.5$)	1.9282	2.3329
Perturbed plant ($M = 8.8$)	22.990	3.8438
Unfailed jet (transient filter)		
Nominal plant	5.0161	2.8658
Perturbed plant	23.718	5.0349
Failed jet (steady-state)		
Nominal plant	6.4095	5.5288
Perturbed plant	23.655	7.3224
Failed jet (transient filter)		
Nominal plant	6.8391	4.4167
Perturbed plant	21.799	6.7050

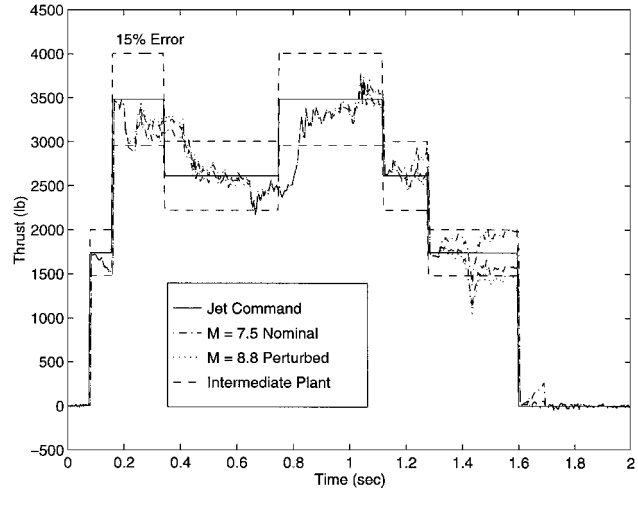

Fig. 5 Jet failure detection using a transient, robust filter F_1.

for both simulations and are shown in part in Figs. 1 and 5. The sum of squared error is the performance measure used for comparison. The term unfailed jet in the table refers to a pattern without jet failure, whereas the term failed jet refers to a firing pattern with jet failures, such as the one shown in Fig. 5. The steady-state filters are reset only when new jets are fired, whereas the transient filters are reset periodically. As described in Sec. III. B, this periodic resetting enhances the ability of the robust estimator to track perturbations quickly.

For all cases shown in Table 1, whether jets fail or not and whether a steady-state or transient filter is used, the Kalman filter estimates are severely degraded in the presence of plant perturbation. By comparison, the robust filter gives similar nominal and robust performance. What the robust filter gives up in nominal performance when compared to the Kalman filter, it more than recovers in robust performance.

Moreover, the transient robust filter outperforms the transient Kalman filter, even for the nominal plant. This is because the high-bandwidth Gauss–Markov model of Eq. (2) used to represent jet thrusts is only approximate, as no linear model can represent a step input. Table 1 also shows that for failure tracking, a transient filter is preferable. When a failure does not occur, however, then the steady-state filter gives slightly better performance because the transient filter is restarted periodically, increasing estimation error while it converges to a solution.

B. Failure Detection and Isolation

We now discuss the performance of the FDI architecture for failure isolation. The objective here is to distinguish between failures in the RCS and elevon control surfaces using both filters in the architecture according to the decision logic described in the preceding section. Figure 6 shows the output residuals for a Kalman filter designed to detect failures in the ailerons. The jet firing patterns is the same as that of the nominal 15-s simulation time. No failure occurs, but elevons are put to use at approximately 6 s. Results for the nominal Kalman filter on the nominal plant show that the residuals remain unchanged, i.e., the filter recognized the elevon command. However, the same filter, when used with the perturbed plant, shows a distinct shift in its residual, pointing to a failure that did not occur: a false alarm! Model mismatch and an actual elevon failure may therefore be difficult for the Kalman filter to distinguish in the presence of model uncertainty. Raising the threshold is not a satisfactory solution to the problem of avoiding false alarms because it will result in missed detections.

The robust filter residuals, on the other hand, remain essentially unchanged for both the nominal and perturbed plants, as shown in Fig. 7. This robust filter design is therefore insensitive to model uncertainty.

Figure 8 shows the results in the presence of failure when using the robust filter design. The elevon is failed on (stuck) for 3.5 s during the simulation period. The robust filter shows similar shifts

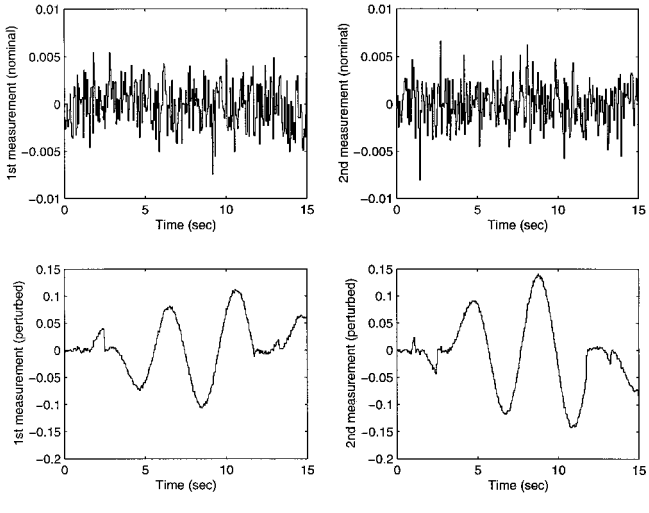

Fig. 6 Output residuals for the Kalman filter in the absence of failures for the nominal plant (above) and perturbed plant (below).

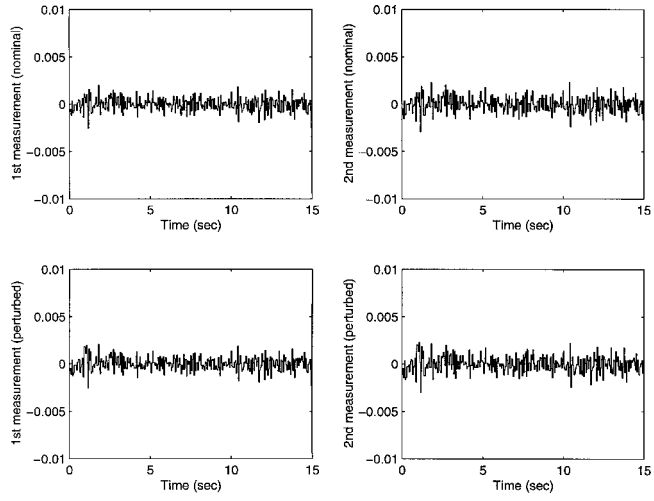

Fig. 7 Output residuals for the robust filter F_2 in the absence of failures for the nominal plant (above) and perturbed plant (below).

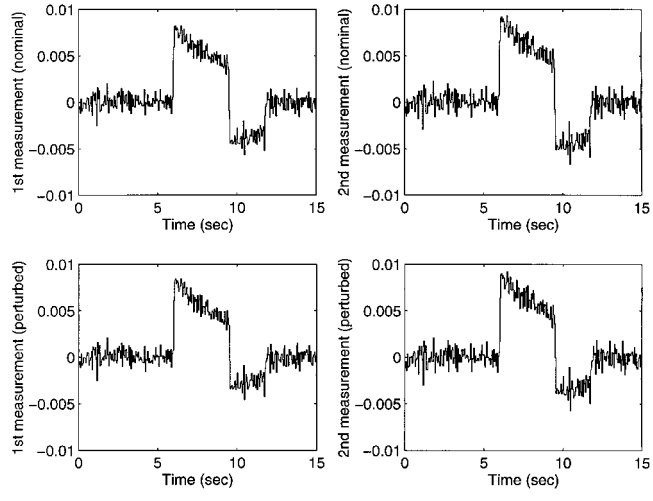

Fig. 8 Output residuals for the robust filter F_2 after an elevon failure is detected by the same filter.

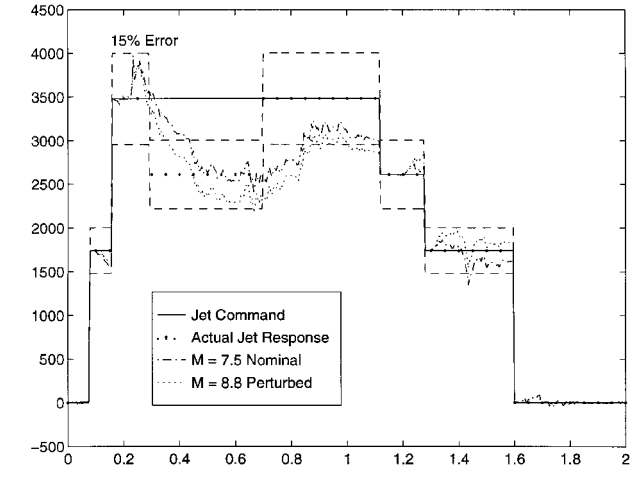

Fig. 9 Jet failure detection using filter F_1 in the presence of elevon failure.

in its residuals for both the nominal and perturbed plants. Figure 8 also shows that the same low threshold may be set for both plants, thus keeping the rate of false alarm low.

Figure 9 shows the jet estimate when an elevon fails on immediately after a jet failure occurs. Between 0.3 and 0.7 s one of four jets fails to fire. Between 0.9 and 1.2 s an elevon remains failed on at 1-deg deflection. The shift in the residual of the robust filter concerned with the ailerons F_2 allows for aileron failure isolation.

Subsequently, the filter concerned with jet thrust estimation F_1 is modified so as to compensate for any sensitivity to the elevon failure. This is accomplished by reparameterizing the filter. The elevon failure has been modeled as an exogenous input, so that the mapping between that input and the thrust estimate is simply reduced. As can be seen, adequate jet performance is maintained.

V. Conclusion

This article addresses the problem of failure detection and isolation for next-generation reentry vehicles' attitude control systems, using the Space Shuttle Orbiter as a model on which to design a prototype system. The objectives are to detect a failure rapidly and to determine accurately whether it occured in the jet RCS or the aerosurfaces. Correct isolation is essential to the second objective, namely, obtaining an accurate estimate of the thrust provided by the jets at all times. The problem is particularly difficult during reentry into the atmosphere because of both the rapid variations in Mach number and the large uncertainties in the vehicle aerodynamic properties.

Achieving the preceding objectives requires filtering algorithms that are, on the one hand, robust to model perturbations, and on the other hand, very sensitive to failures in either the jet firings or the aerosurfaces. A robust FDI architecture that relies on H_∞ filters has been suggested. Implementing this filter in simulation has shown that the robust algorithm performs very well over a wide range of model variations, whereas the nominally optimal Kalman filter can produce false alarms and leave some failures undetected, in the presence of the same variations. Future work would include looking in detail at the RCS jet propulsion model to isolate precisely which jet failed.

Acknowledgment

This research is supported by a Draper Laboratory Internal Research and Development Project. The authors would also like to thank the anonymous reviewers for comments that significantly enhanced the presentation of the work.

References

[1]Jacquemont, C. M., "Aircraft Attitude Determination Using Robust Estimation," M.S. Thesis, Dept. of Aeronautics and Astronautics, Massachusetts Inst. of Technology, Cambridge, MA, Aug. 1997.

[2]Mangoubi, R. S., *Robust Estimation and Failure Detection: A Concise Treatment*, Springer–Verlag, London, Aug. 1998.

[3]Appleby, B., "Robust Estimator Design Using the H_∞ Norm and μ Synthesis," Ph.D. Dissertation, Dept. of Aeronautics and Astronautics, Massachusetts Inst. of Technology, Cambridge, MA, Feb. 1990; also Draper Lab. Rept. T-1065, Cambridge, MA.

[4]De Souza, C., Shaked, U., and Fu, M., "Robust H_∞ Filtering with Parametric Uncertainty and Deterministic Input Signals," *Proceedings of the IEEE Conference on Decision and Control*, Vol. 2, Inst. of Electrical and Electronics Engineers, New York, Dec. 1992, pp. 2305–2310.

[5]Xie, L., De Souza, C., and Fu, M., "H_∞ Estimation for Discrete-Time Linear Uncertain Systems," *International Journal of Robust and Nonlinear Control*, Vol. 1, No. 2, 1992, pp. 111–123.

[6]Mangoubi, R., Appleby, B., and Verghese, G. "Robust Estimation for Discrete-Time Linear Systems," *Proceedings of the American Control Conference*, Vol. 1, American Conference on Automatic Control, Baltimore, MD, May 1994, pp. 656–661.

[7]Mangoubi, R., Appleby, B., and Verghese, G., "Stochastic Interpretation of H_∞ and Robust Estimation," *Proceedings of the IEEE Conference on Decision and Control*, Vol. 3, Inst. of Electrical and Electronics Engineers, New York, 1995, pp. 2377–2832.

[8]Mangoubi, R., Appleby, B., Verghese, G., and Vander Velde, W., "A Robust Failure Detection and Isolation Algorithm," *Proceedings of the IEEE Conference on Decision and Control*, Vol. 4, Inst. of Electrical and Electronics Engineers, New York, 1994, pp. 3943–3948.

[9]Rosello, A. D., "A Vehicle Health Monitoring System for the Space Shuttle Reaction Control System During Reentry," M.S. Thesis, Dept. of Aeronautics and Astronautics, Massachusetts Inst. of Technology, Cambridge, MA, May 1995.

[10]Gertler, J., "Survey of Model-Based Failure Detection and Isolation in Complex Plants," *Control System Magazine*, 1991, pp. 3–11.

[11]Emami-Naeini, A., Akhter, M., and Rock, S., "Effect of Model Uncertainty on Failure Detection," *IEEE Transactions on Automatic Control*, Vol. 33, No. 12, 1988, pp. 1106–1115.

[12]Patton, R. J., and Chen, J., "Techniques in Robust Fault Detection and Isolation (FDI) Systems," *Control and Dynamic Systems*, Vol. 74, Academic International Press, New York, 1996, pp. 171–190.

[13]Frank, P., and Ding, X., "Frequency Domain Approach and Threshold Selector for Robust Model-Based Fault Detection and Isolation," *Proceedings of the IFAC/IMACS Symposium on Fault Detection, Supervision and Safety for Technical Processes (SAFEPROCESS '91)*, Vol. 1, International Federation of Automatic Control/Society for Experimental Mechanics, Baden, Germany, 1991, pp. 299–306.

[14]Lou, X., Willsky, A., and Verghese, G., "Optimally Robust Redundancy Relations for Failure Detection in Uncertain System," *Automatica*, Vol. 22, No. 3, 1986, pp. 333–344.

[15]Gertler, J., and Singer, D., "A New Structural Framework for Parity Equation Based Failure Detection and Isolation," *Automatica*, Vol. 26, No. 2, 1990, pp. 381–388.

[16]Horak, D., "Failure Detection in Dynamic Systems with Modeling Errors," *Journal of Guidance, Control, and Dynamics*, Vol. 11, No. 6, 1988, pp. 508–516.

[17]Emami-Naeini, A., Akhter, M., and Rock, S., "Robust Detection, Isolation, and Accomodation for Sensor Failures," *Proceedings of the American Control Conference*, Vol. 2, 1985, Inst. of Electrical and Electronics Engineers, New York, pp. 1052–1059.

[18]Tsui, C., "A General Failure Detection, Isolation, and Accomodation System with Model Uncertainty and Measurement Noise," *IEEE Transactions on Automatic Control*, Vol. 39, No. 11, 1994, pp. 2318–2321.

[19]Saif, M., and Guan, Y., "A New Approach to Robust Fault Detection and Identification," *IEEE Transaction on Aerospace and Electronic Systems*, Vol. 29, No. 3, 1993, pp. 685–695.

[20]Frank, P., "Enhancement of Robustness in Observer-Based Fault Detection," *International Journal on Control*, Vol. 59, No. 4, 1994, pp. 955–981.

[21]Frank, P., and Ding, X., "Frequency Domain Approach to Optimally Robust Residual Generation and Evaluation for Model-Based Fault Diagnosis," *Automatica*, Vol. 30, No. 5, 1994,pp. 789–804.

[22]Mangoubi, R., Appleby, B., and Farrell, J., "Robust Estimation in Failure Detection," *Proceedings of the IEEE Conference on Decision and Control*, Vol. 2, Inst. of Electrical and Electronics Engineers, New York, 1992, pp. 2317–2322.

[23]Zacharias, G. L., "A Digital Autopilot for the Space Shuttle Vehicle," M.S. Thesis, Dept. of Aeronautics and Astronautics, Massachusetts Inst. of Technology, Cambridge, MA, Feb. 1974.

[24]De Souza, C., "Robust H_∞ Filtering," *Control and Dynamic Systems*, Vol. 65, Academic International Press, New York, 1994, pp. 323–377.

[25]Agustin, R. M., "Robust Estimation and Failure Detection for Reentry Vehicle Attitude Control Systems," M.S. Thesis, Dept. of Mechanical Engineering, Massachusetts Inst. of Technology, Cambridge, MA, June 1998.