Robust Fault Detection for Commercial Transport Air Data Probes

Paul Freeman, Peter Seiler, and Gary J. Balas *

* *Department of Aerospace Engineering and Mechanics, University of Minnesota, Minneapolis, MN 55455, USA (e-mail: {freeman, seiler, balas}@aem.umn.edu)*

Abstract: Air data probes provide essential sensing capabilities to aircraft. The loss or corruption of air data measurements due to sensor faults jeopardizes an aircraft and its passengers. To address such faults, sensor hardware redundancy is typically combined with a voting system to detect and discard erroneous measurements. This approach relies on redundancy, which may lead to unacceptable increases in system weight and cost. This paper presents an alternative, model-based approach to fault detection for a non-redundant air data system. The model-based fault detection strategy uses robust linear filtering methods to reject exogenous disturbances, e.g. wind, and provide robustness to model errors. The proposed algorithm is applied to NASA's Generic Transport Model aircraft with an air data system modeled based on manufacturer data provided by Goodrich Corporation. The fault detection filters are designed using linearized models at one flight condition. The detection performance is evaluated at a particular reference flight condition using linear analysis and nonlinear simulations.

Keywords: Fault detection, air data systems, robust estimation.

1. INTRODUCTION

Stringent safety requirements have driven aircraft system design for decades. The system availability and integrity requirements for commercial flight control electronics are typically no more than 10^{-9} catastrophic failures per flight hour (Bleeg, 1988; Collinson, 2003). The typical industry design solution is based extensively on physical redundancy at all levels of the design. For example, the Boeing 777 has 14 spoilers, 2 outboard ailerons, 2 flaperons, 2 elevators, one rudder and leading/trailing edge flaps (Yeh, 1996, 1998). Each of these surfaces is driven by two or more actuators, all connected to different hydraulic systems. Moreover, the control law software is implemented on three primary flight computing modules. Each computing module contains three dissimilar processors with control law software compiled using dissimilar compilers. The inertial and air data sensors have a similar level of redundancy.

The redundancy-dependant control architectures used in the aircraft industry achieve extraordinarily high levels of availability and integrity. However, the use of physical redundancy dramatically increases system size, complexity, weight, and power consumption. Moreover, such systems are extremely expensive in terms of design and development costs as well as the unit production costs. There is an increasing demand for high-integrity, yet low cost, fault tolerant aerospace systems, e.g. unmanned aerial vehicles and fly-by-wire in lower end business and general aviation aircraft. In such applications, analytical redundancy may be used to limit the number of sensors needed, but the ability to detect sensor failures may also be diminished.

This paper focuses on the use of analytical redundancy to detect faults in an air data system. Nearly all aircraft utilize air data probes to measure total and static pressure to determine airspeed and altitude. For proper operation, the probes must be free of blockages, e.g. due to icing or dirt. Failures of these probes have resulted in numerous fatal accidents of commercial, military, and general aviation aircraft. To address these failures, sensor hardware redundancy is typically combined with voting systems to detect and discard erroneous measurements.

The fault detection problem usually comprises a method to compute residuals and a process to declare faults based on the residuals. It is desired that the generated residual be a good representation of the fault of interest while being insensitive to process and measurement noises. Generation of residuals depends on the information available about the system. If a sufficiently accurate model of the system is available, model based methods can be used to estimate system states and outputs. See (Gertler, 1998), (Isermann, 2005), and (Ding, 2008b) for a detailed treatment of model based and model-free fault detection methods. Based on these methods, this paper uses the H_∞ framework to design an analytical fault detection filter for an air data system.

The paper has the following structure. Models for the system are provided in Section 2. Section 3 describes the H_∞ methods used to design the robust fault detection filter. Simulation results and analysis are given in Section 4. Finally, Section 5 discusses conclusions of this work and directions for future research.

2. MODELING

2.1 Generic Transport Model Longitudinal Dynamics

The NASA Generic Transport Model (GTM) is a remote-controlled 5.5 percent scale commercial aircraft (Murch and Foster, 2007). NASA developed a high fidelity six degree-of-freedom Simulink model of the GTM (Cox, 2010) with the aerodynamic coefficients described as look-up tables. The longitudinal model of the GTM aircraft dynamics is used for this study and is described in more detail in (Freeman et al., 2011).

The longitudinal dynamics of the GTM are described by a standard five-state longitudinal model where V is the air speed (knots), α is the angle of attack (deg), q is the pitch rate (deg/s) and θ is the pitch angle (deg), and h is the altitude (ft). The control inputs are the elevator deflection δ_{elev} (deg) and engine throttle δ_{th} (percent).

The GTM has two engines and equal thrust settings for both engines is assumed. The thrust from a single engine T (lbs) is a function of the throttle setting δ_{th} (percent).

The actuator dynamics are modeled as linear systems. The elevator actuator is a 5Hz bandwidth, first-order system with a 10 ms delay. The engine dynamics are modeled as a second order system.

2.2 Aircraft Trim and Model Linearization

A steady, level reference flight condition is chosen within the GTM flight envelope. The GTM, including actuator dynamics, is trimmed at the following condition:

$$\mathbf{x} = \begin{bmatrix} V \\ \alpha \\ q \\ \theta \\ h \end{bmatrix} = \begin{bmatrix} 75 \text{ knots} \\ 5.63 \text{ deg} \\ 0 \text{ deg/s} \\ 5.63 \text{ deg} \\ 500 \text{ ft} \end{bmatrix}; \quad \mathbf{u} = \begin{bmatrix} \delta_{th} \\ \delta_{elev} \end{bmatrix} = \begin{bmatrix} 33.098\ \% \\ 0.072 \text{ deg} \end{bmatrix} \tag{1}$$

The nonlinear GTM dynamics are linearized about this trim condition. The resulting linear model is used for control law development, filter synthesis, and simulation.

2.3 Control Law

The GTM longitudinal axis flight control law is an airspeed/altitude hold autopilot designed using a combination of classical loop-at-a-time and H_∞ techniques. The design allows for fault detection simulation and analysis while holding the GTM at a cruise condition or performing simple longitudinal maneuvers. These control laws serve merely to provide a closed-loop aircraft model that approximates the flight characteristics of a true aircraft for the purposes of simulation. The full control law interconnection is shown in Figure 1, and details regarding the controllers can be found in (Freeman et al., 2011).

2.4 Inertial and Air Data Sensors

Sensors for angle of attack, pitch rate, and pitch angle are modeled as unity with additive white noise and bias on the true states. Sensor dynamics are neglected in the model.

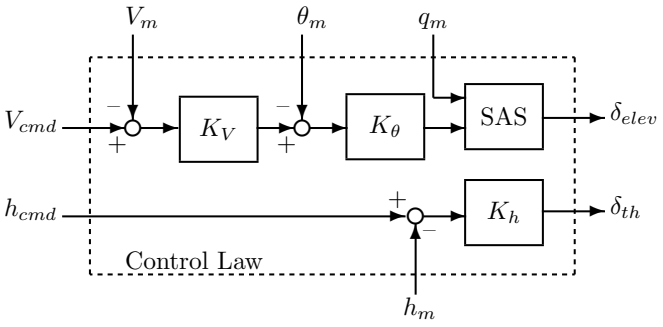

Fig. 1. Autopilot Control Law Architecture

The basic relationships between air data measurements and aircraft states are derived in (Collinson, 2003). For altitudes in the troposphere (up to $\approx$ 36,000 ft), the static pressure p_s is related to altitude h by:

$$h = \frac{T_0}{L}\left(1 - \left(\frac{p_s}{p_{s_0}}\right)^{LR/g}\right) \tag{2}$$

where $T_0 := 518.67^\circ\text{R}$ is the temperature at sea level, $L := 0.00356\frac{^\circ R}{ft}$ is the troposphere lapse rate, $g := 32.17\frac{ft}{s^2}$ is the gravity constant at sea level, $p_{s_0} := 2116.21\frac{lb}{ft^2}$ is the static pressure at sea level, and $R := 1716\frac{ft \cdot lb}{^\circ R \cdot slug}$ is the ideal gas constant.

For compressible air and subsonic speeds, the static and total pressures, p_s and p_t, are related to the calibrated (indicated) airspeed V_c (knots) by:

$$V_c = A_0\sqrt{5\left(\frac{p_t - p_s}{p_{s_0}} + 1\right)^{2/7} - 5}; \tag{3}$$

where $A_0 := 661.48$ knots is the speed of sound at sea level. The calibrated airspeed is equal to the true airspeed at sea level but the two airspeeds differ at altitudes above sea level. A more accurate model at high altitudes would include a model of the total air temperature sensor used to compute true airspeed (Goodrich, 2002). This is neglected in this paper and hence the models are only valid at low altitudes at which the GTM flies.

Pitot-Static Probe Model For typical commercial transport air data sensors, p_s and p_t are measured via independent pressure lines and transducers. Dynamic pressure, $p_{dyn} = p_t - p_s$, is a calculated quantity (Collinson, 2003; Goodrich, 2002). The pressure measurement devices are modeled by inverting the functions in Equations 2-3 to obtain values of static and total pressure from the GTM altitude and airspeed states. To model sensor noise and faults in the pressure measurements, the nominal pressure signals are corrupted by white noise and faults are added to the pressure signals to yield pressure measurements.

Pressure Measurement Processor Model Pressure transducer measurements are used to derive altitude and airspeed measurements for feedback to the control loops and to the pilot. Equations 2 and 3 yield altitude and airspeed measurements from pressure measurements. All pressure signals have units of psi (lbs/in^2) in the GTM simulation and analysis. The air data system architecture is depicted in Figure 2.

Linearizing the air data conversion equations about the trim condition in Equation 1 provides insight into appropriate magnitudes for injected faults:

$$\begin{bmatrix} \mathrm{d}h_m \\ \mathrm{d}V_m \end{bmatrix} = \begin{bmatrix} -1911 & 0 \\ -281.4 & 281.4 \end{bmatrix} \begin{bmatrix} \mathrm{d}p_s \\ \mathrm{d}p_t \end{bmatrix} \tag{4}$$

From Equation 4, faults of magnitude 0.01 psi in p_s will yield an h_m error of -19.11 ft and V_m error of -2.81 knots, respectively. A fault of the same size in p_t will yield an V_m error of 2.81 knots. Additionally, faults injected on p_s and p_t both influence V_m. For simultaneous, equal magnitude faults in p_s and p_t, V_m will be unaffected.

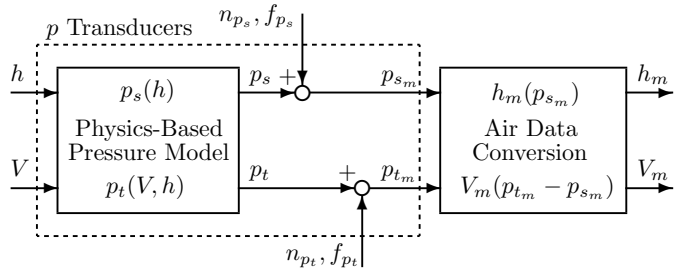

Fig. 2. Air Data Sensor Architecture

3. ROBUST FAULT DETECTION

The H_∞ synthesis framework used to design filters to estimate disturbances, e.g. faults, at the plant input. H_∞ methods offer a number of advantages over traditional Kalman filtering, including superior performance in the presence of model uncertainty and the ability to filter process noise and exogenous disturbances without necessarily having a statistical model of those inputs (Simon, 2006).

A H_∞ filter is synthesized from the linearized GTM model to estimate faults associated with the static and total pressure measurements. Unmodeled dynamics and model uncertainty are neglected for this analysis, but the framework is easily extended to include uncertainty.

3.1 H_∞ Problem Formulation

The linearized aircraft dynamics are connected with the autopilot (Figure 1), inertial sensor models, and air data sensor architecture (Figure 2). The generalized plant *genGTM* has the following inputs: the autopilot reference signals $\tilde{\mathbf{r}} = [V_{cmd}\ h_{cmd}]^\intercal$, the inertial measurement noises, $\tilde{\mathbf{n}} = [n_\alpha\ n_q\ n_\theta]^\intercal$, and the injected pitot faults $\tilde{\mathbf{f}} = [f_{p_s}\ f_{p_t}]^\intercal$. The errors $\tilde{\mathbf{e}}$ are the difference between the injected faults and estimated faults $\hat{\mathbf{f}} = [\hat{f}_{p_s}\ \hat{f}_{p_t}]^\intercal$. The generalized GTM plant has measurement outputs $\mathbf{y} = [p_{s_m}\ p_{t_m}\ \alpha_m\ q_m\ \theta_m]^\intercal$. These are the measurements that will be available to the fault detection filter.

The objective of the H_∞ filter synthesis is to generate a stable filter F which minimizes norm between the disturbances and the errors. Weighting functions are used to describe the expected frequency content of the inputs and the desired frequency content of the errors, the normalized inputs $[\mathbf{r}\ \mathbf{n}\ \mathbf{f}]^\intercal$ and outputs $\mathbf{e}$. Figure 3 shows the desired interconnection of the filter with the generalized plant *genGTM* with signal weights and filter F. Input and output signals with tildes represent their respective normalized signals in physical units.

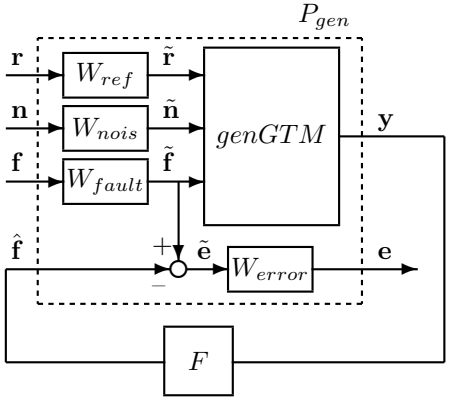

Fig. 3. Interconnection for H_∞ Filter Synthesis

For fault detection, the disturbances are the autopilot reference signals $\mathbf{r}$ and the inertial measurement noises $\mathbf{n}$. The filter seeks to track the injected faults $\mathbf{f}$ with the fault estimates $\hat{\mathbf{f}}$ while rejecting inertial measurement noise $\mathbf{n}$ and reference commands $\mathbf{r}$. Similar H_2 and H_∞ model matching approaches to FDI filter design have been applied in (Varga, 2003, 2009; Marcos et al., 2005; Ding, 2008a; Zolghadri et al., 2006; Mazarsand et al., 2008).

3.2 Signal Weighting Methodology

W_{ref} is a static weight that defines relative size of the autopilot reference commands:

$$W_{ref} = \begin{bmatrix} 1 & 0 \\ 0 & 6.667 \end{bmatrix} \tag{5}$$

W_{nois} reflects the assumption that the inertial sensor noise magnitude is greater at low frequencies and high frequencies while the noise magnitude is reduced at intermediate frequencies. Such behavior can be found in mid-grade inertial sensors that may be utilized on a small UAV. This leads to an unusual choice of weighting function; a simpler model may not account for increased noise magnitudes at low frequency, and would likely have a first-order weighting function. The weight chosen to represent the frequency content associated with these signals is:

$$W_{nois} = \frac{6.119s^2 + 30.76s + 0.1376}{s^2 + 1758s + 0.07861} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 10 & 0 \\ 0 & 0 & 1 \end{bmatrix} \tag{6}$$

The weighting on the inertial sensor noise is chosen iteratively so that the transfer function from $\tilde{\mathbf{n}}$ to $\hat{\mathbf{f}}$ on the unweighted *genGTM* interconnection with F has gain less than 1 for all frequencies. The weighting on the n_q signal is an order of magnitude larger than the other weightings due to the higher noise level on the pitch rate sensor output. Near 10^{-4} rad/s, W_{nois} rolls down 40 dB; it then rolls up approximately 50 dB near 1 rad/s. The transfer functions from inertial sensor noise to pressure fault estimates show stronger attenuation in the $\hat{f}_{p_s}$ channel; this fact is important for analysis of the filter performance.

W_{fault} Faults are injected into the static and total pressure channels to corrupt the measurements. The fault weight, Equation 7, is chosen such that the DC gain represents large faults (-20 dB). The weight is small for frequencies greater than 5 rad/s to penalize tracking of high frequency faults. The aircraft dynamics roll off near

this frequency; hence, the effect of higher frequency faults will not appear in the angle or rate measurements. This diminishes the tracking ability of F at high frequencies. The high frequency gain of W_{fault} is -60 dB.

$$W_{fault} = \frac{0.001s + 3.562}{s + 35.27} \, I_2 \tag{7}$$

W_{error} The error weight represents the inverse of the allowable tracking error at each frequency. Normally, the error weight would be large at low frequency to ensure close tracking. Tracking at high frequency is less desirable and error weightings will roll off to some small high frequency gain. The particular nature of this problem, however, is such that the usual error weighting methodology cannot be adopted for the generalized GTM filter synthesis. Equation 8 shows the DC gain of $P_{\tilde{\mathbf{f}}\tilde{\mathbf{y}}}$, the partition of $genGTM$ from $\tilde{\mathbf{f}}$ to $\tilde{\mathbf{y}}$.

$$P_{\tilde{\mathbf{f}}\tilde{\mathbf{y}}}|_{\omega=0} = \begin{bmatrix} -0.0370 & 0.0370 \\ -0.0370 & 0.0370 \\ -49.984 & 49.984 \\ 0 & 0 \\ -49.984 & 49.984 \end{bmatrix} \tag{8}$$

Note that the matrix representing the DC gain is rank deficient. Thus, faults in the direction $\mathbf{f} = [1\ 1]^\intercal$ are indistinguishable from an unfaulted conditon. The unobservability of this fault direction at DC has a simple physical explanation. As mentioned previously, a simultaneous and equal fault in both pitot probes has no effect on the airspeed measurement. A fault in the $\mathbf{f} = [1\ 1]^\intercal$ direction only causes a bias in the altitude measurement. The model for the longitudinal dynamics is unaffected by a constant offset in altitude. Thus, a fault in the $\mathbf{f} = [1\ 1]^\intercal$ direction will cause the closed loop system to adjust to a biased value of altitude but all measurements will appear, in steady state, to converge back to their original trim conditions.

This rank deficiency places limits on the fault detection performance at low frequencies. For a filter F to ensure perfect fault tracking at low frequency, F must be a pseudoinverse of $P_{\tilde{\mathbf{f}}\tilde{\mathbf{y}}}$ over that frequency range. In particular, a filter F that would make the tracking error arbitarily small at low frequency cannot be synthesized by `hinfsyn` (Balas et al., 2010) because the partition is rank-deficient and its pseudoinverse does not exist.

To circumvent this problem, W_{error} is chosen such that the DC gain is small (-40 dB) and begins to roll up at 10^{-5} rad/s to -12 dB at 10^{-2} rad/s. For frequencies greater than 10^{-2}, the traditional approach of rolling off to a small high frequency gain (-60 dB) is applied. This error weighting has a small DC gain, rolls up at very low frequencies, and rolls down again at higher frequencies; this allows for the best filter performance given the inherent system limitations.

$$W_{error} = \frac{0.0011s^2 + 0.1106s + 4.606 \times 10^{-5}}{s^2 + 0.4683s + 4.606 \times 10^{-3}} \, I_2 \tag{9}$$

3.3 FDI Filters

The weighted interconnection shown in Figure 3 is generalized into the weighted generalized plant P_{gen} (Doyle et al., 1988). The filter F is synthesized with a γ-value of 0.0046 using P_{gen} and `hinfsyn` to meet the objectives described in Sections 3.1-3.2. The `hinfsyn` algorithm synthesizes a filter at the low γ-value for a few reasons. First, the small weight choices scale γ to be small. Next, model uncertainty is not considered in this formulation, allowing for stronger filter attenuation of disturbances.

4. ANALYSIS OF ROBUST FAULT DETECTION

The fault detection performance of the filter F synthesized in Section 3 is examined for three fault scenarios. First, fault detection performance is analyzed for a step faults in the static pressure measurement of the closed-loop linear generalized GTM model, $genGTM$. Detection performance is analyzed in the presence of inertial sensor noise. Similarly, fault detection performance is evaluated for a step fault in the total pressure measurement of $genGTM$ while the static pressure measurement remains unfaulted. Finally, the performance of the synthesized filter is examined for a simultaneous fault in p_s and p_t in the closed-loop nonlinear GTM.

The filter can take advantage of the expected closed-loop system dynamics to generate fault estimates since it was synthesized using a closed-loop generalized plant, These estimates would be more accurate than estimates resulting from the common open-loop synthesis approach that fails to model the dynamics associated with the expected operation of a system in the field.

4.1 Typical Faults and Desired Filter Performance

Numerous common fault scenarios exist for typical air data systems, including blockages, icing, and water accumulation. Particular combinations of blockages in the pitot inlet, drain hole, static port, or pressure lines can lead to significant errors in altitude and/or airspeed measurements. Based on manufacturer data provided by Goodrich Corporation, a simple model for certain air data blockages can be obtained by injecting step faults into the static and total pressure measurements. The performance of the synthesized filter under these scenarios is examined.

The filter should yield estimates that track the generalized fault inputs reasonably quickly with minimal steady-state error. False positives are very undesirable. Any fault detection system implemented with the goal of control reconfiguration must be sufficiently fast as to allow for reconfiguration before undesired aircraft maneuvers become unsafe.

4.2 Step Fault Detection in Linear GTM

The fault tracking performance of F for a small step fault in the static pressure measurement is shown in Figure 4. At time $t = 1$ second, a step fault of magnitude 0.01 psi is injected into the static pressure measurement signal for the linear generalized plant $genGTM$ using the filter F. The simulation has a duration of 4 seconds and inertial and

pitot sensor noise is included. The filter outputs $\hat{\mathbf{f}}$ should track the fault in the faulted static pressure channel and show no fault in the unfaulted total pressure channel.

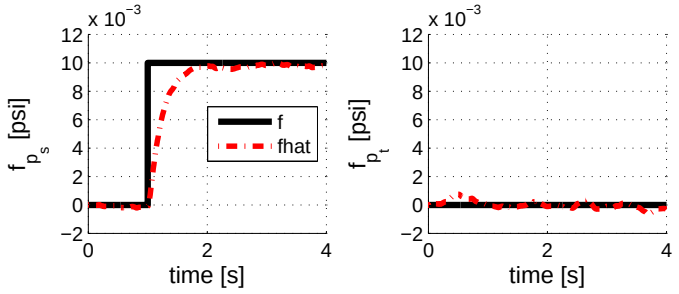

Fig. 4. Fault estimation: p_s step, Linear GTM

The filter detects the static pressure fault in the linear model nearly instantaneously and tracks it effectively. The filter does not yield errorless fault tracking, however, due to the *genGTM* rank deficiency described in Section 3.2. The slowest pole of F has a frequency on the order of 10^{-5} rad/s, so the fault estimation error will grow quite slowly. The fault estimate will eventually decay to zero in the faulted channel and drift away from zero in the unfaulted channel. Measures to combat this drift must be designed into any algorithm that can be implemented on an operational system.

The fault estimate in the unfaulted total pressure fault channel contains considerable noise relative to the static pressure fault estimate. Because noise on the inertial sensors is fed into the filter, the resulting fault estimates will be noisy. As stated in Section 3.2, inertial sensor noise couples to $\hat{f}_{p_t}$ more strongly than $\hat{f}_{p_s}$, accounting for the higher noise levels in the total pressure estimate. Since the inertial measurements are fed into the airspeed hold autopilot, the filter relies on these measurements to detect the presence of a fault in the airspeed measurement much more than for the altitude measurement. Consequently, the noise levels in $\hat{f}_{p_t}$ are larger. While the H_∞ filter is designed to minimize the effect of sensor noise on the fault estimates, it cannot be eliminated entirely.

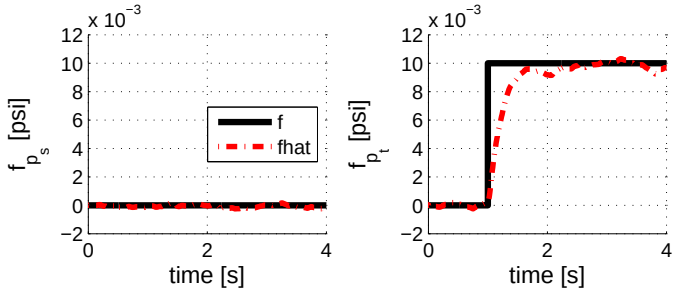

Fig. 5. Fault estimation: p_t step, Linear GTM

Next, the same simulation is conducted with a total pressure fault rather than a static pressure fault. Figure 5 shows that the filter detects the total pressure fault in the linear model and tracks the fault with some residual noise for the reasons explained above. The filter correctly estimates no fault on the static pressure measurement. For long simulations, as with the static fault scenario, the estimates will eventually drift.

4.3 Simultaneous Fault Detection

A fault of equal magnitude injected simultaneously into the static and total pressure channels will not have an impact on the airspeed measurement. Because such a compound fault has limited observability, it is interesting to examine the ability of F to detect such a condition. Given the desirable performance of F on the linear GTM, the performance of F for this more difficult fault will be investigated on the nonlinear longitudinal GTM.

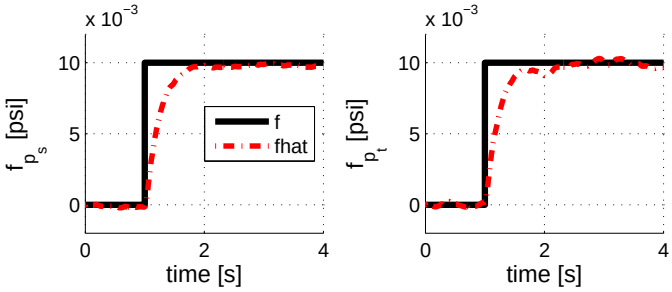

Fig. 6. Fault estimation: Simultaneous p_s and p_t step, Nonlinear GTM

Figure 6 shows the fault estimates for the simultaneous fault in the nonlinear GTM. Figure 7 shows the control inputs, aircraft state responses, and static pressure fault estimate throughout the nonlinear GTM simulation. The simultaneous fault results in a bias in the altitude measurement. All of the control inputs and aircraft states– except the altitude measurement–converge back to the original trim condition. The only effect of the simultaneous fault is that the aircraft converges to an offset altitude. Despite the simultaneous fault that does not appear in the airspeed measurement at zero frequency, the filter is able to detect the initial step in both measurements as faults. The filter uses the inertial state measurements to track the fault by compensating for the dynamic response of the aircraft to the step changes in the measurements. The fault estimates eventually decay to zero due to the unobservability of this fault in steady state.

5. CONCLUSIONS

This paper developed a method to detect faults in a pitot-static probe using H_∞ synthesis of a robust fault detection filter. Signal weights were chosen to circumvent the unobservable fault at low frequency. The performance of the fault detection filter was analyzed on the linear and nonlinear longitudinal GTM dynamics for individual step faults as well as a simultaneous, equal magnitude fault in each pressure measurement. This approach was shown to be effective for detecting step faults and combinations of step faults. By extending modeling to uncertainty in the GTM aerodynamic model and actuator dynamics, a robust fault detection filter can be synthesized using μ-synthesis techniques. Additionally, extending robust fault detection filtering to an aircraft's entire flight envelope using linear parameter-varying techniques would provide a stronger understanding of fault tolerance implications at all parts of the flight regime.

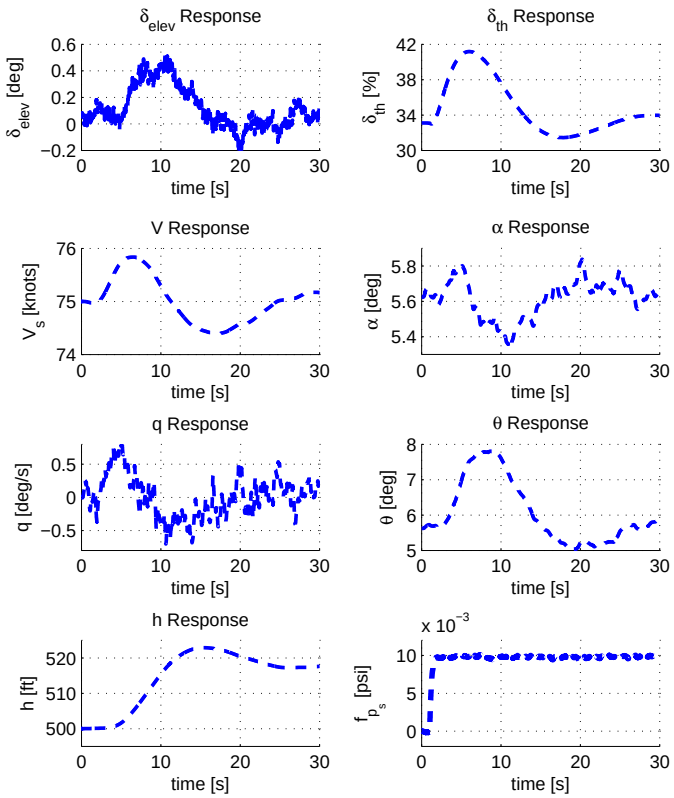

Fig. 7. Control Input and State Response: Simultaneous p_s and p_t step, Nonlinear GTM

ACKNOWLEDGEMENTS

This material is based upon work supported by the National Science Foundation under Grant No. 0931931 entitled CPS: Embedded Fault Detection for Low-Cost, Safety-Critical Systems. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation. The authors acknowledge Andrei Dorobantu for the design of the inner-loop SAS controller.

REFERENCES

Balas, G., Chiang, R., Packard, A., and Safonov, M. (2010). *Robust Control Toolbox*. The MathWorks.

Bleeg, R. (1988). Commercial jet transport fly-by-wire architecture considerations. In *AIAA/IEEE Digital Avionics Systems Conference*, 399–406.

Collinson, R. (2003). *Introduction to Avionics Systems*. Kluwer, 2nd edition.

Cox, D. (2010). *The GTM DesignSim v1006*.

Ding, S.X. (2008a). *Model-based Fault Diagnosis Techniques*. Springer.

Ding, S.X. (2008b). *Model-based Fault Diagnosis Techniques: Design Schemes, Algorithms, and Tools*. Springer, 1st edition.

Doyle, J., Glover, K., Khargonekar, P., and Francis, B. (1988). State-space solutions to standard h-2 and h-infinity control problems. *IEEE Transactions on Automatic Control*, 34(8), 831–847.

Freeman, P., Seiler, P., and Balas, G. (2011). Robust fault detection for commercial transport air data probes: Extended version with modeling. *http://www.aem.umn.edu/ AerospaceControl/CPS.html*.

Gertler, J.J. (1998). *Fault detection and diagnosis in engineering systems*. Marcel Dekker, 1st edition.

Goodrich (2002). *Air Data Handbook*. Goodrich Corportation – Sensors and Integrated Systems.

Isermann, R. (2005). *Fault-diagnosis systems: an introduction from fault detection to fault tolerance*. Springer, 1st edition.

Marcos, A., Ganguli, S., and Balas, G. (2005). An application of H-infinity fault detection and isolation to a transport aircraft. *Control Engineering Practice*, 13(1), 105–119.

Mazarsand, E., Jaimoukha, I., and Li, Z. (2008). Computation of a reference model for robust fault detection and isolation residual generation. *Journal of Control Science and Engineering*, 2008, 1–12.

Murch, A. and Foster, J. (2007). Recent NASA research on aerodynamic modeling of post-stall and spin dynamics of large transport airplanes. In *45th AIAA Aerospace Sciences Meeting and Exhibit, Reno, Nevada*, AIAA-2007-463.

Simon, D. (2006). *Optimal State Estimation: Kalman, H-infinity, and Nonlinear Approaches*. John Wiley & Sons.

Varga, A. (2003). On computing least order fault detectors using rational nullspace bases. In *In Proceedings of the IFAC Symp. SAFEPROCESS'2003, Washington D.C.*

Varga, A. (2009). General computational approach for optimal fault detection. In *SAFEPROCESS '09, Barcelona, Spain*.

Yeh, Y. (1996). Triple-triple redundant 777 primary flight computer. In *Proceedings of the 1996 IEEE Aerospace Applications Conference*, 293–307.

Yeh, Y. (1998). Design considerations in Boeing 777 fly-by-wire computers. In *Third IEEE International High-Assurance Systems Engineering Symposium*, 64–72.

Zolghadri, A., Castang, F., and Henry, D. (2006). Design of robust fault detection filters for multivariable feedback systems. *International journal of modelling and simulation, ACTA Press*, 26(1).