

University of Southampton Research Repository ePrints Soton

Copyright © and Moral Rights for this thesis are retained by the author and/or other copyright owners. A copy can be downloaded for personal non-commercial research or study, without prior permission or charge. This thesis cannot be reproduced or quoted extensively from without first obtaining permission in writing from the copyright holder/s. The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the copyright holders.

When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given e.g.

AUTHOR (year of submission) "Full thesis title", University of Southampton, name of the University School or Department, PhD Thesis, pagination

UNIVERSITY OF SOUTHAMPTON

FACULTY OF ENGINEERING AND THE ENVIRONMENT

Engineering Sciences

**Artificial Intelligence and Mathematical Models for Intelligent Management of
Aircraft Data**

by

Peter Robin Knight

Thesis for degree of Doctor of Philosophy

October 2012



The author receiving the 'Derek George Astridge Safety in Aerospace Award' in 2005

UNIVERSITY OF SOUTHAMPTON
ABSTRACT
FACULTY OF ENGINEERING AND THE ENVIRONMENT
ENGINEERING SCIENCES

Doctor of Philosophy

**ARTIFICIAL INTELLIGENCE AND MATHEMATICAL MODELS FOR INTELLIGENT
MANAGEMENT OF AIRCRAFT DATA**

by Peter Robin Knight

Increasingly, large volumes of aircraft data are being recorded in an effort to adapt aircraft maintenance procedures from being time-based towards condition-based techniques. This study uses techniques of artificial intelligence and develops mathematical models to analyse this data to enable improvements to be made in aircraft management, affordability, availability, airworthiness and performance. In addition, it highlights the need to assess the integrity of data before further analysis and presents the benefits of fusing all relevant data sources together.

The research effort consists of three separate investigations that were undertaken and brought together in order to provide a unified set of methods aimed at providing a safe, reliable, effective and efficient overall procedure. The three investigations are:

1. The management of helicopter Health Usage Monitoring System (HUMS) Condition Indicators (CIs) and their analysis, using a number of techniques, including adaptive thresholds and clustering. These techniques were applied to millions of CI values from Chinook HUMS data.
2. The identification of fixed-wing turbojet engine performance degradation, using anomaly detection techniques, applied to thousands of in-service engine runs from Tornado aircraft.
3. The creation of models to identify unusual aircraft behaviour, such as uncommanded flight control movements. Two Chinook helicopter systems were modelled and the models were applied to over seven hundred in-service flights.

In each case, the existing techniques were directed toward a condition-based maintenance approach, giving improved detection and earlier warning of faults.

CONTENTS	Page
ABSTRACT.....	III
CONTENTS.....	V
DECLARATION OF AUTHORSHIP.....	VII
ACKNOWLEDGEMENTS.....	IX
ACRONYMS AND ABBREVIATIONS.....	XI
CHAPTER 1.....	1
1 INTRODUCTION.....	1
1.1 Background.....	1
1.2 Air Accident Investigation.....	4
1.3 Aircraft Diagnostics.....	6
1.4 Aircraft Prognostics.....	7
1.5 Thesis Overview.....	8
CHAPTER 2.....	11
2 LITERATURE REVIEW.....	11
2.1 Modern Aircraft Data Sources.....	11
2.2 The Need for Data Management.....	21
2.3 Techniques.....	29
2.4 Examples of AI and Modelling Techniques used to Manage Aircraft Data	46
2.5 Summary.....	57
CHAPTER 3.....	59
3 HUMS CI INTELLIGENT HEALTH MANAGEMENT (IHM).....	59
3.1 General.....	59
3.2 Chinook HUMS Data	63
3.3 HUMS CI IHM Architecture	73
3.4 Auditing the Data	74
3.5 Analysis Tools	77
3.6 Decision Process.....	87
3.7 Implementation.....	91
3.8 Assessment.....	96
3.9 HUMS CI IHM Conclusions.....	103
CHAPTER 4.....	105
4 ENGINE PERFORMANCE DEGRADATION MODELS.....	105
4.1 General.....	105
4.2 Data Sourced.....	107
4.3 Data Correction	112

CONTENTS	Page
4.4 ProDAPS Anomaly Modelling	116
4.5 Performance Snapshots	120
4.6 Engine Performance Conclusions	138
CHAPTER 5.....	139
5 IDENTIFYING UNCOMMANDED FLIGHT CONTROL MOVEMENTS	139
5.1 Background – Chinook Control	139
5.2 Review of UFCM incidents.....	142
5.3 Data Analysis to Identify UFCMs.....	143
5.4 UFCM Identification Conclusions	157
CHAPTER 6.....	159
6 CONCLUSIONS	159
APPENDIX A: CHINOOK HUMS INFORMATION	163
A1.1 Chinook Driveshaft Configuration.....	164
A1.2 Schematic of Chinook Component and Accelerometer Locations	165
A1.3 Monitored Components.....	166
A1.4 Accelerometer Locations	167
A1.5 Passbands.....	168
A1.6 Bearings and Gear Terminology	169
A1.7 Data Acquisition and Labelling	170
A1.8 Chinook Health Indicators	171
A1.9 GenHUMS Analysis	173
A1.10 Chinook HUMS CI IHM Results Charts	177
APPENDIX B: TORNADO EHUMS INFORMATION	201
A1.11 Corrections to ISA Conditions	201
A1.12 Tornado EHUMS Anomaly Detection Results Table	202
A1.13 Tornado EHUMS Anomaly Detection Results Charts.....	208
REFERENCES.....	229

DECLARATION OF AUTHORSHIP

I, *Peter Robin Knight*, declare that the thesis entitled '*Artificial Intelligence and Mathematical Models for Intelligent Management of Aircraft Data*' and the work presented in the thesis are both my own, and have been generated by me as the result of my own original research. I confirm that:

- this work was done mainly while in candidature for a research degree at this University;
- where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
- where I have consulted the published work of others, this is always clearly attributed;
- where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
- I have acknowledged all main sources of help;
- where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
- parts of this work have been published as:
 1. P. R. Knight, J. Cook, and H. Azzam, "Intelligent Management of Helicopter Health and Usage Management Systems Data," *Proceedings of the Institute of Mechanical Engineers - Part G: Journal of Aerospace Engineering*, vol. 219, pp. 507-524, 2005.
 2. H. Azzam, J. Cook, P. R. Knight, and E. Moses, "FUMS™ Fusion and Decision Support for Intelligent Management of Aircraft Data," *Proceedings of IEEE Aerospace Conference*, 2006.
 3. N. H. Wakefield, P. R. Knight, K. P. J. Bryant, and H. Azzam, "FUMS™ Artificial Intelligence Technologies Including Fuzzy Logic For Automatic Decision Making," *Annual Meeting of the North-American-Fuzzy-Information-Processing-Society*, 26-28 June 2005.
 4. P. R. Knight, H. Azzam, S. J. Newman, and A. J. Chipperfield, "Artificial Intelligence and Mathematical Models for Intelligent Management of Aircraft Data," *International Conference on Prognostics and Health Management*, 2008.

Signed:

Date:.....

ACKNOWLEDGEMENTS

The author wishes to thank Dr Simon Newman, Dr Andrew Chipperfield and Dr Hesham Azzam for their guidance and supervision throughout this study. This study would not have been possible without the support of the UK MOD who supplied data, information and funds, and GE Aviation who sponsored the study. In addition, the author would like to thank colleagues at GE Aviation for their support, and family members for proof-reading the thesis.

Dedicated to my family: Alice, Aravis and Seth.

With thanks to God for his faithful love and grace.

ACRONYMS AND ABBREVIATIONS

AC	Aircraft
ACT	Actuator
ADC	Automatic Data Correction
ADR	Accident Data Recorder
AE	Acoustic Emission
AFCS	Automatic Flight Control System
AI	Artificial Intelligence
AI	Attitude Indicator
AJACT	Nozzle Shroud Position
ALT	Altitude
ANN	Artificial Neural Network
ART	Adaptive Resonance Theory
ASH	Aircraft Structural Health
AUIG	Auto Ignition
AURA	Advanced Uncertainty Reasoning Architecture
AUW	All Up Weight
BKTSIN	Reverse Buckets Deployed
BLOB	Binary Large Object
BSD	Bulk Storage Device
BT	Bayesian Theory
CAA	Civil Aviation Authority
CB	Combiner Bearing
CBM	Condition-Based Maintenance
CI	Condition Indicator
COR	Corrected
CPT	Control Position Transducer
CSMU	Crash Survivable Memory Unit
CVM	Comparative Vacuum Monitoring
DA	DASH Actuator
DAME	Distributed Aircraft Maintenance Environment
DAPAT	Diagnostic And Prognostic Analysis Tool
DARPA	Defence Advanced Research Projects Agency
DASH	Differential Air Speed Hold
DEF STAN	Defence Standard
DG	Directional Gyro
DMEMB	Degree of Membership
DREA	Defence Research Establishment Atlantic
DSS	Distributed Sensor Systems
DSS	Decision Support System
DS&S	Data Systems and Solutions
DT	Decision Trees
EASA	European Aviation Safety Agency
EB	Energy Band
EDMS	Engine Distress Monitoring System
EFDC	Early Failure Detection Cells – MOD engine oil debris analysis groups
EFH	Engine Flying Hours

EHM	Engine Health Management
EHUMS	Engine Health Usage Monitoring System
EN	Engine
EPSRE	Engineering and Physical Sciences Research Council
ES	Expert Systems
ESA	Enhanced signal
ETE	Enveloped signal
EUMS	Engine Usage Monitoring System
FAA	Federal Aviation Administration
FADEC	Full Authority Digital Engine Control
FDR	Flight Data Recorder
FEA	Finite Element Analysis
FFT	Fast Fourier Transform
FL	Fuzzy Logic
FMECA	Failure Modes Effects and Criticality Analysis
FOD	Foreign Object Damage
FOQA	Flight Operations Quality Assurance
FREEZE	Reheat Fail or Freeze Indication
FS	Fitness Score
FT	Feet/Foot
FTE	Envelope Spectrum
FUACT	Reheat Fuel Actuator Position
FUMS™	Fleet and Usage Management System
FWD	Forward
G	Gravity (or Normal Acceleration)
GA	Genetic Algorithm
GE	General Electric
GFBAL1	Weapons Fire
GMM	Gaussian Mixture Models
GPA	Gas Path Analysis
GSS	Ground Support System
HARDEN	Reheat Enabled
HI	Health Indicator
HJ	Health Judgements
HMI	Human Machine Interface
HMM	Helicopter Mathematical Model
HOMP	Helicopter Operations Monitoring Programme
HPNGV	High Pressure Nozzle Guide Vanes
HRS	HUMS Recording Session
HUMS	Health and Usage Monitoring System
Hz	Hertz
IAS	Indicated Air Speed
ICD	Interface Control Documents
IDMS	Ingested Debris Monitoring System
IEEE	Institute of Electrical and Electronics Engineers
IHM	Intelligent Health Management
ILCA	Integrated Lower Control Actuator
IMATE	Intelligent Maintenance Advisor for Turbine Engines

IMD	Integrated Mechanical Diagnostics
IMP	Impulsiveness
IP	Information Processor
IPT	Integrated Project Team
ISA	International Standard Atmosphere
ISRD	In-Service Reliability Demonstration
JSF	Joint Strike Fighter (F-25/Lightning II)
K	Kelvin or Kilo (1000)
KTS	Knots (nautical mile per hour)
L	Left
LC	Lane in Control
LCC	Life Cycle Costs
LCT	Longitudinal Cyclic Trim
LCTA	Longitudinal Cyclic Trim Actuator
LFAULT	Reheat Dormant Fail
LITS	Logistic Information Technology System
LL	Log Likelihood
LOG	Logistics
LP	Low Pressure
LPCDEM	Lectric Pressure Control Demand (controls fuel flow – Tornado EHUMS)
M6	Normalised 6 th statistical moment
MAAAP	Management, Affordability, Availability, Airworthiness and Performance
Max	Maximum
Mb	Megabyte
MC	Markov Chains
MDAT	Mechanical Diagnostics Analysis Toolkit
MET	Metrics Evaluation Tool
Min	Minimum
MN	Mean
MOD	Ministry of Defence
MODSS	Multi-Objective Decision Support System
MooN	M out of N
MPPBSD	Maximum Permitted Period Between Successful Downloads
MRC	Max Rate of Change
MTBSF	Mean Time Between System Failure
MWO	Maintenance Work Order
NASA	National Aeronautics and Space Administration
NDT	Non-Destructive Testing
NH	High pressure shaft speed (or NHV)
NL	Low pressure shaft speed (or NLV)
null	Invalid data point
OAT	Outside Air Temperature
OBPHM	Off-Board Prognostic Health Management
ODM	Oil Debris Monitor
OLM	Operational Loads Monitoring
ONFTSW	Open Nozzle for Taxi Enable
OPNBAR	IP BOV Solenoid

OSD	Operational Service Dates
ρ	Density
P	Pressure
P	Port
PA	Probability of Anomaly
PASVS	Total Pressure
PAVS	Static Pressure
PCA	Principal Component Analysis
PDSM	Prognostics Decision Support Mechanism
PHM	Prognostic Health Management
PK	Peak
PLA	Pilot Lever Angle
PLBA	Reheat Reset (by Pilot) Enabled
PLRH	Pilot's Lever Angle
POD	Probability Of Detection
POFA	Probability Of False Alarm
PP	Peak-to-Peak
ProDAPS	Probabilistic Diagnostic and Prognostic System
PZT	Piezoelectric Transducers
R	Right
RIPS	RB199 Installed Performance System
RMS	Root Mean Squared
ROC	Receiver Operating Characteristic
RPM	Revolutions Per Minute
RTB	Rotor Track and Balance
S	Starboard
SHM	Structural Health Monitoring
SIG	Original Signal
SO	Shaft Order
SOA	Spectral Oil Analysis
SPHM	Structural Prognostic Heal Monitoring
STD	Standard Deviation
SVD	Singular Value Decomposition
SVM	Support Vector Machines
T	Temperature
T1	Outside air temperature/ Intake Total Temperature
TBM	Time-Based Maintenance
TBT	Turbine Blade Temperature
TBT SWT	TBT Datum Switch
TTACT	Tt1 Actuator Drive Position
UFCM	Uncommanded Flight Control Movements
UI	User Interface
UK	United Kingdom
US	United States (of America)
VAU	Vehicle Airborne Unit
VG	Vertical Gyro
VRML	Virtual Reality Modelling Language
WLOOL	Working Line Limit (WLL) Exceedance

WOW	Weight On Wheels
WRAM	Work Records and Asset Management
WT	Wavelet Transform
XDRV	Gearbox Cross Drive Engaged/Disengaged
XMSN	Transmission

CHAPTER 1

1 INTRODUCTION

1.1 Background

Man has long desired to leave the ground and take flight like a bird, both for the pure emotional pleasure of it or, in the recent past, to provide rapid and convenient transport. In Greek mythology, Icarus was able to fly, wearing wings fashioned from wax and feathers, but flew too near to the sun whereupon the wax melted and he fell into the sea and perished. Although Icarus lost his battle against gravity, eventually mankind was able to find a way to overcome it; until something goes wrong.

On the 17th of December 1903, brothers Orville and Wilber Wright, two bicycle mechanics-turned-inventors, took turns (Orville was the first) to pilot the first powered aeroplane twenty feet above a wind-swept beach in North Carolina, Figure 1.1. Almost inevitably, air accidents soon followed. The first fatal powered aircraft accident was of the Wright Flyer III, in 1908: One of the two propellers separated in flight, tearing loose the wires bracing the rudder and causing the loss of control of the aircraft. Orville Wright, the pilot, broke his leg, pelvis and ribs and injured his back. Lieutenant Tom Selfridge, his passenger, suffered a crushed skull and died a short time later, [1] [2]. Clearly, this should not have happened, but it did, possibly because of ignorance, a weakness in the materials used or simply human mistakes.




Figure 1.1: The Wright brothers' first flight, [2].

The word 'ignorance' may seem unfair, however, aviation was a brand new science then and many factors, which are now well known, had not yet been

encountered and were therefore not considered. One person died, which in statistical terms may not be significant, but for the family of Lieutenant Selfridge it was nothing short of a tragedy.

The worst accident to date, caused by mechanical failure, was a Japan Airlines Boeing 747 that crashed on Mt. Osutaka, Japan, on August the 12th 1985. All but four of the 524 people onboard were killed. Improper repairs, after a tail strike seven years earlier, led to a rupture of the pressure bulkhead and loss of all controls, [3]. Obviously this should not have happened, but it did with colossal ramifications.

The consequences of accidents can be very severe, not only in fatalities but also in the grounding of whole fleets of aircraft. For example, on the 25th July 2000, an Air France Concorde caught fire shortly after takeoff from Paris' Charles de Gaulle Airport on a charter flight to New York, Figure 1.2. The pilots lost control and the plane crashed into a hotel restaurant. Subsequent investigations revealed that the aircraft had run over a metal strip which had earlier dropped from a Continental Airlines DC-10 during its departure roll. This caused the tires on the Concorde to explode and puncture the underwing fuel tanks. All 109 passengers and five people on the ground died and, as a result, all Concorde were removed from service, [4]. Concorde was the result of many years of research, design and engineering development. It pushed the boundaries of supersonic flight further than it had ever gone in civil aviation, and in one fateful day it disappeared forever.



Figure 1.2: Concorde catching fire on takeoff, 2000, [5].

From the relatively simple Wright Flyer we now have aircraft which are extremely complex engineering systems. A human is not capable of continually monitoring all of the data in real-time and determining the integrity

of the vehicle or prognosis of failure and accident so sophisticated automated methods have been developed.

As will be described below, the focus of attention regarding flight safety has changed. Initially, all that could happen would be to visit the accident site, evaluate the wreckage, sweep up the bits and try to learn from what had happened and to try to prevent such an incident ever happening again. This was of small consolation to either the people directly involved in the accident or those left behind to mourn the loss. Accidents are almost always a consequence of several separate incidents colluding to cause the eventual, and often catastrophic, final accident. It is almost impossible to predict such a confluence of events so, if the loss of an aircraft is to be avoided, mechanisms have to be put in place to alert the pilot or ground crew of a potential accident arising. This may be a structural problem, an aerodynamic limitation about to be exceeded or, indeed, an error of judgement. Such monitoring is becoming the norm and has been in place in both commercial and military aviation for some time. Sensors placed in specific locations monitor the performance of critical items and alert the pilot to a potential failure. This requires the characteristics of the aircraft to be well established in order to provide a yardstick from which to base any decision to raise an alert. There will often be several different sensor systems, allowing for the failure of one system to be highlighted and a correct decision to be made. False alarms are detrimental to efficient operation and minimising of operating cost, however, missing a vital change in aircraft component behaviour is an extreme situation to be avoided at all cost.

The way in which a sensor being triggered is acted upon is now being revisited. Traditionally, an alert would mean landing as quickly as possible and grounding the aircraft, or in certain circumstances, the entire fleet. This has profound consequences for the efficiency of the airline and can endanger the goodwill of the paying passenger. In a military scenario, this would mean the removal of a vital aircraft from an operation.

If a move from a diagnostic procedure, where the sensors monitor what *has happened*, can be made to a prognostic method where the sensor can make a prediction when it is *likely to happen*, then safety can be retained, with the

correction of the problem made at a convenient time in the aircraft's schedule or permit the military operation in question to continue without any interruption.

A final comment is that, because of the complex nature of modern aircraft, data produced by the sensors fitted to an aircraft fleet will be considerable in magnitude. The question now facing a system designer is whether you can "see the wood for the trees". The management of this enormous quantity of data, in order to transform it into useful information, is the basis of this thesis.

1.2 Air Accident Investigation

There are moral, commercial and political pressures to reduce the number and severity of accidents because of the cost, both in terms of money and human life. To do this, accidents are investigated to see whether any lessons can be learnt to prevent them in the future.

The roots of the UK Air Accidents Investigation Branch (AAIB) date back to 1915. Its stated purpose is: "*To improve aviation safety by determining the causes of air accidents and serious incidents and making safety recommendations intended to prevent recurrence.*", [6]. After an accident investigation, the AAIB issues a report which may include a number of recommendations. These recommendations are compiled to define the airworthiness requirements for an aircraft to be safe to fly, and so receive airworthiness certification. Figure 1.3 shows the number of AAIB recommendations between 2004 and 2009 and the number of resulting European Aviation Safety Agency (EASA) final recommendations. As more accidents occur, these requirements become more stringent with the aim of making aviation safer. The airworthiness requirements ensure that an aircraft has been designed and manufactured to international standards, has a maintenance programme by which the aircraft can be serviced, checked and repaired and has a continued airworthiness programme that supports the aircraft in operation, [7].

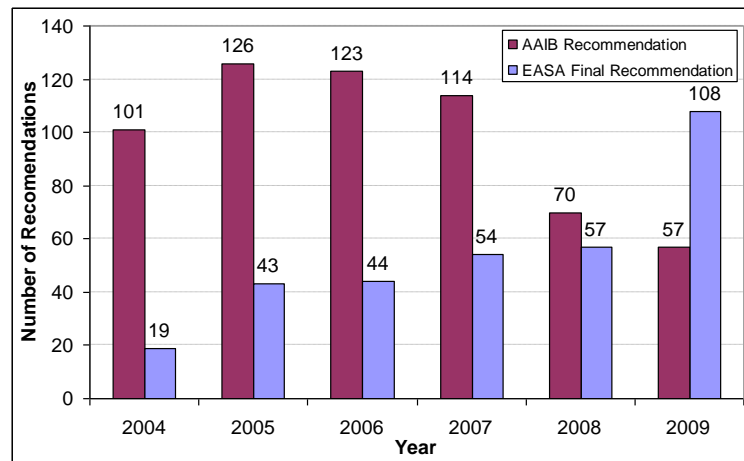


Figure 1.3: AAIB and EASA recommendations per year, [8, 9]

Figure 1.4 depicts the cycle of aircraft crashes or incidents leading to an investigation, which may make recommendations for airworthiness, which in turn affects aircraft design and maintenance policies.

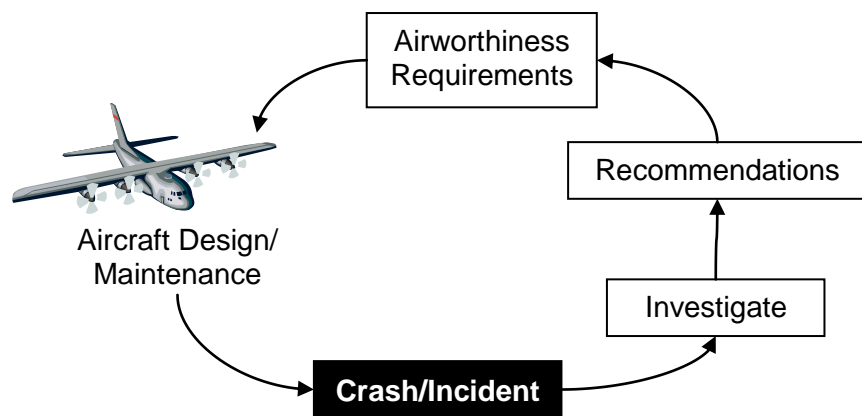


Figure 1.4: Schematic of the crash-investigation cycle

Organisations and laws have been set up in order to enforce adherence to airworthiness regulations. In the United States, the Air Commerce Act became law in 1926 to ensure civil air safety, which led to the Federal Aviation Administration (FAA) being formed in 1958. In the UK, aircraft safety was the responsibility of the Department of Transport from 1919 until the formation of the Civil Aviation Authority (CAA) in 1972.

Despite airworthiness requirement revisions and updates, the number of fatal accidents per flying hour for US airlines has not noticeably decreased in the last twenty years (Figure 1.5). The accident rate has, however, slightly improved over the last ten years, despite the rise in flight hours from about 12 million to nearly 20 million. The percentage of worldwide fatal accidents,

caused by mechanical failures, has not decreased over the past fifty years (Figure 1.6). The figure represents 1,300 fatal accidents, involving commercial aircraft, for which a specific cause is known. Aircraft with ten or less people onboard, military aircraft, private aircraft and helicopters are not included.




Figure 1.5: US airline accidents per 10,000 flight hours, [10]




Figure 1.6: Percentage of world-wide fatal aircraft accidents, caused by mechanical failure, by decade, [11]

1.3 Aircraft Diagnostics

In order to improve safety, designers and airworthiness authorities have been trying to move towards systems that give a 'red light' indication of mechanical faults before they develop and cause accidents. This would allow the pilot to make a safe landing and the fault could be fixed without loss of life.

In order to move to a 'red light' system, the process for designing, operating and maintaining aircraft has to be changed. Traditional aircraft design has

been based either on the “fail-safe” principle (whereby any single failure will permit the total structure to remain “safe” and will not lead to a catastrophe), or on the “safe-life” principle (whereby those components which cannot be designed to be “fail-safe” will have a fixed life, based either on theoretical design principles or testing to destruction). The assigned life-limit is based on the worst anticipated operational conditions and will therefore be excessively conservative (i.e. expensive) for aircraft operating under more normal conditions. Under the ‘red light’ system, onboard sensors can be used to track the actual usage of individual aircraft and components, which can then be used to assess the current condition of the aircraft. Hence, aircraft design and maintenance is transitioning from the current Time-Based Maintenance (TBM), where aircraft are inspected/maintained at regular intervals, to Condition-Based Maintenance (CBM), where maintenance is only performed when it is triggered by, for example, measured usage or degraded performance.

1.4 Aircraft Prognostics

The next progression in the development of flight safety is towards a ‘timeline’ based system, i.e., to make a prognosis of when a current fault may lead to a failure. The more warning that can be given, the cheaper it is to take corrective action (for example, maintenance can be carried out during scheduled down time rather than aborting flights). In order for such a methodology to operate, data analysis algorithms need to produce a fault propagation timeline, with associated uncertainty (Figure 1.7).




Figure 1.7: Hypothetical percentage system health prediction over time showing uncertainty and potential warning times.

Thus, three system types for airworthiness improvement can be described as:

- **Crash/Incident** – no warning of failure; the aircraft or its debris is investigated to try to prevent re-occurrence.
- **Red Light** – faults are diagnosed and the systems give warning of imminent failure, avoiding the crash or incident, but corrective action is still expensive.
- **Timeline** – the propagation of faults to failure is predicted to give a prognosis of when corrective action needs to be taken, when convenient, thus reducing costs and avoiding incidents. This could be described as an orange light, indicating that a failure mode has been detected.

To realise this move to CBM, and the desired improvements in aircraft Management, Affordability (reduced cost), Availability, Airworthiness (safety) and Performance (MAAAP), the data should be intelligently managed.

These are the items on the wish list and the thesis describes a research effort on three fronts, to provide proven methods that can be used on aircraft, both commercial and military, to furnish these aspirations.

1.5 Thesis Overview

This thesis asks the question “can an improved management and analysis of aircraft data lead to an improved detection of faults?” To investigate this question, and to determine which types of analysis are most effective in identifying potential fault conditions, a research programme was undertaken making use of Artificial Intelligence (AI) and mathematical models to develop novel algorithms to manage and analyse aircraft data for improved MAAAP. These developed algorithms were validated against large, in-service aircraft datasets, and have led to the definition of those key elements required in a framework for managing aircraft data.

A survey of both current practice and state-of-the-art aircraft data recording and monitoring techniques was undertaken and is presented in Chapter 2. The review includes an introduction to the typical data sets available for modern aircraft and the challenges faced when analysing this data.

As already highlighted, three separate, but related, investigations were undertaken during the course of this research. In each, the aim was to improve the current system from crash/incident investigation to an automated red-light warning system and, if possible, to a timeline giving early warning of failures.

A summary of each investigation follows:

- 1) **HUMS CI Intelligent Health Management:** The management of helicopter Health and Usage Monitoring System (HUMS) Condition Indicators (CIs) and their analysis using a number of techniques, including data reduction by Principal Component Analysis (PCA), adaptive thresholds, prognostic trends and clustering. The MOD Chinook fleet was used as an example. The HUM system was installed on UK MOD Chinook helicopters in the late nineties and offers a red-light capability to warn of imminent failure. The work carried out by this study aims to move this towards the timeline prognosis, reduce the number of false alarms, and to analyse a wider number of components and data signals than the current system. The initial findings of the analysis were published in the Journal of Aerospace Engineering in 2005, [12], which was awarded the 'Derek George Astridge Safety in Aerospace' award. This work is presented in Chapter 3.
- 2) **Aircraft Turbojet Engine Performance Degradation Models:** The identification of engine performance degradation, using anomaly detection techniques, applied to in-service data from Tornado aircraft engines. Current analysis involves manual data recording and analysis that may result in a red-light warning to ground an aircraft. This study aims to move this to an automated system that identifies failures much earlier. The preliminary analysis was presented at the Institute of Electrical and Electronics Engineers (IEEE) aerospace conference in 2006, [13]. This work is presented in Chapter 4.
- 3) **UFCM Identification:** The use of linear models to identify unusual aircraft behaviour, such as Uncommanded Flight Control Movements (UFCM), applied, as an example, to Chinook helicopters. These occurrences are currently only analysed after the incident has occurred. The work carried

out by this study aims to move towards a red-light warning before the event, and to identify events that are currently missed. This work is presented in Chapter 5.

These three approaches cover different aspects of aircraft operation which could contribute to an accident occurring. They cover both fixed and rotary wing airframes and engines. Two are focussed on engineering issues; the third examines the influence of the pilot. In this way, the ability of the methods, developed in this thesis, to perform in differing situations and under differing influences provides a solid justification for their adoption in the future.

An overview of the work presented in this thesis was also published in the first International Conference on Prognostics and Health Management in 2008, [14]. The thesis was undertaken while the author was in the employment of GE Aviation Systems.

CHAPTER 2

2 LITERATURE REVIEW

A survey of current practice and state-of-the-art aircraft data recording and monitoring techniques is presented in this Chapter. Common sources of modern aircraft data are described and the need for this data to be intelligently managed and analysed is presented. A brief introduction to AI and modelling techniques is also presented. Finally, examples from the literature of managing and analysing aircraft data for improved MAAAP are discussed.

2.1 Modern Aircraft Data Sources

Since the study aims to manage and analyse aircraft data, this section identifies the most important data sources currently recorded for modern aircraft that may require intelligent management and analysis. The following data sources are described:

- Flight Data Recorders (FDR)
- Health and Usage Monitoring Systems (HUMS)
- Operational Loads Monitoring (OLM)
- Engine Health Management (EHM)
- Structural Health Monitoring (SHM)
- Debris Monitoring
- Maintenance Databases

2.1.1 *Flight Data Recorders (FDR)*

Flight Data Recorders (FDR) (also ADR, for Accident Data Recorders) have been used for many years to record flight parameters such as altitude, airspeed, control positions and accelerations. The data is typically recorded between 16 Hertz (Hz) (for parameters that vary quickly, like roll-rate) and 0.25 Hz (for parameters that don't change very often, such as 'weight on wheels' flags). They may also record cockpit voice, radio and background sounds. The first FDRs, developed in 1939, used light shining on photographic film. However, it was not until the 1950s, spurred on by the grounding of the fleet of De Havilland DH106 Comets after a series of accidents, that FDR units, as we know today, were developed, [15, 16]. FDR data are primarily

used for air accident/incident investigation; further use of FDR data to improve MAAAP would be invaluable. FDR data is also used for Flight Operations Quality Assurance (FOQA) applications, such as the Helicopter Operations Monitoring Programme (HOMP), [17, 18], which monitors flight operations by routinely analysing aircraft flight data to detect deviations from normal, expected, or flight manual practice. They provide continuous operational quality control with timely feedback on sub-standard practices, and produce valuable information for the evaluation and improvement of operating procedures.

White and Vaughan discussed how fleet usage monitoring, using FDR data, is essential in improving aging US Army helicopter safety, availability and affordability, [19]. The remaining life of components are determined using assumed worst-case flight profiles across the fleet. The actual flight profiles will vary widely from these assumptions. FDR data can be used for 'flight regime recognition', which can be used to determine the remaining life of components more accurately, improving safety and reducing costs.

2.1.2 Health and Usage Monitoring Systems (HUMS)

Since the 1990s, helicopters have been fitted with Health and Usage Monitoring Systems (HUMS), which particularly focus on measuring and recording vibration of gears and bearings, and on performing Rotor Track and Balance (RTB), along with recording a range of other health and usage data. In fact, HUMS has been mandated for retro-fit to all United Kingdom (UK) Ministry of Defence (MOD) rotary wing aircraft capable of carrying nine or more passengers, with Operational Service Dates (OSD) past 2010, and to all new helicopter procurements, [20]. Typically HUMS vibration data is acquired for a short period at very high frequency (e.g. 100 kHz) at one or more points during a flight. The data is usually summarised by the airborne unit using signal processing and statistics to derive CIs.

Liu and Pines, of the University of Maryland, have analysed US civil rotorcraft accidents caused by vehicle failure or malfunction, between 1998 and 2004, [21]. Figure 2.1 shows a pie chart of the results of their analysis, giving the number and percentage of accidents for each cause.




Figure 2.1: Causes for vehicle factor related US civil rotorcraft accidents (1998-2004), [21]

The study reached the following conclusions:

1. The average ratio of the number of accidents caused by vehicle factors to the total number of accidents, during the period of 1998 to 2004, has been reduced by more than one half, in comparison to the accident data examined between 1963 and 1997.
2. Engine failure or malfunction remains the primary reason for vehicle factor related accidents.
3. The turbine engine has a potential vulnerability to mechanical failures.
4. Human errors account for almost one in five vehicle failure or malfunction related accidents (design, manufacture, maintenance or operation).

To improve the performance of the next generation of integrated mechanical health and usage monitoring systems, the study made the following recommendations:

1. The monitoring of the mechanical properties and thermal conditions of the turbine, the conditions of the bearing and the adaptor/coupling in a turbine engine deserve HUMS researchers' focus.

2. Clutch and coupling are two components in the drive-train which are susceptible to failure. These components have not received enough attention.
3. The diagnostic methodologies of the main rotor, as well as tail rotor fatigue fracture, should be pursued.
4. Tail rotor drive shaft hanger bearings need to be considered carefully in a HUMS system.
5. Better methodologies and systems should be developed for the monitoring of hydraulic, fuel and oil systems. In particular, a reliable fuel monitoring and fuel starvation alarm system is needed.

Draper, of the UK MOD HUMS Integrated Project Team (IPT), discussed the operational benefits of using HUMS on helicopters, [22]. The Chinook GenHUMS completed an In-Service Reliability Demonstration (ISRD), which found a system reliability of over five thousand flying hours, significantly above the reliability requirement for the aircraft to have a 99% probability of successfully completing an eight flying hour working day, without experiencing a system failure (equating to a minimum Mean Time Between System Failure (MTBSF) of about eight hundred flying hours). Use of HUMS has reduced the time to diagnose and rectify Uncommanded Flying Control Movements (UFCM) by 75% and prevents unnecessary removal and replacement of serviceable flight control components from the aircraft. Overall, HUMS has delivered cost savings, resulting from preventing incidents and accidents, equating to approx £4.6m per year. Having outlined these benefits, Draper highlighted that there was the potential for greater improvements by further developing the management of the HUMS data.

Land *et al* also studied the potential savings of HUMS and concluded that it both increases vehicle reliability and safety while significantly reducing maintenance and insurance costs, [23].

2.1.3 Operational Loads Monitoring (OLM)

Traditionally, fatigue damage to critical components was estimated solely using flight time and β -factors (damage, measured in terms of cycles, per hour). However, analysis of in-service flight data shows that there is often little correlation between damage and flying hours. Figure 2.2 gives an example of

in-service flight data showing how there can be little correlation between damage and engine flying hours (EFH).




Figure 2.2: Showing that critical area life consumption (number of cycles) is not well correlated with flight time, [24]

To improve the estimation of fatigue damage, UK military jets started to use a single fatigue load meter on each aircraft, which counts time spent in various normal acceleration (G) bands. Fatigue meters give a better measure of how much damage is likely to have been induced on critical components in each flight. DEF STAN 00-970 section 3.2.24, issued in 1983, stated that “*provision shall be made in every aeroplane for the installation of an RAE fatigue load meter and one or more switches to start and stop the meter automatically in flight.*”, [25].

Operational Loads Monitoring (OLM) programmes give a much improved measure of damage on particular components. A large number of strain gauges (e.g. a legacy UK MOD combat aircraft had about fifty) are used to record the strains on key components during normal operations of a small fleet sample. The fatigue lives of the fleet sample components are calculated from the strain measurements. The fatigue lives of the rest of the fleet are estimated by extrapolating the sample results using recorded parametric data. For example, the Canadian Air Force CC130 fleet has an OLM Individual Aircraft Tracking (IAT) programme and a Data Analysis System (DAS), which was developed to process and analyse both parametric and strain data and assess usage severity, [26]. However, OLM is prohibitively expensive to apply to a whole fleet and comes with its own problems of calibration (especially in

different configurations such as with/without weapons) and signal data corruption, [27].

An emerging technology, developed in the Fleet and Usage Management System (FUMS™) (originally developed by MJA Dynamics, then Smiths Aerospace and now GE Aviation Systems), [28], uses models, combining a number of techniques including Artificial Neural Networks (ANNs), trained using OLM data from one aircraft, to predict the damage across the fleet using FDR data that is always recorded for air accident investigations. These methods have been shown to be very accurate and much cheaper than monitoring strain across a large fleet, [29]. For example, Figure 2.3 shows how well the blind predication of fatigue over three flights, including sortie profiles and configurations not in the training data, match the strain gauge measured values. The author was part of a team that updated DEF STAN 00-970, [30] and the associated guidance material to cover the use and qualification of non-adaptive prediction methods, such as ANNs, to monitor critical component fatigue for airworthiness approval.

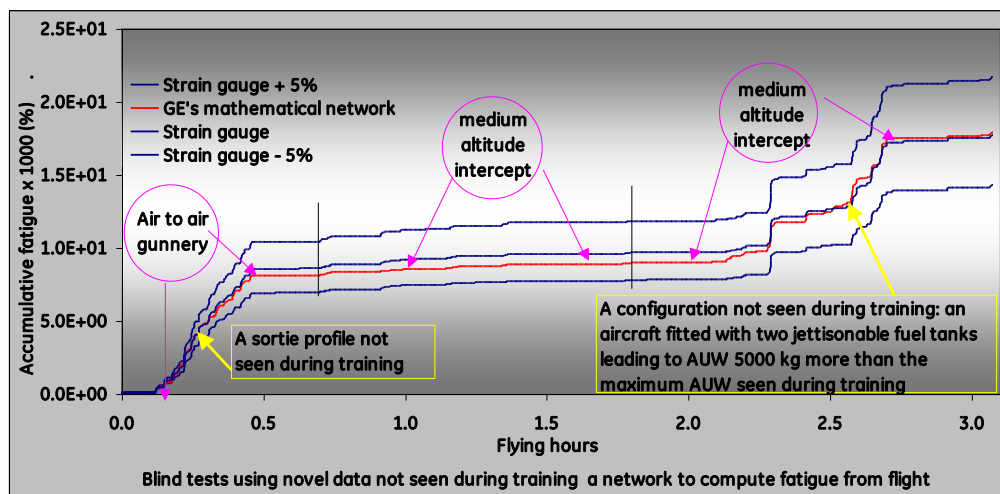


Figure 2.3: Blind prediction of military aircraft fatigue, [29]

2.1.4 Engine Health Management (EHM)

EHM and Engine Usage Monitoring Systems (EUMS) are increasingly being used in both civil and military aircraft to record spool speeds, temperatures, pressures, vibration and usage of turbine engines.

Jaw reviewed recent advancements in aircraft EHM technologies and made recommendations for the future, [31]. He sums up EHM progress as:

“The traditional Engine Health Management approach uses fleet statistical data and signal processing techniques to detect and isolate faults. Modern EHM approaches enhance the traditional approach with physics-based models, individual engine performance tracking, predictive algorithms and decision support capabilities.” [31]

The following EHM functional areas can be identified:

1. Gas path performance monitoring.
2. Oil and debris monitoring (considered in section 2.1.6).
3. Vibration monitoring.
4. Usage and life monitoring.

The major requirements of EHM are:

1. Automated monitoring, analysis and decision support.
2. Accurate results with high confidence.
3. Robust capabilities against noise and faulty information.
4. Wide coverage of fault conditions.
5. Predictive capabilities.
6. Using existing, or as few as possible, sensing instruments.
7. Flexible, modular and open architecture.
8. User friendliness.

EHM systems are used to identify “engine component problems, commonly called faults or failures, including: erosion, corrosion, fouling, dirt build-up, foreign object damage (FOD), worn seals, excessive tip clearance, burned or warped turbine stator or rotor blades, partially or wholly missing blades, plugged fuel nozzles, rotor disk or blade cracks induced by fatigue or operation outside normal intended limits, etc.” [31].

Jaw’s review [31] shows how useful EHM can be in identifying a wide range of engine faults. The requirements listed above can be carried across to any health monitoring program and can form a good basis for the design of more intelligent systems.

2.1.5 Structural Health Monitoring (SHM)

Increasingly, dedicated systems are fitted to aircraft to monitor the health of structures. For example, an array of Acoustic Emission (AE) sensors can be used to detect the sounds of cracks developing in both metal and composite structures. The arrival times at each sensor can be used to triangulate the exact location of the crack, even if it is not visible to inspection. Examples of the types of degradation measured are: corrosion, fatigue, damage, delamination/disbanding, etc., Figure 2.4.



Figure 2.4: Examples of disbond, crack and corrosion damage, [32]

Ever improving sensing and processing technologies are being used to assess the health of structures. For example, Roach and Rackow, of the FAA Airworthiness Assurance Centre in the USA, have used Distributed Sensor Systems (DSS) for Structural Prognostic Health Monitoring (SPHM), [32]. The DSS include the use of the following in-situ sensors: Comparative Vacuum Monitoring (CVM), Piezoelectric Transducers (PZT), fibre optics and remote field eddy currents. Figure 2.5 shows example outputs from PZT sensors. The white dots indicate the position of the sensors, and the colours indicate areas of crack damage after different numbers of cycles. These sensors are combined to create distributed networks to monitor large panels.




Figure 2.5: Colour-coded PZT heat-map images showing crack growth, [32]

2.1.6 Debris Monitoring

As bearings and gears wear, crack, fatigue or are damaged, debris is deposited in the oil that is used as a lubricant. This debris can be collected in oil filters, magnetic plugs, dedicated Oil Debris Monitors (ODMs), or detected by performing Spectral Oil Analysis (SOA) on a sample of the oil to find element concentrations. Debris analysis has been in use since the early 1980s, when the Canadian Defence Research Establishment Atlantic (DREA) developed it to try to evaluate the condition of their Sea King helicopter gearboxes, which had been plagued with main gearbox problems, [33, 34].

Dempsey has used oil debris analysis to investigate tapered roller bearing damage, [35]. Figure 2.6 shows an example of how the debris mass can relate closely to the actual damage occurring. The figure shows how little debris is accumulated over the majority of the running time and that the debris count increases rapidly once the damage starts to progress.

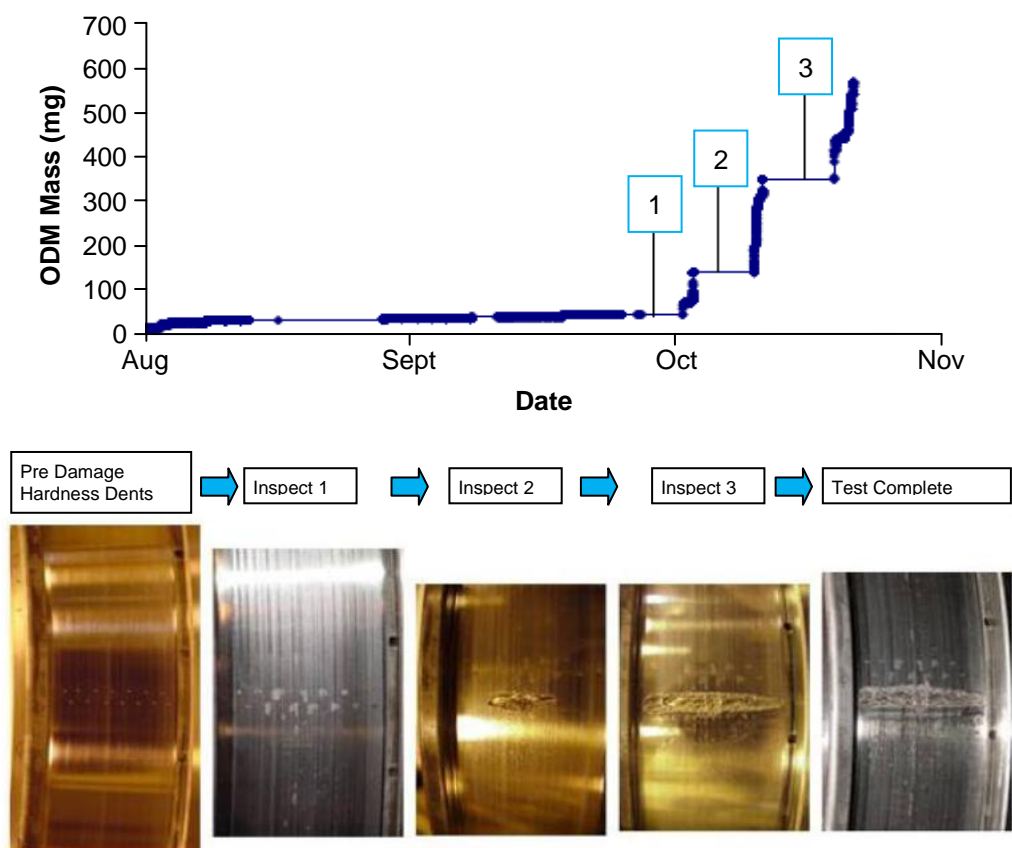


Figure 2.6: ODM mass accumulation as seeded bearing damage progresses, [35]

Gas path debris can also be monitored, as developed for the Joint Strike Fighter (JSF) program, [36, 37]. This monitors the ingestion of debris (foreign objects such as a bolt, or fluids/particles such as salt spray) into the engine, in real time, using electrostatic sensors. A simple schematic of the Engine Distress Monitoring System (EDMS) and Ingested Debris Monitoring System (IDMS) is shown in Figure 2.7.

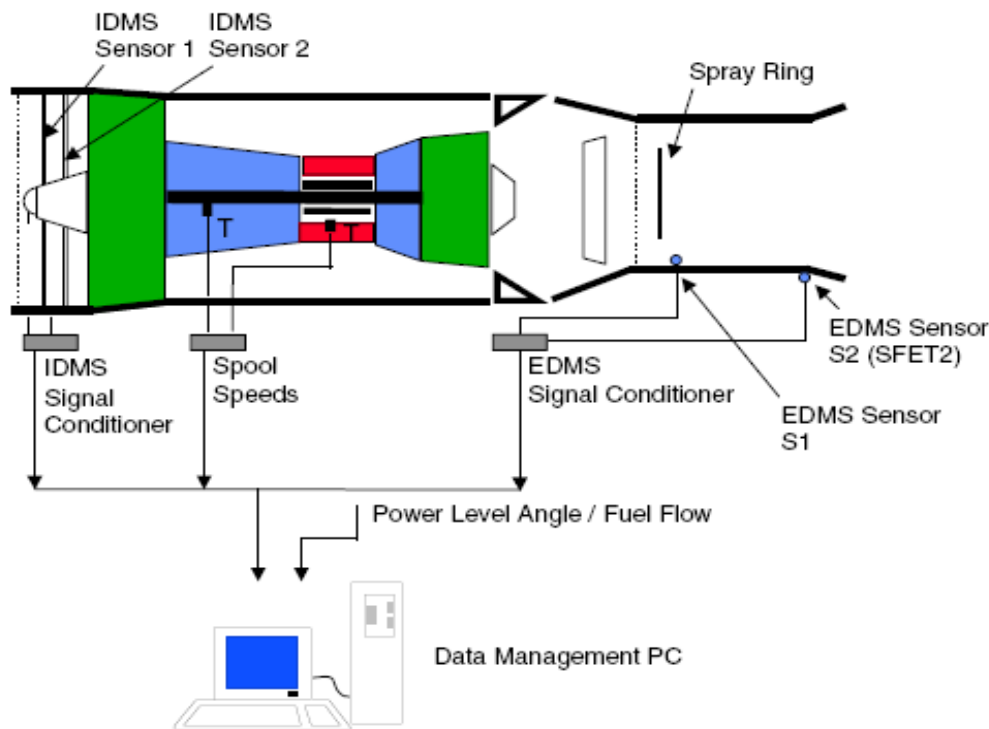


Figure 2.7: Schematic of IDMS & EDMS equipment for the JSF engine, [36]

2.1.7 Maintenance Databases

Maintenance information is also increasingly recorded digitally to aid diagnostics and prognostics. For example, the UK MOD Logistic Information Technology System (LITS) and the Work Records and Asset Management (WRAM) databases store the following kinds of information:

- Component strip reports – the results of investigating faults on components.
- Maintenance Work Orders (MWO) – information on every maintenance job carried out.
- Pilot forms – indicating flight profiles, stores, mass and centre of gravity, etc.
- Asset Tracking – identifies the part(s) fitted to each aircraft.

This information can be very useful in building up a fault database and linking symptoms with faults for improved diagnostic and prognostic capabilities.

2.2 The Need for Data Management

The drive for improved MAAAP has resulted in the large number of different data sources being collected, as described in the previous section. To date, limited advantage is being gained from this data because it is not intelligently managed. Often the data is simply collected and stored but not analysed in any way. Intelligent data management needs to handle the large data volume; assess the data integrity and correct or discard corrupt data as appropriate; fuse diverse sources of data; and finally, process the data to turn it into information used to make decisions. The following paragraphs give various cases from the literature where the need for aircraft data management is emphasised.

Evans, of the CAA, recommended best practice for managing HUMS data, [38]. While recognising that *'HUMS was probably the most significant isolated safety improvement of the last decade'* he also states that *'expert interpretation and/or physical investigation is usually necessary to detect incipient failures that would otherwise remain undetected.'* He emphasises the following areas of HUMS that require management:

- Setting appropriate thresholds and documenting any changes.
- Tracking the frequency with which data is downloaded and analysed.
- Maintaining data continuity (both the Maximum Permitted Period Between Successful Downloads (MPPBSD) and access to a helicopter's past vibration history) when analysing HUMS data.
- Prioritising maintaining the serviceability and system reliability of HUMS.
- Having a structured diagnostic approach. *'Successfully downloading data on a regular basis will only be of safety benefit if there is an effective response to any HUMS warnings.'*
- Establishing clear responsibilities for each aspect of HUMS.
- Ensuring adequate training and staff competency for all those using HUMS.

- Through life monitoring of the performance of HUMS and its management and using this feedback to make improvements.

Evans summarises that all of these aspects should be brought together into a comprehensive HUMS Management System. Such a system is the subject of the work presented in Chapter 3.

Larder *et al* describes how the monitoring function of HUMS needs to be enhanced by management, [39]. Table 2.1 shows the comparisons made between the use of HUMS to simply monitor aircraft and the aim of improving this by managing the data.

Table 2.1: Comparison of using HUMS for monitoring and management, [39]

	HUMS as Monitoring	HUMS as Management
Monitoring	Providing diagnostic data to detect and diagnose faults to indicate a requirement for maintenance action.	Providing diagnostic and prognostic data to predict requirements for future maintenance.
Health	Improving safety by detecting faults which represent a hazard to airworthiness.	Giving early indications of potential problems for maintenance planning purposes.
Usage	Relatively simple and largely limited to an automation of the aircraft logbook.	A key function to manage the usage of the aircraft and control life expired component replacements.
Output	Data: While HUMS has provided a significant advance in the quantity and quality of data available on aircraft, this has largely remained in the form of data rather than being transformed into information.	Information: The HUMS will convert monitoring data into information, which facilitates the effective management of aircraft maintenance.
Use	Stand-alone: HUM systems have mostly been stand-alone systems in two respects - the HUMS ground stations have not been integrated with other aircraft maintenance management or logistics systems and the HUMS outputs have not been fully integrated into the aircraft maintenance policy.	Integrated: The HUMS will be an integral part of the aircraft maintenance management and logistics system and the HUMS outputs will be fully integrated into the aircraft maintenance policy.

Pipe discussed the important features to consider when measuring the performance of HUMS, [40]. Pipe recommends looking at the following building blocks of a health usage management system in turn: sensors, signal

acquisition, signal processing, database management, alarm generation and management. The need to set appropriate thresholds, specific to each group of data, was also emphasised. Incorrect thresholds can result in missed faults or too many false alarms. Methods of assessing fault diagnosis are presented in section 2.3.6.

Draper, of the UK MOD HUMS IPT, commented that lessons could be learnt from the Chinook HUMS programme, where too much attention was paid to the onboard systems at the expense of the ground stations, [22]. For example, the ground station was contracted to operate on a UNIX platform. However, this was antiquated by the time it came into service. It is essential that the two systems are fully integrated: the Chinook HUMS had numerous Interface Control Document (ICD) errors, which were very expensive to rectify. Draper also states that *“there was a degree of naivety regarding the amount of data that HUMS would produce and how it would be handled through the support infrastructure”*.

2.2.1 Data Volume

As already described in section 2.1, the current trend is towards ever-increasing volumes of aircraft data being recorded and stored. The growth of data volume is not limited to the aviation industry and includes industries such as internet search engines, medical, marketing, banking, engineering and science. One poll in 2007 found that 22% of respondents reported mining databases of 1 terabyte or more, about double the number for 2006, [42]. The following paragraphs give some aviation examples of the large volumes of data available and how this has increased over time.

Chinook HUMS data is stored in a relational database, consisting of over four hundred tables as well as Binary Large Object (BLOB) files, of approximately one megabyte (MB) per flight, containing raw vibration signals and FDR parameter snapshots. The vibration signals are summarised by CIs. Nevertheless, the number of summarised CIs could be very large and typically about three thousand CIs are downloaded per flight from the Chinook HUMS (from about forty accelerometers; monitoring about 170 components; with up to four pass-band filters; up to twenty CIs can be computed from each dataset). Since the MOD operate a fleet of over forty aircraft, each flying many sorties per day, very large numbers of CI values must be managed daily.

The Eurofighter Typhoon database contains over 600 tables, covering Aircraft Structural Health (ASH), EHM, Logistics (LOG), Non Destructive Tests (NDT) and SHM. In addition, a Bulk Storage Device (BSD) can be used to record various configurations of flight parameters and the Crash Survivable Memory Unit (CSMU) stores over 6000 flight parameters to be used in incident investigations. This results in a database that is over 100 gigabytes (GB) for only thirty aircraft over three years. The expected UK fleet will be over 200 aircraft, over a life of thirty years; hence, over six terabytes (TB) of data will be stored and processed over the life of the Typhoon in the UK.

It is anticipated that the Lightning II (the F35 Joint Strike Fighter (JSF)) will routinely download five GB of data per sortie. With the anticipated worldwide fleet of over 3000 aircraft, flying many missions a day, hundreds of terabytes of data will be generated.

Table 2.2 gives some indication of the growth of FDR data capture over the past 30 years.

Table 2.2: Number of FDR parameters for various UK MOD aircraft types

Year of service entry	Aircraft	Number of FDR parameters	Frequency of download
1982	Chinook	158	Most flights
1982	Tornado	283	Every sortie on 3 OLM Aircraft
1998	Merlin (EH101)	296	Only for incidents
2000	WAH-64 Apache	1500	Routinely every sortie
2002	Typhoon (Eurofighter)	6059	Only for incidents
2012	Lightning II (F35 JSF)	1204	Routinely every sortie

Due to the large amounts of data, the intelligent management approach must be able to store and process terabytes of data and efficiently turn it into useful information that is clearly presented, to allow decisions to be made. It should also facilitate retrieving and displaying relevant data requested by the user.

2.2.2 Data Integrity

Before processing data it is important to assess its quality/integrity. There are a large number of potential causes for such data corruption. For example, in legacy strain gauge data the following causes of corruptions have been identified [43]: Short term failures in power supplies can cause sudden drops in strain measurements to values close to zero. The strain voltage levels are usually very small (millivolts), and, therefore, the strain signals are amplified hundreds of times (e.g. 200 times). The signal amplifiers can suffer failures leading to intermittent gain changes and signal corruption. Whilst the strain system cables are screened to prevent electromagnetic interference, operating in high radio fields can cause signal corruption, especially for aging strain gauge systems. Dry joints and soldering that fails to achieve perfect contacts can cause erratic strain behaviour characterised by the signals being stuck at a wrong strain levels for a period of time. Changes in resistance due to ingress of moisture can cause corruption. Failures in temperature compensation mechanisms can lead to sensitivity to local temperature changes and signal drifts. Strain gauge sensitivity can be influenced by changes in bonding characteristics. Errors can also occur during strain signals multiplexing, synchronisation and recording. The corruption patterns of a strain gauge system are influenced by the system maintenance status; they can vary with time and can be at variance with those of other systems.

Multiple methods exist for identifying corrupt data and, if possible, replacing it with synthesised values. For example, the Smiths (now GE Aviation Systems) Automatic Data Correction (ADC) algorithm identifies short period corruptions in flight data and strain measurements, [44]. The identified short period corruptions can include: spikes, multi-spikes, spike-step transitions, steps, hesitant steps, step reversals, dropouts, constant signals, complex corruptions and jumps. An example of each corruption type is shown in Figure 2.8. The examples shown come from three different aircraft systems.

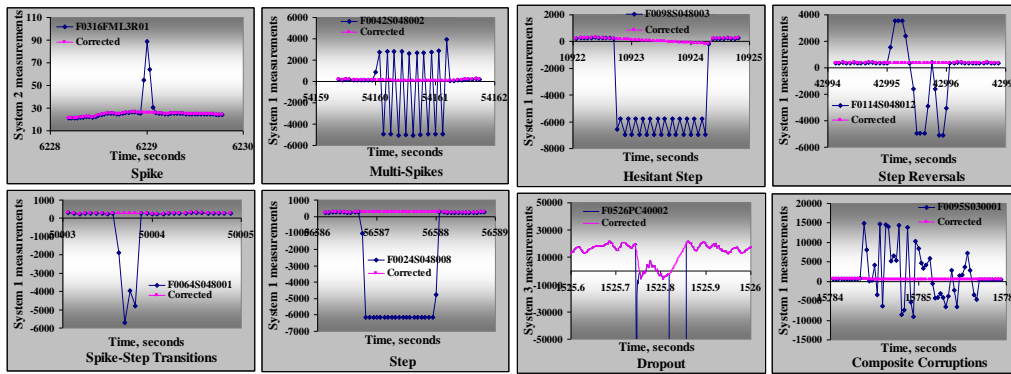


Figure 2.8: Examples of short period corruptions corrected by the Smiths Automatic Data Correction algorithm, [44]

The Smiths data quality algorithms can also identify long period corruptions caused by temporarily or persistently inoperative sensors and/or calibration problems. For a target sensor, long period corruption was identified by comparing the statistics of the sensor data across a number of sorties and by cross-correlating its data with other sensor data and/or with synthetic data generated by neural networks. The statistics of the sensor data were computed over the entire sortie, over its most probable data levels and at a number of predefined flight conditions referred to as ‘hypercubes’ or ‘points in the sky’. A decision making process was implemented to fuse sensor health indicators derived from the above statistics, [44]. Examples of observed strain corruptions are shown in Figure 2.9. The chart on the left shows that the data from the sensor on the right wing gave significantly lower readings than the left wing – the synthesised values are more correlated with the left wing sensor. The right hand chart shows an example where the first thousand seconds of data were lost for both sensors, but could be approximated from other flight parameters.




Figure 2.9: Examples of long period corruptions identified by the Smiths data quality algorithms, [44]

Roemer *et al* stressed the importance of detecting and correcting faulty data:

“Sensor problems such as ground loop faults, sensor drift or electrical noise can often appear as the onset of a performance or vibration fault and must be isolated and detected properly.” [45]

Goebel *et al* identified the following data anomalies, which should be removed by pre-processing before diagnostic or prognostic techniques can be applied: non-linearities, noise and outliers, sensor failures, saturation, disjoint response to events, information smoothing, information fading and reliability scaling, [46]. These data anomalies will be apparent in most data sets, not just within aviation.

Han and Kember in “Data Mining – Concepts and Techniques” explore methods for cleaning data to remove missing values, noisy data and inconsistencies, [47]. For example, when missing values are identified you could choose to: ignore the whole record, fill the missing value manually or fill with a global constant, a mean value or interpolate between adjacent values. All of these techniques will bias the data and so must be carefully considered.

With the increasing capabilities and application of HUMS, the FAA, in conjunction with the US Navy, was interested in assessing the impact of degraded data on the performance of these systems, [48]. They investigated the effect on flight regime recognition of the following data degradations and found that most of them could be identified: data spikes, signal discontinuities, lost signals, constant signal, intermittent signal, erratic signal, band edge signal and signal drift.

Since data may contain corruption or inconsistencies of one sort or another, the intelligent management tools must always assess the data integrity and, where appropriate, correct it or discard it before further analysis. It is also advisable to keep a record of the corruptions identified, to help give confidence in the analysis results and for future improvements to the data acquisition system.

2.2.3 Data Fusion

Given the large number of potential data sources outlined in section 2.1, it is necessary to be able to fuse/merge data from multiple sources for improved analysis. For example, a better indication of faults in helicopter drive-train

components can be identified if vibration data and wear debris data are used together. This data is usually recorded by different systems in different databases and so are often analysed in isolation. However, information about debris can help to confirm the source of increased vibration, and isolate the type of damage. This data is also sampled at different intervals. Typically, vibration data is recorded multiple times per flight, whereas debris may only be recorded every 20 flying hours. To be useful the disparate data sets must be temporally aligned.

Data fusion can be performed on raw data or after some processing or generation of health indicators. For example, FUMS™ offers users the ability to fuse data from many different sources (Figure 2.10) and the flexibility to align data using complex criteria, such as finding the closest match between dates within specific limits such as one week, [13, 49, 50].

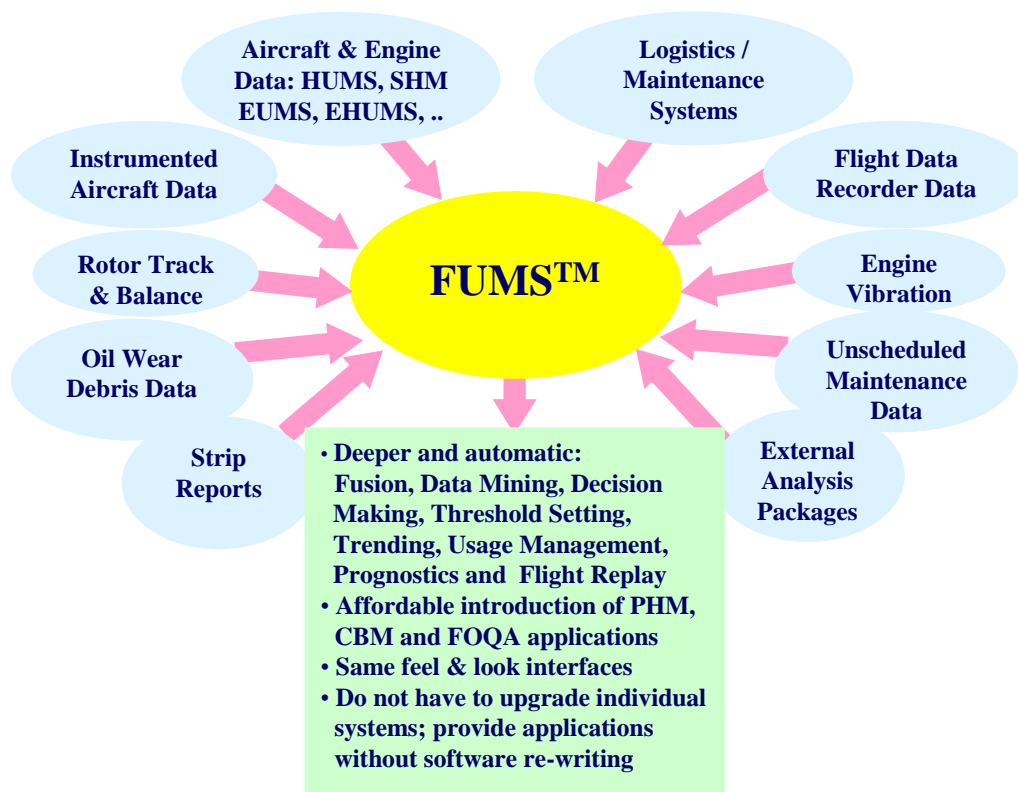


Figure 2.10: Schematic of the FUMS™ fusion platform, [13]

As described in [51], there are at least three levels at which fusion can be performed: sensor level, features level and decision level. Before fusing data from different sources it may be mapped into a single comparable domain, such as a fault likelihood score on a scale of 0 to 1, [52].

Goebel *et al* , [46], have fused data sources such as engine operation parameters, vibration data, oil debris and mathematical spall propagation models, to give improved engine fault diagnostics. As well as fusing data sources they have also investigated fusing the results of two different modelling approaches, [53]. One is to model, from first principles, the physics of fault initiation and propagation; the second is an empirical model of condition-based fault propagation rates using data from experiments in which the conditions are controlled, or otherwise known and the component damage level is carefully measured. These two approaches have competing advantages and disadvantages. However, fusing the results of the two approaches produces a result that is more robust than either approach alone, [54].

Not only do different data sources need to be fused, but each stage of the diagnostics process may use different techniques and technologies from a number of suppliers, which must be integrated to provide a coherent solution. Callan *et al* have investigated the integration of PHM systems for CBM, [55]. Layers of PHM systems such as acquisition, data manipulation, condition monitoring, health assessment, prognostics and automatic decision reasoning, need to be aligned from potentially different suppliers.

The intelligent management strategy should use all available relevant data sources as part of any decision process. It should also be able to fuse the results of multiple analysis techniques performed on each data set to give an improved overall decision.

2.3 Techniques

There are a large number of AI and model based techniques available for analysing data, to determine faults and predict time to failure. The following sections give a brief introduction to some techniques, which are referenced later in the thesis: Mathematical models, signal processing, feature extraction, data mining, reasoning and decision making.

2.3.1 Mathematical Models

Mathematical models attempt to model the real life physical processes behind measured data, in order to predict the expected data that would be recorded from a given system in a particular state (healthy, faulty, under particular loads, etc.). The model's results can be compared with measured data to identify features. In order to be useful, the models must be as realistic as possible, taking into account all the failure modes and operating conditions of the system. This requires a good engineering understanding of the system.

A fundamental problem in the development and validation of PHM technologies is the general shortage of realistic fault signature data. While healthy signatures can be obtained from operational systems, faults are relatively rare and difficult to observe. The PHM community often have to rely on bench level seeded fault test data collected under a limited set of conditions. Models can also be built using healthy data and then used to identify outliers.

Lybeck *et al* have developed a modelling and simulation toolset for the vibration signatures of faulted components in propulsion subsystems such as gearboxes, [56]. Figure 2.11 shows how synthesised vibration data is built up from a forcing function to calculate the response, which is combined to give the output. The figure shows an example simulated time series vibration trace for a bearing, compared with measured data.




Figure 2.11: Simulating vibration data and comparing with measured values

Physics based fault models also use techniques such as regression using Singular Value Decomposition (SVD), state deciders to use different models for different conditions and Finite Element Analysis (FEA) to build up complex relationships within structures.

2.3.2 Signal Processing

When analysing time-variant data, such as vibration, various processing techniques may be used to help understand the content. The following are some common techniques.

Fourier Transforms (usually referred to as Fast Fourier Transforms – FFT) convert time series (such as measured vibration signals) into frequency series. The FFT tells us what frequencies are present in the signal and at what amplitude. Figure 2.12 shows an example signal containing a combination of 10, 20 and 30 Hz sine waves and its associated FFT with peaks at these frequencies.

FFTs are particularly useful when analysing gear vibration data that will have strong frequency components at meshing frequency between the teeth of two gears. FFTs can be used to isolate these peaks and remove them if necessary to emphasise other damaging vibration (filtering the FFT and applying an inverse FFT). For another example, helicopter vibration would contain strong FFT peaks at the main rotor speed (and harmonics, e.g. twice the rotor speed) and at blade passing frequencies.

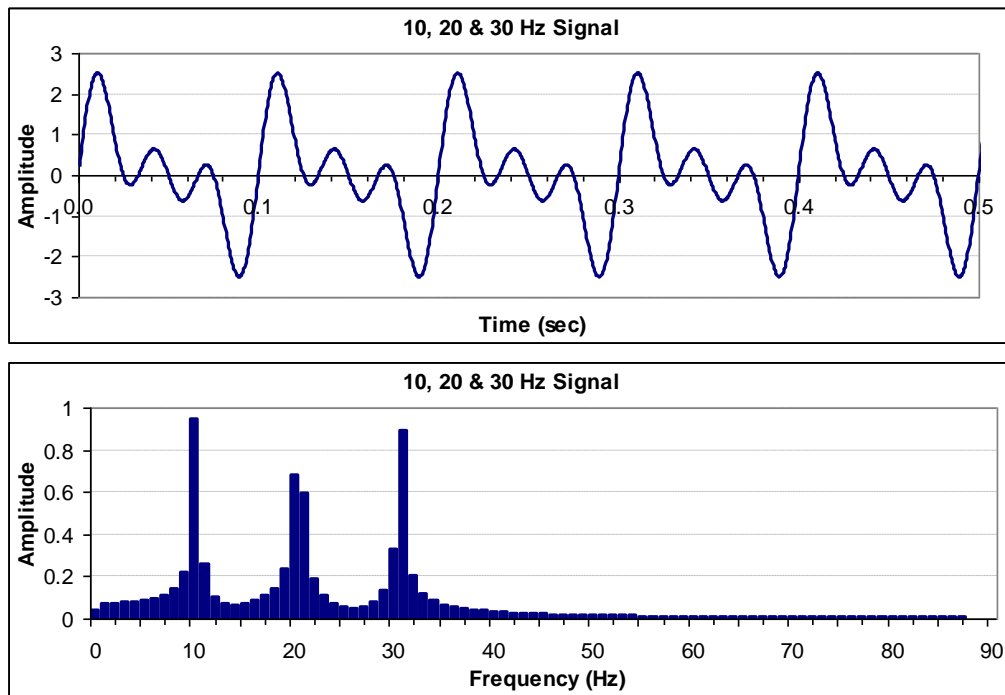


Figure 2.12: Example signal containing 10, 20 and 30 Hz and its FFT

The statistics of non-stationary signals change with time. Such signals can be analysed by techniques such as Wavelet Transforms (WT) and short term FFT; the later applies an FFT to a subset of the signal, and then repeats this, moving the section of the signal along each time. This allows you to analyse how the frequency changes with time.

2.3.3 Feature Extraction

Feature Extraction is a method of compressing large volumes of data into some simple measures, which can be more easily compared with other samples. For example, a vibration signal may be summarised by statistics such as the kurtosis, maximum difference and Root Mean Squared (RMS), etc. As another example, Acoustic Emission (AE) signals may be compared by features such as peak amplitude, duration, rise time and the number of times a threshold is crossed, Figure 2.13.




Figure 2.13: Acoustic Emission signal feature extraction, [57]

2.3.3.1 Statistics

Statistical techniques can be used to identify outliers and trends. For example, data that lies outside three standard deviations from the mean will be the most unusual data points. Other statistics such as RMS and statistical moments (mean (1st), variance (2nd), skewness (3rd), kurtosis (4th), M5 and M6) are all essential in analysing vibration data. Moving averages and medians can be used to filter noise/spikes in the data to see the underlying trend, and least-squared linear trend lines can be used to forecast the value into the future. Correlation and Chi-squared can be used to identify how similar two data sets are. The CIs generated by HUMS are essentially statistics of the original vibration data.

2.3.4 Data Mining

Data Mining is a term that describes the process of turning large quantities of data into useful information or knowledge. As stated by Han and Kember in “Data Mining – Concepts and Techniques”:

“The major reason that data mining has attracted a great deal of attention in the information industry in recent years is due to the wide availability of huge amounts of data and the imminent need for turning such data into useful information and knowledge”, [47]

































Data Mining can be used to find patterns in data in the form of associations, classes, clusters, outliers and evolution. Large data sets can be reduced by methods such as data cube aggregations, dimensionality reduction, data compression and numerosity reduction. Barnathan gives a good review of methods for mining complex high-order datasets, [58].

Data mining is being progressed in many fields. For example, in 2008 a new journal was launched called BioData Mining which is focused on the development of data mining techniques applied to biological data, [59]. Bellazzi and Zupan have published a review of predictive data mining in clinical medicine giving a worked example, discussions on techniques and offers guidance, [60]. Their recommendations are applicable to the wider use of data mining, and include:

- Define the success criteria in advance. Set acceptable ranges of evaluation statistics prior to modelling.
- Model probabilities, not crisp class membership. Prefer methods that report confidence intervals.
- Avoid over-fitting. Never test models on data that was used in their construction. If possible, test the resulting model on an independent separate data set.
- Prefer modelling techniques that expose relations and can present them in a readable form.
- If still of acceptable performance prefer simple modelling techniques, possibly those that derive models that can be reviewed and criticized by experts.

Table 2.3 shows the results of a poll on which techniques are most used for data mining. Decision trees and rules are consistently the most popular.

Table 2.3: Poll on data mining techniques [61]

Data Mining Technique	2005	2006	2005 Percentage	2006 Percentage
Decision Trees/Rules	107	90	 14%	 16%
Clustering	101	70	 13%	 12%
Regression	90	67	 11%	 12%
Statistics	80	64	 10%	 11%
Association rules	54	54	 7%	 9%
Visualization	63	38	 8%	 7%
SVM	31	31	 4%	 5%
Neural Nets	61	31	 8%	 5%
Sequence/Time series analysis	26	24	 3%	 4%
Bayesian	30	24	 4%	 4%
Nearest Neighbor	34	20	 4%	 3%
Boosting	25	17	 3%	 3%
Hybrid methods	23	14	 3%	 2%
Bagging	20	13	 3%	 2%
Genetic algorithms	19	12	 2%	 2%
Other	20	4	 3%	 1%

2.3.4.1 Cluster Algorithms

There are a plethora of cluster algorithms available, each suited to particular data sets or types of information to be gathered. Representative examples are given below.

The FUMS™ anomaly/novelty detector can efficiently recognise, in huge data sets, densities of clusters, anomalies and non-linear patterns, having a wide range of sizes/orientations, [62].

Support Vector Machines (SVM) are a set of related supervised learning methods used for classification and regression. A special property of SVMs is that they simultaneously minimize the empirical classification error and maximize the geometric margin; hence they are also known as maximum margin classifiers. Many linear classifiers (hyper-planes) separate the data. However, only one achieves maximum separation. For example, Figure 2.14 shows two clusters of data (circles and squares) which could be separated by many lines (L1, L2, L3...). L2 gives the best separation.




Figure 2.14: Support Vector Machines – finding the best separation, [63]

Adaptive Resonance Theory (ART) allows clusters to keep learning, by moving the clusters as more data arrives. This means that the cluster model is always up to date, using as much information as possible. However, it may mean that the model learns about unwanted conditions.

The Kohonen map is an example of an unsupervised neural network. The map consists of a number of units, which are arranged in a fixed topology – in this case, a rectangular grid. Each unit has an n -dimensional weight vector associated with it. This weight vector can be interpreted as that unit's position within the data space. A fundamental feature of the Kohonen map is that each unit occupies a position within the n -dimensional data space and also a position within the topology of the map. The Kohonen map's aim is to adapt the weight vectors so that the units of the map are distributed over the data-space where the density of units in any part of the space reflects the density of data in that part of the space. However, the map will also preserve its topology - i.e. neighbouring units in the map's fixed topology will tend to end up next to each other in the data space as well. Imagine a three-dimensional data space with data distributed unevenly within this space, Figure 2.15. Now imagine the Kohonen map as trying to arrange an elastic fishing net within this space so that:

- The relative density of nodes at any point within the space approximately corresponds to the density of data in that point.
- The surface area of the stretchable mesh is minimised.

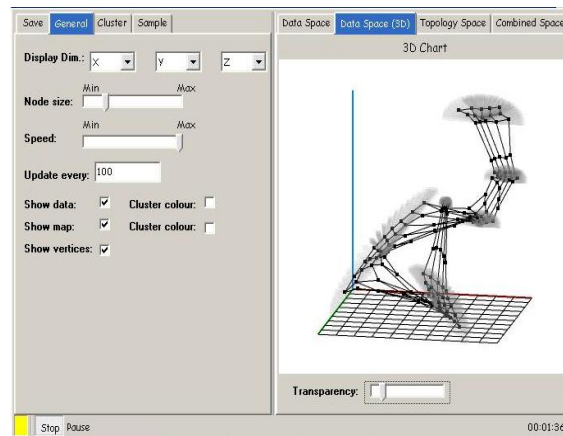


Figure 2.15: FUMS™ Kohonen map example

The Kohonen map has been applied to many problems ranging from the Travelling Salesman Problem (TSP) to robot learning. The method is also useful for reducing the original data down to a set of representative data points that is far smaller than the original data set or can be used for dimensionality reduction.

2.3.4.2 Artificial Neural Networks (ANN)

Haykin in his book “Neural Networks a Comprehensive Foundation”, [64], gives the following description of what a neural network is and how it relates to the network of neurons in the human brain:

“A neural network is a massively parallel distributed processor that has a natural propensity for storing experiential knowledge and making it available for use. It resembles the brain in two respects:

1. *Knowledge is acquired by the network through a learning process.*
2. *Inter-neuron connection strengths known as synaptic weights are used to store the knowledge”, [64]*

An ANN can be trained on a data set with known inputs and required outputs (the target). Once trained, given another set of inputs from the same source, the output can be computed. For example, FDR parameters such as vertical acceleration and air speed from one aircraft can be used as training inputs, and measured wing strain can be used as the target. Once trained, the network could be used to synthesise the wing strain for other aircraft given the FDR data, [29] [44].

An ANN is made up of individual neurons that take in N inputs, sum them together with a weight for each input, and then pass this through an activation function. The activation function can take various forms such as thresholds, linear, Gaussian, sigmoid, etc., as shown in Figure 2.16. Neurons are then stacked together to create a network. This usually takes the form of an input layer, output layer and a number of hidden layers made up of any number of neurons, Figure 2.17.




Figure 2.16: A model of a neuron, [65] and example activation functions, [66]




Figure 2.17: Typical ANN, [67]

2.3.4.3 Genetic Algorithms (GA)

Genetic Algorithm (GA) theory was developed by Holland in the mid-1960s to simulate the adaptive optimisation process of nature, [68]. GA theory simulates the principles of evolution as put forward by Darwin in 1859, [69]. The slow process of adaptation is governed by inheritance from the survived fittest among a population of genetic diversity that also allows adaptation by mutation.

For example, Figure 2.18 shows a set of experimental inputs x_1 , x_2 and x_3 that give a corresponding output y_m . A 'chromosome' containing three 'genes' holds values for three weights to multiply by the input x values. These weights may be picked at random to start with or may be set using some knowledge of the problem. The chromosome is then used to calculate a calculated y value (y_c) from the experiment inputs. y_c is then compared to the experimental result y_m to find the error – representing the fitness of the chromosome. The chromosome is then modified to create offspring, either by mutating a gene at random, or by swapping a gene with another chromosome.




Figure 2.18: Schematic of GA, [50]

2.3.5 Reasoning and Decision Making

Expert Systems (ES) have developed to capture human experience and provide reasoning and decision making capabilities. The first ES were constructed in the late 1960s. These systems attempt to take the place of a human expert, mimicking their decision-making capabilities. ES are intelligent computer programs that use knowledge and inference procedures to solve problems that are difficult enough to require significant human expertise for their solution, [71].

The main techniques that provide reasoning and decision making capabilities include: Crisp Logic, Fuzzy Logic and Bayesian Networks.

2.3.5.1 Crisp Logic

Crisp Logic applies hard and fast rules using Boolean logic of the form:

IF condition THEN action

Multiple conditions can be combined using operators such as OR, AND or NOT, for example:

IF (condition 1 OR (condition 2 AND NOT condition 3)) THEN action

The main shortcoming of Crisp Logic is how to treat uncertainty. This is where Fuzzy Logic and Bayesian Networks can help, by formulating the expression in a form such as:

IF condition with certainty x THEN fact with certainty $f(x)$

2.3.5.2 Fuzzy Logic (FL)

“Humans have a remarkable capability to reason and make decisions in an environment of uncertainty, imprecision, incompleteness of information and partiality of knowledge, truth and class membership. The principal objective of fuzzy logic is formalisation/mechanisation of this capability.”, [72]

FL is the application of the concept of fuzzy sets that was first published in 1965 by Zadeh, [73]. A fuzzy set is a mathematical model of vague qualitative or quantitative data. For example, consider a birds-eye view of a forest in Figure 2.19.

- Is location A in the forest? Certainly yes, $\mu_{forest}(A) = 1$.
- Is location B in the forest? Certainly not, $\mu_{forest}(B) = 0$.
- Is location C in the forest? Maybe yes, maybe not. It depends on a subjective (vague) opinion about the sense of the word "forest". Let us put $\mu_{forest}(C) = 0.6$.




Figure 2.19: Simple fuzzy set: birds-eye view on a forest, [74]

A basic FL application might characterise sub-ranges of a continuous variable. For instance, a temperature measurement for a cooling fan system might have several separate membership functions, defining particular temperature ranges needed to control the fan properly. Each function maps the same temperature value to a truth value in the zero to one range. These truth values can then be used to determine how the fan should be controlled.




Figure 2.20: Simple fuzzy set definition, [75]

In Figure 2.20, *cold*, *warm* and *hot* are functions mapping a temperature scale. A point on that scale has three "truth values" — one for each of the three functions. For the particular temperature shown, the three truth values could be interpreted as describing the temperature as, say, "fairly cold" (*cold* = 0.8), "slightly warm" (*warm* = 0.2) and "not hot" (*hot* = 0), [75].

FL usually uses IF/THEN rules, or constructs that are equivalent, such as fuzzy associative matrices. Rules are usually expressed in the form:

IF *variable* IS *set* THEN *action*

For the example in Figure 2.20, rules may be:

IF temperature IS cold THEN slow down fan

IF temperature IS warm THEN maintain level

IF temperature IS hot THEN speed up fan

The AND, OR and NOT operators of Boolean logic exist in fuzzy logic, usually defined as the minimum, maximum and complement. For example:

IF temperature IS cold OR warm = $\text{MAX}(\text{cold}, \text{warm}) = \text{MAX}(0.8, 0.2) = 0.8$

Ross describes FL and engineering applications in more detail in his book, [76].

2.3.5.3 Bayesian Networks

A Bayesian network (or belief network/Bayesian Theory (BT)) is a probabilistic graphical model that represents a set of variables and their probabilistic dependencies, [77]. Bayesian networks start with known probabilities and progress these to dependants using conditional probability, [78], i.e., given the event B, the probability of the event A is x.

The concept is best explained by a simple example. Suppose that there are two reasons that could cause grass to be wet: either the sprinkler is on, or it's raining. Also, suppose that the rain has a direct effect on the use of the sprinkler (namely that when it rains, the sprinkler is usually not turned on). Then the situation can be modelled with the Bayesian network in Figure 2.21. All three variables ($G = \text{Grass wet}$, $S = \text{Sprinkler}$ and $R = \text{Rain}$) have two possible values T (for true) and F (for false).




Figure 2.21: A simple Bayesian network, [77]

The model can answer questions like "What is the likelihood that it is raining, given the grass is wet?" by using the conditional probability formula and summing over all associated variables:

$$\begin{aligned}
 P(R = T | G = T) &= \frac{P(G = T, R = T)}{P(G = T)} = \frac{\sum_{S \in \{T, F\}} P(G = T, S, R = T)}{\sum_{S, R \in \{T, F\}} P(G = T, S, R)} \\
 &= \frac{(0.99 \times 0.01 \times 0.2 = 0.00198_{TTT}) + (0.8 \times 0.99 \times 0.2 = 0.1584_{TFT})}{0.00198_{TTT} + 0.288_{TTF} + 0.1584_{TFT} + 0_{TFF}} \approx 35.77\%
 \end{aligned}$$

The difficulty in using Bayesian networks for real world engineering applications is finding appropriate probabilities, especially when fault cases are very rare.

2.3.6 Evaluating Techniques

Since a system may have many failure modes, each with individual characteristics, it is often advantageous to use multiple techniques and then use reasoners to evaluate an overall assessment.

To illustrate the performance of the individual model-based reasoners, a matrix presentation, of input versus response, can be used, Figure 2.22. Quantities expressed on the diagonal are correct classifications. Remaining quantities displayed are either false positives (no fault input, but a fault output - top row), mis-classifications (wrong fault output - off diagonal in fault rows) or false negatives (fault input, but no output - no fault column). The goal is for the technique to have as many predictions as possible on the diagonal axis.

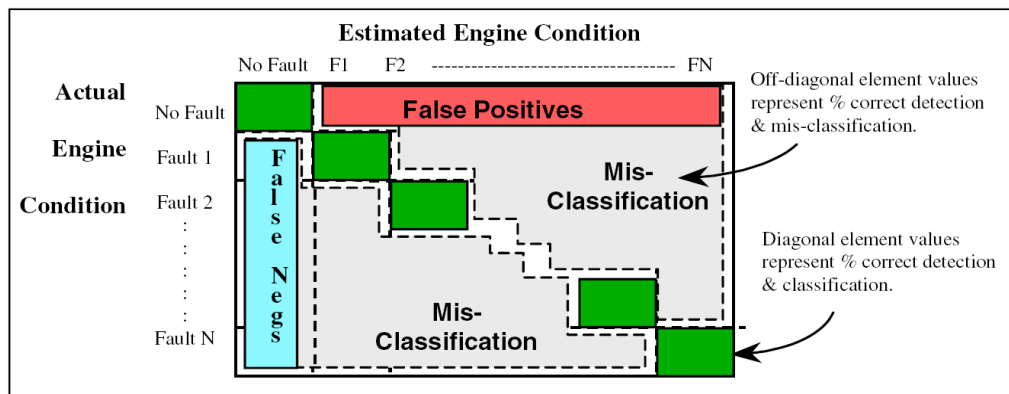


Figure 2.22: Matrix View of Reasoner Performance, [48]

Bock *et al* have developed new algorithms for user defined false alarm rates, stating: “The assignment of a fault detection threshold implies a trade-off between true positives (sensitivity) and false positives (specificity).” [49].

Byington *et al* have developed a Metrics Evaluation Tool (MET) to assess the performance and effectiveness of vibration features typically used in HUMS, [50]. The method is based on a detection decision matrix developed by Dowling of Boeing, [51] shown in Figure 2.23.

Detection Decision Matrix			
Outcome	Fault (F1)	No Fault (F0)	Total
Positive (D1) (detected)	a Number of detected faults	b Number of false alarms	$a+b$ total number of alarms
Negative (D0) (not detected)	c Number of missed faults	d Number of correct rejections	$c+d$ Total number of non-alarms
Total	$a+c$ Total number of faults	$b+d$ Total number of fault-free cases	$a+b+c+d$ Total number of cases

Figure 2.23: Detection Decision Matrix, [51]

From this matrix, the detection metrics can readily be computed. The probability of detection (POD) assesses the detected faults over all potential fault cases:

$$POD = P(D1 / F1) = \frac{a}{a + c} \quad 2.1$$

The probability of false alarm (POFA) considers the proportion of all fault-free cases that trigger a fault detection alarm:

$$POFA = P(D1 / F0) = \frac{b}{b + d} \quad 2.2$$

The accuracy is used to measure the effectiveness of the algorithm in correctly distinguishing between a fault-present and fault-free condition. The metric uses all available data for analysis (both fault and no fault):

$$Accuracy = P(D1 / F1 \& D0 / D0) = \frac{a + d}{a + b + c + d} \quad 2.3$$

2.4 Examples of AI and Modelling Techniques used to Manage Aircraft Data

In this section a number of examples are presented where AI and mathematical models have been applied to aircraft data. In fact, these techniques are applicable to a wide range of applications outside aviation but the focus of this literature survey was within the aviation industry. The examples have been split into the following categories:

- Mathematical Models
- Improving HUMS
- Engine Health Management
- Prognostics
- Managing Uncertainty

2.4.1 *Mathematical Models*

Physics based mathematical models allow real-world laws to be captured and used to compare with measured data to identify anomalies. Models can also be used to synthesise data where it is unavailable. Azzam investigated the use of mathematical models and artificial intelligence techniques to improve HUMS prediction capabilities, [79]. Azzam developed a Helicopter Mathematical Model (HMM) that included a physics based rotor wake model, rotor elastic model, fuselage model and dynamic model. The HMM can be configured to represent a particular helicopter, by setting appropriate coefficients (such as the number of blades, rotor diameter, centre of gravity, mass, moments of inertia, etc.), and can be used to simulate loads and vibration for specified flight conditions. Aerodynamic, structural and hydraulic faults, such as pitch link errors, blade cracks and mass imbalance, can also be simulated and the resulting vibration signatures generated. The theoretical results, generated using the HMM, compare favourably with test data gathered from in-service aircraft. Azzam has also used ANNs to simulate helicopter blade lug fatigue, wing pivot fatigue, tail rotor torque and engine fatigue. Extensive trials have been undertaken that show that fatigue values can be synthesised, without the use of strain gauges, giving improved lifing accuracy and at a significantly lower cost than fitting strain gauges to the whole fleet, [79].

2.4.2 Improving HUMS

Various studies have attempted to improve fault detection using HUMS data, simulations or seeded fault data. The following paragraphs give some examples:

Byington *et al* have developed feature extraction and fault isolation algorithms and applied them to three test rig examples, [80] and [81]. The three examples were: aircraft engine ceramic bearing data, aircraft engine bearing test cell data and T63 turboshaft engine test cell data. The methods focus on the high frequency bands which are most likely to indicate faults. ANNs and tree classifiers were used to isolate faults. As stated in their conclusions, they have yet to apply the techniques to in-service aircraft data, which makes it difficult to assess the usefulness of the methods.

Hochmann and Baringer presented a number of case studies showing how drive train component CIs can be analysed to give improved diagnostics, [82]. The following case studies were presented:

1. A time to spatial domain transformation example, where one index position (tachometer value) is offset relative to the others.
2. A gear component example of how subtle characteristics in the CIs could lead to mis-diagnosis. In the example, what was thought to be a tooth crack issue turned out to be due to tooth positional deviations.
3. A rolling element bearing example, demonstrating how global changes may be misinterpreted for component specific issues.

Feng mathematically simulated rolling element bearing defects in the inner or outer race, [83]. This can be used to understand the likely fault that is causing the vibration patterns that have been measured by HUMS. However, in-service data is always harder to analyse than mathematically simulated faults, so the techniques do not always have direct read across. Appendix A1.6 gives details of typical bearing geometry.

Hamza *et al* investigated the use of knowledge capture, GA optimisation and case-based reasoning (using Bayesian networks) to improve the reliability of HUMS, [84]. Figure 2.24 shows the data flow involved in capturing knowledge of HUM systems and using it to improve the reliability of HUMS diagnostics. The aim of this technique is to learn from past data off-board and then use

Bayesian case based reasoning onboard to give more reliable HUMS fault detection and isolation. The knowledge capture tool combines the fault type, the diagnosis algorithm and the performance metrics such that they can be used by the GA. The GA determines the optimum algorithm to use for each fault type. The database of both fault and no-fault data builds up over time improving the case based reasoning of the on-board system. The system was integrated into the Chadwick-Helmuth Vibration eXPert tool. While the concepts of this approach seem to offer HUMS improvements, it is not clear whether the system has been used on in-service data, or how well the system actually performs.

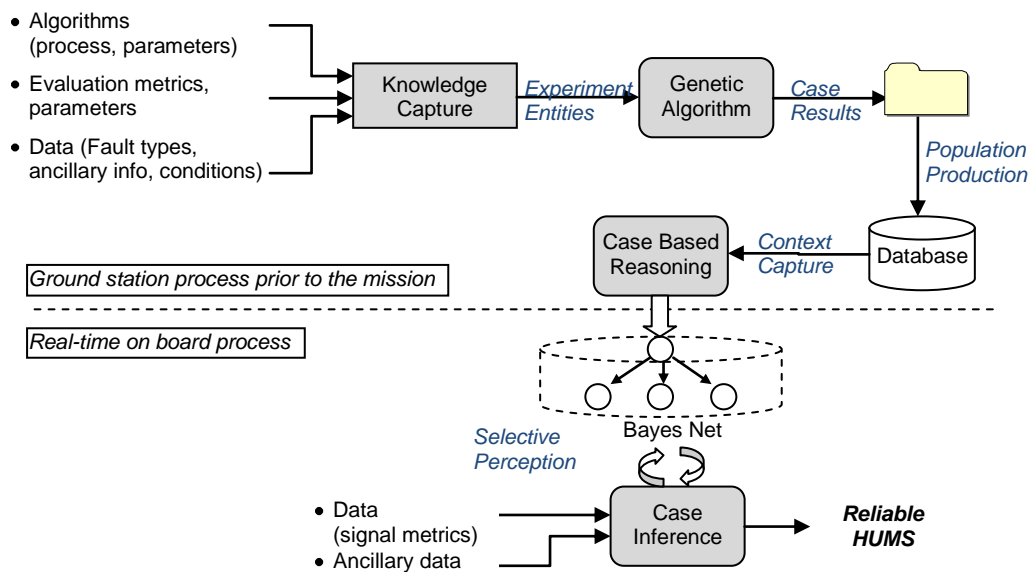


Figure 2.24: Adaptive HUMS Knowledge Management Architecture, [84]

Bechhoefer *et al* have investigated the setting of HUMS CI thresholds by modelling aircraft and torque band variance, i.e. having different thresholds for different levels of torque and for each individual aircraft, [85, 86]. The techniques have been applied to data from thirty Sikorsky UH-60L Black Hawk aircraft. An Integrated Mechanical Diagnostics (IMD) HUMS, was also developed by Bechhoefer *et al*, [87]. IMD-HUMS combines CIs into Health Indicators (HIs) and uses Kalman filtering techniques to reduce the noise in the data. The CIs are fused to reduce the false alarm rate by simply assigning a value of 0 for the CI if it is below the warning level, 1 if it is above the warning level but below the alarm level, and 2 if it is above the alarm level, (2.4).

Figure 2.25 shows 3 shaft order CIs and the HI derived from them. A Kalman filter is then applied to the HI reducing the noise and making the fault condition clear after acquisition 325. The filter does, however, delay the detection of the fault, which is clear in the raw data by acquisition 300.




Figure 2.25: IMD-HUMS example – three shaft order CI values, the HI value (cyan) and the Kalman filter of the HI (black), [87]

$$HI = \frac{\sum_{i=1}^n W_{CI}}{\sum_{i=1}^n 2}, W_{CI} = \begin{cases} 0, CI < warnings \\ 1, warnings \leq CI \leq alarm \\ 2, CI \geq alarm \end{cases} \quad (2.4)$$

An engineering software utility called Mechanical Diagnostics Analysis Toolkit (MDAT) was developed by Goodrich and was used to investigate the propagation of damage in the four bearing defect passage frequencies: cage, ball/roller, outer race and inner race, [88]. Appendix A1.6 gives further details of bearing geometry. Again, MDAT employs the concept of a HI that fuses multiple CIs using both rule-based and statistical approaches.

2.4.3 Engine Health Management

A number of major players in engine health management have been developing a holistic approach that fuses the multiple sources of engine data for improved results. The following paragraphs expand on a few examples:

Ashby, of GE Aircraft Engines, and Scheuren, of the Defence Advanced Research Projects Agency (DARPA), developed a condition-based Intelligent Maintenance Advisor for Turbine Engines (IMATE), [89]. IMATE integrates sensor and model-data from various diagnostic, prognostic and usage sources, with information fusion algorithms to produce a more accurate assessment of engine condition than those available from any individual source of information. Typical inputs to the information fusions are: Full Authority Digital Engine Control (FADEC) sensor values, intermittence tests, trending algorithms, model based reasoners, vibration analysis, engine models and life usage algorithms. The paper presents the results of IMATE using simulated data. The information fusion requires that all diagnostic tools give a HI in the range of 0 to 1, [52]. The tools are also given a weighting based on the reliability of detecting faults in simulated data.

Volponi *et al* (including NASA, Pratt and Whitney, and the US Army) have developed an information fusion architecture for engine diagnostics and health management, [51]. They list a number of data sources for use in fusion including: engine gas path measurements, oil/fuel system measurements, vibration measurements, structural assessment sensors, FADEC codes, onboard engine models, maintenance/analysis history and companion engine data. In the architecture developed, data fusion can take place at one of three levels:

- **Sensor Level:** Align data, for example, by resampling.
- **Features Level:** Using extracted features, Failure Modes Effects and Criticality Analysis (FMECA) and Bayesian networks for fault isolation.
- **Decision Level:** Used fuzzy logic for feature level fusion.

Roemer *et al* have been working on prognostics and diagnostics of gas turbine faults for many years. In 1999 they published work using a probabilistic approach to diagnose gas turbine engine faults, [90]. The probabilistic methods are aimed at robustly assessing the condition of engine sensor signals, performing mechanical and performance diagnostics, and providing critical component prognostics. By 2005 they had reported on an integrated set of turbo-machinery health monitoring techniques that, when implemented, would offer significant potential for reducing current turbo machinery Life Cycle Costs (LCC), [45]. Various techniques were investigated for diagnostics and

prognostics of turbine engine faults, including the use of trained neural networks, fuzzy logic analysis, signal auto and cross-correlation, statistical anomaly detection, high-pass filtering and Kohonen and Bayesian techniques.

The Distributed Aircraft Maintenance Environment (DAME) project, funded by the Engineering and Physical Sciences Research Council (EPSRC) e-Science programme, brought industry and academia together to develop methods of using grid computing to improve engine diagnostics, [91]. DAME was undertaken in partnership with Rolls-Royce (who provided the aero-engine data for the diagnostic system), Data Systems and Solutions (DS&S) and Cybula (a York University spin-off company, managing the data storage technology in the project). The academic partners in the project were the Universities of York, Leeds, Sheffield and Oxford. This architecture removes the need for the expensive transfer of data to a central repository, since the aircraft data remains at the airport where it was downloaded and all processing occurs on site. Only limited amounts of relevant data need to be moved between the end-user and the data-nodes during a search/analysis process. The DAME architecture follows the following steps:

- Download raw data.
- Pre-process and archive the raw data.
- Train the Advanced Uncertainty Reasoning Architecture (AURA), which is a neural network based pattern matching engine. The architecture is designed to be generic, allowing any pattern matching engine to be used.
- Diagnose anomalies using case-based reasoning and integrating model-based fault detection and isolation approaches.
- Interface via web page.

The DAME projects aimed to combine two current analysis methods: ground-based trend analysis of civil engine snapshot data across large fleets with on-wing analysis or high frequency vibration data. This means that the analysis needs access to the large, high frequency data for the whole fleets, amounting to many terabytes of data.

Frith and Karvounis have investigated the use of modelling for the diagnosis and prognosis of helicopter engine gas path health, [92]. These models focus

on estimating power assurance parameters and detecting abnormal engine operations. Fuzzy Logic is used to aid the decision process of engine degradation levels. The capabilities have been demonstrated against simulated degradation applied to 30 flights of operational and trial data.

Yu *et al* have developed an approach to aircraft engine anomaly detection and diagnostics, [93]. This technique uses statistics and fuzzy logic for trend recognition, shift evaluation and failure classification. Figure 2.26 shows an example output of the approach, to identify outliers and shifts in engine performance data. The figure shows that single outliers are recognised, but are not classified as shifts. Only when a subset of new data does not statistically belong to a previous data window is it classed as a shift.




Figure 2.26: Example results from GE shift analyser prototype, [93]

Kobayashi and Simon have developed a hybrid ANN and GA technique for aircraft engine performance diagnostics, [94]. They describe some of the challenges involved in engine performance diagnosis:

“First, the only information available for health parameter estimation is the sensed parameters and the number of health parameters to be estimated is often greater than the number of sensors available. Another issue is that the sensor measurements are often distorted by noise and bias. Finally, the combined effect of system non-linearity and sensor selection may result in multiple health degradation scenarios producing similar measurement shifts.”

Koehl has used FEA to model a reference design mission for aero engine life usage monitoring, [24]. The FE models are used to generate comparison data, including gas path temperatures, metal temperatures and stresses. Sensor data can then be compared with the simulated model results to identify shifts, trends and anomalies. Safe crack initiation and propagation concepts are also used, and consideration is made for lost or missing data.

2.4.4 Prognostics

Historically, most effort has been expended on diagnosing current faults in a system. However, this approach may give little time to take action. Ideally, the operator wants to be able to plan corrective maintenance in advance and know when aircraft or engines will be out of service. To do this, systems need to be able to estimate when a fault will occur in the future, or how long a detected defect will take to propagate to a critical level, whereupon action must be taken. Often the first step is to diagnose a fault, then predict how it will propagate. The following paragraphs give some examples of research into techniques that attempt to predict the future health of aircraft systems.

Byington *et al* have used data-driven neural networks to predict remaining life for aircraft actuator components, [95]. The methodology employed for prognostics and health management is as follows: Flight control data is checked for data quality, diagnostic features are extracted using neural-networks and signal processing (RMS of high frequency pressure data and servo current at the system natural frequency) and then the data is classified to identify damage levels. The various fuzzy logic classification modes are then fused using Dempster-Shafer Combination and Bayesian Inference. Kalman filters are then used to assess the progression of the fault over time. The Remaining Useful Life (RUL) can then be computed, giving the user an estimate of time to failure. The diagnostic features were found to have a good correlation with both simulated and in-service electro-hydraulic servo valve degradation.

Lyer *et al* have developed a Decision Support System (DSS) for use in operational (real-time on-board) decision making with PHM specific data, [53]. Challenges arise from the large number of different information pieces upon which a decision maker has to act. Conflicting information from on-board and off-board PHM (OBPHM) modules, seemingly contradictory and changing

requirements from operations, as well as maintenance for a multitude of different systems within strict time constraints, make operational decision-making a difficult undertaking. Figure 2.27 illustrates the proposed DSS Prognostics Decision Support Mechanism (PDSM). The PDSM communicates with the OBPHM system and is made up of an Information Processor (IP) and a Multi-Objective Decision Support System (MODSS). The IP fuses different data sources, manages uncertainty and checks for data consistency, whilst the MODSS evaluates the data ‘fitness’ and ranks the results. The simulator module models different operational scenarios and produces evaluations for the performance objectives, along which each solution or alternative scenario is to be evaluated. For a given flight schedule assignment and its associated maintenance plan, an evolutionary search is then conducted to find solutions that are optimal to the stated performance objectives. The results can then be presented clearly to the user (operator, mission control, maintainer) via a Human Machine Interface (HMI).

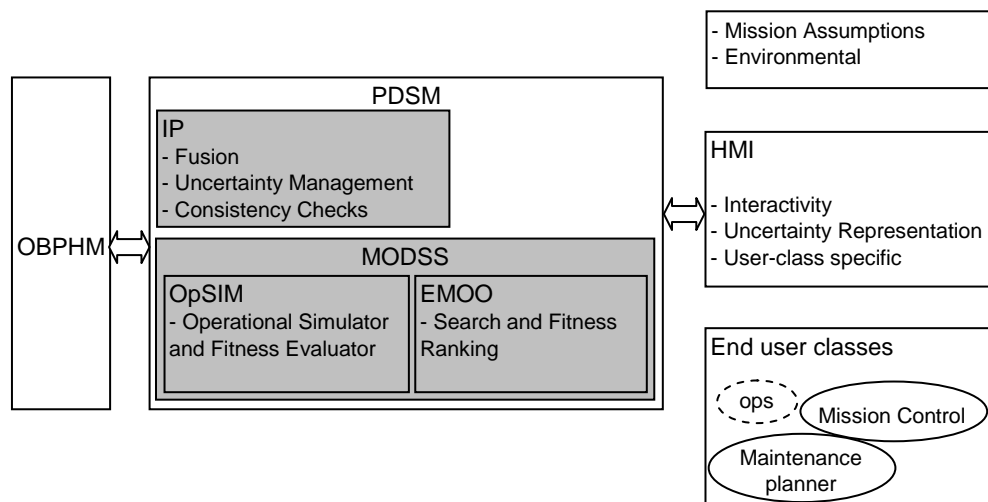


Figure 2.27: Decision Support System, [53]

2.4.5 Managing Uncertainty

Uncertainty representation and management is the Achilles heel of fault prognosis in CBM systems. No prediction will be one hundred percent accurate but usually some measure of how likely a prediction is to be correct can be computed. Khawaja *et al* investigated the use of ANNs for calculating confidence intervals for predictions, [96]. Figure 2.28 shows how the uncertainty in prediction methods can be represented. The historical data is certain and so represented by a solid line. Multiple predictions of how the fault

will progress can be computed and statistically compared to give a probability function of a given health measure and point in time.




Figure 2.28: Example of Representing Uncertainty in Prediction Methods, [96]

Roemer *et al* have used collaborative probabilistic and pattern recognition techniques to diagnose particular fault error patterns with associated confidence and severity levels, [45]. Figure 2.29 illustrates how data propagates away from the regions of normal operation within the fault detection threshold. At each point the probability of a fault being present can be evaluated.

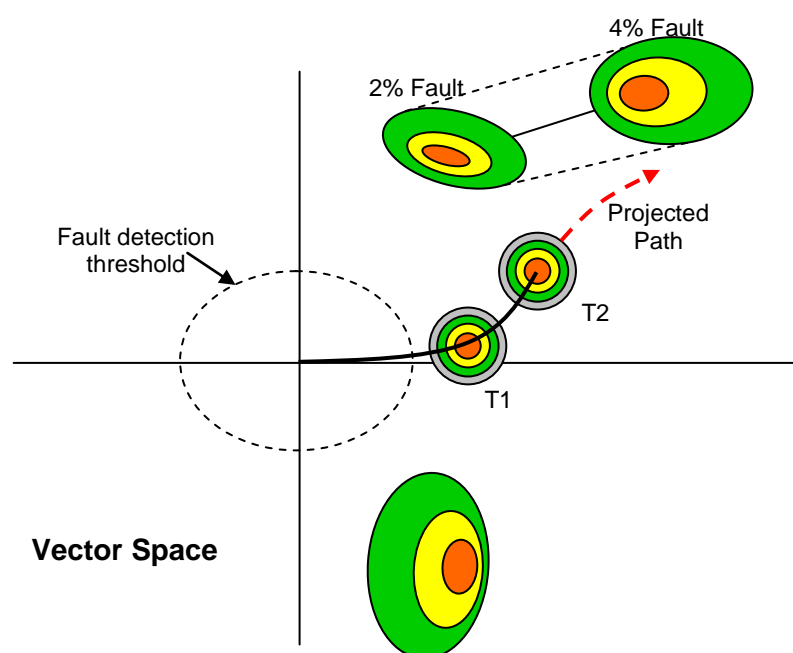


Figure 2.29: Probabilistic fault diagnostics process, [45]

Figure 2.30 and Figure 2.31 show other representations of the effects of uncertainty in prediction methods for CBM: both show how the time when the component would be triggered for replacement by the CBM method varies with improved prediction variance (i.e. a higher confidence with less uncertainty gives improved prediction variance which may lead to component life extension).




Figure 2.30: Prediction uncertainty schematic, [97]




Figure 2.31: Relationship between remaining life uncertainty and mission reliability, [53]

2.5 Summary

This chapter presented a review of some of the available literature with relevance to the study. Sources of modern aircraft data were described and the need for this data to be intelligently managed and analysed was presented. A brief introduction to AI and modelling techniques was also presented. Finally, examples from the literature were given for managing and analysing aircraft data for improved MAAAP using these AI and modelling techniques.

The literature shows the marked progress and advances that have been made in recent years in the areas of diagnostics and prognostics. However, there are still significant challenges:

- **Hard to validate:** For methods that are data driven (such as ANNs), the approach must be validated, proving that sufficient data has been used to cover all current and future aspects of operation and fault cases. The author was part of a team to update DEF STAN 00-970, [30], which includes guidance on the qualification of non-adaptive data driven approaches. Before this update, only fatigue meters were mandated for military aircraft. Validation of adaptive approaches (such as clusters that are updated as new data arrives), is even more difficult.
- **Not proven:** Many of the examples presented have not yet been fielded and so their applicability to in-service situations is not proven.
- **Insufficient fault cases:** Case based reasoners, such as Bayesian networks, require many cases, of all types of fault and healthy condition, in order to build accurate probability models. This may be time consuming and expensive to collect and would require input from experts in the field. For example, in the medical industry, expert systems have been developed to aid in patient diagnosis, but these have required many years of input from experienced doctors to take into account all factors when making a diagnosis.
- **Addressing data integrity issues:** All data driven methods also have to contend with the often poor integrity of the data from sensors. For example, operational strain gauges are affected by factors such as placement/orientation precision, bonding, age, environments (sand,

temperature, humidity) along with the inherent inaccuracy and resolution of the sensor. The same load values can produce different strains because of, for example, strain hysteresis, drift and strain gauge fatigue, [44].

In the next three chapters the work carried out by the author is presented. This study aims to address the shortcomings of data management identified from the literature. Fusion, from the outputs of the various techniques, has been accomplished using crisp logic, which is easier to validate than AI techniques. The methods have been applied to huge volumes of in-service aircraft data, further validating them. Processes have also been developed to address sensor issues.

Whilst attempting to address these challenges, the study's primary aim is to improve current methods, from manual investigations, towards the ultimate goal of automated predictive analysis. Three case studies are used to explore the way aircraft data can be more effectively managed, each using different techniques but all improving on the current state-of-the-art.

CHAPTER 3

3 HUMS CI INTELLIGENT HEALTH MANAGEMENT (IHM)

This chapter presents the analysis performed, as part of this study, to intelligently manage the CIs generated by onboard HUMS. The study used the Chinook helicopter HUMS as an example, but the techniques could be applied to any HUMS. The initial findings of the analysis were published in [12] and have also been reported to the MOD, [98, 99]. Note that throughout this Chapter tail numbers have been changed to make them unidentifiable and have been labelled as AC1 to AC40.

3.1 General


Whilst the MOD HUMS ground stations and maintenance/logistic systems contain a wealth of data, there is a growing interest in automatically analysing the data within these systems, [79, 100]. For example, large volumes of HUMS data are routinely downloaded from MOD Chinook helicopters. The raw vibration data are reduced to CIs that summarise periods of tachometer and accelerometer data, using time and frequency domain analysis. Nevertheless, the number of summarised CIs can be very large; typically about three thousand CIs are downloaded per flight, but many more could be configured. The data comes from as many as forty accelerometers, monitoring about 180 gears, bearings and shafts. Each vibration signal can be filtered to produce up to four pass-band filtered datasets; up to twenty CIs can be computed from each filtered dataset.

Appendix A gives details on the different signal processing techniques performed onboard the Chinook and the CIs that are computed from these vibration signals. The key CIs studied as part of this thesis are shown in Table 3.1. CIs are usually also referenced by prefixing the abbreviation of the data from which they were calculated: “Original signal” (SIG), “Enhanced signal” (ESA), “Enveloped signal” (ETE) and “Envelope spectrum” (FTE).

Table 3.1: Key Chinook CI names and descriptions

Name	Description
PP	“Peak-to-peak”. Difference between the max and min data values.
PK	“Peak”. Maximum absolute data value.
SD	Standard deviation of the data.
MN	Mean of the data.
MRC	“Maximum Rate of Change”. Max absolute difference of adjacent data values.
M6	Normalised 6 th moment of the data.
IMP	“Impulsiveness”. Normalised 4 th moment (kurtosis) of the data.
EB	“Band Energy”. RMS value of the data.
TON	Tonal energy of the band-passed envelope spectrum
WHT	White noise energy of the band-passed envelope spectrum
SO _n	Magnitude of nth shaft order vibration

Figure 3.1 shows the key Chinook drivetrain component groups monitored by HUMS. A detailed schematic of sensors fitted to monitor each component within these groups can be found in Appendix A. The appendix also describes the terminology for parts of bearings and planetary gear systems.

**Figure 3.1: Key Chinook drivetrain component groups monitored by HUMS**

A HUM system, designed to provide adequate health coverage by generating a large number of CIs, would provide a wealth of health information, but would require an intelligent, automatic, off-board, analysis capability to handle the large volume of CIs. A HUM system, designed to produce a small set of CIs, could reduce off-board logistic burdens but would not provide adequate health coverage.

HUMS automatically generates alerts when engine and gearbox vibration exceed their cautionary and threshold levels. These exceedances must be examined regularly and appropriate action taken, such as grounding the aircraft and stripping the part to identify damage. For CI values less than the cautionary/threshold levels, it would be very useful to know when they would reach these levels. The manual assessment of these exceedances, for each helicopter, each engine, each gearbox and each component, can be laborious and will divert the assessor's attention to data analysis tasks, rather than maintenance and airworthiness guideline tasks. Additionally, there may be value in data that has not triggered an alert, which is currently not being exploited. A capability to automatically analyse such a vast dataset is therefore needed.

Various studies have been undertaken to try to improve on HUMS [80] to [88]. However, many of these have not been applied to large in-service aircraft data, so it is difficult to assess their strengths. The following capabilities were therefore proposed, to be part of the HUMS CI IHM, to address the need to automatically analyse the large Chinook HUMS dataset:

- The processing should be performed automatically, allowing the user to assess the updated status of the entire fleet, starting from a high level warning report.
- The warning report should be colour coded to indicate the severity of the warning.
- The user should be able to navigate down further to see more detailed information including trend charts. The trend reports would include prognostic information, indicating the expected time at which a cautionary/threshold level will be reached. The reports would also include information describing the quality of trends and their types.
- It should be possible to configure the number of data points over which the trend would be evaluated and set conditions to restart trends: for example, the user might require the trend to restart after certain maintenance actions. The user should be able to set short-term and long-term trend reports for the same component.

3.1.1 The Features of a Condition Indicator

The measurements, from which an indicator is computed, can capture information about many factors, including the health state of the monitored component (Health). The only information required is the Health of the component and thus the effects of the other factors should be isolated. The other factors include:

Individuality: Differences between vibration characteristics of aircraft of the same type or similar components fitted on different aircraft. Individual characteristics can arise from factors such as age and manufacturing and maintenance tolerances.

Operation: For example, more vibration at a higher speed or a higher loading condition. The effect of operation can be reduced if the FDR data indicates the effect of the current operational condition.

Noise: Background random variations in vibration characteristics.

Sensor Health: For example, high vibration levels can be induced by accelerometer faults and not by monitored component faults.

It is highly desirable to eliminate the effects of these factors leaving only Health by, for example, statistical processes. Figure 3.2 is a schematic that shows the condition indicator ingredients for two fictitious aircraft, A and B, with eleven and eight flights respectively. The figure also shows the effect of ideal processes that have removed the effects of all undesirable factors. One of the objectives of this investigation was to develop such processes.




Figure 3.2: Schematic of data characteristics

3.2 Chinook HUMS Data

This section describes the Chinook helicopter HUMS data used in this study, incidents observed by the MOD and the current approach employed to analyse such incidents. More detailed information about the Chinook HUMS system, including schematic diagrams, accelerometers and monitored components, and descriptions of the CIs can be found in Appendix A.

3.2.1 Data Sourced

HUMS data were sourced from the MOD, covering nearly twenty thousand HUMS recording sessions (HRS), across forty Chinook aircraft, Figure 3.3 and Figure 3.4. The drive-train CI data contains approximately forty-six million CI values.




Figure 3.3: Number of HUMS recording sessions sourced per aircraft




Figure 3.4: Number of HUMS recording sessions sourced by month

3.2.2 *Current CI Management Approach*

HUMS alerts are generated when certain CIs exceed their predetermined thresholds. The thresholds are fixed throughout the fleet, but are specific to particular key components and indicators. Since a large number of automatically generated alerts have been identified by the MOD as false alarms, a sceptical approach is taken. The approach, adopted by the MOD to identify monitored component faults, as described by private communication, [101], is as follows:

- (i) Alerts would trigger further investigation by experts asking the following questions:
 - 1 Are the CI values realistic? If a CI value is unrealistically high (e.g. ETE_M6 value higher than one thousand) or suddenly triples, the investigator should suspect a sensor fault rather than a component failure.
 - 2 Is there a gentle upward trend of CIs over time (positive gradient)? A positive CI trend is likely to indicate a potential component fault more than a CI sudden jump to an exceedance.
 - 3 Are other CIs consistent with the alert? For example, if the signal peak to peak (SIG_PP) value is high, one would expect the signal standard deviation (SIG_STD) value to be also high.
 - 4 Does the raw time history have characteristics indicative of a fault? For example, periodic peaks at the meshing frequency, superimposed on data similar to that seen for a healthy component, could indicate damage to a gear tooth. On the other hand, an accelerometer fault could cause dramatic changes in the time history.
 - 5 Are the peaks in the FFT of the raw data consistent with the monitored component characteristics (geometry, speed, meshing frequency, etc.)?
- (ii) The calculation process performed on-board the aircraft, to calculate the CIs, would be emulated to confirm the correctness of the on-board computation.
- (iii) Oil from the component would be analysed to identify debris that could confirm the potential component fault.

- (iv) An independent tape-recording of data, from an accelerometer fitted close to the monitored component, could also be made during a ground run. The analysis of the recorded data would help to indicate the presence of a component fault or failed HUMS sensor.

The above investigations could confirm the presence of a component fault. In this case, a decision would be made to replace the component. A strip report about the removed component would be requested to gather concrete evidence of the component faults.

3.2.3 Incidents of Data Anomalies Observed by the MOD

Information was also sourced from the MOD concerning observed HUMS CI anomalies, [101]. Ten case studies are presented which cover a spectrum of faults that the system is intended to detect. Some of the observed anomalies were associated with mechanical failures that had occurred. Other CI anomalies had been under investigation by the MOD. The information was used to help understand the types of fault seen in the field and the associated HUMS values. For the first five cases charts are included which show the relevant CI value (amplitude) against data acquisition number (i.e. each data point). Where available, images of the damaged component are also shown. The last five cases are depicted using fleet charts of a CI value, where the x-axis is grouped by aircraft and then date.

3.2.3.1 Case 1

An increasing trend in the values of the impulsive CI 'M6' for the aircraft 'AC17' shown in Figure 3.5 was detected by HUMS threshold exceedances and so the bearing was removed from service on the 23rd of June 2003. The photo shown in Figure 3.5 depicts the damage found to the inner race of the bearing. There was also evidence of damage to the rolling elements.




Figure 3.5: Case 1 - Chart of M6 CI values indicating combiner bearing fault and photo of damage found

3.2.3.2 Case 2

Data from the aircraft AC25 showed an upward trend for M6. For this component monitored by M6, the trends in HUMS data, as shown in Figure 3.6, were not detected. However, the debris monitoring system identified a large chip and so the gearbox was removed in October 2001. Figure 3.7 shows photos of the failed inner race of the bearing.




Figure 3.6: Case 2 - Chart of M6 CI values indicating combiner bearing failure




Figure 3.7: Case 2 - Photos showing failed combiner bearing, [102]

3.2.3.3 Case 3

High M6 values alerted the MOD to a potential engine bearing fault for AC26. In July 2003, the component monitored by M6 was removed and inspected, but no fault was seen. So, the component was fitted to AC22. However, the CI values became even higher so the bearing was removed again in July 2004. The CI values, across both aircraft, are shown in Figure 3.8. The bearing was stripped and close inspection revealed a crack in one of the rolling elements, 4.2mm long and penetrating 3.2mm into the surface, Figure 3.9.




Figure 3.8: Case 3 - Chart of M6 CI values indicating engine input bearing fault




Figure 3.9: Case 3 - Photos showing crack in bearing rolling element

3.2.3.4 Case 4

The M6 used to monitor an aft transmission bearing had a step increase for AC 17, Figure 3.10. The bearing was eventually removed from service in 2005 and stripped. There was damage found to both the inner race and the rolling elements, Figure 3.11.




Figure 3.10: Case 4 - Chart of M6 CI values indicating aft transmission bearing fault



Figure 3.11: Case 4 - Photos showing damage to aft bearing inner race and rolling elements

3.2.3.5 Case 5

The M6 used to monitor a right-hand combiner bearing for AC20 had a step increase, Figure 3.12. Figure 3.13 shows the damage found to one of the rolling elements, after the bearing was removed and inspected. The first increase in the CI values was noted on the 7th of October 2004. On the 5th November an investigation was triggered. The component was eventually removed on the 15th of February 2005.




Figure 3.12: Case 5 - Chart of M6 CI values indicating combiner bearing fault




Figure 3.13: Case 5 - Photos showing damage to combiner bearing rolling element

3.2.3.6 Other Case Studies

Figure 3.14 to Figure 3.17 show CI values for some of the incidents identified by the MOD. Each chart shows M6 indicator values for a particular component across the whole fleet, grouped by aircraft and then sorted by date. These cases had M6 values which were significantly different to fleet values and so should be relatively easy to detect, for example using a fixed threshold.




Figure 3.14: Case 6 (AC2) and case 7 (AC19) - aft gear faults




Figure 3.15: Case 8 - Aft bearing fault




Figure 3.16: Case 9 - Combiner bearing fault




Figure 3.17: Case 10 - Engine bearing fault

3.3 HUMS CI IHM Architecture

Figure 3.18 shows the architecture for the intelligent analysis capability that was developed. The techniques considered were chosen due to their success in the literature and their availability and familiarity to both the author and the end user. The process flow is as follows: Firstly, the vibration data is acquired, processed and stored on-board the Chinook and downloaded into a ground database. After this, the IHM process audits the data to check for inconsistencies. Next, the data is pivoted from 3rd normal form into a format designed for speed of access. The data can then be reduced using PCA. Multiple threads of analysis are then performed to identify anomalies, including adaptive thresholds, trends and clustering. Then the outputs of the analysis are fused and reasoning is used to decide whether to alert the user. Finally, the results are summarised and presented in a user friendly way. Each block is described in more detail in the following sections.

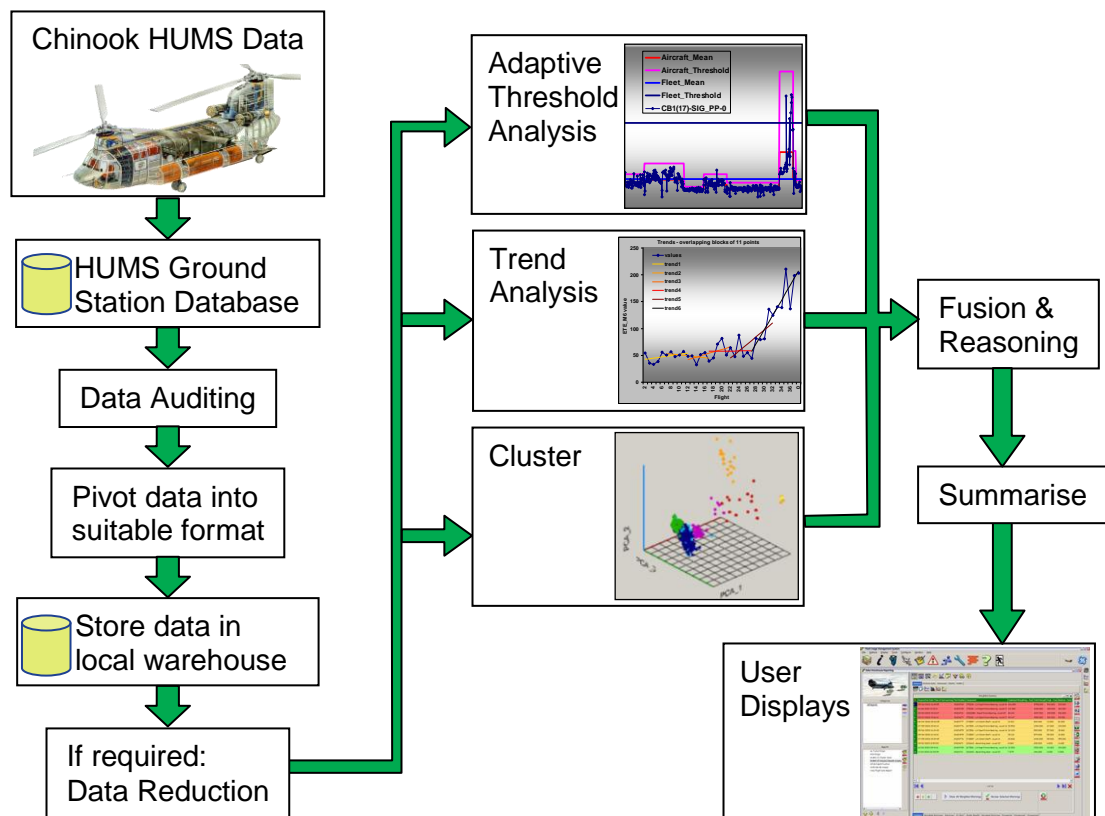


Figure 3.18: HUMS CI IHM architecture

3.4 Auditing the Data

HUMS data is likely to contain some erroneous values, which must be filtered out before further analysis can be performed. Erroneous data can include: unrealistically high CI values, flights with conflicting date/time stamps and monitored components incompatible with accelerometers. An audit algorithm was designed to detect and remove such data, generating reports showing the detected errors. The following sections provide a simplified description of the algorithm.

3.4.1 *Flights with Conflicting Date/Time Stamps*

Occasionally, two different HRS numbers or two different sequence IDs from the same aircraft are stamped with the same date/time. It is not known why this occurs. This inconsistency implies that two different flights occurred at exactly the same time. In this investigation, inconsistent data were filtered out and not considered further, to enable reliable trend analysis. The inconsistent data were identified and removed when the sequence ID and/or the HRS number varied for the same tail number and date/time stamp.

3.4.2 *Monitored Components Incompatible with Accelerometers*

Other examples of database inconsistencies are incompatible combinations of monitored component IDs and accelerometer location IDs. Table 3.2 shows some examples of such inconsistencies. The first example shows an inconsistent combination, where CIs were generated from an accelerometer located in the forward rotor transmission to monitor an input pinion gear on the aft rotor transmission. This implies that a sensor is monitoring a component that is over ten metres away, and so is obviously incorrect. The cause of such errors is unknown. This type of inconsistency was filtered out, using a list of compatible components and accelerometers, as shown in Appendix A1.3.

Table 3.2: Examples of monitored components incompatible with accelerometers

Monitored Component ID	Component Description	Accelerometer ID	Accelerometer Description
9	Aft Gear 1 - Input Pinion Gear	5	Forward Transmission
75	Combiner Bearing 1 – Left Hand Input Pinion Bearing	22	Aft Transmission
66	Synchronous Shaft #6	14	Left Engine Transmission
67	Engine Bearing 1 – Left Hand Input Pinion Bearing	6	Forward Transmission
105	Forward Bearing 3 - 2nd Stage Planet Bearing	19	Combining Transmission

3.4.3 Unrealistically High Values

It was found that the HUMS data contained unrealistically large values. For example, the average of the CI 'ETE_M6', over most components, was found to be approximately 50. A value above 200, for the combiner bearings, could indicate a fault. Nevertheless, the maximum value of the CI, shown in Figure 3.19 for component FB7(29), was found to be 415,000 and sixteen out of 660 values exceeded 5000. However, faults were not reported for these 16 cases. Therefore, it was concluded that excessively high values should be filtered out, before further processing, by setting an upper limit for each indicator. An excessively high CI value would indicate temporary short period sensor data corruption. Successive occurrences of excessively high values could indicate persistent sensor fault.




Figure 3.19: Example of high values needing to be filtered out

3.4.4 Data Gaps

The audit process was also designed to flag up potential data gaps. It was found that aircraft may only capture data for some of the required CIs. This may be due to the time taken to calculate the CIs, for example, if the aircraft has left the cruise condition before the complete capture is finished. If a sensor is not functioning correctly, its data may be missing altogether. If the tachometers are not recording correctly, all the synchronous data may be missing. The audit process checks for missing data and alerts the user. Missing data is identified where data is present for some, but not all, accelerometer locations, or when no data has been recorded at all for a long time period (e.g. three months).

3.4.5 Data Audit Results

The HUMS CI auditing capability was applied to about 46 million CI values and 423246 CIs of these failed the audit. Table 3.3 shows the number of CIs that failed for each audit type.

Table 3.3: Number of CIs to fail audit for each audit type

Audit Code	Description	Number Failed	Percentage of all CIs
1	Value too high.	26,997	0.06%
2	Component and accelerometer location incompatible.	376,254	0.81%
3	Date/time, HRS, sequence conflict.	19,995	0.04%
Total	All.	423,246	0.91%

The most common incompatibilities involved accelerometers 23 (aft transmission) and 6 (forward transmission). The CIs with the highest number of audit failures were the log average magnitude of 1st shaft order (FSA_SE1) and the normalised 6th statistical moment of the enveloped signal (ETE_M6).

3.5 Analysis Tools

This section describes some of the key analysis tools used to intelligently manage the Chinook HUMS CI data.

3.5.1 *Threshold Management*

One of the approaches for identifying potential faults from the CI values is to use thresholds. Setting threshold values, so that all faults are identified with minimum possible false alarms, is difficult. Since the statistics of the various CIs are not the same, a different threshold is required for each CI. For example, the peak to peak of a signal (SIG_PP) and its mean (SIG_MN) are likely to be of a different order of magnitude. The signal characteristics also vary with accelerometer locations due to factors such as loadings of neighbouring bearing and gears. As mentioned in section 3.1.1, age and manufacturing/maintenance tolerances can lead to significant individual aircraft characteristics.

Therefore, it would be beneficial to have three thresholds defined for each CI and component combination: aircraft-specific, fleet-wide and user-defined. This leads to too many thresholds to be defined manually and so automation would be required. The automation method developed in this study sets and updates thresholds using statistics of previous datasets. Figure 3.20 shows fleet-wide and aircraft-specific thresholds, set for one CI. The fleet-wide threshold will indicate that the data is abnormally high compared to the whole fleet. Whereas the aircraft-specific threshold will indicate that the data is abnormally high compared to the previous data for the same aircraft. This can give useful information in diagnosing the cause of the high values.

In this study, the thresholds were set at three standard deviations above the mean, since for a normal Gaussian distribution this means that over 99% of the data would be below the threshold. Some HUMS calculate these learnt thresholds, so, where available, these could be used rather than calculating them off-board.




Figure 3.20: Example of fleet and individual aircraft thresholds set using mean and standard deviations

Exceedances of fleet and aircraft thresholds could then be generated. Exceedances can also be ranked according to how far they were above the threshold. For this study, the number of standard deviations above the threshold indicated the severity of the exceedance.

3.5.1.1 Criteria for Threshold Update

Setting thresholds would require a sufficient number of CI values for the thresholds to be statistically significant. The thresholds should be regularly updated to ensure that they reflect the current fleet status and individual aircraft characteristics. Therefore, an automatic schedule should be implemented to update the thresholds. Criteria would be required to initiate the automatic update. For example, the criteria could be regular update times, such as every night, every weekend or once a month, or could be triggered after accumulating new database records. It could also be necessary to reinitiate the threshold setting process after certain events, such as maintenance actions or after a specified time interval. After such events, previously set fleet/aircraft threshold values should be used until a sufficient number of CI values become available to calculate statistically significant thresholds.

In this investigation, the thresholds were updated on a flight-by-flight basis, using all available data prior to the current flight. Regular updates at short time intervals would require retrieving few records from the database. Re-initiation,

for example, every three months, would require retrieval of three month's worth of data. It is worth noting that the re-initiation periods could overlap. Figure 3.21 illustrates the potential effect of different update criteria for the mean of a CI. The figure shows that all the methods settle down to a fairly constant value after a few months (about fifty data points). If the threshold is restarted every month a more realistic picture of the current data is represented but this may be modelling faults or may mask drifting trends in the data.




Figure 3.21: Effect of different methods of periodically updating means

A decision process would be required to select whether a CI value should be used in the update process. Values that significantly exceed a previously set threshold, for example, by more than one standard deviation, could be excluded from the update process to prevent large exceedance values influencing the normal, healthy CI statistics.

The threshold process also tracks how many points exceed the threshold in a row. This means that multiple spikes can be distinguished from consistent exceedances which are more likely to indicate a component fault.

3.5.2 Automatic Trending

A high indicator value on its own may be insufficient evidence of component failure. It could be down to erroneous values, sensor failure, processing failure, database error, etc. A trend of the indicators can give a more reliable indication as to whether an exceedance is due to component failure. Data trends can also provide prognostics. The trend can be used to predict if and when CIs, that have not reached a threshold, will exceed it.

Figure 3.22 indicates a number of different periods over which trends could be calculated and updated:

- The trend could use all available data for the particular component and CI. This method may hide medium/short-term trends.
- A specified number of flights may be used, such as the last six flights.
- Overlapping flights could be used.
- Shows the gradient of the trends for method a), updated every data point.

One or any combination of these may be used.

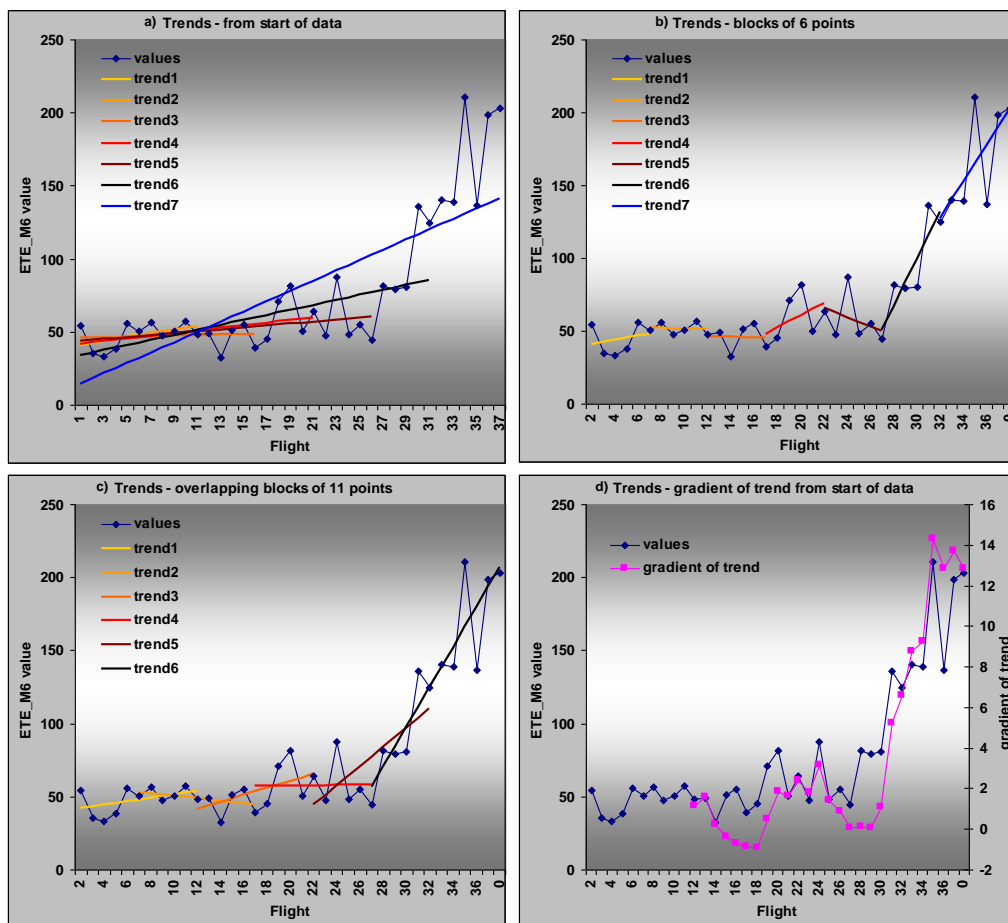


Figure 3.22: Examples of different linear trending methods

The trend gradient could also be used to generate alerts: negative, zero and very high gradients should be ignored, whereas reasonable positive gradients should generate alerts, with an associated severity level. For example: if a gradient of 5 indicates the most severe gradient (100%) and zero values, negative values and gradients over ten should be ignored. (3.1) may be used

to generate a percentage severity. Positive severities could generate an alert. (3.1) is illustrated in Figure 3.23.

$$Severity(\%) = \frac{5 - |5 - Gradient|}{5} \quad (3.1)$$




Figure 3.23: Graphical illustration of (3.1)

In the example shown in Figure 3.24, the CI has exceeded the fleet threshold (fleet mean plus three standard deviations) and the trend is increasing. Using (3.1) and the gradient 4.12, the severity is 82%.




Figure 3.24: Increasing Trend Example

Trends can also be calculated on CIs that have not crossed a threshold. If the gradient is positive, an estimate of the time to exceedance can be calculated. For example, Figure 3.25 shows a trend of signal peak to peak CI for the right hand combiner bearing. In this case, assuming that the trend remains the same and the aircraft flies with the same regularity, the CI will exceed the aircraft threshold in January and the fleet threshold in May.




Figure 3.25: Trend prediction against date

It may be more useful to have this information in terms of flying hours as shown in Figure 3.26. In this case, assuming that the trend remains the same, the CI will exceed the aircraft threshold after 110 flying hours and the fleet threshold in 340 flying hours. The two examples show the importance of trending against the correct time values.




Figure 3.26: Trend prediction against flight hours

A trending algorithm was implemented to calculate the gradient of a 'least squares linear best fit' trend line through the data. (3.2) to (3.4) can be used for calculating the gradient m and y-intercept c using the x and y values and their means \bar{x} and \bar{y} . In the developed algorithm, a short-term (by default fifteen data points) and long-term (by default thirty data points) trend can be calculated. Any data points that are above any of the thresholds are not used as part of the trend calculation. The gradients are stored to be used later as part of the overall decision process. The trend lines are also extended into the future to identify whether the data would cross one of the thresholds within a specified time.

$$y = mx + c \quad (3.2)$$

$$m = \frac{\sum (x - \bar{x})(y - \bar{y})}{\sum (x - \bar{x})^2} \quad (3.3)$$

$$c = \bar{y} - m\bar{x} \quad (3.4)$$

3.5.3 Principal Component Analysis (PCA)

HUMS typically stores twenty CI values but may store as many as 80 CIs for each component. These CIs are often well correlated since they are different statistical transformations on the same original vibration data. Using PCA this large number of CIs could be reduced to fewer principal components, with little loss of information, which would greatly reduce the processing, analysis, storage and time required. The eigenvalues computed by the PCA algorithm give an indication of the variation in each principal component, and hence can be a good guide to choose the appropriate number of components to use. For example, Figure 3.27 and Figure 3.28 show how 75 CIs (17 CIs for 4 passbands and 7 with no passbands) are reduced to 3 principal components, whilst retaining 99.5% of the CI variation. The patterns in the original CI data can also be seen in the 3 principal components.




Figure 3.27: Seventy-five CI values for the left combiner bearing.




Figure 3.28: First three principal components for the left combiner bearing.

3.5.4 Zone Analysis

For each unique combination of monitored component and accelerometer in the diagnostic configuration, CIs are computed. CIs for different monitored components, computed from the data of the same accelerometer, could have similar characteristics, such as those associated with a fault. For example, CB1 to CB7 are all monitored using accelerometer 16 and, therefore, the associated CIs could have comparable values. For example, Figure 3.29 shows that the CI values for CB1 and CB4 are comparable.




Figure 3.29: Combiner bearing component zone comparison

Grouping components and accelerometers into zones can add a further test to distinguish between potentially corrupt data and a genuine fault. Alerts generated from one CI and not from other CIs within a zone could indicate erroneous data, rather than a fault.

3.5.5 Cluster Analysis

The FUMS™ anomaly/novelty detector can efficiently recognise, in huge data sets, densities of clusters, anomalies and non-linear patterns having a wide range of sizes/orientations, [62]. The cluster algorithm is based on ISODATA [103] and uses k-means to attribute each value to a cluster. For m clusters in n dimensional space, the clusters are defined by m vectors of length n defining the centre and variance and m , $n \times n$ covariance matrices to define cluster rotations. The algorithm will find the optimum number of clusters within a specified range.

This tool can be used to separate the data into one or more clusters of normal data and other minority clusters of anomalous data. The minority clusters can be further analysed to determine if they contain consecutive flights from one aircraft, which could indicate a fault.

For example, Figure 3.30 shows how data for two principal component CIs from the left combiner bearing are grouped into two normal clusters and four minority (abnormal) clusters. By selecting the minority clusters, shown in the scatter chart, the associated table data records were highlighted. By examining the highlighted records it was found that they were all consecutive records from the same aircraft. These records were found to be acquired from the failed bearing, reported in [102].




Figure 3.30: Typical cluster results for a progressive fault case

Once identified, the clusters could be labelled according to the type of fault found. For example, in a perfect world, cluster 'Abnormal_3' might be labelled 'crack initiation', cluster 'Abnormal_1' labelled 'crack propagation' and cluster 'Abnormal_2' labelled 'imminent failure'. If further analysis showed that this anomaly was in fact a sensor failure, it could be labelled to that effect. In

practice, fault progression does not always fall in successive clusters. The severity of the fault could be based on the number of points from the particular aircraft in minority clusters and/or the distances of the minority clusters from the normal clusters.

Cluster analysis was applied to each component. Points belonging to large clusters were labelled as 'normal', whereas points in outlying clusters, containing fewer than a configurable proportion of the data points, were labelled as 'abnormal' (a default of five percent was used). Additionally, points lying far from the centres of any of the identified clusters (e.g. more than ten standard deviations) were recorded as 'unassigned'. The Euclidean distance, in standard deviations, from a cluster centre is known as the quality, [62]. For the cases where one or more CIs were missing for an acquisition date-time it was not possible to assign a cluster to the CIs. A 'null' cluster group was used in these cases.

3.6 Decision Process

The previous sections have described both the vast number of CIs to be analysed and the variety of tools that may be required to analyse them. To get more value from the results, each Health Judgement (HJ), produced by a tool, should be fused together and reasoning used to decide whether the combination of HJs indicate a fault. Care must always be taken when fusing data to ensure that information is not lost, and unforeseen fault cases are not masked. In order to be effective, the data and analysis should be presented in a user friendly and quick access manner. As far as possible, the results from the analysis should be summarised, giving the user the option to navigate down to more details and more information.

The HJs can be fused by an AI tool such as a Bayesian belief network, or crisp or fuzzy logic (rules) can be used to derive a final decision confirming or denying the presence of a fault. The following sections describe two decision algorithms that were developed, implemented and investigated by the author; the first was implemented in a final trial system that has been used by MOD engineers.

3.6.1 Rule-Based Decision Process

For this decision process simple logical rules are used to assign each warning a level - "High", "Medium" or "Low". The pseudo-code shown in Figure 3.31 gives an example of how the various HJs (exceedances, trend gradients, cluster type) can be fused to give a warning level:

```

IF(
  (num_exceed_fleet_in_row > 4 OR num_exceed_aircraft_in_row > 4 )
  AND (
    (percentage_exceed_fleet > 0 AND percentage_exceed_aircraft > 0)
    OR (percentage_exceed_fleet > 50 OR percentage_exceed_aircraft > 50)
  )
  AND (long_term_gradient > 0 AND short_term_gradient > 0)
  AND cluster_quality > 4
) THEN warning_level = "High"
ELSE IF (
  (num_exceed_fleet_in_row > 2 OR num_exceed_aircraft_in_row > 2 )
  AND (percentage_exceed_fleet > 0 OR percentage_exceed_aircraft > 0)
  AND (long_term_gradient > 0 OR short_term_gradient > 0)
  OR cluster_quality > 6
) THEN warning_level = "Medium"
ELSE warning_level = "Low"

```

Figure 3.31: Rule based decision logic

HJs were also combined by using a weight factor for each CI (for example, weighting the M6 values above PP, or pass band two over pass band one) and each HJ (for example, weighting clustering above trending) and then combining the weightings to give an overall score. Care was taken to provide robustness in the cases where CIs were missing for a particular acquisition and to ensure that valid warnings generated by the CIs that were present, resulted in appropriate weightings. It was also necessary to ensure that the total weight was factored by the number of relevant CIs for the component.

A 'percentage confidence' was then calculated for each summarised weighted warning to indicate the probability that a fault condition existed. This percentage was a function of the current total weight and a configurable number of previous weights. The confidence could aid in discriminating between possible sensor faults and actual mechanical faults, by assuming that a sensor / recording fault would be most likely to generate a very abrupt change, whereas an actual fault would show a gradual worsening and should generate a number of warnings from multiple CIs.

3.6.2 Fuzzy Logic Decision Process

An alternative decision method was investigated using fuzzy logic, [73]. The concepts of fuzzy logic are presented in section 2.3.5.2. The tools described in section 3.5 generate HJs, which were analysed using fuzzy logic to give the user an indication of the severity of any warnings. The HJs included: the percentage exceedance over the fleet and aircraft thresholds (%Fleet_Ex and %AC_Ex respectively); the number of exceedances in a row for both the fleet and aircraft thresholds (NEXROWFT and NEXROWAC respectively); the short and long-term gradient of a linear trend on the points preceding each exceedance (GRADST and GRADLT respectively); the rate of change of the gradient (GradRate); and the degree of membership of the normal cluster (Cluster DMEMB), [104].

For each HJ a number of fuzzy sets were defined and a fuzzy rule was used to describe the degree of membership of each set. Table 3.4 shows some examples of rules used to give the degree of membership of the fuzzy set (Fuzzy DMEMB). Figure 3.32 graphically illustrates some of these rules. In a final system design, the end user should have the ability to change the rules to achieve improved results.

Table 3.4: Fuzzy rules for HUMS CI IHM

Variable (x)	Set	Rule for Set Membership
NEXROWAC	Persistent	If($x > 3$ then 1 else if($x > 1$ then $(x-1)/2$ else 0))
	Not-Persistent	If($x < 1$ then 1 else if($x < 3$ then $1-(x-1)/2$ else 0))
%AC_Ex	None	If($x < -20$ then 1 else if($x > 20$ then 0 else $1-(x+20)/40$))
	Slight	If($(x > -20$ and $x < 20$) then $1-\text{abs}(x/20)$ else 0)
	Exceedance	If($x < 0$ then 0 else if($x < 20$ then $x/20$ else if($x < 200$ then 1 else if($x < 300$ then $1-(x-200)/100$ else 0))))
	Error	If($x < 200$ then 0 else if($x < 300$ then $(x-200)/100$ else 1))
GradRate	Decreasing	If($x < -0.1$ then 1 else if($x < 0$ then $-(x*10)$ else 0))
	Zero	If($(x > -0.1$ and $x < 0.1)$ then $1-\text{abs}(x*10)$ else 0)
	Increasing	If($x < 0$ then 0 else if($x < 0.1$ then $x*10$ else 1))
Cluster DMEMB	Normal	If($x < 5$ then $1-x/5$ else 0)
	Abnormal	If($x < 5$ then $x/5$ else if($x < 20$ then 1 else if($x < 25$ then $1-(x-20)/5$ else 0))))
	Extreme	If($x < 20$ then 0 else if($x < 25$ then $(x-20)/5$ else 1))

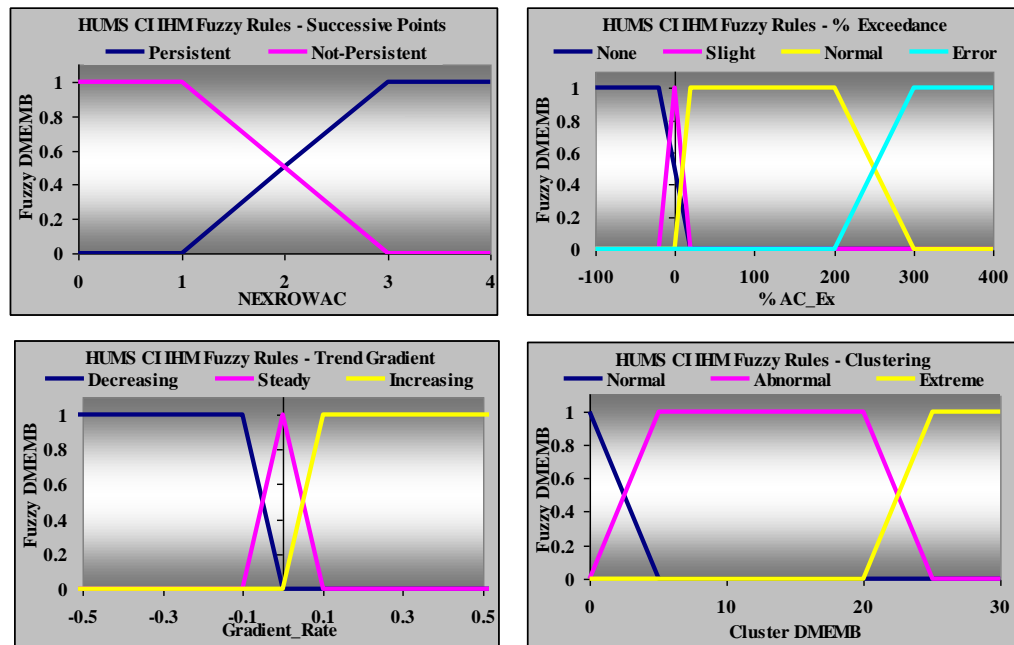


Figure 3.32: Charts of example fuzzy rules for HUMS CI IHM

The following fuzzy rule was then used to evaluate the membership of a component failure alert:

$$\begin{aligned}
 \text{Component_Alert} = & (\text{Persistent}_{\text{NEXROWFT}} \text{ OR } \text{Persistent}_{\text{NEXROWAC}}) \\
 & \text{AND } (\text{Exceedance}_{\% \text{Fleet_Ex}} \text{ OR } \text{Exceedance}_{\% \text{AC_Ex}}) \\
 & \text{AND } \text{Positive}_{\text{GRADST}} \text{ AND } \text{Positive}_{\text{GRADLT}} \\
 & \text{AND } (\text{Abnormal}_{\text{Cluster DMEMB}} \text{ OR } \text{Extreme}_{\text{Cluster DMEMB}})
 \end{aligned}$$

To perform this rule in fuzzy logic, the OR function becomes the maximum of the values (MAX) and the AND function becomes the minimum of the values (MIN):

$$\begin{aligned} \text{Component_Alert} = \text{MIN}(& \text{MAX}(\text{Persistent}_{\text{NEXROWFT}}, \text{Persistent}_{\text{NEXROWAC}}), \\ & \text{MAX}(\text{Exceedance}_{\% \text{Fleet_Ex}}, \text{Exceedance}_{\% \text{AC_Ex}}), \\ & \text{Positive}_{\text{GRADST}}, \text{Positive}_{\text{GRADLT}}, \\ & \text{MAX}(\text{Abnormal}_{\text{Cluster D MEMB}}, \text{Extreme}_{\text{Cluster D MEMB}})) \end{aligned}$$

The user requires an indication of the confidence of the alert generated, so a defuzzification rule was used to give a percentage confidence as to whether the alert is related to a component failure.

It was found that the Fuzzy logic approach did not give significant improvements in fault isolation and introduced more complex configuration that the end-users would find more difficult to adjust. As a result the rule-based method described above was implemented in the final trial system.

3.7 Implementation

It is important that developed methods can be deployed to engineers at the front line and that they facilitate intuitive and efficient access of pertinent information, to enable them to decide whether to take actions, such as grounding aircraft. Therefore, the algorithms described in the previous sections were implemented in the FUMS™ framework using a combination of C++ and the internal scripting language. FUMS™ also allows User Interfaces (UI) (referred to as reports) to be designed by a user without software rewriting. Reports can contain various items, including: data tables, charts, forms (a collection of buttons, edit boxes, images, etc.) and tools (which allow data processing). Report items are interactive and allow the user to filter data, zoom charts, select options for processing, select data in a table, which will then also be highlighted in connected charts, etc. Reports were designed and configured by the author to display the results of the automated analysis, starting with a high level summary and allowing the user to easily navigate down to more detail. These reports are described below:

3.7.1 Configuration

Reports were designed to allow the user to easily set the configuration for the IHM application for optimum performance. The report allows setting the CI and process weights, selecting the number of acquisitions to use for short and long-term trends and how far to extrapolate them, and defining cluster and adaptive threshold parameters such as the minimum and maximum number of clusters, how many acquisitions are required before the statistical thresholds are applied and the number of standard deviations above the mean to use for thresholds.

A report was also developed to allow the user to reset the thresholds for a selected aircraft (or all) and a selected component (or all). The tool was designed to keep a log, which includes the user name and comments, and the date and time when the threshold was reset.

In addition, a report was configured that displays the results of the audit algorithm, giving the user quick access to a list of CIs that were rejected and why.

3.7.2 Summary Display

A report was designed and implemented to display the IHM analysis results with tabs for different levels of detail. The Summary tab summarises all the warnings, showing how many were generated at each level (“High”, “Medium” or “Low”), for each aircraft and for each component, Figure 3.33. The rows are ranked and coloured according to the confidence value. The user can specify the colour coding methodology but, by default, any values with a confidence of more than 50% are coloured red, a confidence of less than 10% is coloured green and all others are coloured amber. Warnings are also tagged according to the date and time that they were generated (i.e. when the process was run). An entry is also made, even if no warnings were generated, to indicate that the data was processed correctly on a particular day.

Each warning can be marked as ‘reviewed’ once it has been analysed by the user. The reviewed warnings can be viewed using the “Reviewed Warnings” tab. The user will be prompted to enter a reason for moving the warning to the reviewed tab. A log is kept with the reviewed warning showing user name, the date and time and the reason it was reviewed.

	Acquisition Date	Time of last warning	Tail Number	Component	Combined Weighting - Sum	Total Weight	High - Sum	Medium - Sum	Low - Sum	Percentage Confidence	Last Value	High - Delta
1	25/06/2003	16:29:55	AC001	[75]CB1 - LH Input Pinion Bearing - accel 16	126.059	5758.000	531.000	319.000	91.000	0.000		531.000
2	16/06/2003	10:15:11	AC004	[75]CB1 - LH Input Pinion Bearing - accel 17	113.088	6149.000	399.000	481.000	479.000	0.936		399.000
3	20/02/2003	19:24:07	AC010	[100]A88 - Input Pinion Bearing - accel 20	80.441	4197.000	339.000	301.000	131.000	20.804		339.000
4	05/07/2003	19:06:31	AC006	[75]CB1 - LH Input Pinion Bearing - accel 17	29.147	1580.000	106.000	93.000	177.000	45.973		106.000
5	06/04/2003	05:32:05	AC007	[74]EB7 - LH Clutch Shaft - accel 14	12.412	844.000	13.000	81.000	125.000	11.104		13.000
6	27/05/2003	15:40:14	AC007	[67]EB1 - LH Input Pinion Bearing - accel 14	20.500	1394.000	67.000	130.000	169.000	0.000		67.000
7	25/06/2003	16:34:30	AC005	[67]EB1 - LH Input Pinion Bearing - accel 14	32.324	590.000	57.000	15.000	24.000	0.000		57.000
8	25/06/2003	16:34:01	AC005	[74]EB7 - LH Clutch Shaft - accel 14	35.118	579.000	59.000	14.000	12.000	0.000		59.000
9	27/05/2003	15:39:45	AC006	[74]EB7 - LH Clutch Shaft - accel 14	20.882	1420.000	59.000	155.000	164.000	0.000		59.000
10	19/06/2003	10:53:09	AC001	[10]A62 - Bevel Ring Gear - accel 20	9.060	208.000	4.000	11.000	64.000	0.000		4.000
11	16/06/2003	09:19:21	AC005	[67]EB1 - LH Input Pinion Bearing - accel 14	22.500	1530.000	43.000	193.000	192.000	0.000		43.000
12	11/03/2003	20:38:35	AC011	[10]A62 - Bevel Ring Gear - accel 20	7.875	246.000	0.000	0.000	111.000	0.000		0.000

Figure 3.33: HUMS CI IHM display: Summary tab

3.7.3 Warnings Display

Depending on which line is selected in the Summary tab, all the warnings, including the HJs generated for the selected aircraft and component, are displayed on the Warnings tab. If no selection was made on the summary, all the HUMS CI warnings will be displayed. Table 3.5 shows all the HJs and other information that is displayed on the Warnings tab.

Table 3.5: The Fields for Warnings

Column Name	Description
generated date time	The time at which the data was processed and the warning generated.
aircraft_serial_number	The tail number of the aircraft, as in the input data.
sequence_id	The sequence_id, as in the input data.
hrs_number	The hrs_number, as in the input data.
display_date_time	The display_date_time, as in the input data.
value	The value of the CI, as in the input data.
monitored_component_id	The monitored_component_id, as in the input data.
accel_location_id	The accel_location_id, as in the input data.
dtrain_indi_type_id	The dtrain_indi_type_id, as in the input data.
magnitude	The magnitude, as in the input data.
filter_passband_id	The filter_passband_id, as in the input data.
num_exceed_fixed_in_row	The number of sequential CIs that are above the fixed threshold, as set in the HUMS database.
num_exceed_fleet_in_row	The number of sequential CIs that are above the automatically generated fleet threshold (based on statistics for the whole fleet).
num_exceed_aircraft_in_row	The number of sequential CIs that are above the automatically generated aircraft threshold (based on statistics for the particular aircraft).
fleet_thresholds	The fleet threshold value used.
aircraft_thresholds	The aircraft threshold value used.
percentage_exceed_fixed	The percentage by which the CI is above the fixed threshold.
percentage_exceed_fleet	The percentage by which the CI is above the fleet threshold.
percentage_exceed_aircraft	The percentage by which the CI is above the aircraft threshold.
long_term_gradient	The gradient of the linear trend through the previous points that have not generated warnings.
long_term_y_intercept	The y-intercept of the linear trend through the previous points that have not generated warnings.
long_term_gradient_normalised	The gradient of the linear trend normalised by dividing by the mean.
short_term_gradient	The gradient of the linear trend through the previous points that have not generated warnings.
short_term_y_intercept	The y-intercept of the linear trend through the previous points that have not generated warnings.
short_term_gradient_normalised	The gradient of the linear trend normalised by dividing by the mean.
cluster	The cluster the point has been assigned to (Abnormal or Normal).
cluster_quality	The distance of the point from the centre of the cluster in terms of standard deviations.
points_to_exceed_fleet_short	The number of points into the future required for the short-term trend to exceed the fleet threshold.
points_to_exceed_aircraft_short	The number of points into the future required for the short-term trend to exceed the aircraft threshold.
points_to_exceed_fixed_short	The number of points into the future required for the short-term trend to exceed the fixed threshold.
points_to_exceed_fleet_long	The number of points into the future required for the long-term trend to exceed the fleet threshold.
points_to_exceed_aircraft_long	The number of points into the future required for the long-term trend to exceed the aircraft threshold.
points_to_exceed_fixed_long	The number of points into the future required for the long-term trend to exceed the fixed threshold.
warning_level	The warning level calculated by the decision process (High, Medium or Low).

3.7.4 CI Display

This page shows a chart plotting all the CI values, the thresholds, the selected point that had exceeded the thresholds and short and long-term trends from the selected point. A dotted line is used to show the trends extended into the future, to help visualise when the thresholds would be exceeded. Where a fixed threshold is available, this is also plotted on the chart. The drop-down boxes, at the top of the display, show the details of the selected warning from the Warnings tab. These can be changed directly to look at another CI for the same component, or at a CI of a different aircraft, or a CI of a different component. This also allows the user to look at data that has not generated any warnings. The user can also choose whether to look at data only for the selected aircraft, or for the whole fleet.

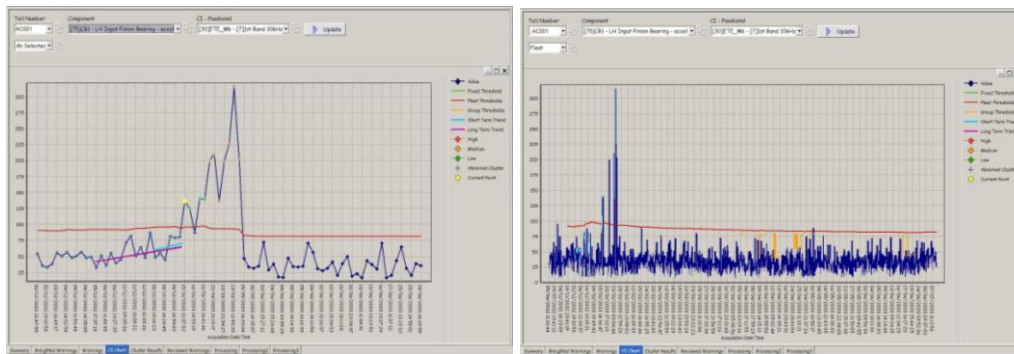


Figure 3.34: HUMS CI IHM display: CI Charts tab – One aircraft (left) and the whole fleet (right)

3.7.5 Cluster Results Display

All available data, for the particular component selected on the Warnings tab, is shown in the cluster results page. The data is converted into three principal axes (using PCA) and plotted on two three-dimensional scatter charts, Figure 3.35. One scatter chart is coloured by the cluster name, the other by the aircraft tail number. At the bottom of the tab a form is available for the user to rename individual clusters. This updates the display and saves the new name for future runs of the process. If insufficient points have been analysed for clusters to be identified, the user will not be able to rename clusters.

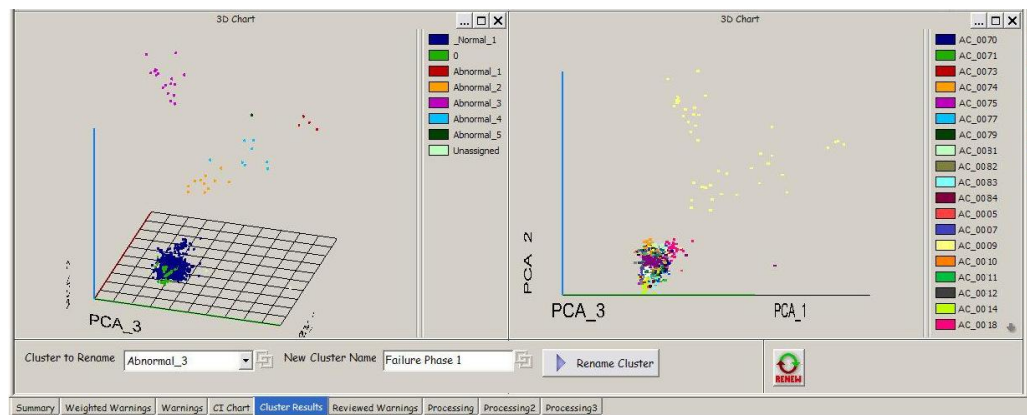


Figure 3.35: HUMS CI IHM display: Cluster Results tab

3.8 Assessment

In order to assess the enhanced health monitoring capability that was developed, it was run on a subset of data and manually inspected. The results were compared to the fault cases described in section 3.2.3 and to the “Chinook Query Register”, which lists Chinook HUMS related investigations made by the MOD.

Of the 46 million CI values sourced, spanning the operation of 40 aircraft between 2002 and 2006, a subset was selected to make the comparisons easier. The Combiner Bearings (CB) and Engine Bearings (EB) were selected, since over half the fault cases are for these components. In addition, only the impulsive indicator (ETE_M6) and energy indicators (total energy (EB), tonal energy (TON) and white noise energy (WHT)) were used, since these are most sensitive to bearing faults. This reduced the data set of CIs, which had passed the audit, to 1.8 million values.

It is also worth noting that there was a configuration change in the onboard HUM system in April 2004. This change grouped some of the bearings, which had the same CI values. For example, there are three input shaft combiner bearings (CB1, CB2 and CB3); CIs, such as M6 and EB, have exactly the same values for all these bearings, but the pattern match CIs are different for each. Originally, these values were duplicated and associated with each bearing. However, after the configuration change, they were grouped into a new component called CB123. For this analysis, for the first configuration, only the first component was used (i.e. CB1), otherwise the group component was used (i.e. CB123). For the subset selected, there are three combiner input pinion bearings (CB123), four combiner gear bearings (CB4567), three engine input pinion bearings (EB123) and four engine output gear bearings (EB4567), each for both the left and right hand engines. Hence, a total of twenty eight components were analysed, broken down into eight groups.

Processing the data took five and a quarter hours, which equates to less than twenty seconds to process each day’s worth of data (PC spec: Dell Precision PWS390 with Intel® Core™2Quad CPU (Q6600@2.4GHz) and 3.5GB RAM). The processing included: pivoting the data to a suitable format, running the cluster algorithm, calculating fleet and individual thresholds and identifying

exceedances, calculating short and long term linear trends, and then calculating an overall confidence of a fault being present.

Once processed, a total of just over 60 thousand warnings were generated (approximately 3% of the data) covering the two-year period (10843 high, 20571 medium and 28820 low). This works out to about 40 warnings per day. Note that a particular fault is usually characterised by many warnings, since historical data anomalies may be present for long time periods and many different CIs and related components will generate warnings.

The summary that is first presented to the user is grouped by aircraft and component, and ranked by the calculated confidence that there is a fault. The summary had a total of 176 out of a possible 320 rows populated, and 21 were given a confidence of over 50%. The key fields of the summary table are presented in Table 3.6. The table shows all the results with a confidence over 50% (coloured red) and also results with a confidence under 50% that were associated with known fault cases, or MOD query register entries.

Out of the ten fault cases described in section 3.2.3, five related to the components and time period covered by this analysis and all were alerted by the system. In addition, all of the related MOD query register entries were alerted by the system. On inspection, only three of the results with a confidence of over 50% did not look like a possible fault, but rather were due to very low initial values causing the learnt threshold to be too low. Figure 3.36 shows the worst-case receiver operating characteristic (ROC) curve for the results. It is based on a threshold on the confidence, assuming that only events identified by the MOD were real faults. It shows that 83% detection can be achieved with only a 13% false positive rate.




Figure 3.36: HUMS CI IHM ROC curve

Table 3.6: Summary of results from the IHM process

#	Aircraft	Component	Confidence	Case #	In Register
1	AC10	LH-CB123	99		
2	AC14	RH-CB1	99		y
3	AC14	RH-CB4	99		y
4	AC17	LH-CB1	99	1	
5	AC17	LH-CB4	99	1	
6	AC6	RH-CB1	99		
7	AC38	RH-CB123	95		
8	AC31	RH-EB4567	82	threshold issue	
9	AC31	RH-EB123	72	threshold issue	
10	AC17	RH-EB123	71		
11	AC2	LH-CB123	70		
12	AC17	RH-EB4567	70		
13	AC29	RH-EB4567	69		
14	AC6	RH-CB123	68	9	
15	AC33	LH-EB1	64		
16	AC7	LH-EB123	62		
17	AC33	LH-EB4	62		
18	AC10	LH-EB123	58		
19	AC32	LH-CB123	55	threshold issue	
20	AC1	RH-EB123	54		
21	AC14	RH-CB123	53		y
...					
40	AC22	LH-EB123	28	3	
...					
46	AC26	LH-EB4	26	3,10	
...					
48	AC33	LH-CB123	25		y
49	AC26	LH-EB1	25	3	
...					
80	AC20	RH-CB123	16	5	y
...					
89	AC2	RH-CB123	14		y
...					
176	AC8	LH-EB4567	0		

Each of the results presented in Table 3.6 is described in detail in Appendix A1.10, showing charts of all the related CI data and descriptions of related MOD findings. Representative case studies are presented in the following sections that illustrate some of the successes and shortcomings of the developed IHM process. All the charts shown are taken from the developed system that was delivered to the MOD.

3.8.1 HUMS CI IHM Results - Case 1

For the aircraft 'AC10' and the left hand combiner bearing group, all the energy CIs were found to be above the calculated fleet threshold for all available data. The data for the energy band CI for passband 4 is shown in Figure 3.37. The values are nearly twice as high as the values for Case 9 described in section 3.2.3. It is therefore likely there was some damage to a combiner bearing. However, the values are fairly constant for over a year which indicates that any damage present is not progressing significantly. The values are highest for February to April 2005. At the end of this period it is likely that some maintenance was performed, as there is a gap in the data between the 15th and 26th of April, but no maintenance information was available to confirm this. The M6 CI values were found to only have a few spikes slightly above the calculated fleet threshold causing a few low level warnings. It seems sensible that the algorithm has given this data a high confidence over 99% as these values should be further investigated by the MOD.




Figure 3.37: HUMS CI IHM chart of aircraft 'AC10' and component 'LH-CB123' showing the CI 'FTE_EB' for passband 4

3.8.2 HUMS CI IHM Results - Case 2

For the aircraft 'AC14' and the right hand combiner bearing group, all the energy CIs were found to be above the calculated fleet threshold from the start of the data in March 2003, until the last but one point on the 29th of January 2004. The last data point is normal (on the 6th of February 2004), as shown in the left hand chart of Figure 3.38. In April 2003 there were found to be a large number of very high spikes in the M6 data, as shown in the right hand chart of Figure 3.38. These were investigated by the MOD, as recorded in the query register, in July and again in August 2004. The conclusions were: *"Analysis confirms impulsive content but cannot be attributed to any bearing frequencies in the Combiner Xmsn. Consider this still to be spurious arisings"*. The data for combiner bearing component group 'CB4567' had similar values. This shows that the system correctly classified this data as a fault case and gave it a high confidence of 99%.




Figure 3.38: HUMS CI IHM chart of aircraft 'AC14' and component 'RH-CB1' showing the 'FTE_EB' CI for passband 4 (left) and the 'ETE_M6' CI for passband 3 (right)

3.8.3 HUMS CI IHM Results - Case 3

For the aircraft 'AC17' and the left hand combiner bearing group, all the energy CIs were found to be above the calculated fleet threshold, from the start of the data on the 2nd of December 2002 until the 23rd of January 2003. There was also found to be a step down in energy CI values for passband 3 between the 6th and 11th of December 2002. The M6 values ramp up from about the 16th of January 2003. Figure 3.39 shows these two CIs for both passband 3 and 4. This data relates to fault case 1, presented in section 3.2.3.1. The combiner bearings were replaced and normal operation can be observed in the data from May onwards. The CB4 group was found to have similar characteristics. The algorithm has given this known fault data a high confidence of 99%, thus demonstrating that known faults can be detected.

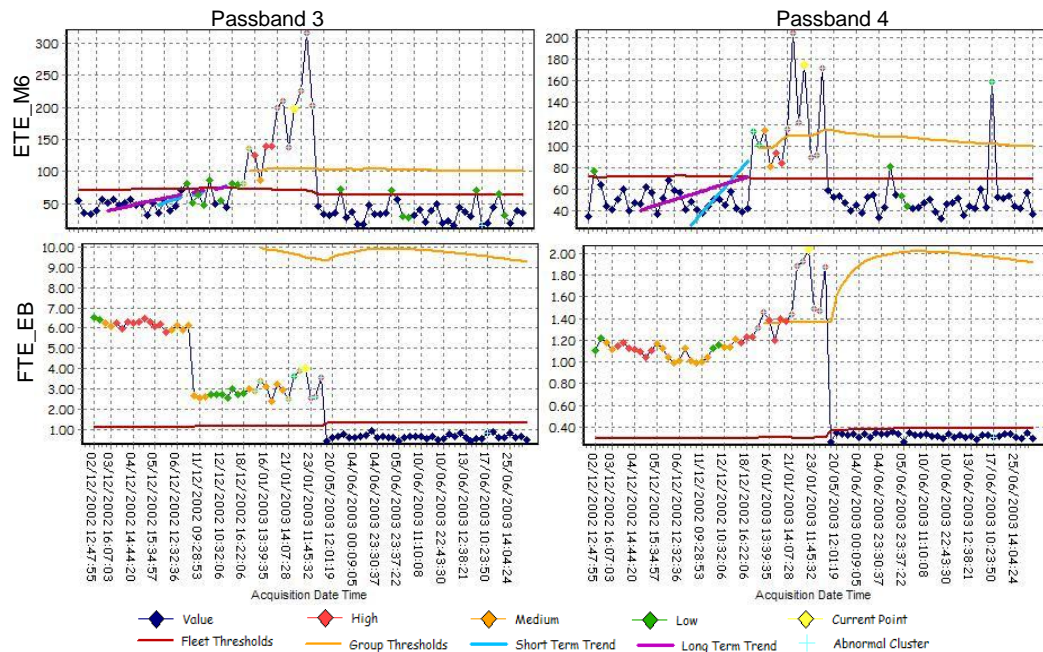


Figure 3.39: HUMS CI IHM chart of aircraft 'AC17' and component 'LH-CB1' for the CIs 'ETE_M6' and 'FTE_EB' for passbands 3 and 4

3.8.4 HUMS CI IHM Results - Case 4

For the aircraft 'AC31' and the right hand combiner bearing group, a large number of acquisitions were found to be above the calculated fleet threshold. However, the fleet threshold ramps up fairly quickly over this period and then the CI values are marked as normal, as shown in Figure 3.40. This doesn't look like a real fault. Further investigations revealed that there were about 380 acquisitions from AC15, AC23 and AC3 in 2003, which had much lower means and standard deviations than later data, causing the fleet threshold to be too low. If this analysis was run regularly, this would be spotted early on and the fleet statistics could be reset to start at the end of January 2004. This demonstrates the importance of the period defined before which the statistics can be relied upon. In this case, it seemed that three aircraft had abnormally low vibration which skewed the initial statistics. This problem was found to occur for two other cases with confidence scores over 50%. Resolving this threshold setting problem would mean that these cases would no longer be flagged to the end user.




Figure 3.40: HUMS CI IHM chart of aircraft 'AC31' and component 'RH-EB123' for the CI 'FTE_EB' for passband 4

3.9 HUMS CI IHM Conclusions

In this section, the problem of analysing huge volumes of Chinook HUMS CIs was described. A suite of techniques, including data cleaning, automatic adaptive threshold evaluation, prognostics using linear trends and data mining by clustering were presented, that have been developed, implemented in the FUMS™ framework and used to automatically and intelligently manage this data and present it to an expert in the field to make airworthiness judgments. Comparisons have been given between the practices currently employed by the MOD and the results of the techniques developed by the author. The tools developed would make analysing anomalies easier for the MOD.

The assessment of the techniques showed that all the related fault cases and query register entries were identified, with the lowest confidence being about 14% and all but two being above 25%. Of all the other events identified with a confidence of more than 50%, only three seemed to be data that is unlikely to relate to a fault. High level warnings were generated well before any of the MOD investigations were initiated, indicating that the methods could alert the users to issues earlier. For example: The first alert for case 1 was issued 52 days before the component was replaced, the first alert for AC14 RH-CB123 was 60 days before the MOD started investigating it, the first alert for case 3 was 49 days before the MOD started investigating it, the first alert for AC33 LH-CB123 was 7 days before it was investigated by the MOD.

The assessment could be improved by quantifying the confidence level that would have been generated at different points in time for each event, since the analysis presented only gave an overall confidence at the end of all the data. In addition, it would improve the analysis if continuous periods of alert could be grouped, since the current summary may group multiple faults together. It would also have been beneficial to have more maintenance information, to confirm whether the other events identified actually related to faults.

The analysis could also be improved by identifying long gaps in the acquisition data or using maintenance data to split the data. The methods could also be enhanced by attempting to fuse information from more data sources to give more confidence in fault predictions. For example, the Chinook has an oil debris monitoring programme, frequent FDR downloads and maintenance data stored in a logistics database.

The next chapter looks at the problem of identifying engine performance degradation. As with the Chinook, the aim of the analysis is to identify anomalous data for further investigation, to allow events to be identified automatically and to give an improved lead time on pending issues.

CHAPTER 4

4 ENGINE PERFORMANCE DEGRADATION MODELS

This section presents analysis carried out to identify engine performance degradation using measured aircraft data. Some of the work was presented at the Institute of Electrical and Electronics Engineers (IEEE) aerospace conference in 2006, [13], and to the MOD in 2004, [105].

4.1 General

Gas turbine performance deteriorates during operation due to degradation of gas path components. The most common causes of degradation are compressor fouling, increased blade tip clearance due to wear and erosion, labyrinth seal damage, foreign and domestic object damage, hot end component damage and corrosion. These physical faults result in changes in gas turbine thermodynamic performance, which lead to changes in observable engine parameters, such as temperature, pressure, rotational speeds and fuel flow-rate. The degraded performance, reflected in the observable engine parameters, can be used to detect component faults. Gas Path Analysis (GPA) was originally coined by Urban in 1975 as a technique for using models of engine performance, compared with engine data, to monitor engine performance, [106]. Figure 4.1 illustrates how physical faults cause degraded performance, which can be detected by changes in measured parameters.

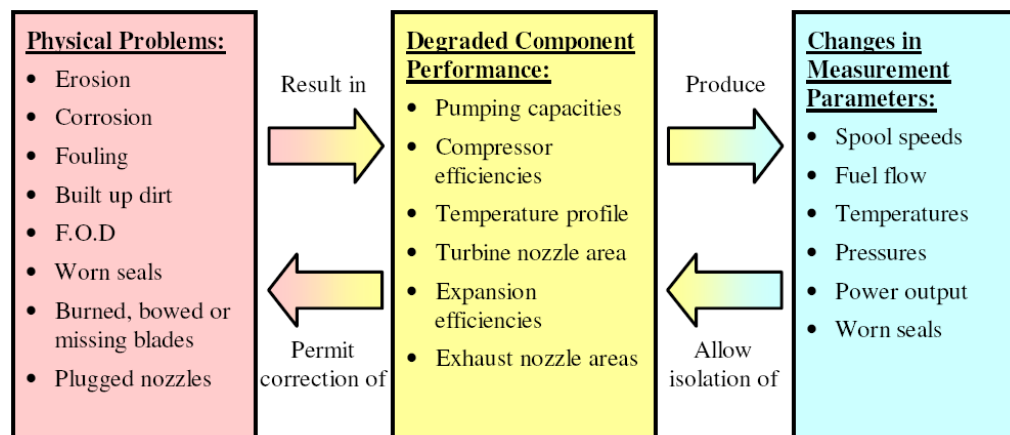


Figure 4.1: Gas turbine fault diagnostics approach, [106]

In 2002 Li reviewed performance-analysis-based gas turbine diagnostics, covering data validation, linear and non-linear, AI (ANNs, GA, Expert Systems, FL) and transient based models, [107]. Li's review also indicated the relative computation speeds and complexities of the various techniques, as shown in Figure 4.2.

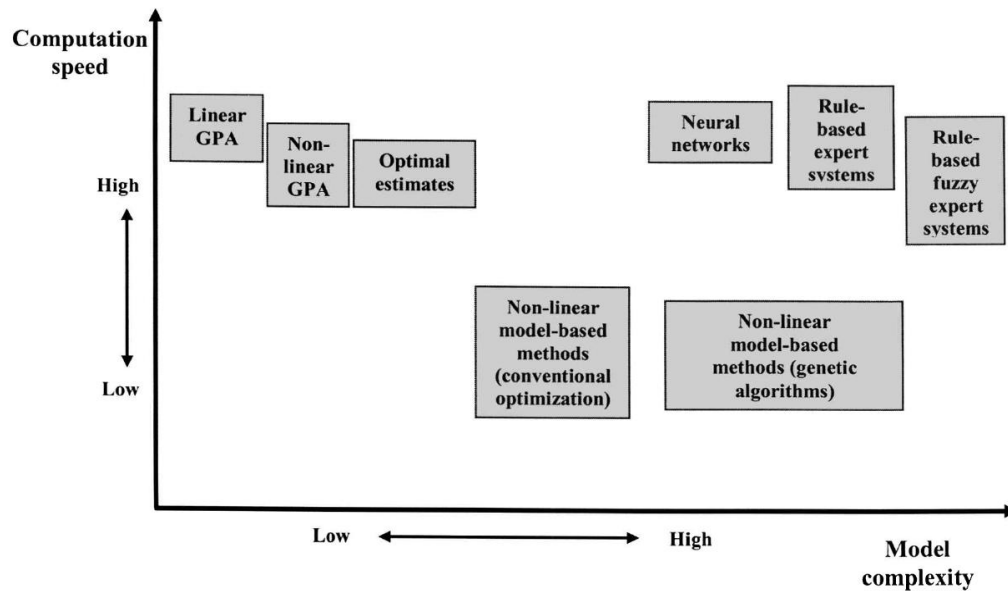


Figure 4.2: Comparison of diagnostic methods in terms of computation speed and model complexity, [107]

Jaw also describes how ANNs, FL, GA, Markov Chains (MC), BT, ES and Decision Trees (DT) are applied to GPA, [31]. Figure 4.3 shows how these methods can be classified according to computation complexity and speed.

(Speed)	High	Linear GPA	Artificial Neural Network	Fuzzy Logic Expert Systems
	Medium	Kalman Filter	Hidden Markov model	Non-linear Estimation
	Low	Bayesian Theory	Parameter Estimation	Genetic Algorithms
		Low	Medium	High
		(Complexity)		

Figure 4.3: GPA algorithmic performance assessment, [31]

In this study, data from the Tornado RB199 engine were used to develop methods for analysing engine performance degradation.

4.2 Data Sourced

Data from Rolls Royce RB199 engines fitted to Tornado aircraft, were sourced from the UK MOD. The data were from the Engine Health and Usage Monitoring System (EHUMS) that is being retro-fitted to the entire UK fleet. By the end of 2009 about fifty percent of the fleet had been fitted with EHUMS. Figure 4.4 shows the layout of the RB199 engine and the location of the significant components.




Figure 4.4: Cut-away diagram of the RB199 engine, [108]

The data sourced contained 2613 EHUMS downloads. Each download can contain multiple flights (sorties). Of these, 450 downloads had corrupt engine and/or tail numbers (e.g. engine serial number 0000, tail number "□ÿ000"). This is a significant fraction of the data (17%), and was mostly due to the details not being entered after main engine control unit (MECU) changes by the maintainers when the system was new. In addition, forty-one sorties had invalid data block markers. A total of 27,470 sorties, from 136 aircraft and 236 engines, from January 2006 to October 2009, were sourced. The average flight was one hour forty minutes long: hence the parameters recorded at 8.33Hz have an average of about fifty thousand points per flight.

The data were read using the FUMS™ system, which took about fifty hours, including calculating statistics on all parameters, and correcting corrupt data (see section 4.3); this corresponds to just seven seconds per flight (PC spec: Dell Precision PWS390 with Intel® Core™2Quad CPU (Q6600@2.4GHz). Some of these flights may potentially be duplicated, since the EHUMS data is recorded in a loop, so the same flight can appear in more than one download. Duplicates were later removed as part of the analysis.

Airspeed (IAS) and altitude parameters were derived from Static Pressure (PAVS) and Total Pressure (PASVS) using the following International Standard Atmosphere (ISA) equations:

$$Altitude = 1 - \left(\frac{T_{AMB}}{T_{LAPS}} \right) \left(\frac{PAVS}{P_{AMB}} \right)^{0.190263} \quad (4.1)$$

$$IAS = \frac{1}{0.514} \sqrt{\frac{2000 \times PASVS}{\rho_{AMB}}} \quad (4.2)$$

Where:

$$T_{AMB} = 288.15 \text{ K}$$

$$P_{AMB} = 101.325 \text{ kPa}$$

$$\rho_{AMB} = 1.2256$$

$$T_{LAPS} = 0.00198122 \text{ K/ft (6.5K/km)}$$

$$1\text{kt} = 0.5144 \text{ m/s}$$

The parameters read are given in Table 4.1. A suffix of “_L” and “_R” was used to indicate whether an engine had been fitted on the left-hand or right-hand position in the aircraft respectively.

Table 4.1: EHUMS parameters

Dry Parameters				
Name	Description	Units	Frequency	Range
NHV	High Pressure Shaft Speed	%	8.333Hz	0 to 125
NLV	Low Pressure Shaft Speed	%	8.333Hz	0 to 125
TBT	Turbine Blade Temperature	Degrees C	8.333Hz	0 to 1000
PLRH	Pilot's Lever Angle	Degrees (Position)	2.083Hz	-30 to 15
T1	Intake Total Temperature Equivalent to OAT	Degrees K	2.083Hz	200 to 495
PAVS	Aircraft Static Pressure	kPa	2.083Hz	0 to 220
PASVS	Aircraft Total Pressure - Ps0	kPa	2.083Hz	-5 to 158
LPCDEM	Lectric Pressure Control Demand (controls fuel flow)	Volts	2.083Hz	2.5 to 7.5
SFLAGA	Lane 1 Fail	Flag	2.083Hz	Good = 0 Fail = 1
SFLAGB	Lane 2 Fail	Flag	2.083Hz	Good = 0 Fail = 1
LC	Lane In Control	L1 or L2	2.083Hz	Lane1 = 0 Lane2 = 1
XDRV	Gearbox Cross Drive Engaged/Disengaged	Flag	2.083Hz	Disengaged = 0 Engaged = 1
OPNBAR	IP BOV Solenoid	Flag	2.083Hz	Energised/ Closed = 1
GFBAL1	Weapons Fire	Flag	2.083Hz	Weapons Fire on = 0
AUIG	Auto Ignition	Flag	2.083Hz	Ignition = 0 Non-Ignition = 1
TBTSWT	TBT Datum Switch	Low/Datum	2.083Hz	Datum = 0 Low = 1
Reheat Parameters				
Name	Description	Units	Frequency	Range
AJACT	Nozzle Shroud Position	Degrees	2Hz	-5 to 12
FUACT	Reheat Fuel Actuator Position	Degrees	2Hz	-5 to 12
TTACT	Tt1 Actuator Drive Position Distributes fuel between 3 reheat fuel manifolds. The colder the OAT, the more fuel goes to the flame holder and primary, with less to the bypass.	Degrees	2Hz	-5 to 12
HARDEN	Reheat Enabled	True/False	2Hz	Reheat OFF = 0 Reheat ON = 1
ONFTSW	Open Nozzle for Taxi Enable	True/False	2Hz	TN Closed = 0 TN Open = 1
BKTSIN	Reverse Buckets Deployed	True/False	2Hz	Stowed = 0 Deployed = 1
PLBA	Reheat Reset (by Pilot) Enabled	True/False	2Hz	Reset = 0
FREEZE	Reheat Fail or Freeze Indication	True/False	2Hz	Freeze/Fail = 1
LFAULT	Reheat Dormant Fail	True/False	2Hz	Good = 1 Fail = 0
WLOOL	Working Line Limit (WLL) Exceedance	True/False	2Hz	Limit Exceed = 1

Figure 4.5 shows the number of sorties and flying hours, per month, recorded in the EHUMS data. The figure shows that there was a marked increase in EHUMS data from August 2008, as the EHUMS units are retrofitted to the fleet.




Figure 4.5: Total flying hours (top) and number of sorties (bottom), per month

Figure 4.6 shows correlations between parameters for a typical flight. While the individual charts are small, the overall figure allows identification of parameters that are well correlated or don't have normal distributions.

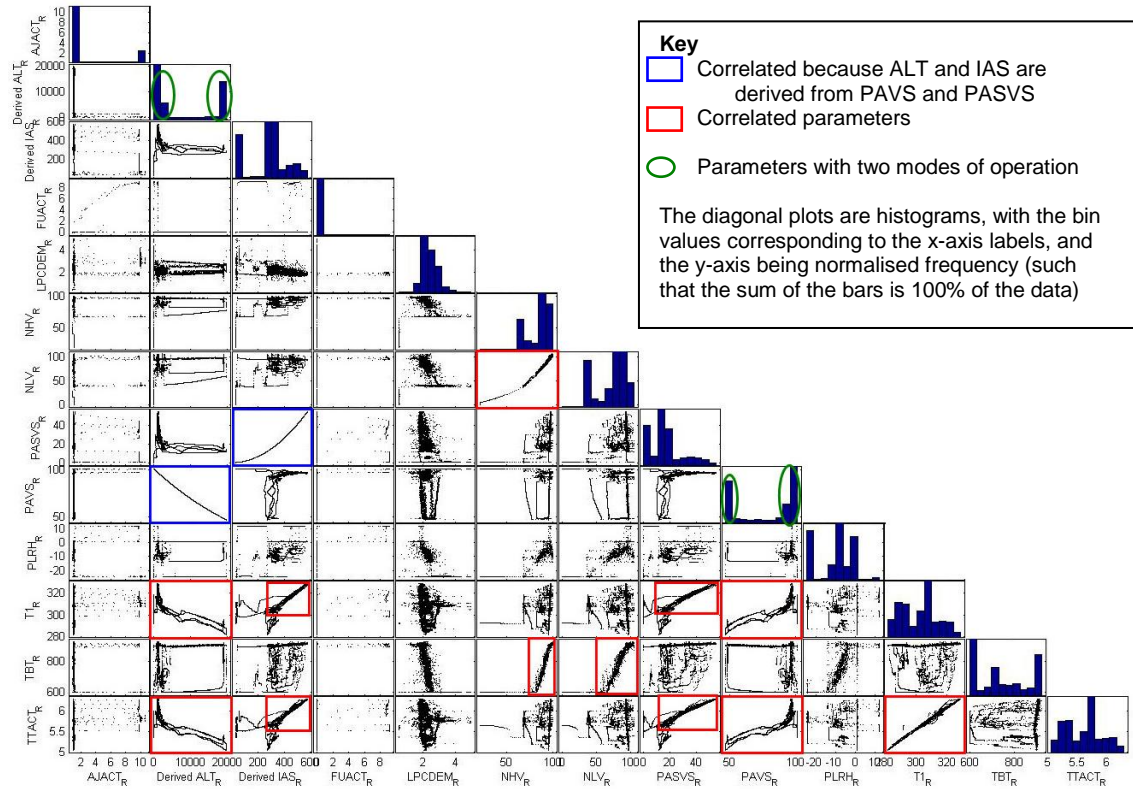


Figure 4.6: Parameter scatter plot matrix with histograms for a typical flight

4.3 Data Correction

The FUMS™ short period automatic data correction (ADC) , [44], algorithm was configured for the EHUMS data. ADC is designed to correct corruptions such as spikes, steps and drop outs, which occur over periods of less than two seconds. It is only applicable to the twelve continuous parameters such as NHV, NLV, TBT, derived Altitude and Airspeed, and not to discrete data, which only has values of true/false or one/zero.

The configuration data of the ADC algorithm are derived by analysing a sample from the data signal under consideration. The configuration data can include both generic data for all signals and configuration data specific to a single signal. Corrections are performed using linear interpolation between valid values, either side of the period of corruption.

The key configuration parameter is DMAX. DMAX is a real number evaluated for each signal and represents the maximum expected rise/fall of the signal. Exceeding the maximum rise/fall DMAX can occur over a number of points denoted by NSEARC; legacy data has indicated that a suitable value of NSEARC is three data points (regardless of data frequency) and can be applied to the majority of signals.

The algorithm marks the data with one of the following flags (note that only type 4 corruption is corrected):

- 0 = Data valid, no corruption.
- 1 = Long period of corruption (data not corrected).
- 2 = DC signal – i.e. successive values are constant (data not corrected).
- 3 = Jump as sensor switched on condition. Data preceding this point can be considered as "sensor off" (data not corrected).
- 4 = Short period corruption (e.g. spikes). Data flagged and corrected with interpolated values.
- 5 = Jump type data corruption (data not corrected).

The 27,470 EHUMS sorties sourced represented a total of 46,298 recorded hours. The twelve continuous parameters therefore contained a total of 555,576 hours of data. Table 4.2 shows the total hours of data that was flagged by the ADC algorithm. The table also shows the percentage of data

identified, and the number of periods of corruption that were identified. A total of 1547 hours were flagged as corrupt, which represents just 0.28 percent of the data. Only 0.068 percent of the data was of type four and so was corrected by linear interpolation.

Table 4.2: Results of ADC by correction code

Code	Type	Hours	%	Count
1	Long period	1133.78	0.2041%	12,065
3	Jump at start	2.56	0.0005%	5,559
4	Short period	378.98	0.0682%	149,257
5	Jump	31.59	0.0057%	4,968
Total		1546.92	0.2784%	171,849

Figure 4.7 shows a FUMS™ report that was configured to view the results of the data correction. It allows the user to select the aircraft, flight and parameter of interest and to see a chart of the raw and corrected data. Figure 4.9 to Figure 4.12 show statistics of the number and duration of corruptions found, by aircraft and parameter. The figures show that: aircraft 6 had significantly more corruptions than other aircraft; TTACT was the parameter with the longest duration of corruption; PLRH and TBT had the largest number of corrupt periods. Figure 4.13 shows some example corruption and the corrected values.

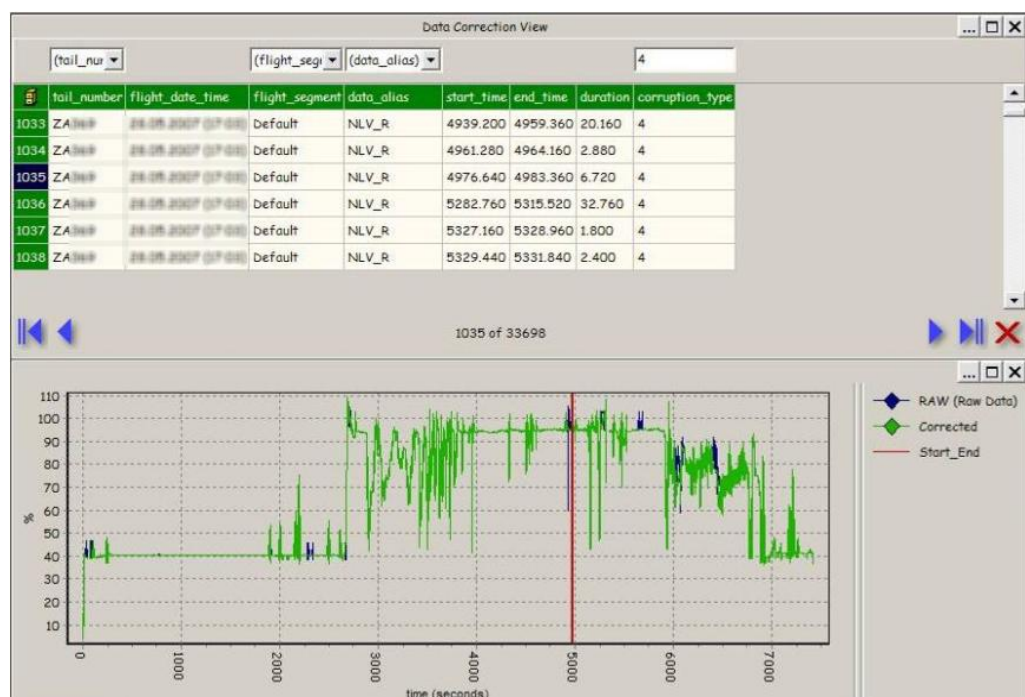


Figure 4.7: Report to view data correction




Figure 4.8: Number of corruptions identified per aircraft




Figure 4.9: Duration, in minutes, of corruptions identified per aircraft




Figure 4.10: Number of corruptions identified per parameter




Figure 4.11: Duration, in minutes, of corruptions identified per parameter




Figure 4.12: Number of corruptions identified per aircraft and parameter




Figure 4.13: Example of data corruption and the corrected values

4.4 ProDAPS Anomaly Modelling

GE Aviation has developed a Probabilistic Diagnostic and Prognostic System (ProDAPS) for anomaly detection, [55]. The objective of anomaly detection is to identify abnormal behaviour that might be indicative of some fault. Anomaly detection is used in a number of applications, with the underlying theme being that there is no large library of tagged fault data with which to train a model. The process is conceptually simple; a model of normal behaviour is built using a training data set, then new data is assessed for its fit against this model. If the fit is not within the model's threshold then it is flagged as anomalous. Nearly all approaches assume that a set of normal data is available to construct a model of normal behaviour. Modelling with in-service data therefore presents significant challenges. Due to issues such as a lack of feedback from the repair and overhaul process, undetected instrumentation problems, maintenance interventions, etc., it must be assumed that any database of historical in-service data will contain unknown anomalies.

Anomaly models are built from a set of input data, with input parameters selected according to the particular monitoring requirements for the model. The anomaly models are based on Gaussian Mixture Models (GMMs) and provide detailed density mapping of the data. GMMs allow complex distributions to be modelled by summing a number of Gaussian distributions. A Gaussian distribution $d(x)$ is governed by (4.3) where μ is the mean (location of the peak) and σ is the variance (the measure of the width of the distribution). Multiple Gaussian distributions can then be summed as shown in (4.4), each with a weight w corresponding to the number of samples represented by that distribution. An example of summing four Gaussian distributions to represent the density of the data is illustrated in Figure 4.14. In multi-dimensional problems, the individual distributions are often called clusters since they represent a subset of the data in terms of density distribution.

$$d(x) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x-\mu)^2}{2\sigma^2}} \quad (4.3)$$

$$f(x) = \sum_{i=1}^n w_i d_i(x) \quad (4.4)$$




Figure 4.14: Single parameter distribution modelled by a Gaussian mixture model

The clusters in a model can rotate to represent correlations between parameters. The rotation is defined by the cluster covariance matrix. In ProDAPS, the models are then adapted so that they reject any abnormalities existing in the training data. A carefully designed automatic ‘model adaptation’ process detects regions in the cluster space that are not representative of normal behaviour and then removes these clusters. The adaptation process is complex but is controlled by a simple tuning parameter. The final model provides a poor fit to samples in the training data that are outliers. A significant amount of effort was expended developing the novel automated model adaptation process, as this is the key to the successful building of models using in-service data, containing various unknown anomalies.

The resulting models are sophisticated statistical representations of the data generated from in-service experience; fusing sets of input parameters to reduce a complex data picture into a single parameter time-history, called a ‘Log Likelihood’ (LL) or ‘Fitness Score’ (FS) trace. The FS measures the degree of abnormality in the input data and mirrors the shape of any significant data trends. It represents a ‘goodness of fit’ criterion, indicating how well data fits a model of normality. Therefore, the FS has a decreasing trend as data becomes increasingly abnormal.

Another novel feature of the modelling process is that it does not require data to be categorised as ‘training’ or ‘test’ (which is a common practice in data modelling, to ensure that built models will generalise to data not used for training). All data can contribute to a model and the standard procedure for building a model is to use all available historical data, apart from cases that are known a-priori to be anomalous. This also has the advantage that online model updates can be performed as new data are acquired. The ability to update models is important, particularly for a new aircraft type where data is initially limited.

The anomaly model FS output is converted into a ‘Probability of Anomaly’ (PA) measure, which is a normalised probability measure that ranges between zero and one. For each model there is a PA distribution which is an extreme value distribution. A FS value is passed to the PA distribution and a PA value (probability) is returned. Most FS values will return a PA of zero because most acquisitions will be normal. The PA values provide a measure that is normalised across models. This allows model outputs to be compared. Such a measure could be fed into a secondary process, such as automated reasoning, to assess the nature of an anomaly. In addition, it would facilitate data mining of anomalous patterns in the search for new knowledge.

ProDAPS models can also be used to predict an input variable from any combination of other inputs. These can be subtracted from the measured value to give deltas, which indicate whether the measured data is higher or lower than expected by the model. Influence Factors (IF) can also be computed, which give an indication (0-1) for each input parameters as to how it is affecting the overall PA. Figure 4.15 shows the data flow in a typical ProDAPS application: the input data (CIs) are pre-processed (e.g. moving median filter, removing of outliers, splitting the data when components were replaced), then used to build anomaly models; these models can then be used to calculate the FS and PA. Where an alert is identified (usually using an m-out-of-n (MoonN) criteria – i.e. if m is 3 and n is 4, 3 out of every 4 PA values must be over a threshold, e.g. 0.9), the IF can be interrogated to help diagnose the fault. Periodically the models can be rebuilt to represent changes in the fleet.

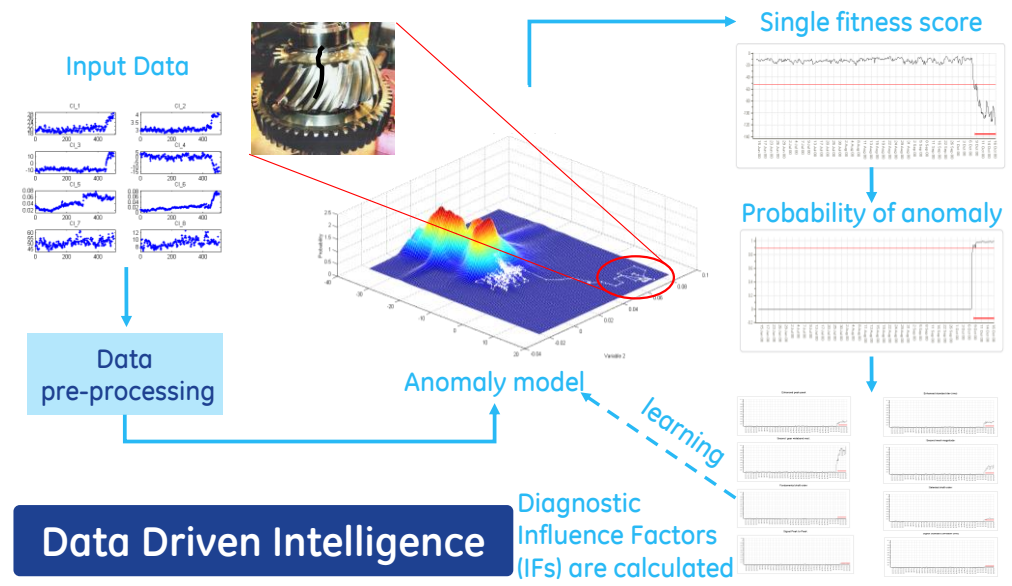


Figure 4.15: ProDAPS data driven intelligence, [55]

4.5 Performance Snapshots

There is an established link between the deterioration of the High Pressure Nozzle Guide Vanes (HPNGVs) in the RB199 engine and degradation in the engine performance, [109]. The study of Accident Data Recorder (ADR) information has revealed that engines suffering uncontained failures had undergone a progressive deterioration in high pressure spool speed, over a number of sorties, and had been operating significantly below their placard values immediately before the failure. This indicates that the deterioration of the HPNGVs, leading to such events, would have been gradual and that this, and similar incipient mechanical integrity problems, could have been detected by monitoring performance parameters. It is believed that early investigation of low NH trends could have prevented the loss of at least two Tornado aircraft. Furthermore, installed performance trending of the engine allows early detection of the following types of fault, [109]:

- Performance deterioration.
- Incorrect placard figures entered by pilots or maintainers.
- Turbine damage.
- Compressor damage.
- Optical pyrometer faults.
- Incorrect TBT datum values.
- Performance below NL Minimum Installed Thrust Limit (MITL).
- High NH/NL.

In order to address this requirement, the aircraft crews are required to record engine performance data prior to take-off. The procedure for recording is as follows:

- Record current OAT.
- Select maximum dry power (no afterburner).
- After ten seconds (minimum) record NHV, NLV and TBT readings.
- On completion of sortie, transfer data onto MOD Form 705C.

The above procedure is taken from reference [109]. These values are currently used in the RB199 Installed Performance System (RIPS) to evaluate the engine performance.

4.5.1 Automatically Identifying Performance Snapshots

An algorithm was developed by the author to automatically identify this RIPS performance run just before take-off and to record values of all parameters at this time. Rather than using a single value, an average is taken over five seconds. Additionally, the spool speeds and temperatures are corrected to ISA conditions using the T1 value and the equations in Appendix A1.11. The snapshot data could then be used for performance modelling, removing the need for the pilot to manually record it. Figure 4.16 shows the twelve continuous parameters for a typical EHUMS data set around the time of the performance run.

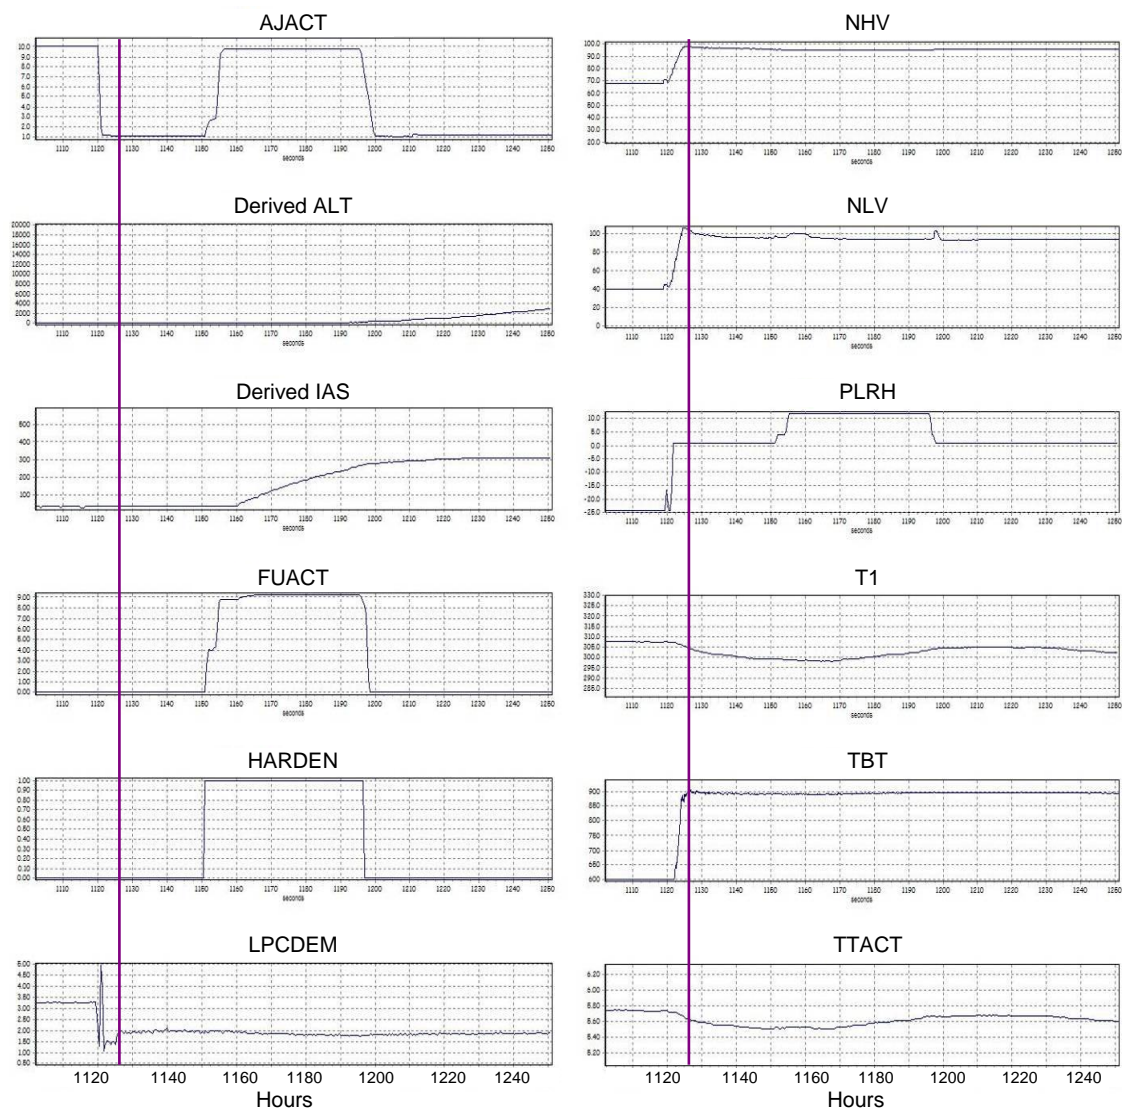


Figure 4.16: Example data around the performance run (from 1125 to 1150 seconds)

Figure 4.17 shows pseudo-code of the algorithm developed to automatically identify snapshots in the EHUMS data.

```

Set points_for_mean = 10
Set points_after_start = 20
Create all the columns required for the snapshot data
Query for a list of all translated flights
Loop round each Flight
    Get the flight data for required parameters
    Skip flights with very short duration (<30 seconds)
    Interpolate any null data points
    Resample to lowest frequency (2Hz)
    Filter Data for IAS > 200 to find the time spent in the air
    Filter Data NHV > 70 to find the engine run time
    Count the number of flights
    Filter the flight data for:
        90< NHV      <120,
        0< IAS      <110,   (These represent max dry power while stationary)
        -1< PLRH    <2
        -0.1< HARDEN <0.1
    Check that there are enough points
        Find mean of required number of points around the snapshot position
End of loop round flights

Check that there are no duplicates
Filter out flights with no snapshots found
Sort by engine_serial, download_date_time, flight_number
Calculate unique_eng_run_number (integer that increments for each engine run)
Calculate FIT (1 if there is an engine change, else 0)
Calculate Accumulative Flight Length, Time In Air, and Engine Run Time
Convert T1 to OAT (Celsius)
Correct NH, NL and TBT to ISA conditions
Save the results

```

Figure 4.17: Pseudo-code for the steady state snapshot algorithm

4.5.2 Performance Snapshots - Dry

Out of a potential 27,470 translated flights, 13,584 snapshots were taken by the algorithm. Note that over six thousand of the flights had less than a minute of engine run-time. Parameter histograms (Figure 4.18) and correlations between parameters (Figure 4.19) were investigated to help with understanding the data and choosing input parameters for the anomaly model. Note that the following discrete parameters always held a constant value during the snapshot: HARDEN=0, BKTSIN=1, PLBA=1, ONFTSW=0. The histograms show that TTACT and LPCDEM appear to have two distinct modes which the MOD revealed to be due to two different marks of engine since Mk104 has a longer jet pipe. A ProDAPS anomaly model was built using NH_COR, NL_COR, and TBT_COR as inputs.

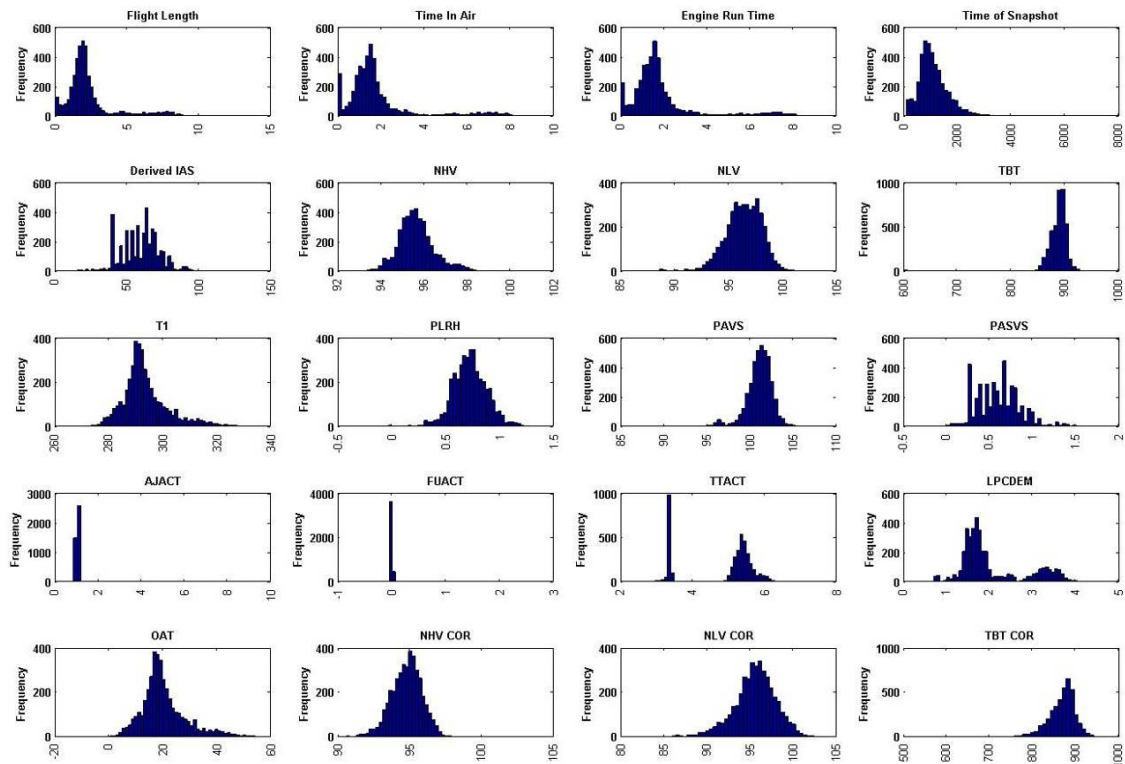


Figure 4.18: RB199 EHUMS snapshot histograms

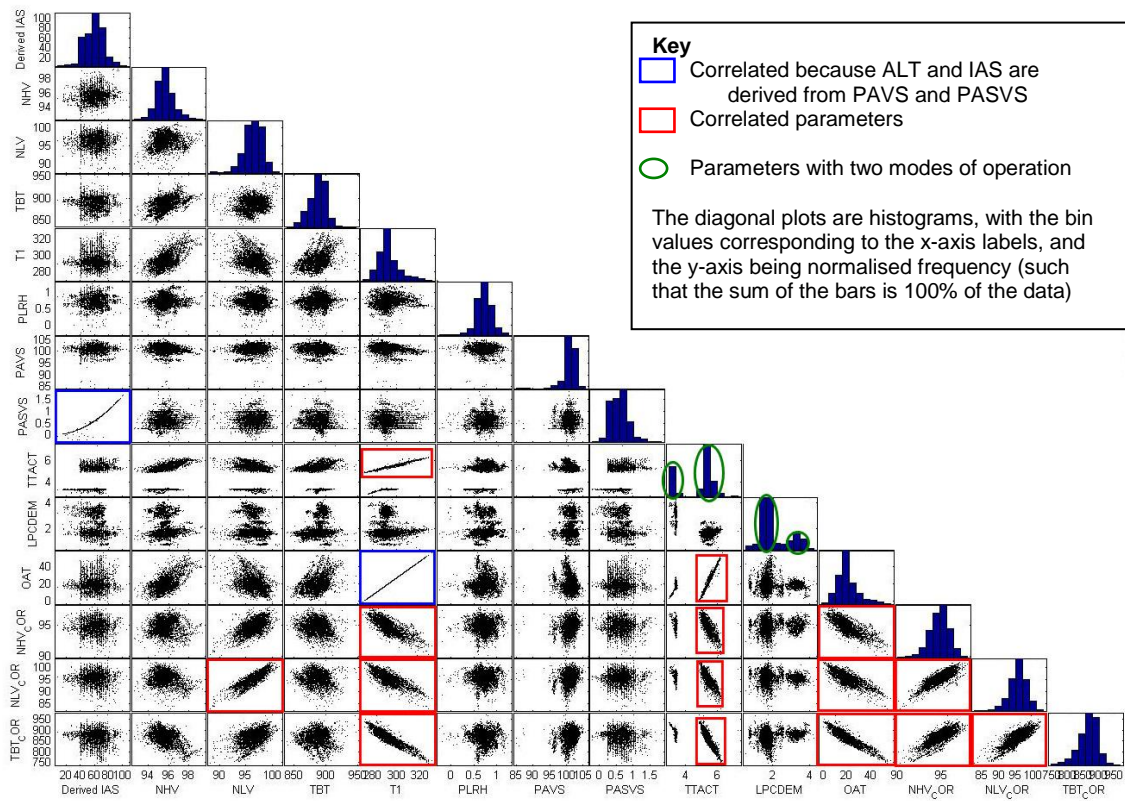


Figure 4.19: RB199 EHUMS snapshot scatter plot matrix with histograms

4.5.3 Performance Snapshots - Reheat

Using the same algorithm, developed for identifying the dry performance snapshot run just before take-off, snapshots were taken during take-off with the afterburner on (reheat). If this condition could give comparable results to the pre-takeoff run, it would remove the need for this extra run to be performed. This in turn would reduce fuel burn and increase sortie turn-around times as well as reducing unnecessary wear and tear on the engine. Figure 4.20 shows the twelve continuous parameters for a typical flight, around the time of the reheat snapshot. The following configuration changes were made to the algorithm (see section 4.5.1): The `points_after_start` value was increased from twenty to thirty and the filter configuration was changed to:

```

90 < NHV      < 120
0  < IAS      < 200
9  < PLRH     < 20
0.9 < HARDEN  < 1.1

```

Out of a potential 27,470 translated flights, 11,892 snapshots were taken by the algorithm. Note that nearly nine thousand of the flights had no time in the air. Quick analysis (looking at parameter distributions) identified the following outliers, which were caused by data corruption and so were removed: one flight had a value for TTACT of 0.2 degrees when the next smallest value was 1.6 degrees; two flights had very low NL values (0.09% and 0.27% respectively) when the next lowest NL value was 87.8%; one flight had an AJACT value of 5.6 degrees, which is 8.5 standard deviations from the mean. Finally, eleven snapshots were removed due to low OAT values (less than -2 °C) as these appeared to represent a different engine configuration, Figure 4.21.

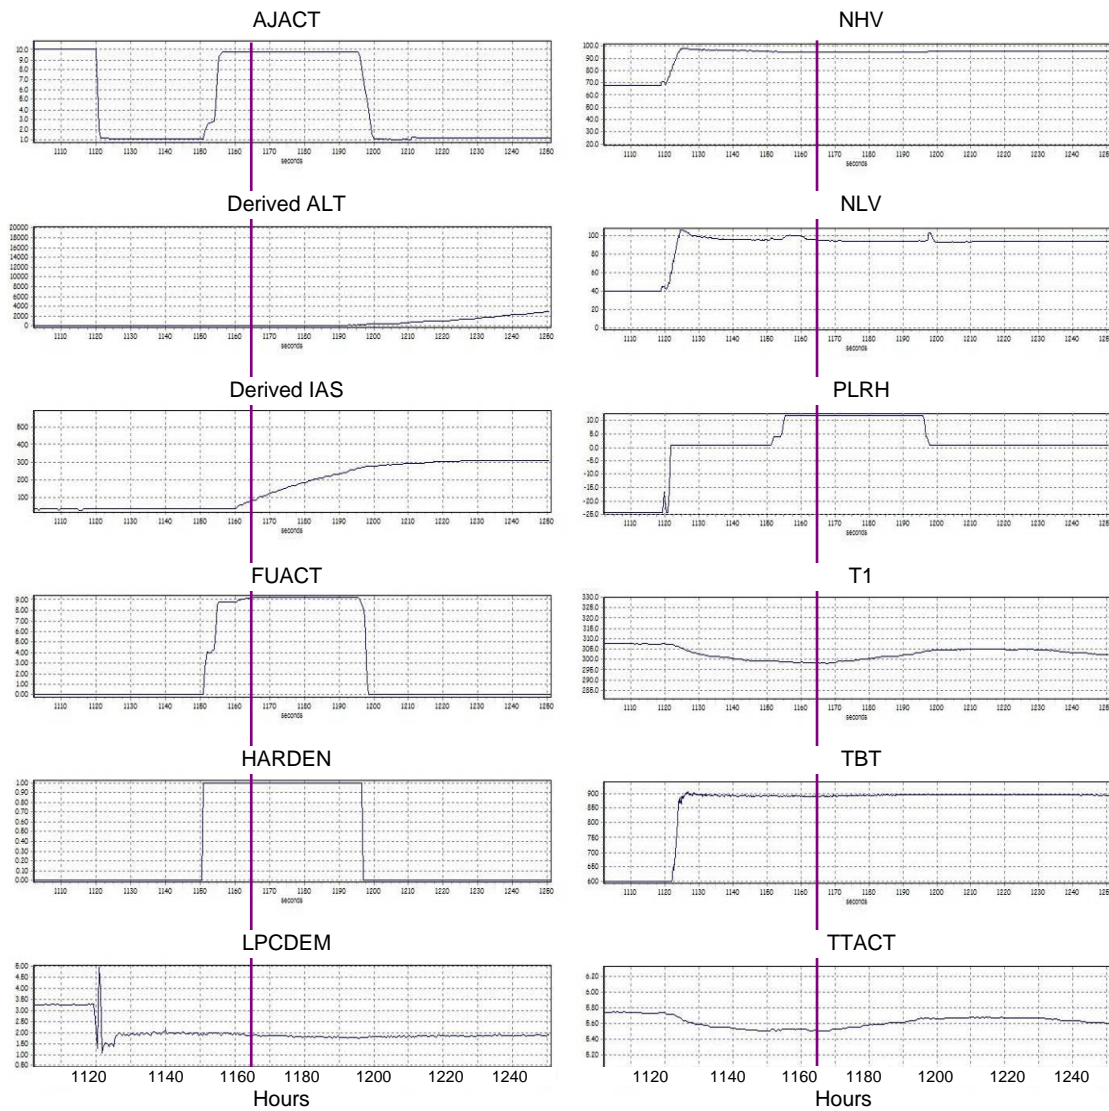


Figure 4.20: Example data around the reheat snapshot (from 1160 to 1195 seconds)




Figure 4.21: Reheat snapshot scatter plot of NH vs. TBT coloured by OAT

Parameter histograms (Figure 4.22) were investigated to aid with understanding the data. Note that the following discrete parameters always held a constant value during the snapshot: HARDEN=1, BKTSIN=1, ONFTSW=0, OPNBAR=1. The histograms show that PLRH, AJACT, FUACTION, TTACT and LPCDEM appear to have more than one distinct mode. Further understanding of the engine may therefore be required to help decide whether it would be beneficial to create two or more anomaly models for these different states.

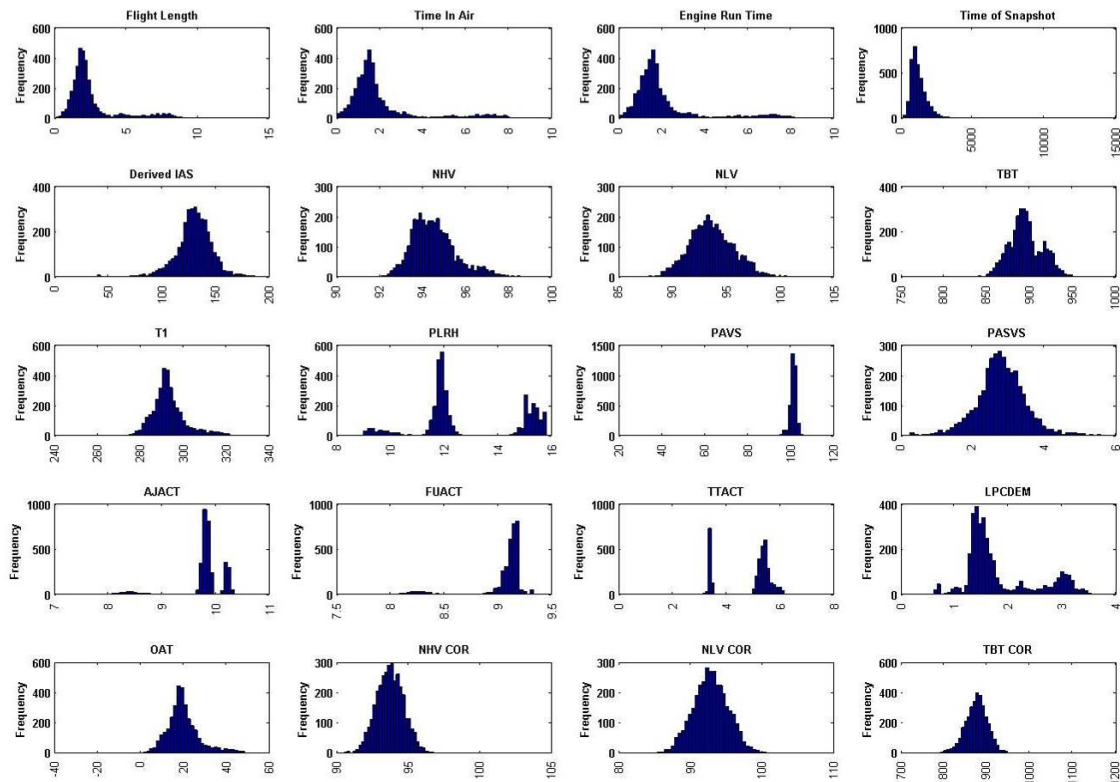


Figure 4.22: Reheat Snapshot histograms

4.5.4 Results

Using the performance anomaly model, the PA and LL values, along with the deltas between the actual and predicted values for NH, NL and TBT, were calculated. A summary table was also calculated that shows, for each engine, the number of points with LL outside two and three standard deviation limits. A MooN calculation was used to highlight where three out of four successive points were outside the threshold. A count was also made where the PA values were above 0.9 and a MooN calculation was performed on the PA data.

A FUMS™ report was configured to enable the results to be viewed. Figure 4.23 shows the summary table on the first page of the report. Colour coding was used to highlight data of interest. An engine can be selected and on the subsequent page the original, corrected and delta values for the parameters can be viewed. Figure 4.24 shows the layout of this display, which is used to present the results in the following sections. The charts show the fleet mean (green) and three standard deviation bounds (light blue). The charts also indicate, via a vertical maroon line, where an engine has been fitted to a different airframe. Some of the charts presented combine the dry and reheat snapshots on one chart – in this case the dry data and its statistics (mean and standard deviation bounds) are in blue, and reheat in green. Note that the engine numbers used are not actual MOD/RR engine serial numbers.

A full table of all the 60 engines that had at least one point in alert is given in Appendix A1.12. The table gives a description of what the data looks like and relates it to any Maintenance Work Orders (MWO). Twenty two of the engines looked like the alerts were likely to relate to an engine degradation of some sort. The table also shows the number of flying hours between the alert being generated and the fault being resolved. On average, the data is in alert for over sixty flying hours before it is resolved. This offers the operator a good window to investigate the fault and perform maintenance. Appendix A1.13 presents charts of all the engines shown in the table. A selection of engines that show the range of data anomalies detected are described in the following paragraphs.

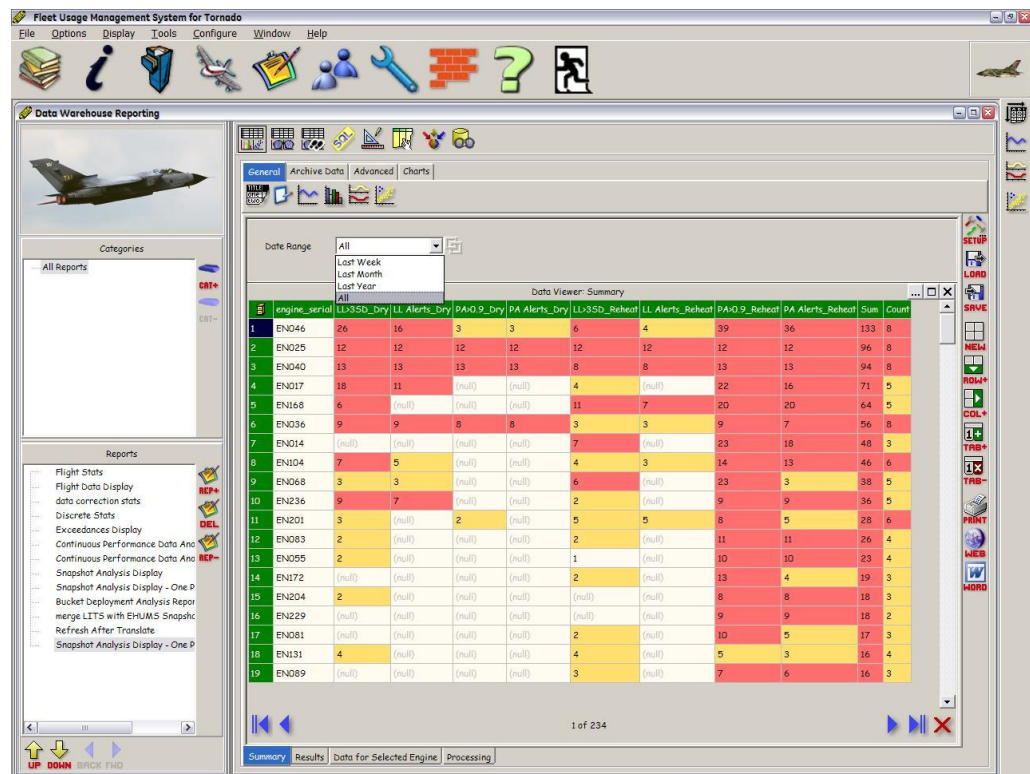


Figure 4.23: RB199 performance data – FUMS™ report to view results

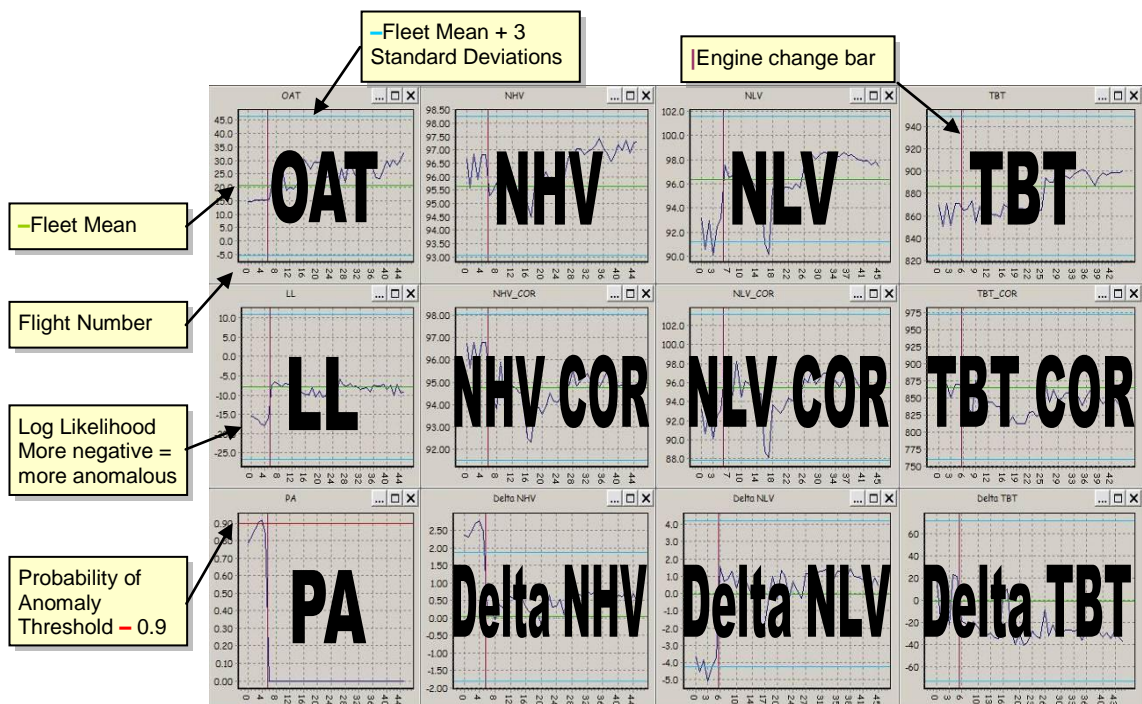


Figure 4.24: EHUMS anomaly results display layout

4.5.4.1 Engine 25 – TBT Sensor Failure

Eleven data snapshots from EN025 were found to be extremely anomalous (more than ten standard deviations from the mean), Figure 4.25. For all of these flights the TBT values were 600 (the minimum TBT value) whilst the other parameters appear to be normal. This suggests a sensor failure. Figure 4.26 shows how the TBT value drops off halfway through the previous flight (note that the raw sensor value after failure was 600; the plot shows this corrected to ISA conditions using OAT). The MWO revealed that the TBT fault was noted and an amplifier was replaced. Hence these snapshots were then removed from the training data and a new model was built. Note that it took the MOD eleven flights to spot this error, which could have been picked up after only one flight using this automated method.




Figure 4.25: EN025 performance results




Figure 4.26: Flight from EN025 where the TBT sensor failed

4.5.4.2 Engine 157 – Incorrect Snapshot Position

EN157 had a single snapshot value that was identified as anomalous, Figure 4.27. The delta TBT values show that the TBT reading was lower than expected. Investigation of the parameters around the snapshot position showed that the engine was not in a steady state condition. The engine was then returned to idle and ramped back up to maximum power a few minutes later, which did reach steady state conditions, Figure 4.28. This is an example where the algorithm, identifying the pre-take-off steady state condition, needs further refinements to ensure the engine is in a steady state condition.

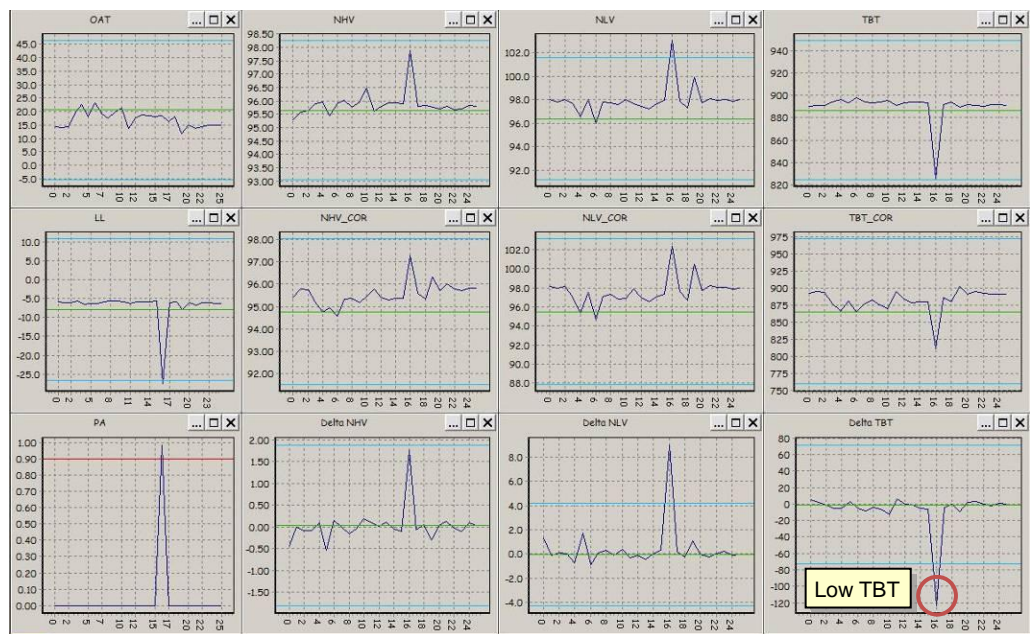


Figure 4.27: EN157 performance results




Figure 4.28: EN157 data showing wrong and correct snapshot positions

4.5.4.3 Engine 201 – Data Corruption

EN201 had two snapshot values that were identified as anomalous, Figure 4.29. The first corresponded to flight ten where the delta values indicated that NH was high and TBT was low. The second corresponded to flight fifteen where the delta values indicated that NH was low and NL was high. Investigating the data around the time of the snapshots revealed that there was severe data corruption, which could not be corrected to sufficient accuracy to avoid skewing the snapshot values, Figure 4.30. Hence this snapshot was removed from the training data. This example highlights the importance of identifying and correctly dealing with corrupt data values.

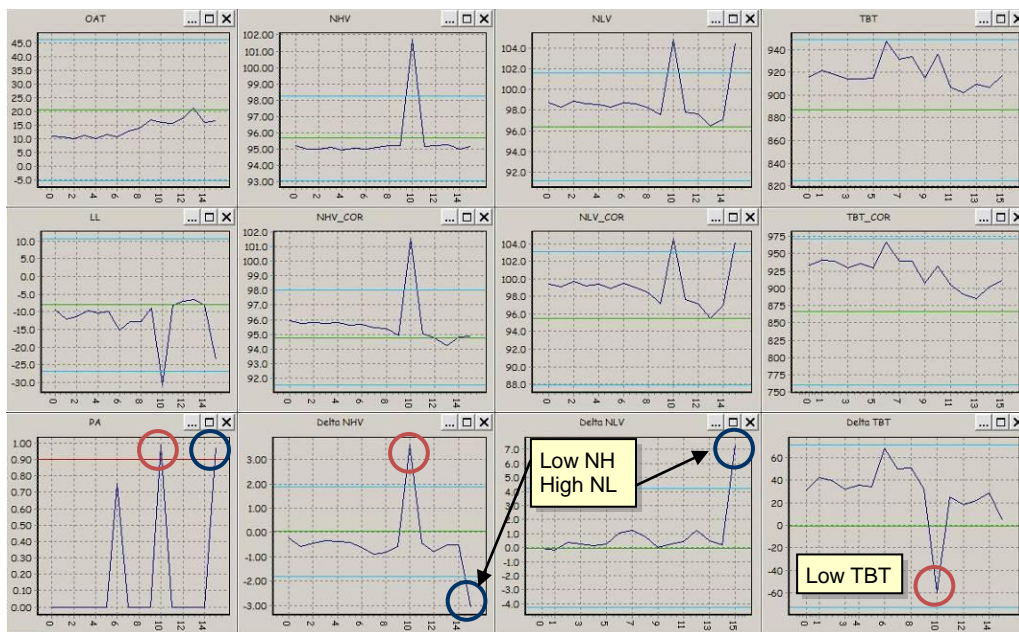


Figure 4.29: EN201 performance results




Figure 4.30: EN201 flight 15 around the snapshot position, showing the raw corrupt data (left) and the corrected data (right)

4.5.4.4 Engine 30

EN030 has a few anomalous points (mainly for the reheat snapshot) before an engine change, Figure 4.31. The deltas appear to show an increase in TBT and drop in NH for the last six flights for this engine. The engine was removed due to debris identified by the Early Failure Detection Cells (EFDC). When fitted to another aircraft, nine months later (presumably after being reconditioned), the data returned to normal. The analysis gave about twenty five hours warning of this fault.

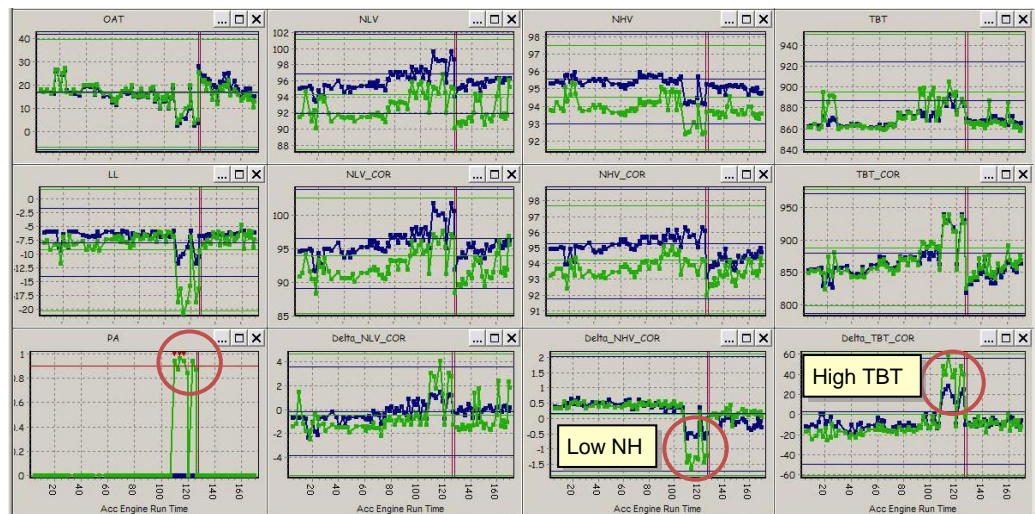


Figure 4.31: EN030 performance results

4.5.4.5 Engine 100

EN100 had a few anomalous points at the start of the data, Figure 4.32. The engine was then fitted to another airframe and the data for the engine returned to normal, implying that maintenance was performed to remedy the problem; however, no related MWOs were identified to confirm this. The deltas indicate that NH is high whilst NL is low.




Figure 4.32: EN100 performance results

4.5.4.6 Engine 204

EN204 performance seems to deteriorate, with delta NH rising sharply after 42 hours, Figure 4.33. The MWO indicated that there had been a FOD event, resulting in the low pressure compressor being damaged. Despite this, it took over 36 flight hours for this fault to be noticed by the MOD.




Figure 4.33: EN204 performance results

4.5.4.7 Engine 36

EN036 performance data becomes anomalous after 60 engine running hours due to high NH, Figure 4.34. The MWO identified FOD as the cause and engine modules 1-3 were replaced.




Figure 4.34: EN036 performance results

4.5.4.8 Engine 104

EN104 performance data has high NH and low NL at the start, Figure 4.35. The values appear to become normal about ten engine run hours before an engine change.

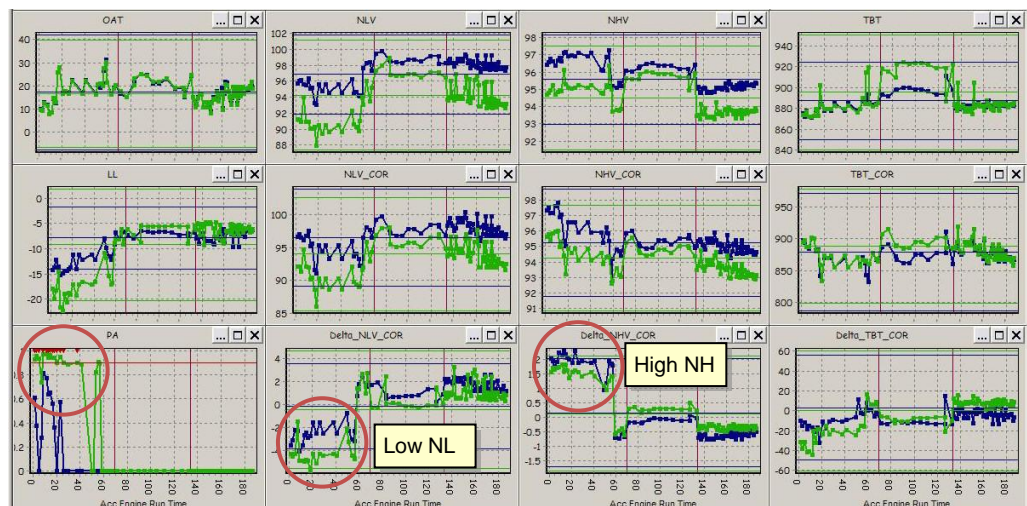


Figure 4.35: EN104 performance results

4.5.4.9 Engine 55

EN055 performance data has high NH and slightly low NL and TBT for ten flights then returns to normal, Figure 4.36. There are two MWO that relate to this period: The first indicated that the engine was replaced due to debris identified by the EFDC; the second states that the main fuel feed supply pipe was damaged. These alerts would have given about twenty flight hours warning of these faults.




Figure 4.36: EN055 performance results

4.5.4.10 Engine 2

EN002 performance data has alerts only for the reheat snapshot data. However, the shapes of the deltas are similar for the dry snapshots also, Figure 4.37. The trend of NH was found to have a negative gradient, with NL and TBT having a positive gradient. The engine was then removed for maintenance and the values returned to normal levels. The MWO indicated that the engine was removed because its performance was below the minimum installed performance tolerance level allowable.




Figure 4.37: EN002 performance results

4.6 Engine Performance Conclusions

In this chapter, anomalies were identified in a large set of RB199 data, by automatically identifying ‘snapshots’ of data during a mandatory maximum dry power engine run before take-off and with reheat applied during the take-off run. The integrity of the data was also analysed using an automatic data correction algorithm, which flagged 0.28 percent of the data as corrupt, and performed linear interpolation to correct 0.068 percent of the data.

The methods identified anomalies in 60 of the 230 engines, of which about half could be attributed to faults identified in maintenance records. On average, alerts were raised over 60 flying hours before the fault was fixed by the operators. This shows that the methods could offer significant improvements in MAAAP. The automation of the snapshot data collection could also remove the need for pilots to perform long ground runs and manually take readings from instruments. These improvements would, however, require rigorous airworthiness qualification that would be costly and time-consuming. In addition, the EHUMS units, which record the data, have yet to be fitted to the entire fleet.

CHAPTER 5

5 IDENTIFYING UNCOMMANDED FLIGHT CONTROL MOVEMENTS

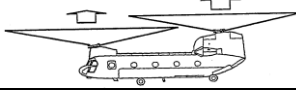
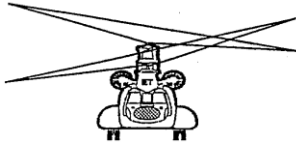
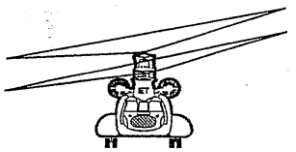
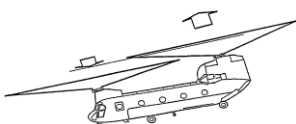
Uncommanded Flight Control Movements (UFCMs) are incidents where the aircraft makes an unexpected manoeuvre (such as a jolt in pitch, roll or yaw) without pilot inputs or external factors such as a gust. UFCMs are usually caused by faults in the Automatic Flight Control System (AFCS) or one of its inputs (i.e. sensors such as gyros) or outputs (i.e. actuators).

This chapter presents analysis carried out to automatically identify UFCM events. Some of the initial findings were presented to the MOD in 2005, [110] and published at the first Prognostic Health Management conference, [14].

5.1 Background – Chinook Control

The Chinook helicopter has a tandem contra-rotating rotor configuration. This cancels the inherent rotor torque, thereby eliminating the need for a conventional tail rotor, meaning that all installed power is available for lift and translation/rotation of the aircraft. This alters the flying control actions as shown in Table 5.1.

Table 5.1: Comparison between Chinook and conventional helicopter controls

Desired Behaviour	Pilot Input	Conventional Helicopter	Chinook (Tandem)	Diagram [111]
Climb/Descent	Collective Leaver	Increase pitch on all blades of the main rotor	Increase pitch on all blades of both rotors	
Yaw	Pedals	Change pitch of tail rotor	Apply cyclic pitch to tilt fore and aft rotors laterally in opposing directions	
Roll	Cyclic Stick left/right	Apply cyclic pitch to main rotor to tilt left/right	Apply cyclic pitch to tilt fore and aft rotors laterally in the same direction	
Pitch	Cyclic Stick fore/aft	Apply cyclic pitch to main rotor to tilt fore/aft	Differential Collective Pitch (DCP)	

The Chinook has a duplex Automatic Flight Control System (AFCS) providing short-term damping and long-term attitude hold. The AFCS receives and processes signals from sensors; it passes control signals to electro-hydraulic and electro-mechanical actuators in the flying control system. Pitch and roll attitude, airspeed, heading, barometric altitude and radar altitude hold functions are incorporated in the system, [112].

As shown in Figure 5.1, the forward rotor shaft has a forward tilt of 9° and the aft rotor a forward tilt of 4° . This was required to provide a ground taxi capability and to eliminate abnormal nose-down attitude during cruise. This configuration causes a number of problems, which are discussed below.




Figure 5.1: Chinook tandem configuration, showing rotor shaft tilts [111]

The first problem is that, as the forward speed increases, because of the different tilts, the forward rotor will produce less lift, creating a pitch-down force. The pilot would have to counteract this by pulling back on the cyclic stick at higher speed, which is counterintuitive. The solution to this problem was to introduce a Differential Air Speed Hold (DASH) system. The DASH system has three main purposes: to provide a positive linear stick plot for increasing speed, to maintain commanded airspeed and to correct the aircraft attitude for any pitch trim changes. Figure 5.2 illustrates the relationship between cyclic stick position and forward speed, with and without the DASH operating. The DASH actuators are controlled by the AFCS, which is based on four signals: airspeed, pitch attitude, longitudinal cyclic stick position and DASH actuator position feedback. The DASH operates in different modes above and below forty knots (kts).




Figure 5.2: Effect of DASH system on the relationship between airspeed and cyclic stick position, [111]

The rotor configuration also causes disc flap-back at higher speeds, requiring increased nose-down pitch, resulting in higher profile drag and an undesirable cockpit environment (i.e. excessive nose down pitch). The force generated by the disc flap-back creates structural stress on the rotor shafts, which could reduce component life. In addition, the flap-back may cause the aft traversing blade, which is flapping down, to contact with the droop stops, causing droop stop pounding. In order to solve these problems, the Chinook is fitted with a Longitudinal Cyclic Trim (LCT) system, which automatically tilts the rotor discs forward, as airspeed increases. A negative side-effect of this system is that the more the rotors are tilted forward, the smaller the separation between the overlapping discs, which causes increased noise. In summary, the LCT system aims to achieve the following: maintain a near level fuselage for any constant airspeed, improve the cockpit environment, reduce profile drag, reduce stress and vibration on the rotor shafts and heads, and prevent droop stop pounding.

The LCT operation varies with flight conditions as follows: When on the ground, the swashplate is aligned at approximately 90° to the rotor shaft, i.e., zero cyclic pitch. When airborne, and travelling at less than 60 kts, the LCT actuators retract, tilting the rotors back. Between 60 and 150 kts the actuators extend in proportion to indicated air speed. At 150 kts, at sea level, the actuators reach their forward stops. Figure 5.3 shows how the LCT position varies with airspeed. Since increased altitude decreases the air density, the LCT actuators also extend by 0.2° per 1,000 ft. This means that the forward stops will be reached earlier at higher altitudes; conversely the LCTAs become

active as early as 40 kts (below 40 kts the IAS is unreliable, so the AFCS assumes 40 kts airspeed).




Figure 5.3: Relationship between the forward and aft LCT positions with airspeed, at sea level [113]

Due to the complex nature of the Chinook control laws, it is easy to see how failures of any one of the systems (AFCS, DASH, and LCT) or their inputs (e.g. sensors such as gyros) or outputs (e.g. actuators) could cause noticeable events. It is these events that this analysis is attempting to detect.

5.2 Review of UFCM incidents

A database of forty UFCM events, between February 1996 and January 2005, was obtained from the MOD. To prevent traceability, the data had the aircraft tail number information removed. FDR data was not available for the UFCMs. The UFCM events were summarised, as described by the pilot, into the following categories: Cockpit Warning, Jolt, Pitch Disturbance, Vertical Bounce, and Yaw Disturbance.

The resulting maintenance actions were also summarised, by listing the items that were replaced:

- AFCS
- Differential Air Speed Hold (DASH) Actuator (DA)
- Directional Gyro (DG) and Vertical Gyro (VG)
- Other actions, which include: Pitch or Roll Control Position Transducer (CPT), Integrated Lower Control Actuator (ILCA), LCT Actuator, Attitude Indicator (AI), chaffed cables, incorrect fitting, etc.

Table 5.2 shows the maintenance actions associated with the various categories of UFCM.

Table 5.2: Summary of components replaced after different types of UFCMs

UFCM	AFCS1	AFCS2	DG	DA	VG1	VG2	Other
Cockpit Warning				1			
Jolt							1
Pitch Disturbance	4	4		12	2	1	3
Vertical Bounce	3	1			3	3	4
Yaw Disturbance	4	5	8	1	2	3	4
Total	11	10	8	14	7	7	12

5.3 Data Analysis to Identify UFCMs

FDR data for 775 Chinook flights from thirty-five aircraft, between October 1999 and January 2006, were also sourced from the MOD. Table 5.3 lists the parameters from these flights used in this analysis and their full descriptions, units, sampling rate (Hz) and statistics.

Table 5.3: Chinook helicopter parameters investigated

Parameter	Description	Units	Hz	Min	Max	Mean	Seconds	% NULL
AFT_LCTA	Aft LCT Actuator Position	Inch	2	-1.9	2.0	0.6	3279599	0.13%
ALT_RATE	Altitude Rate	Ft/Min	16	-8192.0	8128.0	-32.3	3279431	0.03%
COLL_STICK	Collective Stick Position	Inch	2	-128.0	127.5	4.9	3279599	0.13%
DASH_ACT_1	DASH Actuator 1	%	2	-128.0	127.5	30.8	3279599	0.03%
DASH_ACT_2	DASH Actuator 2	%	2	-128.0	127.5	32.7	3279599	0.03%
FWD_LCTA	Forward LCT Actuator Position	Inch	2	0.0	1.0	0.4	3279599	0.14%
IAS	Indicated Airspeed	Kts	1	0.0	163.0	45.3	3295119	0.03%
LH_WOW	Left Hand Weight On Wheels	1=Gnd 0=Air	1	0	1	0.5	3295119	0.03%
NX	Lateral Acceleration	G	4	-1.0	0.6	0.0	3279501	0.14%
NY	Longitudinal Acceleration	G	4	-1.0	0.9	0.0	3279501	0.14%
NZ	Vertical Acceleration	G	4	-3.2	2.8	1.0	6558906	0.14%
PITCH_ANGLE	Pitch Angle	Deg	2	-51.0	386.0	85.1	3279599	0.13%
PITCH_STICK	Pitch Stick Position	Inch	2	-79.5	61.0	2.5	3279599	0.13%
PRESS_ALT	Pressure Altitude	Ft	16	-20112.4	16688.7	738.8	3279431	0.14%
RH_WOW	Right Hand Weight On Wheels	1=Gnd 0=Air	1	0	1	0.5	3295119	0.03%
ROLL_ANGLE	Roll Angle	Deg	2	-65.5	434.0	174.7	3279599	0.13%
ROLL_STICK	Roll Stick Position	Inch	2	-103.5	90.0	-0.3	3279599	0.13%
YAW_PEDAL	Yaw Pedal Position	Inch	2	-82.0	90.5	-1.4	3279599	0.13%

Some initial statistical analysis was performed on the data to check its integrity. As shown in Table 5.3, the number of points (shown as seconds), maximum, minimum, mean, and number of null points (i.e. error in the

airborne system recording the data, shown as a percentage of the data) were calculated. The data contained 910 hours of data, totalling over 230 million data points for the eighteen parameters of interest. The statistics took about one hour to process (PC spec: Dell Precision PWS390 with Intel® Core™2Quad CPU (Q6600@2.4GHz). The following observations were made:

- Of the 775 flights, 130 were ground runs (i.e., both the LH_WOW and RH_WOW flags were true throughout the flight).
- Whilst the normal range of the collective stick is zero to 110, 588 flights (none of which were ground runs) had collective stick values of less than 10. This meant that this parameter could not be relied upon.
- The data frequency was found to be set to zero for all parameters in two flights and some parameters in another three flights. However, for further processing, the data was assumed to be at the frequency listed in Table 5.3.
- On average, 0.08% of data was found to be null. Five flights had more than 1% null, and the highest percentage was 25.9%. No single parameter appeared to have more null points than others.
- Table 5.4 shows that a significant number of flights had constant values (i.e., the maximum was the same as the minimum) for some parameters.

Table 5.4: Number of flights with all values constant (i.e. max = min) by parameter

Parameter	Flights	Ground Runs	Total
AFT_LCTA	13	47	60
ALT_RATE	1	1	2
COLL_STICK	37	44	81
DASH_ACT_1	4	27	31
DASH_ACT_2	4	27	31
FWD_LCTA	5	28	33
IAS	11	15	26
LH_WOW	6	130	136
NX	0	1	1
NY	0	3	3
NZ	0	0	0
PITCH_ANGLE	8	36	44
PITCH_STICK	16	29	45
PRESS_ALT	0	1	1
RH_WOW	13	130	143
ROLL_ANGLE	11	51	62
ROLL_STICK	6	29	35
YAW_PEDAL	16	31	47

5.3.1 Example UFCM

One of the flights for which FDR data was available was identified by the MOD as having a UFCM event. During the flight, incorrect operation of the Forward LCT Actuator (FWD_LCTA) resulted in a pitch oscillation.

Figure 5.4 shows the instant in the flight where the UFCM occurred. The aircraft pitch-down can be seen in the FUMS™ screenshot of the 3D animation and the pitch angle time trace. Note that the FWD_LCTA is approximately proportional to the IAS, except at the point when the UFCM occurred (approximately twenty-three minutes into the flight), where the FWD_LCTA value plummets to zero.

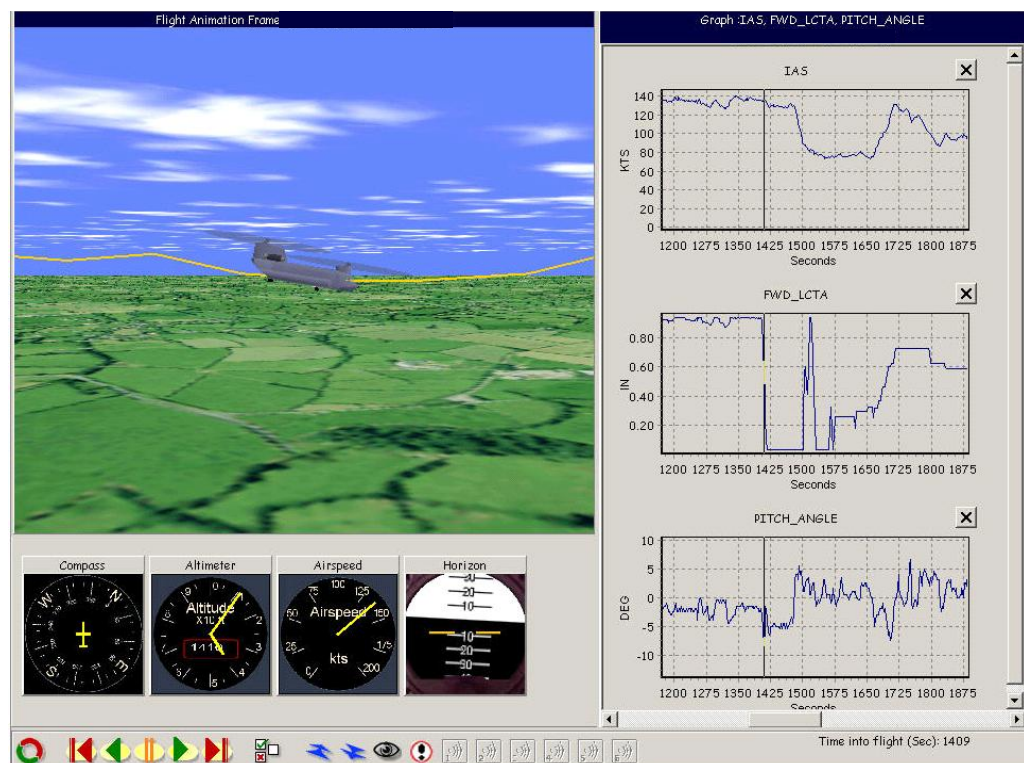


Figure 5.4: FUMS™ flight animation screenshot showing the UFCM

The various flight parameters related to UFCM events, for a typical flight, were plotted in a scatter matrix, to understand the relationships between parameters, Figure 5.5. The plot shows that the fore and aft LCTAs are well correlated, as are the two DASH actuators. In addition, these parameters are well correlated with IAS above 40 kts. These could then be used to build simple linear models, against which each flight could be compared.

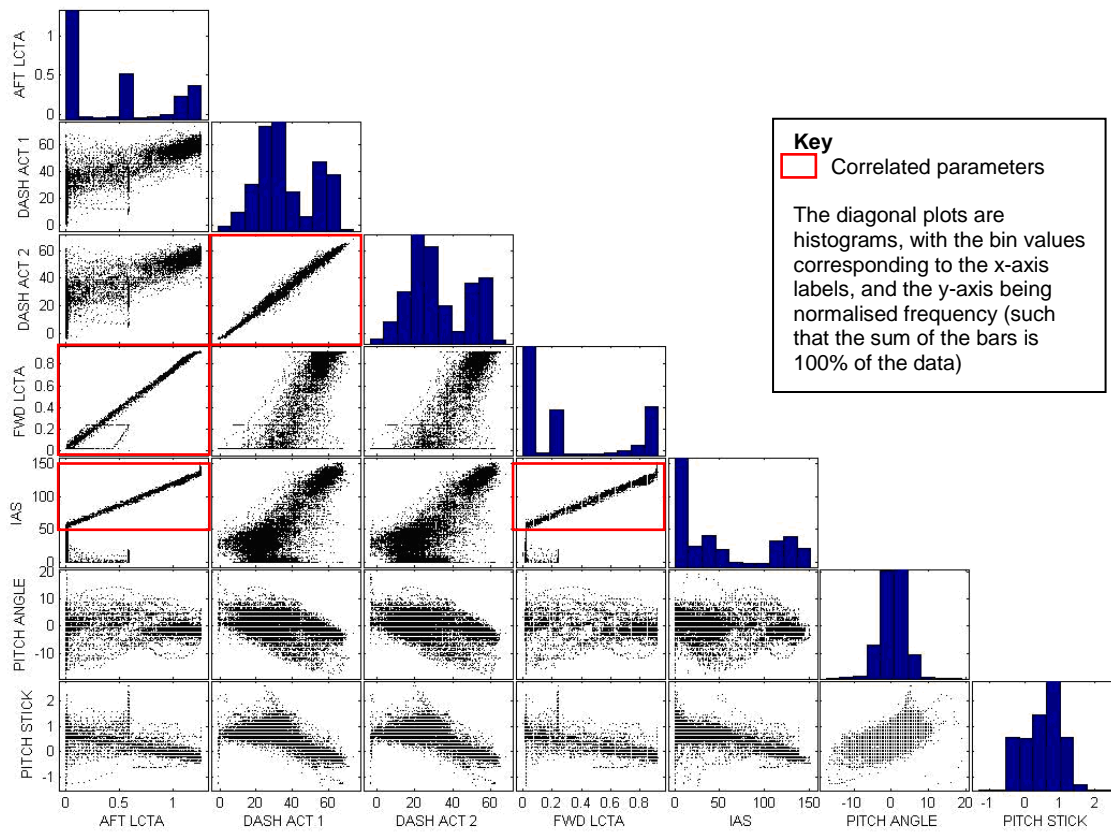


Figure 5.5: Scatter plot matrix for UFCM related Chinook parameters for a typical flight

5.3.2 Modelling the LCT Actuator

From the Chinook control laws (section 5.1), the following relationships can be derived for the LCT Actuators:

$$LCTA_{predicted} = m \times IAS + c + a \frac{PRESS_ALT}{1000} \quad (5.1)$$

Where :

$$m = \frac{LCTA_{max} - LCTA_{min}}{IAS_{max} - IAS_{min}}$$

$$c = -m \times IAS_{min}$$

$$a = 0.2^\circ$$

$$FWD_LCTA_{predicted} = \frac{m_{FWD}}{m_{AFT}} AFT_LCTA \quad (5.2)$$

Table 5.5 shows the constants for (5.1). Note that the Chinook documentation gave the LCT position in degrees; however, the data is recorded in inches. In order to convert between the two, the maximum and minimum values were extracted from the data and compared to the maximum and minimum position in degrees from the documentation. Investigating the data also revealed that the actuator reached their stops at about 140 kts at sea level, rather than the 150 kts stated in the documentation; consequently, this was used in the equations. Figure 5.6 shows pseudo-code of the algorithm that was developed to identify LCT position anomalies.

Table 5.5: Parameter Limits

Parameters	IAS Documented	IAS Used	FWD_LCTA Degrees	AFT_LCTA Degrees	FWD_LCTA Inches	AFT_LCTA Inches
Min	60	60	-1.2	-3.25	0	0
Max	150	140	4	4	0.9453	1.2891
range	90	80	5.2	7.25	0.9453	1.2891


```

Set tolerance = 0.3 Inches
Loop through all flights.
  Read the two LCTA parameters, IAS and PRESS_ALT.
  Record total data duration.
  Up sample IAS to 2Hz from 1Hz.
  Down sample PRESS_ALT to 2Hz from 16Hz.
  Filter where IAS > IASMIN (60 kts).
  Record total time above 60 kts.
  Identify where AFT_LCTA ≠ AFT_LCTApredicted ± tolerance using (5.1)
  Identify where FWD_LCTA ≠ FWD_LCTApredicted ± tolerance using (5.2)
  For each exceedance, calculate the number of events (continuous
  periods outside tolerance) and total duration.
  For each exceedance, count events and time with more than a
  minimum duration (5 seconds).

```

Figure 5.6: Pseudo-code for algorithm to identify anomalous LCTA events

The algorithm was run on all 775 available flights. Table 5.6 gives a summary of the results. The exceedances were further investigated to identify the kinds of event identified. Each figure, presented below, shows the relevant data for IAS greater than 60 kts. The figures each show four plots: Scatter charts are shown for the three relationships modelled (IAS vs. AFT_LCTA (top left), IAS vs. FWD_LCTA (top right), and FWD_LCTA vs. AFT_LCTA (bottom right)), with the actual data (blue), the predicted value (green) and any anomalies (deviations from the predicated value by more than the defined tolerance) marked (red). The bottom left plot show time series plots of these three parameters around the time of the detected anomaly. This plot also has flags indicating the periods that fall outside each model.

Table 5.6: LCTA Results Summary

	Seconds	Hours	Flights
Total	3,297,044	915.8	775
IAS>60	1,717,889	477.2	507
Exceedances	27497	7.6	79
Percent of Total	0.83%	0.83%	10.19%
Percent of IAS>60	1.60%	1.60%	15.58%

The known UFCM, discussed in section 5.3.1, was identified by the developed algorithm and is shown in Figure 5.7. Multiple exceedances were recorded for all three models.

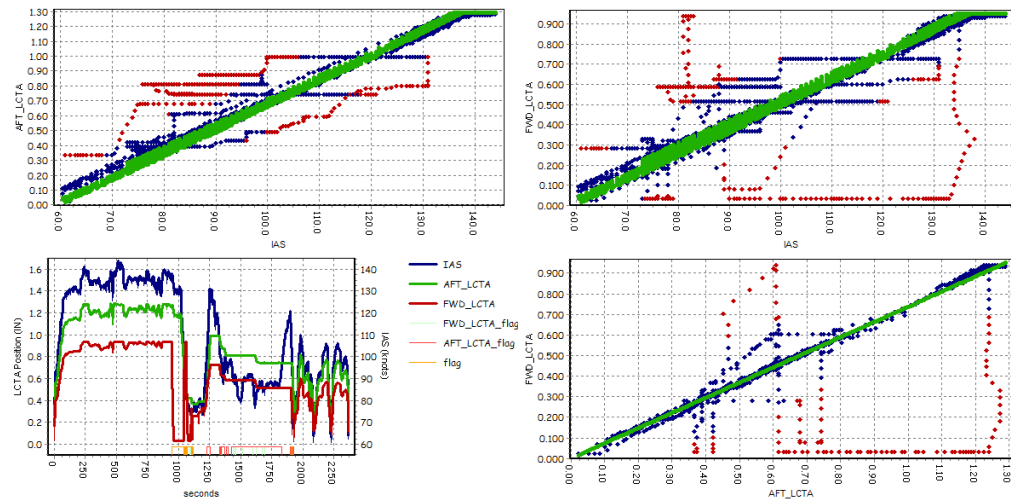


Figure 5.7: Known UFCM

Twenty-five flights, from twenty different aircraft, appeared to have LCTA oscillations similar to the one shown in Figure 5.8. It is not known what causes this kind of oscillation during flight, and they may be classified as UFCMs.

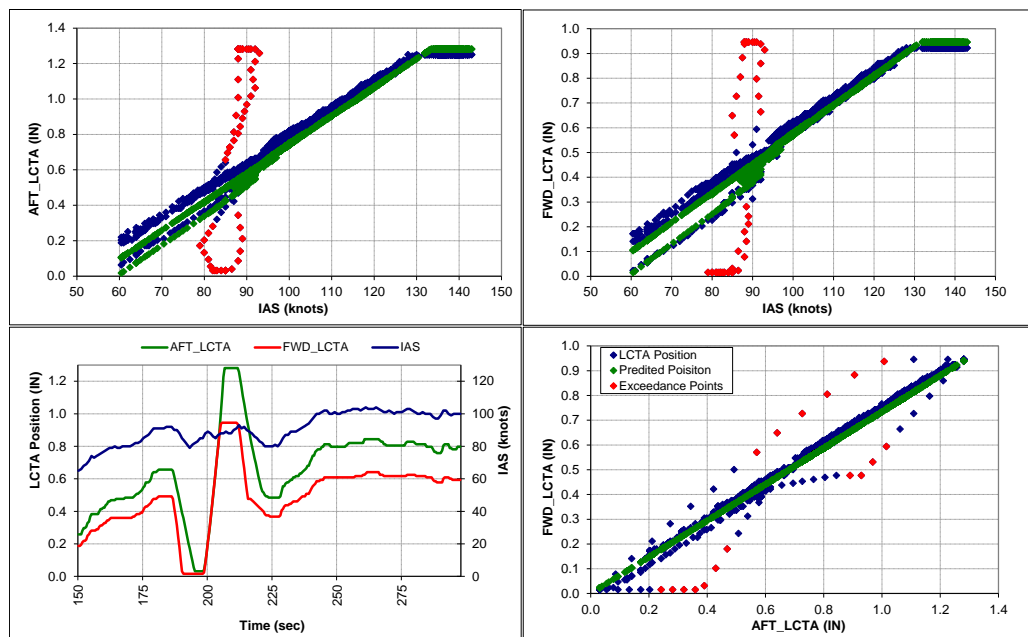


Figure 5.8: Example of possible UFCMs which have an LCTA oscillation

In addition, twelve flights, representing all flights for two aircraft, appear to have a shift in the operation of the aft LCTA. Figure 5.9 shows one of these flights.

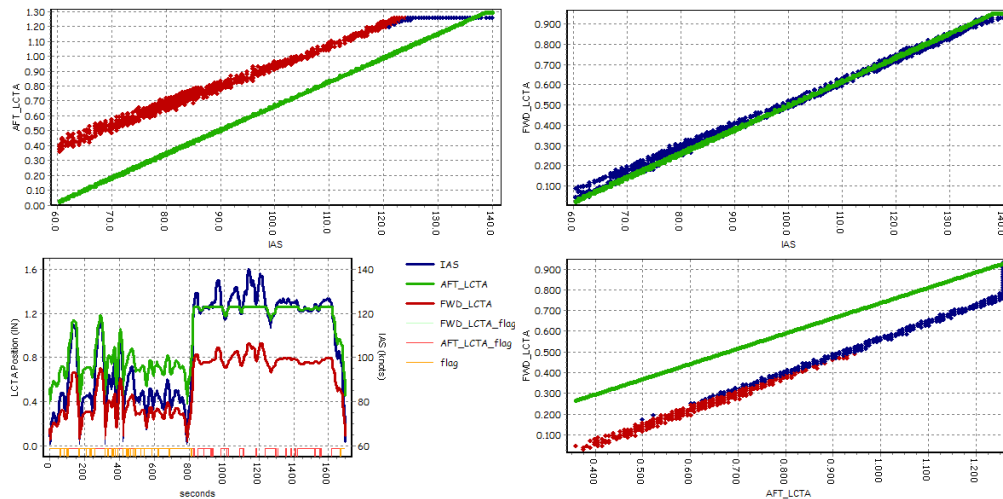


Figure 5.9: Example flight with shift in LCTA relationship

Four flights, all from one aircraft, appeared to have an AFT_LCTA failure, as the values are zero throughout the flight. This could be a sensor issue rather than the actuator being at a fixed position. The FWD_LCTA appeared to be operating correctly. Figure 5.10 shows one of these flights.




Figure 5.10: Example with AFT_LCTA stuck at zero

Twenty-nine flights had exceedances which seemed to relate to periods of transition, where the LCTAs responded more slowly than expected to changes in IAS. In most cases, the events occurred immediately before or after a change in airspeed crossing the 60 kts level. Figure 5.11 shows a flight with some of these events – note that the data shown has been filtered for IAS greater than 60 kts. These are not UFCMs and may indicate that the tolerances used to detect anomalies are too restrictive.

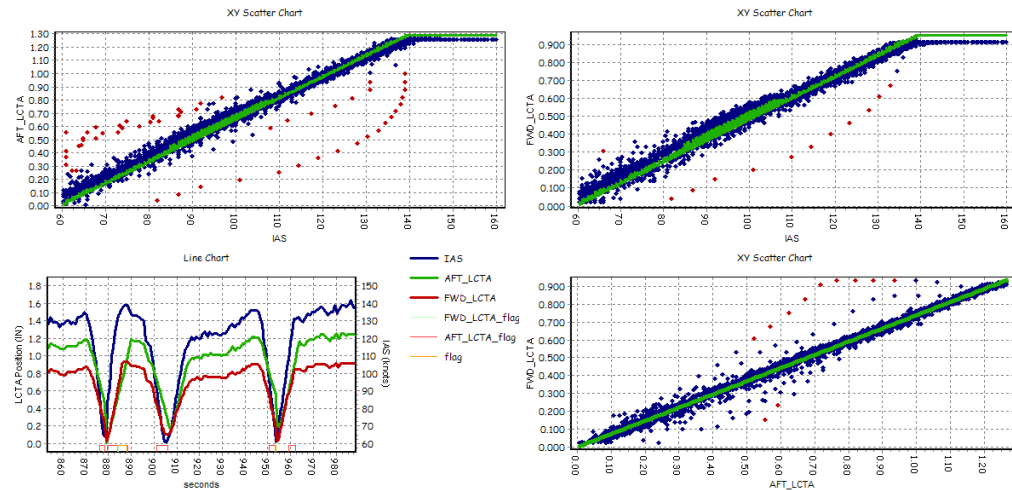


Figure 5.11: Examples of transition events

The eight remaining anomalies detected could potentially relate to UFCM events. The following figures show some of these events. Figure 5.12 and Figure 5.13 show flights where the LCTAs appeared to stick for about a minute. Figure 5.14 shows an example flight where the LCTA positions retract to zero during flight and then return to their expected positions.

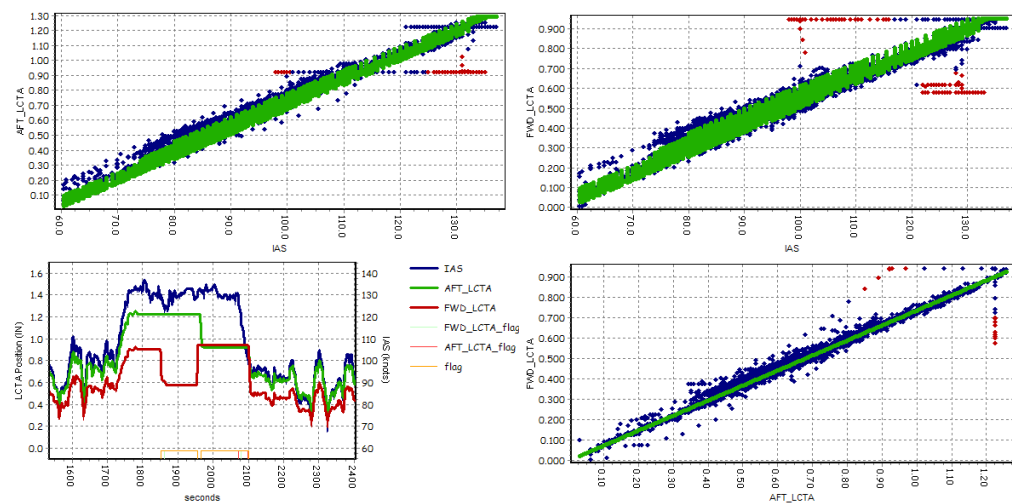


Figure 5.12: First example flight where LCTAs appear to “stick”

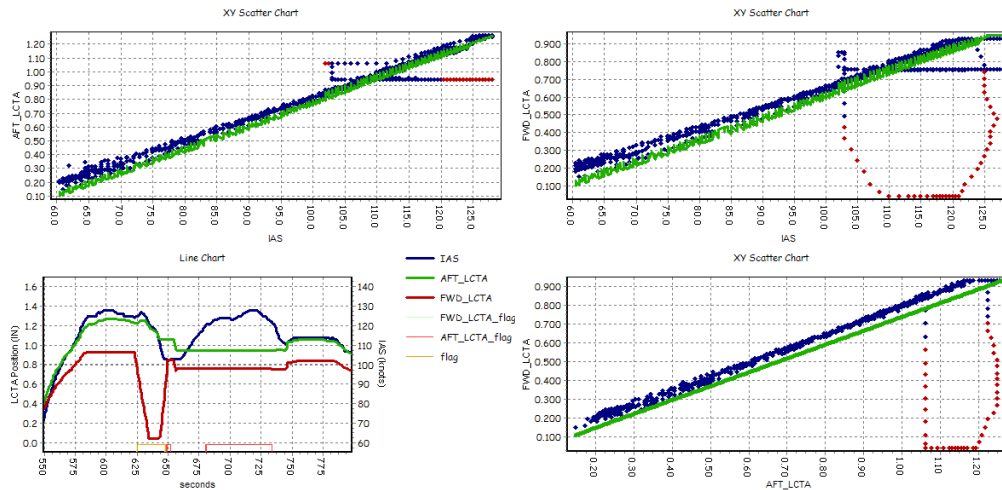


Figure 5.13: Second example flight where the LCTAs appear to “stick”

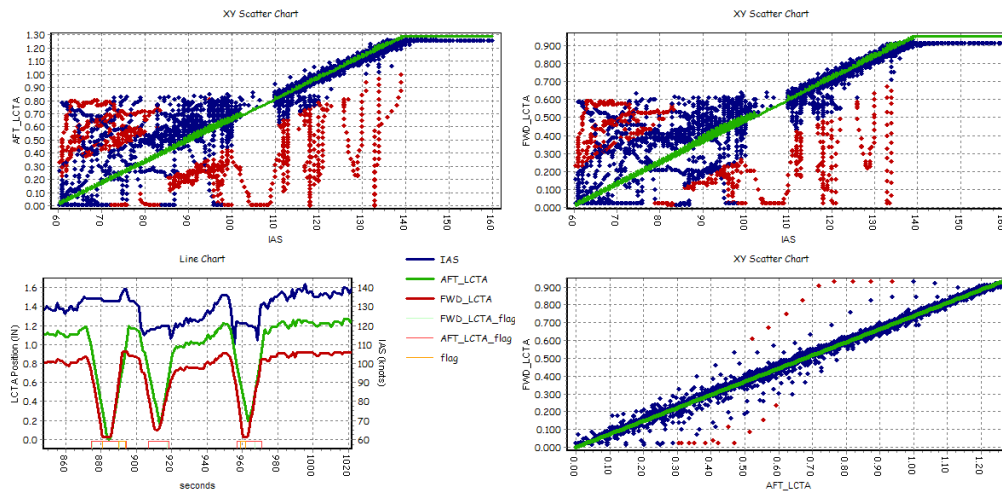


Figure 5.14: Example flight where the LCTAs appear to drop to zero during flight

Table 5.7 shows the number of flights with each of these exceedance event types. The table also shows which of the three models were exceeded. From this it may be possible to extend the algorithm to also predict the type of event that is likely to have occurred. For example, most of the transition events only exceeded the IAS vs. AFT_LCTA model, all the oscillation events exceeded either all three models or both IAS LCTA models, the shift and all zero events occurred when both the IAS vs. AFT_LCTA and the FWD vs. AFT models were exceeded.

Table 5.7: Number of LCTA anomalies detected categorized by the models exceeded and the type of event

Event Type → ↓ Models Exceeded	All zero	Known UFCM	Oscillation	Possible UFCM	Shift	Transition	Total
IAS vs. AFT						22	22
IAS vs. FWD_LCTA				1		1	2
IAS vs. AFT_LCTA FWD vs. AFT_LCTA	4				9	4	17
IAS vs. FWD_LCTA IAS vs. AFT_LCTA			8	3	3	2	16
All Three		1	17	4			22
Total	4	1	25	8	12	29	79

5.3.3 Modelling the DASH Actuator

To understand the distribution of the data for the DASH actuator, two-dimensional histograms were produced, using all 775 flights. Figure 5.15 shows the two dimensional histograms for the position information for two DASH actuators against each other and against IAS. (a – DASH actuator 1 and 2 showing the full data range, b – same as a but zoomed to expected (and most dense) range, c – DASH actuator 1 and IAS, d - DASH actuator 2 and IAS). It can be clearly seen that there is a strong but not exact correlation between the parameters. It is also clear that the DASH actuator position often appears corrupt (DA position is a percentage and should be in the range or 0-100%).

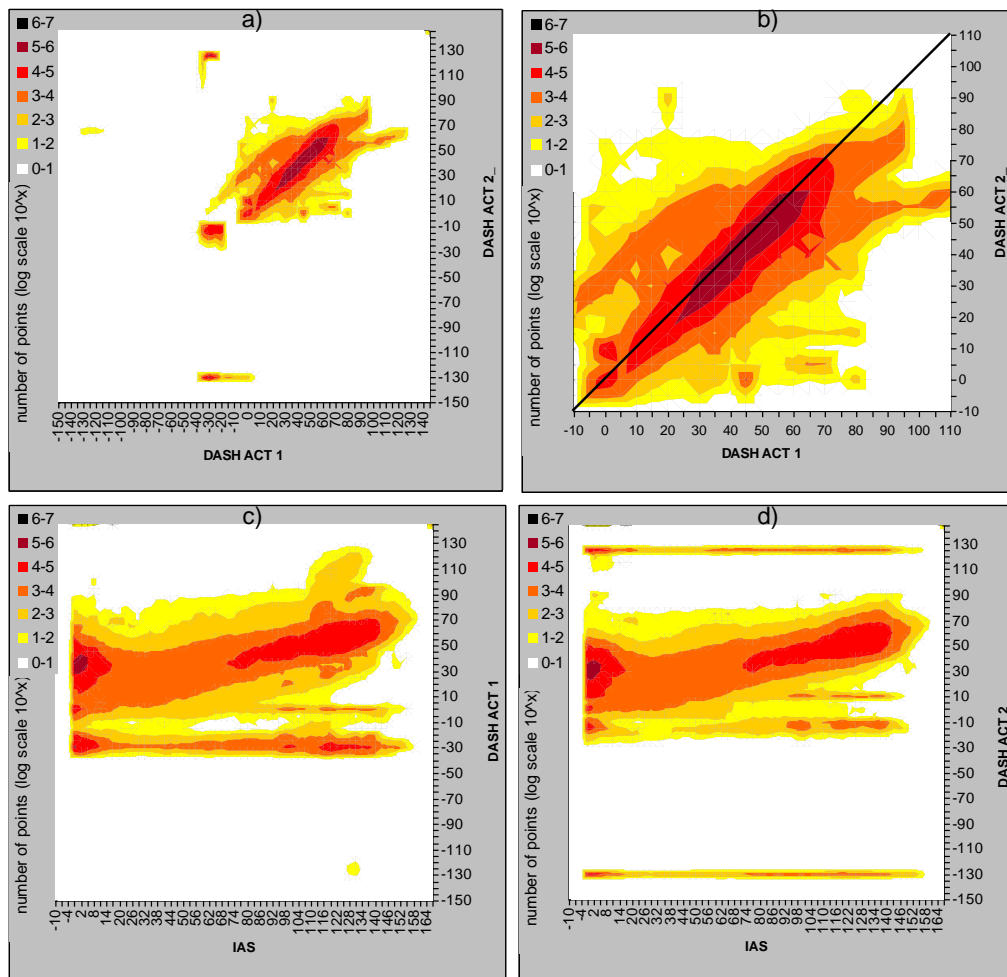


Figure 5.15: 2D histograms of 75 flights for IAS and DASH actuator positions

An algorithm was developed to identify where the DA position is out of range and to identify any points where the two DAs are not aligned within a configurable tolerance. The algorithm developed is described in pseudo-code in Figure 5.16.

```

Loop through all flights.
Read the two DASH Actuator parameters and IAS.
Resample IAS from 1Hz to 2Hz, the same as the DA data.
Filter where IAS > 40 kts (since the DA operates differently and the
    IAS parameter is unreliable below 40 kts).
Identify where DA is out of range (<0% or >100%).
Identify where DA2 ≠ DA1 ± tolerance (20%).

```

Figure 5.16: Pseudo-code for the algorithm developed to identify DA anomalies

The algorithm was run on all 775 available flights. It was found that 8.9% of flights had every single DA position out of range, even when IAS was above forty kts. For four aircraft, every single flight had these erroneous DA readings, implying some sort of data acquisition problem. These four aircraft represented 40% of the out of range flights. It was also noted that all but two of the flights with all points out of range occurred before October 2001, with only 12 flights having valid readings before October 2001. It is therefore likely that a system error caused these and was rectified in October 2001 across the fleet. An additional 4.3% of flights had some points out of range.

11.2% of flights had some points outside the 20% limit, comparing the two DA positions. This represents about 3% of flight time. Figure 5.17 shows the cumulative number of flights with percentage and actual time out of limit. The figure shows that, of the 87 flights, only 30 were out of limit for more than 5% of the flight and only 35 flights were out of limit for more than a minute.




Figure 5.17: Plots of cumulative flights with percentage time (left) and actual time (right) out of DA limit

One aircraft had every single point out of limit. Figure 5.18 shows a scatter plot of the two DAs for all eight flights from this aircraft. The solid blue line shows the expected relationship between DA1 and DA2. It appears that, for

this aircraft, DA2 is operating with approximately a 35% offset compared to DA1.




Figure 5.18: All eight flights from an aircraft with all points out of DA limits (each colour is a separate flight, the solid blue line is the expected relationship)

Five flights from one aircraft were found to have a non-linear relationship between DAs (an example is shown in Figure 5.19). A number of flights were found to have periods where one DA appeared to ‘stick’ in one position (an example is shown in Figure 5.20). Figure 5.21 shows four further examples of the kind of exceedance that was picked up by the algorithm. Further input from the MOD may reveal the causes of these anomalies.




Figure 5.19: Example flight which was found to have a non-linear DA relationship




Figure 5.20: Example flight where DA2 was found to ‘stick’ for about half an hour during flight

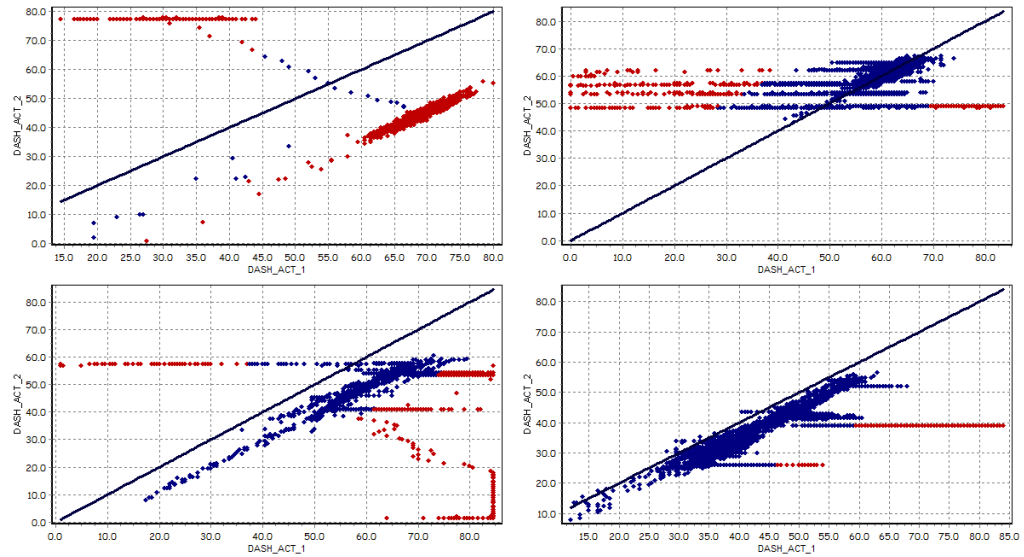


Figure 5.21: Further examples of flights with DA anomalies

5.4 UFCM Identification Conclusions

This chapter presented the analysis of Chinook UFCMs. Forty UFCM events, recorded by the MOD, were analysed for their cause and effect. Two models were developed to identify anomalies. The first model identified anomalies in the LCTA, by modelling the relationship between LCTA position and airspeed, and flagging data that did not fit this model. The resulting anomalies were investigated and categorised as being potential UFCMs, LCTA oscillations, transitions, or shifts. The second model investigated anomalies in the DA; again the results were categorised as shifts or possible events. Whilst the models have been shown to be good at identifying anomalies, feedback from the MOD would be required to tailor these to improve UFCM investigations.

CHAPTER 6

6 CONCLUSIONS

The introduction left the reader with the question of whether a system of techniques can be developed, using sensors monitoring aircraft characteristics, to provide data which can be processed (or managed) to give reliable information on the health of an aircraft, at the present time, and in the future.

This study shows that an improved management and analysis of aircraft data can indeed lead to an improved detection of faults. A study of the literature revealed that a wealth of aircraft data is being recorded and a plethora of techniques and studies are being undertaken, in an attempt to transform the data into useful information. However, it was found that although data are routinely recorded often it is not processed at all. Also, many of the previous studies have only been limited to test-bed sourced data or only a small sample of in-service data, and so need to be further and more widely validated. This research has addressed these shortfalls and makes the following key contributions:

- A new algorithm was developed to identify data integrity problems in Health Usage Monitoring System (HUMS) condition indicator (CI) data. This is a necessary precursor to any further analysis of any HUMS data.
- A HUMS intelligent health management system was developed, which was able to detect all previously documented faults many days earlier and identify new anomalies not previously recorded. The algorithm was validated using 1.8 million records of Chinook HUMS data, and has subsequently been supplied to the MOD for use on the SeaKing and Merlin helicopters. This technique can be extended for use on any HUMS.
- A new algorithm was developed to identify a snapshot of steady state performance parameters from Tornado engine HUMS data. This algorithm has the potential to save the MOD from performing pre-flight engine runs, which waste both fuel and time, and is subject to human

error in data recording. The algorithm is flexible enough to be adapted to any aircraft engine.

- New algorithms were developed to identify anomalies in the operation of Chinook differential airspeed hold and longitudinal cyclic trim actuators. The concepts behind the algorithm could be applied to other aircraft systems.
- These developed algorithms were validated using huge volumes of in-service aircraft data, covering a wide range of operating conditions:
 - The Chinook HUMS audit algorithm successfully identified just under half a million inconsistencies in 46 million CI values.
 - The HUMS intelligent health management system identified 83% of known faults with a false positive rate of just 13%, after analysing 1.8 million CI values from 40 aircraft.
 - The Tornado EHUMS analysis identified 60 anomalies, of which 30 related to known issues, in over 27 thousand engine runs from 236 engines.
 - The Chinook flight system anomaly detection algorithm was applied to 775 flights from 35 aircraft, representing 910 hours of data, totalling over 230 million data points. 79 flights were found to have longitudinal cyclic trim actuator anomalies, and 87 flights had differential airspeed hold actuator anomalies.
- The study has demonstrated the application of a framework for managing aircraft data that includes: data acquisition and storage, data integrity checking, feature extraction, fusion, decision support and a user interface. It is the recommendation of this study that aircraft operators consider all of these aspects when seeking to manage the health of their assets.
- Throughout, algorithm development was not focussed on a specific aircraft. The methodology was applied in a systematic manner enabling the underlying mathematical approach to stand on its own and thereby be capable of extending naturally to other situations and air vehicles.

In a final set of remarks, the thesis describes a research effort directed at various problem areas and components, including the pilot, of both fixed and rotary wing aircraft. It identifies the characteristics of typical data and how the operation of sensors, installed in order to protect the aircraft, must themselves also become part of the data review and subsequent management.

The majority of the research was directed towards a well-established transport helicopter. The transmission of a helicopter is both vital and complicated. Failure is not normally survivable and so the many components must be monitored and accurate and reliable decisions must be made by the system. This requires solid decision making from the algorithms to sift the data for the information required to assure the safe operation of the helicopter.

The aspect of engine testing, on a long-established fighter aircraft, whilst on a sortie, was addressed. The present technique, using full throttle operation before take-off, is both expensive and problematic. It consumes fuel, adds to the structural loads on the airframe and engine and asks an additional duty of the pilot when on the threshold of a military sortie. Data was examined and processed to give reliable and effective information for the successful operation of the engines.

The flight of an aircraft is under the joint control of a pilot and the flight control system. The latter can be analysed, since it is based on mathematical techniques. The former, however, being human, is a different proposition. The vagaries of human intervention in the system need to be isolated in order to make sense of the incoming data. One pilot may fly with soft hands, whilst another may be rather heavy-handed in technique. This disparity is likely in the operation of a military aircraft, particularly a helicopter, in contrast to civil airlines, which largely make use of an autopilot. To be of any use, the data management technique developed must be capable of separating out the differences in piloting technique from a change in airframe characteristics, which may, in turn, indicate a developing fault.


It was considered important not to focus on specific airframes which would render any results rather limited. The study used data from actual aircraft flights but always sought to examine the wider picture. In this way the techniques, methods, software and user interface development would be applicable generally and therefore move the field forward.

As with all research, there are always more questions generated than answered. The goal will always be to eliminate air accidents entirely. A noble aspiration but not guaranteed as yet. The work described in this thesis moves

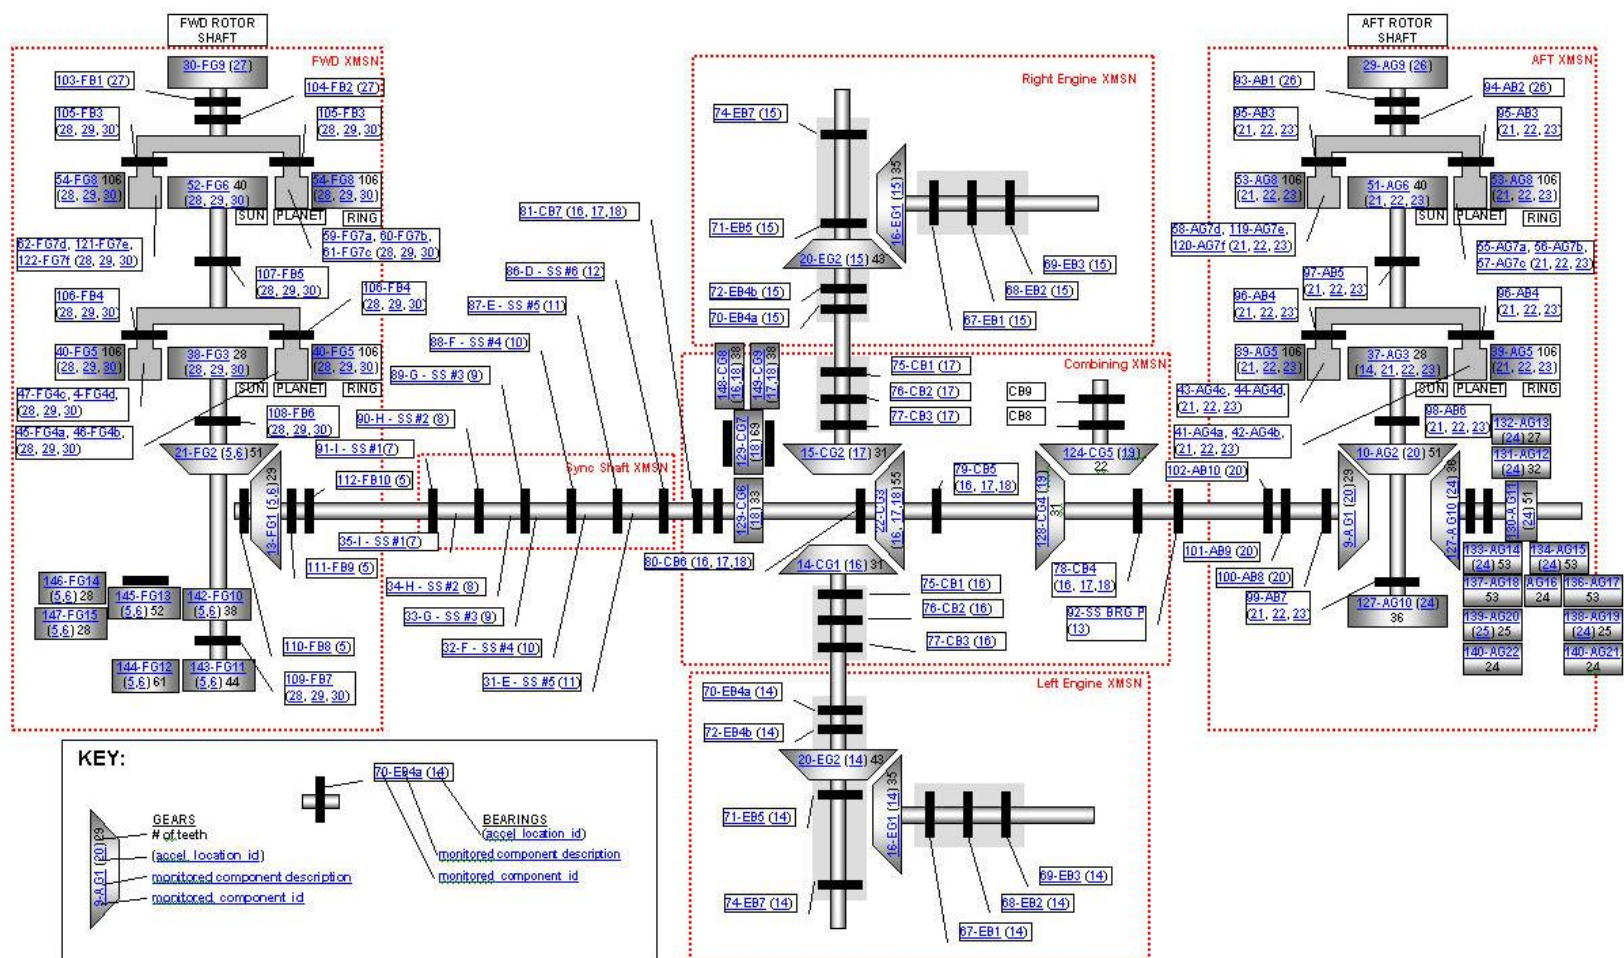
the subject to a point of establishing a set of methods, which are documented and verifiable. Improvements in hardware and the associated developments in software and supporting mathematical techniques will move on from this point towards the final goal. It is hoped that the results and conclusions described in this thesis will provide a solid foundation for this further work.

APPENDIX A: CHINOOK HUMS INFORMATION

A1.1 Chinook Driveshaft Configuration



A1.2 Schematic of Chinook Component and Accelerometer Locations



A1.3 Monitored Components

id	description	Accel Location
9	AG1 - Input Pinion Gear	20
10	AG2 - Bevel Ring Gear	20
13	FG1 - Input Pinion Gear	5,6
14	CG1 - LH Input Pinion Gear	16
15	CG2 - RH Input Pinion Gear	17
16	EG1 - LH Input Pinion Gear	14
17	EG3 - RH Input Pinion Gear	15
20	EG2 - LH Output Gear	14
21	FG2 - Bevel Ring Gear	5,6
22	CG3 - Collector Gear	16,17,18
29	AG9 - AFT Rotor Shaft	26
30	FG9 - FWD Rotor Shaft	27
31	E - Sync Shaft #5	11
32	F - Sync Shaft #4	10
33	G - Sync Shaft #3	9
34	H - Sync Shaft #2	8
35	I - Sync Shaft #1	7
36	P - Sync Shaft #8	13,21,28
37	AG3 - 1st Stage Sun Gear	21,22,23
38	FG3 - 1st Stage Sun Gear	28,29,30
39	AG5 - 1st Stage Ring Gear	21,22,23
40	FG5 - 1st Stage Ring Gear	28,29,30
41	AG4a - 1st Stage Planet Gear 1	21,22,23
42	AG4b - 1st Stage Planet Gear 2	21,22,23
43	AG4c - 1st Stage Planet Gear 3	21,22,23
44	AG4d - 1st Stage Planet Gear 4	21,22,23
45	FG4a - 1st Stage Planet Gear 1	28,29,30
46	FG4b - 1st Stage Planet Gear 2	28,29,30
47	FG4c - 1st Stage Planet Gear 3	28,29,30
48	FG4d - 1st Stage Planet Gear 4	28,29,30
51	AG6 - 2nd Stage Sun Gear	21,22,23
52	FG6 - 2nd Stage Sun Gear	28,29,30
53	AG8 - 2nd Stage Ring Gear	21,22,23
54	FG8 - 2nd Stage Ring Gear	28,29,30
55	AG7a - 2nd Stage Planet Gear 1	21,22,23
56	AG7b - 2nd Stage Planet Gear 2	21,22,23
57	AG7c - 2nd Stage Planet Gear 3	21,22,23
58	AG7d - 2nd Stage Planet Gear 4	21,22,23
59	FG7a - 2nd Stage Planet Gear 1	28,29,30
60	FG7b - 2nd Stage Planet Gear 2	28,29,30
61	FG7c - 2nd Stage Planet Gear 3	28,29,30
62	FG7d - 2nd Stage Planet Gear 4	28,29,30
65	EG4 - RH Output Gear	15
66	D - Sync Shaft #6	12
67	EB1 - LH Input Pinion Bearing	14
68	EB2 - LH Input Pinion Bearing	14
69	EB3 - LH Input Pinion Bearing	14
70	EB4a - LH Output Gear Bearing	14
71	EB5 - LH Output Gear Bearing	14
72	EB4b - LH Output Gear Bearing	14
73	EB6 - LH Output Gear Bearing	14
74	EB7 - LH Clutch Shaft	14
75	CB1 - LH Input Pinion Bearing	16,17,18
76	CB2 - LH Input Pinion Bearing	16,17,18
77	CB3 - LH Input Pinion Bearing	16,17,18
78	CB4 - Combiner Gear	16,17,18
79	CB5 - Combiner Gear	16,17,18
80	CB6 - Combiner Gear	16,17,18
81	CB7 - Combiner Quill	16,17,18
82	EB1,2,3 LH Transmission	14
83	EB4,5,6,7 LH Transmission	14
84	CB1/2/3 - LH Combiner Bearings	16
85	CB4,5,6,7 Combiner Bearings	18
86	D - Sync Shaft #6 Bearing	12
87	E - Sync Shaft #5 Bearing	11
88	F - Sync Shaft #4 Bearing	10
89	G - Sync Shaft #3 Bearing	5,9
90	H - Sync Shaft #2 Bearing	8
91	I - Sync Shaft #1 Bearing	7
92	P - Sync Shaft #8 Bearing	13
93	AB1 - AFT Rotor Shaft Bearing	26
94	AB2 - AFT Rotor Shaft Bearing	26
95	AB3 - 2nd Stage Planet Bearing	21,22,23
96	AB4 - 1st Stage Planet Bearing	21,22,23
97	AB5 - Sun Gear Shaft Bearing	21,22,23
98	AB6 - Sun Gear Shaft Bearing	20,21,22,23

id	description	Accel Location
99	AB7 - Sun Gear Shaft Bearing	20,21,22,23
100	AB8 - Input Pinion Bearing	20
101	AB9 - Input Pinion Bearing	20
102	AB10 - Input Pinion Bearing	20
103	FB1 - FWD Rotor Shaft Bearing	27
104	FB2 - FWD Rotor Shaft Bearing	27
105	FB3 - 2nd Stage Planet Bearing	28,29,30
106	FB4 - 1st Stage Planet Bearing	6,28,29,30
107	FB5 - Sun Gear Shaft Bearing	28,29,30
108	FB6 - Sun Gear Shaft Bearing	6,28,29,30
109	FB7 - Sun Gear Shaft Bearing	6,28,29,30
110	FB8 - Input Pinion Bearing	5
111	FB9 - Input Pinion Bearing	5
112	FB10	5
113	AB1,2 Bearing Group	26
114	AB5,6,7 Bearing Group	20
115	AB8,9,10 Bearing Group	20
116	FB1,2 Bearing Group	27
117	FB5,6,7 Bearing Group	6
118	FB8,9,10 Bearing Group	5
119	AG7e - 2nd Stage Planet Gear 5	21,22,23
120	AG7f - 2nd Stage Planet Gear 6	21,22,23
121	FG7e - 2nd Stage Planet Gear 5	28,29,30
122	FG7f - 2nd Stage Planet Gear 6	28,29,30
124	CG5 - Oil Cooler Blower Fan	19
126	CG7 - Pump Drive Gear	18
127	AG10 - Spiral Bevel Acc Drv	24
128	CG4 - Blower Drive Gear	19
129	CG6 - Pump Drive Gear	18
130	AG11 - Spur Accessory Drive	24
131	AG12/14/17/18 Pumps and Idler	24
132	AG13 - Not Used	24
133	AG14 - Not Used	24
134	AG15/21 Idler Gears	24
135	AG13/16/22 Alternator Gears	24
138	AG19 - Idler Gear	24
139	AG20 - Oil Cooler Fan	24,25
142	FG10 - Aux Lube Drive Gear	5,6
143	FG11 - Aux Lube Pump Drive Gear	5,6
144	FG12 - Drive Gear	5,6
145	FG13 - Idler Gear	5,6
146	FG14 - Main Lube Pump Drive	5,6
147	FG15 - Hyd Flight Cont Pump Drv	5,6
148	CG8 - LH Pump Drive Gear	16,18
150	EB1 - RH Input Pinion Bearing	15
151	EB2 - RH Input Pinion Bearing	15
152	EB3 - RH Input Pinion Bearing	15
153	EB4a - RH Output Gear Bearing	15
154	EB5 - RH Output Gear Bearing	15
155	EB4b - RH Output Gear Bearing	15
156	EB6 - RH Output Gear Bearing	15
157	EB7 - RH Clutch Shaft	15
158	EB1,2,3 RH Transmission	15
159	EB4,5,6,7 RH Transmission	15
160	Main accessory drive AB12 AFT	24
161	Main accessory drive AB11 AFT	24
162	Idler B ball 206 AFT	24
163	Lube pump G ball 106 AFT	24
165	Hydraulic pump D/H ball 107 AFT	24
166	Idler E/K ball 106 AFT	24
167	Idler I ball 107 AFT	24
168	Alternator C/F/L ball 106 AFT	24
169	Alternator C/F/L ball 107 AFT	24
170	Fan AB13 ball 106 AFT	24
171	Fan AB13 ball 107 AFT	24
172	Main Accessory Drive AFT	24
173	Pumps AFT	24
174	Alternators AFT	24
176	Oil Cooler Fan AFT	24
179	Blower gear bearings CB8/9 CMB	18
180	Lube pump drive ball 106 FWD	6
181	Idler ball 107 FWD	6
182	Hyd flt con pump ball 107 FWD	6
183	Lube / Hydraulic pumps FWD	6
186	CB1 - RH Input Pinion Bearing	17
187	CB2 - RH Input Pinion Bearing	17
188	CB3 - RH Input Pinion Bearing	17
189	CB1/2/3 - RH Combiner Bearings	17

A1.4 Accelerometer Locations

id	description	Long description	Orientation	monitored_component_id
1	FWD HEAD (lat)	MForward Head Lateral (ch.28)	Lateral	
2	FWD HEAD (vert)	Forward Head Vertical (ch.27)	Vertical	
3	AFT HEAD (lat)	Aft Head Lateral (ch.30)	Lateral	
4	AFT HEAD (vert)	Aft Head Vertical (ch.29)	Vertical	
5	FWD XMSN (rad ID5)	IFD pad I/P shaft (ch.5)	Radial	13, 21, 89, 110, 111, 112, 118, 142, 143, 144, 145, 146, 147
6	FWD XMSN (rad ID6)	Bell crank mounting bolt (ch.6)	Radial	13, 21, 106, 108, 109, 117, 142, 143, 144, 145, 146, 147, 180, 181, 182, 183
7	Sync Shaft (rad ID7)	Sync shaft #1 (ch.7)	Radial	35, 91
8	Sync Shaft (rad ID8)	Sync Shaft #2 (ch.8)	Radial	34, 90
9	Sync Shaft (rad ID9)	Sync Shaft #3 (ch.9)	Radial	33, 89
10	Sync Shaft (rad ID10)	Sync Shaft #4 (ch.10)	Radial	32, 88
11	Sync Shaft (rad ID11)	Sync Shaft #5 (ch.11)	Radial	31, 87
12	Sync Shaft (rad ID12)	Sync Shaft #6 (ch.12)	Radial	66, 86
13	Sync Shaft (rad ID13)	Sync Shaft #8 (ch.13)	Radial	36, 92
14	LH ENG XMSN (RAD ID14)	Left hand eng xmsn (radial to cross shaft) (ch.14)	Radial	16, 20, 67, 68, 69, 70, 71, 72, 73, 74, 82, 83
15	RH ENG XMSN (RAD ID15)	Right hand eng xmsn (radial to cross shaft) (ch.15)	Radial	17, 65, 150, 151, 152, 153, 154, 155, 156, 157, 158, 159
16	COMB XMSN (Rad ID16)	Combiner Left IFD pad (Radial to cross shaft) (ch.16)	Radial	14, 22, 75, 76, 77, 78, 79, 80, 81, 84, 148
17	COMB XMSN (Rad ID17)	Combiner Right IFD pad (Radial to cross shaft) (ch.17)	Radial	15, 22, 75, 76, 77, 78, 79, 80, 81, 186, 187, 188, 189
18	COMB XMSN (Ax ID18)	Combiner Aft o/p IFD pad (Axial to O/P shaft) (ch.18)	Axial	22, 75, 76, 77, 78, 79, 80, 81, 85, 126, 129, 148, 179
19	COMB XMSN (Rad ID19)	Combiner Cooling Fan rear left (Radial to fan) (ch.19)	Radial	124, 128
20	AFT XMSN (rad ID20)	Aft XMSN IFD pad I/P shaft (Rad to I/P shaft) (ch.20)	Radial	9, 10, 98, 99, 100, 101, 102, 114, 115
21	AFT XMSN (rad ID21)	Aft upper ring gear bolt 12 o'clk (Rad to O/P shaft) (ch.21)	Radial	36, 37, 39, 41, 42, 43, 44, 51, 53, 55, 56, 57, 58, 95, 96, 97, 98, 99, 119, 120
22	AFT XMSN (rad ID22)	Aft upper ring gear bolt 4 o'clk (Rad to O/P shaft) (ch.22)	Radial	37, 39, 41, 42, 43, 44, 51, 53, 55, 56, 57, 58, 95, 96, 97, 98, 99, 119, 120
23	AFT XMSN (rad ID23)	Aft upper ring gear bolt 7 o'clk (Rad to O/P shaft) (ch.23)	Radial	37, 39, 41, 42, 43, 44, 51, 53, 55, 56, 57, 58, 95, 96, 97, 98, 99, 119, 120
24	AFT XMSN (vert ID24)	Aft Lower accessories drv csing 6o'clk (Vert) (ch.24)	Vertical	127, 130, 131, 132, 133, 134, 135, 138, 139, 160, 161, 162, 163, 165, 166, 167, 168, 169, 170, 171, 172, 173, 174, 176
25	AFT XMSN (rad ID25)	Aft XMSN cooling fan (Radial to fan) (ch.25)	Radial	139
26	AFT XMSN (Rad ID26)	Aft vert shaft 12o'clk (Radial to O/P shaft) (ch.26)	Radial	29, 93, 94, 113
27	FWD XMSN (Rad ID1)	Main lift IFD pad (Radial to O/P shaft) (ch.1)	Radial	30, 103, 104, 116
28	FWD XMSN (Rad ID2)	Lower ring gear bolts (11o'clk) (ch.2)	Radial	36, 38, 40, 45, 46, 47, 48, 52, 54, 59, 60, 61, 62, 105, 106, 107, 108, 109, 121, 122
29	FWD XMSN (Rad ID3)	Upper ring gear bolts (3 o'clk) (ch.3)	Radial	38, 40, 45, 46, 47, 48, 52, 54, 59, 60, 61, 62, 105, 106, 107, 108, 109, 121, 122
30	FWD XMSN (Rad ID4)	Upper ring gear bolts 7 o'clk (ch.4)	Radial	38, 40, 45, 46, 47, 48, 52, 54, 59, 60, 61, 62, 105, 106, 107, 108, 109, 121, 122
31	FWD Frame 95 Vertical	Beneath access panel vertical	Vertical	
32	FWD Frame 95 Lateral	Beneath access panel lateral	Lateral	
33	AFT Frame 534 Vertical	STN 534 cabin roof mounted onto bulkhead vertical	Vertical	
34	AFT Frame 534 Lateral	STN 534 cabin roof mounted onto bulkhead lateral	Lateral	
35	STVA 2 (NOSE)	Right frame of STVA mount	Vertical	
36	STVA 3 (STBD)	Beneath pilot seat (right)	Vertical	
37	STVA 1 (PORT)	Beneath co-pilot seat (left)	Vertical	
38	LH ENG Pos 1, Radial	Left Engine, Position 1, Radial	Radial	
39	LH Eng Pos 2 Radial	Left Engine Position 2 Radial	Radial	
40	LH Eng Pos 3 Axial	Left Engine Position 3 Axial	Axial	
41	RH ENG Pos 1, Radial	Right Engine, Position 1, Radial	Radial	
42	RH ENG Pos 2, Radial	Right Engine, Position 2, Radial	Radial	
43	RH ENG Pos 3, Axial	Right Engine, Position 3, Axial	Axial	
44	Unused B1 (20)	Unused Board 1 Accelerometer Mux 20	Unused	
45	Unused B1 (21)	Unused Board 1 Accelerometer Mux 21	Unused	
46	Unused B1 (22)	Unused Board 1 Accelerometer Mux 22	Unused	
47	Unused B1 (23)	Unused Board 1 Accelerometer Mux 23	Unused	
48	Unused B1 (24)	Unused Board 1 Accelerometer Mux 24	Unused	
49	Calib B1 (0)	Calibration Board 1 Accelerometer Mux 0	Calibration	
50	Calib B2 (0)	Calibration Board 2 Accelerometer Mux 0	Calibration	

A1.5 Passbands

passband_id	Description	Short Name
5	Broad Band 100 - 40000 Hz (1638)	1
10	Broad Band 100 - 10000 Hz (6553)	1
15	Broad Band 100 - 40000 Hz (6553)	1
20	Broad Band 100 - 12000 Hz (6553)	1
25	Broad Band 100 - 40000 Hz (3276)	1
6	Second Band 100 - 10000 Hz (163)	2
11	Second Band 100 - 2500 Hz (6553)	2
16	Second Band 100 - 10000 Hz (655)	2
21	Second Band 100 - 2500 Hz (6553)	2
26	Second Band 100 - 10000 Hz (327)	2
7	Third Band 10000 - 20000 Hz	3
12	Third Band 2500 - 5000 Hz	3
17	Third Band 10000 - 20000 Hz	3
22	Third Band 2500 - 5000 Hz	3
27	Third Band 10000 - 20000 Hz	3
8	Fourth Band 20000 - 40000 Hz	4
13	Fourth Band 5000 - 10000 Hz	4
18	Fourth Band 20000 - 40000 Hz	4
23	Fourth Band 5000 - 10000 Hz	4
28	Fourth Band 20000 - 40000 Hz	4

A1.6 Bearings and Gear Terminology

The following figures show the part that make up some of the common gears and bearings monitored by a HUMS system.




Figure 6.1: Rolling bearing components, [114]




Figure 6.2: Front-view cross-section diagram of a planetary gearbox, [115]

A1.7 Data Acquisition and Labelling

The MOD has approximately forty Chinook helicopters, each of which is identified by a unique aircraft tail number. Every time the onboard HUMS system is powered up, a session number, called a HUMS Recording Session (HRS) number (`hrs_number`), is assigned. HUMS vibration is acquired when the aircraft enters pre-defined flight regimes, and the vibration data is then sorted and pre-processed in the Vehicle Airborne Unit (VAU). The datasets recorded during a sequence of HRSs, which can cover up to twenty five flights, are downloaded to the Ground Support System (GSS). After downloading data from HUMS to GSS, the sequence identifier (`sequence_id`) is incremented and the HRS number is reset. The aircraft serial numbers, sequence IDs, HRS numbers and associated datasets are stored in relevant HUMS database tables. The HUMS relational database consists of over four hundred tables. Each data item is also identified within the HUMS database by a date/time stamp (`display_date_time`). In this thesis, the word 'flight' refers to an 'HRS', which can include ground-run data or a complete dataset of a flight that can involve one or more take-offs and landings. Various CIs (identified by the `drivetrain_indicator_type_id`) can be computed, from both raw and band-pass filtered vibration, measured by an accelerometer using information about the monitored component (e.g. a gear meshing frequency). A schematic of the Chinook monitored components is shown in section A1.2 including descriptions of all components and accelerometers. Therefore, eight HUMS IDs are required to uniquely identify each CI value:

- 'aircraft_serial_number'
- 'sequence_id'
- 'hrs_number'
- 'display_date_time'
- 'accel_location_id'
- 'monitored_component_id'
- 'passband_id'
- 'drivetrain_indicator_type_id'

The raw vibration data, from which the CIs are computed, can be stored by HUMS in Binary Large Object (BLOB) files. This data can be used to validate the onboard processing, and/or perform additional processing if required. Snapshots of FDR parameters can also be recorded within the BLOB files, which can give an indication of how the aircraft was flying when the data was acquired.

Component CIs are computed from vibration measurements acquired by accelerometers, using diagnostic configuration information, such as shaft speeds and meshing frequencies. The derived signal is identified by the accelerometer identifier and the monitored component identifier. In this thesis a derived signal is identified using the first word of the description of the monitored component, followed by the accelerometer identifier number in brackets. For example, CB1 (16) refers to a signal derived from the accelerometer sixteen to monitor the first input pinion Combiner Bearing (CB). A particular CI is identified by its abbreviated name (see Table 6.2) and the filter pass-band identifier, for example: CB1(16) ETE_M6-5.

A1.8 Chinook Health Indicators

The following paragraphs provide brief descriptions of the terms and signal processing techniques used in the on-board HUM system to calculate CIs, [98, 116].

Time Series: A series of values representing consecutive samples of a signal equally spaced in time.

Time Average: A way of averaging data, performed by splitting a time series into adjacent blocks of equal length, then computing the average of the corresponding samples in each block. This has the effect of reducing all components of the original data, except those whose frequency is periodic with the length of the block.

Spectrum: A representation of the frequency composition of a signal, usually its Fourier Transform. Each sample (or bin) is a complex number representing magnitude and phase.

Spectral Enhancement: A way of processing data, often applied to a time average, performed by computing its spectrum, modifying the values of individual bins (setting some values to zero), then converting the spectrum back to a time signal, usually via the inverse Fourier Transform. This is useful for reducing or emphasising specific frequency components.

Gear Mesh Frequency: The frequency at which the teeth of meshing gears pass through the region of contact. The vibration signal generated usually contains components at this frequency and its harmonics.

Power Spectrum: A representation of the power of the frequency content of a signal. Each bin is a real number, representing the square of the magnitude of the corresponding bin in the spectrum.

Power Average: A way of averaging data, performed by splitting a time series into adjacent blocks of equal length, calculating the power spectrum of each block, then computing the average of corresponding bins in each power spectrum. This has the effect of reducing all components of the original data, except those whose frequency is relatively constant. The square root of the final result is often used, giving average amplitudes.

Band-pass Filtering: A process of reducing all components of the original data with frequencies outside a specified band.

Envelope: If a signal consists of high frequencies, whose amplitudes change slowly over time, the envelope is a low frequency signal, with values corresponding to the peaks of the high frequency signal when they occur and changing smoothly between them. It is often calculated as the magnitude of the Analytic Signal.

Analytic Signal: A complex-valued representation of a real-valued signal, having an imaginary component in phase quadrature (shifted by ninety degrees) with the real component.

A1.9 GenHUMS Analysis

The Generic HUMS (GenHUMS) bearing vibration analysis process, used on the Chinook aircraft, consists of two elements. In the primary analysis, the acquired vibration data (SIG) is band-pass filtered, the filtered signal is “enveloped”, and the enveloped signal (ETE) is then converted into the frequency domain to produce an envelope power spectrum (FTE). Secondary analysis involves the calculation of a series of Condition Indicators (CIs) at different stages of the primary analysis process. The CI’s include pattern matching indicators which search for defect patterns associated with damage on bearing components: both Energy Index (EI) and Detection Index (DI) (a percentage score) are calculated for the inner race (IN), outer race (OU), cage (CG) and rolling elements (EL). Figure 6.4 shows example frequencies and pattern match weights. Tones are matched if a peak is found within 1% of the defect frequency. The pattern match logic is based on the log of the envelope power spectrum, with a moving median filter, and a threshold of the mean plus two standard deviations. Figure 6.3 illustrates the primary analysis performed by GenHUMS.

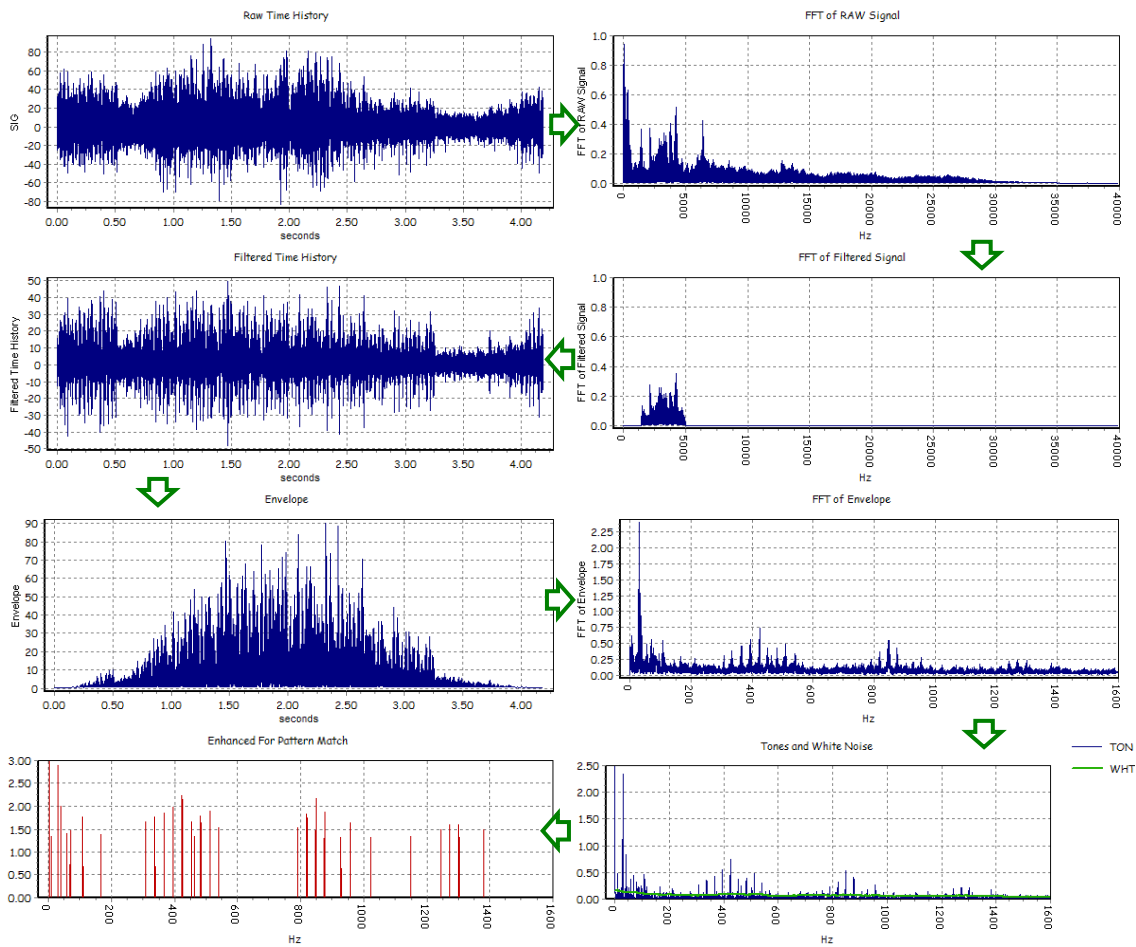


Figure 6.3: Example GenHUMS primary analysis




Figure 6.4: Example bearing defect tone frequencies and pattern match weights

Table 6.1 shows some of the most commonly calculated CIs.

Table 6.1: CI name and descriptions

Name	Description
PP	“Peak-to-peak”. Difference between the maximum and minimum data values.
PK	“Peak”. Maximum absolute data value.
STD	Standard deviation of the data.
MN	Mean of the data.
MRC	“Maximum Rate of Change”. Maximum absolute difference of adjacent data values.
M6	Normalised 6 th moment of the data.
IMP	“Impulsiveness”. Normalised 4 th moment (kurtosis) of the data.
EB	“Band Energy”. RMS value of the data.
TON	Tonal energy of the band-passed envelope spectrum
WHT	White noise energy of the band-passed envelope spectrum
IN	EI for tone pattern generated by inner race damage
ID	DI for tone pattern generated by inner race damage
OU	EI for tone pattern generated by outer race damage
OD	DI for tone pattern generated by outer race damage
EL	EI for tone pattern generated by rolling element damage
ED	DI for tone pattern generated by rolling element damage
CG	EI for tone pattern generated by cage damage
CD	DI for tone pattern generated by cage damage
SOn	Magnitude of nth shaft order vibration

Specific CIs are calculated from the vibration data, after various types of processing. These pre-processing techniques are described below:

“Original signal” (SIG) applies to both gears and bearings. For gears, it refers to the time average of the vibration data. For bearings, it refers to the unprocessed vibration time series.

“Enhanced signal” (ESA) is specific to gears. It refers to the residual component of the time average after spectral enhancement, usually including removal of harmonics of the gear mesh frequency.

“Enveloped signal” (ETE) and “Envelope spectrum” (FTE) are specific to bearings. They refer to the envelope of the band-pass-filtered vibration data (the magnitude of the “analytic signal”) and the power average of the envelope spectrum respectively.

Table 6.2: CI vibration data

CI Type	Original signal	Enhanced signal	Enveloped signal	Envelope spectrum
PP	SIG_PP	ESA_PP	ETE_PP	
PK	SIG_PK			
STD	SIG_STD	ESA_STD	ETE_STD	
MN	SIG_MN		ETE_MN	
MRC	SIG_MRC			
M6		ESA_M6	ETE_M6	
IMP		ESA_IMP		
EB				FTE_EB

A1.10 Chinook HUMS CI IHM Results Charts

This section presents some of the results of running the Chinook HUMS CI intelligent management algorithms on the CB and EB data. Results are presented for each aircraft and component combination shown in Table 3.6, in section 3.8, i.e. all those with a confidence of greater than 50% and those that match cases described in section 3.2.3, or entries in the MOD Chinook HUMS query register. For each aircraft and component combination presented: a description is given of the anomalies in the data and, where applicable, how they relate to a confirmed fault case or investigation; the overall confidence and number of high, medium and low warnings are given; charts for each of the four CIs analysed, for two passbands. The charts show the CI values, colour coded if they have generated high, medium or low warnings or fall in an abnormal cluster, and show the fleet and group (aircraft) thresholds as they develop over time. Some charts also show the short and long term gradients for a selected CI value. The section and figure labels follow the format: Aircraft identifier (e.g. AC01) and component identifier (e.g. LH-CB123 indicates the left hand combiner bearing group comprising bearings 1, 2 and 3). Further details of the naming convention can be found in section 3.8.

AC10 LH-CB123

All the energy CIs were found to be above the fleet threshold for all available data. The values are highest for February to April 2005. At the end of this period it is likely that some maintenance was performed, as there is a gap in the data between the 15th and 26th of April. The M6 values were only found to have a few spikes slightly above the fleet threshold causing a few low level warnings.

Confidence: 99.0 Number of Alerts: High: 392 Medium: 622 Low: 506

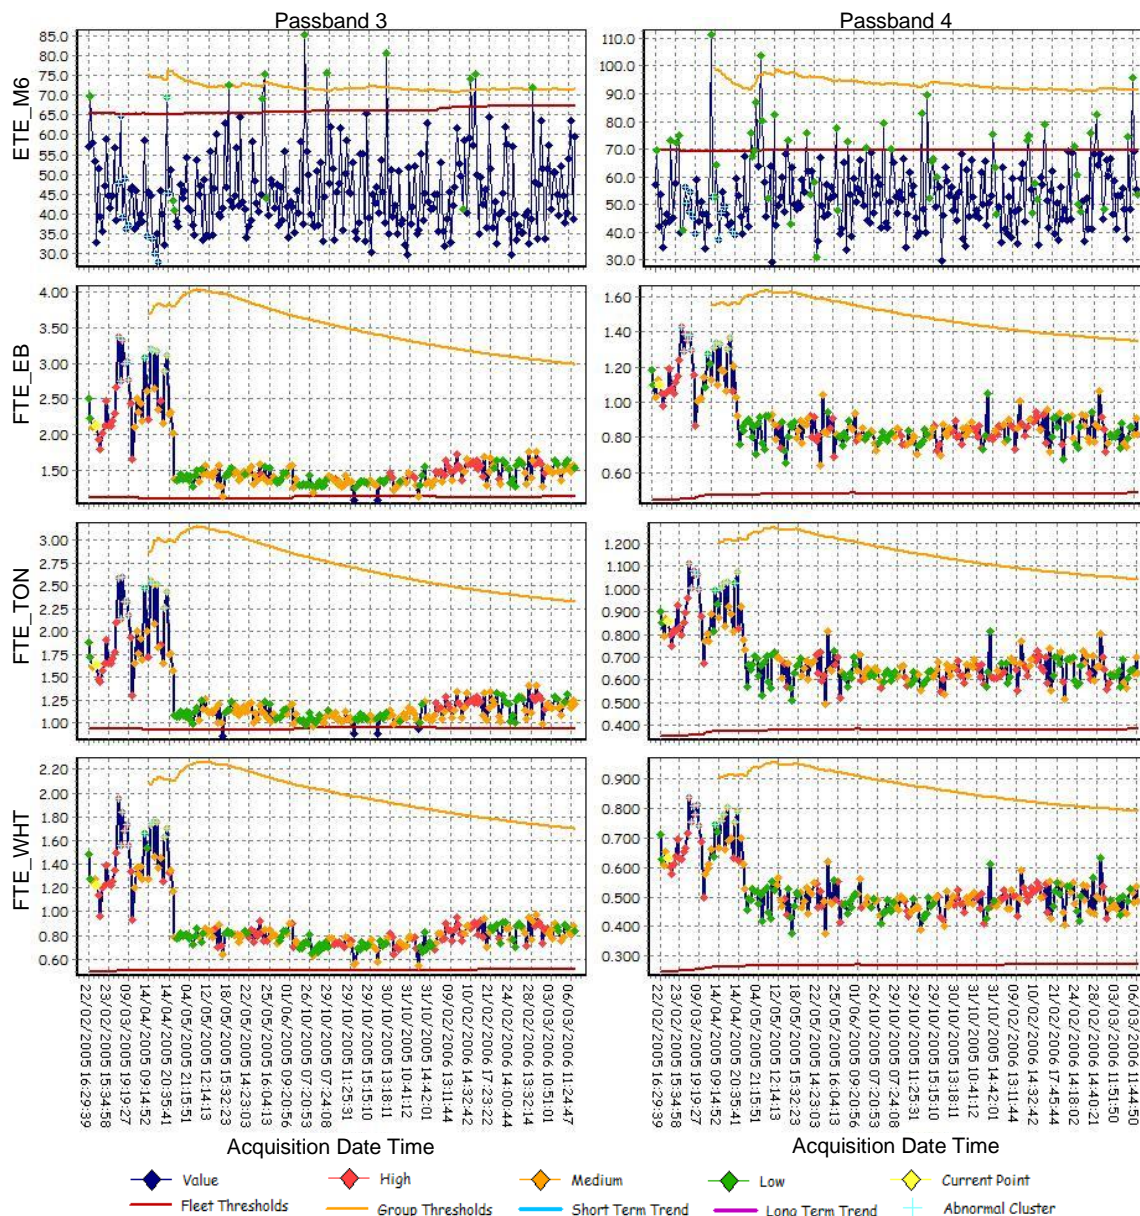


Figure 6.5: AC10 LH-CB123

AC14 RH-CB1

All the energy CIs were found to be above the fleet threshold from the start of the data in March 2003 until the last but one point on the 29th of January 2004. The last data point was normal (on the 6th of February 2004). In April 2003 there were found to be a large number of very high spikes in the data. These were investigated by the MOD, as recorded in the query register, in July and again in August 2004. The MOD conclusions were: “Analysis confirms impulsive content but cannot be attributed to any bearing frequencies in the Combiner Xmsn. Consider this still to be spurious arisings”. The data for CB4567 had similar values.

Confidence: **99.0** Number of Alerts: High: **395** Medium: **390** Low: **491**
 Passband 3 Passband 4

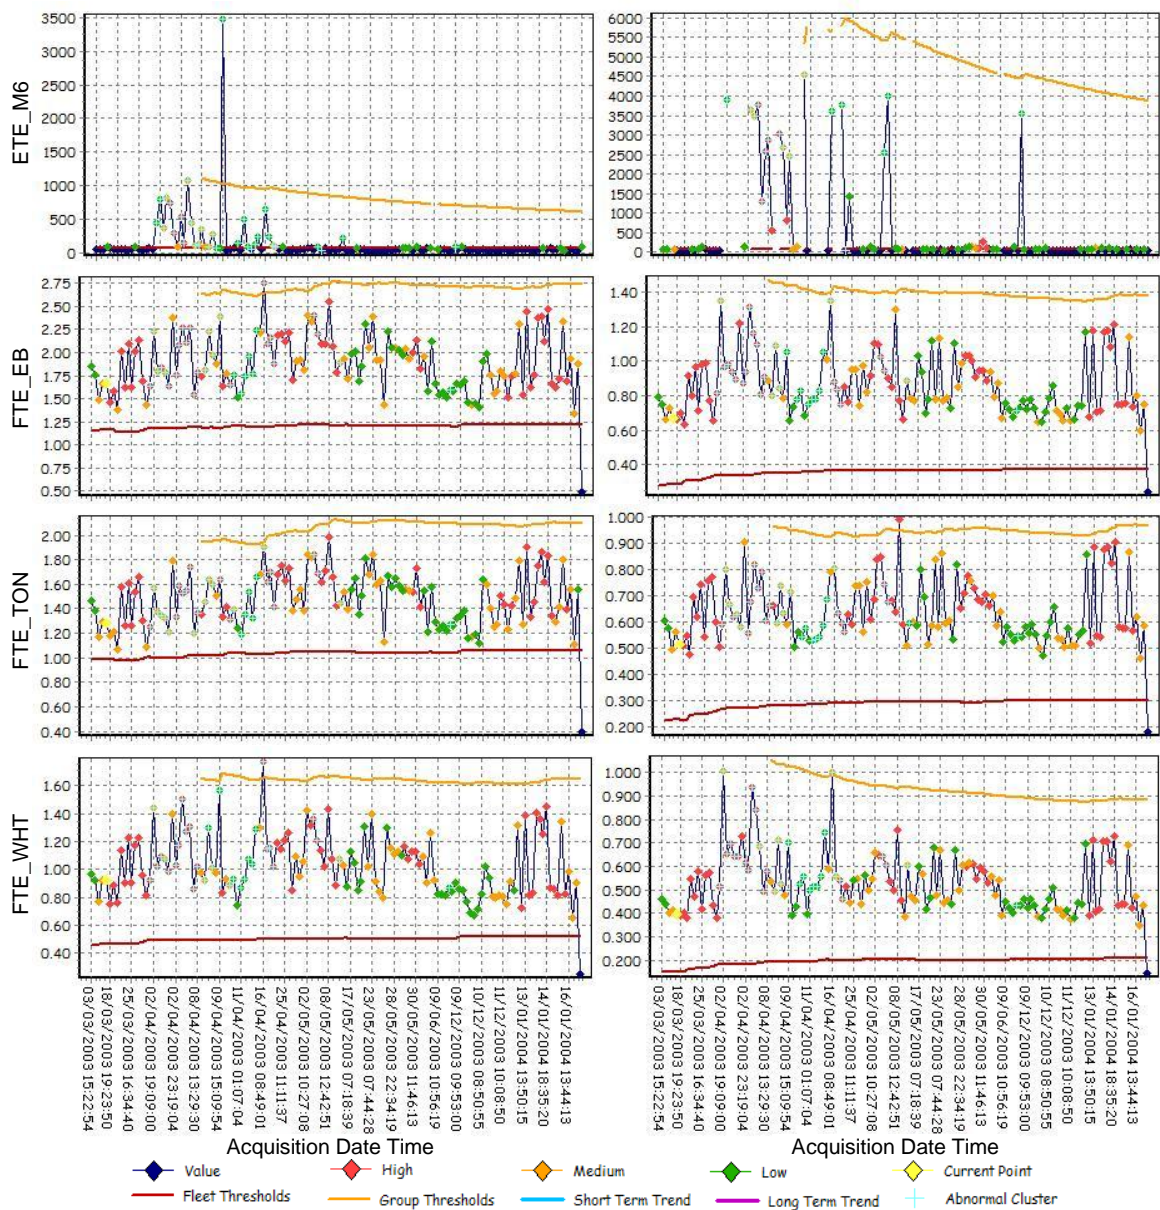


Figure 6.6: AC14 RH-CB1

AC17 LH-CB1

All the energy CIs were found to be above the fleet threshold from the start of the data on the 2nd of December 2002 until the 23rd of January 2003. A step down in energy CI values for passband 3 between the 6th and 11th of December 2002, was also noted. The M6 values were found to ramp up from about the 16th of January 2003. This data relates to fault case 1, presented in section 3.2.3.1. The combiner bearings were replaced and normal operation was observed in the data from May onwards. The data for CB4 is also presented on the next page, and has similar characteristics.

Confidence: **99.0** Number of Alerts: High: **192** Medium: **202** Low: **135**
 Passband 3 Passband 4

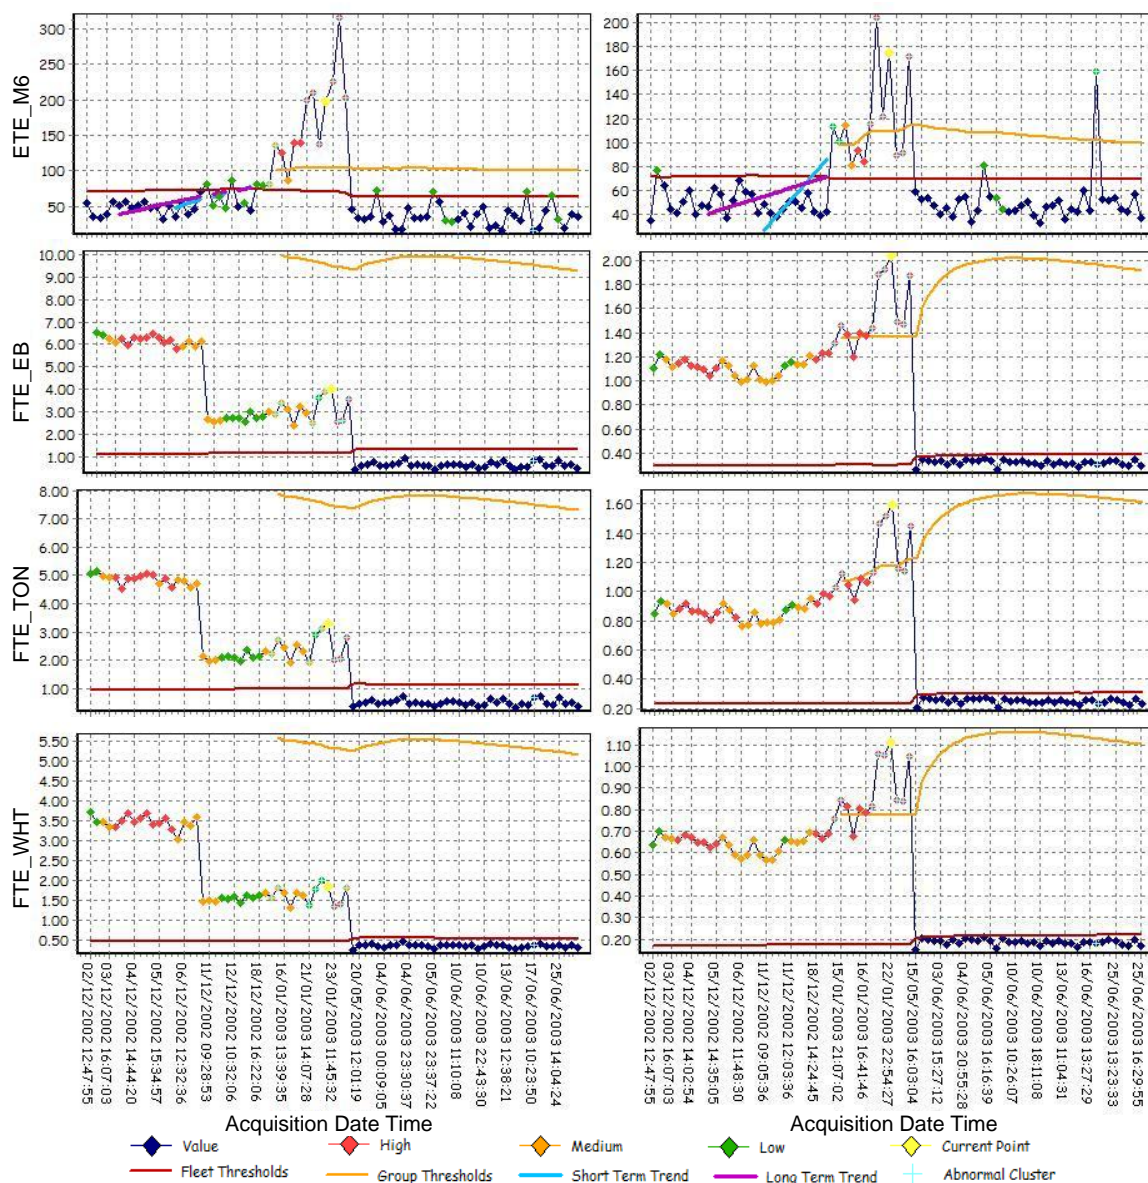


Figure 6.7: AC17 LH-CB1

Confidence: **99.0** Number of Alerts: High: **157** Medium: **224** Low: **131**

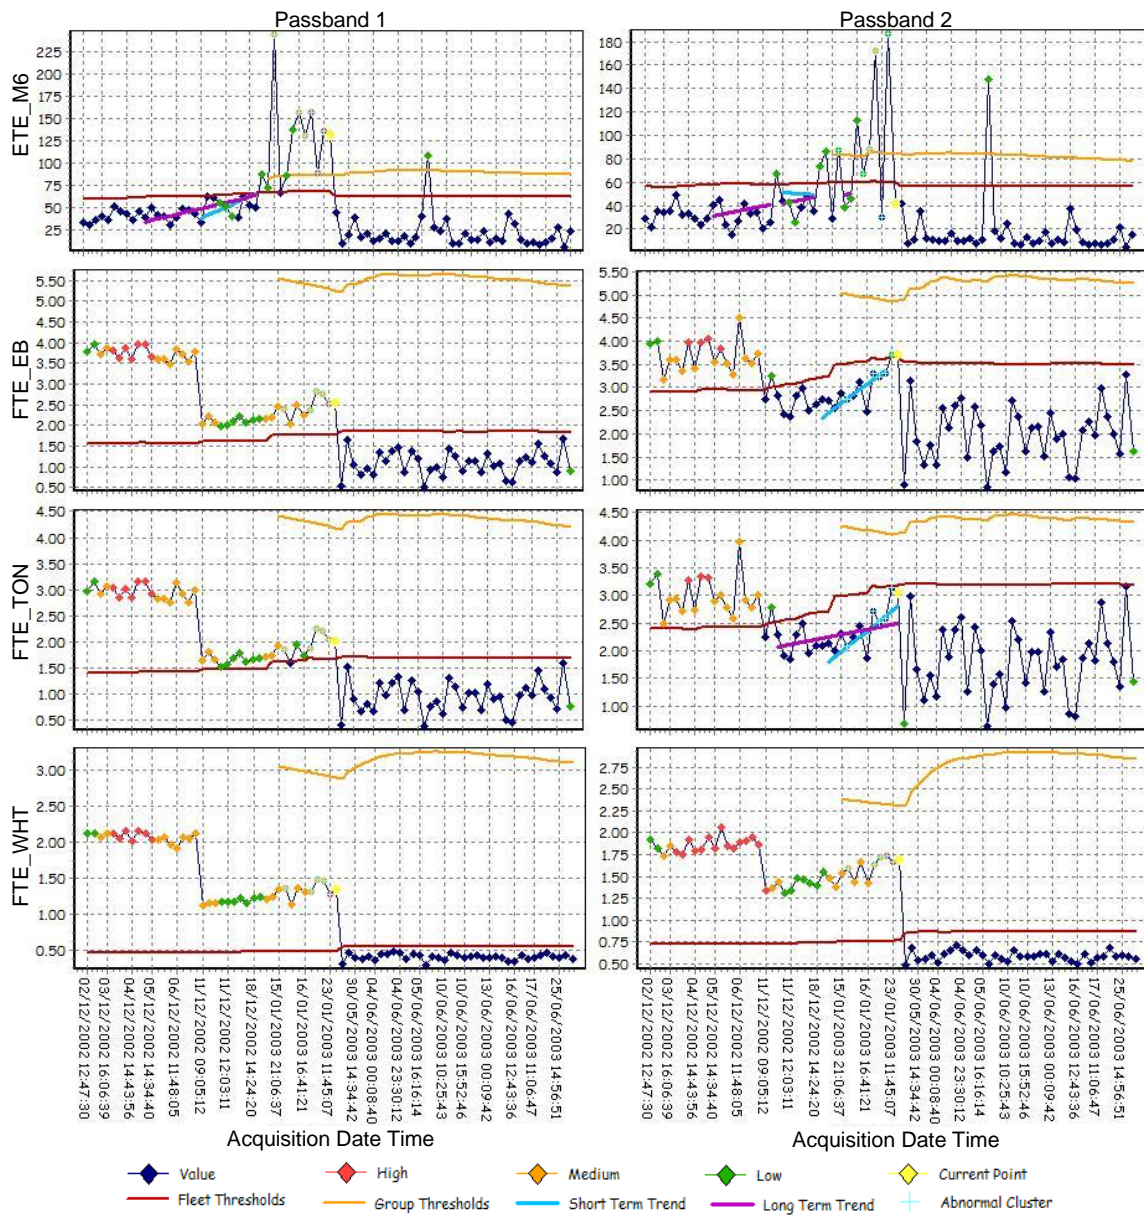


Figure 6.8: AC17 LH-CB4

AC6 RH-CB1

All the energy CIs were found to be above the fleet threshold from the start of the data on the 23rd of March 2003 until the end of data on the 8th of July 2003. The M6 values were found to have a spike on 23rd of April 2003 and then remained high from the 26th of June onwards.

Confidence: **99.0** Number of Alerts: High: **164** Medium: **74** Low: **74**

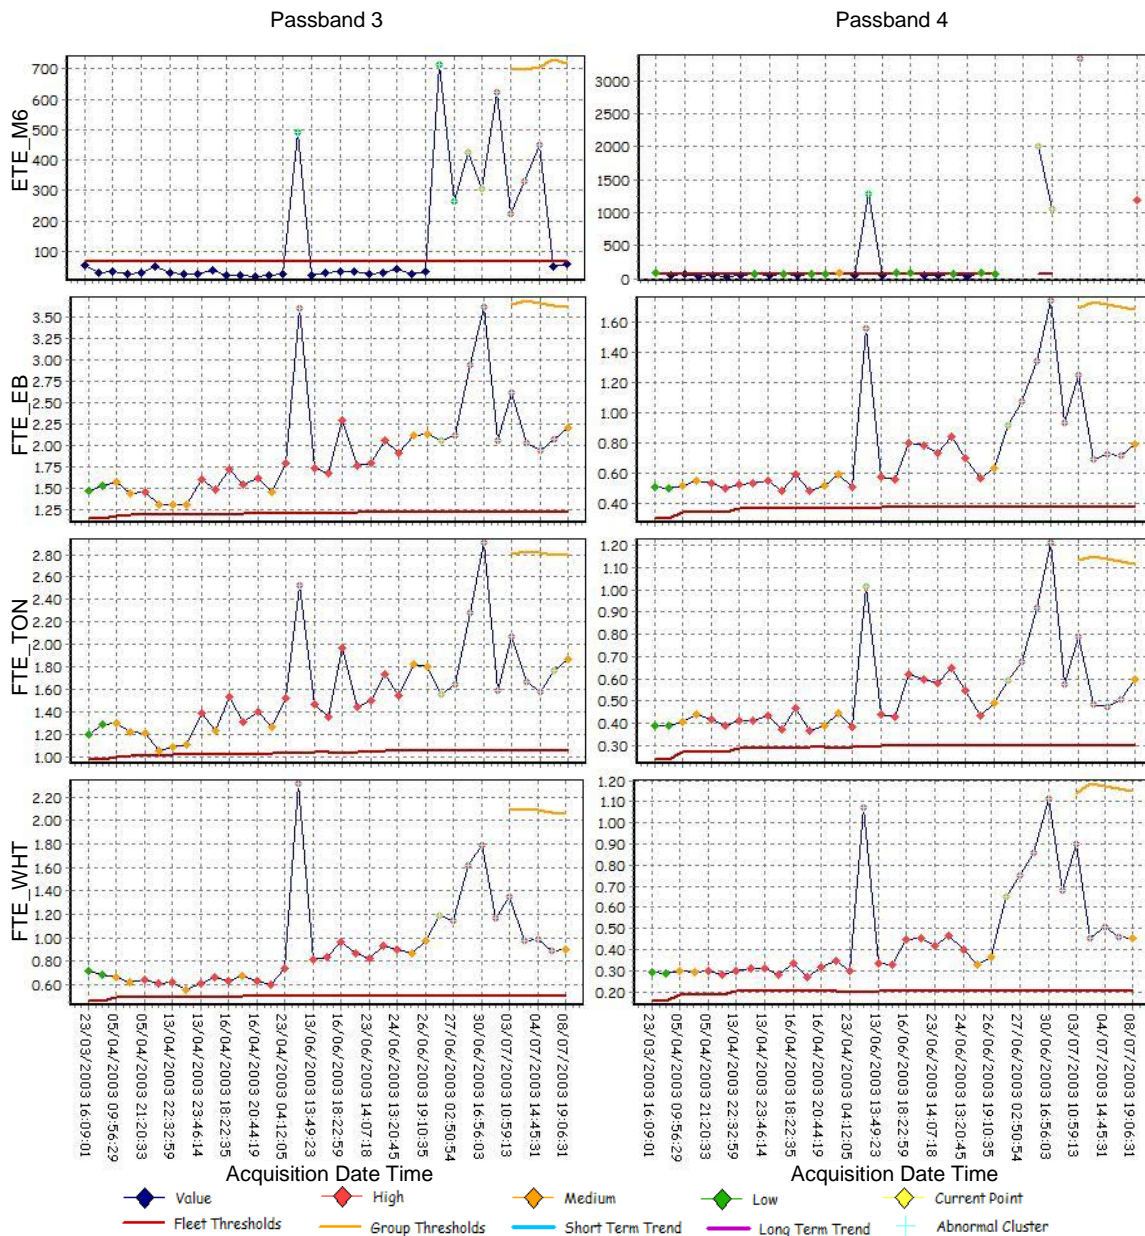


Figure 6.9: AC6 RH-CB1

AC38 RH-CB123

All the energy CIs, for passband 4, were found to be above the aircraft threshold, from the 12th of April 2005 onwards. It is likely that maintenance was carried out before this date, since the previous acquisition was on the 26th of August 2004. For passband 3, the energy levels were found to return to normal on the 18th of April 2005: note that the next acquisition was not until the 6th of May, so it is likely that maintenance was performed in this period. However, the energy levels were found to be well below the fleet limit and there were no significant M6 alerts.

Confidence: 94.8 Number of Alerts: High: 327 Medium: 198 Low: 36

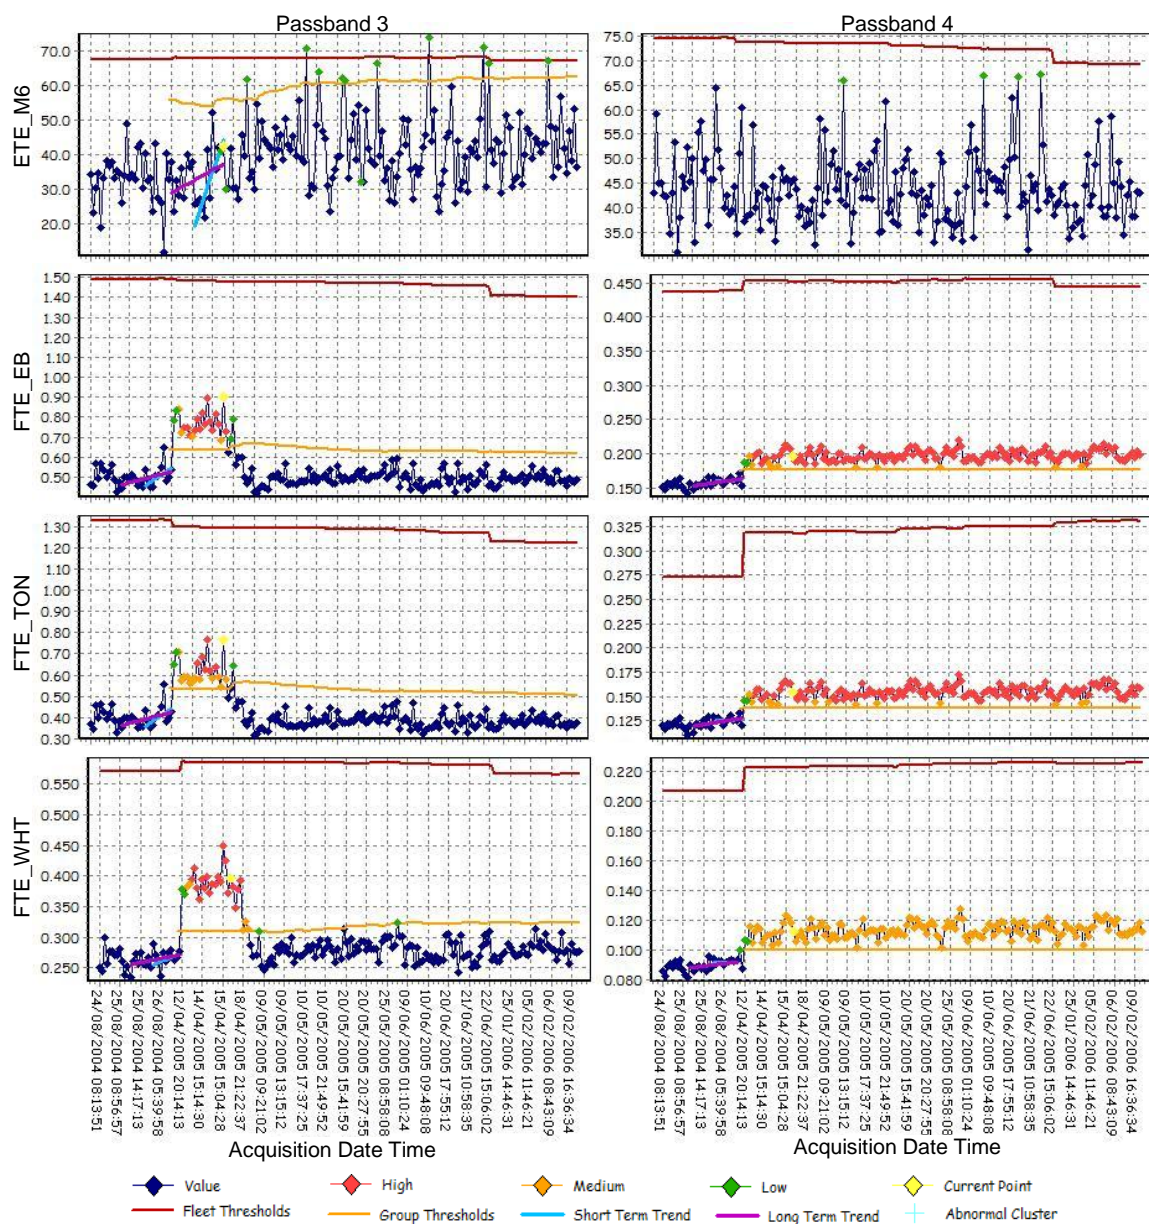


Figure 6.10: AC38 RH-CB123

AC31 RH-EB123

The warnings generated for this component do not appear to indicate a real fault. All the CIs were found to be constant and at a low level, but the fleet threshold was initially very low. Further investigations revealed that there were about 380 acquisitions from AC15, AC23 and AC3 in 2003, which had much lower values of mean and standard deviations than later data, causing the fleet threshold to be too low. If this analysis was run regularly, this would be spotted early on and the fleet statistics could be reset to start at the end of January 2004. The data is very similar for EB4567.

Confidence: **71.8** Number of Alerts: High: **106** Medium: **191** Low: **219**

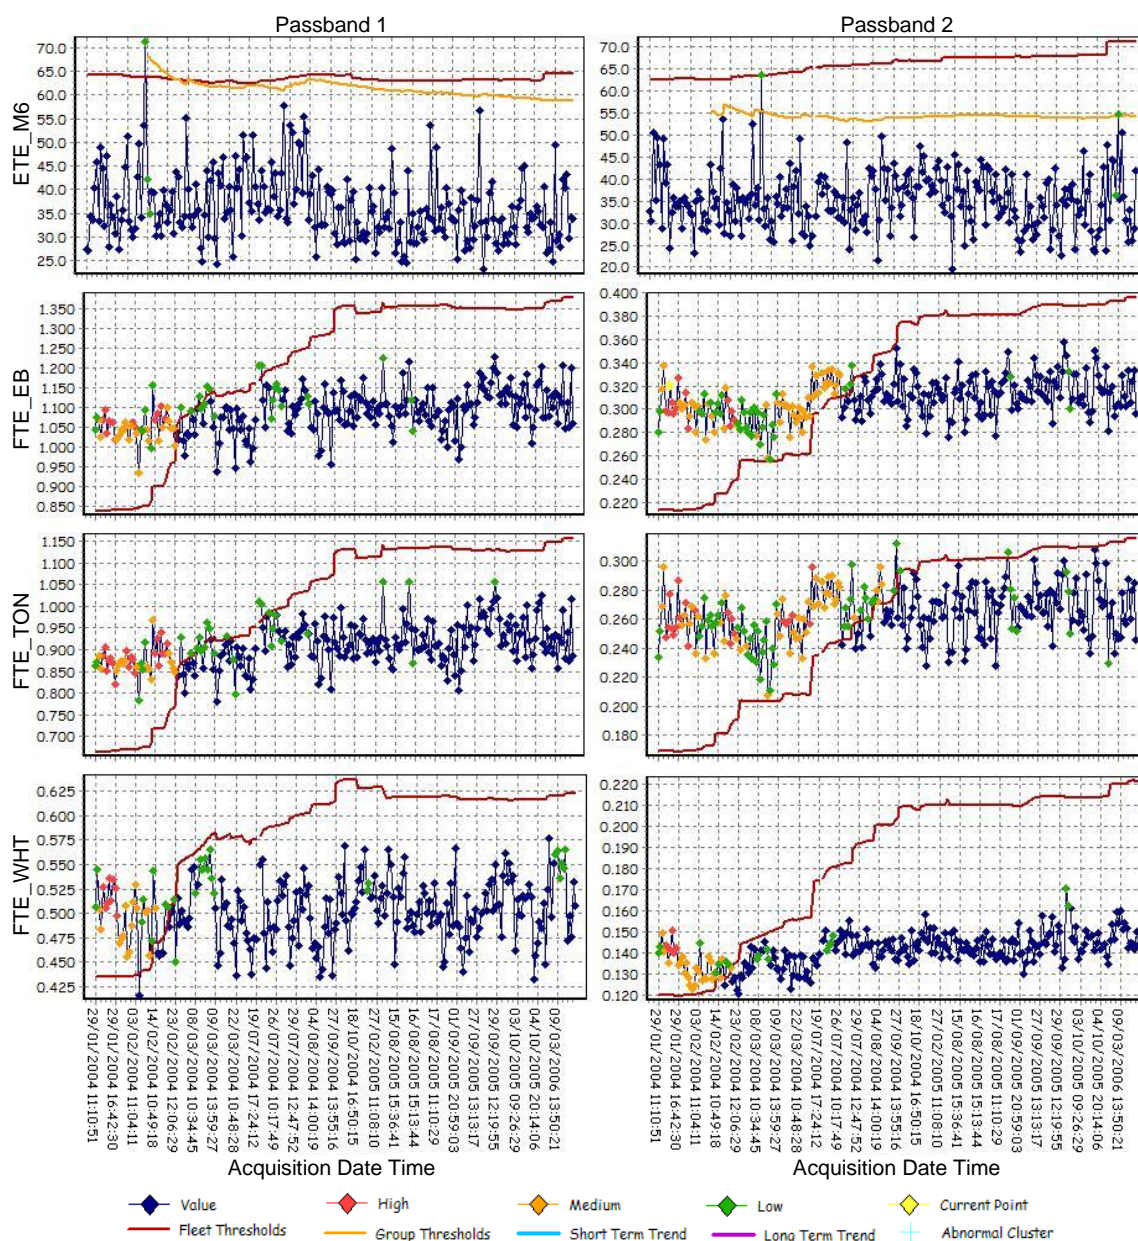


Figure 6.11: AC31 RH-EB123

AC17 RH-EB123

There appear to be two issues identified in this data. Initially, there were found to be M6 passband 1 and 2 alerts from start of data, from the 24th of January until the 24th of March 2004. There was then a big gap in the data, during which this fault seems to have been cleared. However, the energy values were then found to step up on the 23rd of January until the end of the data on the 22nd of March 2006. Although all the energy indicators for both passbands have this step, only the passband 1 data was found to have warnings generated and the values were only just above the fleet threshold. The data is very similar for EB4567.

Confidence: **70.8** Number of Alerts: High: **23** Medium: **87** Low: **261**

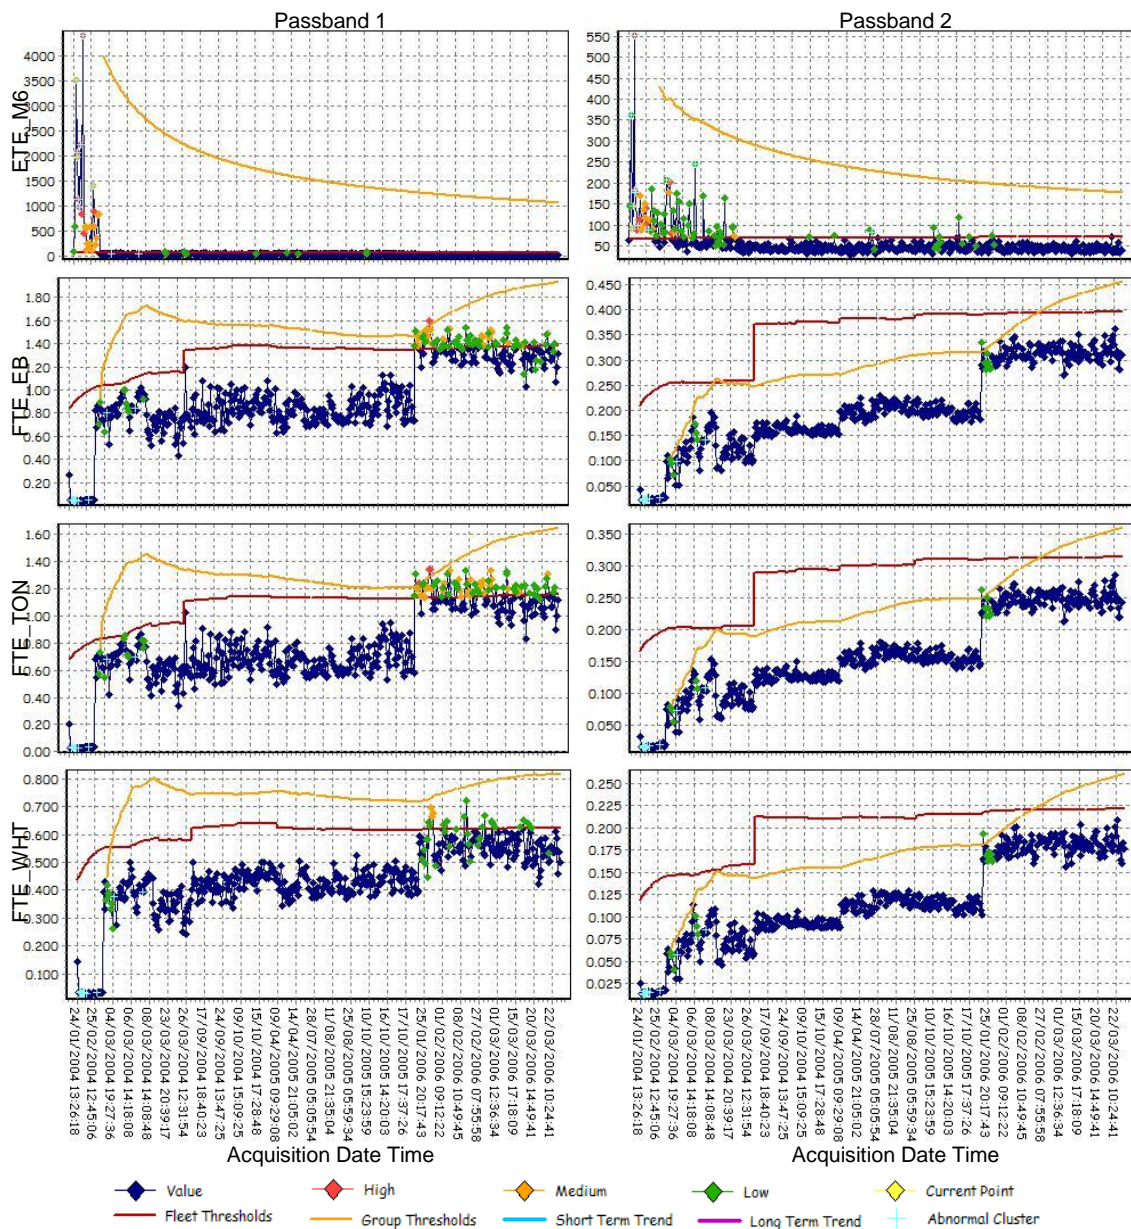


Figure 6.12: AC17 RH-EB123

AC2 LH-CB123

Initial low energy values for this aircraft have caused a low aircraft threshold. The values were found to be below the fleet thresholds for passband 4. A clear shift in the data was identified on the 17th of September 2004, so something has changed which may be of interest to the operator. There were no significant M6 warnings.

Confidence: 70.3 Number of Alerts: High: 524 Medium: 386 Low: 244

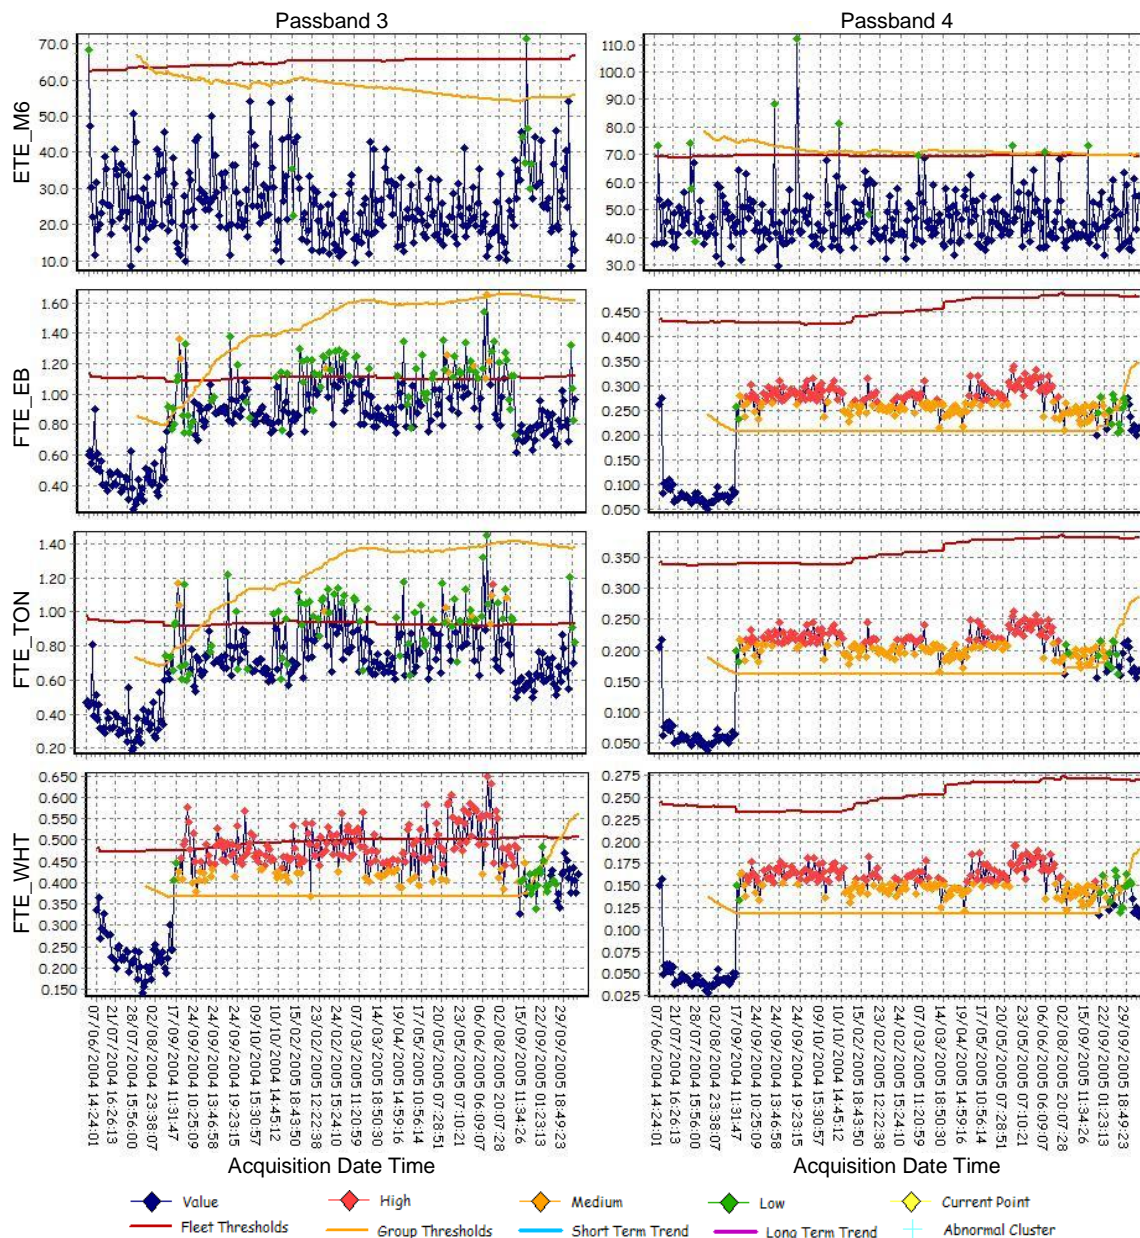


Figure 6.13: AC2 LH-CB123

AC29 RH-EB4567

Passband 2 energy CIs were found to have a clear shift on the 1st of June 2005 until the end of data on the 6th of March 2006. There were also some high M6 values in this period.

Confidence: **68.6** Number of Alerts: High: **145** Medium: **10** Low: **47**

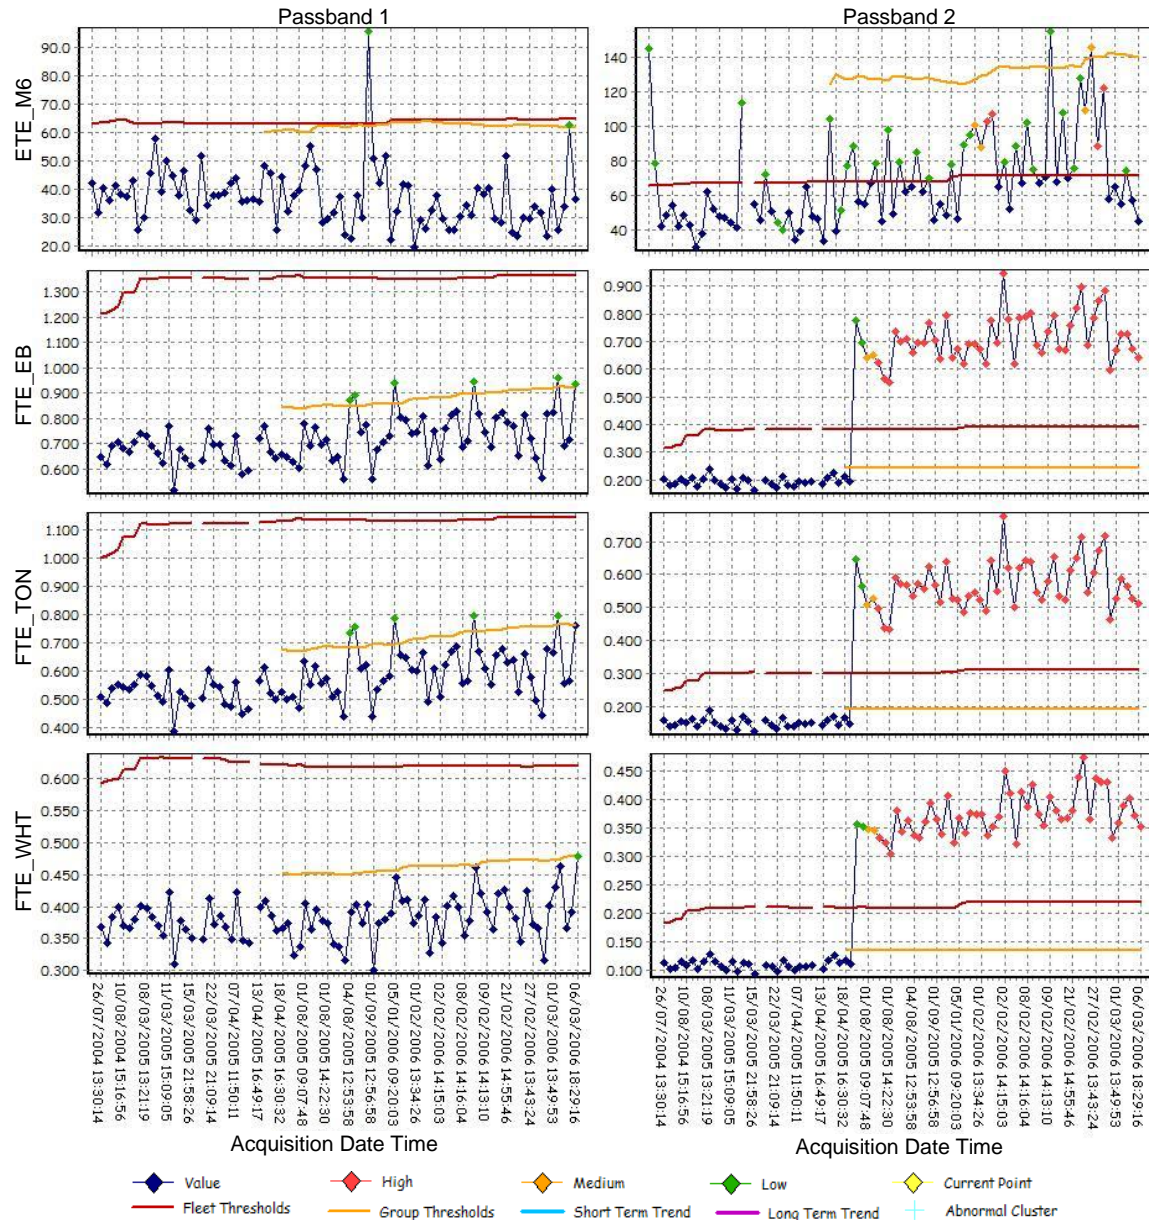


Figure 6.14: AC29 RH-EB4567

AC6 RH-CB123

M6 values were found to be very high for between the 3rd and 22nd of June 2004. This period relates to fault case 9 presented in section 3.2.3.6. The energy values were also fairly high, from the start of data on the 17th of January until the 22nd of June 2004.

Confidence: **67.7** Number of Alerts: High: **91** Medium: **101** Low: **257**




Figure 6.15: AC6 RH-CB123

AC33 LH-EB1

All of the energy CI values were found to be above the fleet thresholds for all available data. This is also the case for the EB4 group.

Confidence: **63.5** Number of Alerts: High: **99** Medium: **93** Low: **78**




Figure 6.16: AC33 LH-EB1

AC7 LH-EB123

The EB and TON CIs for passband 1 were found to be very high, from the start of the data on the 22nd of April until the 22nd of August 2005. There was then a gap in the data until the 29th of November, when the values returned to normal. The M6 levels were found to be very low for the first period, then higher, but still normal for the second, resulting in the aircraft threshold being too low.

Confidence: 62.3 Number of Alerts: High: 86 Medium: 207 Low: 121

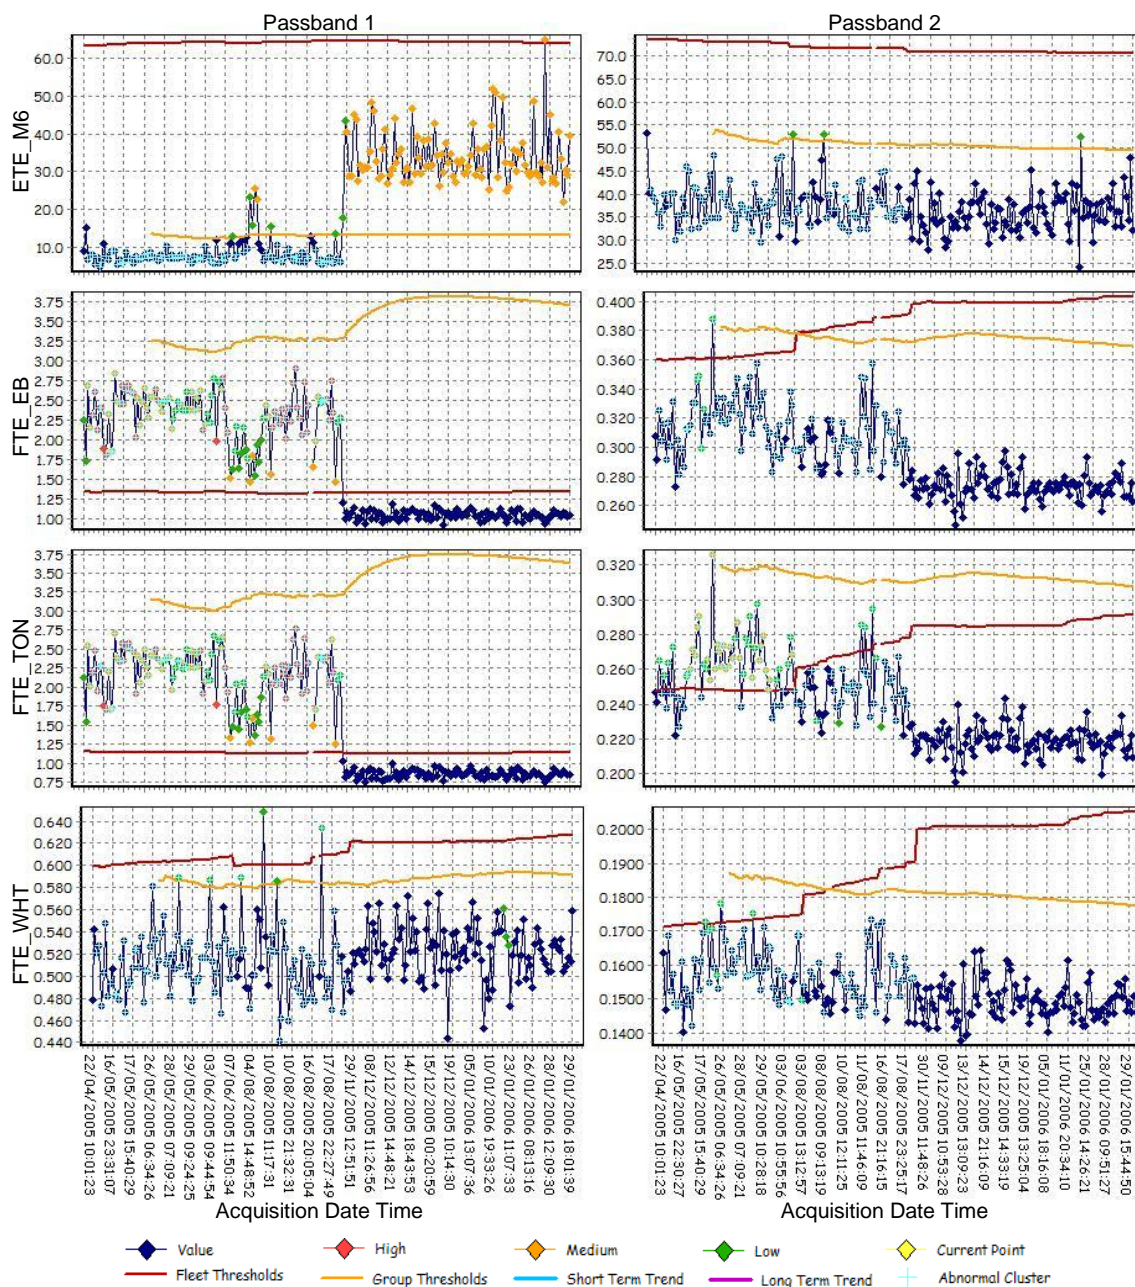


Figure 6.17: AC7 LH-EB123

AC10 LH-EB123

The M6 values were found to be very high between the 8th and 24th of February 2006. The previous good acquisition was on the 31st of October 2005, so there is a big gap, indicating major maintenance may have occurred. The energy levels actually fall for this period. These conditions may indicate an instrumentation fault. The values are very similar for EB4567.

Confidence: **57.7** Number of Alerts: High: **36** Medium: **38** Low: **15**

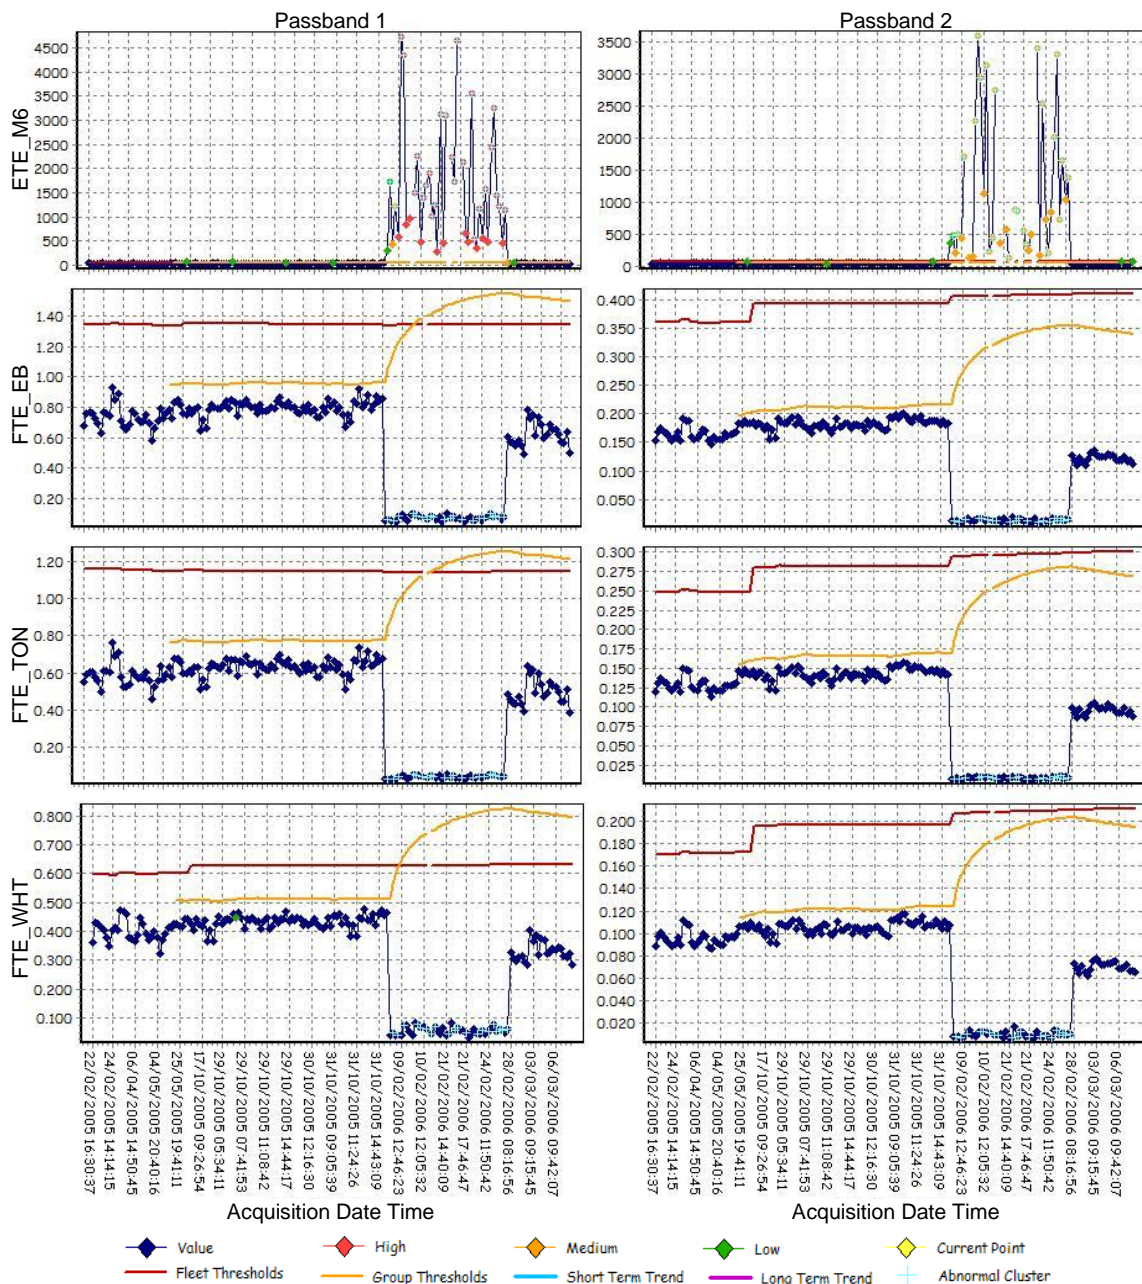


Figure 6.18: AC10 LH-EB123

AC32 LH-CB123

The CIs were found to have very low initial energy values causing a low aircraft threshold resulting in some alerts. This is not considered to represent a real fault case and may indicate that a longer period is required before group thresholds are used.

Confidence: **55.4** Number of Alerts: High: **38** Medium: **63** Low: **198**

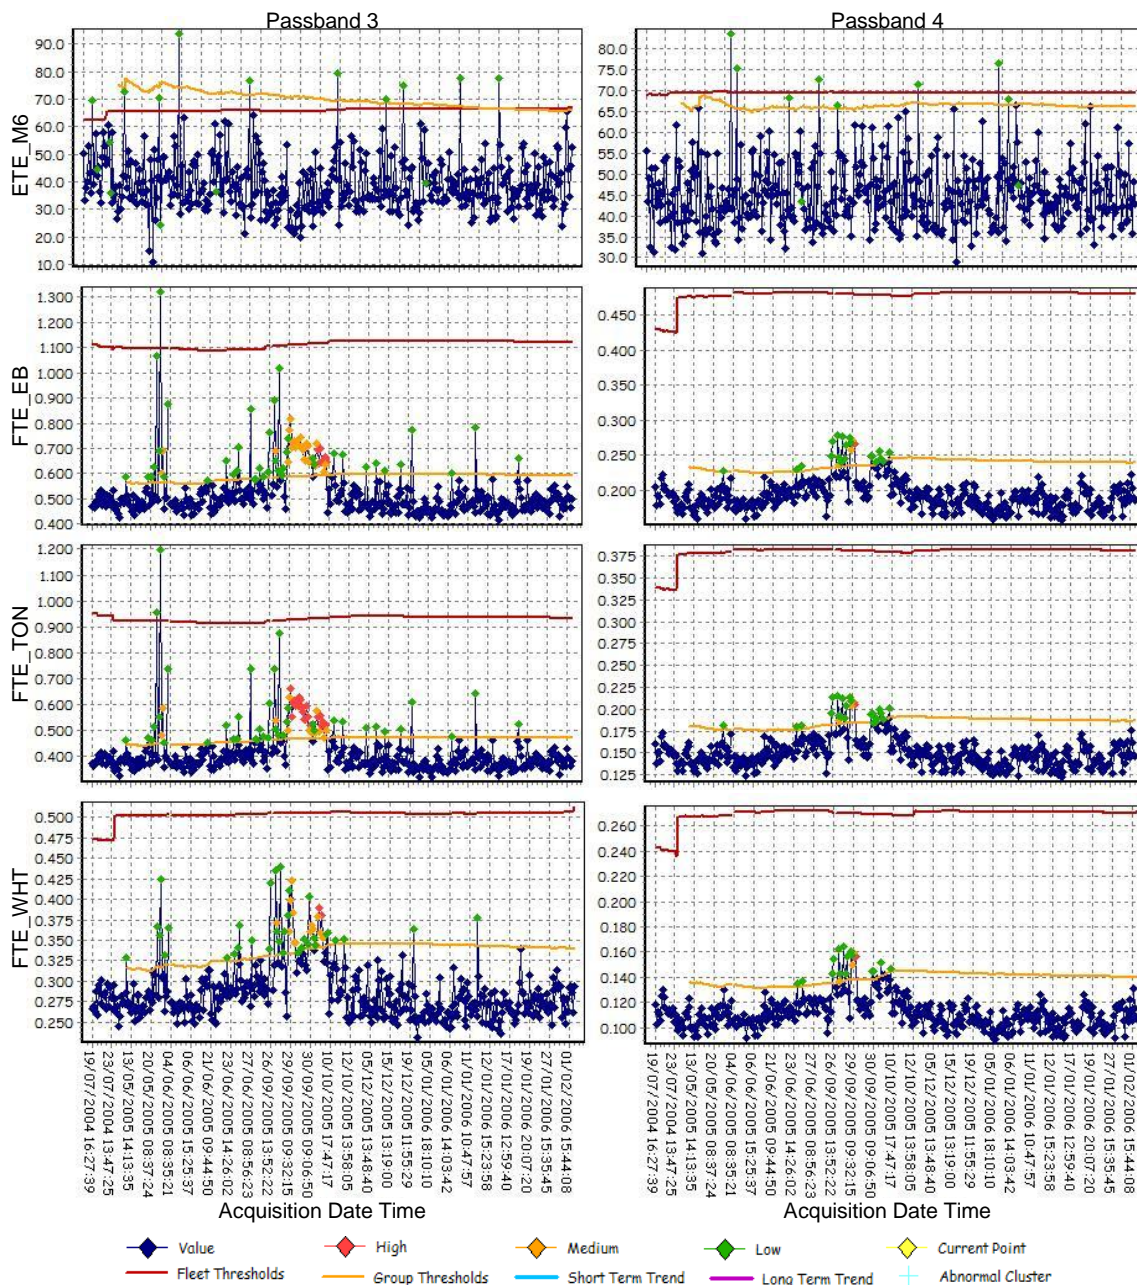


Figure 6.19: AC32 LH-CB123

AC1 RH-EB123

The initial warnings seem to be caused by the fleet threshold being too low. However, the energy CIs for passband 1 continue to be high relative to the fleet until the 25th of November 2005. The levels were found to be only just above the fleet threshold, but there is then a clear shift indicating that maintenance has corrected the issue.

Confidence: **53.6** Number of Alerts: High: **33** Medium: **129** Low: **188**

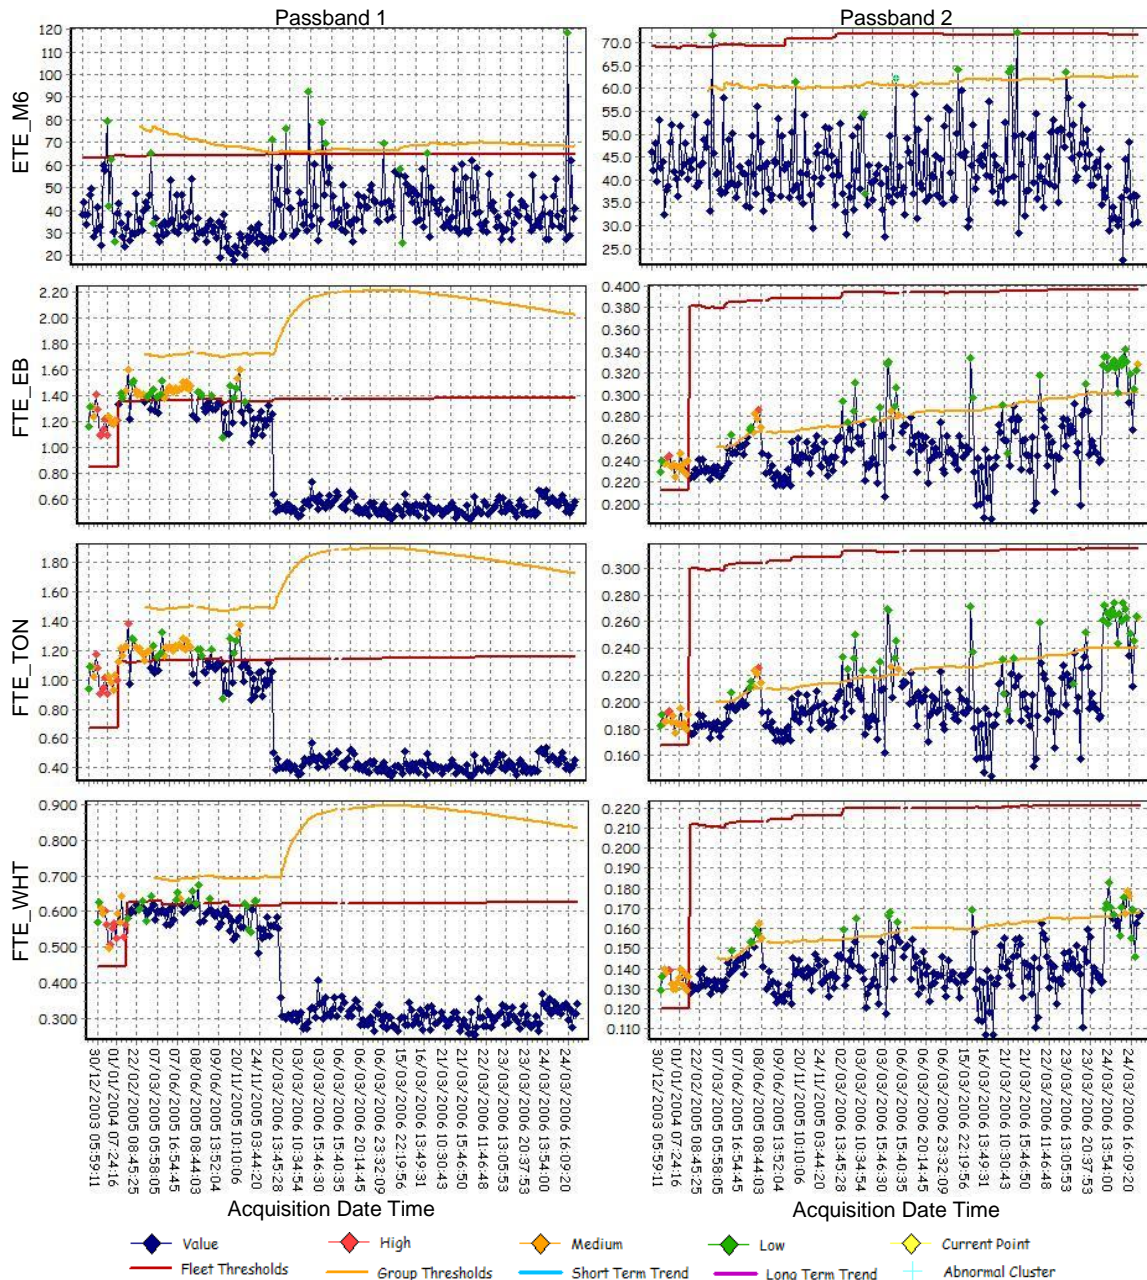


Figure 6.20: AC1 RH-EB123

AC14 RH-CB123

Both the energy and M6 CI values were found to be high for passband 4, from the start of the data on the 19th of July 2004, until the 17th of September 2004. The next data acquisition is on the 1st of April 2005 and the data appears to be normal. This case was investigated by the MOD, on the 17th of September 2004, and the query register conclusions were: “No matches for CB1-3 defect freq.'s, however, freq.'s at 2760hz and 3570hz (oil cooler fan?) couldn't really relate the frequencies to anything.”

Confidence: **52.5** Number of Alerts: High: **87** Medium: **166** Low: **265**

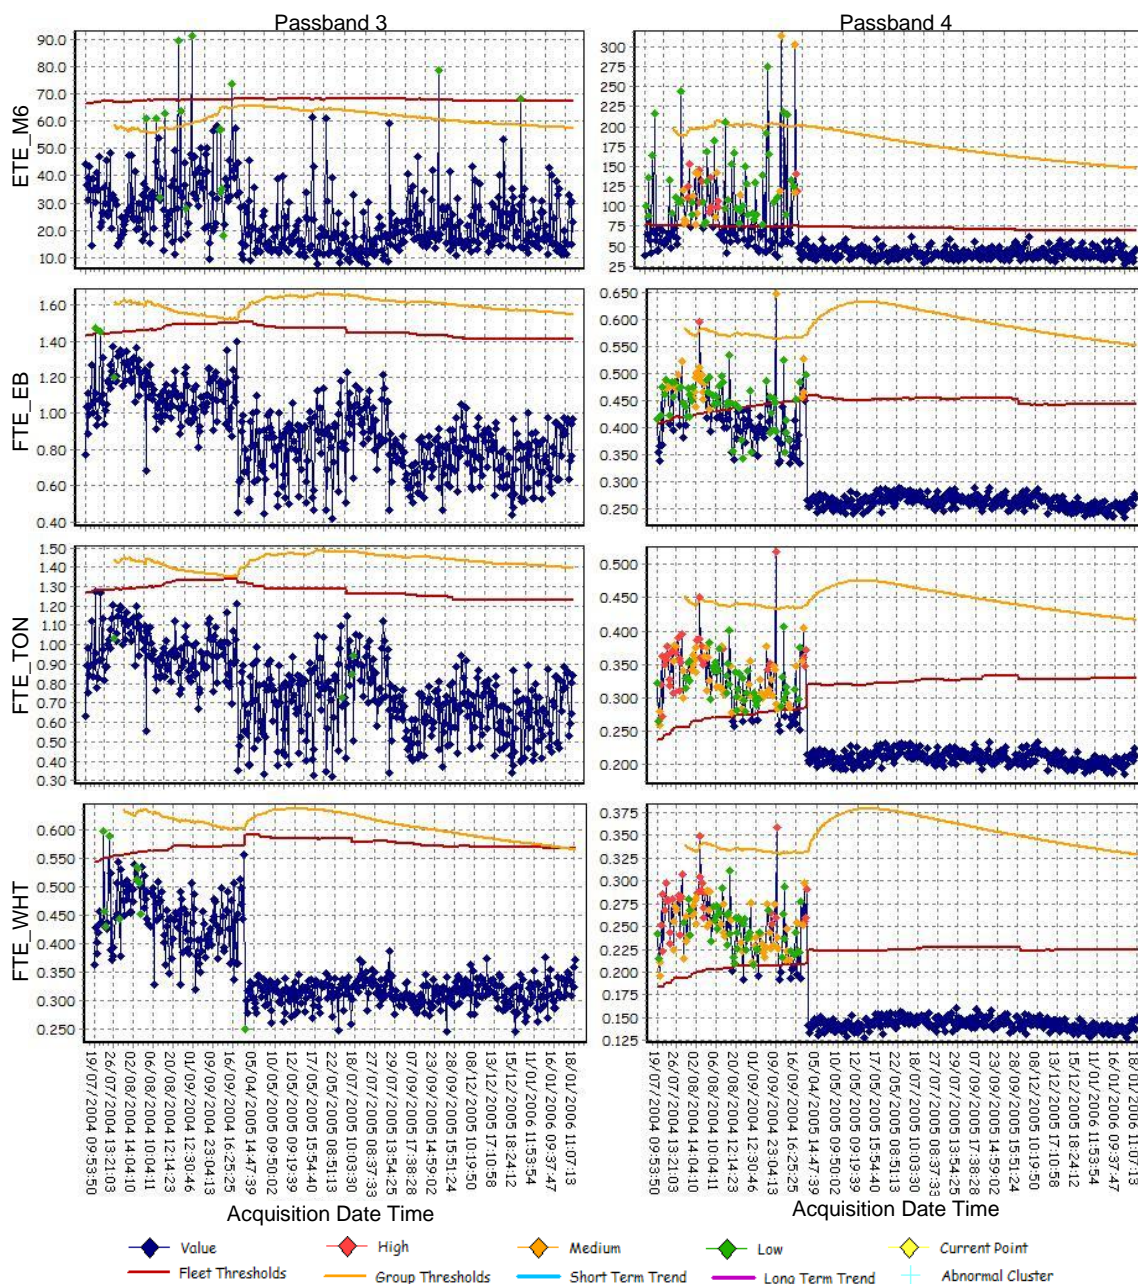


Figure 6.21: AC14 RH-CB123

AC33 LH-EB4567

All the energy CIs for passband 2 were found to be high for all data.

Confidence: 41.8 Number of Alerts: High: 373 Medium: 489 Low: 625

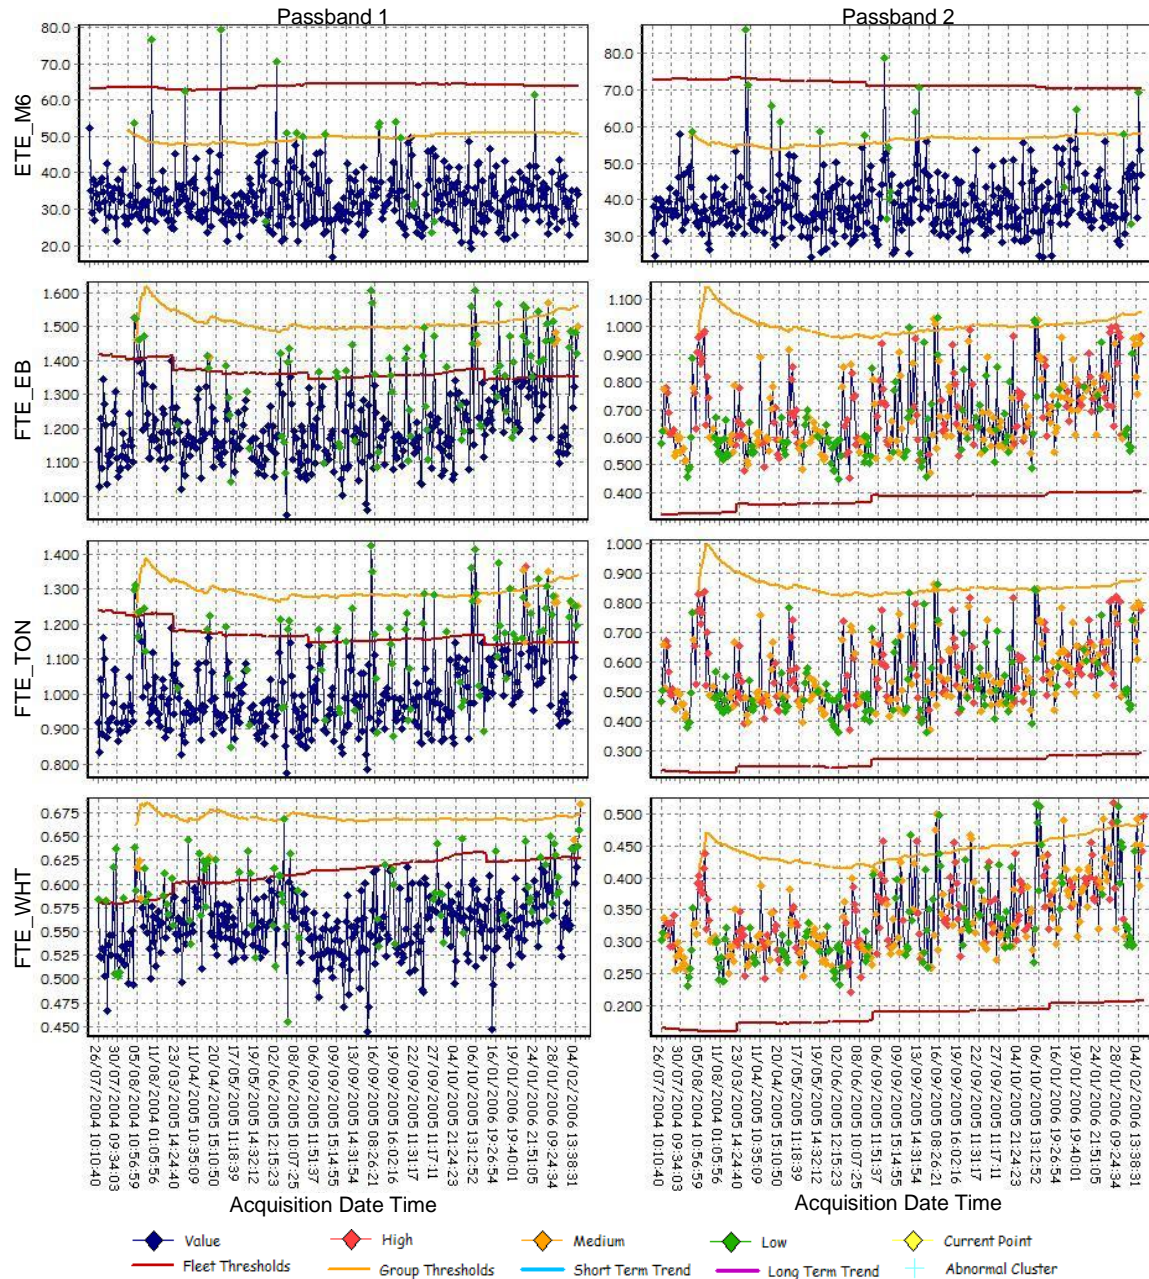


Figure 6.22: AC33 LH-EB4567

AC26 LH-EB1

M6 values for passband 3 were found to be very high for all data sourced data (7th of May to 25th of June 2003). This data relates to case number 3 described in section 3.2.3.3. This gearbox was then fitted to AC22, but the problem persisted. The values were also found to be high for EB4, which is the same fault as that described as case 10 in section 3.2.3.6.

Confidence: 25.7 Number of Alerts: High: 39 Medium: 26 Low: 41

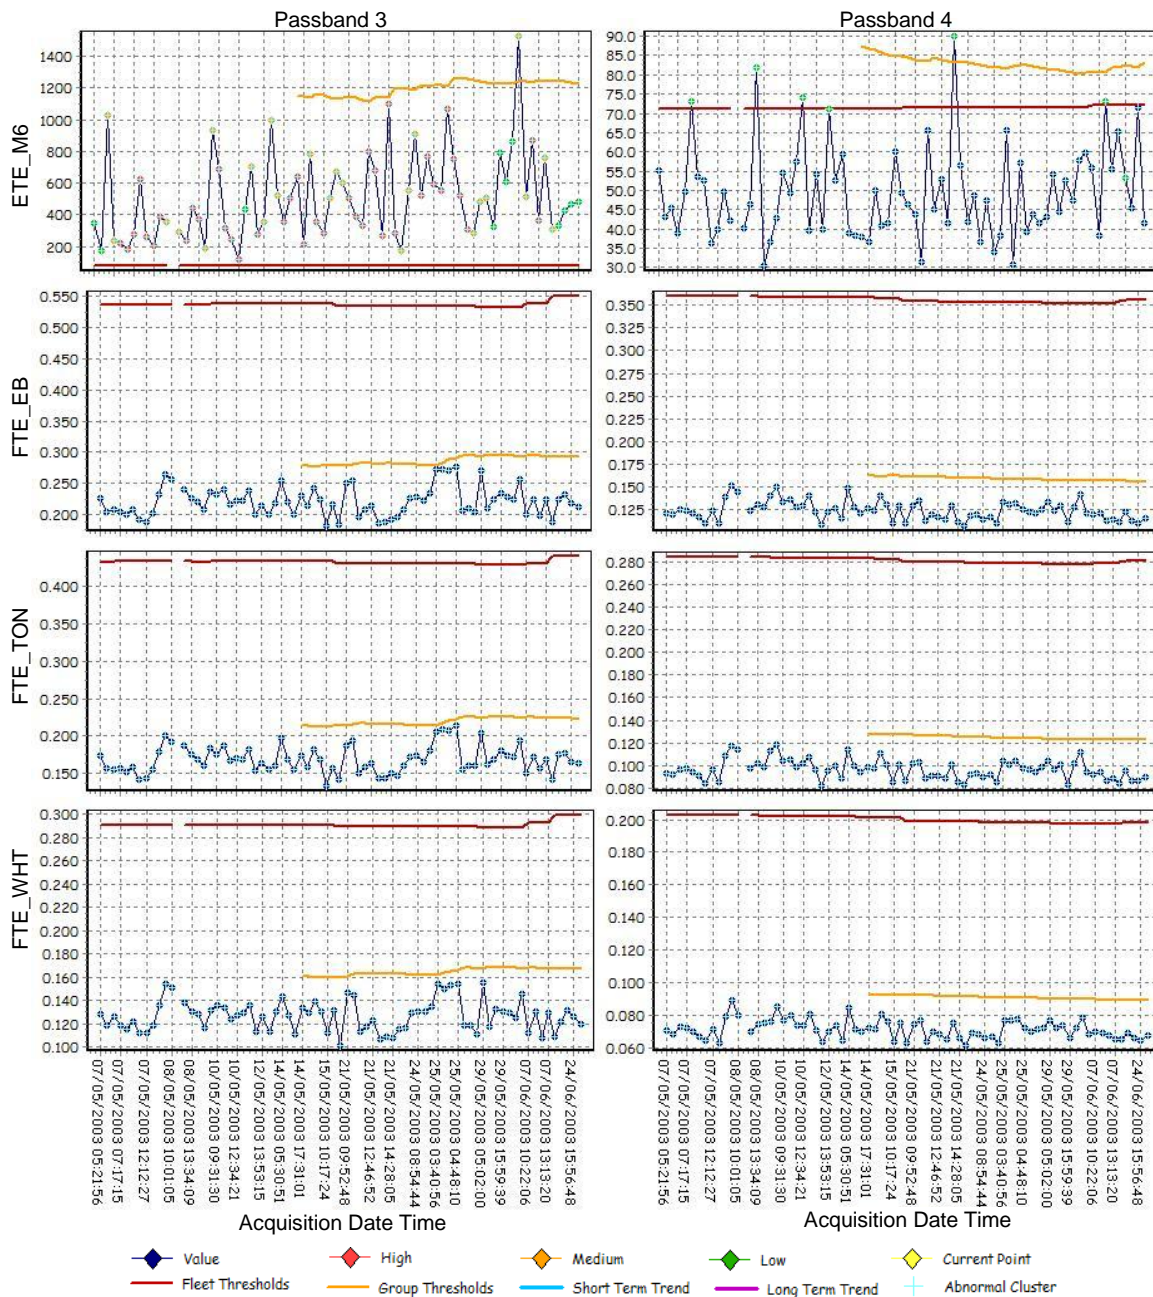


Figure 6.23: AC26 LH-EB4

AC22 LH-EB123

This data also relates to case 3 described in section 3.2.3.3. Initially the M6 values were extremely high for passband 2. The component was then removed on the 26th of July 2004. The M6 values then returned to normal, but the energy CIs became high and were found to be above the fleet threshold.

Confidence: **28.0** Number of Alerts: High: **39** Medium: **17** Low: **15**

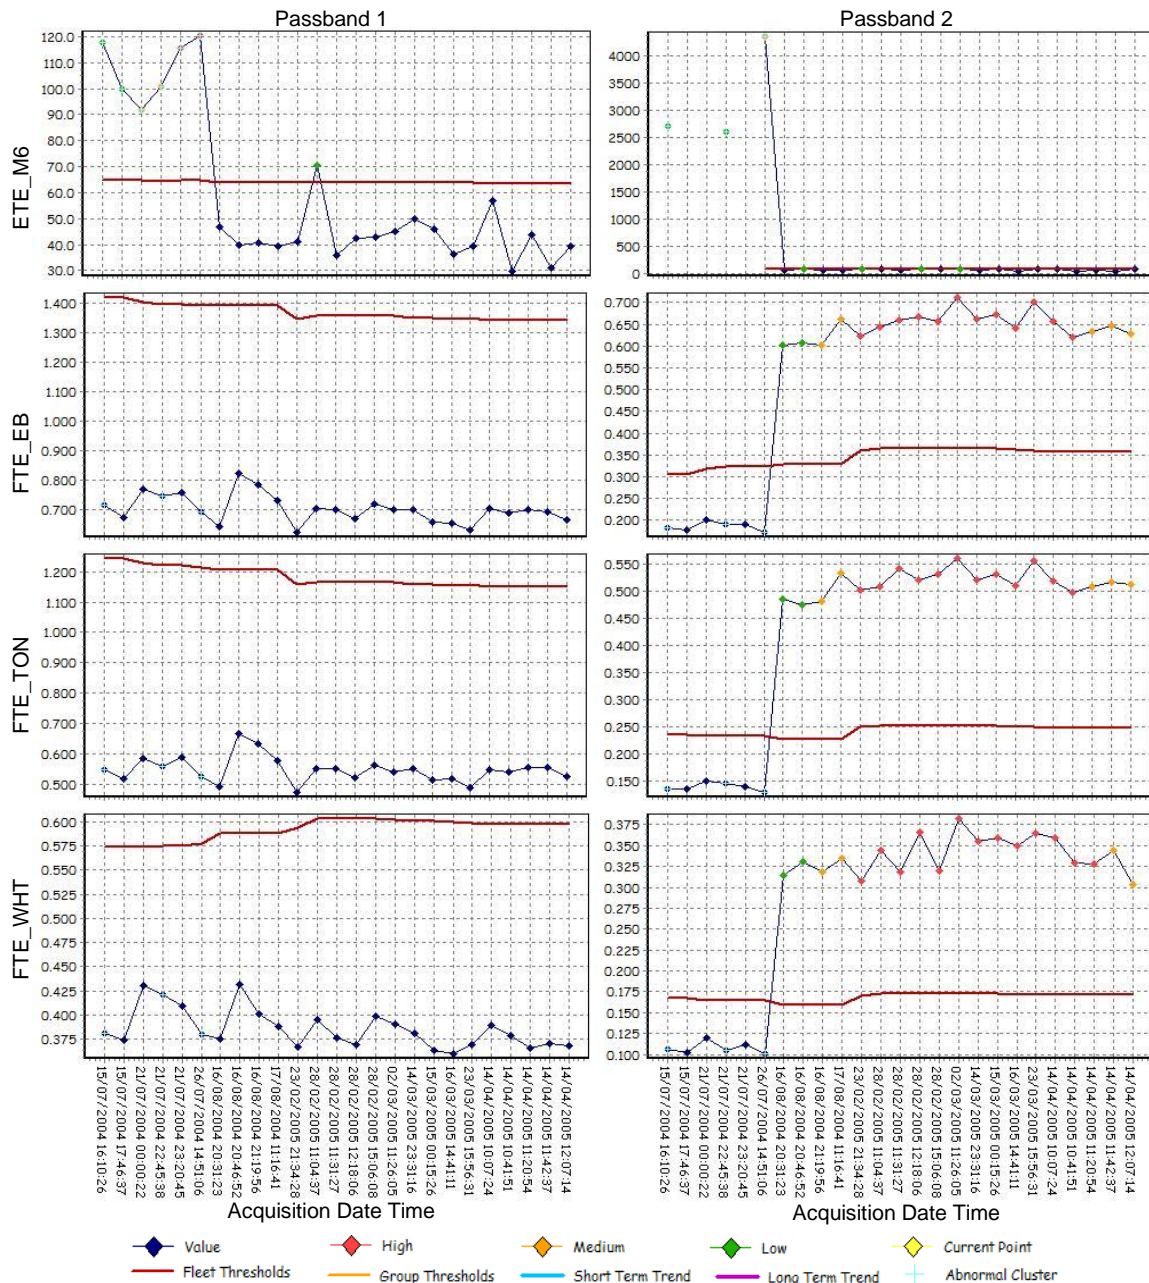


Figure 6.24: AC22 LH-EB123

AC33 LH-CB123

Very high M6 values were recorded for passband 4 from the 22nd of September 2005 onwards. These values were investigated by the MOD, as documented in the query register, at the end of September, and again in January 2006, however, no gear or bearing defects were identified.

Confidence: 25.4 Number of Alerts: High: 3 Medium: 13 Low: 152

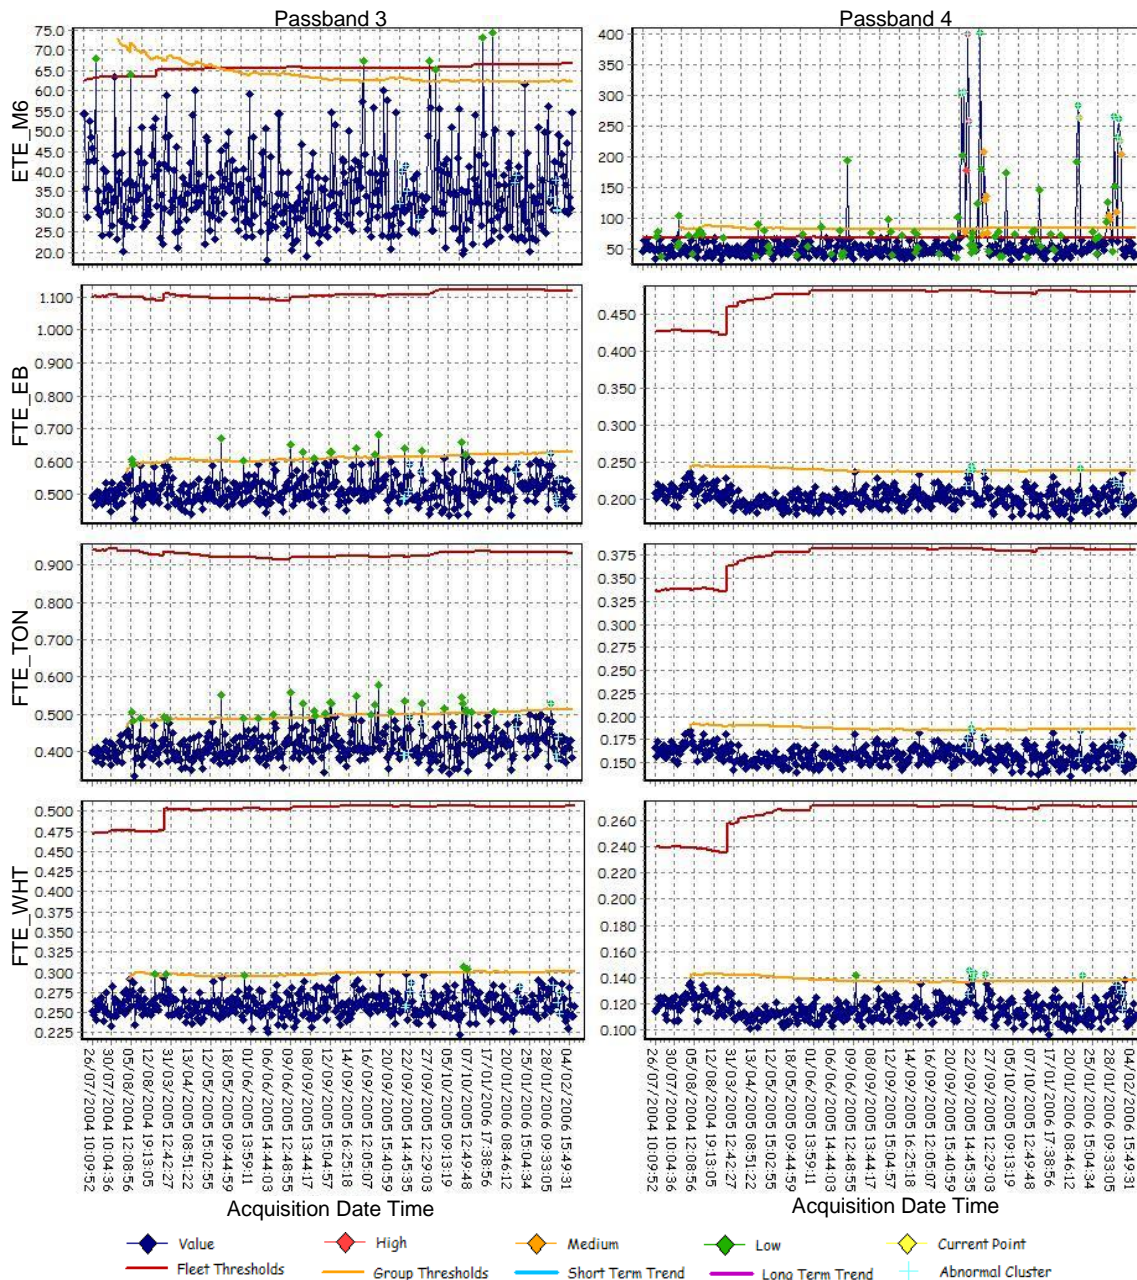


Figure 6.25: AC33 LH-CB123

AC20 RH-CB123

The M6 values for passband 4 were found to step up to high values on the 29th of November 2004. The M6 values remained high until June 2005, when they get even higher before returning to normal levels. These values relate to case 5 described in section 3.2.3.5.

Confidence: 15.6 Number of Alerts: High: 9 Medium: 40 Low: 150




Figure 6.26: AC20 RH-CB123

AC2 RH-CB123

The M6 values for passband 4 were found to be very high from the start of the data in June 2004. These were investigated by the MOD on the 11th of October 2004. Maintenance was then performed after the 20th of October 2004. The next acquisition was on the 15th of February 2005 and the M6 values returned to normal. The maintenance also caused the energy CIs in passband 4 to step up to values above the aircraft threshold, but still below the fleet threshold, resulting in many low level warnings. Note that the confidence level relates to the most recent CI values.

Confidence: 14.0 Number of Alerts: High: 42 Medium: 40 Low: 710




Figure 6.27: AC2 RH-CB123

APPENDIX B: TORNADO EHUMS INFORMATION

A1.11 Corrections to ISA Conditions

Temperatures (6.1), pressures (6.2), fuel-flows (6.3) and spool speeds (6.4) can be corrected to standard-day (International Standard Atmosphere) conditions, in order to compare like with like. The following equations for correction are taken from [117]. They assume that temperature is in Kelvin and pressures in KPascal.

$$T_{COR} = \frac{288.15 \times T}{T_{AMB}} \quad (6.1)$$

$$P_{COR} = \frac{101.325 \times P}{P_{AMB}} \quad (6.2)$$

$$N_{COR} = \frac{N}{\sqrt{\frac{T_{AMB}}{288.15}}} \quad (6.3)$$

$$F_{COR} = \frac{101.325 \times F}{P_{AMB} \sqrt{\frac{T_{AMB}}{288.15}}} \quad (6.4)$$

A1.12 Tornado EHUMS Anomaly Detection Results Table

This section presents the results from running the performance anomaly detection algorithms on Tornado engine data. The table shows all engines that have a PA value greater than 0.9 or an LL value more than three standard deviations from the mean, for the dry or reheat snapshot, for at least one flight. The table has one row for each engine. The engine serial numbers have been de-identified but have a one-to-one mapping to real RB199 engines. The ‘alert’ columns indicate that the data is above the threshold for at least three out of four successive flights. The table also gives the following information:

Column Name	Description
Row	Row number as in the FUMS display
Engine ID	De-identified engine number
LL>3SD	Number of flights more than 3 SD from the LL Mean
LL Alerts	Number of LL>3SD for 3 out of 4 consecutive flights
PA>0.9	Number of flights with PA greater than 0.9
PA Alerts	Number of PA>0.9 for 3 out of 4 consecutive flights
Sum	Sum of all alert columns (the table is sorted by this)
Count	The number of alert columns that are non-zero
Max	The maximum number of flights above a threshold
Total Flights	The total number of flights available for that engine
Num Fits	The number of engine fits (installations on different aircraft)
% (Max/Total Flights)	The percentage of flights above a threshold
flight hours in alert	The approximate number of flight hours that an alert was raised before the issue was resolved
Dry/Reheat/Both	Whether the dry (D), reheat (R) or both (B) snapshots have flights in alert
Alert in Fit	The fit number in which the alerts are present
Engine Position	Left (0) or right (1) engine
Fault?	Whether manual inspection of the data and MWOs indicate that this is likely to be a fault of some sort
Description	Comments on what the data is doing and what the deltas indicate
Related MWO	The text of Maintenance Work Orders immediately before or after the fault, where available.

The section after the table shows the charts of each of the engines listed in the table. The charts show the original data for OAT, NH, NL and TBT, the data after they have been corrected to ISA condition, and the deltas calculated by the ProDAPS anomaly model. The charts also show the LL and PA values. The x axis of each chart is the accumulative engine run time in hours. Each chart shows the dry snapshot values in blue and the reheat values in green. Fleet mean and three standard deviation limits are also shown to make it clear when data is outside the normal bounds. Finally vertical maroon lines indicate when an engine has been fitted to a different aircraft. The charts are labelled using the engine identifier, e.g. ‘EN001’.

Row	Eng ID	Dry				Reheat				Sum	Count	Max	Total flights	Num Fits	% (Max /total flights)	Flight hrs in alert	Dry/Reheat Both	Alert in Fit	Position	Fault?	Description	Related MWO
1	EN046	26	16	3	3	6	4	39	36	133	8	39	133	2	29%		B	1	0	N	Mainly Reheat Alerts. High NL.	None
2	EN025	12	12	12	12	12	12	12	12	96	8	12	144	1	8%	90	B	1	1	Y	TBT stuck at 600 for 13 flights.	25/07/07 - NO TBT ON LANES TEST 08/08/07 - [ENC-9S] ON PLACARD CHECK RIGHT ENGINE ACHIEVED 101.3 HN, 932 DEG C TBT AND STEADY R TBT CAPTION - [E1T] RH MECU REPLACED
3	EN040	13	13	13	13	8	8	13	13	94	8	13	18	1	72%		B	1	0	N	All points in Alert. NL/NH slightly high, TBT slightly low. Doesn't seem like real fault - possible good engine.	19/01/09 - LH BURNER BLOW OUT ON T/O SUBSEQUENTLY RE-LIT NO PROBLEMS AFTER SOAKING FOR 15-20 SECONDS
4	EN017	18	11			4		22	16	71	5	22	146	1	15%	264	B	1	1	Y	All data looks fairly anomalous throughout, with NH being slightly high.	None
5	EN168	6				11	7	20	20	64	5	20	42	1	48%	95	B	1	1	Y	Low NL, high NH for most data. Reheat more severe.	09/09/08 - R/H ECU THROTTLE FAILS MP50-20/1 PARA 6.14 (NO ENGINE RUNAWAY PROTECTION)
6	EN036	9	9	8	8	3	3	9	7	56	8	9	41	1	22%	20	B	1	0	Y	Data anomalous after 60 hours (Nov 08) due to high NH. All one download.	16/10/08 - L/H ECU COMPRESSOR STAGE ONE FODDED DAMAGE OUT OF LIMITS - Modules 1-3 replaced
7	EN014					7		23	18	48	3	23	50	1	46%		R	1	0	N	All points only for reheat - TBT high, NH low.	18/09/08 - LH ECU STUCK IN ENC NONE RESETABLE. 24/10/08 - ECU LIFE EX - MODULE No 10 IPT Rotor
8	EN104	7	5			4	3	14	13	46	6	14	115	3	12%	90	B	1	1	Y	High NH, low NL. Becomes normal just before fit change.	16/04/08 - MECU cannibalised 21/08/08 - RH ECU TO BE REPLACED
9	EN068	3	3			6		23	3	38	5	23	146	1	16%		B	1	0	N	Only Reheat alerts. NL very spiky, NH low.	None
10	EN236	9	7			2		9	9	36	5	9	60	1	15%	21	B	1	1	Y	NH operating high after 100 hours (19/03/08). No data between 23/08/07 and 19/03/08.	30/06/08 - FURTHER INVEST REQUIRED FOR RH ECU PULSING. 04/07/08 - RH MFCU TO BE REPLACED
11	EN201	3		2		5	5	8	5	28	6	8	113	2	7%	49	B	1	1	Y	NL and TBT spike at 32 hours - data looks corrupt. However, the majority of the alerts are caused by high TBT. There is an improvement after an engine change. There is no data for this engine between 30/05/07 and 29/10/08.	29/05/08 - RH REHEAT FAILS TO LIGHT



Row	Eng ID	Dry				Reheat				Sum	Count	Max	Total flights	Num Fits	% (Max /total flights)	Flight hrs in alert	Dry/Reheat Both	Alert in Fit	Position	Fault?	Description	Related MWO
		LL>3SD	LL Alerts	PA>0.9	PA Alerts	LL>3SD	LL Alerts	PA>0.9	PA Alerts													
12	EN083	2				2		11	11	26	4	11	44	2	25%		B	2	0	N	Mostly Reheat alerts with one Dry spike at the end of the data. Seems to be caused by low NL. Data is sporadic with one download from June 2006 (when fitted to ZA447), three from October 2008 and one from February 2009.	Too large date range to find MWO
13	EN055	2				1		10	10	23	4	10	48	1	21%	23	B	1	1	Y	High NH, low NL for ten flights then returns to normal.	02/08/06 - R/H E.C.U. TO BE REPLACED, FAILS E.F.D.C. (MAR/77/06 REFERS) 03/08/06 - R/H ECU RHFCU MAIN FUEL FEED SUPPLY PIPE FOUND DAMAGED
14	EN172					2		13	4	19	3	13	101	1	13%		R	1	1	N	Intermittent Reheat Alerts with slightly high TBT and low NH. NL seems erratic.	N/A
15	EN204	2						8	8	18	3	8	41	1	20%	36	B	1	1	Y	Two periods of high NH from 42 hours onwards - all are from 1 download on 6/7/06.	23/06/06 - R/H LP COMPRESSOR DAMAGED BEYOND LIMITS (FOD)
16	EN229							9	9	18	2	9	113	2	8%		R	1	1	N	Low TBT in a hot environment - perhaps good engine?	N/A
17	EN081					2		10	5	17	3	10	81	1	12%		R	1	0	N	Intermittent Reheat alerts with high NL, slightly high TBT and low NH.	N/A
18	EN131	4				4		5	3	16	4	5	119	1	4%		B	1	0	N	Intermittent Reheat alerts with slightly high NH and low TBT.	N/A
19	EN089					3		7	6	16	3	7	27	1	26%		R	1	0	N	Reheat alerts with high NL and TBT, and low NH.	None
20	EN163					4		7	3	14	3	7	104	1	7%		R	1	1	N	Intermittent Reheat alerts with high NL, slightly high TBT and low NH.	N/A
21	EN018					3		10		13	2	10	122	2	8%		R	1&21	N	All the data is intermittently in Reheat alert, with low TBT - this may be a good operation in a hot climate (OAT is mostly > 40 °C). Unlikely to be a fault since it has persisted for so long (500 hours!). Seems to be clear shift in Delta TBT around 100 hours (20/05/07).	N/A	
22	EN159	7	6							13	2	7	13	1	54%		D	1	0	N	Possibly high NH - PA below 0.9.	N/A
23	EN020					3		5	3	11	3	5	34	1	15%		R	1	0	N	Intermittent Reheat alerts with slightly high TBT and low NH.	N/A
24	EN002							7	4	11	2	7	213	2	3%	197	R	1	0	Y	Reheat alerts with high NL, TBT and low NH - returns to normal after engine change (4/10/06).	08/11/06 - LH ECU RUNS BELOW MITL

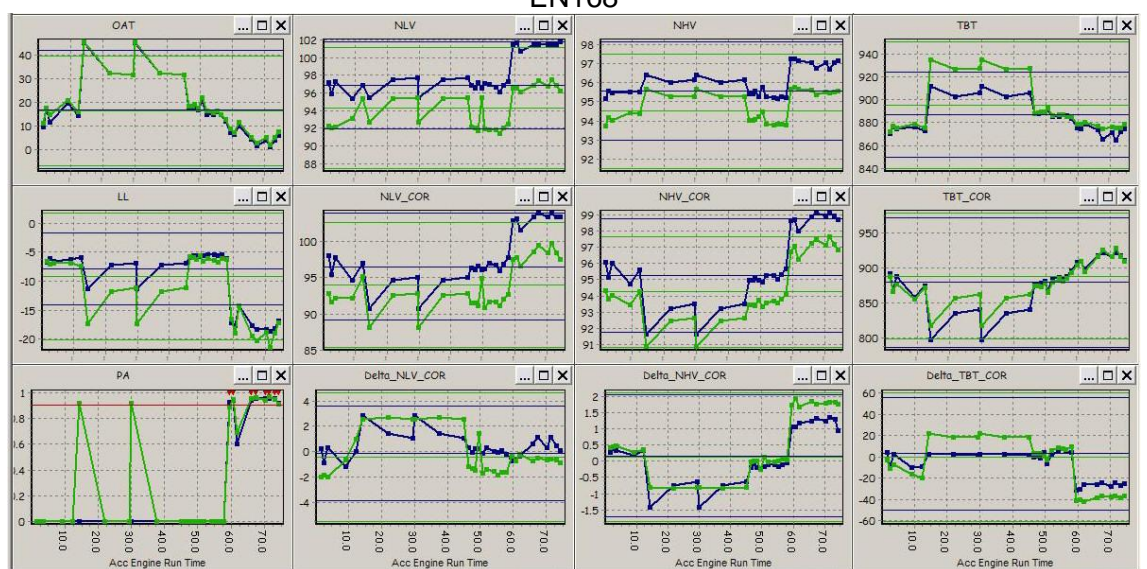


Row	Eng ID	Dry				Reheat				Sum	Count	Max	Total flights	Num Fits % (Max /total flights)	Flight hrs in alert	Dry/Reheat Both	Alert in Fit	Position	Fault?	Description	Related MWO
		LL>3SD	LL Alerts	PA>0.9	PA Alerts	LL>3SD	LL Alerts	PA>0.9	PA Alerts												
25	EN230					1	5	4	10	3	5	14	1	36%		R	1	0	N	Intermittent Reheat alerts with slightly high TBT and low NL/NH.	N/A
26	EN045	1					9		10	2	9	76	2	12%	58	B	2	0	Y	NH trending high, mainly after fit change.	21/01/09 - ENGINE SURGE led to Engine change. 27/01/09 - L/H ECU FAILS TO GOVERN ON NL GOVERNOR TEST
27	EN048						5	4	9	2	5	20	1	25%		R	1	0	N	Only Reheat alerts. TBT low, possible linked to high OAT.	
28	EN027	2		2		2	2		8	4	2	154	3	1%	18	B	3	1	Y	A couple of very high values across all three parameters. In addition, there appears to be shifts in the data for NH and NL at approximately 88, 100 and 138 hours.	14/11/08 - RH ECU TO BE REPLACED (FODDED)
29	EN100	2		2		2	2		8	4	2	61	2	3%	6	B	1	0	Y	Low NL, high NH. Returns to normal after engine change. No data for this engine between 19/04/05 and 18/04/07.	No record of engine 6800 removed/fitted
30	EN030					1	4	3	8	3	4	98	2	4%	25	R	1	0	Y	High NL and TBT and low NH - cleared after engine change.	22/01/07 - EFDC REPORT MAR/11/07 REJECT LH ECU
31	EN043	5	3						8	2	5	145	1	3%		D	1	1	N	Noticeable shift in NH at about 80 hours (July 07).	04/07/07 - R/H ECU TO BE REPLACED, EFDC REPORT MAR 75/07 REFERS
32	EN042					2	5		7	2	5	85	2	6%		R	1	1	N	Intermittent Reheat alerts - but deltas look normal? No data between 8/9/06 and 4/4/09.	29/08/06 - [ENC-SCX-DHW] RH ECU to be replaced because of no NL indication
33	EN157	2		2		1	1		6	4	2	59	1	3%		B	1	1	N	2 spikes: one at 25 hours (Sept 08) and one at 78 hours (Feb 09) with high NH/NL and low TBT. The second appears in both models.	None.
34	EN142					1	5		6	2	5	92	1	5%		R	1	1	N	Intermittent Reheat alerts with slightly high TBT/NL and low NH.	N/A
35	EN165	3	3						6	2	3	44	1	7%		D	1	1	N	A few high Dry PA values. Possible high TBT. Doesn't look significant.	N/A
36	EN034					1	4		5	2	4	89	1	4%	80	R	1	1	Y	Intermittent Reheat Alerts with slightly high NH and low TBT.	N/A
40	EN115	2					1		3	2	2	53	2	4%	7	B	1	1	Y	Six possible anomalies in a row before an engine change. High NL, low NH - possibly linked to high OAT.	06/08/08 - ECU FAILS TO START - INTERNAL FAILURE OF ENGINE GEARBOX. RH ECU NO ROTATION




Row	Eng ID	Dry				Reheat				Sum	Count	Max	Total flights	Num Fits	% (Max /total flights)	Flight hrs in alert	Dry/Reheat Both	Alert in Fit	Position	Fault?	Description	Related MWO
		LL>3SD	LL Alerts	PA>0.9	PA Alerts	LL>3SD	LL Alerts	PA>0.9	PA Alerts													
41	EN008						3		3	1	3	55	2	5%	71	R	1	0	Y	Low NH, high TBT, then switches to high NH, low TBT (around 46 hours Nov 08), then returns to normal after engine change.	04/11/08 - FAILS TO REVERSIONARY LANE WITH REHEAT CAPTION. LH PTO SHFT FOUND DAMAGED 18/03/09 - LH ECU NOZZLE FREES AT 80% AJ ON SELECTION OF TAXI. 02/04/09 - LH BUCKET. GROUNDCREW REPORTED SYMPTOMS WHICH WOULD NOT RESET. DRY CONDITIONS. LH UPPER THRUST REVERSE BUCKET CRACKED	
42	EN113						3		3	1	3	133	1	2%		R	1	1	N	Intermittent alerts, high NL.		
43	EN120						3		3	1	3	89	1	3%		R	1	0	N	Intermittent Reheat alerts with slightly high TBT/NL and low NH.	N/A	
44	EN188				1		1		2	2	1	12	1	8%		R	1	1		Low TBT. Snapshot at 12 hours looks incorrect as it is towards the end of the flight.	17/03/09 - RH ENGINE FUEL FLOW 6-8 KG/MIN LESS THAN LH ENGINE BETWEEN 75-90% NH MATCHED THROTTLES .SEA LEVEL TO 16K FEET. 23/03/09 - AIRTEST- R/H ENGINE WINDMILL RELIGHT E11 FAILED TO LIGHT AFTER 80S - MAX NH 20% MAX T7 290 DEG	
45	EN054				1		1		2	2	1	16	1	6%		R	1	1		Single point, slightly low No, high TBT.	N/A	
46	EN057	1		1					2	2	1	58	2	2%	0	D	1	1	Y	Single point, very high for all parameters - possible corruption issue?	N/A	
47	EN182	2							2	1	2	50	1	4%		D	0	1		High reheat NL for most of the data.	N/A	
48	EN206						2		2	1	2	66	1	3%	96	R	1	0	Y	Reheat high NH, low TBT. Returns to normal about 85 hours (March 08).	08/03/08 - [ENR-P5] LH AJ GAUGE INDICATES 50%. [A1T] LH ECU NPT (TRANSDUCER NOZZLE POSN) REPLACED & FUNCTIONALLY TESTED SATIS	
49	EN194						2		2	1	2	53	2	4%	40	R	2	0	Y	Not clear what is causing this?	None	
50	EN112						2		2	1	2	120	2	2%		R	1	1		A few fairly anomalous reheat points at the start of the data with slightly high NH and low TBT.	23/09/05 - NOZZLE SLOW TO OPEN ON SELECTION OF TAXI NOZZLE. RH MECU TO BE REPLACED	
51	EN122						2		2	1	2	4	1	50%		R	1	0		All four data points have high TBT and low NH.	None	
52	EN155	2							2	1	2	34	1	6%		D	1	0		Two low NL points.	N/A	




Row	Eng ID	Dry				Reheat				Sum	Count	Max	Total flights	Num Fits	% (Max /total flights)	Flight hrs in alert	Dry/Reheat Both	Alert in Fit	Position	Fault?	Description	Related MWO
		LL>3SD	LL Alerts	PA>0.9	PA Alerts	LL>3SD	LL Alerts	PA>0.9	PA Alerts													
53	EN170	2								2	1	2	13	1	15%		D	1	0		Not significant.	N/A
54	EN197	2								2	1	2	7	1	29%		D	1	0		Slightly high TBT, low NL.	N/A
55	EN187	2								2	1	2	17	1	12%		D	1	1		Two points slightly high TBT.	N/A
56	EN126	1								1	1	1	62	1	2%		D	0	1		Reheat NL high.	N/A
57	EN235						1		1	1	1	1	127	2	1%	105	R	2	0	Y	Low NL/TBT high NH after engine change. Steps back again at 140 hours.	None
58	EN145						1		1	1	1	1	74	1	1%		R	1	0		High NL, low NH.	N/A
59	EN212	1							1	1	1	1	44	1	2%	0	D	0	0	Y	Dry NH seems very high for some periods. Only intermittent reheat alerts raised.	18/08/08 - LH ECU #8029 COMBUSTION CHAMBER HOLED AND UNZIPPING
60	EN233						1		1	1	1	1	160	3	1%	19	R	1&20	Y		NL and TBT spike at 135 hours (ZA401 Sept 07). Additional period of data with low TBT and high NH.	17/08/08 - L/H ENGINE SURGE DURING FLIGHT NO CWP CAPTIONS BUT L/H TBT RISE SEEN. INVESTIGATE HP SPOOL STIFF TO TURN. MODULE No 03 (HPC) replaced

A1.13 Tornado EHUMS Anomaly Detection Results Charts

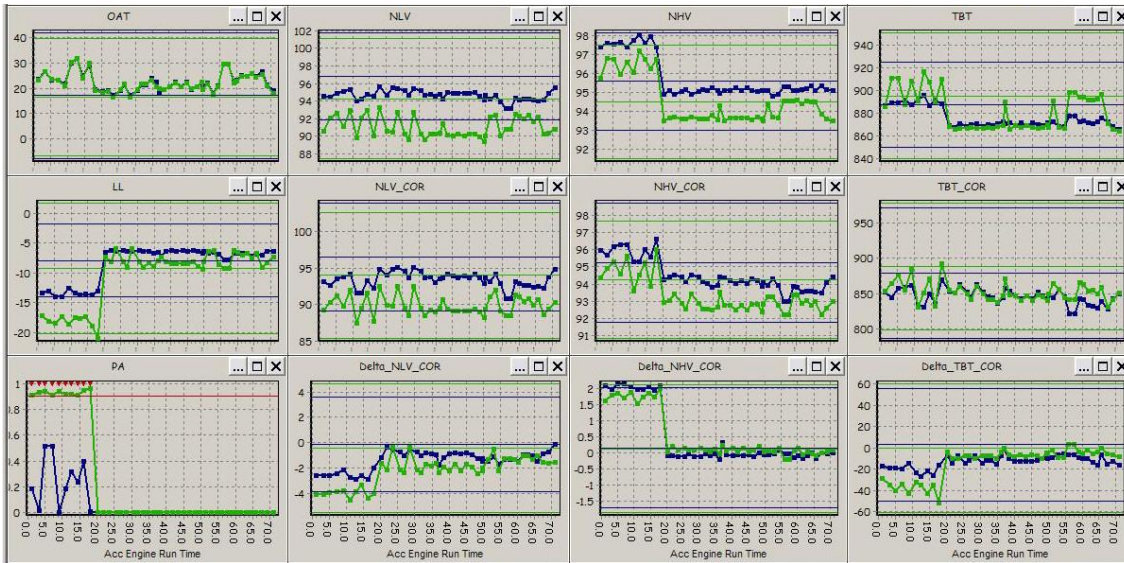










EN083





EN055




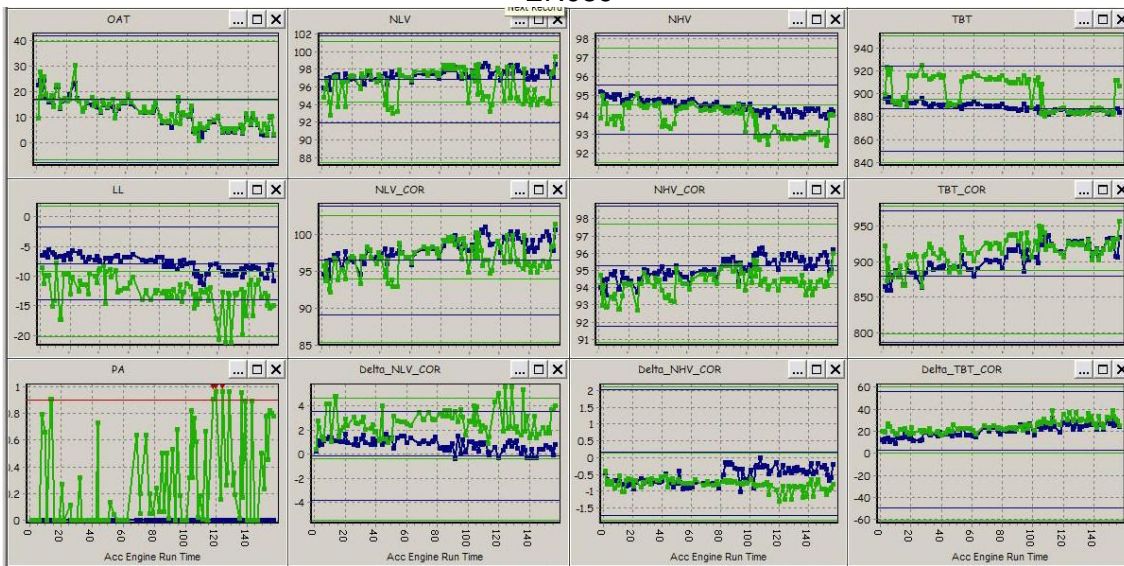

EN172

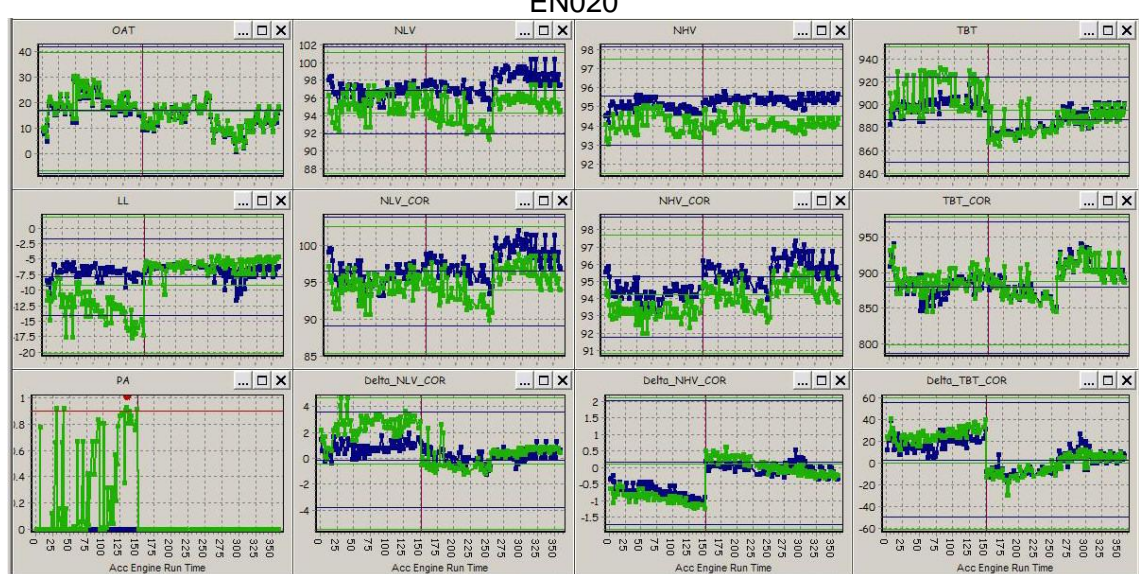
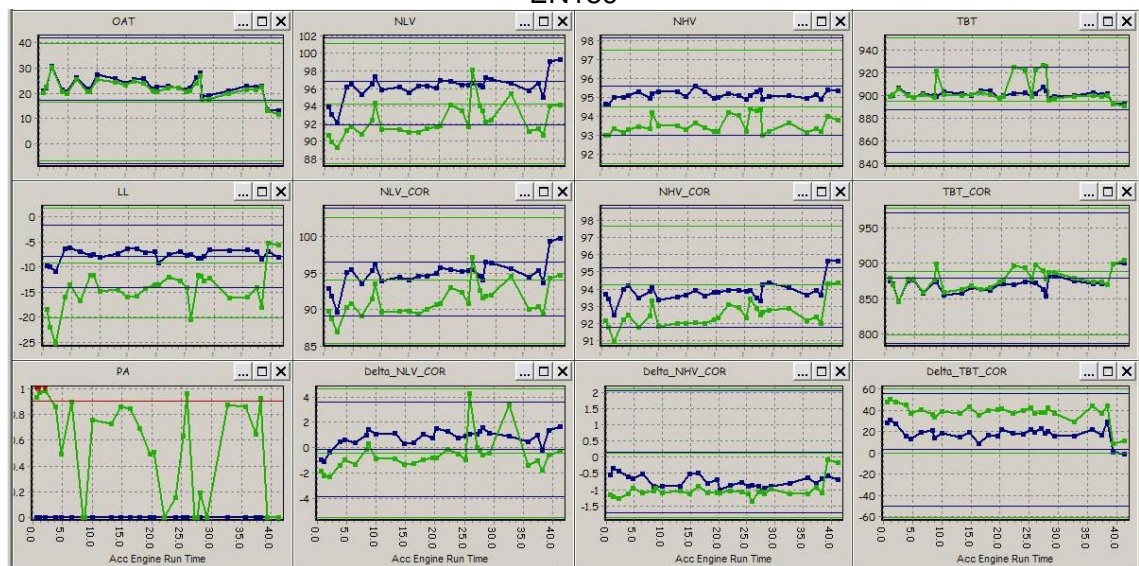



EN204



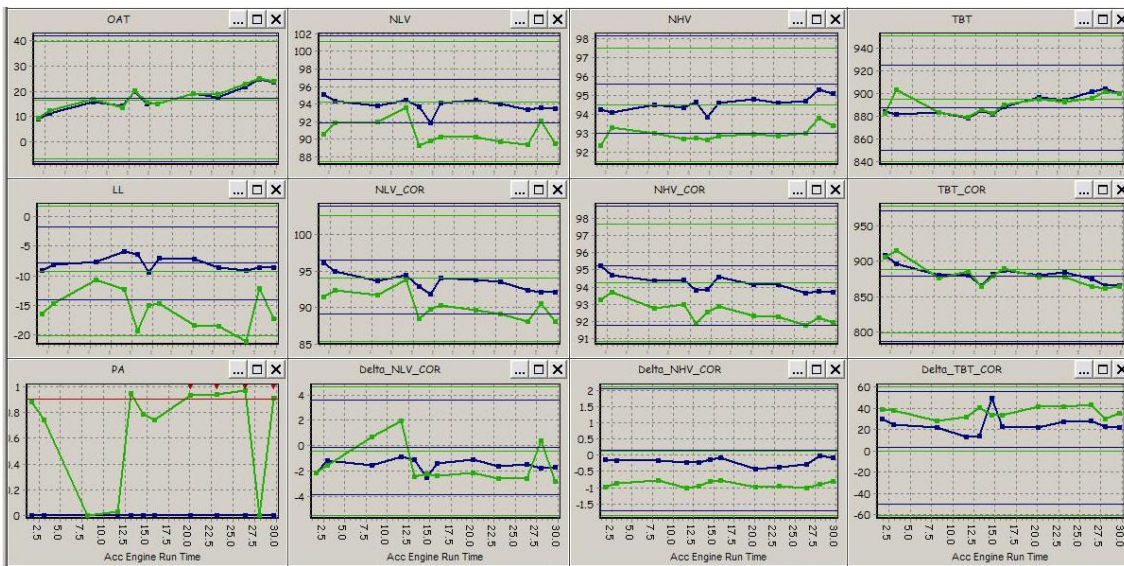


EN131


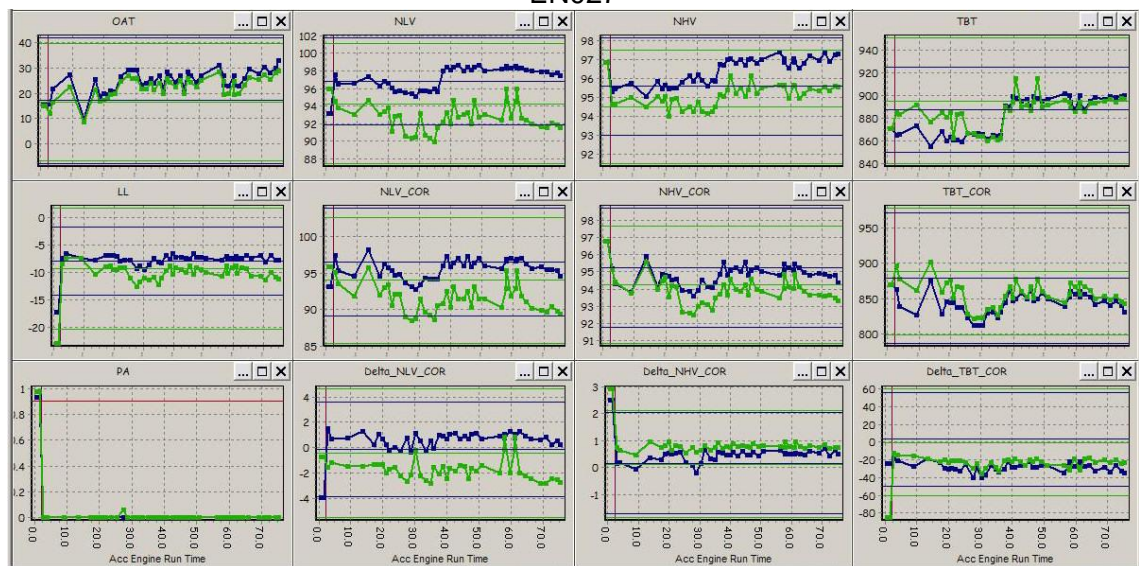
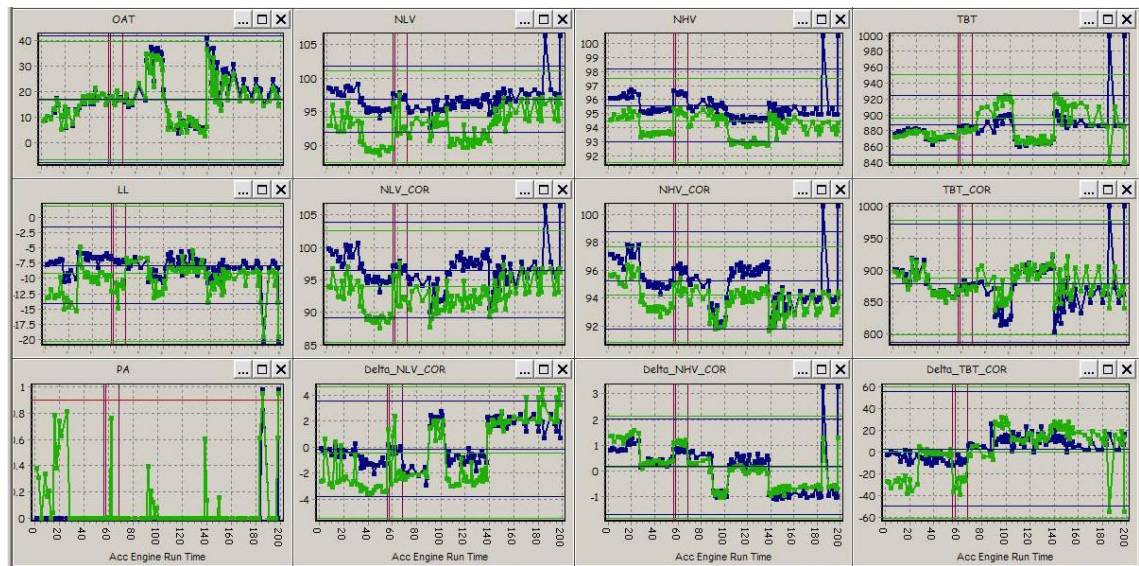







EN002









EN048












EN157







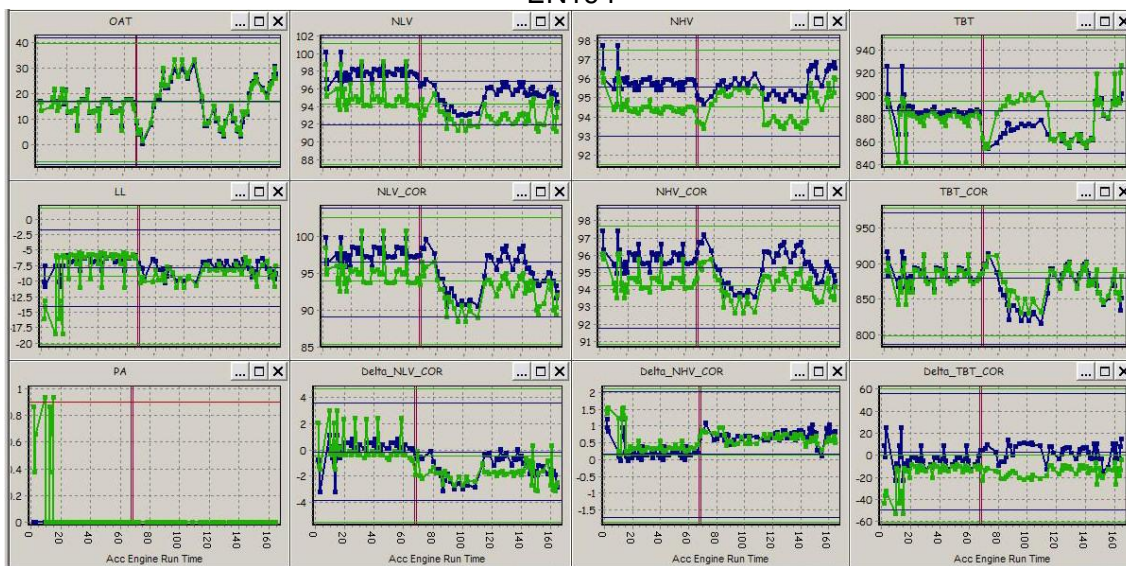



EN119


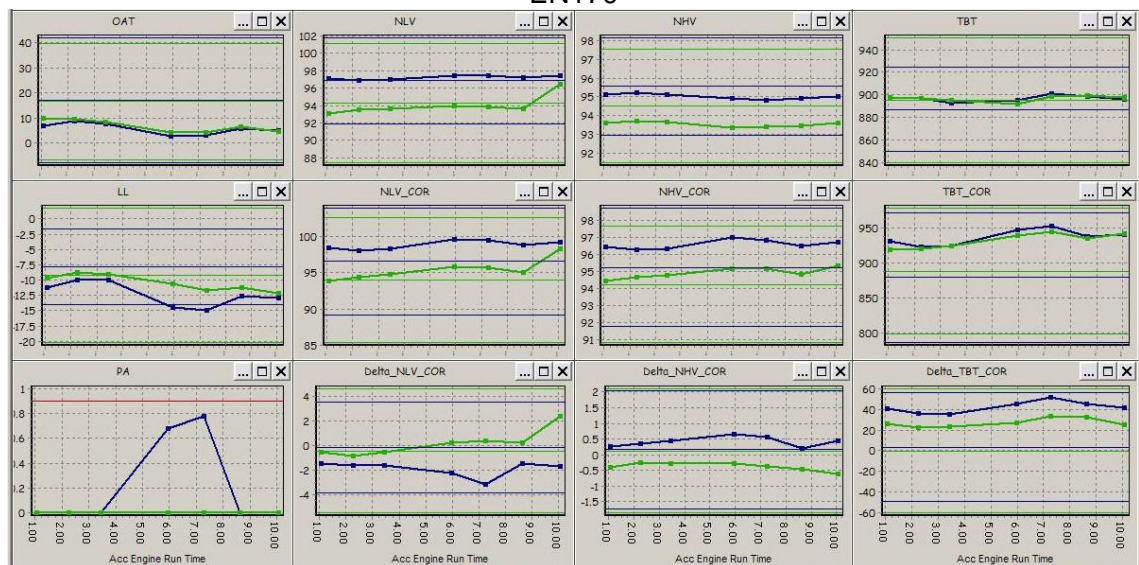
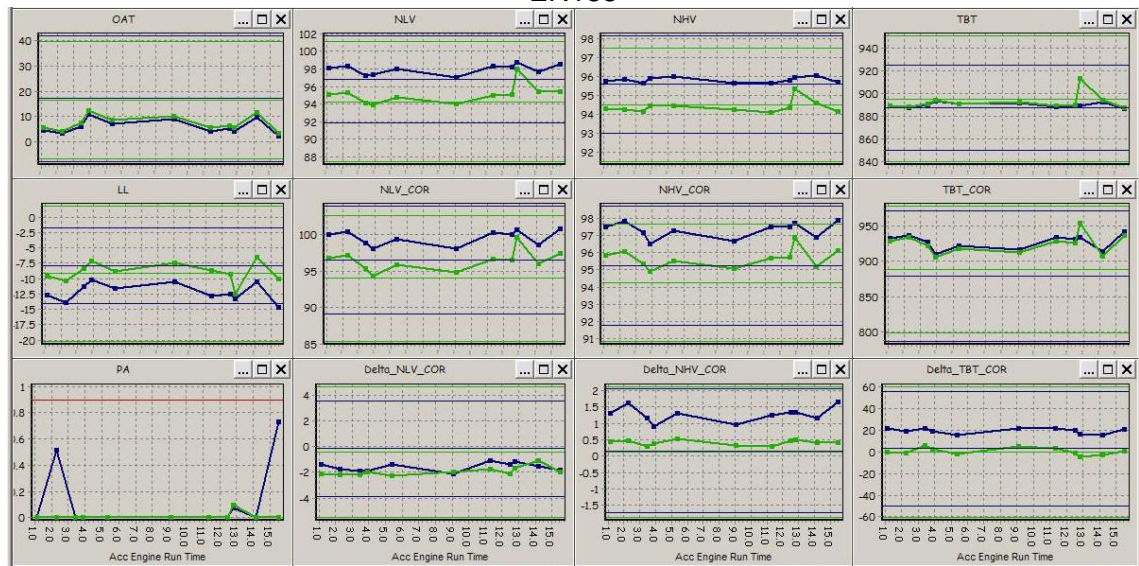



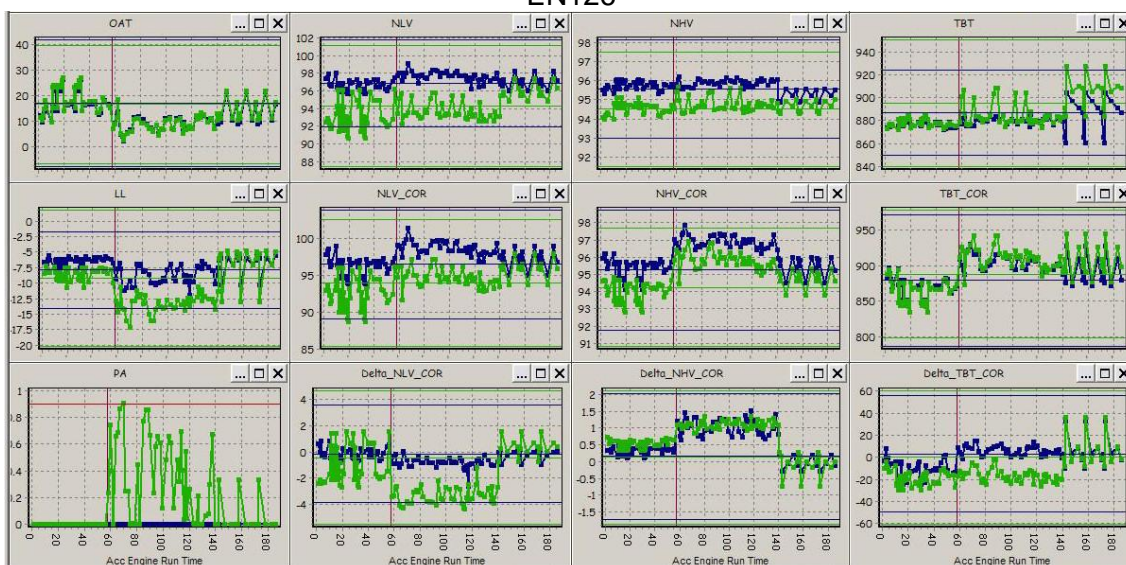










EN122







REFERENCES

- [1] I. Mackersey, *The Wright Brothers*, Little, Brown 2003. 0 316 86144 8
- [2] "Eye Witness History," <http://www.eyewitnesstohistory.com/wright.htm>, accessed: May 2010
- [3] "Plane Crash Info - Japan 747 Crash," <http://www.planecrashinfo.com/w19850812.htm>, accessed: May 2010
- [4] "Air Disaster.com - Air France Crash 2000," http://www.airdisaster.com/cgi-bin/view_details.cgi?date=07252000®=F-BTSC&airline=Air+France
- [5] "Air Disaster.com - Air France Crash 2000 - Photo," <http://www.airdisaster.com/photos/afsst/2.shtml>, accessed: May 2010
- [6] "Air Accidents Investigation Branch " www.aaib.gov.uk, accessed: May 2010
- [7] "Institution of Mechanical Engineers - Airworthiness," <http://www.imeche.org/industries/aero/airworthiness/>, accessed: May 2010
- [8] "Air Accidents Investigation Branch - Publications - Progress Reports," http://www.aaib.gov.uk/publications/progress_reports.cfm, accessed: May 2010
- [9] "European Aviation Safety Agency - Annual Safety Recommendations review 2009 - SAR-002-2010," http://www.easa.eu.int/ws_prod/g/doc/Safety/2009%20-%20Annual%20Safety%20Recommendations%20Review.pdf, accessed: May 2010
- [10] "US National Transport Safety Board - Aviation Accident Statistics," <http://www.nts.gov/aviation/Table5.htm>, accessed: May 2010
- [11] "Plane Crash Info - Cause," <http://www.planecrashinfo.com/cause.htm> accessed: May 2010
- [12] P. R. Knight, J. Cook, and H. Azzam, "Intelligent Management of Helicopter Health and Usage Management Systems Data," *Proceedings of the Institute of Mechanical Engineers - Part G: Journal of Aerospace Engineering*, vol. 219, pp. 507-524, 2005.
- [13] H. Azzam, J. Cook, P. R. Knight, and E. Moses, "FUMS™ Fusion and Decision Support for Intelligent Management of Aircraft Data," *Proceedings of IEEE Aerospace Conference*, 2006.
- [14] P. R. Knight, H. Azzam, S. J. Newman, and A. J. Chipperfield, "Artificial Intelligence and Mathematical Models for Intelligent Management of Aircraft Data," *International Conference on Prognostics and Health Management*, 2008.

-
- [15] Australian Research Laboratories, "Australia invented the Black Box voice and instrument recorder," <http://apc-online.com/austrade/blackbox.htm>, accessed: May 2007
- [16] D. Beaudouin and C. Beaudouin, *Charles Beaudouin: une histoire d'instruments scientifiques*, EDP Sciences Editions, 2005. 2868838073
- [17] B. D. Larder and N. Norman, "A Trial Helicopter Operations Monitoring Programme (HOMP)," *Royal Aeronautical Society Conference*, vol. London, March 2000.
- [18] B. D. Larder and N. Norman, "Application of HOMP to Helicopter Maritime Operations," *Royal Aeronautical Society Conference*, March 2001.
- [19] D. White and R. Vaughan, "Fleet Usage Monitoring Is Essential In Improving Aging U.S. Army Helicopter Safety, Availability, And Affordability," *9th Joint FAA/DoD/NASA Aging Aircraft Conference*, 2006.
- [20] J. Cook, "Personal communication," MOD, Materials and Integrity Group, Fleetlands, Gosport, 2004.
- [21] L. Liu and D. J. Pines, "Analysis of U.S. Civil Rotorcraft Accidents Caused by Vehicle Failure or Malfunction, 1998 - 2004," *American Helicopter Society 61st Annual Forum*, Grapevine, Texas, 2005.
- [22] A. Draper, "The Operational Benefits of Health and Usage Monitoring Systems in UK Military Helicopters," *Third International Conference on Health and Usage Monitoring - HUMS2003*, 2003.
- [23] J. Land and C. Weitzman, "How HUMS systems have the potential of significantly reducing the direct operating cost for modern helicopters through monitoring," *Proceedings of the AHS 51st Annual Forum*, Fort Worth, TX, 1995.
- [24] M. Koehl, "Algorithmic Aero Engine Life Usage Monitoring Based on Reference Analysis of Design Mission," *3rd International Workshop on Structural Health Monitoring*, Stanford University, Stanford, CA, 2001.
- [25] "Defence Standard (Def Stan) 00-970 Issue 1," 1983.
- [26] A. M. Van-Den-Hoeven, "CC130 Data Analysis System for OLM/IAT," *NATO Research and Technology Organisation (RTO) Applied Vehicle Technology (AVT)*, vol. "Exploitation of Structural Loads/Health Data for Reduced Life Cycle Costs", Brussels, Belgium, 11-12 May 1998.
- [27] A. Draper, "Fatigue Usage Monitoring in UK Military Helicopters," *Workshop on Helicopter Health and Usage Monitoring Systems, Australia, February 1999*, vol. DSTO-GD-0197 pages 153-166, D/DHP/16/2/43, 1999.

-
- [28] H. Azzam, F. Beaven, M. Wallace, N. H. Wakefield, and P. R. Knight, "Qualification Guidelines for Non-Adaptive Prediction Methods," Smiths Aerospace Ltd REP1620, 2004.
- [29] H. Azzam, F. Beaven, A. Smith, and M. Wallace, "FUMS™ Fleet Management Applications and Algorithms," *Aging Aircraft Conference*, 2007.
- [30] "Defence Standard (Def Stan) 00-970 Issue 1," originally issued: 1983, revised 2006. 2006.
- [31] L. C. Jaw, "Recent Advancements in Aircraft Engine Health Management (EHM) Technologies and Recommendations for the Next Step," *Proceedings of Turbo Expo 2005: 50th ASME International Gas Turbine & Aeroengine Technical Congress*, 6-9 June 2005.
- [32] D. Roach and K. Rackow, "Health Monitoring Of Aircraft Structures Using Distributed Sensor Systems," *9th Joint FAA/DoD/NASA Aging Aircraft Conference*, 2006.
- [33] A. M. Toms and M. P. Barrett, "Using Filter Debris Analysis To Identify Component Wear In Industrial Applications," *Proceedings of the 61st Meeting of the Society for Machinery Failure Prevention Technology*, vol. Integration of Machinery Failure Prevention Technologies into Systems Health Management, pp. 239–245., 2007.
- [34] D. E. Veinot and G. C. Fisher, "Wear Debris Examination as a Condition Monitoring Technique for the Sikorsky Sea King Helicopter Main Gearbox," *STLE Special Publication SP-27*, pp. 119-131, May 1989.
- [35] P. J. Dempsey, G. Kreider, and T. Fichter, "Investigation of Tapered Roller Bearing Damage Detection Using Oil Debris Analysis," *Proceedings of IEEE Aerospace Conference*, 2005.
- [36] C. E. Fisher, "Gas Path Debris Monitoring – A 21st Century PHM Tool," *Proceedings of IEEE Aerospace Conference*, 2000.
- [37] H. Powrie and A. Novis, "Gas Path Debris Monitoring for F-35 Joint Strike Fighter Propulsion System PHM," *Proceedings of IEEE Aerospace Conference*, 2006.
- [38] A. J. Evans, "Managing a Successful HUMS Operation," *Third International Conference on Health and Usage Monitoring - HUMS2003*, 2003.
- [39] B. Larder, H. Azzam, C. Trammel, and G. Vossler, "Smith Industries HUMS: Changing the M from Monitoring to Management," *Proceedings of IEEE Aerospace Conference*, 2000.
- [40] K. Pipe, "Measuring the Performance of a HUM System - the Features that Count," *Third International Conference on Health and Usage Monitoring - HUMS2003*, 2003.

- [41] J. A. Keller, R. Branhof, D. Dunaway, and P. Grabill, "Examples Of Condition Based Maintenance With The Vibration Management Enhancement Program," *American Helicopter Society 61st Annual Forum, Grapevine, Texas*, 2005.
- [42] "Largest Data Size Data-Mined," http://www.kdnuggets.com/polls/2007/largest_database_data_mined.htm, accessed: Jan 2011
- [43] G. Venn, "Personal communication," Westland Helicopters Limited, 2003.
- [44] H. Azzam, F. Beaven, I. Hebden, L. Gill, and M. Wallace, "Fusion and Decision Making Techniques for Certifiable, Affordable Structural Prognostic Health Management " *Proceedings of IEEE Aerospace Conference*, 2005.
- [45] M. J. Roemer and G. J. Kacprzynski, "Advanced Diagnostics and Prognostics for Gas Turbine Engine Risk Assessment," *Proceedings of IEEE Aerospace Conference*, 2005.
- [46] K. Goebel, P. Bonanni, and N. Eklund, "Towards an Integrated Reasoner for Bearings Prognostics," *Proceedings of IEEE Aerospace Conference*, 2005.
- [47] J. Han and M. Kamber, *Data Mining: Concepts and Techniques*, Morgan Kaufmann Publishers, 2001. 1-55860-489-8
- [48] G. Barndt, S. Sarkar, and S. Maley, "The Effects of Degraded Data on the Performance of Health Usage Monitoring Systems (HUMS) in Rotary Wing Aircraft," *9th Joint FAA/DoD/NASA Aging Aircraft Conference*, 2006.
- [49] H. Azzam, J. Cook, P. R. Knight, and N. H. Wakefield, "FUMS™ Fusion For Improved Aircraft MAAAP," *Proceedings of IEEE Aerospace Conference*, 2005.
- [50] H. Azzam, F. Beaven, M. Wallace, and I. Hebden, "Optimisation of Fusion and Decision Making Techniques for Affordable SPHM," *Proceedings of IEEE Aerospace Conference*, 2006.
- [51] A. J. Volponi, T. Brotherton, R. Luppold, and D. L. Simon, "Development of an Information Fusion System for Engine Diagnostics and Health Management," NASA, vol. NASA/TM—2004-212924, 2004.
- [52] K. Goebel, M. Krok, and H. Sutherland, "Diagnostic Information Fusion: Requirements Flowdown and Interface Issues," *Proceedings of IEEE Aerospace Conference*, 2000.
- [53] N. Iyer, K. Goebel, and P. Bonissone, "Framework for Post-Prognostic Decision Support," *Proceedings of IEEE Aerospace Conference*, 2006.
- [54] K. Goebel, N. Eklund, and P. Bonanni, "Fusing Competing Prediction Algorithms for Prognostics," *Proceedings of IEEE Aerospace Conference*, 2006.

- [55] R. Callan, B. Larder, and J. Sandiford, "An Integrated Approach to the Development of an Intelligent Prognostic Health Management System," *Proceedings of IEEE Aerospace Conference*, 2006.
- [56] N. Lybeck, B. Morton, S. Marble, A. Hess, and J. Kelly, "Modelling and Simulation of Vibration Signatures in Propulsion Subsystems," *Proceedings of IEEE Aerospace Conference*, 2006.
- [57] K. P. J. Bryant and H. Azzam, "Intelligent Analysis Of Damage Acoustic Emission Data," Smiths Aerospace, Electronic Systems - Southampton REP1599, 2004.
- [58] M. Barnathan, "Mining complex high-order datasets," *Dissertation for PhD, Temple University*, 2010.
- [59] J. S. Aguilar-Ruiz, J. H. Moore, and M. D. Ritchie, "Filling the gap between biology and computer science," *BioData Mining*, 2008.
- [60] R. Bellazzi and B. Zupan, "Towards knowledge-based gene expression data mining," *J. Biomed. Inform.*, 2007.
- [61] "Data Mining Methods "
http://www.kdnuggets.com/polls/2006/data_mining_methods.htm, accessed: Jan 2011
- [62] N. H. Wakefield and H. Azzam, "A Novel Cluster Identification Algorithm Using a Minimum Area Method," Smiths Aerospace, Electronic Systems - Southampton MJAD/R/311/02, 2002.
- [63] Wikipedia, "Support vector machine,"
http://en.wikipedia.org/wiki/Support_vector_machine, accessed: May 2007
- [64] S. Haykin, *Neural Networks - A Comprehensive Foundation*, Macmillan Publishing Company, 1994. ISBN: 0-02-352761-7
- [65] N. Nikolaev, "Single-Layer Perceptrons," University of London.
- [66] "Introduction to Neural Networks,"
<http://www.itee.uq.edu.au/~cogs2010/cmc/chapters/Introduction/>, accessed: May 2007
- [67] B. J. Copeland and D. Proudfoot, "Turing's Neural Networks of 1948," 2000.
- [68] J. H. Holland, "Introduction to Genetic Algorithms."
- [69] C. Darwin, *The Origin of Species by Means of Natural Selection: Or, the Preservation of Favored Races in the Struggle for Life*, John Murray, London, 1859.

-
- [70] H. Azzam, F. Beaven, M. Wallace, and I. Hebden, "Optimisation of Fusion and Decision Making Techniques for Affordable SPHM " *Proceedings of IEEE Aerospace Conference*, 2007.
- [71] P. Harmon and D. King, *Expert Systems*, John Willey & Sons, 1985. ISBN: 0-471-80824-5
- [72] Scholarpedia, "Fuzzy Logic," http://www.scholarpedia.org/article/Fuzzy_Logic, accessed: March 2005
- [73] L. A. Zadeh, "Fuzzy Sets," *Information and Control* vol. 8, pp. 338--353, 1965.
- [74] Scholarpedia, "Fuzzy Sets," http://www.scholarpedia.org/article/Fuzzy_Sets, accessed: May 2007
- [75] Wikipedia, "Fuzzy Logic," http://en.wikipedia.org/wiki/Fuzzy_logic, accessed: March 2005
- [76] T. J. Ross, *Fuzzy Logic with Engineering Applications*, Second ed., John Wiley & Sons, 2004. ISBN: 0-470-86075
- [77] Wikipedia, "Bayesian network," http://en.wikipedia.org/wiki/Bayesian_network, accessed: May 2007
- [78] F. V. Jensen, *An Introduction to Bayesian Networks*, UCL Press, 1996. ISBN: 1-85728-332-5
- [79] H. Azzam, "The Use Of Mathematical Models And Artificial Intelligence Techniques To Improve Hums Prediction Capabilities," *Innovation in Rotorcraft Technology Proceedings, The Royal Aeronautical Society*, 1997.
- [80] C. Byington, R. F. Orsagh, P. Kallappa, J. Sheldon, M. DeChristopher, S. Amin, and J. Hines, "Recent Case Studies in Bearing Fault Detection and Prognosis," *Proceedings of IEEE Aerospace Conference*, 2006.
- [81] M. J. Smith, C. S. Byington, P. Kalgren, A. Parulekar, and M. DeChrisopher, "Layered Classification for Improved Diagnostic Isolation in Drivetrain Components," *Proceedings of IEEE Aerospace Conference*, 2006.
- [82] D. Hochmann and G. Baringer, "Analytical Mechanical Diagnostic Benefits: Case Studies," *Proceedings of IEEE Aerospace Conference*, 2006.
- [83] N. S. Feng, E. J. Hahn, and R. B. Randall, "Simulation Of Vibration Signals From A Rolling Element Bearing Defect," *DSTO International Conference on Health and Usage Monitoring, Melbourne, February 19-20*, 2001.
- [84] R. Hamza, S. Menon, and S. McRoberts, "Advanced Knowledge Management For Helicopter HUMS," *HUMS User Conference*, 2001.

-
- [85] E. Bechhoefer and A. P. F. Bernhard, "Setting HUMS Condition Indicator Thresholds by Modelling Aircraft and Torque Band Variance," *Proceedings of IEEE Aerospace Conference*, 2004.
- [86] E. Bechhoefer and A. P. F. Bernhard, "Experience in Setting Thresholds for Mechanical Diagnostics in the UH-60L Fleet Demonstration," *Proceedings of IEEE Aerospace Conference*, 2005.
- [87] E. Bechhoefer and E. Mayhew, "Mechanical Diagnostics System Engineering in IMD HUMS," *Proceedings of IEEE Aerospace Conference*, 2005.
- [88] H. H. Chin, E. Mayhew, and D. L. Green, "Assessing Bearing Health for Helicopter Power Train Systems," *American Helicopter Society 61st Annual Forum, Grapevine, Texas*, 2005.
- [89] M. J. Ashby and W. J. Scheuren, "Intelligent Maintenance Advisor for Turbine Engines," *Proceedings of IEEE Aerospace Conference*, 2000.
- [90] M. J. Roemer and D. M. Ghiocel, "A Probabilistic Approach To The Diagnosis Of Gas Turbine Engine Faults," *International COMADEM Congress, Tasmania, Australia*, 1999.
- [91] T. Jackson, J. Austin, M. Fletcher, M. Jessop, B. Liang, A. Pasley, M. Ong, X. Ren, G. Allan, V. Kadiramanathan, H. A. Thompson, and P. Fleming, "Distributed Health Monitoring for Aero-Engines on the GRID: DAME," *Proceedings of IEEE Aerospace Conference*, 2005.
- [92] P. Frith and G. Karvounis, "Model-Based Decision Support Tools For T700 Engine Health Monitoring," *HUMS User Conference*, 2001.
- [93] L. Yu, D. J. Cleary, and P. E. Cuddihy, "A Novel Approach to Aircraft Engine Anomaly Detection and Diagnostics," *Proceedings of IEEE Aerospace Conference*, 2004.
- [94] T. Kobayashi and D. L. Simon, "A Hybrid Neural Network-Genetic Algorithm Technique for Aircraft Engine Performance Diagnostics," *American Institute of Aeronautics and Astronautics*, 2001.
- [95] C. S. Byington, M. Watson, and D. Edwards, "Data-Driven Neural Network Methodology to Remaining Life Predictions for Aircraft Actuator Components," *Proceedings of IEEE Aerospace Conference*, 2004.
- [96] T. Khawaja, G. Vachtsevanos, and B. Wu., "Reasoning about Uncertainty in Prognosis: A Confidence Prediction Neural Network Approach," *Proceedings of IEEE Aerospace Conference*, 2005.
- [97] J. R. Bock, T. Brotherton, P. Grabill, D. Gass, and J. A. Keller, "On False Alarm Mitigation," *Proceedings of IEEE Aerospace Conference*, 2006.

-
- [98] H. Azzam, P. R. Knight, D. Hart, and I. Grainger, "FUMS™ for MOD DAPAT – Tasks 2b, 3b, 4b, 7a and 7b Contract COMMPSSPT/015," Smiths Aerospace, Electronic Systems - Southampton REP1580(1), 2003.
- [99] P. R. Knight, R. Ellison, R. Horabin, R. Abbott, K. P. J. Bryant, G. Pitt-Nash, and H. Azzam, "FUMS™ for MOD DAPAT phase II, Task 2.2 and Task 4.3 of Contract COMMPS/042," Smiths Aerospace, Electronic Systems - Southampton RES158(1), 26 March 2007.
- [100] H. Azzam and N. Harrison, "Intelligent Management of Helicopter HUMS Data," CAA Paper 99006, 1999.
- [101] M. D. Pryce, "Personal communication," Assistant Directorate Aircraft Integrity Monitoring (AD AIM), Ministry of Defence (MOD), Fleetlands, Gosport, UK, 2003.
- [102] D. N. Symonds, "Combiner Bearing Failure," Assistant Directorate Aircraft Integrity Monitoring (AD AIM), Ministry of Defence (MOD), Fleetlands, Gosport, UK 2003.
- [103] G. H. Ball and D. J. Hall, "ISODATA, an iterative method of multivariate data analysis and pattern classification," *IEEE International Communications Conference*, 1966.
- [104] N. H. Wakefield, P. R. Knight, K. P. J. Bryant, and H. Azzam, "FUMS™ Artificial Intelligence Technologies Including Fuzzy Logic For Automatic Decision Making," *Annual Meeting of the North-American-Fuzzy-Information-Processing-Society*, 26-28 June 2005.
- [105] F. Beaven, N. H. Wakefield, P. R. Knight, K. P. J. Bryant, and H. Azzam, "FUMS™ for DAPAT – Tasks 8a, 9a, 10a, 10b," Smiths Aerospace, Electronic Systems - Southampton RES134, 2004.
- [106] L. A. Urban, "Parameter selection for multiple fault diagnostics of gas turbine engines,," *AGARD-CP-165 ASME Paper 74-GT-62, J. Eng. Power*, pp. 225-230, April 1975.
- [107] Y. G. Li, "Performance-Analysis-Based Gas Turbine Diagnostics: A Review," *Proceedings of the Institute of Mechanical Engineers - Part A: Journal of Power and Energy*, vol. 216, p. 363, 2002.
- [108] Tornado-Data.com, "RB199 Engine Black and White Cutaway Diagram," <http://www.tornado-data.com/History/engine/cutaway.htm>, accessed: March 2009
- [109] "RAF Engineering Authority General Orders and Special Instructions, Leaflet 302, Health Monitoring of Tornado Engines and Accessories, DAP 101B-4100-2(R)1," 2006.

-
- [110] P. R. Knight, R. Abbott, and R. Ellison, "FUMS™ for MoD DAPAT – Phase 2 – Tasks 5.2 and 8.1 of Contract COMMPS/042," Smiths Aerospace, Electronic Systems - Southampton RES141, 2005.
 - [111] "Chinook Maintenance School Training Manual, Section 12, Tandem Rotor Theory of Flight."
 - [112] "Chinook Helicopter Mk2 Aircrew Manual."
 - [113] "Chinook Maintenance School Training Manual, Section 14, Automatic Flight Control System ".
 - [114] NSK, "Introduction to Bearings," <http://www.nsk.com/services/basicknowledge/introduction.html>, accessed: May 2007
 - [115] Stefan Vorkoetter, "Demystifying Gearing and Gearboxes," <http://www.stefanv.com/rcstuff/qf200003.html>, accessed: May 2007
 - [116] D. Hart, "Signal Processing Techniques for HUMS and FUMS™ Applications," Smiths Aerospace, Electronic Systems - Southampton MJAD/RTDN/003/00, 2000.
 - [117] P. P. Walsh and P. Fletcher, *Turbine Performance*. Oxford, Blackwell Science, 1998.