AIAA Guidance, Navigation, and Control Conference
2 - 5 August 2010, Toronto, Ontario Canada

AIAA 2010-7855

Fault Detection and Fail-Safe Operation with a Multiple-Redundancy Air-Data System

Ryan D. Eubank,[1] Ella M. Atkins,[2] and Stephanie Ogura[3]
University of Michigan, Ann Arbor, MI, 48109

Air-data systems (ADS) measure wind speed and direction, the loss of which requires aerodynamic forces to be estimated from inertial measurements and aircraft dynamics and performance models. The nature of ADS measurements require air-data probes be subject to the spectrum of environmental conditions. Even with designs meant to withstand harsh conditions, instances of ADS probe failure have been recorded for diverse platform types and situations. Further, since all ADS probes on a common platform are subject to the same conditions, instances of multiple simultaneous failures are not uncommon. Robust air data measurement therefore becomes a multi-sensor data-fusion problem wherein the system may be subject to failures that effect groups of like sensors, such as pitot-static probes, simultaneously. This paper presents an algorithm for fault detection and data fusion of ADS failures in the framework of an unmanned autonomous seaplane with a heritage of air-data probe failures. The fault detection scheme is based on sensor signal characterization and monitoring and on the comparison and fusion of redundant sensor measurements. A GPS/INS-driven backup will also be proposed that can be used both as an ADS diagnostic tool and to allow safe flight to an emergency landing or until air-data sensor functionality can otherwise be restored. Flight test data from two generations of unmanned seaplanes demonstrates the efficacy of the algorithm for a range of real-world failure cases with varied sensors.

Nomenclature

α	=	angle-of-attach
β	=	sideslip
Λ_{type}	=	results of pass/fail test for a specific *type* of sensor health indicator
ξ	=	decay parameter for exponential weighting determination
σ_i^2	=	signal variance
ϕ, θ, ψ	=	roll, pitch, and yaw Euler angles
C_n	=	confidence rating for the n^{th} sensor signal (0,1)
k	=	time-step, indicating discrete measurement or computation
P_{type}	=	probability of that a specific *type* of sensor failure has accurately been detected
S_n	=	n^{th} sensor signal
T_{type}	=	sensor failure threshold for a certain test *type*
u	=	airspeed
$\vec{u}_A$	=	airspeed vector in reference frame *A*
$v(s)$	=	variance function applied to signal *s*
v_1,v_2	=	left and right pitot-static measurements, respectively
v_a	=	propeller anemometer measurements
v_e, v_{est}	=	wind-model based airspeed estimate output
v_{RESULT}	=	fault-tolerance algorithm output, high-confidence airspeed result
w	=	weighting coefficient
$wind$	=	estimated environmental wind vector

[1] Graduate Research Assistant, Aerospace Engineering, 2016 FXB, Student Member
[2] Associate Professor, Aerospace Engineering, 3009 FXB, Associate Fellow
[3] Undergraduate Research Assistant, Aerospace Engineering and Computer Sci., 1250 FXB

I. Introduction

CONVENTIONAL air-data system (ADS) probes provide direct measurement of body-relative wind represented as airspeed (u), angle-of-attack (α), and sideslip (β) as well as pressure-based altitude (z) estimates. As stability margins are traded for efficiency/performance and levels of automation increase, measurements of wind-relative attitude and speed become more important. As such ADS are among the most basic and critical sensor packages onboard an aircraft and are generally comprised of similar, but not strictly homogeneous, environmental sensors. The nature of ADS measurements require probes to be directly exposed to the flight vehicle's operating environment. In addition to human error and physical damage, foreign material infiltration and atmospheric anomalies (e.g., extremely high water content) are two of the most common causes of ADS probe failures.[1-4] Such failures tend to effect similar sensors identically giving rise to complete sensing loss despite redundancy. These failures affect commercial,[2,4] military,[5] and general aviation aircraft.[1,3,4] Furthermore, though no comprehensive failure statistics are available, it follows that failures would also be experienced in unmanned aerial vehicle (UAV) ADS probes. The immediate consequences of ADS failure are incorrect airspeed/direction and altitude readings. A human pilot will likely recognize the failure if the reading is clearly incorrect, but will be less likely to immediately notice if the reported values are reasonable even if they are based on incorrect measurements. Despite recognition, the human pilot might still have trouble maintaining stable flight without airspeed, particularly in turbulent atmospheric conditions. An autopilot, whether part of a manned or unmanned aircraft system, nominally incorporates airspeed into its flight control laws. Upon failure, if the erroneous data is not detected, control excursions can be substantial and induce unsafe flight conditions (e.g., pitching to a dive given an airspeed approaching stall).

Deployed commercial and UAS autopilots are generally not adaptive to incorrect ADS information, so the pilot of a manned aircraft will typically initiate manual flight control without direct knowledge of airspeed, while a UAS will either execute a safe ditch or be controlled remotely, again without airspeed data. The potential for an autopilot or pilot to react improperly to erroneous wind data introduces appreciable risk, as evidenced by accidents such as Aero Peru Flight 603, in which ground crews failed to remove tape from the pitot-static system after cleaning the aircraft, or the X-31A at NASA Dryden, in which pitot icing compromised system readings. Note that although redundant ADS probes are present on most high-cost aircraft, common failure modes or incorrect failure diagnosis has also resulted in catastrophic accidents, such as Austral Lineas Aeroeas Flight 2553 in which the flight crew improperly referenced the pilot's airspeed indicator and induced structural failure by exceeding safe airspeed limits.

This paper will present an ADS failure mitigation algorithm that fuses data from multiple wind and inertial sensors to diagnose and react to air-data sensor failures. The methods can be applied to a range of systems and sensor types but, for the purposes of this paper, the specific sensor measurements are defined in the context of instrumentation affixed to an unmanned seaplane with a history of pitot failures due to water ingestion. This motivating autonomous seaplane system will be introduced in the next section followed by a review of literature and statistics concerning the frequency and implications of ADS failures. The background section will conclude with a review of pertinent past research in ADS failure mitigation. Following the background materials our sensor fault detection and data-fusion algorithm is presented. Finally, results are provided from two generations of unmanned seaplane flight test programs that have demonstrated the efficacy of this solution. The paper concludes with an analysis of the capabilities and limitations of the presented algorithm.

II. Background

This section introduces a motivating case study from an unmanned seaplane program with a heritage of ADS failures. ADS fault tolerance is then motivated from past incidents, followed with a summary of our investigation into the impact, common causes, and mitigation strategies associated with ADS failures. Finally, fault identification and tolerance, in the context of ADS, are discussed

A. Motivating Example: Autonomous Seaplanes

In the summer of 2007, the Flying Fish autonomous unmanned seaplane[6-8] first began open-water testing. During initial autonomous controller development an unpredictable ADS failure mode manifested through a variety of seemingly unrelated system malfunctions. The initially inexplicable behaviors included: failed automated takeoff, sudden pitch to stall, and sharp dives from cruise. The issues were manually traced to anomalous airspeed values due to water blockage of ADS probe ports. Like most floatplanes, the Flying Fish relied on a pressure-based pitot/static probe for airspeed determination, but unlike a full size floatplane the UAV's scale prevented the probes

from being mounted any further than ~1.0m from the water's surface. Unavoidable physical proximity, in combination with energetic transit over the water resulted in water impingement on the air-data probes. In-flight incidents revealed that ingested water might not immediately cause a failure, but that variations in flight attitude or airspeed might precipitate a delayed failure of a water-compromised probe. Blockages resulted in a wide range of erroneous measurements; while some blockages produced near constant output others displayed damped/biased tracking of actual airspeed. The Flying Fish ADS probe was subsequently relocated to minimize water incursion and a redundant probe and measurement system were added. Despite these modifications, however, water blockage issues still occurred, requiring the vehicle to be recovered and manually cleared of ingested water, often found well inside the pressure tubing where pitot heating is ineffective. A second-generation Flying Fish seaplane developed at the University of Michigan was equipped with a more comprehensive redundant ADS system. This vehicle has also experienced ADS sensor problems, although it is equipped with dual heated pitot probes further separated from the water. The new ADS system has redundancies including duplicate ADS sensors, different types of ADS sensing technologies, probe measurement redundancy, and failed-sensor recovery mechanisms. Specifically, the system includes two 5-hole pressure probes and a propeller-based anemometer. The 5-hole probes combine pitot/static airspeed measurement and barometric altitude with lateral/vertical differential pressures for the determination of angle of attack and side-slip angle. A heating element on each 5-hole-probe allows for cold weather operation and has sufficient heating capacity to rapidly evaporate freshwater blockages. Heat-based pitot clearing has not been evaluated in a marine environment where the mineral content of the water may contraindicate the application of evaporative clearing. The propeller-anemometer uses hall-effect sensors to measure the rotation rate of a small high-pitch propeller in order to determine airspeed. Dual hall-effect sensors within the anemometer head provide redundant measurement of propeller rotation.

B. Exploration of Air-Data System Failures

The rates and impact of ADS failure on commercial, military, general, and unmanned aviation are nontrivial. Records and reports on the subject suggest that the problem may be increasingly prevalent, with growing air travel volume, and while research is being conducted on related topics no uniform solution yet exists. Aviation safety databases provide evidence of significant commercial aviation losses due to ADS failure. The Aviation Safety Network database has records of at least eleven ADS (pitot probe) failures over the past three decades that have resulted in significant damage or loss of life.[4] These examples alone represent a nontrivial financial loss and 342 documented fatalities (339 in the past 15 years). More recently, interim accident reports for Air France Flight 447 indicate air data system anomalies were likely experienced[9] and that Airbus platforms have had 35 recorded incidents of multiple ADS failures since 2003.[10] This is a nontrivial result even over the large number of total flight operations conducted by Airbus airliners given the likelihood that many transient ADS failures were not documented. The effect of ADS failures can also be observed in general aviation (GA) aircraft incident records,[1-3] but concise statistics have proven more difficult to collect. For example, the Aviation Safety Reporting System database[1] contains numerous instances of general aviation ADS-related failures but aggregate results for this specific contributing factor are not readily available. Further complicating analysis, the varied causes, effects, and results of ADS faults can lead to failure statistics being associated with a number of different classifications (e.g. inclement weather, instrument fault, and flight control failure).

C. Air-Data System Failure Mitigation Research

Safety in aviation is paramount, particularly for commercial aviation due to the extreme property and loss-of-life costs associated with aircraft accidents. Industry and government demand highly-validated systems with safety and availability requirements that specify a 10^{-9} maximum probability of critical failure per flight hour.[11] Two major results of these standards are stringent validation requirements that slow the application of state-of-the-art concepts and ever-increasing complexity in the avionics employed on all classes of aircraft. The avionics of most commercial aircraft now feature multi-redundant self-monitoring systems with segregation and purposeful dissimilarities between related/redundant software and hardware; these systems require complex redundancy negotiation and consensus voting strategies to operate.[12,13] While it has been noted that this complexity may induce unexpected and counterintuitive results, even to the point of introducing new failure modes,[14] it is still the case that these systems have demonstrated high reliability. Nevertheless, these systems are still fundamentally vulnerable to failures in their external sensing apparatus. Further, while these complex redundancies and failure mitigation strategies enable the negotiation of failing redundant sensors these systems are still largely unable to handle common failure mechanisms simultaneously disabling entire classes of like sensors. These shortcomings are the fundamental reasons to research new mechanisms for air-data failure mitigation. Most research into ADS fault tolerance and recovery tends to fall

into one of three categories: (1) signal-based diagnostics, (2) alternative sensing mechanisms, and (3) finding ways to operate without traditional ADS sensors.

1. Signal-Based ADS Failure Rejection

Looking first to ADS-specific failure research we find the seminal work of Houck and Atlas which provides insight into fundamental mechanisms for ADS failure diagnosis. Houck and Atlas analyzed failed ADS sensor signals and were amongst the first to propose that probe blockage reduced signal energy levels, that large signal variations were generally sufficient (but not necessary) to demonstrate sensor functionality, and that signal characteristics might be used to indicate air-data probe health.[15] Very few examples of this type of analysis exist for ADS-specific applications. Houck and Atlas ultimately proposed that even at a fixed altitude the nominally-constant static pressure varied slightly as a function of acceleration and that the derivative of the static-port pressure signal would be a good indicator of probe health. Unfortunately, independent static pressure measurements are not always available in UAV applications as the desire for volume, weight, and cost savings make the implementation of a single pressure transducer for pitot-static measurements more likely. Regardless, Houck and Atlas' methods utilize or suggest several of the tools that will be employed in this paper including individual signal characterization and comparison with previous statistics and predetermined operating thresholds.

2. Alternative ADS Sensing Apparatus

A more common means for avoiding failures associated with a particular class of sensors is to employ alternative instrumentation for the same measurements. Examples of variations on the classic ADS pitot-probe include flush air-data sensing (FADS) systems[16] and self-aligning multi-hole conical probes.[17] FADS systems employ pressure ports with openings flush to, and distributed over, a vehicle's aerodynamics surfaces while self-aligning conical probes are driven by pressure forces into alignment with local airflows. In both cases the geometrically-related measurements collected at distributed sensing locations produce an over-defined system from which the ADS states can be re solved. These systems can provide both fault-tolerance and error reduction provided they are designed such that the ADS states are observable from different subsets of probes. Further, the novel structure of the sensors produces changes the potential failure modes and reduces the likelihood of simultaneous failure with the more common traditional ADS probes. The system presented in this paper will, as previously indicated, utilize alternative low-cost ADS sensing technologies to avoid having failure modes common across all of the vehicle's air-data sensors.

3. Circumventing ADS Sensing

A number of researchers have proposed alternatives for flight operation in the absence of ADS measurements. The advent and proliferation of GPS systems and continued improvements to sensors and filtering in inertial navigation systems (INS) have given rise to mechanisms for deriving estimates of ADS states indirectly.[5,18-19] Consequently, flight control laws have emerged that do not require ADS state variables.[20-21] Implementations of the latter generally have performance limitations (e.g., limits on wind speed or variability), must possess particularly wide margins for safe operation, and/or must employ some alternate motion sensing mechanism (e.g. machine vision). More pertinent to this work is the concept of ADS state estimation for which two basic formulations, differing by a time derivative, have been proposed. Starting from known initial conditions, such as the point of ADS failure, the estimation algorithms infer wind from either: a) the difference between a wind-unaware dynamic estimate of inertial velocity and the measure of inertial velocity[5] or, b) the integration of lateral, longitudinal, and vertical accelerations by a dynamic model that includes wind.[18-19] Regardless of the method, trigonometry is applied to the resulting velocity triangles to determine angle of attack and sideslip. The difficulty with these methods is their reliance upon high-accuracy high-rate inertial sensors and upon high-fidelity dynamic models. While these assumptions are reasonable for the military vehicles for which much of the referenced research was intended, they are not necessary appropriate for a slow-flying, low-cost UAV with MEMs-based sensors. With low-cost UAV platforms as the target, the algorithm presented in this paper will extract external wind estimates from the sensor system and leveraging wind estimates to infer wind-relative motion in the absence of trusted ADS measurements.

III. Methods

The fault tolerant ADS algorithm (Fig. 1) is composed of three primary elements: a signal-fault detection scheme, a confidence-discriminate data-fusion procedure, and an inertial measurement driven wind-estimation system. Signal fault detection extracts and tests signal characteristics to estimate the likelihood of sensor failure. The confidence-discriminate data-fusion combines the signal fault results into sensor confidence values. Confidence values are used to judge sensor fitness and eliminate signals from failed sensors before being used in the weighted-average data-fusion. The wind-estimation module utilizes the resulting composite air-data vector to refine a local wind model which is subsequently used, in conjunction with inertial navigation measurements, to estimate the air-

data vector in the event of ADS sensor failures. This wind-estimate also provides a baseline wind vector for judging individual sensor measurements. The three algorithm stages are executed sequentially. Wind estimates are fed back to the signal-fault detection scheme to become one of the confidence-rated signals combined in the data-fusion cycle.

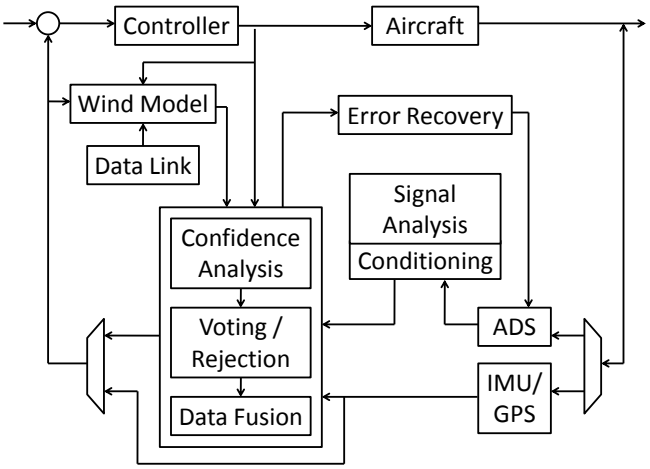

Figure 1. ADS Fault-Tolerant Control System Diagram

A. Signal Fault Detection

The purpose of the signal-fault detection scheme is to discern anomalous operating conditions that may indicate a sensor failure. The number, type, and redundancy of the sensors are not considered at this stage, rather each signal is judged based on individual parameter-based signal models. The signal models are captured from the sensor specifications and the analysis of both failed and operational sensor signals. This stage of the algorithm is composed of three major procedures: signal characteristic extraction, signal model determination, and signal fault testing.

Common methods for signal characterization in fault detection schemes include statistical metrics, such as arithmetic average and variance, and spectral analysis methods, such as the wavelet[22] or Fourier[23] decompositions. Spectral decompositions are typically applied for fault detection in systems with cyclic behavior[22] or harmonic content/excitation.[24] The efficacy of frequency decomposition methods for common classes of air-data sensors are, as of yet, unknown but these systems do not generally have a strong frequency component. The team is currently investigating the application of frequency decomposition methods for propeller anemometer diagnostics, but no conclusions have yet been drawn. For the remaining pressure-based ADS signals we employ an arithmetic average and variance algorithm for signal characteristic extraction. Unlike decomposition strategies, which are focused on frequency-keyed information, the mean and variance are utilized to obtain smoothed, low-pass filtered, characteristics of a signal.[22] In this case we will be using a sliding average and variance formulation. The k^{th} arithmetic-average and variance over m samples of the n^{th} signal ($s_n(k)$) is given, for the set of non-negative/not-all-zero weighting coefficients (w_i), by the formulae:

$$\bar{s}_n(k) = \frac{\sum_{i=1}^{m} w_i s_n(k+1-i)}{\sum_{i=1}^{m} w_i} \tag{1}$$

$$\nu(s_n(k), k) = \frac{\sum_{i=1}^{m} w_i \left(s_n(k+1-i) - \bar{s}_n(k)\right)^2}{\sum_{i=1}^{m} w_i} \tag{2}$$

To reduce the number of computations per iteration, the sums of time-invariant weighting factors are normalized producing the following simplified formulation:

$$\bar{s}_n(k) = \sum_{i=1}^{m} w_i x_{k+1-i}\,,\ \sum_{i=1}^{m} w_i = 1 \tag{3}$$

$$\nu(s_n(k), k) = \sum_{i=1}^{m} w_i \left(s_n(k+1-i) - \bar{s}_n(k)\right)^2 \tag{4}$$

Exponential weighting is utilized to favor the most recent data, reducing phase delay between the raw and filtered signals. Our parameterized exponential weighting formula is given by:

$$\vec{w}(\xi) = [w_1(\xi), w_2(\xi), \ldots, w_m(\xi)]^T, 0 < \xi \leq 1 \tag{5}$$

$$w_i(\xi) = \frac{\xi(1-\xi)^{i-1}}{\sum_{k=1}^{m} \xi(1-\xi)^{k-1}}, i = \{1, \ldots, m\} \tag{6}$$

$$w_i(\xi = 0) = w_1 = \cdots = w_m \tag{7}$$

The decay parameter ξ determines the relative influence of aging data points, defining a continuum between preserving only the most recent sample at one extreme (ξ=1) and approaching equal weighting of all points at the other extreme (ξ arbitrarily close to zero). We define equal weighting for the special case ξ=0. In this case, the un-weighted arithmetic mean and standard deviation are recovered from Eqs. (1)-(2).

To apply Eqs. (1)-(7) the sliding sample-window size, m, and the weight decay parameter, ξ, must be selected. These values were tuned empirically to balance signal tracking against delay and low-pass filter performance for each signal. It may be possible to formulate an optimal tuning of these parameters if a cost function can be formulated based on the signal following characteristics and low-pass filtering requirements. This approach did not prove necessary the Flying Fish sensor systems but may be appropriate for any larger or more complex ADS. The resulting tuned sliding average and variance formulations were then applied to the signals of functional, failed, and failing sensors to extract sensor model parameters for each flight mode (taxi, takeoff, climb, cruise, turning, descent, and landing). Average variance ($\bar{v}(s_n, mode)$), peak change rate ($\dot{s}_{n,peak}(k, mode)$), and peak rate of variance change ($\dot{v}_{peak}(s_n, mode)$) were extracted from functional sensor data for each signal and each flight mode. Performing the same calculations on failed and failing sensor data enabled the determination of tolerances for deviation in the model parameters. During this process it was discovered that the peak rate of change during a failure was generally within the normal dynamic range of our sensors. That is to say, the initial dynamics of probe failures are almost indistinguishable from the dynamic response of the functional system except that the measurements are increasingly incorrect. Subsequently, peak rate of change was discounted as a fault detection metric but tolerances for deviation from average variance ($T_{\bar{v}}(mode)$) and peak rate of variance change ($T_{\dot{v}_{peak}}(mode)$) were recorded for each sensor. Drawing from the manufacturer's specifications and failed sensor data the saturation limits of the sensor ($s_{n,min/max}$) were quantified to give a total of six fault-detection parameters. The collected set of all fault-detection parameters and the averaging window and decay parameter for each filter comprise the parameterized sensor signal model.

The final step in the signal-fault detection process is the evaluation of extracted characteristics using the metrics and tolerance stored for each sensor, for each flight regime. The combination of model parameters produces three distinct sensor-fault tests. The first test determines if the signal variance exceeds the variance deviation threshold for the current flight mode:

$$|v(s_n(k), k) - \bar{v}(s_n, mode)| > T_{\bar{v}}(mode) \tag{8}$$

The result of the variance test is recorded as a binary pass ("1") or fail ("0") vote ($\Lambda_{\mathrm{var}}(n, k)$) for each sensor at each time step, k. The second test determines if the signal variance increases or decreases too quickly and violates the peak-variation-rate tolerance. In order to make this determination we require a smooth baseline measurement of the signal variance for which we compute a sliding average of the variance results ($\bar{v}(s_n(k), k)$). Again the window size, m, and the weight decay parameter, ξ, are selected empirically to find an acceptable tradeoff between signal following, smoothing, and delay. The signal test is formulated as:

$$|v(s_n(k), k) - \bar{v}(s_n(k), k)| > T_{\dot{v}_{peak}}(mode) \tag{9}$$

The result is a binary pass-fail vote for the variance rate test ($\Lambda_{\mathrm{rate}}(n, k)$) of the n^{th} sensor at the k^{th} time step. The final test considers if the sensor has entered a saturation region. Data analysis indicated a threshold of 3% of the saturation limits provided an appropriate balance between missed and false fault detection. The set of all flight data demonstrated that an operational sensor's signal remain further than 3% from the saturation limit for over 99% of all measurements. Conversely, failures that produced saturated signals approached within 3% of the saturation limit for 99% of subsequent incorrect measurements. The primary occurrence of saturation in the un-failed sensor is low-speed saturation for extremely low or zero air speed (or a slight tailwind) prior to initiating takeoff. The tests for saturation take the form:

$$s_n(k) < (1.03)s_{n,min} \tag{10a}$$
$$s_n(k) > (0.97)s_{n,max} \tag{10b}$$

This test produces, as with the previous tests, a binary indication a saturation test failure of the n^{th} sensor at the k^{th} time step ($\Lambda_{\mathrm{sat}}(n, k)$). Finally, the time history of sensor fault votes is output to the confidence-discriminate data-fusion algorithm in order to develop sensor confidence values.

B. Confidence-Discriminate Data Fusion

The goal of confidence-discriminate data-fusion is to leverage signal confidence and redundant-data comparisons to combine like measurements while excluding failed sensors. This sensor discrimination and data fusion process is comprised of three steps: signal confidence determination, failed sensor rejection, and final data fusion.

The signal confidence of the n^{th} signal ($C_n(k)$) is developed from the time history of signal-fault votes for that sensor in a two step process. First we accumulate signal-fault votes into a probability that a sensor has passed a specific fault test. This probability is created by using a large-window moving average which allows failure votes to have a long influence period and also mitigates spurious intermittent false-negative/positive votes. For this we define probabilities for each of the three primary fault types: signal variance exceeds threshold (P_{var}), rate of variance change exceeds threshold (P_{rate}), and signal exceeds saturation tolerance (P_{sat}). The probabilities are created by a moving average of the binary voting history for each fault:

$$P_{var}(n,k) = \overline{\Lambda}_{\mathrm{var}}(n,k) \tag{11a}$$
$$P_{rate}(n,k) = \overline{\Lambda}_{\mathrm{rate}}(n,k) \tag{11b}$$
$$P_{sat}(n,k) = \overline{\Lambda}_{\mathrm{sat}}(n,k) \tag{11c}$$

The second step is to combine the three failure-voting probabilities into the final signal confidence by a weighted average. The two variance probabilities are given even weighting (0.3) while the saturation probability is given a slightly higher weighting factor (0.4).

$$C_n(k) = (0.3)P_{var} + (0.3)P_{rate} + (0.4)P_{sat} \tag{12}$$

This distribution is best explained by examining the features of our fault detection process. Specifically, a complete saturation fault indicates an undeniable failure condition whereas the variance-based faults indicate only a likelihood of sensor failure. By giving a slightly greater weight to the saturation test we can select a confidence threshold (0.7) that is exceeded in the event of a complete saturation failure but that cannot be surpassed by any single variance failure. Armed with confidence values for each signal the error-rejection and data fusion can be completed.

Perhaps the most important requirements for a fault-tolerant system are mechanisms for judging and rejecting questionable signals from the set of all available sensors. Willsky's survey of design methods for failure detection provides a good summary of this field.[25] Common methods include neural networks, voting or outlier rejection, model-based analysis, and filter-based techniques, including recursive least-squares and the Kalman filter.[26-30] We will focused on outlier rejection (OR) and voting; these concepts are closely related but typically utilize different operating principles. Voting schema are often comprised of rule-based judgments. Outlier rejection, on the other hand, generally relies on statistical analysis and, while subjectivity remains in the selection of metrics and thresholds, OR methods are usually based on commonly accepted statistical practices/measures (e.g. using a multiple of the standard deviation to define an outlier). A comprehensive treatment of outlier rejection in statistical data can be found in Barnett and Lewis.[31] For this paper we will utilize sensor confidence as a per-signal voting mechanism and employ a simple outlier rejection scheme whenever three or more redundant signals are available. As previously indicated a confidence threshold of 0.7 was selected; whenever a signal's confidence drops below this threshold it is rejected from data fusion. If three or more redundant measurements remain after the confidence-based elimination they are subject to a consensus-seeking outlier rejection scheme that eliminate signals too dissimilar (subject to a threshold) from any majority amongst all signals. Utilizing Eqs. (1) and (2) confidence-weighted average and variance are computed for the set of redundant sensors. The square root of the variance gives the standard deviation of the set of signals. Any signal that is more than one standard deviation from the average is eliminated.

$$|s(k) - \bar{s}(k)| > \sqrt{v\big(s(k)\big)} \tag{13}$$

The remaining mathematical concepts that must be considered are the class of methods employed for data fusion. Hall and Llinas provide a comprehensive introduction to data fusion and a comparison and classification of data fusion operators can be found in Block's 1994 manuscript.[32-33] It can be shown that a great many of the filter-based data fusion algorithms are based on least-square error concepts.[34] A noisy measurement (z_i) of some value (x_i) subject to zero-mean uncorrelated white noise (v_i) with varianceσ_i^2, can be written as:

$$z_i(k) = x(k) + v_i(k), v_i \sim N(0, \sigma_i^2) \tag{14}$$
$$E\big(v_i v_j\big) = 0, i \neq j \tag{15}$$

We can formulate an unbiased estimator (x^*) as the weighted summation of noisy measurements (z_i):

$$x^*(k) = \sum_{i=1}^{n} w_i z_i(k) \tag{16}$$
$$E(x^*(k) - x(k)) = 0 \tag{17}$$

Minimization of the expected error between the estimator and signal recovers the weighted arithmetic average. Subsequent minimization of the mean square error between the estimator and signal produces an ideal weighting based on variance:[34]

$$w_j = \frac{\prod_{l \neq j}^{n} \sigma_l^2}{\sum_{l=1}^{n} \prod_{k \neq l}^{n} \sigma_k^2} \tag{18}$$

However, since the signal confidence calculations already indirectly consider signal variance it is more useful at this juncture to substitute confidence-based weights as they capture a greater amount of data than variance-based weighting alone. A linear confidence weighting, which reduces to zero at the confidence threshold, can be given for *n* redundant sensors by:

$$w_{i,conf} = \frac{C_i(k) - 0.7}{\left(\sum_{j=1}^{n} C_j(k)\right) - (0.7)n}, i = \{1, \ldots, n\} \tag{19}$$

Recall that the low-confidence signals have already been eliminated, so the weighting strategy above gives a normalized positive weighting that satisfies the requirements for the weighted average. The resulting combined air-data measurements are output to the wind-estimation system for the development of the inertial-measurement based auxiliary air-data estimate.

C. Wind Estimation

The goal of the wind estimator is the generation of an air-data vector that can serve as both a reference for ADS failure detection and a failsafe reading to promote safe pilot/autopilot operation of ADS-dependant flight controls in the event of partial/complete air-data sensor failure. The wind estimation scheme has two effective modes: nominal operation, wherein some number of ADS sensors are functional and wind estimation is dominated by direct measurements, and failsafe operation where, in the absence of ADS inputs, winds are estimated from previously collected wind statics. Refinements to this model would require higher fidelity models for both the aircraft and environmental wind processes. Such models are typically more readily available for commercial aircraft than for small UAS. The primary mechanism for wind estimation and, subsequently, air-data vector estimation is a three step process of extracting inertial wind measurements from body-relative sensors, updating the wind estimate, and recovering air-data measures from the wind model and current inertial measurements.

Wind model relationships to body and inertial measurements are expressed by rotation matrices for pitch (θ), roll (ϕ), and yaw (ψ) Euler angles about the *x*, *y*, and *z* axes, respectively:

$$R_x(\phi), R_y(\theta), R_z(\psi) \tag{20a,b,c}$$

The first step in the wind estimation procedure is to resolve an inertial referenced environmental wind measurement from high-confidence air-data vector and vehicle motions. We first develop the measured airspeed ($u(k)$) into a vector in the aircraft body frame (B) using angle-of-attack ($\alpha(k)$) and sideslip ($\beta(k)$):

$$\vec{u}_B(k) = R_z(\beta(k)) R_y(-\alpha(k)) [u(k) \quad 0 \quad 0]^T \tag{21}$$

This vector is subsequently rotated into the inertial frame (I) using the aircraft's roll ($\phi(k)$), pitch ($\theta(k)$), and yaw ($\psi(k)$) Euler angles:

$$\vec{u}_I(k) = R_z(\psi(k)) R_y(\theta(k)) R_x(-\phi(k)) \vec{u}_B(k) \tag{22}$$

To recover an inertial wind measurement ($wind(k)$) we must add the body's inertial frame velocity ($\vec{v}(k)$) to the body relative airspeed vector:

$$wind(k) = \vec{u}_I(k) + \vec{v}(k) \tag{23}$$

The next step is to update the actual wind model. For the small unmanned seaplane we employ a simple spatially-uniform average-based wind model. The locally measured wind (primarily during drift) is accepted as the global wind estimate with a weighted time average of the wind measurements used as a reasonable estimate of the current steady wind. Further the wind is assumed to only have velocity components in a local horizontal plane, that is, there is no vertical component of wind. This set of assumptions are reasonable for the Flying Fish mission as it will transit through only a small range of altitudes (<100m) and over a fairly short distance (<1000m) and time (2-5min) during each flight. With an updated wind model we can construct the inertial-measure-based ADS estimates. First we recover the estimated inertial-frame airspeed vector ($\vec{u}_I^*(k)$) by differencing the wind estimate with the vehicle velocity:

$$\vec{u}_I^*(k) = \overline{wind}(k-1) - \vec{v}(k) \tag{24}$$

Rotating the inertial airspeed vector estimate by the Euler into the body frame produces a body-frame relative airspeed vector estimate ($\vec{u}_B^*(k)$):

$$\vec{u}_B^*(k) = R_x(\phi(k))R_y(-\theta(k))R_z(-\psi(k))\vec{u}_I^*(k) \tag{25}$$

Trigonometry can then be applied to recover the angle of attack and sideslip values:

$$\alpha(k) = \tan^{-1}\left(\frac{u_{B,z}^*(k)}{\left(\left(u_{B,x}^*(k)\right)^2+\left(u_{B,y}^*(k)\right)^2\right)^{1/2}}\right) \tag{26}$$

$$\beta(k) = \tan^{-1}\left(\frac{u_{B,y}^*(k)}{\left(\left(u_{B,z}^*(k)\right)^2+\left(u_{B,x}^*(k)\right)^2\right)^{1/2}}\right) \tag{27}$$

Airspeed is recovered from the magnitude of the body-frame airspeed vector estimate:

$$u(k) = \|\vec{u}_B^*(k)\| \tag{28}$$

D. Algorithm Summary

Unmanned Seaplane ADS Fault Tolerance Algorithm:

1. Signal Fault Detection Block:
- *Input:* Sensor signals: $s_n(k)$
- *Output:* Fault detection votes: $\Lambda_{var}(n,k)$, $\Lambda_{rate}(n,k)$, $\Lambda_{sat}(n,k)$

a) Extract signal characteristics (Eqn. 3-4)
b) Select test parameters for current flight mode
c) Perform fault detection tests:
i. Variance within expected thresholds, vote Pass/Fail: 1/0 = $\Lambda_{var}(n,k)$ (Eqn. 8)
ii. Rate of variance change within expected thresholds, vote Pass/Fail: 1/0 = $\Lambda_{rate}(n,k)$ (Eqn. 9)
iii. Signal response sufficiently far from saturation, vote Pass/Fail: 1/0= $\Lambda_{sat}(n,k)$ (Eqn. 10a,b)

2. Confidence-Discriminate Data-Fusion Block:
- *Input:* Fault votes: $\Lambda_{var}(n,k)$, $\Lambda_{rate}(n,k)$, $\Lambda_{sat}(n,k)$
- *Output:* High-confidence ADS values: $\alpha(k)$, $\beta(k)$, $u(k)$

a) Sensor confidence assessment:
i. Compute probabilities for each failure type: P_{var}, P_{rate}, P_{sat} (Eqn. 11a,b,c)
ii. Compute confidence for each sensor: $C_n(k)$ (Eqn. 12)
b) Sensor voting / outlier rejection:
i. Reject low confidence signals ($C_n(k) < 0.7$)
ii. Reject outliers (Eqn. 13)
c) Confidence-weighted sensor fusion (Eqn. 3, 19)

3. Wind Estimation Block:

Input: High confidence ADS values: $\alpha(k)$, $\beta(k)$, $u(k)$

Output: Estimated ADS values: $\alpha^*(k)$, $\beta^*(k)$, $u^*(k)$

a) Compute k^{th} Wind Estimate: $wind(k)$

i. Resolve airspeed (u) as body-frame vector (Eqn. 21)

ii. Rotate air-data vector into inertial frame (Eqn. 22)

iii. Compute wind (Eqn. 23)

b) Update wind average: $\overline{wind}(k)$ (Eqn. 3)

c) Construct airspeed, AOA, and sideslip estimates from wind estimate:

i. Compute inertial-frame air-vector estimate (Eqn. 24)

ii. Rotate estimated air-vector into body-frame (Eqn. 25)

iii. Determine estimated ADS values from body-frame air-vector (Eqn. 26-28)

IV. Results

The algorithm was first tuned and tested with pre-recorded flight data from the first-generation Flying Fish. This data provides a very good basis for testing and development as it contains a wide variety of ADS failures. The full algorithm was then applied, after tuning, to the second-generation Flying Fish vehicle with its dual heated 5-hole/pitot-static probes and propeller-anemometer ADS. All of these results will be summarized below.

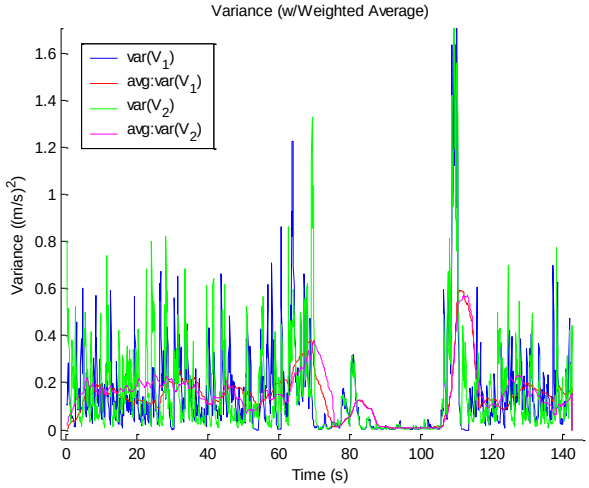

Figure 2. Variance analysis – functional sensor

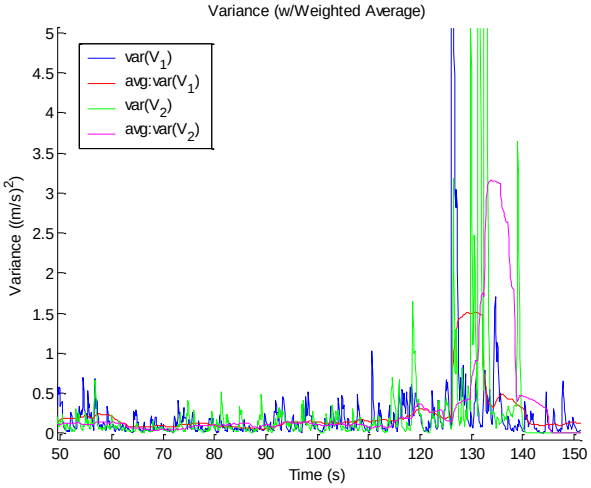

Figure 3. Variance analysis - failed sensor

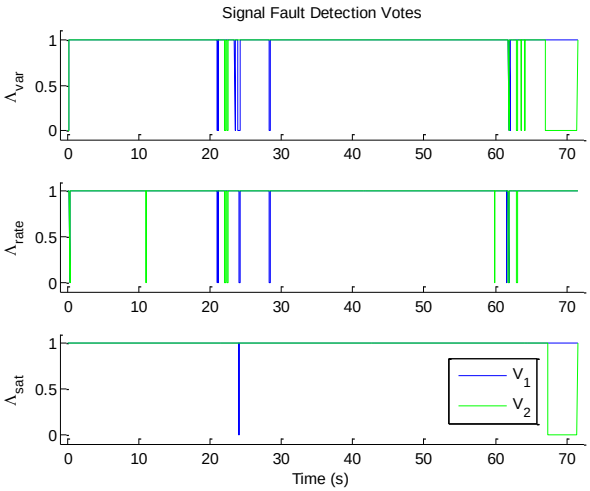

Figure 4. Signal fault detection votes

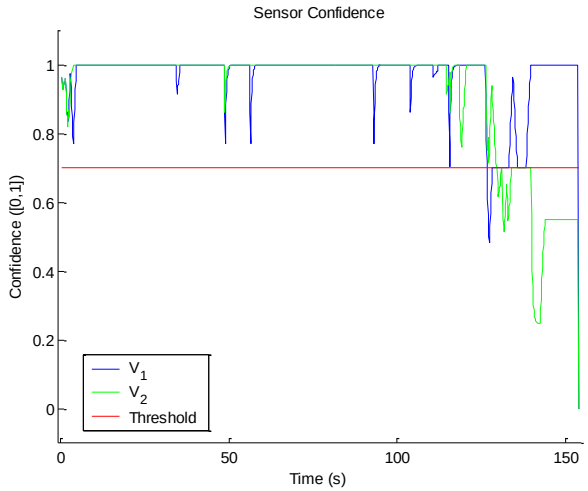

Figure 5. Composite sensor confidence

A. Autonomous Seaplane Example: Generation I

Datasets from our first generation autonomous seaplane provide a wide range of test cases including: single and double in-flight pitot/static failures, intermittent failures, and datasets that begin with failed sensors. As discussed in the algorithm description signal characteristics are extracted from the analysis and comparison of functional (Fig. 2) and failed (Fig. 3) ADS sensors data.

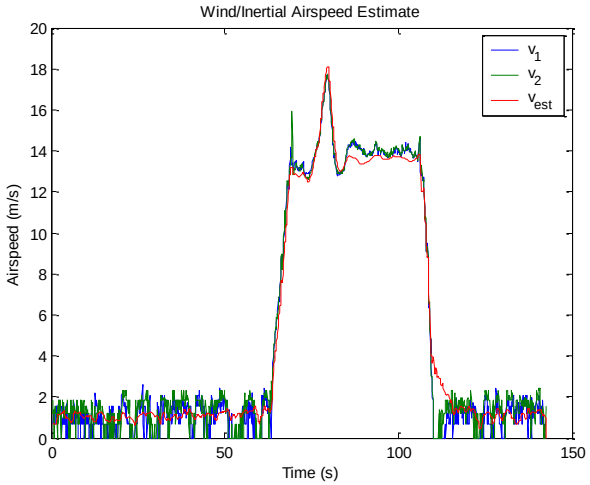

Figure 6. Wind-model airspeed estimate (no-failures)

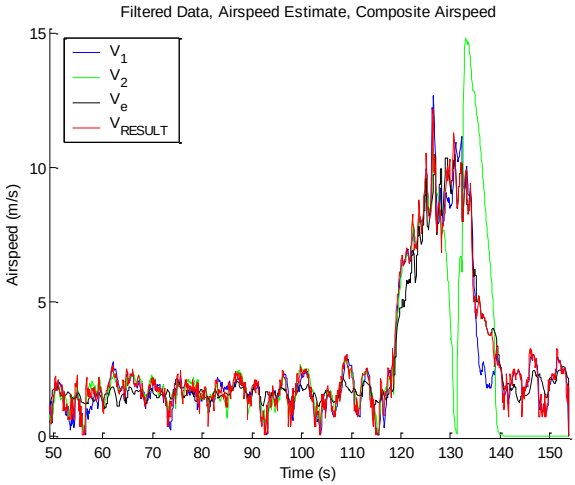

Figure 7. Algorithm results - single probe failure

After tuning the sample the characteristic extraction and signal conditioning parameters the algorithm successfully rendered valid binary votes for the three classes of signal-faults (Fig. 4). The results of the binary fault detection decisions are then combined to determine composite sensor confidence (Fig. 5).

The wind-model air-data estimate shows good correlation with functional ADS sensors (Fig. 6). The complete algorithm, combining the wind model, signal confidence, voting, and data fusion to produce a single high-confidence airspeed has been tested for a wide range of cases including single (Fig. 7) and double (Fig. 8) airspeed sensor failure cases.

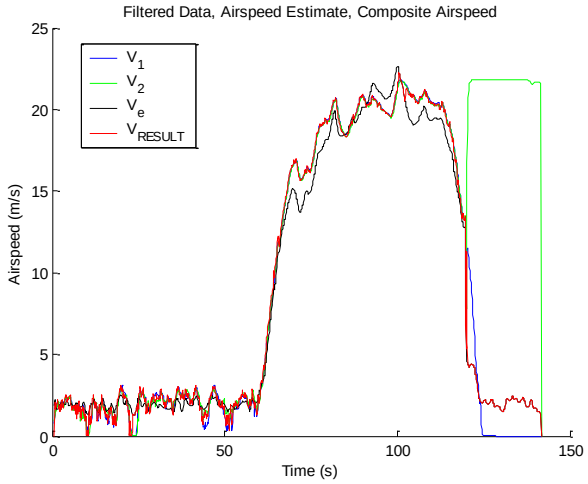

Figure 8. Algorithm results - double probe failure

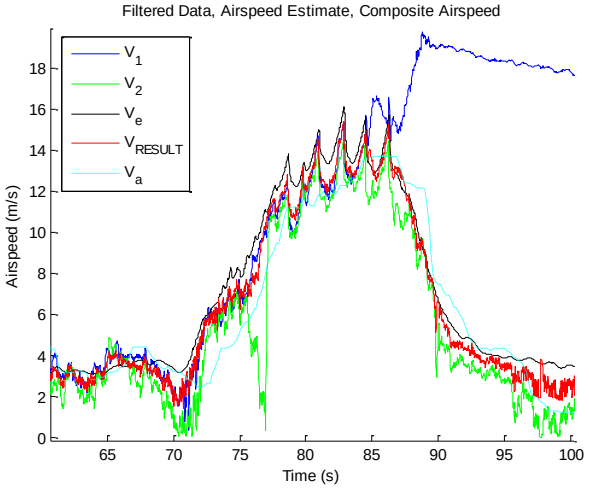

Figure 9. ADS failure during high-speed taxi

B. Autonomous Seaplane Example: Generation II

Following the development of the algorithm with the legacy data, as reviewed above, the algorithm was applied to the second generation flight vehicle. In this section we will present test results from these sensors as well as wind-model airspeed estimates and the final composite high-confidence airspeed results.

The failure mitigation system performed well during preliminary testing. The algorithm accurately handled errors both during high-speed taxi tests (Fig. 9) and during simple flight tests (Fig. 10). During the high-speed taxi test (Fig. 9) the algorithm correctly eliminates an erroneous sensor excursion during the approach to hydroplaning

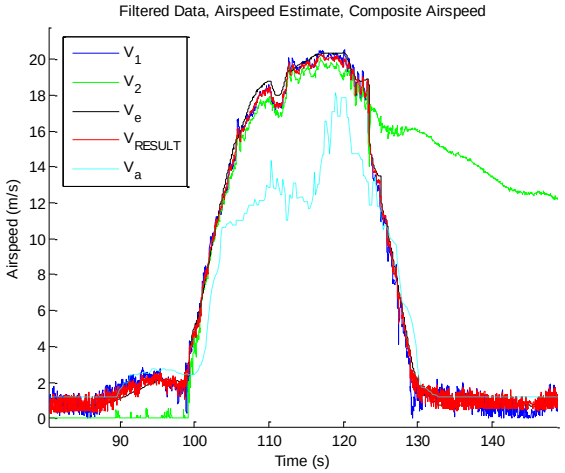

Figure 10. ADS failure during flight testing

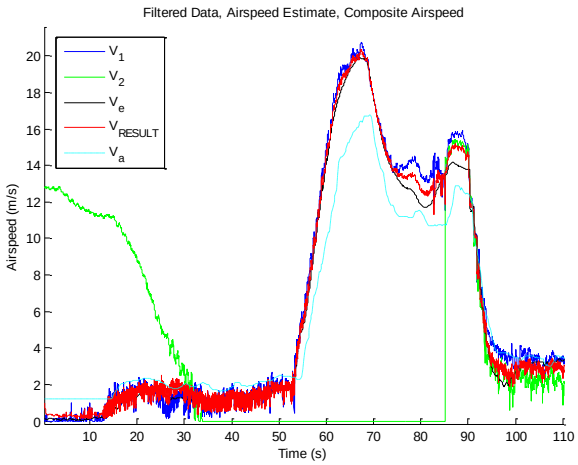

Figure 11. Recovery of initially failed ADS sensor

speeds (75s) and also correctly rejects a high-pressure blockage that biases the one of the airspeeds high after 85s. Flight test results (Fig. 10) show the algorithm correctly rejecting low-speed saturation in (0-100s) and, similar to the taxi test, rejecting a high-pressure biased signal during descent/landing (125s). These tests also demonstrate some interesting dynamics and issues associated with the propeller anemometer. At low airspeed the counting limit of the digital timer and the rotating friction of the propeller produce saturation effects. Conversely at high speeds the sensor response is increasingly non-linear as the small 3cm propeller is driven to rotation speeds in excess of 11000rpm. Furthermore the installed prototype anemometer was subject to greater wear than previously anticipated. Continuous high-speed rotation during flight began to erode the propeller's waterproof bushings resulting in misalignment of the propeller and hall-effect sensors. The team found that while the anemometer could be realigned on shore each morning, giving good results for early flight tests (Fig. 9), the progressive wear and stresses of flight testing resulted in non-negligible signal degradation over the course of a day (Fig. 10). The team anticipates correcting the design issue with an update to the propeller bearing of the miniature anemometer prototype.

Continued testing allowed the algorithm to handle increasingly difficult failures. One of the first major trials for the failure mitigation system was a flight test that began with a blocked pitot-static probe. Consensus voting was able to distinguish the correct signals after sensor confidence was established (Fig. 11) and the signal was successfully reintegrated to the confidence voting procedure when the probes blockage cleared (85s). More impressive however are the results obtained during flight testing in a rainstorm (Fig. 12). The ADS algorithm successfully rejects several erroneous ADS sensor excursions and negotiates a complete sensor failure and two subsequent probe recoveries.

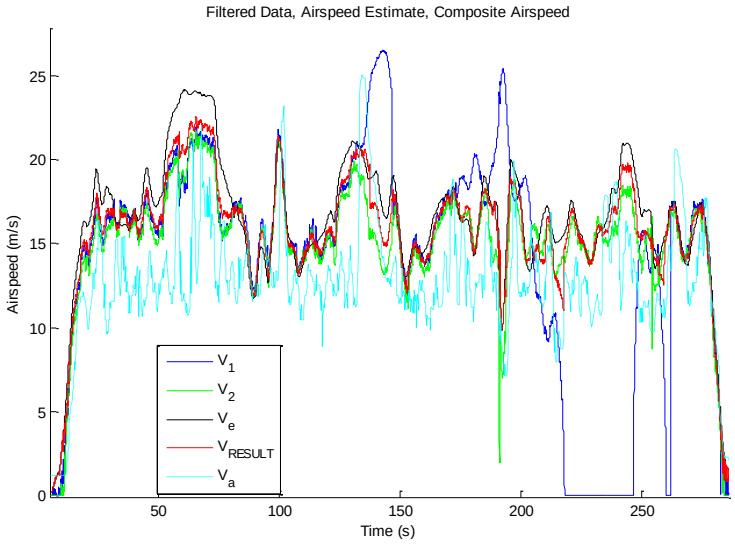

Figure 12. ADS failures during flight through rain

V. Conclusion

Air-data system measurements are critical for robust takeoff performance in an autonomous seaplane and are central to flight control laws across all aircraft classes. Harsh environmental conditions can lead to systematic failures of multiple homogeneous sensors which may not be adequately handled by common fault tolerance

mechanisms designed to handle isolated failures. The algorithm described above has been shown to negotiate a high-confidence air-data vector from sets of redundant air-data sensors subject to failure. The system provides a failsafe capability through the application of an inertial navigation system to a wind model. Algorithm inputs are not limited in number or type except that they must provide a numerically quantifiable signal. Known limitations of the algorithm include heavy reliance on the appropriate selection of fundamental analysis parameters (sample sizes, weighting, and filtering frequencies) and a requirement for accurate knowledge of sensor performance.

Continued efforts are expected to produce a simulation of suitable complexity and fidelity to more completely test the performance of the algorithm. Currently the authors are working to develop appropriate dynamic models for the flight vehicle and atmospheric wind processes in order create a high fidelity wind estimation system with increased estimation capability and fewer constraints. Furthermore we hope to explore a combined implementation of the dynamic-model air-vector estimation mechanisms discussed in references 5, 18, and 19 to produce higher accuracy in-flight estimates of the ADS states and better handle probe failure. Perhaps the most important pending work is the analysis and automation of parameter selection for the tuning of the ADS confidence filter system. While this process was fairly straightforward for the Flying Fish vehicles, for which our algorithm proved robust to variations in the tuning parameters, the broader application of this method would require greater understanding/quantification of the effects of modifying filter parameters. A more robust and general-purpose application of our filter may benefit from automated parameter selection or adaptation to yield optimal environmentally-tuned results.

Acknowledgments

The authors would like to thank the faculty and technical staff of the University of Michigan's Marine Hydrodynamics laboratory and of the faculty and staff of the Department of Aerospace Engineering for collaborating on and supporting the Flying Fish project. We would also like to extend our thanks to our recovery pilot, Derrick Yeo, and to the University of Michigan Biological Station at Douglas Lake for supporting our flight testing program. This work was funded in part under contract N00421-06-C-0094 from the Defense Advanced Research Projects Agency (DARPA).

References

[1]United States. National Aeronautics and Space Administration, Federal Aviation Administration, Incident Report Summaries. Aviation Safety Reporting System (ASRS) Database. January 2010 <http://asrs.arc.nasa.gov/search/database.html>.

[2]United States. National Transportation Safety Board, Investigation Reports. Aviation Accident Database. Jan 2010 <http://www.ntsb.gov/AVIATION/Aviation.htm>.

[3]United States. Federal Aviation Administration, Incident Data Records. Aviation Safety Information Analysis and Sharing: Accident/Incident Data System (AIDS). January 2010 <http://www.asias.faa.gov >.

[4]Flight Safety Foundation. Aviation Safety Network, Accident Descriptions ASN Aviation Safety Database. 18 January 2010 <http://aviation-safety.net/database/>.

[5]Major McLaren, S., "Velocity Estimate Following Air Data System Failure," Thesis. Air Force Institute of Technology. United States Air Force, 2009. Report No.: AFIT/GAE/ENY/08-M21.

[6]Macy, D., Eubank, R., Atkins, E., Bernal, L., Washabaugh, P., Meadows, G., Wild, N., Smith, D., Van Sumeren, H., "Flying Fish: A Persistent Ocean Surveillance Buoy with Autonomous Aerial Repositioning," *AUVSI Conference*. AUVSI. San Diego, CA, June 2008.

[7]Eubank, R., Atkins, E., Macy, D., "Autonomous Guidance and Control of the Flying Fish Ocean Surveillance Platform," *Infotech@Aerospace Conference.* AIAA. Seattle, WA, April 2009.

[8]Meadows, G., Atkins, E., Washabaugh, P., Meadows, L., Bernal, L., Gilchrist, B., Smith, D., VanSumeren, H., Macy, D., Eubank, R., Smith, B., Brown, J., "The Flying Fish Persistent Ocean Surveillance Platform," *Unmanned Unlimited Conference*. AIAA. Seattle, WA, April 2009.

[9]Bureau d'Enquêtes et d'Analyses (BEA, Office of Investigate and Analysis, France), "Interim Report on the Accident on 1st June 2009 to the Airbus A330-203 Registered F-GZCP Operated by Air France Flight AF 447 Rio de Janeiro – Paris," BEA, NTSB, CENIPA, AAIB, BFU, July 2009. Report No.: f-cp090601ae.

[10]Bureau d'Enquêtes et d'Analyses (BEA, Office of Investigate and Analysis, France), "Interim Report n°2 on the Accident on 1st June 2009 to the Airbus A330-203 Registered F-GZCP Operated by Air France Flight AF 447 Rio de Janeiro – Paris," BEA, NTSB, CENIPA, AAIB, BFU, December 2009. Report No.: f-cp090601ae2.

[11]Sghairi, M., de Bonneval, A., Crouzet, Y., Aubert, J., Brot, P., "Challenges in Building Fault-Tolerant Flight Control System for a Civil Aircraft," International Journal of Computer Science. Vol. 35, Is. 4. IAENG. November 2008.

[12]Brière, D., Favre, C., Traverse, P., "A Family of Fault-Tolerant Systems: Electrical Flight Controls, from Airbus A320/330/340 to Future Military Transport Aircraft," *Microprocessors and Microsystems*, Vo. 19, No. 2. Amsterdam: Elsevier, March 1995.

[13]Yeh, Y., "Triple-Triple Redundant 777 Primary Flight Computer," *Aerospace Applications Conference*, Vol.1. IEEE (Feb 1996): 293-307.

[14]Johnson, C. W., "The Dangers of Interaction with Modular and Self-Healing Avionics Applications: Redundancy Considered Harmful," *27th International System Safety Conference.* System Safety Society. Aug 2009.

[15]Houck, D., Atlas, L., "Air Data Sensor Failure Detection," *17th Digital Avionics Systems Conference*, Vol.1. AIAA/IEEE/SAE (October 1998): D17/1-D17/8.

[16] Rohloff, T., Whitmore, S., Catton, I., "Fault-Tolerant Neural Network Algorithm for Flush Air Data Sensing," *Journal of Aircraft*, Vol. 36, No. 3. AIAA. (May–June 1999): 541-549.

[17]Calia, A., Poggi, V., Schettini, F., "Air Data Failure Management in a Full-Authority Fly-By-Wire Control System," *International Conference on Control Applications*. IEEE (October 2006): 3277-3281.

[18]Westhelle, C. H., "X-38 Backup Air Data System (AeroDAD)," *40th AIAA Aerospace Sciences Meeting*. AIAA. Reno, NV. January 2002.

[19]Colgren, R., Frye, M. Olson, W., "A Proposed System Architecture for Estimation of Angle-Of-Attack and Sideslip Angle," *Guidance Navigation and Control Conference.* AIAA (August 1999) 743-750.

[20]Kornfeld, R.,Hansman, R. Deyst, J., "The Impact of GPS Velocity Based Flight Control on Flight Instrumentation Architeccture," Report and Thesis. International Center for Air Transportation. Massachusetts Institute of Technology. 1999. Report No.: ICAT-99-5.

[21]Park, J., Ro, K., "A Prototype Design, Test and Evaluation of a Small Unmanned Aerial Vehicle for Short-range Operations," *3rd "Unmanned Unlimited" Technical Conference*. AIAA. Chicago, IL. September 2004.

[22]Jafarizadeh, M.A., Hassannejad, R., Ettefagh, M.M., Chitsaz, S. "Asynchronous Input Gear Damage Diagnosis Using Time Averaging and Wavelet Filtering," *Mechanical Systems and Signal Processing 22*. Amsterdam: Elsevier (2008): 172-201

[23]Chethan, P., "Blind Fault Detection Using Spectral Signatures," Thesis. Louisiana State University. 2003.

[24]Jiang, J., Yang, J, Lin, Y., Liu, C., Ma, J. "An Adaptive PMU Based Fault Detection/Location Technique for Transmission Lines Part I: Theory and Algorithms," *Transactions on Power Delivery*, Vol. 15, No. 2. IEEE (April 2000): 486-493.

[25]Willsky, A., "A Survey of Design Methods for Failure Detection in Dynamic Systems," *Automatica*, Vol. 12. Oxford: Pergamon Press (1976): 601-611.

[26]Yong Liu, Yi Shen, Hengzhang Hu, "A New Method for Sensor Fault Detection, Isolation and Accommodation," *16th Instrumentation and Measurement Technology Conference*, Vol.1, IEEE (May 1999): 488-492.

[27]Klein, L.A., "A Boolean Algebra Approach to Multiple Sensor Voting Fusion," *Transactions on Aerospace and Electronic Systems*, Vol. 29, No. 2. IEEE (April 1993): 317-327.

[28]Omana, M., Taylor, J.H., "Fault Detection and Isolation Using the Generalized Parity Vector Technique in the Absence of a Mathematical Model," *Conference on Control Applications*, Singapore. IEEE. October 2007.

[29]Alvergue, L., Li, X., Aravena, J., "Overinstrumented Systems," *Region 5 Technical Conference*. IEEE (April 2007): 326-330.

[30]Caliskan, F., Hajiyev, C.M., "Aircraft Sensor Fault Diagnosis Based on Kalman Filter Innovation Sequence," *37th Conference on Decision and Control*, Vol. 2. IEEE (December 1998): 1313-1314.

[31]Barnett, V., Lewis, T. Outliers In Statistical Data. 3rd Ed. New York: John Wiley & Sons, 1994.

[32]Hall, D.L., Llinas, J., "An Introduction to Multisensor Data Fusion," *Proceedings of the IEEE*, Vol. 85, No. 1. (January 1997): 6-23.

[33]Bloch, I., "Information Combination Operators for Data Fusion: A Comparative Review with Classification," *Proceedings, Image and Signal Processing for Remote Sensing*, Vol. 2315. SPIE (December 1994): 148-159.

[34]Thompson III, A.A., "Data Fusion for Least Squares," Aberdeen Proving Ground, MD: Ballistic Research Laboratory - U.S. Army Laboratory Command, December 1991.