

Fault Detection and Isolation for Autonomous Parafoils

Matthew R. Stoeckle¹, Amer Fejzic², Louis S. Breger³, and Jonathan P. How⁴

^{1,4} *Massachusetts Institute of Technology, Cambridge, MA, 02139, USA*
stoke927@mit.edu
jhow@mit.edu

^{2,3} *Draper Laboratory, Cambridge, MA, 02459, USA*
afejzic@draper.com
lbreger@draper.com

ABSTRACT

Autonomous precision airdrop systems are widely used to deliver supplies to remote locations. Payloads that are delivered far from their intended targets or with high impact velocity may be rendered unusable. Faults occurring during flight can severely degrade vehicle performance, effectively nullifying the value of the guided system, or worse. Quickly detecting and identifying faults enables the choice of an appropriate recovery strategy, potentially mitigating the consequences of an out-of-control vehicle and recovering performance. This paper presents a multi-observer, multi-residual fault detection and isolation (FDI) method for an autonomous parafoil system. The detection and isolation processes use residual signals generated from observers and other system models. Statistical methods are applied to evaluate these residuals and determine whether a fault has occurred, given a priori knowledge of system uncertainty characteristics. Several examples are used to illustrate the detection and isolation algorithm online using available navigation and telemetry outputs. Tests of this FDI method on a large number of high-fidelity simulations indicate that it is possible to detect and isolate some common faults with a high percentage of success and a very small chance of raising a false alarm.

1. INTRODUCTION

Autonomous precision airdrop is used to deliver payloads to areas that would be dangerous or difficult to reach through more conventional means. Missions for guided parafoils include military resupply of troops and humanitarian efforts (Hattis & Tavan, 2007). As described in Hattis, Campbell, Carter, McConley, and Tavan (2006), the goal of the system is to land the payload as close as possible to the target

while minimizing ground speed at impact. Flight testing has shown that a variety of faults can occur (Tavan, 2006). These faults increase target miss distances and landing speeds, potentially rendering payloads unusable. In addition, the possibility of an in-flight fault and resulting behavior could preclude delivering supplies to more densely populated areas where an out-of-control vehicle could pose a danger to persons or property. Detecting, isolating, and responding to faults can improve performance and expand the space of missions available for guided parafoils. This work designs and implements a Fault Detection and Isolation (FDI) strategy that is effective in the unique conditions under which the parafoil operates.

Online systems for FDI fall into two categories: those that exploit hardware redundancy and those that rely on analytical redundancy (Hwang, Kim, Kim, & Seah, 2010). Systems with a large number of sensors, actuators, and measurements employ hardware redundancy for FDI or system health management (Figueroa et al., 2009) (Figueroa, Schmalzel, Morris, Turowski, & Franzl, 2010). The parafoil has a minimal number of sensors, and so analytical redundancy methods are used.

Isermann and Ballé (1997) define FDI terminology. A fault is defined as an unpermitted deviation of at least one characteristic property or parameter of the system from the acceptable/usual/standard condition. Fault detection is the determination of the faults present in a system and the time of detection. Fault isolation is the determination of the kind, location, and time of detection of a fault. The process of isolation follows that of detection.

For FDI to be effective, 1) the effects of faults must be distinguishable from the effects of unknown inputs including modeling errors, disturbances, and measurement uncertainty, and 2) faults must be distinguishable from each other (Frank, 1994). This is typically accomplished by considering a resid-

Matthew Stoeckle et al. This is an open-access article distributed under the terms of the Creative Commons Attribution 3.0 United States License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

ual signal (Hwang et al., 2010). The residual signal chosen has approximately zero mean when no fault is present and nonzero mean when a fault has occurred. In this context, a residual signal is the difference between a measurable system output and the corresponding expected output. After the residual has been generated, it is evaluated. The goal of the evaluation process is to determine whether a fault alarm should be raised based on the properties of the residual signal.

A large group of FDI methods are classified as observer-based. These methods use an observer of the nominal system to generate the expected system output. This output is used along with measurements from the actual system to generate the residual signal. Though a simulation of the system with no feedback can also be used to generate the residual signal, an observer is chosen to make the residual generation process robust to differences in initial conditions.

A common method of residual generation that uses observers is called the fault detection filter (FDF). This method generates a residual signal that is projected onto subspaces associated with various faults, so that detection and isolation are both possible (Beard, 1971) (Jones, 1973). See Douglas and Speyer (1995) for a robust implementation of the FDF. For isolation, the FDF requires that each fault under consideration acts on the system in a known, unique way. This is not the case for the parafoil system; many faults act on the control lines and are not distinguishable from each other.

The eigenstructure assignment approach is used to de-couple effects of disturbances from those of faults by nulling the transfer function from the disturbances to the residual signal (Patton & Chen, 2000). A weighting matrix that is used to assign eigenvectors to the closed-loop observer of the system accomplishes this task. In order to construct this weighting matrix, however, there must be more independent outputs of the system than independent disturbances (Patton & Chen, 2000). The parafoil system is a single-output system, so eigenstructure assignment is not possible.

The FDI method presented in this paper is observer-based, but takes a different approach than the FDF. Many existing observer-based methods incorporate isolation into the detection process by exploiting the system property that each fault under consideration is distinguishable from all other faults (Frank, 1994). However, this is not the case for many faults that occur on the parafoil system. As a result, the detection and isolation processes are separate for this work.

For detection, a residual signal is generated using observer-based methods. This residual is evaluated using hypothesis testing. If the magnitude of the residual signal crosses above a predetermined detection threshold, a fault is declared. Sargent et al. (2011) use hypothesis testing with thresholds for FDI on the Orbital Cygnus vehicle. Rossi (2012) uses hypothesis testing for health management of spacecraft.

If a fault is declared, isolation is performed. In this paper, isolation is broken into two phases. The first phase uses a residual signal from actuator data. If evaluation of this signal indicates that an actuator fault has occurred, isolation is complete. However, if the first phase of isolation does not declare an actuator fault, phase two begins.

Phase two of isolation uses a bank of fault-specific observers to differentiate between non-actuator faults. The purpose of these observers is to determine when the system exhibits characteristics of a particular fault (Willsky, 1976). Evaluating residual signals from these observers indicates if a specific non-actuator fault is present. Successful isolation will result in the declaration of a fault on one of the actuators or the declaration of a particular non-actuator fault.

Section 2 gives an overview of how the parafoil and payload system operates. Section 3 describes common faults that have been observed in flight tests. Sections 4 and 5 describe the detection and isolation methods, respectively. Section 6 combines detection and isolation into the full FDI algorithm. Section 7 concludes the paper.

2. PARAFOL SYSTEM OVERVIEW

A typical system consists of a canopy, airborne guidance unit (AGU), and payload (Figure 1). An example of a parafoil system, as well as some performance characteristics, is given in Bergeron, Fejzic, and Tavan (2011).

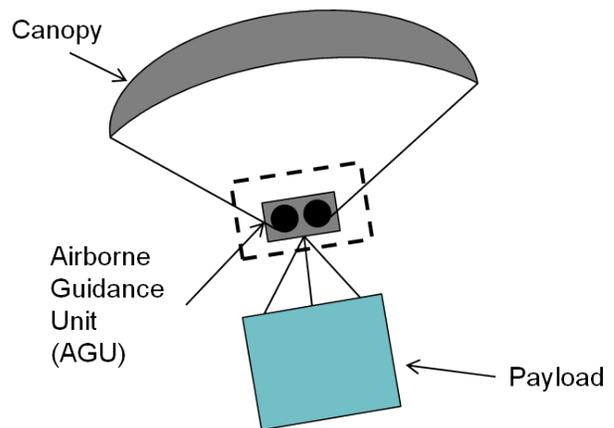


Figure 1. Parafoil canopy, AGU, and payload.

Lateral control of the parafoil is accomplished using two control lines that attach to the left and right trailing edges of the canopy. These lines are wound around two motors on the AGU (Figure 2). Details of guidance, navigation, and control (GN&C) implementation on the parafoil system are described by Carter, George, Hattis, Singh, and Tavan (2005).

The motors retract and extend the control lines, deflecting the trailing edges of the canopy and inducing a nonzero turn rate. The motors on the AGU are equipped with encoders that mea-

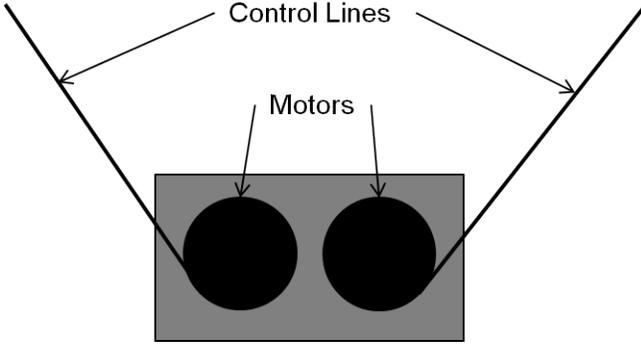


Figure 2. Parafoil AGU and control lines.

sure the deflections of the control lines. The resulting performance of the parafoil is estimated using this information. A payload is attached underneath the AGU by several support lines.

Different dynamical models of varying complexity have been developed for the parafoil and payload system (Barrows, 2009) (Ward, Montalvo, & Costello, 2010). The observer that will be used for FDI is based on a linearized model of the lateral dynamics, derived from nonlinear dynamics in Crimi (1990). Lateral dynamics were chosen because they accurately reflect the heading rate response of the system while allowing for a straightforward linear model. A similar linearized model of lateral parafoil dynamics is used in Slegers and Costello (2004) for model predictive control, and is applicable to observer-based FDI as well. The model used in this paper is described by the linear, time-invariant (LTI) system,

$$\dot{x}(t) = (A + \Delta A)x(t) + (B + \Delta B)u(t) + Wd(t) \quad (1)$$

$$y(t) = Cx(t) + \eta(t) \quad (2)$$

where $A, \Delta A \in \mathbb{R}^{4 \times 4}$, $B, \Delta B \in \mathbb{R}^{4 \times 2}$, $C \in \mathbb{R}^{1 \times 4}$, and $W \in \mathbb{R}^4$. A , B , and C are known dynamics, control, and output matrices, respectively. The matrix B can be written as $[b_1 \quad -b_1]$, where $b_1 \in \mathbb{R}^4$ (i.e., both motors affect the system equally, but in opposite directions). The matrix W determines how the process noise acts on the system states. The matrices ΔA and ΔB represent unknown modeling errors. The matrices A and B are determined from known system parameters. The example parafoil used for simulation in this work has the following parameters: canopy weight 70 lbs, canopy area 900 ft², canopy span 50 ft, canopy chord length 16 ft, and nominal payload weight 1800 lbs. These parameters are among many used to determine the linearized lateral dynamics.

The states of the LTI system in Eqs. (1) and (2) are $x(t) = [\beta(t) \quad \dot{\phi}(t) \quad \phi(t) \quad \dot{\psi}(t)]^T$, where $\beta(t)$ is the sideslip an-

gle, $\dot{\phi}(t)$ is the roll rate, $\phi(t)$ is the roll angle, and $\dot{\psi}(t)$ is the yaw rate. The control input is $u(t) = [\delta_R(t) \quad \delta_L(t)]^T$, where $\delta_R(t)$ and $\delta_L(t)$ are the right and left motor deflections, or motor toggles, respectively. Under healthy conditions, the deflection at each motor will match the corresponding deflection of the control line. This will not be the case when some faults occur. The effects of these faults on the system dynamics are discussed in Section 3. The disturbance term $d(t) \in \mathbb{R}$ is the process noise of the system. The chosen output $y(t) \in \mathbb{R}$ is the heading rate of the system, which is subject to uncertainty that is captured in the navigation error term $\eta(t) \in \mathbb{R}$. Heading rate in the context of the parafoil is defined as the rate at which the airspeed velocity vector of the parafoil rotates with respect to the inertial North axis (see Figure 3).

The effects of the uncertainty and noise terms, ΔA , ΔB , $d(t)$, and $\eta(t)$, on the FDI process cause the residual signals to be nonzero even when no fault has occurred. However, the size of the residual during a healthy flight is small compared to the size of the signal when a fault has occurred. In other words, faults are still observable even if the noise and uncertainty terms are neglected. Therefore, neglecting these four terms and formulating the FDI problem using an observer-based approach as opposed to a Kalman filter-based approach was chosen to minimize the computational complexity.

After ignoring these terms, the lateral system dynamics reduce to,

$$\dot{x}(t) = Ax(t) + Bu(t) \quad (3)$$

$$y(t) = Cx(t) \quad (4)$$

This simpler linearized lateral dynamics model will be used for residual generation.

The parafoil system has no sensors for measuring heading rate directly; instead, this quantity is estimated using an Extended Kalman Filter (EKF). The only state information that the parafoil software has access to is position and translational velocity data from the onboard GPS. The GPS measures the ground speed as well as the sink rate of the parafoil. An EKF is used to estimate the wind velocity, and from this information the airspeed velocity and the heading rate are estimated, similarly to work done in Ward et al. (2010). The system states used in Eqs. (1-4) are not available from the EKF and are unknown.

Figure 3 shows some parafoil states and provides insight into the estimation of heading rate. The GPS measures the ground velocity V_g . The EKF estimates the wind velocity V_w . The airspeed V_a is estimated using vector addition, and from that an estimate of the heading angle χ is obtained. The EKF uses this information to generate an estimate of the heading rate $\dot{\chi}$.

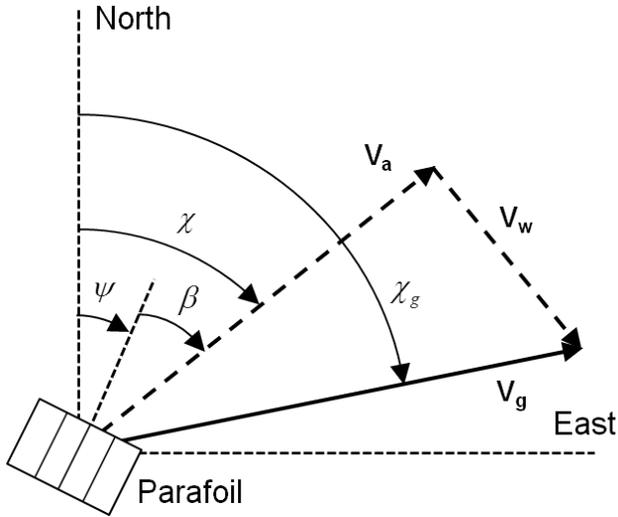


Figure 3. Parafoil states and reference frame.

3. COMMON FAULTS

Flight testing of the parafoil system has shown that some faults occur frequently (Tavan, 2006). A hierarchy of common faults is shown in Figure 4. Four of the faults shown in the hierarchy are considered for FDI. These faults are chosen because they have effects that are both well-defined and well-understood. The faults are: stuck motor, severe saturation, broken control line, and criss-crossed control lines.

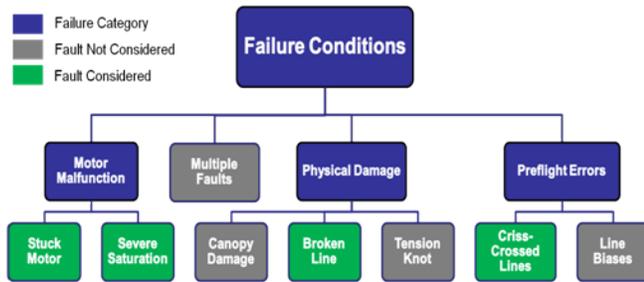


Figure 4. Fault hierarchy.

A stuck motor fault occurs when one of the motors on the AGU fails to respond to commands. When this fault occurs, the line attached to the faulty motor is stuck at an unknown position while the other control line is free to move as commanded. If detection fails, motor commands will be given that the motor will not be able to respond to.

Under nominal conditions, a known motor saturation limit exists. This limit is the maximum line deflection possible for each motor. A severe saturation fault occurs when the maximum possible line deflection is significantly less than expected; a saturation limit of 0.65 meters is considered significant for this work. If this fault is not detected, it is possible

that a large command will be given but the corresponding deflection will not occur. At that point, the system will likely continue to command a large deflection to achieve a desired heading rate that it can never attain.

A broken line fault occurs when one of the control lines that are attached to the motors on the AGU breaks. In this case, the motor is still free to turn, but there is no corresponding response in line deflection. This prevents the parafoil from turning in the direction of the side on which the line is broken. This fault often occurs upon canopy deployment.

It is possible, while rigging the lines to the AGU, that the control line attached to the left trailing edge of the parafoil is spooled around the right motor, and vice versa. In this case, a command to the right motor will yield a deflection in the left control line, and a command to the left motor will yield a deflection in the right control line. This fault is called criss-crossed lines. This is an example of a fault that has a straightforward recovery strategy. No change to the existing guidance strategy is necessary; the controller need only reverse the commands given to each motor to achieve the desired performance. However, this recovery approach cannot be implemented unless FDI successfully detects and isolates the fault.

The non-actuator faults are added to the linearized lateral model of the parafoil dynamics in Eq. (3) as follows,

$$\dot{x}(t) = Ax(t) + (B + \Delta B_f)u(t) \quad (5)$$

where $\Delta B_f \in \mathbb{R}^{4 \times 2}$ are changes to the dynamics that occur when either a broken control line or criss-crossed lines fault is present. When a broken left line occurs, $\Delta B_f = \begin{bmatrix} 0 & b_1 \end{bmatrix}$. When a broken right line occurs, $\Delta B_f = \begin{bmatrix} -b_1 & 0 \end{bmatrix}$. When a criss-crossed lines fault occurs, $\Delta B_f = \begin{bmatrix} -2b_1 & 2b_1 \end{bmatrix}$. When any of these non-actuator faults occur, the deflections of the left and right control lines will not match the motor toggles in $u(t)$.

Actuator faults enter the system in a different way. The motor toggles $u(t)$ result from passing commanded toggles $u_{cmd}(t)$ into the motors on the AGU. Therefore, actuator faults (i.e. stuck motor and severe saturation) are not modeled in Eq. (5) but instead manifest themselves in a value of $u(t)$ that is different from what is expected. Actuator faults are identified by comparing $u(t)$ to $u_{nom}(t)$, the nominal, or expected, motor toggle. This term is introduced in Section 4.

4. DETECTION

Fault detection is the process of determining the faults present in a system and the time at which those faults occurred (Isermann & Ballé, 1997). A detection alarm is raised for any fault; it is not necessary during detection to know which particular fault has occurred. Detection is accom-

plished by comparing a known system output with an expected system output. The difference between these two quantities is a residual signal. This signal should be chosen such that it is large when a fault is present and small otherwise (Frank, 1994). If the residual signal is large, a fault is declared. This paper uses an observer-based fault detection method, where the expected system output is generated using an observer. This observer is designed to model the parafoil and payload system when no faults are present.

Detection is broken into two phases: residual generation and residual evaluation (Hwang et al., 2010). Residual generation is the process of constructing the residual signal. Residual evaluation is the process of taking this signal and using it to either validate or reject a null hypothesis. The null hypothesis is that the system is healthy; a rejection of this hypothesis indicates a fault. Residual evaluation is performed using a threshold, which is designed so that if the residual signal rises above this threshold there is a reasonable probability that a fault is present (Frank, 1994),

$$\text{If } r(t) \leq \lambda_{th}, \text{ null hypothesis confirmed} \quad (6)$$

$$\text{If } r(t) > \lambda_{th}, \text{ null hypothesis rejected; fault} \quad (7)$$

where $r(t) \in \mathbb{R}$ is a time-varying residual signal, and $\lambda_{th} \in \mathbb{R}$ is a mission-specific constant threshold value. The FDI method in this work uses the parafoil heading rate output for residual generation (see Section 2). The parafoil guidance system commands the parafoil by specifying a desired heading rate. If the parafoil is not tracking the heading rate as expected, the system is likely in a faulty condition.

4.1. Residual Generation

To generate the heading rate residual signal, an observer is constructed based on the linearized lateral dynamics described in Eqs. (3) and (4). Figure 5 shows that a motor toggle command u_{cmd} is passed through both the AGU motor, which is subject to actuator faults, and a model of a healthy motor.

The output of the AGU motor is the actual motor toggle u ; the output of the motor model is the nominal motor toggle u_{nom} . The difference between these two signals is small when no actuator faults are present. The nominal input u_{nom} is used as the input to the observer, which is constructed as follows,

$$\dot{\hat{x}}(t) = A\hat{x}(t) + Bu_{nom}(t) + L(y(t) - \hat{y}(t)) \quad (8)$$

$$\hat{y}(t) = C\hat{x}(t) \quad (9)$$

where $\hat{x}(t) \in \mathbb{R}^4$ is an estimate of the system states $x(t)$, $\hat{y}(t) \in \mathbb{R}$ is the observer estimate of the heading rate, and A , B , and C are the matrices from the parafoil dynamics in Eqs. (3) and (4). The feedback gain $L \in \mathbb{R}^4$ is designed to

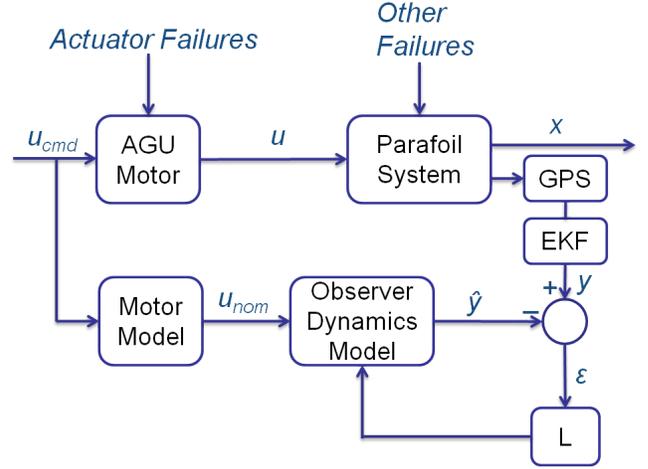


Figure 5. Heading rate observer block diagram.

make $A - LC$ stable. Error terms are defined as,

$$e(t) = \hat{x}(t) - x(t) \quad (10)$$

$$\dot{e}(t) = \dot{\hat{x}}(t) - \dot{x}(t) \quad (11)$$

$$\varepsilon(t) = \hat{y}(t) - y(t) = Ce(t) \quad (12)$$

and the residual signal used for fault detection is chosen as,

$$r(t) = \varepsilon^2(t) \quad (13)$$

Squaring $\varepsilon(t)$ ensures that the residual signal is non-negative. This aids the residual evaluation process.

Error dynamics of the observer/plant system can yield insight into the behavior of the residual signal. Though the error $e(t)$ is not measurable, Eq. (12) shows that $\varepsilon(t)$ is a function of $e(t)$. By substituting Eqs. (5) and (8) into Eq. (11), the error dynamics are shown as,

$$\dot{e}(t) = (A - LC)e(t) + B(u_{nom}(t) - u(t)) - \Delta B_f u(t) \quad (14)$$

Equation (14) shows that there are several instances where $\dot{e}(t)$ can become nonzero. The first term, $(A - LC)e(t)$, will decay to zero exponentially for a stabilizing L . The gain L can always be chosen to stabilize $A - LC$ if (A, C) is observable (Van de Vegte, 1994). The second term in Eq. (14) will be nonzero when $u_{nom}(t)$ is not equal to $u(t)$. There are two expected sources of error between $u_{nom}(t)$ and $u(t)$. The first is modeling errors between the motor model and the actual motor. The second is an actuator fault, where $u(t)$ is not behaving as expected. If the motor model accurately models behavior of the actual motor, only an actuator fault will cause a noticeable increase in the magnitude of that term, and thus

a significant increase in $\dot{e}(t)$. The third term will be nonzero when a broken line or criss-crossed lines fault has occurred.

4.2. Residual Evaluation

A threshold for detection is the main tool used in this work for residual evaluation. This threshold is chosen such that there is a high probability of a fault when the residual is above the threshold and a low probability of a fault when the residual is below the threshold. Statistical methods are used for threshold determination. To improve detection statistics, the residual signal at each time step is smoothed over the previous 25 seconds of flight using a moving average. The time period over which smoothing occurs can be varied according to design needs. A longer period better emphasizes the trend of the signal while filtering out noise, but will cause a lag between the occurrence of a fault and the response of the signal. This parameter was tuned numerically to achieve desired detection characteristics.

When designing a threshold, the goal is to minimize two quantities: probability of missed detection $P(MD)$ and probability of false alarm $P(FA)$ (Rossi, Breger, Benson, Sargent, & Fesq, 2012) (Sturza, 1988). The probability of missed detection is the probability that a fault has occurred and no fault alarm is raised; the probability of false alarm is the probability that an alarm is raised when no fault has occurred. These quantities are predicted by collecting data from simulations of healthy flights and flights in which a fault has occurred. Cumulative density functions (CDFs) of data from simulated healthy flights and flights in which a fault occurs are useful in visualizing how a chosen threshold affects $P(FA)$ and $P(MD)$.

In order to collect the data used in the CDFs, a high-fidelity, nonlinear, 6 degree-of-freedom (DOF) simulation was used, similar to the model described in Ward et al. (2010). Each simulation is generated with different initial conditions. These conditions include not only three-dimensional position, orientation, and velocity, but environmental conditions such as wind profile and system irregularities. System irregularities are variations from parafoil-to-parafoil, turn bias and lift-to-drag ratio for example, that change the flight characteristics. The linearized lateral dynamics used for the observer initialize the state vector to zero, and do not account for wind or differences between each individual parafoil system.

To generate CDFs that will accurately show $P(FA)$ and $P(MD)$ for various thresholds, large data sets were collected that reflect the range of conditions a parafoil system experiences during a healthy flight as well as flights in which faults of varying type and severity have occurred. For the healthy data set, 1000 Monte Carlo simulations were performed that varied the following characteristics: initial position, initial altitude, initial velocity, initial attitude, payload weight, wind conditions, and turn bias (i.e., nonzero heading rate in the

presence of zero control). No fault occurred during any of these simulations. The data of interest from each Monte Carlo run is the maximum value that the smoothed residual signal reaches during the simulation. This maximum bounds the smoothed residual signal expected from healthy flights.

Data collection from the fault cases was treated differently than data collection from healthy flight simulations. Again, 1000 Monte Carlo simulations were performed, but for these simulations random faults were inserted. The nature of each fault was chosen with uniform probability from the four fault cases discussed in Section 3. The time of the fault was chosen with uniform probability to be an integer value between 0 and 500 seconds. This choice ensured sufficient time for detection, because most simulated flights last longer than 600 seconds. The severity of the fault, if applicable, was randomized. For example, the broken line fault severity need only be randomized as a left or right line break, but a stuck motor fault occurs on either the left or right motor and has a particular value (e.g., 0.5 meters) at which the motor is stuck. Multiple faults were not considered.

The relevant data collected from these flights is the maximum value reached by the smoothed residual signal during the first 60 seconds after the fault occurs. In flight, there is a window of time after a fault occurs at which point recovery from the fault is either impossible or impractical. Thus, the detection method is only given a predetermined amount of time to raise an alarm. This time period is a design parameter that should be set based on mission requirements and recovery techniques in use. Using the collected data, CDFs were generated (Figure 6) and the performance of various thresholds were analyzed. The green circles shown on Figure 6 mark $P(FA)$ and $P(MD)$ on the CDF. $P(MD)$ is the intersection of the threshold line with the fault data curve. The intersection of the threshold line with the healthy curve is $1 - P(FA)$.

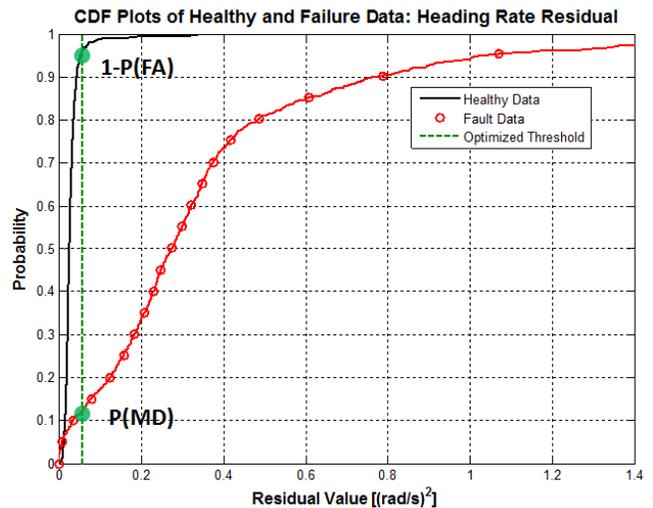


Figure 6. CDFs for heading rate residual.

The selection of an appropriate threshold given the data in Figure 6 depends on the emphasis placed on minimizing $P(FA)$ versus $P(MD)$. The following figure of merit (FOM) is used to penalize $P(FA)$ and $P(MD)$ as desired (Rossi, Benson, Sargent, & Breger, 2012):

$$FOM = 1 - \frac{c_1 P(FA) + c_2 P(MD)}{c_1 + c_2} \quad (15)$$

The weightings c_1 and c_2 can be varied according to design needs, where a higher weighting on either $P(FA)$ or $P(MD)$ indicates that it is more important to minimize that particular probability. Then, a threshold is determined that maximizes the chosen FOM. Table 1 shows thresholds, $P(FA)$, and $P(MD)$ for different weightings. The threshold shown on Figure 6, 0.0574 (rad/s)^2 , as a vertical dotted line was chosen by placing an equal weighting on $P(FA)$ and $P(MD)$.

Table 1. Optimized thresholds for various FOMs.

c_1	c_2	Threshold $(\frac{\text{rad}}{\text{sec}})^2$	$P(FA)$	$P(MD)$	FOM
1.0	0.0	0.3995	0.000	0.733	1.000
0.8	0.2	0.0995	0.012	0.172	0.956
0.6	0.4	0.0607	0.032	0.127	0.927
0.5	0.5	0.0574	0.038	0.119	0.921
0.4	0.6	0.0574	0.038	0.119	0.913
0.2	0.8	0.0573	0.038	0.119	0.897
0.0	1.0	0.0017	0.999	0.017	0.983

The results presented are from simulations and not flight data. However, the same procedure can be applied to actual systems. The detection threshold can be adjusted as needed to achieve the desired performance given anticipated increases in process noise and navigation error.

4.3. Detection Results

Performance evaluation of the detection method consists of comparing predicted $P(MD)$ and $P(FA)$ (Table 1) with the corresponding probabilities resulting from the implementation of the detection method in simulation. The fault detection method presented in this paper was tested on 1000 Monte Carlo simulations of flights with randomized conditions. Each flight was chosen with equal probability to have no fault, a stuck motor fault, severe saturation, a broken line, or criss-crossed lines. The severity of each fault was randomized where applicable. Table 2 shows results from these simulations as well as the predicted values of $P(FA)$ and $P(MD)$ from Table 1.

Figure 7 plots $P(MD)$ versus $P(FA)$ for the data summarized in Table 2. The closer the data are to the origin, the better the performance (Rossi, Benson, et al., 2012). This plot

Table 2. Comparison of $P(MD)$ and $P(FA)$ between predicted and simulated results for detection.

Threshold $(\frac{\text{rad}}{\text{sec}})^2$	$P(FA)$ Predicted	$P(FA)$ Sim	$P(MD)$ Predicted	$P(MD)$ Sim
0.3995	0.000	0.000	0.733	0.594
0.0995	0.012	0.005	0.172	0.201
0.0607	0.032	0.051	0.127	0.154
0.0574	0.038	0.058	0.119	0.149
0.0573	0.038	0.058	0.119	0.149
0.0017	0.999	0.796	0.017	0.064

also indicates an important point about fault detection: there is always a tradeoff between false alarm and missed detection (Rossi, 2012). These quantities are minimized according to design criteria by maximizing the figure of merit in Eq. (15).

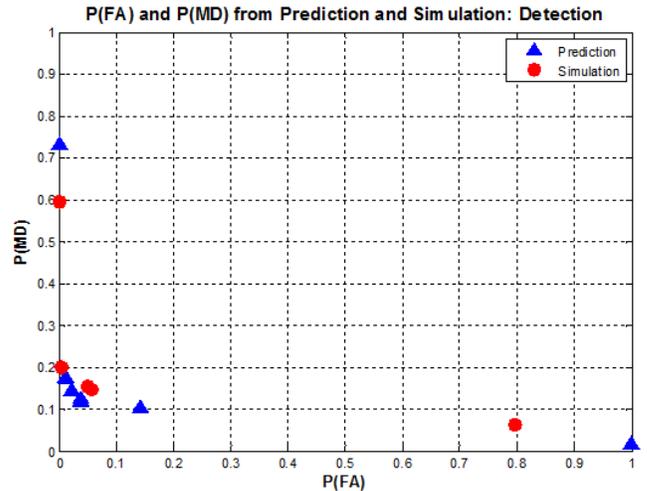


Figure 7. Comparison of predicted and simulated false alarm and missed detection rates.

4.4. Detection Example

Figure 8 shows an example of the detection process on a simulated flight. The fault, a stuck right motor at 0.863 meters, occurs 126 seconds into the flight. The top subplot of Figure 8 indicates nominal tracking of motor commands up until the time of the fault. Once this occurs, there is a large difference between the command and the response in the right motor. This difference is reflected in the smoothed residual signal, which is shown in the bottom subplot. The threshold chosen for detection during this flight is 0.0574 (rad/s)^2 , which is the optimized value when equal weighting is placed on $P(FA)$ and $P(MD)$ (Table 1). The smoothed residual is below the threshold, but not zero, before the fault. However, once the fault occurs, the difference between $u(t)$ and

$u_{nom}(t)$ is large enough to cause the smoothed residual signal to cross the chosen threshold. When the residual crosses the threshold at 134 seconds a detection alarm is raised. After the alarm is raised, the isolation procedure begins.

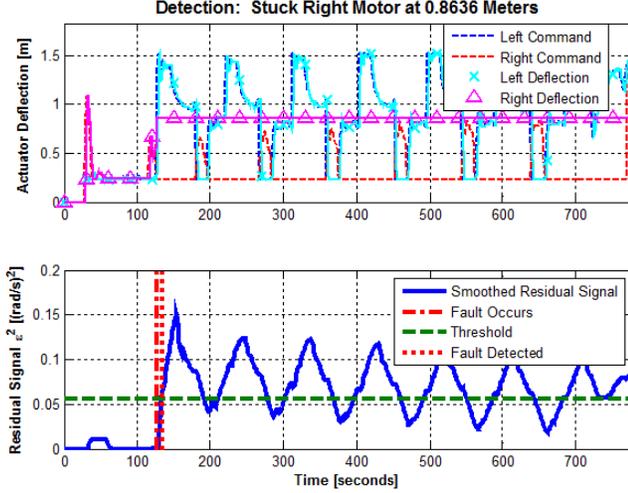


Figure 8. Detection example for stuck right motor fault.

5. ISOLATION

Once the detection algorithm has determined that a fault has occurred, the isolation process attempts to determine which particular fault is present. The data from the AGU motors can be used to determine if the fault is actuator-related. The first phase in the isolation process considers a motor residual signal, similarly to how the heading rate residual was evaluated during detection in Section 4. Each signal, one for each motor, should be small when the motor is behaving well and large when an actuator fault has occurred. These signals are evaluated using hypothesis testing. Appropriate thresholds must be selected such that $P(FA)$ and $P(MD)$ are minimized.

If the results of the first phase indicate that the fault is not actuator-related, phase two of isolation uses a bank of fault-specific observers (Willsky, 1976) to attempt to declare that a particular non-actuator fault has occurred. Some faults, such as a stuck motor and severe saturation, are difficult to model a priori because each of these faults is parameter-dependent (e.g. a stuck left motor at 0.5 meters). Other faults, such as criss-crossed lines and broken line, can be modeled in a straightforward manner as the effects of the faults are well-known. When a residual signal generated from a fault-specific observer is small, it is likely that the system has experienced the fault associated with that particular observer. Successful isolation will result in the declaration of a left motor fault, right motor fault, broken left line, broken right line, or a criss-crossed lines fault.

5.1. Motor Residual

Each motor on the AGU is equipped with an encoder that measures the corresponding control line deflection. This measurement is used as the performance metric for the motors on the AGU. Outputs from a nominal motor model are required to construct the motor residual signals used for isolation. The motor model used in the heading rate observer (Figure 5) is the nominal motor. The actual motor will output the motor deflection as measured by the encoders, and the nominal motor will output an expected motor deflection. The absolute value of the difference between these two signals is the motor residual signal,

$$r_{m,R}(t) = |\varepsilon_{m,R}(t)| = |\delta_{R,nom}(t) - \delta_R(t)| \quad (16)$$

$$r_{m,L}(t) = |\varepsilon_{m,L}(t)| = |\delta_{L,nom}(t) - \delta_L(t)| \quad (17)$$

where $r_{m,R}(t), r_{m,L}(t) \in \mathbb{R}$ are the right and left motor residual signals, respectively. The residual generation process is shown in Figure 9.

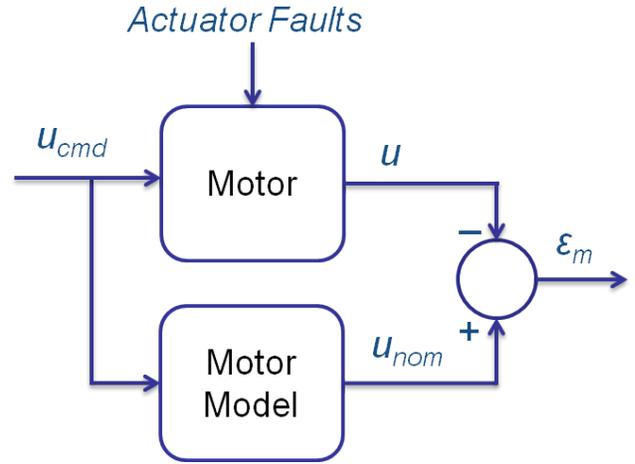


Figure 9. Motor residual block diagram.

Each motor residual is smoothed to emphasize the trend of the signal, similarly to the way the heading rate residual was smoothed in Section 4. Instead of the 25 second smoothing period used for the heading rate residual, each motor residual is smoothed over the previous 10 seconds at each time step. The behavior of the motor residuals is not as erratic as that of the heading rate residual and a longer smoothing period is not necessary. Once each motor residual has been generated and smoothed, it is evaluated by choosing a threshold. If a smoothed motor residual signal crosses above the chosen threshold within 5 seconds after detection occurs, an actuator fault in that particular motor is declared. This 5 second window is sufficient to detect most motor faults, as shown in the results in Section 6.1. This short amount of time also helps to

minimize total isolation time if the fault is not actuator-related and fault-specific observers must be used.

Figure 10 shows CDF plots of healthy data and data from flights in which an actuator fault has occurred. The healthy data represents the maximum value of the smoothed motor residual observed during the entire flight. All of this data is for a left motor fault. Results for right motor faults are almost identical and are not presented in this work. The fault data, which randomizes the type, time, and severity of each actuator fault, represents the maximum value of the smoothed motor residual observed during the first 5 seconds after the fault is declared.

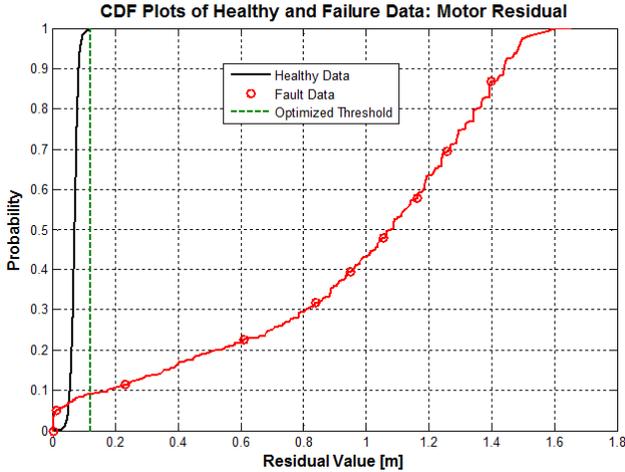


Figure 10. CDFs for motor residual.

As with the heading rate residual in Section 4, weightings on the FOM in Eq. (15) must be chosen for $P(FA)$ and $P(MD)$ to determine an optimized threshold. For equal weightings on each probability, the optimized threshold is 0.1237 meters. This threshold is shown as a vertical dotted line in Figure 10. Table 3 shows optimized thresholds for several different FOMs.

Table 3. Optimized thresholds for various FOMs (motor residual).

c1	c2	Threshold (m)	P(FA)	P(MD)	FOM
1.0	0.0	0.1363	0.000	0.094	1.000
0.8	0.2	0.1363	0.000	0.094	0.981
0.6	0.4	0.1238	0.000	0.093	0.962
0.5	0.5	0.1237	0.001	0.093	0.953
0.4	0.6	0.1198	0.002	0.092	0.944
0.2	0.8	0.1002	0.015	0.086	0.928
0.0	1.0	0.0074	0.999	0.040	0.960

5.2. Fault-Specific Observers

If the evaluation of the motor residual signals from Section 5.1 indicates nominal performance, phase two of the isolation procedure begins, which uses fault-specific observers. As with the heading rate residual, the residuals for fault-specific observers are constructed by differencing the EKF-estimated heading rate from the actual system and the expected heading rate output from the system observer. However, the expected heading rate signal comes from an observer that uses the dynamics of a system with a specific fault implemented. A block diagram of the residual generation process is shown in Figure 11. As with the heading rate and motor residuals, each fault-specific observer residual is smoothed. The residuals are smoothed at each time step over the previous 5 seconds of flight. This is to ensure that the residuals respond quickly to faults in order to minimize total isolation time. As with the heading rate residual, this smoothing parameter was tuned to achieve desired results from the fault-specific observers.

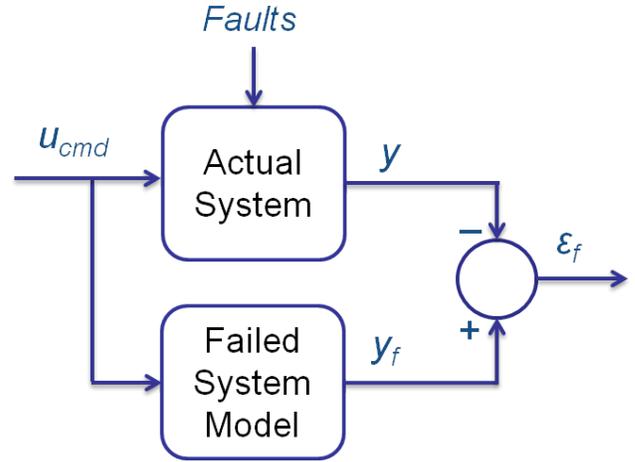


Figure 11. Fault-specific observer block diagram.

Fault-specific observers are constructed for the broken left line, broken right line, and criss-crossed lines faults. First, consider the broken left line fault. A simple modification to the nominal observer in Eq. (8) is required to construct this observer. The broken left line fault-specific observer is,

$$\dot{\hat{x}}_{bl}(t) = A\hat{x}_{bl}(t) + b_1\delta_R(t) + L(y(t) - \hat{y}_{bl}(t)) \quad (18)$$

$$\hat{y}_{bl}(t) = C\hat{x}_{bl}(t) \quad (19)$$

The observer for a broken right line is similar,

$$\dot{\hat{x}}_{br}(t) = A\hat{x}_{br}(t) - b_1\delta_L(t) + L(y(t) - \hat{y}_{br}(t)) \quad (20)$$

$$\hat{y}_{br}(t) = C\hat{x}_{br}(t) \quad (21)$$

where $\hat{x}_{bl}(t), \hat{x}_{br}(t) \in \mathbb{R}^4$ are state estimates for the broken left and right line observers, respectively, and $\hat{y}_{bl}(t), \hat{y}_{br}(t) \in \mathbb{R}$ are the outputs from each observer. The only modification needed to construct these observers is to remove either δ_L or δ_R from u_{nom} in Eq. (8), depending upon which line has broken. The corresponding residual signals for these fault-specific observers are,

$$r_{bl}(t) = (\hat{y}_{bl}(t) - y(t))^2 \quad (22)$$

$$r_{br}(t) = (\hat{y}_{br}(t) - y(t))^2 \quad (23)$$

where $r_{bl}(t), r_{br}(t) \in \mathbb{R}$ and $y(t)$ is the EKF-estimated heading rate for the system.

The observer for the criss-crossed lines fault is,

$$\dot{\hat{x}}_{cc}(t) = A\hat{x}_{cc}(t) - Bu_{nom}(t) + L(y(t) - \hat{y}_{cc}(t)) \quad (24)$$

$$\hat{y}_{cc}(t) = C\hat{x}_{cc}(t) \quad (25)$$

where $\hat{x}_{cc}(t) \in \mathbb{R}^4$ is the state estimate for the criss-crossed line observer, and $\hat{y}_{cc}(t) \in \mathbb{R}$ is the output from the observer. The only modification to Eq. (8) needed to create the observer is to reverse the effect of the input. The residual signal for the criss-crossed lines observer is

$$r_{cc}(t) = (\hat{y}_{cc}(t) - y(t))^2 \quad (26)$$

where $r_{cc}(t) \in \mathbb{R}$.

These residuals represent how well each observer models the current condition of the system, and should be small only when the fault that is modeled in the observer is present in the actual system. Thresholds need to be designed such that when the smoothed residual is above the threshold, it is likely that the corresponding fault is not present. These thresholds are used in the second isolation phase.

The second phase of isolation begins by assuming that all three faults considered in this phase (i.e., broken left line, broken right line, criss-crossed lines) are present in the system. At each time step, the smoothed residual signals associated with these three faults are evaluated. If a certain signal has crossed above its corresponding threshold, the fault associated with the signal is eliminated from consideration. Once two of the three signals have crossed their corresponding thresholds, the fault associated with the residual that remained below its threshold is declared and isolation is successful. The FDI method is given 90 seconds after the fault occurs to both detect and isolate the fault. As described in Section 4.2, this time limit is imposed to force a successful

isolation to occur in a reasonable amount of time such that there is sufficient time remaining in the flight for recovery. This parameter can be changed according to mission requirements.

If, at the end of the 90 second isolation period, one or fewer smoothed fault-specific observer residuals have crossed their corresponding thresholds, a final check is performed to attempt to isolate the correct fault. The relative size of each smoothed residual signal that has not already crossed its threshold is compared and the fault associated with the smallest signal is declared. The relative size R of each smoothed residual signal is given by

$$R = \max(r_s(t))/T \quad (27)$$

where T is the chosen threshold for the fault-specific observer, and the maximum of the corresponding smoothed residual signal $r_s(t) \in \mathbb{R}$ is computed over the 90 second isolation period.

There are several instances where isolation can fail. The first case occurs when all three of the smoothed fault-specific observer residual signals cross above their corresponding thresholds. This results in inconclusive isolation and no further action is taken. The other category of failure during isolation is called false isolation. This occurs when a fault is declared that is different from the actual fault that has occurred. This type of failure can occur during the motor observer phase if a motor fault is incorrectly declared, or during the fault-specific observer phase if the incorrect fault is declared. Several probabilities are used to assess the effectiveness of the isolation method. The probability that the correct fault is isolated given successful detection is $P(ISO)$. The probabilities that, once detection occurs, the isolation phase is inconclusive or declares an incorrect fault are given by $P(INC)$ and $P(FI)$, respectively.

5.2.1. Broken Line

The results for both the broken left line and broken right line faults are almost identical, so only the results from simulations of a broken left line are presented in this paper. In order to determine an appropriate threshold for the smoothed broken left line residual signal, the behavior of the signal was analyzed under three conditions: a broken left line is present, a broken right line is present, and a criss-crossed lines fault is present. CDFs showing the maximum value of the smoothed residual observed during the first 90 seconds after the occurrence of a fault for 1000 Monte Carlo simulations of all three fault cases are shown in Figure 12. A threshold for the broken line fault must be chosen so that the smoothed residual signal remains under the threshold for a large percentage of flights in which a broken left line fault occurs and crosses above the threshold for a large percentage of flights in which a broken

right line or criss-crossed lines fault occurs. The thresholds for the broken line and criss-crossed lines faults were chosen to provide a high rate of successful isolation while minimizing the time at which isolation completes. The threshold chosen for the broken line fault is 0.07 (rad/s)^2 and is shown as a vertical dotted line on Figure 12.

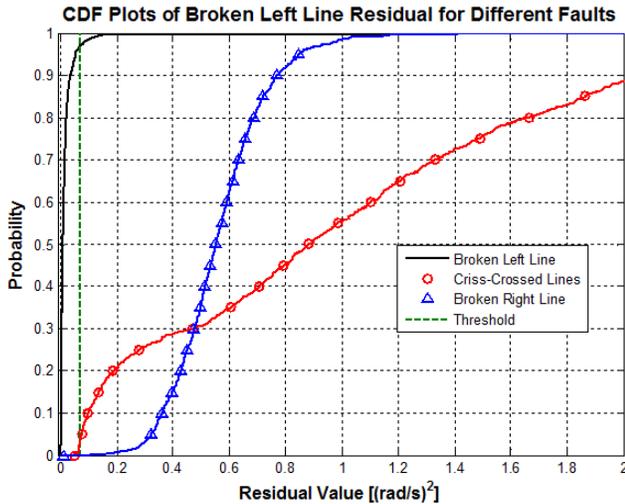


Figure 12. CDFs for broken left line residual.

The data from the CDFs in Figure 12 indicate that, given the chosen threshold, the probability that a broken left line is present but the smoothed residual signal associated with the broken left line still crosses the threshold is 2.9%. The probability that a broken right line is present but the smoothed broken left line residual remains under the threshold for the 90 second isolation period is 0.1%. The probability that a criss-crossed lines fault is present but the smoothed broken left line residual remains under the threshold for the duration of the isolation period is 2.9%. These probabilities can be similarly obtained for the broken right line case.

5.2.2. Criss-Crossed Lines

Figure 13 shows CDFs of data collected from the smoothed criss-crossed lines residual signal in the presence of a criss-crossed lines fault as well as broken left and right line faults. The CDFs show the maximum value reached during the first 90 seconds after the fault occurs for 1000 Monte Carlo simulations of the three fault cases. The threshold chosen for the smoothed criss-crossed lines residual signal is 0.13 (rad/s)^2 and is shown as a vertical dotted line on Figure 13.

The data from the CDFs in Figure 13 indicate that, given the chosen threshold, the probability that a criss-crossed lines fault is present but the residual signal still crosses over the threshold is 13.8%. The probability that a broken right or left line is present but the residual remains under the threshold for the isolation period is 5.9%. Isolation results for both

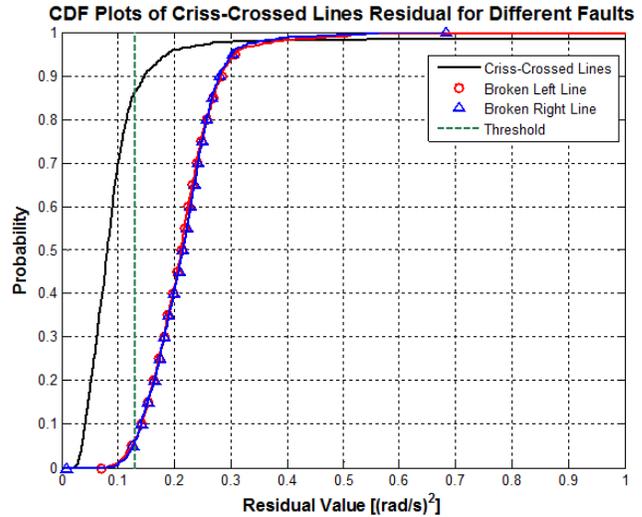


Figure 13. CDFs for criss-crossed lines residual.

the broken line and criss-crossed lines cases are presented in Section 6.1.

6. FULL FDI IMPLEMENTATION AND RESULTS

Sections 4 and 5 outline methods for both detection and isolation of faults for the parafoil and payload system. Figure 14 shows an overview of the integrated process.

The first step is fault detection. The heading rate signal is monitored throughout the entire flight. If, at any point, the smoothed residual signal rises above a predetermined threshold, a fault is declared. Once the alarm is raised, the algorithm progresses to the isolation method.

Isolation begins by evaluating the motor residual signal at the time of fault detection. If the smoothed residual from either the left or right motor is above the predetermined threshold, a fault in the corresponding motor is declared. With the declaration of an actuator fault, the FDI process ends.

Alternately, if the motor residual does not cross the threshold within 5 seconds after detection, the isolation algorithm progresses to the evaluation of residuals from a bank of fault-specific observers. If the smoothed residual signal associated with one of these observers is small, the algorithm declares that the fault corresponding to that particular observer has occurred. If none of the fault-specific observer residuals indicate that the system is exhibiting the characteristics of any known fault, then FDI has failed. In this case, it is likely that a non-actuator fault that does not have an observer associated with it has occurred. If more than one fault-specific observer models the actual system well, FDI is unsuccessful. Multiple faults occurring during the same flight are not considered in this work.

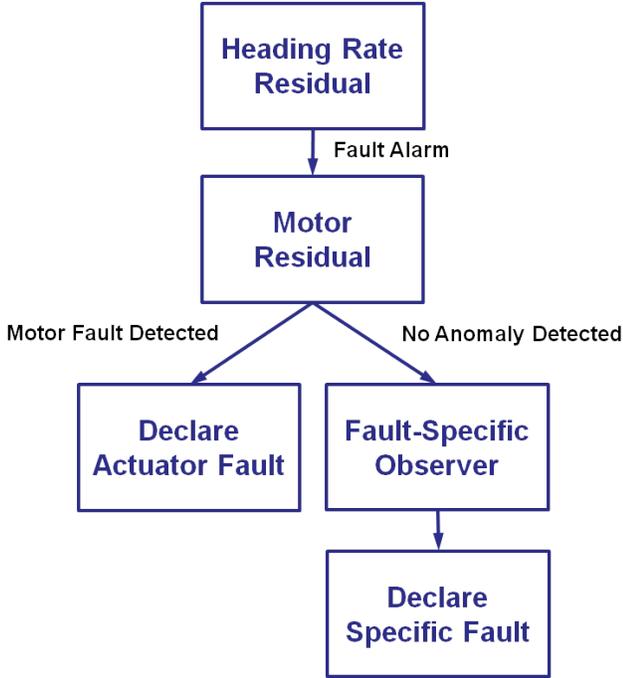


Figure 14. FDI Procedure.

6.1. Results

This section presents results of the full FDI implementation for three different fault categories: motor fault (i.e. stuck motor and unexpected saturation), broken line, and criss-crossed lines. This paper contains no comparison with parafoil FDIR work in the literature, as no such work was identified at the time of publication. Using the detection threshold for an equal weighting on $P(FA)$ and $P(MD)$ from Table 1, the motor residual threshold for an equal weighting on $P(FA)$ and $P(MD)$ from Table 3, the broken line fault-specific observer threshold chosen in Section 5.2.1, and the criss-crossed lines fault-specific observer threshold chosen in Section 5.2.2, the performance of the full FDI method was tested on 1000 Monte Carlo simulations in which a fault in one of the three categories being considered occurred at a random integer time during the first 500 seconds of flight. The performance of the FDI method is evaluated in terms of the probability of detecting the fault, $P(DET)$, as well as $P(ISO)$, $P(INC)$, and $P(FI)$. Table 4 shows FDI results for all three fault categories. The results compiled for motor faults and the broken line fault are compiled for the left motor and control line, respectively. Results from a right motor fault or a broken right line are almost identical to those obtained from the left side and are not presented in this work.

Table 4 indicates that a large portion of missed detections occur when attempting to detect motor faults. This is due in part to severe saturation. It is possible that, during the 60 second detection period, no motor command is given that is greater than the severe saturation limit. In this case, the parafoil be-

Table 4. FDI results for three fault categories.

Fault	$P(DET)$	$P(ISO)$	$P(INC)$	$P(FI)$
Left Motor Fault	0.732	0.969	0.003	0.029
Broken Left Line	0.989	0.990	0.004	0.006
Criss-Crossed Lines	0.997	0.953	0.019	0.028

haves as if no fault has occurred. Cases like this, and other scenarios in which the motor command and motor response are very similar, generally result in unsuccessful detection.

Despite the detection issues for the motor fault, the rate of isolation given a successful fault detection is high. The broken line and criss-crossed line faults have high isolation rates as well, validating the use of both the motor observer and fault-specific observers.

6.2. Examples

The first example of the full FDI implementation is for a stuck left motor at 1.295 meters. Figure 15 shows the heading rate residual and motor residual for a stuck left motor fault.

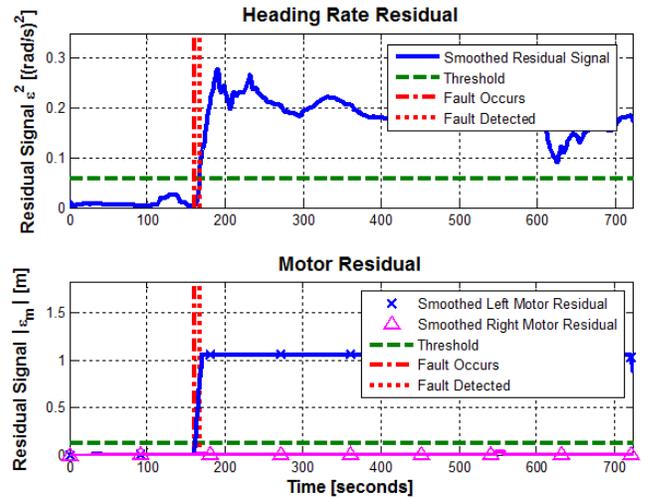


Figure 15. Stuck left motor fault: Heading rate and motor residuals.

The top subplot in Figure 15 shows the heading rate residual for the stuck motor case. A fault alarm is raised at 167 seconds, 7 seconds after the fault occurs. After this alarm is raised, the motor residual, shown in the bottom subplot, is evaluated to determine if the fault is actuator-related. The plot indicates that, at the time of detection, the left motor is showing off-nominal behavior while the right motor is behaving well. A left motor fault is declared and FDI is complete.

The second example demonstrates successful FDI for a broken left line fault. The heading rate and motor residuals are shown in Figure 16, and the fault-specific observer residuals

are shown in Figure 17 and Figure 18. The fault occurs at 193 seconds and a detection alarm is raised 15 seconds later. The motor residual is examined next. This residual indicates that both motors are behaving nominally, meaning that the fault is not related to the actuators. The next step is to consider the bank of fault-specific observers. Successful isolation occurs at 213 seconds, when both the broken right line and criss-crossed line residuals have crossed their corresponding thresholds. The broken left line residual remains under its threshold for the entire flight. The algorithm reports a broken left line 20 seconds after the fault occurs, successfully completing the FDI procedure.

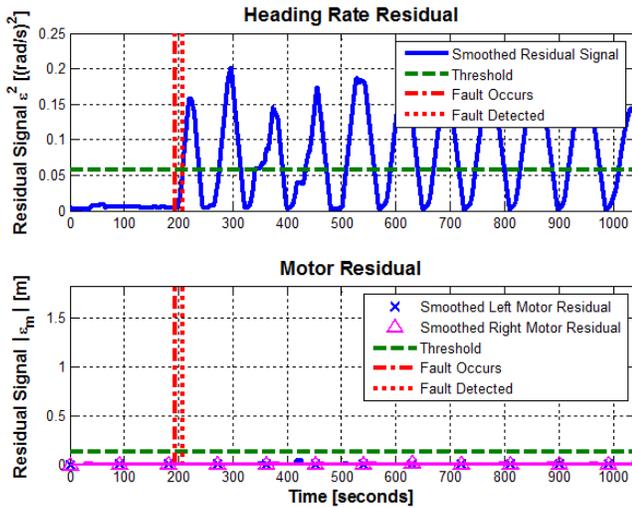


Figure 16. Broken left line fault: Heading rate and motor residuals.

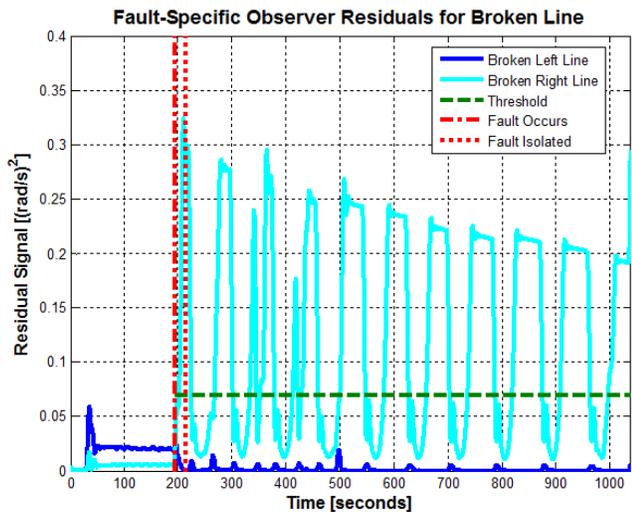


Figure 17. Broken left line fault: Broken line fault-specific observer residuals.

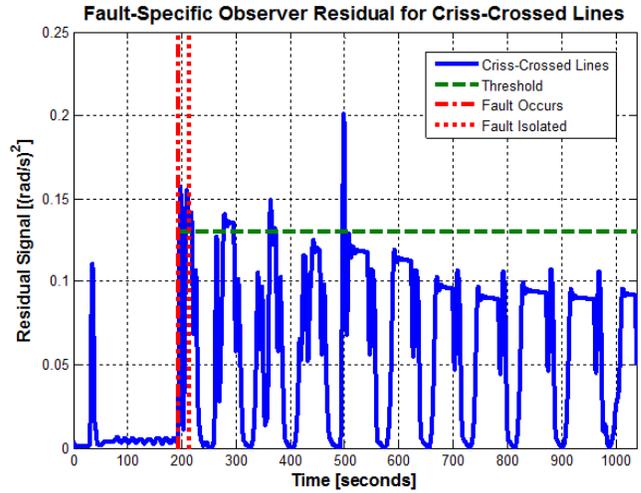


Figure 18. Broken left line fault: Criss-crossed line fault-specific observer residual.

7. CONCLUSION

This paper presented a fault detection and isolation method for an autonomous parafoil system. The detection method evaluates a residual signal generated from navigated heading rate and an observer based on a nominal system model. The isolation method uses separate residual signals generated from motor telemetry to determine whether a given fault is actuator-related. Other faults are isolated using fault-specific observers. We combined these methods and evaluated the approach against four common faults using high-fidelity Monte Carlo simulations. The results of these simulations showed that choosing an appropriate detection threshold allows for a high rate of detection with minimal false alarms. Choosing a threshold of 0.0995 (rad/s)^2 yielded an 82.8% success rate for detecting faults with a 0.5% rate of false alarms. For a broken line fault, successful detection occurs 98.9% of the time. Given successful detection, successful isolation of a broken line fault occurs at a rate of 99.0%. This work introduced a method for parafoil FDI that can detect and isolate common faults in an effective, timely, and predictable manner.

REFERENCES

- Barrows, T. (2009). Multibody parafoil model. In *20th AIAA aerodynamic decelerator systems technology conference and seminar*. Seattle, WA. doi: 10.2514/6.2009-2945
- Beard, R. V. (1971). *Failure accommodation in linear systems through self-reorganization*. Doctoral dissertation, Massachusetts Institute of Technology, Cambridge, MA.
- Bergeron, K., Fejzic, A., & Tavan, S. (2011). Accuglide 100: Precision airdrop guidance and control via glide slope

- control. In *21th AIAA aerodynamic decelerator systems technology conference and seminar*. Dublin, Ireland.
- Carter, D., George, S., Hattis, P. D., Singh, L., & Tavan, S. (2005). Autonomous guidance, navigation, and control of large parafoils. In *18th AIAA aerodynamic decelerator systems technology conference and seminar*. Munich, Germany. doi: 10.2514/6.2005-1643
- Crimi, P. (1990). Lateral stability of gliding parachutes. *Journal of Guidance, Control, and Dynamics*, 13(6), 1060-1063. doi: 10.2514/3.20579
- Douglas, R. K., & Speyer, J. L. (1995). Robust fault detection filter design. In *Proceedings of the american control conference* (Vol. 1, p. 91-96). Seattle, WA. doi: 10.1109/ACC.1995.529214
- Figueroa, F., Schmalzel, J., Morris, J., Turowski, M., & Franzl, R. (2010). Integrated system health management: Pilot operational implementation in a rocket engine test stand. In *AIAA Infotech@Aerospace*. Atlanta, GA. doi: 10.2514/6.2010-3454
- Figueroa, F., Schmalzel, J., Walker, M., Venkatesh, M., Kapadia, R., Morris, J., ... Smith, H. (2009). Integrated system health management: Foundational concepts, approach, and implementation. In *AIAA Infotech@Aerospace conference*. Seattle, WA. doi: 10.2514/6.2009-1915
- Frank, P. M. (1994). Enhancement of robustness in observer-based fault detection. *International Journal of Control*, 59(4), 955-981. doi: 10.1080/00207179408923112
- Hattis, P., Campbell, D. P., Carter, D. W., McConley, M., & Tavan, S. (2006). Providing means for precision airdrop delivery from high altitude. In *AIAA guidance, navigation, and control conference and exhibit*. Keystone, CO.
- Hattis, P., & Tavan, S. (2007). Precision airdrop. *Aerospace America*, April, 38-42.
- Hwang, I., Kim, S., Kim, Y., & Seah, C. E. (2010). A survey of fault detection, isolation, and reconfiguration methods. *IEEE Transactions on Control Systems Technology*, 18(3), 636-653. doi: 10.1109/TCST.2009.2026285
- Isermann, R., & Ballé, P. (1997). Trends in the application of model-based fault detection and diagnosis of technical processes. *Control Engineering Practice*, 5(5), 709-719. doi: 10.1016/S0967-0661(97)00053-1
- Jones, H. L. (1973). *Failure detection in linear systems*. Doctoral dissertation, Massachusetts Institute of Technology, Cambridge, MA.
- Patton, R. J., & Chen, J. (2000). On eigenstructure assignment for robust fault diagnosis. *International Journal of Robust and Nonlinear Control*, 10(14), 1193-1208. doi: 10.1002/1099-1239(20001215)10:14<1193::AID-RNC523>3.0.CO;2-R
- Rossi, C. (2012). *Vehicle health monitoring using stochastic constraint suspension*. Master's thesis, Massachusetts Institute of Technology, Cambridge, MA.
- Rossi, C., Benson, D., Sargent, R., & Breger, L. S. (2012). Model-based design for vehicle health monitoring. In *Infotech@Aerospace*. Garden Grove, CA. doi: 10.2514/6.2012-2577
- Rossi, C., Breger, L., Benson, D., Sargent, R., & Fesq, L. (2012). Vehicle health monitoring using stochastic constraint suspension. In *AIAA guidance, navigation, and control conference*. Minneapolis, MN.
- Sargent, R., Mitchell, I., Breger, L., Benson, D., Bessette, C., Zanetti, R., & Groszkiewicz, J. E. (2011). A fault management strategy for autonomous rendezvous and capture with the ISS. In *Infotech@Aerospace*. St. Louis, MO. doi: 10.2514/6.2011-1497
- Slegers, N., & Costello, M. (2004). Model predictive control of a parafoil and payload system. In *AIAA atmosphere flight mechanics conference and exhibit*. Providence, RI.
- Sturza, M. A. (1988). Navigation system integrity monitoring using redundant measurements. *Journal of the Institute of Navigation*, 35(4), 69-87.
- Tavan, S. (2006). Status and context of high altitude precision airdrop delivery systems. In *AIAA guidance, navigation, and control conference and exhibit*. Keystone, CO.
- Van de Vegte, J. (1994). *Feedback control systems*. Prentice Hall.
- Ward, M., Montalvo, C., & Costello, M. (2010). Performance characteristics of an autonomous airdrop system in realistic wind environments. In *AIAA atmosphere flight mechanics conference*. Toronto, ON, Canada.
- Willsky, A. S. (1976). A survey of design methods for failure detection in dynamic systems. *Automatica*, 12, 601-611.