Australian Government

Australian Transport Safety Bureau

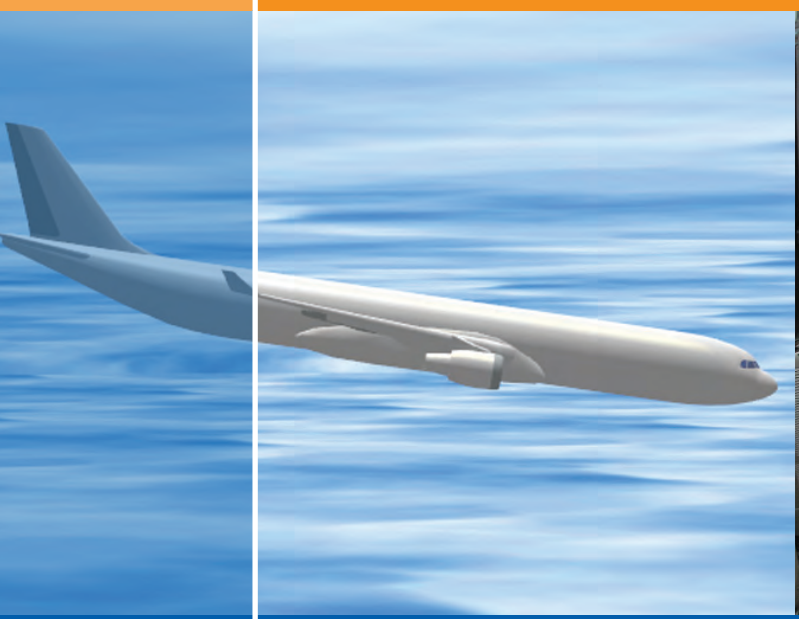

ATSB TRANSPORT SAFETY REPORT
Aviation Occurrence Investigation
AO-2008-070
Final

In-flight upset
154 km west of Learmonth, WA
7 October 2008
VH-QPA
Airbus A330-303

Australian Government

Australian Transport Safety Bureau

ATSB TRANSPORT SAFETY REPORT

Aviation Occurrence Investigation
AO-2008-070
Final

In-flight upset
154 km west of Learmonth, WA
7 October 2008
VH-QPA
Airbus A330-303

Released in accordance with section 25 of the *Transport Safety Investigation Act 2003*

Published by: Australian Transport Safety Bureau
Postal address: PO Box 967, Civic Square ACT 2608
Office: 62 Northbourne Avenue Canberra, Australian Capital Territory 2601
Telephone: 1800 020 616, from overseas +61 2 6257 4150
Accident and incident notification: 1800 011 034 (24 hours)
Facsimile: 02 6247 3117, from overseas +61 2 6247 3117
Email: atsbinfo@atsb.gov.au
Internet: www.atsb.gov.au

ISBN and formal report title: see ‘Document retrieval information’ on page vii.

CONTENTS

DOCUMENT RETRIEVAL INFORMATION

Report No.	Publication date	ISBN
AO-2008-070	December 2011	978-1-74251-231-0

Publication title

In-flight upset, 154 km west of Learmonth, Western Australia, 7 October 2008, VH-QPA, Airbus A330-303

Prepared By

Australian Transport Safety Bureau
PO Box 967, Civic Square ACT 2608 Australia, www.atsb.gov.au

Acknowledgements

Figures 1, 2, 26, F1 and F2: Courtesy of Google Earth
Figures 4, 5, 9, 10, 11, 13 and 18: Courtesy of Airbus
Where figures from other sources are reproduced, the acknowledgement is provided with the figure.

Abstract

On 7 October 2008, an Airbus A330-303 aircraft, registered VH-QPA and operated as Qantas flight 72, departed Singapore on a scheduled passenger transport service to Perth, Western Australia. While the aircraft was in cruise at 37,000 ft, one of the aircraft's three air data inertial reference units (ADIRUs) started outputting intermittent, incorrect values (spikes) on all flight parameters to other aircraft systems. Two minutes later, in response to spikes in angle of attack (AOA) data, the aircraft's flight control primary computers (FCPCs) commanded the aircraft to pitch down. At least 110 of the 303 passengers and nine of the 12 crew members were injured; 12 of the occupants were seriously injured and another 39 received hospital medical treatment.

Although the FCPC algorithm for processing AOA data was generally very effective, it could not manage a scenario where there were multiple spikes in AOA from one ADIRU that were 1.2 seconds apart. The occurrence was the only known example where this design limitation led to a pitch-down command in over 28 million flight hours on A330/A340 aircraft, and the aircraft manufacturer subsequently redesigned the AOA algorithm to prevent the same type of accident from occurring again.

Each of the intermittent data spikes was probably generated when the LTN-101 ADIRU's central processor unit (CPU) module combined the data value from one parameter with the label for another parameter. The failure mode was probably initiated by a single, rare type of internal or external trigger event combined with a marginal susceptibility to that type of event within a hardware component. There were only three known occasions of the failure mode in over 128 million hours of unit operation. At the aircraft manufacturer's request, the ADIRU manufacturer has modified the LTN-101 ADIRU to improve its ability to detect data transmission failures.

At least 60 of the aircraft's passengers were seated without their seat belts fastened at the time of the first pitch-down. The injury rate and injury severity was substantially greater for those who were not seated or seated without their seat belts fastened.
The investigation identified several lessons or reminders for the manufacturers of complex, safety-critical systems.

THE AUSTRALIAN TRANSPORT SAFETY BUREAU

The Australian Transport Safety Bureau (ATSB) is an independent Commonwealth Government statutory agency. The Bureau is governed by a Commission and is entirely separate from transport regulators, policy makers and service providers. The ATSB's function is to improve safety and public confidence in the aviation, marine and rail modes of transport through excellence in: independent investigation of transport accidents and other safety occurrences; safety data recording, analysis and research; fostering safety awareness, knowledge and action.

The ATSB is responsible for investigating accidents and other transport safety matters involving civil aviation, marine and rail operations in Australia that fall within Commonwealth jurisdiction, as well as participating in overseas investigations involving Australian registered aircraft and ships. A primary concern is the safety of commercial transport, with particular regard to fare-paying passenger operations.

The ATSB performs its functions in accordance with the provisions of the *Transport Safety Investigation Act 2003* and Regulations and, where applicable, relevant international agreements.

Purpose of safety investigations

The object of a safety investigation is to identify and reduce safety-related risk. ATSB investigations determine and communicate the safety factors related to the transport safety matter being investigated. The terms the ATSB uses to refer to key safety and risk concepts are set out in the next section: Terminology Used in this Report.

It is not a function of the ATSB to apportion blame or determine liability. At the same time, an investigation report must include factual material of sufficient weight to support the analysis and findings. At all times the ATSB endeavours to balance the use of material that could imply adverse comment with the need to properly explain what happened, and why, in a fair and unbiased manner.

Developing safety action

Central to the ATSB's investigation of transport safety matters is the early identification of safety issues in the transport environment. The ATSB prefers to encourage the relevant organisation(s) to initiate proactive safety action that addresses safety issues. Nevertheless, the ATSB may use its power to make a formal safety recommendation either during or at the end of an investigation, depending on the level of risk associated with a safety issue and the extent of corrective action undertaken by the relevant organisation.

When safety recommendations are issued, they focus on clearly describing the safety issue of concern, rather than providing instructions or opinions on a preferred method of corrective action. As with equivalent overseas organisations, the ATSB has no power to enforce the implementation of its recommendations. It is a matter for the body to which an ATSB recommendation is directed to assess the costs and benefits of any particular means of addressing a safety issue.

When the ATSB issues a safety recommendation to a person, organisation or agency, they must provide a written response within 90 days. That response must indicate whether they accept the recommendation, any reasons for not accepting part or all of the recommendation, and details of any proposed safety action to give effect to the recommendation.

The ATSB can also issue safety advisory notices suggesting that an organisation or an industry sector consider a safety issue and take action where it believes appropriate, or to raise general awareness of important safety information in the industry. There is no requirement for a formal response to an advisory notice, although the ATSB will publish any response it receives.

TERMINOLOGY USED IN THIS REPORT

Occurrence: accident or incident.

Safety factor: an event or condition that increases safety risk. In other words, it is something that, if it occurred in the future, would increase the likelihood of an occurrence, and/or the severity of the adverse consequences associated with an occurrence. Safety factors include the occurrence events (e.g. engine failure, signal passed at danger, grounding), individual actions (e.g. errors and violations), local conditions, current risk controls and organisational influences.

Contributing safety factor: a safety factor that, had it not occurred or existed at the time of an occurrence, then either: (a) the occurrence would probably not have occurred; or (b) the adverse consequences associated with the occurrence would probably not have occurred or have been as serious, or (c) another contributing safety factor would probably not have occurred or existed.

Other safety factor: a safety factor identified during an occurrence investigation which did not meet the definition of contributing safety factor but was still considered to be important to communicate in an investigation report in the interests of improved transport safety.

Other key finding: any finding, other than that associated with safety factors, considered important to include in an investigation report. Such findings may resolve ambiguity or controversy, describe possible scenarios or safety factors when firm safety factor findings were not able to be made, or note events or conditions which 'saved the day' or played an important role in reducing the risk associated with an occurrence.

Safety issue: a safety factor that (a) can reasonably be regarded as having the potential to adversely affect the safety of future operations, and (b) is a characteristic of an organisation or a system, rather than a characteristic of a specific individual, or characteristic of an operational environment at a specific point in time.

Risk level: The ATSB's assessment of the risk level associated with a safety issue is noted in the Findings section of the investigation report. It reflects the risk level as it existed at the time of the occurrence. That risk level may subsequently have been reduced as a result of safety actions taken by individuals or organisations during the course of an investigation.

Safety issues are broadly classified in terms of their level of risk as follows:

- Critical safety issue: associated with an intolerable level of risk and generally leading to the immediate issue of a safety recommendation unless appropriate corrective safety action has already been taken by the relevant organisation.
- Significant safety issue: associated with a risk level regarded as acceptable only if it is kept as low as reasonably practicable. The ATSB may issue a safety recommendation or a safety advisory notice if it assesses that further safety action may be practicable.
- Minor safety issue: associated with a broadly acceptable level of risk, although the ATSB still encourages the relevant organisation(s) to take safety action. The ATSB may sometimes highlight a safety message or make a safety comment.

Safety action: the steps taken or proposed to be taken by a person, organisation or agency in response to a safety issue. When the ATSB has been advised of safety action in response to a safety issue, it is published in the final report.

ABBREVIATIONS

AC	Advisory circular
AC	Alternating current
ACJ	Advisory circular joint
ACARS	Aircraft communications, addressing and reporting system
ACMS	Aircraft condition monitoring system
ADIRS	Air data and inertial reference system
ADIRU	Air data inertial reference unit
ADM	Air data module
ADR	Air data reference
AIRMAN	AIRcraft Maintenance ANalysis database
AOA	Angle of attack
AP	Autopilot
ARINC	Aeronautical Radio Inc.
ARP	Aerospace recommended practice
AS	Aerospace standard
ASIC	Application-specific integrated circuit
ATA	Air Transport Association
ATSB	Australian Transport Safety Bureau
BEA	Bureau d'Enquêtes et d'Analyses pour la sécurité de l'aviation civile (France, Bureau of Investigations and Analysis for the Safety of Civil Aviation)
BITE	Built-in test equipment
CAAP	Civil Aviation Advisory Publication
CAS	Computed airspeed
CASA	Civil Aviation Safety Authority
CASR	Civil Aviation Safety Regulation
c.g.	Centre of gravity
CMS	Central maintenance system
COM	Command
CPU	Central processing unit
CRC	Cyclic redundancy check

CSM	Cabin services manager
CSS	Customer services supervisor
CVR	Cockpit voice recorder
DADS	Digital air data system
DC	Direct current
DGAC	Direction Générale de l'Aviation Civile (France)
DMC	Display management computer
DMU	Data management unit
EASA	European Aviation Safety Agency
ECAM	Electronic centralized aircraft monitor
EDAC	Error detection and correction
EFCS	Electrical flight control system
EIS	Electronic instrument system
EMI	Electromagnetic interference
ESS	Environmental stress screening
ETI	Elapsed time indicated
E/WD	Engine/warning display
FAA	Federal Aviation Administration (US)
FAR	Federal Aviation Regulation
FCOM	Flight crew operating manual
FCPC	Flight control primary computer (also known as PRIM)
FCSC	Flight control secondary computer (also known as SEC)
FDR	Flight data recorder
FHA	Functional hazard assessment
FL	Flight level
FM	Flight management
FMGEC	Flight management, guidance and envelope computer
FMGES	Flight management, guidance and envelope system
FMEA	Failure mode and effects analysis
FO	First officer
FT	Functional test (an SSM indication)
FTA	Fault tree analysis
FW	Failure warning (an SSM indication)

FWS	Flight warning system
F/CTL	Flight control
GPS	Global Positioning System
GPWS	Ground proximity warning system
HASS	Highly accelerated stress screening
HF	High frequency
IR	Inertial reference
ICAO	International Civil Aviation Organization
IEC	International Electrotechnical Commission
ISIS	Integrated standby instrument system
I/O	Input/output
JAR	Joint Aviation Requirement
LGCIU	Landing gear control interface unit
LRU	Line-replaceable unit
MBU	Multiple bit upset
MON	Monitor
MTBF	Mean time between failures
MTBUR	Mean time between unscheduled removals
NASA	National Aeronautics and Space Administration
NCD	No computed data (an SSM indication)
ND	Navigation display
NO	Normal operation (an SSM indication)
NTSB	National Transportation Safety Board (US)
OEB	Operations engineering bulletin
PA	Passenger address
PCB	Printed circuit board
PED	Personal electronic device
PFD	Primary flight display
PFR	Post-flight report
PHC	Probe heat computer
PRIM	Common name for flight control primary computer (FCPC)
PSSA	Preliminary system safety assessment
QAR	Quick access recorder

RAM	Random access memory
RTCA	Radio Technical Commission for Aeronautics
SATCOM	SATellite COMmunications
SAE	Society of Automotive Engineers
SAO	Spécification assistée par ordinateur (computer assisted specification)
SD	System display
SEC	Common name for flight control secondary computer (FCSC)
SEE	Single event effects
SEU	Single event upset
SSA	System safety assessment
SSM	Sign/status matrix
TAT	Total air temperature
THS	Trimmable horizontal stabiliser
TSD	Troubleshooting data
UTC	Universal Time Coordinated
VHF	Very high frequency
VLF	Very low frequency

EXECUTIVE SUMMARY

Key investigation outcomes

The in-flight upset on 7 October 2008 occurred due to the combination of a design limitation in the flight control primary computer (FCPC) software of the Airbus A330/A340, and a failure mode affecting one of the aircraft's three air data inertial reference units (ADIRUs). The design limitation meant that, in a very rare and specific situation, multiple spikes in angle of attack (AOA) data from one of the ADIRUs could result in the FCPCs commanding the aircraft to pitch down.

When the aircraft manufacturer became aware of the problem, it issued flight crew procedures to manage any future occurrence of the same ADIRU failure mode. The aircraft manufacturer subsequently reviewed and improved its FCPC algorithms for processing AOA and other ADIRU parameters. As a result of this redesign, passengers, crew and operators can be confident that the same type of accident will not reoccur.

The investigation identified several lessons or reminders for the manufacturers of complex, safety-critical systems. With the knowledge that systems are becoming increasingly complex, it also identified a need for more research into how design engineers and safety analysts evaluate system designs, and how their tasks, tools, training and guidance materials could be improved to minimise design errors.

Although in-flight upsets are very rare events, the accident on 7 October 2008 also provided a salient reminder to all passengers and crew of the importance of wearing their seat belts during a flight whenever they are seated.

Summary of the occurrence

At 0132 Universal Time Coordinated (0932 local time) on 7 October 2008, an Airbus A330-303 aircraft, registered VH-QPA and operated as Qantas flight 72, departed Singapore on a scheduled passenger transport service to Perth, Western Australia. At 0440:26, while the aircraft was in cruise at 37,000 ft, ADIRU 1 started providing intermittent, incorrect values (spikes) on all flight parameters to other aircraft systems. Soon after, the autopilot disconnected and the crew started receiving numerous warning and caution messages (most of them spurious). The other two ADIRUs performed normally during the flight.

At 0442:27, the aircraft suddenly pitched nose down. The FCPCs commanded the pitch-down in response to AOA data spikes from ADIRU 1. Although the pitch-down command lasted less than 2 seconds, the resulting forces were sufficient for almost all the unrestrained occupants to be thrown to the aircraft's ceiling. At least 110 of the 303 passengers and nine of the 12 crew members were injured; 12 of the occupants were seriously injured and another 39 received hospital medical treatment. The FCPCs commanded a second, less severe pitch-down at 0445:08.

The flight crew's responses to the emergency were timely and appropriate. Due to the serious injuries and their assessment that there was potential for further pitch-downs, the crew diverted the flight to Learmonth, Western Australia and declared a MAYDAY to air traffic control. The aircraft landed as soon as operationally practicable at 0532, and medical assistance was provided to the injured occupants soon after.

FCPC design limitation

AOA is a critically important flight parameter, and full-authority flight control systems such as those equipping A330/A340 aircraft require accurate AOA data to function properly. The aircraft was fitted with three ADIRUs to provide redundancy and enable fault tolerance, and the FCPCs used the three independent AOA values to check their consistency. In the usual case, when all three AOA values were valid and consistent, the average value of AOA 1 and AOA 2 was used by the FCPCs for their computations. If either AOA 1 or AOA 2 significantly deviated from the other two values, the FCPCs used a memorised value for 1.2 seconds. The FCPC algorithm was very effective, but it could not correctly manage a scenario where there were multiple spikes in either AOA 1 or AOA 2 that were 1.2 seconds apart.

Although there were many injuries on the 7 October 2008 flight, it is very unlikely that the FCPC design limitation could have been associated with a more adverse outcome. Accordingly, the occurrence fitted the classification of a 'hazardous' effect rather than a 'catastrophic' effect as described by the relevant certification requirements. As the occurrence was the only known case of the design limitation affecting an aircraft's flightpath in over 28 million flight hours on A330/A340 aircraft, the limitation was within the acceptable probability range defined in the certification requirements for a hazardous effect.

As with other safety-critical systems, the development of the A330/A340 flight control system during 1991 and 1992 had many elements to minimise the risk of a design error. These included peer reviews, a system safety assessment (SSA), and testing and simulations to verify and validate the system requirements. None of these activities identified the design limitation in the FCPC's AOA algorithm.

The ADIRU failure mode had not been previously encountered, or identified by the ADIRU manufacturer in its safety analysis activities. Overall, the design, verification and validation processes used by the aircraft manufacturer did not fully consider the potential effects of frequent spikes in data from an ADIRU.

ADIRU data-spike failure mode

The data-spike failure mode on the LTN-101 model ADIRU involved intermittent spikes (incorrect values) on air data parameters such as airspeed and AOA being sent to other systems as valid data without a relevant fault message being displayed to the crew. The inertial reference parameters (such as pitch attitude) contained more systematic errors as well as data spikes, and the ADIRU generated a fault message and flagged the output data as invalid. Once the failure mode started, the ADIRU's abnormal behaviour continued until the unit was shut down. After its power was cycled (turned OFF and ON), the unit performed normally.

There were three known occurrences of the data-spike failure mode. In addition to the 7 October 2008 occurrence, there was an occurrence on 12 September 2006 involving the same ADIRU (serial number 4167) and the same aircraft. The other occurrence on 27 December 2008 involved another of the same operator's A330 aircraft (VH-QPG) but a different ADIRU (serial number 4122). However, no factors related to the operator's aircraft configuration, operating practices or maintenance practices were found to be associated with the failure mode.

Many of the data spikes were generated when the ADIRU's central processor unit (CPU) module intermittently combined the data value from one parameter with the label for another parameter. The exact mechanism that produced this problem could

not be determined. However, the failure mode was probably initiated by a single, rare type of trigger event combined with a marginal susceptibility to that type of event within the CPU module's hardware. The key components of the two affected units were very similar, and overall it was considered likely that only a small number of units exhibited a similar susceptibility.

Some of the potential triggering events examined by the investigation included a software 'bug', software corruption, a hardware fault, physical environment factors (such as temperature or vibration), and electromagnetic interference (EMI) from other aircraft systems, other on-board sources, or external sources (such as a naval communication station located near Learmonth). Each of these possibilities was found to be unlikely based on multiple sources of evidence. The other potential triggering event was a single event effect (SEE) resulting from a high-energy atmospheric particle striking one of the integrated circuits within the CPU module. There was insufficient evidence available to determine if an SEE was involved, but the investigation identified SEE as an ongoing risk for airborne equipment.

The LTN-101 had built-in test equipment (BITE) to detect almost all potential problems that could occur with the ADIRU, including potential failure modes identified by the aircraft manufacturer. However, none of the BITE tests were designed to detect the type of problem that occurred with the air data parameters.

The failure mode has only been observed three times in over 128 million hours of unit operation, and the unit met the aircraft manufacturer's specifications for reliability and undetected failure rates. Without knowing the exact failure mechanism, there was limited potential for the ADIRU manufacturer to redesign units to prevent the failure mode. However, it will develop a modification to the BITE to improve the probability of detecting the failure mode if it occurs on another unit.

Use of seat belts

At least 60 of the aircraft's passengers were seated without their seat belts fastened at the time of the first pitch-down. Consistent with previous in-flight upset accidents, the injury rate, and injury severity, was substantially greater for those who were not seated or seated without their seat belts fastened.

Passengers are routinely reminded every flight to keep their seat belts fastened during flight whenever they are seated, but it appears some passengers routinely do not follow this advice. This investigation provided some insights into the types of passengers who may be more likely not to wear seat belts, but it also identified that there has been very little research conducted into this topic by the aviation industry.

Investigation process

The Australian Transport Safety Bureau investigation covered a range of complex issues, including some that had rarely been considered in depth by previous aviation investigations. To do this, the investigation required the expertise and cooperation of several external organisations, including the French Bureau d'Enquêtes et d'Analyses pour la sécurité de l'aviation civile, US National Transportation Safety Board, the aircraft and FCPC manufacturer (Airbus), the ADIRU manufacturer (Northrop Grumman Corporation), and the operator.

1 FACTUAL INFORMATION: GENERAL

1.1 History of the flight

Note. Information in this section was obtained from flight crew interviews and data from the aircraft's recorders. Explanations of aircraft systems are provided in section 1.6, further information from the recorders is provided in section 1.11, and further information about cabin-related events is provided in section 4.2.

1.1.1 Prior to the occurrence

At 0132 Universal Time Coordinated (UTC)[1] (0932 local time) on 7 October 2008, an Airbus A330-303 aircraft, registered VH-QPA (QPA), departed Singapore on a scheduled passenger transport service to Perth, Western Australia. The aircraft was operated as Qantas flight 72, and there were 303 passengers, three flight crew (captain, first officer and second officer), and nine cabin crew on board. The captain was the pilot flying.

The flight crew reported that the departure and climb out from Singapore proceeded normally. By 0201, the aircraft was in cruise with autopilot 1 engaged and maintaining the aircraft's altitude at flight level (FL)[2] 370. The autothrust was engaged and managing engine thrust to maintain a cruising speed of Mach 0.82.[3]

The flight crew stated that the weather was fine and clear and there had been no turbulence during the flight. At 0433, the captain returned to the flight deck from a scheduled rest break. At 0439 the first officer left for a rest break, and the second officer then occupied the right control seat.[4]

1.1.2 Start of occurrence sequence (0440:26)

At 0440:26, one of the aircraft's three air data inertial reference units (ADIRU 1) started providing incorrect data to other aircraft systems. At 0440:28, the autopilot automatically disconnected[5], and the captain took manual control of the aircraft.

Within 5 seconds of the autopilot disconnecting, a series of caution messages began appearing on the aircraft's electronic centralized aircraft monitor (ECAM), each

---

1 Local time in Singapore and Western Australia was UTC plus 8 hours.

2 At altitudes above 10,000 ft in Australia, an aircraft's height above mean sea level is referred to as a flight level (FL). FL 370 equates to 37,000 ft.

3 At high altitudes, the aircraft's Mach number, rather than airspeed, is more important for aircraft performance. The Mach number is the ratio of the aircraft's speed relative to the speed of sound. A Mach of 0.82 at FL370 equated to a computed airspeed of about 270 kts.

4 The A330 was designed to be operated by two pilots (captain and first officer). Second officers were carried to relieve the captain and first officer during long sectors on some trips. On this day, the flight crew were rostered to operate the Singapore-Perth flight and then a Perth-Singapore flight. Second officers do not normally occupy either of the control seats during landing or takeoff.

5 Consistent with an automatic disconnection (as opposed to a voluntary disconnection by the flight crew), there was a distinctive (cavalry charge) aural signal and a warning message (AUTO FLT AP OFF) on the ECAM.

associated with a master caution chime. The crew also started receiving aural stall warnings and overspeed warnings, although each warning was only annunciated briefly. These cautions and warnings occurred frequently, and continued for the remainder of the flight.

The crew cancelled the autopilot disconnection warning message (AUTO FLT AP OFF) on the ECAM and then engaged autopilot 2. After cancelling the autopilot message, they noticed a NAV IR 1 FAULT[6] caution message on the ECAM, with an associated IR 1 fault light on the overhead panel. The ECAM was also displaying other caution messages at this time.

In addition to the warnings and cautions, the crew reported that the airspeed and altitude indications on the captain's primary flight display (PFD) were fluctuating. No such fluctuations were occurring on the first officer's PFD or the standby flight instruments. The fluctuations on the captain's PFD appeared to be based on unreliable information, as there was no other indication that the aircraft was actually near a stall or overspeed condition. Because the captain was unsure of the veracity of the information on his PFD, he used the standby instruments and the first officer's PFD when flying the aircraft.

Data from the flight data recorder (FDR) showed that autopilot 2 was engaged for 15 seconds before being disconnected by the crew.[7] The FDR also showed that, during the period between the initial autopilot 1 disconnection and when autopilot 2 was engaged, the aircraft's altitude increased to 37,180 ft.[8] During the short period when autopilot 2 was engaged, the aircraft started to return to the assigned level. Although the crew received numerous ECAM caution messages, none of them required urgent action, and none of them indicated any potential problems with the aircraft's flight control system. However, the captain was not satisfied with the information that the aircraft systems were providing, and he asked the second officer to call the first officer back to the flight deck to help them diagnose and manage the problems.

1.1.3 First in-flight upset[9] (0442:27)

At 0442:27 (1242:27 local time), while the second officer was asking the cabin services manager (CSM) via the cabin interphone to send the first officer to the flight deck, the aircraft abruptly pitched nose down. The FDR showed that the pitch-down movement was due to a sudden change in the position of the aircraft's elevators, and that the aircraft reached a maximum nose-down pitch angle of 8.4°. The flight crew described the pitch-down movement as very abrupt, but smooth. It did not have the characteristics of a turbulence-related event and the aircraft's movement was solely in the pitching plane.

---

6 This message indicated a fault with the inertial reference part of ADIRU 1. The crew did not receive a message at this time indicating a fault with the air data reference part of the ADIRU.

7 The captain reported that he disconnected autopilot 2 as it was a required action in the event of an unreliable airspeed situation. Although the captain's airspeed values were fluctuating, there was no evidence of any problems with the other two airspeed sources during the flight.

8 This slight increase in altitude was associated with small pitch-up inputs from the captain's sidestick.

9 The term 'in-flight upset' is used in this report to refer to an abrupt manoeuvre of an aircraft.

The FDR showed that the captain immediately applied back pressure on his sidestick to arrest the pitch-down movement. The aircraft's flight control system did not initially respond to the captain's sidestick input, but after about 2 seconds the aircraft responded normally and the captain commenced recovery to the assigned altitude. During this 2-second period the aircraft descended about 150 ft. Overall, the aircraft descended 690 ft over 23 seconds before returning to FL370. An animation showing the aircraft's movement during the first upset, based on the FDR data, is provided on the Australian Transport Safety Bureau (ATSB) website.[10]

During the upset, the FDR recorded a peak vertical acceleration of -0.80 g.[11] A significant number of occupants were thrown around the cabin, resulting in injuries and damage to overhead fittings.[12] The second officer activated the seat-belt sign to ON and soon after (0442:43) made a public address announcement for passengers and crew to return to their seats and fasten their seat belts immediately.

The aircraft's position at the time of the first in-flight upset was over the Indian Ocean, 154 km west of Learmonth, Western Australia (Figure 1). This location, and the location of other key events, is shown in Figure 2.

Figure 1: Aircraft track (whole flight)

[10] See www.atsb.gov.au/publications/investigation_reports/2008/aair/ao-2008-070.aspx.

[11] Acceleration values recorded by the FDR were measured near the aircraft's centre of gravity (c.g.). The nominal vertical acceleration that is recorded when the aircraft is on the ground is 1g, and a negative g indicates a downwards acceleration of the aircraft. The vertical acceleration in the rear of the cabin was probably over -1.2 g (Appendix A).

[12] As required by procedures, the captain and second officer were wearing their seat belts at the time of the upset.

Figure 2: Aircraft track and key events

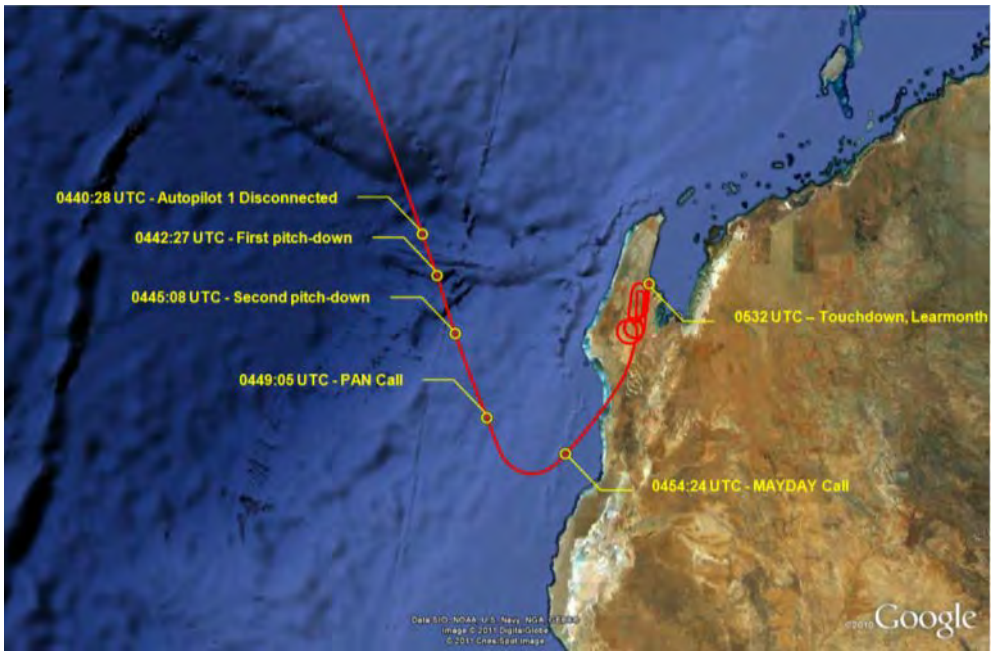

After returning the aircraft to FL370, the flight crew commenced actions to deal with multiple ECAM caution messages. These were:

- NAV IR 1 FAULT: the required action presented on the ECAM was to switch the ATT HDG (attitude heading) switch from the NORM position to the CAPT ON 3 position. The crew completed the required action and cleared the message.
- F/CTL PRIM 3 FAULT[13]: the required action was to select the flight control primary computer 3 (FCPC 3) OFF and then select it back ON. The crew completed the required action and cleared the message.
- NAV IR 1 FAULT: this message had reappeared at the top of the ECAM with no additional required actions, and was cancelled by the crew.
- F/CTL PRIM 1 PITCH FAULT: there was no required action associated with this message. The crew confirmed that there was no associated fault light on the overhead panel and then started reviewing the flight control page on the ECAM's system display.

1.1.4 Second in-flight upset (0445:08)

At 0445:08 (1245:08 local time), while the crew were responding to the ECAM messages, the aircraft commenced a second pitch-down movement, reaching a maximum pitch angle of about 3.5° nose down. The flight crew described the event as being similar in nature to the first event but less severe. The captain promptly applied back pressure on his sidestick to arrest the pitch-down movement. He said that, consistent with the first event, this action initially had no effect, but soon after the aircraft responded normally. FDR data showed that the flight control system did not respond to flight crew inputs for at least 2 seconds, and that the aircraft descended 400 ft over 15 seconds before returning to FL370.

---

[13] F/CTL is used for fault messages associated with the flight control system. PRIM is the term used in fault messages and operational procedures to refer to a flight control primary computer (FCPC).

Following the second in-flight upset, the crew continued to review the ECAM messages and other indications. The first ECAM message they noticed was F/CTL ALTN LAW (PROT LOST).[14] The next messages were recurrences of the NAV IR 1 FAULT and F/CTL PRIM 3 FAULT messages. The crew reported that the IR1 FAULT light and the PRIM 3 FAULT light on the overhead panel were illuminated. No other fault lights were illuminated.

The crew reported that by this time ECAM messages were frequently scrolling, with each new caution message being placed at the top of the list. The NAV IR 1 FAULT message kept recurring, together with several other messages, such as NAV GPS FAULT, and they could not effectively interact with the ECAM to action and/or clear the messages. Master caution chimes associated with the ECAM messages were frequently occurring, together with aural stall warnings and overspeed warnings. The crew stated that these constant aural alerts, and the inability to silence them, were a significant source of distraction.

1.1.5 Diversion to Learmonth

At 0446:10, the captain made a public announcement to the cabin, advising that the crew were dealing with flight control problems, and telling everyone to remain seated with their seat belts fastened. The second officer contacted the cabin again by interphone to ask a flight attendant to send the first officer to the flight deck.

The captain reported that, after the second upset event, he observed that the automatic pitch trim (autotrim) was not functioning[15] and he began trimming the aircraft manually. The crew advised that, because the autotrim was not working, they thought the flight control system was in direct law.[16]

With the exception of the loss of autotrim, the captain reported that the aircraft was flying normally. At 0447:25, he disconnected the autothrust to minimise any potential problems associated with the erroneous air data information affecting the electronic engine control units. He then flew the aircraft without the autopilot or autothrust engaged, and using the standby instruments, for the remainder of the flight.

The first officer returned to the flight deck at 0447:40, taking over from the second officer in the right control seat while the second officer moved to the third occupant seat. The crew discussed the situation, and the captain stated that he would continue flying the aircraft.

---

[14] The crew reported that they did not recall seeing amber crosses on the PFDs, which were meant to be displayed if the flight control system was in alternate law or direct law. The aircraft manufacturer advised that there was no technical reason why these amber crosses would not have been displayed on the occurrence flight.

[15] A ‘USE MAN PITCH TRIM’ message was not displayed to the crew on their PFDs as, at the time of the occurrence, this message was only displayed if the flight control system was in direct law. The aircraft manufacturer advised that this problem was being addressed with a new design standard, which was certified in 2011.

[16] The flight control system was in alternate law from 0445:11 until the end of the flight. Autotrim was generally available in alternate law, but it was lost in this case due to the sequence of fault messages associated with the FCPCs (PRIMs) (section 2.2.2).

The crew decided that they needed to land the aircraft as soon as possible. They were concerned that further pitch-down movements could occur, and they were aware that Learmonth was relatively close and that it was a suitable destination for an A330 landing.[17] In addition, when the first officer returned to the flight deck, he advised the other flight crew that there had been some injuries in the cabin.

At 0449:05, the first officer made a PAN[18] broadcast to air traffic control, stating that they had experienced 'flight control computer problems' and that some of the aircraft's occupants had been injured. He requested a clearance to divert to and track direct to Learmonth. The controller cleared the crew to descend to FL350.

The captain told the second officer to obtain further information from the cabin while the captain and first officer started preparing for the descent. At 0450:40, the second officer contacted the flight attendant at the Left 1 door position in the cabin (section 4.1.1) to get further information on the extent of the injuries.

At 0451:25, the first officer requested further descent from air traffic control. The controller cleared the crew to leave controlled airspace and proceed direct to Learmonth.

After receiving advice from the cabin of several serious injuries, the captain asked the first officer to declare a MAYDAY.[19] At 0454:25, the first officer declared the MAYDAY and advised air traffic control that they had multiple injuries on board.

1.1.6 Remainder of the flight

During the process of organising the diversion to Learmonth, the flight crew again reviewed what had happened, their current situation, and the ECAM messages. They noted that the NAV IR 1 FAULT and F/CTL PRIM 3 FAULT messages were still occurring, together with several other caution messages. They concluded that the ECAM was not providing them with useful information or recommended actions. Consequently, at 0456:05, the first officer contacted the operator's maintenance watch unit[20], located in Sydney, New South Wales, by a satellite communications system (SATPHONE) to brief them on the situation and to seek assistance.

There were subsequently several communications between the flight crew and maintenance watch about the fault messages and other flight deck indications. In one discussion at 0510, maintenance watch advised the flight crew that 'ADIRU 1' appeared to be common to the fault messages being displayed, but that there was also some conflicting information regarding elevator control. They provided no recommended actions at that stage. In a subsequent discussion, maintenance watch recommended that, at the crew's discretion, they could select PRIM 3 (or flight

---

[17] The first upset event occurred when the aircraft was 154 km (83 NM) west of Learmonth. Learmonth was the closest airport suitable for an A330 landing.

[18] A PAN transmission is made in the case of an urgency condition which concerns the safety of an aircraft or its occupants, but where the flight crew does not require immediate assistance.

[19] A MAYDAY transmission is made in the case of a distress condition and where the flight crew requires immediate assistance.

[20] Maintenance watch provided 24-hour assistance to en-route flight crews regarding technical issues. The aircraft's fault messages were automatically sent in real time to ground receivers, and could be accessed by maintenance watch (section 1.9).

control primary computer 3) OFF. The crew discussed this recommendation and, at 0520, switched that computer OFF. This action had no effect on the scrolling ECAM messages, stall warnings or overspeed warnings.

During the descent, the flight crew also had several communications with air traffic control regarding descent procedures, whether the approach and landing would be normal, whether there were any dangerous goods on board, and the availability of emergency services at Learmonth. The flight crew also had multiple communications with the cabin services manager, and made public announcements to the cabin regarding the situation and the diversion to Learmonth (section 4.2).

In addition to communications with maintenance watch, air traffic control and the cabin, the flight crew worked together to provide the captain with all the information he needed to fly the aircraft. They also needed to manage a range of problems with aircraft systems. For example, the flight crew were unable to enter an RNAV (GNSS)[21] approach into the flight management computer due to fault messages associated with the Global Positioning System (GPS) units. The second officer had to manually control the cabin pressure during the descent due to a pressurisation system fault, and the crew noted that they would need to use manual braking during landing due to an autobrake fault.

In order to lose altitude for landing, the captain conducted a series of wide left orbits to maintain the aircraft's speed below 330 kts (maximum operating speed). He reported that he descended cautiously in order to prevent any potential problems associated with another unexpected pitch-down event.

The crew completed the approach checklist and conducted a flight control check above 10,000 ft. After further descent, the aircraft was positioned at about 15 NM (28 km) for a straight-in visual approach to runway 36. The precision approach path indicator was sighted at about 10 NM (16 km), and the aircraft landed at Learmonth at 0532 (1332 local time).

1.2 Injuries to persons

Injuries	Crew	Passengers	Total
Fatal	-	-	-
Serious	1	11	12
Minor	8	99	107
None / unknown	3	193	196
Total	12	303	315

As some of the occupants received serious injuries, the occurrence was classified as an accident.[22] Further injury information is presented in section 4.6.

[21] RNAV (GNSS) approach: area navigation global navigation satellite system non-precision approach. Previously termed a GPS approach.

[22] An accident is defined in the Australian *Transport Safety Investigation Act 2003* as an investigable matter involving an aircraft where a person dies or suffers a serious injury, or the aircraft is destroyed or seriously damaged. Under the *Transport Safety Investigation Regulations 2003*, a

1.3 Damage to aircraft

There was significant damage to overhead fittings in the cabin, consistent with passengers or crewmembers being thrown around the cabin during the first in-flight upset. A more detailed description of the cabin damage is provided in section 4.3. No other damage to the aircraft was identified during an inspection at Learmonth (see also section 1.12.1).

1.4 Other damage

No structures or objects external to the aircraft were damaged.

1.5 Personnel information

	Captain	First Officer	Second Officer
Licence[23]	ATPL, issued 14 May 1992	ATPL, issued 10 Oct 2001	CPL, issued 8 Dec 2004
Total flying hours	13,592	11,650	2,070
Total command hours	7,505	2,020	1,400
Total A330 hours	2,453	1,870	480
Hours last 90 days	165	198	188
Hours last 30 days	64	78	62
Hours last 7 days	29.7	29.7	29.7
Hours last 24 hours	4.4	4.4	4.4
A330 endorsement	17 Feb 2004 (command)	7 Oct 2005 (command)	4 Jan 2008 (copilot)
Last cyclic proficiency check	26 Jul 2008	6 Aug 2008	15 Aug 2008
Class 1 Medical Certificate expiry	23 Jun 2009	29 Jul 2009	18 Nov 2008

All of the crew were appropriately qualified and licensed to conduct the flight.

In addition to 2,453 hours in command on A330 aircraft, the captain also had 3,272 hours in command on Boeing 767 aircraft, and experience as a first officer on Boeing 747 and 767 aircraft. He had command endorsements on the A330 (2004) and Boeing 757/767 (1993), and first officer endorsements on the 747 (1991), 757/767 (1993), and 747-400 (1993).

The first officer had a command endorsement on the A330 (2005) and 1,870 hours as a first officer on the aircraft type. He also had a command endorsement on Boeing 737-300/800 aircraft (2001), a first officer endorsement on 747-400 aircraft (1996), and previous experience as a first officer on both aircraft.

serious injury is defined as 'an injury that requires, or would usually require, admission to hospital within 7 days after the day when the injury is suffered'.

[23] ATPL: Air Transport Pilot (Aeroplane) Licence. CPL: Commercial Pilot (Aeroplane) Licence.

The second officer had a first officer endorsement on the A330 and 480 hours experience as a second officer on the aircraft type.

All of the flight crew operated a flight from Brisbane to Singapore on 5 October 2008, with 36 hours rest before commencing duty for the QF72 flight on 7 October 2008. Each of the crew reported that they had at least 7 hours sleep the night before the occurrence flight and a similar amount of sleep the night before.

None of the flight crew reported any medical or physiological problems, and no such problems were evident on the cockpit voice recording for the flight. A review of medical records held by the Civil Aviation Safety Authority (CASA) identified no problems.

Information on the qualifications and experience of the cabin crew is provided in section 4.1.3.

1.6 Aircraft information

1.6.1 General information

Aircraft type	Airbus A330-303
Manufacturer's serial number (MSN)	0553
Year of manufacture	2003
Registration	VH-QPA
Certificate of Registration	31 October 2003
Certificate of Airworthiness	26 November 2003[24]
Total airframe hours	20,040
Total airframe cycles	3,740

The Airbus A330 is a large capacity, wide-body, twin-engine aircraft, which is used for medium-to-long-range air transport operations. The A330 was developed at the same time as the four-engine A340, and most of the aircraft systems were common to both aircraft types. The A340 was first certified in Europe in December 1992 and the A330 was first certified in Europe in October 1993 (see also section 2.3.1). As of October 2008, about 570 A330s and 360 A340s had been manufactured.

1.6.2 Overview of aircraft systems

Figure 3 provides a very simplistic overview of the relationship between some of the aircraft systems used for controlling the aircraft's flightpath. In basic terms:

- The electrical flight control system (EFCS) controlled the operation of the aircraft's flight control surfaces (such as the elevators and ailerons).

[24] The aircraft was delivered new to the operator as an A330-301 model in November 2003. The original Certificate of Airworthiness was dated 26 November 2003. In December 2004 the aircraft was modified from a -301 to a -303 model, and a new Certificate of Airworthiness was issued on 10 December 2004.

- The flight crew provided inputs to the EFCS, either indirectly using the autopilot or directly using the flight crew controls. The controls included the sidestick controllers (or sidesticks) to manoeuvre the aircraft in pitch and roll, and foot-operated rudder pedals to control yaw.
- The air data and inertial reference system (ADIRS) provided information on important flight parameters such as airspeed, altitude, angle of attack (AOA), and attitude to the EFCS and autopilot, as well as to the flight displays used by the crew.

Figure 3: Simplified overview of the relationship between aircraft systems

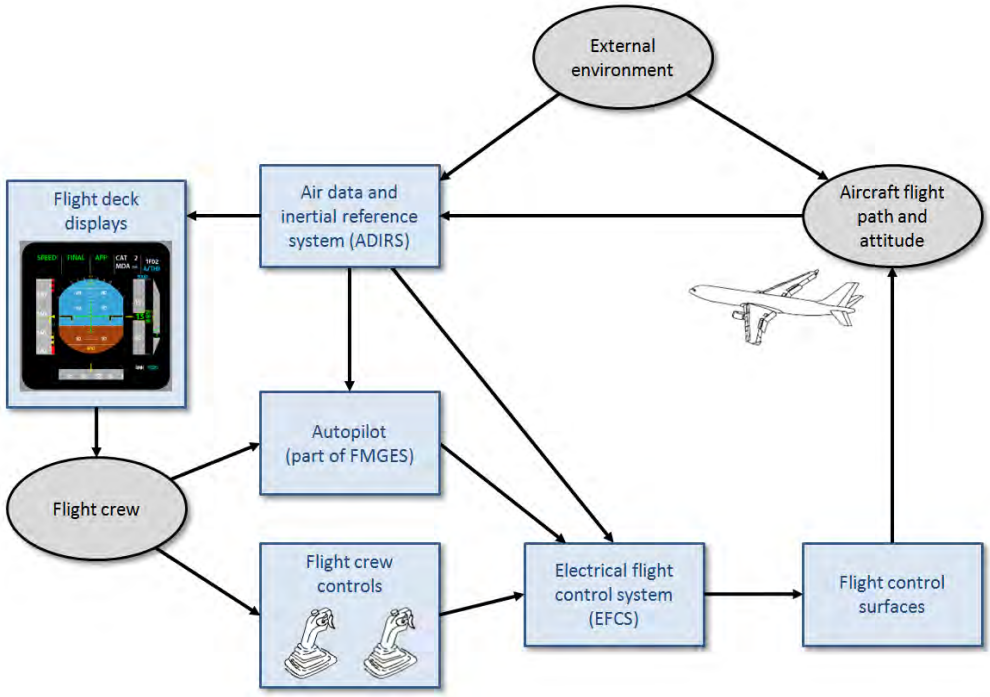

1.6.3 Electrical flight control system (EFCS)

System overview

The A330's electrical flight control system (EFCS) was a 'fly-by-wire' system. That is, there was no direct mechanical linkage between most of the flight crew's controls and the flight control surfaces.[25] Flight control computers sent movement commands via electrical signals to hydraulic actuators that were connected to the control surfaces.[26] The computers sensed the response of the control surfaces to these commands, and adjusted the commands as required.

Figure 4 provides an overview of the Airbus fly-by-wire system, and Figure 5 shows the flight control surfaces on the A330.

[25] A330s are described as either 'basic' or 'enhanced' models (QPA was an enhanced model). Basic models have a mechanically-controlled rudder while the later enhanced models have fly-by-wire rudder control. For both models, elevator trim included a mechanical backup.

[26] On a conventional airplane, inputs from the pilots or the autopilot were transmitted to the hydraulic actuators by an arrangement of mechanical components.

Figure 4: Overview of a fly-by-wire flight control system

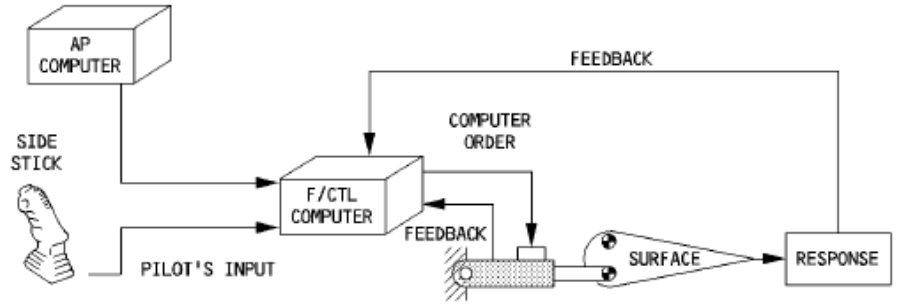

Figure 5: A330 flight control surfaces

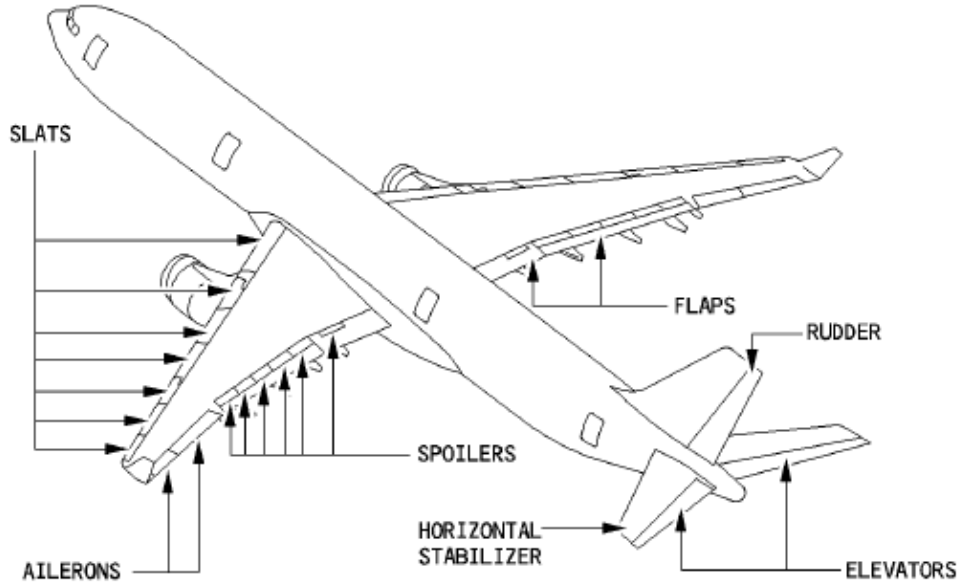

The A330 EFCS had three flight control primary computers (FCPCs, commonly known as PRIMs) and two flight control secondary computers (FCSCs, commonly known as SECs). One of the FCPCs (normally FCPC 1) acted as the 'master' FCPC. It computed the appropriate control orders, and sent these orders to the other computers to action. More detailed information regarding the functioning of the FCPCs is provided in section 2.1.

Overall, the A330's EFCS provided many advantages relative to a conventional flight control system, including stability augmentation, reduced crew workload, and flight-envelope protection.

Flight control laws

The master FCPC computed the control orders according to a 'control law', with different functionality provided depending on the law being used. There were three levels of control law, and each level provided different functionality as follows:

- Normal law. The EFCS detected when the aircraft was approaching the limits of certain flight parameters, and commanded control surface movements to prevent the aircraft from exceeding these limits (that is, it prevented the aircraft from

exceeding a predefined safe flight envelope). Automatic flight-envelope protections included high AOA protection, load factor limitation, pitch attitude protection, roll attitude protection, and high speed protection.

- Alternate law. The EFCS switched to alternate law if there were certain types or combinations of failures within the flight control system or related systems. Some types of protection, such as high AOA protection, were not provided, and others were provided using alternate logic.
- Direct law. The EFCS switched to direct law in situations where there were more failures of relevant, redundant systems in addition to those that led to the reversion to alternate law. No flight-envelope protections were provided, and control surface deflection was proportional to sidestick and rudder pedal movement by the flight crew.

Pitch control

The EFCS achieved control of the aircraft's pitch by using two elevators and the trimmable horizontal stabiliser (THS). The elevators provided short-term changes in pitch, and the THS provided longer-term changes where needed so that continuous elevator deflections were not required.

Maximum elevator deflection was 30° nose-up and 15° nose-down. FCPC 1 normally controlled the elevators (see section 2.1.1 for further details).

Maximum THS deflection was 14° nose-up and 2° nose-down. In normal law and in most cases in alternate law, the EFCS automatically controlled THS movement (known as 'autotrim'). Autotrim was not available in direct law and with some types of failures in alternate law. However, the flight crew were always able to manually control the trim by using the pitch trim wheels on the flight deck's centre pedestal.

Figure 6: Trimmable horizontal stabiliser and elevators on QPA

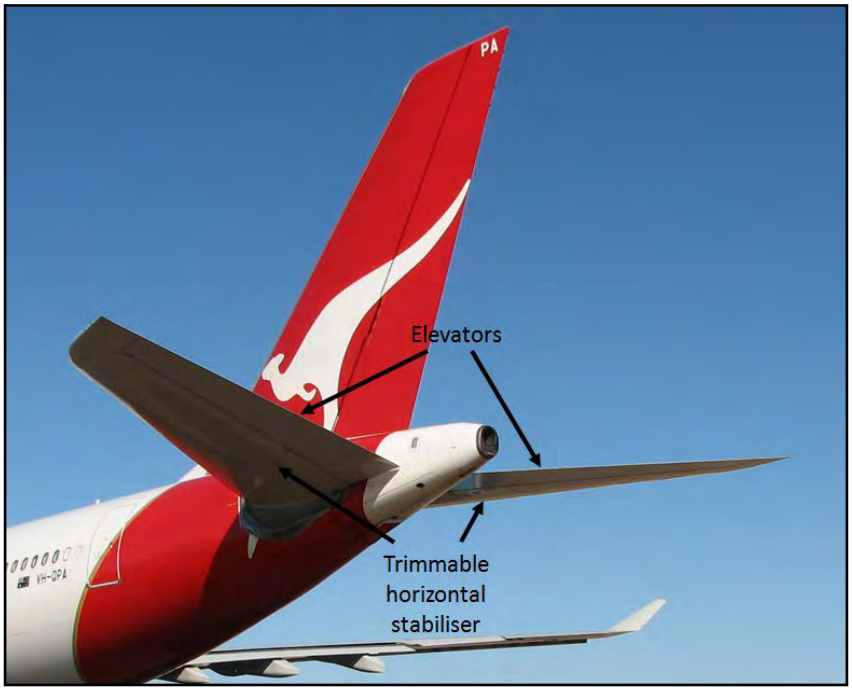

1.6.4 Air data and inertial reference system (ADIRS)

System overview

The air data and inertial reference system (ADIRS) provided important information about the outside environment (such as air pressure and temperature), the aircraft's state relative to the outside air (such as airspeed, altitude and angle of attack), and the aircraft's state relative to the Earth (position, motion and orientation).

To provide redundancy, the ADIRS included three air data inertial reference units (ADIRU 1, ADIRU 2, and ADIRU 3). Each was of the same design, provided the same information, and operated independently of the other two. ADIRU 1 from QPA is shown in Figure 7.

Figure 7: ADIRU 1 (ADIRU 4167) from QPA

Each ADIRU had two parts, an air data reference (ADR) part and an inertial reference (IR) part, which were integrated into a single unit. The two parts shared some common modules, such as the central processing unit module. In most cases, if one of the two parts failed, the other could still operate.

Overall, the ADR outputted about 30 flight data parameters and the IR outputted about 60 flight data parameters. Examples are provided in Table 1. Apart from the flight data, the ADR and IR also transmitted documentary data (such as the ADIRU part number and serial number), status data (such as the operating mode), and fault data (of the ADR, IR and system inputs).

Table 1: Examples of ADIRU flight data parameters

	Output frequency (times per second)	Recording systems
Air data parameters		
Standard altitude[27]	16	FDR, QAR
Altitude rate	16	Not recorded
Computed airspeed	8	FDR, QAR
True airspeed	8	Not recorded
Mach	8	FDR, QAR
Corrected AOA	16	FDR, QAR
Indicated AOA	16	Not recorded
Static air temperature	2	FDR, QAR
Total air temperature	2	QAR
Total air pressure	8	Not recorded
Inertial reference parameters		
Pitch attitude (angle)	50	FDR, QAR
Pitch rate	50	Not recorded
Pitch acceleration	50	Not recorded
Roll attitude (angle)	50	FDR, QAR
Roll rate	50	Not recorded
Roll acceleration	50	Not recorded
Flightpath angle	25	QAR
Flightpath acceleration	50	QAR
Groundspeed	25	FDR, QAR
Magnetic heading	25	FDR, QAR
Wind speed	10	FDR, QAR
Wind direction	10	FDR, QAR
Inertial latitude	5	Not recorded
Inertial longitude	5	Not recorded

The ADIRUs' outputs were transmitted to several other aircraft systems, including the flight displays, autopilot, and EFCS. All three ADIRUs provided data to each of the FCPCs.

The overall structure of the ADIRS is summarised in Figure 8. More detailed information on the ADIRU architecture is provided in section 3.2.

[27] The ADIRU outputted several different altitude parameters, including standard altitude (referred to as 'altitude' in this report). The other parameters are discussed in section 3.3.4.

Figure 8: Overview of the ADIRS

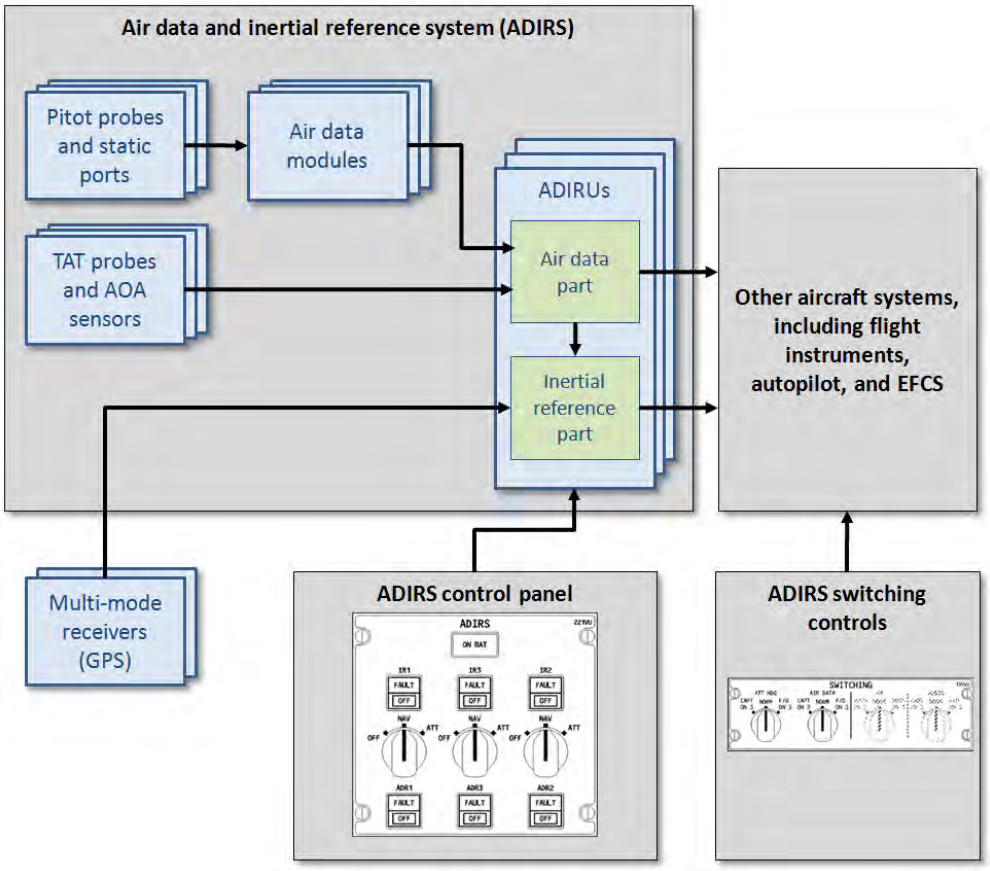

Air data reference part

The ADR part of the ADIRU provided information about the aircraft's movement through the air and atmospheric information. It obtained its inputs from sensors mounted on the aircraft's fuselage.

Each ADIRU had its own, independent sensors.[28] An AOA sensor and a total air temperature (TAT) probe provided data via analogue electrical signals directly to the ADIRU. In addition, a pitot probe and two static ports provided data to the ADIRU via air data modules (ADMs), which converted air pressure signals to digital signals.

The ADR parameters used instantaneous measurements; that is, each measurement was completely independent of previous measurements.[29] As a result, any corruption of the data did not have an ongoing effect on subsequent calculations.

---

[28] There were three AOA sensors, three pitot probes and six static ports on the aircraft. There were only two TAT probes; one provided data direct to ADIRUs 1 and 3, the other provided data direct to ADIRU 2.

[29] The one exception was the altitude rate (or vertical speed), which was dependent on the rate of change of the previous few measurements of altitude.

Inertial reference part

The IR part of the ADIRU provided information about the aircraft's position, orientation, and velocity with respect to the Earth. It obtained its data from a set of inertial instruments within each ADIRU, which continually measured acceleration in all three axes (pitch, roll and yaw) as well as rotational movement.

The IR part constantly updated the aircraft's three-dimensional position and orientation based on the movement it sensed from a known starting position and orientation. The process of determining this starting position was known as 'inertial alignment', which occurred at the beginning of each flight when the aircraft was on the ground and stationary. Subsequent inertial measurements changed the calculated position, orientation and velocity by a very small amount for each measurement cycle. As the IR parameters were dependent on previous values, an introduced error would affect subsequent values. The IR parameters were also highly interdependent, and an error in one parameter would affect other parameters.

Each ADIRU received GPS data from one of two multi-mode receivers in order to augment the inertial reference computations.

ADIRS control panel

The ADIRS control panel provided local fault indications for each part of each ADIRU (Figure 9). If there was a fault with the IR part of an ADIRU, an amber 'FAULT' light illuminated. The relevant part of the ADIRU could be deactivated by pressing the OFF pushbutton below the fault light. The ADR part of the ADIRU operated in the same manner.

Figure 9: ADIRS control panel

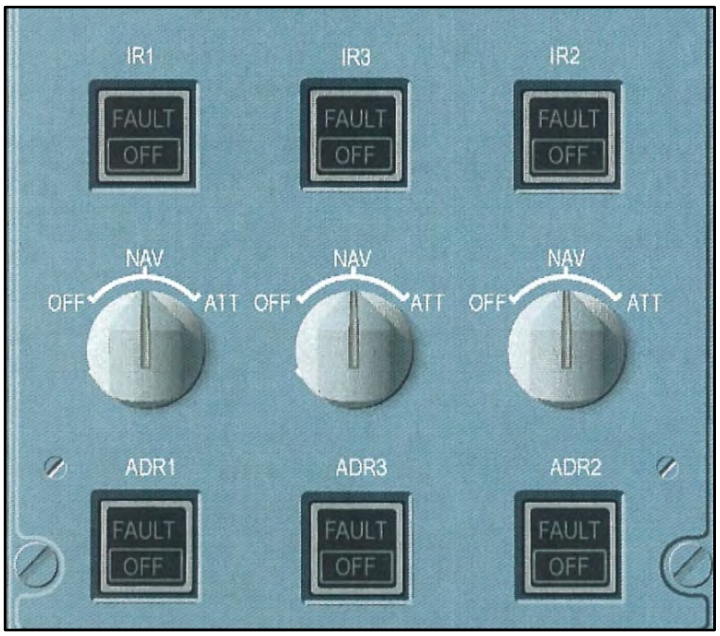

The panel also had an IR mode rotary selector for each ADIRU that allowed the flight crew to select one of three modes:

- OFF; the ADIRU was not energised and the IR and ADR parts were not available

- NAV; the ADIRU supplied full inertial data and air data to other systems (normal mode of operation)
- ATT; the ADIRU supplied full air data but limited inertial data (only attitude and heading information) to other systems.

The ADIRS control panel was located on the flight deck's overhead panel (see the upper part of Figure 10).

Figure 10: Location of relevant displays and controls

ADIRS switching controls

In normal operation, ADIRU 1 provided information for the captain's flight displays and ADIRU 2 provided information for the first officer's flight displays. In the event of a failure of ADIRU 1 (or ADIRU 2), the flight crew could manually switch the source of the information for the captain's (or first officer's) displays to ADIRU 3. This was achieved using either the ATT HDG switch for IR parameters (see the left control in Figure 11) and/or the AIR DATA switch for ADR parameters (see the second left control in Figure 11).

The ADIRS switches were located on the pedestal panel in the flight deck (see the lower part of Figure 10).

Figure 11: ATT HDG and AIR DATA switches

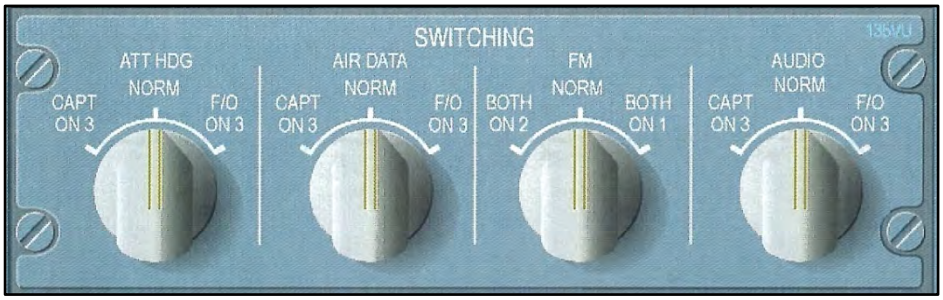

1.6.5 Processing of angle of attack information

Angle of attack

Angle of attack (AOA) is an important air data parameter that the EFCS uses to control the aircraft's pitch. AOA is a measurement of the vertical angle of a wing (using a nominal reference line known as the 'chord line') relative to the airflow (Figure 12). The AOA is not the same as the aircraft's pitch angle, which is the angle of the aircraft's body relative to the horizon.[30]

Figure 12: Angle of attack

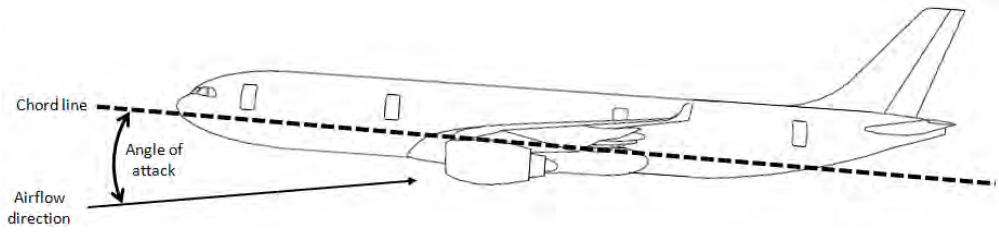

As the aircraft's AOA increases, the airflow over the wing eventually becomes more turbulent, reducing the amount of lift produced by the wing. The lift reduces rapidly as the angle is further increased; a condition known as 'aerodynamic stall'. The angle at which the amount of lift starts to reduce is known as the 'critical angle' or 'stall angle'. Stall prevention is accomplished by limiting the AOA, and stall recovery can be accomplished by reducing the AOA (by pitching the aircraft nose down).

[30] The chord line is typically inclined slightly relative to the aircraft's body.

For an A330, the typical operational range of AOA was 1 to 10° during all phases of flight. During normal cruise flight, AOA was typically about 2 to 3°.

Angle of attack sensors

The three AOA sensors (AOA 1, AOA 2, and AOA 3) were installed on the forward fuselage. AOA 1 and AOA 2 were installed on the left and right sides of the fuselage respectively, and AOA 3 was installed below AOA 2 (Figure 13). The AOA 1 sensor sent data to ADIRU 1, the AOA 2 sensor sent data to ADIRU 2, and the AOA 3 sensor sent data to ADIRU 3.

Figure 13: AOA sensor locations

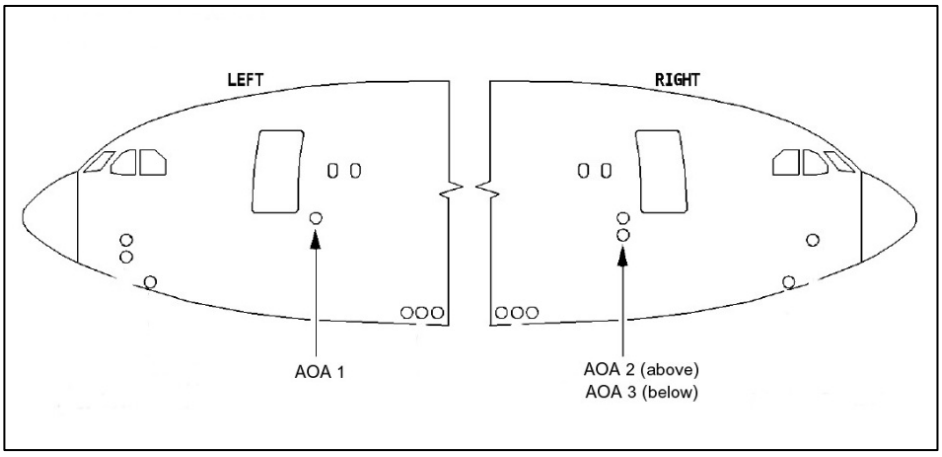

Each sensor had a vane that aligned with the airflow past the aircraft, and measured the vertical angle of this airflow relative to the aircraft's body. The vane angle had a range limit of +65 to -65°. Figure 14 shows the AOA sensors of QPA.

Figure 14: AOA sensors of QPA

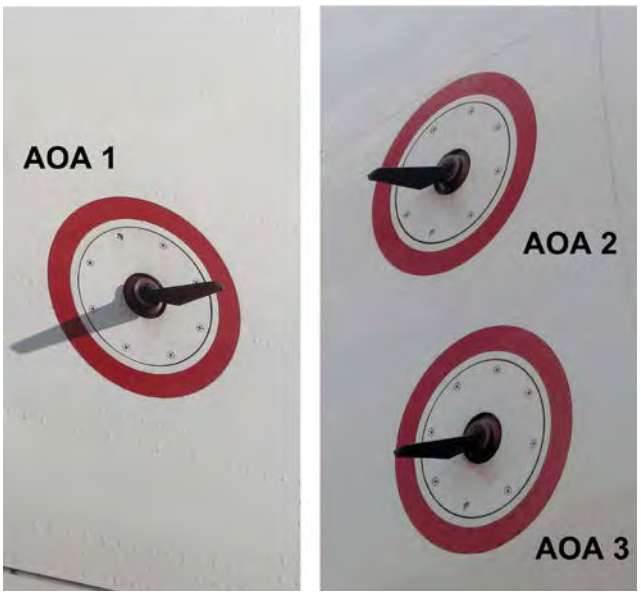

Processing by ADIRUs and FCPCs

Each AOA sensor utilised two independent outputs ('A' and 'B'). The relevant ADIRU compared the A and B signals and, if they agreed, processed the data.

To calculate 'indicated AOA', the ADIRU corrected the AOA vane angle for a 25° offset. To derive the 'corrected AOA', the ADIRU adjusted indicated AOA for the aircraft's configuration (such as slats/flaps position). The corrected AOA value was passed on to other aircraft systems, and the output range could vary from -40 to +90°.

Each ADIRU sent its AOA outputs to all three FCPCs. Figure 15 provides a simplified representation of the relationship between the AOA sensors, ADIRUs and FCPCs. Each FCPC continually monitored the data from all the ADIRUs to check for validity and consistency. Further information on the algorithm used by the FCPCs to process AOA data is provided in sections 2.1.4 and 2.1.5.

Figure 15: AOA inputs to the ADIRUs and FCPCs

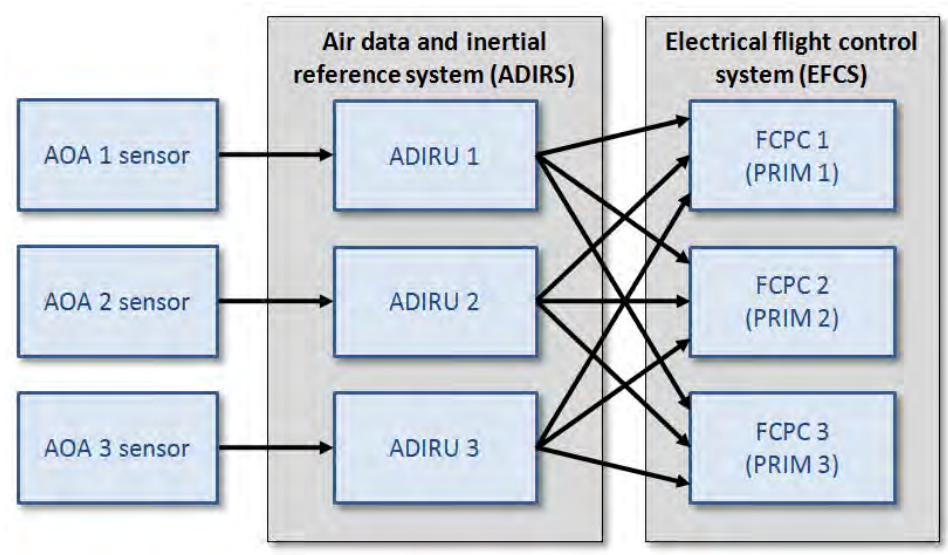

Aircraft response to high angles of attack

Under normal law on the A330, the EFCS protected the aircraft against a stall condition. If the EFCS detected that the aircraft's AOA exceeded a threshold value, it would command a pitch-down movement using the aircraft's elevators (section 2.1.7).

The aircraft's flight warning system (FWS, see section 1.6.8) also monitored the AOA data from the three ADIRUs. If it detected that the AOA was above a threshold value (which was different to that used by the EFCS), it would trigger an aural stall warning. In normal law, the threshold value of the warning was set at a high level of 23° to prevent unwarranted activations. In alternate or direct law, high AOA protection was lost and the stall warning was triggered when the highest of the valid AOA values exceeded the threshold for the flight conditions at the time.

In common with most aircraft types, AOA was not displayed to the flight crew.

1.6.6 Autopilot

The aircraft's flight management, guidance and envelope system (FMGES) provided a range of autoflight functions that were designed to minimise crew workload, increase efficiency, and eliminate many routine flight crew tasks. The FMGES consisted of two identical flight management, guidance and envelope computers (FMGECs). The flight guidance part of the FMGECs included the autopilot, autothrust and flight director functions.

The autopilot stabilised the aircraft around its centre of gravity, and acquired and tracked a flightpath. It commanded the position of the flight control surfaces (via the EFCS) for pitch, roll, and yaw.

The A330 had two autopilots. The flight crew could engage either autopilot 1 or autopilot 2 by pressing the corresponding pushbutton. Except during some runway approach situations, only one autopilot could be engaged at any time. Autopilot 1 was part of FMGEC 1 and autopilot 2 was part of FMGEC 2. Whichever autopilot was engaged, the associated FMGEC was the 'master' and in charge of autopilot computations.

FMGEC 1 computed autopilot 1 commands based on ADIRU 1 data, although this data was checked against that from the other two ADIRUs. FMGEC 2 operated in a similar way but used ADIRU 2 as its main source of data.

The autopilot could be intentionally disconnected by the crew, or it could automatically disconnect due to a number of different conditions. These conditions included a discrepancy between the values provided by the ADIRUs on a relevant parameter, such as airspeed or altitude. In most cases, the autopilot could be re-engaged by the crew after it had been automatically disconnected. Operation of autopilot 2 was unaffected by a problem with ADIRU 1 data, and autopilot 1 was unaffected by a problem with ADIRU 2 data.

1.6.7 Flight displays

The location of the captain's and first officer's primary flight displays (PFDs) are shown in Figure 10. Each PFD displayed flight information such as aircraft attitude, airspeed, altitude, vertical speed, heading, track, and autoflight information.

In normal operations, the captain's PFD displayed information from ADIRU 1 and the first officer's PFD displayed information from ADIRU 2. Each PFD could be manually switched to ADIRU 3 as a substitute information source using the ATT HDG and AIR DATA switches (Figure 11). Various 'flags' or messages were displayed on the PFD if a parameter was unavailable or invalid.

In addition to the two PFDs, primary flight information such as attitude, airspeed and altitude was also provided on the integrated standby instrument system (ISIS), mounted in the centre of the instrument panel (see to the left of the ECAM in Figure 10). Most of the information displayed by the ISIS was measured directly using air pressures sourced from the standby pitot probe and static ports, as well as an inertial measurement module internal to the ISIS.

The navigational displays (NDs), located next to the PFDs, presented aircraft position information referenced to the flight-planned route and ground-based navigation aids. They could also display weather radar information.

The information presented on the flight displays was generated by two display management computers (DMCs), based on inputs from the ADIRUs, FWS and other sources.

1.6.8 Overview of avionics fault-detection processes

Types of faults

Modern aircraft systems are composed of line-replaceable units (LRUs)[31] and other items of equipment, each of which includes many components. These units will occasionally develop physical faults or failures.[32] The faults can be permanent, transient (appear for a short time then disappear), or intermittent (occur from time to time). Transient and intermittent faults are often known as 'soft' faults.

Hard faults can be reliably reproduced during operation or testing, whereas soft faults are more difficult to reproduce unless the circumstances that induced them are well known. With soft faults, the equipment behaviour will generally return to normal after the equipment is rebooted or power is cycled.

Self-detected faults

Aircraft systems are designed so that the probability of an equipment fault is below a specified level.[33] In addition, systems are designed so that they can manage the occurrence of a fault using methods such as fault detection and equipment redundancy.

The fault-detection processes on the A330 are summarised in Figure 16. In general terms, the aircraft's systems were designed so that the primary means of detecting faults was self-detection (that is, internal detection by the unit or system itself). If a fault was self-detected, then it was easier to manage. For example, the system could flag incorrect output data as invalid and other aircraft systems could then ignore it. Alternatively, some systems did not provide outputs to other systems, and could not rely on external fault detection.

---

[31] An LRU is a unit that, due to its size, weight and connections, could be easily removed during line maintenance. It would be replaced with a serviceable unit, while the original was dispatched for repair. The ADIRUs and FCPCs are examples of LRUs.

[32] In some situations, the term 'fault' is used to refer to an anomaly in the system or unit, whereas the term 'failure' is used to refer to when the system or unit is unable to perform its functions (due to a malfunction or loss of function). In practice, manufacturers and operators use the terms interchangeably, and they are used interchangeably in this report.

[33] Regulatory requirements for aircraft systems such as the EFCS are discussed in section 2.3, and the aircraft manufacturer's equipment specification for the ADIRUs is discussed in section 3.1.1.

Figure 16: Overview of system fault reporting processes

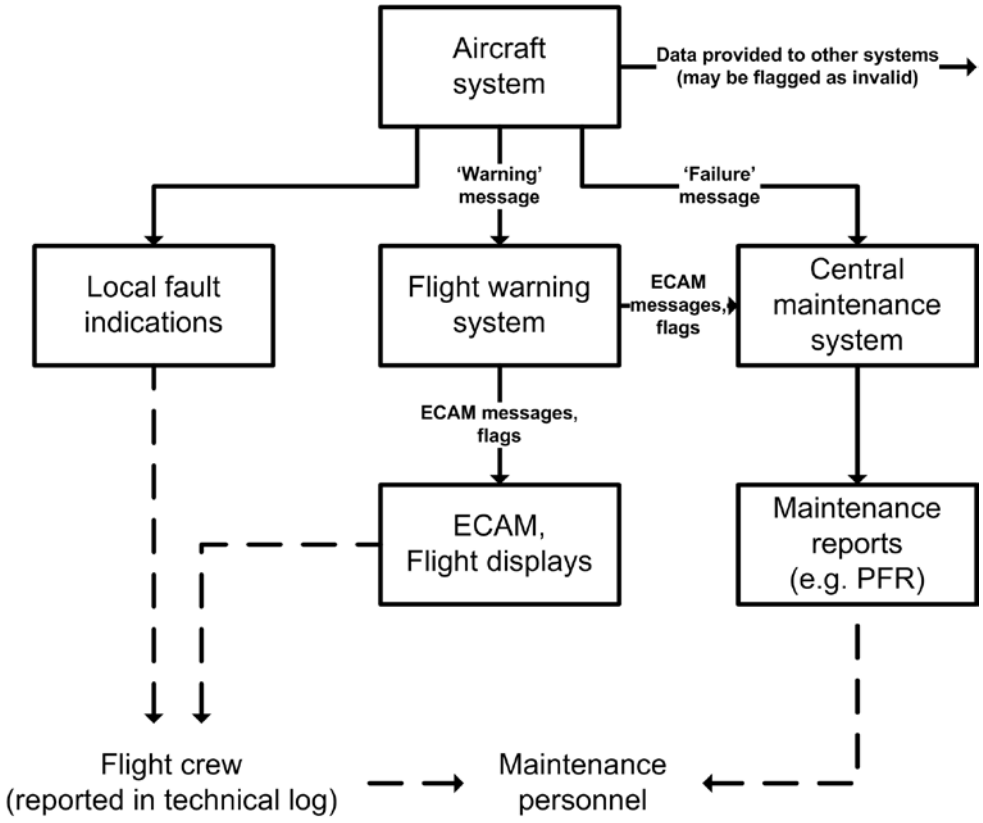

On the A330, systems could self detect faults by using built-in test equipment (BITE)[34] and, depending on the system, a wide range of tests were conducted. If a fault was self-detected, it was managed using the following actions:

- depending on the severity of the fault:
  - flag any affected output data as 'invalid' so that other systems would not use it (see below)
  - stop the transmission of any output data
  - shut down the system.
- send a message to the flight warning system (FWS), which generated ECAM messages and other associated fault indications that were provided to the flight crew
- send a signal to illuminate a local fault light on the overhead panel (if applicable)
- send a message to the central maintenance system (CMS), which included the message in maintenance reports such as the post-flight report (section 1.12.2)
- record a fault message in the system's BITE memory, which could be used in subsequent maintenance troubleshooting activities.

[34] In general, some of the LRUs within a system actually contained the BITE and conducted the fault detection, rather than the system as a whole.

A system could flag its output data as invalid by using a specific value in the sign/status matrix (SSM) field of the data. The available SSM values are listed in Table 2.

Table 2: SSM values for output data

Status	Validity	Description
Failure warning (FW)	Invalid	The transmitting (source) system detected a failure that made one or more of its output data words unreliable.
No computed data (NCD)	Invalid	The transmitting system was unable to compute reliable output data for reasons other than its own failure.
Functional test (FT)	Valid (on ground) Invalid (in-flight)	The transmitting system conducted some functional tests while the aircraft was on the ground. If this SSM value occurred in flight, then the data was considered to be invalid.
Normal operation (NO)	Valid	The transmitting system detected no problems with the output data.

Externally-detected faults

A receiving system could detect when a source system either stopped providing data, or flagged its output data as invalid (Table 2).

If a source system (such as the ADIRUs) provided incorrect data to a receiving system (such as the EFCS or FMGES), and this data was not flagged as invalid, then this fault could have safety consequences. Accordingly, some receiving systems had additional processes for monitoring input data. For example, the EFCS and FMGES compared the values of some flight data parameters that were provided by the three ADIRUs.

If a receiving system detected a problem with a source system, then it could record a fault message and provide it to the CMS.

1.6.9 Flight warning system

The aircraft's FWS monitored other aircraft systems, detected failures and unsafe flight conditions, and provided the flight crew with operational assistance for normal and abnormal aircraft system configurations. It performed these functions by:

- receiving 'failure' messages from other systems
- monitoring the data outputs of some systems (such as the ADIRUs)
- generating ECAM warning and caution messages
- activating master warning and master caution lights
- generating aural alerts and synthetic voice messages.

The FWS provided warning and caution indications that were classified at three 'failure levels', with each level based on the consequences of the fault. More serious conditions were provided with more salient aural alerts and visual indications, as outlined in Table 3.

Table 3: Failure level classifications and associated indications

Level	Significance	Aural alert	Visual indication
3	Red warning: configuration or failure that required immediate flight crew action. Included aircraft in dangerous configuration or flight condition, or a system failure affecting flight safety.	Continuous repetitive chime, specific sound or synthetic voice	Master warning light Warning message on ECAM (red) Automatic presentation of the relevant system page on the ECAM's situation display
2	Amber caution: configuration or failure that did not require immediate action, but the flight crew should be made aware. Time and situation permitting, the crew should consider the cautions without delay to prevent further degradation of the affected system. Included system failures without any direct consequence on flight safety.	Single chime	Master caution light Caution message on ECAM (amber) Automatic presentation of the relevant system page on the ECAM's situation display
1	Amber caution: situation that required crew monitoring. Included system failures leading to a loss of redundancy or system degradation.	None	Caution message on ECAM (amber), generally without any required actions

For a level 3 failure, the FWS illuminated the red, flashing master warning lights that were located on both sides of the glareshield. The FWS also produced a continuous repetitive chime or other continuous aural alert. For example, a stall warning was associated with a synthetic voice stating 'STALL STALL' followed by a 'cricket' noise. The aural alert for a level 3 failure continued until the failure condition no longer existed or the crew had cancelled the warning. Some level 3 failures, such as stall and overspeed warnings, were not associated with an ECAM message.

For a level 2 failure, the FWS illuminated the amber, steady master caution lights on both sides of the glareshield. It also produced a single aural chime. The master caution chimes could not be cancelled by the crew.

Some types of system faults were also associated with the presentation of a local fault light on the overhead panel. In that case, the light was illuminated by the relevant system itself, rather than the FWS. If the underlying condition was not resolved, the fault light generally remained illuminated, even after the associated ECAM message was cancelled by the crew.

Table 4 provides examples of the aural alerts and visual indications for faults relevant to the 7 October 2008 occurrence.

Table 4: Summary of indications for selected types of faults

Event	Aural alert	ECAM message	Other visual indication
Warnings			
Autopilot disconnect	Cavalry charge	AUTO FLT AP OFF	Master warning light
Stall warning	'Stall' synthetic voice and cricket	None	Master warning light
Overspeed warning	Repetitive chime	None	Master warning light
Cautions			
IR 1 fault	Single chime	NAV IR 1 FAULT	Master caution light, local IR fault light
ADR 1 fault	Single chime	NAV ADR 1 FAULT	Master caution light, local ADR fault light
FCPC 1 pitch fault	None	F/CTL PRIM 1 PITCH FAULT	None
FCPC 3 fault	Single chime	F/CTL PRIM 3 FAULT	Master caution light, local PRIM 3 fault light
Reversion to alternate law	Single chime	F/CTL ALTN LAW	Master caution light, indications on PFDs

1.6.10 Central maintenance system

Each electronic system that had a BITE capability sent fault information to the aircraft's central maintenance system (CMS). In addition, the FWS sent information to the CMS regarding the warning and caution messages presented to the flight crew on the ECAM or by other means.

Based on the information received, the CMS produced various reports to aid in maintenance troubleshooting and in return-to-service testing. These reports included the post-flight report (PFR), which was normally produced and printed at the end of a flight (section 1.12.2).

In addition to obtaining the PFR, maintenance personnel could interrogate the BITE information from the various aircraft systems for the most recent flight or for previous flights (section 1.12.3), and initiate functional tests of those systems (section 1.12.4). The CMS could also send fault information from the aircraft to the ground using ACARS[35] so that the messages were accessible in real time by airline maintenance personnel.

The systems that provided fault information to the CMS classified the system faults as one of three 'failure classes', as follows:

- Class 1, or faults that had a direct effect on the operation of the flight and were displayed to the flight crew. These include faults that resulted in a message on the ECAM, or warning flags on the flight displays.

[35] ACARS: Aircraft communications, addressing and reporting system. ACARS transmitted maintenance and operational messages at intervals throughout a flight.

- Class 2, or faults that did not have a direct effect on the operation of the flight but may have had an effect if there was a subsequent fault. They were accompanied by one or more 'MAINTENANCE STATUS' messages that were only brought to the attention of the flight crew via the ECAM's 'status' page once on the ground.
- Class 3, or faults that had no effect on the operation of the aircraft and were not indicated to the flight crew. These messages were therefore not included in the PFR, but could be found by interrogating a system's BITE data.

The three failure levels used by the FWS were not the same as the three failure classes used by the CMS. The relationship between failure levels and failure classes was based on the operational consequences of the fault, as shown in Table 5.

Table 5: Relationship between failure classes and failure levels

Operational consequence for the flight	Failure class (CMS)	Failure level (FWS)
May have consequences	Class 1	Level 3 Level 2 Level 1
No consequences	Class 2 Class 3	Nil

1.6.11 Electronic centralized aircraft monitor (ECAM)

System overview

The electronic centralized aircraft monitor (ECAM) provided the flight crew with information on the status of the aircraft and its systems, including any system faults. It also displayed the required flight crew actions for most normal, abnormal and emergency situations. Overall, it was a very important tool for enabling the flight crew to identify, diagnose and respond to system faults.

The ECAM provided its information on two display units that were located in the centre of the instrument panel (Figure 10):

- The upper unit, or engine/warning display (E/WD), presented information such as engine primary indications, fuel quantity information and slats/flap positions. The bottom part of the E/WD presented warning or caution messages when a system fault occurred and memo messages when there were no faults. A representation of the ECAM's E/WD display during takeoff is shown in Figure 17.
- The lower unit, or system display (SD), presented synoptic information for different systems. In some cases, the ECAM would automatically provide the relevant system's information following a system fault. The flight crew could also select different system pages. In addition, the SD could present a 'status page', which provided an operational summary of the aircraft status, including a list of inoperative systems, cancelled cautions, approach procedures and limitations (such as speed).

Figure 17: ECAM engine / warning display (E/WD)

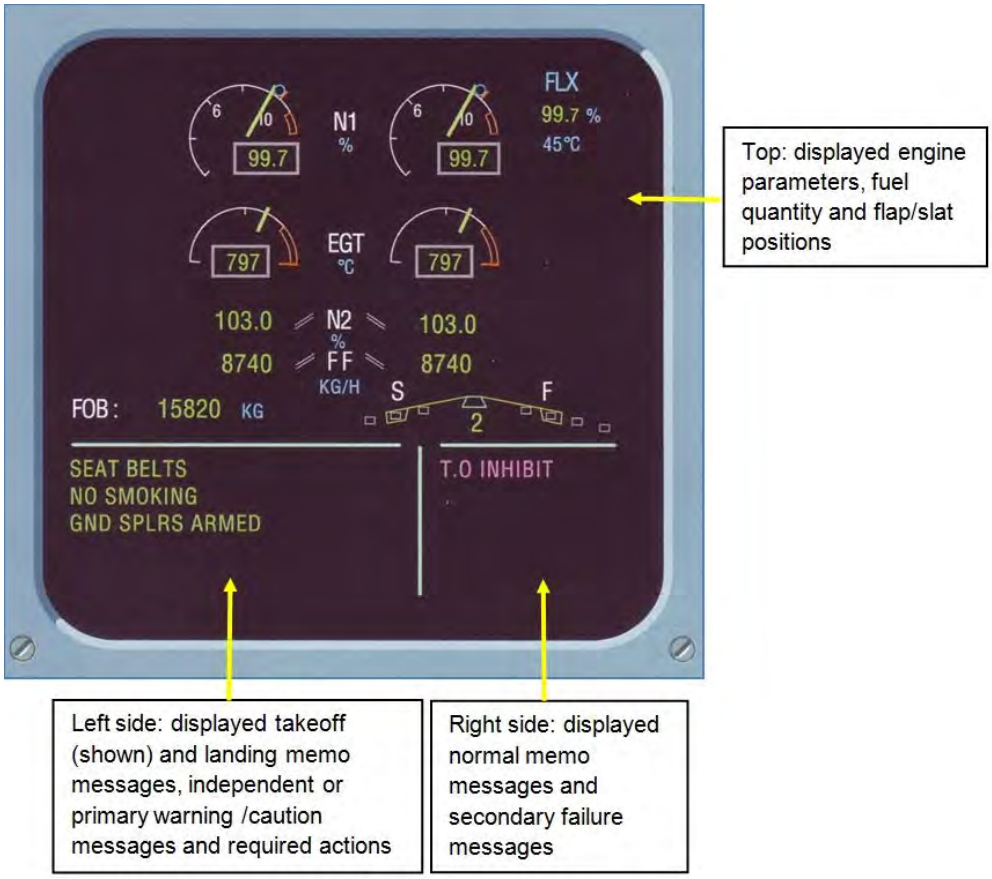

A key principle of the ECAM's design philosophy was to present information on an 'as needed' basis. For example, when displaying fault messages, it provided the appropriate emergency/abnormal procedures in addition to associated synoptic information.

Presentation of ECAM warning and caution messages

The ECAM presented a short message indicating the nature of a warning or caution in red or amber, depending on the failure level. Any required crew actions were displayed in blue text on separate lines below the relevant message. Figure 18 provides more detail on the presentation of warning and caution messages and their associated required actions.

Figure 18: Example of ECAM warning and caution messages

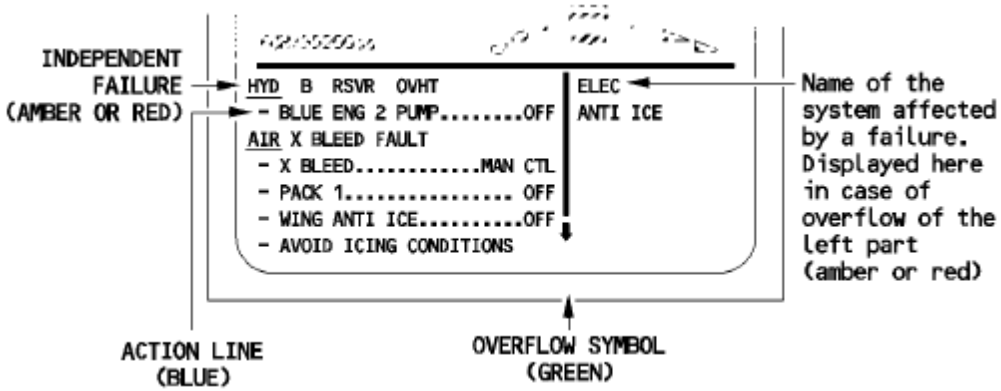

There were seven lines available at the bottom of the E/WD to display warning and caution messages.[36] The messages were displayed in a priority order, with the most important messages displayed at the top. Level 3 messages were displayed above level 2 messages, which were displayed above level 1 messages. When there were multiple messages at the same level, the most recent message had the highest priority.

If the flight crew completed a displayed action, the ECAM automatically removed the action line below the relevant message. The flight crew could also clear a message by pressing the 'clear' pushbutton. If the conditions that led to the presentation of a warning or caution message were no longer present, the ECAM automatically removed the message. If the conditions for the message returned, the message was again displayed.

Relevant flight crew procedures

The operator's A330 *Flight Crew Operating Manual* (FCOM) was based on the aircraft manufacturer's manual. Volume 3 of the manual contained standard operating procedures and abnormal and emergency procedures. For most of the abnormal and emergency procedures, the ECAM and the FCOM presentations were consistent, with the procedures organised under the relevant ECAM warning or caution message.[37]

Required actions associated with some of the ECAM messages relevant to the occurrence flight are presented in Table 6.

---

[36] As indicated in Figure 18, a green arrow at the bottom of the screen indicated whether there were more messages than could be displayed in the space available.

[37] For a small number of abnormal or emergency situations, the relevant procedure was not associated with an ECAM message. For some of those situations, flight crew were required to complete the relevant procedure from memory and not refer to the manuals (for example, emergency descent or unreliable airspeed indications). In other situations, temporary procedures of significant importance could be promulgated as Operations Engineering Bulletins (OEBs). All of the procedures relevant to this occurrence were available for display via the ECAM.

Table 6: Required actions associated with selected ECAM messages

ECAM message	Required actions
AUTO FLT AP OFF	No required actions (for crew awareness only).
NAV IR 1 FAULT	Move ATT HDG switch to CAPT ON 3 position [Note. In some cases additional required actions could be displayed. These included selecting the IR pushbutton switch to OFF. However, there was no indication on the cockpit voice recording or from crew interviews that this action was displayed on the occurrence flight.]
NAV ADR 1 FAULT	Move AIR DATA switch to the CAPT ON 3 position Select ADR 1 pushbutton OFF [Note. This ECAM caution did not appear until 0513, and the length of time that it was displayed could not be determined. The local ADR fault light on the overhead panel was not illuminated (section 1.12.6).]
F/CTL PRIM 1 PITCH FAULT	No required actions (for crew awareness only).
F/CTL PRIM 3 FAULT	Select PRIM 3 [FCPC 3] OFF then ON [Note. If this procedure was unsuccessful, the required action was to select FCPC 3 OFF.]
F/CTL ALTN LAW (PROT LOST)	Do not use speed brake Maximum speed limited to 330 kts / Mach 0.82

1.6.12 Weight and balance

The take-off weight of the aircraft on the 7 October 2008 flight was 207,065 kg, which was below the maximum take-off weight of 233,000 kg. The aircraft's centre of gravity (c.g.) at the time of the occurrence was within normal operating limits.[38]

Inspection of the aircraft's cargo area after landing found that all cargo was loaded in the same position as that recorded on the load manifest for the flight, and no load shift was evident. All of the cargo containers and palletised cargo remained properly secured by the integral cargo restraint systems built into the floor of the cargo holds. Each individual freight container and pallet was also examined for load shift or break out of individual items from within each unit, and no such problems were evident.

1.6.13 Aircraft maintenance

QPA underwent a C-check[39] in March 2008 (prior to the occurrence) and in May/June 2009 (after the occurrence). These checks found nothing of relevance to the investigation.

---

[38] The FDR recorded the aircraft gross weight at the time of the first in-flight upset as 185,280 kg and that the c.g. was 28.2% of the reference distance. For this gross weight, the allowable in-flight c.g. range was 14–41% of the reference distance.

[39] Aircraft checks included checks conducted before the first flight each day and before each flight (described as line maintenance), as well as more detailed scheduled maintenance checks that were called A, B, C and D checks. The tasks involved and intervals between checks increase from A through to D checks. An A check could occur every few days, while a D check could be required

The aircraft had a technical log that contained records of in-service maintenance, including reports from the flight crew. A review of technical log entries for the aircraft's relevant systems identified a previous event on 12 September 2006 that involved similar warnings and caution messages but no pitch-down events (section 1.16.2).

The investigation also reviewed the aircraft's maintenance records, focusing on the ADIRUs, FCPCs, FCSCs, FMGECs, AOA sensors and probe heat computers (PHCs). The review covered component and defect histories, modification status, service bulletins, task cards and maintenance schedules. Nothing else of relevance to the investigation was found. Further details on the service history of the relevant systems are provided in section 1.12.

1.7 Meteorological information

The flight crew reported that, at the time of the occurrence, the weather was fine and clear and there was no turbulence. Cabin crew and passengers provided similar reports.

An assessment of the weather conditions by the Australian Bureau of Meteorology stated that, at the time of the occurrence, the aircraft appeared to be in the vicinity of the sub-tropical jet stream and well south of any significant convection activity. Turbulence at a moderate or greater level was unlikely to have influenced the aircraft at the time of the occurrence.

An examination of information from the aircraft's FDR found that the vertical acceleration data prior to and during the two in-flight upsets was not consistent with the effects of moderate or severe turbulence (section 1.11.5).

1.8 Aids to navigation

Not applicable to this occurrence.

1.9 Communications

For external voice communications, the aircraft was equipped with two high-frequency (HF) radios, three very high frequency (VHF) radios, and a satellite communications system (SATCOM). The aircraft could also transmit its position automatically using the automatic dependent surveillance-broadcast (ADS-B) system.

The aircraft was also equipped with an aircraft communications, addressing and reporting system (ACARS). ACARS used VHF radio or SATCOM to transmit routine flight operations and engineering data to the operator's maintenance watch personnel. ACARS transmitted data intermittently as required. Fault messages recorded by the CMS were included in ACARS reports, and maintenance watch was able to view these reports during their communications with a flight crew.

---

every 10 years, commensurate with the increasing complexity of those checks. A and B checks were generally performed at airport gates while C and D checks were performed in maintenance hangars.

For internal aircraft communications, the aircraft had a cabin intercommunication data system. Its functions included a passenger address (PA) system for public announcements, and a cabin interphone system for communication between the flight and cabin crew stations.

No problems were reported with the serviceability of any of the external or internal communication systems.

1.10 Airport information

Learmonth Airport, which was near Exmouth in Western Australia, was routinely used by air transport aircraft, and was listed as an 'alternate' for international flights.[40] The airport had a single, sealed runway (runway 18/36), which was 3,047 m long and 45 m wide. Facilities included runway and approach lighting, a precision approach path indicator, and non-precision navigational aids.

1.11 Flight recorders

The aircraft was fitted with three flight recorders: a cockpit voice recorder (CVR), a flight data recorder (FDR), and a quick access recorder (QAR).

1.11.1 Cockpit voice recorder (CVR)

System description

The aircraft's CVR recorded the total audio environment in the cockpit area, which included crew conversations, radio transmissions, aural alarms, switch activations, and engine and airflow noise.[41] It retained the last 2 hours of information in solid-state memory, operating on an endless-loop principle.

The CVR data was successfully downloaded by the ATSB. Analysis of the audio showed that power was removed from the CVR soon after the aircraft arrived at the terminal at Learmonth. As a consequence, the CVR retained the recorded audio from prior to the initial autopilot disconnection and including both pitch-down events.

Recorded aural alerts

There were no recorded warnings or cautions in the period prior to the autopilot disconnection alert at 0440:28. The disconnection alert lasted 3 seconds before it was cancelled by the crew. After that point, there were frequent occurrences of the following aural alerts:

- Master caution chimes. The first master caution chime occurred at 0440:33. During the period 0440:28 (autopilot disconnection) to 0442:27 (first

---

[40] An alternate is an aerodrome specified in a flight plan to which a flight may proceed when it becomes inadvisable to land at, or continue toward, the original destination.

[41] The aircraft's CVR was manufactured by L3 Communications (model FA2100; part number 2100-1020-02; serial number 000252164).

pitch-down) there were at least 22 master caution chimes.[42] The chimes continued for the remainder of the flight.

- Stall warnings. The first stall warning occurred at 0440:45, and there were at least 10 stall warnings in the 2-minute period between the autopilot disconnection and the first pitch-down. The warnings continued for the remainder of the flight. In all cases the stall warnings were brief; on some occasions they were truncated before the first 'STALL' was annunciated, and in all cases they were truncated before a full cycle of the warning (that is, 'STALL STALL' followed by a cricket noise).
- Continuous repetitive chimes. This type of alert was used for several different types of warnings (that is, level 3 failures). Based on a comparison with the FDR data and other information sources, the only warning conditions present on the flight, other than the autopilot disconnection and the stall warnings, were overspeed warnings (also discussed in section 3.3.4). The first aural overspeed warning occurred at 0440:37, and there were at least seven warnings in the 2-minute period between the autopilot disconnection and the first pitch-down. The warnings continued for the remainder of the flight, but they were less frequent than the caution chimes and stall warnings. In all cases the aural signals were brief and they were truncated after two or less chimes.

1.11.2 Flight data recorder (FDR)

System description

The aircraft's FDR recorded approximately 1,100 parameters of aircraft flight data.[43] It used solid-state memory as the recording medium and operated on an endless-loop principle.

The FDR data was successfully downloaded by the ATSB, and it included data for over 217 hours of aircraft operation, comprising the occurrence flight and 24 previous flights.

Overview of FDR data

The FDR data showed that the flight was uneventful until 0440:26, when anomalies in the recorded data for the ADIRU 1 parameters commenced. The first recorded anomaly was a deviation in the pitch attitude. At 0440:28 the autopilot disengaged and data spikes[44] became evident on all of the recorded ADIRU 1 parameters.

Table 7 provides a summary of the key events recorded on the FDR. Further details on ADIRU-related data are provided in the rest of this section. Further information relating to AOA data is provided in section 1.11.4, and further information about the pitch-down events is provided in section 1.11.5.

---

[42] Not all of the master caution chimes would have resulted in a new message on the ECAM (section 1.6.11).

[43] The aircraft's FDR was manufactured by L3 Communications (model FA2100; part number 2100-4043-02; serial number 000428627).

[44] A spike is a short-duration change in the value of a parameter that exceeds (or is below) the normal value by a large amount.

Table 7: Sequence of events (from the FDR)

Time	Time to event[45]	Event
0132:02	-0310:25	Takeoff from Singapore
0201:16	-0241:11	Aircraft reached top of climb (37,000 ft)
0440:26	**-0002:01**	**Start of incorrect ADIRU 1 data** (oscillation in pitch attitude)
0440:28	-0001:59	Autopilot 1 disconnected (involuntary)
0440:28	-0001:59	First of many master warnings
0440:29	-0001:58	First of many master cautions
0440:31	-0001:56	IR 1 Fail indication commenced (remained for rest of the flight)
0440:34	-0001:53	First of many AOA 1 spikes (+50.6°)
0440:41	-0001:46	First ADR 1 Fail indication (less than 4 seconds duration)
0440:50	-0001:37	First of many stall warnings
0440:54	-0001:33	First of many overspeed warnings
0441:12	-0001:15	Autopilot 2 engaged
0441:14	-0001:13	Aircraft reached 37,180 ft then began to descend to 37,000 ft
0441:28	-0000:59	Autopilot 2 disconnected
0442:27	**0000:00**	**First pitch-down event commenced**
0442:28	0000:01	Captain applied back pressure to the sidestick
0442:28	0000:01	Maximum nose-down elevator position +10.3°
0442:29	0000:02	Minimum vertical acceleration -0.80 g, pitch angle -8.4°
0442:30	0000:03	FCPC master changed from FCPC 1 to FCPC 2
0442:31	0000:04	Maximum vertical acceleration +1.56 g recorded
0442:31	0000:04	FCPC 3 (PRIM 3) Fault (remained for 120 seconds)
0443:45	0001:18	Crew switched IR source for captain's displays from IR 1 to IR 3
0444:31	0002:04	Crew reset FCPC 3
0445:08	**0002:41**	**Second pitch-down event commenced**
0445:09	0002:42	Captain applied back pressure to the sidestick
0445:10	0002:43	FCPC master changed from FCPC 2 to FCPC 1
0445:11	0002:44	FCPC 3 Fault (remained for rest of the flight)
0445:11	0002:44	Flight control law changed from 'normal law' to 'alternate law'
0445:11	0002:44	Maximum nose-down elevator position +5.4°
0445:12	0002:45	Minimum vertical acceleration +0.20 g, pitch angle -3.5°
0445:13	0002:46	Maximum vertical acceleration +1.54 g
0447:25	0004:58	Autothrust disengaged
0450:24	0007:57	Aircraft changed heading (to divert to Learmonth)
0532:08	0049:41	Aircraft touched down at Learmonth
0542:12	0059:45	Aircraft stopped at terminal

[45] Total time prior to or after the first pitch-down (hours minutes:seconds).

FDR process for recording ADIRU data

The FDR recorded 11 ADIRU parameters (Table 1), and sampled most of these parameters once per second.

For all of the parameters except AOA, the FDR only recorded the data from one of the ADIRUs. The priority source of the recorded parameters was the same as that selected for the captain's flight displays (using the ADIRS switching controls discussed in section 1.6.4). More specifically:

- If the captain's AIR DATA switch was set to NORM, then ADR 1 was the priority source of the ADR parameters recorded on the FDR. If the switch was set to CAPT ON 3, ADR 3 was the priority source. The switch position was not changed during the occurrence flight.
- If the captain's ATT HDG switch was set to NORM, then IR 1 was the priority source of the IR parameters recorded on the FDR. If the switch was set to CAPT ON 3, IR 3 was the priority source. The crew selected the switch to the CAPT ON 3 position at 0443:45.

In addition, if ADIRU 1 flagged one of its output parameters as 'invalid' (section 1.6.8), then the FDR recorded that parameter from a different ADIRU. This switching was done on a parameter-by-parameter basis, and the FDR returned to using ADIRU 1 when that ADIRU again flagged its output data for the parameter as valid.

The only exception to this parameter switching was for AOA.[46] The FDR explicitly recorded two AOA values: one from the captain's ADR source and one from the first officer's AOA source. For this occurrence, as the ADR switch remained in the NORM position, the FDR recorded AOA 1 and AOA 2 for the entire flight. Because two values were recorded, the recording logic was configured so that there was no switching of the source even if the data was flagged as 'invalid'.[47]

The FDR also recorded a separate 'Fail' parameter for ADR 1, IR 1, ADR 2, IR 2, ADR 3, and IR 3. Each of these FDR parameters showed a fail value if a fault was indicated by the relevant ADR or IR itself, or if the ADIRU flagged key parameters as invalid for more than 500 msec.[48] In other words, the presence of a fail value for ADR 1 or IR 1 indicated that the FDR had switched to a new source for at least some parameters.[49]

---

[46] As discussed in section 1.16.4, each ADIRU outputted indicated and corrected AOA. The FDR and QAR only recorded corrected AOA, and all AOA data discussed in this report is corrected AOA.

[47] As the AIR DATA switch remained in the NORM position for the entire flight, the source of AOA 1 was always ADR 1 (that is, the ADR part of ADIRU 1), and the source of AOA 2 was always ADR 2.

[48] The ADR parameters were standard altitude and/or computed airspeed, and the IR parameters were pitch attitude and/or roll attitude.

[49] The fail indications on the FDR did not correlate exactly with the NAV IR 1 and NAV ADR 1 ECAM caution messages that were generated by the FWS. The inputs for the FDR's fail parameters were provided by the aircraft's DMCs rather than from the FWS. The triggering logic and conditions (for example confirmation times) for the FWCs and DMCs were not identical. Further information on the DMCs is provided in Appendix B.

In summary, the source of most of the ADR and IR data recorded by the FDR could vary. Table 8 provides a summary of the sources for the occurrence flight, based on the actual FDR data (discussed below).

Table 8: FDR's source of ADR and IR parameters (except AOA)

Time	AIR DATA switch position	Source of ADR data	ATT HDG switch position	Source of IR data
Takeoff to 0440:28	NORM	ADR 1	NORM	IR 1
0440:28 to 0441:37	NORM	ADR 1/ADR 2	NORM	IR 2
0441:37 to 0443:45	NORM	ADR 1	NORM	IR 2
0443:45 to 0506:48	NORM	ADR 1	CAPT ON 3	IR 3
0506:48 to landing	NORM	ADR 1/ADR 2	CAPT ON 3	IR 3

Summary of ADR data for the occurrence flight (from the FDR)

Soon after 0440:26, spikes became evident on all of the recorded ADR parameters from ADIRU 1 (Figure 19). The recorded spikes occurred at different times for each parameter, and the spikes for many of the parameters contained repetitions of a small number of values. As the data output rate of the ADIRUs was much higher than the recording rate of the FDR[50], the data available to the investigation was incomplete, and it was not possible to determine the full nature, magnitude and frequency of the data spikes. Further discussion of the data-spike characteristics is provided in section 3.3.

From 0440:41 to 0441:37, the ADR 1 Fail parameter intermittently recorded a fail value (Figure 20), which indicated that at least one of the ADR 1 parameters (either standard altitude or computed airspeed) was flagged as invalid during this period. From 0441:38 to 0506:48, the ADR 1 Fail parameter did not record any fail indications; however, the FDR recorded many spikes in all recorded ADR parameters (including standard altitude and computed airspeed) during that time. Consequently, because no source switching had occurred, it was evident that ADIRU 1 had transmitted data spikes to other systems as valid data.

The FDR did not record any fail values for ADR 2 or ADR 3.

[50] As examples, the ADIRU outputted Mach eight times per second and the FDR recorded it once per second, and the ADIRU outputted pitch attitude 50 times per second and the FDR recorded it four times per second.

Figure 19: FDR plot showing key ADIRU parameters

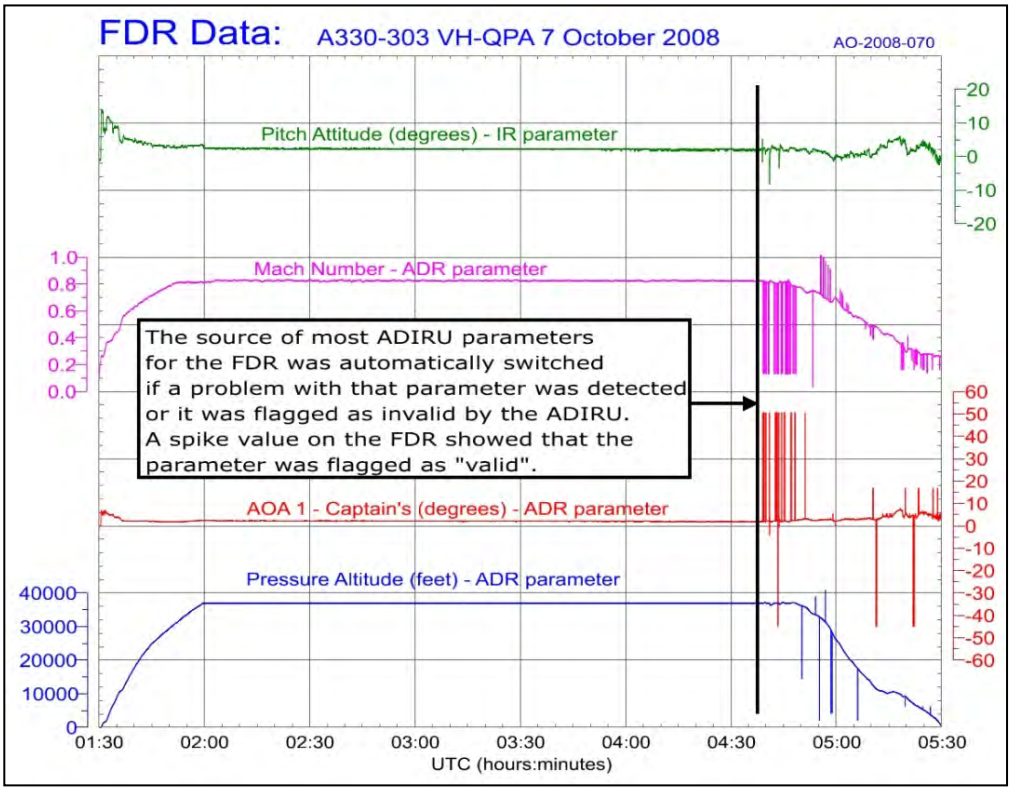

Figure 20: ADR 1 and IR 1 Fail parameters (from the FDR)

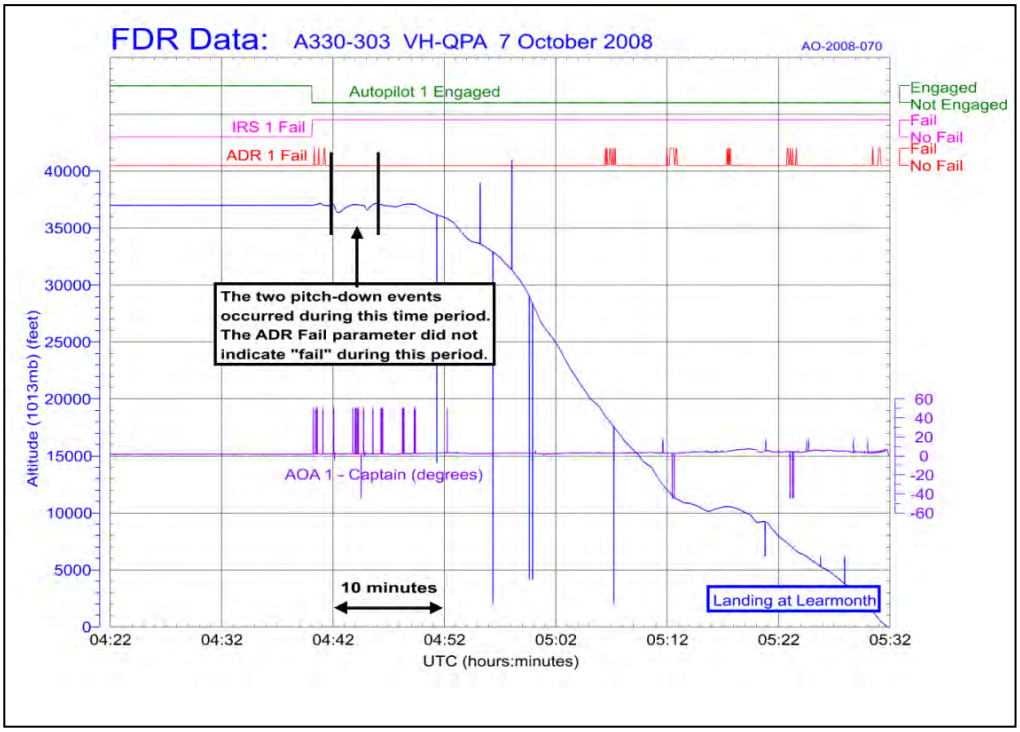

Summary of IR data for the occurrence flight (from the FDR)

The FDR only recorded two data spikes for the IR parameters, both for magnetic heading (at 0441:38 and 0442:04).

From 0440:26 to 0440:31, some of the IR parameters showed deviations from their correct values. After 0440:31, the IR 1 Fail parameter continuously recorded a fail value (Figure 20). This result indicated that, after this time, ADIRU 1 flagged all of its IR 1 data for at least one of its parameters (either pitch attitude or roll attitude) as invalid, and that the FDR's source of IR data had quickly switched from IR 1 to IR 2 for one or more parameters. From 0443:45, when the crew selected the ATT HDG switch to CAPT ON 3, the IR 3 became the source of all the IR parameters recorded by the FDR, and no deviations or spikes were recorded.

The FDR did not record any fail values for IR 2 or IR 3.

1.11.3 Quick access recorder (QAR)

System description

The aircraft's QAR recorded approximately 250 aircraft flight data parameters.[51] The data was stored on a removable magneto-optical disk with a capacity of 230 megabytes.

Files stored on the QAR disk were successfully recovered by the ATSB. They contained flight data from the occurrence flight and six previous flights.

The QAR and FDR obtained their data via different signal paths (Appendix B).

QAR process for recording ADIRU data

The QAR recorded 16 ADIRU parameters, including the 11 parameters that were recorded by the FDR (Table 1). As with the FDR, the QAR sampled most of the parameters once per second. However, the FDR and QAR were independent systems and they sampled the parameters at different times. As with the FDR, both AOA 1 and AOA 2 were recorded, but only one source of the other parameters was recorded.

The QAR received its data from a data management unit (DMU), and the source of the ADIRU parameters was configurable when the DMU was programmed. For the occurrence aircraft, the source of most ADIRU parameters was fixed to ADIRU 1. Unlike with the FDR, there was no switching of the source recorded by the QAR even if ADIRU 1 flagged a parameter as being invalid. The only exceptions were computed airspeed and magnetic heading; if the QAR data for either of those parameters was flagged as invalid by ADIRU 1, then the QAR's source switched to ADIRU 2.

---

[51] Unlike the CVR and FDR, the QAR was not required to be installed by regulation. However, operators elected to install the recorder to enable routine and easily accessible monitoring of aircraft and flight crew performance. As the parameters recorded by the QAR were configurable by an operator, it was described as a Digital ACMS Recorder (DAR) in Airbus terminology. To avoid confusion, the generic term QAR is used in this report. ACMS is an abbreviation for Aircraft Condition Monitoring System.

Summary of ADR and IR data for the occurrence flight (from the QAR)

The QAR data showed that the flight was uneventful until 0440:26 when spikes became evident on all of the recorded ADR and IR parameters from ADIRU 1. The data spikes continued until the aircraft landed. The QAR data for several key ADIRU parameters are shown in Figure 21. As indicated, frequent spikes were recorded in IR parameters, and intermittent spikes were recorded for ADR parameters.

When there were no spikes, the ADR parameters appeared to be correct. In addition to the data spikes, all of the IR parameters showed persistent deviations from their expected values. For most of these parameters, the deviations showed oscillatory behaviour (Figure 22). The groundspeed parameter showed a different characteristic, gradually diverging from the actual value and increasing to 1,000 kts, where the value remained for the rest of the flight, except for a number of lower-value data spikes.

A comparison between the overall FDR and QAR data patterns found:

- The presence of spikes for all ADR parameters on both the FDR and QAR throughout the flight after 0440:26, indicating that ADIRU 1 flagged most of its ADR output data (including the data spikes) as valid.
- The presence of spikes for all IR parameters on the QAR but not the FDR during the period from 0440:26 to 0443:45 (when the ATT HDG switch position was changed), indicating that ADIRU 1 flagged most of its IR output data as invalid.

Figure 21: QAR plot showing key ADIRU parameters

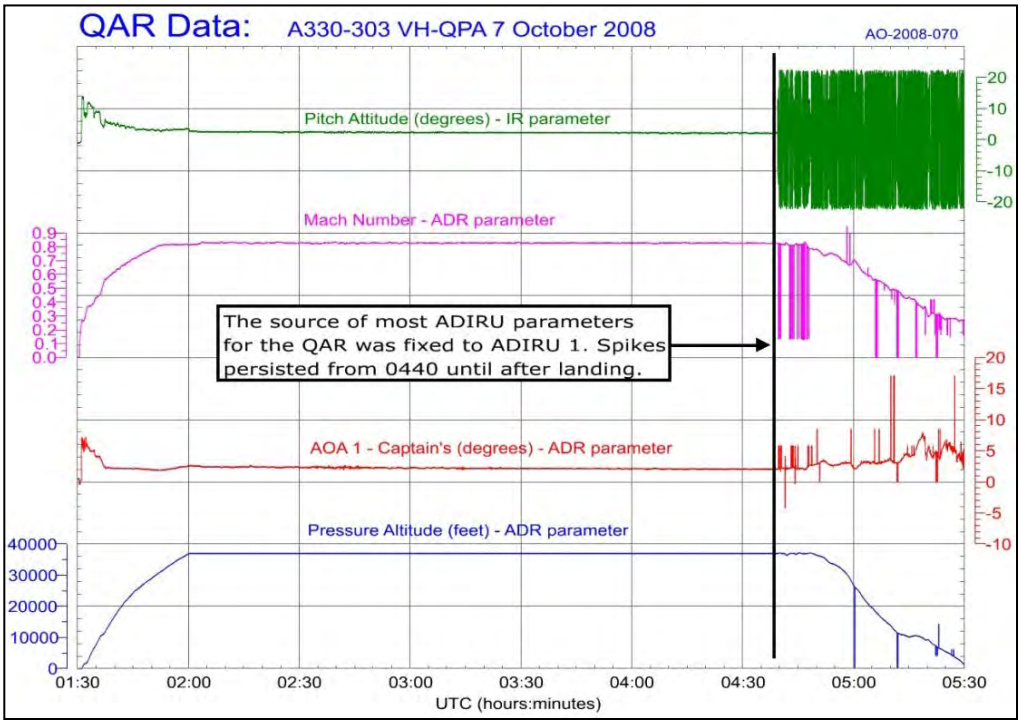

Figure 22: QAR plot showing oscillations in IR parameters

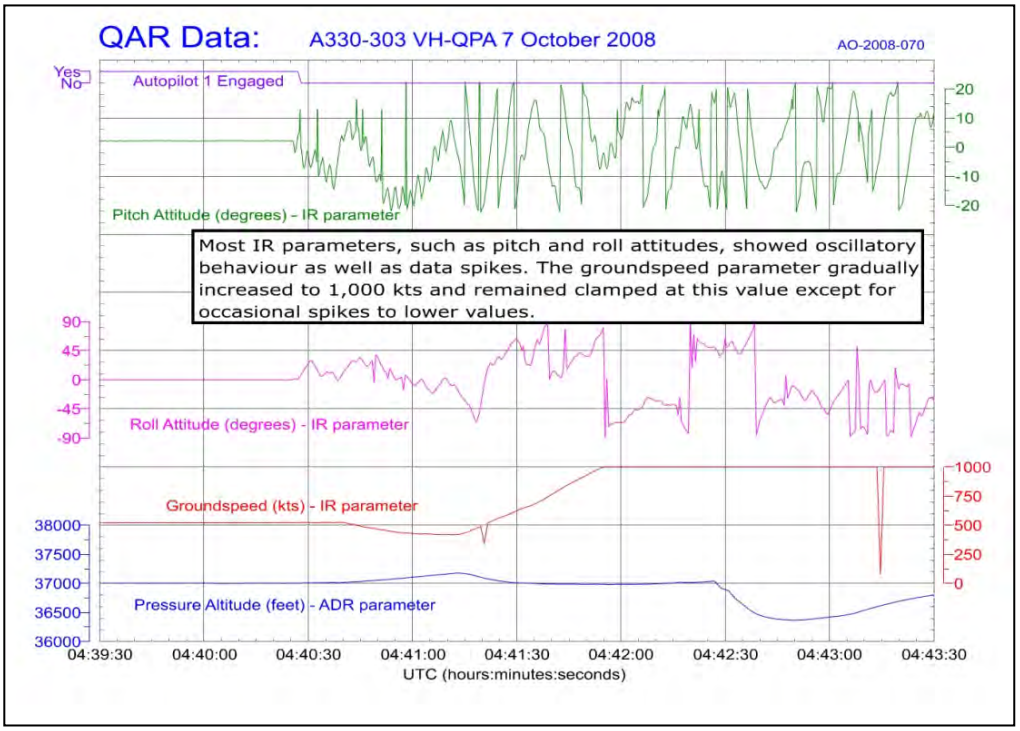

1.11.4 Recorded angle of attack data

FDR data

The FDR recorded both AOA 1 and AOA 2 values once per second. The ADIRU transmitted corrected AOA data16 times per second but the FDR only sampled it once per second. As a result, it was not possible to determine the nature, magnitude and frequency of all the erroneous AOA data transmitted by the ADIRU. However, no spikes or problems were evident at any time for AOA 2.

As the captain's AIR DATA switch remained set to the NORM position, the source of AOA 1 data was always ADR 1. Key results for the AOA 1 data were as follows:

- The FDR recorded 42 AOA 1 spikes during the period from 0440:26 until the aircraft landed at Learmonth.
- The first AOA 1 spike occurred at 0440:34. AOA 1 values changed from 2.1° to 50.625° and back to 2.1° over three successive (1-second) samples.
- One of the recorded AOA spikes occurred at 0442:26, immediately prior to the first pitch-down (0442:27). The AOA value was 50.625°.
- Another of the recorded AOA spikes occurred at 0445:08, immediately prior to the second pitch-down (0445:09). The AOA value was 50.625°.
- Most of the AOA spikes that were recorded by the FDR were 50.625° in magnitude and occurred during the initial 12 minutes after the first spike. Two other values (16.875° and 5.625°) were recorded later in the flight.

Table 9 details the number of spikes for each value recorded by the FDR (see also section 3.3.4).[52]

Table 9: AOA 1 spike details (as recorded by the FDR)

Magnitude (°)	Number	Time of first spike	Time of last spike
50.625	29	0440:34	0452:37
16.875	8	0511:57	0530:22
5.625	18	0500:32	0531:15

Figure 23 shows the AOA spikes recorded by the FDR during the period covering the two pitch-downs. Also shown are the recorded stall warnings, which were based on AOA values, overspeed warnings and master warnings.

Only a small number of situations could trigger a master warning, including an involuntary (uncommanded) autopilot disconnection, stall, overspeed, engine fire, or a high cabin altitude. The first master warning on the occurrence flight was due to the involuntary autopilot disconnection while the later master warnings were due to either stall warnings or overspeed warnings. Although there was a general correlation between the times of the master warnings and the stall or overspeed warnings, it was not exact. This result was explained by the sampling rates and the computation time of the FWS, the sampling rates of the acquisition unit for the FDR, and the brevity of the stall and overspeed warnings.

Figure 23: FDR plot showing recorded AOA spikes for both pitch-downs

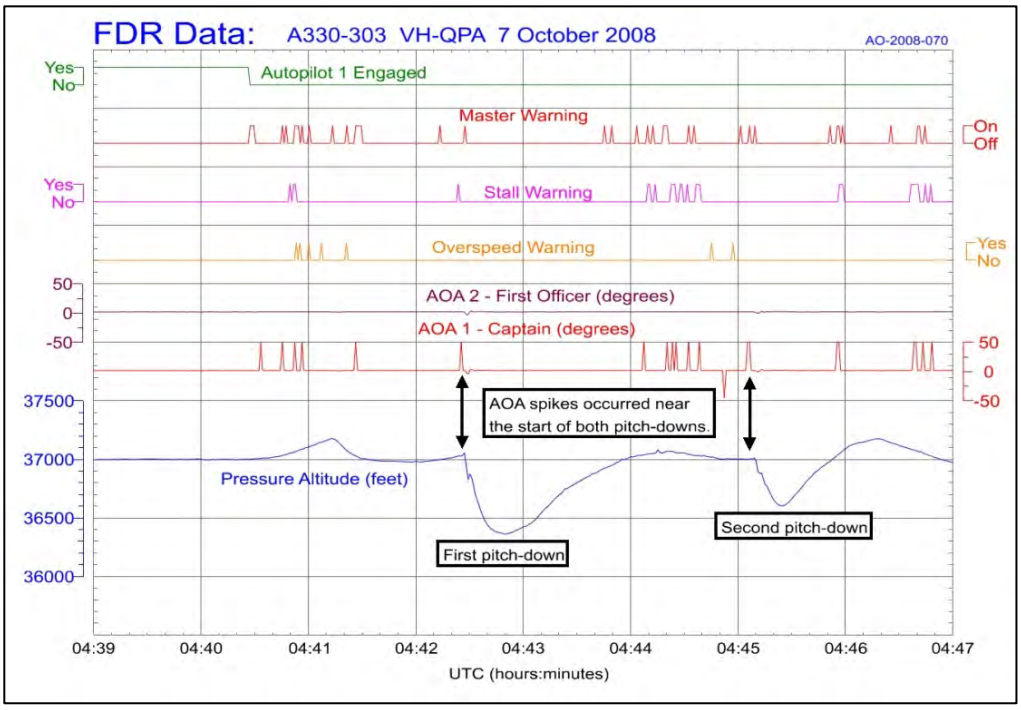

[52] There were also several instances of invalid AOA 1 data being provided by the ADIRU. When this occurred, it was indicated on the FDR with a value of -45° (Appendix B).

QAR data

Overall, the QAR recorded 46 spikes in AOA 1 during the period from 0440:26 to when the aircraft landed. No spikes or problems were evident at any time for the AOA 2 values.

Although not exact, there was a reasonable correlation in the timing between the AOA 1 spikes recorded by the QAR and the FDR. The FDR and QAR systems were not synchronised, so the time at which both recorders sampled the AOA values was not always the same. As a spike may have affected one ADIRU output value but not the next, it was possible that the FDR may have recorded a spike when the QAR did not, and vice versa.

Even when the FDR and QAR recorded AOA 1 spikes at about the same time, they were not always the same value. This difference was due to the two recorders having a different range and resolution for AOA. In short, an AOA value of more than 45° was recorded incorrectly by the QAR as a value of about 45° less than the actual value. For example, a 50.625° spike would be recorded by the QAR as a value of 5.801°.

CVR data

The signal path for the CVR was different to that for the FDR, so the timing of the stall warnings recorded by the two systems would not be exactly the same. Allowing for this difference, the CVR and FDR data were in general agreement. The timing of the stall warnings recorded by the CVR was also generally consistent with the intermittent AOA spikes recorded on the FDR.

1.11.5 Data associated with the pitch-down events

Elevator movements

The range of elevator movement was +15° (nose-down) to -30° (nose-up) and the elevators could move rapidly. Both the left and right elevator positions were recorded on the FDR, and each was sampled twice per second.

During both pitch-downs, both elevators' positions changed abruptly in the nose-down direction (Figure 24 and Figure 25). The magnitude of the change for the first pitch-down was larger, and the rate of change was approximately 15° per second. The duration of the second pitch-down was slightly longer than the first. The abrupt changes for both pitch-downs were followed soon after by a slower recovery period back to normal values.

The movements of the left and right elevators were consistent, and the magnitude of the elevator movements was also consistent with the magnitude of the aircraft's pitch changes.

Figure 24: Acceleration, elevator and sidestick inputs for the first pitch-down

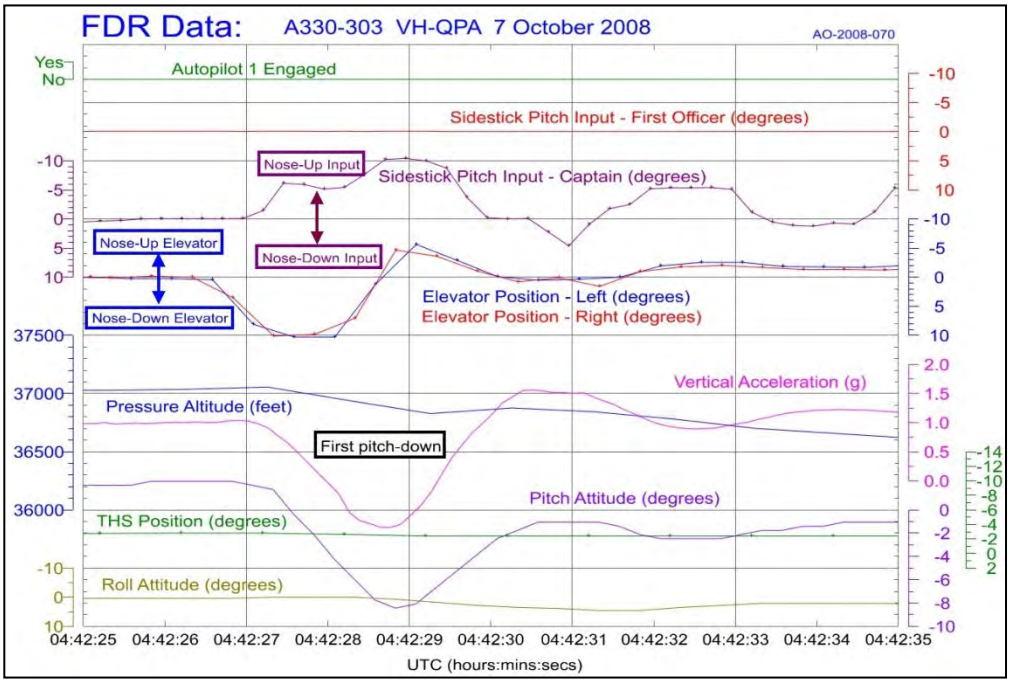

Figure 25: Acceleration, elevator and sidestick inputs for second pitch-down

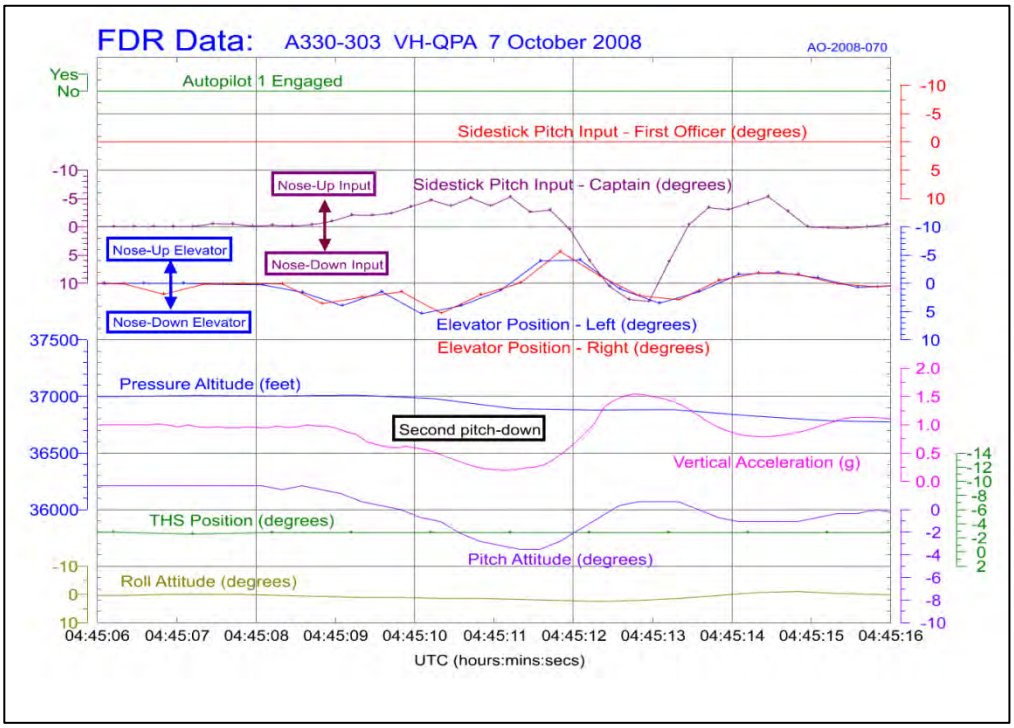

Flight crew pitch inputs

During manual flight, the flight crew make pitch control inputs using their sidesticks. Both the captain's and first officer's sidestick movements were recorded by the FDR.

After the autopilot disengaged (0440:28), the captain's sidestick showed activity that was consistent with small control inputs. The first officer's sidestick was inactive.

There was no sidestick input that correlated with the initiation of either pitch-down. In both cases, very soon after the pitch-down commenced, the captain applied back pressure to the sidestick to produce a nose-up input. Each was moderate and commensurate with the response required to correct for the pitch-downs. The inputs did not have any immediate effect on the recorded elevator position, but within 1 to 2 seconds the elevator position started to correlate with the sidestick inputs. Figure 24 and Figure 25 show the sidestick pitch input parameters for the 10-second period around each pitch-down.

Trimmable horizontal stabiliser movements

The range of the THS position was +2° (nose-down) to -14° (nose-up), and movement was limited to approximately 1° per second. The FDR recorded the THS position once per second, and the data showed that there was no movement of the THS immediately prior to either of the two pitch-downs.

Trim tank changes

A fuel tank was located in the THS to allow for a reduction in the amount of nose-up trim that the THS needed to provide during flight. This in turn reduced the drag produced by the THS and consequently the aircraft's fuel consumption. The trim tank capacity was 6,230 L, which corresponded to a weight of 4,891 kg.[53] Fuel was transferred automatically to and from the trim tank to maintain the aircraft's c.g. within a target range.

The FDR recorded the trim tank fuel quantity, and the data showed that 500 kg of fuel was present at takeoff, reducing to a constant value of 440 kg during cruise. There was no change in the trim tank fuel quantity before either of the pitch-downs.

Turbulence

Recorded flight data can help determine whether an aircraft has experienced turbulence, with the vertical acceleration parameter being the key indication. In flight, vertical acceleration values represent the combined effects of flight manoeuvring loads and turbulence. Turbulence is indicated by a higher frequency scatter in the trace while flight manoeuvring loads are more prolonged and correlate with pitch and roll attitude changes.

Figure 24 shows the vertical acceleration during the first pitch-down, and Figure 25 shows the vertical acceleration during the second pitch-down. There was no evidence of any turbulence in the recorded data during these periods.

---

[53] For a fuel specific gravity of 0.785 kg/L.

Electrical flight control system commands

As discussed above, the FDR showed coincident movement of both elevators, and that the response of the elevators was consistent with flight crew inputs, prior to and after (but not during) each pitch-down. In addition, the FDR showed that the autopilot was not connected during the two pitch-downs, and was therefore not providing inputs to the EFCS at that time.

EFCS pitch commands were not recorded on the FDR. However, the FDR showed that there were recorded spikes in AOA 1 data at the time of the two pitch-downs. As discussed in section 1.6.5, the EFCS could initiate a pitch-down command if it detected that the aircraft's AOA was too high. The EFCS algorithm for determining the AOA value to use when computing flight control commands is based on the three ADIRUs' inputs (section 2.1.5).

1.12 Aircraft and component examinations

1.12.1 General aircraft inspection

Visual inspection of the aircraft at Learmonth found significant damage to overhead fittings in the passenger cabin (section 4.3). No other aircraft damage was identified. There were no missing or loose fasteners, no creases or folds in the fuselage skin, and no signs of damage or distress to any of the fuselage, wing or empennage skin, fairing panels or flight controls.

The FDR data showed that the peak vertical accelerations during the flight were -0.80 g and +1.56 g, with almost no lateral acceleration. The aircraft maintenance manual defined the normal flight operating range as -1.0 g to +2.5 g. Aircraft operation within this range did not require additional inspections. Based on a review of the FDR data, the aircraft manufacturer asked for a visual inspection of the elevator servo-control attachment fittings. This inspection found no problems.

After removal of the cargo, the aircraft hold's structure and restraint systems were inspected for damage which might be attributed to the event. No problems were found.

1.12.2 Post-flight report

Types of messages

The post-flight report (PFR) was produced by the CMS at the end of each flight. The PFR contained fault information received by the CMS from other aircraft systems' BITE and the FWS.

PFR messages were of two main types:

- Cockpit effect messages. These level 3, 2 and 1 'failure level' messages were generated by the FWS and presented to the flight crew on the ECAM and/or other displays (section 1.6.9). The messages enabled the flight crew to know which operational functions were no longer available.
- Maintenance fault messages. These class 1 and class 2 'failure' messages were generated by the CMS, based on inputs from other systems (section 1.6.10).

They provided information for maintenance personnel on the status or functioning of the aircraft's systems. A class 1 message was usually associated with one or more cockpit effect messages that were presented to the flight crew during the flight. A class 2 message was not associated with any messages provided to the flight crew.[54]

The PFR had a number of limitations, including:

- it only recorded information to the nearest minute
- it only showed the first occurrence of a cockpit effect or a maintenance fault message (that is, a repeat occurrence of the same message would not be shown)
- a correlation function performed by the CMS grouped together all of the maintenance fault messages associated with the same system[55] at the same time (within 1 minute), and it would only record the first fault that was detected (along with a list of other systems that detected the fault).

The PFR from the flight recorded 22 cockpit effect messages that were presented to the crew during the period from 0440 until the end of the flight.[56] These included 21 different ECAM messages and a 'flag' message.[57] All were caution messages except the AUTO FLT AP OFF warning message.

Appendix C provides the complete PFR from the flight.

ADIRU fault messages

The two types of caution messages that could be displayed to the flight crew about an ADIRU problem were a NAV IR [1, 2 or 3] FAULT and a NAV ADR [1, 2 or 3] FAULT. These messages could be generated by the ADIRU self-detecting a fault and transmitting that information to the FWS, by the FWS detecting that the IR or ADR parts had stopped transmitting data (that is, the data was not 'refreshed'), or by another ADIRU detecting a problem (for the ADR only) and transmitting that information to the FWS.

The PFR for the flight showed two caution messages associated with the operation of the ADIRUs:

- NAV IR1 FAULT (at 0440). Given that ADIRU 1 consistently flagged the IR parameters as invalid, it is likely that it self-detected the problem.
- NAV ADR 1 FAULT (at 0513). Subsequent analysis determined that this message was generated by the IR part of ADIRU 3 (section 1.12.6) and not by ADIRU 1 itself.

---

[54] A class 2 maintenance message was associated with one or more 'maintenance status' messages, which were placed in the 'cockpit effects' column of the PFR even though they were not presented to the flight crew. Class 3 messages were not recorded on the PFR but could be obtained from the CMS or the relevant system's BITE.

[55] Fault messages were grouped by the Air Transport Association (ATA) chapter reference, with each chapter referring to a different system. This categorisation is widely used in aircraft documentation.

[56] The 'cockpit effects' column of the PFR also included six 'maintenance status' messages. These were class 2 maintenance fault messages and were not presented to the crew during flight.

[57] A 'flag' was displayed on the flight crew's flight displays. When a flag was displayed there was no associated ECAM message.

In addition to these two caution messages, the PFR also recorded a series of maintenance faults associated with ADIRU 1 as follows:

- IR 1 class 2 fault (at 0440): this fault was detected and reported by ADIRU 1.
- ADR 1 class 2 fault (at 0440): this fault was detected and reported by ADIRU 1 and probably associated with some ADR data not being refreshed at a sufficient rate (section 3.7.3).
- ADIRU 1 class 1 fault (at 0440): this message was reported by the FMGES, EFCS, FWS, DMC 1 and ground proximity warning system (GPWS).

Subsequent maintenance fault messages associated with ADIRU 1 were detected by the EFCS at 0452, 0455 and 0500, the IR part of ADIRU 3 at 0506 and 0513, and the landing gear control interface units (LGCIUs) at 0531.

The PFR showed no recorded faults for ADIRU 2 or ADIRU 3.

EFCS fault messages

The PFR recorded four caution messages that were associated with the operation of the EFCS, as shown in Table 10. Further explanation of these messages is provided in section 2.2.2.

Table 10: PFR cockpit effect messages for EFCS

Time	Cockpit effect message	Interpretation
0442	F/CTL PRIM 1 PITCH FAULT	The independent command and monitor channels of FCPC 1 detected a discrepancy between the actual and commanded elevator position.
0442	F/CTL PRIM 3 FAULT	A discrepancy was detected between the control orders that were independently computed by the command and monitor channels of FCPC 3. [Note. FDR data showed that a second F/CTL PRIM 3 FAULT occurred at 0445.]
0445	F/CTL PRIM 2 PITCH FAULT	The independent command and monitor channels of FCPC 2 detected a discrepancy between the actual and commanded elevator position.
0445	F/CTL ALTN LAW	Due to multiple FCPC faults, the EFCS control law reverted from normal to alternate law.

The PFR also recorded class 1 maintenance fault messages associated with the two pitch faults at 0442 and 0445.

Other cockpit effect messages

Several of the cockpit effect messages on the PFR were attributable to ADIRU 1 providing incorrect or insufficient output data (see Table 11).

Table 11: Cockpit effect messages due to problems with ADIRU output data

Time	Cockpit effect message	Interpretation
0442[58]	AUTO FLT AP OFF	The autopilot disconnected and the disconnection was not commanded by the crew. This was due to a discrepancy between data that was sourced from the different ADIRUs (section 1.12.8).
0440	NAV GPWS FAULT	The GPWS required ADIRU parameters such as true track angle, computed airspeed, true airspeed and roll attitude, and it was only connected to ADIRU 1. A problem with ADIRU 1 resulted in a loss of the GPWS function.
0440	FLAG ON CAPT ND MAP NOT AVAIL	Due to a loss of heading information, the map that was usually shown on the captain's navigation display (ND) was unable to be displayed. Consequently a red warning flag was displayed on the ND but there was no ECAM message.
0440	NAV FM/GPS POS DISAGREE	There was a latitude and longitude cross-check error between the GPS and FMGES data (based on IR information).
0441	NAV GPWS TERR DET FAULT	There was a loss of GPWS functionality due to a loss of data from ADIRU 1 (see NAV GPWS FAULT above).
0442	BRAKES AUTO BRK FAULT	The autobrake required ADIRU parameters to perform its function (that is, providing a set deceleration level during landing). A loss of ADIRU 1 data resulted in a loss of the autobrake function.
0451	EIS DISPLAY DISCREPANCY	One of the DMCs detected a problem with IR data (pitch, roll or heading) or ADR data (altitude) which therefore affected the electronic instrument system (EIS).
0528	CAB PR LO DIFF PR	A low differential pressure between the cabin and external-to-aircraft conditions was detected by the cabin pressure controllers. The automatic cabin pressure control was lost as ADR data (altitude and vertical speed) was required from ADIRU 1.

The remaining cockpit effect messages were considered to be spurious. That is, ADIRU 1 provided incorrect information on its fault parameters regarding several, related systems (see also section 3.7.4). These messages are listed in Table 12.

[58] Although the time shown on the PFR was 0442, the FMGES BITE and FDR data showed that this message would have first occurred at 0440.

Table 12: Cockpit effect messages due to spurious fault messages from ADIRU 1

Time	Cockpit effect message	Interpretation
0440	NAV GPS 1 FAULT	ADIRU 1 detected the GPS function to be faulty. No faults were recorded in the BITE data for the multi-mode receiver that included the GPS unit.
0440	NAV GPS 2 FAULT	As above for NAV GPS 1 FAULT.
0440	A.ICE L CAPT STAT HEAT	This message was due to corrupted data output from ADR 1 (section 1.12.9).
0442	NAV IR NOT ALIGNED	This message indicated that IR 1, IR 2 or IR 3 was 'not aligned'. The aircraft manufacturer advised that it was a spurious message resulting from the anomalous behaviour of ADIRU 1.
0445	A.ICE R CAPT STAT HEAT	As above for A.ICE L CAPT STAT HEAT.
0448	A.ICE CAPT PROBES HEAT	As above for A.ICE L CAPT STAT HEAT.
0456	A.ICE CAPT PITOT HEAT	As above for A.ICE L CAPT STAT HEAT.
0508	A.ICE CAPT AOA HEAT	As above for A.ICE L CAPT STAT HEAT.

1.12.3 Troubleshooting data from the aircraft systems

The PFR only showed a summary of the warning/caution and maintenance fault messages. The most detailed level of maintenance fault data that could be obtained from the CMS was troubleshooting data (TSD).[59] The TSD showed each maintenance message that was received by the CMS in a raw numerical format that could be printed out and decoded using relevant documentation.

Based on an initial examination of the FDR data, and to minimise the possibility that data might be lost when power was reapplied to the aircraft, the aircraft manufacturer recommended removing ADIRU 1 and probe heat computer (PHC) 1 before conducting any data downloads or testing of the aircraft's systems. These units were replaced with spare units before data downloads or functional testing commenced.

The TSD for several aircraft systems was downloaded, and subsequent analysis found it to be generally consistent with the PFR data. The key result was that several aircraft systems detected problems with the data being provided by ADIRU 1, but no problems were detected with the data being provided by ADIRU 2 or ADIRU 3 (Table 13). The TSD was generally not accurate enough to determine which ADIRU parameters were involved or the exact nature of the problem that triggered the fault messages.

[59] To obtain all possible BITE data, some LRUs were removed from the aircraft and downloaded by the unit manufacturer.

Table 13: Troubleshooting data relating to the ADIRUs

Units	ADIRU connections	Results
FCPC 1, FCPC 2, FCPC 3	All 3 ADIRUs	At 0440, problem detected with ADIRU 1 data (IR and ADR). Subsequent problems also detected. No problems with ADIRU 2 or ADIRU 3 data.
FMGEC 1, FMGEC 2	All 3 ADIRUs	At 0440, problem detected with ADIRU 1 data. No problems with ADIRU 2 or ADIRU 3 data.
ADIRU 2, ADIRU 3	All 3 ADIRUs	At 0440, problem detected with altitude and/or true airspeed data not being received from ADIRU 1 at the designed rate (section 1.12.6). No problems with ADIRU 2 or ADIRU 3 data.
GPWS	ADIRU 1	At 0440, problem detected with ADIRU 1 data.
LGCIU 1, LGCIU 2	ADR1 and ADR 3	At 0531, problem detected with loss of computed airspeed from ADR 1. No problems with ADR 3 data.

1.12.4 Functional testing of aircraft systems

After the PFR and TSD data were downloaded, functional tests were performed on several aircraft systems at Learmonth in accordance with the aircraft manufacturer's recommendations. The systems included the EFCS, ADIRS, FMGES, PHCs, multi-mode receivers (which included the GPS receivers), and the electrical power generation system.

All of the functional tests were successfully completed except for a single EFCS task that involved the reconfiguration of the elevator servo-controller to another computer. The aircraft manufacturer advised that this was a previously identified anomaly that was only triggered under a very specific set of circumstances and was not related to the occurrence.[60]

1.12.5 Aircraft wiring examinations

The data-spike failure mode affected data that was transmitted on multiple, segregated wires out of the ADIRU (and based on information obtained from multiple, segregated wires into the ADIRU). A simultaneous problem with multiple, segregated wires was considered very unlikely.

Nevertheless, a range of testing was conducted on the aircraft wiring, and no problems were found. More specifically:

- Due to the extent of damage to the ceiling panels in the cabin, all the panels were removed and the wiring looms were visually inspected while the aircraft was at Learmonth. No defects were observed.
- After the aircraft was ferried to a maintenance base in Sydney, the operator conducted precautionary checks of the aircraft's ADIRU interface wiring that

[60] The anomaly only occurred during the performance of functional tests. It appeared with the introduction of a particular FCSC software version.

involved continuity, short circuit, electrical bonding and shielding tests. No problems were found.

- In November 2008, the operator conducted integrity and time domain reflectometry tests on the aircraft's ADIRU 1 databus wiring, and verification of databus signals using a bus analyser. No problems were found.
- In May 2009, the aircraft manufacturer and the operator conducted additional checks of the ADIRU 1 installation and configuration. The checks included ventilation, electrical bonding to the aircraft structure, input and output databus wiring isolation, input and output databus waveforms, discrete input signals (power and mode selection), AOA input signal waveforms, and alternating current (AC) and direct current (DC) power supply waveforms. No faults were found that were relevant to understanding the occurrence or otherwise considered significant.
- In May/June 2009, the aircraft underwent a scheduled C-check. No aircraft wiring or configuration problems were identified during that check.

1.12.6 ADIRU examinations

Service history

The aircraft was fitted with three LTN-101 ADIRUs, manufactured by Northrop Grumman Corporation. The basic details of the ADIRUs are provided in Table 14.

Table 14: Details of the aircraft's ADIRUs

	ADIRU 1	ADIRU 2	ADIRU 3
Model number	LTN-101	LTN-101	LTN-101
Part number	465020-0303	465020-0303	465020-0303
Serial number	4167	4687	4663
Date of manufacture	Aug 2002	Jan 2004	Dec 2003
Date installed in QPA	Feb 2004	Apr 2006	Jul 2008
Date installed in position	Mar 2006	Apr 2006	Jul 2008
Software version	0316	0316	0316
Total operating hours	30,282	26,985	25,423

All three units had been removed from aircraft during their operational life due to reported faults, which was not abnormal for units with similar operating hours (see section 3.9 for more details on ADIRU reliability). More specifically:

- unit 4167 was examined and repaired in May 2003
- unit 4687 was examined and repaired in February 2006
- unit 4663 was removed for examination on two previous occasions (October 2007 and May 2008) but, on both occasions, no fault was found.

None of those reported problems were related to the 7 October 2008 occurrence.

A review of all the operator's technical log entries related to ADIRUs identified one previous event involving similar ADIRU behaviour as that on 7 October 2008. That event occurred on 12 September 2006 and involved the same ADIRU (unit 4167) in

position 1 on QPA. Further details of that event are discussed in section 1.16.2 and Appendix D, and further information on the service history of unit 4167 is discussed in section 3.5.4.

BITE data from ADIRU 1

The aircraft's three ADIRUs were removed to download the units' BITE data and conduct examination and testing.[61]

The BITE data from ADIRU 1 showed no fault messages from the occurrence flight. Given the fault messages recorded by other systems related to ADIRU 1, some fault messages should have been recorded. In addition, several routine messages normally stored in BITE memory were either not recorded or had anomalies. These included:

- An alignment record should have been recorded after the ADIRU was turned on in Singapore. It was not recorded.
- A routine NAV update record should have been recorded when the unit was shut down at Learmonth. It was not recorded.
- Routine elapsed time interval (ETI)[62] timestamps should have been recorded during the flight. The ADIRUs were on for 14.8 hours before being shut down at 1525. However, the ETI observed at turn on at the manufacturer's test facility was about 0.7 hours after takeoff.
- Routine temperature records should have been recorded every hour. None were recorded after the start of the event (0440).

Subsequent analysis indicated that the absence of recorded fault messages was associated with a problem in storing of BITE data rather than a problem with the execution of the BITE tests themselves (section 3.7).

BITE data from ADIRUs 2 and 3

The BITE data from ADIRUs 2 and 3 showed that all the routine BITE messages were correctly recorded. The data did not show any fault messages related to ADIRUs 2 and 3, but did show fault messages related to the way ADIRU 1 transmitted data to other aircraft systems.

Although the three ADIRUs were essentially independent units, they exchanged some ADR data and each unit monitored the others' transmission of that data. More specifically:

- The IR part of an ADIRU required certain ADR parameters for its computations (for example, true airspeed data was required in conjunction with groundspeed to determine the wind speed and wind direction).

---

[61] ADIRU 1 (unit 4167) was removed from the aircraft at Learmonth prior to any data downloads or functional testing of the other units on the aircraft. The other two units were removed after the aircraft was ferried back to Sydney. All three units were sent to the ADIRU manufacturer's facilities in Los Angeles in the US for data download and testing under the supervision of the ATSB and other investigation agencies.

[62] The ETI was the total operating time of the ADIRU, from turn on to turn off.

- Normally the IR part would use the ADR part from the same ADIRU to obtain this data. If this ADR part was not available, then the IR part could source these parameters from another ADR.
- In order to obtain ADR information from another ADIRU, each unit had a digital air data system (DADS) input from the two other ADIRUs. The DADS inputs supplied true airspeed and altitude data.
- If there was a problem with the ADR data received by an IR part, then the affected ADIRU would record a fault in its BITE and generate a fault message.
- Each ADIRU also performed input range monitoring of the parameters that were outputted over the DADS databuses. The acceptable ranges were 0 to 599 kts for true airspeed and -2,000 to 50,000 ft for altitude.

For the 7 October 2008 flight, ADIRU 2 and ADIRU 3 both recorded fault messages that indicated problems with the true airspeed or altitude data outputted from ADIRU 1 to the other ADIRUs. These messages are summarised in Table 15.

Table 15: BITE summary from ADIRUs 2 and 3 for 7 October 2008 occurrence

Time[63]	ADIRU	Fault description	Comment
0440:30	ADIRU 2	Input databus refresh rate failed (class 3)	Altitude and/or true airspeed were not being received from ADIRU 1 at the designed rate.
	ADIRU 3	Input databus refresh rate failed (class 3)	
0506:36	ADIRU 3	Input databus refresh rate failed (class 1)	Altitude and/or true airspeed were not being received from ADIRU 1 at the designed rate.
0512:00	ADIRU 3	Failure detection and exclusion still GPSSU 1 (class 3)	The GPS satellite failure detection and exclusion function failed; the monitor determined that there were several data inconsistencies but was unable to isolate them to a specific source.
0512:54	ADIRU 2	Input data SSM failed (class 3)	Altitude and/or true airspeed were received from ADIRU 1 with the SSM field set to 'failure warning' or 'no computed data'.[64]
	ADIRU 3	Input data SSM failed (class 1)	
0523:24	ADIRU 2	Input range failed (class 3)	Altitude and/or true airspeed received from ADIRU 1 were out of range.

Normally these types of fault messages would be class 3 maintenance faults. However, at 0443:45 the flight crew switched the ATT HDG switch to the CAPT ON 3 position. IR 3 then provided the data for the captain's flight displays. As the captain's AIR DATA switch remained on the NORM position, ADR 1 was still the primary source of air data for the captain's systems, and therefore became the primary source of air data for IR 3. Consequently, any problem IR 3 detected with ADR 1 after 0443:45 became more significant and resulted in a class 1 failure. The

[63] The ADIRU BITE recorded fault messages to a resolution of 6 seconds.

[64] A review of the QAR data showed that at 0513 there were spikes in standard altitude, Mach and wind speed (derived from true airspeed). True airspeed was not recorded by the FDR or QAR.

class 1 fault messages shown in Table 15 correlated with the maintenance fault messages on the PFR at 0506 and 0513 (section 1.12.2).

The class 1 message at 0513 resulted in the 'NAV ADR 1 FAULT' caution message that was recorded on the PFR. In other words, this caution message was generated by ADIRU 3 and not ADIRU 1. As the local ADR 1 fault light located on the overhead panel could only be activated by ADIRU 1, it would not have illuminated on the occurrence flight.

Unit testing

The testing of ADIRU 1 (unit 4167) included visual inspections, functional testing of the unit and its modules, functional checking of the software, and a range of environmental tests (including temperature, vibration and electromagnetic interference). Although the unit had transmitted a significant amount of incorrect data to other systems, and was associated with several fault messages, extensive testing did not identify any problems relevant to the occurrence.

More extreme testing was also conducted on an exemplar unit.[65] No problems relevant to the investigation were identified. Routine acceptance testing was conducted on ADIRU 2 and ADIRU 3, and no problems were found.

Further details on the ADIRU testing are provided in Appendix E.

1.12.7 FCPC examinations

Service history

The aircraft's three FCPCs were manufactured by Airbus. Basic details for these units are provided in Table 16.

Table 16: Details of FCPCs installed on QPA

	FCPC 1	FCPC 2	FCPC 3
Part number	LA2K2B100D80000	LA2K2B100D80000	LA2K2B100D80000
Serial number	2K2007270	2K2006165	2K2006170
Date of manufacture	Nov 2007	Nov 2003	Nov 2003
Date installed in QPA	Jun 2008	Nov 2003	Nov 2003
Date installed in position	Jun 2008	Jun 2008	Nov 2003
Software version	P7/M16	P7/M16	P7/M16
Total operating hours	2,349	33,007	35,150

All three of the FCPCs were installed new in the aircraft, and none had been removed for repair. Unit 2K2006165 (FCPC 2) was initially installed on the aircraft as FCPC 1. However, the previous FCPC 2 developed a fault in June 2008 and it was removed for examination. Unit 2K2006165 was swapped to the FCPC 2

[65] A unit that was functionally identical to ADIRU 4167 and had the same hardware and software modification status. To minimise the chance of losing perishable data, or that the test equipment might damage ADIRU 4167, testing was performed on the exemplar unit before being performed on ADIRU 4167.

position, and the faulty unit was swapped to the FCPC 1 position and then replaced by 2K2007270. At the time of the replacement (13 June 2008), maintenance personnel conducted a series of tests and inspections to confirm that the FCPCs were operating normally.

BITE download and unit testing

Following the occurrence, the three FCPCs were removed from the aircraft and examined by an authorised agency. The key results of this examination were:

- Each FCPC was loaded with identical, uncorrupted operational software.
- The BITE data was downloaded from each FCPC, and no faults relating to the occurrence flight were found.[66] Both FCPC 2 and FCPC 3 contained faults from earlier flights that were unrelated to the pitch-downs.
- Each of the computers was subsequently tested, and no fault was found with FCPC 1 or FCPC 2. FCPC 3 failed a lightning protection test on one input. The aircraft manufacturer advised that the relevant input was not used when an FCPC was installed as FCPC 3, and therefore the problem was unrelated to the pitch-downs.

A review of the FCPC operational logic found that the EFCS faults that were recorded on the PFR were due to self-monitoring discrepancies detected by the FCPCs during the pitch-downs, and that they were not associated with any physical faults of the computers (section 2.2.2). However, a problem was identified with how the FCPC software was designed to manage incorrect AOA data (section 2.1).

1.12.8 FMGEC performance review

The FMGECs were manufactured by Thales Honeywell (part number C12858BA02). Service history details for the two units were as follows:

- FMGEC 1 (serial number Q00173002571) was manufactured in February 2003. It had undergone a repair in April 2007 before being fitted to the aircraft.
- FMGEC 2 (serial number Q00173003903) was manufactured in July 2007 and fitted new to the aircraft in April 2008.

A review of the aircraft's technical log entries related to the FMGES system for the 12 months prior to 7 October 2008 identified no faults or problems related to the occurrence.

Each FMGEC had two channels as follows:

- Command (COM) channel. The COM channel issued autopilot control commands to the EFCS, and it based its computations on all three ADIRUs (using the median or middle value of each relevant parameter).
- Monitor (MON) channel. The MON channel also computed autopilot control commands, but based its computations on one ADIRU. FMGEC 1's MON channel used ADIRU 1, and FMGEC 2's MON channel used ADIRU 2.[67]

---

[66] The FCPCs were interfaced to the CMS and FWS through two flight control data concentrators (FCDCs). Fault messages such as PRIM 1 PITCH FAULT, PRIM 3 FAULT and PRIM 2 PITCH FAULT were stored in the FCDC BITE rather than the in FCPC BITE.

If there was a discrepancy between the COM and MON channels' computations, and the discrepancy was confirmed after 100 msec, then the autopilot was disconnected and a fault message was recorded.

On the 7 October 2008 flight, the BITE of FMGEC 1 recorded a fault message at 0440 that stated there was a discrepancy between the COM and MON channels. This message indicated that the autopilot disconnection was due to differences between the value of a parameter from ADIRU 1 and the values of the same parameter from the other ADIRUs. It was not possible to determine which ADIRU parameter was involved.[68]

Given that the autopilot 1 disconnection was explained by the autopilot logic, and that neither autopilot was engaged at the time of the two pitch-downs, there was no need for the FMGECs to be tested.

1.12.9 Probe heat computer examination

As discussed in section 1.6.4, sensors on the outside of the aircraft provided the source information for the ADIRUs' ADR output parameters. To prevent them from being affected by ice, each sensor had electrical heating (anti-icing) that was controlled by a probe heat computer (PHC). The aircraft had three PHCs; PHC 1 provided heating for ADIRU 1's sensors, PHC 2 provided heating for ADIRU 2's sensors, and PHC 3 provided heating for ADIRU 3's sensors.

If a heating fault occurred with any of the aircraft's sensors, then the relevant PHC would send a fault message to the associated ADIRU, which in turn would send a fault message to the FWS. The FWS would then send a message to the ECAM to alert the flight crew.

PHC 1 could send five anti-ice fault indications to the FWS via ADIRU 1, as detailed in Table 17. All but one of these messages were recorded on the PFR for the 7 October 2008 flight.

The PHCs were manufactured by Intertechnique (part number 785620-3). PHC 1 (serial number 785620IN2083) was manufactured in March 2004. Following a repair, it was fitted to the aircraft in September 2007.

There was no maintenance fault message on the PFR indicating a problem with the PHCs, and no fault message was recorded in the PHCs' BITE data. PHC 1 was removed from the aircraft in Learmonth and subsequently tested by an authorised agency. No fault was found. A review of the aircraft's technical log entries related to anti-icing systems for the 12 months prior to the occurrence identified no faults or problems related to the occurrence.

---

[67] The COM channel compared the values of relevant parameters from each ADIRU against the median value. If there was a discrepancy with ADIRU 1 (or ADIRU 2) for more than 450 msec, then the FMGECs would reject the IR or ADR part of that ADIRU for the remainder of the flight. The MON channel would switch to ADIRU 3 if the associated COM channel detected a problem with the normal ADIRU used by the MON channel.

[68] This was a different fault message to the one recorded by both FMGEC 1 and FMGEC 2 that reported a problem with the data being provided by ADIRU 1 (section 1.12.3).

Table 17: PHC 1 anti-ice fault indications

Sensor	Cockpit effect message	Recorded on PFR
Left static port	A.ICE L CAPT STAT HEAT	Yes, at 0440
Right static port	A.ICE R CAPT STAT HEAT	Yes, at 0445
All (5) sensors simultaneously	A.ICE CAPT PROBES HEAT[69]	Yes, at 0448
Pitot probe	A.ICE CAPT PITOT HEAT	Yes, at 0456
AOA 1 sensor	A.ICE CAPT AOA HEAT	Yes, at 0508
TAT sensor	A.ICE CAPT TAT HEAT	No

The potential reasons for the anti-ice fault messages included either a PHC 1 fault, an ADIRU 1 fault, or multiple faults with independent sensors. Based on a review of the available information, the 'A.ICE' cockpit effect messages on the PFR were considered to be a result of erroneous ADIRU 1 outputs (see also section 3.3.4).

1.12.10 Angle of attack sensor examination

The aircraft's AOA sensors were manufactured by Goodrich (part number 0861ED). All three of the aircraft's AOA sensors were installed new on the aircraft in 2003 and had never been replaced. A review of the aircraft's relevant technical log entries related to the ADIRS identified no problems related to AOA sensors throughout the aircraft's operating history.

The FDR data showed problems with AOA 1 (as well as other data from ADIRU 1). Accordingly, the AOA 1 sensor (serial number 0861ED-972) was removed from the aircraft following the occurrence and tested by an authorised agency. No fault was found with the sensor and all test parameters were within limits. The wiring between the AOA 1 sensor and the ADIRU was also tested, and no problems were identified.

Given that the FDR and QAR data only showed problems with AOA 1 data, there was no need for the other two AOA sensors to be examined.

1.13 Medical and pathological information

As discussed in section 1.5, there was no evidence that medical or physiological factors affected the flight crew's performance. Information on passenger and crew injuries is provided in section 4.6.

1.14 Fire

There was no evidence of fire on the aircraft.

---

[69] This message was not sent from the ADIRU but was generated by the FWS when all five probe heat faults were active simultaneously. The FWS prioritised fault messages so that an A. ICE CAPT PROBES HEAT message would suppress individual probe heat messages. As a result, it was possible for an individual probe heat message to appear on the PFR at a later time than the A. ICE CAPT PROBES HEAT message or not at all.

1.15 Survival aspects

Information on cabin safety matters is provided in Part 4. Of note in this occurrence was that over 60 passengers were seated but not wearing seat belts, and that these occupants had a significantly higher rate of injury than those who were wearing their seat belts.

1.16 Tests and research

A wide range of tests, examinations and simulations were conducted as part of the investigation, as reported in Part 2 (FCPC-related) and Part 3 (ADIRU-related). The present section focuses on a review of potentially-related occurrences.

1.16.1 Previous flight control occurrences associated with ADIRU failures

Most of the systems on modern aircraft are highly reliable. Although equipment faults do occur, they rarely have a significant effect on the safety of a flight due to system design features such as fault detection and the use of multiple units for redundancy.

All models of ADIRUs develop occasional faults, and section 3.9 provides an overview of faults and reliability associated with the LTN-101 ADIRU. However, it is extremely rare for any ADIRU failures to have an undesirable effect on an aircraft's flight controls. Airbus advised that it is unaware of any previous occurrences where an ADIRU failure on one of its aircraft has resulted in undesirable elevator commands.

There has been one previous case where an ADIRU failure led to an in-flight upset of a civilian aircraft. That occurrence involved an ADIRU failure on a Boeing 777-200 aircraft, which occurred on 1 August 2005, 240 km north-west of Perth, Western Australia.[70] The ADIRU model used on that aircraft was made by a different manufacturer and was of a different system design to the model used on QPA; rather than three separate ADIRUs, the 777 had one ADIRU with redundant components. The aircraft experienced an uncommanded pitch-up, problems with indicated airspeed, and activation of the stall warning and stickshaker devices. The occurrence involved hardware failures to two accelerometers within the ADIRU, and inputs from one of the faulty accelerometers being treated as valid data due to a software design problem within the ADIRU. The occurrence was unrelated to the occurrence involving QPA on 7 October 2008.

1.16.2 Other ADIRU data-spike occurrences

Search for other data-spike occurrences

The ATSB and the operator conducted a detailed review of the operator's maintenance records for its A330 fleet for events with similar ADIRU behaviour as occurred on the 7 October 2008 flight. Only one event was identified (12 September

[70] See ATSB investigation report AAIR200503722 available at www.atsb.gov.au.

2006, see below). During the investigation, a third event was reported on 27 December 2008 (see below).

The aircraft manufacturer conducted searches of the PFRs from a significant proportion of Airbus aircraft flights to identify any PFRs with a similar set of messages as were recorded during the three known occurrences. No other events were identified. Further details of these searches are provided in Appendix D.

The aircraft and the ADIRU manufacturers advised that they were not aware of any other occurrence involving similar anomalous ADIRU behaviour. They also advised that, if such a problem had occurred and no fault was found in a subsequent ground test of the unit, then the event would probably not be reported to them. Communications by the operator and ADIRU manufacturer with other A330 operators since 7 October 2008 also did not identify any similar occurrences. Another ADIRU manufacturer also advised that it was not aware of any similar types of events.

12 September 2006 data-spike occurrence

The 12 September 2006 occurrence involved the same aircraft (QPA) and the same ADIRU (4167) as the 7 October 2008 occurrence. No recorded data was available for that flight other than the PFR. However, the PFR and the flight crew's description of numerous warning and caution messages provided sufficient evidence to conclude that the occurrence involved similar ADIRU data output anomalies as that which occurred on the 7 October 2008 flight. Following the 2006 flight, line testing of the ADIRU was conducted in accordance with the manufacturer's procedures; no fault was found and the aircraft was returned to service.

Further details of the 2006 occurrence are provided in Appendix D.

27 December 2008 data-spike occurrence

The 27 December 2008 occurrence involved another of the operator's A330-303 aircraft, registered VH-QPG (QPG), and a different LTN-101 ADIRU (serial number 4122). In that occurrence, the crew reported receiving numerous caution messages and that the messages on the ECAM were constantly changing. The crew followed a new procedure that was introduced after the 7 October 2008 occurrence, which was successful in shutting down the ADR part 28 seconds after the failure mode started (and 24 seconds after the autopilot disconnected). The new procedure was not effective in shutting down the IR part.[71]

The FDR and QAR data showed evidence of data spikes on all IR parameters and some of the ADR parameters (in the 28 seconds during which it operated in the failure mode). General observations about the spike timing and magnitudes were similar to those for the 7 October 2008 event. In addition to data spikes, the IR data for the 27 December 2008 event also showed similar patterns as the IR data on the 7 October 2008 flight.

ADIRU 4122 was removed from the aircraft and sent to the manufacturer's facilities for examination and testing and to download the BITE data. The testing

---

[71] Following the 27 December 2008 occurrence, the procedure was modified further to eliminate this problem (section 7.1.1).

did not identify any faults. The BITE data was very similar to that recovered from unit 4167 after the 7 October 2008 occurrence. That is, there were no faults recorded during the occurrence flight, and several routine messages normally stored in BITE were not recorded.

Further details of the 27 December 2008 occurrence are provided in Appendix D.

Comparison between the three occurrences

In both the 12 September 2006 and the 27 December 2008 occurrences there was no effect on the aircraft's flight controls, and consequently there were no FCPC (PRIM) faults or PITCH faults.

The three occurrences all happened off the coast of Western Australia (Figure 26), although there were significant distances between the three events. Summary details of the three occurrences are provided in Table 18.

Figure 26: Locations of the three ADIRU data-spike occurrences[72]

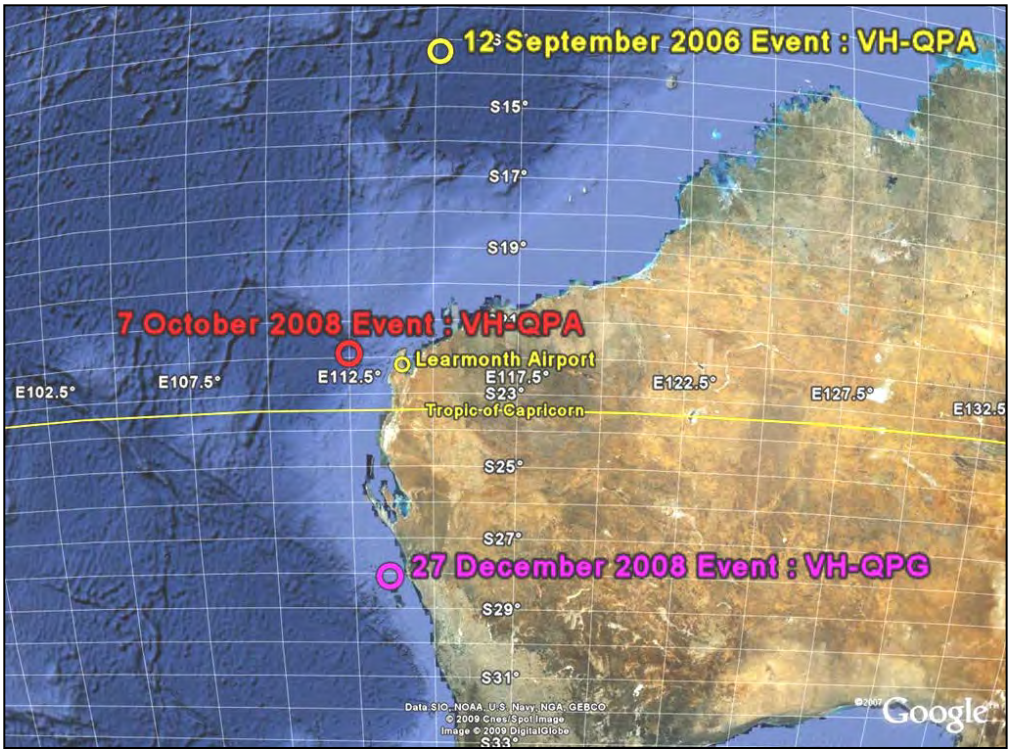

[72] The point shown for each event is where the anomalous ADIRU behaviour commenced.

Table 18: Summary of the three similar ADIRU-related occurrences

	Occurrence 1	**Occurrence 2**	**Occurrence 3**
Date	12 Sep 2006	7 Oct 2008	27 Dec 2008
Time (UTC)	2052	0440	0829
Aircraft model	A330-303	A330-303	A330-303
Aircraft serial number	0553	0553	0603
Aircraft registration	QPA	QPA	QPG
Departure	Hong Kong	Singapore	Perth
Destination	Perth	Perth	Singapore
Altitude	FL410	FL370	FL360
Latitude	13.3712 South	21.9227 South	28.1040 South
Longitude	115.1204 East	112.4983 East	113.6230 East
Distance from Learmonth	980 km north	154 km west	650 km south
ADIRU model	LTN101	LTN101	LTN101
ADIRU serial number	4167	4167	4122
ADIRU position	ADIRU 1	ADIRU 1	ADIRU 1
Software version	0315	0316	0316
Autopilot engaged	Number 2	Number 1	Number 1
Autopilot disconnect	No	Yes	Yes
PFR	NAV IR 1 FAULT, NAV GPS FAULTs, A.ICE FAULTS and other messages	NAV IR 1 FAULT, NAV GPS FAULTs, A.ICE FAULTS and other messages	NAV IR 1 FAULT, NAV GPS FAULTs, A.ICE FAULTS and other messages
ADIRU 1 BITE data	Not available	No faults, problems with routine messages	No faults, problems with routine messages
Other systems' BITE data	Not available	Problems with ADIRU 1	Problems with ADIRU 1
Crew description	Numerous ECAM messages, constantly changing; overspeed warnings, stall warnings	Numerous ECAM messages, constantly changing; overspeed warnings, stall warnings	Numerous ECAM messages, constantly changing
Crew actions	ADR 1 selected OFF after 30 minutes (after ADR fault light observed)	ADR 1 left ON (no fault light)	ADR 1 selected OFF after 28 seconds (due to new crew procedure)
FDR/QAR data	Not available	Spikes on all IR and ADR parameters	Spikes on all IR and some ADR parameters
Effect on flight control system	No	Yes	No

Comparison of the three occurrences to other operations

The three known data-spike occurrences occurred on two aircraft in the same operator's fleet. The investigation examined if there were any unique aspects of the operator's aircraft, maintenance practices or operational practices that could have been related to the ADIRUs' anomalous behaviour. No related factors were found. More specifically:

- Aircraft QPA and QPG were manufactured at different times. The operator advised that QPA and QPG did not have any unique configurations or types of aircraft systems that made them different to the operator's other A330-303 aircraft.
- The aircraft manufacturer advised that the operator's A330 aircraft did not have any unique aircraft systems compared with the rest of the world A330/A340 fleet. That is, each aircraft system used by the operator was also used by at least one other operator that had LTN-101 ADIRUs fitted to its aircraft.
- The aircraft manufacturer, operator and ADIRU manufacturer advised that there was nothing unique about the operator's processes for operating or maintaining its ADIRUs. All repairs and software upgrades were conducted by the ADIRU manufacturer. All system tests conducted on the aircraft were conducted in accordance with the manufacturer's procedures.The aircraft manufacturer and operator conducted a detailed examination of QPA's ADIRS installation and configuration, focusing on the ADIRU 1 wiring and connections. No problems were found.
- A detailed review of the sequence of events in the 7 October 2008 occurrence identified that nothing unique or unusual occurred during the flight, either on the flight deck or in the cabin. Nothing unique or unusual was reported for either of the other two occurrences. The FDR data for both the 7 October 2008 and 27 December 2008 flights showed nothing unique or unusual prior to the start of the anomalous ADIRU behaviour. There was no problematic cargo carried on any of the three flights.
- The three occurrences happened during cruise; however, with medium or long distance flights the majority of flight time involves cruise. All three occurred at different times during each flight.
- The three occurrences all took place on routes between Perth and Asia (either Singapore or Hong Kong), although there were significant distances between the occurrence locations. The operator advised that about 19% of its A330 sectors in 2008 were on flights between Perth and Singapore or Hong Kong and passed in relatively close proximity to Learmonth. In addition, about 29% of its A330 flights passed within 1,500 km of Learmonth. The investigation also identified that other A330/A340 operators, including operators with LTN-101 ADIRUs fitted, conducted regular flights between Asian locations and Perth.

1.16.3 Comparison to other A330/A340 occurrences

Between 1992 and 2009, A330/A340 aircraft conducted over 28 million flying hours with very few accidents or serious incidents. There have been no other accidents on A330/A340 aircraft associated with the flight control system providing pitch-down commands in response to incorrect ADIRU data.

Up until 2009, the A330 had not been associated with any accidents resulting in fatalities. On 1 June 2009, an Airbus 330-200, operated as flight AF447, impacted the Atlantic Ocean on a flight from Rio de Janeiro, Brazil to Paris, France. An investigation by the French Bureau d'Enquêtes et d'Analyses pour la sécurité de l'aviation civile (BEA) is ongoing, and all the factors contributing to the accident have yet to be determined.[73]

The AF447 accident and the 7 October 2008 accident on flight QF72 involved a very different set of events and conditions. More specifically:

- On the AF447 flight, there were a series of fault messages that showed inconsistencies between the measured airspeeds calculated by the three ADIRUs, as well as consequences on other aircraft systems. Such messages have occurred on several flights near large cumulous cloud build-ups, and they have been associated with pitot probe obstruction due to icing. During the QF72 flight on 7 October 2008, the aircraft did not encounter weather conditions associated with pitot probe icing, there was no problem with the performance of the pitot probes, and there were no fault messages that showed airspeed inconsistencies. Although ADIRU 1 provided some incorrect airspeed data, there was no problem with the data provided by the other two ADIRUs.
- In contrast to the QF72 event, the AF447 event did not involve an ADIRU failure. The ADIRUs on the two aircraft were different models made by different manufacturers, and the model fitted to the A330 that was involved in AF447 has not been associated with any ADIRU data-spike occurrences.
- During both flights, stall warnings occurred. However, during the QF72 flight, the stall warnings were brief, nuisance warnings associated with incorrect AOA data from an ADIRU, and the aircraft did not approach or enter an aerodynamic stall during the flight.

1.17 Organisational and management information

1.17.1 A330 operations

The operator commenced A330 operations in late 2002. At the time of the 7 October 2008 occurrence, the operator's A330 fleet included 10 A330-303 aircraft and six A330-202 aircraft. An associated Australian operator (Jetstar) had five A330-202 aircraft, with a sixth joining its fleet in late 2008. The operator of QPA controlled the maintenance for all 22 A330 aircraft. All the aircraft were obtained new from the manufacturer, and all were fitted with the same types of ADIRUs and FCPCs as fitted to QPA.

Throughout the period from 2003 to 2008, the Qantas/Jetstar A330 fleet conducted 60,973 sectors and accumulated 312,834 flight hours. In 2008, the fleet conducted 9,149 sectors and accumulated 93,406 flight hours.

---

[73] The BEA has released three Interim Reports on the accident which are available at www.bea.aero/en.

1.17.2 Processes for reporting and monitoring aircraft faults

If one of the operator's flight crew noticed a fault or problem with an aircraft system or item of equipment, they were required to complete a technical log entry at the end of each flight.[74] Faults or problems could also be detected by maintenance personnel when reviewing the aircraft's PFR.

Following the report of faults or problems, line maintenance personnel would assess the available information, and conduct inspections and tests as required, using the procedures and troubleshooting guidelines provided by the aircraft and/or equipment manufacturer. If the unit passed the required system tests, then generally it would remain on the aircraft. If the unit did not pass the system tests, or there were other reasons for concern (such as a recurring problem), the unit would be removed from the aircraft and sent to the equipment manufacturer or an authorised organisation for examination.

The operator used a database to record the technical log entries and the remedial actions taken by maintenance personnel. The database contained fields for recording related information, including the aircraft registration, flight details, and a number associated with the affected system or equipment.[75] However, the serial number of a unit was not recorded unless it was removed from the aircraft for examination.

The operator tracked the history of each unit in other databases. The main events that were tracked were modifications (such as software updates), repairs, the location of the unit, and any movements of the unit (aircraft and position number). However, the history for each unit did not include reported faults or problems that did not result in its removal from an aircraft.

The operator's engineering department monitored the reliability performance of each type of system or unit across the fleet. This monitoring was based on statistics such as the mean time between unscheduled removals (MTBUR) and the mean time between failures (MTBF) (section 3.9.1). In addition to monitoring MTBUR and MTBF, the operator's engineering department reviewed technical log entries to identify repeating or related problems with a specific unit on an aircraft. Based on this monitoring, a unit could be removed for more detailed examination.

The operator's procedures for reporting and processing system and equipment faults were consistent with general industry practice.

The aircraft manufacturer reported that, in its experience, recurrent issues were identified by operators, who then contacted the aircraft and/or system manufacturers. It also advised that there are a variety of forums and review meetings that regularly occur involving operators and manufacturers which enable the identification of potential issues or trends.

In addition, the aircraft manufacturer noted that a specific cockpit effect or maintenance fault message could result from an internal problem or a problem external to the relevant system. Accordingly, there can be problems tracking specific types of messages (such as IR or ADR faults) and associating them with a

---

[74] As noted previously, the flight crew could report a problem of a more serious nature to maintenance watch during a flight.

[75] The number was the Air Transport Association (ATA) chapter reference, with each chapter referring to a different system.

specific item of equipment (such as the ADIRU) as not all such fault messages will mean there was actually a problem with that item of equipment.

1.18 Additional information

The format of this report is generally consistent with that recommended by the International Civil Aviation Organization (ICAO). However, to aid readability, detailed factual information on some topics has been included in additional parts rather than contained within the recommended structure of Part 1. More specifically:

- Part 2 discusses the design of the FCPCs, focussing on their algorithm for processing AOA data and how it was developed.
- Part 3 discusses the design of the LTN101 ADIRU and the results of analyses to determine the nature and origins of the data-spike failure mode.
- Part 4 provides a detailed description of cabin safety matters, focussing on passengers' use of seat belts and on the injuries associated with the in-flight upsets.

2 FACTUAL INFORMATION: ELECTRICAL FLIGHT CONTROL SYSTEM

Data from the flight data recorder (FDR) showed that the two pitch-downs on the 7 October 2008 flight were due to elevator movements, and that these movements were not due to flight crew commands or turbulence. The electrical flight control system (EFCS) was designed to command pitch-down movements if it detected that the aircraft's angle of attack (AOA) was too high, and FDR data showed that there were very high (50.625°) spikes in one of the aircraft's three AOA values at about the time of both pitch-downs.

Accordingly, the investigation examined in detail the design of the EFCS, the role that the AOA spikes may have had on the pitch-downs, the suitability of the EFCS' algorithm for processing AOA data, and the processes used to develop that algorithm.

2.1 A330/A340 flight control system design

2.1.1 Design overview

Computers

As discussed in section 1.6.3, the A330/A340 EFCS contained five computers:

- Three flight control primary computers (FCPCs, commonly known as PRIMs). The FCPCs generated the control orders in normal, alternate and direct laws, and also directly controlled some of the control surfaces. The FCPCs took inputs from several other systems, including the flight management, guidance and envelope system (FMGES) and the air data inertial reference units (ADIRUs), as well as the flight crew controls.
- Two flight control secondary computers (FCSCs, commonly known as SECs). The FCSCs directly controlled some of the flight control surfaces, and could take inputs from the flight crew controls in direct law.

The computers generated flight control commands at two levels: control orders and servo signals.

Control orders

One of the three FCPCs acted as the 'master'. Based on all of its inputs, the master computed the control orders for each control surface and sent these orders to the other computers to be executed. FCPC 1 was normally the master, and the other FCPCs monitored the master's operation and could take over the role of master if a fault was detected.

Servo signals

All of the flight control surfaces, including the elevators and ailerons, contained servo-controlled[76] hydraulic actuators and position sensors. To execute the control orders from the master FCPC, and change the position of a control surface, a computer sent servo signals to the relevant actuators.

Each of the five computers acted as the servo-controller for different control surfaces. For example, FCPC 1 normally generated the servo signals for the elevators and the trimmable horizontal stabiliser (THS). With regard to the ailerons, FCPC 1 normally controlled the left inboard aileron servos, FCPC 2 normally controlled the right inboard aileron servos, and FCPC 3 normally controlled the outboard aileron servos.

For the purpose of redundancy, multiple computers were connected to each control surface. If a computer was unable to execute the master FCPC's orders for a particular control surface due to a fault, then another computer would take over that servo-controller role. For example, the priority sequence for acting as the servo-controller for the elevators was FCPC 1, FCPC 2, FCSC 1 then FCSC 2.

Not all of the computers could send servo signals to each of the flight control surfaces. For example, FCPC 3 could not perform the servo-control function for the elevators.

2.1.2 Fault-tolerant design features

'Fault tolerance' refers to a system's ability to maintain its functionality in the presence of faults. Fault-tolerant design features are used extensively in the design of hardware and software for safety-critical systems such as a flight control system. A fundamental assumption is that faults can never be fully eliminated, but their probability and consequences can be managed to an acceptable level.

The A330/A340 EFCS included a range of design features to provide fault tolerance. These included:

- Redundancy. The use of five different computers provided redundancy in the event of a failure of one or more computers. In the presence of certain types of faults or processing problems, the role of master switched from one FCPC to another FCPC. In addition, the servo-controller for a control surface could switch to another FCPC or FCSC.
- Self-checking pairs. Each computer had two physically independent channels. The command (COM) channel computed the control orders and/or servo signals, and the monitor (MON) channel conducted the same computations and compared the results. The two channels had their own processor, power supply, memories, and input/output circuits. The use of two channels helped to identify hardware or processing problems (section 2.1.3).
- Monitoring. Each computer had built-in test equipment (BITE) to monitor its own performance and that of the other computers, as well as to monitor other elements of the system such as actuators and sensors. The FCPCs also

[76] A servo, or servo-mechanism, is a control device that uses an automated feedback loop to improve accuracy. Servo-mechanisms are often used to allow a low-power signal to drive a higher-power device.

monitored external systems that provided data to them, such as the ADIRUs, to check the validity and consistency of this data (section 2.1.4).

- Data diversity. The computers' clocks were not synchronised, and the COM and MON channels' clocks in each computer were not synchronised. The computers and channels therefore used data sampled from the sensors and external systems (such as the ADIRUs) at slightly different times, which added robustness to the monitoring processes.
- Equipment dissimilarity. The hardware and software for the FCPCs and FCSCs were different. In addition, the software for the COM and the MON channels were developed by different teams using the same specification. The use of separate design implementations reduced the potential influence of common-mode failures or software coding errors.
- Flight control law reconfiguration. If there were certain types of faults or processing problems, the EFCS reverted to a lower level of control law because it could not provide flight envelope protections with the appropriate level of reliability.
- Physical segregation. The computers were installed in separate locations on the aircraft, which helped prevent a total loss of functionality in the event of some types of damaging events. Hydraulic and electrical system routes were also segregated.

2.1.3 FCPC self-monitoring logic

As illustrated in Figure 27, each FCPC consisted of two physically independent channels; a COM channel and a MON channel. Both channels obtained separate inputs from other systems, such as the ADIRUs and the position sensors in the control surfaces.

Figure 27: Self-monitoring processes for elevator movement (simplified)

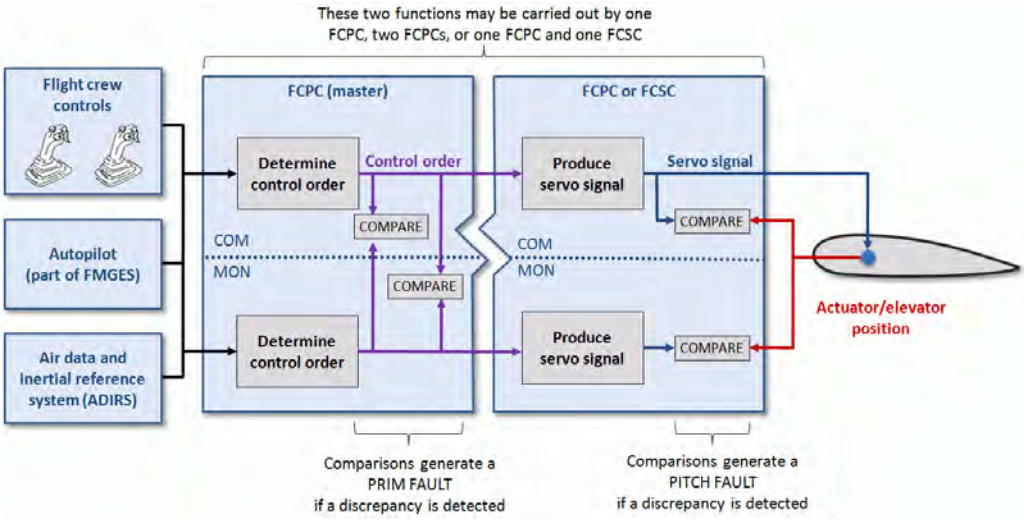

There were two primary comparisons between the COM and MON channels' processes. The first comparison was conducted on the control orders generated by the FCPC when it was the master. More specifically:

- The COM channel computed the control orders, and the MON channel conducted the same computations and compared the results.
- If a difference was detected that exceeded a threshold value for a predetermined period of time, then:
  - the master FCPC was rendered inoperative[77]
  - a F/CTL PRIM [1, 2 or 3] FAULT message was generated
  - the FCPC with the next highest priority took over as the master.

The second comparison was conducted on the servo-control loop with each control surface. More specifically:

- The computer's COM channel sent servo signals to the hydraulic actuators to move the control surface to the appropriate position (based on orders from the master FCPC). The MON channel also computed the appropriate control surface position and compared the results.
- If a difference was detected between the computations and the position of the actuator or the control surface, and the difference exceeded a threshold value for a predetermined period of time, then:
  - the computer no longer performed that control function (but the rest of the computer's functions were still available)
  - a fault message related to the computer and the control surface was generated
  - the computer with the next highest priority for the affected control surface took over as the servo-controller.

For example, if a discrepancy was detected between the elevator position and the computed elevator position, and this difference was confirmed after the specified time period, then a F/CTL PRIM 1 PITCH FAULT was generated. FCPC 2 then became the servo-controller for the elevators.

The predetermined time period for the servo-control loops comparison was significantly shorter than that for the control order comparison.

2.1.4 Monitoring of ADIRU parameters

General monitoring logic

As the EFCS had full authority over the aircraft's flightpath in normal law, it was important that an FCPC's control orders were based on the most accurate flight data parameters from the ADIRUs.

[77] In some cases the flight crew could reactivate the affected FCPC by following the recommended ECAM actions.

Each FCPC used a number of parameters from each ADIRU. The FCPC's software monitored the parameters as follows:

- If an ADIRU flagged the data from one of its parameters as invalid in terms of its sign/status matrix (SSM) (section 3.3), then the FCPC ignored it.
- The FCPC compared the three ADIRUs' values of each parameter for consistency. If any of the values differed from the median (middle) value by more than a threshold amount for longer than a set period of time, then the FCPC rejected the relevant part of the associated ADIRU (that is, ADR or IR) for the remainder of the flight. This rejection did not result in a warning or caution message for the flight crew.

AOA monitoring logic

The basic process used by the FCPCs for checking the validity and consistency of the AOA data is shown in the upper portion of Figure 28. The FCPCs monitored the three ADIRUs' output values every 40 msec (or 25 times per second). If any of the three values deviated from the median by more than a predetermined threshold for more than 1 second, then the FCPC rejected the relevant ADR for the remainder of the flight.

The FCPC conducted additional comparisons when determining the AOA value to use when computing the control orders (section 2.1.5).

2.1.5 Determining ADIRU values for computing control orders

General computation logic

For most of the ADIRU parameters, the FCPC software used the median value when computing the control orders. The use of the median of three values as the system input was a common technique to ensure that significant discrepancies in one value would not influence system performance. Although the use of the median values as the input is generally a robust process, there are still many aspects that need to be considered to ensure that an algorithm works effectively (Butler 2006).

Figure 28: FCPC algorithm for processing AOA data[78]

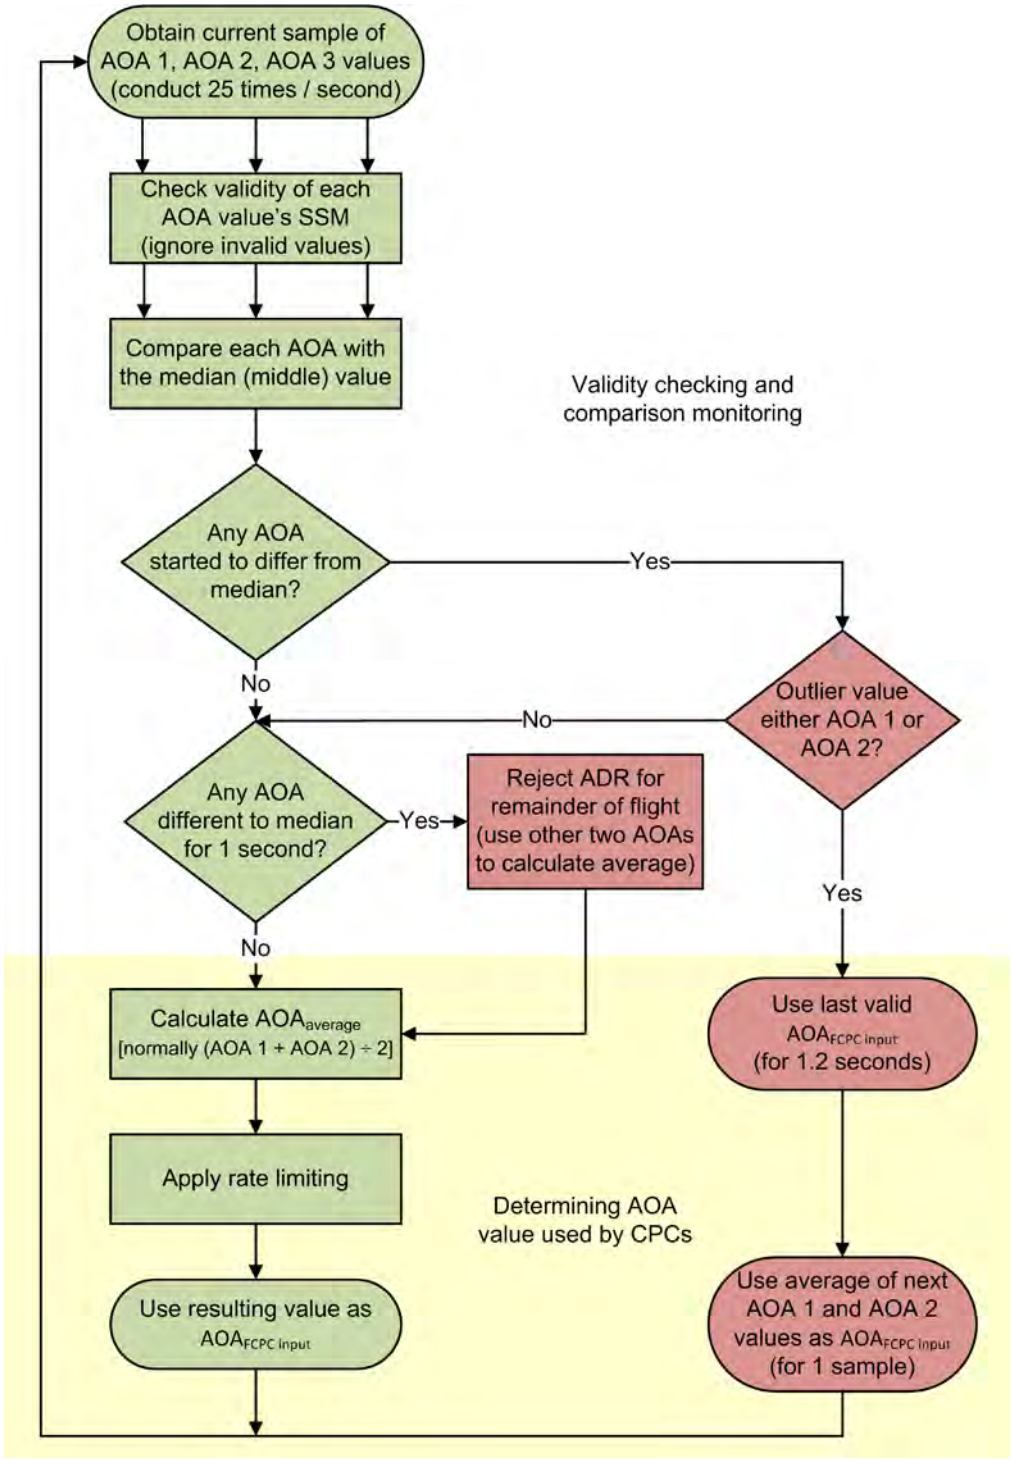

[78] This flowchart is a simplification of the actual algorithm, and is also represented in a different format to that used by the system designers.

AOA computation logic

The FCPC software did not use the median as the AOA value for computing flight control orders ($AOA_{FCPC\ input}$) because of the physical location of the AOA sensors. The AOA 1 sensor was located on the left side of the fuselage and the AOA 2 and AOA 3 sensors were located close together on the right side of the fuselage (section 1.6.4). As a result, there was a potential for the AOA 2 and AOA 3 sensors to provide values that were significantly different to the AOA 1 sensor in some situations such as aircraft sideslip.[79] If both AOA 2 and AOA 3 varied from the correct value in a consistent manner, then AOA 1 would be rejected even if it was closest to the correct value.

In order to minimise these effects, the FCPCs used the average value of AOA 1 and AOA 2 ($AOA_{average}$) to calculate $AOA_{FCPC\ input}$. In addition to the monitoring logic discussed in section 2.1.4, the FCPCs used other mechanisms to prevent discrepancies in either AOA 1 or AOA 2 from influencing $AOA_{FCPC\ input}$ as follows (see also the lower portion of Figure 28):

- $AOA_{FCPC\ input}$ was rate limited to ensure that any rapid changes did not have a significant effect on the FCPC's computations.
- If either AOA 1 or AOA 2 deviated from the median value (of all three AOA values) by more than a predetermined value (or 'monitoring theshold'), the most recent valid $AOA_{FCPC\ input}$ value was memorised and used for 1.2 seconds. During this memorisation period, the current values for AOA 1 and AOA 2 were not used for determining $AOA_{FCPC\ input}$.

At the end of a 1.2-second memorisation period, the FCPCs used the average of the current AOA 1 and AOA 2 values as the $AOA_{FCPC\ input}$ for one sample. No rate limiting was applied to this value, and there was no comparison between the three AOA values (other than the 1-second consistency monitoring). After using this one sample, the FCPCs returned to the normal operating mode (that is, they used the current $AOA_{average}$ value with rate limiting applied).

The monitoring processes discussed in section 2.1.4 occurred at the same time as the process to calculate $AOA_{FCPC\ input}$. Therefore, if there was a discrepancy in AOA 1 (or AOA 2) such that it was significantly different from the median value, the FCPC started the 1-second monitoring period as well as the 1.2-second memorisation period. If the value of AOA 1 remained above the monitoring threshold, the FCPC rejected ADR 1 after 1 second. Then, at the end of the 1.2-second period, it used the average of the two remaining AOA values (AOA 2 and AOA 3) for subsequent computations of $AOA_{FCPC\ input}$.

The FCPC algorithm for processing AOA was unique to the A330/A340 aircraft. Section 2.5.1 provides information on the development of the algorithm.

2.1.6 FCPCs' ability to manage incorrect AOA data

The FCPC algorithm for processing AOA data was able to detect and manage almost all situations involving incorrect or inconsistent AOA data being sent from the ADIRUs as 'valid' data. Some typical examples are presented in Figure 29, and described further below.

---

[79] Sideslip is a condition in which an aircraft's flightpath is displaced right or left from the longitudinal axis.

Step-change

A significant step-change of either AOA 1 or AOA 2 would trigger a 1.2-second memorisation period and have no effect on $AOA_{FCPC\ input}$ (see scenario A in Figure 29). If the change lasted for more than 1 second, the FCPC would reject the relevant ADR and, following the 1.2-second memorisation period, $AOA_{FCPC\ input}$ would be based on the average of AOA 3 and the remaining AOA value.

A step-change less than the monitoring threshold would have a constant but very minor effect on $AOA_{FCPC\ input}$ computations.

Runaway

A runaway (consistently increasing or decreasing value) of either AOA 1 or AOA 2 would trigger a 1.2-second memorisation period when it reached the monitoring threshold. The runaway value would subsequently have no effect on $AOA_{FCPC\ input}$ (see scenario B in Figure 29). Following the memorisation period, $AOA_{FCPC\ input}$ would be the average of AOA 3 and the remaining AOA value.

Prior to reaching the monitoring threshold, the runaway value would have a brief, minor effect on $AOA_{FCPC\ input}$.

Spike(s)

A single, short-duration spike in AOA 1 or AOA 2 would trigger a 1.2-second memorisation period, with the last valid $AOA_{FCPC\ input}$ used during that period. Following the memorisation period, $AOA_{FCPC\ input}$ would again be based on current values of AOA 1 and AOA 2.

The occurrence of any other spikes during the 1.2-second memorisation period would not trigger another memorisation period (see scenario C in Figure 29). The occurrence of a spike after the end of the memorisation period would trigger a new memorisation period.

Spike values 1.2 seconds apart

Following the 7 October 2008 occurrence, the aircraft manufacturer identified a scenario in which deviations in only one of the three AOA values could significantly influence $AOA_{FCPC\ input}$. The scenario involved two spikes in AOA 1 (or AOA 2) with the following properties:

- The first spike was different to the median of the three AOA values, triggering a 1.2-second memorisation period. This spike lasted for less than 1 second.
- The second spike was present at the end of the memorisation period (1.2 seconds after the start of the first spike).

Figure 29: FCPC processing of several AOA input scenarios

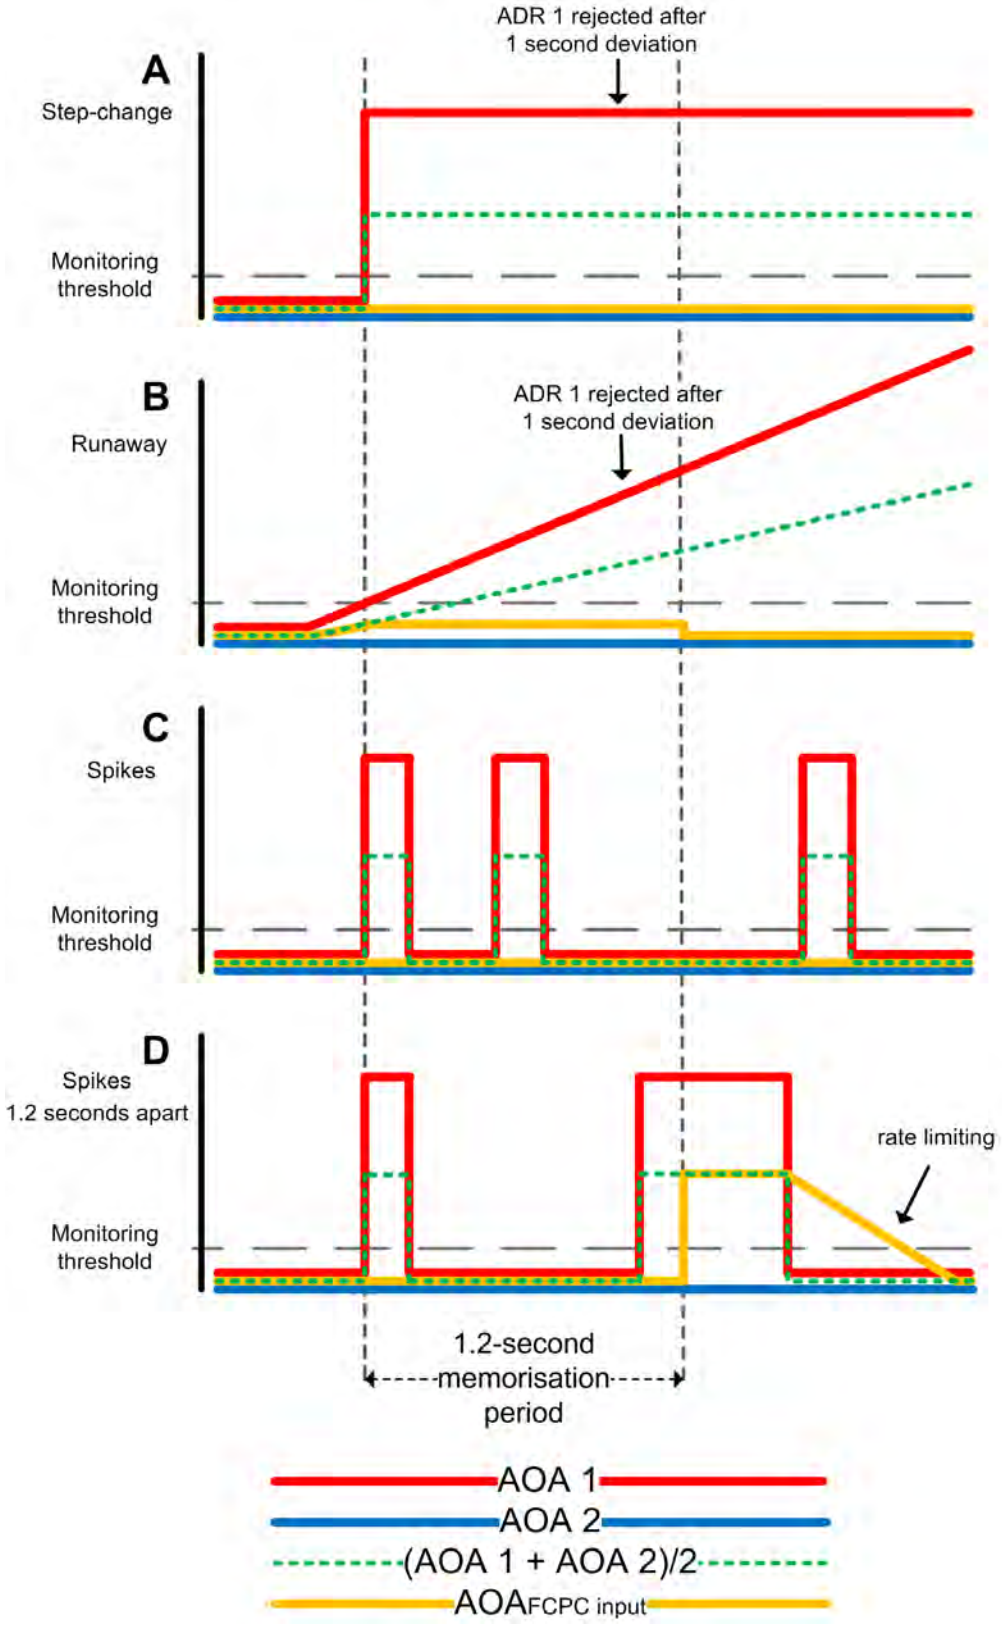

If this scenario occurred, the FCPC would treat the second AOA spike as valid data after the end of the memorisation period.

Scenario D in Figure 29 illustrates an example of this scenario involving incorrect inputs from AOA 1. AOA 2 remains at the actual value, while AOA 1 has two large-value spikes with the second present 1.2 seconds after the start of the first . The calculation of $AOA_{FCPC\ input}$ is as follows:

- Before the first spike, the average of AOA 1 and AOA 2 ($AOA_{average}$) is used as the basis of $AOA_{FCPC\ input}$.
- After the first spike is detected, the last valid value of $AOA_{FCPC\ input}$ is used for 1.2 seconds.
- After the 1.2-second memorisation period, the next values of AOA 1 and AOA 2 are assumed to be valid, and they are used to calculate $AOA_{FCPC\ input}$.
- Following this initial sample, the normal mode of calculating $AOA_{FCPC\ input}$ is again used (that is, the current $AOA_{average}$ followed by rate limiting). This means that, while the second AOA 1 spike is still present, it directly affects $AOA_{FCPC\ input}$.
- After the second spike stops, the rate limiting causes the $AOA_{FCPC\ input}$ value to steadily decrease back to the correct value (that is, the current average of AOA 1 and AOA 2).

Simultaneous AOA deviations

Another scenario where incorrect AOA values from the ADIRUs could significantly influence the value of $AOA_{FCPC\ input}$ involved simultaneous deviations of two AOA values. The deviations had to be of similar magnitude throughout their period of deviation. If this occurred, the algorithm would reject the correct AOA and base the $AOA_{FCPC\ input}$ on the average of the two remaining, incorrect values.

The aircraft manufacturer had identified this scenario when it developed the A330/A340 aircraft. However, a simultaneous failure of two ADRs (or related components) was considered to be 'extremely improbable' (section 2.5.3).[80]

2.1.7 Effects of AOA changes on elevator control orders

Two of the EFCS's flight envelope mechanisms could respond to high $AOA_{FCPC\ input}$ values and initiate a nose-down elevator command: high AOA protection and anti pitch-up compensation. If both corrective mechanisms were triggered at the same time, their contributions were added. The characteristics of the two EFCS mechanisms are summarised in Table 19.

[80] By extension, another potential (and less likely) scenario could involve simultaneous deviations of all three AOA values.

Table 19: Characteristics of elevator control mechanisms

	High AOA protection	Anti pitch-up compensation
Control law	Normal law only	Normal or alternate law
Speed	Any	Mach 0.65 or more
Altitude	Any (must exceed threshold for at least 2 seconds when aircraft below 500 ft)	Any
Configuration	Any	Landing gear retracted, flaps up
Maximum authority	4° elevator movement (at time of pitch-downs)	6° elevator movement

High angle of attack protection

Aerodynamic stall in large aircraft is a potentially dangerous condition and aircraft manufacturers incorporate design techniques to prevent it. On the A330/340, the FCPCs continually monitored the $AOA_{FCPC\ input}$. If the master FCPC detected that this value exceeded a predefined threshold (alpha max), then it issued control orders for a nose-down elevator movement to reduce the AOA and prevent a stall.

High AOA protection was only available when the aircraft was in normal law. If $AOA_{FCPC\ input}$ was outside the range of -10° to +30°, the control law reverted from normal law to alternate law, and the protection was therefore no longer available.[81] In addition, when the aircraft was more than 500 ft above ground level, the protection was effective immediately; when the aircraft was below 500 ft, it was only active after $AOA_{FCPC\ input}$ exceeded the threshold for 2 seconds or more.

The maximum authority or change in elevator movement that could result from the the high AOA protection varied depending on several factors. The aircraft manufacturer reported that, at the time of the two in-flight upsets, the maximum authority was about 4° of elevator movement. The protection would be applied until the aircraft's AOA was reduced below the stall angle.

Anti pitch-up compensation

Anti pitch-up was a mechanism included in the A330's control laws to compensate for a pitch-up[82] tendency at high Mach numbers and high AOA. The compensation was only available above Mach 0.65 and when the aircraft was in a 'clean' configuration (that is, with the landing gear and flaps retracted). Its maximum authority was 6° of elevator movement.

---

[81] A number of different conditions could lead to the control law reverting to alternate law. The condition that led to the reversion on the 7 October 2008 flight is discussed in section 2.2.2.

[82] On a statically stable aircraft, the centre of lift is situated behind the aircraft's centre of gravity and an increase in AOA would lead to an increase in lift and a restorative tendency to pitch the aircraft nose down. However, at higher Mach numbers and AOAs, it is possible to stall the wing tips. On a swept-wing aircraft, the centre of lift will then move forward, leading to a reduced nose-down reaction to increasing AOA, reduced stability, and a tendency to pitch up.

Potential effects during different phases of flight

The two flight envelope mechanisms would not issue pitch-down commands when the aircraft was close to the ground. More specifically:

- The high AOA protection had a confirmation time of 2 seconds when the aircraft was below 500 ft above ground level. As AOA spikes longer than 1 second would result in the relevant ADR being rejected, this protection could not be activated by a multiple AOA spike scenario that occurred below 500 ft.
- The anti pitch-up compensation was only available above Mach 0.65 and in the clean configuration. Therefore, this mechanism would not realistically occur during final approach or initial climb situations, when the aircraft was operating in close proximity to the ground.

Following the 7 October 2008 occurrence, the aircraft manufacturer conducted flight simulations with the aircraft just above 500 ft and at typical approach speeds. Based on the high AOA protection command alone,[83] the decrease in altitude due to a pitch-down without a flight crew response was not significant (less than 100 ft). When the flight crew were asked to respond to the pitch-down, the decrease in altitude was much less. During these simulated events, there was an autopilot disconnection, but the flight crew easily recovered control.

2.2 Examination of FCPC performance on 7 October 2008

2.2.1 Simulations to determine role and origin of elevator deflections

The aircraft manufacturer conducted a series of simulations to determine the nature of the factors that could have contributed to the pitch-downs on the 7 October 2008 flight. These activities were performed using an engineering simulation tool developed during the aircraft development process (see section 2.4.4).

Effect of elevator deflections

The first simulation study used the elevator deflections recorded on the FDR and the flight conditions present at the time of the first pitch-down to confirm the aircraft's pitch movements during the event. The study found that the recorded elevator deflections were sufficient to explain the aircraft movement, which confirmed that turbulence did not contribute to the pitch-downs.

Role of sidestick inputs

The second study used the flight crew's sidestick pitch inputs that were recorded on the FDR as inputs into the engineering simulation tool. It confirmed that the sidestick inputs were not sufficient to explain the recorded elevator deflections at the start of the pitch-downs, and therefore another factor was involved in initiating these deflections.

---

[83] Anti pitch-up compensation was not included as it would not realistically be available at that altitude.

Simulation of AOA values for the first pitch-down

The third study determined if values of $AOA_{FCPC\ input}$ could lead to the recorded elevator deflections for the first pitch-down (at 0442:27), after taking account of the sidestick pitch inputs. The simulations used a two-spike scenario with the following properties:

- the first AOA 1 spike triggered a 1.2-second memorisation period with the value of $AOA_{FCPC\ input}$ being 2.3° (based on the recorded AOA 1 and AOA 2 values before the spike)
- the second AOA 1 spike was present at the end of the memorisation period, leading to a step-change of $AOA_{FCPC\ input}$ to 26° (based on an average of an AOA 1 spike of 50.6° and an AOA 2 value of 2.3°)
- a steady decrease in the 26° $AOA_{FCPC\ input}$ value after the spike stopped (due to the rate limiting function).

The time period that the second spike existed after the end of the memorisation period was then varied. This time period is represented by the value of 'X' in Figure 30.

Figure 30: AOA values used in the simulation study for first pitch-down[84]

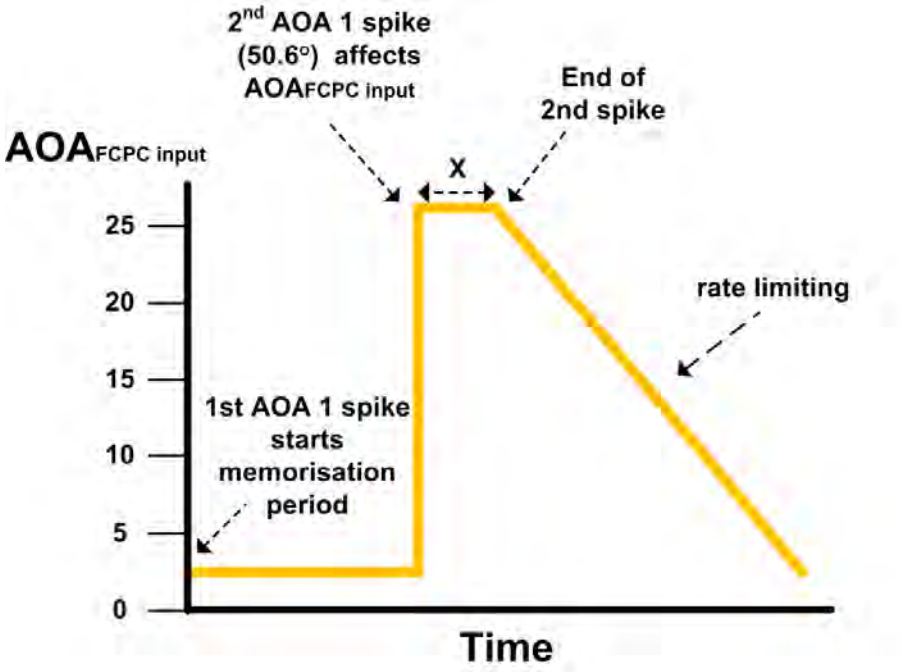

This study confirmed that the recorded elevator deflections during the first pitch-down could be produced using a 26° value of $AOA_{FCPC\ input}$ for a duration of about 400 msec at the end of the memorisation period (that is, the value of X was 400 msec). Based on this simulation, it could be concluded that the AOA 1 spike of 50.6° lasted at least 400 msec but less than 1 second.[85]

---

[84] Scenario D in Figure 29 illustrates how AOA 1 spikes could produce these $AOA_{FCPC\ input}$ values.

[85] If the duration of the spike was greater than 1 second, then the ADR 1 would have been rejected for the remainder of the flight (and a relevant fault message recorded).

A subsequent simulation confirmed that the 10° elevator deflection recorded on the FDR was the result of 4° of high AOA protection and 6° of anti pitch-up compensation.

The aircraft manufacturer advised that, for the situation that occurred during the two pitch-downs, the two corrective mechanisms remained active as long as the $AOA_{FCPC\ input}$ was above the triggering threshold. This period was about 1.8 seconds for the first pitch-down. During this period, the flight crew's sidestick inputs would have had no influence on the aircraft's pitch.

Simulation of AOA values for second pitch-down

The same process was used to determine the AOA values required to replicate the second pitch-down (at 0445:08). A single memorisation period with different spike values and durations did not match the recorded elevator movements. Further studies identified a more complex scenario that did match these movements, and this scenario involved at least four spikes in succession and triggering two memorisation periods in $AOA_{FCPC\ input}$, as shown in Figure 31.[86]

Figure 31: AOA values used in simulation study for second pitch-down

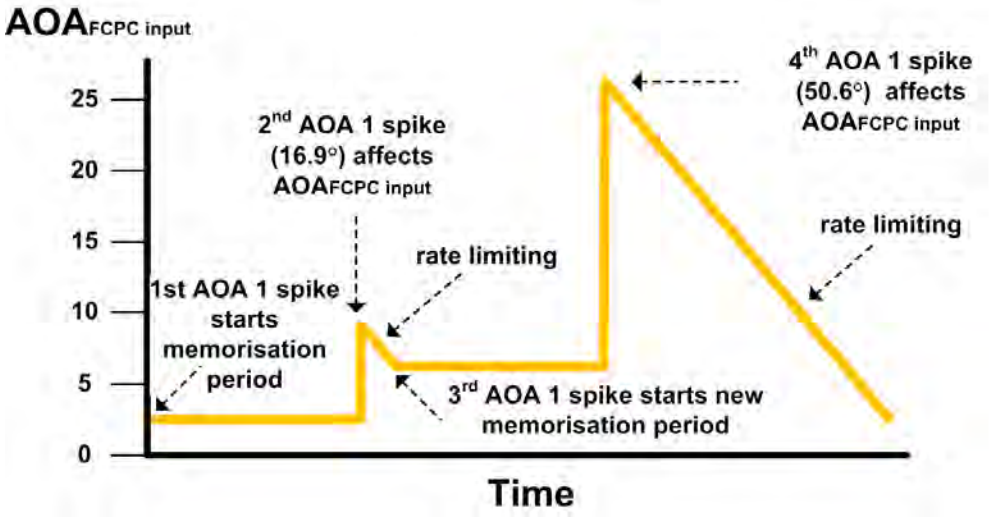

The aircraft manufacturer also advised that the matching scenario required that the control law reverted to alternate law at about the time of the fourth spike. Consequently the high AOA protection was not available after this time, and there was only 6° of elevator deflection from the anti pitch-up compensation. This figure matched the maximum nose-down elevator position of 5.4° that was recorded during the second pitch-down.

Based on this scenario, the flight crew's sidestick inputs would have had no influence on the aircraft's pitch for about 2.8 seconds during the second pitch-down.

---

[86] One of the spikes in the matching scenario was 16.9°. The recorded values of AOA 1 spikes on the FDR at about this time were 50.625°, but values of 16.9° were recorded later in the flight (section 1.11.4).

2.2.2 Review of recorded flight control system fault messages

EFCS event sequence

The sequence of fault messages recorded for the 7 October 2008 flight by the FDR and the post-flight report (PFR) was reviewed to determine their consistency with the designed operating logic of the EFCS computers. The rest of this section discusses these messages, and Table 20 summarises the sequence of events related to the computers' operation.

Table 20: FCPC sequence of events

Time[87]	Event	Master FCPC	Elevator controller	Active law
Prior to 0442:27	Uneventful flight	FCPC 1	FCPC 1	Normal
0442:27	First pitch-down (10° elevator deflection)	FCPC 1	FCPC 1	Normal
0442:29	F/CTL PRIM 1 PITCH FAULT	FCPC 3	FCPC 2	Normal
0442:30	F/CTL FCPC 3 FAULT	FCPC 2	FCPC 2	Normal
0444:31	FCPC 3 status changed from Fault to No Fault (reset by the flight crew)	FCPC 2	FCPC 2	Normal
0445:08	Second pitch-down (6° elevator deflection)	FCPC 2	FCPC 2	Normal
0445:09	F/CTL PRIM 2 PITCH FAULT	FCPC 3	FCSC 1	Normal
0445:10	F/CTL PRIM 3 FAULT	FCPC 1	FCSC 1	Alternate

F/CTL PRIM 1 PITCH FAULT

The first recorded EFCS fault message was the F/CTL PRIM 1 PITCH FAULT, recorded on the PFR at 0442.[88] This message meant that the MON channel of FCPC 1 had detected a difference in the actual elevator position (based on commands from the COM channel) and the expected elevator position calculated by the MON channel. To generate the fault, the difference needed to exceed a threshold value for the predetermined period of time.

The COM and MON channels sampled the ADIRUs' AOA values at different times. Consequently, there was a brief difference in their calculations of $AOA_{FCPC\ input}$ values at the end of the memorisation period when a second AOA spike was present. There was also a longer difference in their $AOA_{FCPC\ input}$ values during the subsequent rate-limiting period due to the two channels starting these periods at different times (Figure 32). These slight differences in the timing of $AOA_{FCPC\ input}$ values were sufficient for the MON channel to detect a difference between its computation of the appropriate elevator position and the actual elevator position commanded by the COM channel.

---

[87] The FDR only recorded some events, such as FCPC FAULTS and the FCPC master, every 4 seconds.

[88] PITCH FAULT messages were not recorded on the FDR or QAR. They were recorded on the PFR to the nearest minute. Although its exact time was not recorded, the message was consistent with what occurred during the first pitch-down.

Following the F/CTL PRIM 1 PITCH FAULT message, FCPC 1 could no longer act as the servo-controller for the elevators, or as the master FCPC in normal law. Based on the system design, another FCPC (in this case FCPC 3) became the master, and FCPC 2 took over as the servo-controller for the elevators.

Figure 32: Simplified representation of the COM and MON differences during the pitch-downs

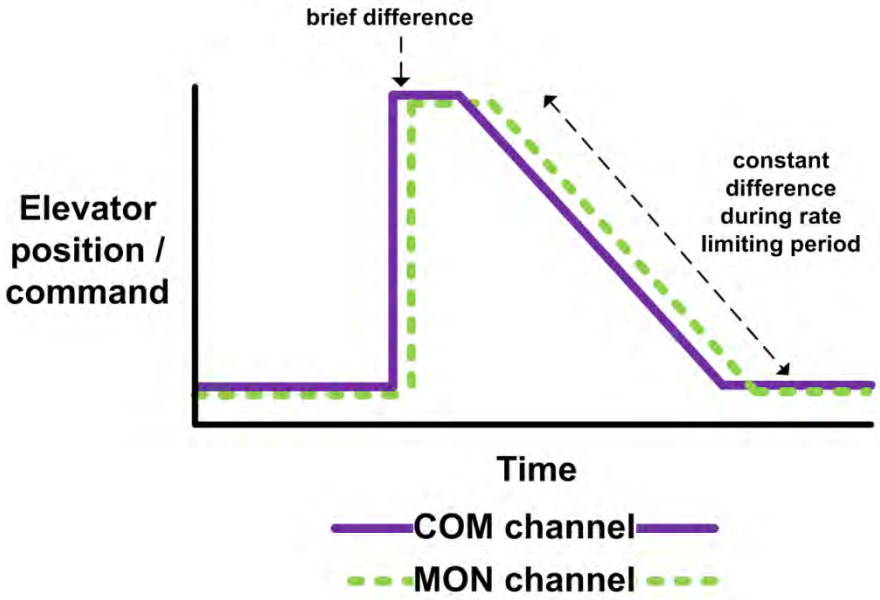

F/CTL PRIM 3 FAULT

A few seconds after the first pitch-down started, there was a F/CTL PRIM 3 FAULT message recorded on the FDR (0442:31). This message meant that the FCPC 3 detected a difference in the control orders calculated by the COM and MON channels, and that this difference was confirmed after the predetermined period of time, which was significantly longer than that for the PITCH FAULT.

As FCPC 3 was not a servo-controller for the elevators, it could not generate a PRIM 3 PITCH FAULT. However, because it was acting as the master FCPC, it was sending control orders for elevator movements to another computer (in this case FCPC 2). These control orders for elevator movement were based on $AOA_{FCPC\ input}$ values. Consequently, there was a discrepancy between the COM and MON channels' computation of the elevator control orders over time by FCPC 3, for similar reasons as for the pitch fault (Figure 32).

Following the F/CTL PRIM 3 FAULT, FCPC 3 was no longer operational and FCPC 2 therefore became the master. After the flight crew followed the relevant ECAM actions to reset FCPC 3 (0444:31), it again became operational.

Subsequent fault messages

At about the time of the second pitch-down, a F/CTL PRIM 2 PITCH FAULT was recorded on the PFR and a F/CTL PRIM 3 FAULT was recorded on the FDR. These fault messages occurred in a similar way as the messages during the first pitch-down. Because FCPC 2 had experienced a pitch fault, it could no longer act as the servo-controller for the elevators, and FCSC 1 took over that function. The role of master FCPC switched initially from FCPC 2 to FCPC 3, which generated a

second PRIM 3 FAULT. The role of master then switched back to FCPC 1. However, because FCPC 1 had already recorded a pitch fault, it was not able to act as the master under normal law. The control law therefore reverted to alternate law.

In addition to generating the messages and changing the control law, the sequence of faults also affected the autotrim function. The priority order for performing the servo-controller role for the THS was FCPC 1, FCPC 2 and FCPC 3. After both FCPC 1 and FCPC 2 had experienced a pitch fault, and FCPC 3 had a PRIM fault, none of them could manage the autotrim function associated with the THS.

2.2.3 Summary

The simulation studies showed that spikes in AOA 1 values 1.2 seconds apart could lead to the FCPCs sending pitch-down commands to the elevators, and that these commands were consistent with the elevator deflections observed during the two pitch-downs. In addition, the studies confirmed that flight crew inputs and turbulence did not contribute to the pitch-down commands. The EFCS fault messages recorded during the flight were also consistent with the operational logic of the system in response to the pitch-down commands.

Given that the pitch-down commands were consistent with the operational logic of the EFCS, the investigation examined the requirements and activities involved in developing the FCPC algorithm for processing AOA data.

2.3 Requirements for designing flight control systems

2.3.1 Certification basis for the A330/A340

Airbus applied for the certification of the A330/A340 aircraft types in June 1988.[89] The applicable certification basis was the European Joint Aviation Requirement (JAR) 25 (change 13, effective 5 October 1989), with some exceptions and special conditions. The A330/A340 aircraft were originally certified by the Direction Générale de l'Aviation Civile (DGAC) of France. The A340-211 was the first model certified in December 1992 (Type Certificate TC 183), with the first A330 model (A330-301) certified in October 1993 (Type Certificate TC 184).

The A330/A340 aircraft were also jointly certified in the United States (US) under the US Federal Aviation Regulation (FAR) 25. The US Federal Aviation Administration (FAA) validated the DGAC certification for the A340 in May 1993 and the A330 in October 1993.[90]

---

[89] Type certification is the process used by a regulatory authority to ensure that a new aircraft type complies with the applicable airworthiness requirements.

[90] Under the provisions of *Civil Aviation Safety Regulation 21.29A*, which was in place at the time the first A330 was certified in Australia (2002), the Australian Civil Aviation Safety Authority (CASA) issued a Type Acceptance Certificate for the A330 based on the fact that it had already been issued with a Type Certificate by the national aviation authority of a recognised foreign country.

2.3.2 Regulatory requirements

The relevant certification requirements for a flight control system were specified in JAR 25.671 and JAR 25.1309. The FARs had the same requirements.

JAR 25.671 (*Control Systems: General*) dealt with specific types of failures that could affect the functioning of the control surfaces. It effectively stated that the aircraft had to be capable of continued safe flight and landing following specific types of 'failures or jamming' in the flight control system or associated control surfaces. A specific failure, or combination of failures, that could affect the continued safe flight and landing had to be demonstrated, by analysis or test, to be 'extremely improbable' (see section 2.3.3 for the definitions of probability terms).

JAR 25.1309 (*Equipment, systems and installations*) applied to a range of different aircraft systems, including the flight control system. It outlined more detailed requirements than JAR 25.671, including the following:

> (b) The aeroplane systems and associated components, considered separately and in relation to other systems, must be designed so that
>
> (1) The occurrence of any failure condition which would prevent the continued safe flight and landing of the aeroplane is extremely improbable; and
>
> (2) The occurrence of any other failure condition which would reduce the capability of the aeroplane or the ability of the crew to cope with adverse operating conditions is improbable...
>
> (d) Compliance with the requirements of sub-paragraph (b) of this paragraph must be shown by analysis, and where necessary, by appropriate ground, flight, or simulator tests. The analysis must consider
>
> (1) Possible modes of failure, including malfunctions and damage from external sources;
>
> (2) The probability of multiple failures and undetected failures;
>
> (3) The resulting effects on the aeroplane and occupants, considering the stage of flight and operating conditions; and
>
> (4) The crew warning cues, corrective action required, and the capability of detecting faults.

2.3.3 European Advisory Circular Joint No. 1 to 25.1309

Advisory Circular Joint (ACJ) No. 1 to JAR 25.1309 outlined guidance for demonstrating compliance with JAR 25.1309. A key concept in the ACJ was a 'failure condition', which referred to a condition that resulted from a failure in the aircraft system, and caused or contributed to an undesirable effect on the functioning of the system or the complete aircraft.[91] The failure condition could also occur due to a scenario involving a combination of failures, including failures in related aircraft systems.

---

[91] An example of a failure condition is 'total loss of wheel braking' and an example of a failure that could lead to this condition is 'brake system control unit – power supply failure'. For further details see ARP4761 (discussed in section 2.6.2).

The ACJ was built around the principle that systems should be designed so that there was an inverse relationship between the severity of the consequences of a failure condition and the condition's probability of occurrence. This concept was described in terms of a series of consequence (or 'effect') levels and probability levels.

The four consequence levels were catastrophic, hazardous, major and minor. They were defined using a range of criteria, as presented in Table 21. The probability levels were probable, improbable (divided into remote and extremely remote), and extremely improbable. In addition to verbal descriptions, the ACJ provided numerical indicators of each probability level (Table 22).

Consistent with JAR 25.1309(b), the ACJ stated that failure conditions associated with a catastrophic effect should be 'extremely improbable', and failure conditions associated with a hazardous effect should be no more likely than 'extremely remote'.

The ACJ advised that the methods of demonstrating compliance with JAR 25.1309(d) would depend on the complexity of the system. In addition, it noted that the assessment of the system should consider a range of factors, including the possible failure modes that could lead to the failure condition, operation of related systems, operating conditions, phase of flight, capability of detecting failures and maintenance procedures. The ACJ also stated that assessments could take account of previous experience using similar systems.

Table 21: Effect levels described in ACJ No. 1 to JAR 25.1309

Effect level	Definition
Catastrophic	• loss of the aeroplane and/or fatalities
Hazardous	• a large reduction of safety margins; • physical distress or workload such that the flight crew cannot be relied upon to perform their activities accurately or completely; or • serious injury to, or death of, a relatively small proportion of the occupants
Major	• significant reduction in safety margins; • reduction in the ability of the flight crew to cope with adverse operating conditions as a result of the increase in workload or as a result of conditions impairing their efficiency; or • injury to occupants
Minor	• airworthiness is not significantly affected and any actions are well within the capability of the crew, such as o slight reduction of safety margins o slight increase in workload o physical effects but no injury to occupants

Table 22: Probability levels described in ACJ No. 1 to JAR 25.1309

Probability level	Qualitative definition	Quantitative description[92]
Probable[93]	may occur once or several times during the total operational life of each aeroplane of the same type	> 10^{-5} per flight hour
Remote (category of improbable)	unlikely to each aeroplane during its total operational life but which may occur several times when considering the total operational life of a number of aeroplanes of the same type	10^{-5} to 10^{-7} per flight hour
Extremely remote (category of improbable)	unlikely to occur when considering the total operational life of all aeroplanes of the same type, but nevertheless, has to be considered as being possible	10^{-7} to 10^{-9} per flight hour
Extremely improbable	so extremely remote that it does not have to be considered as possible to occur	< 10^{-9} per flight hour

2.3.4 United States Advisory Circular 25.1309-1A

The FAA released Advisory Circular 25.1309-1A (*System design and analysis*) in June 1988. When it was released, it provided more detailed guidance than the ACJ regarding the methods that could be used to satisfy the requirements of FAR 25.1309. The background section of the document noted that there had been 'an increase in the degree of system complexity and integration, and in the number of safety critical functions performed by systems' in the years prior to the AC's release. It also stated that due to difficulties in assessing hazards for such systems, more structured approaches were being used for such assessments, which therefore required more detailed guidance.

Some key features of the AC included the following:

- It provided the same guidance as the European ACJ on the concept of an inverse relationship between the severity of the effects of a failure condition and the probability of its occurrence, and this concept was illustrated with the diagram shown in Figure 33. However, the AC used the term 'major' to refer to both the 'hazardous' and 'major' effect levels described in the ACJ, and it occasionally used the term 'severe major' to refer to more serious conditions within this major level.
- It used the term 'fail-safe design', which meant that no single failure should result in a catastrophic failure condition. Although the European ACJ current at the time did not include this term, the FAA and the European Aviation Safety Agency (EASA)[94] advised that it was a commonly held principle in both the FAA and the European certifying authorities at the time.[95]

---

[92] For example, a value of 10^{-3} per flight hour equated to once every 1,000 flight hours, and a value of 10^{-7} per flight hour equated to once per 10,000,000 flight hours.

[93] The range for 'probable' was also split into two, with 'frequent' described as more than 10^{-3} per flight hour and 'reasonably probable' described as 10^{-3} to 10^{-5} per flight hour.

[94] EASA took over the role of aircraft certification in Europe in 2003.

[95] The fail-safe design principle was explicitly stated in later versions of the European ACJ.

- It stated that the identification and classification of failure conditions was necessarily qualitative. However, the assessment of the associated probability level could be either qualitative or quantitative, and the analysis could range widely in scope depending on factors such as the severity of the failure condition and the complexity of the system.
- It outlined brief guidance on the use of specific analysis techniques. It noted that functional hazard analysis (FHA) was a useful technique to identify and classify potentially-hazardous failure conditions, and it also referred to other techniques for identifying the causes and probabilities of failure conditions, including fault tree analysis and failure mode effects analysis (FMEA).
- It noted that the means of compliance described in the AC were not directly applicable to software assessments because it was 'not feasible to assess the number or kinds of software errors, if any, that may remain after the completion of system design, development, and test'. The AC stated that design objective (DO) 178A provided an acceptable means of compliance for assessing and developing the software used in computer-based systems.

Figure 33: Probability versus consequences graph (from AC25.1309-1A)

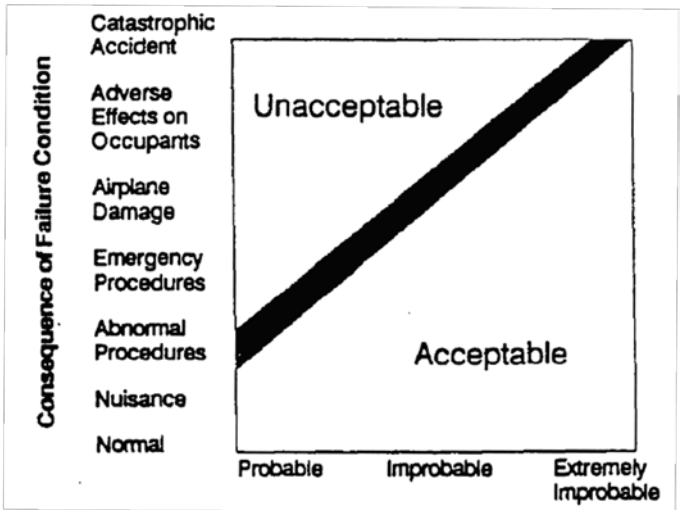

2.3.5 Design objective 178A

DO-178A (*Software considerations in airborne systems and equipment certification*) was produced by the Radio Technical Commission for Aeronautics (RTCA)[96] in March 1985. The purpose of the document was to 'describe techniques and methods that may be used for the orderly development and management of software for airborne digital computer-based equipment and systems'.

The design objective outlined three software levels that enabled the development process to be tailored in accordance with a system's criticality. The levels referred to the degree of stringency or thoroughness required by the manufacturer's development processes to provide design assurance, with Level 1 software requiring the highest standard.

[96] RTCA is a private, not-for-profit corporation that develops consensus-based recommendations regarding communications, navigation, surveillance, and air traffic management system issues.

System criticality was also defined at three levels:

- Critical. Functions for which a failure condition or design error would prevent continued safe flight or landing. Generally associated with Level 1 software.
- Essential. Functions for which a failure condition or design error would reduce the capability of the aircraft or the ability of the crew to cope with adverse operating conditions. Generally associated with Level 2 software.
- Non-essential. Functions for which a failure condition or design error could not significantly degrade aircraft capability or crew ability. Generally associated with Level 3 software.

DO-178A provided high-level guidance for the generation of software requirements, the verification that the resulting design met the requirements, and validation that the requirements were adequate. It also noted that for systems that performed certain critical and essential functions:

> ...it may not be possible to demonstrate an acceptably low level of software errors without the use of specific design techniques. These techniques, which may include monitoring, redundancy, functional partitioning or other concepts, will strongly influence the software development program, particularly the depth and quality of the verification and validation effort...
>
> NOTE: It is appreciated that, with the current state of knowledge, the software disciplines described in this document may not, in themselves, be sufficient to ensure that the overall system safety and reliability targets have been achieved. This is particularly true for certain critical systems such as full authority fly-by-wire. In such cases it is accepted that other measures, usually within the system, in addition to a high level of software discipline may be necessary to achieve these safety objectives and demonstrate that they have been met.

2.4 Development of the A330/A340 flight control system

2.4.1 Overview of aircraft systems development

The process for developing a safety-critical system needs to provide assurance that hardware failures and design errors are minimised, and that the likelihood that these failures and errors lead to failure conditions that affect safety is also minimised. Aircraft manufacturers provide that assurance through their processes for generating, verifying and validating the requirements for the system's design.

The basic structure of a system development process is usually represented as a V-cycle (Figure 34), where time is represented horizontally (left to right) and system hierarchy represented vertically (with the whole aircraft at the top). Initially (top left), the top-level design requirements for the whole aircraft are developed. The aircraft is then decomposed into systems and the requirements for each system are specified. Each system is then decomposed into parts or items of equipment (including software), and the requirements for each item are specified. As the requirements become more detailed from the aircraft level down to the equipment level, they represent design decisions and essentially become part of the system design.

The second part of the V-cycle involves a series of evaluation activities to ensure the suitability of the final product. These activities are known as ‘verification’ and ‘validation’. Verification is the process of ensuring that the final product meets the requirements (that is, the product was built correctly). Validation is the process of ensuring that the requirements are sufficiently correct and complete (that is, the right product has been built).

Figure 34: Simplistic representation of aircraft development process

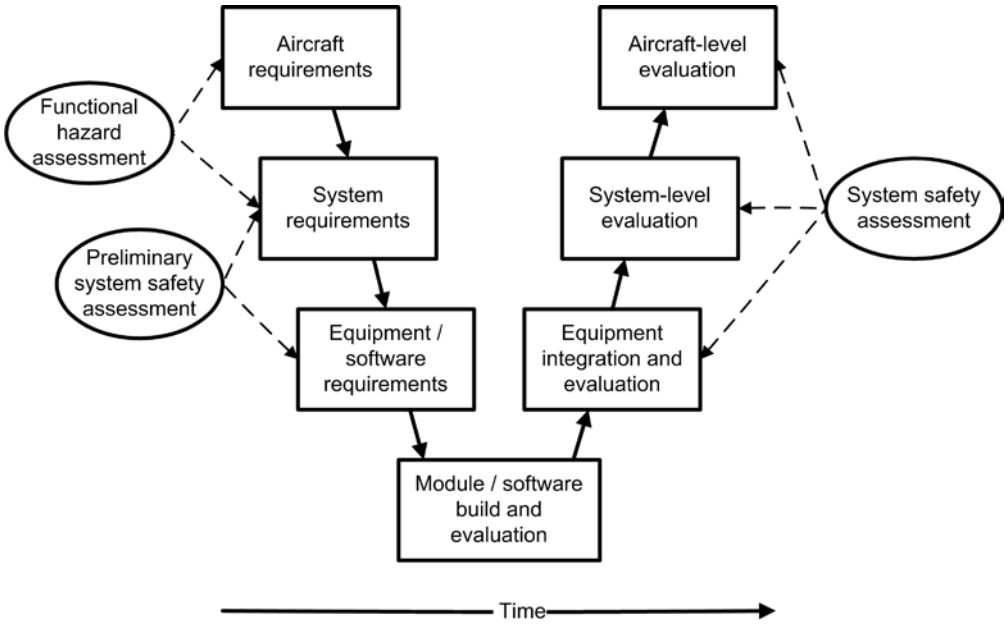

Verification and validation activities include testing the individual items of equipment (and software), and then progressively integrating the equipment into systems for more sophisticated testing activities, until the aircraft is evaluated as a complete entity. Verification and validation also include methods such as peer reviews, modelling and other analyses. The V-cycle is an iterative rather than a fixed process as the verification and validation activities can lead to design changes throughout the cycle.

As indicated in Figure 34, safety assessment activities are a key part of a good system development process. Initial safety assessment activities help evaluate the initial design and derive requirements, and a final system safety assessment provides assurance that the resulting system meets the safety requirements.

It is generally accepted that, for all but the simplest systems, it is impossible to guarantee the correctness of all the system requirements and associated assumptions. In order to reduce the potential effect of errors in the requirements or subsequent design implementation, systems are designed with fault-tolerant features in their architecture, such as redundancy and dissimilarity.

2.4.2 Overview of the development process for the A330/A340 EFCS

The aircraft manufacturer's process for developing the A330/A340 EFCS is summarised in Figure 35.[97] It followed the general structure of a typical V-cycle, although it included the significant use of modelling and simulation tools to help validate the system requirements before the initial version of the equipment was built.

Figure 35: V-cycle representing the A330/A340 EFCS development process[98]

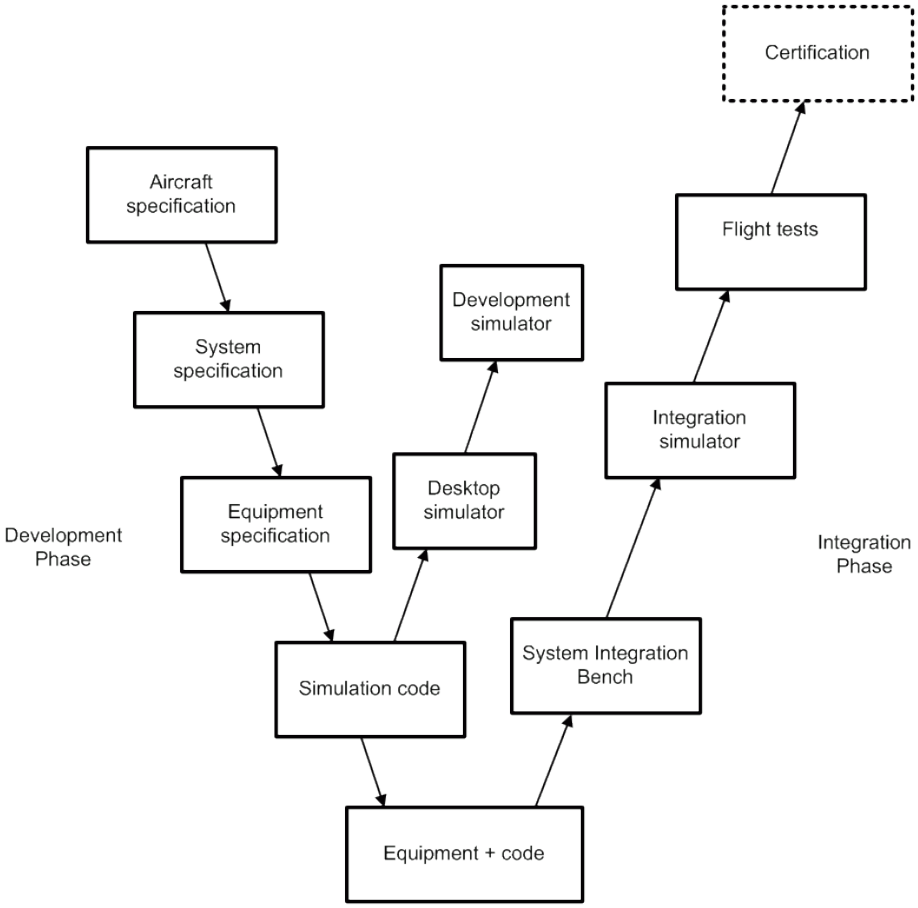

2.4.3 System requirements

The A330/A340 EFCS specification (or set of requirements) included:

- Safety requirements. Due to the EFCS's safety-critical function, the aircraft manufacturer designed it to meet stringent safety requirements. These requirements were based on the certification requirements and associated advisory material (section 2.3), as well as requirements that were identified by the manufacturer during its safety assessment activities (section 2.4.5).

---

[97] The information in sections 2.4.2, 2.4.3 and 2.4.4 about the Airbus design process for the EFCS on the A320 and A330/A340 was outlined in Briere and Traverse (1993), Briere, Favre and Traverse (1995), Fauve (1994), Goupil (2010) and Traverse, Lacaze and Souyris (2004).

[98] Adapted from Goupil (2010).

- Functional requirements. These requirements included the flight control laws, data monitoring, operation of control surfaces, reconfigurations, operating logic, and equipment input/output information. They were written using a graphic computer-assisted method, known as SAO[99], which ensured that each of the language's 'symbols' was formally defined and that there were strict rules to govern their interconnections. By using a formal language, the manufacturer was able to reduce the potential for errors when the functional requirements were translated into the software requirements used by software engineers.[100]
- Related information. Additional requirements included physical characteristics, such as size and weight, and environmental requirements.

Many of the requirements for the A330/A340 EFCS were based on the Airbus A320 EFCS. The A320, which was certified in 1988, was the first civil aircraft to be equipped with a full-authority electrical flight control system.[101] When developing the A330/A340, one of the manufacturer's objectives was to reproduce the architecture and principles of the A320 as much as possible for the sake of commonality and design efficiency, taking into account the different performance objectives of the A330 and A340 in terms of the increased range and payload. In addition, the design of the A330/A340 took advantage of technological improvements.

The software used in the EFCS computers was classified as 'critical' in accordance with DO-178A, and therefore the required software level was Level 1.

2.4.4 System verification and validation

Initial activities

The first validation activities during the development of the A330/A340 EFCS included peer reviews of the requirements for the whole system and the associated equipment, together with the initial safety assessments (section 2.4.5). Where relevant, any identified problems led to changes in the requirements.

Engineering simulations (prior to building the equipment)

The manufacturer's target was to validate the system requirements at the earliest possible stage in order to minimise the problems associated with redesigning systems after they were built. To achieve this target, it developed simulation tools that enabled the EFCS design to be comprehensively tested prior to building the system.[102] Because the functional requirements were written in a formal language, they were able to be used to automatically develop the software code for these simulation tools.

---

[99] Spécification assistée par ordinateur (computer assisted specification).

[100] The use of a formal language also enabled much of the software to be automatically coded, thereby reducing the likelihood of coding errors.

[101] Airbus and its predecessor organisations had substantial experience in developing aircraft flight control systems over several decades, with the level of sophistication and the use of digital computers gradually increasing throughout that period.

[102] Simulation tools can enable products to be tested well outside of normal operating parameters, or for a larger range of operational situations to be tested in a shorter period of time.

The first simulation tool, OCAS[103], used a real-time computer to link the control laws with an aircraft movement simulation. It accepted inputs from simplified controls including a sidestick controller, and the results were displayed on a simplified primary flight display. Design engineers used the tool to assess the quality of the control laws and their effects.

Another simulation tool, OSIME[104], simulated the complete flight control system, including the computers, actuators and sensors. It linked the SAO definition of the whole system to the complete servo-control modes and to the simulation of aircraft movement. There were several hundred inputs into the system, including parameters such as weight, centre of gravity, wind and turbulence, and ADIRU inputs such as altitude and airspeed. This tool allowed the engineers to inject values of the parameters and examine their effects.

In addition to helping validate the design prior to building the software and associated equipment, the simulation tools allowed design changes to be efficiently developed and evaluated. They also enabled 'non-regression testing'[105] to be conducted early in the change process to help ensure that any proposed changes would not introduce new problems.

Testing and simulations (after building the equipment)

After the EFCS hardware and software was built, it was subjected to a series of further verification and validation activities, including:

- Test of the equipment on test benches. The flight control computers were tested by providing simulated inputs and observing the internal parameters. Test benches were also used to tune the servo-controls for each control surface.
- Tests on the 'iron bird' and flight simulator. The iron bird was a test platform with systems installed and powered as on an actual aircraft, and the flight simulator incorporated an aircraft flight deck and flight control computers. For some tests, the iron bird and the flight simulator were coupled.
- Flight tests. Several aircraft were fitted with comprehensive flight test instrumentation, recording more than 10,000 flight control parameters for later analysis.

The manufacturer conducted the tests and simulations according to specified test programs. If the behaviour of the system was not satisfactory, a problem report was raised, registered and investigated. If any of the testing identified a need to modify the system design, the resulting modification was again subject to safety assessment, simulations, testing and other verification and validation activities, in addition to non-regression testing to ensure that no new problems were introduced as a result of the change. The design was not 'frozen' until all of the design evaluation activities were completed.

---

[103] Outil de conception assistée de spécification.

[104] Outil de simulation multi-équipement.

[105] Non-regression testing (also known as regression testing) determines whether any changes made to a system have led to any problems with the existing system functionality.

2.4.5 Safety assessment activities

General process

The safety assessment activities conducted by manufacturers when developing aircraft systems occur in three phases[106]:

- Functional hazard assessment (FHA). The FHA identifies the failure conditions associated with a system that could have repercussions at the aircraft level.[107] More specifically, it identifies the functions of the system and the failure conditions for each of the functions[108], determines the adverse effects of each failure condition, and classifies the level of the effects (that is, catastrophic, hazardous, major, or minor). Based on these results, the FHA generates safety requirements in terms of the maximum allowed probability of the failure condition (for example, a hazardous failure condition should not occur at a probability higher than 'extremely remote', see section 2.3.3).
- Preliminary system safety assessment (PSSA). The PSSA evaluates the proposed system design and determines how failures within the existing design could lead to the failure conditions, and whether the FHA's probability-based safety requirements can be met by the proposed design. Additional requirements could be introduced to ensure the safety requirements will be met.
- System safety assessment (SSA). The SSA is a systematic examination of the system, its architecture and installation. It summarises all of the significant failure conditions and their effects on the aircraft, and is based on the FHA and PSSA. Whereas the PSSA is conducted to derive design requirements and determine whether the system design could reasonably be expected to meet the requirements, the SSA is conducted to demonstrate that the safety requirements have been met. Results of simulation and testing activities conducted for verification and validation purposes are also included in the SSA where relevant.

A key part of the PSSA is the use of fault tree analysis[109] or other similar top-down methods for identifying the failure scenarios, or combinations of failures and/or other factors, that could lead to each of the failure conditions. Where quantitative estimates of failures are derived, a fault tree analysis is also used to determine whether the relevant probability-based safety requirements can be met. In addition to a top-down approach, a bottom-up approach is also used to determine how equipment failures could potentially lead to the failure conditions of concern. This

---

[106] A detailed description of the three phases is provided in ARP4761, released in December 1995 (section 2.6.2).

[107] A FHA is firstly done at the aircraft level. Based on this analysis, decisions are made regarding the required aircraft systems. The results of the aircraft-level FHA flow down to the system-level FHA.

[108] In a FHA, failure conditions were generally described in terms of basic ways in which the function may not be adequately performed. Typical examples were loss of function, undetected loss of function, function not performed when required, or malfunction (function not performed correctly).

[109] Fault tree analysis is a very widely used top-down method for determining system reliability in many industries. It involves reasoning backwards from a specific event (known as a 'top event') to the combinations of factors that can lead to that event, and representing these factors in a graphical format. It often involves determining the overall probabilities that the top event will occur. Further details are provided in NASA (2002) and ARP 4761.

generally involves conducting a failure mode and effect analysis (FMEA)[110] on each item of equipment and determining equipment failure rates. Further discussion of PSSA methods such as fault tree analysis and FMEA is provided in section 2.6.3.

Safety assessment activities are generally conducted by different engineers (safety analysts) than those who design the system (design engineers).

Safety assessment for the A330/A340 EFCS

Overall, the aircraft manufacturer's methodology for conducting safety analysis activities was consistent with the guidance provided in the European ACJ to JAR 25.1309 and the FAA's AC25.1309.

In terms of the FHA, the aircraft manufacturer advised that the identification and classification of the failure conditions for the A330/A340 EFCS was based on engineering analysis, knowledge that it had from its previous experience, and the FMEAs provided by the manufacturers of related equipment. The classifications were based on the effects of the failure condition on the system as well as other factors such as handling qualities, aircraft performance, and aerodynamic loads on the aircraft structure. The FHA documentation included the description of the failure condition (including its repercussion or effect on the aircraft), the classification (or level of effect), and the rationale used to justify the classification.

The range of EFCS functions considered during the FHA included the processing of ADIRU parameters. For each ADIRU parameter used by the FCPCs, the FHA generated a list of failure conditions. The failure conditions related to the FCPCs' processing AOA data are discussed in section 2.5.3.

With regard to the PSSA, the aircraft manufacturer advised that its identification of the failure scenarios leading to the failure conditions and the determination of their probability levels were also based on engineering analysis, knowledge that it had from its previous experience, and the FMEAs provided by the manufacturers of the related equipment. It used qualitative methods to assess design problems, environmental hazards and human factors aspects, and both qualitative and quantitative methods for assessing physical or hardware failures. The depth of the required assessment depended on the classification of the failure condition (that is, more detailed analysis was conducted for catastrophic failure conditions than for minor or major failure conditions).

The EFCS PSSA for the first A330/A340 model was finalised in June 1991, and the results of relevant PSSA activities for the FCPC algorithm for processing AOA data are discussed in section 2.5.3. The FMEA for the LTN-101 was finalised in September 1992 and the results are discussed in section 3.8.

The EFCS SSA was finalised in November 1992. The aircraft manufacturer reported that it verified that the safety requirements were met, and ensured that all the necessary design features were incorporated into the system architecture,

---

[110] FMEA is a very widely used bottom-up method for determining system reliability in many industries. It involves reasoning forwards from a specific failure mode to the effects of the failure mode. For each component of interest, it involves identifying the function(s) of the component, the ways in which the component can fail (or failure modes), and the effects of each failure mode on the item of equipment or the system. It often involves determining failure rates for each failure mode. Further details are provided in the US Military Standard MIL-STD-1629A (*Procedures for performing a failure mode, effects and criticality analysis*) and ARP 4761.

equipment software and aircraft installation. The final assessment contained over 2,000 pages of documentation, and data from the SSA was included in the certification dossier that was provided to the certifying authority (that is, the DGAC).

2.5 Development of the algorithm for processing AOA

2.5.1 Preliminary A330/A340 FCPC algorithm

The FCPC algorithm for processing AOA data was a small but important element of the overall EFCS. The general safety objectives for the algorithm (and the rest of the EFCS) were based on JAR 25.671, JAR 25.1309 and associated guidance material, and included the safety objective that no single failure should result in a catastrophic failure condition. The aircraft manufacturer advised that other requirements for the algorithm were that the:

- AOA values from the three ADIRUs were to be acquired by each channel (COM and MON) of each FCPC
- aircraft behaviour must be acceptable in the case of a runaway of one of the three AOA values.

Consistent with the aircraft manufacturer's objective (section 2.4.3), the preliminary design of the FCPC algorithm was the same as that certified for the A320 in 1988. It did not include a memorisation period and had several other differences to the algorithm that was ultimately used on the production A330/A340 aircraft (described in sections 2.1.4 and 2.1.5).

The preliminary FCPC algorithm for processing AOA data was subject to the manufacturer's processes for generating requirements, verification and validation, including safety assessment. None of these activities, up until flight testing, identified any need to change the algorithm's design.

In December 1991, during a test flight on an A340, a problem was identified with the operation of the preliminary algorithm. The algorithm did not effectively manage a specific situation where AOA 2 and AOA 3 on one side of the aircraft were temporarily incorrect and AOA 1 on the other side of the aircraft was correct, resulting in ADR 1 being rejected.

2.5.2 Redesign of the preliminary algorithm

After reviewing the problem that was identified during flight testing, the aircraft manufacturer redesigned the FCPC algorithm. This led to the final algorithm for processing AOA data, which included the 1.2-second memorisation period and several other changes (described in sections 2.1.4 and 2.1.5).

The manufacturer advised that the 1.2-second duration of the memorisation period was considered to be the maximum time period that the aircraft could use a memorised AOA value in dynamic manoeuvres. Accordingly, the FCPCs required a current AOA value to be obtained at the end of the memorisation period, and the final algorithm was designed to ensure that a second memorisation period would not be triggered without a new value first being obtained.

Consistent with the aircraft manufacturer's normal system development processes, the design change was the subject of additional verification and validation activities, including safety assessment, simulations and testing.

2.5.3 Safety assessment of the FCPC algorithm

The SSA for the EFCS identified the following types of failure conditions associated with AOA values:

A. Loss of information (or false failure detection)

B. Loss of redundancy (2 ADR off) or equivalent false failure detection

C. Use of erroneous value, greater than the actual one

D. Use of erroneous value, lower than the actual one

In terms of the third failure condition, three potential consequences or effects were identified. One of these was defined as:

> Runaway of the AOA protection: pitch-down command on the elevators that cannot be counteracted by the crew.

This consequence was essentially what occurred during the 7 October 2008 occurrence flight.[111] Given this potential consequence, the SSA classified the failure condition as 'hazardous' based on engineering judgement. Accordingly the safety objective derived from the relevant regulatory requirements was that the probability of the failure condition should be 'extremely remote'. The manufacturer advised that its safety objective for the failure condition was actually 'extremely improbable', which was more stringent.

The SSA for the final algorithm identified one scenario that could lead to the failure condition involving a pitch-down command due to an erroneous $AOA_{FCPC\ input}$ value. This scenario involved incorrect AOA data being simultaneously provided by two ADRs. The assessed probability of this failure combination, based on the FMEAs provided by the ADIRU manufacturers, was calculated to be less than the safety objective of 10^{-9} per flight hour. The manufacturer advised that this failure scenario had previously been identified for the initial (A320) algorithm design and also included in the PSSA for the A330/A340 EFCS.

The manufacturer reported that when it evaluated the final algorithm during the redesign process in 1991-1992, it did not identify any other failure scenarios that could lead to the failure condition. It stated that it considered failure scenarios involving a single ADIRU, such as a runaway AOA and AOA spikes, during the assessment activities, and had concluded that the algorithm was robust to any single ADIRU failure.

In summary, the SSA for the final FCPC algorithm for processing AOA did not identify a failure scenario involving multiple spikes in AOA 1 (or AOA 2) that were 1.2-seconds apart, and which subsequently occurred on 7 October 2008.

---

[111] Although the anti pitch-up compensation was not specifically referred to in the SSA, the manufacturer advised that both high AOA protection and anti pitch-up compensation were triggered by the same condition and both produced similar effects. Therefore, to include one mechanism implied that the other was also included. It also noted that the classification of the effect level would not change if both mechanisms were specifically mentioned.

2.5.4 Factors associated with the identification of failure scenarios

In terms of reasons why the failure scenario was not identified, the aircraft manufacturer noted that the FMEA supplied by the LTN-101 ADIRU manufacturer did not identify the type of ADIRU failure mode that occurred on the 7 October 2008 flight (that is, multiple, undetected spikes in AOA and other ADR parameters) (see also section 3.8).

The manufacturer also stated that, prior to and during the design process in 1991-1992, it was aware that AOA spikes could occur on many flights, but in its experience only a very small number of spikes (if any) occurred on any particular flight. Prior to the 7 October 2008 occurrence, it was not aware of any events involving the ADIRU data-spike failure mode or similar ADIRU behaviour. Between 1992 and 2009, A330/A340 aircraft conducted over 28 million flying hours and the 7 October 2008 occurrence was the only known example where multiple AOA spikes from one ADIRU had led to an undesired pitch-down command.

In addition, the manufacturer advised that simulation and testing activities conducted during verification and validation could not identify all the failure modes for a complex system such as the EFCS. It was also not practical to simulate or test the effects of all possible types and combinations of inputs into such a complex system.

The manufacturer was asked to provide information on the risk controls, such as procedures and training, that it had in place to help maximise the likelihood that failure scenarios would be detected during the system development process. It advised that, at the time of the A330/A340 design, the development process was based on conception reviews, its own experience, FMEAs provided by equipment suppliers, the SSA process, various tests and simulations, and discussion with airworthiness authorities. The manufacturer issued a detailed guidance document for its system designers in 1996 (section 2.6.2).

2.5.5 Involvement of regulatory authorities

EASA advised that, for a design change of the type that occurred to the FCPC algorithm for processing AOA following the initial flight testing in December 1991, a certifying authority would not generally look closely at the change unless it arose due to a known or high profile problem. EASA and the aircraft manufacturer also advised that it was unlikely that there was any discussion regarding the design change at the time it occurred. In general, only the final design would be presented to the certifying authority.

Following the 7 October 2008 occurrence, there were regular meetings between the manufacturer and the regulator regarding the EFCS design. During these discussions, the classification of the failure condition involving a pitch-down command due to an incorrect $AOA_{FCPC\ input}$ was confirmed to be ‘hazardous’ (and not ‘catastrophic’).

2.5.6 Scope of the design limitation

The design limitation with the flight control system algorithm for processing AOA data only applied to the A330/A340 aircraft types. Other Airbus aircraft (including the A320 and A380) would not have been significantly affected by the ADIRU

data-spike failure mode because they used different algorithms for processing AOA, and these algorithms did not include a memorisation period.

Following the 7 October 2008 occurrence, the aircraft manufacturer conducted a detailed review of the A330/A340 FCPC algorithms for processing other ADIRU flight data parameters. Some limitations were identified with the algorithms for a number of other parameters, in terms of their ability to handle a multiple data-spike situation. None of these limitations were significant, and all related to the ability to manage situations when one of the three ADIRUs was already unavailable.

2.6 Developments in the design and assessment of safety-critical systems

The development of safety-critical systems is an evolving field and there have been many significant enhancements in the field, after the A330/A340 was certified in 1992, to help minimise the risk associated with system design errors. These include industry standards and guidelines for system development processes, and research into different ways of conducting safety assessment activities.

The investigation reviewed these developments to determine their applicability to preventing future design errors such as the limitation with the FCPC algorithm (and, to some extent, the ADIRU data-spike failure mode discussed in Part 3). However, the issues discussed in this section apply to the design of all complex systems, rather than just to the design processes of any specific organisation.

2.6.1 Role of software in safety-critical systems

Prior to reviewing recent developments, it is useful to consider some general aspects associated with the use of software in safety-critical systems. Software by itself does not necessarily make a system more complex, and some systems with minimal software can also contain significant complexity. However, due to its flexibility and functionality, software has become widely used in aircraft systems, and its increased use has added to the overall complexity of many system designs. Rushby (1995) noted:

> Complexity is a source of design faults... Design faults can occur in any system, independently of the technologies used in its construction... but, because design faults are often due to a failure to anticipate certain interactions among the components of the system, or between the system and the environment, they become more likely as the number and complexity of possible behaviours and interactions increases.
>
> Individual software components perform complex functions in modern systems, and collectively they provide the focus for the interactions among all parts of the system, and between the system and its environment and operators. Furthermore, software, because of its mutability, is also the target for most of the changes that are generated in requirements and constraints as the overall design of a system evolves. Thus, software carries the burden of overall system complexity and volatility, and it is to be expected that design faults will most commonly be expressed in software.

Mechanical devices generally fail in known and predictable ways, whereas the manifestation of software problems is more difficult to predict. Due to the complexity of some modern, software-based systems, many experts have stated that

it is very difficult if not impossible to identify all the design errors or possible ways in which the system may not perform as desired. For example, a UK Health and Safety Commission (1998) report stated:

> Computer systems are vulnerable because they almost invariably contain design faults in their software (and perhaps in their hardware) that are triggered when the computer system receives appropriate inputs. Many of these faults will have been present from inception, and others will have been introduced during any changes that have taken place throughout the system lifetime. The reality is that even programs of surprisingly modest size and complexity must be assumed to contain design faults. It is the responsibility of the designer, having done whatever is possible to minimise the number of residual faults, to try to ensure that any remaining ones do not have an unacceptable effect upon other systems with which the computer system interacts: in particular, that they do not compromise the safety of the wider system.

There is widespread agreement that when software problems contribute to accidents and serious incidents, the problems are usually due to flaws in the design requirements. In other words, the software worked according to its design, but there was a problem with the design itself rather than the way the software code was written. For aerospace systems, the problems generally involve incompleteness in the requirements, particular in terms of the interactions between systems and the inability of the software to handle certain states or conditions (Lutz 1993, Leveson 2004a).

There has been no systematic evaluation of the contribution of software design problems leading to aircraft flight control system occurrences. Although problems with software requirements have previously contributed to such occurrences (for example, see section 1.16.1), the investigation found no salient evidence to suggest that such systems have not, to date, generally performed at appropriate safety levels.

The aircraft manufacturer noted that the accident rate for modern aircraft with more complex system designs (such as the A320, A330 and A340) is lower than that for previous generations of aircraft. It also stated that, even though systems continue to develop in line with technological advances, SSA and other system development processes also continue to develop and become more sophisticated.

2.6.2 Developments in industry standards and guidance

Several industry standards and guidance documents were issued in the 1990s for developing complex aircraft systems and meeting the requirements of JAR/FAR 25.1309. Airbus and/or its related organisations were involved in the development of these documents.[112]

Design objective 178B

A major revision of DO-178A (section 2.3.5) was issued in December 1992. The new version (DO-178B) was developed as a result of the rapid advances in software

[112] Many other standards have been issued in recent years for the development of safety-critical systems or software in various industries, but only those specifically applicable to aircraft systems have been included in this report.

technology and differing interpretations being applied to some areas of the previous version.

DO-178B revised the criticality and software levels used in DO178A. Instead of criticality categories, DO178B used failure condition categories consistent with the ACJ to JAR 25.1309; that is, catastrophic, hazardous (or severe-major), major, minor, and no effect. The associated software levels were A, B, C, D and E. The required software level was based on the contribution of the software to the potential failure condition (for example, Level A software was required for a catastrophic failure condition).[113] The use of five levels enabled the software development processes to be more finely tuned to the potential impact of the system on safety.

The new version also provided objectives for software life cycle processes (such as software planning, software development, and software verification), descriptions of activities and design considerations for achieving the objectives, and descriptions of the evidence that indicated that the objectives were satisfied. It also provided more detailed guidance in some areas than DO-178A. For example, DO-178B included a list of robustness test cases to demonstrate the ability of software to respond to abnormal inputs and conditions. The robustness test cases included guidance that 'the possible failure modes of the incoming data should be determined, especially complex, digital data strings from an external system'. No further guidance was provided on the types of failure modes that should be considered.

A third revision of DO-178, DO-178C, was scheduled to be issued during 2011. It was expected to provide further enhancement and clarification of the core processes in DO-178B, as well as to address topics dealing with advances in complex avionics software development. These topics include the use of formal methods and model-based development (section 2.6.3), and the use of object-oriented software.[114]

In 2000, the RTCA issued DO-254 (*Design Assurance Guidance for Airborne Electronic Hardware*) as a complementary document to DO-178 for hardware components. Its purpose was to provide 'design assurance guidance for the development of airborne electronic hardware such that it safely performs its intended function, in its specified environments.' DO-254 defined five hardware design assurance levels, analogous to those in DO-178B.

Aerospace recommended practice 4754

In November 1996, the Society of Automotive Engineers (SAE)[115] issued aerospace recommended practice (ARP) 4754 (*Certification considerations for highly-integrated or complex aircraft systems*). The ARP was 'intended to provide designers, manufacturers, installers, and certification authorities a common international basis for demonstrating compliance with airworthiness requirements

---

[113] Revisions of the FCPC software were developed as Level A software.

[114] Object-oriented programming is one of many software programming paradigms and is increasingly used in the development of aircraft software. It groups information into objects with associated properties and functions.

[115] The SAE developed standards for the design of road, marine and aircraft vehicles. These standards were often adopted by regulatory agencies as design requirements.

applicable to highly-integrated or complex systems'. It provided a high-level, integrated overview of processes such as requirements development, safety assessment, requirements validation, implementation verification, configuration management, and process assurance.

ARP4754 stated that highly-integrated and complex systems presented greater opportunities for errors in requirements development and design. It also noted that it was generally not practical to develop a finite test suite for such systems that would conclusively demonstrate that there were no residual development errors. Consequently development assurance was also required.[116] It outlined five development assurance levels, which were consistent with the software levels outlined in DO-178B.

In addition to providing guidance on validating the correctness and completeness of the requirements, ARP4754 provided guidance on validating assumptions. This process involved ensuring that assumptions were explicitly stated, appropriately disseminated, and justified by supporting data.

The ARP noted that several methods may be needed to support validation, including traceability, analysis, modelling, testing, service experience with similar systems, and engineering judgement. In terms of traceability, it stated that each requirement should be traceable to a parent requirement or a specific design decision, and each assumption should be traceable to a standard, practice, analysis or test. The ARP also stated that the verification of a system may be supported by inspection and review, analysis, testing and service experience.

A new version of the ARP (ARP 4754A, *Guidelines for the development of civil aircraft and systems*) was issued in late 2010. The new version took account of developments within the industry over the preceding period, and strengthened the emphasis on the development assurance concept.

Aerospace recommended practice 4761

In December 1996, SAE issued ARP 4761 (G*uidelines and methods for conducting the safety assessment process on civil airborne systems and equipment*). The guidelines and methods described in ARP4761 provided a 'systematic means, but not the only means' to showing compliance with JAR/FAR 25.1309. The guidelines applied to both the hardware and software of a system.

The overview section of the ARP stated:

> The safety assessment process includes requirements generation and verification which supports the aircraft development activities. The process provides a methodology to evaluate aircraft functions and the design of systems performing these functions to determine that the associated hazards have been properly addressed. The safety assessment process is qualitative and can be quantitative.

---

[116] ARP4754 defined development assurance as 'a process involving specific planned and systematic actions that together provide confidence that errors or omissions in requirements or design have been identified and corrected to the degree that the system, as implemented, satisfies applicable certification requirements.'

> The safety assessment process should be planned and managed to provide the necessary assurance that all relevant failure conditions have been identified and that all significant combinations of failures which could cause those failure conditions have been considered.

ARP 4761 provided a detailed description of the safety assessment process, based around the three main activities of functional hazard assessment (FHA), preliminary system safety assessment (PSSA), and system safety assessment (SSA).

Developments in Airbus guidance materials

In 1996, the aircraft manufacturer issued ABD0200 (*Guidelines and requirements for the system designers*) as the reference document for its system designers. It advised that the manual had a modular structure so that the 'continued evolution of common Airbus practices could be captured and incorporated'.

In addition to ABD0200, the manufacturer developed a higher-level manual AP2288 (*Requirements for system and equipment development*) to define its system development process. This manual described each process by providing details on objectives, main 'actors', activities, and related inputs and outputs.

The manufacturer advised that these two manuals contained a large volume of guidance material and incorporated the principles of ARP4754, ARP4761 and other relevant industry standards. It also noted that that the manuals were updated in response to in-service feedback, such as the results of accident and incident investigations, or regulatory changes.

Additional guidance material

Standards such as DO-178 and ARP4754 provided general guidance for developing safety-critical software; they were not designed to provide detailed guidance or checklists of specific issues to consider when developing or reviewing requirements for a safety-critical system. However, a number of textbooks and guidance manuals have been published that do provide more detailed guidance[117], and several institutions now provide training courses for design engineers and safety analysts.

Checklists are often used when developing and reviewing system requirements, and research has shown that checklists focussing on safety-related aspects can increase the chances of detecting safety-related design problems (Lutz, 1996). However, given the wide range and complexity of system designs, it is unreasonable to expect that every specific, potential problem with every type of software design could be specified in the form of checklists or guidance material.

A review of a sample of guidance manuals and checklists did not identify any specific guidance that was directly applicable to the design limitation associated with the A330/A340 FCPC algorithm for processing AOA data. Although many referred to the importance of checking input values, and referred to intermittent faults, none appeared to specifically refer to a multiple data-spike situation.

---

[117] Examples include NASA (2004), Joint Software System Safety Committee of the US (1999), and Storey (1996).

Ongoing developments

A range of research and development activities have been taking place in the field of developing safety-critical systems to address various challenges and look at different ways of providing safety assurance. A detailed review of these activities is beyond the scope of this report, but some examples are worth noting.

NASA recently requested responses from the aviation industry on the topic of 'verification and validation of flight-critical systems' and approaches to improve these activities. The summary of responses document (Graves and Jacobsen 2010) noted that emerging challenges associated with system verification and validation included:

> increasing system complexity; an exponential increase in software requirements; an increase in tests required using current V&V [verification and validation] methods, safety, cost and schedule impacts associated with V&V; emergence of distributed architectures & trends in microprocessor technologies...

The document also stated:

> The current aviation system has an enviable safety record; however, advances in technology are placing an increasing strain on our ability to assure the integrity of new and anticipated systems. Additionally, there is a perception that current approaches for the assurance of complex flight-critical systems impose a barrier to innovation.

In addition, several experts have discussed the suitability of different types of evidence for demonstrating the safety of a system or its software. For example, Hawkins and Kelley (2010) have noted that although various standards require specific types of evidence, there has been little guidance to date for determining how these types of evidence are sufficient or trustworthy in a particular context. Various options have been proposed or are being developed to address this issue.

2.6.3 Developments in safety assessment methods

Recognition of limitations using traditional methods

As previously noted, the development of a safety-critical system needs to provide assurance, through many different processes, that the resulting system meets appropriate safety requirements. An SSA is a key element of this development assurance, and ARP 4761 provided significant industry guidance for safety assessment activities. However, there has been recognition that the methods advocated in the ARP, and used in the industry for many years prior to the ARP, have some limitations when applied to complex systems and consequently there has been significant interest in developing new approaches to safety assessment. Prior to discussing some of these approaches, it is useful to review some of the limitations of the traditional methods.

Fault tree analysis, FMEA and similar methods have been widely recommended and widely used for identifying the scenarios that can lead to a failure condition. These methods were originally developed for evaluating hardware-based systems. They have also been adapted and widely used for systems that are software-based, but many experts have argued that they are not necessarily well suited to this purpose due to the inherent complexity of the systems being designed.

For example, experts have argued that for large, complex systems involving software it is difficult for design engineers and safety analysts to comprehend all the ways in which the system could respond to different events (Rushby 1995). More specifically, Bozzano et al. (2003) stated:

> One of the most challenging issues in system development today is to take into consideration, during development, all possible failure modes of a system and to ensure safe operation of a system under all conditions. Current informal methodologies, like manual fault tree analysis (FTA) and failure mode and effect analysis (FMEA)..., that rely on the ability of the safety engineer to understand and to foresee the system behaviour are not ideal when dealing with highly complex systems, due to the difficulty in understanding the system under development and in anticipating all its possible behaviours.

A significant amount of guidance material is available on how to develop fault trees and use fault trees to calculate failure probabilities (Lisagor et al. 2010).[118] However, very little of this guidance discusses exactly how to identify failure scenarios. Leveson (1995) stated that much of the focus of fault tree analysis is directed towards generating probabilities of failure, whereas most of the errors in hazard analysis are due 'to the failure to foresee all the ways in which the hazard could occur'.

Furthermore, Leveson (2009a) noted that traditional safety analysis techniques like fault tree analysis provide little guidance to analysts about the actual analysis process, and that the quality of the resulting analyses for complex systems varies significantly depending on the analysts' skill. Redmill (2002) also noted that the construction of fault trees involves a significant degree of subjectivity and variability between users, and Manion (2007) concluded that a variety of biases can influence the performance of each step in the process of developing a fault tree. Similarly, Papadopoulos et al. (2001) stated:

> ...the safety case usually fails to offer a coherent and complete picture of the ways in which low-level component failures contribute to hazardous malfunctions of the system. Although fault trees are built for this purpose, the traditional process of constructing these fault trees relies heavily on expert knowledge, and lacks a systematic or structured algorithm which the analyst can apply on a system model in order to derive the tree. In the context of a complex system this process becomes tedious, time consuming and error prone...

FMEA are also frequently described as being time consuming and tedious to conduct for complex items of equipment, and consequently they are often not completed and able to be used until late in the PSSA process. In addition, an FMEA generally only deals with single failures rather than more complex failure modes involving multiple failures (Leveson, 1995).

---

[118] Examples of standards providing guidance on fault tree analysis include ARP4761 and the *Fault tree handbook with aerospace applications* (NASA, 2002).

To conduct an effective fault tree analysis or FMEA requires safety analysts to have a very good understanding of the system and all its components. Experts have noted that design engineers and safety analysts generally work with different views of the system's design. This can lead to potential problems with consistency and completeness of understanding between the two groups. A recent US National Aeronautics and Space Administration (NASA) research report (Joshi et al. 2006) noted:

> Safety engineers traditionally perform analysis, such as fault tree analysis, based on information synthesized from several sources, including informal design models and requirements documents. Unfortunately, these analyses are highly subjective and dependent on the skill of the engineer. Fault trees are one of the most common techniques used by safety engineers, yet different safety engineers will often produce fault trees for the same system that differ in substantive ways. The final fault tree is often produced only through a process of review and consensus building between the system and safety engineers. Even after a consensus is reached, it is unlikely that the analysis results will be complete, consistent, and error free due in part to the informal models used as the basis of the analysis. In fact, the lack of precise models of the system architecture and its failure modes often forces the safety analysts to devote much of their effort to gathering information about the system architecture and system behavior and embedding this information in the safety artifacts such as the fault trees.

In summary, the task of identifying scenarios that lead to failure conditions requires expertise and a detailed knowledge of the proposed system design. However, as system designs get more complex, then traditional safety assessment activities become inefficient and difficult to use successfully.

Formal methods and model-based safety analysis

The term 'formal methods' refers to the use of mathematical techniques in the design and evaluation of complex systems. NASA (1995) stated:

> Formal Methods (FM) consist of a set of techniques and tools based on mathematical modeling and formal logic that are used to specify and verify requirements and designs for computer systems and software. The use of FM on a project can assume various forms, ranging from occasional mathematical notation embedded in English specifications, to fully formal specifications using specification languages with a precise semantics. At their most rigorous, FM involve computer-assisted proofs of key properties regarding the behavior of the system.

The use of formal methods in the development of complex systems has been growing for over 20 years. For aviation applications, the use of formal methods was initially focused on the development of requirements. For example, the development of the functional requirements for the A320 and A330/A340 were based on formal methods, and these were used to develop models that could be used to conduct simulations for validating the EFCS system design (section 2.4.3).

In recent years, there has been considerable interest in using formal methods to create a system model, and then using the model to automate some of the safety assessment tasks. Proponents have argued that this approach enables the design engineers and safety analysts to work with the same view of the system, and to provide better assurance that design errors will be identified in a more efficient manner (Pumphrey 2001, Joshi et al. 2006).

One type of automated safety assessment involves fully defining the failure logic of each component's inputs and outputs, and then integrating the components into a full system model. The system model can then be used to automatically generate fault trees of the system (Papadopoulos et al. 2001).

A second type of automated safety analysis involves using static analysis tools such as 'model checkers' and 'theorem provers'. This approach involves building a model of the system with various failure modes included, and then applying the tools to the model to automatically generate a list of the failure modes which violate a specific, formally-defined requirement.

In 2001 to 2003, a group of European aircraft manufacturers and research institutions, including Airbus, conducted a project titled 'Enhanced Safety Assessment of Complex Systems' (ESACS) to examine the utility of model-based safety analysis activities (Bozanno et al. 2003). ESACS was followed by another project conducted by the same organisations, titled 'Improvement of Safety Activities on Aeronautical Complex Systems' (ISAAC), to further expand the scope and maturity of the methodology developed by ESACS (Akerland et al. 2006). Similar projects have also been facilitated by NASA (Joshi et al. 2006, Tribble et al. 2004).[119]

Automated safety analysis projects have reported some promising results, and they have started being accepted as part of the basis for the certification of new systems (Akerland et al. 2006). However, there are still limitations with the approach. For example, it relies heavily on having an accurate model of the system and the environment in which it operates, and uses this model for all the associated analyses. As it is not computationally possible to fully model a complex system, analysts therefore need to make assumptions and decisions about what to include in the model. Any limitations in the model will have an influence on all the derived safety analyses, and there is limited guidance available to date to best determine how to ensure the system model is adequate (Lisagor et al. 2010).

Lisagor et al. (2010) also noted that the outputs of some automated analyses can be 'unmanageably' large and difficult to interpret, and that there are inherent dangers with trying to understand the results as being equivalent to those from traditional analysis methods when they are actually based on quite different processes. In addition, the projects to date have focused on simpler types of failure modes, such as discrete and permanent faults rather than transient or intermittent faults or the timing-related aspects of faults (Lisagor et al. 2006; Tribble et al. 2004).

Alternative methods of safety assessment

Leveson has proposed that traditional safety assessment techniques can be adapted to some extent to handle new technology, they are not that well suited for this purpose as they are based on an inappropriate model of accident causation which focuses primarily on hardware faults and failures. She has proposed a different type of model derived from systems theory called 'system-theoretic accident modelling and processes' (STAMP) (Leveson 2004a, 2009a).

---

[119] The NASA request for industry responses on the topic of 'verification and validation of flight-critical systems' (Graves and Jacobsen, 2010) contained many responses advocating modelling and formal methods.

According to this model, accidents are conceived as resulting from inadequate control or inadequate enforcement of safety-related constraints on the design, development and operation of a system. Safety is viewed as a control problem, and accidents occur when component failures, external disturbances, and/or dysfunctional interactions between system components are not adequately managed. Other safety experts have proposed similar concepts (Saleh et al. 2010).

Based on her model, Leveson has proposed a hazard analysis approach known as system-theoretic process analysis (STPA) (Leveson 2009a). The first step of STPA is to identify the potential for inadequate control of the system that could lead to a hazardous state.[120] The second step is to determine how the unsafe control actions could occur. This step uses a top-down process to identify causal scenarios that could lead to hazardous control actions. It involves developing a (control) process diagram for the relevant components and considering a generic list of causal factors, based on a simple control model, to guide the process.

Some initial results have shown that STPA identifies more hazardous scenarios than a traditional approach that is based on fault tree analysis. However, it is a new method and its application and evaluation to date has been limited.

2.6.4 Research into the performance of engineers and analysts

As previously noted, modern safety-critical systems are complex and therefore their designs are difficult to evaluate. As there appears to be limited guidance material to conduct such evaluations, the conduct of peer reviews or SSAs to evaluate a system design will always heavily rely on the judgement and expertise of the design engineers and safety analysts involved.

A significant body of research has examined 'engineering judgement', but most of this research has focused on describing and exploring ways to improve how engineers assess the probability of relatively rare failure events (for example, Goossens et al 2008). Nevertheless, there has been some research that has looked at how experienced design engineers conduct their tasks. For example:

- Klein (1998) reported that design engineers, like many other experts, often made 'recognition-primed decisions'. That is, engineers generally considered the situation and compared it with their experience, and then selected a solution and mentally evaluated that solution rather than comparing all the available options.
- Cross (2004) reviewed several studies that studied expert designers. He noted that such designers were often solution-focused, and often persisted with an initial solution or used 'satisfying' behaviour (that is, they selected a solution that worked rather than looking for the best option available).
- Ahmed et al. (2003) compared novice and experienced design engineers and found that novices used a 'trial and error' approach, whereas experienced designers evaluated decisions prior to implementing them, and used their experience of previous projects to identify potential problems and solutions.

[120] The model notes that control actions can be hazardous in four ways: a control action required for safety is not provided or not followed; an unsafe control action is provided that leads to a hazard; a potentially safe control action is provided too late, too early, or out of sequence; and a safe control action is stopped too soon.

Researchers have stated that much more research needs to be done to fully understand how expert design engineers conduct their work, and how to effectively train novices to become experts. The research to date in this area has examined a range of design tasks, but limited research has focused on the task of reviewing system designs for safety-related problems. In addition, very little research has looked at the ways that safety analysts conduct activities such as FHA, fault tree analysis or FMEA.

As well as understanding how design engineers and safety analysts conduct their evaluations of system designs, it is important to understand the factors that can affect the performance of these personnel. The aviation industry has applied an enormous amount of effort into studying the human factors[121] issues for safety-critical personnel such as flight crew, air traffic controllers and aircraft maintenance personnel, and ensuring their work tasks are appropriately designed. However, the investigation found very little research that has examined the human factors issues affecting design engineers and safety analysts, or systematically examined the types of factors most likely to lead to design errors.

The large amount of human factors research conducted in other domains could provide insights into the types of issues that may be relevant for design evaluation tasks. Strigini (1996) noted that research into human performance limitations has identified many factors could adversely affect engineering judgements, including judgements such as whether a fault tree is complete. For example, research has shown that human decision-making ability can be affected by a range of different factors, such as task complexity, previous experiences, and the available information or cues about the decision or problem (Klein 1998). In addition, system design and evaluation is generally a team activity, with all design decisions being reviewed by others. Research has shown that detecting errors in other people's performance can vary widely depending on the context, and that errors of omission are relatively difficult to detect (Reason 1990).[122]

Although general human factors research can be applied to the work of design engineers and safety analysts, it would be more useful if research was specifically conducted for design evaluation tasks. A key principle of human factors is that the specific context in which work tasks are conducted needs to be well understood in order to determine how the design of the tasks, tools, training and guidance material can be improved to minimise the likelihood of errors or increase the ability to detect such errors.

---

[121] Human factors is the scientific discipline concerned with understanding the interactions among humans and other elements of a system, and applying theory, principles, data and methods to design in order to optimise safety, human well-being and overall system performance.

[122] Fischoff et al. (1978) and some other researchers (see Silvera 2005) have examined some biases that occur with fault tree analysis when information about some failure scenarios is omitted, but none of this research has been conducted with experienced safety analysts.

3 FACTUAL INFORMATION: AIR DATA INERTIAL REFERENCE UNITS

The recorded information from the 7 October 2008 flight showed that air data inertial reference unit (ADIRU) 1 (unit 4167) provided some incorrect flight data to the aircraft's other systems. This included data spikes on air data reference (ADR) parameters that were not flagged as invalid by the ADIRU. Several systems detected problems with ADIRU 1's performance, but the unit itself did not record any faults, and extensive testing of the unit did not identify any relevant problems. The same failure mode has occurred on two other occasions.

Given that the reasons for the failure were not identified by unit testing or recorded fault messages, a more detailed examination of the ADIRU design, recorded data, and in-service performance was conducted.

3.1 LTN-101 ADIRU history

3.1.1 Equipment specification

During the development of the A330/A340 aircraft, the aircraft manufacturer developed equipment specification *SPE A 3410 A002* for the air data and inertial reference system (ADIRS). The version applicable to the LTN-101 ADIRU design was issued in February 1990.

The specification outlined the technical requirements that the air data inertial reference system (ADIRS) and its major components (including the ADIRUs) needed to satisfy. These included requirements for the ADIRU's functions, operating conditions, data formats, performance levels, interface design, and general characteristics. It also outlined general safety objectives, general reliability objectives, built-in test equipment (BITE) requirements, safety analysis requirements, and a software criticality level. Where applicable, the specification referred to regulatory requirements, industry standards and Airbus directives.

The integration of the ADIRS with other systems was the responsibility of the aircraft manufacturer. Consequently, the ADIRS equipment specification did not state how the data transmitted from the ADIRUs would be used by other aircraft systems.

3.1.2 Overview of LTN-101 usage

The LTN-101 ADIRU was initially designed to meet the Airbus equipment specification for use on A330/A340 aircraft, and it was certified in 1993. In addition to being used on A330/A340 aircraft, it was also certified for use on the Airbus A320, the Saab 2000, and the Bombardier CRJ, Q400 and CL604 models.

The LTN-101 was one of two ADIRU models certified for the A330/A340. By the end of 2008, LTN-101 ADIRUs were fitted to about 44% of the A330/A340 fleet.

As of April 2010, LTN-101 units had accrued over 128 million hours of operation, including 113 million hours in Airbus aircraft. Over 8,000 units had been manufactured by July 2010, and the units were used by over 60 different operators.

3.2 LTN-101 ADIRU design

3.2.1 Design overview

As noted in section 1.6.4, the LTN-101 contained an air data reference (ADR) part and an inertial reference (IR) part. Overall, the two parts contained seven modules as shown Figure 36.

Figure 36: ADIRU functional architecture (simplified)

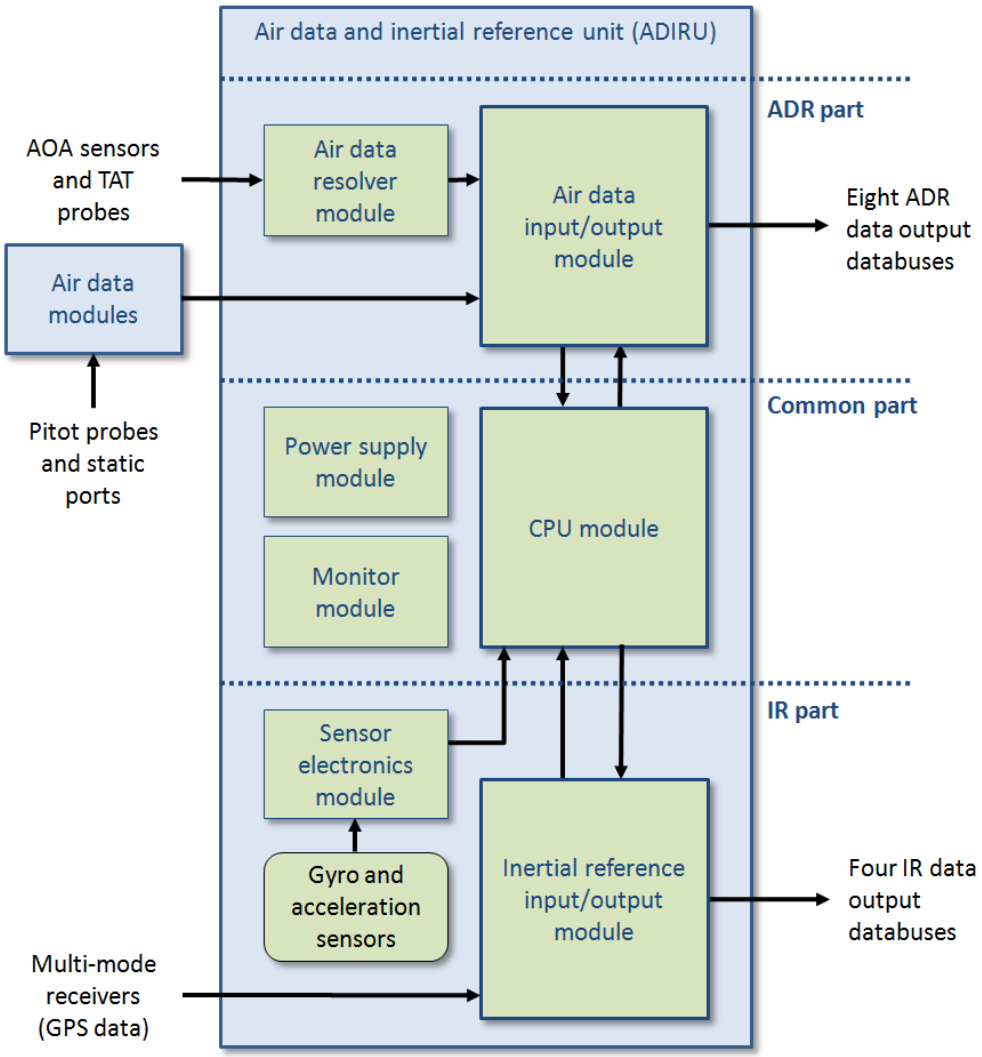

Further details of the modules are as follows:

- air data resolver module, which interfaced directly with the angle of attack (AOA) sensors and total air temperature (TAT) probes
- air data input/output module, which interfaced with the external air data modules (ADMs) and the air data resolver module, and sent the final processed ADR data to other systems
- sensor electronics module, which provided digital signal processing for the inertial instruments

- inertial reference input/output module, which interfaced with the sensor electronics module, received GPS data, and sent the final processed IR data to other systems
- central processor unit (CPU) module, which was the main controlling and processing module (section 3.2.2)
- monitor module, which provided some basic BITE logic and common discrete interfaces (such as warning, fail and on/off indications)
- power supply module, which converted the aircraft power to the various voltages used by the other modules.

The ADIRU also contained:

- a set of inertial instruments, comprising rotation (gyro) and acceleration sensors
- a motherboard to electrically connect the internal modules and other components
- a rear connector panel, including multiple separate output databuses (eight for ADR data and four for IR data).
- To achieve a high degree of fault tolerance, the ADIRU included BITE that provided monitoring and for isolation of faults within the unit. The BITE was designed so that a wide range of ADIRU functions were monitored and a high proportion of faults would be detected by the ADIRU itself. In addition to fault detection, the ADIRU's BITE recorded routine maintenance information. BITE records were stored in memory chips within the various modules. Further information on the operation of BITE functions of potential relevance to the data-spike failure mode is provided in section 3.7.

3.2.2 CPU module

The CPU[123] module performed the following functions:

- monitored and controlled the other modules
- conducted BITE tests, and managed and recorded BITE data
- received data from the unit's other modules, including flight data from the air data input/output module and the inertial reference input/output module
- processed the data (such as performing corrective calculations)
- packaged the data in a format suitable for outputting to other aircraft systems (section 3.3.2)
- sent the resulting data to the two input/output modules, which sent the data to the other systems.

ADR data was not passed through the inertial reference input/output module, and IR data was not passed through the air data input/output module. Consequently, the CPU module was the only area within the ADIRU with the potential to directly affect ADR data, IR data and BITE memory.[124]

---

[123] A CPU is the part of a digital computer system that executes instructions from a computer program or programs. It normally refers to a single integrated circuit.

[124] Other areas common to the ADR and IR parts were the power supply module, monitor module, rear connector panel and the motherboard. Although the power supply module, monitor module,

To perform its functions, the CPU module contained:

- four read-only memory (ROM) chips[125] that stored the ADIRU's software (known as the 'operational flight program')
- a CPU chip that executed the software instructions
- four random access memory (RAM) chips for general CPU use (CPU RAM)
- a 'companion' application-specific integrated circuit (ASIC)[126] that performed various interfacing and monitoring tasks, such as controlling access to the different areas of memory, controlling memory addressing, and buffering[127] and passing data and instructions between the CPU and memory chips
- a 'wait-state' RAM chip for storing parameters used by the ASIC for memory partitioning and access timing (see section 3.2.3 and *Wait-state tests* in section 3.4.9)
- a single chip of non-volatile[128] memory for storing BITE data.

3.2.3 Software

The Airbus ADIRS equipment specification required that the ADIRU software be designed to the highest level of stringency or thoroughness; that is Level 1 according to DO-178A (section 2.3.5).

A software program is a set of instructions that a computer processor executes to perform a certain task or set of tasks. In the case of the LTN-101 ADIRU, the CPU module contained software that carried out or controlled almost all of the ADIRU's tasks, including the input, processing and transmission of data, and BITE monitoring.

When the ADIRU was turned on, the software instructions were transferred from the ROM chips to the CPU chip for the ADIRU to operate. The CPU chip cached and loaded the relevant instructions as they were required in the program sequence. Sections of the program that needed high-rate access by the CPU chip were temporarily loaded in RAM, which had a faster access time.

To increase the reliability of the software, it was divided into 12 separate partitions according to criticality and functionality. Each partition used rigidly defined memory locations and could only access the system resources it required. If a software partition attempted to access an input/output device or memory location that was not assigned to it, an error message was generated.

When the ADIRU was initially turned on, and periodically afterwards, a cyclic redundancy check (CRC)[129] was performed on the software for each partition to ensure that it was not corrupted.

---

and motherboard had the potential to indirectly affect both ADR and IR data, subsequent analysis showed that there were internal processing problems within the CPU module (section 3.4).

[125] The chips were a type of electrically erasable programmable read-only memory (EEPROM) that is known as 'flash memory'.

[126] An integrated circuit, often also known as a microchip or chip, is an electronic circuit comprising numerous electronic elements manufactured as a single device.

[127] A buffer is a region of memory that temporarily stores data while it is being transferred.

[128] Non-volatile memory retains its stored data even when power is removed.

3.3 Examination of data-spike patterns

3.3.1 Background

As discussed in section 1.11, the ADIRU output data for all parameters that were recorded by the FDR and QAR contained a large number of data spikes on both ADR and IR parameters. In addition, the magnitude of the ADR spikes did not appear to be random. As noted in section 1.11.4, the AOA spikes only had three different values. Some other parameters also had spikes with repeating values. Accordingly, the investigation examined the data-spike values to determine the nature of any patterns.

Prior to discussing the patterns, it is important to understand the format in which the ADIRU outputted its data to other systems. This format is known as the ARINC 429[130] digital information transfer system. The FDR and QAR recorded data in a different format.

3.3.2 ARINC 429 digital information transfer system

The ARINC 429 system used 32-bit 'words' to represent information, with each bit containing a binary value (that is, a '0' or a '1'). The word was divided into five fields as shown in Figure 37.

Figure 37: ARINC 429 data word

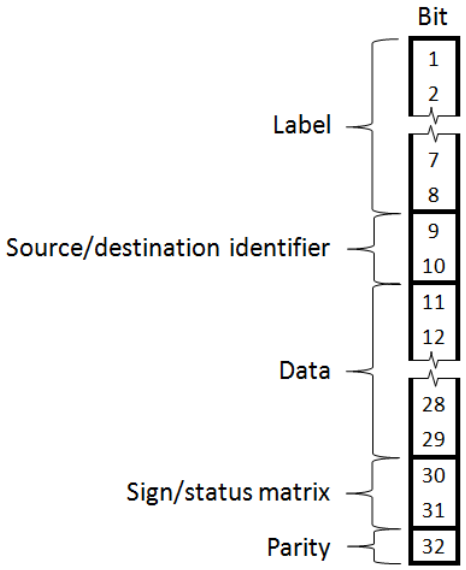

[129] A CRC is a function that takes a data stream as an input and calculates an output value. A CRC can be used as a 'checksum' to detect accidental alteration of data during transmission or storage.

[130] ARINC 429 is an aviation industry standard for the transfer of digital data between avionics system elements. ARINC (Aeronautical Radio Inc.) sponsors aviation industry committees that produce standards that allow avionics interchangeability.

The fields had the following functions:

- Label. An eight-bit field that uniquely identified the parameter. As the label was transmitted first, other systems could ignore the parameters they did not require.
- Source/destination identifier. A two-bit field that identified either the source of the transmission or the intended destination.
- Data. A 19-bit field that contained the output parameter's value.
  - Some of the output parameters did not use all of the 19 bits as they were not necessary for the required range and resolution. For example, AOA used 13 of the 19 bits (see Figure 38 for an example). Any unused bits were set to '0'.
  - Each data bit represented a value that was double that of the preceding bit, and all values were summed to produce the actual parameter value.
  - Some of the parameters also used one of the bits as a sign bit, indicating whether the value was positive or negative.
- Sign/status matrix (SSM). A two-bit field used to indicate the status (or validity) information of the transmitted word. The available values had the following meanings: failure warning, no computed data, functional test, and normal operation (section 1.6.8).
- Parity. A single bit that enabled a receiving system to check the integrity of a data word's transmission. The bit was set to '1' if the rest of the ARINC 429 data word contained an even number of '1's, or was set to '0' otherwise. The word was only accepted by the receiving system if it had an odd number of '1's.

The ARINC 429 word was constructed or packaged within the CPU module. However, the parity bit was calculated and added by either the air data input/output module or the inertial reference input/output module (depending on the databus over which the word would be eventually transmitted).

3.3.3 Recording system data format

Although the ADIRU outputted data using a 32-bit word with 19 bits available for the parameter value, the FDR (and QAR) used 12-bit words for the parameter value. When extra resolution was required, the FDR used two words. For example, altitude was split into a 'coarse' word (with 12 bits of data) and a 'fine' word (with 10 bits of data, one sign bit and one unused bit). Other FDR parameters did not require 12 bits; for example, AOA was recorded using nine bits (eight data bits and one sign bit), and Mach was recorded using eight bits (eight data bits and no sign bit).

The FDR recorded different ARINC 429 data bits for each parameter. For example, the FDR recorded altitude data bits 14 to 27 and the sign bit (29), AOA data bits 20 to 27 and the sign bit, and Mach data bits 19 to 26.

Figure 38: ARINC 429 word for an AOA value of 50.625°

Arinc 429 binary word	Bit number	Bit value	Value of each data bit	Example value
Label	1	1		This sequence corresponds to the label for corrected AOA 241 (octal).
	2	0		
	3	1		
	4	0		
	5	0		
	6	0		
	7	0		
	8	1		
SDI	9	1		Denotes ADIRU 1
	10	0		
Data	11	0	Data bits 11 to 16 were not used for AOA and were set to zero.	0.000
	12	0		0.000
	13	0		0.000
	14	0		0.000
	15	0		0.000
	16	0		0.000
	17	0	0.044	0.000
	18	0	0.088	0.000
	19	0	0.176	0.000
	20	0	0.352	0.000
	21	0	0.703	0.000
	22	0	1.406	0.000
	23	0	2.813	0.000
	24	1	5.625	5.625
	25	0	11.250	0.000
	26	0	22.500	0.000
	27	1	45.000	45.000
	28	0	90.000	0.000
	29	0	Sign bit	0 = positive
SSM	30	1		Normal operating (valid)
	31	1		
Parity	32	1		Odd number of 1's
			Total:	**50.625°**

The value of each data bit also varied for each parameter. For example, a '1' for data bit 27 meant 45° for AOA and 32,768 ft for altitude. Figure 39 and Figure 40 show how the data fields for AOA, Mach, and altitude represented real values in the ARINC 429 format. These figures also show the range of data bits used for each parameter in the ARINC 429 format, and as recorded on the FDR. The red-highlighted figures are discussed in the following sections.

Figure 39: AOA and altitude bit mapping

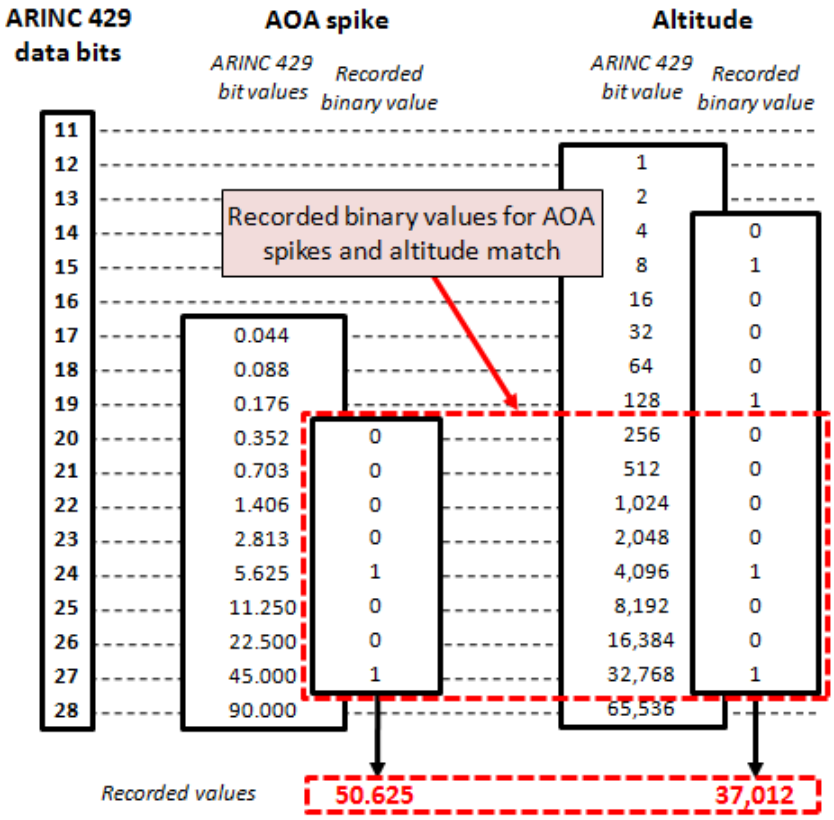

Figure 40: Mach and altitude bit mapping

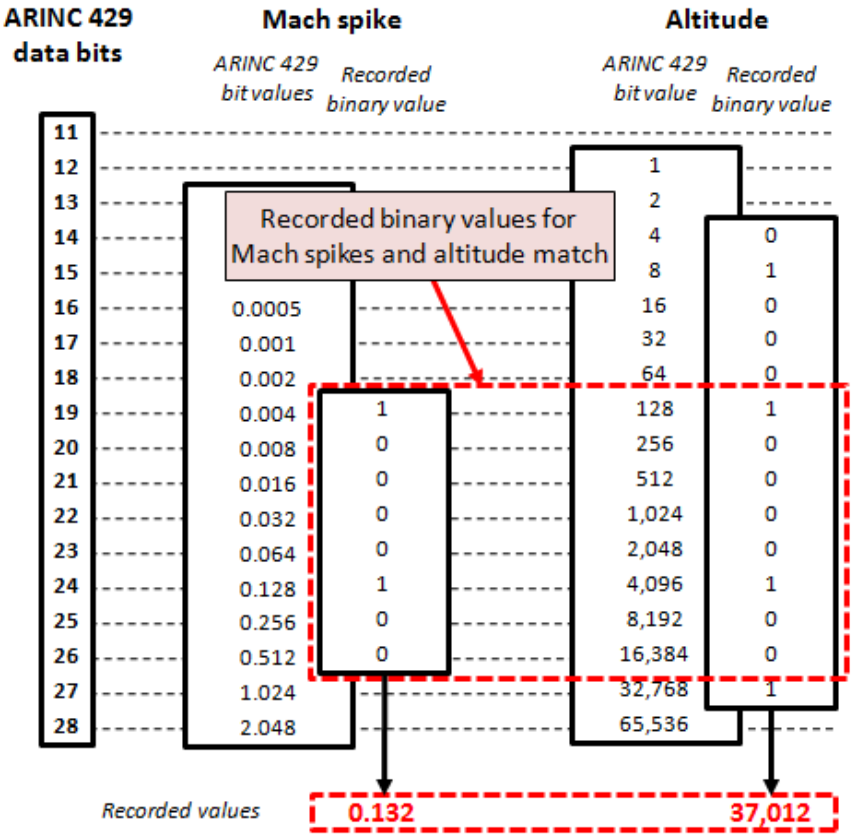

3.3.4 Data-spike patterns for the 7 October 2008 flight

Mach versus altitude

A review of the data-spike values for Mach indicated that many of them appeared to vary with altitude. A more detailed examination identified the following relationship (see also Figure 41 and Figure 42):

- While the aircraft was at 37,000 ft, the value of most of the Mach spikes was 0.128. The Mach binary values corresponded to the altitude binary values for 37,000 ft (see also Figure 39).[131]
- During the aircraft's descent to Learmonth, the binary values for most of the Mach spikes corresponded to the altitude value at the same time.
- Overall, 88% of the ARINC 429 binary values for the Mach spikes matched exactly with the corresponding altitude values at the corresponding time.

Computed airspeed versus altitude

Some binary values for computed airspeed corresponded to the binary values for altitude during the aircraft's descent. However, the proportion of the binary values for the computed airspeed spikes that matched altitude values (28%) was lower than that for Mach spikes (88%).

AOA versus altitude

There was also some evidence of matching between the corrected AOA data spikes and the altitude values. This evidence included the following (see also Figure 43 and Figure 44):

- The binary value of the first AOA spike of 50.625° corresponded to the binary value for altitude at that time (37,000 ft). Although most of the 50.625° spikes occurred when the aircraft was at 37,000 ft, there was one 50.625° spike after the aircraft had started descending.
- The binary value of the first recorded AOA spike of 16.875° corresponded to the binary altitude value at that time (12,492 ft). Subsequent 16.875° spike values did not correspond to altitude values.

[131] Strictly speaking, a Mach value of 0.128 could correspond to any altitude in the range 36,864 ft to 36,991 ft.

Figure 41: Qualitative correlation between Mach spikes and altitude

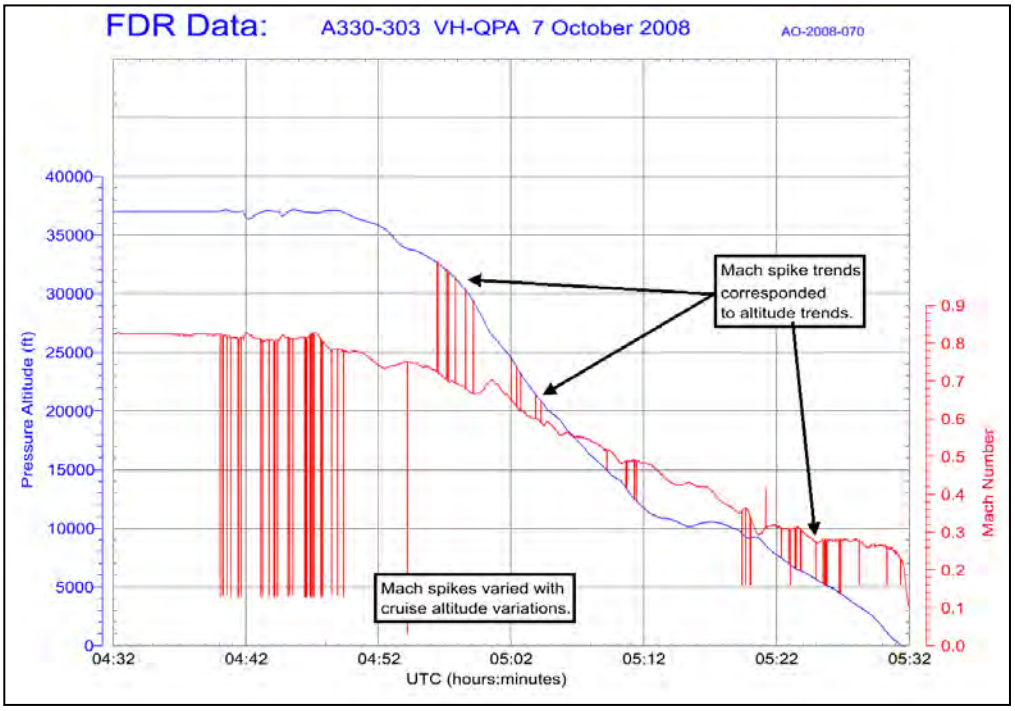

Figure 42: Quantitative correlation between Mach spikes and altitude

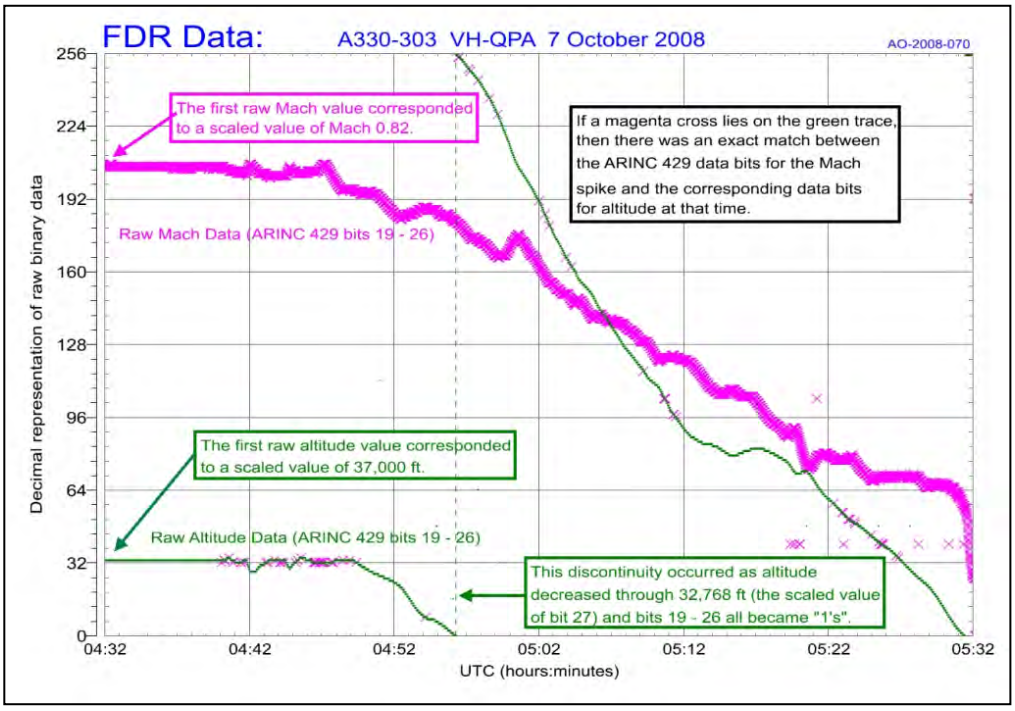

Figure 43: Qualitative correlation between AOA 1 spike values and altitude

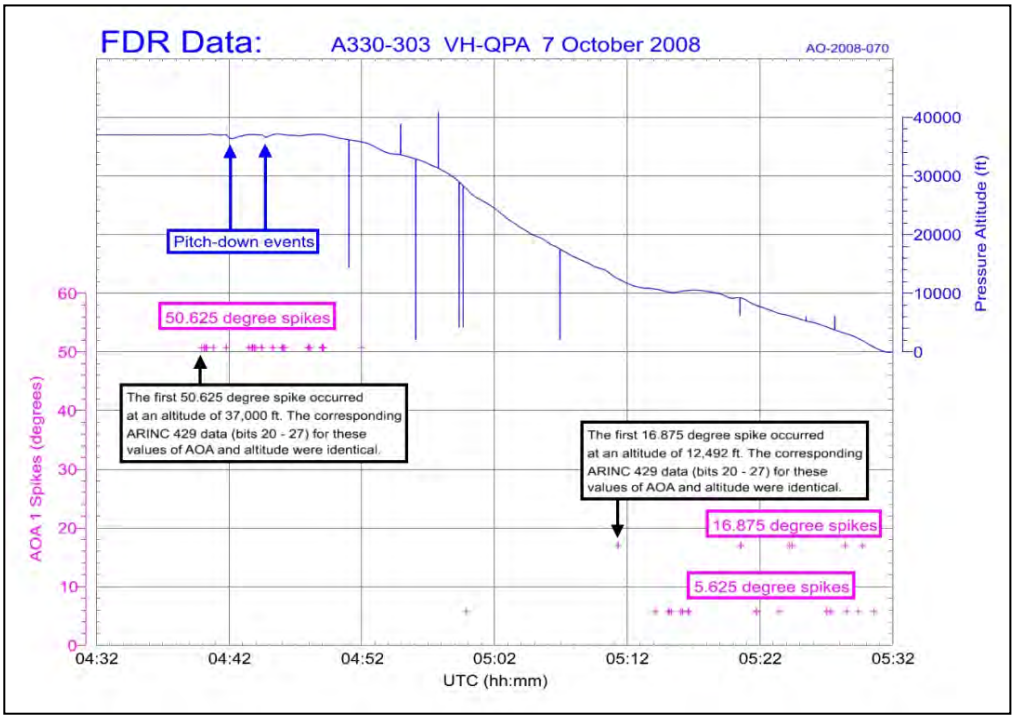

Figure 44: Quantitative correlation between AOA 1 spikes and altitude

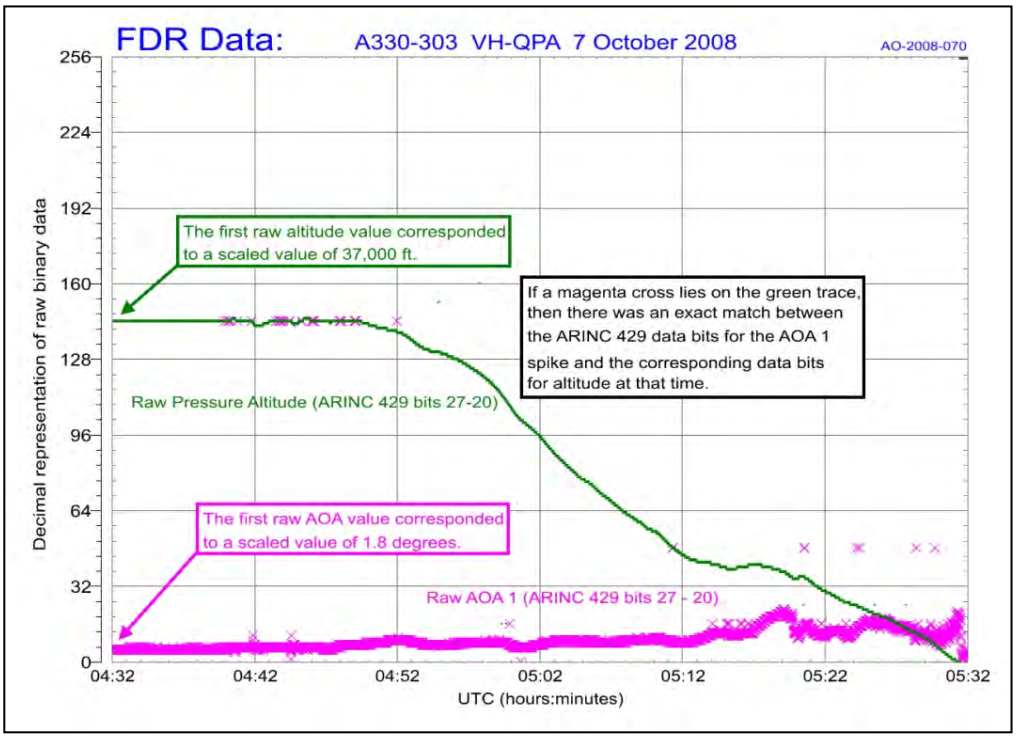

ADR discrete word #1

In addition to flight data parameters, the ADIRU outputted words containing documentary, status and fault data. This data was outputted in the ARINC 429 format on the same databuses that outputted the flight data parameters. Some of these discrete words contained multiple parameters; for example, ADR discrete word #1 included 17 separate status flags, including probe heat status information. Figure C1 in Appendix C lists the parameters and their corresponding bit locations for ADR discrete word #1.

As discussed in section 1.12.9, the probe heat faults resulting in the 'A.ICE' cockpit effect messages on the PFR were considered to be spurious messages as a result of incorrect ADIRU 1 outputs. The probe heat discretes were located in the lower order bits of the ARINC 429 word's data field (that is, bits 12, 14, 15, 16 and 17), with a value of '0' indicating a fault and '1' indicating no fault. If the ADR discrete word #1 parameter received an incorrect data value, similar to a spike value on a flight data parameter such as Mach or AOA, then this could account for the spurious activation of the probe heat faults.

The probe heat discretes correlated with the least significant bits (smallest values) of flight data parameters such as altitude (that is, 1, 4, 8, 16 and 32 ft). As a result, there were many altitude values that could have triggered probe heat faults if they had coincided with a data spike and been mapped into the discrete word. Due to the limitations in the recorded data, the investigation was unable to determine whether the probe heat faults were due to altitude data or some other parameter being mapped into the discrete word.

ADR discrete word #1 also contained the 'overspeed warning' bit (bit 19) that was the trigger for the FWS to issue an overspeed warning to the crew. In other words, the FWS triggered overspeed warnings based on the value of bit 19 on the discrete parameter and not the actual computed airspeed or Mach value. As with the probe heat faults, the overspeed warnings were considered to be spurious warnings as a result of incorrect ADIRU 1 outputs. That is, each overspeed warning appeared to equate to a data spike on the ADR discrete word # 1 parameter, but the origin of the incorrect data on the parameter could not be determined. Unlike overspeed warnings, the FWS determined whether a stall warning would be generated by using actual values of corrected AOA, together with other parameters such as flap/slat position (see sections 1.6.5 and 1.11.4).

Altitude parameters

The ADR part of the ADIRU outputted four altitude parameters: standard altitude, baro-corrected altitude #1, baro-corrected altitude #2, and baro-corrected altitude #3. The baro-corrected altitudes corresponded to the values from the captain's, first officer's and standby altimeters respectively. Above the transition altitude (typically 10,000 ft), all altimeters were referenced to 1013.2 hPa, so standard altitude and baro-corrected altitude values were identical (within normal system tolerances). Below the transition altitude, the crew adjusted their altimeters using the current local atmospheric pressure setting. The baro-corrected altitude parameters reflected these settings, and standard altitude was always referenced to 1013.2 hPa.

For the accident flight, the pressure setting at Learmonth was 1014 hPa. As this was close to the standard setting of 1013.2 hPa, no significant difference was expected

between the values of any of the four altitude parameters. The FDR recorded standard altitude, while the QAR recorded standard altitude and baro-corrected altitude #1. In summary, the altitude values that correlated with some data spikes could have originated from any of the four altitude parameters.

Other parameters

A review of some of the other recorded data did not identify any other salient matches between the binary values of the different ADIRU parameters. However, most of the ADIRU parameters were not recorded, and most of the data for each parameter was not recorded. In addition, the nature of the IR data made it difficult for the investigation to identify spikes in that data and to establish whether any patterns existed. As discussed in section 1.11, the IR parameters rapidly diverged from realistic values and oscillated between unrealistically large values at varying rates for the remainder of the flight.

3.3.5 Data-spike patterns for the 27 December 2008 flight

For the 27 December 2008 occurrence, there was only a 24-second period when spikes were outputted from ADIRU 1 before the ADR was switched off. The computed airspeed spikes appeared to correlate with altitude data in a similar way as for 7 October 2008 event, but there were too few spikes recorded to form a definitive conclusion. There were no Mach or AOA spikes recorded. Further information on the FDR and QAR data are provided in Appendix D.

3.3.6 Summary

For the 7 October 2008 occurrence, there was a clear match between the values of most of the Mach data spikes and the concurrent values for altitude, both at cruise and during the descent. There was also some evidence of data spikes for computed airspeed matching with altitude values in both the 7 October 2008 and 27 December 2008 occurrences, and a possible relationship between the AOA spikes and altitude values in the 7 October 2008 occurrence.

The investigation was unable to determine the extent to which other patterns existed in the data spikes because of limitations in the recorded data. The investigation was also not able to determine whether the repeated values of some data spikes were due to the spike values being stored in memory, or due to the spike values resulting from the exchange of data with another parameter that did not vary during the period of interest (such as a documentary data, status data or fault data parameter).

3.4 Data flow analyses

3.4.1 Background

The investigation examined the LTN-101 ADIRU's data processing stages to determine if there were indications of which stage or stages problems may have occurred. A simplified summary of the data flow used for most of the output

parameters is provided in Figure 45.[132] The processing stages considered likely to be affected by the data-spike failure mode are highlighted in red in Figure 45 and discussed in the remainder of this section.

Figure 45: Summary of the data flow for the ADR and IR parameters

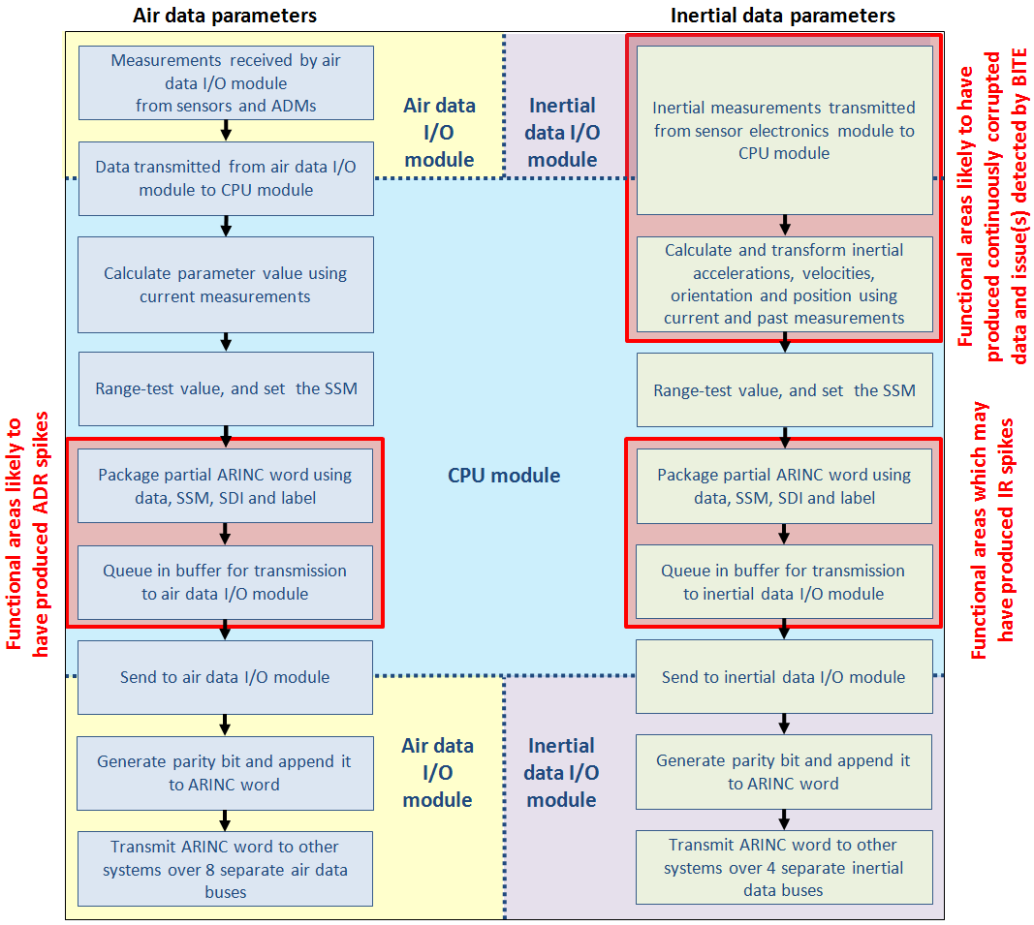

3.4.2 Receiving input data

The ADIRU received its input data from several external and internal sensors, which provided data to the ADIRU's input/output modules over multiple, independent wires. The two input/output modules were also independent. Overall, the investigation concluded that it was very unlikely that there would be simultaneous problems with several independent sensors or multiple independent wires, or even two independent input/output modules. In addition, the ADIRU's input/output modules and input and output wiring were tested, and no problems were identified (see section 1.12.5 and Appendix F). Analysis of subsequent processing stages identified that data processing problems also originated within the ADIRU itself.

[132] Some processing steps were carried out by separate software functions for different parameters. For example, the software functions for range-testing computed airspeed and Mach were different, although both were conducted within the CPU module.

3.4.3 Calculating parameter values

ADR parameters

As noted in section 1.11, the recorded values between data spikes for each ADR flight parameter did not show any anomalies. That is, the underlying data (or data between the spike values) was correct.

The calculation of some of the ADR parameters required the use of other parameters. Accordingly, if the data spikes were introduced at the calculation stage, it would be expected that a data-spike value in a given parameter would produce a concurrent change in the value of any dependent parameters. A review of the data for some of the key parameters for the 7 October 2008 flight found the following:

- The calculation of altitude involved applying a correction for the aircraft's AOA. There was no correlation between the AOA spikes and the concurrent values of altitude.
- The calculation of Mach involved applying a correction for altitude. There was no correlation between the altitude spikes and the concurrent values of Mach.

The aircraft's recorders would not have sampled the relevant parameters at the same time on each occasion. However, given the large number of data spikes, it would be expected that if these spikes were being introduced at or before the calculation stage, a number of the altitude values would have shown a predictable change in response to the AOA spikes. Similarly, a number of the Mach values would have shown a predictable change in response to the altitude spikes. Accordingly it appeared that the data spikes were introduced after the calculation stage.

IR parameters

In addition to containing data spikes, the output data for the IR parameters included significant variations from the correct value. In other words, the underlying data was incorrect.

As discussed in section 1.6.4, the calculation of most IR parameters was heavily dependent on previous measurements. Therefore, if an incorrect data value was introduced into the calculations, its effect would be cumulative and result in a 'drift' from the correct value over time. The rapidity of the drift would depend on the magnitude of the incorrect value and the dependence of the subsequent values on that value.

The ADIRU temporarily stored IR measurements and the results of intermediate and final calculations in RAM. If incorrect data was being stored, or data was read from the wrong location in memory, the IR calculations would use an erroneous input and result in a variation from the correct values for multiple IR parameters. The ADIRU manufacturer advised that this mechanism could explain the underlying corruption of the IR data that was observed on the occurrence flights.

The IR parameter values were also highly interdependent; that is, each parameter was influenced by some of the other IR parameters. Therefore, if one or more parameters were corrupted prior to being used for the calculation of other parameters, these other parameters would also become corrupted. These other parameters would then be used in subsequent calculations, resulting in an increasingly erroneous set of computations.

3.4.4 Range checking

Twenty-eight ADIRU output parameters, including altitude, AOA, groundspeed and Mach, were range tested by the ADIRU's BITE before being sent to the input/output modules. Any values that exceeded the predefined limits resulted in the parameter's data being set as invalid (for as long as the data exceeded the limits).

The investigation examined the data for the 7 October 2008 occurrence for parameters where the range check boundaries were exceeded. The recorded data (including data spikes) for most of the parameters were within the range-check limits. For example, the range check limits for AOA were -40 to +90°, and none of the AOA spikes were outside that range. However, there was relevant data involving computed airspeed (an ADR parameter), Mach (an ADR parameter) and groundspeed (an IR parameter) that is summarised in Table 23.

Table 23: Range check information for computed airspeed, Mach and groundspeed

Parameter	Relevant range check	Recorded data (7 October 2008)
Computed airspeed (CAS)	The allowable CAS range was 0 to 450 kts. If CAS was outside this range then it was set to 0 and the SSM was marked as invalid by setting it to 'failure warning' (FW). If the range check passed, then a further check was performed to see if CAS was less than 30 kts. If so, then it was set to 0 and the SSM was marked as invalid by setting it to 'no computed data' (NCD).	Several CAS spike values were recorded between 0 and 30 kts without being flagged as invalid. During landing at Learmonth, the computed airspeed was set to 0 after the speed decreased through 30 kts. This indicated that the range check function was operating correctly at that time.
Mach	The allowable Mach range was 0 to 1. If Mach was outside this range then it was set to 0 and the SSM was marked as invalid by setting it to 'failure warning' (FW). If the range check passed, then a further check was performed to see if Mach was less than 0.1. If so, then it was set to 0 and the SSM was marked as invalid by setting it to 'no computed data' (NCD).	Two Mach spike values of 1.016 were recorded without being flagged as invalid. During landing at Learmonth, the Mach was set to 0 after the speed decreased below 0.1. This indicated that the range check function was operating correctly at that time
Groundspeed	If the calculated groundspeed exceeded 1,000 kts, then the output value was set to 1,000 kts and the SSM was marked as invalid by setting it to 'failure warning' (FW).	No values of groundspeed were recorded that were greater than 1,000 kts. The underlying (and incorrect) groundspeed data showed a gradual increase up to the value of 1,000 kts and remained at that value with intermittent spikes to lower values.

Overall, the pattern of data indicated that the range-check functions were operating normally when the underlying data exceeded the predefined limits. However, in the case where data spikes exceeded the limits, the range checks were not applied, indicating that the data spikes were introduced after the range checks were conducted.

3.4.5 Packaging the ARINC 429 word

The process of combining the ARINC 429 fields into a 32-bit data word occurred in the CPU module and just prior to the word's transmission to one of the input/output modules. As described in section 3.3, many of the data spikes appeared to result from the output data word containing the binary data (or data field) from another parameter, indicating that either the data field was combined with the wrong label field, or the label field was combined with the wrong data field.

As previously noted, the ADR part of the ADIRU generated four altitude parameters, and each of them would have provided very similar information during the 7 October 2008 flight. The ADIRU also outputted these parameters at 16 times per second, which was a higher rate than for most of the other ADR parameters (Table 1). Consequently, if there was a problem with the packaging process, it would probably be most clearly evident with altitude values being used for other parameters, as was observed.

In summary, the available evidence indicated that a packaging process that combined a data field with the wrong label field, or a label field with the wrong data field was consistent with many of the recorded data spikes. This could have occurred due to the label field or the data field being retrieved from the wrong location, or being retrieved in an incorrect order. However, any such packaging problem was intermittent rather than consistent.

3.4.6 Queuing data for transmission

After an ADR parameter data word was packaged, it was queued and then transmitted to the air data input/output module. The ADR parameters were organised in 10 groups for the queuing and transmission, with there being 32 transmission phases from each group outputted per second.

The parameters were grouped according to their required output rates. For example, 10 parameters (including AOA and altitude) were included in every second phase, so that those parameters would be outputted 16 times per second. Another 10 parameters (including Mach and computed airspeed) were included in every fourth phase and were outputted eight times per second. Some parameters (including TAT, SAT, and status messages) were included in every sixteenth phase and so were outputted twice per second.

The memory location used to buffer (temporarily store) the parameter data prior to transmission was the same for all of the parameters in each group, and the parameters would always be sequenced in the same order within the group. For example, each of the following sets of parameters was in the same queue and shared a memory location:

- baro-corrected altitude #1, Mach and right static pressure
- baro-corrected altitude #2, computed airspeed and average static pressure
- corrected AOA, ADR discrete word #2[133], average static pressure (corrected), SAT, and barometric correction (in inches of mercury)

[133] ADR discrete word #2 did not contain any parameters that could be analysed using FDR, QAR or PFR data and is not discussed further in this report.

- indicated AOA, ADR discrete word #1, impact pressure, TAT and barometric correction (in millibars).

Two of the parameter pairs that were found to exhibit some correlation in the recorded data (section 3.3.4) shared a memory location. These pairs were Mach and altitude, and computed airspeed and altitude. Although corrected AOA showed some correlation with altitude, it did not share a memory location with any of the altitude parameters. Similarly, ADR discrete #1, which incorporated static heat, probe heat, and other status information, showed some correlation with altitude but did not share a memory location with any of the altitude parameters.

The ARINC 429 words were queued in RAM in two 16-bit halves. The least significant half comprised the label (bits 1-8) and data bits 9-16 while the most significant half included data bits 17-29. In other words, the label was in the same 16-bit half as the least significant data bits and the most significant data bits were in the other half. For most parameters, the FDR recorded bits from the most significant half of the ARINC 429 word; for example, only Mach bits 16-27 were recorded by the FDR.

In summary, a queuing process that combined the least significant half of one ARINC 429 word with the most significant half of a different word would be consistent with many of the recorded data spikes (particularly for Mach and computed airspeed). However this process would not explain the spurious probe heat faults that were recorded on the PFR, as four of the five probe heat discretes were located in the least significant 16 bits of the discrete word, which included the label. In addition, the IR parameters did not queue the data for multiple parameters in the same location, and therefore a problem with queuing within the same memory location would not explain the IR data spikes.

3.4.7 Generating parity bits

The parity bit was calculated and appended to the ARINC 429 word in hardware within the air data and the inertial reference input/output modules. Receiving systems used the parity bit to check whether the data had been corrupted during or after transmission, and they would reject any such corrupted data.

Depending on the design of the receiving systems, numerous fault messages could be generated. For example, the FCPCs would report an ADIRU problem if insufficient valid data was received. In addition, for some parameters the other ADIRUs would also detect any words that had an incorrect parity using the digital air data system (section 1.12.6). No such faults were recorded on the occurrence flights.

The FDR and QAR systems both checked the parity of the data words, and any data words with incorrect parity were not recorded. As many data spikes were recorded by both systems, there appeared to be no parity bit problems with the data outputted by the ADIRU.

In summary, the data problems associated with the data-spike failure mode occurred prior to the parity bit being added by the input/output modules.

3.4.8 Transmitting data to other systems

The ADIRU failure mode affected data that was transmitted on multiple, segregated wires out of the ADIRU. A simultaneous problem with multiple, physically-segregated wires was considered very unlikely. A range of testing conducted on the wires also found no problems. A wiring problem would also be expected to produce parity bit errors, and none were observed (section 3.4.7).

In addition to problems with the wiring itself, another possible mechanism for producing the erroneous output data was electromagnetic interference (EMI) on the output databuses (see also section 3.6.5). EMI would typically affect a word on a databus in a random manner; that is, some or all of the binary values would be changed in any given word. The recorded information from the 7 October 2008 event did not exhibit this behaviour. More specifically:

- If EMI was directly affecting the transmitted data, there would be extensive parity bit problems, but no such problems were indicated in the recorded data (section 3.4.7).
- The ARINC 429 fields (section 3.3.2) were predominantly valid; that is, almost all had a valid label field, SSM field, and SDI field. EMI on the databus would be expected to produce many invalid values for each field.
- Many of the data spikes for some parameters showed a consistent value, which would be very unlikely if they were produced via EMI on an output databus.

Overall, there was no evidence to indicate that the ADIRU's data was affected after it was transmitted.

3.4.9 Additional analyses

Memory address comparisons

The ADIRU's input and output data parameters were all stored in the RAM chips, along with other information such as the intermediate results of processor computations. Each data item was allocated a particular fixed memory address (location within memory) and could be accessed using 'address lines', a set of interfaces that each carried one bit of the memory address to be accessed at any particular point in time.

The investigation examined the potential for data to have been written to and read from the incorrect locations in memory. This could result in some parameters being mislabelled or containing data from other parameters, as well as information from other locations in memory. However, an examination of the memory addresses found that the data spikes could not be explained by data being misread from an adjacent address line. The investigation also could not find any consistent patterns in the addresses that could explain the observed data spikes.

Parameter label comparisons

As noted in section 3.3.2, each parameter had an 8-bit field to represent its label in the ARINC 429 system. Examples of these labels are provided in Table 24 for several key parameters.

Table 24: Binary values for the label field for some key parameters

Parameter	Label							
	1	**2**	**3**	**4**	**5**	**6**	**7**	**8**
Standard altitude	1	0	0	0	0	0	1	1
Baro-corrected altitude #1	1	0	0	0	0	1	0	0
Baro-corrected altitude #2	1	0	0	1	0	0	0	0
Baro-corrected altitude #3	1	0	1	0	1	0	0	1
Mach	1	0	0	0	0	1	0	1
Computed airspeed	1	0	0	0	0	1	1	0
Corrected AOA	1	0	1	0	0	0	0	1
ADR discrete word #1	1	0	1	1	1	0	0	0

As indicated in the table, Mach and baro-corrected altitude #1 differed by only one bit, and that computed airspeed and baro-corrected altitude #1 also differed by only one bit. However, the bits involved in both cases were different, indicating that there was probably not a single 'stuck bit' problem[134] with the data interfaces between the CPU chip and the RAM chips.

In addition, the corrected AOA's label was at least two bits different from each of the altitude labels, indicating that its data spikes did not appear to be produced simply by one bit in the label field being changed (or 'flipped').

Overall, bit flips in the parameter label did not appear to provide a viable explanation for the data spikes, as the investigation could not conceive of a mechanism that would frequently produce such bit flips without the mechanism being applied consistently to all labels. In addition, only 39 of the 256 possible label values were used for ADR data. Corruption of the label field would be expected to produce a significant amount of data with invalid labels, and consequently a significant amount of missing data for the actual parameters. This pattern was not observed in the recorded data.

Wait-state tests

The ADIRU's CPU chip accessed memory and other modules through an internal databus. Due to the different speeds of the various circuits, the CPU module incorporated 'wait-state' timing that ensured that it waited a predetermined amount of time for the memory or other modules to respond. If the timing was incorrect, the CPU chip could write or read memory to an incorrect location. The wait-state timings were stored in the RAM chip associated with the ASIC.

The ADIRU manufacturer performed a test on an LTN-101 unit that varied the duration of each wait state to establish the effect of these variations. A number of failures occurred, which would normally be expected when conducting such a test. However, no effects relevant to the occurrence were observed.

---

[134] A stuck bit in such an interface would normally be expected to affect the same bit in each word that is sent across the interface. It would be the result of a single faulty part in either the transmitting or receiving chip, or in the physical data line between them (that is, a track on the circuit board).

ARINC 429 packaging analysis

The ADIRU manufacturer attempted to identify all of the potential mechanisms on the ADR part by which ARINC 429 words could have been 'cross-contaminated'. Each mechanism was then examined in detail to determine whether its expected effects were consistent with the characteristics of the data-spike failure mode.

Twenty-four potential mechanisms were identified, including incorrect software branching and data read/write fails. Nineteen of the mechanisms were assessed as having effects different to those of the data-spike failure mode. Another four were considered unlikely or improbable for various reasons, such as some of the expected effects differing significantly from the characteristics of the data-spike failure mode.

The only mechanism identified as a 'likely or plausible cause' by the ADIRU manufacturer involved the ADIRU's method for queuing ADR data prior to sending it to the air data input/output module. This potential mechanism was previously discussed in section 3.4.6, and was noted as being consistent with many of the recorded data spikes; however, it did not explain the spurious probe heat faults or the IR data spikes.

3.4.10 Summary

Based on the various analyses, the investigation was able to limit the possible functional areas that could have produced the observed output data patterns (as summarised in Figure 45). More specifically:

- The spikes in the ADR parameters were probably introduced within the CPU module after the parameters were range tested and before the ARINC 429 messages were queued in the CPU module's output buffer.
- There was insufficient evidence to determine the origin of the IR spikes, although there was no evidence to suggest that they were generated in a different manner to the ADR spikes.
- Another form of data corruption was introduced into the IR parameters at some point before the inertial calculations were completed.

Overall, the data flow analyses confirmed that the incorrect data was due to a disruption of some processing activities within the CPU module. Although the exact nature of the disruption was not able to be determined, it was notable that many of the observed effects were consistent with problems with storing or retrieving data from memory. For example:

- The CPU module stored IR data in RAM when it was being transferred from the sensor electronics module, and stored the intermediate and final calculation results in RAM. Incorrect storage or retrieval of this data could lead to significant corruption of IR data due to the feedback and interdependence of the parameters.
- Packaging the ADR and IR data into ARINC 429 words involved reading data from the RAM and then combining it with the data label and other information. A problem with these functions could explain the apparent mixing of parameter data.
- The storage of fault data and routine messages in the ADIRU's BITE memory involved the use of the RAM and associated interfaces. A problem with this

storage could result in BITE information not being recorded correctly (see also section 3.7.5).

The components of the CPU module that were involved in storing and retrieving data in memory were the CPU chip, companion ASIC, the wait-state RAM chip, and the other four RAM chips. As noted, testing of the wait-states did not reveal any effects relevant to the occurrence.

3.5 Review of ADIRU configuration and service history

3.5.1 Background

Two of the three known data-spike occurrences involved the same LTN-101 ADIRU (unit 4167), and the other occurrence involved a unit with a similar serial number (unit 4122). Accordingly, the investigation examined the configuration and service history of the two units to determine if there were specific aspects of these units that could be associated with the data-spike failure mode.

3.5.2 Hardware configuration

Overview of module and component batches

Each module of the ADIRU consisted of a printed circuit board (PCB) containing electronic components. External contractors manufactured and populated the PCBs with components. When completed, the PCBs were delivered to the ADIRU manufacturer, who tested the individual modules and then assembled and tested the ADIRU as a complete unit.

The ADIRU manufacturer reported that modules were ordered in lots of 50. Each module was annotated with a part number and a serial number, and detailed records were kept so that a module's service history could be tracked. In general, modules from the same batch contained components from identical or similar batches.

Most components were ordered in lots of approximately 500. For the CPU chip[135], a life-time purchase was made to cover future production and spares. The CPU chip was critical to the design of the ADIRU and a life-time purchase mitigated the risk of a future lack of supply or obsolescence. However, other components were not life-time purchases, and they could be replaced with equivalent components over time. These changes would not result in a design change or a corresponding part number change.

Complex components within each module, such as memory chips, were annotated with a part number and usually had additional markings such as the production batch number, a manufacturing date code, and sometimes a serial number. Less complex components, such as resistors, were not annotated with a part number, although they had a part number when they were ordered.

The ADIRU manufacturer did not have records of the serial numbers, batch numbers or date codes of the components within each module. Therefore, it was not possible to track the service history of the components, or to easily compare the

[135] Part number 80960MC, manufactured by Intel Corporation.

component configurations of different units. Another ADIRU manufacturer advised that its tracking capability was also at the module level and not the component level.

In general, there can be variations between components that occur during manufacture or over their service life, even though they are built to a specific standard and have the same part number. Components from the same batch will normally be more similar than components from different batches.

Module and component batches for ADIRUs 4167 and 4122

Both ADIRUs 4167 and 4122 were manufactured in August 2002. At the time of the data-spike occurrences (12 September 2006, 7 October 2008 and 27 December 2008 respectively), their configurations were very similar. The two units had the same part numbers for each module except for two sub-assemblies of the sensor electronics module.

The CPU modules that were used in both units 4167 and 4122 had the part number 465474-03, and the serial numbers were 9-7315 and 9-7273 respectively (a difference of 42 in the build sequence). The two modules were built in adjacent batches of 50 (9-7300 to 9-7349 versus 9-7250 to 9-7299).

In addition to the CPU chip, the other key components in the two units' CPU modules had the same part numbers and came from the same or similar batches. More specifically:

- The companion ASICs had the same part number[136], and all other markings were identical.
- The wait-state RAM chips had the same part number[137], and a visual examination showed that all markings were identical for both units.
- The RAM chips used for CPU data all had the same part number.[138] For unit 4167, each of the four chips had identical markings and contained a '-9' in the production code. For unit 4122 three devices had a '-10' in a production code while one had a '-9'.

Subsequent CPU modules

The LTN-101 ADIRU's CPU module was later redesigned to reduce costs and to include error detection and correction (EDAC). EDAC is used for detecting and correcting single-bit errors in RAM chips to give protection from single event effects (SEEs, see section 3.6.6). This change was a significant redesign and resulted in a new CPU module part number (466871-01). The EDAC was performed by a new ASIC, and all of the RAM chips used on the CPU module were replaced with a different chip.[139]

---

[136] Part number 5962-99A2001QXC, manufactured by AMI Semiconductor.

[137] Part number MT5C2564C, manufactured by Austin Semiconductor.

[138] Part number MSM8128JMB, manufactured by Mosaic Semiconductor.

[139] EDAC was functionally located between the CPU chip and the RAM and required more memory to be available than in a design without EDAC.

The redesigned CPU module was incorporated in all ADIRUs from serial number 4385 (late 2002) onwards. The ADIRU manufacturer advised that it was also used in some ADIRUs prior to serial number 4385, and it was incorporated when older units were returned for repair.

3.5.3 Software version

The ADIRU software was changed from time to time as updates and improvements were incorporated. At the time of manufacture, units 4167 and 4122 had software version -0312 installed. Updated versions were usually promulgated as optional service bulletins[140], and operators could decide whether the advantages of installing an updated version of the software were sufficient to justify the logistics of upgrading each aircraft (three ADIRUs per aircraft). The operator of QPA elected not to load software versions -0313 and -0314 in any of their ADIRUs.

Software version -0315 was loaded on unit 4167 on 20 July 2005 and was the version installed at the time of the 12 September 2006 occurrence. Software version -0316 was released in August 2008; it was the version installed on unit 4167 at the time of the 7 October 2008 occurrence and it was also installed on unit 4122 at the time of the 27 December 2008 occurrence. As far as could be determined, most of the LTN-101 ADIRUs had software versions -0315 and/or -0316 installed.

3.5.4 Service histories

Details of the service histories of units 4167 and 4122 are summarised in Figure 46. Notable features include:

- Unit 4122 was removed from service three times for examination (two times for reported drift problems[141]), and unit 4167 was removed once. The removal frequency was normal for this type of unit.
- Both units experienced a similar type of fault early in their service life that involved temperature sensor failures. In both cases, the manufacturer replaced the sensor electronics module and performed a thermal calibration. These were the only faults of this type experienced on the operator's fleet; however, the manufacturer advised that it was not an uncommon fault, with 24 cases occurring across the world fleet of LTN-101s during the period from 2006 to 2008. Both units had many hours of service following these repairs.
- In addition to removals, both units had been associated with reported IR and/or ADR faults that did not require removal. Such faults were not unusual (section 3.9.2).
- Unit 4122 experienced an IR 1 fault in July 2008, 5 months prior to the 27 December 2008 data-spike occurrence. Available BITE data was consistent with the fault being due to a type of SEE (section 3.6.6). The ADIRU

---

[140] Service bulletins are issued by aircraft, component, or engine manufacturers to provide operators with relevant service information. Not all service bulletins are safety-related, and compliance with a particular service bulletin can only be mandated by the State of Registry of an aircraft.

[141] After many months or years of service, the drift rate of laser-ring gyro inertial systems can exceed the allowable threshold and the unit needs to be re-calibrated. This is a relatively routine maintenance requirement.

manufacturer reported such BITE data was not uncommon, with over 100 similar events reported across the world fleet since 2000.

- Unit 4167 experienced an IR 1 fault in June 2006, 3 months prior to its first data-spike occurrence on 12 September 2006. Very little recorded data in the unit's BITE was available for this period, but the available information indicated that the fault was not consistent with a similar type of SEE as occurred for unit 4122 in July 2008.
- Unit 4167 had operated for over 2 years after its first data-spike occurrence without any reported problems, and its software was replaced with a new version during this period (as part of a routine upgrade).

3.5.5 Summary

As discussed in section 1.16.2, there was nothing unusual about the operational environment associated with the three data-spike occurrences. Therefore, the fact that two of the three data-spike occurrences involved one unit (out of about 8,000) provided a strong indication that some aspect specific to the affected units contributed to the data-spike failure mode.

The software installed on the two affected units was common to most of the units, and no software-related problems were found on the two affected units during testing (section 3.6.2). Therefore, it is probable that some aspect of the two affected units' hardware was associated with the failure mode.

The key hardware components in the CPU modules of units 4167 and 4122 were very similar, but the number of other units that had the same level of similarity could not be determined as component details were not recorded.

Figure 46: Service history timeline for units 4122 and 4167

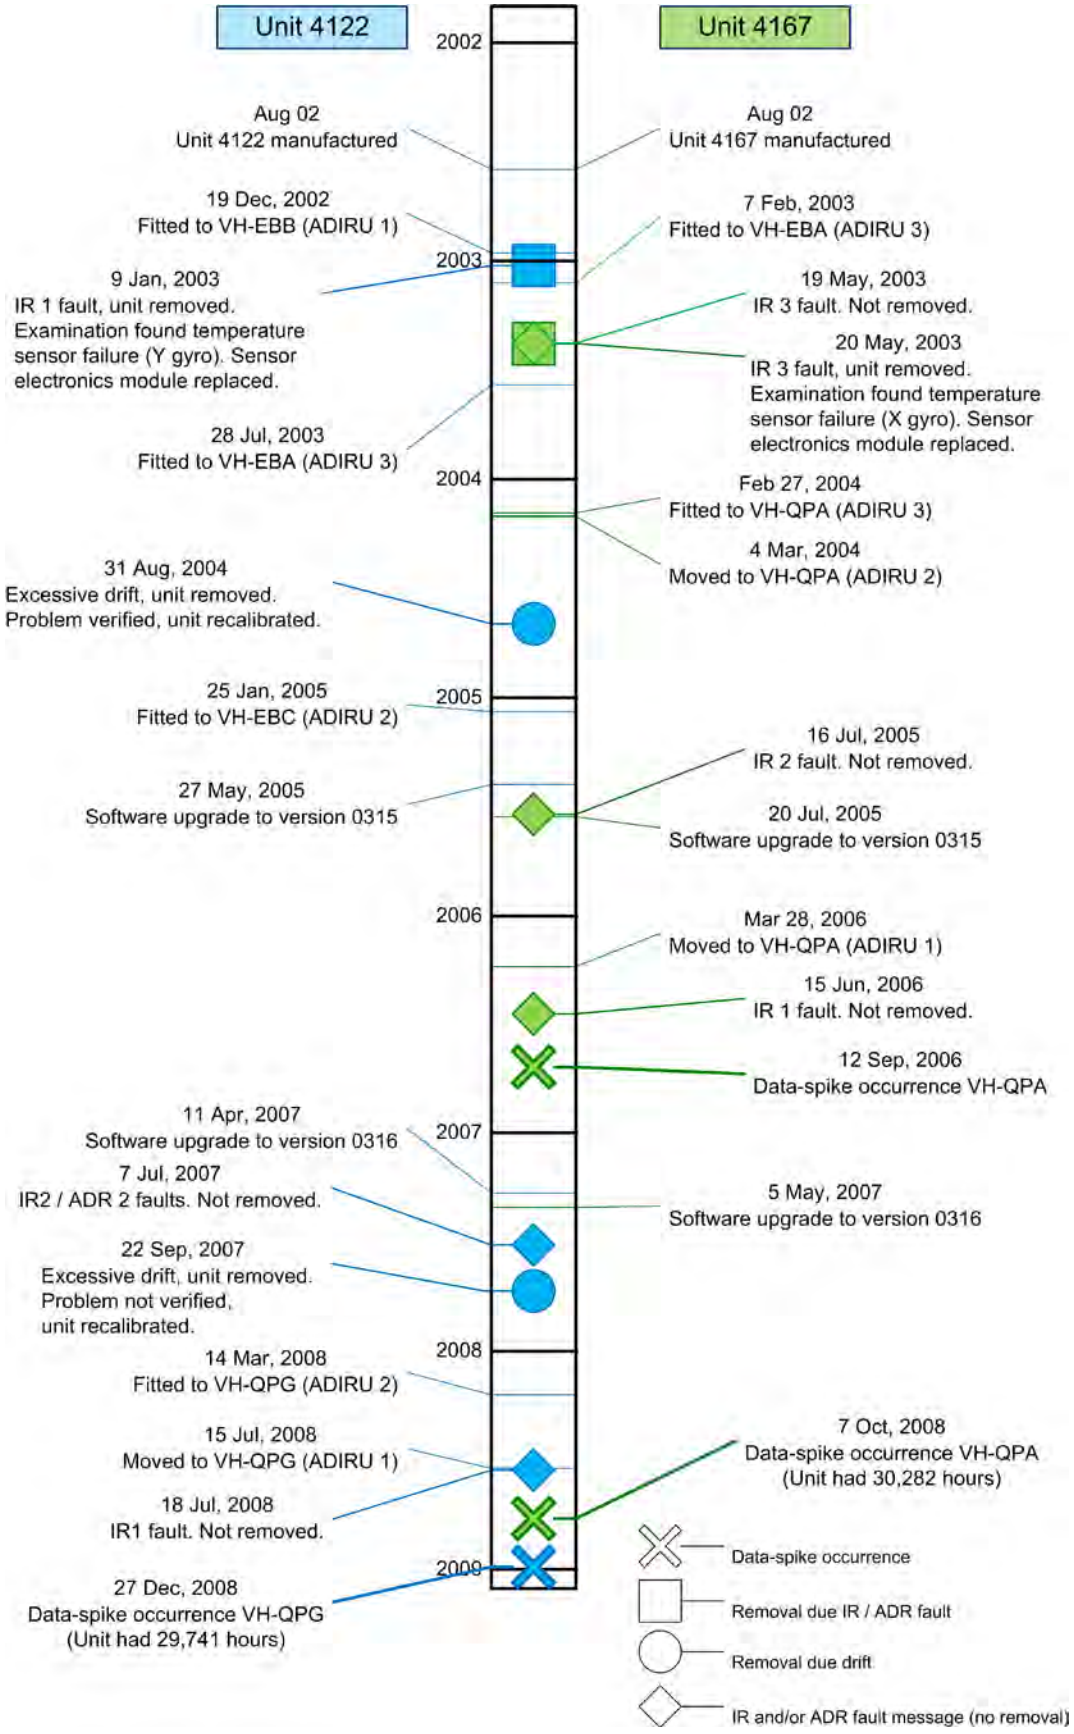

3.6 Potential trigger types

3.6.1 Background

As discussed in section 3.4.10, the data-spike failure mode involved a disruption of processing activities within the ADIRU's CPU module. The failure mode was also a soft fault; that is, it ceased to exist after the unit was powered off.

Soft faults can be triggered by a range of factors, such as a software corruption or 'bug', or a hardware anomaly triggered by some form of environmental factor (including physical environment factors, electromagnetic interference, or single event effects). Evidence relevant to each of these possibilities is provided in the remainder of this section.

None of the effects of the data-spike failure mode were observed until the failure mode was triggered, and the effects recurred continuously until the unit was shut down. In addition, the failure mode could not be replicated during subsequent testing using a wide variety of potential triggers, indicating that the failure mode could only be triggered by either a rare event, or under rare circumstances. This behaviour indicated that the disruption to the CPU module's processing was triggered by a single event rather than a series of events (that is, it was very unlikely that a separate event triggered each data spike or other processing problem).

3.6.2 Software

Software 'bug'

A software 'bug' is a flaw or mistake in a program that causes it to behave in a way that was not intended by its designers. Such bugs only manifest themselves under a specific set of circumstances that produce the same fault each time they occur. However, it may be extremely difficult to reproduce the program state that existed at the time a problem occurred because complex software has a very large number of data items and functions, all of which interact with each other and most of which are time-dependent.

Overall, if the failure mode was due to a software bug, it would be expected that it would not reoccur more than once on any given unit without also occurring on many other units. This was not the case with the LTN-101 data-spike failure mode, where the problem occurred twice on one unit and only once on another unit out of more than 8,000 units in operation.

Software corruption

The LTN-101 ADIRU software was stored in ROM. If software corruption were to occur in this type of memory, it would generally result in a hard fault because the erroneous software would be loaded every time the unit was powered up, and the fault would be reproduced as soon as the corrupted instruction was executed.

Nevertheless, software corruption can sometimes affect a system's behaviour in less obvious ways, such as if the corruption occurred in a set of instructions that were only executed on rare occasions. In such cases the problem would only reoccur when those instructions were executed. As discussed in section 1.16.2, there was

nothing unusual regarding the flights associated with the three data-spike occurrences.

The ADIRU's BITE included repeated software integrity checks using a common and reliable method called a cyclic redundancy check (CRC, see also section 3.2.3). These checks involved each software partition being verified during initial loading, with any corruption being immediately detected. Each software partition was also verified at intervals while the software was running, and any corruption resulted in the ADIRU shutting down the affected partition(s) and generating ADR and/or IR fault messages.

The two ADIRUs (4167 and 4122) involved in the three data-spike occurrences had the integrity of their software checked during the post-occurrence unit tests described in section 1.12.6. The software was found to be correct, showing that no software corruption had taken place. In addition, functional testing of both ADIRUs found no software-related problems.

Relevant software BITE

In addition to the CRCs, the LTN-101 ADIRU included several mechanisms to detect when its software was not operating correctly, and no problems were noted during the 7 October 2008 or 27 December 2008 flights. These mechanisms including the following:

- Watchdog timers. Three timers were used to detect if the CPU was 'hung', or if there were other failures related to the CPU and software operation. If each timer was not refreshed with a new value that was written every 50 msec, then dedicated electronic circuits would shut the ADIRU down and generate ADR and IR fault messages. The software also monitored the value of the timers at intervals to ensure that they were operating and being refreshed correctly. As this did not occur during the data-spike occurrences, it was highly likely that the software was correctly refreshing the watchdog timers and that the CPU did not 'hang'.
- Memory partitions. The ADIRU's internal memory was divided into several partitions, which included the executive (main software), inertial reference input, inertial reference output, inertial reference instrument data, and air data partitions. If a software function attempted to access a partition for which it did not have access, a fault message was triggered. If the executive partition was violated, the ADIRU would stop servicing either of the watchdog timers, triggering the effects discussed above. Although the NAV IR 1 FAULT message was generated during the data-spike occurrences, the NAV ADR 1 FAULT message was not generated, indicating that there were no partition violations for the ADR and executive partitions.

3.6.3 Hardware fault

There are many potential reasons for hardware faults.[142] An example of a basic fault is a wiring connection that has broken or worked loose. The connection can remain 'open' permanently, or transition between 'open' and 'closed' with vibration or other movement. Similar effects can be produced with conductive foreign materials,

[142] For more information about the mechanisms of hardware failure in electronic equipment, see Martin (1999).

such as liquids or small loose fragments of metal that can cause short circuits between components on a circuit board as they move around.

Hardware faults can be reproduced in testing, either because they are a hard fault (such as a broken wire) or are a soft fault that can be reproduced using environmental influences (such as temperature extremes, see section 3.6.4). Extensive visual examination and functional testing was conducted on ADIRUs 4167 and 4122 and on all of the modules from each ADIRU without any relevant problems being identified (Appendix E).

A more complex fault might involve a component that has a characteristic that varies from its nominal value for some reason, such as age, environmental and other factors. For example, a resistor might exhibit a change in its resistance as it ages, or as the ambient temperature changes. These variations can change a component's behaviour in such a way that it is no longer completely compatible with other components, resulting in the reduced reliability of their interactions. These types of faults can only be triggered by a specific set of operational and/or environmental conditions, and the system will operate normally until these conditions exist.

In summary, although some aspect of the affected units' hardware was probably associated with the failure mode (section 3.5.5), it was very unlikely that a hardware fault triggered the failure mode. A much more likely scenario was that a marginal hardware weakness of some form made the units susceptible to the effects of some type of environmental factor, which triggered the failure mode.

3.6.4 Physical environment factors

The properties of electronic components are susceptible to change as a result of various mechanisms, including:

- temperature extremes and rapid temperature changes
- vibration
- contamination by liquids or particles (such as dust)
- chemical changes (including corrosion)
- mechanical stress
- electromagnetic interference (discussed in section 3.6.5)
- single event effects (discussed in section 3.6.6)

These factors may cause a temporary change (such as when an integrated circuit stops operating due to excessive temperature) or a permanent change (such as when that temperature is high enough to cause physical damage). Over time, combinations of these factors can lead to degradation of the electronic components through processes such as fatigue and corrosion.

The LTN-101 was designed to meet the requirements of various specifications (including DO-160C[143] and Airbus specifications), and units were tested during

[143] DO-160C (*Environmental Conditions and Test Procedures for Airborne Equipment*) was an industry standard produced by the RTCA for aircraft equipment manufacturers. It covered a range of different types of environmental effects on equipment (see also Appendix E).

development to ensure compliance with the specifications. The specifications included resilience to factors such as:

- physical forces (including accelerations, shock, vibration, explosion and air pressure differentials)
- temperature (including extremes, variation and loss of cooling air)
- contaminants (including sand and dust, fungus, salt spray, fluids and humidity).

Short-term environmental factors, including temperature and vibration may be reproduced in a test laboratory, and others such as corrosion, contamination and damage from mechanical stress, may be identified by detailed visual inspection. The investigation carried out a series of inspections and tests on ADIRU 4167 (and, to some extent, ADIRU 4122) to identify whether it was susceptible to, or had been damaged by, environmental factors. These activities included (see also Appendix E):

- detailed visual inspection of the units' interior and exterior, including microscope examination of the modules after their removal from the units
- vibration testing
- environmental stress screening, incorporating 15 hours of temperature cycles with power cycling at temperature extremes to attempt to induce a malfunction
- highly accelerated stress screening, incorporating temperature and vibration extremes close to the design limits of the unit
- a test in which the unit was covered in a heat shroud and run continuously while in a heat-stressed state.

No relevant faults or signs of environmental effects were observed on the ADIRUs and there was no recurrence of the data-spike failure mode throughout the testing.

Operational details of the occurrence flights were also examined and no factors were identified that would have produced any unusual environmental conditions.

3.6.5 Electromagnetic interference

Background

Electromagnetic interference (EMI) is an undesired disturbance in the normal operation of an electrical system as a result of electromagnetic emissions (often termed 'noise') from another source. Those emissions could originate from:

- other aircraft systems
- other onboard sources, including cargo and personal electronic devices (PED) carried by passengers
- external artificial sources, such as ground-based radar sites and communications facilities
- natural sources, such as electrical storms and electrostatic discharge.

Electromagnetic emissions can take the form of radiated emissions (disturbances in the surrounding electromagnetic field, such as radio waves) or conducted emissions (undesired variations or 'noise' in the voltage or current carried by a wire or set of

wires). The electric[144] field strength of radiated emissions reduces in proportion to the distance from the emitting system; that is, a doubling in distance results in the field strength being halved.

Complex computer systems may be affected by EMI in ways that are difficult to characterise, predict or replicate. EMI tends to affect interfaces between computer chips, or between systems, rather than the internal interfaces within a computer chip. However, those interfaces can disrupt the normal operation of a chip, such as if the chip's power supply is abnormal or if the interface between a CPU chip and memory chips is subjected to interference. For example, a computer may shut down or 'hang', or otherwise behave unpredictably, if one of its internal interfaces is disrupted (Hanson 1987).

A unit's susceptibility to EMI is highly dependent on the characteristics of the emitting and receiving systems, the aircraft itself, and other aspects of the electromagnetic environment. As such it can be difficult, if not impossible, to exactly reproduce those conditions during testing or to completely evaluate the potential for EMI to occur.

Further background information on EMI is provided in Appendix G.

Design requirements and certification

European and US airworthiness requirements included requirements for an aircraft and its equipment to be designed to limit their level of electromagnetic emissions and to ensure that they were resistant to EMI. The applicable industry standard during the development of the A330/A340 was DO-160C, and the aircraft and its systems (including the ADIRUs) were designed, tested, and qualified in accordance with this standard.

DO-160C specified EMI resilience in terms of conducted current or radiated electromagnetic fields at different frequencies. The ADIRU and its associated wiring were required to withstand 150 milliamperes (mA) of conducted currents between 10 kHz and 400 MHz and 100 V/m radiated electric field strength between 400 MHz and 18 GHz.

Potential sources from other aircraft systems

Measurements were taken of the electromagnetic environment on the aircraft that was involved in the 7 October 2008 occurrence (QPA) on the ground and in flight. These measurements showed that the combined emissions of the other systems did not exceed the DO-160C susceptibility limits in the area around the ADIRUs or in the ADIRU wiring (Appendix F).

ADIRU 4167 and an exemplar unit were subjected to EMI testing at a range of different frequencies, and no problems were noted (Appendix E). The testing covered the frequencies of other aircraft systems that were considered to be a possible source of EMI, including the aircraft's electrical power supply (which had the potential to produce stronger radiated and conducted emissions than most other

---

[144] An electromagnetic wave consists of a changing electric field associated with a changing magnetic field. The two fields are always in the same proportion when in the same medium and when not very close to an emitting source, so in a known medium such as in air, only one of the two fields generally needs to be specified. For example, DO-160C specifies only the electric field strength.

systems), radio transmitters, and in-flight entertainment system. In addition, there were also no faults reported with these systems during the occurrence flights, or notable problems reported in the occurrence aircraft's recent maintenance records.

Other potential onboard sources of EMI

The cargo manifests for the three occurrence flights were examined for items that might have been possible sources of electromagnetic interference and none were identified. Following the 7 October 2008 occurrence, all of the cargo was removed from the aircraft and inspected for items that might have been possible sources of electromagnetic interference. None were identified. The investigation did not determine whether any of the passengers' luggage could have included powered electronic devices, although it was considered unlikely that such devices could have provided a significant level of emissions.

The investigation surveyed passengers and flight crew from the 7 October 2008 flight about whether any PEDs were in use during the flight, with a particular focus on devices such as mobile phones and laptops that were capable of transmitting signals.[145] Although any electronic device can cause interference, those which transmit via radio waves create the highest risk. Overall, only a small number[146] of PEDs were reported to be on at the time of the occurrence. All were reported to be in the appropriate flight mode, and no problems were reported with the performance of the devices.

Expert advice was obtained from a major telecommunications company regarding the aircraft locations at the time of the three events, the ranges from the mobile phone base stations[147], the performance characteristics of mobile phones when in- and out-of-range of a base station, and the network transmission technology in use at the time (that is, CDMA and GSM[148]). The company advised that the location of the 12 September 2006 occurrence was too far from the nearest base station for a mobile phone to have been useable on board the aircraft, and the 7 October 2008 occurrence was possibly but unlikely to have been within range. Mobile phone activity through base stations closest to the location of the 7 October 2008 occurrence were reviewed to identify any evidence that a mobile phone on the aircraft may have been making or receiving transmissions, and no such evidence was identified.[149]

---

[145] All the crew were interviewed about a range of topics, including PEDs. The passenger questionnaire asked passengers about any PEDs they were using, or that were in use by other passengers nearby. Further details on the questionnaire are provided in Appendix I.

[146] Based on passenger questionnaires and interviews, there were at least seven laptops in use during the flight, three mobile phones (in the appropriate flight mode), and at least two other electronic devices. Almost all of these devices were being used in the centre or rear sections of the cabin.

[147] Mobile telephones generally do not transmit unless they are within range of a base station, or if they are connecting to a nearby device via a short-range wireless technology such as Bluetooth. Other PEDs, including laptop computers and hand-held games, can also produce significant emissions through similar technologies.

[148] CDMA: Code division multiple access. GSM: Global System for Mobile Communications.

[149] This search would not have identified mobile phones that were turned on and not in flight mode, and which did not or receive a call or message during the period of interest.

Although the investigation could not eliminate the possibility that an unreported, transmitting PED was been on board the aircraft at the time, ADIRU testing (Appendix E) and aircraft testing (Appendix F) covered all of the identified frequencies of concern at much higher power levels than a PED would normally be capable of producing. No problems were noted. The ADIRUs were also located in the avionics bay, approximately under the front galley area and a significant distance away from most of the passengers.

If a significant EMI source was on the aircraft and producing emissions sufficient enough to affect the ADIRU, it may have had observable effects on other systems. However, there were no faults or problems reported by any other system indicating that it had been affected by EMI.

Potential external artificial sources of EMI

As shown in Figure 26, the three known data-spike occurrences all took place within 1,000 km of Learmonth. Given that the locations were in a broadly similar geographical area, the investigation examined potential sources of EMI in the area. Two potential sources were identified: the Harold E. Holt Naval Communication Station located at Exmouth (near Learmonth), and a nearby high frequency (HF) radio communications site.

The naval communication station transmitted at a very low frequency (VLF) of 19.8 kHz, and the transmission power was about 1 megawatt using an omni-directional antenna. The station was transmitting at the time of the three data-spike occurrences. However, it was considered extremely unlikely that these transmissions had any effect on the ADIRUs for several reasons:

- Estimated electric field strengths as a result of station's transmissions at the locations of the three occurrences were several orders of magnitude below the level at which the ADIRU model was tested during certification (Table 25).[150]
- The station had used the 19.8 kHz frequency for over 10 years, and it transmitted almost continuously with the exception of weekly maintenance periods. A large number of flights involving aircraft fitted with LTN-101 ADIRUs have been conducted near the station, and no problems have been noted.
- The Australian Department of Defence advised that there were no problems or malfunctions with the transmitter near to or during the times of the three occurrences, and that there had been no changes in the nature of the station's transmissions in recent years.
- Similar VLF transmitters are also located in other countries such as the US, UK, China, France, India, Japan and Russia. A large number of flights involving aircraft fitted with LTN-101 ADIRUs have been conducted near those stations, with no interference problems known to the investigation.
- ADIRU 4167 was tested for conducted susceptibility at the relevant frequency (19.8 kHz) and at levels that far exceeded those to which it was exposed during

[150] FAA advisory circular AC 20-158 gives an equivalent bulk current of 0.1 mA per V/m at 19.8 kHz. The equivalent bulk current at the time and location of the 7 October 2008 occurrence was therefore 0.0059 mA. To comply with DO160C, aircraft systems were designed and tested to withstand 150 mA, a level 25,000 times the current to which an ADIRU would have been subjected on 7 October 2008 due to the Harold E. Holt station.

either of its data-spike occurrences, and no problems occurred (Appendix E). An exemplar unit was also tested to a higher level of EMI at the same frequency and no problems were noted.

- The aircraft that was involved in the 7 October 2008 occurrence was flown over the station while it was transmitting. Measurements of the EMI environment on the aircraft near the ADIRUs were undertaken, and no measurable influence from the station was observed (Appendix F).
- EMI effects on avionics are generally associated with higher frequencies than those emitted by the station.

Table 25: Estimated electric field strengths as a result of transmissions from the Harold E. Holt Naval Communication Station

Occurrence	Aircraft	Approximate distance to station (km)	Approximate electric field strength (V/m)
12 Sep 2006	VH-QPA	950	0.011
7 Oct 2008	VH-QPA	170	0.059
27 Dec 2008	VH-QPG	700	0.014

The high frequency (HF) radio communications site located on North West Cape near Learmonth can transmit signals in the HF frequency band (3 to 30 MHz) at a signal power of 10 kilowatts or less. Records indicated that the site was transmitting at the time of the 12 September 2006 and 27 December 2008 events. It was not transmitting at the time of the 7 October 2008 event. In addition, no problems were noted with that frequency range during the ADIRU testing.

The investigation also considered the potential effect of EMI from a transmitting satellite that could have passed above the aircraft at around the time of the events. Calculations were performed using a hypothetical satellite in the lowest feasible orbit and having a relatively high transmitted power. For several reasons, it was considered extremely unlikely that such satellite transmissions could have had any effect on the ADIRUs.

Potential natural sources of EMI

The investigation did not identify any natural sources of electromagnetic radiation particular to the regions in which the three events occurred. For example, during the 7 October 2008 occurrence, the aircraft was well south of any convective (potentially lightning-inducing) meteorological activity, there was no reported turbulence, and there was no lightning activity in the vicinity.

Static charge can build up between the aircraft and its environment or between different parts of an aircraft. If the charge is sufficiently high, an ‘electrostatic discharge’ (such as arcing) may occur, producing electromagnetic fields that can disrupt electronic equipment. The investigation could not determine whether any significant electrostatic discharging had occurred during the flights. However, as previously noted, there was no evidence that EMI affected any other aircraft systems, and no unusual conditions were identified that could have been present during the events and contributed to a significantly abnormal electromagnetic environment.

3.6.6 Single event effects

Background

There is a constant stream of high-energy galactic and occasional bursts of solar radiation interacting with the Earth's upper atmosphere. That interaction creates a cascade of secondary particles, and some of those particles (particularly neutrons) can affect aircraft avionics systems. A single event effect (SEE) is the response of an electronic component to the impact of a single particle. Unlike EMI, which affects interfaces between chips, SEE affects a chip directly.

The most common type of SEE, a single event upset (SEU), occurs when the particle deposits an electric charge inside a memory cell that is sufficient to change the cell's logic state (known as a 'bit flip'). A range of other types of SEE can also occur. Although some SEEs can result in permanent damage to components, most are soft faults.[151]

High-density integrated circuits, such as memory and CPU chips, can be particularly susceptible to SEEs due to their relatively large number of memory locations and the reduced 'feature' size (that is, functional parts in newer chips generally become smaller and more sensitive). SEEs have been confirmed in space-, air- and even ground-based computer systems, and have been shown to generate soft errors in a wide range of different aircraft systems.

The probability of an SEE occurring to a particular component at any particular point in time is dependent on a range of factors, including the number, energy, and type of particles in the area of the component's operation and the component's sensitivity to SEE. Other relevant factors include the direction and location of the particle strike, and the time in the device's program cycle at which the strike occurs.

Most SEEs have no adverse effects on a system's performance. For example, changing the logic state of a bit within a data word would change the value of that word, but only until the word was refreshed with a new value.

More adverse effects occur when the cells that are affected include program data; that is, information critical to the successful execution of the core software, such as software instructions, CPU register information or program pointers. For example, software might 'jump' to the wrong location in a program, or even to a location outside of program memory and treat the data there as valid program instructions. Most of this critical information is stored and used internally within the chip. As the successful execution of software depends heavily on program 'states' (that is, the tasks being performed at any given point in time), such a disruption can have wide-ranging and complex outcomes. The most common type of adverse effect in these cases results in the system 'hanging', although many other adverse states are also possible and may not be predictable in practice.

System manufacturers use a variety of techniques to reduce the effects of unintended changes of logic states, whether they are due to an SEE or other source. These include hardware and software design features such as redundancy, monitoring and partitioning. A specific technique to reduce the effects of SEE is

---

[151] In the SEE field, a 'soft fault' occurs when the logic state of a cell is changed to an invalid value, and a 'firm fault' refers to a failure that can only be reset by rebooting or cycling the power to the system. In this report, both of these phenomena are termed soft faults.

error detection and correction (EDAC), which involves using redundant memory information so that a 'bit flip' can be detected and the data either ignored or corrected.

Further background information about SEEs is provided in Appendix H.

Design standards and certification

SEE is a particularly serious concern to designers of space-based systems such as satellites, mostly due to the extreme reliability requirements of such systems. Until recently, there was limited formal recognition of the effects of SEE on airborne systems. For example, during the development of the A330/A340, there were no specific regulatory or aircraft manufacturer requirements for airborne systems to be resilient to SEE.

In December 1995, the aircraft manufacturer issued ABD100 (*Equipment – Design – General requirements for suppliers*), which incorporated numerous design requirements. The document stated that 'hardware and software implementation solutions shall consider the failure rate possibility of SEU (Single Event Upset) due to particle environment (radiations as for example: protons, neutrons, heavy ions...) at high flight altitude'. Recommended design activities to limit the SEU rate included limiting the RAM solutions, and assessing the 'risk of RAM memories' and the 'register part of microprocessors' (that is, memory locations internal to the processor chip that are common to all software functions and can therefore have a very high probability of an adverse outcome if they are corrupted). The 1995 version stated that the SEU rate risk was to be included in assessments of the equipment's mean time between failures (MTBF). A subsequent version (December 1996) stated that both the reliability requirements and safety requirements needed to met after taking account of the SEU risk.

In 2006, the International Electrotechnical Commission (IEC)[152] published Technical Specification (TS) 62396-1 (*Process management for avionics – Atmospheric radiation effects – Part 1: Accommodation of atmospheric radiation effects via single event effects within avionics electronic equipment*). The specification provided guidance on atmospheric radiation effects on avionics systems, and design considerations for the accommodation of those effects within avionics systems.[153]

In December 2007, the aircraft manufacturer updated ABD100 to include reference to TS 62396. It also specified that an equipment supplier must perform various SEE analyses and provide the aircraft manufacturer with certain information about the SEE behaviour of the system.

During the investigation, the European Aviation Safety Agency (EASA) advised that SEU and multiple bit upset[154] (MBU) were relatively new phenomena that, up

---

[152] The IEC is a worldwide organisation with the object of promoting international cooperation on standardisation in the electrical and electronic fields. It publishes a range of documents including International Standards and Technical Specifications.

[153] The IEC also issued subsequent parts to TS 62936 that dealt with more specific topics relating to SEE.

[154] A MBU occurs when the energy deposited by a single particle causes an upset to more than one bit in a device.

until a few years ago, were not directly considered in the aircraft certification process. It stated that during the certification of the A330/A340, there was no reference to SEU or MBU and that it was unlikely that the phenomena were directly covered. However, it noted that the system architectures were designed to protect the aircraft from single-point failures by the use of design techniques, such as command and monitor channels and voting mechanisms.[155]

The US Federal Aviation Administration (FAA) advised during the investigation that there were no current standards specifically targeting SEE faults. However, it stated that design assurance, fault tolerance and system safety assessments were required under FAR 25.1309, and that these requirements were intended to provide tolerance to all foreseeable single failures regardless of their origin, and such failures would implicitly include SEE.

As noted in section 3.5.2, the CPU module on units 4167 and 4122 did not incorporate EDAC, nor was it required by the aircraft manufacturer's specification. However, the ADIRU did include some features that could detect and mitigate the effects of a wide range of failures, including many failures that could have been triggered by SEE. For example, the ADIRU's program memory was periodically checked for integrity, and that check would have failed if the program memory had been corrupted. This and other BITE functions are discussed in section 3.7.

Space weather at the time of the occurrences

The ATSB requested the Bureau of Meteorology's (BOM) Ionospheric Prediction Service to examine the space weather[156] at the time and location of the three occurrences (12 September 2006, 7 October 2008 and 27 December 2008). Geomagnetic observations from Learmonth and Darwin and cosmic radiation observations from Hobart were examined. The conclusion by BOM was that the space weather for the three occurrences was within the normal ranges for the relevant locations.

Although the neutron fluxes[157] for the three occurrences were at around the normal levels, this still meant there were a significant number of particles present. The nominal high-energy (greater than 10 million electron-volts) neutron flux is about 5,600 neutrons per cm^2 per hour at 40,000 ft and 45° latitude (North or South).[158] There is a very high variation of neutron flux with altitude (about 300 times higher at 40,000 ft than at ground level) and a minor variation with latitude (about four times higher at the poles than at the latitude of the 7 October 2008 occurrence). The

---

[155] EASA advised that during the development of the A350 and A380 it requested Airbus to consider the effects of SEU and MBU on systems and equipment, and Airbus required equipment manufacturers to consider the effects of SEU and MBU and to mitigate these effects.

[156] 'Space weather' refers to the variable environmental conditions in near-Earth space. It is distinct from the concept of weather within the atmosphere, and deals with phenomena involving plasma, magnetic fields, radiation and other matter in space. Space weather can have consequences at ground level as well as effects in the upper atmosphere.

[157] Neutron flux is the number of neutrons passing through an area during a period of time. The unit of neutron flux used in this report gives the approximate number of neutrons that pass through an area of one square centimetre (cm2) in 1 hour.

[158] This figure was stated in IEC TS 62396 (discussed later in this section). High-energy neutrons are those with an energy of 10 million electron-volts or more. The flux for neutrons between 1 to 10 million electron volts is about 3,000 per cm^2 per hour for the same location.

estimated high-energy neutron fluxes present at the time and location of the three occurrences were 700 (12 September 2006), 1,000 (7 October 2008) and 1,700 (27 December 2008) neutrons per cm^2 per hour.[159]

The most useful method for ascertaining whether a type of equipment behaviour is related to SEE is to use correlations of the rate of events compared with the predicted neutron flux, as well as taking into account other hypotheses and information about the system of interest. This approach was of limited value in this case due to the small number of occurrences involved.

Component testing

Similar types of RAM chips to those used in the LTN-101 ADIRU were tested for their susceptibility to SEU in 1993 by specialists contracted by the ADIRU manufacturer. Both types of RAM were found to be susceptible to SEU. For the CPU RAM the equivalent[160] estimated mean time between SEU operated in a polar region at 45,000 feet was about 75 days. For the wait-state RAM, the mean time was 175 days. The test report recommended that the manufacturer consider various means of improving resilience to SEE, including the substitution of more resilient RAM and the use of EDAC.

Three variants of a chip[161] similar to the CPU RAM type and two variants of the wait-state chip type[162] were tested for proton[163] and heavy-ion[164] SEU susceptibility in 1994. All chips tested were reported to be 'very susceptible to both heavy ions and protons' though they exhibited different levels of susceptibility.

The wait-state RAM type was tested in 1999 for susceptibility to heavy-ion bombardment, and was compared with two other types of RAM. Of the three types of RAM tested, that used in the LTN-101 was the most resilient to heavy-ion SEU.

There were no SEE tests of the CPU chip or ASIC chip known to the investigation.

During the investigation, the ATSB received expert advice that the susceptibility of a specific component to SEE could vary significantly between components with the same part number from the same batch (up to a factor of 2 or 3), and even more for components from different batches (up to a factor of 10). The expert advised the ATSB that the component-level test results were typical of devices of the period.

---

[159] These figures were calculated by establishing the level of solar modulation of galactic cosmic rays from ground-level monitors for the time of the occurrences and applying the QinetiQ Atmospheric Radiation Model (QARM) to calculate the in-flight levels. See http://qarm.space.qinetiq.com.

[160] The mean time between SEU has been adjusted here to account for the different amount of RAM in the study versus the amount of RAM in the LTN-101 ADIRU program memory.

[161] The chips tested were Micron Tech MT5C2568, MT5C2568-35 and MT5C2568-70, with somewhat similar characteristics to the Austin chips that were used in the LTN-101.

[162] The chips tested were Mosaic MSM8128K and MSM8128KL, with very similar characteristics to the Mosaic chips that were used in the LTN-101.

[163] The effects of proton and neutron SEE are comparable, so SEE testing using one is generally representative of the other.

[164] A heavy ion is an ionic (charged) particle heavier than a helium nucleus; that is, more than about twice as heavy as a neutron. High-energy heavy ions are present in the atmosphere and can cause SEE, but do not penetrate the atmosphere or aircraft skin as deeply as high-energy neutrons. At normal aircraft altitudes, neutrons are the predominant form of high-energy particle.

Unit testing

In 2005, as part of the ADIRU manufacturer's investigation of another LTN-101 ADIRU failure mode (known as 'dozing', see section 3.9.3), testing was conducted to determine the LTN-101 model's susceptibility to neutron SEE. The neutron fluxes used in the test were very high (billions of neutrons per cm² per second) to emulate neutron exposure over long periods of normal operation. The following LTN-101 ADIRU versions were tested:

- ADIRU with the newer CPU module incorporating EDAC (part number 466871) and software version -0315. Four test runs were performed that resulted in several different anomalies, most commonly a 'system functional shutdown' (or hanging). The calculated mean time between failures (MTBF) was about 6,200 hours, with failure defined as the system shutting down.
- ADIRU with the same type of CPU module as on units 4167 and 4122 (part number 465474) and software version -0315. One test run was performed that demonstrated a corruption of the ADIRU software after about 1 minute of neutron bombardment. The calculated MTBF for this sample was about 30,300 hours.

The test report concluded that SEE had 'a real and varied effect' on the LTN-101's operation. During the investigation, the ADIRU manufacturer advised that the functional shutdowns that were generated during the testing were similar to dozing events but had different BITE data signatures (see section 3.9.3 for more details on dozing events). An expert on SEE advised the ATSB that the results were typical of systems of the period.

None of the tests generated data spikes or other data output anomalies. However, only limited testing was done, with only one test run on a CPU module with the same part number as unit 4167 and 4122. The similarity of the CPU module to the modules in the affected units (in terms of the batch numbers of key components) could not be determined. The testing was conducted with neutrons at 14 million electron-volts, which was in the 'high-energy' range as defined by IEC 62396 but towards the low end of that range. Consequently, the test was not sufficient to examine the model's susceptibility to the full range of neutrons at the higher energy levels that exist in the atmosphere. For example, higher energy particles are more likely to trigger a MBU, which can produce different system effects than a SEU. In addition, the data-spike failure mode might not have been triggered due to its relative rarity compared with other types of effects.

During the investigation, the parties to the investigation of the 7 October 2008 occurrence discussed the advantages and disadvantages of additional SEE testing of LTN-101 ADIRUs, particularly unit 4167. The ATSB received expert advice that the best way of determining if SEE could have produced the data-spike failure mode was to test the affected units at a test facility that could produce a broad spectrum of neutron energies. However, the ADIRU manufacturer and aircraft manufacturer did not consider that such testing would be worthwhile for several reasons, including that:

- A level of susceptibility to SEE had already been demonstrated through previous unit and component testing.
- Design changes had been made to the later production ADIRUs to incorporate EDAC in the CPU, which would mitigate the effects of SEE.

- There were significant logistical difficulties in obtaining access to appropriate test facilities and developing test software and procedures, and another option to assess the resilience of the unit to SEE was considered more practicable (see *Theoretical analysis* below).
- The repeated occurrence of a more likely failure mode such as a system functional shutdown might prevent a less likely failure mode such as a data-spike occurrence from being observed over a limited time period.

Theoretical analysis

As a result of the data-spike occurrences, the ADIRU manufacturer reviewed the effects of SEU on the reliability predictions given in the LTN-101 FMEA (section 3.8.2) and the safety objectives listed in the aircraft manufacturer's equipment specification (section 3.8.1). The review focused on the effects of SEU on the RAM components of the CPU module, and was conducted on the LTN-101 version with the same CPU module part number as units 4167 and 4122 (that is, with no EDAC).

The theoretical analysis concluded that the SEU rate was 1.96×10^{-5} per hour (or one upset in every 51,020 hours). In addition, the analysis concluded that, after considering the SEU rate, the overall predicted failure rates for the ADIRU would not change significantly. More specifically:

- The predicted MTBF for the ADIRU decreased from 15,212 hours to 11,948 hours (see also section 3.9.1 for further discussion of MTBFs).
- A SEU would increase the detected failure rate but not affect the undetected failure rate (that is, all of the relevant potential failures identified in the LTN-101 FMEA were found to be detected by BITE).

With the inclusion of SEU as a potential failure mechanism, all of the safety objectives (section 3.8.1) were still satisfied. However some were only marginally satisfied and the predictions relied on several assumptions and large tolerances in the test data.[165]

Some types of SEE can produce quite different and more serious effects on system performance than other SEEs. Although the theoretical analysis involved reviewing the ADIRU's FMEA to identify specific failure modes that could be initiated by an SEU, it did not consider other types of SEE (like a MBU). In addition, none of the specific failure modes identified by the FMEA were considered to be relevant to the data-spike failure mode. Without knowing the actual mechanism involved in initiating the data-spike failure mode, the extent to which it could have been triggered by SEE could not be determined.

In-service SEE history

LTN-101 ADIRUs had previously exhibited anomalous behaviour that was attributed to a SEU. For example, the operator of QPA had experienced 116 events

---

[165] For example, there were significant differences in the SEU rates given in the component test data (for example there was a factor of 3 difference in the case of the Mosaic RAM). In addition, as SEU susceptibility can vary from component to component, there was no guarantee that the SEU test values were representative of the susceptibilities of the actual components used in ADIRUs 4167 and 4122.

that involved a NAV IR and/or NAV ADR fault message during the period from 2003 to 2008, excluding the data-spike occurrences (section 3.9.2). Of these, the ADIRU had been removed for examination on 24 occasions. The ADIRU manufacturer had attributed the problem to a SEU on two occasions.

One of these occasions involved a NAV IR 1 FAULT message on ADIRU 4122 on 18 July 2008. The BITE data showed that a checksum fault had occurred (that is, the BITE detected that the copy of operational software stored in read-only memory was different to the version loaded into RAM when the ADIRU was in operation). The ADIRU manufacturer reported that they had records of about 100 similar events on LTN-101 units since 2000, and that similar errors had resulted during the unit testing in 2005.

A high proportion of the 116 events were soft faults that resulted in the ADIRU shutting down. Insufficient information was available for most of these events to determine the origin of the faults.

The investigation considered the likelihood of two events occurring on a single unit. A typical computer silicon chip, or integrated circuit, is primarily made of millions of tiny devices called transistors, each of which operates by accumulating and distributing electric charges within the chip. At a nominal rate of 6,000 neutrons per cm^2 per hour, a device with similar susceptibility and capacity to the CPU RAM on the LTN-101 ADIRU was estimated to have an SEE every 75,000 operating hours[166], with the effects depending on the transistor or transistors affected. Accordingly, it is extremely unlikely that two separate SEEs involving the same transistor would occur on a single unit. However, more than one transistor could produce the same failure mode if affected by SEE, and this would increase the likelihood of the effect occurring more than once in the life of a unit. In addition, a particular unit or group of units may have greater susceptibility relative to other units. Overall, the likelihood of two instances of a rare type of SEE occurring on one unit due to separate strikes on a small number of transistors was difficult to accurately estimate but it was not considered negligible.

3.6.7 Summary

The investigation examined a range of potential triggers that may have initiated the data-spike failure mode within the CPU module, and key points for each are summarised in Table 26. Although a definitive conclusion could not be reached, there was sufficient information from multiple sources to conclude that most of the potential triggers were very unlikely to have been involved.

---

[166] Typically, an A330 might fly 4,500 hours in a year (with 70% of the total time spent at cruise levels where neutron flux is highest). At that rate, 75,000 flight hours would equate to about 17 years of aircraft operation. This is a 'best case' analysis as the flux value does not take into account particles other than neutrons or neutrons at lower energy levels, or that a single neutron may be able to pass through multiple transistors (that is, a side-on strike). It is possible that the combined effects of these tolerances could reduce this period of aircraft operation by at least one or two orders of magnitude.

Table 26: Evaluation of potential triggers

Trigger	Key points	Assessment
Software corruption	ADIRU software was verified as intact after the occurrences. Unit 4167's software was reloaded and verified between the two occurrences involving this unit.	Very unlikely
Software 'bug'	Would not be expected to occur twice on one unit without many other occurrences on other units. Functional testing of software found no problems. No unique circumstances identified with the occurrence flights that could trigger a rare bug.	Very unlikely
Hardware fault	Extensive unit and module testing found no problems. Visual examination of the units did not identify any physical damage or other abnormalities. Not consistent with a 'soft fault'.	Very unlikely
Physical environment	Unit testing beyond relevant standards found no problems. Visual examination of the units did not identify any physical damage or other abnormalities that could result in a relevant equipment fault when exposed to normal or abnormal environmental conditions. The physical environment was normal during the three flights. Nothing unusual found with aircraft environment during testing.	Very unlikely
EMI from aircraft systems	Extensive unit testing found no problems. Measurement of the electromagnetic environment within the aircraft during ground and flight tests showed nothing unusual or excessive. It was not possible to reproduce the exact conditions of the occurrence flights during testing. Wiring integrity tests found no problems. The aircraft configuration was not unique or unusual. No problems with the other ADIRUs installed on same aircraft.	Unlikely
EMI from other onboard sources	No sources of concern were identified. Extensive unit testing found no problems. Measurement within the aircraft while PEDs were in use showed very minor effects on the electromagnetic environment.	Very unlikely
EMI from external sources	No sources of concern were identified. Extensive unit testing found no problems. The electromagnetic environment during flight tests showed nothing unusual or excessive. No problems with other systems during the occurrence flights.	Very unlikely
SEE	The intensity of high-energy particles for the three occurrences was not unusual. The ADIRU had limited mechanisms to detect and manage SEE (that is, no EDAC). No SEE testing was performed on the occurrence units. SEE testing on another unit did not induce the data-spike failure mode (although the testing was limited in scope). Difficult to accurately estimate the likelihood of two SEEs occurring on the same ADIRU twice in its operational life.	Insufficient evidence to estimate likelihood

EMI from other aircraft systems (such as the aircraft's power supply or radio transmitters) was more difficult to discount than EMI from other sources. This is because it was not possible to determine or reproduce the exact electromagnetic environment that existed within the aircraft at the time of the occurrences, and because other aircraft systems had the potential to produce much stronger emissions than the other sources. However, there was still sufficient information available from the unit testing, aircraft testing and other sources to conclude that this option was still unlikely.

The other potential trigger that was relatively difficult to discount was SEE. In the absence of testing information from the affected units, there was insufficient information to make a definitive conclusion.

3.7 ADIRU built-in test equipment operation

3.7.1 Background

A wide range of ADIRU functions and components were monitored and/or tested by the BITE, including:

- ADM status
- ADM and external sensor data integrity
- inertial sensor environment and data integrity
- input databus integrity
- ADR output data consistency (between ADIRUs)
- air data and inertial data analogue-to-digital conversion
- volatile and non-volatile memory integrity
- BITE memory integrity
- power supply quality and status
- CPU operation and arithmetic
- output parameter range
- output databus integrity.

As part of the safety analysis that was conducted during the development of the LTN-101 (section 3.8), the ADIRU manufacturer estimated that the BITE provided the following failure detection rates (that is, the percentage of failures that the ADIRU would detect):

- 98.5% of CPU module failures
- 94.5% of air data input/output module failures
- 94.8% of inertial reference input/output module failures
- 96.0% of power supply failures
- 92.1% of monitor module failures
- 93.5% of failures overall.

The ADIRU's response to a detected fault depended on the nature of the fault, but usually involved isolating the affected function(s), sending a fault message to the flight warning system (FWS), sending a fault message to the central maintenance system (CMS), and flagging any affected parameters with an invalid SSM.

The BITE tests were performed during initialisation of the system (power-on) and most were performed continuously during normal operation.

BITE information was designed to be recorded in non-volatile memory so that it could be analysed if required. In addition to fault detection, the BITE also stored routine maintenance information, such as system temperature, time in operation, inertial alignment records, flight leg records, instrument trends, and navigation updates and performance.

3.7.2 Flight data parameter tests

The investigation reviewed the tests that were conducted by the ADIRU on the flight data parameters to determine if any of the tests could have detected the data-spike failure mode. The most relevant tests were the range tests conducted on many of the parameters prior to them being transmitted to other systems. The available evidence showed that these range checks were working correctly, with groundspeed clamped to 1,000 kts and computed airspeed set to 0 kts when the speed reduced to less than 30 kts (section 3.4.3). No other reasonableness checks were performed on the output data, nor were they required in the aircraft manufacturer's specification.

Other tests were also conducted on various parameters. For example, tests conducted on AOA values were:

- AOA integrity monitor. When Mach was greater than or equal to 0.75, and with the aircraft straight and level (determined using two IR parameters), the corrected AOA was compared to the pitch attitude. If the difference between them was greater than or equal to 0.6°, a class 3 maintenance message was transmitted to the CMS. No such messages were recorded during the data-spike occurrence flights. However, after the data-spike failure mode was initiated, the IR parameters were flagged as invalid (section 1.11.3), which would have inhibited the operation of this check. In any event, a class 3 fault would not result in a warning or caution message being provided to the flight crew (section 1.6.10).
- Air data sensor input comparison monitor. During takeoff, each ADIRU compared its air data parameters, including indicated AOA, with the values from the other two ADIRUs using information shared across the digital air data system (DADS) buses (section 1.12.6). For AOA, if the difference between the indicated AOA values exceeded 3°, then a class 2 maintenance fault message was generated. The test was not conducted when the aircraft was airborne.

In common with the AOA checks, none of the specific tests that were conducted for the other ADR flight data parameters would be expected to generate a warning or caution message to the flight crew in the event of a data-spike occurrence.

There were many tests on IR parameters that, had one failed, would have resulted in the IR parameters being flagged as invalid. For example, the output range monitor for groundspeed would flag IR parameters as invalid if groundspeed exceeded the range-limit value of 1,000 kts. The QAR data showed that groundspeed had exceeded 1,000 kts and been clamped to that value (section 1.11.3). As BITE data

from ADIRU 1 was not correctly recorded, the investigation was unable to determine all of the failed tests and when they initially failed.

3.7.3 ARINC 429 databus wraparound tests

The ADIRU included a test (or dummy) parameter among the normal operational data that was outputted on the databuses. The test parameter's data consisted of an alternating pattern of '1's and '0's that was reversed in every second transmission. The ADIRU then read the data from the databus and compared it with what was sent. The test failed if the transmitted and received data disagreed or if no data was received.

The wraparound test was effectively a hardware test of the ARINC 429 transmitters and receivers in the air data and inertial reference input/output modules; it was not a check of whether the CPU module calculated, temporarily stored, and then outputted data correctly to the input/output modules.

The eight ADR output buses were tested in turn at a rate of eight tests per second, and the four IR output buses were tested in turn at a rate of 50 tests per second. Detected problems could result in various consequences as follows:

- If three tests failed in a row on any one bus, a class 2 maintenance fault message was transmitted to the CMS (either 'MAINTENANCE STATUS ADR 1' or 'MAINTENANCE STATUS IR 1'). As discussed in section 1.6.10, class 2 messages were not displayed to the crew.
- If three tests failed in a row on all of the ADR buses, then a class 1 maintenance fault message was transmitted to the CMS resulting in a NAV ADR [1, 2 or 3] FAULT message being displayed to the crew. The ADIRU also set the SSM of all the ADR outputs to 'failure warning' and illuminated the ADR fault light on the overhead panel.
- If three tests failed in a row on all IR buses, a class 1 maintenance fault message was transmitted to the CMS resulting in a NAV IR [1, 2 or 3] FAULT message being displayed to the crew. The ADIRU also set the SSM of all the IR outputs to 'failure warning' and illuminated the local IR fault light.

The LTN-101's wraparound tests were only conducted on the test parameter. Wraparound tests were not conducted on the flight data parameters used by other systems, nor were they required in the aircraft manufacturer's specification. Another ADIRU manufacturer advised that its ADIRUs' wraparound tests were also only conducted on a test parameter.

For the 7 October 2008 occurrence, a class 2 maintenance fault message for ADR 1 occurred at 0440 as soon as the data spikes commenced (section 1.12.2). It is likely that this message occurred due to the wraparound test self-detecting a temporary problem on some, but not all, ADR databuses. BITE data from ADIRUs 2 and 3 (section 1.12.6) indicated that there was a 'refresh' problem with data from ADR 1 at 0440. In addition, FDR data indicated that the ADR Fail was intermittently active at 0440, consistent with a refresh problem. Therefore, it is likely that the wraparound test also detected a problem with the refreshing of data. The problem must have cleared before the conditions for a class 1 message (and a NAV ADR 1 FAULT) were satisfied.

3.7.4 False fault messages

The ADIRU provided several fault messages that were considered to be spurious (section 1.12.2), including fault messages associated with the probe heat system (section 1.12.9). In addition to flight data parameters, the ADIRU outputted words containing fault data (such as the ADR discrete word #1 parameter), and these words were outputted on the same ADR databuses that outputted the flight data parameters. As indicated in section 3.3.4, a data exchange between the ADR discrete word #1 parameter and altitude or another parameter could account for the probe heat faults and overspeed warnings transmitted by the ADIRU.

An alternative explanation for the class 2 maintenance fault message for ADR 1 was that it was a spurious fault message similar to the probe heat faults. The ADIRU reported the class 2 message to the FWS using a bit on the ADR discrete word #1 parameter (see section 3.3.4 and Figure C.1). If the label from this parameter was combined with the data from another parameter, the message may have been triggered. However, in all three occurrences the class 2 maintenance fault message was generated at the beginning of the failure mode and prior to the probe heat fault messages, suggesting that they were caused by different mechanisms.

3.7.5 Summary

There was evidence that during the data-spike failure mode, at least some of the ADIRU's BITE tests were working as per their design. For example, the IR parameters were flagged as invalid and there was a corresponding IR 1 fault message, and some of the range checks were operating effectively.

There was no evidence that any of the BITE tests did not operate as designed. The storage of BITE data did not function correctly, as expected fault messages (such as those that caused IR parameters to be flagged as invalid) and some routine messages were not recorded. However, this was a characteristic of the data-spike failure mode and did not mean that there was any problem with the operation of the BITE tests. There were also some false fault messages, which appeared to be attributable to the same mechanism that produced the data spikes.

3.8 ADIRU safety analysis

3.8.1 Safety analysis requirements

The Airbus ADIRS equipment specification outlined minimum safety objectives for the ADIRU, which were based on the Airbus system safety assessment for the aircraft (section 2.4.5). Most of the safety objectives related to the loss of more than one ADR and/or IR. In terms of an individual unit, one of the objectives stated that an 'undetected' failure of one IR or one ADR must be less than 3×10^{-5} per flight hour.[167]

The specification included the following definitions:

- Failure. The unit no longer performed 'the functions for which it was designed within the specified performance'.

[167] This equated to a rate of no more than one 'undetected' failure in every 33,000 flight hours.

- Detected failure. Any failure where the ADIRU flagged its output data as invalid, or there was a complete cessation of output data.
- Undetected failure. 'Any failure which can result in incorrect data transmission on one or several ARINC word(s)..., and which [was] not self-detected by the ADIRU'. For a given parameter, 'incorrect data' meant that the output value, rate of change, or transmission characteristics were out of tolerance.

In addition, the equipment specification required the ADIRU manufacturer to conduct a failure mode and effects analysis (FMEA, see section 2.4.5) to demonstrate that the unit satisfied the relevant safety objectives. It also required the ADIRU manufacturer to determine the rate of different types of failure conditions for each of the ADIRU's output data parameters.

3.8.2 Failure modes and effects analysis (FMEA)

The ADIRU manufacturer's safety analysis document was issued in September 1992. The document stated that the primary purpose of the FMEA was:

> ...to assess the effects of various failure modes of components (single failures) on the circuit in which the component was used and then the effect on the equipment as a whole, and to precisely identify the means of failure detection for each case...

The manufacturer developed the FMEA by reviewing the design documentation, and partitioning each module into the 'lowest possible functional circuit group'. The failure rate for each component or 'piece part' (such as capacitors, transistors, RAM chips) was determined using an industry standard.[168] The failure modes, or types of failures, for each component were also identified, and the proportion of failures associated with each failure mode was estimated based on the manufacturer's service experience. The effects of each failure mode at the local module level and the ADIRU level were then determined. These determinations considered the effectiveness of the unit's BITE.

The result of the analysis was a table that listed the functional circuit, the function performed, the failure mode, the failure effects, the detection method, the calculated detected failure rate, and the calculated undetected failure rate.

The ADIRU FMEA used a systematic method of identifying the failure modes of a component and determining the effects on the system level. However, as with most FMEAs, it focused on single failures at a time and it only dealt with hardware failures. FMEA is a widely used technique by equipment manufacturers for analysing reliability and safety performance. As noted in section 2.6.3, traditional safety analysis methods such as FMEAs are associated with a number of limitations when conducted for complex systems.

None of the FMEA results for any of the components identified a failure mode resulting in frequent data spikes on output parameters (see also section 3.8.4). Following the 7 October 2008 occurrence, the ADIRU manufacturer reviewed its FMEA to determine if it indicated potential failure modes that could be analysed further as part of the investigation. None were identified.

---

[168] US Department of Defense, MIL-HDBK-217E, *Military Handbook: Reliability prediction of electronic equipment*, October 1986.

3.8.3 Summary analyses

The FMEA results were grouped together in a failure mode and effects summary (FMES) for each module. Based on this information, higher level summaries were developed for each module in terms of the overall failure rate (for the IR and ADR parts separately), and the proportion of the failures that were detected and undetected. For example, the estimated detected failure rate for the CPU module was 6.34 failures per million flight hours, and the undetected failure rate was 0.10 failures per million flight hours. The overall BITE effectiveness for the CPU module was calculated to be 98.52%.

The ADIRU manufacturer's safety analysis document compared the summary data with the Airbus safety objectives. In terms of the safety objective for an 'undetected failure of one IR or one ADR', it was calculated to be 0.72×10^{-5} per flight hour, which was less than the objective of 3×10^{-5} per flight hour.

3.8.4 Failure analysis of output parameters

The Airbus ADIRS equipment specification required that the ADIRU manufacturer determine the detected and undetected failure rate of 10 failure conditions for 53 output data parameters (including AOA) (Table 27). It did not state any requirements regarding the probability levels of the failure conditions.

Table 27: Output parameter failure conditions

Failure condition	Description
Loss of transmission	A 'dead' output on one or more outgoing databuses.
Oscillatory	Repetitive variations around a central value with an output amplitude out of tolerance.
Gain error	Output value is X times the actual physical value, with X being such that the output parameter may be out of tolerance.
Transport delay error	Delays in transmitting output data.
Frozen output	A constant value the same as that previously transmitted.
Limited error	A parameter out of tolerance in terms of its value but still within the defined range (that is, within specified limits).
Limited drift	A parameter drifting from its actual value at a rate that was less than the defined limit per minute.
Large error	Same as 'limited error', but the value is outside the defined range.
Large drift	Same as 'limited drift', but the value drifts at a rate greater than the define limit per minute.
Indeterminate output	Output could have any value with respect to the correct value.

The ADIRU manufacturer's analysis of the potential failure conditions was based on the FMEA results and its knowledge of the ADIRU's design. The safety analysis document noted that many of the conditions overlapped. With regard to 'indeterminate output', the safety analysis document stated the following:

> An indeterminate output can have any value with respect to the correct value, and may, therefore, greatly exceed tolerance limits and cause failure warnings. It can also be radically variable with respect to the correct value. It appears to be a combination of all the above defined failure categories, except Loss of Transmission, plus additional factors, such as indeterminate processing faults. It generally will affect more than one output label.

Based on this description, a data-spike occurrence could be considered as consistent with an indeterminate output scenario. In addition, data spikes could also be considered to be consistent with a 'large error' scenario.

The equipment specification did not state whether the failure conditions of interest were permanent, transitory or intermittent. The discussion of the failure conditions in the ADIRU manufacturer's safety analysis document was consistent with either a permanent or transitory effect on the value of an output parameter, and there was no discussion of the potential for frequent or intermittent data spikes.

The ADIRU manufacturer's safety analysis provided estimated probabilities for each of the 10 failure conditions for each parameter. For example, the estimated indeterminate output failure rate for corrected AOA was 8.72 failures per million flight hours (detected) and 0.11 failures per million flight hours (undetected). This estimation included a range of different failure types that were unrelated to the data spike failure mode.

3.9 ADIRU in-service performance

3.9.1 Reliability levels

The Airbus ADIRS equipment specification outlined minimum reliability requirements. For the ADIRU, these were:

- mean time between unscheduled removals (MTBUR[169]) of 6,000 hours
- mean time between failures (MTBF[170]) of 6,315 hours.

Units subject to an unscheduled removal were sent to the ADIRU manufacturer for examination. If the reported problem was verified, then the fault was considered a failure. MTBUR data was based on the number of units returned for examination, and MTBF data was based on the number of returned units for which verified failures were found.[171]

---

[169] The MTBUR was obtained by dividing the total number of flight hours logged by all units over a certain period of time by the number of units removed during that same period.

[170] The MTBF was obtained by dividing the total number of flight hours logged by all units over a certain period of time by the number of the units that failed during that the same period.

[171] Removal of a unit to update software or move the unit to another location were not counted as 'removals' for MTBUR purposes.

The average MTBUR for LTN-101 units on all Airbus A330/A340 aircraft was about 8,700 hours in 2004, increasing from the end of 2006 and reaching about 14,900 hours in 2008. The MTBF was about 11,500 hours in 2004 and 19,600 hours in 2008.

The ADIRU manufacturer reported that it was only aware of in-service problems associated with units when those units were sent to it for examination. Operators did not provide it with information on ADIRU fault messages that occurred but did not require the unit to be removed. Another ADIRU manufacturer advised that a similar situation existed for its units.

3.9.2 Operator's in-service ADIRU performance

The MTBUR and MTBF for the operator's A330 fleet during the period from 2004 to 2007 was generally better than the worldwide Airbus A330/A340 fleet with LTN-101 ADIRUs fitted. During 2008, the operator's MTBUR was about 11,100 hours and the MTBF was about 17,100 hours, which was slightly below the rest of the worldwide fleet.

The investigation reviewed the operator's maintenance records to identify events involving a reported NAV ADR [1, 2 or 3] FAULT and/or a NAV IR [1, 2 or 3] FAULT message from 2003 to 2008.[172] In addition to the three data-spike occurrences, the search identified 116 other events during the 2003 to 2008 period, which involved 312,834 aircraft flight hours. More detailed results were as follows:

- Most (83) of the events involved both IR and ADR fault messages, with 21 involving only an IR fault message and 12 only an ADR fault message.
- Eighteen events occurred during the engine start phase and the rest in cruise.
- Following 24 of the events, the ADIRU was removed for examination, and an actual (hard) fault was verified by the ADIRU manufacturer on 12 occasions.[173]
- For the majority of events, there were no fault messages on the flight's post-flight report (PFR) other than those directly related to the ADIRU. None of the events involved the same PFR pattern as the three known data-spike occurrences. For some of the events that involved other messages on the PFR (such as a NAV GPS FAULT or an autopilot disconnection), the QAR data was reviewed and no data spikes or other similar anomalies were observed.
- Unless the BITE data was obtained and a detailed investigation performed, it was not possible to know the reasons for the reported fault messages on any specific occasion. However, the review found that no other events involving data spikes or effects on the flight control system were identified.

The aircraft manufacturer advised that it was not practicable to compare the rate of reported IR and ADR faults between the operator and other operators.[174] However,

---

[172] The review looked at technical log entries and the AIRMAN database (discussed in Appendix D). Technical log entries were also reviewed to identify flights involving abnormal ECAM behaviour; none were identified.

[173] In addition to IR or ADR faults, units could also be removed for other reported performance problems that did not result in a fault message, such as inertial drift.

the aircraft and ADIRU manufacturers believed that, other than the three data-spike occurrences, the general pattern of faults reported on the operator's flights were not different to what other operators had experienced.

3.9.3 Other ADIRU failure modes

The ADIRU manufacturer reported that it had not encountered any other failure modes that produced a similar type of incorrect data output as occurred during the data-spike events. However, it reported that the LTN-101 had been associated with another failure mode that produced a similar BITE signature as the data-spike occurrences. That failure mode was known as 'dozing'.

The key features of a dozing event were as follows:

- Similar to a data-spike occurrence, there were no fault messages recorded in the ADIRU's BITE memory, even though such messages should have been present. In addition, some routine BITE messages were not recorded, and these messages were similar to those not recorded during the 7 October 2008 and 27 December 2008 occurrences.
- In contrast to a data-spike occurrence, the ADIRU ceased to provide any output data (either ADR or IR) to other systems.[175] It therefore had no safety impact, other than the loss of one of the three sources of IR and ADR information for the remainder of a flight. There was no effect on the autopilot or the flight control system.
- In contrast to a data-spike occurrence, ADR and IR fault messages were generated at the start of the event. In addition, neither of the fault lights on the overhead panel illuminated.
- Similar to a data-spike occurrence, the problem was a 'soft' fault; subsequent line testing identified no problem with an affected unit and a power cycle would restore the unit without any further problems. For those units that were removed and sent to the manufacturer for testing, no related problems were found.

Following a series of reported dozing events in 2005, the ADIRU manufacturer conducted a detailed investigation. Several affected units were subjected to examinations that included the normal manufacturer test procedure, together with EMI testing and environmental testing. The environmental testing included operation and power cycling at high and low temperature extremes. Testing for SEE was also conducted (section 3.6.6). The origins of the problem were not identified.

The ADIRU manufacturer also conducted a statistical analysis of dozing events to identify any potential patterns. This involved reviewing the BITE records for all units that had been returned for examination, regardless of the reason for the removal, to identify instances of the typical BITE signature. The BITE review identified that dozing events had occurred throughout the period 1994 to 2009, and that the number of events had generally increased each year (consistent with more units in service), with about 100 identified events each year in both 2008 and 2009.

---

[174] Although the AIRMAN database (Appendix D) contained information on such fault messages for most Airbus operators, the database was difficult to use and it was not practicable to make such comparisons.

[175] As it resulted in a complete cessation of output data, a dozing event was a 'detected' failure according to the Airbus equipment specification.

The review identified no apparent pattern associated with factors such as hours of service, unit configuration or software version. However, a higher proportion of units in the serial number range 4,000 to 5,000 had experienced dozing events compared to other serial number ranges.

Because the problem was transient, the aircraft manufacturer's advice to operators was to not remove the unit if it passed line testing on the aircraft. Therefore, many units that experienced a dozing event would not have been removed and no dozing signature would have been observed in the BITE data.[176] In other words, the number of dozing events was likely to have been much higher than 100 per year in 2008.

Of the 119 events involving reported ADR and/or IR fault on the operator's fleet during the period 2003 to 2008, 22 appeared to be dozing events. The technical log entries of 19 of the events were consistent with a dozing event.[177] A review of available BITE records for the operators' other units identified three additional events in which a dozing signature had occurred and there was no associated technical log entry for the relevant flight (although there were reported problems on adjacent flights). Most of the operator's ADIRUs were in the serial number range 4,000 to 5,000, and all but one of the probable dozing event unit serial numbers were in this range.

3.9.4 Safety levels

The aircraft manufacturer advised that it could not provide a figure for the LTN-101 ADIRU's undetected failure rate as it depended on the extent to which observed performance problems were identified and reported by operators. However, almost all of the problems that were reported by operators had been self-detected by the ADIRU.

Of the 119 events involving an ADR and/or IR fault on the operator's fleet during the period 2003 to 2008, the three data-spike occurrences were the only events in 312,834 flight hours that appeared to be classifiable as undetected failures. The data-spike occurrences were 'undetected' failures, as in each case the ADIRU provided incorrect ADR data to other systems without flagging this data as invalid.

As previously noted, there were only three known data-spike occurrences on LTN-101 ADIRUs, and the model had accrued over 128 million hours of operation.[178]

---

[176] The storage capacity for routine BITE messages was limited, and the effects of a dozing event on the BITE would be overwritten within about a week of normal line operations. Therefore, the number of dozing events that were detected through a review of BITE data would underestimate the actual number of events.

[177] Relevant BITE data was only available for five of these events, and a dozing BITE signature was identified in four of them. In the other case, the technical log entry was fully consistent with a dozing event but the BITE signature was not fully consistent.

[178] It should be noted that 'flight hours' refers to the aircraft flight time. Each ADIRU operates for a longer period than the flight time for each flight, and there are three ADIRU's operating on most flights.

4 FACTUAL INFORMATION: CABIN SAFETY

The first pitch-down resulted in many of the occupants in the cabin being thrown around the aircraft. The investigation consequently evaluated the extent to which the passengers and crew were wearing seat belts, the extent to which appropriate advice had been provided to passengers regarding the use of seat belts, and the effect of wearing seat belts on injury rates and severity. The investigation also examined the extent to which any other factors in the cabin may have influenced the frequency or severity of injuries.

4.1 Overview of cabin, crew and passengers

4.1.1 Layout of the cabin

The layout of the aircraft cabin is shown in Figure 47. There were eight exits, named according to their side (left or right) and position relative to the front of the aircraft (1 to 4). Ten cabin crew seats, which were used by cabin crew during takeoff and landing, were located near to each exit. A cabin interphone was located near each exit and was reachable from each of the cabin crew seats. A cabin crew rest area consisting of four seats was located in rows 40 to 41.

There were 297 passenger seats on the aircraft: 30 located at the front of the aircraft in business class (rows 1 to 5, between doors 1 and 2), 148 in the centre of the aircraft (rows 23 to 41, between doors 2 and 3), and 119 in the rear of the aircraft (rows 45 to 60, between doors 3 and 4).

4.1.2 Cabin crew requirements

Under Australian civil aviation requirements, a minimum of nine flight attendants were required to be carried on the A330-300 with a full passenger load.[179] Ten flight attendants were rostered to operate the flight, but one became unavailable and was not able to be replaced before the flight.

The operator's *Flight Administration Manual* stated that the primary responsibility of all cabin crew was passenger safety. The cabin crew included a customer services manager (CSM), who was responsible for supervising the cabin crew and the administration of the in-cabin service, and a customer services supervisor (CSS), who had a direct accountability for safety and service in economy class.

4.1.3 Cabin crew information

All of the cabin crew held endorsements for the Airbus A330 and Boeing 747 and 767 aircraft. Flight attendants and flight crew were required to complete emergency procedures (EP) training annually, and all the crew on the aircraft had completed the training within 7 months of the occurrence.

---

[179] The relevant CASA instrument required at least one flight attendant for each unit of 36 passengers (which equated to nine cabin crew for 303 passengers), and at least one flight attendant for each cabin exit (which equated to at least eight cabin crew).

Figure 47: Cabin layout

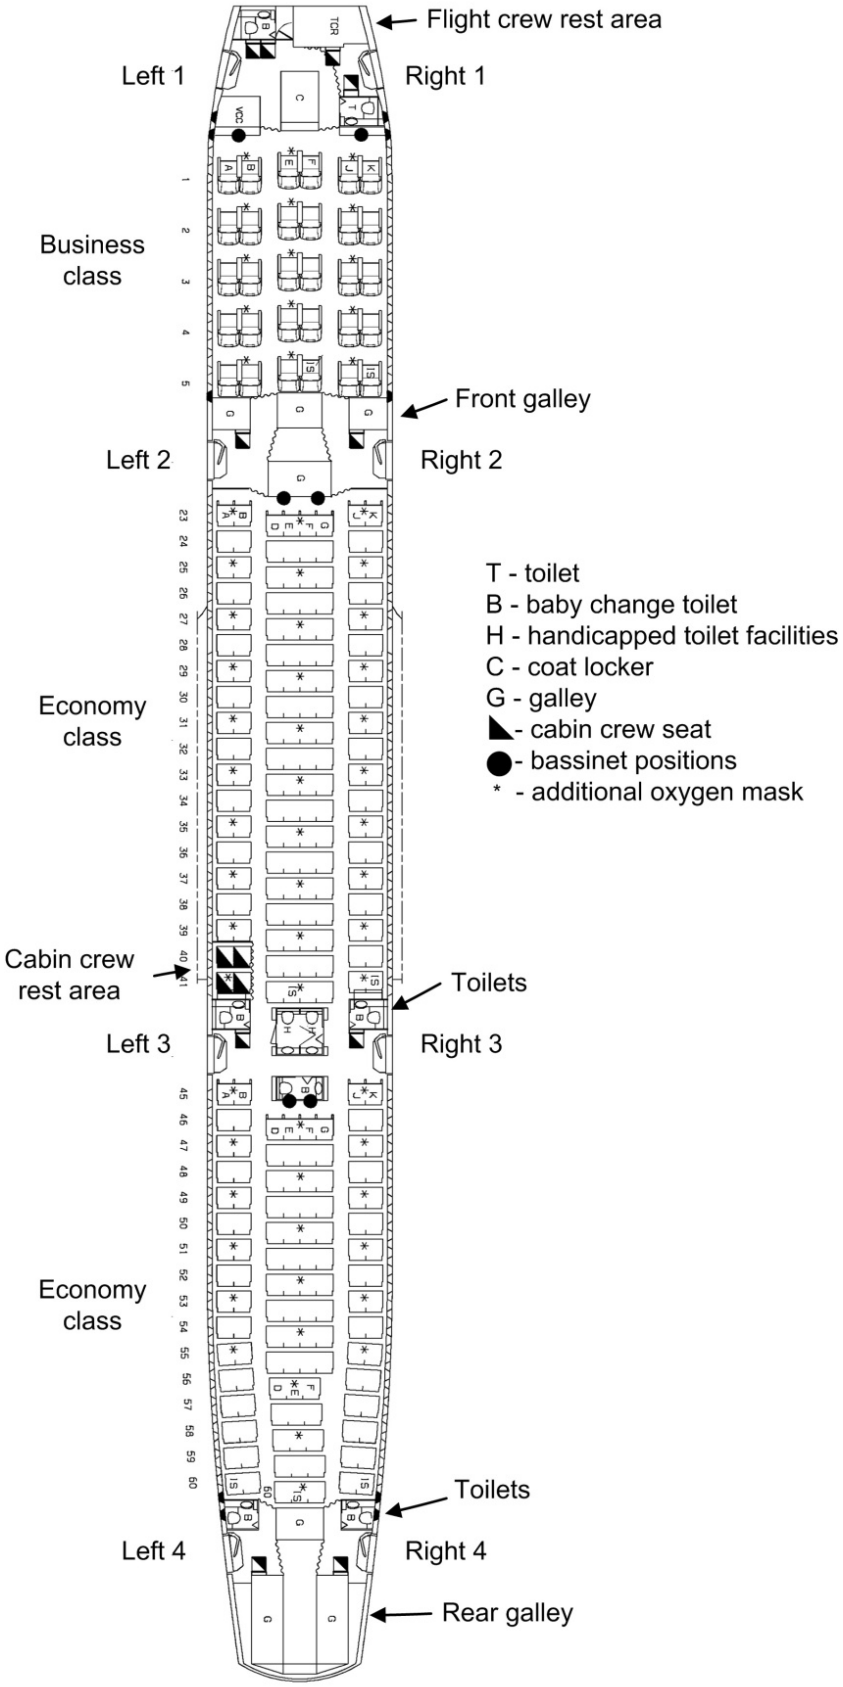

Summary details of the cabin crew's experience and relevant qualifications are provided in Table 28). The CSM had 6 years experience in that role, and the CSS was acting in the role for the second time.

Table 28: Cabin crew qualifications and experience

Crew position[180] (takeoff)	A330 endorsement	Last EP training	Flight attendant experience
Left 1 (CSM)	3 May 2006	14 May 2008	12 years
Left 1 assist	25 May 2004	30 June 2008	8 years
Right 1	27 February 2006	16 August 2008	5 years
Left 2	29 November 2006	1 April 2008	2 years
Right 2	15 January 2005	24 June 2008	4 years
Left 3	15 January 2005	1 April 2008	5 years
Right 3	2 March 2005	6 August 2008	37 years
Left 4 (CSS)	2 March 2005	23 April 2008	12 years
Right 4	4 May 2004	9 July 2008	6 years

4.1.4 Passenger information

The 303 passengers included six infants (less than 2 years old), 20 children (aged 2 to 17), 138 passengers aged between 18 and 45, and 138 passengers aged over 45. The age of one passenger was unknown. Most of the passengers (129) were from Australia. Of the remaining passengers, 65 were from European countries (mainly the United Kingdom) and 97 from Asian countries (mainly India and Singapore).

All of the 297 passenger seats were occupied during the takeoff. The other six passengers included three infants who were seated with a parent (two in the centre section and one in the rear section of the aircraft). In addition, three passengers on staff travel arrangements were allocated non-passenger seats for takeoff; one in the fourth occupant seat on the flight deck and two in the cabin crew rest area.[181]

The ATSB distributed a questionnaire to the passengers to obtain information about their experiences and observations during the flight. It also included questions on safety information, posture at the time of the in-flight upsets (seated versus standing), use of seat belts, and any injuries. A total of 98 questionnaires were returned to the ATSB, which equated to a sample of 35% of the 277 adult passengers. In addition, the investigation also obtained some information by interview or correspondence from 21 other passengers. The information from questionnaires, interviews and correspondence included details on pertinent topics such as injuries and seat belt use for many other passengers.

---

[180] The crew position refers to the seat allocated to the flight attendant for the takeoff and landing. Due to the accident on this flight, the cabin crew were not seated in the same positions for the landing.

[181] The operator's procedures stated that, in exceptional circumstances, cabin crew rest seats could be used for passengers if there was agreement to do so by all the cabin crew and the captain.

Overall, the investigation obtained information on key topics from a reasonable proportion of passengers from each section of the aircraft and from each major demographic group. Further details on the questionnaire and the questionnaire respondents are provided in Appendix I.

4.2 Sequence of events in the cabin

4.2.1 Events prior to the first in-flight upset

The cabin crew reported that they provided the passengers with the operator's standard pre-flight safety demonstration and subsequent public announcements (section 4.4.2). These included a public announcement during the aircraft's climb on departure from Singapore, which stated '... for your safety keep your seat belt fastened whenever you are seated'. The captain also advised that he made a public announcement at the beginning of the flight reminding the passengers to keep their seat belts fastened when seated. Passengers confirmed that these announcements were made.

As no turbulence was experienced during the flight, the crew did not illuminate the seat-belt sign or make any subsequent public announcements regarding the use of seat belts prior to the first in-flight upset.

Shortly before the first in-flight upset, the cabin crew completed the meal service and secured the service carts in the galleys, and four flight attendants had commenced their rostered rest break in the crew rest area.[182] The first officer left the flight deck at 0440 to commence his rest break.

Most of the cabin crew and 13 of the passenger questionnaire respondents noticed a slight change in the aircraft's flight profile or a slight reduction in thrust a few minutes prior to the first upset. These observations were consistent with the aircraft descending from 37,180 ft back down to 37,000 ft following the autopilot disconnection at 0440:27. Other than that event, none of the cabin crew or passengers noticed anything unusual with the flight prior to the first upset.

4.2.2 First in-flight upset (0442:27)

The two passengers who were seated in the cabin crew rest area for the takeoff had moved to cabin crew jump seats that were located at the front of the aircraft during the flight, and another passenger from the centre section had moved to the rear galley shortly before the first upset. Therefore, at the time of the first upset, there were 33 passengers in the front of the aircraft (one on the flight deck and 32 in the cabin), 149 in the centre section and 121 in the rear section. Several passengers were on their way to or from the toilets.

The first officer, CSM and two other flight attendants were standing in the forward galley at the time of the first upset. The CSM was answering a call on the cabin interphone from the second officer when the upset occurred. The four flight attendants in the cabin crew rest area were preparing to leave the area as their rest break was about to end. Another flight attendant was walking through the centre

[182] The operator's procedures required all cabin crew to have a 20-minute break free of duties in a cabin crew rest area within 6 hours of commencing duty.

section, and the ninth flight attendant was standing in the rear galley talking with two passengers. The two passengers were staff members, and one was an off-duty CSM.

Passengers and crew reported that the first upset occurred without any warning. They first noticed a sudden movement of the aircraft, generally described as a 'drop' or a 'pitch-down'. A significant number of passengers and crew were not seated or were seated without seat belts fastened (section 4.5), and they were thrown upwards. There was a loud bang as most of these occupants hit the ceiling, followed by the sound of many passengers screaming. Many of the occupants were injured when they hit the ceiling, or when they landed back on the floor or seats. Several of the overhead storage compartments opened, and there were some reports that bags fell out. Loose objects were also thrown upwards and around the cabin.

Immediately following the first upset, the second officer turned the seat-belt sign ON and made a public address announcement, requiring 'all passengers and crew be seated and fasten seat belts immediately'.

4.2.3 Second in-flight upset (0445:08)

Soon after the first upset, some of the passengers were moving about the cabin, helping injured passengers, retrieving items, or closing overhead compartments that had opened. Some of the cabin crew instructed passengers to be seated with their seat belts fastened.

Most of the passengers who had not been seated during the first upset were seated with their seat belt fastened before the second upset. Ten passengers reported that they were still in the process of getting back to their seat, and two passengers reported that they were seated during the first upset but were then helping an injured family member back to her seat when the second upset occurred.

Three of the cabin crew from the crew rest area and a passenger who had been in the aisle during the first upset were seated in the cabin crew rest area with seat belts fastened before the second upset. Another flight attendant had secured herself in the Right 2 cabin crew seat. Two other flight attendants had started moving through the centre section of the cabin, helping passengers into their seats, and one secured herself at the Left 3 seat before the second upset. When the second upset occurred, the first officer and two flight attendants were holding on to fittings in the forward galley, and the flight attendant and two passengers in the rear galley were holding on to fittings in that galley.

Passenger and crew descriptions of the second upset were similar to the descriptions of the first event.[183] However, the crew all reported that the second upset was less severe than the first. Although the questionnaire respondents were not specifically asked to compare the two events, 18 passengers stated that the second event seemed to be less severe and three said that it seemed to be more severe. Some passengers described the second event as more disturbing as it suggested that the upsets would keep occurring.

---

[183] Seven passenger questionnaire respondents reported that they did not notice a second event. These respondents included passengers who were seriously injured during the first upset, as well as passengers who were seated with seat belts on and received no injury.

4.2.4 Events after the second upset

Immediately following the second upset, one of the flight attendants in the front galley secured herself into the Left 1 cabin crew seat and the CSM secured herself into the Left 2 seat. The CSM helped the first officer remain secured near the Left 2 seat until he was requested by the other flight crew to return to the flight deck. One flight attendant in the centre section continued to provide some assistance to passengers before securing himself in the Right 3 seat. The off-duty CSM in the rear galley provided medical assistance to the injured flight attendant and passenger in the rear galley before securing the injured flight attendant into the Right 4 seat and the injured passenger into the Left 4 seat. She subsequently secured herself into a passenger seat at the rear of the aircraft.

During the remainder of the flight, the flight crew made several public announcements to the passengers and crew, instructing them to remain seated with their seat belts fastened. The operating CSM ensured that relevant information was provided to the flight crew regarding the status of the cabin and also ensured that other cabin crew were kept informed of the situation. The flight crew advised the cabin crew that they had to remain seated due to the potential risk of another upset event.

Prior to and after the second event, the cabin crew provided instructions to passengers to be seated and to keep their seat belts fastened. Some passengers requested medical assistance but the cabin crew advised them that they were unable to leave their seats. Some passengers provided medical attention to other passengers seated close to them, and cabin crew provided advice from their seats to some passengers about medical treatment. Towards the end of the flight, cabin crew and passengers in some areas cleared the aisles of debris and bags within their reach and without leaving their seats.

Details of the public announcements and other significant communications to and from the flight crew regarding cabin safety issues are provided in Table 29. These events do not include communications between cabin crew.

4.2.5 Events after landing

After the aircraft landed at 0532, the cabin crew moved through the cabin to assess the situation and start providing assistance to injured passengers. The CSM informed the flight crew about the extent of the injuries, and the first officer then informed air traffic control (and, via the controller, emergency response personnel) that there was at least one broken arm, one broken leg, several concussions and a significant number of head lacerations on board.

The crew agreed that the best course of action was for medical personnel to board the aircraft and start treating the most seriously-injured passengers and crew before disembarking the other passengers. At 0541, the first officer made a public announcement for the passengers to remain seated and to keep the aisles clear so that medical personnel could attend to the most seriously injured. Shortly after the aircraft arrived at the terminal at 0542, medical personnel boarded the aircraft. Passengers and crew reported that medical treatment was promptly provided to those who were injured.

Table 29: Significant cabin communications (based on the CVR)

Time	Event
0442	First pitch-down event. Second officer made a public announcement, telling all passengers and crew to be seated and to fasten seat belts immediately.
0445	Second pitch-down event.
0446	Captain made public announcement, advising that the crew were dealing with flight control problems, and telling everyone to remain seated with their seat belts fastened.
0447	Second officer called the flight attendant who was seated at Left 1, and asked her to send the first officer to flight deck. The flight attendant contacted the CSM at Left 2 to pass on the request. The CSM called the second officer to advise that the first officer was on his way.
0448	First officer arrived on the flight deck, and advised that people in the cabin were injured. Flight crew decided to make a PAN broadcast.
0449	First officer made a PAN broadcast to air traffic control, advising that the aircraft had flight control computer problems and that some people were injured. He requested a diversion to Learmonth.
0451	Second officer called the flight attendant at Left 1, asking her to obtain information about the injuries in the cabin. The flight attendant advised that she could not see all of the cabin. The second officer told her not to get out of her seat but to call the mid-section to obtain more information.
0452	First officer contacted air traffic control, requesting the controller to organise medical assistance at Learmonth.
0453	Off-duty CSM in rear galley called the flight deck, advising the second officer that a flight attendant and two passengers were seriously injured. Second officer told the off-duty CSM to be seated with her seat belt fastened as the crew could not guarantee that another upset would not occur. Based on the injury information, the captain asked the first officer to declare a MAYDAY.
0454	First officer made a MAYDAY broadcast to air traffic control, advising that they had at least one broken leg and some cases of severe lacerations.
0457	CSM (at Left 2) called the flight deck, advising the second officer that there was extensive damage to the ceiling and some injuries. Second officer advised the CSM that they were diverting to Learmonth. CSM stated that she would brief the other cabin crew and gather further information.
0500	Captain made public announcement, advising that the flight crew knew people were injured. He also advised that they were diverting to Learmonth, expecting to arrive in 10 to 15 minutes, and that medical assistance would be waiting. He told everyone to remain seated with their seat belts fastened.
0502	CSM contacted the flight deck to ask whether they should close the overhead lockers. The first officer advised her that it was safest to stay seated with seat belts fastened. He also advised it was not an emergency landing, and that no cabin preparation was needed and no public announcement was required about preparing for landing.
0527	CSM made a public announcement, telling passengers to follow any directions provided by the crew.
0527	First officer made a public announcement, advising that the aircraft would land in 3 to 4 minutes, the crew did not anticipate any further problems, and that passengers should listen for any crew instructions.
0532	Aircraft touched down at Learmonth.

After the passengers disembarked, there were some difficulties in the liaison between the organisations that were involved in managing and processing the passengers. Those issues were reviewed during a multi-agency debrief after the occurrence that was coordinated by the Westralia Airports Corporation and included representatives from the private, government and non-government organisations that were involved in the emergency response. As the identified problems were not relevant to the safety-related focus of the ATSB investigation, they are not discussed further in this report.

4.3 Cabin examinations

4.3.1 Post-accident inspection

An inspection of the aircraft's interior found significant damage to the overhead fittings (passenger service units, overhead stowage compartments and ceiling panels), mainly in the centre and rear sections of the passenger cabin. The damage was consistent with impact by persons or objects.

There was evidence of damage or impact to overhead fittings above one seat in the front section, 16 seats in the centre section (9% of the seats), and 27 seats in the rear section of the cabin (22% of the seats). Examples of some of the more significant damage are shown in Figure 48 and Figure 49.

None of the ceiling panels above the cabin aisle-ways in the front section exhibited any damage or movement from their fixed position. However, 11 of the 28 ceiling panels in the centre section, three of the six ceiling panels in the toilet area between the centre and rear sections, and 14 of the 22 ceiling panels in the rear section showed some damage or movement. Examples of some of the more significant damage are shown in Figure 50 and Figure 51.

The doors of three of the overhead compartments in the centre section and two of the doors in the rear section were not attached, and a small number of other doors in these two sections were partially dislodged.

Damage to overhead panels and fittings was evident in two of the toilets between the centre and rear sections, and one of the toilets at the rear of the aircraft.

Other notable findings from the cabin inspection included that:

- There was no apparent damage to any of the seat belts. A more detailed examination of a sample of seat belts was subsequently conducted (section 4.3.2).
- The seat squabs (horizontal cushions or pads that a passenger sits on) for three of the seats in the centre section and two of the seats in the rear section were not on the seats. Passengers advised that the squabs became detached during the first upset.
- Oxygen masks had deployed from above nine of the seats, and also in the rear galley. These masks were deployed during the first upset as a result of impact damage.
- Some of the cabin portable oxygen cylinders and some of the aircraft's first aid kits had been deployed. Cabin crew advised that this equipment was used for treating the passengers.

Figure 48: Example of damage to the fittings above passenger seats (centre section)

Figure 49: Example of damage to the fittings above passenger seats (rear section)

Figure 50: Example of damage to the ceiling panels in the aisle (rear section)

Figure 51: Example of damage to the ceiling panels in the aisle (centre section)

4.3.2 Seat belt examinations

Six passengers reported that they were seated with their seat belt fastened at the time of the first in-flight upset, but that the belt became unfastened and did not restrain them in their seats. None of the six passengers could provide details of how their belts released, and no other passengers reported any problems with the operation of their seat belts.

The seat belts on the aircraft were a very common type of lap belt with a lift-lever buckle. The investigation examined a sample of 51 belts on the occurrence aircraft. The sample included the belts of four of the passengers who had reported seat belt problems before the examination was conducted (November 2008).[184] It also included the belts of passengers that the investigation knew had received hospital treatment but did not know if they were wearing a seat belt. No problems were identified with the condition of the webbing, buckle or connector of any of the 51 belts examined.

During the seat belt examinations, investigators identified a scenario that could result in a seat belt being inadvertently unfastened. The scenario involved the buckle of a very loosely-fastened belt catching on the underside of the seat's right armrest. The examinations also noted that belts that were fastened this loosely posed a significant injury risk, even if they remained fastened. Only three of the six passengers who reported a problem advised that they had their belts loosely fastened.

The inadvertent release scenario had not been identified in previous investigations involving in-flight upsets. It was not possible to determine whether the scenario actually occurred to the three passengers who reported having loosely-fastened seat belts on the occurrence flight.

Further information on the seat belt design and the examination of the inadvertent release scenario are provided in Appendix J.

4.4 Seat belt requirements

4.4.1 Regulatory requirements and guidance

Use of seat belts

Australian Civil Aviation Regulation (CAR) 251(1) stated:

> ...seat belts shall be worn by all crew members and passengers:
>
> (a) during take-off and landing;
>
> (b) during an instrument approach;
>
> (c) when the aircraft is flying at a height of less than 1,000 feet above the terrain; and
>
> (d) at all times in turbulent conditions.

[184] Two of the six passengers who reported that they were wearing their seat belts but that their belts did not restrain them provided their reports after November 2008.

Requirements in the US and many other countries state that seat belts shall be worn by all passengers and crew during takeoff, landing, and when the seat-belt sign has been illuminated.[185] The Civil Aviation Safety Authority (CASA) has proposed moving to a similar set of requirements in Australia.[186]

Safety instructions

With regard to passenger briefings, the Australian Civil Aviation Order 20.11 stated that operators were required to orally brief passengers on several matters, including the 'use and adjustment' of seat belts. The CASA Civil Aviation Advisory Publication (CAAP) 253-2 (*Passenger safety information: Guidelines on content and standard of safety information to be provided to passengers by aircraft operators*) provided more detailed guidance to operators. It included the following statement about seat belts in a section on safety briefings:

> Passengers must be briefed on the use and adjustment of seat belts, ie. the method of fastening, tightening and unfastening.
>
> The briefing should include that seatbelts must be fastened anytime the "seatbelt" sign is illuminated and that any instruction from crew members in relation to the seatbelt must be obeyed.
>
> Passengers should be informed that seatbelts are to be worn low and tight, and kept fastened anytime they are seated.

The CAAP also included the following statement about safety information cards:

> The card should have instructions for fastening, tightening, and unfastening seatbelts and indicate they must be fastened during takeoff, landing and whenever the fasten seatbelt sign is on.

4.4.2 Operator requirements and passenger briefings

Normal procedures

The operator's *Flight Administration Manual* outlined policies, procedures and standards for crew members. In addition to reiterating the Australian regulatory requirements for seat belts, the manual stated:

> Seat belts (including full harnesses where fitted) shall be worn by all passengers and Cabin Crew whenever the Seat Belts sign is illuminated. The only exception to this requirement is when Cabin Crew are performing safety related duties. Seat belts, when worn, shall be properly adjusted and securely fastened.

---

[185] See US Federal Aviation Regulations 121.311 and 121.317.

[186] Notice of Proposed Rulemaking 9809RP (1998), *Proposed Regulations Relating to Passenger and Crew Member Safety.*

The flight crew's procedures required them to illuminate the seat-belt sign prior to taxiing the aircraft. The cabin crew's procedures required them to provide a pre-flight safety demonstration to passengers, and for the A330 they provided the demonstration by video. The audio track of the video stated:

> Having your seatbelt done up low and tight is absolutely essential during takeoff, landing and turbulence. It is a Qantas requirement that you keep it on at all other times.

After the flight crew turned off the seat-belt sign following the takeoff, the cabin crew were required to provide the following public announcement to passengers:

> The Seat Belt sign is now off, however, for your safety keep your seat belt fastened whenever you are seated.

The operator's procedures also recommended that the flight crew provide passenger briefings at various times during the flight. It was common practice early in a flight for the flight crew to remind passengers to wear their seat belts when seated.

When the seat-belt sign was illuminated during the descent, the cabin crew were required to provide the following announcement:

> The cabin lights will be dimmed for landing. Passengers and crew must now be seated with their seat belts fastened.

The operator's cabin crew were required to check that passengers were wearing their seat belts before takeoff and during descent. Where possible, they were also required to check passenger compliance if the seat-belt sign was illuminated during flight. However, during flight when the seat-belt sign was not illuminated, there was no policy or procedure requiring cabin crew to check or enforce passenger seat belt use.

The operator's safety information card located at each seat for the A330-300 contained important information for passengers. It included a diagram showing how to fasten, tighten and unfasten the seat belt (Figure 52). In addition to containing regulatory requirements, the operator's safety information cards also contained additional safety advice. There was no additional advice relating to seat belts on the operator's A330-300 safety information card.

Figure 52: Extract from the operator's A330-303 safety information card

Abnormal or emergency procedures

In the case of unanticipated turbulence with an immediate safety hazard, the flight crew were required to illuminate the seat-belt sign, and to provide a public address announcement stating 'All passengers and crew be seated and fasten seat belts

IMMEDIATELY'. Cabin crew were required to secure themselves in the nearest available seat or to wedge themselves in the aisle.[187]

If the situation permitted, the CSM was to initiate the 'call back procedure' and ascertain the condition of the cabin and provide a report to the flight crew. The call back procedure involved calling all the cabin crew stations simultaneously on the cabin interphone. Where possible, the flight attendants located at each station were to respond to the call and provide a report on the condition of the cabin. Cabin crew located in the crew rest area, or otherwise unable to reach the interphone without unfastening their seat belts, were not required to answer the call. If the CSM was not in the cabin, then the CSS or another senior flight attendant was to initiate the call.

4.4.3 Other operator's procedures and passenger briefings

The requirements and guidance outlined in CAAP 253-2 relating to seat belts was consistent with overseas requirements and guidance.[188] The operator's procedures regarding passenger seat belt use were consistent with these requirements, and the same basic procedures were followed by other operators in Australia, as well as in many international airlines.

4.4.4 Guidance on how seat belts should be worn

For seat belts to be effective, they need to be worn correctly. Lap belts are designed to be worn across the passenger's hips, and pre-flight safety demonstrations inform passengers that seat belts need to be worn 'low and tight'. The reason for this requirement is that the pelvic bones are best able to withstand loads during impacts. More specifically[189]:

> The safety belt should be placed low on your hipbones so that the belt loads will be taken by the strong skeleton of your body. If the safety belt is improperly positioned on your abdomen, it can cause internal injuries. If the safety belt is positioned on your thighs, rather than the hipbones, it cannot effectively limit your body's forward motion.[190]

Loose seat belts also do not effectively limit the body's motion during vertical forces, and also increase the likelihood of a person being injured due to being thrown against armrests or other fixtures.

In their pre-flight safety demonstrations, operators tell passengers to keep their seat belts fastened low and tight, but this instruction is usually provided in the context of

---

[187] In situations not involving an immediate safety hazard, the processes were similar. However, the public address announcement did not contain the word 'immediately'. Cabin crew were required to prioritise their duties to secure service carts and other equipment, and be seated within 1 minute of the seat-belt sign being illuminated.

[188] The CASA requirements were similar to those outlined by the US FAA in Advisory Circular (AC) 121.24B (*Passenger safety information briefing and briefing cards*), effective July 2003.

[189] Federal Aviation Administration, *Seat belts and shoulder harnesses: Smart protection in small airplanes*. AM-400-91/2, revised May 2004.

[190] Belts placed below the pelvic joint can allow 'submarining' during certain types of impacts. During submarining, the occupant slides forward under the seat belt, leading to additional injuries due to being unrestrained and being squeezed between the seat and the belt.

takeoff, landing or when the seat-belt sign is illuminated. Operators also advise passengers to keep their seat belts fastened at all times when seated, but when this advice is provided, passengers are generally not reminded of the need to keep the belts fastened low and tight (or at least relatively firm).

In some cases, passengers are provided with information that they can loosen their seat belts after takeoff. For example, the operator's safety information cards for some of its aircraft types (but not the A330-300) contained the following information[191]:

> FASTEN YOUR SEAT BELT AT ALL TIMES TO PREVENT INJURY
>
> You must keep your seat belt fastened at all times. Make sure that it is low and tight over your hips. Practice opening and closing it. The few seconds you spend fumbling for your seatbelt during an emergency or in turbulence can determine whether or not you are injured. The seat belt can be loosened after take off, but pilots cannot predict clear air turbulence so please keep your seat belt on. Tighten your seat belt again before landing and remember on arrival to remain in your seat with your seat belt fastened until the seat belt sign is turned off.

4.5 Posture and seat belt use

Reliable information on posture and seat belt use at the time of the first upset was obtained for 164 passengers. Of these passengers, 81 were seated with their seat belts fastened, 60 seated without their seat belts fastened, and 23 were not seated. Details for the three sections of the aircraft are presented in Table 30.

Table 30: Posture and seat belt use (where reliable information was available)

Category		Seated, belt on	Seated, belt off	Seat belt use rate	Not seated	Not known	Total
Location	Front	14	7	67%	1	11	33
	Centre	43	26	62%	11	69	149
	Rear	24	27	47%	11	59	121
Total		81	60	57%	23	139	303

Of the 22 adults who were not seated, two were in the toilets, at least 10 were standing near their seats, and at least five were in an aisle. The other non-seated passenger was an infant who had just been picked up by a parent.

The overall seat belt use rate for seated passengers was 57% (81 out of 141). This rate was significantly higher in the front (67%) and centre (62%) sections than in the rear section (47%).

The passenger questionnaire asked these passengers who had their seat belt fastened during the first upset whether their belt was 'tightly fastened' or 'loosely fastened'. Similar information was obtained from a small number of other passengers. Of the 81 passengers known to be wearing seat belts, 50 said their belts were 'tightly

[191] The operator advised that the safety information cards for different aircraft types contained a different set of additional safety information because they were developed at different times. In all cases the additional safety information was included due to an initiative by the operator rather than due to a regulatory requirement.

fastened' and 13 said they were 'loosely fastened'. The information was not available for the other 18 passengers who wore seat belts.

Reliable information on posture and seat belt use was not available for 139 passengers. However, information from several sources suggested that there were more than 60 passengers seated without their belts fastened. More specifically:

- Six passengers reported to the ATSB that they were seated with their seat belts fastened at the time of the first upset, but that the belt became unfastened and did not restrain them in their seats. As there was some doubt regarding whether or not their belts were fastened (Appendix J), these passengers were not included in the totals discussed above.
- Eight passengers were known to be injured but information on their posture and seat belt use was not provided. The injuries received by these passengers were consistent with not wearing a belt or with not being seated (section 4.6.7).
- Although most of the damage to the fittings above passenger seats was consistent with the information obtained from passengers and crew members, there was damage above two seats in the centre section and eight seats in the rear section that was not able to be accounted for by the information provided (section 4.6.9).

Although more than 60 passengers were likely to have been seated without their seat belts fastened, the overall seat belt use rate was likely to have been higher than 57%. The investigation's processes for obtaining passenger information focused to some extent on identifying the passengers not wearing seat belts.[192]

4.6 Injuries

4.6.1 Injury levels

The Western Australia Department of Health reported that 51 passengers and two crew members received medical treatment at a hospital, either in Learmonth or in Perth. Eleven of these passengers and one of the flight attendants were admitted to hospital. Injury information was also obtained from the passenger questionnaire, as well as from interviews and correspondence with the passengers and crew.

Based on the available information, nine of the 12 crew members and 110 of the 303 passengers were known to be injured. Three of the crew members and 58 passengers were known to have not been injured. The injury statistics are summarised in Table 31.

[192] For example, the questionnaire asked passengers to provide information on passengers near them who were not wearing their seat belts, but did not specifically ask about other passengers wearing seat belts. In addition, the investigation attempted to contact all passengers who had received hospital medical treatment and who had not completed the passenger questionnaire; most of those contacted were seated without their seat belts fastened.

Table 31: Levels of injury

Injury level	Crew	Passengers	Total
Total injuries	9	110	119
Serious (hospital admission)	1	11	12
Hospital treatment (not admitted)	1	40	41
Other minor injury	7	59	66
Not injured	3	58	61
Unknown	0	135	135
Total occupants	12	303	315

Of the 119 injuries, 12 were classified as serious. Under the Australian *Transport Safety Investigation Regulations (2003)*, a serious injury is defined as 'an injury that requires, or would usually require, admission to hospital within 7 days after the day when the injury is suffered'. Many countries use the International Civil Aviation Organisation (ICAO) definition of serious injury.[193] This definition includes several conditions, such as hospitalisation for more than 48 hours, fracture of any bone (except simple fractures of fingers, toes and nose), and lacerations that cause severe haemorrhage. The two definitions produced the same result for this occurrence.[194]

To best assess the factors associated with the injuries, the injury numbers for the passengers and crew were combined. The number of serious injuries was still too low to meaningfully compare injury rates for factors such as location in the aircraft, posture or seat belt use.[195] Consequently, the criterion of whether a person received medical treatment at a hospital soon after the occurrence provided a more useful indicator of the more significant injuries.

The number of aircraft occupants who received hospital treatment was also a more reliable indicator of injury potential than the total number of known injuries. Injury information was not available on the injury status of 135 of the passengers. It is likely that the investigation identified most if not all of the passengers who received hospital treatment, but it is unlikely that the investigation identified all of the passengers who received minor injuries but did not attend a hospital.

This report only discusses physical injuries. Many of the occupants of the aircraft reported that they also experienced significant stress or anxiety as a result of the in-flight upsets.

4.6.2 Injury levels classified by posture and seat belt use

Reviews of injuries experienced during turbulence events and other in-flight upsets have also shown that injury rate and severity is much lower for aircraft occupants

---

[193] ICAO, *Annex 13 to the Convention on International Civil Aviation: Aircraft Accident and Incident Investigation*, 9th Edition, July 2001.

[194] The use of the ICAO definition resulted in 12 serious injuries. However, four of these occupants were different to the 12 who were admitted to hospital and classified as serious injuries under the *Transport Safety Investigation Regulations*..

[195] It is also worth noting that some of the occupants were admitted to hospital for observational purposes. The severity of injuries appeared to vary significantly for those who were seriously injured, as well as for those who received minor injuries.

wearing seat belts (Appendix K), and the data for the 7 October 2008 occurrence was the same.

Table 32 provides injury information for the 7 October 2008 occurrence according to the known information about posture and seat belt use. The key results were as follows:

- The overall injury rate for occupants wearing seat belts (31%) was significantly lower than for occupants who were seated but not wearing seat belts (93%) or those who were not seated (97%).
- The hospital treatment rate for occupants wearing seat belts (7%) was significantly lower than for occupants who were seated but not wearing seat belts (32%) or those who were not seated (52%). The difference between those who were seated and not wearing seat belts and those not seated was not statistically significant.
- It is likely that there were more passengers injured due to not wearing seat belts than was reported (section 4.6.7).

Table 32: Injury details for all occupants by posture and seat belt use[196]

	Posture and seat belt use				
Injury level	**Seated, belt on**	**Seated, belt off**	**Not seated**	**Not known**	**Total**
Serious (hospital admission)	2	3	6	1	12
Minor (hospital treatment)	4	17	10	10	41
Other minor injuries	19	30	14	3	66
Total injuries	25	50	30	14	119
Not injured	56	4	1	0	61
Unknown	2	8	0	125	135
Total occupants	83	62	31	139	315
Hospital treatment rate	7%	32%	52%		
Total injury rate	31%	93%	97%		

4.6.3 Injuries to occupants not seated

Except for one of the cabin crew, all of the 31 occupants who were not seated were injured as a result of impacting parts of the aircraft. At least 24 of these occupants hit overhead fittings, and some were also injured as a result of landing on the floor. Two of the occupants were also hit by other occupants.

Most of these occupants received multiple injuries. The primary injury types were spinal injury (five), neck or back injury (10), head injury (7), arm injury (four), leg injury (two), or bruising over whole body (two).

Six of the non-seated occupants were admitted to hospital. These included four passengers with spinal injuries, a passenger with a head injury, and a flight attendant with a leg injury.

[196] Note that the total number of occupants known to be seated with belts on, seated with belts off, and not seated is higher in this table as it includes passengers and crew members.

4.6.4 Injuries to seated occupants not wearing seat belts

For the 50 injured occupants who were seated but not wearing seat belts, most (42) were injured as a result of impacting parts of the aircraft during the first in-flight upset. At least 36 of these occupants hit overhead fittings, with some also being injured when landing on the floor or seats. Two of these passengers were also hit by other occupants. Of the remaining eight occupants, one reported being injured as a result of reaching for an armrest, and the injury mechanism for the other seven was not reported.

Most of these occupants received multiple injuries. The primary injury types were spinal injury (one), neck or back injury (21), head injury (15), arm injury (seven), and leg injury (five).

Three of the seated passengers not wearing seat belts were admitted to hospital. These included an adult with a spinal injury, another adult with a neck/back injury, and an infant with minor head injuries who was admitted for observation.

4.6.5 Injuries to occupants wearing seat belts

The only crewmembers wearing seat belts were the two flight crew on the flight deck, and neither was injured. Passengers wearing seat belts experienced the following types of injuries:

- Twelve passengers reported experiencing whiplash or being jolted during the upset. Eleven of these passengers reported strain/sprain injuries or pain to the neck or back, and another reported a shoulder strain. Two passengers reported that reaching for another passenger at the time of the upset contributed to their injuries.
- Four passengers reported being hit by other passengers or objects. The primary injuries were shoulder/chest bruising (and cracked ribs), shoulder bruising, neck/back pain and leg bruising.
- One passenger reported hitting the seat in front, resulting in neck/back pain.
- Two passengers reported hitting the armrests of their seats, resulting in injuries to their sides. One of these passengers was reaching for another passenger at the time.
- One passenger received a wrist injury when reaching for another passenger.
- One child received abdominal contusions from a seat belt.
- Four passengers reported neck/back injuries but the injury mechanism was not clear.

Two of the passengers wearing seat belts received serious injuries. The child who received abdominal contusions was admitted to hospital for observation. The other passenger experienced neck pain at the time of the upset, and was admitted to hospital 3 days after the occurrence after experiencing a stroke.

4.6.6 Tightly- versus loosely-fastened seat belts

The injury rate for the 18 passengers with loosely-fastened seat belts (46%) was not significantly different to the rate for the 50 passengers with tightly-fastened seat belts (34%).

Of the 25 passengers wearing seat belts who were injured, six stated that their belts were loosely fastened. One of these passengers reported a neck/back injury when she was thrown back into her seat as the aircraft stabilised, and she believed that the loose seat belt contributed to the injury. For the other five passengers, it did not appear that a loose seat belt contributed to the injuries. Three of these passengers were injured because they were hit by other people or objects, and one hit an armrest when reaching for another passenger. The other passenger received a shoulder injury but the injury mechanism was not known.

For the child who received abdominal contusions, the child's parent reported that the seat belt was firmly fastened.

4.6.7 Other injured passengers

As previously mentioned, six passengers reported that they were wearing their seat belts but were not restrained in their seats. All of these passengers were injured as a result of impacting parts of the aircraft, and at least four hit overhead fittings. Three received hospital medical treatment. Most of these six passengers received multiple injuries, and the primary injury types were neck or back injury (three), head injury (one), arm injury (one), and leg injury (one). Four of these passengers were located in the centre section and two in the rear section.

Eight other passengers were injured, but their posture and seat belt use was not reported. Seven received hospital medical treatment but none were admitted to hospital. The primary injury types were spinal injury (one passenger), neck or back injury (three), and head injury (four). One of these passengers was reported to have hit overhead fittings, and there was damage to the overhead fittings above the assigned seats of that and another passenger. The injury mechanism for the others was not reported. Four of these eight passengers were located in the centre section and four in the rear section.

4.6.8 Injuries by location in the aircraft

Table 33 provides data on the injury levels for each section of the aircraft. Crew members in the front galley were included in the front section.

Table 33: Injury levels for each section of the aircraft

Injury level	Front	Centre	Rear	Total
Serious (hospital admission)	0	7	5	12
Hospital treatment (not admitted)	1	25	15	41
Other minor injuries	9	29	28	66
Total injuries	10	61	48	119
Not injured	17	28	16	61
Unknown	12	65	58	135
Total occupants	39	154	122	315
Hospital treatment rate	3%	21%	16%	17%
Total injury rate	26%	40%	39%	38%

The percentage of occupants who were injured appeared to be lower in the front section (26%) than the centre (40%) or rear (39%) sections, but the difference was not statistically significant. However, the percentage of occupants who required hospital treatment was significantly lower in the front section (3%) than the centre (21%) or rear (16%) sections.

As discussed in section 4.5, the seat belt use rate for passengers was lower in the rear section than in the front or the centre sections. In addition, the acceleration forces during to the in-flight upsets were highest in the rear section of the aircraft (Appendix A). However, a review of the cabin damage indicated that there were more passengers injured in the rear section than was reported (section 4.6.9).

4.6.9 Comparison of injuries with cabin damage

The information provided by the questionnaire respondents and other passengers about posture, seat belt use and injury mechanisms accounted for most of the damage observed in the cabin. Not all of the occupants who hit overhead fittings caused damage or movement to the fittings, but almost all of the observed damage or movement was consistent with passengers or crew members at those locations being unseated or seated without their seat belts fastened.

The available information did not account for damage above two seats in the centre section and 10 seats in the rear section. Some of these seats in the rear section were adjacent to each other, and the damage may therefore have been due to the same passenger. Overall, it was considered likely that there were at least two passengers from the centre section and eight passengers from the rear section who hit and damaged overhead fittings, but no information was reported to the investigation about their posture or seat belt use. Two passengers with assigned seats at these locations had injuries consistent with hitting the ceiling, but no information was available on whether the other passengers were injured.

There was damage in two of the four toilets that were located between the centre and rear sections, and one of the two toilets at the rear of the aircraft. The damage to the toilet at the rear of the aircraft was not consistent with other information. Therefore, in addition to the 23 passengers known to have been unseated, there was probably another passenger who was unseated and injured.

4.7 Factors influencing the use of seat belts

4.7.1 Previous research

Reviews of injuries from previous in-flight upsets have shown that not all seated passengers wear their seat belts when seated (Appendix K). The associated investigation reports have rarely discussed the reasons why passengers were not wearing their seat belts.

The investigation into the 7 October 2008 occurrence identified only one research study that examined the factors associated with seat belt use on aircraft. In that study, researchers asked passengers waiting in an airport in the United States how often they wore seat belts when the seat-belt sign was turned on and when the sign was turned off and the aircraft was not in the takeoff or landing phase (Girasek and Olsen 2007). Overall, 7% reported that they would ‘rarely’ or ‘never’ wear their

seat belt when the seat-belt sign was off. There was no difference in reported seat belt use between male and female passengers. However, younger adults were more likely to not wear seat belts compared with older adults, and Asian passengers were more likely to not wear seat belts compared with other groups. Other factors associated with a lower reported seat belt use included lower frequency of air travel, lower household income, travelling with friends, and factors relating to higher levels of alcohol use.

The results from that study were generally consistent with the extensive research that has been conducted into the factors affecting seat belt use in road vehicles (Appendix L).

The rest of this section reviews information obtained during the investigation to help understand factors associated with the use of seat belts during the occurrence flight.

4.7.2 Demographic factors

Table 34 provides data on posture and seat belt use at the time of the first upset for those passengers for whom reliable information was obtained. The data is presented in terms of gender, nationality and age.

Table 34: Posture and seat belt use (where reliable information available)

Category		Seated, belt on	Seated, belt off	Seat belt use rate
Gender	Male	43	29	60%
	Female	38	31	55%
Nationality	Australia	45	29	61%
	Europe	18	9	67%
	Asia	14	19	42%
	Other	4	3	57%
Age (years)	Infant (< 2)	0	3	0%
	2 - 17	11	4	73%
	18 - 30	8	13	38%
	31 - 45	14	16	47%
	46 - 60	28	15	65%
	Over 60	20	9	69%
Total		81	60	57%

Data was only available for four of the six infants. Of these, three were seated with a parent or in a bassinet but not restrained, and another had just been picked up by a parent. Data was available for 15 children aged from 2 to 17, and their seat belt use rate (73%) was higher than for adult passengers (57%).

Statistical analyses[197] were done to compare the adult passengers on several variables. The main results were:

- There was no difference between the seat belt use rate of males (59%) and females (55%).
- The seat belt use rate appeared to increase with age, and the difference between passengers aged 18 to 45 (43%) and passengers over 45 (67%) was statistically significant.
- Adult passengers from Asian countries had a significantly lower seat belt use rate (39%) than other adult passengers (63%).
- Adult passengers in the rear section had a lower seat belt use rate (44%) than passengers from the front and centre sections (64%). The proportion of adult passengers aged 18 to 45 was higher in the rear section (63%) compared to the front and centre sections (41%). The proportion of adult passengers from Asian countries was also higher in the rear section (46%) compared to the front and centre sections (26%).

4.7.3 Situational factors

The investigation attempted to obtain seat belt use data from other flights to determine if seat belt use varied depending on situational factors, such as the time of day or duration of the flight. The operator reported that it did not have any data regarding passenger seat belt use rates during flights when the seat-belt sign was not illuminated. Two other operators in the Asia-Pacific region were also contacted, and both reported that they had not collected any such information on seat belt compliance. No relevant research studies were identified.

Passenger surveys are occasionally conducted during aircraft occurrence investigations. One previous survey conducted by the ATSB asked about seat belt use during the cruise phase of flight. The occurrence was a depressurisation event involving a Boeing 747-438, 475 km north-west of Manila Airport in the Philippines on 25 July 2008.[198] The depressurisation event occurred about 55 minutes into the flight. At the time a meal service was being conducted in the economy section and a drink service in other sections of the cabin. From the 346 passengers on board, 152 questionnaire responses were obtained. Of these passengers, 127 reported they were seated with their seat belt fastened, 22 said they were seated without their seat belt fastened, and two were not seated. The seat belt use rate for these passengers was therefore 85%.

The seat belt use rate during the 25 July 2008 occurrence was higher than for the 7 October 2008 occurrence. This could be attributable to the time of the event in the flight, occurring soon after takeoff and prior to the meal service being completed. In contrast, the 7 October 2008 occurrence occurred after the meal service and over 3 hours into the flight.

In addition to factors such as the time since takeoff, it could be expected that seat belt use would be higher on flights where turbulence had been experienced. On the

---

[197] Statistical comparisons were done using the χ^2 (Chi squared) test for independent groups. In this report, 'statistically significant' means that the chance of the difference being present due to chance alone was less than 5%.

[198] See ATSB investigation report AO-2008-053 available at www.atsb.gov.au.

7 October 2008 flight, there had been no turbulence and no need for the flight crew to illuminate the seat-belt sign prior to the first upset.

4.7.4 Previous flying experience

The questionnaire asked passengers to provide information on how many commercial airline flights they had been on before the occurrence flight. For the 98 questionnaire respondents, seated passengers who wore seat belts had broadly the same previous flying experience as seated passengers who did not wear seat belts. About 85% of both groups had over 20 commercial airline flights prior to the occurrence flight.

The seat belt use rate for passengers who had flown on 20 or less flights (62%) was no different to the use rate for passengers who had flown more than 20 previous flights (67%).

4.7.5 Attention to safety information

Previous research has noted that, although passengers believe that cabin safety communications are important, they generally pay a low amount of attention to such communications. Reasons for the low attention include overconfidence, high message recognition (rather than recall), issues related to the content presentation, and social norms within the cabin (Parker 2006).

The investigation's questionnaire asked passengers 'how much attention did you give to the pre-flight safety demonstration either given by the flight attendants or presented on video'. The available answers were 'no attention' (one response), 'a little attention' (15), 'some attention' (33) or 'full attention' (48). The percentage who reported paying some or full attention was not significantly different between passengers wearing seat belts (78%) and passengers seated but not wearing seat belts (93%).

Passengers who said that they paid no attention or a little attention to the demonstration were asked to provide a reason why. Most (14) of these passengers reported that they were familiar with the briefing from previous flights, and one passenger reported falling asleep soon after boarding. Three of the passengers reported that they primarily attended to the location of the emergency exits.

The questionnaire also asked passengers about the extent that they 'read the safety card (in your seat pocket) prior to the event'. The available answers were 'not at all' (29 responses), 'some parts quickly' (22), 'some thoroughly' (26) and 'all thoroughly' (19). The percentage who reported that they read some or all of the card thoroughly was not significantly different between passengers wearing seat belts (42%) and passengers seated but not wearing seat belts (52%).

Passengers who did not read the card or read 'some parts quickly' were asked to provide a reason why. Many (22) of these passengers reported that they were familiar with the contents, either from many previous flights or from specific flights in the recent past. Twelve of the passengers reported that they believed that the safety demonstration had provided sufficient information, and another six reported that they had attended only to the location of the emergency exits.

The questionnaire asked passengers whether they had any personal characteristics that would have affected their ability to understand or read instructions. Two

passengers reported that they had hearing difficulties and one reported an English language difficulty. Two of these passengers were wearing seat belts at the time of the first upset, and the other passenger reported that he normally wore his seat belt when seated but on this occasion had only just returned from the toilet.

4.7.6 Understanding of seat belt requirements

Research studies and safety investigations have previously found that passengers have a limited understanding and recall of important cabin safety information (for example, Flight Safety Foundation, 2006). These findings have been associated with emergency response procedures that are rarely required, such as emergency evacuation, brace position, using an oxygen mask, and wearing a life jacket. In contrast, passengers know how to fasten the seat belts, as they are required to do this at the beginning and end of each flight. As far as could be determined, none of the previous research examining the effectiveness of cabin safety communications has specifically examined passengers' understanding of when seat belts should be worn.

The investigation's questionnaire asked the passengers to answer the following question: 'Prior to the in-flight upset events, what was your understanding of when you should wear your seat belt?' The free-text responses were coded as follows:

- at all times when seated (64 responses)
- preferably or desirably at all times when seated (8)
- only during takeoff, landing or when the seat-belt sign was illuminated (13)
- unclear response or no response (13).

Table 35 provides more detailed results according to whether the respondents were wearing seat belts, seated without their seat belt on, or not seated. Passengers who were wearing seat belts were significantly more likely to 'understand' that seat belts should be worn at all times when seated, and passengers who were not wearing seat belts were significantly more likely to understand that seat belts should be worn only during takeoff, landing or when the seat-belt sign was illuminated.

Table 35: Passenger understanding of when seat belts should be worn

Passenger understanding	Seated, belt on	Seated, belt off	Seat belt use rate	Not seated
All times when seated	42	12	78%	9
Desirable when seated	4	4	50%	
Takeoff, landing, seat-belt sign illuminated	1	11	8%	
Unclear or no response	7	2		3

Note. Three questionnaire respondents who reported that they were wearing seat belts but that the seat belts became unfastened are not included in this table.

The proportion of passengers aged over 45 who understood that seat belts should be worn at all times when seated (87%) was higher than the proportion of passengers aged 18 to 45 who had this understanding (58%). The percentage of passengers from Australia who understood that seat belts should be worn at all times when seated (88%) was higher than the percentage from European countries (67%) or Asian countries (60%).

4.7.7 Previous seat belt use

The questionnaire asked passengers to state whether, on previous flights, they normally wore their seat belts. The question was asked for six different activities or phases of flight, as shown in Table 36.

Table 36: Passengers previous seat belt use

Activity	Worn, tightly fastened	Worn, loosely fastened	Not worn
Takeoff	93	2	0
Meal service	54	32	8
Reading, using in-flight entertainment system	56	30	8
Sleeping	46	42	7
Descent for landing	87	8	0
Landing	93	1	0

All passengers reported that they wore seat belts during takeoff, descent, and landing. Ten passengers (10% of respondents) reported that they did not normally wear their seat belts during one or more of the three cruise activities: meal service, reading/in-flight entertainment, and sleeping. All of these 10 passengers were seated at the time of the first upset and, in terms of their seat belt use:

- Two of the passengers were wearing seat belts. Both of these passengers also reported that they understood seat belts should be worn at all times when seated.
- Eight of the passengers were not wearing seat belts. Six of these passengers reported that they understood seat belts should be worn only during takeoff, landing or when the seat-belt sign was illuminated, and the other two reported that it was desirable to wear seat belts when seated.

A significant proportion of the questionnaire respondents (48%) reported that they normally wore their seat belts 'loosely-fastened' during one or more of the cruise activities. The investigation had no information to help determine how loose a seat belt was before a passenger regarded it as 'loosely fastened'. It is likely that there were significant differences in how passengers interpreted this term.

4.7.8 Reasons for not using seat belts

Of the 98 passengers who completed the questionnaire, 29 reported that they were seated without their seat belt fastened. The questionnaire asked those passengers 'why you were not wearing your seat belt'.

Ten of the passengers provided the following reasons:

- about to get up to go to the toilet or just returned from the toilet (seven responses)
- been to the toilet and then forgot to refasten their seat belt (three responses).

All of these 10 passengers reported that they normally wore their seat belts during the cruise activities. Nine of them also stated that they understood seat belts should be worn at all times when seated, and the other passenger reported understanding

that seat belts should be worn only during takeoff, landing or when the seat-belt sign was illuminated.

The other 19 questionnaire respondents who were not wearing seat belts provided the following reasons:

- recently returned from the toilet or walking around the cabin but then engaged in other activities, such as watching entertainment or sleeping (five responses)
- about to get up to go to the toilet but then engaged in other activities (one response)
- trying to get to sleep and finding the seat belt uncomfortable (one response)
- holding and playing with a baby in order to keep the baby still (one response)
- knowing that they would be returning to the toilet in the future (two responses)
- no particular reason (nine responses). Two of these passengers noted that it had been a smooth flight until that point, and another reported being fatigued and complacent.

Ten of these 19 passengers stated that they understood that seat belts should be worn during takeoff, landing or when the seat-belt sign was illuminated, four stated that they understood it was desirable to wear seat belts when seated, and three stated that they understood seat belts should be worn at all times when seated. Eight of these 19 passengers also reported that they did not normally wear their seat belts during one or more of the cruise activities.

Many of the non-seated passengers (13) were reportedly on their way to or returning from a toilet. The remainder were attending to their children, requesting items from a flight attendant, or letting a passenger seated next to them out of their seat.

4.7.9 Approaches for increasing passenger use of seat belts

Turbulence-related injuries have been a significant source of concern to the aviation industry. Consequently, there have been significant efforts to develop means of providing flight crews with more advanced warning of potential turbulence. There have also been efforts by regulatory authorities and other agencies to improve the crew procedures used by operators in preparing for and responding to turbulence-related events. For example, the US FAA released advisory circular (AC) 120-88A (*Preventing injuries caused by turbulence*) in 1996, which outlined suggested measures to operators. The operator's procedures as outlined in section 4.4.2 were generally consistent with the AC.

AC 120-88A also recommended activities 'to improve passenger compliance with seating and seat-belt instructions from crewmembers'. These included:

- Video presentations incorporated as part of a flight attendant's safety demonstration can illustrate the benefits of using effective turbulence practices.
- Articles in airline publications, pamphlets in seat back pockets or information on safety information cards can encourage passengers to engage in effective practices such as keeping seatbelts fastened at all times.

Regulatory authorities and other safety organisations have also provided publicly-available information for passengers emphasising the importance of wearing seat belts when seated. Examples of this information are provided in Appendix M.

Passengers from the 7 October 2008 occurrence provided suggestions on the passenger questionnaire for improving future safety, and most of the suggestions related to increasing passengers' use of seat belts. These suggestions included:

- Many passengers stated that it should be compulsory to wear a seat belt when seated. There have been proposals in the past to mandate that passengers wear seat belts when seated, or to require that the seat-belt sign is illuminated at all times. Regulatory authorities have expressed the view that such measures are impractical to enforce, and reduce the effectiveness of the seat-belt sign when it is most useful (Flight Safety Foundation 2001). Cabin crew have also advised that there are significant difficulties for them associated with attempting to enforce seat belt requirements when the seat-belt sign is not illuminated.
- Several passengers suggested placing more emphasis on the importance of wearing a seat belt during safety demonstrations and briefings, or providing more frequent reminders during the flight of the importance of wearing seat belts.
- One passenger suggested that seats should have sensors to detect when a passenger was seated and provide appropriate reminders to fasten the seat belt. Although such measures have been adopted for many new road vehicles, where seat belt use is a more significant problem, the extent to which it is a practical or cost-beneficial solution in aircraft is unknown.

As previously discussed, there has been very little research on the factors affecting the use of seat belts. There has also been very little research examining the effectiveness of different approaches to encouraging passengers to keep their seat belts fastened when seated.

4.8 Additional cabin safety matters

4.8.1 Cabin baggage

Several passengers commented that some of the cabin baggage placed in the overhead storage compartments was too large and too heavy for the compartments.

The operator had carry-on baggage limits which applied to all passengers in terms of the number, size and weight of bags. Although the size of a bag can readily be observed and checked, it is not easy for cabin crew to check the weight of baggage that passengers bring into the cabin.

Although the doors of several overhead storage compartments were dislodged during the occurrence, and other compartments opened, the investigation was not able to establish how much baggage actually fell out of the compartments. The investigation had no reports of passengers being injured by cabin baggage, although some injuries may have occurred. The investigation also could not determine the extent to which the damage to the storage compartment doors was due to loads shifting within the compartments or to external impact.

4.8.2 Handholds in the cabin

The FAA AC 120-88A recommended that aircraft cabins have handholds placed in appropriate locations around the cabin for cabin crew and passengers to hold on to in the event of unexpected turbulence. Recommended locations included the galleys, areas where passengers may be standing near toilets, and under overhead storage compartments. Some passengers on the 7 October 2008 flight stated that there should be handholds provided in the toilets.

The operator's A330-303 aircraft had handrails located in the galleys and handholds located throughout the cabin underneath the overhead storage compartments. There were no handholds located in or outside the toilets.

Although more handholds or handrails could help minimise injuries in some types of in-flight upsets, the first upset during the 7 October 2008 flight happened so quickly that it is unlikely that passengers or crew members would have been able to use a handhold or handrail unless they had been holding on to it prior to the event. In addition, the design and placement of additional handholds or handrails would need to be carefully considered, as they could potentially become a source of injury if hit by an occupant during an upset.[199]

[199] The flight attendant who was standing in the rear cabin impacted a handrail in the galley after hitting the ceiling, which exacerbated his injuries and also damaged the handrail.

5 ANALYSIS

5.1 Overview

The 7 October 2008 occurrence involving the Airbus A330-303 aircraft, registered VH-QPA (QPA), occurred when the aircraft suddenly pitched nose down while in cruise at FL370 (37,000 ft). A second, less significant pitch-down occurred 2 minutes later.

Data from the aircraft's recorders and simulations by the aircraft manufacturer showed that the pitch-downs were due to nose-down movements of the aircraft's elevators. The evidence also showed that the elevator movements were not initiated by turbulence, flight crew inputs, autopilot inputs, problems with the aircraft's weight or balance, or a technical fault with the elevators or other relevant parts of the electrical flight control system (EFCS).

The elevator movements were in fact commanded by the EFCS's flight control primary computers (FCPCs). More specifically:

- The FCPCs were designed to command a pitch-down if they detected that the aircraft's angle of attack (AOA) was too high. The relevant corrective mechanisms were high AOA protection and anti pitch-up compensation.
- A subsequent review of the FCPC algorithm for processing AOA data identified a very specific (and unintended) scenario in which incorrect AOA data from only one of the aircraft's three air data inertial reference units (ADIRUs) could trigger a pitch-down command. The scenario required two AOA spikes, with the second being present 1.2 seconds after the start of the first.
- Two minutes before the first pitch-down, ADIRU 1 started outputting spikes in AOA data (and other ADIRU parameters), and the spikes were present at the time of both pitch-downs.
- Simulations by the aircraft manufacturer confirmed that AOA spikes of the magnitudes recorded during the flight could initiate the elevator movements observed during the pitch-downs.

A summary of the main factors involved in the occurrence is presented in Figure 53. In essence, a design limitation with the FCPC software combined with an ADIRU failure to falsely activate the corrective mechanisms and produce the pitch-downs. The subsequent vertical accelerations led to a large number of injuries to the aircraft's occupants, with the number and extent of these injuries being exacerbated by many of the occupants not wearing seat belts.

This analysis discusses each of these factors, and several other topics of interest, including the reporting and recording of technical faults, flight crew performance, and cabin safety aspects.

Figure 53: Overview of the 7 October 2008 occurrence

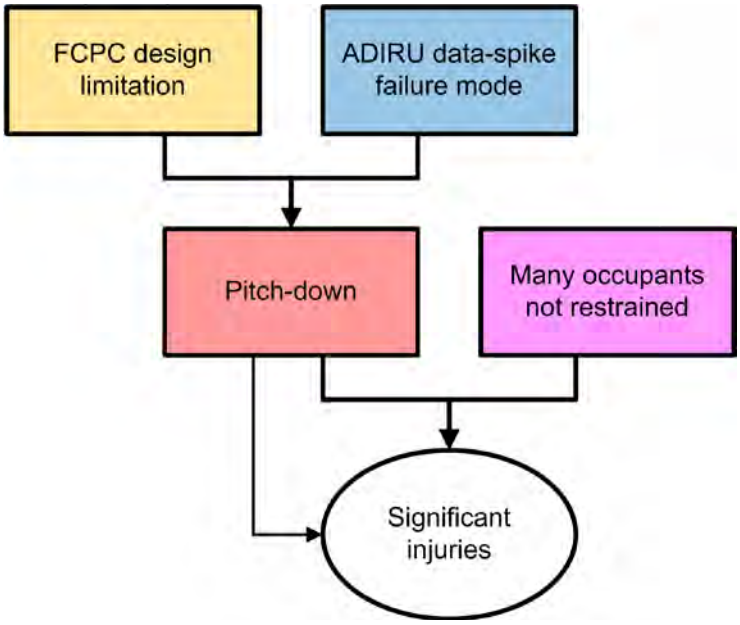

5.2 FCPC design limitation

5.2.1 Nature of the design limitation

AOA is a critically important flight parameter, and an aircraft with a full-authority flight control system (such as that on the A330 and A340) needs to be designed so that it obtains and uses accurate AOA information. The primary means of defence against an ADIRU providing incorrect AOA data to the FCPCs was the ADIRU itself, but this was not effective on the occurrence flight (section 5.3).

However, aircraft systems are designed with the expectation that technical faults will occasionally occur. Accordingly, the aircraft had three ADIRUs to provide redundancy and fault tolerance. Using the median of three values for a parameter as the system input is a common and generally robust algorithm, and the A330/A340 EFCS used this approach for most parameters. However, in order to address aerodynamic issues associated with the locations of the three AOA sensors, the FCPCs based the system input on the average value of AOA 1 and AOA 2. Nevertheless, they still used all three AOA values to check for consistency, as a basis for filtering out deviating values of AOA 1 and AOA 2, and for triggering a 1.2-second memorisation period using the previous value if an errant value of AOA 1 or AOA 2 was detected.

The FCPC algorithm was generally very effective, and could deal with almost all possible situations involving incorrect AOA data being provided by one ADIRU. It could manage step-changes, runaways, single spikes, and most situations involving multiple spikes or intermittently incorrect data. For example, the ADIRU data-spike failure mode occurred on 12 September 2006 with spurious stall warnings (and therefore AOA spikes) occurring over a 30-minute period with no reported effect on the aircraft's flightpath. On the 7 October 2008 flight, there were a large number of AOA spikes transmitted by ADIRU 1, and almost all of these were effectively filtered by the FCPCs.

Nevertheless, the FCPC's AOA algorithm could not effectively manage a scenario where there were multiple spikes such that one triggered a memorisation period and another was present 1.2 seconds later. The problem was that, if a 1.2-second memorisation period was triggered, the FCPCs accepted the next values of AOA 1 and AOA 2 after the end of the memorisation period as valid. In other words, the algorithm did not effectively handle the transition from the end of a memorisation period back to the normal operating mode when a second data spike was present.

5.2.2 Risk associated with the design limitation

The first in-flight upset resulted in a large number of injuries, some of them serious, and it was very distressing to many of the aircraft's occupants. However, it is very unlikely that the FCPC design limitation could have been associated with a more adverse outcome. More specifically:

- The 10° nose-down elevator command was very close to the highest magnitude possible from the EFCS's two corrective mechanisms. The second AOA spike of 50.6° resulted in the AOA value used by the FCPCs ($AOA_{FCPC\ input}$) being 26°. If the $AOA_{FCPC\ input}$ had been over 30°, the EFCS would have reverted to alternate law, which would have resulted in one of its corrective mechanisms (high AOA protection) not being active.
- There was limited potential for multiple pitch-downs of the same magnitude. As demonstrated during the occurrence flight, the fault-detection processes of the FCPCs would be expected to lead to the EFCS reverting to alternate law after two pitch-downs.
- The aircraft only descended a total of 690 ft during the first pitch-down. Although this was due in part to prompt action by the flight crew, the magnitude of the pitch-down would have been much less if the same AOA spike pattern had occurred when the aircraft was closer to the ground. Anti pitch-up compensation was not available when the aircraft was in the approach configuration or the speed was less than 0.65 Mach (which occurs during descent and initial climb). In addition, high AOA protection would have had no effect when the aircraft was below 500 ft above ground level. Flight simulations also showed that an undesired pitch-down just above 500 ft would be easily recoverable by a flight crew.
- If a pitch-down had occurred during climb or descent, more of the aircraft's occupants would have had their seat belts fastened (as the seat-belt sign would have been illuminated).

It is possible to conceive of situations where a flight crew could overreact to a significant nose-down command, which could result in more significant accelerations experienced in the cabin. In addition, if the cabin crew had been using service carts at the time, this could have led to more serious injuries. However, it would seem very unlikely that the pitch-down could have led to the loss of the aircraft or a large number of fatalities. Accordingly, the 7 October 2008 accident fitted the classification of a 'hazardous' effect rather than a 'catastrophic' effect, as defined by the relevant European certification requirements.

Determining the risk level of a failure condition, such as an undesired pitch-down command, involves considering the probability as well as the consequences of the condition. It is impossible to eliminate all potential hazards, and a system design

process needs to ensure that there is an inverse relationship between the severity of any adverse consequences and the probability of those consequences.

For a 'hazardous' effect level, the certification guidance material stated that the probability should be no more than 'extremely remote', which was nominally equivalent to a probability of 10^{-7} to 10^{-9} per flight hour. The FCPC design limitation only existed on A330/A340 aircraft, and it had existed since the aircraft types commenced operations in 1992. However, the 7 October 2008 occurrence was the only known case of a pitch-down command due to incorrect AOA data from one ADIRU in over 28 million flying hours on A330/A340 aircraft. This equates to a probability of less than 3.6×10^{-8} per flight hour, which is within the recommended range for 'hazardous' effects.[200]

Although the observed risk level was lower than the minimum certification requirement, the design limitation was still very undesirable and posed a significant threat to the safety of those on board the aircraft. The aircraft manufacturer took prompt action to address the problem, and subsequently redesigned its algorithm to eliminate the problem (section 7.1). Nevertheless, given that the design limitation existed, the investigation examined the reasons why it occurred in order to determine the lessons for future design processes.

5.2.3 Reasons for the design limitation

Non-awareness of the failure scenario

The development process for a safety-critical system has many elements to minimise the risk (probability and consequences) of a design error. These include peer reviews of design requirements, and system safety assessment (SSA), testing and simulation activities that are done as part of the verification and validation processes. It is widely accepted that not all the potential failure modes and failure scenarios for complex systems can be identified in practice, and fault-tolerant design features are included in a system to reduce the risk of such problems.

The A330/A340 FCPC algorithm for processing AOA data was redesigned after a problem was found with the initial algorithm during flight testing that was conducted before the aircraft type was certified. The redesign unintentionally introduced the design limitation in the algorithm, and the fault-tolerant features of the system were not able to fully mitigate the problem. The design limitation was not identified during the redesign activities. Although the SSA identified the relevant failure condition (incorrect, high AOA data leading to a pitch-down command), it did not identify the scenario that led to this condition on the 7 October 2008 flight. The results of the SSA and other design evaluation activities can be summarised as shown in Figure 54.

[200] The second pitch-down consisted of two commands very close together in time, and was considered to be a single event for the purposes of this type of assessment.

Figure 54: Summary of results of SSA activities for FCPC algorithm

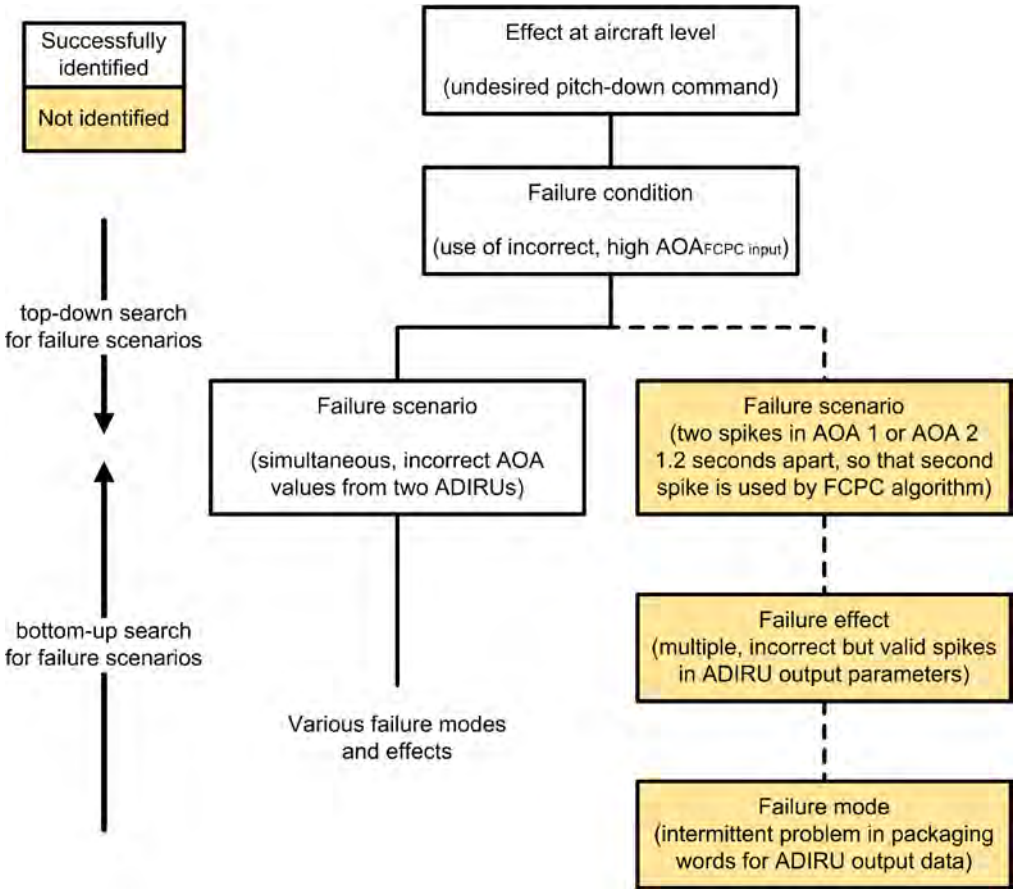

There were alternative algorithm designs or additional features that, in hindsight, would have prevented the design limitation or reduced its influence. Examples include rate limiting on the three AOA input values, range checks on the input values, or reasonableness checks involving comparisons between pitch and AOA. However, a system development process needs to balance many competing requirements, such as minimising the risk of introducing new failure conditions and design errors, minimising the data processing resources required, and unnecessary complexity. The inclusion of specific design features must be based on an identified need, and in this case the system development process did not identify the design limitation and therefore the need for any additional features.

The development of the new algorithm occurred in the period from 1991 to 1992, and determining the exact reasons why all of the development activities at that time did not identify the design limitation was made difficult by the amount of information available nearly 20 years later. Nevertheless, it is possible to discuss some contextual factors and inherent limitations associated with the system development process.

Limitations of using identified equipment failure modes

During design reviews and SSA activities, design engineers and safety analysts use a variety of approaches and sources of information to identify design problems (or failure scenarios) that will lead to the failure conditions of concern. A key approach is to use knowledge of how relevant items of equipment can fail or produce incorrect outputs, examine the effects of these failure modes in a particular design,

and determine whether they could lead to the failure condition of interest (a bottom-up approach). Important sources of information for this approach are the failure mode and effects analyses (FMEAs) conducted on relevant items of equipment, and past experience with similar equipment.

There are problems with the bottom-up approach. In the case of the FCPC algorithm for processing AOA data, the FMEA for the LTN-101 ADIRU did not identify the data-spike failure mode. In response to the aircraft manufacturer's specification, the ADIRU safety analysis included estimated probability data for particular types of incorrect ADIRU outputs. Although one of these outputs (indeterminate output) could be interpreted as being consistent with a multiple data-spike scenarios, it was also consistent with a wide range of other scenarios, and the ADIRU manufacturer's safety analysis did not discuss the possibility of multiple data spikes.

Although FMEAs use a systematic approach to identify component failure modes and their effects, they have limitations and make assumptions that affect their ability to identify all of the potential failure modes for complex systems. The reasons why the ADIRU manufacturer's FMEA and other development activities did not identify the data-spike failure mode could not be determined without knowing the exact nature of the failure mechanism involved (see also section 5.3.3).

As well as the ADIRU FMEA, previous experience was also not useful in this case. The aircraft and the ADIRU manufacturers, as well as the manufacturer of another type of ADIRU, reported that they had not observed a multiple data-spike scenario before. With continual changes and increasing complexity in equipment design, new types of failure modes will occasionally occur. The LTN-101 was a new model that was developed for the A330/A340, and therefore the validity of previous in-service experience was somewhat limited.

If the ADIRU FMEA or previous experience had specifically identified a realistic potential for multiple or frequent spikes in output data during a flight, then it is likely that the EFCS safety assessment activities would have looked more closely at the potential for data-spike patterns to create problems. It is worth noting that, once aware of the data-spike failure mode, the aircraft manufacturer reviewed its algorithms for processing other ADIRU parameters and identified limitations with some of these algorithms.

In summary, FMEAs and past experience are important sources of information but they have limitations, and they cannot be relied upon when identifying scenarios that lead to failure conditions, particularly for new, complex and highly-integrated safety-critical systems. As even relatively rare equipment failure modes can lead to problems for such systems, significant attention must also be devoted to other approaches to ensure that the system design is robust.

Other limitations of system safety assessment activities

In addition to examining the identified equipment failure modes, and determining whether they could produce the failure conditions of concern or be effectively managed, design engineers and safety analysts also need to search for 'theoretical' constraints or weaknesses in a system's design and use them to identify actual, specific failure scenarios. That is, they need to examine the proposed design itself to identify the scenarios or combinations of factors that could lead to the failure

conditions using a top-down approach. If such failure scenarios are identified, their probability and consequences can then be analysed.

When the FCPC algorithm for processing AOA data was developed in 1991 to1992, the aircraft manufacturer's top-down search for weaknesses was not completely effective. With complex, highly-integrated systems, detecting failure scenarios is difficult. The traditional processes for identifying failure scenarios use methods such as fault tree analysis that rely heavily on expert judgements, and there appears to have been limited guidance available to design engineers and safety analysts to assist with these judgements. This situation was not specific to any manufacturer.

In this case, the design limitation could probably have been identified if the designer engineers and safety analysts had conducted a systematic examination of the effects of all types of input values on the algorithm, for each of its modes of operation, and especially for the transitions between the modes. Based on the evidence available to the investigation, the extent that this was done could not be determined. Its thoroughness would have been adversely affected if any assumptions were made on the types of potential AOA input values that were expected. It should be noted that such a systematic examination would not be a simple task for a whole, complex system. However, for a specific design change to a small part of a system, or for safety-critical functions, it is more justifiable. It is also possible to conduct a partial examination of the effects of different input values on an algorithm, if an exhaustive examination is not feasible.

In recent years there have been many efforts directed at improving the efficiency and effectiveness of safety assessment activities, and the aircraft manufacturer has been significantly involved in many of these activities. These efforts include the development of improved guidance material for design engineers and safety analysts. However, it is also worth noting that system designs are generally becoming more complex over time in an effort to meet multiple competing objectives, including safety.

In addition to the development of general guidance material, a focus of recent development work has been on model-based development and automated safety analysis. These approaches will undoubtedly assist design engineers and safety analysts, and help simplify the nature of their tasks with complex systems in the future. However, it is not clear that they would have been effective for identifying the design limitation in this case. For example, automated safety analysis techniques again focus on the effects of known equipment failure modes (a bottom-up approach), and to date they have only dealt with relatively simple types of failure modes. A failure mode involving multiple incorrect inputs a specific time apart would probably be beyond the scope of this approach, at least at this time.

Limitations of design requirements and assumptions

Past research has shown that most software design problems for safety-critical systems are due to incomplete requirements, particularly with regard to the interaction between different systems. The development of complete and correct requirements is a very important part of a system development process, and SSA activities are one of the necessary activities for ensuring that the requirements are appropriate.

The design limitation with the FCPC's AOA algorithm appears to be an example of incomplete requirements. However, this characterisation may not be particularly

useful in this case. If the EFCS specification included a requirement for the AOA algorithm to be robust to all types of failures of a single ADIRU (and not just an AOA runaway), this would not have necessarily led to any additional analysis, simulation or testing of the design. According to the manufacturer, the development process had already considered other types of incorrect ADIRU outputs. A requirement that the design should be robust against multiple data spikes would probably have been more effective, but such a specific requirement would realistically only have been included if the design limitation had already been identified, or similar design problems had been identified in the past. In other words, it seems more useful to consider the incomplete requirements of the FCPC algorithm as a problem with the SSA and other design evaluation activities, rather than simply a problem with the requirements themselves.

The aircraft manufacturer advised that it had assumed during the development process that the algorithm was robust to any problem on a single ADIRU. There was no evidence that this assumption was formally stated in the SSA or the system specification. However, as already noted, the development process considered other types of incorrect ADIRU outputs that were known or expected, and including a formal assumption would not necessarily have led to any additional analysis, simulation or testing.

Limitations of simulation and testing activities

Another means of detecting a design problem is through the use of the simulation and testing activities conducted during the verification and validation processes. However, the selection of the simulations and tests needs to be prioritised based on an identified need, and this will usually focus on confirming that the design meets the specified requirements, and that it effectively manages identified failure modes or specific types of incorrect inputs. Any activities beyond the scope of verifying the explicitly-defined design requirements must rely on the expertise of those involved, which is as fallible as any other human activity.

Due to the wide range of potential inputs into a complex system such as the EFCS, simulation and testing programs cannot exhaustively examine all the possible patterns of inputs. In the case of the FCPC algorithm for processing AOA, the simulation and testing activities examined the new design's ability to handle the situation that led to the redesign. They also included previously identified tests to ensure there were no regression problems with the system design. However, they would not realistically have included a scenario involving multiple AOA data-spikes 1.2 seconds apart unless the potential problem had previously been identified.

Summary

Overall, the manufacturer's development process for the A330/A340 EFCS in the early 1990s included many appropriate, state-of-the-art safety assurance methodologies, and its SSA process was consistent with industry standards at the time. Nevertheless, a design limitation was inadvertently introduced during the redesign of the FCPC algorithm for processing AOA data.

The aircraft manufacturer's bottom-up search for failure scenarios was unlikely to be effective in identifying the design limitation because the ADIRU failure mode had not been previously encountered, or identified by the ADIRU manufacturer in its FMEA. The exact reasons why the top-down search processes did not detect the

problem could not be determined based on the available information. However, overall it can be concluded that the design, verification and validation processes (including safety assessment) used by the aircraft manufacturer did not fully consider the potential effects of frequent spikes in the data from an ADIRU.

This occurrence provided several lessons or reminders for the manufacturers of complex, safety-critical systems, and these lessons are discussed further in section 5.6.

5.3 ADIRU data-spike failure mode

5.3.1 Nature of the failure mode

The AOA data spikes that occurred during the 7 October 2008 flight were just one aspect of a specific failure mode involving the LTN-101 model ADIRU. The key features of the failure mode were as follows:

- The ADIRU outputted numerous spikes on air data reference (ADR) parameters. The spikes had short durations (less than 1 second), occurred at different times and frequencies for each parameter, and had a limited number of values for many of the parameters. With the exception of the data spikes, all of the ADR output data appeared to be correct. The ADIRU outputted the data spikes to other systems as valid data.
- The ADIRU also outputted numerous spikes on inertial reference (IR) parameters, with similar characteristics to the ADR spikes. In addition, the rest of the IR data varied from the expected values, usually showing some oscillatory characteristics. The ADIRU outputted almost all of its IR data to other systems as invalid data.
- The ADIRU generated an IR fault caution message but it did not generate an ADR fault message.
- Although some of the ADIRU's fault detection processes had worked (for example, to flag the IR data as invalid and generate an IR fault), no fault messages were recorded in the unit's built-in test equipment (BITE) memory. In addition, some routine BITE information was not recorded.
- Once the failure mode started, the ADIRU's abnormal behaviour occurred at a relatively constant rate until it was shut down. After its power was cycled (turned off and on), the unit performed normally (that is, the problem was a 'soft' fault).

Overall, the data-spike failure mode affected a wide range of the ADIRU's functional areas. The failure mode is only known to have occurred on three occasions, with very similar behaviour occurring on each occasion. Two of those occasions involved the same ADIRU (serial number 4167).

Although there were instances of ADIRU failures on other occasions, they did not exhibit the same effects on the output data. However, the effects on the recorded BITE data were similar to another LTN-101 failure mode known as dozing, which was also a soft fault. In contrast to the data-spike events, dozing was considered a benign failure mode since all data output from the ADIRU had ceased. There was insufficient evidence to determine whether any common contributing factors were involved in producing the two failure modes.

5.3.2 Risk associated with the failure mode

The main problem associated with the failure mode was that the ADR data spikes were not flagged as invalid (that is, some aspects of the failure were 'undetected' by the ADIRU), and the ADIRU did not shut down the ADR part or generate an ADR fault message advising the crew to select the ADR part OFF. Consequently, other systems would treat the data spikes as valid unless they had their own means of detecting and managing the incorrect data. The aircraft was fitted with three ADIRUs in order to provide the redundancy necessary to minimise any problems.

As a result, the potential consequences of the failure mode depended on the design of the aircraft on which the ADIRU was installed, particularly the capabilities and functions of the other systems that used ADIRU data. The most serious, known consequence of the failure mode on an A330 or A340 was what occurred on the 7 October 2008 flight: the FCPCs commanding a pitch-down movement because they did not effectively filter the AOA spikes. A review of the FCPC algorithms for processing other ADIRU parameters identified some other limitations, but these were of much less significance and only applied if another ADR was already unavailable.

In addition to the potential for a pitch-down command, other known effects associated with the data-spike failure mode on the A330/A340 included:

- a significant number of nuisance warning and caution messages, which could significantly increase crew workload (section 5.5.5)
- incorrect information being presented intermittently on one of the primary flight displays (PFD)
- disconnection of the autopilot (in those cases where the ADIRU involved was associated with the engaged autopilot)
- unavailability of some other aircraft systems, depending on the ADIRU involved. For example, in cases where ADIRU 1 was affected, the ground proximity warning system (GPWS) was no longer available.

In terms of the overall probability, the data-spike failure mode was very rare (or 'extremely remote' in terms of the language used in certification requirements), having only been observed on three occasions in over 128 million hours of unit operation. The LTN-101 ADIRU therefore met the relevant reliability requirements in terms of mean time between failures (MTBF), and it appeared to meet the relevant safety requirements in terms of the 'undetected failure' rate.

The LTN-101 was also installed on other aircraft types, and of these only the A320 had a full-authority fly-by-wire flight control system. However, the A320's algorithm for processing AOA data was different to that of the A330/A340 and was not affected by the ADIRU failure mode.

In summary, the primary hazard associated with the failure mode was a pitch-down command on the A330/A340 due to the limitation in the FCPC algorithm for processing AOA data. However, the ADIRU failure mode had the potential to create more problems than most undetected failures, and therefore investigating the reasons for the failure mode, or how to mitigate the effects of future occurrences, was important.

5.3.3 Reasons for the failure mode

Overview

The most significant aspect of the failure mode associated with the pitch-downs was that the ADR data spikes (particularly for AOA) were sent to the FCPCs and other systems as valid data. The elements involved in producing these data spikes are summarised in Figure 55. In simple terms, a trigger event combined with a susceptibility within the ADIRU to that type of trigger event to initiate a failure mechanism that disrupted the ADIRU's internal processing and generated the data spikes. A limitation in the coverage of the ADIRU's built-in test equipment (BITE) meant that the ADR data spikes were transmitted to other systems as valid data.

Figure 55: Simplified model of the ADIRU data-spike failure mode

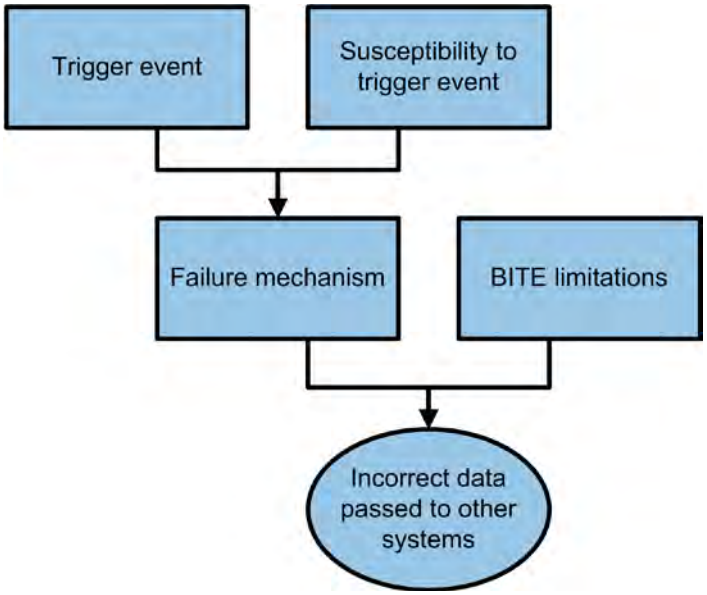

Similarities between the events

Prior to discussing each of the elements of the failure mode, it is important to consider the implications of some of the similarities between the three known data-spike events (12 September 2006, 7 October 2008, and 27 December 2008). These similarities include:

- Two events occurred on the same aircraft (QPA), and the other event occurred on another of the operator's aircraft with the same equipment configuration (QPG). However, no evidence was found to indicate that there was anything unusual with the operator's aircraft configuration, operating practices or maintenance practices. In addition, potentially relevant features of the aircraft, such as aircraft wiring or electromagnetic interference (EMI) from other aircraft systems, were tested and no problems were identified.
- All three events occurred to ADIRUs in the ADIRU 1 position. However, there was nothing significantly different about the wiring or connections for this position relative to the other two ADIRU positions. In addition, detailed examination and testing of the ADIRU wiring and connections on QPA did not identify any problems.

- All three events occurred in a broadly similar, geographical area (within a radius of 760 km). There was significant public interest during the investigation in the potential effects of transmissions from the Harold E. Holt Naval Communication Station, which was located in this area. However, based on several different types of evidence, the investigation found that it was very unlikely that the station's transmissions would have adversely affected the ADIRU. This includes evidence from ADIRU and aircraft testing, and the fact that the magnitude of the station's transmissions at the location of the three occurrences was several orders of magnitude below what the ADIRU (and aircraft) were designed and tested to tolerate.
- Two events involved the same ADIRU (unit 4167), and the other event involved an ADIRU with a similar configuration and serial number (unit 4122). There was nothing unique or anomalous with the affected units' software, and there were over 8,000 LTN-101 units in service. Accordingly, it was reasonable to conclude that some aspect of the affected units' hardware was probably associated with the failure mode.

Failure mechanism

A series of analyses determined that the failure mode almost certainly occurred with the ADIRU's central processing unit (CPU) module. More specifically, evidence indicated that many of the data spikes for ADR parameters were produced when the CPU module packaged the 32-bit output data words. These data spikes were found to be the result of the data word being packaged with either the wrong label field, or the wrong data field. Evidence also indicated that the CPU module's data processing stages before and after the data packaging operated normally for the ADR data.

The packaging problem was intermittent rather than consistent. In addition, it was very unlikely that each data spike was due to a separate fault, failure or trigger event. A much more likely scenario is that the failure initiated a problem in a higher-order process for organising the storage or retrieval of buffered data within the CPU module. Other symptoms of the failure mode, such as the corruption of the IR data and the storing of BITE data, also involved buffering data in memory. The investigation was not able to identify the precise mechanism involved, although some possibilities were able to be excluded (such as corruption of the data labels or problems with the wait states).

Unit susceptibility

As discussed above, the failure mode probably involved some form of hardware problem. The data-packaging process involved several components within the CPU module, including the CPU chip, application-specific integrated circuit (ASIC), a wait-state random access memory (RAM) chip, and four RAM chips for general CPU use. The investigation was not able to identify which of these component(s) were directly involved.

There can be variations in the properties of components with the same part number, including those manufactured within the same batch but even more so for those manufactured in different batches. The two affected ADIRUs had CPU-module components that were manufactured in the same or adjacent batches. Due to limitations in the way that component details were recorded, the investigation was unable to determine how many other ADIRUs contained components that were

manufactured in the same or adjacent batches to these units. However, it was known that most of the units manufactured after these units (after the end of 2002) had a redesigned CPU module with some different components.

In addition to not being common to all ADIRUs of the same design, the nature of the probable hardware problem was such that it only manifested itself on rare occasions or when presented with a rare form of triggering event. Since the LTN-101 model design met relevant design specifications, and the two affected units passed extensive testing, then it would also seem that the hardware limitation was marginal in nature.

Trigger types

Some of the possible triggering events identified by the investigation that could have initiated the failure mode included a software 'bug', software corruption, hardware fault, physical environment factors (such as temperature or vibration), and EMI (from other aircraft systems, other on-board sources, or external sources). Each of these possibilities was found to be unlikely or very unlikely based on multiple sources of evidence, which included ADIRU testing and an absence of evidence of the existence of the trigger event (or at least existence at a magnitude that would cause concern). The unit was also specifically manufactured to be resistant to all of these trigger types.

The other trigger type considered by the investigation was a single event effect (SEE). Although the intensity of high-energy particles was not unusual at the time of the three data-spike occurrences, such particles are always present. The CPU modules for the two affected units did not have error detection and correction (EDAC), which decreased their resilience to SEE compared to units manufactured after 2002 (which had EDAC installed). Previous testing of LTN-101 units, and components within the CPU module, had shown a level of susceptibility that was not unusual for systems manufactured at about the same time. Although the previous testing had not identified the data-spike failure mode, this testing was fairly limited in nature.

It would seem very unlikely that an SEE could occur at the same location within the same unit, and produce the same effect, without also occurring on many other units of the same type. However, susceptibility to SEE can vary significantly between components with the same part number, and there may have been more than one location that could produce the same effect from an SEE. In addition, having a particle strike in the same area on the same unit is conceivable given the level of exposure to high-energy particles that occurs at cruise altitudes.

Overall, the probability that the failure mode was triggered by SEE could not be reliably estimated without knowing the exact mechanism involved in the failure mode, or by demonstrating that the failure mode could occur during testing of the affected units. It was unfortunately not practicable for the investigation to test the units at an appropriate facility.

In summary, the investigation had sufficient evidence to conclude that most of the potential types of triggers were probably not associated with the data-spike failure mode. However, there was insufficient evidence available to determine whether SEE could have triggered the failure mode.

With the decreasing size of electronic devices, there is an increased risk of SEEs unless the device, or the overall system, is designed with appropriate mitigations in

place. Although aviation manufacturers have been paying more attention to SEE in recent years, it was only recently that formal guidance for such manufacturers was developed. The major certification authorities have stated that they expect manufacturers to address SEE hazards during system development processes, but there are at present no specific regulatory requirements in place. Overall, it would seem that more work is required to ensure that SEE is specifically and adequately considered in the development of all safety-critical aircraft systems.

Limitations with built-in test equipment

The LTN-101 ADIRU was designed so that almost all problems would be detected and, depending on the severity of the problem, appropriate action taken (such as sending a fault message, informing the flight crew, flagging the output data as invalid, or shutting the system down). The available evidence indicated that the BITE tests functioned as designed during the three data-spike occurrences. These tests resulted in the ADIRU flagging the incorrect IR data as invalid and generating an IR fault. Although no fault messages were recorded, this appeared to be a problem with the buffering of data within the CPU module rather than the execution of the BITE itself.

It was clear that the BITE did not successfully detect and manage the problem with the ADR data spikes. The unit's wraparound checks probably detected an ADR problem and sent a class 2 maintenance message to the central maintenance system (CMS). However, this response was not sufficient to generate a caution message for the flight crew. In addition, the ADIRU did not flag the ADR data as being invalid. The BITE included output parameter range checking, but the failure mode occurred after the range-checking tests had been performed. Even if the problem had occurred earlier in the processing sequence, most of the data spikes, including the AOA spikes, were within the allowable range and would not have failed a range test.

Overall, it would not have been practical to test every step of ADIRU processing, as the BITE complexity would increase substantially, resulting in possible adverse effects on ADIRU processing performance and reliability. The selection of BITE tests depends on the equipment specification and the safety assessment and other evaluation activities conducted during the system development process. In the case of the LTN-101, the FMEA and other system development processes did not identify the data-spike failure mode, and consequently did not introduce specific mitigators such as BITE tests to manage its occurrence.

Summary

The ADIRU data-spike failure mode occurred due to a combination of some form of trigger event, either external or internal to the unit, with a marginal susceptibility to that type of event within the CPU module of a limited number of units. This combination caused the ADIRU to enter a state that intermittently disrupted the CPU's processes for managing the storage and retrieval of temporary data, and the unit's BITE was not sufficient to detect some aspects of the failure mode, particularly the transmission of data spikes on ADR parameters.

Operationally, the LTN-101 ADIRU met the aircraft manufacturer's equipment specification in terms of its overall reliability rate and undetected failure rate. However, the data-spike failure mode had the potential to cause significant difficulties for other systems, and therefore lessons for preventing or mitigating the

effects of such failure modes need to be carefully considered for future systems. Some of these lessons are discussed in section 5.6.

5.4 Seat belts

5.4.1 Use of seat belts

In-flight upsets with the potential to cause injuries are relatively rare events. Although most often due to turbulence, they can also occur due to technical problems and/or flight crew actions. Regardless of the reason for an in-flight upset, the evidence from this accident, as well as previous accidents, shows conclusively that wearing a seat belt significantly decreases the likelihood of being injured and the severity of any injuries.

There are obviously legitimate reasons why passengers need to move around the cabin when the seat-belt signs are not illuminated. Therefore, any significant upset that occurs without warning during cruise on most flights will probably result in some injuries. On the occurrence flight, there were at least 23 passengers who were not seated at the time of the first in-flight upset.

However, there were also more than 60 passengers who were seated without their seat belts fastened. Although some of these passengers were in the process of getting in and out of their seats, the majority did not appear to have a valid reason for not wearing their seat belts. Some of these passengers appeared to routinely not wear their seat belts, and others said they normally wore them but forgot on this occasion. Some groups, such as younger adults, had lower seat belt use rates than other groups. The rate of non-wearing may have been higher than normal due to the timing of the upset, occurring after the meal service on a 5-hour flight during the day.

Previous research has indicated that some passengers do not comprehend some of the safety messages provided in passenger briefings, but these problems were found for more complex and very rarely required actions, such as emergency evacuations and using oxygen masks. Wearing seat belts whenever seated is a relatively simple action in comparison, and applicable on every flight.

In accordance with the operator's procedures, the crew provided multiple announcements at the beginning of the flight that advised passengers to keep their seat belts fastened when seated. Similar announcements had been used by most airlines for many years, and almost all the passengers would have heard the same messages many times before.

A small proportion of passengers stated that their 'understanding' of when they 'should' wear seat belts was during takeoff, landing and when the seat-belt sign was illuminated. Although this finding may be consistent with these passengers not comprehending the crew's safety announcements, it is more likely the case that they were aware of the recommendation but had a different perspective on the importance of following the recommendation.

Improvements in airline procedures have led to an increased emphasis on seat belt reminders by crews in recent years. Practicable techniques to further increase seat belt use when the seat-belt sign is not illuminated are limited. Illuminating the sign all the time would reduce the effectiveness of its use in higher-risk situations, and

requiring cabin crew to enforce the use of seat belts when the sign is not illuminated would create difficulties in the relationship between the crew and passengers. Engineering solutions, such as automated reminders built into the seats, would be difficult to implement and justify, particularly for existing seats.

The aviation industry needs to conduct further research into the reasons why some passengers do not wear seat belts, and the effectiveness of different communication techniques for increasing seat belt use. Limited research has been done in these areas to date. More frequent reminders during a flight, more variety in the communications, or messages targeted for specific demographic groups all have the potential to increase compliance. Using examples such as the present accident may also be useful in some types of communications.

5.4.2 How seat belts are worn

To be effective, seat belts need to be worn low and tight across the hips. Keeping a seat belt tight is most important during takeoff and landing, and it does not necessarily need to be as tight during cruise except during turbulent conditions. However, a seat belt should still be relatively firmly fastened during cruise, as a significant degree of slack will increase the risk of injury in the event of an unexpected upset.

A significant proportion (48%) of the 98 passengers who completed the ATSB questionnaire indicated that they wore their seat belt 'loosely' during cruise activities, and 16 of the 81 passengers that were known to be wearing seat belts at the time of the first upset reported that their belts were loosely fastened. However, there was little evidence that incorrect seat belt wearing contributed to injuries. The injury rate for passengers who said their seat belts were loosely fastened was no different to the rate for those who said their belts were tightly fastened, although only a small sample size was involved. It is possible that many passengers' understanding of the term 'loosely' simply meant that their seat belt was looser than during takeoff, even though it was still relatively firm.

Overall, the issue of how seat belts are worn during the cruise phase of flight appears to be a less important issue than whether the seat belts are actually worn in the first place. It is worth noting that operators typically instruct passengers to wear their seat belts 'low and tight' for the takeoff, but generally provide no reminders of how they should be worn after the takeoff. In some cases the advice to passengers indicates that the seat belts can be loosened after takeoff, but does not reinforce the need for the belts to still be relatively firm and worn across the hips. More detailed guidance for passengers on this topic would be useful.

5.5 Other aspects

5.5.1 Response to the 12 September 2006 event

In theory, the 12 September 2006 occurrence provided an opportunity for the ADIRU data-spike failure mode and the design limitation of the A330/A340 flight control system to be identified before the 7 October 2008 occurrence. In reality, based on the available information, there were good reasons for not conducting any further investigation at the time.

The flight crew of the 12 September 2006 flight discussed their event with maintenance watch, and completed a technical log entry. At the time, the event appeared to be primarily associated with nuisance ECAM messages. There was no autopilot disconnection, and no effect on the flight control system. The crew turned the ADIRU off in flight, and the anomalous behaviour ceased. Line maintenance personnel realigned the unit and conducted a system test, and no further problems were identified.

Incident reporting, and incident investigation in order to prevent accidents, are vitally important components of a safety management system. However, equipment faults on modern aircraft are not uncommon. In a system comprising multiple redundant units, the failure of one unit typically does not lead to any effect on the operation of a flight. Thorough investigations cannot be conducted into any incident or fault unless there are indications of an underlying problem that could influence future safety. In this case, the benefits of a detailed investigation were not obvious.

5.5.2 Systems for recording the in-service performance of equipment

An issue encountered by the investigation was that the full performance histories of the major line-replaceable units (LRUs) on an aircraft, such as the ADIRUs or FCPCs, were not able to be easily reviewed or evaluated. Equipment manufacturers generally only have details of in-service problems when a unit is removed from the aircraft and sent to them for examination. They do not normally receive information by operators on other in-service problems that do not warrant removal of the unit.

Operators generally record these in-service problems in a database, and the database record for each problem or event includes the aircraft and the type of equipment involved. However, it generally does not include the specific unit(s) that may be associated with the problem. Problems that recur over a short period could be readily identified with this system by reviewing recent entries. However, problems that recur over a longer period would not be readily identified, particularly if the units had been moved to a different position or another aircraft during that period.

The existing situation meant that the full extent to which specific units, or groups of units of a particular type, were experiencing or reporting faults was generally not being assessed in a systemic way by either operators or manufacturers. Manufacturers also did not have an accurate picture of the extent to which units were shut down during a flight due to reported performance problems. A more detailed recording system could therefore provide benefits for assessing the overall safety, reliability and availability of each type of unit, and better identify potentially problematic units.

Even if the operator's technical log database had recorded the unit associated with each log entry, this enhanced recording would not have prevented the 7 October 2008 occurrence. The history of previously reported faults for this unit was not frequent enough or atypical enough to have warranted further investigation.

5.5.3 Initial flight crew response

There was a 2-minute period between the commencement of the ADIRU failure and the flight control system's first pitch-down command. The only flight crew action that would have prevented this pitch-down command was to select the ADR part of ADIRU 1 OFF.

With the information available to the flight crew at the time of the occurrence, it was not reasonable to expect that they, or most crews in the same situation, would have made that selection in the available time. With very few exceptions, abnormal and emergency procedures on the A330 were carried out in response to ECAM messages. There was no ECAM message advising the flight crew of a NAV ADR 1 FAULT, or otherwise requiring the crew to select the ADR OFF. There were also no other procedures available for the situation they were experiencing.

If no specific action was recommended by the ECAM, or provided in other procedures, then the crew needed to evaluate the situation carefully before responding. The autopilot had disconnected, and there were nuisance stall and overspeed warnings, numerous caution messages on the ECAM, and problems with the air data information that were being provided on the captain's primary flight display. The crew had not encountered such a situation before, and the underlying source of all the problems was not obvious. More importantly, based on their available systems knowledge, they had no reason to suspect that any of the problems they were experiencing were associated with a condition that could adversely influence the flight control system.

The flight crew on the 12 September 2006 flight did select the ADR OFF. However, they only took this action after 30 minutes, and only after they detected an intermittent ADR 1 fault light on the overhead panel. On the 7 October 2008 flight, no fault light was illuminated. The post-flight report did indicate that a NAV ADR 1 FAULT was recorded 33 minutes after the ADIRU failure, but the amount of time that it was displayed on the ECAM was not clear. More importantly, it was not present prior to the two pitch-down commands.

The flight crew of the 27 December 2008 flight involving similar, anomalous ADIRU behaviour had new procedures available from the aircraft manufacturer as a result of the 7 October 2008 occurrence. When presented with the ADIRU failure mode, the crew promptly executed the procedures and selected the ADR OFF after 28 seconds, and this removed the potential for an inadvertent pitch-down associated with the incorrect ADIRU data.

5.5.4 Flight crew response to the pitch-downs

The captain's sidestick responses to both pitch-downs were prompt. The FCPC pitch-down commands were each present for about 2 seconds, and the captain's response to both occurred prior to the FCPCs being able to respond to these inputs. Given that the situation was sudden and unexpected, there was a risk that the flying pilot could have overcorrected (that is, provided an excessive sidestick response), which would have led to more severe vertical accelerations during the recovery. However, in addition to being timely, the captain's sidestick responses were also of the appropriate magnitude.

After the first pitch-down, the flight crew were presented with a situation that was even more confusing and much more serious, with several aircraft systems not functioning correctly. They continued to evaluate the available information and followed the recommended ECAM actions, but none of these actions were effective in preventing the second pitch-down.

Given the nature of their situation, the crew's decision to divert to Learmonth was appropriate. They had good reasons to consider that further undesired pitch-down commands could occur, and it was therefore safest to get the aircraft on the ground

as soon as possible. The diversion was also the quickest way to get medical assistance to the seriously injured passengers and crew. Learmonth was relatively close to the aircraft's position at the time, and it was a suitable aerodrome for an A330 landing.

Following the decision to divert, the crew continued their attempts to diagnose the problems, and also requested assistance from the operator's maintenance watch. These efforts did not resolve the situation and, as the flight progressed, they needed to focus more of their efforts on identifying and managing the logistical issues and threats associated with the diversion and landing. They also needed to maintain communications with the cabin, air traffic control and maintenance watch. These tasks were performed with a high degree of coordination and effectiveness by the flight crew.

Although the decision to divert was made at 0447, the aircraft did not land until 0532. This period of time was consistent with the time needed to cautiously descend the aircraft to minimise problems associated with another pitch-down.

5.5.5 Flight crew workload

Warning and caution messages

The crew experienced a significant workload during the occurrence. One of the main factors contributing to the workload was the frequently changing ECAM messages, together with the distracting noises from the caution chimes and stall warnings. The ECAM was designed to help manage abnormal and emergency situations by providing relevant, synoptic information and recommended crew actions. Due to the frequent and repetitive nature of the messages, the crew could not effectively interact with the ECAM to determine which messages were important. As well as increasing workload, this situation would have increased the crew's unwillingness to trust the aircraft's systems.

The ECAM performed as it was designed to perform. It had rules to prioritise the presentation of messages by failure level and recency, and for almost all abnormal and emergency situations these priority rules would work well. However, the rules could not effectively cater for the ADIRU data-spike failure mode, which involved a large number of messages at the same level and many of them frequently repeated. Redesigning the system to cater for this specific situation would be a major undertaking. In addition, any design change that allowed the flight crew to stop the presentation of fault messages to minimise distractions would introduce a risk of removing potentially important information.

The key problem with the fault messages for the data-spike failure mode was not the presentation of excessive or nuisance messages, but that the key message that would have resolved the problem (that is, NAV ADR 1 FAULT) was not presented in a timely manner. Improving the fault detection properties of the ADIRU would therefore seem to be more important than redesigning the ECAM to cater for a particular failure mode.

Use of autopilot

Operating the aircraft without an autopilot for the remainder of the flight also increased the captain's workload. The data-spike failure mode involving

ADIRU 1 affected the operation of autopilot 1 but it would not have affected autopilot 2. After autopilot 1 disconnected, the captain engaged autopilot 2, but he disconnected it shortly after.

The crew made no further attempts to use autopilot 2, which was understandable given their decreasing level of trust in the aircraft's systems. The use of manual control also enabled the captain to more quickly respond to any further pitch-downs. Although re-engaging autopilot 2 would have reduced the captain's workload, it would not have prevented the pitch-down commands from the FCPCs. In addition, autopilot 2 would have automatically disconnected during each of the pitch-downs.

5.5.6 Cabin communications

Following the first upset, the flight crew promptly advised passengers and crew to be seated and to fasten their seat belts, and they also repeated this message soon after the second upset.

Initially the flight crew were focused on evaluating and managing the problems with the aircraft's systems. They realised that the pitch-down was serious and would have resulted in injuries, and did not need any additional information from the cabin at that stage. After the first officer returned to the flight deck they had additional information about the extent of injuries and damage, and decided to divert to Learmonth.

Soon after deciding to divert, the flight crew requested further information from the cabin. Some of the initial information was provided by an off-duty cabin services manager (CSM) who contacted the flight deck directly. Ideally, communications from the cabin in this type of situation should occur through the on-duty cabin services manager (where available), as this will enable the manager to integrate the information and minimise unnecessary distractions for the flight crew. In this case, the off-duty CSM's communications provided useful information to the flight crew.

Overall, there was regular and effective communication between the flight deck and the cabin, and within the cabin crew. The flight crew also regularly kept the passengers informed of their situation.

Some of the passengers and crew were seriously injured and needed medical attention. The flight crew decided that they could not allow the cabin crew to leave their seats because of the risk of further injuries if another upset occurred. Given the uncertain situation they were experiencing, their rationale was justified. Diverting the aircraft, and landing as soon as possible, was the safest way of getting medical attention to those who were seriously injured.

5.6 Final comments and lessons for new systems

The investigation into the in-flight upset occurrence involving QPA on 7 October 2008 was difficult and took an extensive amount of time. It covered a range of complicated issues, including some that had rarely been considered in depth by previous aircraft accident investigations (such as system safety assessments and single event effects).

Ultimately, the occurrence involved a design limitation in the flight control system that had not been previously identified by the aircraft manufacturer, and a failure

mode with the ADIRU that had not been previously identified by the ADIRU manufacturer. Given the increasing complexity of such systems, this investigation has offered an insight into the types of issues that will become relevant for future investigations. It also identified a number of specific lessons or reminders for the manufacturers of new complex, safety-critical systems to consider. These include:

- System safety assessments (SSAs) and other design evaluation activities should recognise that ADIRUs and similar types of equipment can generate a wide range of patterns of incorrect data, including patterns not previously experienced.
- Failure mode effects analyses (FMEAs) have a limited ability to identify all equipment failure modes, particularly for complex, highly-integrated systems.
- Where practicable for safety-critical functions, SSA and other design evaluation activities should consider the effects of different values of system inputs in each mode of operation, particularly during transitions between modes.
- The BITE for ADIRUs and similar types of equipment should check the results of each key stage in the processing of output data.
- SEEs are a potential hazard to aircraft systems that contain high-density integrated circuits. Designers should consider the risk of SEE and include specific features in the system design to mitigate the effects of such events, especially in systems with a potentially significant influence on flight safety.
- The in-service performance records for safety-critical line-replaceable units should include all reported performance problems, not just those that result in the removal of the unit from the aircraft.
- The records for the key components within safety-critical systems should include details such as production or batch codes as well as the part number where practicable.

A broader lesson concerns the safety assessment activities needed for complex systems. In recent years there have been developments in the guidance material for system development processes and research into new approaches for SSA. However, design engineers and safety analysts also perform a safety-critical function, yet the investigation found little published research that has examined the human factors issues affecting such personnel. In other words, there has been limited research that has systematically evaluated how these personnel conduct their evaluations of systems, and how the design of their tasks, tools, training and guidance material can be improved so that the likelihood of design errors is minimised. The need for further research and development in this area will become more important as system complexity increases over time.

6 FINDINGS

From the evidence available, the following findings are made with respect to the in-flight upset involving the Airbus A330-303 aircraft registered VH-QPA that occurred 154 km west of Learmonth, Western Australia, on 7 October 2008. They should not be read as apportioning blame or liability to any particular organisation or individual.

Note: 'Safety factors' are events or conditions that increase risk. If a safety factor refers to a characteristic of an organisation or a system that has the potential to affect future safety, it is called a 'safety issue'. The ATSB classifies safety issues as critical, significant or minor depending on the level of associated risk, and it encourages relevant organisations to take safety action to address these issues. Further descriptions of these terms are provided in TERMINOLOGY USED IN THIS REPORT on page ix.

6.1 Contributing safety factors

- There was a limitation in the algorithm used by the A330/A340 flight control primary computers for processing angle of attack (AOA) data. This limitation meant that, in a very specific situation, multiple AOA spikes from only one of the three air data inertial reference units could result in a nose-down elevator command. *[Significant safety issue]*
- When developing the A330/A340 flight control primary computer software in the early 1990s, the aircraft manufacturer's system safety assessment and other development processes did not fully consider the potential effects of frequent spikes in the data from an air data inertial reference unit. *[Minor safety issue]*
- One of the aircraft's three air data inertial reference units (ADIRU 1) exhibited a data-spike failure mode, during which it transmitted a significant amount of incorrect data on air data parameters to other aircraft systems, without flagging that this data was invalid. The invalid data included frequent spikes in angle of attack data. Including the 7 October 2008 occurrence, there have been three occurrences of the same failure mode on LTN-101 ADIRUs, all on A330 aircraft. *[Minor safety issue]*
- The LTN-101 air data inertial reference unit involved in the occurrence (serial number 4167) also had a previous instance of the data-spike failure mode, indicating that it probably contained a marginal weakness in its hardware, which reduced the resilience of the unit to some form of triggering event.
- For the data-spike failure mode, the built-in test equipment of the LTN-101 air data inertial reference unit was not effective, for air data parameters, in detecting the problem, communicating appropriate fault information, and flagging affected data as invalid. *[Minor safety issue]*
- The air data inertial reference unit manufacturer's failure mode effects analysis and other development processes for the LTN-101 ADIRU did not identify the data-spike failure mode.

- Although passengers are routinely reminded to keep their seat belts fastened during flight whenever they are seated, a significant number of passengers have not followed this advice. At the time of the first in-flight upset, more than 60 of the 303 passengers were seated without their seat belts fastened. *[Minor safety issue]*

6.2 Other safety factors

- In recent years there have been developments in guidance materials for system development processes and research into new approaches for system safety assessments. However, there has been limited research that has systematically evaluated how design engineers and safety analysts conduct their evaluations of systems, and how the design of their tasks, tools, training and guidance material can be improved so that the likelihood of design errors is minimised. *[Minor safety issue]*
- The large number of spurious warnings and caution messages that resulted from the anomalous air data inertial reference unit behaviour created a significant amount of workload and distraction for the flight crew.
- Single event effects (SEE) have the potential to adversely affect avionics systems that have not been specifically designed to be resilient to this hazard. There were no specific certification requirements for SEE, and until recently there was no formal guidance material available for addressing SEE during the design process. *[Minor safety issue]*
- The LTN-101 air data inertial reference unit (ADIRU) model had a demonstrated susceptibility to single event effects (SEE). The consideration of SEE during the design process was consistent with industry practice at the time the unit was developed, and the overall fault rates of the ADIRU were within the relevant design objectives. *[Minor safety issue]*
- Industry practices for tracking faults or performance problems with line-replaceable units were limited, unless the units are removed for examination. Consequently, the manufacturers of aircraft equipment have incomplete information for identifying patterns or trends that can be used to improve the safety, availability or reliability of the units. *[Minor safety issue]*
- There has been very little research conducted into the factors influencing passengers' use of seat belts when the seat-belt sign is not illuminated, and the effectiveness of different techniques to increase the use of seat belts. *[Minor safety issue]*
- Although passengers are routinely advised after takeoff to wear their seat belts when seated, this advice typically does not reinforce how the seat belts should be worn. *[Minor safety issue]*

6.3 Other key findings

- As of the end of 2009, A330/A340 aircraft had accumulated over 28 million flight hours. The occurrence on 7 October 2008 was the only occasion when incorrect data from an air data inertial reference unit had resulted in inadvertent elevator commands. This in-service performance was consistent with the relevant certification requirements.

- As of April 2010, the LTN-101 air data inertial reference unit (ADIRU) had accumulated over 128 million hours of operation. The data-spike failure mode had only been observed on three occasions. The ADIRU's in-service performance met the aircraft manufacturer's safety and reliability objectives.
- Air data inertial reference unit (ADIRU) 2 and ADIRU 3 operated normally throughout the flight.
- Although air data inertial reference unit 1 transmitted a significant amount of incorrect data on inertial reference parameters to other aircraft systems, almost all of this data was flagged as invalid.
- It is very likely that the air data inertial reference unit (ADIRU) data-spike failure mode involved a problem with the data packaging and queuing within the ADIRU's central processing unit module. This fault resulted in numerous data anomalies, including air data reference parameters being intermittently transmitted with the data or label of another parameter. Despite extensive testing and analysis, the exact origins of the failure mode could not be determined.
- Tests and analyses showed that the air data inertial reference unit data-spike failure mode was probably not triggered by a software bug, software corruption, hardware fault, physical environment factors (such as temperature of vibration), or from electromagnetic interference.
- The three known occurrences of the air data inertial reference unit data-spike failure mode occurred on two A330 aircraft operated by the same operator; however, no factors related to the operator's aircraft configuration, operating practices, or maintenance practices were identified that were associated with the failure mode.
- The flight crew's responses to the warnings and cautions, the pitch-down events, and the consequences of the pitch-down events, demonstrated sound judgement and a professional approach.
- Wearing a seat belt during all phases of a flight, and having the seat belt fastened low and firm, will significantly minimise the risk of injury in the unlikely event of an in-flight upset.

7 SAFETY ACTION

The safety issues identified during this investigation were communicated to the relevant organisations during the investigation. In addition, these organisations were given a draft report and asked to communicate what safety actions, if any, they had carried out or were planning to carry out in relation to each safety issue.

For a critical or significant safety issue, the Australian Transport Safety Bureau (ATSB) expects the relevant organisation(s) to take safety action to address the issue. If appropriate safety action is not taken, the ATSB may issue a formal safety recommendation or a safety advisory notice.

For a minor safety issue, the ATSB notes that the associated risk is considered broadly acceptable. The ATSB still encourages the relevant organisation(s) to take safety action, but it does not issue a formal recommendation or a safety advisory notice.

When the ATSB has been advised of safety action in response to a safety issue, it is published in the final report.

7.1 Flight control primary computer issues

7.1.1 Algorithm for processing AOA data

Significant safety issue

There was a limitation in the algorithm used by the A330/A340 flight control primary computers for processing angle of attack (AOA) data. This limitation meant that, in a very specific situation, multiple AOA spikes from only one of the three air data inertial reference units could result in a nose-down elevator command.

Procedural changes issued by Airbus

On 15 October 2008, the aircraft manufacturer issued Operations Engineering Bulletin (OEB) OEB-A330-74-1, which was applicable to all A330 aircraft fitted with Northrop Grumman ADIRUs.[201] The OEB stated that, in the event of a NAV IR [1, 2 or 3] FAULT (or an ATT red flag being displayed on either the captain's or first officer's primary flight display), the flight crew were required to select the air data reference (ADR) part of the relevant ADIRU OFF and then select the relevant inertial reference (IR) part of the relevant ADIRU OFF. The problem was described as a 'significant operational issue' and operators were advised to inform their pilots of the OEB without delay and insert the procedure in the Flight Crew Operations Manual. A compatible temporary revision was issued to the Minimum Master Equipment List at the same time.

The OEB procedure was subsequently amended in December 2008 to cater for a situation where the IR and ADR pushbuttons were selected OFF and the OFF lights

[201] Operators were advised of the OEB and the associated problem in an operator information telex that was issued on 14 October 2008.

did not illuminate. The new OEB (A330-74-3) required crews to select the IR mode rotary selector to the OFF position if the lights did not illuminate.

Following the 27 December 2008 occurrence, the aircraft manufacturer issued another OEB (A330-74-4, 4 January 2009). This OEB provided a revised procedure for responding to a similar ADIRU-related event to ensure incorrect data would not be used by other aircraft systems.[202] The procedure required the crew to select the relevant IR OFF, select the relevant ADR OFF, and then turn the IR mode rotary selector to the OFF position.

The manufacturer issued similar OEBs to operators of A340 aircraft.

Procedural changes mandated by certification authorities

The OEBs detailed above were subsequently issued as airworthiness directives by the relevant regulatory agencies. More specifically:

- The European Aviation Safety Agency (EASA) issued the procedure in OEB-A330-74-1 as Emergency Airworthiness Directive 2008-0203-E, effective on 19 November 2008. The Australian Civil Aviation Safety Authority (CASA) subsequently issued Airworthiness Directive AD/A330/95 for Australian operators, effective on 20 November 2008.
- EASA issued the procedure in OEB-A330-74-3 as Emergency Airworthiness Directive 2008-0225-E, effective on 22 December 2008, and CASA issued AD/A330/95 Amendment 1, effective on 22 December 2008.
- EASA issued the procedure in OEB-A330-74-4 as Emergency Airworthiness Directive 2009-0012-E, effective on 19 January 2009, and CASA issued AD/A330/95 Amendment 2, effective on 19 January 2009.

The EASA directives also applied to A340 aircraft.

Procedural changes taken by the operator

On 15 October 2008, in response to the initial advice from the aircraft manufacturer, the operator issued Flight Standing Order 134/08 for its A330 operations. On 24 October 2008, this order was replaced by Flight Standing Order 136/08, which incorporated the material from the initial Airbus OEB. In addition, a program of focused training during simulator sessions and route checks was initiated to ensure that flight crew undertaking recurrent or endorsement training were aware of the contents of the Flight Standing Order. Subsequent Flight Standing Orders were issued in response to the modified OEBs in December 2008 and January 2009.

Redesign of FCPC algorithms

The aircraft manufacturer introduced an interim modification to the flight control primary computer (FCPC) software standard (P9A/M18A) that was promulgated using a Service Bulletin. The interim standard incorporated the modified monitoring

---

[202] During the 27 December 2008 occurrence, the flight crew promptly applied the OEB 74-3 procedure. This procedure successfully stopped the transmission of ADR data from the affected ADIRU, but it did not stop the transmission of IR data. The revised procedure addressed this problem.

and filtering of five parameters, including AOA. The standard was retrofitted to the operator's fleet of A330 aircraft, and completed in November 2009.

The revised algorithm for processing AOA data again based $AOA_{FCPC\ input}$ on the average of AOA 1 and AOA 2, but it did not include the 1.2-second memorisation period. Additional processes were used to monitor the consistency of the three AOA values The new algorithm also introduced a mechanism to monitor the overall consistency or oscillation of each AOA, with the associated ADR being rejected for the remainder of the flight if a problem was detected. In the event that an ADR was rejected due to a problem with AOA data, the flight warning system (FWS) would not issue any spurious stall warnings associated with that ADR's data.

Subsequent FCPC software standards were developed for use on all A330/A340 aircraft. These later standards included the redesign of the AOA algorithm (as discussed above), as well as modified algorithms for a number of other ADIRU parameters used by the FCPCs. During 2011, the new software standards were certified by EASA for all but one of the A330/A340 models. The standard for the last model was expected to be certified in February 2012.

When retrofit action has been completed, the aircraft manufacturer (in consultation with EASA) will cancel the relevant OEBs.

ATSB assessment

The ATSB is satisfied that the action being taken by the aircraft manufacturer will satisfactorily address the safety issue.

7.1.2 Processes for developing flight control computer algorithms

Minor safety issue

When developing the A330/A340 flight control primary computer software in the early 1990s, the aircraft manufacturer's system safety assessment and other development processes did not fully consider the potential effects of frequent spikes in the data from an air data inertial reference unit.

Action taken by Airbus

Following the 7 October 2008 occurrence, the aircraft manufacturer reviewed the FCPC algorithms for processing each ADIRU parameter on the A330/A340 aircraft. The review examined the amplitude, duration and frequency of data spikes that could potentially affect the FCPC's control of the aircraft's flightpath. In addition to data spikes, other potential incorrect data patterns were also considered. Based on this review, modifications were made to the algorithms for processing a number of the ADIRU parameters used by the FCPCs.

In addition to the A330/A340, the manufacturer also reviewed the algorithms used for processing ADIRU parameters by the flight control computers on the A320 and A380 aircraft.

The manufacturer also advised that it will apply the lessons learnt from the 7 October 2008 occurrence in terms of the types of incorrect data patterns to be taken into account during future design definition and modification. Accordingly,

modifications were made to its guidance document ABD0200 (*Guidelines and requirements for the system designers*).

7.2 Air data inertial reference unit issues

7.2.1 Data-spike failure mode occurrences

Minor safety issue

One of the aircraft's three air data inertial reference units (ADIRU 1) exhibited a data-spike failure mode, during which it transmitted a significant amount of incorrect data on air data parameters to other aircraft systems, without flagging that this data was invalid. The invalid data included frequent spikes in angle of attack data. Including the 7 October 2008 occurrence, there have been three occurrences of the same failure mode on LTN-101 ADIRUs, all on A330 aircraft.

Action taken by Northrop Grumman Corporation

The ADIRU manufacturer advised in November 2010 that it had examined a wide range of possible mechanisms within the LTN-101's central processing unit (CPU) module that may have produced the air data reference data spikes. Although the exact mechanism could not be identified, it was considering options to improve the robustness of some of the CPU module's processing activities.

7.2.2 Design of built-in test equipment

Minor safety issue

For the data-spike failure mode, the built-in test equipment of the LTN-101 air data inertial reference unit was not effective, for air data reference, in detecting the problem, communicating appropriate fault information, and flagging affected data as invalid.

Action taken by Northrop Grumman Corporation

The ADIRU manufacturer conducted a detailed review of options for improving the LTN-101 BITE so that another instance of the data-spike failure mode would be detected. The options included comparisons of the data and/or label fields that were expected to be outputted on certain parameters with the data and/or label fields actually outputted. However, during testing it was found that these modification options adversely affected the performance of high-frequency tasks.

Subsequently, at the aircraft manufacturer's request, the ADIRU manufacturer enhanced the wraparound monitor test on the dummy label (section 3.7.3) to improve the ADIRU's ability to detect data transmission failures. More specifically:

- The filter was reduced from three consecutive failures to two.
- A class 1 maintenance fault message will be declared if two output databuses fail from either the IR part or the ADR part (previously all databuses had to fail from one part to reach a class 1 level).

- In the event of a class 1 message, all data transmissions from the failed part will be terminated.
- A class 1 failure message will also be reported from the surviving partition to aid in fault identification.

The ADIRU manufacturer advised in October 2011 that the enhancements had passed testing, with certification expected by the end of 2011.

7.2.3 Susceptibility to single event effects

Minor safety issue

The LTN-101 air data inertial reference unit (ADIRU) model had a demonstrated susceptibility to single event effects (SEE). The consideration of SEE during the design process was consistent with industry practice at the time the unit was developed, and the overall fault rates of the ADIRU were within the relevant design objectives.

Action taken by Northrop Grumman Corporation

As discussed in section 3.6.6, the ADIRU manufacturer conducted a 'theoretical analysis' of the potential for a single event upset (SEU) on the LTN-101 ADIRU. The overall result of this analysis was that, after considering the influence of SEEs, the ADIRU still met the aircraft manufacturer's safety objectives.

7.2.4 Other related safety action

The aircraft manufacturer, the ADIRU manufacturer and the operator examined the feasibility of recording additional parameters from ADIRU 1 on the aircraft's quick access recorder (QAR). Such parameters would provide useful information in the event of a data-spike occurrence or other types of occurrences. The required wiring changes to the operator's A330 aircraft to access an additional ADIRU 1 databus and modifications to the aircraft condition monitoring system (ACMS) software. The examination showed that the modification was feasible and, at the time of the publication of this report, the additional parameters were being recorded on all Qantas and Jetstar A330 aircraft.

The ADIRU manufacturer also advised that it was adding a function to the LTN-101 ADIRU's executive partition to record if the fault monitor partition shut down, together with the elapsed time interval (ETI) and the specific monitor that failed. This change would provide more diagnostic information in the event of a failure involving the fault monitor partition.

7.3 Use of seat belts

Minor safety issue

Although passengers are routinely reminded to keep their seat belts fastened during flight whenever they are seated, a significant number of passengers do not follow

this advice. At the time of the first in-flight upset, more than 60 of the 303 passengers were seated without their seat belts fastened.

Action taken by the operator

In October 2011, the operator advised that it was evaluating the feasibility of measuring the seat belt use rate of passengers on a sample of flights. The measurements would include various stages of flight, and situations when the seat belt sign is illuminated and also when it is not illuminated. Based on the results, the operator could then determine if additional specific communications or actions were warranted.

Other safety action regarding seat belt use

As outlined in Appendix M, regulatory authorities and other safety organisations have provided publicly-available information for passengers emphasising the importance of wearing seat belts when seated. With specific relevance to the 7 October 2008 occurrence, additional reminders have been provided as follows:

- In its media statements providing updates on the investigation on 8 and 10 October 2008, the ATSB noted that the 7 October 2008 accident served as a reminder to all people who travel by air of the importance of keeping their seat belts fastened at all times when seated in an aircraft.
- On 27 October 2008, CASA issued a media release that stated that the occurrence was as a timely reminder to passengers to 'remain buckled up when seated at all stages of flight'. The media release also highlighted the importance of passengers following safety instructions issued by flight and cabin crew, including watching and actively listening to the safety briefing given by the cabin crew at the start of each flight.
- With the release of the final ATSB report into the 7 October 2008 occurrence, the ATSB web site highlighted the importance of wearing seat belts when seated and provided links to relevant reference sources for passengers.

7.4 Single event effects

Minor safety issue

Single event effects (SEE) have the potential to adversely affect avionics systems that have not been specifically designed to be resilient to this hazard. There were no specific certification requirements for SEE, and until recently there was no formal guidance material available for addressing SEE during the design process.

Action taken by regulatory authorities

As discussed in section 3.6.6, the International Electrotechnical Commission (IEC) published Technical Specification (TS) 62396-1 in 2006 and subsequent parts soon after.

In March 2010, the FAA advised that IEC TS 62396 was being considered for inclusion in future Aerospace Recommended Practices (ARPs), as well as FAA Issue Papers and Special Conditions. The FAA also advised that the Aerospace

Vehicle Systems Institute (AVSI) helped develop TS 62396, and that one of AVSI's goals is to encourage the establishment of additional test facilities to enable more cost-effective determination of the SEE immunity of systems or units. Additional facilities would further the understanding of new electronic parts' susceptibility and for developing regulations and guidance material.

In May 2010, EASA advised that on recent aircraft certifications (including the A380 and A350), it had requested that the applicant consider the effects of single event upsets (SEUs) and multiple bit upsets (MBUs) on systems and equipment. It also advised that Airbus had required equipment manufacturers to consider the effects of SEU and MBU and to mitigate these effects. As discussed in section 3.6.6, Airbus had included IEC TS 62396 in its requirements for manufacturers since 2007.

APPENDIX A: VERTICAL ACCELERATIONS

The acceleration of an aircraft can be resolved into three components along the vertical, lateral and longitudinal axes. In the case of the in-flight upsets on 7 October 2008, there was virtually no lateral acceleration involved and the changes in longitudinal acceleration were small compared to those for vertical acceleration. Therefore this analysis has only considered the effects of vertical acceleration.

Acceleration is measured in units of 'g', which equates to about 9.8 m/second^2. The nominal value for vertical acceleration is 1 g, which is the value recorded when an aircraft is stationary on the ground. In flight, vertical acceleration represents the combined effects of flight manoeuvring loads and turbulence. A negative vertical g indicates a downwards acceleration of the aircraft.

Acceleration values were sensed by a triaxial accelerometer and were recorded by the aircraft's flight data recorder (FDR). The accelerometer was mounted in a fixed location under the floor in the centre fuselage area, which was within the aircraft's normal centre of gravity (c.g.) range. The aircraft manufacturer advised that the c.g. at the time of the two events was about 25%, which corresponded to a location about 30 m to the rear of the aircraft from the aircraft's nose.

Figure A1 shows the position at which the FDR recorded vertical acceleration data (at about row 35). Based on the FDR data, the aircraft manufacturer also calculated the accelerations at a position 29 m forward of the c.g. (corresponding to the flight crew seats), and a position 18 m behind the c.g. (corresponding to about row 59).

Figure A1: Reference positions for the vertical acceleration data

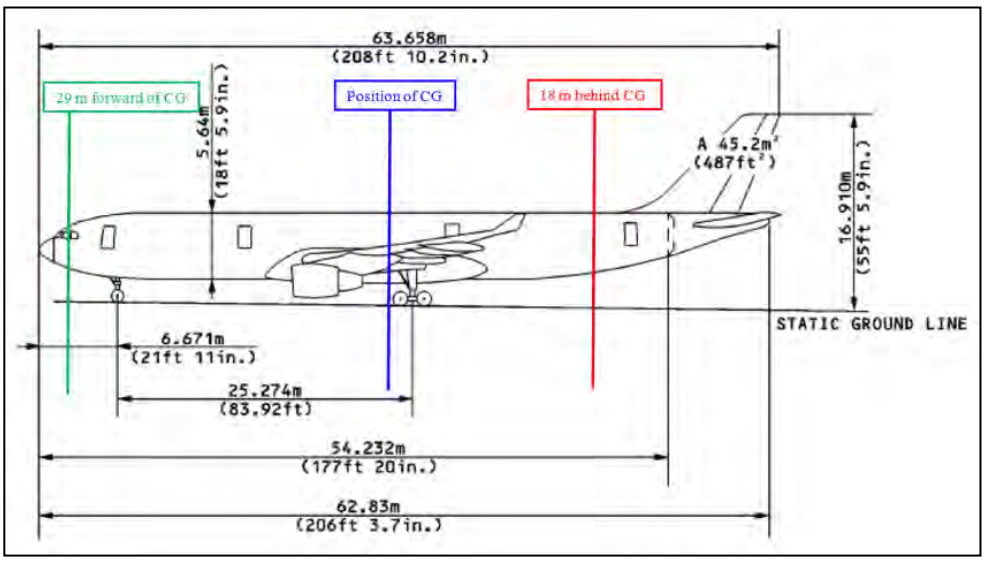

Graphs of the accelerations for both upsets are provided in Figure A2. As shown in the figures, the accelerations were more significant during the first upset. The accelerations were also more significant at the rear of the aircraft than in the centre or front of the aircraft.

Passengers who were seated with their seat belts fastened would have experienced the same accelerations, appropriate to their location, as the calculated values for the aircraft. The situation for passengers who were not restrained was more complex. In this case, the relative motion between the aircraft and the passengers would have

been important. The situation appeared to be more severe at the rear of the aircraft, where passengers who were unrestrained would have been initially accelerated upwards (as the nose came down and the tail went up) but the aircraft would have then started to descend. As a result, at the rear of the aircraft, there would have been relative motion between the aircraft and the passengers that could have combined to increase the force of impact with the aircraft cabin.

Figure A2: Vertical accelerations during the first (top) and second in-flight upsets (bottom)

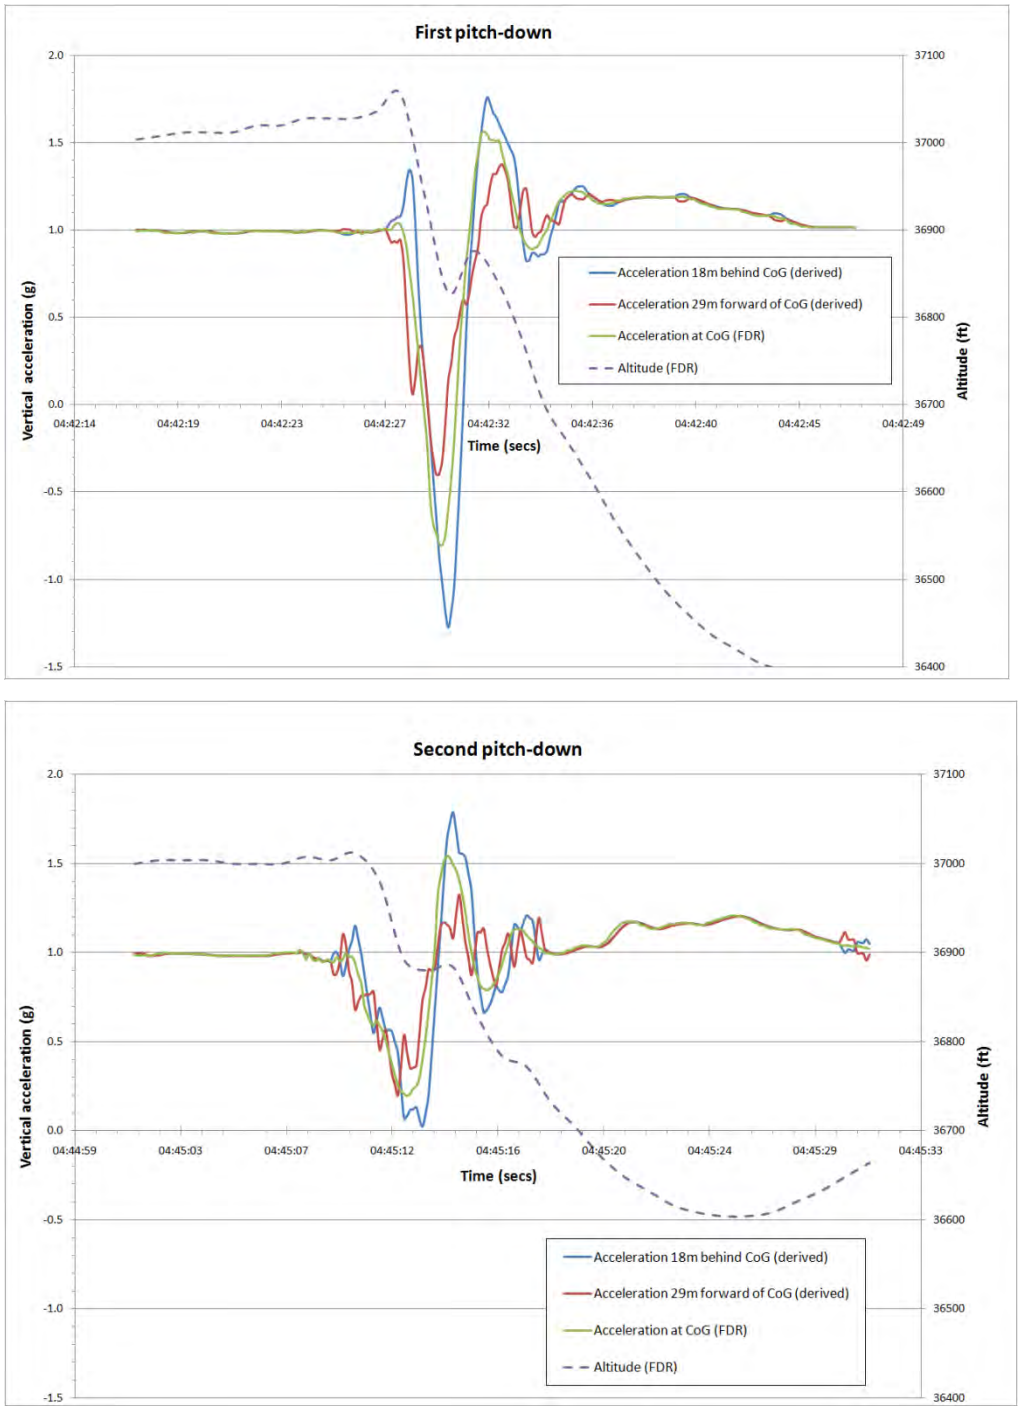

APPENDIX B: FLIGHT RECORDER INFORMATION

Recorded data signal paths

The flight data recorder (FDR) and quick access recorder (QAR) obtained data on air data inertial reference unit (ADIRU) parameters through independent signal paths, as shown in Figure B1.

Figure B1: Signal paths for the FDR and QAR systems

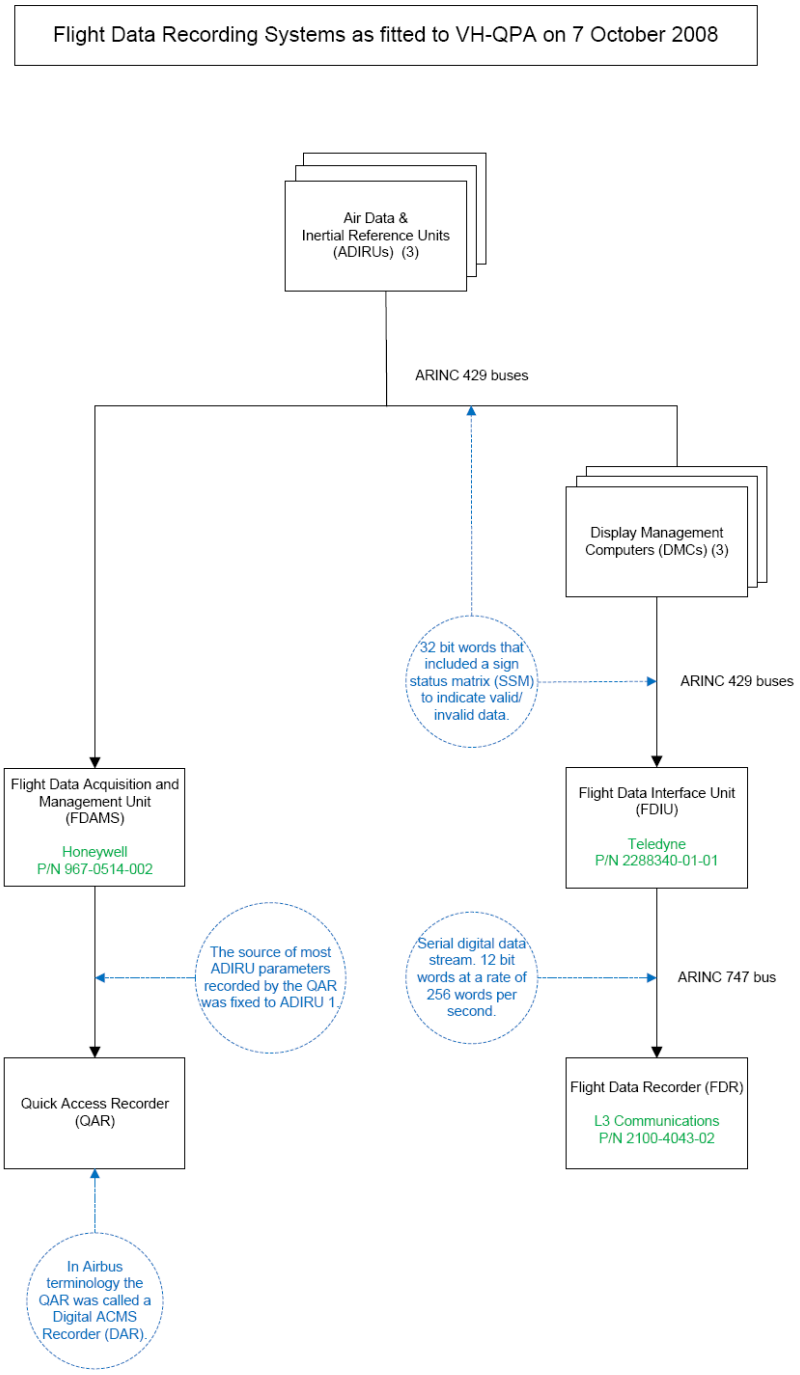

Means of recording parameter invalidity

The FDR obtained its ADIRU data through the flight data interface unit (FDIU) (Figure B1). The FDIU extracted the data bits from the ARINC 429 data words that had been programmed to be recorded by the FDR. Only the data bits were sent to the FDR for recording, and not the sign/status matrix (SSM) bits, which indicated the validity of the data word. To provide some indication of when the data bits for a particular word had been flagged by the ADIRU as invalid, or if the word was missing (or not refreshed), the FDIU cycled the data bits depending on the error.

A typical cycle for a parameter that was recorded once per second, when the FDIU had detected an error, was:

- 1st second: data
- 2nd second: error code (1,536 for a parity error, 2,048 for a missing word, 2,560 for a failure warning, and 3,072 for no computed data)
- 3rd second: data
- 4th second: zero.

This cycle was then repeated until the reason for the error code disappeared.

The FDR processing software scaled these error codes as if they were real data. For example, in the case of angle of attack (AOA), an error code of 3072 (no computed data) was scaled as -45°. A routine example of this behaviour can be seen during takeoff and landing at low airspeeds (Figure B2).

As noted during the 7 October 2008 flight, there were some corrected AOA values of -45°. These values were brief indications of no computed data rather than of data spikes with a -45° value.

Figure B2: Parameter cycling due to FDIU error code

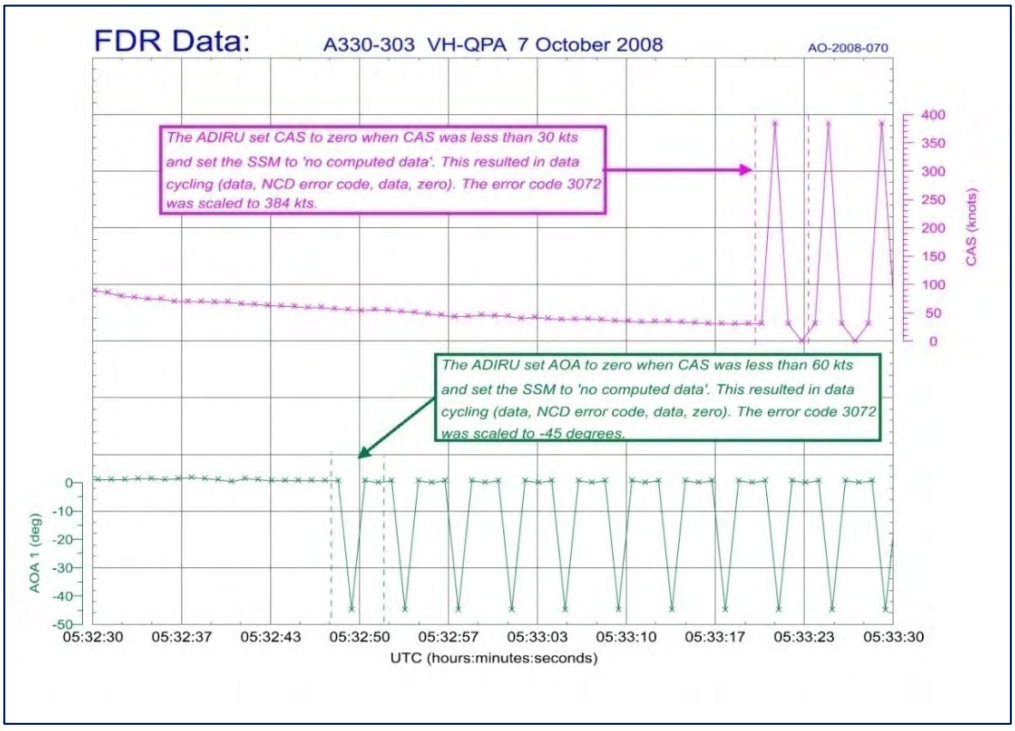

APPENDIX C: POST-FLIGHT REPORT

The following pages show the post-flight report (PFR) in its original format, together with some interpretive comments.

A/C IDENT .VH-QPA DATE OCT07 FLT NBR QFA72 FROM/TO WSSS/YPLM START/END 0125/0535	MAINTENANCE POST FLIGHT REPORT LEG 00	CMC1 PRINTING PAGE 01/05 DATE OCT07 UTC 0953

32 COCKPIT EFFECTS	UTC FLIGHT PHASE	15 FAULTS
ATA 2400 Not Displayed ELEC C/B TRIPPED	0125 Engine Start 02	
ATA 2821 Not Displayed FUEL LEFT PUMP 1 LO PR	0125 Engine Start 02	
	0125 Engine Start 02	ATA 383159 Source *VSC Class 2 Hard GWDU 4 MD DOOR 2 RH (FLUSH SW/COVER SW/IND)
ATA 3621 AIR ENG 1 BLEED FAULT	0133 Cruise 06	ATA 361152 Source BMC1 Class 1 Hard PRESS REG-V(E1-4001HA)
ATA 3831 MAINTENANCE STATUS TOILET	0133 Cruise 06	
	0152 Cruise 06	ATA 341234 Source WXR2 Class 1 Hard ADIRU2+3(1FP2+3) ADRBUS/ WXR2(1SQ2)
ATA 3448 NAV GPWS FAULT	0440 Cruise 06	ATA 341234 Source AFS Class 1 Identifiers Hard *EFCS1 *EFCS2 ADIRU1(1FP1) DMC1 GPWC FWS
ATA 3410 NAV IR 1 FAULT	0440 Cruise 06	

CONTINUED

The first message that was related to the data-spike event. The ground proximity warning system (GPWS) required certain air data inertial reference unit (ADIRU) parameters and was only connected to ADIRU 1. A problem with ADIRU 1 resulted in a loss of the GPWS function.

Analysis determined that this message was probably generated by ADIRU 1 itself, so it was a self-detected failure.

A/C IDENT . VH-QPA DATE OCT07 FLT NBR QFA72 FROM/TO WSSS/YPLM START/END 0125/0535	MAINTENANCE POST FLIGHT REPORT LEG 00	CMC1 PRINTING PAGE 02/05 DATE OCT07 UTC 0953

32 COCKPIT EFFECTS	UTC FLIGHT PHASE	15 FAULTS
ATA 2283 FLAG ON CAPT ND MAP NOT AVAIL	0440 Cruise 06	A red warning flag was displayed on the Captain's navigation display (ND) due to a loss of heading information from inertial reference (IR) 1.
ATA 3414 NAV FM/GPS POS DISAGREE	0440 Cruise 06	There was a cross-check error for latitude and longitude between global positioning system (GPS) data and flight management, guidance and envelope system (FMGES) data (based on IR 1 information).
ATA 3031 A ICE L CAPT STAT HEAT	0440 Cruise 06	This message was due to corrupted data output from air data reference (ADR) 1 (see section 1.12.9).
ATA 3458 NAV GPS 1 FAULT	0440 Cruise 06	No faults were recorded in the built-in test equipment (BITE) data for the multi-mode receivers that included the GPS units. These were spurious messages that were due to a problem with ADIRU 1.
ATA 3458 NAV GPS 2 FAULT	0440 Cruise 06	
ATA 3410 MAINTENANCE STATUS IR 1	0440 Cruise 06	This message was genuinely triggered by ADIRU 1 after it self-detected a failure.
ATA 3410 MAINTENANCE STATUS ADR 1	0440 Cruise 06	This message was probably due to the self-detection of an ADR 1 output wraparound failure by ADIRU 1 (see section 3.7.3).
ATA 2790 MAINTENANCE STATUS EFCS 1	0440 Cruise 06	Flight control primary computer (FCPC) BITE showed that electrical flight control system (EFCS) 1 (that is FCPC 1) had detected a problem with data from ADIRU 1.

A/C IDENT VH-QPA DATE OCT07 FLT NBR QFA72 FROM/TO WSSS/YPLM START/END 0125/0535	MAINTENANCE POST FLIGHT REPORT LEG 00	CMC1 PRINTING PAGE 03/05 DATE OCT07 UTC 0954

32 COCKPIT EFFECTS	UTC FLIGHT PHASE	15 FAULTS
ATA 2790 MAINTENANCE STATUS EFCS 2	0440 Cruise 06	FCPC BITE showed that EFCS 2 (FCPC 2) had detected a problem with data from ADIRU 1.
ATA 3448 NAV GPWS TERR DET FAULT	0441 Cruise 06	There was a loss of GPWS functions due to a loss of data from ADIRU 1 (see NAV GPWS FAULT above).
ATA 2210 AUTO FLT AP OFF	0442 Cruise 06	ATA 279334 Source EFCS1 Class 1 Identifiers Hard EFCS2 BSCU-C2 FCPC1(2CE1) / R G ELEV BSCU-C1 AFS SERVOCTL SV (2CS2) MON The autopilot disconnected and it was not commanded by the crew. This was due to a discrepancy between data sourced from the different ADIRUs.
ATA 3410 NAV IR NOT ALIGNED	0442 Cruise 06	
ATA 2790 F/CTL PRIM 1 PITCH FAULT	0442 Cruise 06	The independent command and monitor channels of FCPC 1 (PRIM 1) detected a discrepancy in the elevator position.
ATA 2790 F/CTL PRIM 3 FAULT	0442 Cruise 06	The independent command and monitor channels of FCPC 3 (PRIM 3) detected a discrepancy in the control orders.
ATA 3240 BRAKES AUTO BRK FAULT	0442 Cruise 06	The independent command and monitor channels of FCPC 2 (PRIM 2) detected a discrepancy in the elevator position.
ATA 2790 F/CTL PRIM 2 PITCH FAULT	0445 Cruise 06	ATA 279334 Source EFCS1 Class 1 Identifiers Hard EFCS2 AFS FCPC2(2CE2) / R Y ELEV SERVOCTL SV (3CS2) MON

CONTINUED

A/C IDENT VH-QPA DATE OCT07 FLT NBR QFA72 FROM/TO WSSS/YPLM START/END 0125/0535	MAINTENANCE POST FLIGHT REPORT LEG 00	DMC1 PRINTING PAGE 04/05 DATE OCT07 UTC 0955

32 COCKPIT EFFECTS	UTC FLIGHT PHASE	15 FAULTS
ATA 3031 A ICE R CAPT STAT HEAT	0445 Cruise 06	This message was due to corrupted data output from ADR 1 (see section 1.12.9).
ATA 2791 F/CTL ALTN LAW	0445 Cruise 06	Due to multiple FCPC faults, the EFCS control law reverted from normal law to alternate law.
ATA 3031 A ICE CAPT PROBES HEAT	0448 Cruise 06	This message was due to corrupted data output from ADR 1 (see section 1.12.9).
ATA 3160 EIS DISPLAY DISCREPANCY	0451 Cruise 06	ATA 316234 Source DMC2 Class 1 Identifiers Hard DMC3, DMC1 DMC1(1WT1)/DU PFD CAPT (1WK1)/DU ND FO(3WK2)
	0452 Cruise 06	ATA 279334 Source *EFCS1 Class 2 Hard FCPC2 (2CE2)/WRG:ADIRU1 BUS ADR1-2 TO FCPC2
	0455 Cruise 06	ATA 341234 Source *EFCS2 Class 2 Identifiers Hard *EFCS1 ADIRU1 (1FP1) ADR1 / FCPC1 (2CE1)
	0455 Cruise 06	ATA 279334 Source *EFCS2 Class 2 Hard FCPC2 (2CE2)/WRG:ADIRU1 BUS ADR1-2 TO FCPC2
ATA 3031 A ICE CAPT PITOT HEAT	0456 Cruise 06	This message was due to corrupted data output from ADR 1 (see section 1.12.9).

CONTINUED

A/C IDENT VH-QPA DATE OCT07 FLT NBR QFA72 FROM/TO WSSS/YPLM START/END 0125/0535	MAINTENANCE POST FLIGHT REPORT LEG 00	CMC1 PRINTING PAGE 05/05 DATE OCT07 UTC 0955

32 COCKPIT EFFECTS	UTC FLIGHT PHASE	15 FAULTS
	0500 Cruise 06	ATA 341234 Source EFCS2 Class 1 Hard ADIRU1 (1FP1)
	0506 Cruise 06	ATA 341234 Source IR3 Class 1 Hard ADIRU1 (1FP1)/ ADIRU3 (1FP3)
ATA 3031 A ICE CAPT AOA HEAT	0508 Cruise 06	This message was due to corrupted data output from ADR 1 (see section 1.12.9).
ATA 3410 NAV ADR 1 FAULT	0513 Cruise 06	ADIRU 3 (IR part) detected a problem with some ADR parameters from ADIRU 1 (see section 1.12.6).
ATA 2131 CAB PR LO DIFF PR	0528 Cruise 06	A low differential pressure between the cabin and outside the aircraft was detected by the cabin pressure controllers. The automatic cabin pressure control was lost as ADR data (altitude and vertical speed) was required from ADIRU 1.
ATA 3231 MAINTENANCE STATUS LGCIU 1	0531 Approach 07	ATA 341234 Source *LGCIU1 Class 2 Identifiers Intermittent *LGCIU2 ADIRU 1 (1FP1)/ LGCIU 1 (05GA1)
ATA 3231 MAINTENANCE STATUS LGCIU 2	0531 Approach 07	Analysis showed that this message was due to a loss of airspeed data from ADIRU 1.
	0531 Approach 07	ATA 232833 Source SATCOM Class 1 Hard RFU1(6RV1) TX COAX

Supplementary information on probe heat faults

Figure C1 shows the specific bits used in the ADR discrete word #1 parameter to indicate specific probe heat faults. To trigger the message on the PFR indicating all probes had heating faults (A.ICE CAPT PROBES HEAT) required bits 12, 14, 15, 16 and 17 all to be '0' (that is, all of the individual probe heat fault discretes needed to indicate a 'Fault').

Figure C1: Bit allocation for ADR discrete word #1 (label 270)

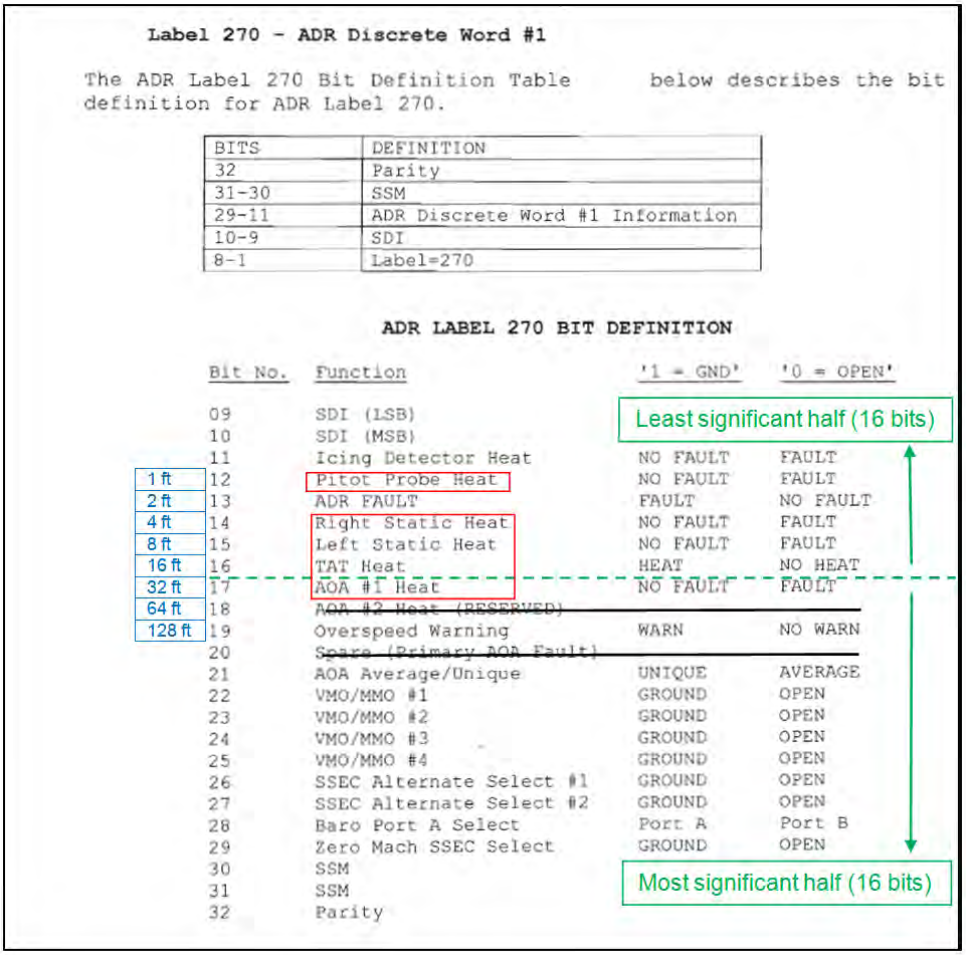

APPENDIX D: OTHER DATA-SPIKE OCCURRENCES

Occurrence on 12 September 2006

On 12 September 2006, an A330-303 aircraft, registered VH-QPA (QPA) and being operated as Qantas flight 68, was on a scheduled passenger transport service from Hong Kong to Perth, Western Australia. At 2052 Universal Time Coordinated (UTC), or 0452 local time, while the aircraft was in cruise at FL410, there was a failure of air data inertial reference unit (ADIRU) 1. The ADIRU was the same unit (serial number 4167) as on the 7 October 2008 flight.

The operator's maintenance watch log recorded that the flight crew contacted maintenance watch and reported that they had a NAV IR 1 FAULT and 'continuous ECAM messages' that could not be stopped, and that the problem was eventually resolved when the crew selected the air data reference (ADR) 1 pushbutton OFF.

The flight crew entered the problem into the aircraft's technical log, noting that there had been a NAV ADR 1 FAULT and that they had received numerous electronic centralized aircraft monitor (ECAM) messages. Maintenance records stated that, in accordance with the manufacturer's maintenance procedures for the relevant post-flight report (PFR) fault messages, an ADIRU re-alignment was conducted and a system test of both the inertial reference (IR) and ADR parts was conducted. No faults were found and the aircraft was returned to service.

Following the 7 October 2008 occurrence, further information was provided by the flight crew regarding the 12 September 2006 occurrence. The crew reported that:

- The event occurred at night and the aircraft was in clear conditions at the time of the event. The first officer was the pilot flying and autopilot 2 was engaged. The autopilot and autothrust remained engaged throughout the event.
- There were numerous ECAM messages, and the messages changed rapidly and could not be properly read or actioned. There were also numerous stall warnings and overspeed warnings.
- Discussions with maintenance watch could not resolve the issue. However, a scan of the overhead panel identified a very weak and intermittent ADR 1 fault light, and the crew decided to turn the ADR 1 off. Following that action, the warning and caution messages ceased and the flight continued without further incident. At no stage was there any effect on the aircraft's flight controls.

The post-flight report (PFR) was very similar to the PFR for the 7 October 2008 occurrence. The cockpit effect messages from both occurrences are summarised in Table D1, together with the cockpit effect messages for the third occurrence on 27 December 2008. Key aspects of the PFR were:

- a NAV IR1 FAULT at 2052, together with a series of other messages starting at the same time
- a NAV ADR 1 FAULT at 2122, probably associated with the crew selecting the ADR 1 pushbutton OFF
- no autopilot disconnection message, consistent with autopilot 2 being engaged rather than autopilot 1, and no flight control primary computer (FCPC, commonly known as PRIM) faults, consistent with there being no in-flight upset.

No flight data was available for the 12 September 2006 flight. The location of the aircraft at the time of the NAV IR 1 FAULT was estimated using positions reported by aircraft communications, addressing and reporting system (ACARS) messages transmitted before and after the fault occurred. That technique gave a position 980 km (530 NM) north of Learmonth, Western Australia (Figure 26).

Table D1: Cockpit effect messages for the three data-spike occurrences[203]

Cockpit effect message	12 September 2006	7 October 2008	27 December 2008
AUTO FLT AP OFF	-	+02	+00
NAV IR 1 FAULT	+00	+00	+00
NAV ADR 1 FAULT	+30	+33	+01
NAV GPWS FAULT	+00	+00	+00
NAV GPWS TERR DET FAULT	-	+01	-
NAV GPS 1 FAULT	+00	+00	+01
NAV GPS 2 FAULT	+03	+00	+01
NAV FM/GPS POS DISAGREE	-	+00	+08
NAV IR NOT ALIGNED	+05	+02	+01
FLAG ON CAPT ND MAP NOT AVAIL	+00	+00	+00
EIS DISPLAY DISCREPANCY	-	+05	-
A. ICE L CAPT STAT HEAT	+03	+00	+01
A. ICE R CAPT STAT HEAT	+03	+05	+01
A.ICE CAPT PITOT HEAT	+03	+16	-
A.ICE CAPT AOA HEAT	-	+28	-
A.ICE CAPT PROBES HEAT	-	+08	-
BRAKES AUTO BRK FAULT	-	+02	-
CAB PR LPO DIFF PR	-	+48	-
MAINTENANCE STATUS IR 1	+00	+00	+00
MAINTENANCE STATUS ADR 1	+00	+00	+00
MAINTENANCE STATUS LGCIU1	+30	+51	+01
MAINTENANCE STATUS LGCIU2	+30	+51	+01
MAINTENANCE STATUS EFCS1	+00	+00	+00
MAINTENANCE STATUS EFCS2	+00	+00	+00
F/CTL PRIM 1 PITCH FAULT	-	+02	-
F/CTL PRIM 3 FAULT	-	+02	-
F/CTL PRIM 2 PITCH FAULT	-	+05	-
F/CTL ALTN LAW	-	+05	-

[203] Numbers indicate the time (in minutes) that the message occurred after the first related message.

Occurrence on 27 December 2008

Overview

On 27 December 2008, an Airbus A330-303 aircraft, registered VH-QPG (QPG) and being operated as Qantas flight 71, was on a scheduled passenger transport service from Perth to Singapore. The aircraft departed Perth at 0750 UTC (1550 local time) and reached its cruise altitude of FL360 at 0814. At 0829 (1729 local time), ADIRU 1 failed.

The flight crew reported that:

- The autopilot (autopilot 1) disconnected and the ECAM started providing a series of caution messages, including a NAV IR 1 FAULT.
- They actioned the relevant operational procedure[204] by selecting the IR 1 pushbutton to OFF and the ADR 1 pushbutton to OFF, and both OFF lights illuminated.
- They continued to receive multiple ECAM messages, and those messages were constantly scrolling on the display. One of the ECAM messages was a NAV IR 1 FAULT, and the recommended crew action was to switch the ATT HDG switch to the CAPT ON 3 position. They completed this action, but it did not stop the ECAM messages.
- They returned to Perth and conducted an uneventful landing. At no stage was there any effect on the aircraft's flight controls.

At the time that the autopilot disconnected, the aircraft was about 480 km (260 NM) north-west of Perth and about 650 km (350 NM) south of Learmonth (Figure 26).

The PFR for the 27 December 2008 flight contained a series of messages associated with ADIRU 1 that were very similar to those for the PFR from the 7 October 2008 occurrence (Table D1). Consistent with there being no in-flight upset, there were no PRIM FAULTS or PRIM PITCH FAULTS.

Flight recorders

The flight data recording system fitted to QPG on 27 December 2008 differed from the system fitted to QPA. QPG had a single box system (Teledyne FDIMU part number 2234340-02-02) with two separate processors to handle both the FDR and ACMS (QAR) functions. However, the source of ADIRU parameters for the FDR and QAR was the same as for QPA (section 1.11).

The overall pattern of ADR and IR data from the recorders was very similar to that for the 7 October 2008 event. A summary of key events from the FDR is presented in Table D2. The FDR data is presented in Figure D1 and D2, and QAR data is presented in Figure D3.

---

[204] This procedure was different to that which applied at the time of the 7 October 2008 occurrence. The relevant procedure at the time of the 27 December 2008 occurrence was based on Airbus Operations Engineering Bulletin (OEB) 74-3 issued in December 2008 in response to the 7 October 2008 accident (section 7.1.1).

Table D2: 27 December 2008 occurrence sequence of events

Time	Time to event	Event
0749:55	-0038:57	Takeoff from Perth
0814:01	-0014:51	Aircraft reached top of climb (36,000 ft or FL360)
0828:52	0000:00	First indication of IR parameters providing erroneous values
0828:55	0000:03	IR 1 Fail indication commenced (sampled every 4 seconds)
0828:56	0000:04	Autopilot 1 disconnected (involuntary)
0829:20	0000:28	ADR 1 Fail indication commenced (indicating crew had selected ADR 1 OFF)
0830:21	0001:29	Autopilot 1 re-engaged
0831:48	0002:56	Crew switched ADR source for captain's displays from ADR1 to ADR 3 (FDR's priority source for ADR parameters switched to ADR 3)
0832:25	0003:33	Crew switched IR source for captain's displays from IR1 to IR 3 (FDR's priority source for IR parameters switched to IR 3)
0925:45	0056:53	Touchdown at Perth (aircraft gross weight was 195.3 tonnes)

Figure D1: FDR data for the 27 December 2008 occurrence (entire flight)

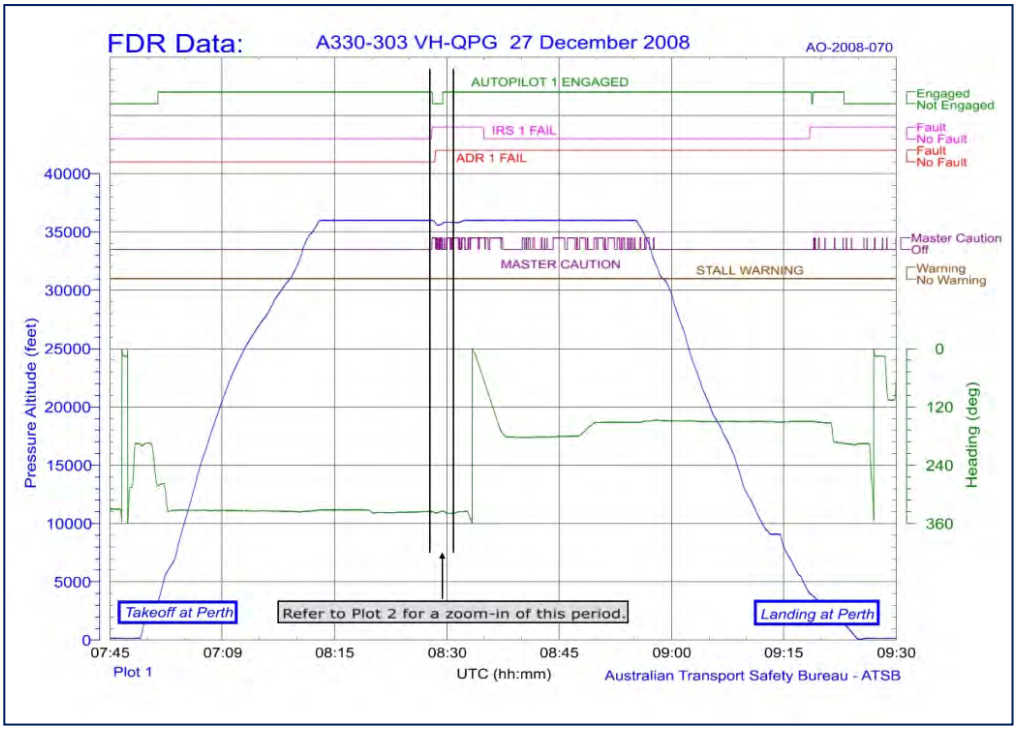

Figure D2: FDR data for the 27 December 2008 occurrence (0626 to 0636 UTC)

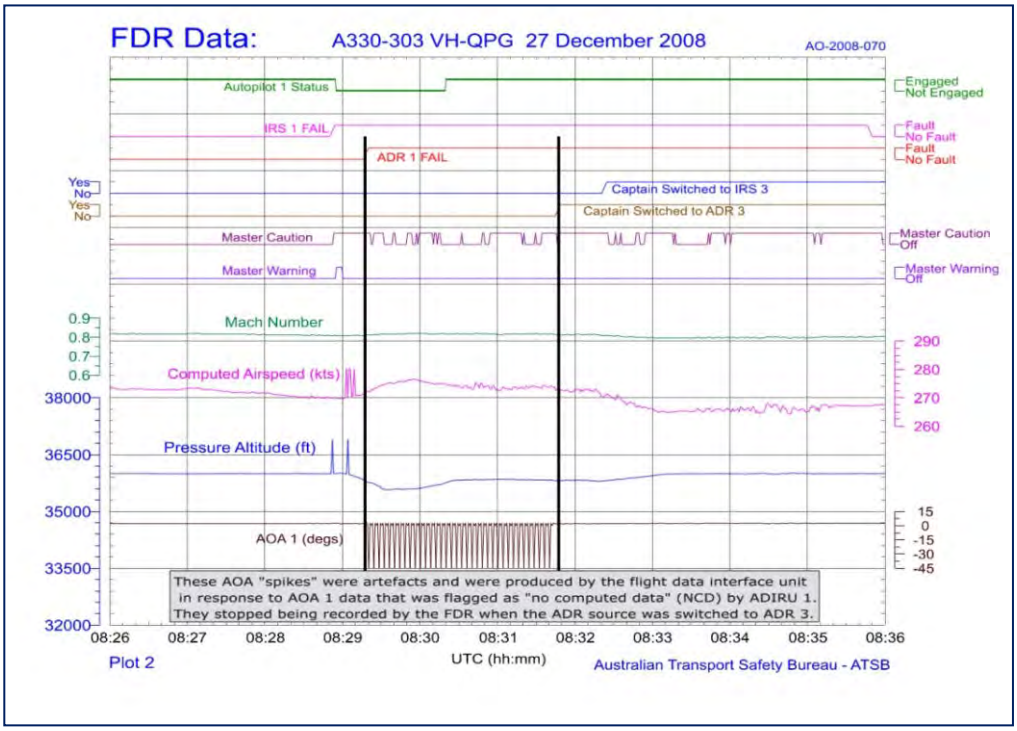

Figure D3: QAR data for the 27 December 2008 occurrence

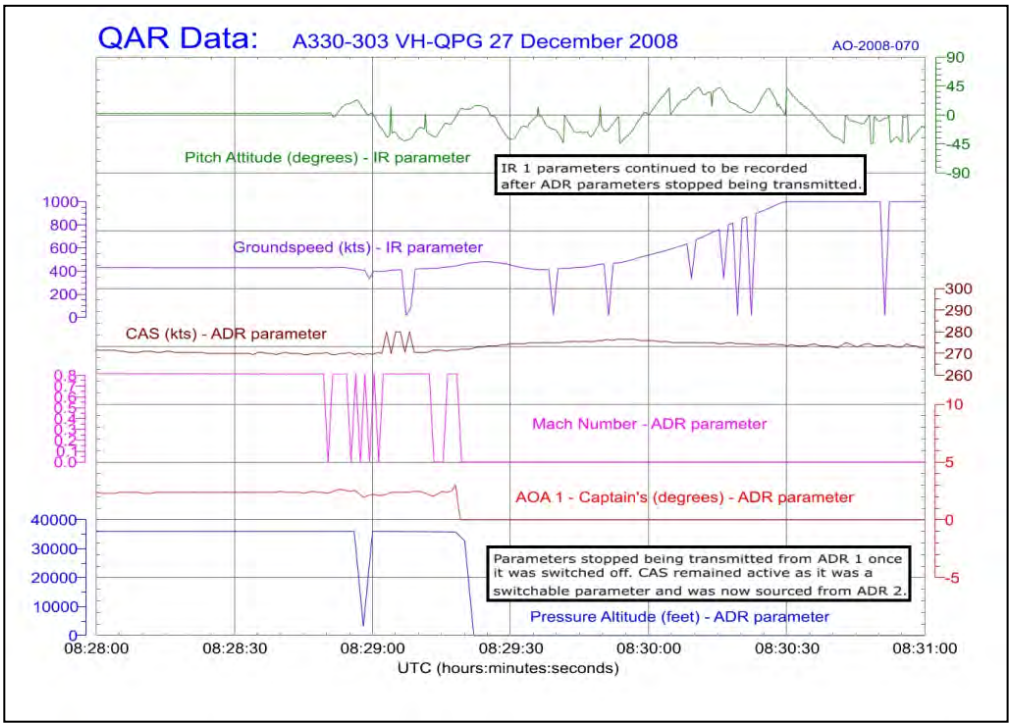

The recorded data showed that the failure mode started at 0828:52. In terms of the ADR data, the crew selected ADR 1 OFF 28 seconds after the failure mode started (at 0829:20). Consequently, ADR 1 ceased transmitting data and the FDR switched its source for the ADR parameters, except angle of attack (AOA) 1 and AOA 2, from ADR 1 to ADR 3.[205] When the crew switched the source of ADR parameters for the captain's displays to ADR 3 (0831:25), the FDR's source of AOA 1 also switched to ADR 3. During the brief period that ADR 1 data was recorded:

- There were no spikes in AOA[206] or Mach.
- There were two spikes for altitude, and two spikes for static air temperature. There were also four small changes (10 kts) in computed airspeed. ADIRU 1 flagged all of these anomalies as valid data.

In terms of the IR data:

- Although the crew reported switching IR 1 off, the QAR data showed that the IR 1 parameters continued to be transmitted by the ADIRU until after landing.
- From the start of the event until after the aircraft landed, the QAR recorded numerous spikes in pitch attitude, drift angle, body lateral acceleration, inertial vertical speed, and magnetic heading.[207]
- In addition to the data spikes, all of the IR parameters showed persistent deviations from their expected values. For most of these parameters, the deviations showed oscillatory behaviour. The groundspeed parameter increased to 1,000 kts, where it remained for the rest of the flight, except for a number of lower-value data spikes.
- With only brief exceptions, ADIRU 1 flagged all of the IR 1 data after the start of the event (0828:52) as invalid.

There were no stall warnings recorded (consistent with no spikes on the captain's AOA) and no overspeed warnings recorded (consistent with only small changes on computed airspeed). However, there were repetitive master cautions recorded after the failure mode commenced, consistent with the occurrence of ECAM messages.

---

[205] At this time the QAR stopped recording any ADR parameters from ADIRU 1. Computed airspeed was defined as a switchable parameter, and its source switched to ADIRU 2 at this time.

[206] Some oscillations were observed in the recorded data for AOA 1 after the ADR Fail parameter commenced indicating Fail. These oscillations were examined and matched the no computed data cycling behaviour expected when ADR 1 was no longer outputting data, consistent with it being switched off (Appendix B).

[207] As with the 7 October 2008 occurrence, the FDR only recorded a small number of spikes at the beginning of the event. For most of the IR parameters, the FDR switched its source if the ADIRU flagged its data as invalid. However, the QAR always sourced its data from ADIRU 1.

ADIRU information

ADIRU 1 was the same model (LTN-101) but a different unit (serial number 4122) to that involved in the 12 September 2006 and 7 October 2008 events.

Following the incident on 27 December 2008, the unit was removed from the aircraft and sent to the manufacturer's facility for downloading of the BITE data and testing. The results of the tests were effectively the same as for unit 4167 from the 7 October 2008 occurrence. That is:

- The BITE data showed no recorded faults, even though several fault messages should have been recorded.
- Several routine messages normally stored in BITE were either not recorded or had anomalies. The pattern was slightly different to that for the 7 October 2008 flight (Table D3).
- Subsequent examination and testing identified no problems with the unit (Appendix E).

Table D3: Comparison of the BITE results for ADIRUs 4167 and 4122

BITE data	7 October 2008 flight ADIRU 4167	27 December 2008 flight ADIRU 4122
Fault records	None recorded during the flight.	None recorded during the flight.
Routine NAV Update record on shutdown	Not recorded.	Not recorded.
Routine elapsed time interval (ETI) timestamps	The ETI observed at turn on at the manufacturer's test facility was about 0.7 hours after takeoff. The ADIRUs were in fact on for 14.8 hours.	After landing, the system was given a full alignment; however, the ETI reverted to the value for the previous alignment in Singapore (the previous flight).
Routine temperature records (every hour)	None recorded after the start of the event.	None recorded but the incident flight was only about one hour long

Search for other data-spike occurrences

Following the 27 December 2008 occurrence, Airbus conducted a review of PFRs using its AIRcraft Maintenance ANalysis (AIRMAN) database. AIRMAN is a ground-based software tool that assists operators of Airbus aircraft to optimise line maintenance and troubleshooting of aircraft. Fault data is downloaded in real-time, and PFRs are stored and available for subsequent analysis.

The use of AIRMAN was an operator-based decision, with most Airbus operators electing to use the tool. Airbus advised that at the end of 2008, there were about 900 A330/A340 aircraft in operation, and 397 had Northrop Grumman LTN101 ADIRUs. AIRMAN data was available for 248 of those aircraft in the 2005 to 2008 period. The sample of 248 aircraft included 48 operators, and included several airlines that operated flights to and from Australia.

Airbus searched the AIRMAN database for PFRs that contained a similar pattern of fault messages as those recorded on the 12 September 2006, 7 October 2008 and 27 December 2008 flights. The search looked for all PFRs with 'NAV GPS *

FAULT' and 'A.ICE *' fault messages, and then reviewed all relevant PFRs in detail.[208] The samples searched were as follows:

- 1 January 2002 to 31 December 2008, long range aircraft (A330, A340) with LTN101 ADIRUs: four events identified (three known events and one other potentially related event)
- 1 January 2005 to 31 December 2008, long range aircraft (A330, A340) with another ADIRU model: no related events
- 1 January 2005 to 31 December 2008, single aisle aircraft (A310, A318, A319, A320, A321) with LTN101 ADIRUs: no related events.

This search identified only one other potentially relevant event. That event involved an A330-200 aircraft (registered VH-EBC) on 7 February 2008. The aircraft operator (Jetstar) was associated with the operator of QPA, and the aircraft's maintenance was conducted by the operator QPA. Although the PFR contained some similarities with the three known events, a detailed examination of the available evidence concluded that this occurrence was unrelated to the three known data-spike occurrences.[209]

The aircraft manufacturer also conducted additional searches of the AIRMAN database for all A330s that were operated by the operators of QPA and EBA for the period 1 January 2002 to 31 December 2008. Those searches involved a variety of different combinations of faults. No other potentially related events were identified.

---

[208] These messages were considered to be the most effective search as A.ICE messages were not normally associated with an IR or GPS problem. The common link between these different systems was an ADIRU and a problem with the ADIRU could be incorrectly indicated as both A.ICE and GPS faults. Searching only for IR faults was unmanageable as there were too many positive hits.

[209] For the VH-EBC event, a key difference was that the ADIRU BITE recorded an actual fault, and subsequent examination of the ADIRU by the ADIRU manufacturer identified a hardware fault. In addition, the crew reported that they did not receive multiple or repeated ECAM warnings or cautions, and that the PFR did not contain distinctive messages associated with the failure mode (such as ADIRU maintenance status or NAV GPWS messages. No FDR or QAR data was available.

APPENDIX E: ADIRU TESTING

Test plan development

Following the 7 October 2008 occurrence, air data inertial reference unit (ADIRU) 1 (unit 4167) was removed from VH-QPA at Learmonth, prior to downloading any data or functional testing of other units on the aircraft. ADIRU 2 (unit 4687) and ADIRU 3 (unit 4663) were removed after the aircraft was ferried back to Sydney, New South Wales. All three units were dispatched to the ADIRU manufacturer's facilities in Los Angeles, United States (US), and quarantined until the relevant investigation agencies had developed an agreed test plan.

The test plan was developed collaboratively between the Australian Transport Safety Bureau (ATSB), the US National Transportation Safety Board (NTSB), the French Bureau d'Enquêtes et d'Analyses pour la sécurité de l'aviation civile (BEA), the ADIRU manufacturer, the aircraft manufacturer, and the operator. The initial testing was conducted with all organisations present, and subsequent testing was conducted wherever possible under the supervision of the ATSB and/or the NTSB.

When developing the test plan and conducting the testing, the following principles were used:

- All parties reached agreement on the necessary testing and the order of testing prior to test commencement.
- Analysis of the available evidence was regularly performed and reviewed prior to proceeding with further tests.
- Tests that could potentially result in the disturbance or loss of evidence (such as damage to or reconfiguration of a unit) were conducted after other tests wherever possible. For example, the removal of the ADIRU covers and re-seating of the circuit modules were considered disturbances of evidence as they could affect the unit's response to electromagnetic interference (EMI), so tests requiring such disassembly were performed after EMI testing.
- The ADIRUs were tested as a complete unit prior to any module testing.
- Where the validation of a test was required, the tests were performed on an exemplar ADIRU (serial number 4461) with the same configuration as the subject ADIRU (4167).
- All tests were conducted within the applicable quality assurance framework (normally that of the organisation performing the test).

The final series of examinations and tests that were conducted on ADIRU 1 (serial number 4167) and the exemplar unit (serial number 4461) are summarised in Table E1 and discussed in the remainder of this appendix. ADIRU 2 (serial number 4687) and ADIRU 3 (serial number 4663) were subjected to a basic series of tests normally conducted on ADIRUs that were returned from service. As no problems were identified during these tests, and there was no other evidence of a problem with these units, (section 1.12.2 and 1.12.6) no further testing was conducted.

Table E1: Summary of testing performed on the ADIRUs

Test	ADIRU 4461 (exemplar unit)	ADIRU 4167	ADIRU 4122
Visual inspections			
External visual inspection	Passed	Passed	Passed
Internal visual inspection	Passed	No relevant failures	Passed
Detailed microscopic and X-ray internal inspection	Not conducted	Passed	Not conducted
Functional tests			
Unpowered ground integrity tests	Passed	Passed	Not conducted
Operational flight program (software) version and integrity verification	Passed	Passed	Not conducted
Basic power-on test with BITE download	Passed	Passed	Passed
Maintenance test procedure	Passed	Passed	Passed
Acceptance test procedure	Not conducted	Passed	Passed
Databus output waveform measurement and data monitoring	Passed	No relevant failures	Not conducted
Module-level testing	Not conducted	No relevant failures	No relevant failures
Environmental tests			
Continuous maintenance test procedure with heat shroud	Not conducted	No relevant failures	Not conducted
DO-160C environmental tests	No relevant failures	No relevant failures	Not conducted
Extended EMI tests	No relevant failures	Not conducted	Not conducted
Environmental stress screening	Passed	No relevant failures	Not conducted
Highly accelerated stress screening	Passed	Passed	Not conducted

Following the 27 December 2008 occurrence, ADIRU 1 (serial number 4122) was removed and sent to the ADIRU manufacturer's facilities for testing. That unit was subjected to many of the same examinations as unit 4167, although some tests were not considered necessary given the results already obtained through testing unit 4167. ADIRU 2 (serial number 5275) and ADIRU 3 (serial number 4354) were not removed or tested as there was no evidence that there were any problems with either of these units (Appendix D).

Visual inspections

External visual inspection

An initial inspection was conducted to establish the shipping container condition and unit identification using part and serial numbers. The units were then inspected externally for signs of damage, corrosion, contamination, correct fit and sealing, and quality control seal integrity. The connector pins were inspected for contamination and damage. No problems were identified.

Internal visual inspection

The units were inspected internally following the removal of the top and bottom covers to determine whether the modules and motherboard were correctly seated and the mounting fasteners were correctly torqued (tightened), and to check for signs of internal damage due to overheating, foreign objects, or other causes. No problems were identified.

Detailed microscopic and X-ray internal inspection

Each module and motherboard was subjected to a detailed, visual inspection using a microscope, as well as internal inspection using X-ray imagery. No problems were identified.

Functional tests

Unpowered ground integrity tests

The units' connector pins for angle of attack (AOA) AOA input and ARINC 429 databus output were measured for direct current (DC) resistance to power supplies and ground. No problems were identified.

Operational flight program (software) version and integrity verification

The units were installed on a dedicated test station that powered up the system and performed a cyclic redundancy check (CRC) on the unit's operational flight program (software) to verify its integrity. In addition, ARINC 429 databus parameters were analysed for anomalies. No problems were identified.

Basic power-on test with built-in test equipment download

The units were installed on a dedicated test station that powered up the system, performed some basic integrity checks, and performed a non-destructive download of the system built-in test equipment (BITE) data. A second download of the system BITE data was conducted following the maintenance test procedure described below. Numerous other BITE data downloads were performed following many of the other tests. The BITE data was then decoded for study and evaluation.

All of the BITE downloads were completed successfully. As discussed in section 1.12.6, the BITE for ADIRU 4167 contained no fault messages from the occurrence flight and there were anomalies in the recorded data for routine

messages. After other tests had been conducted, further BITE data was successfully written by the ADIRU and then downloaded with no anomalies or recorded faults.

ADIRU 4687 recorded three fault messages and ADIRU 4663 contained five fault messages, and all of these messages related to the performance of ADIRU 4167 (section 1.12.6). ADIRU 4122 (from the 27 December 2008 occurrence) exhibited the same type of anomalies in the BITE data as ADIRU 4167 from the 7 October 2008 occurrence (Appendix D).

Maintenance test procedure

The maintenance test procedure was a standard procedure that was performed on all production and service ADIRUs. It tested hardware functionality, including navigation performance and the download of BITE data. It consisted of a total of 177 specific subtests. Results from these tests were compared with the acceptance test procedure results from the most recent service on the units.

As the procedure loaded special test software to replace the normal flight software, it was performed after the operational flight program integrity check. The flight software was then re-installed at the end of the procedure.

A system functional test was also conducted as part of the standard maintenance test procedure with a navigation performance ('drift') test. The navigation performance test was run on a moving platform to exercise the unit's inertial sensors, firstly for 1 hour from a cold turn-on and another hour from a warm turn-on. An additional BITE data download was performed after the navigation performance test.

No problems were identified with any of the units.

Acceptance test procedure

The acceptance test procedure was normally conducted on production and service ADIRUs after a unit had been disassembled to ensure its proper reassembly. It tested hardware functionality including navigation performance and download of BITE data. It was similar to the maintenance test procedure, although less comprehensive as it did not overwrite the unit's operational flight program (software). Since the acceptance test procedure could be conducted without changing a unit's configuration, the unit could subsequently be tested while still in its 'original' condition. No problems were identified.

Databus output waveform measurement and data monitoring

This test focused on the ARINC output transmission circuitry. The goal was to establish if data anomalies could be observed that were similar to spike anomalies recorded during the occurrence. The procedure was performed on the engineering manual integration test station (MITS). This was the same station used for all software integration and formal software testing. The station capabilities allowed real-time monitoring of any BITE failures that may occur. It also separated out all ARINC databuses so that they could be loaded or connected to any receiving equipment desired.

The ARINC transmission levels were measured using an oscilloscope, inspected for indications of noise or abnormal waveforms, and compared to the ARINC specification. This was repeated on all inertial reference (IR) and air data reference

(ADR) output databuses. Loading on each output bus (that is, the number of devices on the bus) was varied from none to the maximum specified by ARINC 429 (20 devices).

For unit 4167, two ADR output databuses (ADR 7 and ADR 8) exceeded the no-load negative voltage limit by a very small amount. The failures were considered to be a result of measurement equipment tolerance and not related to the occurrence.[210] Some output data values exceeded the test's data variation tolerances, which were used to detect excessive variations in the ADIRU's output data. These were attributed to variations arising from the ADIRU being switched ON and OFF during the test, and to scaling errors in the test software. Accordingly, the variations were considered a result of minor problems with the test procedure design.

These issues were not considered relevant to the occurrences under investigation. No other problems were identified.

Module-level testing

The units underwent a detailed bench test that involved dismantling the unit and conducting functional tests of each module to identify any malfunctioning circuitry that was not evident at the system level. These activities included:

- the removal and inspection of modules for damage, heat damage, foreign material or residue, connector contamination, loose components, and solder joint integrity;
- motherboard and motherboard connector inspection;
- a rear panel EMI protection diode conduction check; and
- individual module functional tests.

For unit 4167, the following aspects were noted:

- Some minor failures occurred as a result of problems with the design of the test procedure.
- An analogue-to-digital converter within the (inertial) sensor electronics module was found to be out of tolerance. The ADIRU would normally detect and correct for the error when the module was installed.
- A failure occurred within the power supply of the inertial sensor electronics, but did not reoccur when the test was repeated several times.

For unit 4122, the following aspects were noted:

- Some minor failures occurred as a result of problems with the design of the test procedure.
- Some minor physical anomalies not related to the data-spike failure mode (for example, insulation damage on a wire used only by test equipment) were observed. A waxy residue was found on the connector sockets on the sensor electronics module. The residue had been deposited during manufacture as the

[210] ADR 7 was always connected to a load in the A330-300 installation, so it would be very unlikely to exceed tolerances when in actual operation. ADR 8 was not connected to any other systems in the A330-300 installation and any fault with that databus would have no effect on the aircraft.

result of a problem that was subsequently rectified and was not considered a result of, or cause of, any electrical continuity or other problems.

- The coating on the monitor module circuit board was cracked in one location. The board was tested while undergoing deformation (flexing) and with moisture applied without any functional effects.
- An input discrete (single wire interface) to the monitor module failed a voltage range test. The discrete was used to identify the aircraft type on which the ADIRU would be installed, to disable the ADR part for some aircraft types. The test failed by a marginal amount (0.018 V) due to a redesign of the associated circuitry.
- A power supply diode was found to be cracked. The associated circuitry was not affected significantly. In flight, failure of this diode would result in a number of ADIRU power supply BITE tests being triggered and the ADIRU shutting down.
- The power supply BITE circuitry exhibited a fault, which was eventually traced to a faulty resistor that had passed a check at the beginning of the test. In flight, failure of the resistor would result in the ADIRU shutting down. The problem was considered to have arisen during the testing.

None of the above issues for either unit 4167 or 4122 were considered relevant to the occurrences under investigation. No other problems were identified.

Environmental tests

Continuous maintenance test procedure with heat shroud

The maintenance test procedure described above was run continuously in a loop mode for 11 full cycles with the unit covered in a shroud to provide heat stress (Figure E2). The temperature inside the unit reached 56°C with the shroud on.

Figure E2: ADIRU undergoing test under heat shroud

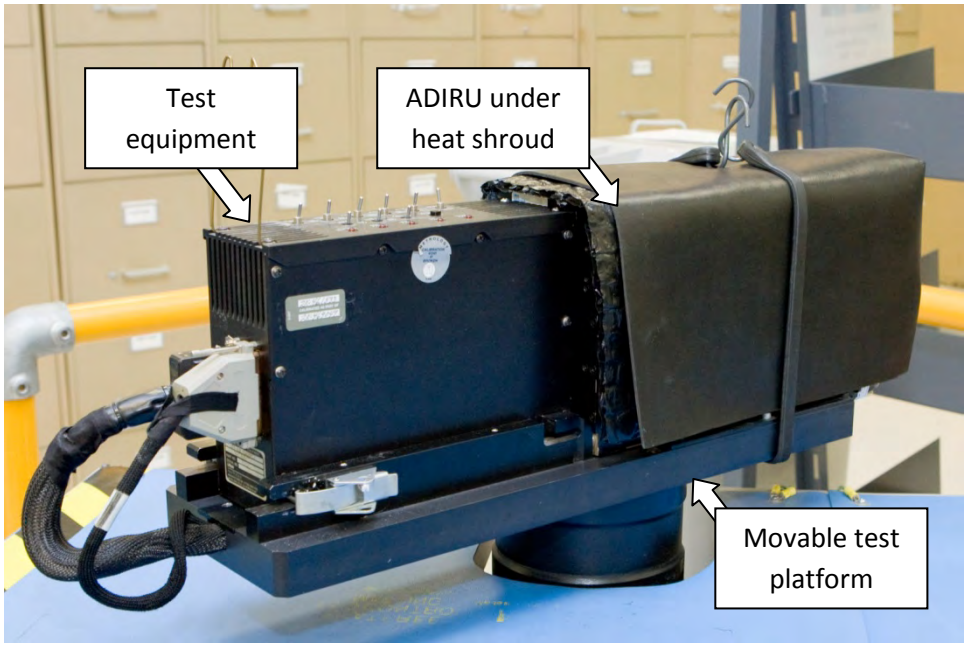

Unit 4167 passed 2,223 of 2,272 tests, with the remainder failing due to navigation drift that was considered normal in a high-temperature environment. This issue was not considered relevant to the occurrences under investigation. No other problems were identified.

DO-160C environmental tests

The units were subjected to certain environmental tests in accordance with DO-160C. These tests were originally used for qualification of the ADIRU design. Only those parts of DO-160C considered relevant to the investigation were undertaken; for example, the units' resilience to liquids was not considered an issue because there was no visible sign of liquid contamination.

The DO-160C tests assessed ADIRU behaviour under each of the following environmental conditions[211]:

- Random vibration along each of the three major orthogonal axes of the ADIRU (DO-160C section 8).
- Power input variation under normal and abnormal conditions, consisting of variations in alternating current (AC) power (voltage, frequency, modulation of voltage and frequency, power interrupt, surge, and under-voltage), and DC power (voltage, voltage ripple, power interrupt, surge, and under-voltage) (DO-160C section 16).
- Power voltage spike (DO-160C section 17).
- Power input audio frequency conducted susceptibility, consisting of 750 Hz to 22 kHz injected current for AC power and 10 Hz to 180 kHz injected current for DC power (DO-160C section 18).
- Induced signal susceptibility, consisting of magnetic fields at 400 Hz induced into the equipment, magnetic fields between 400 Hz and 15 kHz induced into interconnecting cables, electric fields between 380 Hz and 420 Hz induced into interconnecting cables, and voltage spikes[212] induced into interconnecting cables (DO-160C section 19).
- Radio frequency susceptibility, consisting of conducted susceptibility 10 khz to 400 Mhz at 150 mA (nominally representative of a 100 V/m electric field) and radiated susceptibility 30 Mhz to 18 Ghz at 100 V/m (DO-160C section 20).
- Emission[213] of electromagnetic energy between 150 kHz and 1.215 GHz (DO-160C section 21).

---

[211] DO-160C sections 16, 17, 18, 19, and 20 theoretically test for various forms of EMI, although in practice sections 16 and 17 typically simulate variability of the power supply rather than variation as a result of EMI. DO-160C section 21 tests for electromagnetic emissions and can sometimes detect faulty hardware.

[212] Voltage spikes are not the same as data spikes. Voltage spikes (transients) are commonly introduced to a power supply as a result of power switching. The voltage spikes for this test had a duration of 0.05 to 1.00 milliseconds, repetition rate of 0.2 to 10.0 microseconds, burst rate of 8 to 10 Hz, and amplitude of 600 volts peak-to-peak.

[213] This test is normally used to determine a unit's electromagnetic emissions. For the purposes of this investigation, the test was used to detect any hardware fault that would produce abnormal electromagnetic emissions.

During the tests, artificial inputs were provided to the ADIRU. A set of 22 ADR and IR data output values were monitored for variation from the known static value. Figures E3 and E4 show the testing apparatus for vibration and EMI testing respectively.

Figure E3: ADIRU 4167 on the vibration test platform

Figure E4: ADIRU undergoing EMI susceptibility testing

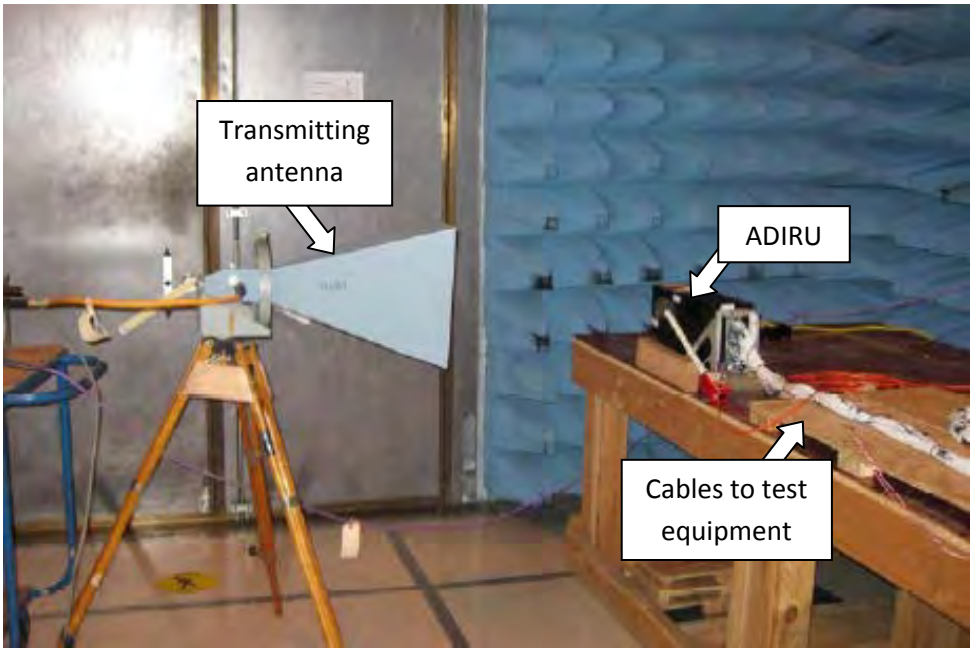

In terms of EMI, the range of frequencies tested incorporated the frequency ranges associated with all identified EMI-inducing sources, including other aircraft systems such as power supply and radio transmitters, the in-flight entertainment system, personal electronic devices (PEDs), and an Australian Defence Force communication station located near Learmonth (Harold E. Holt Naval Communication Station). The EMI testing incorporated additional slow sweeps and

dwells (sustained testing at a fixed frequency) at frequencies of interest, including: the transmission frequency range of the naval communication station, several internal clock and data frequencies used by the ADIRU, various aircraft radio transmission and reception frequencies, and 19 in-flight entertainment (IFE) system operating frequencies.

The naval communication station transmitted at a frequency of 19.8 kHz. It was impracticable to test this frequency for radiated susceptibility. However, if emissions at this frequency were to affect any aircraft system, they would most likely couple onto the aircraft wiring rather than directly to any line-replaceable unit (LRU) due to the very long wavelength of low-frequency transmissions. Accordingly, conducted susceptibility testing was the most appropriate way of assessing the potential impact of such emissions on an LRU such as an ADIRU.

For ADIRU 4167 and the exemplar unit (4461):

- During radiated susceptibility testing in the 40 to 60 MHz range and at 100 V/m, both units exhibited an ARINC 429 databus transmission fault. Inspections revealed that databus shield terminations had broken loose on both units. The testing sequence was completed with no further failures. The equipment was subsequently repaired and testing at the same frequency range and signal strength was repeated. The failure did not repeat on either unit.
- During vibration testing, some IR output data values exceeded the test's data variation tolerances, which were used to detect excessive variations in the ADIRU's output data. These exceedances were attributed to normal inertial drift, which was exacerbated by the vibration imposed on the ADIRU, as well as scaling errors in the test software. Accordingly, the variations were considered a result of minor problems with the test procedure design.

These issues were not considered relevant to the occurrences under investigation. No other problems were identified.

The full range of EMI testing required significant resources. ADIRU 4122 was not subjected to EMI testing because all involved parties agreed that such testing would not provide additional useful information in view of the results of the EMI tests that were performed on ADIRUs 4167 and 4461.

Extended EMI tests

ADIRU 4461 was also subjected to additional environmental tests related to DO-160C and designed to establish typical ADIRU behaviour at high radiated electric field strengths and conducted currents, particularly at certain frequencies of interest. These tests consisted of:

- conducted susceptibility tests up to 375 mA, or 2.5 times the DO-160C level
- radiated susceptibility tests from 30 MHz to 100 MHz at 200 V/m, or twice the DO-160C level
- radiated susceptibility tests from 100 MHz to 18 GHz at 250 V/m, or 2.5 times the DO-160C level
- measurement of emissions between 150 kHz and 1.215 GHz with the ADIRU enclosure lid removed, to detect any abnormal emissions resulting from a hardware defect.

During the tests, artificial inputs were provided to the ADIRU. A set of 22 air and inertial data output values were monitored for variation from the known static value.

Results included the following:

- The ADIRU's BITE detected failures of the input ARINC 429 databuses that would normally provide data to the ADIRU from the aircraft's flight control unit (a control panel on the flight deck).
- The ADIRU's BITE detected one minor gyro fault that had no operational effect.
- The total air temperature input was found to vary by an insignificant amount at very high radiated field strengths.
- Some output data values exceeded the test's data variation tolerances, which were used to detect excessive variations in the ADIRU's output data. These were attributed to data variations arising from the ADIRU being switched ON and OFF during the test, and to changes in the local atmospheric pressure throughout the test. The variations were considered a result of minor problems with the test procedure design.

None of these problems were considered relevant to the investigation.

Environmental stress screen

The environmental stress screen was a test in which an ADIRU was subjected to 15 hours of vibration and temperature cycles with power cycling at temperature extremes to attempt to induce a malfunction. The unit was connected to an ARINC 429 databus analyser to assess output data integrity. For the tests that were conducted as part of the investigation, an extra high temperature test was also conducted to identify whether the units were susceptible to a failure mode known to have occurred on some other ADIRUs.

For ADIRU 4167, a fault in the ADR module was detected by the test equipment, but this was attributed to a fault with the test station and not with the ADIRU itself. No faults were recorded in the BITE data. Five further environmental stress screening tests did not replicate the fault. No other problems were identified.

Highly-accelerated stress screening

The highly-accelerated stress screening test was normally used by the ADIRU manufacturer to simulate long-term temperature and vibration effects on the ADIRU to trigger defects that would otherwise emerge later in the ADIRU's service life. The test subjected the ADIRU to temperature and vibration extremes close to the design limits of the unit while monitoring the system status. The unit's BITE data was downloaded at the conclusion of the test. No problems were identified.

APPENDIX F: AIRCRAFT LEVEL TESTING

In May 2009, an assessment of the aircraft (VH-QPA) was conducted to characterise the electromagnetic environment around air data inertial reference unit (ADIRU) 1 during normal operation and in relative proximity to potential sources of electromagnetic interference (EMI).[214] Measurements included the following:[215]

- 'Conducted' bulk current[216] measurements were taken in the cable bundles connected to ADIRU 1 during ground and flight testing, and in the cable bundles connected to ADIRU 3 during ground testing. The measurements characterised the current in the cables at frequencies between 10 kHz and 500 MHz.
- 'Radiated' electric fields measurements were taken at a location immediately to the rear of ADIRU 1 during ground and flight testing, and immediately to the rear of ADIRU 3 during ground testing. These measurements characterised the electric fields near the ADIRUs at frequencies between 500 MHz and 7 GHz.

Other frequencies could not be tested due to limitations of the test equipment available. Measurements from the sensors were logged by computer.

The measurements were taken in the following stages:

- Measurements taken with the aircraft on the ground and without any aircraft power applied, including with several mobile phones being operated on board the aircraft while the measurements were taken.
- Measurements taken with the aircraft on the ground and with auxiliary power unit (APU) power applied.
- Measurements taken with the aircraft in flight, including with various aircraft systems in operation (such as the passenger in-flight entertainment system and various radio communication systems).

The assessment flight was conducted on 17 May 2009. The aircraft departed from Sydney, New South Wales, flew to Learmonth, Western Australia, and operated in the vicinity of the Harold E. Holt Naval Communication Station near Learmonth, Western Australia, for about 2 hours (Figures F1 and F2). The aircraft was flown directly over the facility in a northerly and easterly direction, along a 90° quadrant at a constant distance to the north-west of the facility, past the facility at a tangent, and for some distance along the 7 October 2008 flightpath before returning to Sydney. The duration of the flight was 10 hours and 48 minutes, during which both radiated field and conducted current measurements were observed continuously and recorded at intervals.

All aircraft systems, including the ADIRUs and electronic flight control system (EFCS), operated normally throughout the testing activities.

After the flight, the flight data recorder (FDR), quick access recorder (QAR) and logged test data were analysed. The analysis did not reveal any anomalous results.

---

[214] At the same time, a range of other tests and checks of the wiring and configuration associated with the aircraft's ADIRUs was also undertaken (section 1.12.5).

[215] Refer to Appendix G for an explanation of technical terms relating to EMI.

[216] Bulk current measurements are taken from an entire cable bundle at once, as distinct from measurements taken from each cable individually.

All measured currents and electric fields were substantially below the DO-160 design requirements for electromagnetic compatibility at all measured frequencies. Conducted current measurements were between about one tenth and one millionth of the DO-160C limits, while radiated measurements were between one third (for the highest measurement, which occurred at a single frequency) and one thousandth of the DO-160C limits.

The naval communications station was later confirmed to be transmitting at the time of the assessment flight. No significant conducted current at the station's transmission frequency of 19.8 kHz was observed at any point throughout the flight (the level of coupling onto the ADIRU wiring was too small to be measurable).

Figure F1: 17 May 2009 test flightpath (overview)

Figure F2: 17 May 2009 test flightpath (detail)

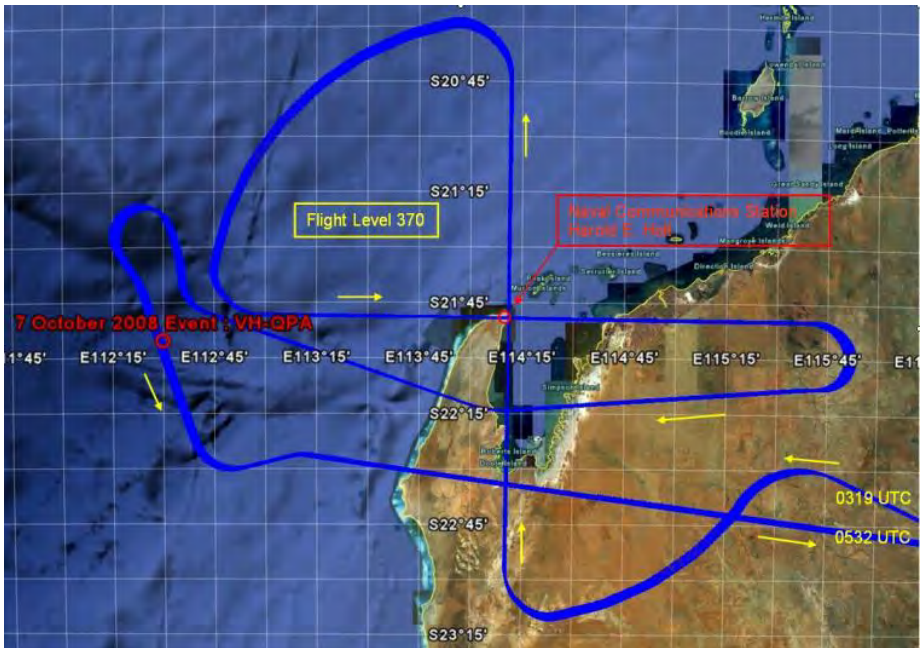

APPENDIX G: ELECTROMAGNETIC RADIATION

Introduction

All electrical systems generate some electromagnetic emissions, commonly called radio waves, either as an intended function of the system or as an unintended consequence of the physical properties of its electrical circuits. All systems can also be disturbed, to varying levels, by emissions from another source. Electromagnetic interference (EMI) is an undesired disturbance in the function of an electrical system as a result of electromagnetic emissions from another source.

Electrical systems, particularly aircraft avionics, are designed to be resilient to undesirable disturbances as a result of emissions from other systems, and also to minimise emissions that may cause disturbances in other systems. For example, the signal strengths within a system are designed to be far greater than the amount of currents and voltages that are expected to be induced by other systems, so that the magnitude of any unintentionally induced signal is too low to have an undesirable effect.

Types of EMI

Figure G1 provides an overview of the different types and sources of EMI.

Figure G1: Types and sources of EMI

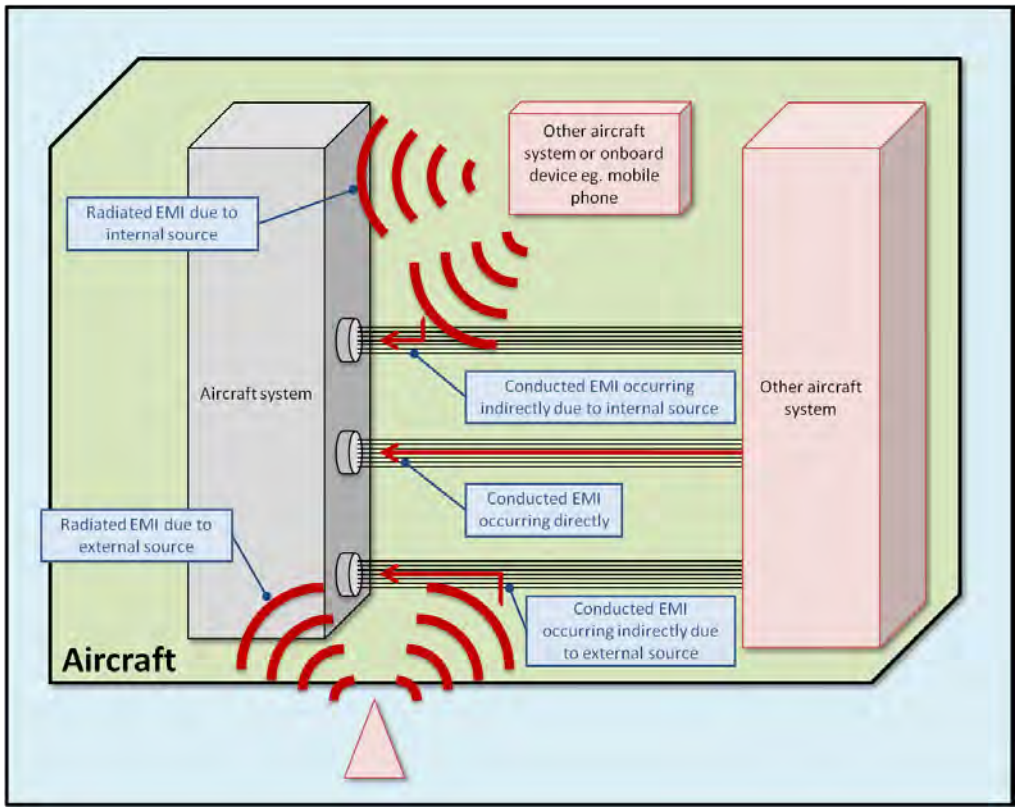

Emissions can be divided into two types: conducted and radiated. Conducted emissions travel along wire interconnects between systems or parts of a system. Radiated emissions travel through free space and are generated by time-varying electrical signals in a conductor. However, radiated emissions can induce currents in electrical wiring, and currents in wiring can emit electromagnetic radiation. In general, the potential for EMI increases with the amount of emitted energy, and decreases with distance.

A system may produce conducted and radiated emissions over a range of frequencies and varying magnitudes. A system may also be susceptible to conducted or radiated interference over a range of frequencies and magnitudes. Those characteristics of a system's emissions and susceptibilities are a consequence of the physical properties of the system, mostly by design. For example, an item of electrical equipment may be enclosed in a metal housing that prevents internal radiated emissions from emanating outside the unit and also shields the unit from external radiated emissions. However, the system's physical properties may change over time as a result of environmental effects and ageing, or if there is a hardware fault. For example, an electrical connection may degrade and result in undesired emissions.

Potential sources of EMI

In an aircraft, electromagnetic emissions may originate from a number of sources:

- other aircraft systems, especially 'transients' (temporary disruptions usually due to switching) from power supplies and other high voltage or high current systems, or digital noise between computer systems
- other onboard sources, including personal electronic devices (PEDs) such as mobile telephones, laptop computers, or handheld gaming devices carried by passengers or crew, or powered electronic devices in the aircraft's cargo
- external artificial sources, such as ground-based radar sites and communications facilities
- natural sources, such as electrical storms and electrostatic discharge.

Emissions may reach a system through various means and not always directly. For example, small electric currents may be induced in an aircraft's skin that creates electromagnetic fields inside the aircraft. In turn, these fields could then induce currents in a bundle of wires which then pass energy into an electronic system.

Effects of EMI

A device cannot always distinguish between real signals and spuriously induced voltages and currents. High levels of EMI may cause an interface between two devices to be severely degraded, while lower levels might cause a receiving device to react to false signals.

An example of this is a radio that picks up electromagnetic waves emitted from devices other than radio transmitters. These false signals cannot be easily distinguished from actual radio transmissions and so are converted into electrical signals as though they represented sound waves, and the radio speaker will emit audio noise that can impede or even drown out the intended transmission. Alternatively, the false signals may disrupt the real signals to an extent where they

cannot be ‘decoded’ by the receiver. Radio systems are usually the most susceptible to EMI as they are specifically designed to receive electromagnetic signals (that is, radio waves), but any electrical or electronic system may be affected by EMI, and these effects depend on the characteristics and function of the affected system. For example, an analogue electrical instrument may show incorrect readings or oscillations.

Digital signals, including those on digital databuses, can also be affected by EMI but in a different way. A low level of interference can be eliminated because of its typically different nature to the nature of the real signals, while higher levels can disrupt the flow of data in a way that is generally detectable by the receiving system. This is because the digital signals need to adhere to very strict criteria including the waveform shape (that is, the timing of pulses, voltage levels, and voltage rise/fall times) and often the validity of the data itself. This is the case with the ARINC 429 databus, which in addition to strict waveform criteria, uses a parity bit to aid the receiving system to detect any corruption to the signals (section 3.4.7). Any significant disruption due to the direct effects of EMI would result in a large number of invalid ‘words’ and random effects on the data values.

Typically, a high-power system such as a power supply would not be affected by noise from a relatively low-power computer system, whereas a computer system might be vulnerable to noise from the power supply, electric motors, other computer systems, and so on.

Electrical and electronic systems operate by sending signals between components, devices and other systems, and these interfaces can be similarly affected by spurious signals. The system could react to these spurious signals as though they were actual signals, or be unable to decode them, with resulting abnormal behaviour. For example, EMI occurring to a set of wires carrying electrical current for an aircraft’s electronic flight control system may be treated as real inputs, resulting in spurious movement of the flight control surfaces.

Similar effects may occur within a single device if one or more of its internal interfaces are exposed to EMI. A computer may shut down or ‘hang’ (or otherwise behave unpredictably) if a faulty electrical circuit generates spurious radio signals that are inadvertently picked up by another circuit.

Different aircraft, systems or components of the same design may exhibit varying levels of susceptibility to EMI due to factors such as shape, slight variations in manufacture, or ageing. For example, a particular wiring bundle may be routed slightly differently from one aircraft to another, which could change both its length (which in turn affects its impedance and resonant frequency) and its location relative to the aircraft’s structure and other equipment (which in turn affects the amount and type of electromagnetic emissions to which the bundle would be exposed). Other factors such as the location and electrical properties of seams and apertures, as well as the shape of electronics bays, can affect the electromagnetic environment around electronic units and wiring.

Although EMI has been an ongoing issue in aircraft design and has been known to regularly cause anomalous behaviour in airborne electrical systems, relatively few aircraft accidents have been attributed to EMI. This is in part due to the difficulty of proving whether EMI was a factor if the affected system is damaged or destroyed, and in part due to the usually indirect effect of EMI on safety (that is, communication and navigation systems may be disrupted, but the integrity of the aircraft’s structure and control systems are rarely affected).

Electromagnetic compatibility design

Modern electronic systems, particularly those installed in large-capacity commercial aircraft, are specifically designed to be resistant to EMI and to minimise emissions that could affect other systems.

Design features that can reduce both susceptibility and emissions include dedicated current return wiring, low-resistance single-point grounding, conductive shielding of circuits, the use of shielded and/or twisted pair wiring, categorisation and physical separation of wiring, capacitive and inductive filters on system inputs and outputs, and other circuit and hardware design techniques. Digital systems may also implement parity, checksums and other integrity checks on data.

EMI testing

The various means by which electromagnetic energy can pass from one place to another, and the number of factors which affect it, give rise to very complex relationships that can be difficult to characterise. For this reason, EMI testing (as opposed to various forms of analysis) has generally been considered the most representative means by which an aircraft or system may be evaluated for resilience to EMI.

Testing of a system's susceptibility to EMI can take two forms:

- Source-victim testing, which is specific to a particular 'victim' (affected) system and one or more 'source' (emitting) systems, often carried out when particular systems are known or suspected to have affected another system.
- Susceptibility testing, which subjects a system to an electromagnetic environment that comprises a range of frequencies. The frequency ranges, field strengths, and conducted currents are chosen to be representative of the environment to which the system is expected to be exposed in service. This form of testing is normally carried out as part of a system's qualification to ensure that it is resilient to a defined electromagnetic environment.

An advantage of susceptibility testing is that if a system is later subjected to emissions from an unanticipated source, it will most likely have previously been shown to be resilient to those emissions as long as the susceptibility testing encompassed the relevant frequency ranges and field strengths.

Conducted susceptibility tests induce interference on the wiring interfaces of the equipment, while radiated susceptibility tests subject the equipment to high strength radio waves. Typically, conducted susceptibility tests covered the 10 kHz to 400 MHz range, such as audio and very low frequency (VLF) frequencies. Radiated susceptibility tests covered the 30 MHz to 18 GHz range, such as high frequency radio and radar frequencies.

The strength of an electromagnetic field can vary significantly with factors such as locations and distances from the energy source, the presence or absence of metallic or other conductive objects, the quality of the electrical connections (such as grounding of electrical units), the presence of any corrosion, manufacturing variability and so on. Because of this, EMI testing will not always reveal susceptibilities, although most confounding factors can be mitigated by using high-power signals during susceptibility testing.

APPENDIX H: SINGLE EVENT EFFECTS

Note: the information in this section is primarily sourced from the International Electrotechnical Commission (IEC)[217] Technical Specification (TS) 62396 *Process Management For Avionics – Atmospheric Radiation Effects* and from advice from single event effects (SEE) specialists.

Origins of SEEs

There is a constant stream of high-energy galactic and occasional bursts of solar radiation entering Earth's upper atmosphere. The radiation consists of subatomic particles (mostly protons) travelling at extremely high speeds. These particles collide with molecules in the Earth's upper atmosphere and generate secondary particles (Figure H1), which then collide with more molecules in the atmosphere, creating a cascade of particles.

A particle can be either charged or uncharged. Charged particles are more readily slowed down by the atmosphere and by solid objects. Neutrons penetrate more deeply due to their absence of charge.

A typical computer silicon chip, or integrated circuit, is primarily made of a large number of tiny devices called transistors, each of which operates by accumulating and distributing electric charges within the chip. A subatomic particle that collides with an atomic nucleus inside a chip can change the amount of electric charge within a particular area (a phenomenon known as localised ionisation). If sufficient charge accumulates in a particular area, a transistor's behaviour can change.

The subsequent operation of the surrounding digital system will change according to the function of the transistor. When this occurs, the change in behaviour is known as an SEE.

Types of particle

Particles that can cause SEEs include:

- Neutrons, which are almost entirely generated through particle collisions in the atmosphere and have the greatest intensity (or 'flux'[218]) at altitudes used by commercial aircraft. Neutrons are difficult to shield against due to their neutral charge.
- Protons, also mostly generated through particle collisions but which have relatively low flux compared with neutrons (about 20 to 30%). Protons can cause similar effects as neutrons but are more readily absorbed by solid materials such as an aircraft's skin.

---

[217] The IEC is a worldwide organisation with the object of promoting international cooperation on standardisation in the electrical and electronic fields. It publishes a range of documents including International Standards and Technical Specifications.

[218] The unit of neutron flux used in this report gives the approximate number of neutrons that pass through an area of one square centimetre (cm^2) in 1 hour.

- Other particles such as heavy ions, pions, and electrons, which either have insufficient flux or insufficient energy to have significant effects on electronic equipment below 60,000 feet.

Figure H1: Single event effects on aircraft electronic equipment

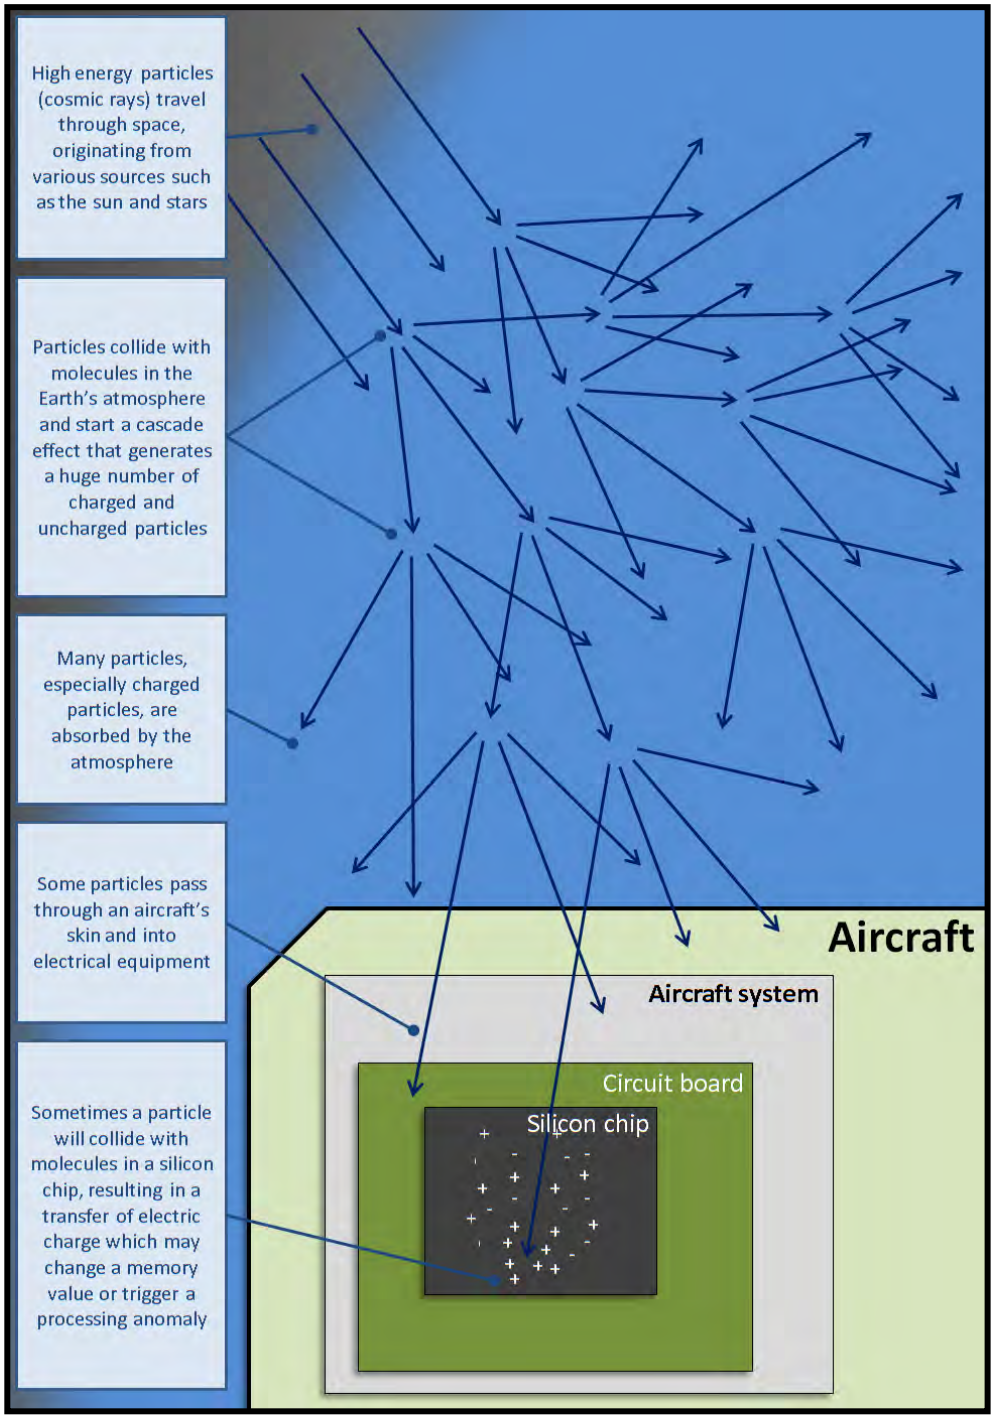

Factors affecting SEE exposure

An aircraft is constantly exposed to galactic and solar radiation, but the amount of exposure varies depending on a number of factors, including the following:[219]

- The flux of particles varies significantly with altitude and, for neutrons, is about 300 times higher at 40,000 ft than at ground level. At moderate latitudes (see below), the neutron flux peaks at about 60,000 ft, dropping to about 5,600 neutrons per cm^2 per hour at 40,000 ft, and less than 20 at the Earth's surface. The flux is somewhat lower at altitudes above 60,000 ft due to the reduced number of nuclear interactions with the Earth's atmosphere, although particles tend to possess higher energies at higher altitudes.
- Particle flux is higher at the Earth's magnetic poles than at equatorial latitudes due to the influence of the Earth's magnetic field. At 40,000 ft, the neutron flux is less than about 2,100 neutrons per cm^2 per hour in the tropical zone[220], and is about 8,400 between the 60° latitudes and the poles.
- There are natural variations in solar activity, which predominantly follow an approximately 11-year cycle (Dyer et al. 2003). In general, this cycle has minimal effect on SEE rates except during the occurrence of 'solar flares', which mostly occur during the solar maxima phase of the cycle. Solar flares can last from a few hours to several days, and in very rare cases can increase fluxes by a factor of up to 300 in polar regions (Dyer and Lei 2001).[221]

The low-energy neutron flux within an aircraft is higher than outside due to interactions with the aircraft's structure and contents (Normand et al. 2006, Dyer et al. 2006). These very low-energy neutrons (called thermal neutrons) are more likely to react with certain isotopes commonly used in the manufacture of silicon chips, resulting in an increase in the SEE rates.

Types of SEE

A particle collision within an integrated circuit can have one of four types of effect on a system:

- No effect. Not all collisions have an effect on the state or operation of a device and this is not characterised as an SEE.
- A non-destructive, temporary or recoverable change is known as single event upset (SEU), or a soft error. This is the case where the problem can be resolved without cycling the power (turning the device OFF and ON).

---

[219] The neutron fluxes in this section are based on the approximate levels given in IEC TS 62396 for neutrons with energy levels of greater than 10 million electronvolts, which is the energy range most likely to be associated with SEEs.

[220] Between about 23° north and south of the equator.

[221] Not all solar flares produce particles. With the exception of solar flares that produce particles, the rate of SEE at aircraft cruise altitudes decreases during the solar maxima because the solar activity reduces the presence of high-energy neutron particles from outside the solar system. However, the reduction is small. During a solar minima (last occurred in late 2008), the neutron flux is about 1.2 times higher than during a solar maxima.

- A non-destructive, unrecoverable change, or firm error.[222] These include single event functional interrupts (SEFI), in which the device ceases to function properly until its power is cycled.
- A destructive, permanent change, or hard error. These include single event latchups (SEL) and are the result of permanent damage to a component that is not recoverable even by cycling the power OFF and ON.

In common usage, both soft and firm errors can be referred to as SEU.

About 1% of SEEs affect more than one binary digit (bit) with a single particle interaction, known as a multiple bit upset (MBU).

Sensitivity to SEE

The probability of an SEE event occurring to any given system is dependent on numerous factors, including the:

- flux, energy and type of particles present at the system's position
- energy, location and direction of a striking particle
- sensitivity of a struck chip, including the number and size of internal transistors, and the amount of charge accumulation required to affect a transistor.

Any active electronic component or device can be susceptible to SEE, especially digital devices such as CPUs, memory, and other digital integrated circuits. As the chip density of integrated circuit components has increased greatly in recent decades, digital systems have generally become more sensitive to SEEs than previously. In particular, memory chips and central processing units (CPUs) are typically most sensitive as they have the highest transistor densities, although any kind of silicon chip may experience SEEs. Chip designs with higher component densities can also be more susceptible to MBU.

An integrated circuit's sensitivity to SEE varies throughout production due to normal variations in materials and dimensions from time to time. For integrated circuits produced in the late 1980s and early 1990s, the sensitivity to SEE between batches can be as much as a factor of 10, and a factor of two to three between integrated circuits within the same batch. Sometimes a specific variation will be unintentionally introduced into a batch that will make the components significantly more or less susceptible. A component that has had an SEE can sometimes become more vulnerable to subsequent SEE due to physical changes in its internal structure.

Passive devices such as analogue filters, capacitors, and resistors are not generally considered susceptible since their behaviour does not change significantly with the small charge transfers typical of SEE events.

Effects of SEE

Computer systems are generally very complex, incorporating large amounts of processing, transmission/reception, memory, and other devices. Within a single chip such as a CPU, thousands of computations are conducted every second and different

[222] Elsewhere in this report, usage of the term 'soft error' includes firm errors.

portions of memory are constantly being written, read, and rewritten. As a result of this complexity, it can be difficult to predict or trace the effect of an SEE.

Often, an affected transistor is part of a memory device and the SEE changes the data stored at one particular location in memory. The consequence of such a memory change depends on the type of data stored in that memory location, the point during the software program sequence at which the event occurs, and the value that the memory takes on as a result of the event.

Many SEE events result in a near-complete loss of system functionality (a 'crash' or 'hang') since the devices in a CPU are critical to nearly every function of the system. The effects of SEE on memory devices are usually more limited in scope. There have been SEEs on other systems where units provided ongoing data output errors without a complete loss of functionality. These were primarily associated with issues with a CPU's program counter (a segment of memory that points to the next instruction to be executed).

The potential for, and impact of SEE is very dependent on the architecture, components, and software of each specific device or system. A particle could affect any one (or more) of billions of transistors and at any particular moment in time in a system that changes its state millions of times a second, so it can be very hard to predict the outcome. For example, the downstream effect of an SEE on a memory cell is highly dependent on the purpose of the memory cell and whether it is read and used prior to it being overwritten.

Effects on spacecraft and aircraft

The risk of SEE is a primary concern for designers of spacecraft, partly due to the very high levels of reliability required by such autonomous and costly systems. Accordingly, spacecraft SEE is better understood and more widely studied than ground or atmospheric SEE.

A number of studies have examined SEE using test equipment (comprising memory chips and equipment to continually monitor the state of those chips) either carried on aircraft or in a simulated aircraft environment. Studies[223] in the 1990s reported rates equivalent to about 10^{-4} to 10^{-2} per flight hour for 544 kB of random-access memory (RAM).[224]

The prevalence of transient SEE on commercial avionics equipment is not well studied or recorded, for several reasons:

- A fault detected by an error detection and correction (EDAC; see *Single event effects mitigation strategies* below) system would not normally be recorded or analysed.
- A transient fault occurring on a non-EDAC avionic system in flight would not reoccur when the system is subsequently tested on the ground. In most circumstances, the engineers would then record 'no fault found' in the maintenance records and return the system to service. Although this action could

[223] Normand (1996) summarised a number of tests conducted up to the mid-1990s and correlated the results with theoretical models. Johansson et al. (1998) and Olsen et al. (1993) also conducted relevant studies.

[224] The figures were adjusted for the amount of RAM used by the LTN-101 ADIRU.

be regarded as appropriate since the system is serviceable, it does not enable easy collection and analysis of the rates and effects of SEE.

- It can be difficult to determine if a transient fault in a non-EDAC system was the result of SEE or other factors. It is therefore likely that most SEE are not recognised or reported.
- Many transient faults would not be recognised by other aircraft systems or by flight and maintenance crews. For example:
  - a corrupted data parameter in a computer's memory may be overwritten before it is used elsewhere
  - a single corrupted databus message would simply be ignored by other systems
  - a short-duration erroneous parameter would normally be filtered out by downstream systems, or have effects that would not be noticeable (such as a momentary change in a cockpit display or an engine's fuel flow).
- The rates of occurrence of SEE in commercial avionic systems, if known, are generally considered proprietary information and are not normally shared.

SEE mitigation strategies

Hardware and software design features can be used to mitigate the effects of SEEs, including:

- Radiation hardening uses less sensitive internal circuit components and designs, including integrated circuits that use higher voltage levels or larger transistors. A less common form of hardening involves the use of a physical shield to absorb some energetic particles before they can pass through onto the chip, but this technique is generally impractical for airborne systems where weight is a limitation. An aircraft's physical structure can shield against low-energy charged particles but can actually increase the fluxes of certain energy neutrons as a result of nuclear interactions.
- Redundancy provides duplicate systems, subsystems, or components which enables a fault (regardless of whether triggered by an SEE) to be detected. For example, a set of transistors can be duplicated within a chip, an entire block of memory may be duplicated on multiple chips, or normally idle CPU computation cycles can be used to execute duplicate instructions. These methods generally require additional hardware and complexity.
- Databuses usually use simple forms of data redundancy methods such as parity checks and checksums, where extra information is transmitted to enable data corruption to be detected.
- Partitioning enables part of a faulty system to be isolated from other parts, which can then continue to operate correctly.
- EDAC is a form of redundancy that stores additional ('redundant') information[225] in memory, to enable data to be checked when it is read.

[225] Redundancy in this context refers to the storage and/or transmission of more information than would be required by the system in the absence of data corruption. Methods of producing this redundant information include duplication of data, parity, checksums, and other more complex algorithms.

Depending on the form of EDAC, single-bit and minor multiple-bit errors can be corrected using the extra information, and more serious multiple-bit errors can often be detected and disregarded. The degree to which EDAC can be retroactively implemented on extant systems is limited and the more effective forms of EDAC require specially designed hardware.

- Many forms of built-in test equipment (BITE) are capable of detecting SEE even if not specifically intended to do so. For example, loopback tests read data back from the databus and compare it with the original data. This method is not effective if the data is corrupted prior to transmission but is very effective in detecting corruptions that occur in the transmission/reception data paths.

SEE testing and diagnosis

The extent to which a type of fault could be due to SEE is usually diagnosed by comparing the relative likelihood of the fault with factors that affect the neutron flux (particularly altitude and latitude), and taking into account other hypotheses, information about the system, and the local environment at the time of the event. However, this approach was not useful for the data-spike failure mode due to there being only a very small number of events.

Depending on the device's sensitivity, only a small number of particles that pass through a single memory location will cause an SEE. Therefore, a practical amount of testing will not be likely to generate all potential problematic effects in a system that could have millions of memory locations.

Often a system's behaviour returns to normal after it is restarted after an SEE or other disrupting event. Because of this, it is difficult to definitively ascertain whether a particular fault could have been caused by SEE, except by exclusion of other potential causes.

A system's resilience and response to SEE can be evaluated through testing, which typically involves bombarding the unit with high-energy particles. The system's resilience can be examined by running special test software that continually logs changes in RAM data, and the responses can be determined by running the normal operational software and checking for observable effects. Both types of test have limitations, particularly the latter for which a common or obvious effect might mask a less common or obvious one.

APPENDIX I: PASSENGER QUESTIONNAIRE

Questionnaire design

Given the nature of the occurrence and the large number of injuries, the investigation wanted to obtain information from as many of the passengers as possible. The most effective way of obtaining the information was the use of a questionnaire.

A draft questionnaire was developed by Australian Transport Safety Bureau (ATSB) investigators based on previous ATSB passenger questionnaires and information about the occurrence that was obtained from sources such as interviews with the flight crew and cabin crew. The draft was distributed for feedback within the ATSB and external parties on 23 October 2008. Minor changes were made based on the comments received.

The final 15-page questionnaire contained a mixture of questions requiring the selection of a response from a list and questions requiring a free-text response. The questions were provided in the following sections:

- general information: including the passenger's name, gender, age range, seat number, and number of previous flights
- safety information: including the amount of attention the passenger gave to the pre-flight safety demonstration and the safety information card
- in-flight upset events: including, for both upset events, the passenger's posture and location, what the passenger was doing at the time, what they saw, heard and felt during the upset, and what they did following the upset
- seat belt use: including the passenger's understanding of when seat belts should be worn, previous use of seat belts during different phases of flight, seat belt use during both upset events, reasons for not wearing a seat belt (if not worn), any problems with the seat belt, recollection of any crew reminders prior to the upsets for passengers to use seat belts, and the location of other passengers observed to be not wearing their seat belts
- injuries: whether the passenger was injured (and if so a description of the injury, how the injury happened, and the nature of any medical treatment)
- children: if the passenger was travelling with children, the name and age of the child, the posture and location of the child during the first upset, the adequacy of the child's seat and restraint, and whether the child was injured (and if so, a description of the injury, how the injury happened, and the nature of any medical treatment)
- other passengers: including the passenger's recollection of other passengers' injuries and whether the passenger provided any assistance to other passengers
- use of electronic equipment: whether the passenger was operating personal electronic equipment at the time of the upsets (and if so a description of the equipment, its operating mode and whether there were any problems with the equipment) and whether other passengers nearby were operating personal electronic equipment

- crew actions: the passenger's recollection of any instructions from the flight crew following the upsets, or any instructions or actions by the cabin crew following the upsets
- suggestions and additional comments.

The questionnaire could be completed electronically, on a hard copy form, or by interview if requested by the passenger.

Distribution of the questionnaire

Contact details were obtained for the passengers from the operator. For many passengers, this included an email address, telephone number, and physical address. However, for many passengers one or more of these details were missing. Further contact details were obtained from the passengers' immigration cards.

The questionnaire was distributed electronically from 28 October 2008 to adult passengers with available email contact addresses. For passengers with only a postal address available, the questionnaire was mailed. Many of the passengers without an email address but with a telephone number were contacted by telephone to obtain an email address or postal addresses to distribute the survey.

Most passengers who completed the questionnaire were asked to pass a blank questionnaire on to other passengers they knew or ask those passengers to contact the ATSB. In addition, reminders to complete the questionnaire were sent out to passengers with known email addresses prior to the release of the preliminary investigation report (18 November 2008) and the first interim factual report (6 March 2009). The preliminary report also asked passengers who had not received a questionnaire to contact the ATSB.

Ultimately, most of the adult passengers were contacted about the questionnaire, either by email or by telephone. However, attempts to contact about 24 adult passengers were not successful and a questionnaire was not able to be distributed. Most of these passengers were from Asian (12) or European countries (11 passengers).

Summary information on questionnaire respondents

The final number of completed questionnaires was 98. The response rate for adults was 35% (98 out of 277). As 24 of the passengers were not contacted, the response rate for adults who were able to be contacted was 39% (98 out of 253). In addition, the investigation also obtained some information on the key questionnaire topics by interview or correspondence from 21 other passengers. The information from questionnaires, interviews and correspondence also included details on some pertinent topics for many other passengers.

The demographic information for all the passengers is provided in Table I1, and the demographic information for the questionnaire respondents only is provided in Table I2.

Table I1: Demographic information for all passengers

Category		Front	Centre	Rear	Total	Proportion
Gender	Male	17	76	62	155	51%
	Female	16	73	59	148	49%
Nationality	Australia	26	61	42	129	43%
	Europe	6	38	21	65	21%
	Asia	0	44	53	97	32%
	Other	1	6	5	12	4%
Age	Infant (0-1)	0	4	2	6	2%
	2 to 12	5	4	8	17	6%
	13 to 17	1	0	2	3	1%
	18 to 30	2	25	38	66	21%
	31 - 45	9	34	30	72	24%
	46 to 60	8	45	29	82	27%
	61 to 75	7	34	8	49	16%
	Over 75	1	3	3	7	2%
	Unknown	0	0	1	1	0%
Total		33	149	121	303	

Table I2: Demographic information for questionnaire respondents

Category		Front	Centre	Rear	Total	Proportion
Gender	Male	8	21	20	49	50%
	Female	4	27	18	49	50%
Nationality	Australia	9	22	16	47	48%
	Europe	2	15	8	25	26%
	Asia	0	10	14	24	24%
	Other	1	2	0	1	2%
Age	18 - 30	0	7	9	16	16%
	31 - 45	3	10	7	20	20%
	46 - 60	5	16	16	37	38%
	Over 60	4	15	6	25	26%
Total		12	48	38	98	

The questionnaire respondents were compared to the other adult passengers to determine the extent to which the survey sample was representative.[226] The comparison indicated that:

- There was no apparent difference between survey respondents and other adult passengers in terms of the proportion of males and females.
- There was no statistical difference between the proportion of adult passengers who responded to the questionnaire from the front (12/27, 44%), centre (48/141, 34%) or rear sections (38/109, 35%) of the aircraft. After accounting for the 24 adults who were not contacted (1, 13 and 10 in front, centre and rear), the response rates were 46%, 38% and 38%.
- A higher proportion of the adult passengers completed the questionnaire from Australia (43%) and European countries (39%) compared to passengers from Asian countries (26%). After accounting for the 24 adults who were not contacted, the response rates were 42%, 47% and 29% respectively.
- A higher proportion of passengers aged over 45 (44%) completed the questionnaire compared to passengers aged 18 to 45 (26%). Most of those who were not contacted were above 45.
- Overall, there were reasonable proportions of survey respondents from each section of the aircraft, and there were reasonable proportions of respondents from the major demographic groups present on the occurrence flight.

There was no evidence that passengers who were injured were more likely to complete the questionnaire. The proportion of questionnaire respondents who received hospital medical treatment (18%) was no different to the proportion of other adult passengers who received hospital medical treatment (17%).

[226] Statistical comparisons were done using the χ^2 (Chi squared) test for independent groups. In this report, 'statistically significant' means that the chance of the difference being present due to chance alone was less than 5%.

APPENDIX J: EXAMINATION OF POTENTIAL FOR INADVERTENT RELEASE OF SEAT BELTS

Passenger reports

Six passengers reported that they were seated with their seat belt fastened at the time of the first upset, but that the belt became unfastened and did not restrain them in their seats. Three of those passengers advised that they had their belts tightly fastened, and three advised that they had their belts loosely fastened. None of the six passengers could provide details of how their seat belts released.

Seat belt description

The passenger seat belts on the operator's A330 aircraft were manufactured by Amsafe (part number 2011-1-661-2258). The belts were a very common type of lap belt with a lift-lever buckle (Figure J1).

Figure J1: Seat belt design

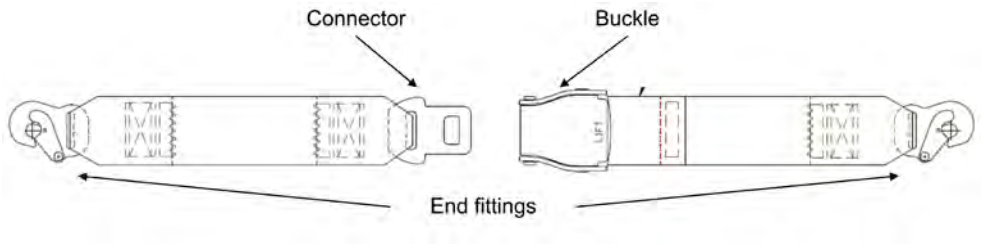

The buckle was on the passenger's left side and the connector on the right side, with the end fittings of both parts being attached to anchorage points on the aircraft seat. To fasten the seat belt, the connector was inserted into the buckle. To release the belt, the buckle cover (or flap) was lifted 30° or more (Figure J2).

Figure J2: Seat-belt buckle on VH-QPA with buckle cover open

The seat belt was designed so that the belt passed over the passenger's pelvis. A passenger could adjust the tightness of the belt by adjusting the distance of the buckle from its end fitting. In general, when the belt was firmly fastened, the buckle

was centred across the passenger's hips. If the belt was loosely fastened, the buckle half of the belt would be longer than the connector half.

Seat belt design requirements

Certification requirements for seat belts were specified in United States (US) Federal Aviation Regulation 25.785 and in European Certification Specification (CS) 25.785. These requirements stated that each seat and seat belt '...must be designed so that a person making proper use of these facilities will not suffer serious injury in an emergency landing...'

More detailed design and testing requirements for aircraft seat belts were outlined in the SAE Aerospace Standard (AS) 8043 (*Restraint systems for civil aircraft)* and AS8049 (*Performance standards for seats in civil rotorcraft, transport aircraft, and general aviation aircraft*).

AS8049 noted that the seat and the restraint system needed to be considered as a total system. It also required that a lap belt be fitted so that the belt passed over the occupant's pelvis. It also stated that, during qualification tests for seats, the seat belt should be 'snug, but not excessively tight' and that all slack be removed.

The seat belt manufacturer advised that seat belts were not designed to be worn improperly adjusted. It also reported that it would not be possible to ensure proper placement of the belt on the occupant's body when belts were worn 'extremely loosely fastened'.

AS8043 outlined a series of requirements for seat belts in terms of their strength and materials. It also included the following requirements for ease of use and inadvertent release:

> A restraint system shall be provided with a single buckle having a single motion release which is readily accessible to the occupant to permit easy and rapid egress by the occupant from the assembly. The buckle release mechanism shall be designed to minimize the possibility of inadvertent release.

For a lift-lever buckle, AS8043 required that the handle provide access for two or more fingers of either hand to actuate the release. However, the standard did not specify a release angle (that is, the angle the buckle cover was required to be lifted up before the belt would release). DeWeese and Gowdy (2002) noted that the United Kingdom required that the release angle be between 70 and 95° whereas most US manufacturers used a release angle of between 30 and 45°.[227] They also noted that, while a larger release angle may decrease the potential for inadvertent release, it could also increase the difficulty of releasing the seat belt if the occupant was in a folded position due to post-crash injuries, debris or aircraft inversion.

[227] In their research, they found no difference in the time to egress an aircraft for different buckle release angles (30°, 60° or 90°).

Potential for inadvertent seat belt release

During examinations of the aircraft's seat belts, investigators identified a scenario by which a loosely-fastened belt could inadvertently release. The scenario involved the following:

- The seat belt had to be very loosely fastened. The buckle of the belt could then slide down off a passenger's right hip.
- The buckle had to be in a vertical orientation under the passenger's right armrest.
- A vertical force needed to be applied such that the buckle cover would get caught on the underside of the armrest or a horizontal ridge on the armrest. If the buckle cover was caught and a vertical force continued to be applied, the buckle would release the connector.
- When the buckle of a very loosely-fastened belt was placed in a position underneath the armrest, investigators could consistently make the seat belt release by positioning the buckle underneath the armrest and then standing up.

Subsequent examinations showed that the applicability of this inadvertent release scenario was not restricted to seats on A330 aircraft or to the operator's aircraft. The scenario was replicated on aircraft from multiple other operators and manufacturers, although it was more difficult to do on some seats than others. The potential for the scenario depended on a range of factors, including the design of the seat and armrest.

- For the scenario to occur on the operator's A330 aircraft, the seat belt had to be adjusted so that there was at least 25 cm of slack in the belt, regardless of the size of the occupant.[228] This meant the belt had to be adjusted to be at or near the end of its adjustment range. The amount of slack was determined by comparing the length of the buckle half of the belt for a firmly-fastened seat belt with one that was loosely-fastened to the minimum extent necessary to enable the inadvertent release scenario to occur.

The seat belt examinations also noted the following:

- A seat belt with 25 cm of slack was very loose and would be difficult to keep in place over a person's hips during a flight. A very loose belt would therefore increase injury risk even if it remained fastened.
- Even if the seat belt had sufficient slack to enable the inadvertent release scenario to occur, the buckle would often not naturally position itself below the armrest. Sometimes a significant amount of manoeuvring in the seat was required before the buckle would be in the necessary position.
- When the inadvertent release scenario did occur, the buckle cover would move to an angle significantly greater than 30°. Therefore, increasing the buckle release angle to more than 30° would not significantly decrease the likelihood of the scenario.
- Overall, the seat belt was simple in design, easy to use and, more importantly, easy to unfasten in the case of an emergency evacuation.

---

[228] The same result was achieved when examining two other aircraft types with slightly different seats but similar seat belts.

Previous occurrences

The seat belt manufacturer, aircraft manufacturer, aircraft operator, and investigation and regulatory agencies associated with the investigation[229], had not previously been aware of this inadvertent release scenario associated with the seat armrest.

A review of a sample of investigation reports into turbulence accidents identified only two cases where a passenger reported being seated with the seat belt fastened and that the seat belt became unfastened during the event. In those cases:

- Passenger one reported that he checked that his belt was fastened during the flight because it did not seem to tighten well, but no information was provided regarding whether the belt was fastened loose or tight at the time of the turbulence. No problems were found with the belt in a subsequent inspection.[230]
- Passenger two was carrying a bag on her lap, and examinations found that such a bag could, if moved sideways, contact the buckle cover and release the seat belt.[231]

Relevance to the occurrence flight

The inadvertent release scenario could not occur to a tightly-fastened seat belt, and therefore the scenario did not occur for at least three of the six passengers who reported having their belts fastened but were not restrained. However, whether the inadvertent release scenario occurred for any of the other three passengers could not be determined based on the available information. It is worth noting that there was no previous knowledge of this scenario actually occurring during an in-flight upset, despite the widespread use of this type of seat belt throughout the aviation industry.

It is also worth noting that a seat belt had to be very loosely fastened before the scenario could occur. Seat belts that were loosely fastened would pose a significant injury risk even if they remained fastened. They would be difficult to keep positioned across the passenger's pelvis, and they would also allow significant movement of the passenger before restraint.

Although the six passengers reported that they had their seat belts fastened at the time of the first in-flight upset, the investigation could not confirm that this was the case. It is possible that their seat belts may have been inadvertently released prior to the upset during their movement in the seat during the flight. It is also possible that the seat belts inadvertently released due to the movement of their arms or other objects located close to the buckle at the time.

---

[229] The investigation agencies included the ATSB, the US National Transportation Safety Board, and the French Bureau d'Enquêtes et d'Analyses pour la sécurité de l'aviation civile. The regulatory agencies included the Australian Civil Aviation Safety Authority (CASA), the United States Federal Aviation Administration (FAA), and the European Aviation Safety Agency (EASA).

[230] Irish Air Accident Investigation Unit report 2002/007 (ATR 42-300, EI-CPT, Mt Errigal, 2 December 2001).

[231] Japan Transport Safety Board report AA2008-01 (Boeing 767-300, JA611J, 27 km south-east of Narita, 27 October 2007).

APPENDIX K: INJURIES DURING IN-FLIGHT UPSETS

Australian in-flight upsets with serious injuries

During the period 1991 to 2009, there were seven in-flight upset events resulting in at least one serious injury in Australia or that involved Australian operators.[232] Details of these occurrences are provided in Table K1. The A330 accident near Learmonth on 7 October 2008 was the most serious event. Consistent with the 7 October 2008 accident, the injuries in the other six events occurred to occupants who were not seated or who were seated and not wearing their seat belts.

Table K1: Australian in-flight upsets resulting in serious injuries, 1991-2009

Date	Aircraft type	Details
30 March 1992	Boeing 747	Turbulence event. Seat-belt sign not on. One flight attendant (unseated) seriously injured.
6 July 1996	Boeing 747	Turbulence event. Seat-belt sign not on. Six flight attendants injured (one seriously) and 24 passengers injured (two seriously). All of the injured occupants were unrestrained, although the available information did not state how many of the passengers were seated or not seated.
5 December 1996	Airbus A340	Autopilot disconnection and aircraft pitch-up (following incorrect control selection by the crew). Seat-belt sign not on. Three flight attendants injured and eight passengers injured (one seriously). All of the injured occupants were unrestrained, although the available information did not state how many were seated or not seated.
27 October 2000	Boeing 747	Turbulence event. Seat-belt sign not on. Two passengers (not seated) seriously injured.
8 April 2002	Boeing 767	Autopilot disconnection and aircraft pitch-up (associated with windshear). Seat-belt sign not on. One flight attendant (not seated) seriously injured.
23 May 2007	Boeing 747	Turbulence event. Seat-belt sign not on. One passenger (not seated) seriously injured.
7 October 2008	Airbus A330	Autopilot disconnect and aircraft pitch-down (due to system problems). Seat-belt sign not on. Twelve occupants seriously injured and at least 107 received minor injuries. Most of those injured were not seated or seated without their seat belts fastened.

[232] All of the events involved high capacity aircraft (that is, more than 34 seats). In addition to in-flight upsets, there were four other events involving high capacity operators during the period that resulted in serious injuries; all occurred on the ground. Three of the events involved one injury only, and in the other event four occupants were seriously injured during an emergency evacuation.

Australian turbulence events

The most common type of in-flight upset is due to turbulence. There were 37 turbulence incidents that involved Australian operators or occurred in Australia during the period 2007 to 2009 and resulted in injuries. Those occurrences resulted in a total of 41 injuries to cabin crew, one injury to a flight crew member, and 13 injuries to passengers. One of the injuries was serious.

For many of the occurrences, information about the occupants' posture and seat belt use were not available. However, at least 32 of the flight attendants and four of the passengers were known not to be seated. Only one flight attendant was known to be seated; that attendant was wearing a seat belt but was injured by an unsecured service cart.

In one occurrence, five passengers were injured; all were either not seated or were seated but not wearing a seat belt. Four of the other injured passengers were known to not be seated, and three were known to be seated. One of the seated passengers was not wearing a seat belt even though the seat-belt sign was on. The other two injured passengers who were known to be seated were wearing their seat belts; one was injured by hot liquid from a spilt urn, and another received a neck injury.

In almost all cases, the injuries occurred when the seat-belt sign was not on, or very soon after the seat-belt sign came on and before action could be taken.

Turbulence events involving US operators

A number of reviews of turbulence-related accidents have been conducted on flights operated by US airlines. Statistics from those reviews include:

- During the period 1980 to 2003, only four people received serious injuries during turbulence that were seated with their seat belts fastened (excluding cases where occupants were hit by other occupants who were not secured).[233]
- During the period 1980 to 2008, there were 234 turbulence accidents, resulting in 298 serious injuries and three fatalities. Of these injuries, 184 involved flight attendants and 114 involved passengers. At least two of the three fatalities involved passengers who were not wearing their seat belts while the seat-belt sign was on.[234] A previous version of the same document stated that, between 1981 and 1997, 73 of the 80 passengers who were seriously injured did not have their seat belts on when the seat-belt sign was on.
- During the period 1982 to 1991, there were 55 accidents resulting in serious injuries (Flight Safety Foundation, 1994). Most (60%) of the serious injuries occurred after the seat-belt sign had been turned on in time for passengers to comply. In all except one case, passengers who were injured had not complied with the seat-belt sign and verbal instructions from the crew. The majority of injuries to flight attendants occurred when they were conducting normal duties or attempting to secure the cabin after the seat-belt sign was turned on.

---

[233] Federal Aviation Administration (2006). *Preventing Injuries caused by turbulence*. Advisory Circular 120-88A. Washington DC: FAA.

[234] Federal Aviation Administration (2009). *Turbulence: Staying Safe*. Updated 4 August 2009. Retrieved 6 August 2009 from: http://www.faa.gov/passengers/fly_safe/turbulence/

- During the period 1992 to 2001, there were 92 accidents involving injuries to flight attendants, and there were a total of 82 serious injuries and 97 minor injuries.[235] Only five of the flight attendants were known to be seated at the time, and only one of these had their seat belt fastened. In that case, the injury occurred when the flight attendant attempted to stop a loose service cart with their foot. Overall, 70% of the injuries occurred when the seat-belt sign was on.
- During the period 1984 to 1999, there were 131 turbulence-related accidents. Flight attendants accounted for 4% of the occupants, 52% of the fatal and serious injuries, and 22% of minor injuries (Tvaryanas 2003). The difference in risk was attributed to the requirement that flight attendants be unrestrained while performing the majority of their crew duties.
- During the period 2002 to 2008, turbulence encounters accounted for more serious injuries than all other types of occurrences (Matthews 2009). Flight attendants accounted for 87% of the serious injuries and 4% of the occupants. The author noted that possible behavioural changes resulting from increased security concerns since 11 September 2001 may have reduced passenger injury risk.

In summary, injuries associated with turbulence encounters and other in-flight upsets have demonstrated that injuries are much more likely to occur when the occupants are not seated or are seated without their seat belt fastened.

[235] Commercial Aviation Safety Team (2001).*Turbulence Joint Safety Analysis Team: Analysis and Results*. Retrieved 6 August 2009 from http://www.cast-safety.org/pdf/jsat_turbulence.pdf

APPENDIX L: SEAT BELT USE IN ROAD VEHICLES

Research has demonstrated that seat belts are very effective in reducing the frequency and severity of injuries in road vehicles. For example, one study concluded that wearing seat belts reduced the probability of being fatally injured in a road vehicle accident by about 50% for front-seat occupants and 25% for rear-seat occupants (Elvik and Vaa, cited by Austroads 2009).

Given the effectiveness of seat belts, there has been a substantial amount of research conducted into the factors influencing the use of seat belts. Austroads[236] (2009) recently conducted an extensive review of the research literature, with much of the research being done in the US but also including research done in Australia and in other countries. The review made the following conclusions:

> While vehicle occupants' reasons for not wearing seat belts are many and varied, and at the risk of over-simplifying a problem that is associated with multiple antecedents and causes, the prevalence of seat belt use is noted to be somewhat lower among:
>
> - males (particularly young to middle age 'adult' males)
> - younger rather than older age 'adult' occupants
> - those of lower socioeconomic status and education and those occupying non-professional or 'blue collar' positions
> - Indigenous or non-Caucasian occupants
> - passengers, particularly in the back or rear seat, rather than drivers
> - travellers on non-freeway/highway/multiple lane roads
> - travellers in rural areas (except perhaps when travelling on highways or freeways)
> - those travelling short distances, at low speeds, at night
> - those driving older vehicles and certain vehicle types such as 4WD and utilities.
>
> In addition to the above, non-use is likely to be greater among persons who:
>
> - are disposed to risk taking
> - have less positive attitudes to belt use, perceive greater barriers to use, and perhaps perceive less normative pressure to wear a belt
> - engage in other risky driving practices such as speeding, drink-driving and have a general disposition to violate road rules
> - are more often involved in crashes.

The report also noted that the evidence relating to ethnicity and seat belt use was inconsistent, particularly when comparing groups other than Indigenous groups in the US. Limited research has been done in Australia.

---

[236] Austroads is the association of Australian and New Zealand road transport and traffic authorities and aims to promote improved road transport outcomes.

In terms of measures to increase seat belt use, the report concluded:

> Seat belt legislation and enforcement have been in existence in Australia and other countries for well over four decades. These appear to be the simplest and most cost-effective long-term approaches to increasing seat belt use rates, and can be enhanced by publicity and education campaigns to raise the awareness of the benefits of wearing seat belts.
>
> In addition, much research effort has been applied to the development and implementation of seat belt reminder and interlock technologies. Unfortunately, only a few evaluations of their effectiveness have been conducted, but of those available, it appears that reminder systems offer overall high protective value, can achieve high levels of seat belt wearing, and target all groups of wearers from full-time and part-time wearers to consistent non wearers. Reminder systems that are aggressive and adaptive (changing characteristics during a trip) appear to be optimal. Interlock systems appear to benefit those most resistant to wearing seat belts in particular.

The International Traffic Safety Data and Analysis Group (IRTAD) recently published a report summarising trends in road safety in 27 member countries.[237] In that report, each country included information regarding several aspects of road safety, including seat belt use. The Australian report noted that seat belt use was estimated to be greater than 95% for front-seat occupants and more than 80% for rear-seat occupants. The Australian report also noted that:

> Despite high general usage rates, the rates of non-use among fatally injured vehicle occupants are still estimated at 28%. Analysis indicates that this high figure is the result of a high crash involvement rate among those who do not wear belts, as well as the fact that they are more likely to be killed if involved in a crash.

The seat belt use rates for several countries are included in Table L1. Caution should be used when comparing countries because data may have been collected in different ways and in different conditions. Overall, seat belt use in Australia and major European countries was higher than for the three Asian countries included in the report (Japan, Korea and Malaysia), particularly for rear-seat passengers. It should be noted that seat belts in rear seats became mandatory in these Asian countries since 2008 whereas they became mandatory in other countries earlier. The extent to which these results apply to other Asian countries (such as India and Singapore), or to seat belt use on aircraft, is not clear.

---

[237] IRTAD *Annual Report 2009*. IRTAD is a permanent working group of the Joint Research centre of the Organisation for Economic Co-operation and Development and the International Traffic Forum.

Table L1: Recent seat belt use rates (from IRTAD Annual Report 2009)

Country	Front seats	Rear seats	Comments
Australia	> 95%	> 80%	Mandatory all seats since 1970s.
Canada	92%	-	Data is for drivers only.
France	98-99%	-	Data is for drivers only.
Germany	96-98%		Data is for drivers only. Mandatory all seats since 1984.
Ireland	90%		Mandatory all seats since 1979.
Japan	97%	46%	Mandatory in rear seats in 2008.
Korea	74-88%	12%	Mandatory in front seats in 1990, rear seats (motorways) in 2008. Lower result in front seats for passengers.
Malaysia	76-92%	19%	Mandatory in front seats in 1978, rear seats in 2009. Results are latest figures; rate for rear seats was higher in earlier 2009 period. Lower result in front seats for passengers.
Netherlands	94-95%	81%	Mandatory all seats since 1992.
New Zealand	95%	87%	Mandatory all seats since 1970s.
United Kingdom	94-95%	86%	Mandatory all seats since 1991.
United States	83%	74%	Laws vary between states.

APPENDIX M: PUBLIC SAFETY INFORMATION ABOUT WEARING SEAT BELTS ON AIRCRAFT

Public safety advice about the use of seat belts on aircraft has been provided by various safety agencies, including the Australian Civil Aviation Safety Authority, the US Federal Aviation Administration, and the Australian Transport Safety Bureau.

Civil Aviation Safety Authority

The Australian Civil Aviation Safety Authority (CASA) website contained the following information for passengers on a webpage about turbulence[238]:

> **Injury prevention**
>
> In-flight turbulence is the leading cause of injuries to passengers and crew. Occupants injured during turbulence are usually not wearing seatbelts, ignoring recommendations to keep seatbelts fastened even when the signs are not illuminated. It is recognised that passengers need to move around the cabin to use restroom facilities or to exercise on long flights. However you should keep your seatbelt fastened at all times when seated.
>
> From 1981 through 1997 there were 342 reports of turbulence affecting major air carriers. Three passengers died, two of these fatalities were not wearing their seat belt while the sign was on. 80 suffered serious injuries, 73 of these passengers were also not wearing their seat belts.
>
> **Turbulence related incidents**
>
> The following are recent jet airliner mishaps from around the world. In each event, at least one passenger/flight attendant was injured during an unexpected turbulence encounter.
>
> - During a flight from Singapore to Sydney with 236 passengers and 16 crew, the airplane encountered turbulence over central Australia. The plane hit an "air pocket" which caused it to drop 300 feet. Nine passengers including one pregnant woman and three crew members suffered various neck, back and hip injuries, with one of the passengers requiring surgery. Those who were injured were not wearing seat belts.
> - During a flight from Japan to Brisbane 16 passengers were injured when a large aircraft encountered turbulence. Passengers had been advised to keep their seatbelts fastened while seated. The pilot in command reported that flight conditions were smooth prior to encountering the turbulence. The weather radar did not indicate adverse weather, so the crew did not turn on the seatbelt signs. A number of the passengers who were not wearing their seatbelts were injured when they were thrown from their seats.
> - A jet hit air turbulence shortly before it landed at a Hong Kong airport, injuring 47 people, seven of them seriously. "It happened very suddenly and everything was very chaotic," one of the 160 passengers aboard the flight said. "The plane just dropped and I saw things flying all over."

[238] Obtained from www.casa.gov.au/scripts/nc.dll?WCMS:STANDARD::pc=PC_91477.

US Federal Aviation Administration

The United States Federal Aviation Administration (FAA) website also contained seat belt information on a webpage regarding turbulence, including the following:[239]

> While turbulence is normal and happens often, it can be dangerous. Its bumpy ride can cause passengers who are not wearing their seat belts to be thrown from their seats without warning. But, by following the guidelines suggested on this site, you can help keep yourself and your loved ones safe when traveling by air.
>
> To keep you and your family as safe as possible during flight, FAA regulations require passengers to be seated with their seat belts fastened:
>
> - When the airplane leaves the gate and as it climbs after take-off.
> - During landing and taxi.
> - Whenever the seat belt sign is illuminated during flight.
>
> Why is it important to follow these safety regulations? Consider this:
>
> - In nonfatal accidents, in-flight turbulence is the leading cause of injuries to airline passengers and flight attendants.
> - Each year, approximately 58 people in the United States are injured by turbulence while not wearing their seat belts.
> - From 1980 through 2008, U.S. air carriers had 234 turbulence accidents*, resulting in 298 serious injuries and three fatalities.
> - Of the 298 serious injuries, 184 involved flight attendants and 114 involved passengers.
> - At least two of the three fatalities involved passengers who were not wearing their seat belts while the seat belt sign was illuminated.
> - Generally, two-thirds of turbulence-related accidents occur at or above 30,000 feet.

The FAA website also contained a video that simulated the effects of turbulence on seated passengers who were wearing and not wearing seat belts. Although the g-forces associated with the simulated event were not available, the video provided a very useful depiction of the benefits of wearing a seat belt and also having other objects restrained in the cabin.[240]

---

[239] Obtained from www.faa.gov/passengers/fly_safe/turbulence (updated 4 August 2009).

[240] The video was available at the following link: www.faa.gov/other_visit/aviation_industry/airline_operators/airline_safety/turbulence/media/cabin_turbulence.wmv

Australian Transport Safety Bureau

In June 2008, the ATSB published an Aviation Safety Bulletin for passengers about keeping safe during turbulence. [241] The document included the following advice for passengers:

> **1. Put your seatbelt on, and keep it fastened when you are seated**
>
> Your seat belt is the best defence against injuries. Keep it fastened low and tight around your waist.
>
> Almost all turbulence injuries involve people who are not properly seated and do not have their seat belt fastened.
>
> When the seat belt sign is on, you are required by law to have your seat belt fastened for your own safety. The pilots or cabin crew will not always have enough time to warn you to put your seat belt on before turbulence hits.
>
> When the seat belt sign is off, you should continue to keep your seat belt fastened. When moving around the cabin to use the restroom facilities and to exercise during long flights, hold on the seat backs as you walk. This will help secure you if the aircraft moves unexpectedly.

[241] ATSB Research and Analysis Report AR-2008-034, *Staying safe against turbulence*. Available from www.atsb.gov.au/media/27791/ar2008034.pdf.

APPENDIX N: SOURCES AND SUBMISSIONS

Sources of information

The sources of information during the investigation included:

- the flight crew, cabin crew and many of the passengers on VH-QPA (QPA) on 7 October 2008
- the flight crew of QPA on 12 September 2008 and VH-QPG on 27 December 2008
- the aircraft operator
- the aircraft manufacturer, the air data inertial reference unit (ADIRU) manufacturer and the seat-belt manufacturer
- the US Federal Aviation Administration, the European Aviation Safety Agency and the Australian Civil Aviation Safety Authority
- the Bureau of Meteorology
- the Australian Department of Defence
- independent experts on single event effects
- independent experts on system safety assessments.

References

Ahmed, S Wallace, KM Blessing, LTM 2003, 'Understanding the differences between how novice and experienced designers approach design tasks', *Research in Engineering Design*, vol. 14, pp. 1-11.

Akerlund, O Bieber, P Boede, E Bozzano, M Bretschneider, M Castel, C Cavallo, A Cifaldi, M Gauthier, J Griffault, A Lisagor, O Lüdtke, A Metge, S Papadopoulos, C Peikenkamp, T Sagaspe, L Seguin, C Trivedi, H Valacca, L 2006, *ISAAC, a framework for integrated safety analysis of functional, geometrical and human aspects*, Proceedings of 3rd European Congress on Embedded Real Rime Systems (ERTS), Toulouse, France.

Austroads 2009, *Non-wearing of adult seat belts in Australia: Where to Next?*, Research Report AP-R346/09.

Bozzano, M Villafiorita, A Akerlund, O Bieber, P Bougnol, C Böde, E Bretschneider, M Cavallo, A Castel, C Cifaldi, M Cimatti, A. Griffault, A Kehren, C Lawrence, B Lüdtke, A Metge, S Papadopoulos, C Passarello, R Peikenkamp, T Persson, P Seguin, C Trotta, L. Valacca, L & Zacco, G 2003, *ESACS: an integrated methodology for design and safety analysis of complex systems*, Proceedings of the European Safety and Reliability Conference (ESREL), Balkema Publishers.

Briere, B, Favre, C & Traverse, P 1995, 'A family of fault-tolerant systems: electrical flight controls, from A320/330/340 to future military transport aircraft', *Microprocessors and Microsystems*, vol. 19, pp. 75-82.

Briere, D & Traverse, P 1993, *Airbus A320/A330/A340 electrical flight controls-a family of fault tolerant systems*, Proceedings of 23rd IEEE International Symposium on Fault Tolerant Computing, pp. 616-623.

Butler, RW 2008, *A primer on architectural level fault tolerance*, Technical Memorandum TM-2008-215108, NASA Langley Research Centre, Hampton, Virginia.

Cross, N 2004, 'Expertise in design: An overview', *Design Studies*, vol. 25, pp. 427-441.

DeWeese, R & Gowdy, V 2002, *Human factors associated with the certification of airplane passenger seats: Seat belt adjustment and release*. Federal Aviation Administration Report DOT/FAA/AM-02/11.

Dyer, C & Lei, F 2001, 'Monte-Carlo calculations of the influence on aircraft radiation environments of structures and solar particle events', *IEEE Transactions on Nuclear Science*, vol. 48, pp. 1987-1995.

Dyer, C Hands, A Ford, K Frydland, A and Truscott, P 2006, Neutron-induced single event effects testing across a wide range of energies and facilities and implications for standards, *IEEE Transactions on Nuclear Science*, vol. 53, pp. 3596-3601.

Dyer, CS Lei, F Clucas, SN Smart, DF and Shea, MA 2003, Solar particle enhancements of single-event effect rates at aircraft altitudes, *IEEE Transactions on Nuclear Science*, vol. 50, pp. 2038-2045.

Favre, C 1994, 'Fly-by-wire for commercial aircraft: The Airbus experience', *International Journal of Control*, vol. 59, pp. 139-157.

Fischhoff, B Slovic, P & Lichtenstein, S 1978, 'Fault trees: Sensitivity of estimated failure probabilities to problem representations', *Journal of Experimental Psychology: Human Perception and Performance*, vol. 3, pp. 330–344.

Flight Safety Foundation 1994, 'Turbulence-related injuries pose continued risk to passengers and cabin crew', *Cabin Crew Safety*, vol. 29.

Flight Safety Foundation 2001, 'Strategies target turbulence-related injuries to flight attendants and passengers' *Cabin Crew Safety*, vol. 36.

Flight Safety Foundation 2006, 'Crew efforts help passengers comprehend safety information', *Cabin Crew Safety*, vol. 39.

Girasek, DC & Olsen, CH 2007, 'Usual seat belt practices reported by airline passengers surveyed in gate areas of a U.S. airport', *Aviation, Space, and Environmental Medicine*, vol. 78, pp. 1050-1054.

Goupil, P 2010, 'Industrial practices in fault tolerant control', *Lecture Notes in Control and Information Sciences*, vol. 399, pp. 157-167.

Goosens, LHJ Cooke, RM Hale, AR & Rodic-Wiersma, Lj 2008, 'Fifteen years of expert judgement at TUDelft', *Safety Science*, vol. 46, pp. 234-244.

Graves, SS & Jacobsen, RA 2010, Verification and validation for flight-critical systems (VVFCS): Summary of responses to solicitation number NNH09ZEA001L, April 2009, Technical Memorandum TM-2010-216715, NASA Langley Research Centre, Hampton, Virginia.

Hanson, RJ 1987, *Conducted electromagnetic transient-induced upset mechanisms: Microprocessor and subsystem level effects*, Proceedings of the Electrical Overstress-Electrostatic Discharge Symposium, pp. 104-109.

Hawkins, R & Kelly, T 2010, *A structured approach to selecting and justifying software safety evidence*, Proceedings of the 5th IET International System safety Conference, Manchester UK.

Health and Safety Commission, 1998, *The use of computers in safety-critical applications*, Final report of the study group on the safety of operational computer systems, UK.

Klein, G 1998, *Sources of power: How people make decisions*, Massachusetts Institute of Technology.

Johansson, K et al. 1998, reported in Hands, A Dyer, CS & Lei, F, 'SEU rates in atmospheric environments: Variations due to cross-section fits and environment models', *IEEE Transactions on Nuclear Science*, vol. 56, pp. 2026-2034.

Joshi, A Heimdahl, MPE Miller, SP & Whalen, MW 2006, *Model-based safety analysis*, Contractor report CR-2006-213953, NASA Langley Research Centre, Hampton, Virginia.

Leveson, N 1995, *Safeware*, Addison-Wesley, Boston.

Leveson, N 2004a, 'The role of software in spacecraft accidents', *Journal of Spacecraft and Rockets*, vol. 41, pp. 564-575.

Leveson, N 2004b, 'A new accident model for engineering safer systems', *Safety Science*, vol. 42, pp. 237-270.

Leveson, N 2009a, *Engineering a safer world: Systems thinking applied to safety*, unpublished manuscript dated July 2009.

Leveson, N 2009b, 'Software challenges in achieving space safety', *Journal of the British Interplanetary Society*, vol. 62.

Lisagor, O McDermid, JA & Pumphrey, DJ 2006, *Toward a practicable process for automated safety analysis*, Proceedings of the 24th International System Safety Conference, Albuquerque, New Mexico.

Lutz, R 1992, *Analyzing software requirements errors in safety-critical, embedded systems*, Proceedings of the IEEE International Conference on Requirements Engineering, IEEE Computer Society Press, pp. 53-65.

Lutz, R 1996, 'Targeting safety-related errors during software requirements analysis', *Journal of Systems and Software*, vol. 34, pp. 223-230.

Manion, M 2007, 'The epistemology of fault tree analysis: an ethical critique', *International Journal of Risk Assessment and Management*, vol. 7, pp. 382-430.

Martin, PL 1999, *Electronic failure analysis handbook*, McGraw-Hill Professional, New York.

Matthews, R 2009, *Targets of additional safety improvements to cabin occupants.* Paper presented at the Annual International Aircraft Cabin Safety Symposium, Torrance California.

NASA 1995, *Formal methods specification and verification guidebook for software and computer systems, Volume 1: Planning and technology insertion*, Washington DC.

Normand, E 1996, 'Single-event effects in avionics', *IEEE Transactions on Nuclear Science*, vol.43, pp.461-474.

Normand, E Vranish, K Sheets, A Stitt, M and Kim, R 2006, Quantifying the double-sided neutron SEU threat, from low energy (thermal) and high energy (>10 MeV) neutrons, *IEEE Transactions on Nuclear Science*, vol. 53, pp. 3587-3595.

Olsen, J Becher, PE Fynbo, PB Raaby, P & Schultz, J 1993, 'Neutron-induced single event upsets in static RAMs observed at 10 km flight altitude', *IEEE Transactions on Nuclear Science*, vol. 40, pp. 74-77.

Papadopoulous, Y McDermid, J Sasse, R & Heiner, G 2001, 'Analysis and synthesis of the behaviour of complex programmable electronic systems in conditions of failure', *Reliability Engineering and System Safety*, vol. 71, pp. 229-247.

Parker, A 2006. *Public attitudes, perceptions and behaviours towards cabin safety communications*, Research and Analysis Report B2004/0238, Australian Transport Safety Bureau.

Pumphrey, D 2001, *Requirements for improving the safety process on complex systems and definition of the tools integration concepts*, Technical Report WP1 for Enhanced Safety Assessment for Complex Systems (ESACS).

Redmill, F 2002, 'Exploring subjectivity in hazard analysis', *Engineering Management Journal*, vol. 12, pp. 139-144.

Rushby, J 1995, *Formal methods and their role in digital systems validation for airborne systems*, Contractor report 4673, NASA Langley Research Centre, Hampton, Virginia.

Saleh, JH Marais, KB Bakolas, E & Cowlagi, RV 2009, 'Highlights from the literature on accident causation and system safety: Review of major ideas, recent contributions, and challenges', *Reliability Engineering and System Safety*, vol. 95, pp. 1105-1116.

Silvera, DH Kardes, FR Harvey, N & Cronley, ML 2005, 'Contextual influences on omission neglect in the fault tree paradigm', *Journal of Consumer Psychology*, vol. 15, pp. 117-126.

Strigini, L 1996, *Engineering judgement in reliability and safety and its limits: What can we learn from research in psychology*, Centre for Software Reliability Technical Report, City University, London.

Storey, N 1996, *Safety-critical computer systems*, Addison-Wesley, Harlow UK.

Traverse, P Lacaze, I & Souyris, J 2004, Airbus fly-by-wire: A total approach to dependability, Proceedings of the 18th Congress of the International Federation for Information Processing, pp. 191-212.

Tribble, AC Miller, SP & Lempia, DL 2004, *Software safety analysis of a flight guidance system*, Contractor report CR-2004-213004, NASA Langley Research Centre, Hampton, Virginia.

Tvaryanas, AP 2003, 'Epidemiology of turbulence-related injuries in airline cabin crew, 1992-2001', *Aviation, Space, and Environmental Medicine*, vol. 74, pp. 970-976.

Submissions

Under Part 4, Division 2 (Investigation Reports), Section 26 of the *Transport Safety Investigation Act 2003* (the Act), the Australian Transport Safety Bureau (ATSB) may provide a draft report, on a confidential basis, to any person whom the ATSB considers appropriate. Section 26 (1) (a) of the Act allows a person receiving a draft report to make submissions to the ATSB about the draft report.

A draft of this report was provided to the:

- Bureau d'Enquêtes et d'Analyses pour la sécurité de l'aviation civile (BEA, France)
- National Transportation Safety Board (NTSB, US)
- aircraft operator
- aircraft manufacturer
- ADIRU manufacturer
- seat belt manufacturer
- flight crew and cabin crew on VH-QPA on the 7 October 2008 occurrence
- Civil Aviation Safety Authority (Australia)
- European Aviation Safety Agency
- Federal Aviation Administration (US)
- Australian Department of Defence.

Submissions were received from the BEA and aircraft manufacturer, the NTSB and ADIRU manufacturer, the operator, and some of the crew members. The submissions were reviewed and where considered appropriate, the text of the report was amended accordingly.

BEA comment

In accordance with International Civil Aviation Organization (ICAO) requirements, if the State conducting the investigation receives comments on the draft final report from other States involved in the investigation, it is required to include the substance of the comments in the final report or, if a State providing the comments desires, append the comments to the final report.

In this case, the BEA requested the following comment relating to a finding in section 6.1 to be appended to the report:

> The ADIRU manufacturer's failure mode effects analysis and other development processes did not identify the data-spike failure mode.
>
> The BEA disagrees with this statement, as it considers that this safety factor refers to a characteristic of an organisation which has the potential to affect future safety. As consequence, this safety factor should be considered as a "safety issue", with a level of associated risk classified as "Minor".

The ATSB did not consider this safety factor to be a safety issue as the available evidence did not enable the investigation to conclude that there was a problem with the process used by the ADIRU manufacturer to conduct the analysis.

In-flight upset - 154 km west of Learmonth, WA, 7 October 2008, VH-QPA , Airbus A330-303