The Dangers of Interaction with Modular and Self-Healing Avionics Applications: Redundancy Considered Harmful

C.W. Johnson, Ph.D.; Department of Computing Science, University of Glasgow, Scotland.

Keywords: Redundancy, Accident Analysis, Digital Avionics, Software Systems Engineering.

Abstract

Redundancy is one of the primary techniques for the engineering of safety-critical systems. Back-up resources can be called upon to mitigate the failure of primary systems. Traditionally, operator intervention can be required to manually switch between a failed unit and redundant resources. However, programmable systems are increasingly used to automatically detect failures and reconfigure underlying systems excluding faulty components. This creates problems if operators do not notice that their underlying systems have been reconfigured. In this paper, we examine a number of additional concerns that arise in the present generation of redundant, safety-critical applications. A range of innovative 'self-healing' avionics applications are providing new benefits through the application of redundancy. They are also raising serious questions about the operators' ability to maintain situation awareness as control passes from primary to secondary and tertiary applications. Two recent in-flight interruptions involving a Boeing 777 and an Airbus A330 are used to illustrate the argument.

Introduction

Redundancy is one of the most widespread techniques for reducing the impact of failure in safety-critical systems. Additional components can be deployed to provide secondary resources that maintain functionality following the failure of primary systems. Over the years, many different variations have been developed:

- **Static redundancy** assumes that several redundant systems are operating at the same time. Although any one component may provide sufficient functionality for an application, majority voting is used to determine the outcome of the duplicated operation or computation. The intention is that if one element fails then it will be outvoted by the majority of other 'healthy' peers. A key strength of this approach is that there is no need to continually detect the possibility of a failure – unless the number of failed components may exceed the number of healthy peers.

- **Dynamic redundancy** relies upon one master system – if it fails then control passes to a redundant resource. In **hot dynamic techniques,** the secondary system is running in the background and may quickly resume primary operation. In **cold dynamic techniques,** the backup application must be restarted and brought up to the point at which the primary application failed. A key issue in all of these dynamic techniques is that it must be possible to determine when a primary system has failed so that the back-up resource can take over operations.

Although these distinctions will be familiar to many safety-critical engineers, it is worth reiterating them here because they represent the starting point for a range of more advanced techniques that are being deployed in complex applications. For example, commercial aircraft are equipped with inertial reference units to supply attitude, flight path vector, track, heading, accelerations, angular rates, ground speed, vertical speed and aircraft position. These have traditionally been separate from the air data computers that calculate barometric altitude, speed, Mach, angle of attack and temperature. However, the close interconnection between these data sources has led to the integration of inertial reference units and air data systems into single units known as air data inertial reference unit (ADIRU). The critical nature of these components has led to the application of redundancy within aircraft such as the Airbus A330 which provides three ADIRUs within the air data and inertial reference system (ADIRS). Each part of each of the three ADIRUs can operate separately in the case of the failure of the other part hence it is possible for the inertial reference unit of ADIRU 1 to continue to operate even though a fault has been noted for the air data component of that ADIRU (Ref. 1). Figure 1 shows how multi-level redundancy is also being used as a key architectural tool within the Boeing 777 ADIRU. The unit is divided into seven fault containment modules or areas. Each of these is physically and electrically isolated from the others. This feature enables operators to continue flying until the number of fault-free modules falls below a minimum specified by the component manufacturer. This fault tolerance supported lower operating costs, for instance, by reducing the potential disruption to aircraft operations from unscheduled maintenance.

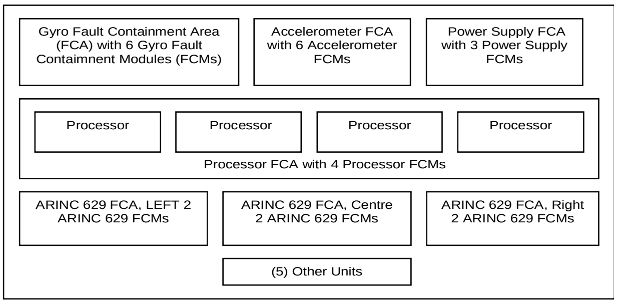
Gyro Fault Containment Area (FCA) with 6 Gyro Fault Contaimnent Modules (FCMs)
Accelerometer FCA with 6 Accelerometer FCMs
Power Supply FCA with 3 Power Supply FCMs
Processor
Processor
Processor
Processor
Processor FCA with 4 Processor FCMs
ARINC 629 FCA, LEFT 2 ARINC 629 FCMs
ARINC 629 FCA, Centre 2 ARINC 629 FCMs
ARINC 629 FCA, Right 2 ARINC 629 FCMs
(5) Other Units

Figure 1 —Air Data Inertial Reference Unit (ADIRU) Architecture (ATSB, Ref. 2, p.5)

Case Study One: A330 In-flight Interruption

This paper uses two recent in-flight 'upsets' to illustrate some of the increasingly complex failure modes that affect the engineering of redundant avionics. Both focus on the ADIRU architectures mentioned above. The first case study occurred in October 2008 when an Airbus A330-303 departed Singapore for Perth, Australia (Ref. 1). The problems for the crew began when the autopilot with a range of aircraft system failure indications during the cruise. While the crew was evaluating the situation, the aircraft abruptly pitched nose-down and descended 650 ft. After returning the aircraft to 37,000 ft, the crew commenced actions to deal with multiple failure messages. A second uncommanded pitch-down event occurred, and the aircraft lost 400 ft. The crew made an emergency broadcast to air traffic control diverted to make a successful landing. The initial report into the incident indicates that, at the time the autopilot disconnected, there was a fault with the inertial reference (IR) part of the air data inertial reference unit (ADIRU) number 1.

A key theme in this paper is that recent applications of redundancy have increased the complexity of avionics. This presents considerable barriers to accident investigators who must piece together the complex events and contributory factors that led towards an adverse event. For this reason, we use Events and Causal Factors (ECF) diagrams to provide an overview of the case study incident. This approach helps to map key areas of the many pages of text that are used in the official reports of the two case studies. ECF diagrams were originally developed by the US Department of Energy. It is important to stress, however, that this is only one of several different notations that might be used to provide a similar overview (Ref. 3). The focus here is less on the technique used for the analysis than on the role that redundancy played in the causes of potential accidents (Ref. 4).

Figure 1 illustrates the initial events leading to the first case study incident. As can be seen, the upset is believed to have been triggered by a discrepancy between channels of PRIM1 – this is the first of three flight control primary computers. The PRIMs help to ensure that the aircraft remains within a 'safe' flight envelope by automatically monitoring and commanding control surface movements. In normal operation, one of these three PRIMs acts as a master sending commands to the others that execute them using servo-controls. This illustrates the level of redundancy in complex avionics systems because each of the three redundant PRIMs can be allocated to receive input from the redundant ADIRU architectures describe above. At the start of the incident, PRIM1 was acting as the master when the discrepancy was detected; this led to the disconnection of autopilot 1 and the inertial reference system function of ADIRU 1 began to indicate 'Fail'. ADIRU 2 and ADIRU 3 seem to operate normally throughout the flight.

A second key theme in this paper is that partly as a consequence of the engineering complexity of redundant systems, it is increasingly difficult for the crews to identify and respond to adverse events. As can be seen from figure 1, the loss of the autopilot led to an Electronic centralized aircraft monitor (ECAM) warning message and a series of master caution chimes. The captain then attempted to engage autopilot 2 and then autopilot 1 without success. The crew confirmed and cleared the message from the ECAM but were then presented with a NAV IR1 FAULT message. Meanwhile, the airspeed and altitude readings on the captain's primary flight display provided fluctuating and, at times, contradictory information with stall and over-speed warnings. Uncertainty over the indications on his own display led the Captain to rely on standby instrumentation and then to use the First Officer's Primary Flight Display.

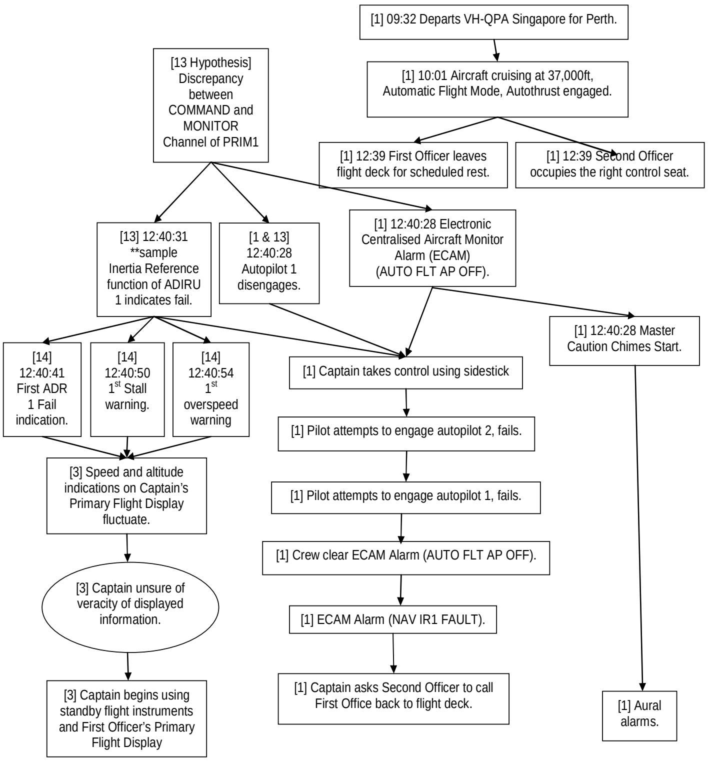
[1] 09:32 Departs VH-QPA Singapore for Perth.
[1] 10:01 Aircraft cruising at 37,000ft, Automatic Flight Mode, Autothrust engaged.
[13 Hypothesis] Discrepancy between COMMAND and MONITOR Channel of PRIM1
[1] 12:39 First Officer leaves flight deck for scheduled rest.
[1] 12:39 Second Officer occupies the right control seat.
[13] 12:40:31 **sample Inertia Reference function of ADIRU 1 indicates fail.
[1 & 13] 12:40:28 Autopilot 1 disengages.
[1] 12:40:28 Electronic Centralised Aircraft Monitor Alarm (ECAM) (AUTO FLT AP OFF).
[1] 12:40:28 Master Caution Chimes Start.
[14] 12:40:41 First ADR 1 Fail indication.
[14] 12:40:50 1st Stall warning.
[14] 12:40:54 1st overspeed warning
[1] Captain takes control using sidestick
[1] Pilot attempts to engage autopilot 2, fails.
[1] Pilot attempts to engage autopilot 1, fails.
[1] Crew clear ECAM Alarm (AUTO FLT AP OFF).
[1] ECAM Alarm (NAV IR1 FAULT).
[1] Captain asks Second Officer to call First Office back to flight deck.
[3] Speed and altitude indications on Captain's Primary Flight Display fluctuate.
[3] Captain unsure of veracity of displayed information.
[3] Captain begins using standby flight instruments and First Officer's Primary Flight Display
[1] Aural alarms.

Figure 2 — Initial Problems Affecting a Flight Control Primary Computer (PRIM1)

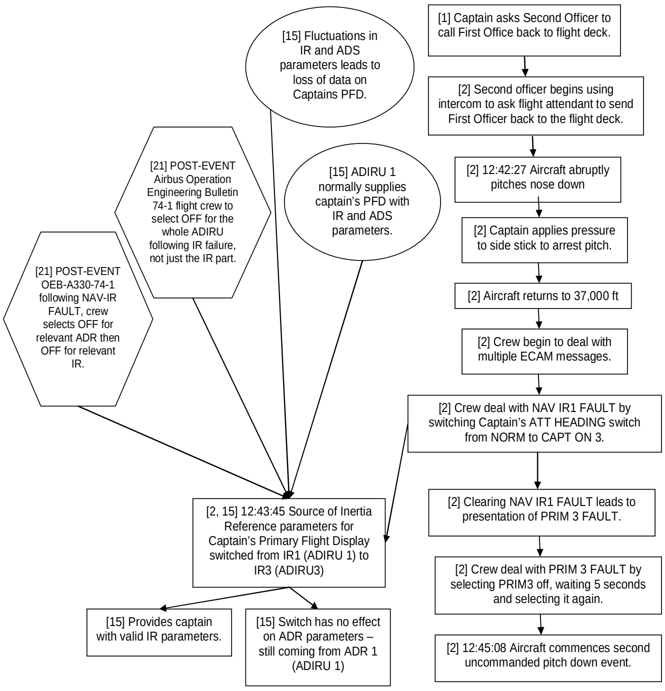
[15] Fluctuations in IR and ADS parameters leads to loss of data on Captains PFD.
[1] Captain asks Second Officer to call First Office back to flight deck.
[2] Second officer begins using intercom to ask flight attendant to send First Officer back to the flight deck.
[21] POST-EVENT Airbus Operation Engineering Bulletin 74-1 flight crew to select OFF for the whole ADIRU following IR failure, not just the IR part.
[15] ADIRU 1 normally supplies captain's PFD with IR and ADS parameters.
[2] 12:42:27 Aircraft abruptly pitches nose down
[2] Captain applies pressure to side stick to arrest pitch.
[21] POST-EVENT OEB-A330-74-1 following NAV-IR FAULT, crew selects OFF for relevant ADR then OFF for relevant IR.
[2] Aircraft returns to 37,000 ft
[2] Crew begin to deal with multiple ECAM messages.
[2] Crew deal with NAV IR1 FAULT by switching Captain's ATT HEADING switch from NORM to CAPT ON 3.
[2] Clearing NAV IR1 FAULT leads to presentation of PRIM 3 FAULT.
[2, 15] 12:43:45 Source of Inertia Reference parameters for Captain's Primary Flight Display switched from IR1 (ADIRU 1) to IR3 (ADIRU3)
[2] Crew deal with PRIM 3 FAULT by selecting PRIM3 off, waiting 5 seconds and selecting it again.
[15] Provides captain with valid IR parameters.
[15] Switch has no effect on ADR parameters – still coming from ADR 1 (ADIRU 1)
[2] 12:45:08 Aircraft commences second uncommanded pitch down event.

Figure 3 — Second Event and Temporary Solutions

Figure 3 extends the initial analysis from the previous ECF diagram to consider subsequent events and also to analyze factors that affected the reliability of the Captain's Primary Flight Display. As can be seen, the Captain's Primary Flight Display usually presented data from ADIRU 1, following the presentation of the NAV IR1 FAULT; the source of data was switched to Inertial Reference unit 3 on ADIRU 3. However, this did not automatically switch the source for Air Data References, which continued to be ADIRU 1. This illustrates the complexity of interaction with redundant systems as crews struggle to ensure that they receive data from a reliable source without knowing for sure which of the alternate ADIRU's is providing reliable information. The problems with the Captain's Primary Flight Display could not be resolved before a second uncommanded pitch down. The uncertainty created by crew interaction with their redundant systems was exacerbated by the way in which the master flight control primary computer was switched from PRIM1 to PRIM 2 following the first pitch down event. The subsequent indication of a fault on PRIM 3 then triggered a further change in the master from PRIM 2 back to PRIM 1 and it was only in subsequent discussions with the operator's maintenance watch unit in Sydney, while the flight was still in the air that the crew decided to switch off PRIM 3. At the time of writing this paper, there is continued uncertainty over the precise ways in which interactions between these different layers of redundancy, between the ADIRUs and the PRIMs, contributed to the symptoms that faced the crew during this incident.

Case Study Two: B777 In-flight Interruption

The second incident occurred during August 2005 and involved a Boeing Company 777-200 aircraft on an international passenger flight from Australia to Malaysia. As with the A330, this resulted in a significant upset while flying on autopilot. The Australian Transport Safety Bureau's investigation again focused on the role of the ADIRU. Although the units were made by a different manufacturer and had a different architecture, as explained in the opening sections, both case incidents were exacerbated by the use of redundancy that is intended to mitigate the impact of failure in safety critical systems. The anomaly in the B-777 had existed in original versions of the ADIRU software, and had not been detected in the testing and certification process for the unit.

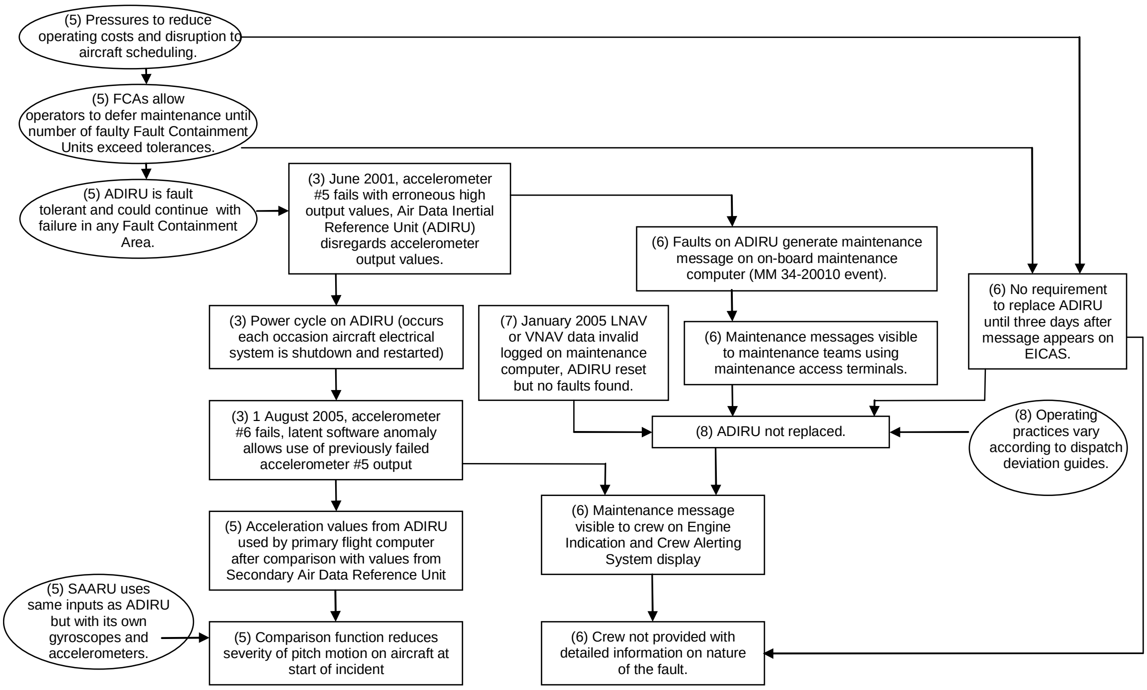
(5) Pressures to reduce operating costs and disruption to aircraft scheduling.
(5) FCAs allow operators to defer maintenance until number of faulty Fault Containment Units exceed tolerances.
(5) ADIRU is fault tolerant and could continue with failure in any Fault Containment Area.
(3) June 2001, accelerometer #5 fails with erroneous high output values, Air Data Inertial Reference Unit (ADIRU) disregards accelerometer output values.
(6) Faults on ADIRU generate maintenance message on on-board maintenance computer (MM 34-20010 event).
(6) No requirement to replace ADIRU until three days after message appears on EICAS.
(3) Power cycle on ADIRU (occurs each occasion aircraft electrical system is shutdown and restarted)
(7) January 2005 LNAV or VNAV data invalid logged on maintenance computer, ADIRU reset but no faults found.
(6) Maintenance messages visible to maintenance teams using maintenance access terminals.
(3) 1 August 2005, accelerometer #6 fails, latent software anomaly allows use of previously failed accelerometer #5 output
(8) ADIRU not replaced.
(8) Operating practices vary according to dispatch deviation guides.
(5) Acceleration values from ADIRU used by primary flight computer after comparison with values from Secondary Air Data Reference Unit
(6) Maintenance message visible to crew on Engine Indication and Crew Alerting System display
(5) SAARU uses same inputs as ADIRU but with its own gyroscopes and accelerometers.
(5) Comparison function reduces severity of pitch motion on aircraft at start of incident
(6) Crew not provided with detailed information on nature of the fault.

Figure 4 — Deferred Maintenance on Fault Tolerant ADIRU's

The opening sections of this paper described how designers of the B777 ADIRU enabled aircraft operators to continue flying and defer maintenance even after a failure had been logged against a fault containment area. Figure

4 shows how an initial failure of accelerometer number 5 in 2001 triggered a maintenance message for the on-board maintenance computer; this was known as an ADIRU MM 34-20010 event. Such messages could be read by maintenance teams using a ground-based terminal. However, these messages were not directly visible to the crew. Some ADIRU events can, however, be displayed in-flight on the Engine Indication and Crew Alerting System (EICAS). If such an in-flight warning occurs, then the ADIRU must be replaced with a serviceable unit within three days. As can be seen in Figure 4, the crew did not receive such a warning following the 2001 accelerometer failure and so the ADIRU was not replaced. The aircraft manufacturer noted that: “the ADIRU can be dispatched with MM 34-20010 present until such time that the operator deems it prudent to remove the ADIRU to avoid a schedule interruption due to occurrence of the ADIRU Status message. The decision to remove the ADIRU based on the presence of MM 34-20010 only is made by the operators on an economic basis, not a safety basis” (Ref. 2, p.8).

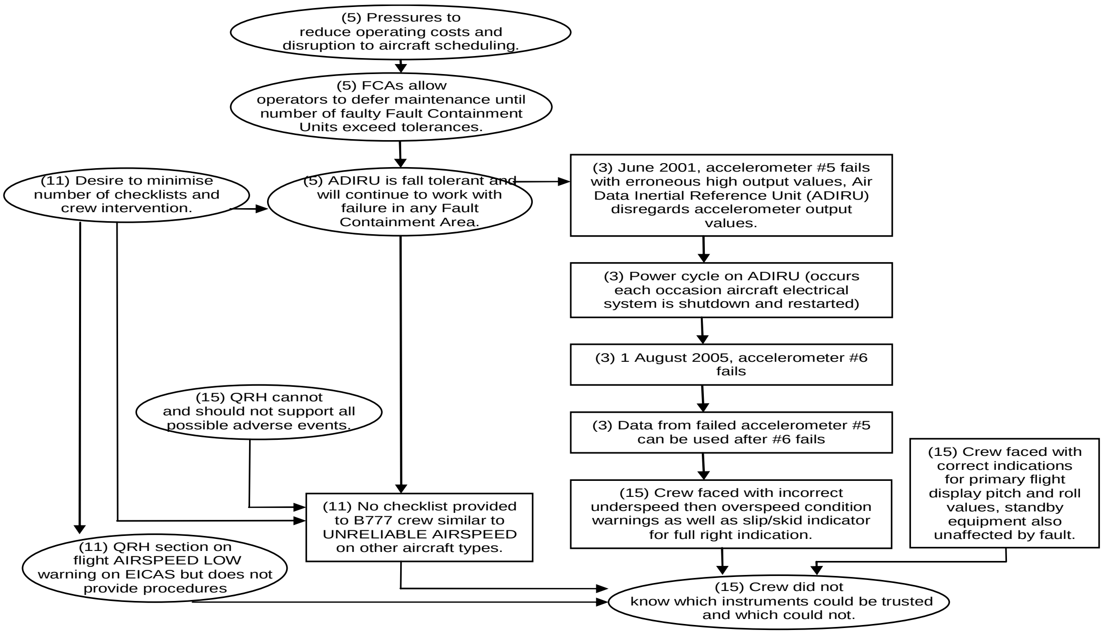
(5) Pressures to reduce operating costs and disruption to aircraft scheduling.
(5) FCAs allow operators to defer maintenance until number of faulty Fault Containment Units exceed tolerances.
(11) Desire to minimise number of checklists and crew intervention.
(5) ADIRU is fall tolerant and will continue to work with failure in any Fault Containment Area.
(3) June 2001, accelerometer #5 fails with erroneous high output values, Air Data Inertial Reference Unit (ADIRU) disregards accelerometer output values.
(3) Power cycle on ADIRU (occurs each occasion aircraft electrical system is shutdown and restarted)
(3) 1 August 2005, accelerometer #6 fails
(15) QRH cannot and should not support all possible adverse events.
(3) Data from failed accelerometer #5 can be used after #6 fails
(15) Crew faced with correct indications for primary flight display pitch and roll values, standby equipment also unaffected by fault.
(11) No checklist provided to B777 crew similar to UNRELIABLE AIRSPEED on other aircraft types.
(15) Crew faced with incorrect underspeed then overspeed condition warnings as well as slip/skid indicator for full right indication.
(11) QRH section on flight AIRSPEED LOW warning on EICAS but does not provide procedures
(15) Crew did not know which instruments could be trusted and which could not.

Figure 5 — Fault Masking, Redundancy and the Impact on Crew Interaction

Figure 5 shows how ADIRU software was developed to check the status of critical components and to allow the unit to continue operation if minimum criteria for the availability of FCAs were met. The implementation of this requirement was flawed in early versions of the code. However, up to version v-03 this problem was mitigated by additional checks in other areas of the application. A renewed requirement to improve shop repair capability led to the flaw being exposed again in OPS v-04. The OPS software up to and including v-07, therefore, contained a bug such that after a power cycle the ADIRU would not recognize that accelerometer number 5 was unserviceable. Figure5 also denotes how a maintenance message was generated following the 2001 fault on accelerometer number 5. However, the fault status was not checked by the ADIRU following a power up for the reasons presented previously. This combination of events led to use of data from the failed accelerometer when a further fault was detected in 2005 with accelerometer number 6.

Figure 5 illustrates key similarities between the A330 and the B-777 incidents. The masking of redundant failures undermined the ability of human operators to diagnose and respond to the problems that confronted them. As can be seen, in the following ECF diagram the crew of the B-777 was faced with incorrect underspeed then over-speed condition warnings as well as slip/skid indicator for full right indication following the failure of accelerometer number 6. At the same time, correct indications were presented for the pitch and roll values on the primary flight display. Standby equipment was also unaffected by the fault. The crew was, therefore, unsure which of the instruments to trust. Their uncertainty was exacerbated by design decisions that stemmed from the underlying philosophy of marking redundant failures during continued operations, mentioned in previous sections. The pilot was faced with a situation that designers had not considered to be possible. The auto throttle system remained active and the underspeed/over-speed warnings suggested that the malfunction may have been related to these functions. In consequence, the pilot attempted to disconnect the autothrottle by pressing the thrust lever disconnect switch and pushing the autothrottle engage switch to toggle it off. However, these attempts were ineffective because the crew failed to switch the autothrottle arm switch from ARMED to OFF. In consequence, the autothrottle continued to increase thrust in response to the low-speed data that was erroneously being supplied from the ADIRU and the fault accelerometer.

Conclusions and Further Work

Redundancy continues to provide significant benefits to the engineering and operation of complex, safety-critical systems. However, this paper has used two recent incidents to illustrate potential hazards as static and dynamic techniques are being extended to support multiple levels of redundancy. The A330 mishap shows that crews may not be able to easily determine the source of a problem or conversely to identify reliable data sources when multiple redundant processing units, such as PRIM 1 to 3, call on data from multiple redundant sources, specifically ADIRU 1 to 3, each of which provides access to multiple duplicated components. We have also used the case study incidents to illustrate concerns about 'self healing' systems in which redundancy is used to justify increased maintenance intervals as routine operations are conducted in the presence of failed components. Again, hazards focus on human interaction with complex redundant architectures – in this case maintenance teams must identify failed components when they are eventually able to work on an application. If they do not correct faults then there is a danger that they will continue to be masked from the crew until further failures eventually undermine redundant architectures.

Afterword

Since writing the initial draft of this paper, a further incident has occurred. The autopilot of an A330 again disconnected with a NAV IR 1 Fault ECAM message denoting a problem with ADIRU Number 1. In this incident the crew followed the advice that the manufacturer issued following the first case study in this paper; the crew selected the IR 1 push-button to OFF and the ADR 1 push-button to OFF and landed successfully. The ATSB interim press release states; "It is very early in the investigation and too soon to draw any conclusions as to specific causal factors involved in this incident. As it appears to be a similar event to a previous event involving an A330 aircraft (AO-2008-070 on 7 Oct 2008) it will be included as part of the earlier investigation" (Ref. 5). It would appear that the crew was able to benefit from the lessons learned in the previous incident; however, it is also clear that we have further lessons to learn in the application of advanced redundancy techniques for safety-critical software.

References

1. Australian Transport Safety Bureau. In-Flight Upset Event 240Km North-West of Perth, WA, Boeing Company 777-2000, 9M-MRG. Aviation Occurrence Report 200503722, Canberra, Australia, 2007.
2. Australian Transport Safety Bureau, In-Flight Upset 154km West of Learmonth, WA, VH-QPA, Airbus A330-303, Aviation Occurrence Investigation AO-2008-070 Preliminary, 2008.
3. C.W. Johnson. *Failure in Safety-Critical Systems: A Handbook of Accident and Incident Reporting*, University of Glasgow Press, Glasgow, Scotland, ISBN 0-85261-784-4, 2003.
4. C. M. Holloway and C.W. Johnson. "Why System Safety Professionals Should Read Accident Reports." In T. Kelly (ed.), *The First IET International Conference on System Safety*, Institute of Engineering and Technology, Savoy Place, London, 325-331, 6-8th June 2006, ISBN 0-86341-646-2, 2006.
5. Australian Transport Safety Bureau. Qantas Airbus A330 incident, 480km North West of Perth on 27 December 2008, Media Release, 2nd January 2009.

Biography

Chris.W. Johnson, DPhil, MA, MSc, FBCS, CEng, CITP, Department of Computing Science, University of Glasgow, Glasgow, G12 8RZ, Scotland, UK, telephone +44 (141) 330 6053, facsimile +44 (141) 330 4913, e-mail – Johnson@dcs.gla.ac.uk, web page http://www.dcs.gla.ac.uk/~johnson

Chris Johnson is Professor of Computing Science at the University of Glasgow in Scotland. He heads a small research group devoted to improving the reporting and analysis of incidents and accidents across safety-critical domains ranging from healthcare, to the military to aviation and rail.