Design for Resilience in Autonomous Systems: Lessons Learned from Controlled Flight into Stall Accidents

Lance Sherry[1]
Center for Air Transportation Systems Research at George Mason University, Fairfax, VA, 22031, USA

Robert Mauro[2]
University of Oregon. Eugene, Oregon, 97403, USA

Resilience is the means to extend the design reliability of a system beyond the design assurance achieved by the technology. In the case of supervised autonomous systems, resilience can be achieved by the intervention of a human operator when the autonomous system creates an undesired state. This paper describes a detailed analysis of the requirements and the design for intervention in the operation of an autonomous function on a modern airliner. The analysis derives the requirements for the intervention task from an analysis of the six step scenario leading to the Controlled Flight into Stall (CFIS) accidents. Some of the requirements that cannot be performed adequately by the human operator (i.e. monitoring for rare events, complex calculations, and correlation of disparate data) are automated in a stand-alone device on the flight deck. The functional design of this intervention support device, known as the Paranoid Pilot Associate, is described. Limitations and implications of the design are discussed.

I. Introduction

For the foreseeable future, humans will remain responsible for the behavior of autonomous systems, no matter how distant in time or space they are from the system. At a minimum, humans will be expected to curtail or abort operations should the autonomous system endanger human life or property.

As the trend towards increased autonomy grows, fundamental questions about the proper role of the human "supervisor" in autonomous systems operation, and the design of the system in support of the supervisory task must be answered.

This paper examines the role of the flight crew in a modern airliner (i.e. Part 121) in monitoring and intervening in operations performed autonomously by the flight deck automation. Specifically, the analysis examines, the design of the "flight deck system" (i.e. flight crew and automation) in protecting the aircraft on the low end of the speed envelope from an aerodynamic stall[1]. An analysis of accidents, classified as Controlled Flight into Stall (CFIS), revealed that the accidents were in part the result of autonomous functions on the flight deck (i.e. speed envelope protection) that failed to perform their intended functions in rare circumstances. These failures were not the result solely of mechanical, electrical, or structural failures, but rather were Functional Complexity Failures. That is, the complexity of the system architecture and behavior of its functions.led to the CFIS. Due to the complex interactions involved in these failures, they are difficult to detect during operations, and difficult to avoid during the design phase. To mitigate this class of failure, and close the gap between the desired operational hazard reliability (i.e. 10^{-9}) and the design assurance for the automation (10^{-5}), human operators (i.e. flight crew) are expected to intervene. However, the human intervention task is neither explicitly trained nor supported by the design of the automation. As a result, it cannot be performed in the manner required to achieve the desired operational hazard reliability.

The analysis in this paper identified the following functions required to perform the intervention task for the CFIS scenario:

---

[1] Associate professor, System Engineering and Operations Research Department, George Mason University, 4400 University Drive, MS: 4A6, Fairfax, VA. 22030, AIAA Member

[2] Associate Professor, Department of Psychology, University of Oregon, 1585 E 13th Ave, Eugene, Oregon, 97403.

1. Identify equipment limitations
2. Correlate equipment limitations with aircraft and environment conditions
3. Compare flight crew intent with automation intent
4. Anticipate thrust command and thrust change points
5. Anticipate changes in airspeed
6. Annunciate flight deck data integrity

Humans are poorly suited to performing many of the functions identified in these requirements (e.g. monitoring for rare events, non-linear calculations, and correlation of disparate data); hence, automation is required to support the human operator. This paper describes the functional design of a Paranoid Pilot Associate (PPA) with four functions: (1) Multi-source Information Correlator: (MsIC), (2) Intention Comparator (IC), (3) Thrust Advance and Airspeed Change Indicator (TA & AC I), and (4) Flight deck Data Integrity Monitor (FDIM).

The implication of this analysis is that autonomous systems must be designed to meet *resilience* standards in addition to traditional *reliability* standards. Whereas reliability standards define the ability of a system or component to perform its required functions under ***stated*** *(i.e. expected)* ***conditions*** for a specified period of time, resilience standards define the intrinsic ability of a system to adjust its functioning prior to, during, or following changes and disturbances, so that it can sustain required operations under both ***expected*** and ***unexpected*** conditions.

In this way, *Design-for-Resilience* is a critical design practice that must be developed to enable the next generation of autonomous systems to handle expected and unexpected conditions. Design-for-Resilience must explicitly address the possibility that the automation may generate inappropriate commands as a result of Functional Complexity Failures (FCF) and define explicit intervention tasks that are appropriately allocated between human and machine. Furthermore, if the tasks allocated to humans require them to perform activities that humans do not perform reliably, then the tasks must be explicitly supported by the automation. For example, humans should not be allocated tasks that require them to monitor for rare events, decode disparate data sources without information on the data integrity, or correlate disparate data without automation support. These functions can be automated by Paranoid Associate (PA) technology. The Paranoid Pilot Associate (PPA) described in this paper is an instantiation of PA technology for the flight deck but it could be applied to any autonomous system.

This paper is organized as follows: Section II provides a summary of the scenarios leading to Controlled Flight Into Stall accidents. Section III describes the tasks required for CFIS interventions. Section IV describes the requirements for and design of flight deck automation in support of CFIS Interventions. Section V discusses lessons learned and makes the case for Resilience-by-Design to achieve the desired safety levels for autonomous systems.

II. Controlled Flight into Stall (CFIS) Accidents

We analyzed twenty-two accidents and incidents that occurred over the last two decades in which a modern airliner experienced a deceleration through 1.3 V_{Stall} all the way to V_{Stall}. In some cases, the energy-state could not be recovered and the aircraft crashed. In other cases, the stall recovery procedure brought the aircraft back to a safe energy state. These accidents and incidents were characterized by a structurally, mechanically, and electronically sound commercial airliner decelerating through the minimum safe operating speed (1.3 V_{Stall}) to the stick-shaker stall speed[1]. In these cases, the automation did not fail. Rather, the automation generated an inappropriate command that flew the aircraft into the stall.

The scenarios that lead to the inappropriate commands in these accidents, followed the pattern shown in Figure 1: (1) a triggering event (e.g., sensor discrepancy, flight crew entry) (2) before or during a deceleration to the minimum safe operating airspeed,(3) resulted in a mode change or change in automation engagement status, that (4) lead to an inappropriate command, that in turn (5) resulted in an inappropriate trajectory that (6) lead to a CFIS.

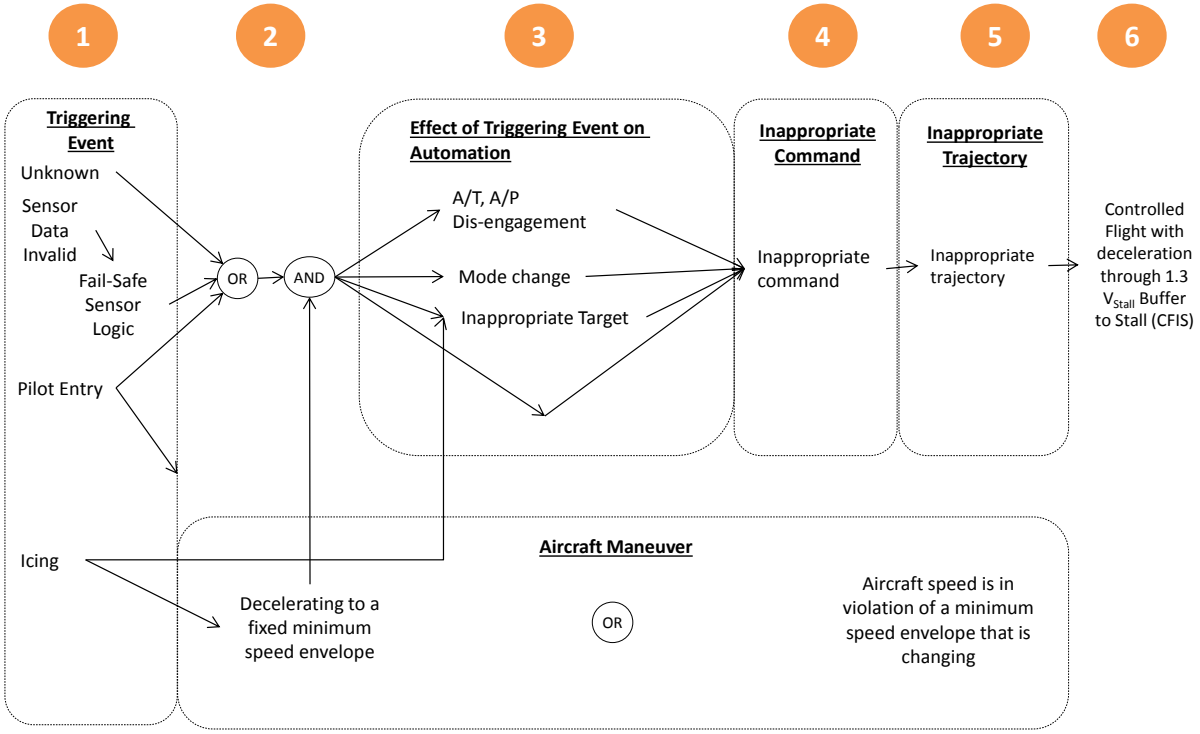

Figure 1. Scenario for Controlled Flight into Stall (CFIS).

The subset of CFIS accidents not related to changes in aerodynamic properties of the aircraft (e.g. icing) and related errors in the computation of the aircraft speed envelope, were characterized by **the automation no longer actively controlling to the airspeed target**. There were two effects of the triggering events on the automation that

Table 1: CFIS Accidents When the Automation Was No Longer Actively Controlling Airspeed

Effects of Triggering Events on Automation	Accidents and Incidents
Automation Engaged, but No Active Speed Control	OZ 214 (dormant mode) TA 1951 (fixed thrust mode)
Automation No Longer Coupled (i.e. "Engaged") for Speed Control	AAL 903 ThomsonFly - Bournemouth ThompsonFly - Belfast ThomsonFly, no location specified Provincial Airlines Midwest 490 Air France 447

yielded this condition. The triggering event led the automation to either: 1) disengage speed control or 2) transition to a mode that did not control to the speed target. These conditions are summarized in Table 1 and described below.

A. Automation Not Engaged

In several accidents, the automation autonomously disengaged, no longer coupling the automation commands to the control surfaces and engines. American Airlines (AAL 903) is an example of this scenario. AAL 903 was instructed to hold at 16,000' due to a weather cell on the arrival procedure. The aircraft decelerated and leveled-off as it made the turn into the holding pattern. During this period, with the Autothrottle commanding idle thrust, the Autothrottle, for reasons unknown, autonomously disengaged. Without the Autothrottle coupled to the engines, and the thrust set to idle, the aircraft decelerated well below the minimum safe operating speed (1.3 V_{Stall}).

B. Automation Engaged, but no Active Speed Control (Inappropriate Mode)

In several accidents, the effect on the automation of the triggering events was to cause a transition to an Autopilot or Autothrottle control mode combination that did not directly control airspeed. The Turkish Airlines (TA)

1951 accident is an example of this phenomenon. The aircraft was vectored for an instrument approach that resulted in a late localizer capture and subsequent clearance to descend that required the aircraft to decelerate to the landing airspeed from a condition "high and fast" above the glideslope.

After receiving a clearance to land, the flight crew instructed the automation to descend to recapture the glideslope from above *and* to decelerate to the landing speed. When the aircraft started the descent form 2000', a latent failure in the Captain's side Radio Altimeter (RA) sensor erroneously caused the Autothrottle to assume the aircraft had landed (the RA read -8 feet) while the First Officer's RA showed the correct altitude. This caused the Autothrottle to transition to the Landing Flare mode. In this mode the throttles are locked at the idle position and *do not control airspeed*. In this configuration, the aircraft decelerated well below the minimum safe operating speed ($1.3\ V_{Stall}$).

Similarly, in the Asiana Air 214 (OZ 214) accident, the flight was vectored for a visual approach at 14 nm to the runway at an altitude that left the aircraft above the desired three-degree glideslope. At 5nm to the runway, the aircraft remained well above the desired glide path. In an attempt to increase the airplane's descent rate and capture the desired glide path, the flight crew selected an Autopilot mode (Flight Level Change Speed [FLCH SPD]) that due to setting the clearance altitude above the current aircraft altitude (in preparation for a missed approach), resulted in the autoflight system initiating a climb. The flight crew disconnected the Autopilot and moved the thrust levers to idle. This action signaled the Autothrottle to transition to a "dormant" mode in which the Autothrottle does not directly control airspeed, but allows the flight crew to manually adjust the flight trajectory by setting the thrust. In this configuration, the aircraft decelerated well below the minimum safe operating speed ($1.3\ V_{Stall}$).

III. Flight Crew Intervention in CFIS

In each of the CFIS accidents, the flight crew, was *not* able to successfully monitor for the rare event and hence the pilots were unable to detect the problem and intervene in a timely manner. There are two main reasons for this. First, humans are unreliable monitors of rare events under any conditions. In these cases, the ability of the flight crews to monitor successfully was further impaired by the high workload., The flight crews were busy performing the multiple complex tasks related to the approach and landing. Second, due to the design of the flight deck automation, the information needed to detect, identify the problem, and intervene was not clearly displayed on the flight deck. To detect and identify the CFIS scenarios in current aircraft cockpits, the pilots must have a deep understanding of automation behavior that is not explicitly provided in standard training material and they must retrieve rarely used facts from memory.

A. Recognizing CFIS Trajectory

In all of these cases, the aircraft trajectory (e.g., inappropriate deceleration) associated with the inappropriate automation commands could not be detected or recognized as it was masked by an appropriate trajectory (i.e. deceleration) leading up to the speed envelope violation.

B. Recognizing the CFIS Modes and Triggering Events

Further, the cues identifying the triggering events and their effect on automation modes and engagement status did not lead to detection, recognition, or diagnosis as the available were hidden from the flight crew by the absence of explicit annunciation. To correctly recognize the scenario required inferences based on memorized rules for interpreting data on the displays and discerning subtle differences associated with overloaded mode labels.

As described above, to detect and diagnose the CFIS scenario answers to the following questions must be obtained:

1) Is airspeed control active? If airspeed control is <u>not</u> active, the pilots must then determine why not:
2) What is the control strategy (i.e. mode) of the airspeed?
3) What is automation configuration?

On the flight deck, the primary source of information on engagement status, target, mode, and aircraft energy-state is the Primary Flight Display (PFD). See Figure 2. This information is distributed across the PFD as summarized in Table 2.

Figure 2: Primary Flight Display (PFD) Displays Information that is an Integrated Form of Formerly Federated Gauges.

Table 2: Summary of Location on PFD for Information Required to Detect, Recognize, and Diagnose CFIS Scenarios.

Information	Location on Traditional PFD	Memorized Rules Required to Interpret
Control active?	FMA	Label interpretation Location of Autothrottle and Autothrottle labels and presence/absence of Autothrottle and Autothrottle Label. Color of FMA and Targets
Engagement Status	Below FMA for Autothrottle and Autothrottle FMA and Targets for VNAV and LNAV	Location of Autothrottle and Autothrottle labels and presence/absence of Autothrottle and Autothrottle Label. Color of FMA and Targets
Mode selected	FMA	Label interpretation
Target selected	Tape	None
Actual	Tape	None

1. Is Airspeed Control Active?

The PFD does not explicitly address this question. The aircraft automation can control airspeed by pitch using the Autopilot or by thrust using the Autothrottle. To determine whether the Autopilot and/or Autothrottle are engaged, the pilot must read the Flight Mode Annunciator (FMA). Typically, when the Autopilot or Autothrottle is engaged, the corresponding mode annunciation on the FMA will be displayed in green. When the Autopilot or Autothrottle is engaged, some FMA's will also display "A/P" and/or "A/T" in green. If the Autopilot or Autothrottle is armed or engaged, the respective buttons on the Mode Control Panel (MCP) may be lit.

However, the automation may be engaged (i.e. Autopilot and Autothrottle coupled), but not in a mode that actively controls to the displayed target. For example, the TA 1951 accident the Autothrottle was in Landing flare mode that locked the throttles to idle thrust. Likewise in the OZ 214 scenario, the Autothrottle was in a "dormant"mode allowing the flight crew to set thrust directly via manual adjustment of the Throttle Lever.

Furthermore, the speed *target* is displayed on the left hand side of the PFD – but without any indication whether or not the automation is actively controlling to that target.

2. What is the Control Strategy (i.e. mode)?

The Autopilot and Autothrottle modes selected are displayed on the FMA. The center FMA label identifies the mode selected to control the lateral path (i.e. course, heading, VOR, localizer). The left and right labels identify the modes controlling the vertical axis. There are two designs used in modern airline flight decks for the vertical labels. One design indicates which control surface is controlling speed or altitude. For example, speed can be controlled by

PITCH or by THRUST. The other design indicates what parameter is controlled by the thrust and pitch. For example, pitch can control SPEED, ALTITUDE, Vertical Speed (VS), or PATH. The labeling of the modes and exclusive use of labels for each mode is critical to recognition and interpretation of the mode.

3. What is the Automation Configuration?

The engagement status of the Autopilot and Autothrottle is displayed under the FMA. This information is also duplicated on the MCP by backlit push buttons. When the Autopilot and Autothrottle disengage, an aural alert is generated, and the Altitude, Heading, and Speed labels are surrounded by an A/P OFF rectangle, or an A/T OFF rectangle, for approximately 10 seconds (not shown in Figure 2). For example, in the AAL 903 scenario, absence of A/T under the FMA, flashing rectangle in the FMA, and absence of backlit on A/T button would have indicated this condition.

IV Requirements and Design of Flight Deck for Intervention in CFIS Tasks

Analysis of the CFIS accidents scenarios identified a sequence of intervention opportunities for the flight crew. As discussed above, even if the flight crew was able to devote sufficient attention to the monitoring task, the flight deck automation does not directly support the flight crew in a meaningful way to detect and respond to the CFIS scenario. This section identifies the requirements for an explicit design for intervention in the CFIS scenario. The requirements and the proposed design solution are summarized in Table 3.

A. Requirements

The set of flight crew tasks that would have to be performed to detect each step in the CFIS scenario is listed below and summarized in Table 3.

1. Identify operational limitations: To determine whether there are operational limitations that could affect a flight, the crew must check for manufacturers' bulletins, advisory circulars, NOTAMs, and other documentation describing limitations on the navigation procedures to be performed (e.g. ILS approach) and the flight deck functions and aircraft equipment that will be used in these procedures. For example, in the case of TA 1951, there would be a limitation in the use of the Autothrottle in final approach with a Radio Altimeter discrepancy. For the case of AF 447, there would be a limitation in relying on air data for pitots operating in Ice Super Saturated Regions (ISSR).
2. Correlate equipment limitations with aircraft and environment conditions: To effectively use the bulletins and other documentation, the flight crew must cross correlate data from aircraft maintenance logs, weather forecasts, etc. with recorded limitations. For example the aircraft in TA 1951 had a long history of Radio Altimeter problems. In the case of AF 447, the flight crew would have had to correlate data from weather forecasts with information on limitations to determine the effect of ISSR along their intended route of flight on their equipment.
3. Compare flight crew intent with automation intent: To avoid divergence in intentions and the resulting flight trajectory, the flight crew must monitor changes in automation targets and modes and detect when the automation migrates away from the flight crew intentions and/or the navigation procedure. For example, in TA 1951 the automation migrated to a "land" mode while the aircraft was at 2000' AGL. In the OZ 214 accident, the autothrottle transitioned to a dormant mode while the pilots expected it to be active.
4. Anticipate Thrust Command and Thrust Change Points. Due to the importance of thrust commands and thrust settings, the flight crew should monitor thrust settings at the time when the thrust is required to change to achieve the desired flight trajectory. In all of the CFIS scenarios involving the Autothrottle, the flight crew failed to check whether the thrust advanced as the airspeed approached the target airspeed.
5. Anticipate Changes in Airspeed: Due to the importance of airspeed, flight crews should monitor airspeed and more importantly changes in airspeed (i.e. deceleration rates). A change in deceleration would be evident as the airspeed reached the target airspeed.
6. Flight deck Data Integrity: To assess the overall state of the flight deck automation, the flight crew should check the integrity of the data used by the automation. If erratic data (e.g. AF 447), or discrepant data (e.g. TA 1951) were detected, the flight crew's assessment of the situation would change and trigger an intervention.

Table 3: Summary of Flight crew tasks to detect CFIS Scenarios and the Current and Proposed Flight deck support required.

Flight crew Task to Intervene in CFIS	Category of Flight crew Task	Current Flight Deck Support on Flight crew Task	Proposed Solution
1. Identify equipment limitations	Monitor	None	Automated monitoring
	Recall disparate/remote information from multiple sources	None	Access to information not available on aircraft
			Algorithms to correlate relevant information from multiple sources
2. Correlate equipment limitations with aircraft and environment conditions	Monitor	None	Automated monitoring
	Recall and recognize complex interactions from multiple disparate/remote sources	Training/ "Airmanship"	Access to information not available on aircraft
			Algorithms to correlate relevant information from multiple sources
3. Compare flight crew intent with automation intent	Monitor	None	Automated monitoring
	Correlate FMA and other information with deep knowledge of automation behavior and overloaded labels	Training/ "Airmanship"	Algorithms to correlate information from multiple sources
4. Anticipate Thrust Command and Thrust Change Points	Monitor	None	Automated monitoring
	Identify thrust change points	Training/ "Airmanship"	Algorithms to identify thrust change points
5. Anticipate Changes in Airspeed	Monitor	Some low speed alerting	Automated monitoring
	Identify airspeed change points		Algorithms to identify airspeed change points and their associated decelerations
6. Flight deck Data Integrity	Monitor	None	
	Calculate data integrity from multiple sources on the aircraft	Training/ "Airmanship"	Status of all sensors and fail-safe sensor logic

As shown in Table 3, all the tasks require monitoring for rare events. It is a well known fact that humans are poor monitors especially for rare events. Several of the tasks require retrieval of information that is not always available on the aircraft (e.g. maintenance logs, bulletins, weather data). Others tasks require complex non-linear calculations (e.g. checking data integrity, determining thrust and airspeed change points) or correlation of disparate data sets (e.g. correlating automation and flight crew intent, correlating limitations with aircraft equipment and flight plan status). All of these tasks are prohibitively difficult to perform while simultaneously performing all of the complex tasks required to configure the aircraft and shoot an approach.

B. Design of Automation to Support Operator Intention: The Paranoid Pilot Associate (PPA)

A clean sheet design of the flight deck system would naturally automate the tasks that cannot be performed by the human flight crew:

- Monitoring for rare events
- Retrieval of off-aircraft data

- Correlation of disparate data from multiple sources
- Complex non-linear calculations

The Paranoid Pilot Associate (PPA) is an instantiation of the Paranoid Associate for the purpose of the CFIS intervention. The PPA is a stand-alone automated device that follows the aircraft trajectory and based on a data-base of scenarios would provide a probabilistic alert to the flight crew with specific instructions on identifying the issues that require attention at that time. For example, any time the vehicle is decelerating to a minimum safe operating airspeed, the PPA would alert the flight crew to the CFIS scenario. Another condition that would trigger an alert is any change in sensor status or sensor selection performed by fail-safe sensor logic. This capability provides automated monitoring and alerting, performs complex non-linear calculations, and complex correlations from disparate data sets.

The PPA, illustrated in Figure 3 would be connected to multiple off-aircraft data-bases including: aircraft maintenance logs, manufacturers bulletins/advisory circulars, NOTAMs, airport/airspace specific blogs, as well as weather and other pertinent flight data. The PPA also would have access to the full suite of sensor data and the associated fail-safe logic. The PPA also would have direct access to the flight plan, the mode logic, the pitch and thrust commands, and the flight trajectory data.

The PPA functions to support the requirements in Table 3, are illustrated in the functional component diagram in Figure 3: (1) Multi-source Information Correlator (MsIC), (2) Intention Comparator (IC), (3) Thrust Advance and Airspeed Change Indicator (TA & AC I), and (4) Flight deck Data Integrity Monitor (FDIM).

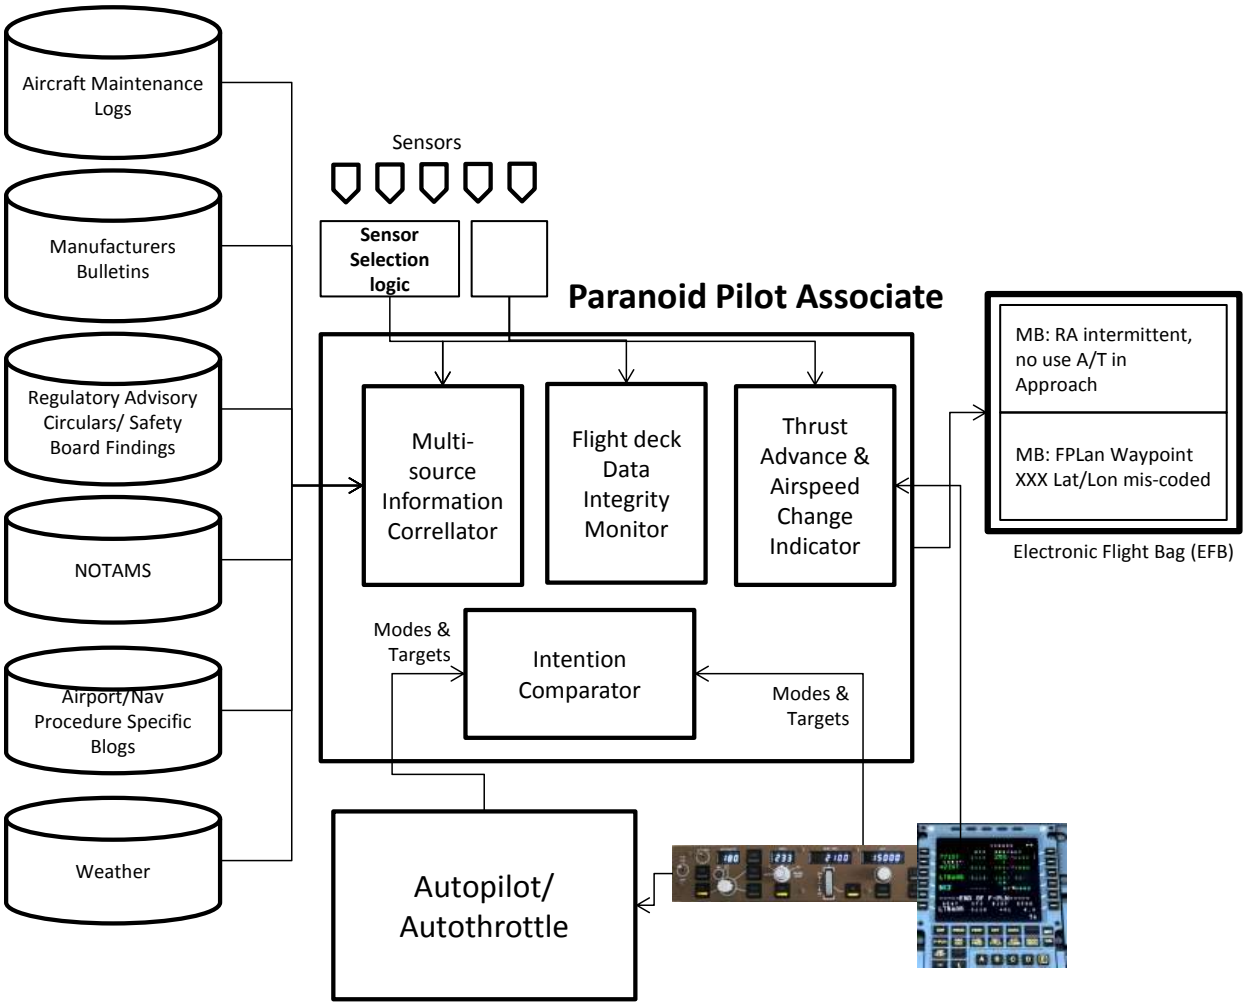

Figure 3: Functional Components of the Paranoid Pilot Associate

1. Multi-source Information Correlator: (MsIC)

The purpose of the MSIC is to identify equipment limitations and correlate equipment limitations with aircraft and environment conditions. To avoid using automation functions and or equipment that is inappropriate, the PPA shall check for manufacturers bulletins, advisory circulars, NOTAMs, and other documentation describing limitations on the navigation procedures to be performed (e.g. ILS approach) or the flight deck functions and aircraft equipment that will be used in the procedures. This massive data-mining/correlation analysis would:

1. Check manufacturer's bulletins and advisory circulars along with maintenance logs for any equipment limitations. For example, in the case of the TA 1951, there could be a limitation in the use of the Autothrottle on final approach with a Radio Altimeter discrepancy. This information coupled with a long history of Radio Altimeter issues on this tail number (i.e. specific aircraft) would provide an added level of attention.
2. Check manufacturer's bulletins and advisory circulars along with maintenance logs along with weather data to identify incompatibilities of equipment with weather. For the case of AF 447, there would be a limitation on relying on air data for pitots operating in Ice Super Saturated Regions (ISSR). The PPA would correlate the flight plan route weather data and equipment limitations.
3. Check NOTAMS with flight plan route for equipment outages or other pertinent information that might limit flight operations
4. Check flight crew "blogs" that are created and maintained for specific airports and airspace along the flight plan. These correlations could include information on ATC procedures, ATC preferences, and "gotchas" while operating in specific airspace.

2. Intention Comparator (IC):

There is no single display component that reliably indicates to the pilot whether airspeed is or is not being controlled. For example in the OZ 214 CFIS scenario, the pilot must recall from memory that airspeed is not being controlled when the Autothrottle is engaged in the "HOLD" mode despite the other cues that suggest that airspeed is being controlled.

Researchers have proposed designs for the display of automation intentions on the flight deck[2, 3] (see example in Figure 3) and have demonstrated improved situational awareness using these designs. However, even if this information were made available to the flight crew, there is no guarantee that it would be used at the time of its greatest utility – when there is a divergence between flight crew intentions and automation intentions.

To avoid divergence in intentions and a resulting divergence from the planned flight trajectory, the PPA shall monitor changes in automation targets and modes that migrate away from the flight crew intentions and/or the navigation procedure (that derived from MCP and MCDU inputs). The primary function is for the PPA to monitor circumstances when the automation is not actively controlling to airspeed. A subset of this category of checking is the status of the coupling between Autopilot/Autothrottle and flight control surface and engine actuators. PPA shall monitor all changes in engagement status, in particular autonomous changes in engagement status.

The PPA could be supported by the display of "no active speed control" directly on the PFD, MCP and MCDU as described in Sherry & Mauro [1] (see Figure 4).[3]

[3] Another type of deception is the use of ***overloaded labels*** for FMA. Exclusive use of labels can be addressed in the design process and trained.

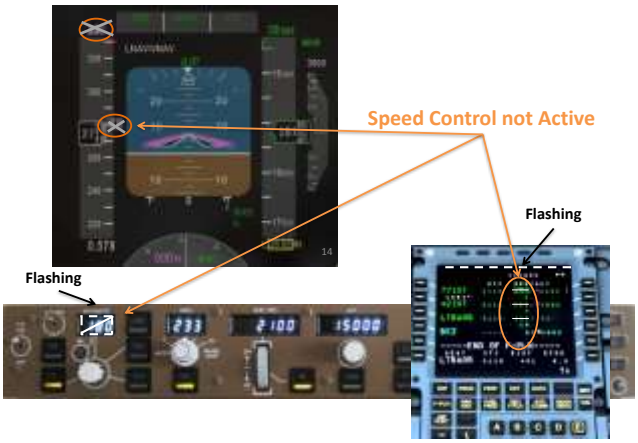

FIGURE 4: Inactive speed control indications on the PFD Airspeed tape, MCP and FMS.

3. Thrust Advance and Airspeed Change Indicator (TA & AC I)

The PPA shall monitor the airspeed relative to the target airspeed along with rate of change of airspeed, pitch rate, thrust setting, and rate of change in thrust setting to determine the airspeed at which the throttles should advance to acquire and maintain the target airspeed. An example display of the thrust advance location can be shown on the airspeed tape on the PFD as illustrated in Figure 5. This calculation exhibits some non-linearities that make it difficult to estimate with a simple rule-of-thumb.

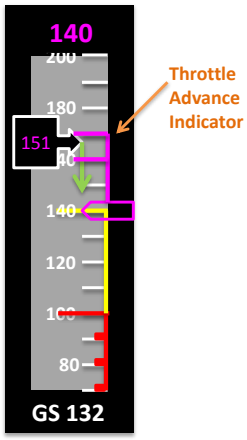

FIGURE 5: During deceleration, Thrust Advance Indicator (TAI) identifies airspeed at which throttles should advance

4. Flight deck Data Integrity Monitor (FDIM)

There are several ways in which the automation can create confusion. First, the sensor data may be erratic or discrepant. Attempting to guess in flight which sensors are valid is a difficult task (e.g. AF 447, XL Germany). In some cases, when confronted with confusing sensor indications, it may be prudent to take actions not based on having accurate sensor data such as flying by "pitch-and-power."

Sherry & Mauro [1] proposed a *Data Integrity* (DI) function that monitors sensor status and calculates an index representing the integrity of the sensor data. Sensor discrepancies, sensor fail-safe logic selection of sensors, and/or erratic sensor data are ways in which the DI score would be downgraded. Once the DI score drops below a threshold, flight crews would know to what degree the automation could be trusted.

This idea is embodied in the *Cynefin* framework[4], developed by Snowden, which characterizes the environment in which a firm operates. In the Simple state, Best Practices are used to sense, categorize and respond. In the Complicated state, Good Practice is used to sense, analyze and respond. In the Complex/Chaotic state, Emergent/Novel practices are used to probe/action, sense, analyze, and respond. The flight deck procedures are not trained this way. The assumption is that the flight crew should always be using Best Practices. However, when the situation is no longer a Simple state (e.g. erratic or discrepant sensor data) one should recognize that maintaining the

procedure is not an option and the flight deck operations should revert to flying the aircraft, not following the procedure.

To assess the overall state of the flight deck automation, the FDIM shall check the integrity of data used by the automation. Erratic data (e.g. AF 447), discrepant data (e.g. TA 1951) would alter the flight crew assessment of the situation and the resulting intervention response.

V Discussion

The Functional Complexity Failures (FCFs) inherent in the CFIS scenario represent a class of failures that are present in many engineered systems. Rooting these out in the design process to meet reliability requirements may yield diminishing returns as it may not be possible to measure the ability of a system or component to perform its required functions under the ***stated*** *(i.e. expected)* ***conditions*** for a specified period of time (i.e. reliability).

An alternative approach is to complement the design for reliability with a design for resilience that may include a human operator or supervisor. Resilience is defined as the ability of a system to adjust its functioning prior to, during, or following changes and disturbances, so that it can sustain required operations under both ***expected*** and ***unexpected*** conditions.

To achieve resilience, the intervention task must be explicitly designed into the system. Where human operators are assigned functional responsibility for the intervention, their tasks must be supported by automation explicitly designed to indicate the need for attention and to provide the information needed to perform the intervention.

The detailed analysis of the CFIS scenario identified six tasks that must be performed to achieve resilience: analysis of the tasks identified four functions that are not the strengths of human operator's (e.g. monitoring rare events). These four functions, automated in the design of a Paranoid Pilot Associate (PPA), were described.

A. Research Questions

There are two research issues that must be addressed in the development of the PPA. First and foremost is the integration of the PPA into the flight deck operations. The primary purpose of the PPA is to draw attention to matters that are considered critical by the PPA. In certain circumstances, it may be appropriate to alert a flight crew to issues of less importance. However, the PPA could interupt the flight crew while the pilots are performing a critical task. This would be counterproductive. Second is the issue of false-positives/nuisance alerts. Nuisance alerts would have the effect of reducing the perceived value of the PPA to the point where the alerts might be ignored.

One of the underlying premises of the PPA is the probabilistic nature of the alerts for inherently uncertain events. The modern flight deck is designed and operated under a deterministic *modus operandi*. To achieve the high reliability and efficiency required, all the uncertainty in operations has been necessarily eliminated. However, it may not be possible to treat rare events requiring intervention in the same manner. Because failure to intervene in these events has high costs, and because these events are rare and inherently difficult to identify with certainty, it may be necessary to allow a probabilistic approach to be used on the flight deck.

Acknowledgments

This work was funded in part by NASA NRA NNX12AP14A and internal GMU Research Foundation funds. Thank you for technical suggestions from Immanuel Barshi, Michael Feary, Randy Bailey, Paul Krasa, Steve Jacklin, Houda Kourdali, Julia Trippe, George Dononhue, Akshay Belle, John Shortle, Mike Hieb, Paulo Costa, and Yong Tian.

References

[1] Sherry, L., R. Mauro (2014) Controlled Flight into Stall: Functional Complexity Failures and Automation Surprises. In Proceedings 2014 Integrated Communications Navigation and Surveillance (ICNS) Conference, Dulles, Va. April, 2014. Page D1-1

[2] Feary, M., D. McCrobie, M. Alkin, L. Sherry, P. Polson, E. Palmer (1998) Aiding Vertical Guidance Understanding. NASA/TM—1998-112217

[3] Sherry, L. & P.Polson (1999) Shared models of flight management systems vertical guidance. The International Journal of Aviation Psychology – Special Issue: Aircraft Automation. L. Erlbaum: N.Y

[4] Snowden, D. & M. Boone (2007). A Leader's Framework for Decision Making. Harvard Business Review: 69–76