DEFENCE
SCIENCE & TECHNOLOGY

Third International Conference on Health and Usage Monitoring-HUMS2003

Graham F. Forsyth (editor)

DSTO-GD-0348

Third International Conference on Health and Usage Monitoring - HUMS2003

Graham F Forsyth (editor)

Air Vehicles Division
Platforms Sciences Laboratory

DSTO-GD-0348

ABSTRACT

This document includes formal papers for the Third International Conference on Health and Usage Monitoring, HUMS2003, which will be held in Melbourne in February 2003. The scope of papers covers a wide range of monitoring issues with focus on the application to helicopters, military aircraft and gas turbine engines.

This document includes only the formal papers made available for publication during December 2002. Other papers will be available from the conference CD-ROM which includes this document, other formal papers, presentations not available as formal papers and related software.

For a short period, all these will also be available from the conference website at URL: http://www.dsto.defence.gov.au/corporate/conferences/hums.

RELEASE LIMITATION

Approved for public release

AQ F03-06-1402

Published by

DSTO Platforms Sciences Laboratory
506 Lorimer St
Fishermans Bend, Victoria 3207 Australia

Telephone: (03) 9626 7000
Fax: (03) 9626 7999

AR-012-516
December 2002

APPROVED FOR PUBLIC RELEASE

Contents

1. Introduction

While the term Health and Usage Monitoring System (HUMS) is comparatively new, the concept of monitoring the health of structural and mechanical components has been around since the industrial revolution and the concept of usage monitoring for about half a century. Apart from the development of new fault identification algorithms, what is new is the last word in the term - System. The combination of the data collection, data verification, health trending and sometimes usage calculation into an integrated system (partly on-aircraft and partly off-aircraft) only originated in the 1980s and is still developing.

The most restrictive definitions of the term HUMS includes only those units which utilise the space, computing power and data associated with the Flight Data Recorder (FDR) on helicopters. These are the so-called "North Sea HUMS" since they originated on helicopters servicing the North Sea oil platforms.

The same concept can however be extended to jet aircraft where the unit is often quite separate from the FDR and more recently to other vehicles such as ships, transport trucks and military land vehicles. The HUMS work in DSTO Air Vehicles Division is primarily focussed on monitoring systems for military aircraft (including helicopters and Unmanned Air Vehicles (UAVs)) and for gas turbine engines. However, the HUMS developments in other vehicles and other related developments (such as in data management and certification) are of interest to us in terms of the feedback into the aircraft and engine areas.

The Third DSTO International Conference on Health and Usage Monitoring Systems, HUMS2003, was scheduled for Melbourne in February 2003 to coincide with the Australian International Airshow. The scope of papers covers the wider area noted above and includes both formal and informal papers. This document includes only those of the formal papers which the authors made available for publication during December 2002.

Other papers will be available from the conference CD-ROM and/or the conference website at URL: http://www.dsto.defence.gov.au/corporate/conferences/hums.

2. Acknowledgments

The HUMS2003 Conference has been organised by a committee comprising:-

Graham Forsyth	Convenor
FLTLT Adam Grey	Representing RAAF Williams
David Blunt	Machine Dynamics
Ben Parmington	Applied Combustion
Joanna Kappas	Mechanical Integrity
George Karvounis	Engine Performance
Eric Lee	Helicopter Life Assessment
Loris Molent	Representing Structural Integrity
Nik Rajic	Representing Smart Structures

As well, Defence Corporate Communications have managed our website and produced publications (including graphics) and Lisa Torrance has helped with financial management.

Contractural arrangements have been entered into with a number of firms including International Conference Management Services (ICMS), Duxton Hotel, Melbourne Aquarium and Aerospace Australia Limited (trading as Airshows Downunder).

3. Conference Sponsors

Sponsors are a critical part of any conference, in fact running a conference would be even more difficult without them since they provide the initial funding necessary to get started as well as contributions to papers and attendance. We are fortunate to have a list of valued sponsors in three categories, shown alphabetically within each category:-

GOLD

DEFENCE SCIENCE & TECHNOLOGY

DSTO
Key Initiative in Smart Structures and Materials
http://www.dsto.defence.gov.au

SILVER

ETMC technologies

ETMC Technologies
http://www.ETMC.com.au

Honeywell

Honeywell
http://www.honeywell.com

QinetiQ

QinetiQ
http://www.QinetiQ.com

semia

Our purpose : To measure and to control vibrations

Semia
http://www.semia.fr

smiths

Smiths
Grand Rapids, Mn USA
http://www.smiths-aerospace.com

VAMtec
VHUMS & ASSET MANAGEMENT TECHNOLOGIES
PARTNERS IN MONITORING, TRACKING & ANALYSIS

VAMtec *Consortium*
GPS Online (AUST) & Dytecna (UK)
http://www.vhums.com

BRONZE

Aerospace Monitoring & Systems, South Africa
http://www.ams.co.za

AMS

Helitune

Helitune Ltd
Torrington, Devon
UK
http://www.helitune.com

4. Papers Included in This Document

Formal papers will be published in this section as received from authors apart from any minor changes in layout necessary to meet printing requirements. Please note that individual or corporate authors may have retained copyright over these works. You should check with authors if you wish to make copies of any paper for any use other than "within your organisation for study or education purposes".

Not all papers presented at the conference appear in this document. Those which are included are included in the order received rather than as any indication of timetable placement, priority or preference.

Measuring the Performance of a HUM System - the Features that Count

Kenneth Pipe
Principal Consultant
Humaware

ABSTRACT©

To enable the operator to discriminate between wide price/performance variations in the range of systems marketed as HUMS requires that the systems' performance can be measured against recognised standards that relate to the HUMS primary function of enhancing helicopter airworthiness. This paper explores the criteria for measuring the performance of a HUMS, the features of the system that are important and the standards that can be applied.

Ultimately, the performance of a HUM system is measured actuarial terms i.e. the impact of the technology in reducing the number of catastrophic and hazardous events on helicopters. The Regulatory Authorities and Manufactures maintain appropriate statistical records on helicopter safety, but there is not a consensus on how to factor in the contribution of HUM to the overall safety improvements achieved since the introduction of the technology. Given this difficulty the paper discusses the other performance criteria that are meaningful in establishing a systems efficacy.

The components of a HUM system form a hierarchy, from the sensors, through processing, alarm generation and finally the operation of the system. The paper discusses how the components of a HUM system contribute to its overall performance and the practical performance criteria that can be applied. Health as well as usage functions are covered as well as functions that relate to maintenance benefits.

HUM systems are unusual in terms of the close interaction of the airborne and ground based components of the system. Because of this design feature maintaining the system's integrity when operating in the helicopter environment is a major issue, the paper outlines where standards need to be established.

Finally the airborne components of a HUMS have to perform as any other item of avionics equipment and the standards that are peculiar to HUM systems are discussed

INTRODUCTION

The first HUM system was certified in November 1991 for operation in the North Sea. The system was one of two competing designs developed for meeting the oil companies' requirements for HUM following the HARP revue of helicopter airworthiness. Both systems were designed to meet the same requirements but each operated with significantly different functionality and man prints. In the following ten years HUM systems for civil helicopters have appeared from Eurocopter and latterly Bell, again ostensibly designed to meet the same North Sea requirements but they are also distinctly different systems. Lately in the military arena Smiths and Goodrich have produced HUM systems which represent yet further variations in design.

There is guidance on what constitutes a HUM system published by the UK CAA Helicopter Health Monitoring Advisory Group[1], but this guidance has not led to a consistent system's design philosophy emerging nor does it establish criteria for the measurement of a HUM system's performance.

This paper aims to explore how the performance of these systems of differing design can be established and their relative attributes compared.

The purpose of a HUM system is to provide a timely indication of the deterioration of the continued airworthiness of a component in order that maintenance can intervene and rectify the defect. It is a supplementary method to the various prescribed inspections and preventative maintence actions contained in the helicopter's maintenance handbook to ensure the helicopter's continued airworthiness and thus enhance the margin of safety of helicopter operations. To achieve this HUM systems' monitor a) the usage of the lifed components b) any exceedances from the operational envelope and c) the health of the power train components. The health monitoring of components provides a goal keeper function, guarding against any failure of the maintenance procedures to preserve the airworthiness of the helicopter. The most important health monitoring feature is the vibration monitoring function which utilise techniques of varying degrees of sophistication to identify defects developing in the power train components.

There are the secondary functions in HUM which relate to obtaining Maintenance Credits. Principal amongst these is the balancing of shafts and rotors without employing specialised maintence test equipment and in the case of the main rotor eliminate the need for maintenance test flights following the rotor adjustments.

It is, though, the primary requirement of a HUM system to enhance airworthiness, which requires that the contribution of HUM to a helicopter's airworthiness can be measured.

Measuring Airworthiness

Airworthiness can be defined as the risk to an aircraft of a catastrophic incident posed by its design or operation. The design criterion set in the JAA's JAR 29.547(b) & 917(b) is that for any failure mode in the dynamic train the risk of it causing a catastrophic failure of the helicopter must be less than 1 catastrophic failure in 10^6 flight hours, which is a probabilistic measurement of a very rare event. The only practical estimates of this probability are from large fleets of aircraft operated over many years. Relating the helicopter's overall risk of failure to the performance of a particular HUM system function is not possible unless statistics from many billions of flight hours more than the fleet will fly can be gathered.

The Regulatory Authorities and the Accident Investigation bodies do look at the contribution, or otherwise, of HUM in the case of a catastrophic failure. Design assessments for certification of new helicopter design can assess the contribution of HUM to overall aircraft airworthiness. Neither of the above processes has produced criteria around which HUM performance standards can be determined.

The Hienrich Principle states that accidents have incidents as their precursors. Analysis of the causes of the much more frequent occurrence of incidents can be used to provide data that would allow measures of risk to be made. This allows performance measures to be established that will by inference ensure the much less frequent accident rate is impacted in the desired direction.
The Regulatory Authorities maintain records of mandatory occurrence reports and the CAA, which is the only Regulatory Authority to mandate HUM[2], does record its contribution to reported events. Therefore there is an experience base of HUM operation that can provide some actuarial evidence of performance. The CAA has published data on HUM performance in the North Sea[4] as summarised in Fig 1. This data was informally updated in 2002 and the success rate was basically the same[6]. This North Sea experience documented by the CAA shows that a success rate in detecting defects of 70%, this can be used as a performance measure for the contribution of HUM to airworthiness.

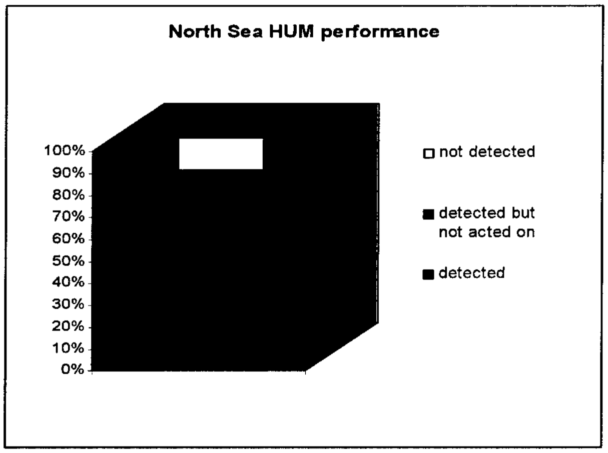

Figure 1

The distribution of faults detected by North Sea HUM systems is shown in Fig 2.

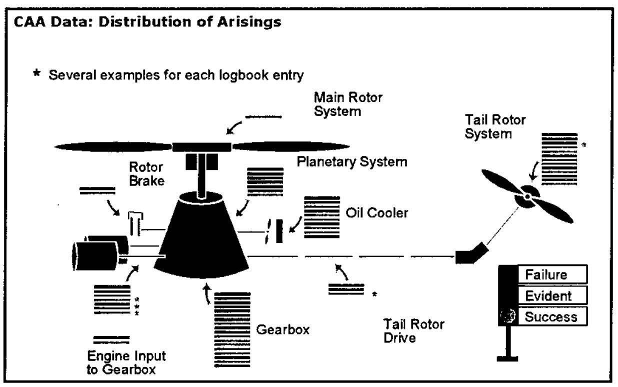

Figure 2#

The diagram shows that, with the exception of the planetary gear system, there is no justification for applying different performance standards for individual components of the dynamic train. Clearly components where the technology has not been shown to be effective should be excluded from the performance measure or a measure more applicable than the North Sea standard should be employed.

This actuarial measurement of performance is only viable for large fleets generating a large number of hours and where there is independent verification of the statistics. This is possible in aviation because of the monitoring of operations imposed by the Regulatory Authorities. It is unlikely that any vendor would accept such a performance criterion unless he is offering a mature system with a track record of success. This overall performance criterion does represent a differentiator for separating the wheat from the chaff in the market place

The corollary to the success rate is the false alarm rate. Specifying a demanding target for the success rate could always be met in principle, but with a high cost in nugatory maintence activity. Fear of this potential problem with HUM has been a major impediment to helicopter manufacturers incorporating HUM into their maintenance practices.

An acceptable definition of false alarm in the HUM context is contentious. For much of the industry any single alarm generated by a HUM system not shown to result in a defect being found is judged to be a false alarm, even though the alarm did not lead to any nugatory maintenance. This is a very demanding definition, particularly for alarms being set on random data such as those generated by vibration. It is usually found that in order to set alarm levels with sufficient sensitivity to fault conditions then some spurious alarms will be generated. If these are shown to be spurious and dealt with without incurring unnecessary maintenance actions, such as inspections, then this should be treated as part of the alarm processing, all be it processing by the engineer rather than the system. A better definition of a false alarm is an alarm that has caused nugatory maintenance action. Dealing with spurious alarms in an HUM environment can be regarded as an operating overhead. This overhead cost should be measured and the system should have features that ensure that it is kept to a minimum. The time spent at the ground station and skill levels for alarm processing needs to be specified by the vendor.

The inverse of the successes are the missed faults, it is an important measure also for controlling the performance of the system. Determining the reason for HUM failing to detect a fault is necessary for quality control of HUM and ensuring that any gaps in performance that are within the scope of the technology are closed

These actuarial measures of HUM performance can be monitored by any operator operating under the CAA's CAP 693 guidelines[3].

Any actuarial measurement of performance relies on ample evidence being available to generate the statistical measures at an acceptable confidence level. For a new HUM technology or a new helicopter design these data are not going to be available and other approaches have to be taken.

Building blocks of a HUM system

The approach to a new system has to be to identify the components of a HUM system and establish the performance criteria for these components such that an estimate of the expected overall performance can be made. Fig 3 shows the processes that have to be performed for an alarm to be produced by an HUM system.
Sensors with the correct characteristics have to be fitted, the signals are then acquired and processed to produce variables that are stored in a data base. There may be more than one database

that has to store without loss the variables and then the variables are processed to produce alarms which have to be managed until either a defect or a false alarm is determined.

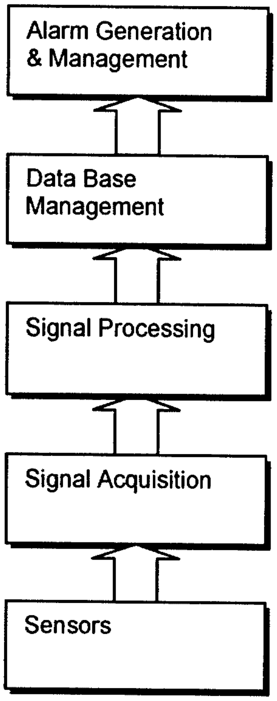

Figure 3

For the monitoring of data such as airspeed the HUM system has to emulate the processing used for acquiring and displaying the data in the cockpit, with application of the same performance measurement criteria. There can be difficulties here in that usually the HUM signal acquisition and processing is not identical to the cockpit instrumentation and can produce results of differing accuracy and resolution. Establishing what the tolerable difference between the HUM data and the cockpit data, which is usually less accurate and of a lower resolution, can be a more strenuous exercise than first envisaged. This problem can be further compounded when data is used in calculations such as those use to compute Power Assurance. These differences are often small but can cost the credibility of the system unless they are understood and accepted by pilots and engineers.

There are sensors and processing used in HUM for different purposes than those for which they are used in the aircraft. Tachometers are typical of these. They are usually pulse devises, integrated to produce shaft speed, for HUM they are also often used as a phase reference for vibration processing, requiring that the performance of the sensor in terms of producing a regular and well defined pulse shape is more tightly specified. In most HUM systems the processing of vibration data is highly dependent on reliable tachometer data requiring that the performance and integrity of this system component is very thoroughly engineered.

Then there are the specialised sensors and processing that are peculiar to HUM systems, these are principally the vibration processing for health monitoring of the shafts, gears and bearings and also the blade tracking systems combined with the vibration processing used for rotor track and balance. Vibration processing adds a great deal of cost and complexity to HUM system and needs to perform well in order to justify the burden it imposes.

Vibration Processing in HUM

Vibration has been shown to be a very powerful tool in fault finding in the helicopter's dynamic train[5]. Any deformation of a rotor or shaft will produce an increase in its vibration. Faults in gearing actions, cracks in splines and spalls in bearings all produce discernable changes in the vibration characteristics of the helicopter.

Shaft induced vibration has been shown to be relatively straight forward to deal with reliably by HUM systems and it produces a high probability of success at an acceptable false alarm rate. This is not so for the more complex analysis techniques where a more variable performance has been experienced in the field.

The difficulty with vibration monitoring is that one sensor integrates a large number of vibration signals from the components of the machinery it is monitoring. A typical vibration signature for a component on a helicopter transmission is shown in Fig 4.

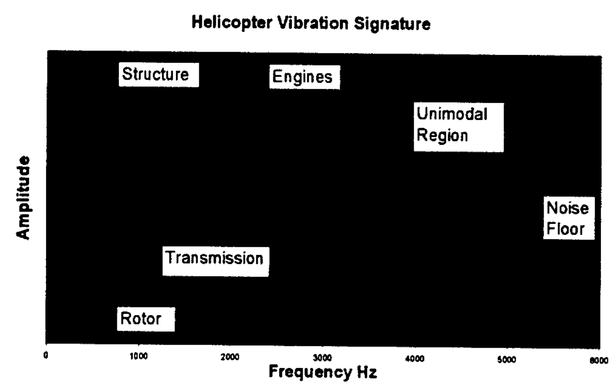

Figure 4

The vibration spectrum for a helicopter can be divided into six regions. The lowest frequencies are quite discrete and are primarily due to the helicopter's rotors and their harmonics. As frequency increases the transmission shafts, engine shafts and meshing frequencies come into play. All of these machinery generated vibrations cause the structure to vibrate and as the sensors are mounted on the structure the structural resonances or modes generate a range of tones throughout the vibration signature. All of these vibrations produce harmonics and the energy in the increasingly complex signature builds up into the characteristic bell shape as shown in the diagram, with the frequency discrimination increasingly becoming less distinct. After the primary shaft orders and major structural resonances frequencies have passed the spectrum is then made up only of the higher harmonics, the energy in the spectrum then falls off and any frequency discrimination disappears altogether. The lack of frequency discrimination leads to the term unimodal to describe this region. Finally no spectral energy is discernable and the energy levels off, this is the noise floor of the sensor and processing.

Each of these regions makes different demands on the vibration processing and different technologies have been produced to provide useful analysis of the data. This has impacts on the sensors as well as the processing capability of HUM. These will now be considered in turn.

Sensors

Performance of the sensors is clearly critical to the performance of the vibration processing in a HUM system. Vibration sensors originally measured position and were very restrictive in frequency range or bandwidth. Technology progressed to velocity sensors which offered higher bandwidths and tracking filters as shown in Fig 5 could be used to allow the analysis of higher frequencies such as engine spool vibration. It was the development of the accelerometer that brought sufficient bandwidth to cope with the whole of the useful spectrum.

Accelerometers do have a problem of roll off at DC as well as at high frequency. Specifying an accelerometer that can measure the low frequency rotor orders as well as the higher frequency harmonics of the mesh tones can prove to be a demanding requirement. Often accelerometers with different sensitivity and bandwidth have to be used on the helicopter, one for structural vibration monitoring giving the required performance at lower rotor and structural resonance frequencies and another type for the transmission and engine frequencies. Environmental conditions also play an important part in accelerometer specification.

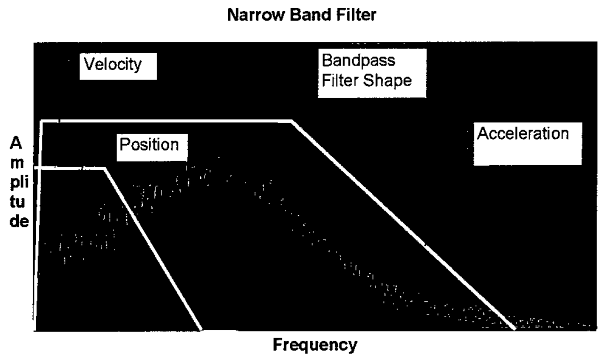

Figure 5

Accelerometers at the start of the HUM system development were largely laboratory devices but great strides have been made in their design and constriction to make them suitable for use in the aviation environment.

The siting and installation of an accelerometer is as critical as its bandwidth, range and sensitivity for determining its performance. If the sensor is acquiring the lower frequency rotor and structural vibration then the sensor positioning is less critical as these vibrations can be sensed almost anywhere in the helicopter. For the higher frequencies or for when the shaft frequency is the same sensor positioning is more critical to produce the necessary discrimination between components. The sensor has to be installed as close as possible to the component so that its vibration dominates the spectrum.

Great care has to be taken over the installation of sensors to ensure that the sensor orientation is appropriate the vibration that is being measured and that the installation bracket does not introduce structural responses that either swamp the signal or interfere with it.

Adequacy of the sensor poisoning can only be demonstrated by the production of spectra that show that the characteristics of the monitored component's vibration signal is clearly discriminated in its normal operation such as the once per rev vibration and harmonics from the shaft and the mesh tones in the gearing.

The problem comes in determining whether components that in normal operation do not produce a characteristic vibration can be monitored. Bearings and splines fall into the class of components that only produce vibration when there is a fault present. The only practical method to determine if the component can be discriminated is by trials using seeded 'marker' faults in the component. Only the helicopter manufacturer has the resources to perform this exercise, but it is expensive and is probably only practical to carry out during development programme's rig tests of components.

In the absence of this evidence what can be done to establish performance? For external bearing such as tail rotor drive bearings, placing the sensor on the bearing housing and checking the signature for the lack of other vibration apart from the supporting shaft is a reasonable approach. It is important that where evidence cannot be produced of the component's characteristic vibration being discerned, then alternate techniques should be available in the HUM system to monitor the component.

The best evidence that he correct sensors have been selected, positioned and installed correctly is to perform trials to establish that the vibration signals can unambiguously discriminate the features required to substantiate the claims made for the defect monitoring. One sensor type is unlikely to meet all monitoring requirements. Specifying the characteristics of the sensor independent of its installation is not sufficient. Positioning and installation of the sensor are much more critical to a HUM system's overall performance and this aspect of the sensors performance cannot be determined from the drawings.

Signal Acquisition

Unlike process variables such as pressure and temperature vibration does not have a single mode of operation to model to develop its performance criteria. Signal acquisition techniques for vibration again follow the development of the technology from RMS measurements of the overall signal energy, through tracking filters and finally the modern digital signal processing techniques. In terms of technology the digital acquisition of signals represents a complete break from the past and using digital techniques to emulate the older RMS and tracking filter techniques is not generally possible.

Current processor technology and very low cost memory means that signal processing is now entirely software driven and signal acquisition is a matter of anti-aliasing filtering, A/D conversion at a high frequency to a resolution of 16 bits or higher and then piping the data into memory. This approach can have performance consequences downstream. For example if a spectral resolution of 2Hz is required to discriminate the rotor frequencies and a maximum frequency of 50 KHz is required to acquire the higher frequencies and there are 16 sensors acquiring data once every 20 minutes then the amount of data accumulated per flight hour is in excess of 150 Mbytes. It is now not a problem to store or process this data, but it could take some time to download it from the aircraft.

If HUM data is going to be useful for identifying meaningful trends then most data needs to be gathered in steady state conditions or in specific regimes. For a reasonable acquisition rate to be achieved then specifying the regime recognition parameters that match the regimes the helicopter actually flies is important. This is a more difficult issue than it appears as there is always a conflict between reducing the range of the regime to minimise the variation in the data, and increasing the range so that it does not restrict the time for data acquisition and hence the number of acquisitions.

How many acquisitions are required per flight hour is the key parameter to specify for data acquisition, not the number of channels. The acquisition rate needs to be sufficient to allow for the time the helicopter spends in manoeuvre. A single channel of acquisition can mean that for the number of sensors fitted in HUM applications the number of acquisitions per flight hour may not be sufficient to generate meaningful trends.

Quality and quantity of data are the features that drive the data acquisition requirements in HUM, and these are the items that need to be specified with clear performance criteria that relate to trending and other uses of the data.

Signal Processing

There is no single signal processing technology that is effective for all types of vibration analysis. The FFT based spectrum analysis techniques that are effective at low frequencies where resonances are pronounced are ineffective in the mid frequencies where the spectral energy is dense and complex in structure and they are totally ineffective in the unimodal region where there is no tonal energy at all. In these cases more sophisticated types of processing have to be employed. Key to the performance of a HUM system is that appropriate vibration analysis techniques are employed for the range of fault detection capability claimed.

All processing is now digital using the FFT in varying degrees of sophistication on the data piped to memory from the acquisition process. There is a problem in FFT processing of matching frequency resolution with frequency range as shown in Fig 5. It shows the number of spectral lines in each decade of frequency range for an analysis with a minimum frequency of 5Hz and a maximum of 10 KHz, plotted on logarithmic scale. In the lower frequencies which contain most of the spectral data there are the least number of lines, 2 in the 1-10 Hz range and 18 in the 10 - 100 Hz range. In the unimodal region where there is no spectral information there are 1800 lines. FFT processing covering a wide frequency range will not produce the frequency resolution where it is required. There are 'zoom' FFT processors that can locally increase resolution in the frequency band and this processing is essential if spectrum analysis is going to be used over a wide frequency range.

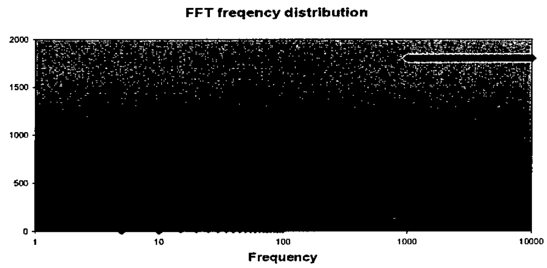

Figure 6

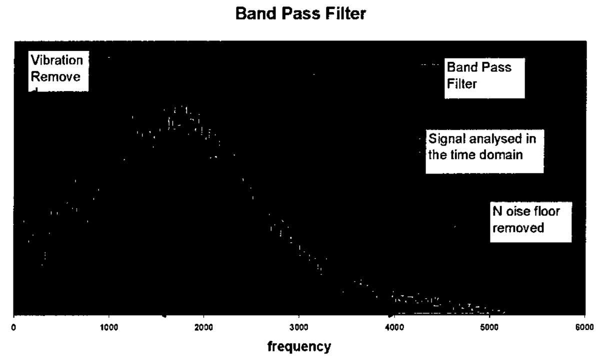

Figure 7

The most powerful technique developed for frequency analysis in the complex environment of the helicopter is Synchronous Signal, or Time Domain, Averaging. This produces spectra where the resolution is a multiple of the period of one cycle of revolution of the shaft, and the frequencies are a function of the shaft speed. This separates the contributions of each shaft to the overall signal, and automatically controls the FFT processings resolution and bandwidth conflict. HUM vibration processing without Signal Averaging is unlikely to be effective for the complex signals generated by the main transmission. This processing is more complex than asynchronous spectrum analysis and a system has to be specifically designed to accommodate it. It cannot be implemented post acquisition without utilising very sophisticated signal processing techniques. Signal Averaging represents a cost driver in HUM system design but it is an essential feature of the system if gear box analysis is to be attempted.

The highest frequencies of the spectrum contain the vibration that relates to fatigue cracking and to bearings pitting or seizing. To analyse this component of the signal the lower frequency information has to be removed. This is normally achieved by band pass filtering as shown in Fig 6. The filter removes both the lower frequency vibration and the noise floor. The energy of the resulting signal is measured or it is subjected to statistical techniques such as Kutorisis estimation to identify the impulsive characteristics of the signal that are related to faults. These techniques are sufficient where a single component's vibration dominates the signal. There is a number of proprietary techniques[5] utilising pattern recognition that are necessary in the more complex areas of the transmission.

Establishing clearly the functionality of the vibration processing is essential in establishing its performance, too many vendors try to hide behind inadequate descriptors such as 'advanced' or 'state of the art'. The standard performance criteria for vibration processing of bandwidth, frequency and amplitude resolution are meaningful for the basic spectral measurements of low frequency features, but are not meaningful for the more complex analysis techniques frequently used in HUM. The criteria for acceptable performance of the vibration processing are that the resulting spectra and other indicators clearly demonstrate that each of the shafts can be discerned from the background signal. For gearboxes where signal averaging is implemented the shaft frequency and mesh tone (dual mesh tone for epicyclics) should be clearly discerned. For systems claiming bearing, fatigue and similar defects then the signal in the unimodal range should be free of tones and well above the noise floor. It is also important to demonstrate insensitivity of the data to the normal small variations in shaft speed, weight and centre of gravity that can be experienced during and between flights

It should then be possible for any HUM system to establish a matrix of which components are monitored and for what type of fault, clearly delineating the signal processing techniques employed to accomplish the task as a basis for establishing performance criteria for the processing.

Data Base Management

HUM systems have to store large amounts of data per flight which over time build up into huge amounts of data. Very little of the data is ever reviewed. Avoiding data corruption and maintaining reasonable access times is not a straight forward design issue. The problems are compounded if the data is to be accessed by a number of users over a network.

The first process in managing the data is the post-flight downloading of the data from the aircraft and in some cases uploading configuration data pre-flight. Clearly, performing these tasks should not interfere with the normal operation of the helicopter. Specifying the routine download times and pilot/engineer interaction with the system is critical. Nothing is more fatal to the acceptance of HUM in the field than the system being perceived as obstructing the normal efficiency of flight operations.

At the ground station there is a similar issues of the acceptable latency in transferring the data from the download device and the subsequent updating of the data base. If care is not taken in the design then after a few hundred flight hours of operation these latencies can become excessive. It is essential that the download times and latency in updating the database have limits specified. Data from a number of aircraft being downloaded simultaneously and non-sequentially can cause chaos. Constipation is a real problem in HUM ground station operation. Flexibility of download modes is as important as latency if the system is to gain acceptance at the flight line. Latency in retrieving data, in particular data for post-flight pilot review, also needs specifying. It is important that these latencies are specified for mature databases containing hundreds, not tens, of flights for a representative number of aircraft being supported by the ground station.

Failure to put practical hard limits on these issues has resulted in HUM systems being designated as unusable. Software engineers tend to concentrate on functionality and make light of these performance issues. They must have performance standards to achieve if setting to work and operating a HUM system is to be a tolerable experience for the operator.

Database integrity is a difficult feature to specify but if the data is to be trusted then integrity is essential. Error detection and correction features need to be included in the specifications, particularly for data bases which are moved with the aircraft. Data loss need to be very rare if confidence is to be achieved in the data if its eventual application is to provide Usage Credits.

Alarm generation and management

Most HUM systems use a threshold method to generated alarms. In setting a threshold level there is always the pressure to set the thresholds conservatively to reduce false alarms resulting in the threshold being set too high to be sensitive to a developing fault as shown in Fig 8. As the occurrence of faults is very rare then there is a lack of a counter pressure to balance the pressure from engineers to minimise the false alarm rate.

To guard against this a target acceptable false alarm rate could be set across the fleet. If the observed rate is less than this target then the thresholds should be reviewed to ensure that fault sensitivity is not being sacrificed. If it is more than the target then the suitability as a HUM indicator of the parameter needs to be reviewed and if possible an alternate more stable parameter used.

Vibration signals and noise on other signal types always produces 'wild' points in the data. No system can maintain sensible threshold levels without an appropriate form of wild point rejection in the processing. All of the wild points in Fig 7 exceeding the lower more sensitive threshold level would be removed by this processing and detection of the trend would remain unaffected.

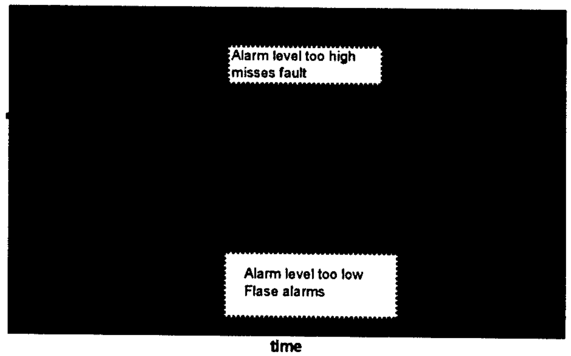

Figure 8

The use of fleet statistics to set alarm levels can also have an impact on performance. In general the higher the frequency of analysis the less likely that a group threshold is going to be valid. If in the group of helicopters the differences in the statistical means and variances of the data monitored are small, then the group threshold level is likely to perform well for all. If as in Fig 8 the means are well separated or one data trend has a dominant variance then a threshold set by the group statistics is only going to be valid for either the data with the highest mean or largest variance.

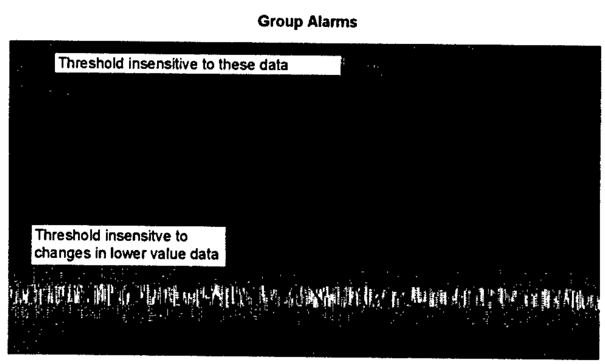

Figure 9

Facilities to corroborate data are also essential to assist in determining whether an alarm is a result of a faulty sensor or just noise. It is rare in vibration that a shaft vibrating abnormally will be discerned on just one sensor.

Sensitivity to faults can only be determined if the HUM has actually identified any. Establishing the range that is known to be sensitive to detecting a fault is important for assurance that the threshold is set appropriately. These data is only likely to be available for mature systems with sufficient flight hours behind them to have developed a track record of success.

Alarm management in HUM as a proactive task the goes beyond the processing of checking and following through on any actions caused by an alarm being transgressed. In order to keep the system adequately sensitive the alarm levels they need to be subject to management scrutiny. Changes in weather, pilot, weight and centre of gravity of the helicopter all change the characteristics of the HUM data and can compromise some of the thresholds levels. Monitoring the performance of the alarms to determine the validity of all of the thresholds being applied is critical to maximising the overall performance of the system.

The key performance measure for managing the threshold setting is the false alarm rate. Too few alarms and there is a problem with fault detection sensitivity, two many alarms and there is a problem with suitability of the data. Also the probability of success is only maintained if the arisings that are identified by other means are investigated to determine if the fault was present in the HUM data and why the alarm threshold was not triggered.

Maintenance Monitoring

Vibration levels on shafts and performance monitoring of the engines are common maintenance functions where a HUM system replaces the existing maintenance test equipment. It needs to be ascertained that the HUM system really does replace the test equipment. For vibration some test equipment utilises older analogue vibration measurements such as RMS or swept band filtering. It is not possible without a great deal of effort to reproduce these data with the digital signal processing techniques found in HUMS. Establishing the correspondence between handbook vibration limits and HUM data is a key performance attribute for a system and this correspondence can usually only be established by the Design Authority.

The ability of HUM systems to routinely monitor the rotor and produce recommendations for adjustment that keep the rotor in smoothest possible condition is an important function. Rotor track and balance systems usually contain proprietary software for determining the rotor adjustments. Specifying performance is not possible without the disclosure of the processing. Accuracy of sensors is not a guide to the overall performance of a rotor smoothing system. For rotor smoothing the only effective check is to mal-adjust a smooth rotor, well within its safety limits, and check the track and balance system predicts the correct adjustment to make to return the rotor to the smooth condition. A correctly performing system should not confuse a weight adjustment with a track rod or a tab adjustment.

The key performance features here are a) the average vibration level achieved and b) the reduction of effort required to tune the rotor. It is not unreasonable to require that the rotor vibration in all axes and flight conditions is maintained to be less than specified percentage of the handbook limits for 95% of the time. This has to be on the condition that the adjustment recommendations produced by the HUM system are implemented. Also this measure should exclude the adjustments following blade removal where a specification should be to bring the system to below the target vibration level in less than two test flights.

Equipment Standards

HUM systems are peculiar in three respects when compared to other avionics, a) the system's complexity, b) the split of functionality between airborne equipment and ground equipment and c) the interaction of the system with the helicopter's design and operation.

If the equipment is specified as part of the helicopter's MEL then many, but not all, of the large number of sensors that make up a HUM system greatly reduces its overall reliability. Long MTBF's are essential on individual equipments in order that the overall system MTBF is reasonable, or the MTTR's are very short and adequate spares are provisioned so that the MEL requirements can be met.

The split of monitoring functions between the airborne and ground components of the system can cause difficulty in certifying the system even if the certification standard is No Hazard/No Benefit. The FAA's HUMS AC[7] requires different software standards to be met for the ground station than for the airborne because of the use of COTS software in most ground station applications. For HUM systems other than No Hazard/No Benefit the ground station software has to be to a higher software standard than the airborne, or independence of the HUM processing to the COTS software has to be demonstrated. Further hurdles have to be cleared of independent verification of data and mitigation of the credit for credits to be possible. This level of functionality is beyond the scope of current HUM system design.

Where maintenance credits are to be sought it usually requires that the Design Authority approves the changes to procedures. Only the Design Authority has the necessary knowledge to determine if an anomaly detected by a HUM is damaging to the extent of risking the airworthiness of the helicopter. This requires that the airframer acknowledges and supports the data analysis produced by HUM. Design Authority approval, or at the very least the HUM has non-objection status, is an essential attribute of any HUM system.

Conclusions

Divining performance measures that relate directly to the purpose of HUM requires that we can assess the contribution of the technology to helicopter airworthiness. The nearest approximation to this is achieved by building up a statistical knowledge of more frequent events such as maintenance arisings and other events reported to the Regulatory Authorities. The monitoring performed by the CAA shows that generally a success rate of 70% in defect detection is achievable, and would provide an overall performance standard for HUM

Measuring the performance of the HUM system in service is most easily achieved by monitoring the false alarm rate. This should be randomly distributed amongst the components of the dynamic train and within a range that is acceptable to the operator. It is not the case for HUM that no news is good news.

For new HUM systems type or a system that is to be fitted to a new aircraft deign then the actuarial data for determining the performance of a HUM does not exist. The performance of the individual building blocks of the system has to be specified such that confidence in the overall system performance will be adequate. For flight data then the normal system engineering practices of modelling the process so that the required performance of each component of the system can be determined will suffice. But the major component of a HUM system is the vibration processing and it cannot be dealt with in such a straight forward manner.

Measuring the performance of a modern vibration analysis system based on digital signal processing is essentially a process of the vendor demonstrating that the features used for monitoring purposes can be clearly discerned in the measured spectra of the normal operation of the helicopter, or at the very least the data contains no phenomena that prevent the monitoring function from being effective. The features that count if the performance of the system is to be unequivocally determined are a) a clear statement of the components that are to be monitored from the sensors b) the techniques to be used for the monitoring c) the acquisition rate d) the positioning of the accelerometers.

The suitability of sensors and processing in terms of resolution, range and bandwidth needs to match the monitoring techniques used. A clear statement of these attributes and their specific purpose in terms of the fault modes monitored in the helicopter's dynamic train components should be produced with evidence that the processing can discern the relevant attributes. Without this any attempt to substantiate the efficacy of the processing cannot be made in any quantitative way.

Measurement and processing are not sufficient in themselves to determine the likely performance of a HUM system. The alarm processing and management functions of the system need to flexible enough that the threshold can be tuned to achieve the desired false alarm rate and maximise the systems sensitivity to faults. The features that count in the alarm system are the availability of management functions for setting and assessing the threshold levels so that the sensitivity to faults is optimised with respect to the acceptable false alarm rate and the cost of managing the spurious alarms by the engineer. A system for removing singleton alarms is essential is viable threshold levels are going to be set.

A HUM is not a fit and forget system, management of the alarm threshold system is essential if the expected improvement in safety margin is to be achieved. For convenience it is all too easy for alarms levels to be too conservatively set making the system insensitive to faults. The other performance check to determine that the thresholds are sensitive to faults is to investigate arising that are not indicated by the HUM and identify whether the threshold is set at a suitable level or on the most appropriate data.

The systems operation has to be acceptable to both pilots and line engineers, and their management. Latency and integrity of the operation of ground station and its data bases are the key features here. These need to be unambiguously specified and performance should be measured on mature databases.

Finally care need to be taken that the system can support the MEL requirements and that the software and HUMS AC requirements are met if maintenance credits are to be sought.

In short do not take anything for granted, ensure that the system is adequately and clearly specified and that acceptance testing includes performance measures that relate directly to the purpose of the HUM system and that the man print of HUM dos not interfere with operational efficiency. As well as a specification that calls out meaningful performance measures the contract has to be sufficiently robust to ensure they are met and maintained.

Glossary

AC	Advisory Circular
A/D	Analogue to Digital
CAA	Civil Aviation Authority (UK)
COTS	Commercial Off The Shelf
FAA	Federal Aviation Authority (US)
FFT	Fast Fourier Transform
HARP	Helicopter Airworthiness Review Panel
HUM	Health and Usage Monitoring
JAA	Joint Airworthiness Authority (Europe)
MEL	Minimum Equipment List
MTBF	Mean Time Between Failure
MTTR	Mean Time To Repair
RMS	Root Mean Square

References

1-'A guide to Health Monitoring in Helicopters', compiled by the United Kingdom Helicopter Health Monitoring Advisory Group, published by the UK CAA, Sept 1990

2 - AAD 001-05-99, May 1999, CAA London

3 - 'CAP 693 - Acceptable Means of Compliance, Helicopter Health Monitoring CAA AAD 001-05-99', May 1999, CAA London

4 - 'The Way Forward on Helicopter Safety' by D.J. Chapman. Presentation to the European Helicopter Operators Committee, May 1996, Estoril, Portugal.

5 - 'Helicopter HUM/FDR: Benefits and Developments' by Brian Larder, Stewart Hughes Ltd, Presented at the American Helicopter Society 55th annual convention, May 1999, Quebec.

6 - Report by Andy Evans, Minuted in 35th HHMAG, June 2002, Gatwick UK

7 - 'Airworthiness Approval of Rotorcraft Health Usage Monitoring Systems (HUMS), FAA Advisory Circular Material to be incorporated in AC-27-1 AND AC-29-2', July 15, 1999

Acknowledgments

Reproduced with the permission of the UK Civil Aviation Authority.

Collins Class Submarine Systems Analysis through Data Mining

Anthony Woolley
Maritime Platforms Division, DSTO

ABSTRACT©

System health and usage monitoring aboard the Collins Class submarines plays an important role in maintaining the operational availability of the submarine systems and the asset as a whole. This is partly achieved through the use of the Integrated Ship Control Management and Monitoring System (ISCMMS). ISCMMS provides the means for submarine system control and monitoring, with information relating to submarine system states being provided in "real time" to the operators, and recorded at regular intervals, for future analysis. The recorded data includes cooling pressures and temperatures, fuel pressures, scavenging air temperatures, periscope or diesel engine usage and submarine depths and speeds. Employing a data analysis methodology known as Knowledge Discovery in Databases (KDD) - also commonly known as "Data Mining" - the ISCMMS data is being analysed to identify new information and relationships that describe the operational aspect of the submarine systems. This information could then be used in predictive condition-monitoring systems, fault diagnosis tools and in the optimisation of the submarine operational profile. The data analysis will also aid in the development of the historic mission profile for each submarine and help identify improved data recording practices. Potential outcomes of the ISCMMS data analysis include optimised maintenance scheduling, monetary savings and an increased submarine operational availability. In this paper, an overview of some of the more common KDD techniques will be described, along with their application to various aspects of submarine system data analysis.

1. INTRODUCTION

Control of the systems of a Collins Class submarine is managed through the Integrated Ship Control Management and Monitoring System (ISCMMS). The ISCMMS backbone, the bus, is an ethernet local area network that allows for submarine control and monitoring to be performed from distributed workstations throughout the submarine. These workstations, known as General Management Units (GMUs) and Local Positions of Operation (LPOs), provide the means for the crew to monitor the state of the majority of the submarine systems, excluding the combat and weapons system and navigation. A basic ISCMMS systems block diagram is shown in Figure 1. Special Purpose Units (SPUs) provide the interface between the ISCMMS network and the submarine systems by passing information back to the operator, and also processing commands issued by the operator or issued automatically by the ISCMMS software. The GMUs are the primary operator positions, while the LPOs are generally a secondary control and monitoring position for all the equipment interfaced through a host SPU.

Since there are many more submarine systems than SPUs, system control is divided between the SPUs. Some systems, such as each of the three diesel engines, have dedicated SPUs but the majority of the SPUs are multi-functional. That is, they provide the interfaces for systems including air conditioning, power distribution, masts, bilge water, gas and fire detection and weight and trim compensation. The multifunction SPUs are denoted by the auxiliary (AUX) SPUs in Figure 1.

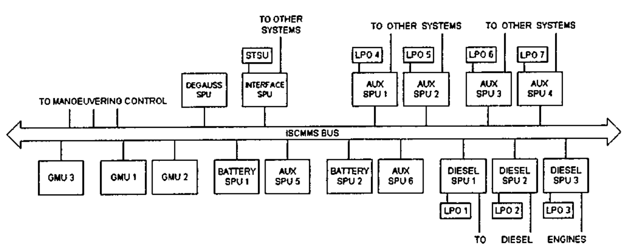

Figure 1. Simplified ISCMMS block diagram.

The actual monitoring of the submarine systems is achieved through the use of analog and digital sensors. Analog sensor data provides information such as temperatures, pressures, depth and speed. The digital sensor data provides an indication of the functional state of a component or system. That is, the digital sensor data indicates whether a system is running or stopped, open or closed or up or down. At regular intervals the sensor data is recorded in the Short Term Storage Unit (STSU) and is later transferred to the Ships Information System (SIS). The SIS is the sea going version of the Ships Information Management System (SIMS), which is a logistics database used in the management of the Collins Class submarines. When a submarine returns to port, the SIS data (including ISCMMS data) is transferred to SIMS.

At the time of writing this paper, there are five operational Collins Class submarines within the Royal Australian Navy (RAN) submarine squadron with a sixth soon to be accepted into RAN service. The first submarine, HMAS Collins, was accepted into service in 1996 followed in succession by HMAS Farncomb, HMAS Waller, HMAS Dechaineux and HMAS Sheean. The sixth submarine, yet to be accepted into RAN service, is NUSHIP Rankin. From the time each submarine was accepted into service ISCMMS data has been recorded on an ongoing basis. In an attempt to identify submarine system trends and/or dependencies, a methodology known as Knowledge Discovery in Databases (KDD) is being employed.

KDD arose from the need for fast, minimally supervised data analysis of large databases. The primary goal is the extraction of useful, previously unknown information describing the data or the processes from which it came. Fayyad *et. al.* [1], defined KDD as:

...the non-trivial process of identifying valid, novel, potentially useful and ultimately understandable patterns in data.

The KDD methodology has successfully been applied in commercial industry where standard statistical techniques may not prove useful due to the analysis tasks being cost and/or time prohibitive or impossible to perform. The KDD process evolved to perform the function that could not be solved by standard statistical analysis.

In this paper, an overview of the more common KDD techniques is given, followed by an overview of ISCMMS data logging and condition-monitoring. Examples of KDD analysis relating to submarine systems are then presented, including sample results. Current research has already identified a need for improved data capture techniques and condition-monitoring and these will also be discussed.

2. THE KDD METHODOLOGY

As mentioned, the KDD methodology arose from the need for fast and minimally supervised data analysis of large databases. The primary goal is the extraction of useful, previously unknown information describing the data or the processes from which it came. The methodology initially had its roots in a process known as "data mining", which consists of many different techniques used for the extraction of information from large databases. As data mining evolved, the process was formalised, resulting in what is now referred to as KDD.

It is common for the KDD and data mining terms to be used interchangeably to mean the same thing, however, data mining is one particular step in the KDD process as shown in Figure 2 [2].

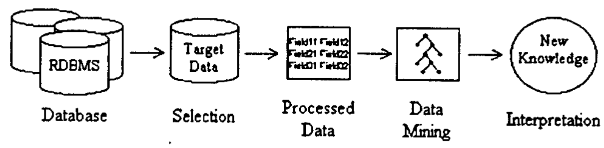

Figure 2. The knowledge discovery process, beginning with the database and ending with the discovery of potentially new information describing the data.

Prior to carrying out any work on the data itself, the analyst must have an understanding of the organisation collecting the data and the business practices employed. For the remainder of this paper, the term "system" will be utilised to describe the processes from which the data is collected. Therefore, a system may be some form of mechanical machinery or the processes involved in inventory control or an organisation collecting data relating to it clients.

After obtaining an understanding of the system processes, the database design must be studied by examining the database table data structures. Data is then extracted from the database and this is followed by data cleansing and transformation. That is, the data is pre-processed and this may involve such tasks as the removal of redundant fields, the removal of records containing missing values or the extrapolation of data to fill in for missing values and/or the removal of outliers. The data may also be subjected to various functional transformations such as normalising or functions combining several fields in order to generate new data for analysis. Once the pre-processing phase is complete, the data is ready for mining.

The algorithms used in the data mining phase will depend on the type of information being sought and some example data mining techniques are presented in Section 2.1.

Once information has been "discovered", it is then necessary to interpret its meaning and usefulness and this invariably includes discussing the results with relevant domain experts. Any new and relevant knowledge can then be used in such tools as expert systems for aiding in the decision making process; or in the development of condition-monitoring tools for monitoring system health and predicting future events that may lead to system failures; or it can be used to help optimise the performance of a system, whether it is a diesel engine or the arrangement of products in a supermarket layout.

The KDD methodology has many examples of successful applications in commercial industry, such as detecting various forms of fraud [3, 4]; increasing sales and improving customer satisfaction [5]; oil sample analysis [6]; medicine [7]; forecasting of inventory control [8]; and correlating alarms in a telecommunications alarm database [9]. Even though these applications are greatly different, it is possible to apply the underlying principles and ideas to a variety of tasks, such as in the analysis of Collins Class submarine system data. By applying the KDD methodology the analysis will delve into the ISCMMS data in greater detail than can otherwise be performed.

2.1 DATA MINING TOOLS

The KDD process is primarily used to achieve the following goals [10]:

1. To verify a hypothesis put forward by the analyst. This approach does not involve the discovery of new information but instead, with each iteration, a new set of questions are put forward in order to refine the search until a user specified limit is reached.

2. To discover new information that describes the data. This is (generally) an automatic process that searches for patterns, trends and generalisations. The two aims of this approach are (i) to try and predict future values of specific variables in a system; or (ii) to describe the data in a way that presents useful information about the system that was not previously known.

A number of data mining techniques exist to achieve these goals but only the more common techniques will be described here.

2.1.1 ASSOCIATION RULES

An association rule is a descriptive pattern of the form "X⇒Y", where X and Y are both instances in a database. The symbol "⇒" is an implication and so the pattern "X⇒Y" can be read as "X implies Y", meaning that if X occurs then Y is also likely to occur. The idea is to scan through the data for re-occurring patterns that occur often enough to be of interest. Associated with each rule is a confidence and support level. The rule has confidence, c, if c% of the X entries in the database have Y instances associated with them. The rule has support, s, if s% of all the entries contain both X and Y. That is, out of the X related items, the probability of an associated entry being Y is c%. Out of all the entries, the probability of X and Y occurring together is s%.

2.1.2 CLUSTERING

Clustering is another form of descriptive data mining. In the clustering process, the data is segmented into a finite set of clusters (or groups or classes) so that instances within each cluster have similar properties. Clustering is primarily an unsupervised task in which the clustering algorithm identifies distinguishing characteristics in the data set.

2.1.3 CLASSIFICATION

The predictive task of classification commences with a pre-classified set of data that is then presented to a classification algorithm. That is, the data has already been divided into classes, either through automatic means or performed manually by a domain expert. The algorithm learns to distinguish between each of the classes the data represents so that rules defining the data classes can be generated. These rules would then be available for use in classifying future, unclassified data.

2.1.4 DEPENDENCY ANALYSIS

Dependency analysis is used to determine those system parameters that are important in predicting the behaviour of another system parameter. The algorithm provides the accuracy with which each parameter can predict the system parameter under investigation and upon choosing specific parameters a cumulative accuracy is then determined. The results can be used in a predictive modelling tool to test the effects of changes made to those parameters and so help in the optimisation of system performance. For example, it may be necessary to predict the output temperatures of diesel engine cylinders. If it is known that the cylinder heads crack (that is, fail) under certain conditions, the model can be used to adjust the parameters found to have an important effect on the cylinders. This means the optimal operating conditions can be determined without having an adverse effect on the actual diesel engine.

Examples of these data mining techniques as applied in the analysis of Collins Class submarine system data are presented in Section 4.

3. ISCMMS

ISCMMS is the control and management system through which the submarine crew controls the majority of the submarine operations. At regular intervals, data relating to the current submarine system states are recorded in one of several logs. A brief description of each log follows:

1. Low Rate Data Log - records data at frequent intervals with the data being provided by analog sensors. The type of information recorded includes pressures, depth, speed and temperatures. For example:
   a. Diesel Engine 1 Cylinder Exhaust Temperature
   b. Main Generator Room Bilge Level
   c. Induction Mast Height

2. Event Log - records data when a digital sensor registers a value change from one defined state to another, for example: on/off.

3. Machinery Runtime Log - records the starting and stopping events of machinery as monitored by a digital sensor.

4. Alert Log - records the alerts raised by a system when operational parameters exceed predefined thresholds.

5. Trace Log - records high frequency data relating to a specific system. The recording is activated/deactivated when an operator specified alert is raised.

6. Weapons Data Log - records information received from the weapons system launch controller relating to weapons discharges.

Built into ISCMMS is a limited form of automatic condition-monitoring. For example, alerts are raised when various conditions are met or exceeded. These conditions are predefined by the supplier of the system and/or determined through operational use. At various stages of system operations it is known that some alerts will be raised by the systems. For example, during diesel engine start-up a low lube oil pressure alert may be raised. The alert is generally not due to system malfunction but due to the system not having reached its operational speeds. In these situations, some alerts are "blocked" so that the submarine's crew is not overwhelmed with unnecessary information. However, there is no predictive condition-monitoring based on the current and historic state of the system, nor is there any health monitoring based on system usage.

The aim of the KDD submarine system research is to determine the possibility of developing predictive condition-monitoring tools for use in the Collins Class submarines. The tools should also aid in fault diagnosis and further logistics engineering analysis, such as maintenance scheduling and the stocking of stores.

4. ISCMMS KDD ANALYSIS

The KDD methodology is currently being used to analyse submarine diesel engine performance and periscope and mast usage. The following sub-sections provide an overview of the research being carried out.

4.1 DATA CLEANSING

As indicated in Section 2, the KDD process may require some data pre-processing. In the case of submarine data mining, the data requires some form of sequential sorting and the removal of unnecessary information. The unnecessary information in this case was primarily duplicate data. The data was also segregated according to the submarine from which it came and whether the submarine was alongside (that is, tied up in port) or at sea. The pre-processing was achieved by laboriously "combing" through the data by hand and also through the use of custom developed software. Further pre-processing was dependent on the data mining techniques employed.

4.2 DIESEL ENGINE ALERT RELATIONSHIPS

The Collins Class submarines are each fitted with three Hedemora V18B14Sub series diesel engines arranged side by side and located towards the aft of the vessel, as shown in Figure 3. The primary purpose of the diesel engine is to charge the submarine battery, which then allows the submarine to remain submerged. However, the diesel engines are only operated when the submarine is on the surface or at periscope depth. Whenever a diesel engine operating parameter exceeds a pre-defined threshold, an alert is raised to inform the operator of the current state of the engine. The alert is recorded in the alert log, in conjunction with the date and time of the occurrence.

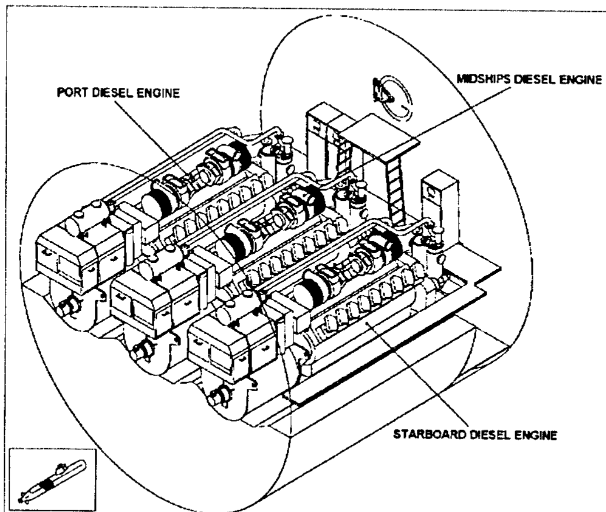

Figure 3. The typical diesel engine layout in a COLLINS Class submarine.

The aim of the diesel engine alert relationship study was to determine the existence of any alert sequences that may be useful in fault diagnosis and condition-monitoring and for the reduction of operator information overload. The discovered alert sequences were presented in the form of association rules.

A condition placed on the analysis required that the various alerts had to be raised within specified time frames. For example, it may have been of interest to find all the alert sequences for alerts raised within 10 minutes of each other.

An example of the type of discovered alert association rule is shown in equation (2).

$$(\text{LOW LUBE OIL PRESSURE}) \Rightarrow (\text{HIGH FRESH WATER TEMPERATURE}) \quad (2)$$
$$\text{confidence} = 15{\cdot}7\%$$
$$\text{support} = 10{\cdot}3\%$$

Note: This association rule is for example purposes only and not a real output.

The association rule of equation (2) states that if an alert is raised relating to "LOW LUBE OIL PRESSURE" then there is a probability of a "HIGH FRESH WATER TEMPERATURE ALERT" alert being raised within the following 10 minutes. The confidence of the rule indicates that there is a 15% probability of the second alert occurring once the first alert is raised. The support indicates that there is a 10% probability that the entire alert sequence will occur during diesel engine operations.

If it is known that the diesel engines have a problem operating during times of high fresh water temperatures, especially when there is a lube oil pressure problem, a situation could then be avoided before it occurred. By generating many association rules for the diesel engine, the basis for an intelligent condition-monitoring system, able to predict future system states and problematic events, would be established. The condition-monitoring system could either automatically resolve the situation, by performing such duties as shutting down or starting up machinery or increasing water flow, or it could alert the operators to the fact that within in the next 10 minutes there is a 15% probability that the diesel engine will cease to function. The operators would then have the responsibility of deciding upon the best course of action.

The effects of such condition-monitoring ultimately have a positive effect relating to the operational availability of the submarine. For example, unnecessary down-time of the diesel engine may be avoided, allowing for the battery to be charged in a timely manner. This in turn allows the submarine to remain submerged for a greater length of time and therefore have a greater probability in successfully achieving its mission.

For each diesel engine, alert sequences were determined covering time frames of 5, 10, 15 and 25 minutes. Alert rules were not limited to two items, such as:

$$\text{Alert\_A} \Rightarrow \text{Alert\_B} \quad (3)$$

Using custom scripts, based on algorithms developed by Houtsma and Swami [11], rules of three, four or more items were obtained. The number of items within a rule was determined by meeting various association rule algorithm limits. Therefore, rules of the following form were also discovered:

$$\text{Alert\_A AND Alert\_B} \Rightarrow \text{Alert\_C} \quad (4)$$
$$\text{Alert\_A AND Alert\_B AND Alert\_C} \Rightarrow \text{Alert\_D} \quad (5)$$

Rules of the form presented in equations (4) and (5) have more use than that presented in equation (3). This is because there may be several instances when Alert_A would appear on the left hand side of the set of rules. This could possibly lead to confusion regarding what might occur next (as indicated by the right hand side of the rule). With more items on the left hand side of a rule, it would be easier to predict what might occur next, though it may take longer to make the prediction. However, a condition-monitoring system could include the knowledge of the association rules of two, three or more elements in conjunction with logic that would provide the operator with an informed assessment of the situation.

The support and confidence levels for each of the association rules were extracted and plotted in bar graphs, like that shown in Figure 4, in order to determine the time frames that would provide the optimum results.

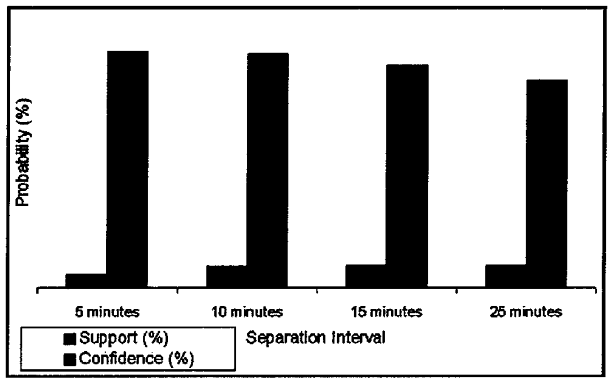

Figure 4. Confidence and support levels obtained for an association rule of the form "`Alert_A AND Alert_B ⇒ Alert_C`" over time frames of 5, 10, 15 and 25 minutes.

However, it was found in many of the results that the confidence level would decrease and the support level would increase over time. Confidence is the ratio of all the alerts in a given association rule occurring together over all the possible alert association rules raised by the diesel engine for a specific time frame. As time progresses, it was expected that the total number of alerts raised by the system would increase. If the confidence level remained constant it could be assumed that the alerts were occurring during every stage of diesel engine operations. Since the confidence was decreasing it was believed that the sequence was occurring only during particular stages in diesel engine operations. Examining the alerts with respect to diesel engine start-up and shutdown events and during normal running speeds, it was found that many alert sequences were associated with diesel engine start-up or shutdown sequences. The next step in the analysis could be to mine for association rules based on particular stages of diesel engine operations.

Since each diesel engine is unique and is used in different ways according to submarine standard operating procedures, there was no significant commonality between alert sequences for the diesel engines of one submarine or across several submarines.

Based on these results it was concluded that diesel engine alert association rules, on their own, would not prove useful in a condition-monitoring system. Further knowledge would be required. However, alert sequences were found during the start-up and shutdown procedures and these might still prove useful for reducing operator information overload. That is, the alerts may be typical of diesel engine start-up and shutdown procedures. If so, they can be "captured" before being presented to the operators and therefore the possibility of a "false alarm" situation could be avoided. Care obviously needs to be exercised so that alerts resulting from a real situation that may lead to diesel engine failure are not captured. This needs to be investigated further.

Further studies are also required to determine if the rules are beneficial in the use of fault diagnosis. This can be achieved by mapping the alert association rules to times of diesel engine "events" (for example, diesel engine cylinder failures).

4.3 DIESEL ENGINE DEPENDENCY ANALYSIS

A preliminary study was carried out to determine if any currently recorded diesel engine parameters could be related to the exhaust temperature of the diesel engine cylinders. The data mining algorithm currently employed in this study is referred to as "Column Importance" found in the data mining software package MineSet 2.5, developed by Silicon Graphics, Inc. (SGI) [12].

Prior to commencing the data mining, the algorithm called for the subject variable, the cylinder exhaust temperatures, to be separated into temperature "bins" of equal size. Similarly, the possible dependant variables were also separated into bins. The set of possible dependent variables includes oil pressures, cooling water temperatures and pressures, generator running speeds and air temperatures and pressures. At this early stage, the number of bins for each possible dependent variable was auto-generated by the mining algorithm, with the bin width varying from variable to variable.

The data mining phase consisted of several passes over the data and on each pass another exhaust temperature dependent variable was discovered. Dependency was determined based on a percentage value indicating the importance a particular variable had in predicting the exhaust temperatures.

There are 18 cylinders per diesel engine and a set of four dependant variables were found for use in predicting the temperature of each individual cylinder. Excluded from the mining was the effect one cylinder had on another cylinder.

Once four dependent variables were discovered, the cumulative predictive power of those variables was determined. Next, the set of four variables for each cylinder was compared to the variables of the other cylinders. The three most common re-occurring variables were then taken to be representative of the entire diesel engine and were plotted in a three dimensional scatter graph like that shown in Figure 5. In the graph, the three dependent variables form the X, Y and Z axes, respectively, and the cylinder exhaust temperature is represented by colour. Here, the Z variable is along the vertical axis, the X variable is along the "diagonal" axis on the left and the Y variable is the "diagonal" axis on the right. Due to the nature of what the information may represent, the graph remains dimensionless. Cylinder exhaust temperature is given in the legend, shown in the lower left corner of Figure 5, and represents an increase in temperature looking from left to right. Due to the limitations of the two dimensional representation of a three-dimensional image, it appears that in Figure 5 the blue spheres are mainly influenced by the X and Z variables while the red spheres are influenced by all three variables. In reality, the blue spheres are also influenced by the Y variables. The problem then becomes one of classification and the need for partitions to define rules governing each class.

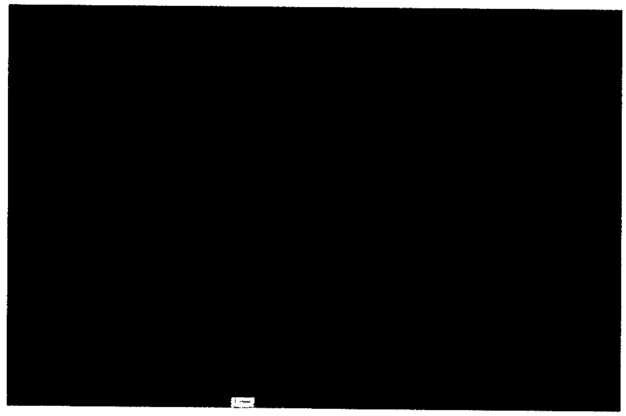

Figure 5. Three-dimensional scatter graph representing the three dependent variables for use in predicting the exhaust temperatures of the diesel engine cylinders.

The cumulative predictive power for the three common re-occurring variables for each of the 18 cylinders was calculated and compared to the predicting power previously calculated. While differences did occur, they were mostly within a few percent.

The graph in Figure 5 represents variable values for one particular cylinder. However, graphs for the remaining 17 cylinders showed similar data distribution patterns. As with the association rules, the effects of various stages in diesel engine operations need to be considered.

By examining the graph, with the aid of the data mining software, it is possible to determine the representative values of the various classes evident in the graph. It may then be possible to develop a mathematical model of exhaust temperature behaviour during the various stages of diesel engine operations.

A diesel engine mathematical model could be useful in the optimisation of diesel engine performance and also in condition-monitoring to determine when operational values are deviating from the normal levels or to predict when the operational values will deviate from the normal.

4.4 MAST AND PERISCOPE USAGE ANALYSIS

The Collins Class submarine is fitted with five masts, a high frequency antenna and two periscopes. More specifically, they are the:

- Electronic Surveillance Measures (ESM) Microwave (MW) Mast
- ESM Sub-Microwave (SMW) Mast
- Induction (or Snort) Mast
- Communications Mast
- Radar Mast
- High Frequency (HF) Whip Antenna
- Search Periscope
- Attack Periscope

The mast and periscope configuration is shown in Figure 6. The figure represents a cut-away of the fin section of the Collins Class submarine.

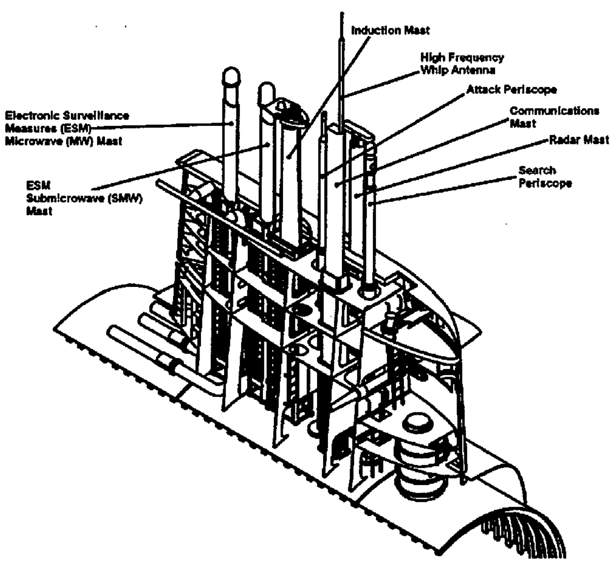

Figure 6. The periscopes and masts of the COLLINS Class submarine.

The ESM masts are for the early detection and direction finding of radio frequency emissions; the induction mast is for the venting of exhaust gases and the intake of clean air; the communications mast and HF Whip antenna allow the submarine to communicate with external sources; the radar is primarily used as a navigation aid for the detection of land masses and surface targets; the search periscope is used for surveillance of the "above-water scene" including target detection and navigation support; and the attack periscope allows for target detection and target identification.

The five masts have their usage recorded in the low rate data log in terms of mast height. The usage of the periscopes and the HF Whip antenna is recorded in the event log in terms of two unique identifiers to indicate the system being in the raised state (represented by '1' and '0' respectively, that is, TRUE and FALSE) and the system being in the lowered state (again represented by '1' and '0', respectively).

Since the required usage information is spread across numerous database fields the pre-processing had to merge the relevant information. For the mast height, anything greater than 0·5m indicated that the mast was raised and was assigned the value of '1'. Anything less than 0·5m indicated that the mast was lowered and was assigned the value '0'. This resulted in a series of consecutive '1s' and '0s'.

Duplicates were then removed so that the remaining entries represented an alternating series of '1' and '0' as shown in Table 1.

Table 1. An example representation of induction mast height pre-processing. TRUE is represented by '1' (mast raised) and FALSE is represented by '0' (mast lowered).

Date/Time Stamp	Induction Raised
2-Jan-2002 12:33:34	1
2-Jan-2002 12:45:44	0
2-Jan-2002 13:03:10	1
2-Jan-2002 13:15:10	0
2-Jan-2002 13:35:10	1
2-Jan-2002 13:43:10	0
2-Jan-2002 14:10:10	1

The periscopes and HF Whip antenna data required pre-processing in a different manner. If either of the periscopes or HF Whip were raised, then the value of the raised identifier would be '1' and the value of the lowered identifier would be '0'. If the system was lowered, the corresponding value of the raised identifier would be '0' and the value of the lowered identifier would be '1'. Since an indication of raised and lowered was required, it was not a simple matter of merging the data from the two variables. A script was developed in order to take the '1' state (that is, TRUE) from the lowered state, convert it to a '0' (that is FALSE to indicate not raised) and then merge it with the '1' state of the raised state. The result should then be an alternating list of '1' and '0' to indicate whether the periscope was raised ('1') or lowered ('0'), similar to that shown in Table 1 for the induction mast.

With the usage information pre-processed in such a manner it was then possible to determine how the masts were used over a period of 24 hours, or, say, a month. The next step would be to cross-reference the information with failures or faults within the associated systems. Alternatively, the usage information could be used in operational analysis giving an indication of the length time masts spend breaching the surface (of the water). This is important because a mast breaching the surface produces a wake trail and radar signature and therefore gives away the submarine's presence. The usage information will also lend support in maintenance scheduling and mission planning.

5. POTENTIAL APPLICATIONS BENEFITING FROM KDD ANALYSIS

There are several areas where the results of applying KDD to ISCMMS could have potential benefits. These are:

- Submarine system models
- Decision aids and condition-monitoring systems
- Identification of improved recording practices
- Submarine historical mission profile
- On-line analysis tools in support of submarine logistics

5.1 SUBMARINE SYSTEM MODELS, DECISION AIDS AND CONDITION-MONITORING SYSTEMS

Submarine system models may provide land-based test facilities a method to reliably reproduce submarine actions to given events. The test facilities would provide the means for testing and tuning new methods and tools prior to being assessed in the operational environment.

The submarine system models would also provide the basis for decision aids and condition-monitoring systems, since the models would form the knowledge base required by such systems. It is hoped that through the KDD analysis of ISCMMS data enough information will be discovered to aid in the development of tools able to predict potential system faults and/or failures and also aid in fault diagnosis. It is expected that being able to predict system faults will result in time and monetary savings and lead to an increased operational availability. Similar savings should also be observed in fast fault diagnosis.

It is ultimately envisaged that the submarine systems would each have their own intelligent condition-monitoring system with the condition-monitoring system taking the place of the SPUs. Each condition-monitoring system would be responsible for recording the data relating to their respective system and from the data, in conjunction with the built-in system models, an intelligent assessment of the current situation would be made. The results would then be presented to the operators at the GMUs and LPOs and they would then decide upon the appropriate course of action.

5.2 IMPROVED RECORDING PRACTICES

Currently ISCMMS data is recorded based on the occurrence of an event, for example, the starting of machinery, or it is time based with submarine system states being recorded at regular intervals. For much of the logistics engineering analysis, this form of recording would suffice. However, in submarine KDD analysis one of the aims is to determine cause and effect. That is, to find the inter-relationships between system parameters - how a change in one parameter will affect the efficiency of another parameter. Alternatively, it may be used to determine the cause of system faults. The current recording practices of the system states do not provide sufficient information to achieve these aims. Therefore, two suggested improvements are being put forward:

1. On demand logging

Instead of recording all system values at regular intervals, the proposed new method would be to record data when a change in state is observed. For example, if the mast height changed by 0.25m then record the change, otherwise ignore it. Each system would have its own specific trigger level that would instigate recording the new value of the system parameter. The trigger level should also be configurable so that the fidelity of the data recording can be increased or decreased as required.

2. High frequency data logging

This data logging function would record data at high frequencies over a rolling window of a given time frame, for example, 10 minutes [13]. If an event occurs within that window then the data would be stored, otherwise it would be discarded. Furthermore, should an event occur at the beginning of a new recording window, the previous window of data should also be retained. Therefore, while recording the current window of high frequency data, the previous window of data is retained. If an event occurs in the current window, the previous and current data windows are retained. Otherwise, the previous window is discarded and then the current window would become the previous window.

A further enhancement to this suggestion would be to record data at a slower frequency, or on demand, and once a system spike is observed, the system would automatically switch to high frequency data logging [14]. This functionality does have some similarities to the Trace Logging function already available in ISCMMS, however more knowledge may need to be incorporated in the system in order to identify the point at which to switch over to high frequency logging.

The high frequency data logging function would also be system dependent and some systems may not require the function.

These two suggestions would only affect data currently recorded by the low rate data log and would possibly replace the high frequency data log. The functionality of the remaining logs would be unchanged, except that some parameters currently being recorded by the event log would in future be recorded by the new on demand logger. These suggestions would also call for changes in sensor types for some systems.

5.3 SUBMARINE HISTORICAL MISSION PROFILE

The data pre-processing phase of the KDD analysis not only results in data ready for data mining but it also generates submarine system usage information. This usage information will be used to generate the submarine squadron's historical mission profile.

A submarine mission may be broken into the following stages [15]:

1. Entering and Leaving Harbour
2. Transit
3. Patrol

During each of the stages the submarine may be in one or more operating modes:

1. Surfaced
2. Dived
    a. Snorting
    b. Dived (0-50)
    c. Dived (>50m)
    d. Periscope depth
    e. Deep

Submarine system usage is generally dependent on the requirements for the current operating stage/mode. For example, the diesel engines would only be used during surface running or when the submarine is at periscope depth. Depending on whether the submarine is on the surface or at periscope depth, different exhaust valves may be employed to vent the battery and diesel engine exhaust gases. When the submarine is at a deep diving depth the periscopes and masts would be stowed and the submarine would be running on battery power.

With the information recorded via ISCMMS it is possible to determine the equipment usage at each stage of an operation. This usage information would aid in the development of improved logistics engineering practices, ranging from maintenance scheduling through to the stocking of stores and also in the development of future submarine systems. It may also aid in operational analysis studies leading to the prediction of mission success.

5.4 ON-LINE ANALYSIS TOOLS

Currently being considered within the submarine community of the RAN and Australian Submarine Corporation (ASC) is the development of a data warehouse for submarine logistics data. Very basically, a data warehouse is a database designed around the need for data analysis. The submarine logistics database, SIMS, is a transactional database designed for the input and extraction of data for managing the submarines. The prototype data pre-processing tools and the data mining techniques being developed in the KDD ISCMMS studies should aid in the development of formal tools for use in the on-line analysis of the SIMS data warehouse.

6. CONCLUDING COMMENTS

While the initial studies have not yielded any new (that is, previously unknown) information about the Collins Class submarines, the potential still exists for the discovery of new and valuable information that may ultimately provide insights into the submarine system processes. This new information could aid in the development of intelligent condition-monitoring systems, fault diagnosis tools and provide information for further logistics engineering analysis.

The analysis has provided an understanding of the current data collection techniques and helped identify where improvements can be made in those techniques.

The data pre-processing will provide the historical mission profile for each submarine and that can also aid in logistics engineering analysis and submarine operational analysis studies.

ACKNOWLEDGEMENTS

The author would like to thank staff of the Submarine Systems Support Facility, Royal Australian Navy and staff of Australian Submarine Corporation for providing assistance with issues relating to ISCMMS.

The author wishes to acknowledge Director General - Maritime Support, Maritime Systems Division, Defence Materiel Organisation, Department of Defence, as the sponsor of this study.

REFERENCES

[1] Fayyad, U., Piatetsky-Shapiro, G., Smyth, P. and Uthurusamy, R., eds. 1996, *Advances in Knowledge Discovery and Data Mining*, MIT Press, Cambridge, MA

[2] Fayyad, U. and Stolorz, P., 1997 Data Mining and KDD: Promise and challenges, *Future Generation Computer Systems*, **13** (2-3), Elsevier Science B. V.

[3] Casarin, P., 1997, Using Data Mining Techniques in Auditing, *Audits and Control Journal*, **5**, Inf. Syst. Audit and Control Assoc.

[4] Williams, G., Commonwealth Scientific and Industrial Research Organisation, 2000, The Data Mining Enterprise, *The Inaugural Information Technology Summer School: Data Mining and Heuristics*, School of Computer Science, Australian Defence Force Academy, The University of New South Wales, 11 November - 8 December 2000

[5] Ho Ha, S., and Park, S. C., 1998, Application of data mining tools to hotel data mart in the Intranet for database marketing, *Expert Systems with Applications*, **15** (1), Elsevier Science Ltd

[6] Esp, D. G., Carillo, M., and McGrail, A. J., 1998, Data Mining Applied to Transformer Oil Analysis Data, *Record of the 1998 IEEE International Symposium on Electrical Insulation*, Arlington, Virginia, USA, June 7-10, 1998, IEEE

[7] Saraee, M., Koundourakis, G. and Theodoulidis, B., 1998, Easyminer: Data Mining in Medical Databases, *IEE Colloquium Intelligent Methods in Healthcare and Medical Applications*, York, UK, October 20, 1998, IEE, UK

[8] Bansal, K., Vadhavkar, S., and Gupta, A., 1998, Neural Networks Based Forecasting Techniques for Inventory Control Applications, *Data Mining and Knowledge Discovery*, **2** (1), Kluwer Academic Publishers

[9] Klemettinen, M., Mannila, H., and Toivonen, H., 1997, A Data Mining Methodology and Its Application to Semi-Automatic Knowledge Acquisition, *Proceedings of the 1997 8th International Workshop on Database and Expert Systems Applications, DEXA'97*, Toulouse, France, September 1-2, 1997, IEEE, Los Alamitos, California, USA

[10] Queens University of Belfast, *Data Mining Notes*, http://www.pcc.qub.ac.uk/tec/courses/datamining/stu_notes/dm_book_2.html

[11] Houtsma, M. and Swami, A., 1996, Mining Quantitative Association Rules in Large Relational Tables, *SIGMOD 96*

[12] Rathjens, D. and Vandenberg, H., 1998, *MineSet User's Guide*, Silicon Graphics, Inc., California, USA

[13] Private discussions with Goodwin, G., Maritime Platforms Disivison, Defence Science and Technology Organisation, October 2002

[14] Private discussions with Boere, B., Maritime Platforms Disivison, Defence Science and Technology Organisation, November 2002

[15] Used with permission, Australian Submarine Corporation, Pty. Ltd., November 2002

COPYRIGHT AND TRADEMARKS

Model Updating and Determination Of Structural Dynamic Properties In The Presence Of A Second Source Of Excitation

Stephen Moore, Zhang Biaobiao, Joseph Lai, and Krishna Shankar
School of Aerospace and Mechanical Engineering,UNSW@ADFA,
Australian Defence Force Academy.

ABSTRACT©

Predictions of the dynamic behaviour of helicopters usually involve finite element (FE) models of the structure. These can be easily modified to determine the effect of different stores configuration or structural modification. The FE models are necessarily simplified models and are often validated or updated on the basis of experimental data. Experimental modal testing of a grounded helicopter or a helicopter suspended from the rotor hub neglects the effects of rotating components, for example the main rotor, and aerodynamic loading. Therefore it is preferable to obtain accurate experimental data during flight. This paper discusses the use of experimental data to update an FE model of a helicopter-like structure. In order to examine the feasibility of obtaining dynamic properties of a helicopter during flight, periodic impulsive excitation and synchronous averaging in the presence of a second source of excitation was applied to a simply supported beam. Model updating has been found to increase the accuracy of dynamic properties predicted by the FE model. Periodic impulsive excitation and synchronous averaging has been shown to improve the quality of modal properties extracted by experimental modal analysis in the presence of a second source of excitation.

INTRODUCTION

The accurate prediction of structural dynamic properties of helicopters is a critical requirement in effective design and modification of helicopter structures. Experimental and analytical techniques are currently used to construct models describing the dynamic properties of helicopter structures. These models can then be used to aid the prediction of component loads, which in turn can be used for calculation of component retirement times. The models also serve as powerful tools when trying to prevent adverse dynamic behaviour resulting from non-standard flight regimes or modification of the helicopter structures.

Finite element (FE) models are often used in the design process of helicopter structures, however, it has been shown that dynamic analysis carried out on these models produces results of limited accuracy when compared to experimental data[1, 2]. Some factors that contribute to poor accuracy include the difficulty of modelling some components, for example welded joints, as well as the accuracy of element parameters used in the FE model. In addition, the time and computational power required to solve complex models can often be prohibitive, consequently simple accurate models are most desirable. Nevertheless, FE models present several advantages over experimentation. Experimentation is usually case specific and therefore many experiments are required to obtain results that fully describe the behaviour of the helicopter structure in a number of configurations. Furthermore, experiments are often expensive, time consuming and in some cases not practical due to safety issues or availability of a suitable test platform.

A solution that aims to improve the accuracy of FE models involves the use of experimental results to modify or update the FE models. This can be carried out manually by using engineering judgement to modify model parameters. Recently, a number of automated techniques have been developed based on optimisation algorithms [3-5] allowing a more systematic model updating procedure to be carried out. Model parameters are modified within physically realistic bounds such that the dynamic behaviour determined by the FE model correlates with that obtained from experimental results. Validation of the model updating procedure can be achieved by testing a modified physical structure and comparing results with those obtained from a similarly modified FE model. A successful outcome to the model updating procedure is based on the ability of the FE model to produce accurate results for a modified structure.

A critical component in the model updating process is the accuracy of the experimental data. Experimental modal analysis (EMA) is a procedure used to obtain the dynamic properties (modal parameters) of a structure. It has been shown [6] that dynamic properties obtained from EMA of a grounded helicopter structure, or those suspended from the main rotor head, differ from results obtained from helicopters in-flight. This suggests that the effects of rotating components and aerodynamic loading are significant. Therefore, a suitable method that allows modal parameters to be determined accurately while the helicopter is in flight is required. A major difficulty involved with carrying out EMA on helicopters during flight is accurately measuring the excitation force. Previous research [1, 6] has attempted to use response only measurements of helicopter vibration, due to ambient aerodynamic excitation, or control inputs applied by the pilot, to determine the dynamic properties of the structure. While these studies have been able to extract modal natural frequencies, mode shape information is usually less accurate.

The use of periodic impulsive excitation and synchronous averaging is aimed at overcoming one of the problems of carrying out EMA on a helicopter during flight. Synchronous averaging of impulse-response records measured at a point on the helicopter airframe can filter out the effects of other sources of excitation that are effectively random, for example aerodynamic loading. The averaged records can be used in the curve fitting process to determine the modal parameters.

This paper discusses two experiments used to demonstrate model updating and the determination of structural dynamic behaviour in the presence of a second source of excitation. The first experiment involves EMA of a helicopter-like structure and model updating of a FE model of the structure using experimental results. The model updating procedure is then evaluated by adding weights to the physical structure, carrying out another set of modal tests, and then comparing the results with those obtained from the similarly modified updated FE model.

The second experiment involves carrying out modal testing on a simply supported beam using periodic impulsive excitation, in the presence of a second source of excitation. Results of the tests are shown for different numbers of averages and relative noise levels.

MODAL ANALYSIS OF HELICOPTER-LIKE STRUCTURE

The helicopter-like structure used for these experiments, shown in figure 1, consisted of a primary structure made up of bar and tube sections, with steel sheet forming the secondary elements of the floor and back panels, roof and upper side plates. The structure was suspended as shown in the figure using elastic cord and the resulting rigid body modes were found to be less than 20% of the first modal frequency, indicated by preliminary experimental testing of the helicopter-like structure. Impact hammer excitation was employed for EMA and fixed response measurements in the direction of the three principal axes were made using B&K type 4374 mono-axial accelerometers mounted on the horizontal and vertical tail spars. There were 134 excitation points throughout the structure with emphasis on the primary structural elements, rather than the secondary steel plate elements. A Hewlett-Packard HP 3566A FFT analyser was used for data acquisition and preliminary signal processing; measurement parameters for the test were as follows: analysis bandwidth 400Hz; frequency resolution 1 Hz; transient and exponential weighting on excitation and response signals respectively and 5 averages per measurement. STARmodal v. 5.23 was used to curve fit the frequency response function (FRF) data and extract the modal parameters. Frequency results are listed in table 1.

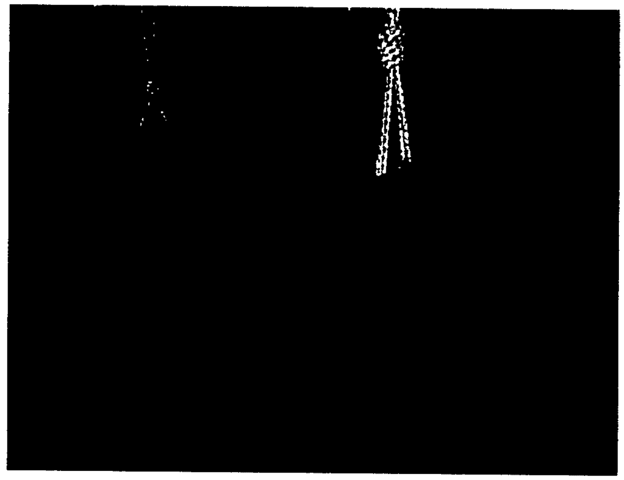

Figure 1 Helicopter-like structure.

Table 1. EMA frequency results.

Mode	Frequency (Hz)	Mode	Frequency (Hz)
1	100.1	11	303.6
2	171.6	12	326.5
3	195.3	13	328.5
4	197.3	14	347.8
5	22.6	15	352.9
6	237.3	16	358.8
7	247.0	17	371.4
8	267.5	18	378.3
9	273.4	19	387.5
10	290.1		

FINITE ELEMENT MODEL OF HELICOPTER-LIKE STRUCTURE

A preliminary FE model [7]was used as the basis for model updating. Modifications were made to the initial FE model such that characteristics of the physical structure were more accurately represented. This included improved modelling of the interface between the beam elements and plates used for the floor, roof and side panels. The ANSYS FE model is shown in Figure 2.

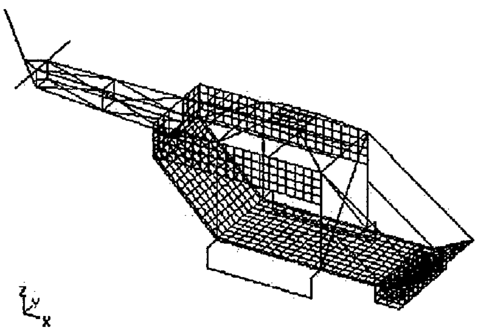

Figure 2 ANSYS FE model of helicopter-like structure.

Beam and solid elements were used for the primary structure. Shell elements were used to model the floor, roof and side plates of the cargo bay and zero damping was applied to the structure. The block lanczos solution method was used and the frequency results are listed in table 2. Comparison of mode shapes obtained from experimental and FE results is carried out during the model updating procedure and discussed below.

FE-Model Updating

The modal frequency results from the initial FE model can be seen to vary significantly from those obtained experimentally. The aim of the model updating procedure is to improve the correlation of FE model and EMA results. The model updating software package FEMtools v2.2 was used for this process.

Table 2. FE model results.

Mode	Frequency (Hz)	Mode	Frequency (Hz)
1	98.304	11	278.53
2	117.76	12	287.92
3	181.52	13	291.44
4	181.69	14	314.89
5	227.31	15	316.35
6	233.35	16	317.77
7	259.24	17	319.21
8	259.93	18	332.87
9	261.00	19	333.11
10	263.25		

The model updating procedure involves a number of steps:

1. Spatial correlation of nodes and points: Nodes from the finite element model are paired with measurement points identified during EMA. Figure 3 shows the FE model of the helicopter-like structure with coincident node/point pairs indicated by dots.

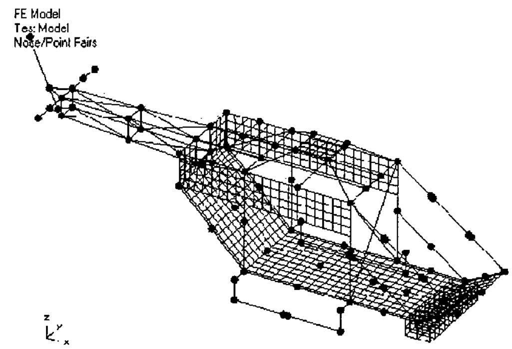

Figure 3 Node/point pairs used in mode updating.

2. Shape correlation: This procedure compares the FE and EMA mode shapes. Numerical shape correlation is an effective method for comparing mode shapes for complex structures. The Modal Assurance Criterion (MAC) was used and is defined in equation 1. Automatic shape pairing was carried out with the following constraints: mode shape pairs have a minimum MAC value of 19% and an allowable frequency difference of 80%. A matrix of MAC values is shown in Figure 5a. A diagonal line of MAC values equal to 100 would indicate perfectly correlated modes.

$$MAC(\psi_a,\psi_e)=\frac{\left|\left(\{\psi_a\}^t\{\psi_e\}\right)\right|^2}{\left(\{\psi_a\}^t\{\psi_a\}\right)\left(\{\psi_e\}^t\{\psi_e\}\right)} \qquad \text{(Eqn. 1)}$$

Where: ψ_a = FE eigenvector
ψ_e = EMA eigenvector

3. Sensitivity analysis: Sensitivity analysis defines the rate of change of a FE model response property, in this case modal frequency, as a function of the change of a FE model parameter. Sensitivity analysis identifies the most efficient parameters to modify during model updating, and the inverse of the sensitivity matrix, the gain matrix, is used during model updating to calculate the magnitudes of parameter changes. Normalised relative sensitivities are independent of units for both the response properties and the model parameters, allowing comparison of several different parameters. Sensitivity analysis identified eight parameters to be modified during model updating: Elasticity matrix scaling; Young's modulus; density; cross-sectional area; moment of

inertia about the local x,y and z axes of elements; and membrane thickness. Parameters can also be modified locally or globally; local parameter changes occur for individual elements while global parameter changes apply to sets of elements.

4. Model updating iterative process: Figure 4 shows the block diagram of the model updating procedure and table 3 lists details of each iteration. The confidence value indicates the estimated error for a parameter. It is calculated as the inverse of the relative error of the parameter, multiplied by 100. For example, an estimated error of 25% corresponds to a confidence value of 400. It should be noted that the internal FE solver was used rather than recalculating response values in ANSYS for each iteration. Because of this, the initial FE results calculated in FEMtools during model updating differ slightly from those obtained from ANSYS. The convergence criterion defines the objective function that is minimised during model updating. The absolute difference between EMA and FE results for each mode was used as the convergence criterion for this analysis.

The final model updating step is the MAC contribution analysis (MCA). This procedure ranks DOFs according to how they affect the MAC value, allowing DOFs with an adverse affect on MAC values to be assessed in terms of the quality of experimental data. In some cases it is beneficial to remove the deflections at particular DOFs to improve MAC values. Eight DOFs out of a total of 134 were removed after MCA.

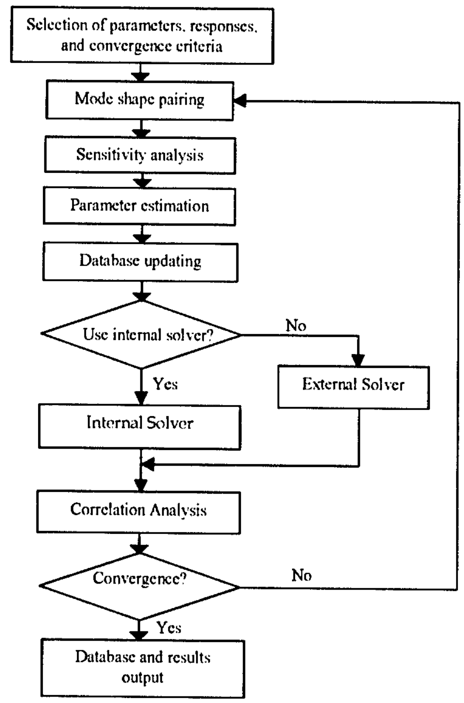

Figure 4. Model updating procedure.

Table 3. Model updating steps.

Model Updating Step#	Parameter Variation	Parameter Bounds and Confidence
1	Elasticity matrix scaling: D (Global and local) Total Iteration # = 5	(-10%<D<10%) Confidence=400
2	Young's Modulus: E (Global and local) Total Iteration # = 5	(-10%<E<10%) Confidence = 400
3	Mass parameter: RHO (Global and local) Total Iteration # = 5	(-10%<RHO<10%) Confidence = 400
4	Cross-sectional area: AX (Global and local) Total Iteration # = 5	(-10%<AX<10%) Confidence = 400
5	Moment of inertia: Ix,Iy,Iz (Global and Local) Total Iteration # = 3	(-10%<X<10%) Confidence = 400
6	Membrane thickness: H	(-10%<H<10%) Confidence = 400
7	Adjusting DOFs Pairing, fine-tuning D,E,H together	(-10%<X<10%) Confidence = 400
8	Adjusting DOFs Pairing, fine-tuning D,E together .	(-10%<X<10%) Confidence = 400
9	MAC Contribution Analysis	

RESULTS OF MODAL UPDATING

Table 4 compares modal frequencies and lists MAC values for mode pairs from the initial FE model and EMA results. Details of mode pairs after model updating are shown in table 5. The results after model updating show improved correlation between FE and EMA results, with less than 5% difference between modal frequencies for eight of the nine mode pairs. Large improvements in MAC values can also be seen for each mode pair, and the averaged MAC value for all mode pairs has approximately doubled. In addition, comparison of the MAC matrices before and after model updating, shown in Figure 5, reveals an improvement in the diagonalisation of the matrix. These results are promising, however, a more practical evaluation of model updating involves modification of the physical structure and FE model and a comparison of the results.

Table 4 Mode shape pairs before model updating.

Pair no.	EMA Mode	Frequency (Hz)	FE Mode	Frequency (Hz)	% Error (Re EMA)	MAC
1	1	100.08	2	100.14	0.06	19.2
2	2	171.58	3	124.48	-37.84	22.8
3	3	195.35	7	253.41	22.91	65.5
4	8	267.47	10	263.00	-1.70	59.5
5	6	237.26	11	271.74	12.69	38.2
6	9	273.43	13	288.14	5.11	29.4
7	10	290.08	14	308.87	6.078	23.2
8	12	326.46	16	317.48	-2.83	21.4
9	18	378.28	20	341.11	-10.90	31.5
10	16	358.84	21	343.80	-4.37	59.5
11	17	371.37	23	370.53	-0.23	21.0
					Average MAC Value	35.6

Table 5 Mode shape pairs after model updating.

Pair no.	EMA Mode	Frequency (Hz)	FE Mode	Frequency (Hz)	% Error (Re EMA)	MAC
1	1	100.08	4	100.05	-0.03	43.2
2	3	195.35	8	194.11	-0.64	81.6
3	4	197.33	9	196.68	-0.33	83.2
4	7	247.03	12	243.51	-1.45	76.0
5	6	237.26	13	247.4	4.1	84.9
6	8	267.47	14	266.92	-0.21	94.4
7	10	290.08	19	300.22	3.38	69.8
8	18	378.28	20	347.29	-8.92	44.9
9	16	358.84	21	357.32	-0.43	69.7
					Average MAC Value	71.9

EVALUATION OF UPDATED FINITE ELEMENT MODEL

The modifications applied to the physical structure and the updated FE model involved clamping additional mass to the rear of the skids, as shown in figure 6. The total additional mass was just over 5.5kg, including the two G-clamps used to secure the weights. Modal testing, curve fitting and extraction of the modal parameters was carried out in the same manner as the initial EMA described above.

The additional mass and G-clamps were modelled with mass elements in the FE model, and a block lanczos solution method was used to extract frequency results and mode shapes. In order to compare the two sets of results, the FEMtools software package was used to pair modes with constraints on pairs of 50% allowable variation of modal frequencies and minimum MAC values of 37%.

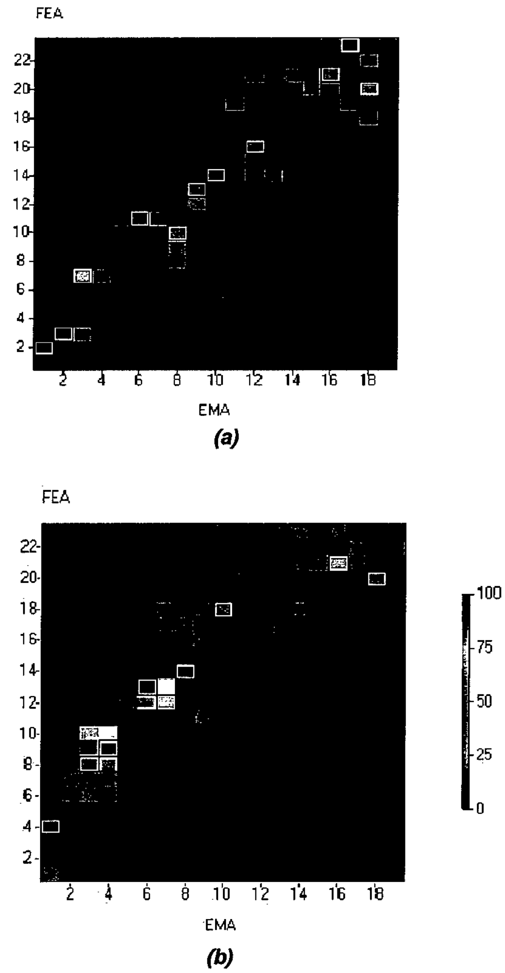

Figure 5 MAC matrix, before (a) and after (b) model updating.

The results for the modified helicopter-like structure, listed in table 6, indicate some improvement over those obtained before model updating. Only two MAC values were above 70%, with six below 50%. The average MAC value of 52.36% indicates some improvement compared to 35.6% for the results before model updating. However, the paired modes for the modified structure typically have marginally larger frequency differences. Comparison of the matrix of MAC values before model updating and for the modified structure, shown in figures 5a and 7 respectively, reveal a greater spread of MAC values greater than 20% away from the ideal diagonal. This suggests greater correlation between high order modes from one set of results and low order modes from the other set of results and vice-versa. Insufficient measurement points or noisy experimental data contribute to innacurate calculation of mode shape data, particlularly for higher order modes and this may account for the greater spread of MAC values seen in figure 7.

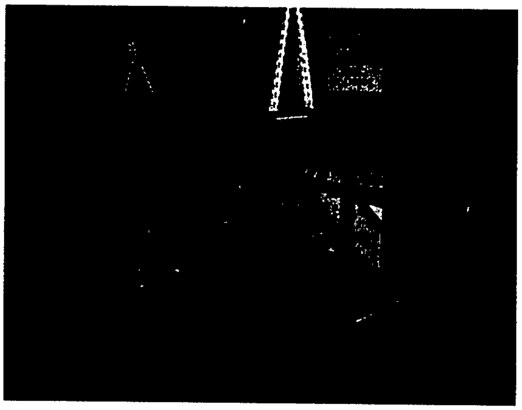

Figure 6 Helicopter-like structure with additional weights attached to skids.

Table 6. Mode shape pairs for modified helicopter-like structure.

Pair no.	EMA Mode	Frequency (Hz)	FE Mode	Frequency (Hz)	% Error (Re EMA)	MAC
1	2	85.33	1	72.29	-15.28	51.8
2	1	75.64	3	85.19	12.63	46.8
3	3	93.18	4	113.87	22.20	43.8
4	9	189.12	6	146.87	-22.34	39.3
5	7	155.84	7	162.42	4.22	75.2
6	15	280.81	9	210.80	-24.93	43.8
7	10	194.37	10	220.96	13.68	69.1
8	12	229.71	11	243.75	6.11	37.8
9	21	326.42	20	327.31	0.27	71.0
10	23	355.29	25	365.47	2.87	45
					Average MAC Value	52.36

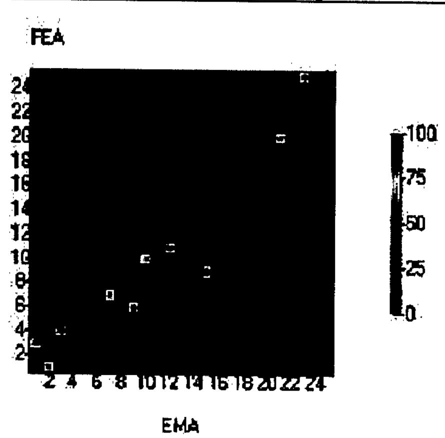

Figure 7 MAC values for modified structure.

MODAL TESTING IN THE PRESENCE OF A SECOND SOURCE OF EXCITATION

A simply supported beam was employed to investigate the use of periodic impulsive excitation and synchronous averaging for EMA in the presence of a second source of excitation. Figure 8 shows the aluminium beam supported by string, approximating free-free conditions in the transverse direction, as transverse elastic vibrations were only considered for these tests. Two piezoceramic plates were bonded to both major surfaces of the beam, close to one end. The primary impulsive excitation was applied using one piezoceramic plate and the second source of excitation (background noise) was applied using the second piezoceramic plate. Response measurements were made with one B&K 4374 monoaxial accelerometer and thirty measurement points were selected along the beam to provide reasonable mode shape resolution for the first seven transverse modes. A B&K 2032 FFT analyser was used for data collection and signal processing with the following measurement parameters: analysis bandwidth 400Hz; frequency resolution 0.5 Hz; transient and exponential weighting on excitation and response signals respectively. STARMODAL v. 5.23 was used to curve fit the frequency response function (FRF) data and extract the modal parameters. Different primary and background excitation levels were used for each experiment, with the level of background noise being defined as ratio of the RMS value of the random noise signal and the amplitude of the impulse signal. The excitation force applied by the piezoceramic plates could not be measured directly; therefore the impulse excitation signal was used for the calculation of FRFs. The number of averages per measurement record was also varied for each experiment.

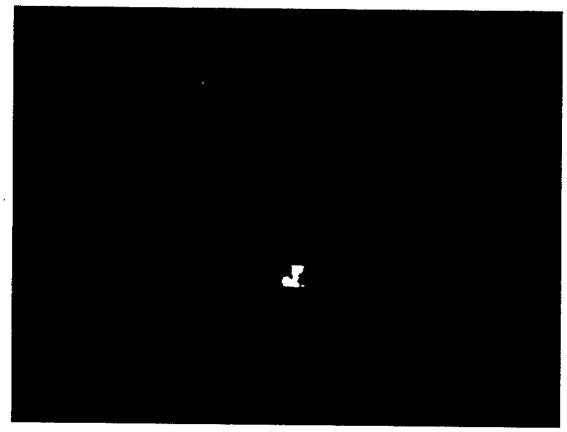

Figure 8 Simply supported aluminium beam. One piezoceramic plate can be seen at the left of the beam.

Table 7. Details of EMA carried out on beam.

Experiment	Excitation	Noise Level	No. Averages
1	Impact hammer	N/A	< 20
2	Impulse excitation (piezo) & random noise (piezo)	0.125	20
3	Impulse excitation (piezo) & random noise (piezo)	0.125	40
4	Impulse excitation (piezo) & random noise (piezo)	0.125	80
5	Impulse excitation (piezo) & random noise (piezo)	0.250	80

An initial test employing impact hammer excitation was used to establish modal frequencies and mode shapes, with up to twenty averages taken per measurement record. Four tests were carried out to investigate the effect of background noise level. A train of impulses was used for primary excitation, and random noise with a bandwidth of 400Hz was used for background excitation. Table 7 lists details of each experiment.

RESULTS

Figure 9 shows modal frequency error for experiments 2-5 compared to experiment 1. These results demonstrate that the addition of background noise does not affect the extraction of modal frequency significantly. The results for the first mode were omitted as it was found that the piezoceramic plates did not effectively excite the first mode, therefore the results for the first mode were poor for all tests employing the piezoceramic plates for excitation. Mode shapes extracted from experimental data were observed to be more sensitive to the level of background noise. Table 8 shows extracted mode shapes of the second mode for each experiment. The results are representative of the other mode shapes except those for the first mode, due to reasons mentioned above. The effect of increasing numbers of averages can clearly be seen with more accurate mode shapes being extracted for greater numbers of averages in experiments 2,3 and 4. The quality of mode shapes extracted from experiments 4 and 5 are similar even though background noise used in experiment 5 was twice the amplitude of that used in experiment 4.

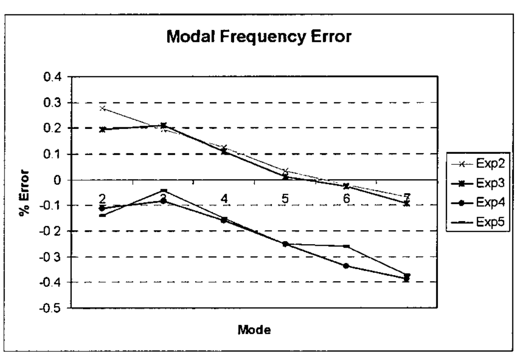

Figure 9. Modal frequency error.

A difficulty involved with periodic impulsive excitation and synchronous averaging is the length of time record required for large numbers of averages when the frequency range of interest is low. Each average is triggered by the impulsive excitation and sufficient time is required to capture the low frequency vibrations. Therefore the time required for large numbers of averages, say up to 500, may become prohibitive in some cases.

Piezoceramic plates are an effective means of exciting the beam structure; however, there are considerations that need to be taken into account. The piezoceramic plates used for these experiments were unable to excite the first mode. Further investigation needs to be carried out to determine whether this is a function of the plate's position on the beam or due to the dimensions and piezoelectric capabilities of the plate. Another limitation of using piezoceramic plates for modal testing is the difficulty involved with obtaining an accurate measure of the excitation force applied to the structure by the piezoceramic plate. For these experiments, the excitation signal was used as the force input, however, this is not an accurate measure of the applied force and can lead to poor experimental results, particularly modal damping values.

CONCLUSION

The use of analytical models is an efficient and flexible method for determining the dynamic properties of structures. The accuracy of the results can be improved by updating the analytical model using experimental results. This procedure was carried out for a helicopter-like structure and model

Table 8. Mode shapes extracted by EMA.

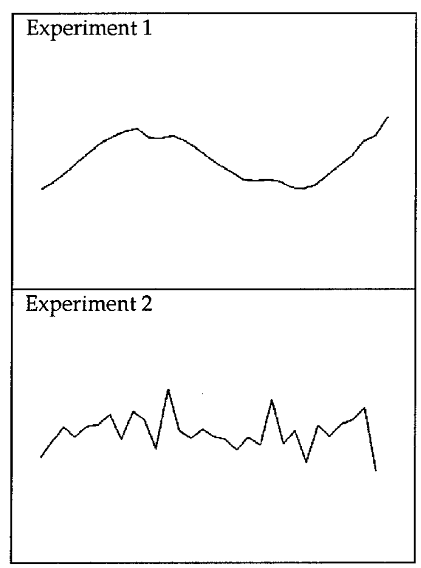

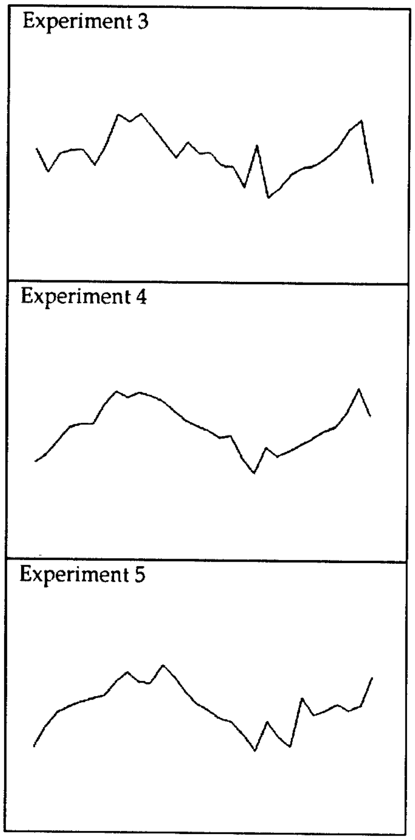

updating was shown to improve the correlation between experimental and finite element model results. Further evaluation of the model updating procedure was carried out by modifying the physical helicopter-like structure and comparing experimental results with those obtained from the similarly modified updated finite-element model. The results show some improvement in the accuracy of the finite element results. More investigations are required to yield more significant improvement.

A critical aspect of this procedure is the accuracy of experimental results. The second series of experiments demonstrated the use of periodic impulsive excitation and synchronous averaging in modal analysis of a beam with a second source of excitation. Varying levels of background noise were added to the beam during modal testing and the results show that synchronous averaging does improve the quality of the extracted mode shapes. Further work is required to determine the full capabilities of synchronous averaging and the use of piezoceramic plates for excitation of structures.

ACKNOWLEDGEMENTS

This project is being supported by a grant under the DSTO Centre of Expertise in Helicopter Structures and Diagnostics.

The advice provided by Professor Bob Randall, as well as the encouragement of Dr Frank Polanco and Dr Albert Wong is greatly appreciated.

REFERENCES

1. Toossi, M., C.B. Callahan, and J.R. Peters. *Determination of AH-64 Airframe In-Flight Modal Frequencies Using Rotor Control-Induced Excitations*. in *American Helicopter Society Annual Forum*. 1994. Washington, DC.: American Helicopter Society.
2. Durno, J.A., G.R. Howland, and W.J. Twomey. *Comparison of Black Hawke Shake Test Results with Nastran Finite Element Analysis*. in *Annual Forum Proceedings - American Helicopter Society*. 1987. St. Louis, Missouri.
3. Twomey, W.J., et al., *General Method for Modifying a Finite-Element Model to Correlate with Modal Test Data*. Journal of the American Helicopter Society, 1991. **36**(3): p. 48-58.
4. Dunn, S.A. *Optimised Structural Dynamic Aircraft Finite Element Models Involving Mass, Stiffness, and Damping Elements*. in *International Forum on Aeroelasticity and Structural Dynamics*. 2001. Madrid, Spain.
5. Uhl, T., et al. *FE-Model Correlation of a Helicopter Using Ground Test Results*. in *Proceedings of the International Modal Analysis Conference -IMAC*. 1999.
6. Hermans, L., et al. *In-Flight Modal Testing and Analysis of a Helicopter*. in *Proceedings of the International Modal Analysis Conference - IMAC*. 1999.
7 Endo, H. and R. Randall, *Model Updating of Helicopter-Like Structure*. 2002, Defence Science and Technology Organisation Aeronautical and Maritime Research Laboratory, Australia.

Metrics Evaluation And Tool Development For Health And Usage Monitoring System Technology

Carl S. Byington, P.E.
Matthew Watson
Patrick Kalgren
Impact Technologies, LLC

Robab Safa-Bakhsh
The Boeing Company
Integrated Defence System

ABSTRACT©

Diagnostic algorithms are typically qualified on specific types of faults with limited test data and engineering judgment of intended applicability. Mechanical system ailments in helicopter gearboxes may vary from root cause conditions such as shaft misalignment; to fatigue events such as spalled bearings; to slower wear processes related to scuffing wear on gears. The accuracy of the fault detection and diagnostic processes for such a wide array of problems will not only depend on the algorithm's sensitivity to signal- to-noise ratio but also load level, failure type and flight condition. Diagnostics algorithms that are sensitive to faulted conditions yet relatively insensitive to confounding conditions are desirable for a broader range of application in helicopter health monitoring. Such generalized algorithms, though, may be less sensitive to early fault detection. In order to assess the risk associated with using certain diagnostic algorithms, qualify them for a range of use, and determine desirable thresholds to produce known false alarm rates, we need to evaluate the detection performance and diagnostic accuracy using established performance metrics.

The current paper describes the development of a Metrics Evaluation Tool (MET) to evaluate the performance and effectiveness of vibration features typically used in Health and Usage Monitoring Systems (HUMS). An overview of the candidate algorithms and method of metrics evaluation is provided. The MET is being demonstrated using a prototype database with seeded fault data from an H-46 aft transmission exercised with typical diagnostic algorithms for gear, shaft and bearing faults. The vibration feature results are analysed using probability of detection and false alarm metrics as well as diagnostic accuracy metrics. In addition, the effects of signal-to-noise ratio and threshold settings on the detection and false alarm metrics are evaluated within the tool. This will allow helicopter manufacturers and HUMS end users to evaluate effectiveness of an algorithm by statistical analysis of seeded fault test data. Moreover, it directly evaluates the risk associated with changing detection thresholds or relying on data containing varying noise levels. The MET tool is being developed under Rotorcraft Industry Technology Association HUMS Technology program. The tool along with Diagnostics database will provide an enhanced capability to assess and record the performance of vibration based diagnostics algorithms

INTRODUCTION

Rotorcraft Health and Usage Monitoring Systems (HUMS) monitor the condition and fault observables in critical mechanical components and systems by detecting abnormalities in processed sensor data and assessing fatigue life usage. Detection of faults in the drivetrain, engine, or rotor system is accomplished via vibration monitoring, oil debris monitoring, and exceedance monitoring of temperatures, pressures, shaft speed, torque, and strain. Usage Monitoring keeps track of time and cycle usage of aircraft, engine, Auxiliary Power Unit and rotor system. Besides OEM Structural Usage monitoring estimates the consumed life of dynamic components based on specific flight regimes and environmental conditions to which the helicopter has been subjected. The combination of fault detection and diagnosis coupled with reliable usage monitoring provides the basis for health and degradation assessments in HUMS systems.

Because drivetrain failures are the second leading cause of helicopter accidents, health monitoring of drivetrain systems can significantly increase safety (Byington 1997). A 1960's-70's study of primarily British rotorcraft experience showed that 22% of all airworthiness related fatal accidents on rotorcraft arise from transmission system problems [Vinall]. HHMAG studied 27 incidents in the UK involving increased levels of vibration and concluded that some 20 of these could have been prevented using a vibration health monitor as a maintenance tool only [HHMAG, 1999]. Drivetrain health monitoring is often considered more critical than engine health monitoring because there is no redundancy in the drivetrain system, and fracture propagation occurs more quickly (Aviation Week, 1993). Drivetrain health monitoring systems are designed to detect significant debris generation caused by surface fatigue faults (pitting, spalling) and vibration events related to material fractures (tooth breakage, shaft breakage) and poor operational conditions (misalignment, imbalance) that will lead to machinery damage. Accelerometers mounted on the casing record the vibrations of various helicopter components, and the vibration signal is usually preprocessed by signal averaging or using other signal conditioning techniques to eliminate unwanted signal energies before evaluating features or figures of merit. Algorithms make comparisons with test data by applying thresholds or seeking pattern differences that may indicate problems in the drivetrain or other helicopter systems. A typical process from energy to decision is shown in Figure 1.

Figure 1 HUMS Feature Processing Concept and Typical Functional Steps

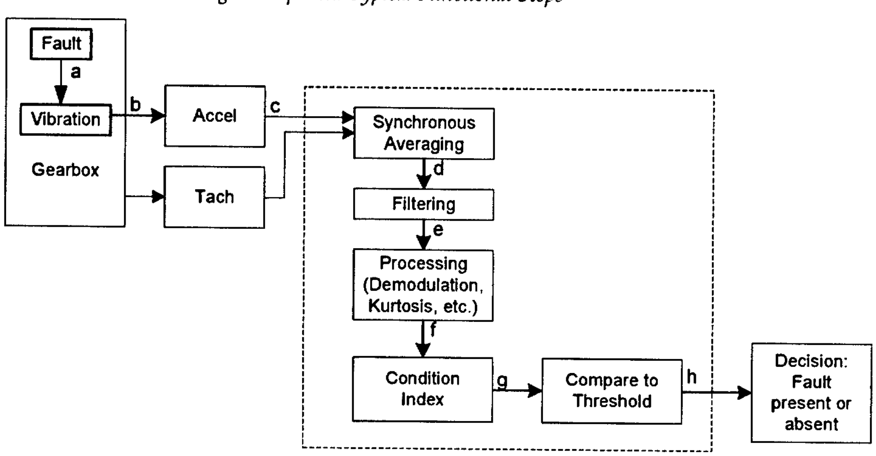

While most researchers and HUMS implementers are quick to acknowledge that drivetrain vibration monitoring is a valuable and effective part of HUMS, it is also felt that better assessment and correlation of algorithms with failure conditions is necessary. Better understanding of probability of detection and false alarm rates under various operating regimes is desired to reduce the HUMS vibration alerts, while maintaining a desirable detection rate. In addition to quantifying detection performance, more accurate diagnostics (localization and severity assessment) methods are needed. Significant data analysis and tools to evaluate the algorithms will be required to evolve from the current understanding to the desired state. The current work describes a step towards providing engineering tools to enable diagnostics algorithm performance analysis and assist in the maturation of HUMS.

HUMS FEATURES

Numerous vibration processing techniques have been identified for mechanical component fault detection in helicopter drivelines. Ritterbush (1997) provides an extensive summary of published techniques starting from Stewart's original paper (1977) to some of the more recent work by Zakrajsek (1989), Nicks and Krishnappa (1991), and Dousis (1994), Maynard (1999), Campbell (2000), McClintic (2000), Begg (2000), Lebold (2000), Reichard (2002) have led to a better understanding of the physics associated with some of the signal processing involved. In addition, Decker (2002), Dempsey (2001), and Mosher (2002) at NASA have refreshed technical interest in the features and their evaluation under varying load, speed and flight conditions.

A survey of the recent studies listed indicates over 40 commonly available with about 20 highly popular features that receive the most attention. Within the current work, we limited our consideration to that which was available in the RITA HUMS database set of algorithms developed by United Technologies Research Center. Rozak (2001) provides a list and description of these algorithms and the organization of the RITA HUMS Database. These are summarized in Figure 2.

Figure 2. UTRC Algorithms currently in Database

GEAR ALGORITHMS	DESCRIPTION
FM0	Scalar, ratio of peak-peak of signal average to sum of rms of mesh freq and gear harmonics
FM2A	Vector, ratio of kurtosis of envelopes after and before cross correlation with the mating gear.
FM4	Vector comprised of Scalars FM4A & FM4B 1) kurtosis of residue signal 2) ratio of power of signal average to residue
ESA_PP	Peak-peak of enhanced signal average (residual)
ESA_SD	Std dev. of enhanced signal average (residual)
RMS	Root Mean Square of average of 0 to 1.5 X max mesh frequency
SHAFT ALGORITHMS	**DESCRIPTION**
S01	The RMS of signal average at once per rev.
S02	The RMS of signal average at twice per rev.
BEARING ALGORITHMS	**DESCRIPTION**
Env_stat	Calculate statistical parameters on envelope
Envextr	Narrow band envelope extraction
Envp	Envelope of Hilbert transformation
Envpsd	Narrow band envelop power spectrum (hanning windowing)

TEST DATA

The seeded fault test data from a H-46 aft transmission, performed by Westland is used to evaluate the algorithms. The 22-second (100Ksamples/Sec) is parsed into 4-seconds segments to enhance the statistical analysis.

CLIENT-SERVER ARCHITECTURE FOR METRICS ASSESSMENT

The development of the RITA HUMS Metrics Evaluation Tool (MET) was motivated by the need to evaluate the performance and effectiveness of algorithms used to detect and diagnose faults. This first required the development of a suitable architecture to allow users to gain access, explore, calculate, and evaluate the data and information. We deemed the most useful form to be a client-server model for requests, processing, and viewing of results. The MET needed to allow its users remote access to the RITA HUMS database, which contains amounts of data much too large to be stored locally by the client application. Figure 3 illustrates the architecture and interface between the MET client and the database server. In addition to eliminating client-side data storage, users also gain the advantage of access to useful applications like MATLAB® and the Oracle 8i database, which is the relational database used within RITA HUMS.

TCL (Tool Command Language) was used to create the customized interface required to use existing algorithms, and new algorithms, from their storage in the database to their execution in engineering software (MATLAB®). The advantages of using TCL include its ability to communicate with a client-side application (like the current MET) using the TCP/IP protocol, as well as its easily available and highly usable interface to Oracle. TCL also allows the flexibility of interfacing with different algorithm platforms, including MATLAB®, C++, and Java. With this client-server design and application program, the complex data manipulations are managed on the server side of the application.

Figure 3. RITA HUMS MET Client-Server Architecture

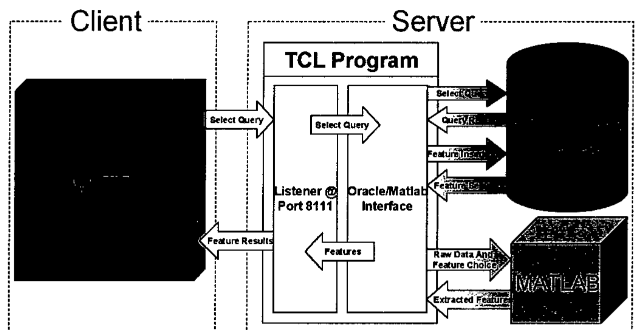

DATABASE DEVELOPMENT AND DATA EVALUATED

The RITA HUMS database (Rozak 2001) provide a common environment to allow an end-user (or client) a means to query available test data and algorithms. The database is structured as a client-server relationship between the Oracle database server and the primary Oracle form. The Oracle application is utilized as the database engine, because it can handle and store large binary objects (BLOB), provides an interface to external C++ coding, and provides comprehensive table/user security. Algorithms are stored in the database in a similar fashion as raw data. They are also treated as large-binary objects.

To place data or algorithms into the database, the user may use containers embedded into the Oracle Form or write SQL (Sequential querying language) scripts to automate the process of data upload. For the current MET demonstration, Impact Technologies populated and analyzed data from the Westland Seeded Fault Data Set (Cameron 1993). Lessons learned from this process are offered at the end of this paper.

Several conversions were performed on the data before it was added to the RITA HUMS database. The 22-second (100kSamples/sec) Westland data was parsed into 4-second segments to allow for additional analysis, and enhance the accuracy of algorithm assessment by the MET. To facilitate automated script upload, the subdivided Westland data was first entered into Microsoft Access using database population scripts written in Visual Basic. The final version of the Westland data resulted in 98,000 individual required field updates to the raw_data table alone in the RITA HUMS database. The magnitude of this effort clearly requires automated means to populate the database. Other relational fields in associated sensor and test tables also required updates to maintain both connectivity and database integrity. Sequential querying language (SQL) scripts were written to accomplish the task of entering this data in the RITA HUMS Oracle. A summary of this process is shown.

Figure 4. Westland Data Population Process Used

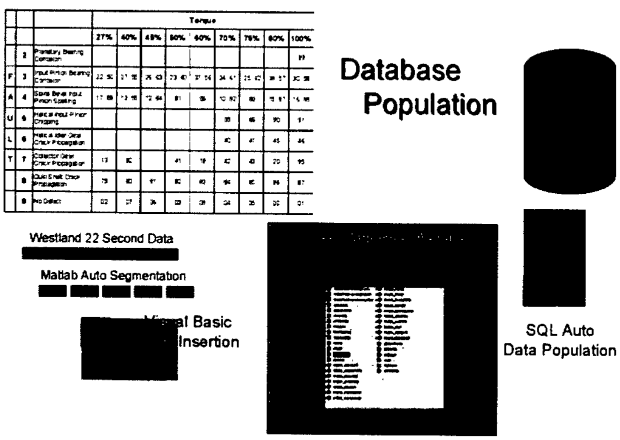

Within the design of the RITA HUMS database, the user has the option to either insert the data into the Oracle database directly as a BLOB or indirectly as a pointer to its storage location external to the actual RITA HUMS database. This flexibility does not prevent the data from being readily retrieved and used for subsequent analysis, and using pointers to external files was found to be much faster and more robust as the database size increases. Therefore the current implementation of the RITA HUMS database uses an external file referencing system. To complete the data storage, a file structure was created consistent with the identifying fields. The use of such an approach requires a logical assignment of the stored high-bandwidth data on the hard drive or like storage media. The current method uses a directory structure that incorporates the platform, subsystem, data, and test id information stored in the database:

../CH46-Chinook/Aft_Transmission/6_1_1992/Test_ID/Raw

This directory structure is consistent with the RITA HUMS fields in the database. This identification schema may need to be revisited as new data is organized into the database. The ability to automatically link the appropriate baseline data with faulted data and a nominal threshold with each data file are two additional future processing capabilities that aid metrics evaluation from the database.

METRICS APPLICATION

Statistical detection and diagnostic isolation are identified as key metrics for vibration algorithm evaluation (Dowling 2001).

Detection Metrics

The following Decision Matrix (Dowling 2001) defines the cases used to evaluate fault detection. It is based on hypothesis testing methodology and represents the possible fault-detection combinations that may occur.

Figure 5. Decision Matrix for Detection Evaluation

Detection Decision Matrix			
Outcome	Fault (F1)	No Fault (F0)	Total
Positive (D1) (detected)	*a* Number of defected 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
	a+c Total number of faults	*b+d* Total number of fault-free cases	*a+b+c+d* Total number of cases

From this matrix, the detection metrics can readily be computed. The Probability of detection given a fault (a.k.a. sensitivity) assess the detected faults over all potential fault cases:

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

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

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

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):

$$\text{Accuracy} = P(D1/F1 \,\&\, D0/F0) = \frac{a+d}{a+b+c+d}$$

Diagnostics Metrics

Diagnostic metrics are used to evaluate classification algorithms, typically consider multiple fault cases, and are based upon the confusion matrix concept. The matrix, as shown below, illustrates the results of classifying data into several categories. The table shows actual defects (as headings across the top of the table) and how they were classified (as headings down the first column). The shaded diagonal represents the number of correct classifications for each fault in the column and subsequent numbers in the columns represent incorrect classifications. Ideally, the numbers along the diagonal should dominate. The confusion matrix can be constructed using percentages or actual cases witnessed.

The probability of isolation (Fault Isolation Rate - FIR) is the percentage of all component failures that the classifier is able to unambiguously isolate. It is calculated using:

$$FIR = \frac{A_T}{A_T + C_T} * 100$$

$$A_T = \sum_i A_i$$

$$C_T = \sum_i C_i$$

Ai = the number of detected faults in component i that the monitor is able to isolate unambiguously as due to any failure mode (Numbers on the diagonal of the Confusion Matrix).

Ci = the number of detected faults in component i that the monitor is unable to isolate unambiguously as due to any failure mode (Number off the diagonal of the Confusion Matrix).

An alternative metric that is also useful is the Kappa Coefficient, which represents how well an algorithm is able to correctly classify a fault with a correction for chance agreement.

Figure 6. Confusion Matrix used for Diagnostic Metrics

Confusion Matrix – Showing Number of Correct and Incorrect Classifications					
Problem ⇨ / ⇩ Classified as	Collector Gear Crack Prop.	Quill Shaft Crack Prop.	Spiral Bevel Input Pinion Spalling	Input Pinion Bearing Corrosion	Total
Collector Gear Crack Propagation		1	1	0	7
Quill Shaft Crack Propagation	0		0	1	7
Spiral Bevel Input Pinion Spalling	1	0		0	6
Input Pinion Bearing Corrosion	0	0	1		7
Total	6	7	7	7	27

$$kappa = \frac{N(observed\ in\ agreement) - N(expected\ in\ agreement)}{N(total) - N(expected\ in\ agreement)}$$

where:

N(observed in agreement) = sum of diagonals in confusion matrix

N(expected in agreement) = sum{[sum of row)/N]*(sum of column)} for diagonals

N(total) = total number of observations

These metrics form the basis of top-level detection and diagnostic effectiveness. A challenge for the research community is to generate realistic estimates to these metrics given the currently limited baseline data and faulted data available. The basis for the statistical analysis pursued in this work is offered next.

STATISTICAL DETECTION CONCEPT

Statistical detection is based upon separability of features between the no-fault and faulted conditions. While a detection probability of 0.9 or higher is desirable, this may not always be possible. In the first place, POFA and POD are correlated because both are measured on the high side of the threshold. If the threshold is raised to decrease the probability of false alarm, the probability of detection is decreased accordingly.

This can be seen in the following figure, in which the signal for no-fault is on the left and faulted feature values is on the right side. The probability of detection can be seen on the right side of the threshold. The lower figure emphasizes the no fault (detection noise) distribution. POFA is also on the right side of the threshold. To decrease POFA, we raise the threshold (move it to the right), and this has the unfortunate effect of decreasing the probability of detection. It can be readily appreciated from this figure that moving the threshold to the right can decrease the probability of false alarm, but this will also decrease the probability of detection.

Figure 7. Relationship of Statistical Detection Metrics

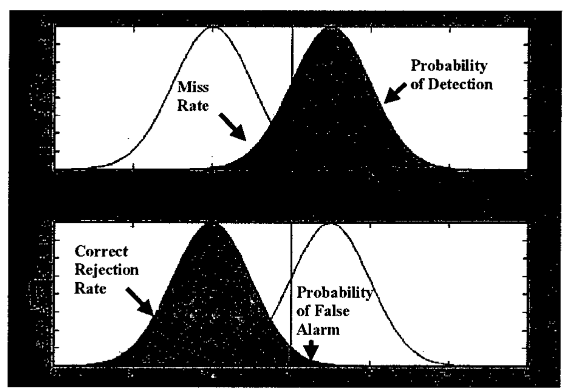

STATISTICAL ANALYSIS

The fault and no-fault conditions can be estimated using probability density functions (PDF). Obviously, with sufficient data these distributions can be identified directly. Within the current limited data environment, we implemented a Rician distribution with a PDF described by

$$p(x)=\begin{cases}\dfrac{x}{\sigma^2}\exp\left[-\dfrac{1}{2\sigma^2}\left(x^2+\alpha^2\right)\right]I_0\left(\dfrac{\alpha x}{\sigma^2}\right) & x>0\\ 0 & x<0\end{cases}$$

Where $I_0(u)$ is the modified Bessel function of the first kind and *zeroth* order and given by

$$I_0(u)=\frac{1}{\pi}\int_0^{\pi}\exp(u\cos\theta)d\theta$$

Given the limited data sets available for metrics assessment (for Westland we have 5 fault and no-fault each for each sensor), an estimate of the statistical variance is evaluated using a Chi-Squared approximation with user defined confidence bounds. This approximation of variance is represented by the interval:

$$\left(s^2\frac{v}{\chi_v^2\left(1-\frac{\alpha}{2}\right)},s^2\frac{v}{\chi_v^2\left(\frac{\alpha}{2}\right)}\right)$$

where s^2 is the calculated variance from the sample and σ^2, the true variance lies within the interval:

$$s^2\frac{v}{\chi_v^2\left(1-\frac{\alpha}{2}\right)}\le\sigma^2\le s^2\frac{v}{\chi_v^2\left(\frac{\alpha}{2}\right)}$$

The PDF for the Chi-Squared (non-central) distribution is given by:

$$p(x)=\begin{cases}\dfrac{1}{2}\left(\dfrac{x}{\lambda}\right)^{\frac{v-2}{4}}\exp\left[-\frac{1}{2}(x+\lambda)\right]I_{\frac{v}{2}-1}\sqrt{\lambda x} & x>0\\ 0 & x<0\end{cases}$$

where $I_r(u)$ is the modified Bessel function of the first kind and order *r*. Defined as:

$$I_r(u)=\frac{\left(\frac{1}{2}u\right)^r}{\sqrt{\pi}\,\Gamma\left(r+\frac{1}{2}\right)}\int_0^{\pi}\exp(u\cos\theta)\sin^{2r}\theta\,d\theta$$

Which has the series representation:

$$I_r(u)=\sum_{k=0}^{\infty}\frac{\left(\frac{1}{2}u\right)^{2k+r}}{k!\,\Gamma(r+k+1)}$$

Using a series expansion, the PDF can be restated as:

$$p(x)=\frac{x^{\frac{v}{2}-1}\exp\left[-\frac{1}{2}(x+\lambda)\right]}{2^{\frac{v}{2}}}\sum_{k=0}^{\infty}\frac{\left(\frac{\lambda x}{4}\right)^k}{k!\,\Gamma\left(\frac{v}{2}+k\right)}$$

Impact implemented the probability density functions given above using C++ code to perform computations. A numeric approximation of the zeroth order modified Bessel function was adapted from code in "Numerical Recipes in C". In order to simplify computation, the zeroth order Bessel function is used in the computation of the Chi-Squared PDF. The very small error introduced by this simplifying approximation is less than the expected finite arithmetic errors in the computation. This code module is embedded within the Metrics Evaluation Tool.

Figure 8. Visualization of Variance Estimate using Chi-Squared Confidence Interval

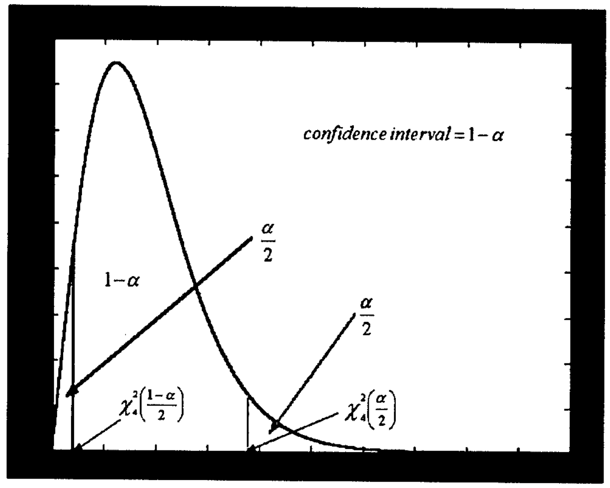

MET CLIENT APPLICATION AND FUNCTION

As described in the sections above, one of the strengths of the MET is its provision of convenient access to a central repository of data and algorithms to its users. This is accomplished through the use of a native TCP/IP interface to Oracle. The client side application interfaces with the user and directly computes the detection and diagnostic metrics from the server-returned feature results.

METRIC CALCULATION WINDOW

The current MET client application is in development and incorporates a tabbed GUI with five separate tabs/panels for a user to view. The primary panel to start the analysis is the Metrics Calculation. Data selection choices in this panel include algorithms, aircraft, systems and components. These selections lists are dynamically populated from Oracle SQL queries to RITA HUMS Database.

Figure 9. Metrics Calculation Panel

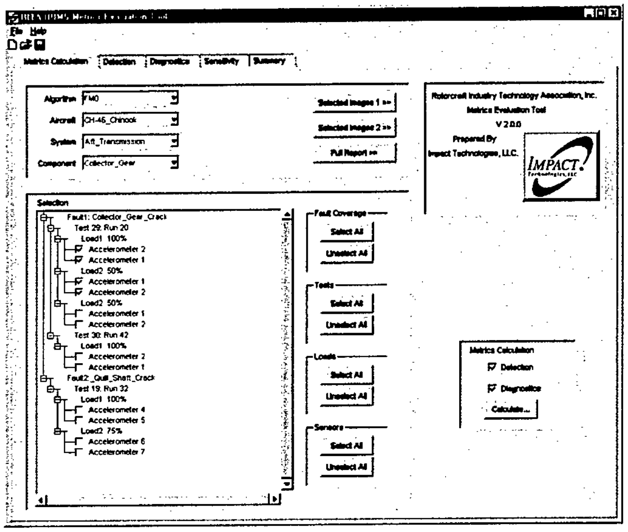

This screen also provides pictures of the system being analyzed, faults and full test report when available in the database. Within this tab, users can also select and unselect individual faults, features, loads, and sensors. The hierarchical choices are dynamically generated from available data in database. The user also chooses detection metrics, diagnostic metrics, or both, according to their preferences.

MET DETECTION WINDOW

Figure 10. Detection Results Panel

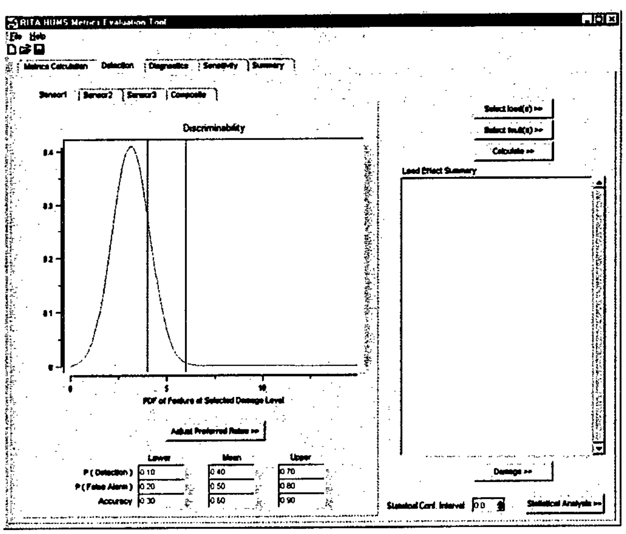

The MET Detection Window not only displays the feature discriminability plot at normal and faulted condition, but also displays results by individual sensor or as composite. Sorts can also be done by load and fault conditions. Within this window, thresholds are determined from false alarm criteria and the statistical confidence bounds provided by the user. The resulting Probability of Detection, False Alarms, and Accuracy are also calculated for the identified threshold.

The Diagnostics panel is still currently under development at the time of this writing.

MET SENSITIVITY WINDOW

Users can view a Receiver Operating Characteristic curve in the MET Sensitivity panel. The ROC plots the Probability of Detection vs. Prob. of False Alarm for a range of thresholds. In the figure shown, the blue lines represent the theoretical values based upon notional Rician distributions for various signal–to-noise ratios. The yellow represent the results for the data being analysed. The results are fictitious in the figure. Within this screen, the user can investigate the effects of additional noise on the data or changing threshold settings to better characterize performance. Additionally, the actual predicted POD and POFA are displayed.

Figure 11. MET Sensitivity Panel

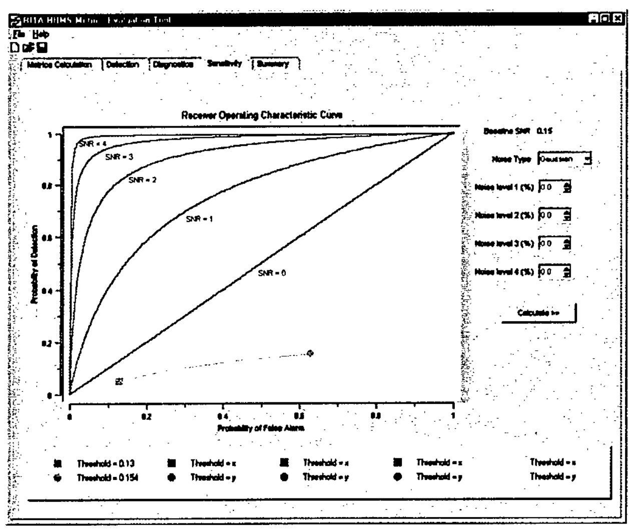

The Summary panel is the final tab in the user interface. A summary of the metrics evaluation project is displayed. Users are provided with "Save As" and "Print" buttons for their future review of the summary file. A user can also choose to save the entire project by using the icon at top left in the File menu.

Figure 12. Summary Panel

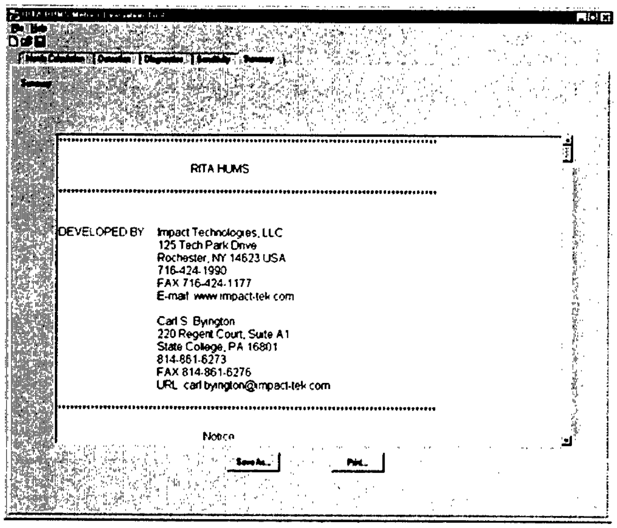

LESSONS LEARNED AND CHALLENGES

Availability of well-documented and easily implemented data continues to be an issue. Identifying standards for the raw formats and headers on these data sets and developing ways to automate the data population will greatly improve the likelihood of data making it into the database for use in Metrics Evaluation. For this specific application, an auto loader routine was developed, using Visual Basic to automate the entry of the Westland helicopter data into Access and its transfer into Oracle. These routines are specific to the Westland data population, and future expansion on the RITA HUMS database will most likely require examination of choices in developing a tool that can automate the entry of large amounts of collected data into the database.

Within the population of the database, a design decision was made to store references to the data within the database field as opposed to inserting the full vibration data files. The handling of raw data as binary large objects (BLOBs) can become problematic in a large and continual growing database. The primary problem comes from the vast size discrepancies that can exist between collections of sampled data. Allocation of storage by the database to accommodate the varying BLOB sizes and the subsequent indexing of this space can create unnecessary overhead, significantly decrease retrieval times, and produce potential corruptions in the database. The latter was chosen as the most efficient method for two reasons. First, storing the data internally as BLOBs would drastically increase the size of the database and possibly cause basic database operations (select, insert, update, and delete queries) to decrease in speed. Second, storing the data externally will aid in the rapid development of the application.

It should be noted that this path and location could be automatically generated from the fields entered within the database. Thus, this approach is amenable to an automated archival of the data on the storage media using a self-generated, standardized directory structure.

CONCLUSIONS AND FUTURE WORK

The current paper describes the development of a Metrics Evaluation Tool (MET) and associated implementation architecture to evaluate the performance and effectiveness of vibration features typically used in Health and Usage Monitoring Systems (HUMS). An overview of the candidate algorithms and method of metrics evaluation is provided. The MET is being demonstrated using a prototype database with seeded fault data from an H-46 aft transmission exercised with typical diagnostic algorithms for gear, shaft and bearing faults. The vibration feature results are analysed using probability of detection and false alarm metrics as well as diagnostic accuracy metrics. In addition, the effects of signal-to-noise ratio and threshold settings on the detection and false alarm metrics can be evaluated within the tool. This evaluation capability in such a metrics tool will allow helicopter manufacturers and HUMS end users to more directly evaluate the risk associated with changing detection thresholds or relying on data containing varying noise levels.

For the future, additional data and features should be added to the database. In addition, specifying the format of the raw data and associated header files is necessary to reduce the amount of work associated with populating data into the database. This will aid in the use of automated scripts to upload new data. Verification of the algorithms prior to using them as a baseline for new feature comparisons is also required. In addition, incorporating fusion and classification algorithms based upon the base features is necessary to achieve diagnostic and fault isolation evaluations. Ultimately, this tool could be adapted from its current engineering analysis use to others such as a ground station aid.

ACKNOWLEDGMENTS

This Tool is developed under RITA HUMS Technology project by Boeing Company. Impact Technology LCC is the prime contractor for this work. The support of other RITA OEMs especially Dr. James Rozak is gratefully acknowledged. Dr. Rozak from Sikorsky Aircraft is also acknowledged for his initial development of the RITA Database relational table structure. Researchers (Dr. Karl Reichard, Dr. Colin Begg, and Mr. Robert Walter) at the Penn State Applied Research Laboratory provided significant assistance with earlier phases of the project.

REFERENCES

Aviation Week and Space Technology, "Health and Monitoring System Enhances Helicopter Safety", (New York), v 139, n 12, p. 94, September 20, 1993.

Begg, Byington "Dynamic Simulation of Mechanical Fault Transition", 54th Meeting of the Society for Machinery Failure Prevention Technology (MFPT), May 1-4, 2000.

Byington, Nickerson, George "Prognostic Issues for Rotorcraft Health and Usage Monitoring Systems", 51st Meeting of the Society for MFPT, April 1997.

Byington, et al. "Evaluation of Human-Machine Interfaces for Aircrew Fault Diagnosis and Management", International Conference on Condition Monitoring at Swansea, April 12-15, 1999.

Cameron, "Final Report on CH46 AFT Transmission Seeded Fault Testing", Westland Helicopter Mechanical Research Paper RP907, September 1993.

Campbell, Banks, Begg, Byington "Diagnostic Feature Comparisons for Experimental and Theoretical Gearbox Failures", 55th Meeting of the Society for MFPT, April 2-5, 2001.

Campbell, Byington, Lebold "Generation of HUMS Diagnostic Estimates Using Transitional Data", 13th International Congress and Exhibition on Condition Monitoring and Diagnostic Engineering Management, December 3-8, 2000.

Decker, "Gear Crack Detection using Tooth Analysis", 56th Meeting of the Society for MFPT, April 15-19, 2002.

Dempsey, Zakrajsek, "Minimizing Load Effects on NA4 Gear Vibration Diagnostic Parameter", NASA TM-2001-210671, Glenn Research Center, Cleveland, OH, February 2001.

Dousis, "Development and Integration of Onboard Drive Train Vibration Monitoring Systems", Bell Helicopter Textron presentation to the American Helicopter Society 50th Annual Forum, Washington, D.C., May 11-13, 1994.

Dowling, Boeing Company, "HUMS METRICS, Monitoring the Monitor", Rotorcraft Industry Technology Association Final Report, January 3, 2000.

HHMAG, The Report of the Flight Deck Health Monitoring Indications Working Group, 87 pages. Submitted for Ratification by the *Helicopter Health Monitoring Advisory Group*, March 25, 1999.

Kay, S.,"Fundamentals of Digital Signal Processing", Volume II, Prentice Hall, ISBN 0-13-504135-X, 1998.

Lebold, McClintic, Campbell, Byington, Maynard "Review of Vibration Analysis Methods for Gearbox Diagnostics and Prognostics", 54th Meeting of the Society for MFPT, May 1-4, 2000.

Maynard, Byington, Nickerson, Van Dyke "Validation of Helicopter Nominal and Faulted Conditions Using Fleet Data Sets, International Conference on Condition Monitoring at Swansea, April 12-15, 1999.

McClintic, Lebold, Maynard, Byington, Campbell "Residual and Difference Feature Analysis with Transitional Gearbox Data", 54th MFPT, May 1-4, 2000.

Mosher, Pryor, Huff, "Evaluation of Standard Gear Metrics in Helicopter Flight Operation", 56th Meeting of the Society for MFPT, April 15-19, 2002.

Nicks, Krishnappa, "Evaluation of Vibration Analysis for Fault Detection Using a Gear Fatigue Test Rig", C404/022, IMEC E Proceedings, 1991.

Reichard, McClintic, Lebold, "A Comparison of Residual and Signal Extraction Techniques for Gear Fault Detection - Physics and Signal Processing", 56th Meeting of the Society for MFPT, April 15-19, 2002.

Ritterbush, "Gearbox Vibration Data Feature Extraction and Analysis for Condition-Based Maintenance", PSU MS Thesis, December 1997.

Rozak, RITA HUMS Diagnostic Database, RHDDB 00-06-1.1(4), Draft 4.0, Sikorsky Aircraft Corporation, March 2001.

Stewart, "Some Useful Data Analysis Technique for Gearbox Diagnostic", Stewart Hughes Ltd., pp. 18.1-18.19, 1977.

Vinall, P.D., "Airworthiness of Helicopter Transmissions," Civil Aviation Authority, *Royal Aeronautical Society Conference: Helicopter Transmissions*, Proceedings of the Symposium, London, England, 30 p., Feb. 6, 1980.

Zakrajsek, "An investigation of Gear Mesh Failure Prediction Techniques", NASA Technical Memorandum 102340 (89-C-005), November 1989.

COPYRIGHT AND TRADEMARKS

Physics of Failure-Based Approach to Life-Consumption Monitoring For Military Vehicles
Part 1 Life-Consumption Monitoring Process

Thomas J. Stadterman
Physics of Failure Team Leader, US Army Materiel Systems Analysis Activity

Alan Peltz
Physics of Failure Team, US Army Materiel Systems Analysis Activity

Roger Harris
Principle R&M Consultant, UK Defence Science and Technology Laboratories

Ian McNee
Technology Chief, Future Systems Technology, QinetiQ

William (Skip) Connon
Senior Mechanical Engineer, Automotive Core, US Army Aberdeen Test Center

Dr. K.K. Choi
Director, Center for Computer Aided Design, University of Iowa

Dr. Jeffery Freeman
Associate Professor, Mechanical and Aerospace and Biomedical Engineering Department
University of Tennessee

ABSTRACT

This is the first part of a two- part presentation on the results of the vehicles portion of The Technical Cooperation Program (TTCP) Joint Systems and Analysis Group, Technical Panel-1 Land Systems, Task-1-03. The overall objective of this task was to demonstrate that life-consumption monitoring could be used to improve the sustainment of military ground vehicles by identifying remaining useful life.

The physics-of-failure modeling of the trailer included the following: developed rigid and flexible-bodied Dynamic Analysis Design System (DADS) model of the trailer; developed a finite element analysis of the US Army trailer; performed various physical measurements of the US Army trailer for model input and verification; completed dynamic analysis of the Army trailer to determine acceleration and stress history for four terrain classes selected; and completed life prediction on the trailer drawbar (weakest part of trailer) for each of the four terrain classes.

Using the physics-of-failure modeling of the trailer, an algorithm was developed to calculate trailer damage per time on an equivalent terrain class using results from the trailer modeling effort. The algorithm also calculates the damage for each terrain and the percentage of damage each terrain contributes. This project has demonstrated the potential for life-consumption monitoring systems to improve system availability and sustainment.

INTRODUCTION

This paper is based on the final report for The Technical Cooperation Program, Joint Systems and Analysis Group, Technical Panel-1 Land Systems, Task 1-03 Physics of Failure Based Approach to Life Consumption Monitoring. This task demonstrated that life-consumption monitoring (LCM) could be used to improve the sustainment of military ground vehicles and munitions. For this project, the UK and the US jointly developed a LCM system for use on vehicles, while Canada developed a LCM system for use on munitions. The overall theme of this task was to demonstrate the principle that LCM could be used to identify remaining useful life and thus provide information needed to improve the sustainment of military equipment. This paper will only address the vehicle's portion of Task 1-03.

For the military ground vehicles portion of the project, a process was developed to calculate life consumed (or damage accumulated) on a critical component of a US Army trailer, based on real-time monitoring of the severity of the vibration experienced by the trailer. Specific steps for this process were developed and included performing a Physics-of-Failure (PoF) analysis to calculate damage of the trailer and designing a sensor system to determine the vibration severity experienced by the trailer. The UK developed a sensor system called Terrain Sensor System/Vibration Severity Sensor (TSS/VSS) which monitors the vibration severity of the trailer in real time, while the US completed the PoF modeling of the trailer. Using the PoF modeling of the trailer, the US and UK jointly developed an algorithm to calculate trailer damage per time on an equivalent vibration severity using results from the trailer modeling effort. Inputs to the algorithm included the time spent in each vibration severity category. For real-time monitoring of vibration severity, the TSS/VSS was mounted on the US Army trailer/vehicle. The vibration severity data were transferred to a remote computer, where software was developed to implement the damage algorithm.

BACKGROUND

Future military land systems require an increased level of reliability and availability to meet their operational requirements. Some new weapon systems have requirements to operate days without maintenance or re-supply during high-tempo missions. These requirements are commonly called Maintenance-Free Operating Periods (MFOP) and are similar to Failure-Free Operating Period (FFOP) requirements. These new systems will not only have to be very reliable, they will have to employ some type of prognostics systems to detect and repair impending failure before the mission. LCM is one type of prognostics system where the environment experienced by a system is monitored, then this information is used with failure models to forecast impending failure. These failure models, commonly called PoF models, would relate time spent in a particular environment to the amount of damage accumulated.

This project uses a PoF-based approach to LCM. PoF is a science-based approach to reliability that uses modeling and simulation to design-in reliability. This approach models the root causes of failure such as fatigue, fracture, wear and corrosion. Computer-Aided Design (CAD) tools have been developed to address various failure mechanisms and sites. An example of a failure mechanism is the fatigue cracking of electronic solder joints. PoF provides the ability to predict and improve the reliability before a system is built. The use of physics of failure leads to an increase in system reliability and a decrease in both acquisition and sustainment costs.

Designing for MFOP or FFOP is the latest initiative aimed at improving the reliability of land systems. This design philosophy, which promises to improve system availability and reduce system operating and support cost, is potentially very attractive to systems operators. However, to apply an MFOP/FFOP strategy effectively, an operator will need to know not only the relevant failure mechanism to be predicted by PoF based models, but also the actual in-service load conditions the system is experiencing. It is, therefore, necessary to develop a process methodology that will insure:

- The right operational and environmental data are being collected for PoF fatigue and wear models; and
- The proper parameters are being collected by the health monitoring systems, enabling accurate life predictions.

In order to apply a strategy for tracking life consumption, the relevant dominant failure mechanisms must be identified and applied in PoF models, with on-board monitoring systems utilized to track the relevant parameters.

APPROACH

The general approach for the Vehicles portion of this project was to define categories of terrain roughness (i.e., vibration severity) over which vehicles would actually be used, and to determine a metric to represent the mission scenarios. It was also necessary to develop a means to determine which terrain roughness a vehicle was on at any particular time, to determine the type of information required (e.g., miles, cycles and units) and to devise a method for information storage. Critical components to be monitored for life consumption had to be selected based on a combination of both inherent reliability and functional criticality. A major portion of the project involved a mechanical PoF modeling and simulation analysis of a vehicle system with a focus on identified critical components. An algorithm to determine component-level damage from mission scenario time, mile, or unit data recorded on the vehicle also had to be constructed. Finally, using information downloaded from the vehicle, the algorithm was used to estimate damage accumulation for the critical component, and to develop an estimate for remaining useful life.

ACCOMPLISHMENTS

A process was developed to calculate life consumed (i.e., damage accumulated) on a critical component of a military trailer, based on real-time monitoring of the severity of the vibration experienced by the trailer. A demonstration system was developed for a US Army trailer pulled by a military vehicle. Specific steps for this process were developed and included performing a PoF analysis to calculate damage of the trailer and designing a sensor system to determine the vibration severity. The UK developed the TSS/VSS, which monitors the vibration severity of the trailer in real time. The US completed the PoF modeling of the trailer, which included the following:

- Development of rigid and flexible-bodied Dynamic Analysis and Design System (DADS) models of the US Army trailer;
- Development of Finite Element Model (FEM) of the US Army trailer;
- Various physical measurements of the US Army trailer for model input and verification;
- Dynamic analysis of Army trailer to determine acceleration and stress history for four terrain classes selected; and
- Life prediction on trailer drawbar (weakest part of trailer) for each of the four terrain classes.

Using results of the PoF modeling, the US and UK jointly developed an algorithm to calculate trailer damage per time on an equivalent terrain class. For real-time monitoring of vibration severity, the TSS/VSS was mounted on the Army trailer/vehicle. The vibration severity data were transferred to a remote computer, where software was developed to implement the damage algorithm. A graphical user interface to facilitate data input to the damage algorithm was also developed. Time on each terrain is entered and the damage accumulated is calculated. The algorithm also calculates the damage for each terrain and the percentage of damage each terrain contributes.

LIFE CONSUMPTION MONITORING PROCESS APPLIED TO ARMY TRAILER

A general process for LCM of vehicle components was developed and applied to a US Army trailer system. The following paragraphs relate the generic steps to the specific application.

Step 1: Define categories of terrain roughness and vehicle speeds (vibration severity)

Four vibration severity levels were selected and related to actual physical test courses and potential vehicle speeds, as shown in Table 1. The four levels, in increasing order of severity, and damage potential, are 'smooth road', 'rough road', 'off-road' and 'X-country'. These levels are related to the following US Army Aberdeen Test Center courses, respectively, 3-mile straightaway paved road, Perryman Cross Country #1, Perryman Cross Country #2, and Perryman Cross Country #3. The chosen speeds for the courses were 50, 35, 25, and 15 miles per hour, respectively. Both x-y terrain data and PSD data for each course were available, except for the 3-mile straight course. This course only had PSD data, and x-y terrain data were computed from the PSD data.

Table 1. Vibration Severity Levels

Vibration Severity Level	1.1.1 Associated Terrain	**Vehicle Speed**
Smooth Road	3-mile straight course	50 mph
Rough Road	Perryman Cross Country #1	35 mph
Off Road	Perryman Cross Country #2	25 mph
Cross Country	Perryman Cross Country #3	15 mph

Step 2. Determine vehicle vibration severity in real time

To determine the vibration severity in real time, the UK developed the TSS/VSS. The TSS/VSS is comprised of an accelerometer mounted on the vehicle's axle, Figure 1, and an electronic box, Figure 2. The measured acceleration levels are amplified and filtered to determine which vibration severity level the vehicle is experiencing. A running total of time spent in each level is recorded.

Figure. 1 TSS/VSS Accelerometer on Axle

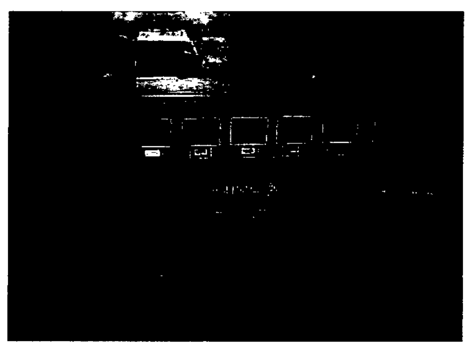

Figure. 2 TSS/VSS Electronics Box

Step 3. Determine what type of information is required from the vibration severity level (e.g., miles, cycles and units) and how to store this information.

It was determined that the time spent in each vibration severity level should be recorded and saved. The TSS/VSS output signal can be fed into a data collection device such as the

Datataker 600 data logger, VeMIS (commercial data logger), or the ADMAS (Aberdeen Test Center data logger), which will convert the time into a simple histogram of vehicle usage. Figure 3 shows the VeMIS system.

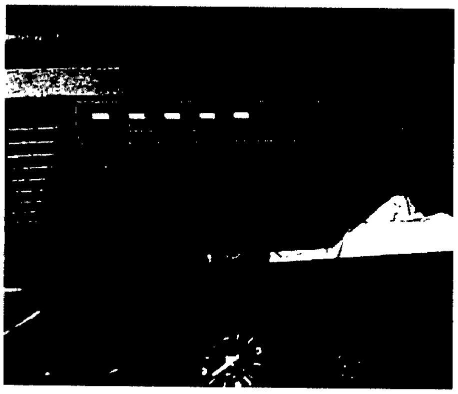

Figure. 3 VeMIS Data Logger

Step 4. Determine the critical components on the vehicle for which life-consumption will be monitored.

The selection of the critical components is a combination of what components are the least reliable and which components are functionally critical to the mission. A combination of general engineering analysis and past field and test data can be used for the selection process. For the US Army trailer, the drawbar was selected as the critical component because of past test data. The trailer's drawbar was known to fail while traversing cross-country terrain. Figure 4 shows a failure of the drawbar. A drawbar failure is also critical to the system and results in a safety risk. If no past data were available, an engineering analysis could be performed using the known environmental stresses (i.e., vehicle vibration, vehicle shock, water, and ground conditions). A simplified finite element analysis could have been performed to locate areas of the trailer with high stress.

Figure. 4 Drawbar Failure

Step 5. Perform mechanical Physics of Failure (PoF) analysis on the system with a focus on the critical components.

Objective of this step is to determine how damage is accumulated over time for the vibration severity levels. A fatigue analysis of the drawbar was completed using dynamic loads generated from computer models. This PoF analysis on the US Army's trailer drawbar was performed by the Army Mechanical Physics of Failure Team, which included the US Army Material Systems Analysis Activity (AMSAA), the US Army Aberdeen Test Center (ATC), the University of Iowa, the University of Tennessee, the US Army Tank-automotive and Armament Command (TACOM) and the Army Evaluation Center (AEC). The PoF analysis used solid models, rigid-body dynamic models, finite element models, flexible-body dynamic models, and fatigue models. The DRAW software program, developed by the University of Iowa, was used to generate life predictions based on inputs from the CAD models. More details of this analysis are provided in the following section.

Step 6. Develop an algorithm to determine component level life consumption.

The PoF modeling was used to simulate trailer life for each terrain/speed (vibration severity level). For each vibration severity level, a simulated time to failure was calculated. Table 2 shows the simulated times to failure. The TSS/VSS determines the time at each vibration severity level. Periodically, the information on the amount of time on each vibration severity level is downloaded or transferred from the vehicle to a computer. The percentage of damage accumulated for each vibration severity level is calculated by dividing the time spent in each vibration severity level by the trailer life (calculated from the PoF modeling) at that level. These percentages are added to determine the total life of the drawbar consumed over all levels. The algorithm for this calculation is as follows:

Life Consumed = SR/SR_Life + RR/RR_Life + OR/OR_Life + XC/XC_life, where:

SR = Time spent in smooth road vibration severity category;
RR = Time spent in rough road vibration severity category;
OR = Time spent in off road vibration severity category;
XC = Time spent in cross country vibration severity category;
SR_Life = Simulated time-to-failure in smooth road vibration severity category;
RR_Life = Simulated time-to-failure in rough road vibration severity category;
OR_Life = Simulated time-to-failure in off road vibration severity category; and
XC_Life = Simulated time-to-failure in cross country vibration severity category.

A running count of the percentage of damage accumulated can be stored in a computer. Maintenance personnel can download the data from the computer to determine if the percentage of damage is over 80 - 90 percent. If the damage is in this range, the trailer will be sent for inspection and repair. This simple algorithm could be expanded to address the age of the shock absorbers, the age of the axles, the average temperature, the cargo weight, and tire pressure.

Table 2. Life Estimates for Each Vibration Severity Level

Vibration Severity Level	Life (time to failure in seconds)
Smooth Road	12000000000
Rough Road	1512617152
Off Road	472247957
Cross Country	43162

Step 7. Estimate damage accumulations based on real time information.

A TSS/VSS was installed on a UK Land Rover and a US Army vehicle/trailer combination. A Visual Basic program was developed to calculate the percentage of life consumed and the percentage of life consumed in each vibration severity level. Figure 5 shows the LCM program. Manual transfer of information from TSS/VSS to the computer program was used, but any of the data logger devices could be used for remote transfer of data and information. Estimates of life consumption based on actual time were calculated.

The following sections summarize the PoF modeling approach, the experimental model validation and model input determination, and the TSS/VSS development.

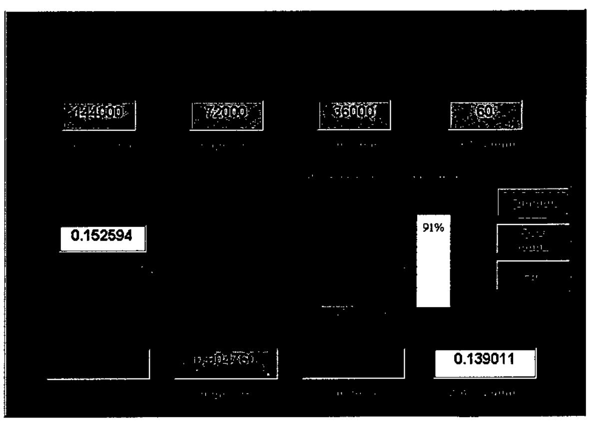

Figure.5 Life-Consumption Monitoring Program

Modeling Process and Models

This modeling and simulation effort involved several engineering disciplines, including:

- Computer-Aided Design (CAD) solid modeling;
- Finite element analysis for deformation and stress;
- Multi-body dynamic analysis for dynamic load history simulation;
- Fatigue life simulation for durability analysis; and
- Experimental testing for model validation.

The modeling approach began by using terrain data gathered by the Aberdeen Test Center for use in the dynamic models. The dynamic model used in this project was the multi-body model Dynamic Analysis and Design System (DADS). DADS was used for a rigid-body analysis and a flexible-body analysis using finite element analysis. The flexible-body DADS model was used to determine the dynamic accelerations at all points on the trailer during a simulated run of the trailer over a given terrain. NASTRAN was used as the finite element model in this project. Finally, the University of Iowa DRAW software tool was used to integrate results from dynamic modeling and finite element modeling to determine the stress history, strain history, and fatigue life. The following sections outline the process of how these modeling approaches were applied to the Army trailer. In each section, the modeling completed for the Army trailer will be addressed.

Computer Aided Design model in PROE

The process of computationally determining the fatigue life of a mechanical system started by developing a three-dimensional parameterized CAD model. In this case, the CAD model, shown in Figure 6, was developed using the Pro/Engineer software. Mass and material properties were assigned to each part of the CAD model, and comparison was made to the physical model. This comparison included the total mass, center-of-gravity location, and total inertia. It is very important that the CAD model approximates the physical vehicle as closely as possible, since both the finite element and multi-body dynamics computational models will be derived from it. The CAD model also formed the basis for consistency for these other models.

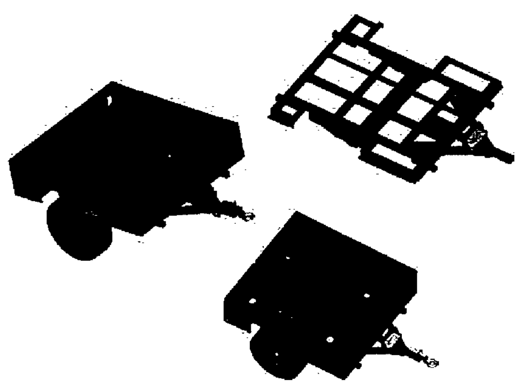

Figure. 6 ProE Solid Model of Trailer

Finite Element Model (FEM) in Hypermesh/NASTRAN

In this PoF approach, the finite element model was used for several purposes. The first purpose was to evaluate the vibrational modes, which were used to create the flexible multi-body dynamics model. These simulated modes were compared to an experimental modal analysis survey of the physical trailer to validate the analysis. The second purpose of the finite element model was to evaluate deformations of the trailer, which are used as static correction modes in the flexible multi-body model. The third purpose was to evaluate the stress and strain under dynamic loading conditions, which was used to evaluate the fatigue life of each component. For the Army trailer, a finite element meshed representation of the trailer, shown in Figure 7, was developed from the CAD model by exporting the frame and cargo box as an Initial Graphics Exchange Standard (IGES) representation. The IGES model was meshed using the HyperMesh software.

The vibrational response of the trailer was determined by performing an eigenmode analysis using NASTRAN. The first eigenmode analysis was done unconstrained and an experimental modal analysis was completed approximating an unconstrained trailer. This allowed for a direct comparison of the experimental and analytical mode frequencies and shapes. The second eigenmode analysis of the trailer was completed by constraining the positions where the trailer body is attached to the lunette and the axle assemblies. Three static deformation mode shapes were also determined from this constrained model by sequentially applying a unit force to each constrained location. The combination of the constrained eigenmodes and static deformation modes form the basis of the Craig-Bampton mode set which is applied by the flexible multi-body dynamics model.

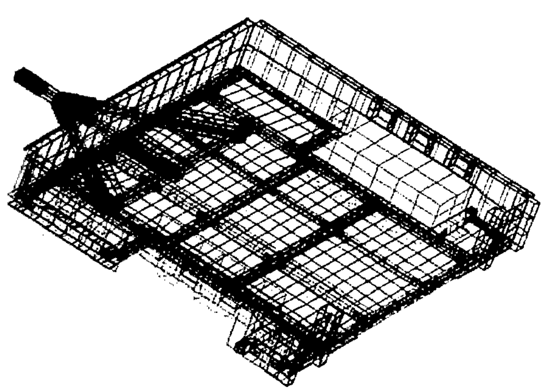

Figure. 7 Finite Element Model of Trailer

Dynamic model in Dynamic Analysis Design System (DADS)

The initial, rigid-body multi-body dynamics model was developed, using the DADS software, directly from the CAD model, and thus had matching mass and inertia properties. It should be noted that the rigid-body DADS model was a preliminary step needed to create the DADS flexible-body model. Output from the DADS flexible-body model was used as input into the DRAW analysis tool. Dynamic properties for the shock absorbers, axles, and tires were required to create a workable model. Experimental tests were undertaken at ATC to determine these properties. Once these properties were experimentally determined, the rigid-body model simulated the trailer assembly as being dropped from a short height above ground and allowing it to dynamically settled. This allowed for comparison of the ride heights and final center-of-gravity locations between the analytical and physical trailers.

The rigid-body, multi-body model was then combined with an existing rigid-body model of the Army tow vehicle, as shown in Figure 8. Using terrain data supplied by the ATC, the combined trailer/vehicle model was traversed across a simulated test course at constant speed. This enabled the evaluation of any simulation difficulties prior to complicating the trailer model with the addition of flexibility effects.

The flexible-body dynamics model differs from the rigid-body model only in that the frame and cargo box of the trailer were made flexible by incorporation of the Craig-Bampton mode set.

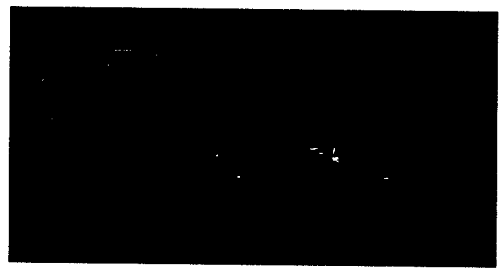

Figure. 8 Combined Vehicle and Trailer Multi-body Dynamics Model

As mentioned previously, this involves the solution of the constrained finite element eigenmode problem and a series of deformation modes, which are referred to as the static correction modes. The DADS software uses an assumed-modes approach to flexible multi-body dynamics. As such, it is left to the user to determine which selection of eigenmodes and static correction modes are represented in the model. In this case, the eigenmodes with the three lowest natural frequencies were selected, along with three static correction modes. The static correction modes helped to compensate for higher frequency, low energy, modes that were not modeled. In order to be useful, the six modes need to be orthogonalized prior to use in simulation. Then the flexible-body trailer frame/cargo bed was substituted for what was previously a rigid-body in the DADS simulation. Proper selection of flexible-body modes is essential when trying to calculate dynamic load history. An identical simulation to the previous rigid-body model simulation was then run, which captured the load history at the axles and lunette and also the distributed inertial loading for the flexible body.

Fatigue analysis in Durability and Reliability Analysis Workspace (DRAW)

The University of Iowa Center for Computer-Aided Design (CCAD)-developed DRAW software was used to determine the durability of the trailer, based on its dynamic load history calculated from the DADS model. The dynamic load history, determined through simulation, was applied as the boundary condition to the finite element model to quasi-statically determine the stress/strain history of the trailer frame and cargo box. Figure 9 shows the flow of information within DRAW. In the DRAW software, the dynamic stress/strain analysis procedure coupled the nonlinear gross motion with the linear elastic deformation using a new quasi-static hybrid method. The procedure estimated multi-axial, elastic-plastic stress-strain history in selected local regions. Two well-known methods, the Neuber method and the Glinka method, which are both based on matching known elastic and unknown elastic-plastic solutions in terms of energy, were used in DRAW to obtain the dynamic strain history.

The DRAW software then used a rain-flow counting procedure to determine the cyclic loading. Since the dynamic stress history contained a very large amount of data, with variable amplitudes, it was necessary to use peak-valley editing before the fatigue life could be determined. This editing procedure removed small cyclic changes which would appear as straight lines on a stress-strain diagram, while leaving those cycles which result in closed hysteresis loops. These cycles were then transformed into a number of constant amplitude strain histories which were used in the fatigue life prediction process.

The DRAW software used the Palmgren-Miner linear damage summation law (Miner's Rule) to assess the damage caused by each loading cycle. Miner's Rule simply states that failure will occur when the summation of the damage caused by individual cycles exceeds unity.

For the trailer, a von Mises strain approach was initially used to determine "hot-spots," or those finite element nodes which have the lowest fatigue life in a given region. Around these nodes, the finite element mesh was refined. From this refined finite element model, an updated flexible-body DADS model was redefined and simulated in order to calculate the proper inertial loading for the dynamic stress computation procedure. DRAW was then used to apply a critical plane method (either tensile-based or shear-based) to more accurately assess the fatigue life in those regions. This two-stage procedure was required, since it was not feasible to simulate or to analyze the whole trailer using a fully refined finite element mesh.

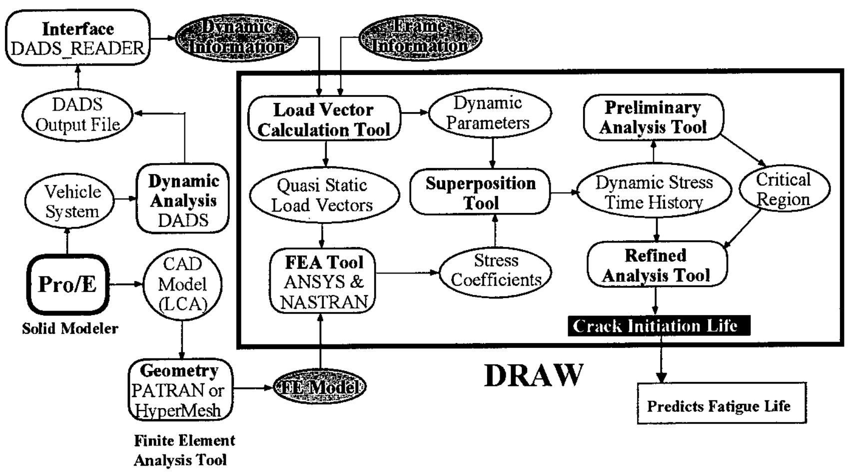

Figure. 9 Flow of Information within DRAW

Results of the DRAW analysis for the Army trailer are summarized in Table 3. The data show the predicted miles which can be driven prior to crack initiation at the most critical node in the model, and are based on a spherical joint between the trailer and the towing vehicle.

Table 3. Predicted Life from DRAW Analysis

COURSE	SPEED (mph)	LIFE (miles)
Perryman 3	**15.0**	**180**
Perryman 2	**25.0**	**3279500**
Perryman 1	**35.0**	**14706000**
Belgian block	**25.0**	**2181900**
Paved road	50.0	unlimited

A fatigue life contour of the trailer is shown if Figure 10. The color white is used depict the areas of lowest life, and indicates an area of concern on the trailer hitch. Areas of white on the cargo bed are related to the test payload, and as such, are not indicative of field fatigue failures.

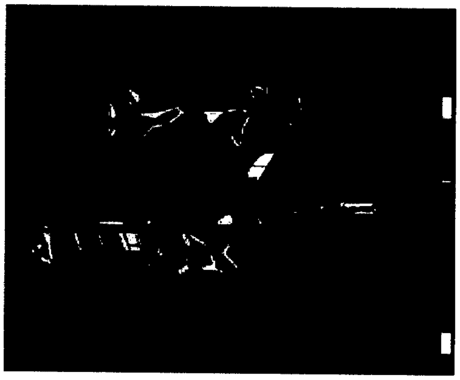

Figure 10. Fatigue Life Contour of Army Trailer

Summary of Experimental Validation

In this study a physical trailer was available to test and compare to the virtual prototype models, which provided a distinct advantage. In general, when applying any computational tool, model verification and validation needs to be incorporated into the overall development and modeling scheme. Rather than decrease the amount of testing required to field a system, the virtual prototyping process and, in particular, computational durability analysis focuses the types of experimental tests to be performed and the types of instrumentation involved in those tests. In the case of the trailer involved in this study, the experimental tests included the following:

- Profilometer measurements of the test courses used for terrain input for DADS;
- Physical measurements of the trailer components which differed or were not included in the blueprints;
- Mass/inertia measurements of the trailer, its subassemblies, and its individual parts;
- Dynamic tests of the trailer spring rates, shock absorbers and tires;
- Modal analysis tests of the trailer frame and cargo box; and
- Instrumented tests of the vehicle/trailer combination on the test courses used in the simulations.

The US Army ATC performed all tests, and each was used in a different phase of the modeling effort.

Physical Measurements

The physical measurements and the mass/inertia measurements were used to validate the CAD model. Validation involved matching the component masses, the overall center-of-gravity location, and the yaw moment of inertia of the trailer. The moments of inertia of the wheel and axle assembly and the entire vehicle were determined using a pendulum platform as shown in Figures 11 and 12.

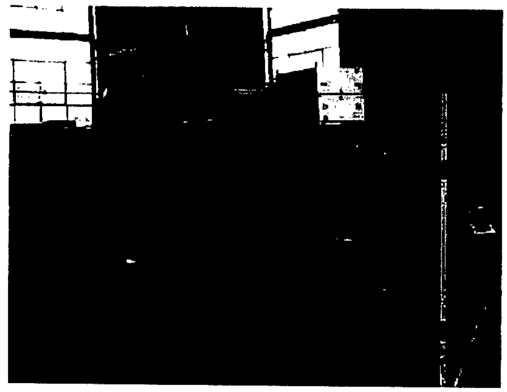

Figure. 11 Moment of Inertia Measurement of Wheel and Axle Assembly

Figure. 12 Yaw Moment of Inertia Measurement of Entire Trailer

Since the other computational models were derived from the CAD model, these properties were required before any additional modeling could proceed. At this stage, it was found that the largest number of questions and uncertainties dealt with items such as the tires and hitch assembly, which were not unique to this trailer. The tow vehicle used the same wheels and tires.

Once the CAD model of the trailer was completed and validated, development of both the finite element and multi-body dynamic models proceeded. Each model required different data and validation procedures. In particular, the multi-body dynamic model required data for components which could not be represented with strictly geometric information. Spring and damping rates were required for both the trailer suspension and tires. Stiffness of both the suspension and the tires were determined independently by static force-deflection measurements. Damping and natural frequency of both were determined by a quick release drop method. Cycle counting of the resultant displacement time history was used to compute natural frequency, and a log decrement and exponential curve fit technique of the same time history was used to compute the damping ratio. The trailer suspension was rendered inactive during the tire measurement phase, and tests were conducted at tire pressures of 17 psi (trailer inflation pressure), 26 psi, 35 psi (vehicle front tire inflation pressure) and 40 psi (vehicle rear tire inflation pressure). A typical decay time history is shown in Figure 13.

In addition to drop tests for tire natural frequencies, rolling tire tests were performed on the Belgium Block test course at four different speeds and at four tire pressures.

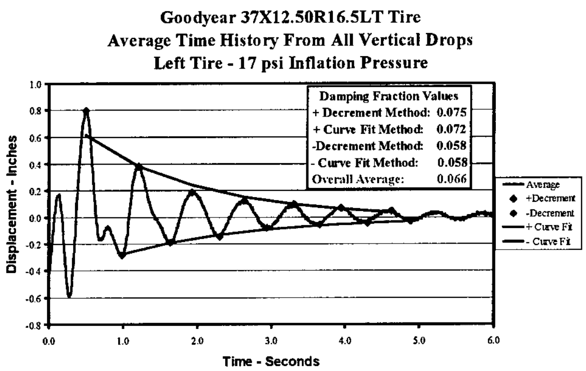

Figure. 13 Determination of Damping Ratio from Decay Time History

Mode Frequencies and Shapes

While the finite element model could be developed strictly from the material property and geometric information, validation of the finite element model required an experimental modal test. The experimental modal analysis was conducted using an instrumented hammer-impact technique. Response was measured at 54 locations on the frame under the trailer frame/bed assembly. The trailer was tested with the wheels and axle removed (frame and bed only) and was suspended from bungee cords connected from the trailer lifting points on each corner to an overhead frame. This configuration represented an unconstrained set of boundary conditions for the trailer. A photograph of the basic setup is shown in Figure 14, and the geometrical distribution of response measurement locations is shown in Figure 15. The first three modes from the experimental procedures and the FEA showed close agreement in shape and frequency.

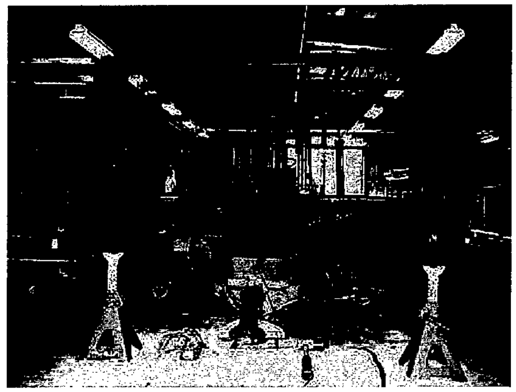

Figure. 14 Experimental Modal Analysis Setup Showing Bungee Cord Suspension and Impact Hammer

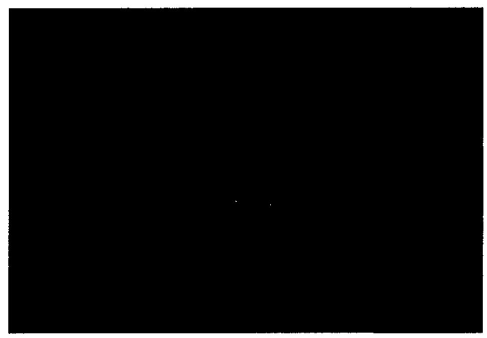

Figure. 15 Experimental Modal Analysis Response Locations

Strain Measurements

Final validation was accomplished with instrumented, on-course testing. This included both accelerometers and strain gauges at locations determined by the computational analysis. The test course used for both the computational and physical testing is more than two miles long. Vertical elevations were measured by a profilometer at 3-inch increments to provide a digital representation of the course. Test course elevation data were presented as vertical elevation as a function of course distance and as a wave number spectrum (power spectral density in the spatial domain). An example of test course elevation measurement data is shown in Figure 16. Results from the flexible multi-body dynamic analysis can be validated with the accelerometer data, while the dynamic strains determined by the DRAW program were validated using the strain gauge data.

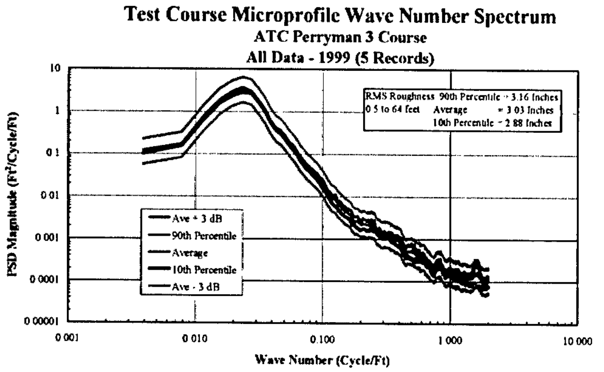

Figure. 16 Test Course Wave Number Spectrum

Applications and Military Benefits

The use of LCM will help enable systems to meet their MFOP requirements and increase mission reliability and availability. Without LCM, systems will not be able to meet the high reliability and availability requirements because that level of reliability cannot be designed in economically. The systems developed in this project provide examples of how to implement LCM for vehicles and munitions.

For military vehicles, the LCM systems can be used to track the damage accumulated (life used) on the dominant failure mechanisms. The life of particular components (such as the drawbar in the trailer example) can be tracked and when the life used exceeds 80 to 90 percent, the component can be replaced. By replacing the components before they fail, the systems remain operational and their availability is increased. Another aspect of LCM is that a system's critical components can be checked before the system goes into a high op-tempo mission. Particular systems can be selected that have components with the most life remaining, which will enable the best vehicles to participate in a critical mission.

The growing demand for improvements in system reliability, rapidly growing system complexity, and the users need for better confidence in the operational availability of systems at any given time, have all resulted in the need for more knowledge of cumulative in-service use than has traditionally been available. The data collection and monitoring systems developed by this project address these needs. They will provide output data applicable to both current operations and to future design programs, and will insure that the right operational and environmental data is being collected for PoF based fatigue models. They will also insure that proper parameters are being monitored thus allowing accurate prediction of life usage.

Conclusions

This project demonstrated that LCM systems could be employed onto military vehicle and munitions systems. LCM systems provide a way of prognosticating failures, so they can be repaired before a weapons system starts a mission. The use of LCM systems will enable systems to meet their MFOP requirements and increase mission reliability and availability.

LCM systems require analysis of the failure mechanisms to determine the stress-life relationship. The vehicle LCM system was based on a detailed PoF analysis of the US Army trailer. This vehicle modeling effort integrated existing tools to demonstrate a simulation-based methodology to assess fatigue of vehicle components. This methodology will provide the designers of vehicles the tools and methodology to evaluate fatigue early on in the design process. Early identification and correction of fatigue problems will reduce the need for equipment redesign, and reduce the overall maintenance costs of the vehicle. It was also learned in this modeling effort that physical tests and measurements were critical for the analysis portion to obtain material and system characteristics and for model validation. The munitions LCM system used probabilistic failure mechanism models for increased accuracy in predicting impending failures. Probabilistic failure mechanism models require a greater understanding of the failure process. The ability to carry out a probabilistic analysis will provide a broader perspective on the problem.

Recommendations

The design of the LCM systems should be started early in the system design. Life-consumption monitoring systems will need to be integrated in the weapons system because they require sensors, data storage devices, data transfer, and possibly computational software. Since PoF analyses are a critical part of LCM design, and will produce more reliable products, the analyses should be routinely completed on new designs. The use of PoF analysis will determine the dominant failure mechanisms and sites, the stresses that affect these failure mechanisms, and the stress-life relationship. Because it would be impossible to implement a LCM system for every failure mechanism, it is recommended that the dominant wearout failure mechanisms be targeted for LCM systems.

REFERENCES

DRAW, Durability and Reliability Analysis Workspace, Center for Computer-Aided Design, College of Engineering, The University of Iowa, Iowa City, IA, **1994**.

DRAW Concept Manual, Center for Computer-Aided Design, College of Engineering, The University of Iowa, Iowa City, IA, **1998**

COPYRIGHT AND TRADEMARKS

Physics Of Failure (Pof) Approach To Life Consumption Monitoring (Lcm) For Military Vehicles
Part 2 Sensors Used For Terrain And Vibration Severity Monitoring

Roger Harris
Principal R&M Consultant, UK Defence Science and Technology Laboratories

Ian McNee
Technology Chief, Future Systems Technology, QinetiQ

ABSTRACT©

This is the second of two Papers that summarise The Technical Cooperation Programme (TTCP) Joint Systems and Analysis (JSA) group Technical Panel (TP)-1 Land Systems Task-1-03 which reported in June 2002[1]. The work conducted by the US Army Materiel Systems Analysis Activity (AMSAA) on PoF analysis is described in Part 1.

This paper describes the work done by the UK to provide a measure of the vibration experienced by a vehicle in use that could be used to support the PoF analysis in the LCM process. Vibration data was collected in real time in a way that was affordable to military vehicle programmes through a Terrain Sensing System (TSS) developed by QinetiQ. The TSS converts the raw vibration data into simple counts of the time period spent on each terrain surface of different roughness in order to determine the damage to the vehicle being accumulated over time for each surface.

INTRODUCTION

Task-1-03 aimed to demonstrate that Life-Consumption Monitoring (LCM) could be used to improve the sustainability of military land vehicles and munitions. The USA and UK undertook the implementation of LCM on land vehicles and developed the 5-step concept shown in Figure 1. The usage, representing the operational duty cycle, combined with detailed fatigue life predictive information obtained in the PoF-based modelling and simulation, is used in a simplistic cumulative damage algorithm to provide life consumption information. The process can be used in the virtual mode; for design evaluation as an aid in optimising expected life in future vehicle design, or in the logistic mode; using real usage data as an aid to in-service vehicle fleet management. The development of the suite of PoF analysis tools used in Steps 4 and 5 was led by AMSAA using an original concept developed by the Computer Aided Life Cycle Engineering (CALCE) group at the University of Maryland[2]. The tools that comprise the suite, their use and the 5-step process are all described in detail in Part 1. For completeness, the 5 Steps are summarised below.

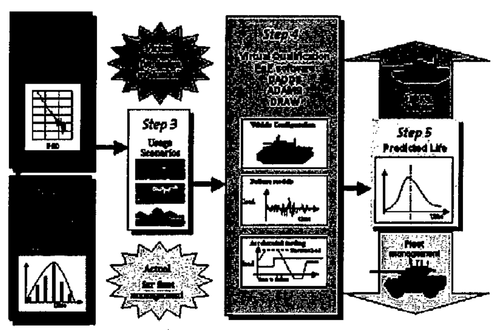

Figure 1. Vehicle Life Consumption Monitoring Concept

This Paper describes the work conducted in the UK to develop a cost effective way of achieving Steps 1 to 3 and goes on to discuss the requirements of a Health and Usage Monitoring System (HUMS) for military vehicles and the importance of data reduction in making the process cost-effective. The 5-step vehicle LCM concept comprises:

- Step 1: identify the time period spent by the vehicle on different terrain surfaces;
- Step 2: map the terrain surfaces to known acceleration power spectral density (PSD) signatures of the surfaces;
- Step 3: define the "usage spectrum" in terms of expected (virtual use) or actual (real usage monitoring) time and speed on the different terrain surfaces;
- Step 4: using the suitably modified outputs from Step 3, perform the analysis using the PoF modelling and simulation software tools;
- Step 5: apply the life consumption algorithm to determine the life consumed/life remaining.

BACKGROUND

Expectations are that the future battlefield environment will be of a higher tempo requiring fewer vehicles to operate for longer periods with less logistic support. This will increase the criticality of failures and place a premium on the availability of good Reliability and Maintainability (R&M) and usage data. Future vehicles will be required to operate for specified periods of time without maintenance - Maintenance Free Operating Periods (MFOPs). These MFOPS will be a challenging requirement that will require a paradigm shift in the approach to maintenance from reactive to proactive as shown in Figure 2.

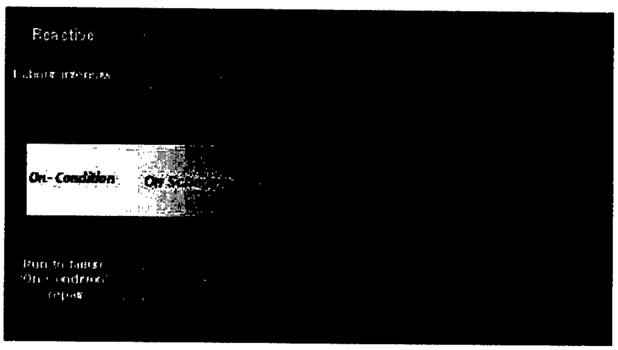

Figure 2. Evolutionary Approach to Maintenance

Furthermore, in the Digitized Battlespace (Land) (DB(L)) vehicles will be expected to operate in an information-rich environment that will require considerable improvements in the timely acquisition, processing, distribution and presentation of **all** information, tailored to the needs of battlefield commanders, weapons systems operators and support staffs. In summary, future systems will not only have to be very reliable, but will also need to employ prognostics to detect impending failures in order that the appropriate action can be taken before a mission begins and to be capable of distributing the associated information efficiently and effectively. The raw data needed to support prognostics would normally require high data sampling rates and large data storage facilities; factors that potentially could drive up the costs to levels beyond the affordability of military land vehicle projects. Consequently, data collection must form part of an overarching top-down strategy that not only fully exploits the data reduction techniques that are addressed later in this paper, but also ensures that data is only collected if a 'customer' for the data can be identified.

PoF-BASED RELIABILITY PREDICTIONS

Traditional methods of estimating reliability, which rely very heavily on statistics, are insufficiently robust to support the accuracy necessary for on-board prognostics. Moreover, predictions based on historical failure rates are only valid for extrapolation into the future as long as the usage patterns remain unchanged: a change in usage pattern is likely to change the failure rates and might even change the dominant failure modes. The PoF concept offers a new scientific approach to reliability that, at the design stage, identifies dominant failure modes as well as the associated failure mechanisms and the environmental exposure that will cause the failure. This provides an opportunity to monitor the exposure of systems to the critical environment in actual use, determine how the damage being accumulated is consuming life and predict the time to first failure.

LOW-LEVEL DATA COLLECTION

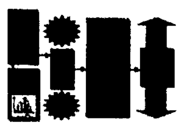

The first step is to identify, in real time, the terrain surface that a vehicle is travelling over in a way that can be related to the loads and stresses being imposed. This on-board measurement is termed the "low-level data collection" input to the LCM process and represents real vehicle usage. The process can be likened to monitoring engine speed by choosing revolution bands and monitoring time in each band. Early work under the UK MoD Applied Research Package (ARP) 26b[3] on the availability of Combat Support Vehicles (CSVs), showed that, although the practice of monitoring vibration loads in service to manage the fatigue lives of aircraft was well established, the high data sampling rates and large data storage requirements drove costs well above the levels that would be acceptable for land vehicles. Mainly to facilitate ease of maintenance in the field, UK Army vehicles were rugged and simple, and there are no automatic data collection systems in service. Commercial-off-the-shelf (COTS) data collection systems have generally been developed around highly predictable and benign commercial usage patterns which do not readily translate to military use. Furthermore, only the most sophisticated COTS systems could handle high data sampling rates, and none were designed to collect terrain data. As a result, a Terrain Sensing System (TSS) was developed for the TTCP LCM project, using a standard beam-axle Land Rover 110.

Military vehicles display highly individual operating profiles that defy any attempt to be represented by a single vehicle in a fleet. Consequently, the ultimate goal of LCM is to instrument **all** vehicles in a fleet; therefore, the TSS is designed to be as simple and as inexpensive as practicable.

Figure 3. TSS Accelerometer Mounted on Land Rover Axle

The TSS comprises an accelerometer that is mounted on the axle of the vehicle - see Figure 3 - and connected by a cable to an electronic box mounted inside the vehicle. The electronic box processes the raw acceleration signals induced by the terrain surface on the vehicle and simplifies them into discrete values. These values are grouped into bands representing terrain surface types and passed to an on-board COTS data-logger for storage and download when required. It is important to note that, while the indicated "terrain surface band" will vary with vehicle speed, the TSS does not directly measure the actual terrain surface profile, surface roughness, or vehicle speed but measures the axle accelerations induced as a function of all three factors. In the development of the TSS, considerable thought was devoted to data simplification to enable processed/filtered data to be retrieved in histogram form. This approach only requires five cumulative registers to store data and removes the need to store real-time vibration data. For the purposes of the TTCP task, the five bands defined in this paper were used but later versions of the TSS have seven channels for use in supporting vehicle trials when the TSS is used to establish a usage profile from the trial plan. Histograms of the actual usage of the vehicles during the trial are then collected for comparison to ensure that the trial plan is being followed faithfully.

HIGH-LEVEL DATA COLLECTION

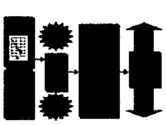

The simplistic "time on terrain surface band" measure of usage which is collected by the TSS needs to be converted to a form that can be used in vehicle damage accumulation calculations. In the case of the TTCP LCM process, this takes the form of acceleration PSD, which can be used to drive the PoF-based failure models. The PSD function reorganizes the complex combination of waveforms that make up a terrain surface profile and roughness signature and presents the results in the frequency domain. A PSD can exist in displacement or acceleration format and is the preferred convention for characterizing terrain surfaces.

The simple TSS output is mapped to the actual TSS terrain surface bands derived from high-level data. Figure 4 shows how low-level data inputs can be linked to road surface profile standards or calibrated surface data.

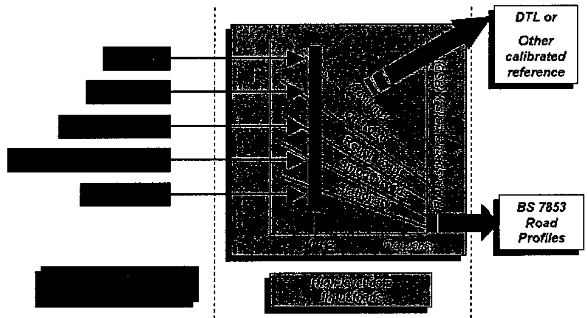

Figure 4. Linking of Low- and High-level Usage Data

A good example of high-level data is BS 7853[4], which specifies a uniform method of reporting measured vertical road profile data on eight classes of roads, streets, highways and off-road surfaces. The profiles are presented as displacement PSDs against spatial or angular frequency. The Standard does not cover cross-country surfaces; however, these can be separately provided as part of a test track standard calibration programme, such as that run by the US Army Test Center (ATC) or data from surfaces calibrated to order using the QinetiQ Dynamic Tyre Load (DTL) vehicle[5] in the UK which is described later in this paper.

USAGE SCENARIOS

To provide benefits to the military, the definition of terrain categories needs to be compatible with those used to define

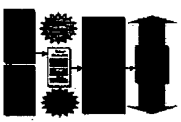

military vehicle requirements and have known or measurable characteristics that can give repeatable results. This information provides the scenario inputs to the LCM process shown in Figure 1. Whilst the assumptions used to develop "Battlefield Missions" vary from nation to nation and project to project, the descriptions of surface categories and speeds used to define usage scenarios in USA and UK military vehicle acquisition are sufficiently alike to allow a common approach. Accordingly, five classes of terrain surface bands were chosen for the original TTCP trial, which in ascending order of severity are:

- "**Stopped**": representing a stationary vehicle with engine running;
- "**Smooth-road**": representing operations on motorways and A-class roads;
- "**Rough-road**": representing operation over B-class road and country lanes;
- "**Off-road**": representing operations over tracks and badly damaged road surfaces;
- "**Cross-country**": overland use over unmade surfaces that represent the upper limit of the vehicle usage spectrum.

All of these categories can be represented by surfaces on test tracks that are used for trials in support of vehicle procurement programmes in the USA and UK.

DTL MEASUREMENT IN WHEELED VEHICLES

In parallel with the CSV availability research under ARP 26b, a further study was conducted in the UK[5] into the measurement of DTL to allow a better understanding of operational duty cycles and how they impact on vehicle reliability. The broad approach of the work was to develop a specification for a data-logging system that would take inputs from sensors on the vehicle and calculate the forces between the vehicle tyres and the ground using an algorithm developed in the programme. An experimental DTL system was installed in a beam-axle, long-wheelbase Land Rover very similar to the TSS vehicle but more comprehensively instrumented. Like the TTCP/TSS, the DTL concept is based on the assumption that the acceleration of the un-sprung mass (Land Rover axle) can be related to the terrain surface profile/roughness; however, the DTL study approaches the monitoring of load inputs in a different way.

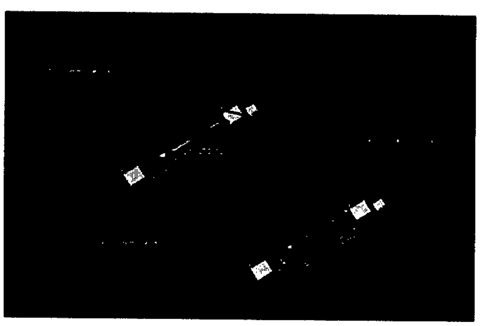

Figure 5. DTL Vehicle Rig

The work conducted under the DTL study is summarised here insofar as it relates to the joint DTL/TSS activities described elsewhere in this Paper.

The main sources of DTL system input are from 5 accelerometers and 4 strain gauges - as shown in Figure 5 - and. 4 thermocouples The remainder of the system comprises a laptop computer, a data logger and analyser capable of logging and processing DTL. The system is installed on a beam-axle Land Rover 110 (similar to the TSS vehicle) and the aim is for the data to be available for download to give design engineers a better understanding of real vehicle operating environments. The DTL measures accelerations in real time but the data is downloaded on tape and retrospectively processed off-board.

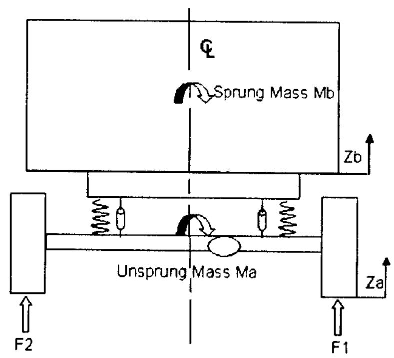

Figure 6. DTL Measurement Technique

Where tyre load = spring force + damper force + un-sprung mass acceleration and equations of motion were derived for full 3D model of vehicle

The algorithm assumes a mass-spring-damper-vehicle model as shown in Figure 6. The inputs of component displacement, velocity and accelerations are derived in two parallel ways - the integration of accelerometer outputs and the differentiation of spring displacement outputs. This allows a comparison to maximise confidence in the algorithm inputs. The assumptions used and the derivation of the two associated equations are beyond the scope of this Paper. However, the study concluded that there was reasonable correlation between the two methods and, compared with strain gauges, accelerometers are more reliable, require less calibration and offered a greater potential for development.

DTL VALIDATION

To establish the accuracy of the DTL system outputs - both instrumentation and algorithm - a validation of the system was conducted on a 4-poster test facility at the Royal Military College of Science at Shrivenham where the vehicle was excited in its three main body vibration modes: bounce, pitch and roll to ensure that the system functioned in all harmonic states of the vehicle. The directly measured tyre loads were compared with the calculated values.

The results of the comparison between the measured and calculated loads show a reasonable match as shown in Figure 7. This confirms that the equations are producing realistic magnitudes and implies that the fundamentals of the equations are correct.

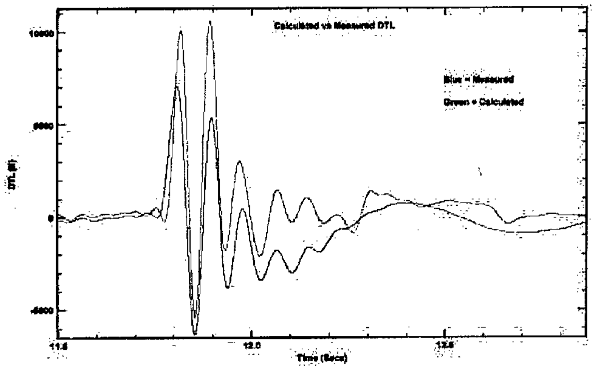

Figure 7: System Response to Step Input

The DTL is capable of supplying the high-level data inputs necessary to drive the TTCP PoF-based LCM modelling and simulation tools by providing vibration PSD data on various terrain surface types.

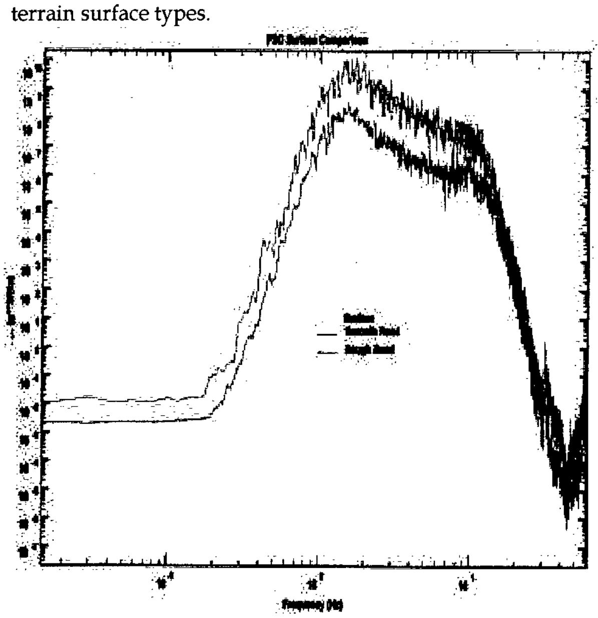

Figure.8: Frequency Analysis

Figure 8 shows an example of an acceleration-PSD output of the DTL, which compares a smooth-road surface and a rough-road surface. It was concluded that the DTL system is too complex and insufficiently robust to be fitted widely across vehicle fleets and that the TSS was a more suitable means of providing the low-level data input as a simple measure of vehicle usage in service. Nevertheless, the DTL could have a potentially important role in providing the high-level data profiles of surfaces that are outside the scope of BS 7853[4].

CALIBRATION OF TSS WITH DTL VEHICLE

The TSS Land Rover (one acceleration signal input) and the DTL Land Rover (37 acceleration signal inputs) were run together on the QinetiQ Longcross test track at Chertsey in the UK. The aims were: to check the suitability of the selected threshold values of the terrain category bands; to investigate the effect of speed on the indicated terrain category bands and to establish potential relationships between the simplified TSS usage data and calibrated terrain surfaces.

TSS LOW-LEVEL DATA

The acceleration thresholds of the terrain category bands were set to the levels listed in Figure 9.

ROAD SENSOR INDICATOR	ACCELERATION LEVEL (UNSPRUNG MASS 10 SEC AVERAGE OF PEAKS)
STOPPED	< 0.1G
SMOOTH	0.1 - 0.4G
ROUGH	0.4 - 1.1G
OFF-ROAD	1.1 - 2.0G
CROSS COUNTRY	> 2G

Figure 9. Acceleration Thresholds for Terrain Bands

To represent the terrain categories above, the DTL Land Rover, followed by the TSS Land Rover, was driven in a series of runs over the four test surfaces measuring the speeds using a calibrated digital speedometer. On completion of the runs, the data from the TSS elapsed time indicators were input manually into a Microsoft Excel spreadsheet, converted into a percentage of the total running time and plotted onto a histogram. Typical results are shown in Figure 10. The trial results show clearly that the indicated terrain band is very dependent on speed; however, this is not an issue in the TSS measurement of accumulated damage, which is dependent on the amplitude of the induced vibrations.

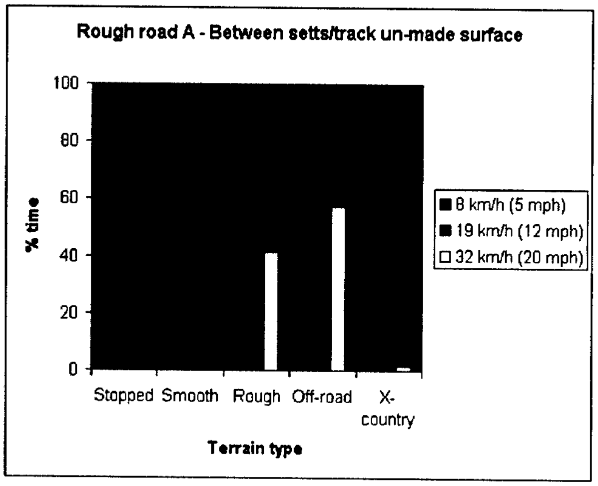

Figure 10. Time on Rough Road

DTL HIGH-LEVEL DATA

Data collected by the DTL Land Rover, and selected as the most representative, was processed and plotted to provide PSD signatures for each of the four terrain surface bands. The readings were taken from the accelerometer fitted to the left rear chassis, to coincide with the TSS accelerometer position. The non-smoothed acceleration PSD is shown in the PSD of Land Rover body plot in Figure 11 overlaid by smoothed curves.

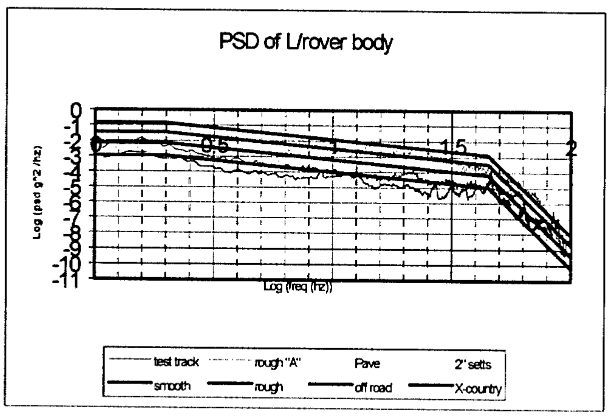

Figure 11. Land Rover Body PSD

This original data was smoothed using the Microsoft Excel best-fit curve facility to provide straight line plots for each selected terrain surface category. These smoothed curves are reproduced in Figure 12.

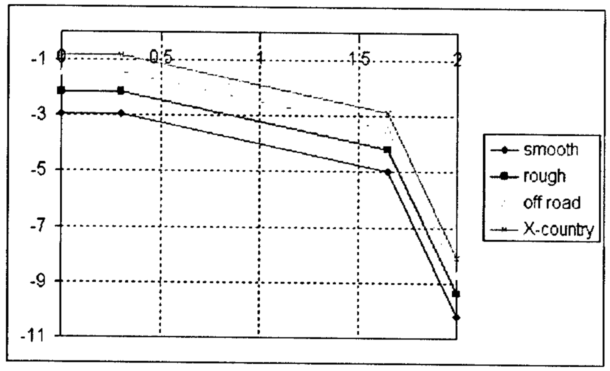

Figure 12. PSD Data for Land Rover Body PSD

DATA SIMPLIFICATION FOR LCM

Temperature, humidity and vibration sensors can be used to monitor, characterise and measure the various environmental loads present during a system's use. The sensors typically provide an electrical output in response to the system's usage conditions and environments, the history of which can be processed, accumulated, stored and used to predict the system's remaining life. Experience has shown that even the simplest data collection systems can accumulate vast amounts of data quickly, requiring either a frequent download procedure or large capacity storage device. Provided that the chronological time history can be dispensed with, temperature and humidity are relatively simple to handle. A series of time samples can be reclassified into temperature change band bins as shown in Figure 13 and then the results can be presented as an accumulated frequency in each bin. This type of output can be readily input into fatigue damage accumulation models such as the CALCE PWA physics of failure analysis tool[2].

If the materials concerned have various transition states at specified temperature bands then the data can be split and stored in bins against each transition state temperature range.

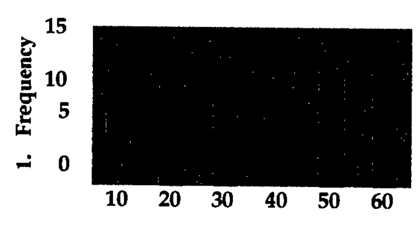

Figure 13: Temperature Change Bands

Output from vibration sensors are far more complex and are often required to be translated into different domains, such as frequency, time, amplitude, and power before they can be input into fatigue analysis software. If a smart sensor system such as the TSS is used to perform the transition from one domain to another in real time then a similar procedure to that used for storing temperature history can be used to capture and store stress levels due to vibration. Whatever information is collected and stored the user must continually check that the sampling rate, frequency range, amplitude etc. remain compatible with the input requirements of the damage/life prediction models that are to be used in subsequent processing.

DATA REDUCTION

The main reasons for using data reduction in life consumption monitoring are:

- reduction of storage space;
- reduction in datalogger CPU load, and;
- alignment with life prediction models.

Efficiency measures of data reduction methods should consider:

- gains in computing speed and testing time;
- the ability to condense load histories without sacrificing important damage characteristics;
- the ability to preserve the sequence of the most damaging reversals in the original loading;
- estimate of the error introduced by omitting data points.

The CALCE group, based at Maryland University[6] are currently studying the accuracy associated with a number of data reduction methods such as:

- Ordered Overall Range (OOR);
- Cycle Counting;
- Rainflow Counting, Range-pair Counting, Peak Counting, Level Crossing Counting, Fatigue Meter Counting, Range Counting, etc.

Two of which are described in this Paper.

THE ORDERED OVERALL RANGE (OOR)

This method of data reduction allows the user to convert an irregular history into a regular sequence of peaks and valleys and also specify the range (i.e., upper limit - lower limit) of the reversals to be eliminated. A screening level (S), expressed as a fraction of the overall range, is used to define the range of the reversals to be eliminated.

Moving across Figure 14 from left to right, taking the shortest possible route through the course[7] a flag marks each location where the direction changes from north to south (or vice versa), and corresponding reversals in the original loading sequence are noted. Peaks and valleys that were originally separated by smaller interrupting ranges now become adjacent, creating larger overall ranges. This allows smaller load fluctuations to be screened out, but preserves the sequence of the most damaging reversals. The OOR method can be applied to a load history by selecting one of the extreme reversals in the history (either the largest peak or the smallest valley), and choosing the next reversal that differs from the selected reversal by more than the

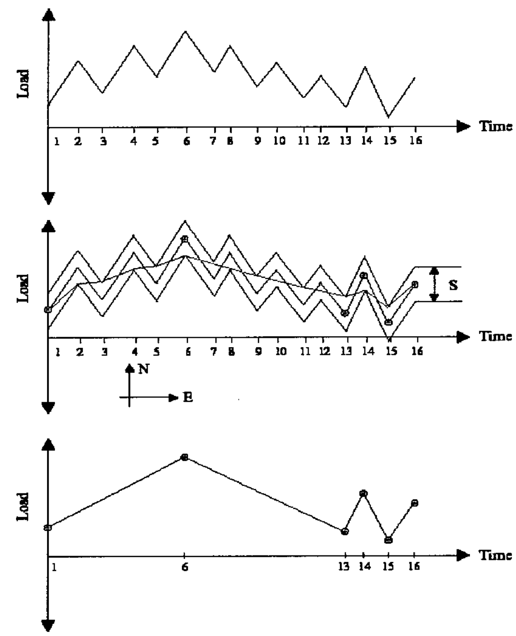

Figure 14: The Ordered Overall Range Method

screening level as the first tentative candidate. The approach for selecting the tentative candidate[7][8] is as follows:

- If the selected starting reversal is a peak, the next valley that differs by more than the screening level from the peak should be selected as the first tentative candidate. Succeeding reversals are then checked. Peaks are checked to see if they differ from the candidate by more than the screening level (event 'x'), and valleys are checked to see if they are lower than the candidate (event 'y'). If event 'y' occurs first (i.e., before event 'x'), then the candidate is rejected and the new valley becomes a candidate. If event 'x' occurs first, the candidate is validated and the newly found peak becomes the next candidate.
- A reversal is defined as the point at which the first derivative of the time history changes sign. A reversal can either be a peak (where the first derivative changes from a positive to a negative sign) or a valley (where the first derivative changes from a negative to a positive sign).
- The screening level is a fraction of the maximum range (i.e., largest peak - smallest valley) in the history.
- If the selected starting reversal is a valley, the next peak that differs by more than the screening level from the peak should be selected as the first tentative candidate. Succeeding

reversals are then checked. Valleys are checked to see if they differ from the candidate by more than the screening level (event 'x'), and peaks are checked to see if they are higher than the candidate (event 'y'). If event 'y' occurs first (before event 'x'), then the candidate is rejected and the new peak becomes a candidate. If event 'x' occurs first, the candidate is validated and the newly found valley becomes the next candidate.

This process continues until the last reversal is counted. Also, since the counting process starts from an extreme reversal (which may or may not be the first reversal in the history), the method has to be applied to both sides of the extreme reversal to take the entire history into account.

CYCLE COUNTING

Cycle counting methods are used in fatigue analysis to transform a time history consisting of several reversals (peaks and valleys) into an equivalent cyclic history[9]. A cycle is a condition when the applied load returns the material to the state it was before the load excursion occurred (see Figure 15). The rainflow counting method[10] has been used in this paper to count cycles.

In the rainflow method, the load-time history is plotted in such a way that the time axis is vertically downward, and the lines connecting the load peaks are imagined to be a series of sloping roofs. The rain flow is initiated by placing drops successively at the inside of each reversal. Cycles and half cycles are identified by imposing certain rules on the rain dripping down the roof[9]:

- The rain is allowed to flow on the roof and drip down to the next slope except that, if it initiates at a valley, it must be terminated when it comes opposite a valley equal to or more negative than the valley from which it initiated. A half cycle is then defined between the starting valley and the next peak.

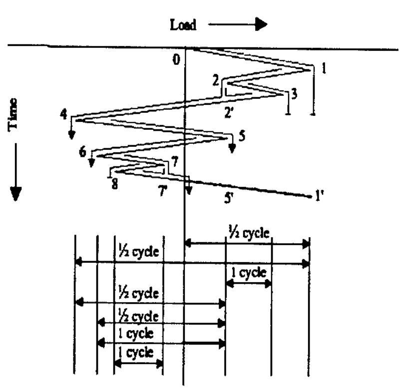

Figure 15: Identifying cycles in a load history

- Similarly, if the rain flow is initiated at a peak, it must be terminated when it comes opposite a peak equal to or more positive than the peak from which it initiated. A half cycle is then counted between the starting peak and the next valley.
- The rain flow must also stop if it meets rain from a roof above. This ensures that every part of the load history is counted once and only once.
- Cycles are counted when a counted range can be paired with a subsequent range of equal magnitude in the opposite direction.

The only shortcoming of the rainflow method is that it can only be used for simple fatigue analysis, since the method does not provide any information about the mean load or the time taken for the load to increase from its peak to its valley (and vice versa).

THREE PARAMETER RAINFLOW CYCLE COUNTING

A modified method called 3-parameter rainflow cycle counting[11] can be used to handle the mean load and time taken for the load to increase from peak to valley. The method accepts a sequence of successive differences between peak and valley values (P/V ranges) in the time history as an input, and determines the range of the cycle, the mean of the cycle, and the half-cycle time. The dwell time at the extremes of the cycle is assumed to be 25% of the half-cycle time. The method identifies cycles as follows[11]: Consider three successive P/V differences d1, d2, and d3, as shown in Figure 15. A cycle is identified only if the following condition is true:

This condition is called the 'loop condition.' For a given sequence of P/V ranges, if the loop condition exists, the method picks the

$$d_1 > d_2 \leq d_3$$

loop corresponding to size d2 off the cycle, leaving only the residual wave 1-2-4-5, corresponding to a half-loop size (d3-d2+d1) in the load plot. This operation is called 'loop-reaping.' For a given sequence of P/V ranges, the 3-parameter rainflow method "reaps" the smaller cycles that occur during a larger cycle.

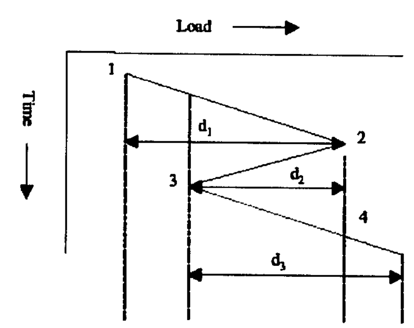

Figure 16: Loop Condition and Loop Reaping[12]

The range, mean, and half-cycle time of the residual half cycle is adjusted according to the loop-reaping condition, and the process is applied until the last P/V range is read. The counted cycles and residual half-cycles are placed in separate arrays that store the cyclic range, the cyclic mean, and the time taken to complete the half cycle.

EXAMPLES OF DATA REDUCTION

Figure 16 below shows the temperature record collected on a field experiment at CALCE[13]. Temperature readings were taken automatically every 10 minutes over a period in excess of 65 days. Applying the OOR algorithm with a 'S' value of "7" reduced the number of points from 9547 to 823.

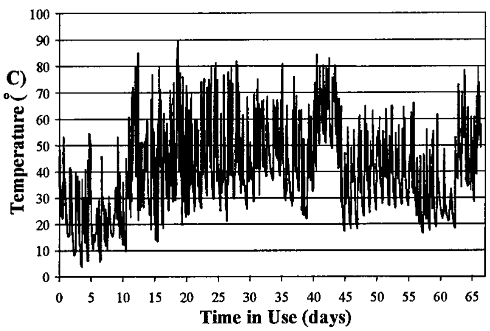

Figure 17. Number of Data Points: 9547 - Time Difference Between Two Data Points: 10 Minutes

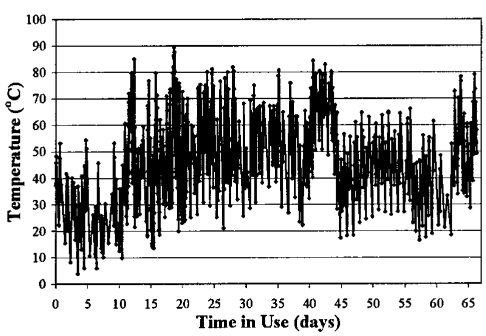

Figure 18: Temperature Data Simplified Using OOR- Number of Data Points: 823

Figure 18 shows the output from the Maryland University Study, which concluded that it was possible to reduce data by up to ninety percent whilst still maintaining less than a two percent error in the system's predicted life.

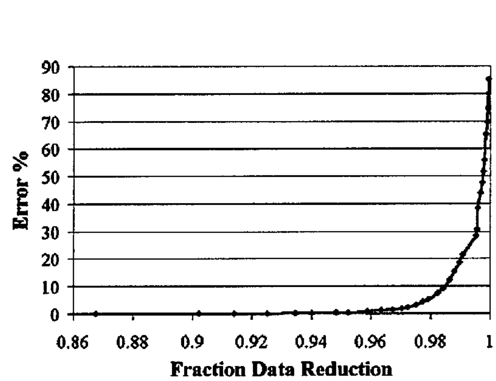

Figure 19. Effects of Data Reduction on Life Prediction Error.

Where Percentage error and Fraction Data Reduction (FDR) were calculated using the formula described below:

$$\text{FRD} = \frac{\text{No. of Remaining Data Points}}{\text{Total No. of Data Points}}$$

$$\text{Error } \% = \left[\frac{\text{Damage Accumulation}}{\text{Damage Accumulation } (S = 0)}\right] * 100 \quad \textbf{2.}$$

THE WAY AHEAD

Several factors are changing the military perspective on the need for HUMS in respect of military vehicle fleets. Of particular note are:

- the drive for improvements in operational availability and reduced costs;
- the impending modernisation of practices in the UK Army to manage vehicle fleets centrally;
- the growing use of HUMS as standard in commercial vehicles (buying anything else comes at an extra cost);
- new European pollution standards which can only be met by electronically controlled engines (which will provide a rich source of health and usage data as a by-product)
- prognostic systems will be necessary to support the MFOP capability for future military vehicles.

The importance of cost effectiveness has been stressed throughout this paper; however, equally important is the effective processing of data and distribution of information to the 'customer's' in the DB(L). Several steps are needed to shift military vehicle fleets from where they are now to where they need to be to support prognostics. Accordingly, a HUMS framework is required, based on open systems architectures that is capable of supporting the little legacy data that is available now while having the potential to accept the outputs from future sophisticated data collection technologies on a 'plug and play' basis. The principle for the UK using a CANbus architecture is shown in Figure 19.

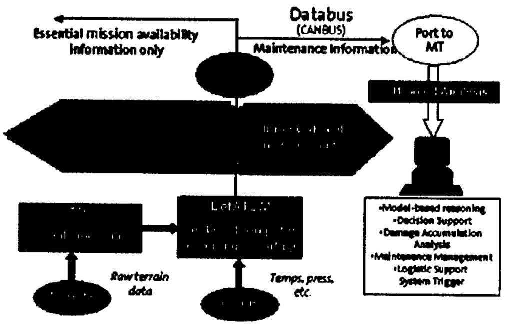

Figure 20. HUMS in the DB(L)

This principle has also been recognised by the international Object Management Group and is reflected in the White Paper[14] on Open Systems Architecture for Condition Based Monitoring (OSA-CBM), which describes a developing standard for open, integrated, condition-based monitoring systems. The top-level overview of the seven-step architecture is shown in Figure 20.

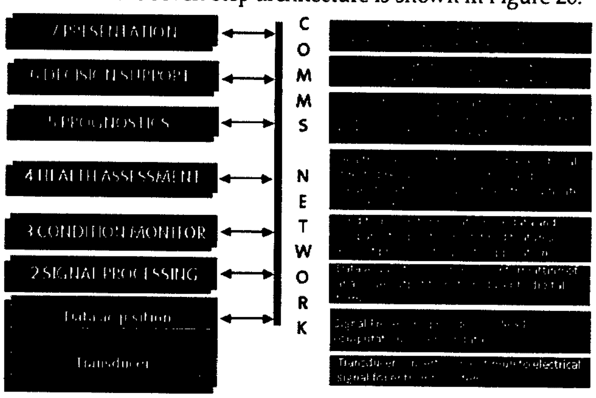

Figure 21: Seven-step OSA-CBM Overview

CONCLUSIONS

To be affordable for military vehicles, data collection for HUMS needs to be carefully planned and designed with the knowledge of how collected data will be analysed and used. The simplest forms of data collection can create storage problems and data simplification is the key to keeping the data to a minimum while still providing the required level of analysis accuracy. Moreover, only data that has a recognised customer should be collected and this data needs to be effectively distributed via an integrating framework constructed on open architectures.

REFERENCES

[1]"*Physics of Failure based Approach to Life Consumption Monitoring*", The Technical Cooperation Programme, Joint Systems and Analysis Group, Technical Panel ` Task 1.03 Final Report ", 30 June 2002.

[2] CALCE PWA Software, CALCE Electronic Products and Systems Center, University of Maryland, 2002

[3] McNee.I.R. *"Report on CSV reliability research in support of ARP26b milestone M9R501 (UC)"* DERA/LSA3/CR990185/1.0 April 1999.

[4]"*Mechanical Vibration - Road Surface Profiles - Reporting of Measured Data*", BS 7853, 1996; ISO 8606, 1995.

[5] Bayes. M.J. *"Progress Report for an In-service Data Logging System for Dynamic Tyre Load"* DERA/LSA2/CR980550 October 1998.

[6] Ramakrishnan A., and Pecht M. *"Life Consumption Monitoring Methodology for Electronic Systems"*, CALCE Electronic Products and Systems Center, University of Maryland, 2000

[7] Fuchs, H.O., Nelson, D.V., Burke, M.A., and T.L. Toomay, "Shortcuts in Cumulative Damage Analysis," *SAE National Automobile Engineering Meeting*, Detroit, 1973.

[8] Tucker, L., Downing, S., and L. Camillo, "Accuracy of Simplified Fatigue Prediction Methods,"*SAE Automotive Engineering Congress and Exposition*, Detroit, MI, February 24-28, 1975.

[9] Collins, J., *Failure of Materials in Mechanical Design*, John Wiley & Sons, New York, 1993.

[10] Socie, D., "Rainflow Cycle Counting: A Historical Perspective," pp. 3-10, *The Rainflow Method in Fatigue*, Butterworth-Heinemann, Oxford, 1991.

[11] Cluff, K.D., "Characterizing the Commercial Avionics Thermal Environment for Field Reliability Assessment," *Journal of the Institute of Environmental Sciences*, pp. 22-28, Vol. 40, No. 4, July/August 1997.

[12] Anzai, H., "Algorithm of the Rainflow Method," pp. 11-20, *The Rainflow Method in Fatigue*, Butterworth-Heinemann, Oxford, 1991.

[13] Satchidananda Mishra, Micheal Pecht, Ted Smith, Ian McNee, Roger Harris, *"Remaining Life Prediction of Electronic Products Using Life Consumption Monitoring Techniques"*, CALCE Electronics Products and Systems Center, Universtiy of Maryland, 2002.

[14] Fred M. Discenzo, William Nickerson, Charles E. Mitchell, Kirby J. Keller *"Open Systems Architecture Enables Health Management for Next Generation System Monitoring and Maintenance.* OSA-CBM Development Group

COPYRIGHT AND TRADEMARKS

Damage Detection Of Gas Turbine Engines By Analysing Blade Tip Timing Data

Wenyi Wang
Air Vehicles Division, DSTO Australia

ABSTRACT©

In this paper, a novel method of detecting damage in gas turbine engines is proposed. Based on the commonly used blade tip timing data, a pulse position/amplitude modulated (PPM/PAM) signal is firstly constructed. The changes in the PPM/PAM signal, which are caused by mechanical faults such as blade/disk cracks and foreign object damage, are identified by demodulating the PPM/PAM signal and averaging. Preliminary results using numerically simulated data show that this method is potentially effective in detecting initial damage. Because of the use of high-resolution numbers in the blade tip timing data, it appears feasible to achieve high detection sensitivity to blade damage in gas turbine engines using the proposed method.

INTRODUCTION

It has been envisaged that prognostics and health management (PHM) will be a built-in capability for the next generation fighter aircraft, e.g., Joint Strike Fighter. Engine health monitoring will be an essential component of the PHM systems. Over the last decade, some sensor technologies have been developed for the investigation of high cycle fatigue and health monitoring of gas turbine engines. Among these technologies, the measurement of engine blade tip timing has been widely used and is becoming an industry standard for non-intrusive stress measurement systems.

When the blades of a bladed disc of a gas turbine engine pass a casing-mounted sensor (of capacitive or eddy-current type) one after the other, a pulse train will be generated. The blade tip timing data are the fused version of the pulse train, which may contain some high resolution (e.g., 32-bit) numbers for each blade-passing event, such as pulse zero crossing time (or blade arrival time - BAT), pulse amplitude (PA) and etc. These numbers are mainly determined by the rotating speed, the material/manufacturing tolerances or in-service degradation, blade vibration, and damage to the assembly. The damage may include foreign object damage and blade/disk cracks.

Several techniques [1~5] have been developed recently in analysing the blade tip timing data, with the emphasis on extracting features of synchronous or asynchronous vibrations. This paper presents a novel method of engine damage detection by concentrating on the extraction of subtle signal changes induced by the damage. Using this method, a pulse position- and amplitude-modulated (PPM/PAM) signal can be formed based on the blade tip timing data. The standard demodulation techniques [6] are then applied to the PPM/PAM signal, in conjunction with signal averaging and residual signal techniques commonly seen in gear fault diagnostics, to extract the signal changes. Preliminary results using numerically simulated data show that the proposed method can be effectively used to detect engine damage. Because blade tip timing data is inherently of high resolution, it appears feasible to achieve a very high sensitivity of damage detection using the proposed method.

THE BLADE TIP TIMING DATA

Blade tip sensors have been used for decades to measure engine blade vibrations [7~10]. The typical sensors, e.g., capacitive, eddy current and optical sensors, usually generate a pulse train when engine blades pass the casing-mounted sensor one after the other. For engine health monitoring, continuous data recordings will be required to represent the progression of health status. However, the continuous recording of the raw blade passage data can be extremely demanding on storage capacities due to high rotating speeds (e.g., 18000 rpm for a F404 disk) and large number of blades (e.g., 64 for the F404 disk) on an engine disk. Hence, the raw blade passage data need to be fused for the ease of storage and archiving. A data fusion technique was described by von Flotow *et al* of Hood Technologies Corporation [1], which characterised each blade- passing event using two 32-bit numbers - the BAT and the PA. The PA is usually associated with the clearance between blade tips and the casing, i.e., tip clearance (TC).

The nominal BAT can be calculated from the nominal rotating speed and the number of blades on the disk. For healthy bladed discs, it is expected that the actual BAT will be affected by speed fluctuations, the material/manufacturing tolerances, and vibrations of the blades and the disk. Hence, the measured BAT may contain random and periodic vibratory (or dynamic) components on top of the static values. When a root crack is developed in one of the blades, the cracked blade under load will deflect more, which can cause static changes to its arrival time.

Moreover, the blade cracking may cause changes to the vibratory components of the BAT of all blades with the possibility of mode localisation (or localised dynamic response) [11] around the cracked blade, which can be referred to as dynamic changes. On the other hand, the damaged blade may also alter the centre of gravity and centrifugal forces of the assembly, which will in turn cause subtle changes to the measured TC. Essentially, both the measured BAT and TC will consist of a static component, a vibratory component, and a static change. The composition of the BAT and TC for a bladed disc with a cracked blade can be summarised by the following diagram.

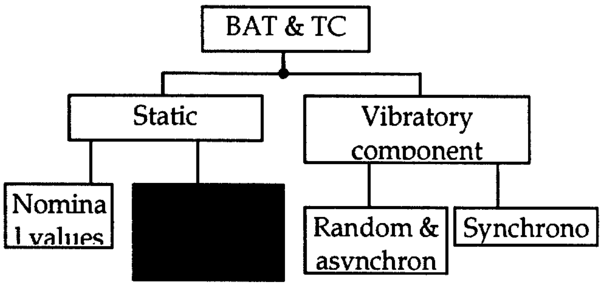

For damage detection and health monitoring of bladed discs, we are only interested in the damage-induced static change; hence, we must be able to separate it from the vibratory component. The separation may be achieved through signal averaging and filtering (see Section 0), which will depend on the amplitude of the static change and the frequency content of the vibratory component.

In order to explore the signal processing algorithms of detecting the damage-induced changes in the BAT and TC, a pulse train signal with ideal delta-pulses can be constructed based on the blade tip timing data. The positions and amplitudes of the delta-pulses are determined by the BAT and PA (or TC) in the blade tip timing data. Hence, the constructed pulse train is a position- and amplitude-modulated signal, i.e., pulse position modulated (PPM) and pulse amplitude modulated (PAM) signal. *Figure. 1* shows a numerically simulated PPM/PAM signal for a complete revolution of a bladed disc with 62 blades and a crack at the 23rd blade.

A three-step procedure was followed during the simulation shown at *Figure. 1*. It started by forming a pulse train with uniform interval and unit amplitude, which is the nominal static component and is known as the carrier signal in modulation theory. Secondly, the pulse position and amplitude were added by random variables with Weibull [10] and Gaussian distributions, respectively, and periodic variations, which corresponds to the vibratory components. To simulate a crack at the 23rd blade, static changes were then introduced to the position and amplitude of the 23rd pulse. The vibratory components and the static changes introduced in step two and three constituted the modulating signal, which can be separated from the carrier signal via demodulation algorithms.

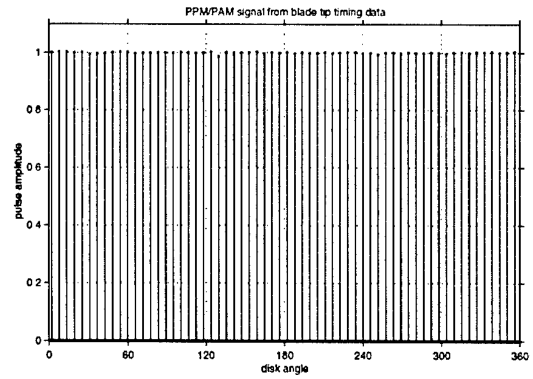

Figure. 1. A simulated PPM/PAM signal of a bladed disc with 62 blades and a crack at blade #23.

DAMAGE DETECTION TECHNIQUE

Based on the PPM/PAM signals, the detection of damage-induced static changes can involve two separations: the separation of modulating signal from the carrier - demodulation, and the separation of static changes from the vibratory components - identification. Both PPM and PAM signals are widely used in digital telecommunications. Standard demodulation algorithms exist in obtaining the modulating signals. The identification can be conducted using signal averaging and filtering techniques. In this section, we present a novel method for damage detection of gas turbine engines, which combines the pulse position demodulation (PPD) and pulse amplitude demodulation (PAD) techniques, the signal averaging and residual signal techniques.

The PPM/PAM signal shown *Figure. 1* represents the BAT and TC information in one complete revolution of the bladed disc. Such PPM/PAM signals from consecutive revolutions can be repeatedly generated using the blade tip timing data. The PPM/PAM signals for each revolution are processed by the PPD and PAD independently and down-sampled by a factor of N_s/N_b, where N_s and N_b are the number of samples in the PPM/PAM signal and the number of blades on the disk, respectively. The re-sampled PPD's and PAD's, with N_b samples, are then averaged to reduce the effect of random variation. Each sample in the averaged PPD (APPD) represents the BAT of each blade relative to the nominal BAT, and each sample in the averaged PAD (APAD) indicates the actual TC with respect to the nominal TC. Further enhancements, e.g., residual signal or band-stop filtering, may be necessary if periodic patterns exist in the APPD and APAD. The method can be abbreviated by **DEMARS** - DEModulation, Averaging and Residual Signal, and it is outlined by the following diagram.

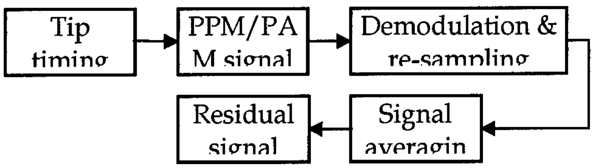

ANALYSIS RESULTS USING SIMULATED DATA

Using the three-step procedure described in Section 2, four possible scenarios of the PPM/PAM signal were numerically simulated:

A. Relatively large (compared to the level of random component) damage-induced static changes with a random vibratory component;
B. Medium static changes with a random component;
C. Medium static changes in random and large periodic vibratory components;
D. Small static changes in random and very large periodic vibratory components.

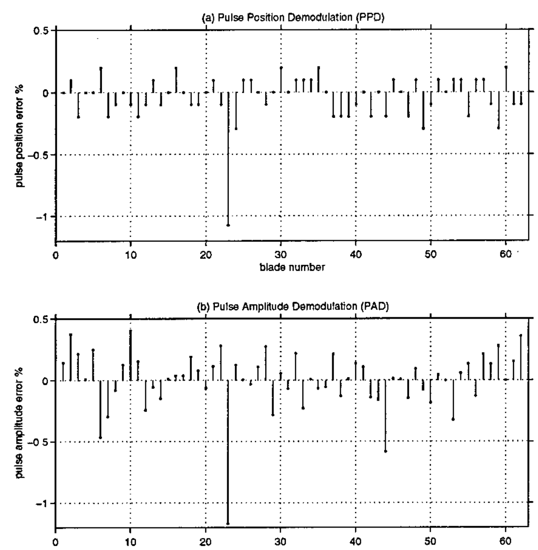

Figure. 2. PPD and PAD (no average) for scenario A – the PPM/PAM signal shown in Figure. 1 with large damage-induced changes at blade #23, plus random modulation.

The parameters used for the numerical simulations are summarised in *Table 1*. These simulated signals were analysed using the DEMARS technique given in Section 3. The analysis results are presented in *Figure. 2* to *Figure. 7*.

Table 1. Parameters used in the simulated PPD/PAM signals

Scenarios	Damage-induced static changes (%)		Vibratory components (%): random/periodic			
	Pulse Position (PP)	Pulse Amp. (PA)	Random (σ)		Synchronous (amp.)	
			PP	PA	PP	PA
A *(Figure. 1)*	1	1	0.16	0.2		
B *(Figure. 3a)*	0.5	0.5	0.16	0.2		
C *(Figure. 4a)*	0.5	0.5	0.16	0.2	1	1
D *(Figure. 5a)*	0.3	0.3	0.16	0.2	3	3

Note: All the percentage numbers are relative to the corresponding nominal PP or PA. The σ is the standard deviation of the random components.

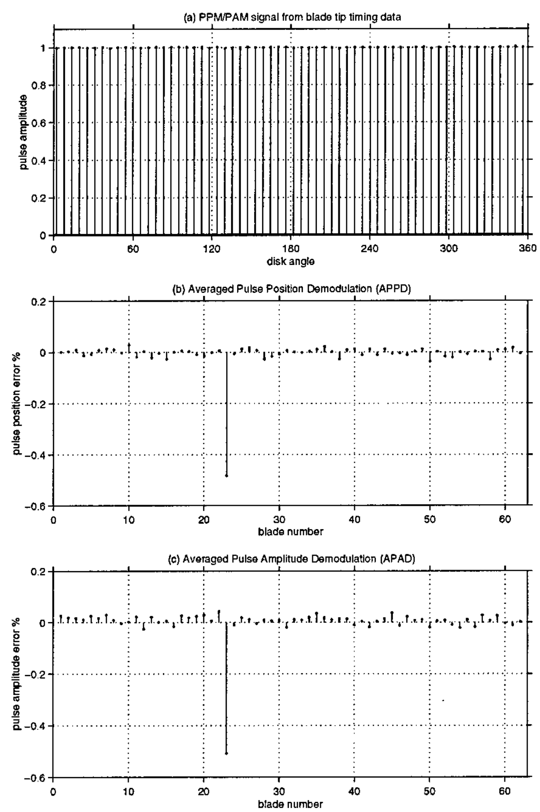

Figure. 3. APPD and APAD (200 average) for scenario B – the PPM/PAM signal with medium damage-induced changes at blade #23, plus random modulation (no periodic modulation).

PPM/PAM with random vibratory component: the signal can be generated by considering the BAT and TC as random variables with the means equal to their nominal values. The variations of the BAT and TC are assumed associated with the random vibration of blades. Static changes of position and amplitude are introduced to a particular pulse in the PPM/PAM signal to simulate a blade cracking. *Figure.* 2 depicts the PPD and PAD of the PPM/PAM signal shown in *Figure. 1*.

In this case, the damage-induced static changes are relatively large (1 percent), which made the detection possible using a single PPD and PAD (no averaging), as seen in *Figure.* 2a and 2b. However, with smaller changes produced by initial damage, it is unlikely that the damage can be detected without averaging the PPD and PAD. Sufficient averages of PPD and PAD can significantly reduce the effects of the random variations to expose the damage-induced static changes. *Figure.* 3a depicts a simulated signal with medium damage-induced changes, i.e., 0.5 percent in pulse position and 0.5 percent in pulse amplitude. The changes were identified after 200 averages, as clearly seen in *Figure.* 3b and 3c.

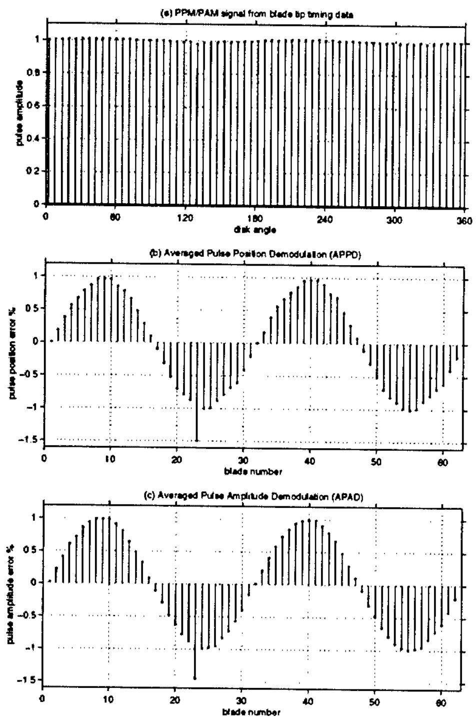

Figure. 4. APPD and APAD (200 average) for scenario C – the PPM/PAM signal with medium damage-induced changes at blade #23, plus random and large sinusoidal (2nd EO) modulations.

PPM/PAM with random and periodic vibratory components: the variations of the BAT and TC can follow some periodic patterns whose frequencies are synchronous with the rotating speed, i.e., integer multiples of the rotating speed (engine orders - EO). An averaging process will attenuate the random and asynchronous components, but not the synchronous components. Hence, the synchronous components will exist in the final APPD and APAD (seen as integer number of cycles). If the damage-induced static changes are comparable to the amplitude of the synchronous components, the changes will be distinguishable in the APPD and APAD, which can be identified by a distinct local deviation from the periodic pattern. This is demonstrated in Figure. 4 where the second EO sinusoidal components were introduced to the pulse position (1 percent amplitude) and pulse amplitude (also 1 percent). However, the static changes might be much smaller than the periodic component, so that the APPD and APAD will need to be further enhanced to allow identification of the damage.

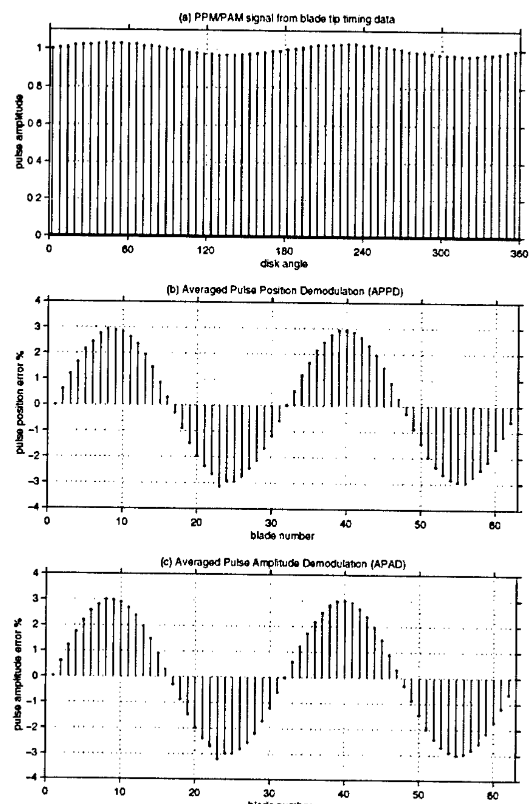

Figure. 5. APPD and APAD (200 average) for scenario D – the PPM/PAM signal with small damage-induced changes at blade #23, plus random and very large sinusoidal (2nd EO) modulations.

PPM/PAM with small damage-induced changes and very large periodic components: this simulates the case where the damage-induced static changes are much smaller than the amplitude of the periodic components. In this case, the static changes will not be obvious in the APPD and APAD. Therefore, we must remove the periodic component using a band-stop (multiple band-stop if necessary) filter to identify the damage-induced changes. The filtered APPD will be similar to the residual signal in gear fault diagnostics, which has been widely used for the last two decades.

Figure. 5 depicts a case where the second EO sinusoidal component with 3 percent amplitude was added to both the pulse position and amplitude. This 3 percent variation was ten times larger than the 0.3 percent crack-induced changes. As can be seen in *Figure. 5*b and 5c, the APPD and APAD have a strong presence of the periodic component whereas the crack-induced changes at the 23rd blade are barely visible. Having removed the sinusoidal component, we can clearly see the 0.3 percent crack-induced changes at the 23rd blade in *Figure. 6*a and 6b.

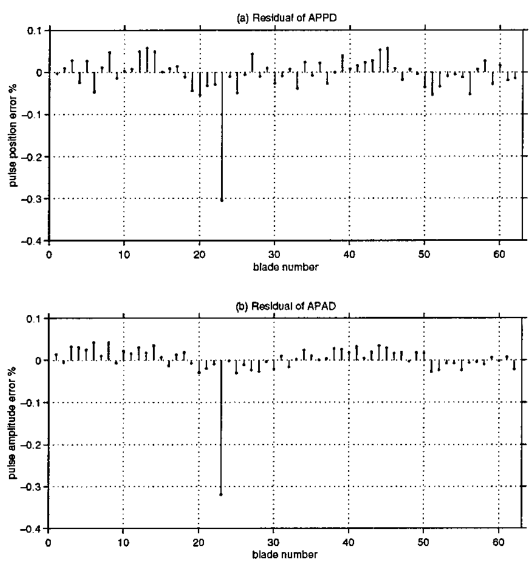

Figure. 6. Identification of small changes imbedded in APPD and APAD shown in Figure. 5b,c (Scenario D) using band-stop filtering.

Another approach to removing the periodic components is to find the difference signal from the APPD and APAD. Depending on the frequencies of the periodic components, we may need to calculate the higher order (> 1) differences. *Figure. 7*a and 7b show the second order difference of the APPD and APAD in *Figure. 5*a and 5b, where the changes are identifiable. However, the largest spikes have shifted to the 24th pulse and we have lost the absolute scale on errors in pulse positions and amplitudes. Furthermore, the differencing method may not work with multi-frequency periodic components. Therefore, use of the APPD or APAD residual signals is recommended for damage detection, and use of the differencing signals is only recommended for reference purpose.

Based on the above discussion and the common nature of extracting subtle signal changes while attempting to detect damage in rotating machinery, it is believed that some other advanced techniques recently developed in DSTO for gearbox fault diagnostics could well be utilised in the damage detection and diagnosis of gas turbine engines. These techniques include the autoregressive (AR) model-based method [12] and the blind deconvolution method [13] that employed the linear prediction error filter, the non-minimum phase signal system, higher order statistics criteria and non-linear optimisation algorithms. These techniques have the potential to provide extremely high detection sensitivity to incipient damage in rotating machinery, which could be vital to the future PHM systems of the single-engine Joint Strike Fighter.

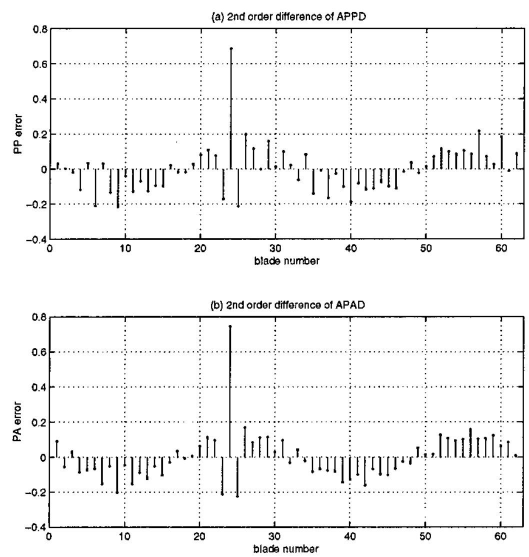

Figure. 7. Identification of small changes imbedded in APPD and APAD shown in Figure. 5b,c (Scenario D) using differencing.

CONCLUSION

A novel method (DEMARS) of damage detection for gas turbine engines has been proposed in this paper. The method combines the pulse position demodulation (PPD) and pulse amplitude demodulation (PAD) techniques commonly used in digital telecommunications, and the signal averaging and residual signal techniques employed in gear fault diagnostics.

Analysis results based on numerically simulated data have shown a promising potential of the DEMARS method. It has been found that the DEMARS can be capable of identifying minor signal changes (induced by incipient damage) embedded in the various background variations that is either random or periodic, or both. Due to the use of high-resolution numbers in blade tip timing data, the detectable changes by the DEMARS can be as small as one minute in the blade rotation angle and 0.3 percent of the nominal tip clearance. This sensitivity may be further improved by using longer PPM/PAM signals constructed from the raw tip timing data and by increasing the number of averages to the PPD and PAD signals.

Future work may include the application of the DEMARS method to spin-rig test data, real engine test data, and ultimately, to aircraft in-flight data. The foreseeable challenge would be those cases that involve mode localisations due to blade cracking (especially for weak inter-blade coupling) [10], rotating stall and other aerodynamic transients, and cases with large background variations (both random and periodic). The transient phenomena can produce vibration energy spikes to a few blades, or even a single blade. If the spikes are synchronous to the rotating speed, i.e., having the same phase for every revolution of the bladed disc, it will make the damage detection more complicated. Large background variations in the blade tip timing data can occur in the situation of incipient damage, where the damage-induced static changes are minimal. In these cases, more advanced signal processing techniques should be explored.

ACKNOWLEDGMENTS

The author would like to thank Dr. J.F. Hou, Mr. G. Stocks, Dr. B. Wicks, Mr. B. Rebbechi and Dr. B.D. Forrester of DSTO for supporting the work and for their comments on the paper. We also acknowledge Dr. A. von Flotow of Hood Technologies Corporation and Mr. G.C. Muschlitz of the U.S. Naval Air Warfare Center for their support to the work.

REFERENCES

1. **von Flotow, A**, Mercadal, M and Tappert, P. 2000 "Health Monitoring and Prognostics of Blades and Disks with Blade Tip Sensors." Proceedings of IEEE Aerospace Conference 2000, vol.6, p433-440, 18-25 March 2000, Big Sky, Montana, USA.
2. **Sonnichsen, H.E.**, 2000 "Real-time Detection of Developing Cracks in Jet Engine Rotors." Proceedings of IEEE Aerospace Conference 2000, p173-183, 18-25 March 2000, Big Sky, Montana, USA.
3. **Carrington, I.B.**, Wright, J.R., Cooper, J.E. and Dimitriadis, G. 2001 "A Comparison of Blade Tip Timing Data Analysis Methods." Proceedings of IMechE Part G - Journal of Aerospace Engineering, Vol. 215, p301-312.
4. **Dimitriadis, G.**, Carrington, I.B., Wright, J.R., and Cooper, J.E. 2002 "Blade-tip Timing Measurement of Synchronous Vibrations of Rotating Bladed Assemblies." Mechanical Systems and Signal Processing, Vol. 16(4), p599-622.
5. **Heath, S.**, 2000 "A New Technique for Identifying Synchronous Resonances Using Tip-Timing." ASME Journal of Engineering for Gas Turbines and Power, Vol. 122, p219-225.
6. **Haykin, S.**, 1994 (3rd edition). Communication Systems. John Wiley & Sons, Inc.
7. **Heath, S.** and Imregun, M. 1998 "A Survey of Blade Tip-Timing Measurement Techniques for Turbomachinery Vibration." ASME Journal of Engineering for Gas Turbines and Power, Vol. 120, p784-791, October 1998.
8. **Petrov, E.**, Sanliturk, K., Ewins, D. and Elliott R. 2000 "Quantitative Prediction of the Effects of Mistuning Arrangement on Resonant Response of a Practical Turbine Bladed Disc." Proceedings of the 5th National Turbine Engine High Cycle Fatigue Conference, 7-9 March 2000, Chandler (Phoenix), Arizona, USA.
9. **Dowell, M.** and Sylvester, G. 1999 "Turbomachinery Prognostics and Health Management Via Eddy Current Sensing: Current Developments." Proceedings of IEEE Aerospace Conference, Aspen, Colorado, March 6-13, 1999, USA, v3, p1-9.
10. **Bladh, R.**, Pierre, C., Castanier, M.P., and Kruse, M.J. 2002 "Dynamic Response Predictions for a Mistuned Industrial Turbomachinery Rotor Using Reduced-Order Modeling." ASME Journal of Engineering for Gas Turbines and Power, Vol. 124, p311-324, April 2002.
11. **Kuang, J.H.** and Huang, B.W. 1999 "Mode Localisation of a Cracked Blade Disk." ASME Journal of Engineering for Gas Turbines and Power, Vol. 121, p335-341, April 1999.
12. **Wang, W.** and Wong, A.K., 2002 "Autoregressive Model-Based Gear Fault Diagnosis." ASME Journal of Vibration and Acoustics Vol.124, p172-179, April 2002.
13. **Wang, W.**, 2001 "Identification of Gear Mesh Signals Using Kurtosis Maximisation and its Application to CH46 Helicopter Gearbox Data." Proceedings of the 11th IEEE Workshop on Statistical Signal Processing, p369-372, 6-8 August 2001, Singapore.

The Operational Benefits of Health and Usage Monitoring Systems in UK Military Helicopters

Alan Draper
HUMS IPT
Defence Logistics Organisation
Ministry of Defence
United Kingdom

ABSTRACT

The potential benefits of HUMS have been lauded for some years. GenHUMS has now been in operational service in the UK Chinook fleet for the last two years and is progressing through its maturity programme. Benefits from the system are being realised and are being exploited by the operators. This paper reviews some of the lessons learnt from the Chinook programme and how they are being incorporated in the Sea King, Puma and Lynx programmes. The paper also outlines the benefits to date and compares against the original claims of the project definition study.

INTRODUCTION

1. The UK military accepted the policy of installing Cockpit Voice / Flight Data Recorders (CV/FDR) as early as 1990. This was to mirror to mandatory civilian airworthiness regulations. CV/FDRs have proven to be extremely beneficial in determining the cause of air incidents. They are however reactive and only prevent incidents of the same type from being repeated. Having undertaken a number of helicopter vibration monitoring trials and having observed developments in the North Sea, it became apparent that the introduction of CV/FDR with full Health and Usage Monitoring System (HUMS) functionality could produce cost benefits as well as proactive improvements in airworthiness.

2. A Project Definition Study (PDS) was commissioned and its results and were incorporated into the Minor Equipment Requirement (MER) documents for the Chinook, Sea King, Puma and Lynx. The positive benefits were used to help secure funding for the project with a return on investment predicted within an average of 6¼ years. Significant benefits were also anticipated from procuring generic HUMS for all legacy aircraft, (hence the term GenHUMS) and a competition was held in 1995 to fulfil the requirement.

GENHUMS PROGRESS

3. Smiths were awarded the contract and work started on the Chinook programme in 1996. The Chinook system entered service in October 2000, embodiment was complete in June 2002. To date the system has amassed some 15,000 hours of operational service.

4. The activity is now centred on the Sea King programme. The Sea King contract placed in July 01 was to install the system to five different marks of aircraft (Mk3 & 3a Search and Rescue, Mk4 Commando, Mk5 Utility and Mk7 Airborne Early Warning). The system has progressed through the preliminary and critical design reviews and the trials installation is expected to commence in early 2003. The first mark will be in service in 2005, all marks will be in service by 2007 and all 90 aircraft will be embodied by late 2008.

5. The Sea King programme draws very heavily on the lessons learnt for the Chinook programme, the most significant of which are discussed in a later section, and will become the baseline standard for the GenHUMS hardware.

6. It has taken time to get the Lynx onto contract. The MoD is currently considering options as to how it is going to replace both Army and Naval variants of Lynx. The Battlefield Light Utility Helicopter (BLUH) and Surface Combatant Maritime Rotorcraft (SCMR) are the titles of the two new staff requirements that have been issued to cover the replacement aircraft. With potential new aircraft being introduced, the MoD has re-examined the value for money of some of the capability upgrades planned for existing Lynx fleets. This has drastically altered the recording requirement of the proposed Lynx HUMS and it has only been possible over the last few months to establish a firm system specification and issue a tender. The Lynx Mk8 contract was placed in November 2002. The Mk8 system will be delivered in 2005. The contract has been designed to provide maximum flexibility should aircraft numbers be affected by BLUH and SCMR procurement decisions. If the BLUH and SCMR requirement is satisfied by a future Lynx derivative then GenHUMS will be directed as the aircraft HUMS.

7. Puma risk reduction activities are now complete, specifications finalised and the company is currently costing the tender. It is the intention to have the contract in place by the end of the fiscal year. The embodiment of HUMS on this aircraft will be part of the Puma Integrated Modification Programme (PIMP). Puma HUMS will be in service by 2006.

LESSONS LEARNT

8. The technical difficulties encountered introducing HUMS into the Chinook are well documented [1,2], most can be attributed to inaccuracies in the Interface Control Document (ICD). To ensure the success of the Sea King, Puma and Lynx programmes, it was essential to learn from the Chinook experience.

9. A significant lesson learnt was to undertake a risk reduction survey on each aircraft type before letting the contract. A draft ICD was prepared for each aircraft from information taken from the aircraft design authorities data pack. The ICD was then scrutinised to identify those signals, which were not fully defined, or there was some issue on the integrity of the signal concerned. Then, at the aircraft, the source, range and definition were established for each parameter. The minimum cost to rectify a discrepancy in the ICD for the Chinook was £350k. At least 7 major discrepancies were identified at each subsequent survey, one survey revealed 20 discrepancies.

10. The surveys also allowed the user community to have an early input as to where the HUMS components should be located. As GenHUMS is being installed into legacy aircraft, some of which have been in service for over 30 years, finding suitable cockpit real estate can be difficult. Having had early visibility of the proposed locations, the user community has been able to comment on their suitability and suggest means for rationalising consoles to make sufficient space. A minor lesson learnt from the Chinook programme was to position the cockpit control unit outside the loci of the restraint belt buckle as a number have been smashed by aircrew eager to leave the aircraft.

11. A painful lesson learnt from the Chinook programme was the importance of the HUMS ground support system design and ensuring it is fully operational as it is introduced into service. In retrospect too much effort was devoted to the implementation of the airborne system at the expense of the ground station. This resulted in the system entering service initially with CV/FDR functionality only.

12. The hardware for the ground support system was specified and procured in 1995. Unix was specified as the operating system as Windows was considered unstable at the time and had not been security accredited. Therefore when the system came in to service it was already antiquated. To prevent this occurring on the follow on programmes, the ground station will be procured as a software application only. When introduced into service the software will be installed onto the latest available hardware.

13. When the HUMS was first considered for installation in UK military aircraft 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. HUMS data is downloaded using a 64Mb PCMCIA card (referred to as a Data Transfer Device (DTD)). An hour-long sortie typically generates 4Mb of HUMS data. The ground stations on which the cards are downloaded are very slow due to the limited processing capability and inefficient software. The majority of the processing is undertaken on a central server. The modus operandi is for each aircraft to upload a DTD at the beginning of the day and remove it on cessation of the day's flying. This can cause substantial delays in downloading and processing HUMS data.

14. For future systems (and the planned upgrade of the Chinook Ground station) it is intended that the ground station will have greater processing power and only retain a limited amount of data locally. All data will be sent to a large central data warehouse where it will be available for fleetwide analysis.

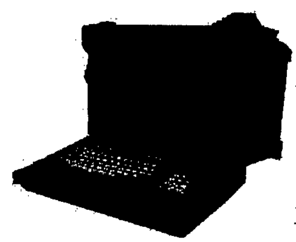

Figure 1: Chinook Ruggedised Portable Ground Station

15. Difficulties have also been encountered with supporting HUMS on deployed operations. The original portable ground station specification was to satisfy use in a nuclear, biological or chemical environment. The only Ruggedised Portable Ground Stations (RPGS) qualified to meet such arduous condition were built like brick outhouses and the lengthy qualification process meant the RPGS processors were even more outdated, on entry into service, than the fixed ground stations. The Chinook squadrons have taken the RPGS, shown in Figure 1, with them on deployments and have been able to take advantage of HUMS data. Admittedly ground crews have had to archive the database daily to free up memory. This is clearly unacceptable but gives an indication of the value of HUMS data to the engineers in the field if they were prepared to take on this administrative burden.

16. Means of upgrading the RPGS are currently being investigated. The options being considered are to replace the RPGS processor with one of a higher specification or replace the RPGS entirely with a commercial tough book, compromising ruggedness against cost. For example the current cost of sending the RPGS to a repairer just for an estimate is equal to the cost of buying a commercial tough book with a three-year warranty.

17. Despite the documented problems with the ground support system [3], the Chinook HUMS has remained operational and useful data is being generated and presented to the engineers to act on. This has only really been possible due to the structured maturity programme the MoD has followed. If all of the functionality had been switched-on on the first day of operation,

the ground support system would have collapsed and system credibility would have been lost. The maturity process took each major function in turn, configured it as necessary so that it was working at full efficiency, before that function took precedence over the existing technique. The maturity process now is drawing to a conclusion with the fleet wide implementation of the Rotor Track and Balance (RTB) functionality. A similar maturity period has been scheduled into the programme for the follow on programmes.

AVAILABILITY, RELIABILITY AND MAINTAINABILITY

18. UK military regulation dictates that if an aircraft is fitted with an accident data recorder and if it is unserviceable then the aircraft cannot fly. When the system first became operational, a number of early Chinook sorties were cancelled as the Start-up Built In Test (SBIT) reported that the Data Acquisition & Processing Unit (DAPU) was unserviceable. In every case the DAPU unserviceable caption was generated from missing HUMS components rather than the failure of CV/FDR. A change was made to the SBIT software to differentiate between HUMS functionality failures and CV/FDR functionality failures. Following the implementation of the change no sortie has been cancelled from unserviceability.

19. The reliability requirement for GenHUMS is for the aircraft to have a 99% probability of successfully completing an 8 flying hour working day without experiencing a system failure. This equates to a minimum Mean Time Between System Failure (MTBSF) of 795 flying hours. The Chinook GenHUMS has just completed its In Service Reliability Demonstration (ISRD) [4]. Every HUMS maintenance arising from the selected aircraft was reviewed and assessed against the reliability requirements. A number of the arisings were clearly attributable to physical abuse and were discounted. All Line Replaceable Units (LRUs) satisfied their individual reliability targets, and the demonstrated system reliability was 5149 flying hours.

20. The only item of significance highlighted during the ISRD was reliability of the accelerometers. The accelerometer life was advertised as in excess of 100,000 hours. To date there have been 12 failures in service indicating an accelerometers life of only 19,000 hours. This appears to be a substantial shortfall, however, when taken into the context that each aircraft has 44 accelerometers and there are 40 aircraft in service, the failure rate becomes acceptable. Notwithstanding this acceptability, an investigation is currently underway to determine common factors contributing to accelerometer failures. It is perceived that many of the failures are related to one or two locations that are prone to abuse when exchanging transmission components.

21. The maintainability targets have also been met. The demonstrated average Mean Time To Repair (MTTR) for all HUMS LRUs and Line Replaceable Items (LRIs) is 41 minutes. The target was set at 45 minutes. It was also demonstrated that 95% of HUMS components take less than 2 hours to replace.

OPERATIONAL BENEFITS

22. Operational benefits from operating HUMS are collated by the Chinook HUMS Implementation Group (CHIG). A procedure has been established for squadron engineers to report all incidents where use of HUMS data has aided their diagnosis and rectification of faults. In the early days, the CHIG were able to capture most incidents but now the use of HUMS in fault diagnosis has become so commonplace that only the major or unusual incidents are reported. Examples of how HUMS is benefiting the Chinook fleet are detailed below:

DIRECT PARAMETER DISPLAY

23. HUMS is an advisory system only. The system is configured so that exceedences, cautions and warnings are not displayed in flight. It is however possible to scroll through the pages of the CCU to access vibration data and directly display the values of the flight parameters being recorded. This functionality has been used frequently by aircrew to determine the correct value when left and right hand gauges read significantly different thus also allowing a decision to be made to continue flying. Whilst this does actually create more maintenance activity in recalibrating gauges, the aircrew have increased confidence in the information presented to them.

UNDEMANDED FLYING CONTROL MOVEMENTS

24. An Undemanded Flying Control Movement (UFCM) occurs when a component of the flight control system experiences a glitch or a failure that results in the flight controls being restricted or forced into doing something other than what has been demanded. It is a phenomenon on all rotary wing aircraft (whether publicised or not) and on average there are 4 occurrences in the Chinook fleet each year. The standard procedure for a UFCM on a Chinook is to land and undertake a comprehensive diagnosis routine. The routine essentially starts at one end of the flight control system and works through to the far end. If nothing obvious is found, as is normally the case, flight control components get exchanged until the problem is rectified. It can be a very hit and miss affair and typically the routine takes 2 to 3 days to release the aircraft back to operational service.

25. Following a UFCM on a HUMS aircraft, the routine is very much simplified. The first action is to extract the flight data file from the system. If the pilot pressed the event marker when the UFCM was experienced, 10 seconds of pre and post event flight data will be captured on the DTD. If the event marker is not pressed or the engineers prefer to see data from over a longer period, the full flight data file can be extracted from the Crash Protected Memory (CPM). The flight data trace is then scrutinised to see if it is possible to identify which component initiated the UFCM. In many cases the engineer is simply looking for a change in the normal operating behaviour of a parameter, such as a flat line or a step increase.

26. Use of HUMS has reduced the time to diagnose and rectify UFCMs by 75%. It also prevents unnecessary removal and replacement of serviceable flight control components from the aircraft. The requirement to extract HUMS data now a formal part UFCM procedure.

EXCEEDENCE MONITORING

27. The exceedence monitoring thresholds were initially set to those called up in the Aircraft Maintenance Manuals (AMMs). Through maturity these have been refined to give the engineers the earliest indication of a potential problem. By introducing HUMS, an excursion over the limits of the AMM can be accurately reported rather than relying on the aircrew to observe the exceedence and report it. Accurate reporting of exceedences can add to the maintenance burden as well as reduce it as the following two examples illustrate.

28. A Chinook was operating as part of a four-ship deployment to support a major exercise in the Middle East. It was heavily laden with troops and cargo. The aircrew were concerned that when making a rapid ascent that they over-torqued the rotor system for a considerable time. They were sufficiently concerned that at the next refuelling point they grounded the aircraft for an investigation. The current Chinook over-torque routine takes a minimum of 5 hours to complete properly. This increases if a component is suspected to be damaged and requires replacement. Examination of the exceedence log on the HUMS ground station did show that the aircraft had indeed over-torqued but for a shorter period of time than perceived by the aircrew. The length of time in over-torque conditions was permissible by the AMM and the aircraft was released back on task by the time the refuelling was complete. Subsequent investigation of the full flight data file from the crash protected memory revealed that the aircraft was flying at the edge of the torque envelope for the majority of that sortie, but had only breached the limit the one time.

29. With only 4 aircraft supporting a major exercise the loss of one aircraft (25% reduction in lift capacity), even for a short period of time, can have serious ramifications. Without the HUMS data, the engineers would have had to make a decision based on the verbal description of the incident by the aircrew. The only option available to them would be to take the aircraft off task and conduct the full corrective routine. HUMS provided accurate information to aid the engineer's decision-making process, which resulted in the aircraft being released back on task.

30. A second Chinook was supporting another exercise. On cessation of the day's activities the HUMS data was being downloaded. Whilst waiting for the exceedence report to be generated the aircrew advised that they had an over-temperature caption on one of the engines. The aircrew advised that they had caught it immediately, recovered the engine and kept a close watch on it for the rest of the day whilst they carried on with their tasking.

31. The HUMS report clearly indicated that the engine had exceeded its maximum temperature limit. In fact the incident had raised 3 separate exceedences. The first was that Power Turbine Inlet Temperature (PTIT) was greater than 870°C for longer than 10 minutes, the second that PTIT was greater than 905°C for longer than 10 seconds and finally that PTIT was greater than 938°C. As a result an engineering investigation was carried out.

32. The full flight data file was extracted and scrutinised. The engine temperature was steadily climbing above its normal operating temperature for approximately 27 minutes before the third exceedence was initiated. The logic for the third exceedence replicates the logic used for the caption on the Cautions & Warning Panel (CWP). The data trace showed that the third exceedence was flagged at the same time as the caption illuminated. The trace also indicated that the caption was not acknowledged for 24 seconds and that the engine was recovered shortly thereafter. The maximum temperature reached was 1200°C. The AMM only permits excursions above the maximum temperature for 12 seconds so the engine was replaced. The subsequent strip report from the engine overhaul facility revealed that a number of components, bearings and seals were significantly damaged from the high temperature.

33. Without HUMS and based on the aircrew report, the engineers would have given a cursory look over the engine for signs of damage and then release the aircraft back into service. It is possible that the engine would have failed completely on a later flight and could have caused a catastrophic accident. With HUMS, the engineers were fully able to determine what actually happened to the engine and were able to make the correct decision to replace it. This instigated a expensive and unscheduled engine change and overhaul but increased the airworthiness of the aircraft significantly.

ROTOR TRACK AND BALANCE AND VIBRATION MONITORING

34. There have always been high expectations from the HUMS automated Rotor Track and Balance (RTB) functionality. It has been disappointing that RTB has had to be left until the end of the maturity process particularly as the balance of investment calculations was centred on full RTB functionality from day one. The primary reason behind the delay was the redesign of the nose-mounted camera to avoid potential airflow problems over the pitot static ports. The redesigned nose-mounted camera was not approved until late 2000 and retrofit activities have only just been completed. In addition the RTB functionality has been held back to enable ground station hardware upgrades to be implemented to ensure RTB processing times are acceptable to the user.

35. Currently the existing procedures using the carried on board Rotortuner equipment are still enforced. A complementary set of HUMS RTB data is recorded at the same time and used to prove out the efficiency of the system. The swap over to HUMS RTB is scheduled for early 2003.

36. To gain familiarity and experience with the system, aircraft coming out of maintenance periods are tuned using HUMS RTB directly. Historically it has taken up to 8 flights to tune the aircraft and clear any maintenance arisings. Of the aircraft tuned using HUMS, only one has required more than three flights to bring the rotor system within limits. The design aim for aircraft coming out of maintenance is that tuning the rotor system should take no more than 2 flights. This is considered achievable, once hardware upgrades have been implemented. Squadron requirements for RTB check test flights should also be limited to post major component exchange. Dedicated check test flights to resolve vibration defects will be eliminated as the equipment is permanently carried on board and vibration is monitored every flight.

37. Though the HUMS RTB and vibration monitoring functionality has yet to take precedence over the existing system, the number of call outs for the RAF Odiham vibration specialists to undertake dedicated check test flights on the squadrons has dropped significantly. The specialists have been able to refer to the previous flight's HUMS RTB data, and using their experience, they have been able to look at the vibration characteristic and determine the cause of the reported problem. Previously they would be required to install the carried on board RTB equipment and schedule a dedicated RTB check test flight. Examples of faults rectified in this manner are inactive Self-Tuning Vibration Absorbers and faulty lag dampers. With the introduction of the HUMS rule based reasoning for RTB and vibration monitoring the squadrons will be able to determine the cause of such failures for themselves.

TRANSMISSION BEARING FAILURE

38. One of the most notable benefits from HUMS to date was the use of the system to perform fleet wide health check monitoring following a break up of a combiner transmission bearing [5].

39. During an exercise in Oman, an aircraft was on a sortie when a CWP caption flashed. No other indications of anything untoward were given and after a cursory check of instruments the flight continued. Half hour later the transmission debris light flashed sporadically for 10 seconds, then came on permanently. The aircrew, following the instructions in their flight reference cards, were attempting to land when the left hand transmission warning light came on. This was re-set but immediately latched on again. In the final turn into the landing, the combiner chip and right hand debris warning lights illuminated.

40. Safely on the ground, an investigation was carried out. Oil levels inside the transmission system following the incident were well within limits. The combiner transmission debris screen was removed and large pieces of debris were discovered.

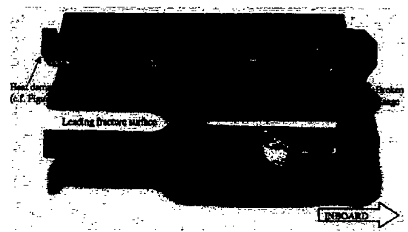

Figure 2: Failed Transmission Inner Bearing Race

41. Due to the serious nature of the incident, the combiner transmission was sent immediately to DARA Almondbank for investigation and overhaul. DARA stripped the transmission and discovered that large quantities of debris had come from the bearing supporting the inner end of the left hand input pinion. The inner race had undergone a complete cross-sectional failure (Figure 2).

42. The aircraft was subject to both Spectrometric Oil Analysis Programme and Wear Debris Sampling but neither had identified the impending failure.

43. The aircraft was fitted with HUMS and though the transmission vibrations threshold and conditions indicators had not yet been matured, raw transmission data was being recorded. The raw HUMS data from that aircraft was evaluated and the failure characteristic was identified. It became apparent from the data that the spalling had initiated at least 95 hours prior to the overload failure. The failure vibration characteristic was converted into a HUMS condition indicator and it was possible to screen all of the other HUMS embodied aircraft within 12 hours.

44. The screening established that no other transmissions displayed similar failure characteristics thus allowing GenHUMS embodied aircraft to remain available for operations. It was necessary for those aircraft not embodied with HUMS to be fitted with dedicated vibration sensors every 25 hours to check and monitor the failed bearing until HUMS was fitted.

45. If HUMS had not been fitted to the aircraft with the failed transmission it would not be possible to identify the failure vibration characteristic. Therefore the only option available to the engineers, to ensure airworthiness, would be to ground the fleet and remove, inspect and replace all combiner transmissions. The last time the fleet was grounded (due to a similar failure) the remove, inspect and replace programme significantly impacted fleet availability for at least 9 weeks.

BENEFITS COMPARISON

46. The original 1993 Project Definition Study (PDS) [6] indicated that, on average, each platform would make a return on investment within 6¼ years. To achieve this return, HUMS would need to generate approximately £13m worth of savings each year. Table 1 illustrates the percentage saving for each user requirement that make up the annual saving.

Table 1: Benefit breakdown per user requirement

Engine Health Monitoring	6%
Transmission Health Monitoring	32%
Rotor Track and Balance and Airframe Health monitoring	45%
Aircraft Usage Monitoring	17%

47. Each element above includes cost savings resulting from HUMS preventing incidents and accidents. In total this equates to approx £4.6m per year. This is a significant amount (36%) and overpowers the other benefits.

48. The savings from accident prevention were determined by an analysis of historical accident records and by assessing whether HUMS would have played a part in preventing the accident from occurring. The PDS analysis indicated that HUMS would have prevented 1 aircraft that required to be repaired on site by specialist personnel, 6 aircraft that required to be returned to works for repair and 7 aircraft that were lost or damaged beyond economical repair. The cost to buy replacement aircraft or to undertake the repairs were averaged and used as an annual HUMS saving.

49. The investment appraisal was reworked in 2000 to support the approval exercise for the Sea King, Puma and Lynx programmes. At this time it was agreed to remove the preventing incidents/accidents factor from the savings equation. Though making a significant contribution to the investment appraisal, it was felt that this factor was impossible to realise. Nobody will give the MoD money back for not crashing an aircraft.

50. The investment appraisal for the 2000 approval focused on the savings identified from the maintenance benefits of HUMS alone that are achievable internally in the MoD. The revised appraisal predicted a total annual saving of £8.84m (£6.85m if deflated to 1993 pricing levels). The breakdown for each user requirement was broadly similar to those shown in Table 1.

51. The indicative implementation costs from the project definition study were approx 50% less than those achieved by the HUMS contract competition. The increased cost along with the reduced annual benefit resulted in the return on investment extending from an average 6¼ years to 19 years. The business case for Sea King, Puma and Lynx was primarily centred on satisfying duty of care and airworthiness requirements.

52. The Chinook element of the revised annual saving is £3.3m. By applying the percentages from Table 1, the MoD should be expecting a saving each year of £200k from engine health monitoring, just under £1.5m for rotor track and balance, £1m for transmission health monitoring and £600k from aircraft usage monitoring.

53. The question is whether the Chinook system is providing the predicted returns. It has proven difficult to ascertain the exact savings for each HUMS arising. The CHIG however has made a financial assessment of the reported arisings and attempted to compare them against the predictions. The word attempted has been used because the predictions were broken down into four distinct key user requirements. In practice it is not so simple to break down the arisings into a single category and the benefit is often shared.

54. Use of HUMS data to resolve defects have prevented engines from being rejected unnecessarily. The unscheduled overhauls cost has been saved so the engine health monitoring element of the annual saving is currently greater than the prediction.

55. The benefits from the RTB functionality is only just being realised due to the redesign of the nose tracker. The consensus from the engineers is that HUMS will significantly reduce the need for dedicated RTB test check flights so the predictions are realistic. The problem is that the majority of the benefit comes from a reduction in flying hours for maintenance activities. Whilst this is may be achieved, it is not a cash benefit to the MoD overall as the saved flying hours are being consumed operationally. The net affect is not a direct cash saving but an increase in operating efficiency and availability.

56. The use of HUMS to check the health of the transmission bearings in the fleet was a good example of how transmission health monitoring benefits can be realised. By itself, this example alone has exceeded the prediction. (The prevention of the accident itself would be recorded as an airworthiness benefit.)

57. Accurate recording of aircraft usage is also identifying benefits. The aircraft maintenance schedule is being driven by the flying hours recorded by the aircrew. The average 20% difference between aircrew and HUMS recorded flying hours will permit a 13-week extension between minor servicing periods. As this will reduce the total number of minor and major servicing periods each year the AUM prediction is most likely to be exceeded. Again with the aircraft in maintenance less frequently there is an increase in operating efficiency and availability.

58. The airworthiness benefits from HUMS have been quite clear even though their potential savings are very subjective. Two instances have been flagged by HUMS, which if left unchecked,

would have resulted in the loss of the aircraft. Assuming a purchase cost of £20m each, it could be argued that HUMS has saved the MoD £40m and that is before the operational impact on the loss of two airframes is taken into account. Admittedly the argument fails the realisable check but gives an insight to the enormous potential of HUMS.

59. In conclusion the Chinook system, once the maturity programme is complete, has every possibility of meeting its predicted savings target. It is very evident that HUMS is saving engineering man-hours and improving the operational availability of the aircraft. These savings will not necessarily result in a realisable cash sum but most likely as a percentage improvement in engineering team productivity and efficient aircraft use. This percentage may be a more accurate indicator of the success of HUMS.

FUTURE ENHANCEMENTS

60. Even though the system has only been in operational use for the last two years, enhancements to the baseline functionality are already being considered:

CV/FDR Replay Station (CRS)

61. The CRS is used to download the full cockpit voice and flight data file from the HUMS CPM. The original HUMS strategy envisioned the use of the CRS after incidents only and not on a regular basis. As such the authority to download the information could only be obtained from the Station Commanding Officer. In operational service the engineers have found the full flight data to be extremely beneficial when diagnosing defects and the authority for download has since been delegated to the squadron senior engineer. (Authority for voice data download remains with the Station CO.) Due to its increased use at the main operating base and during deployments, it has been necessary to increase the numbers of Chinook CRS'S from 3 to 9.

Full FDR download

62. The increased number of CRS's is seen only as a temporary measure. It is the intention to modify the airborne system software such that the full flight data file is downloaded along with the HUMS data each flight. The exact means for download have to be worked out, with the current poor processing performance of the system, the last thing that is required is to slow it down further by trying to process up to 8 hrs of flight data. The initial arrangement will be to have the information on the card and only downloaded to the ground station, if there is a need to, following the maintainers report. If the system is enhanced to include fatigue damage accumulation there will be a need to record flight data continuously. For this the ground station will need to be configured so that it automatically processes the data and archives it correctly to the right tail number without effecting normal HUMS operation.

Fatigue and Usage Management

63. The individual fatigue and usage management system tools developed by MJA Dynamic (Now Smiths Aerospace - Electronic Systems Southampton) for the MoD and presented at HUMS 1999 [7] have now been collated into a single comprehensive toolset referred to now as the Flight Usage Management Software (FUMS) [8]. This name change reflects the change in direction of the tools from direct fatigue calculation to transforming aircraft measurements into diagnostic/prognostic usage information. The aim of FUMS is to provide further improvements in aircraft management, affordability, airworthiness, availability and performance. At the same time FUMS aims to reduce the logistic burden of handling large volumes of data downloaded from individual aeroplanes. Examples of the FUMS toolset are:

Generation and accumulation of Usage indices

64. The usage indices provide concise summaries of aircraft data and at the same time indicate the impact of usage on component condition and life. The usage indices can provide operational management and safety benefits by informing the user about:

- Missions that cause severe usage
- Aircraft configurations that cause severe usage
- Flight conditions that cause high aircrew workloads
- Flight events that cause sever usage; and
- Flying practices that cause severe usage.

By accumulating the usage indices the remaining component life can be assessed.

Monitoring of Operational Exceedences

65. HUMS at present looks at individual parameters and advises when individual physical thresholds have been breached. With FUMS, the software will look across all parameters and identify operational exceedences and significant flight events, by comparing recorded measurements with operationally determined thresholds and flight envelopes. For example, having picked up a heavy underslung load the aircraft might breach its permitted flight envelope whilst trying to navigate through a steep sided valley. FUMS will trigger a warning when the aircraft is flown close to the extreme of its flight envelope.

66. The FUMS prognostic functions can either be implemented as part of the airborne system or as part of the ground station. FUMS will also introduce a prognostics architecture that facilitates the integration of technologies developed by third parties and harmonises their information with the FUMS operational infrastructure. The current FUMS toolset works from the basic data outputs from the installed generic HUMS system and has been developed to enable the technologies to be demonstrated and evaluated quickly and at low cost. As FUMS elements mature they can be introduced into the mainstream HUMS programme.

Improved Data Manipulation & Flight Animation

67. Two elements of FUMS that have already matured and could be introduced into the HUMS programme within a short timescale, are the improved data manipulation and flight animation tools. Trying to reconstruct what is happening with the aircraft from 2 dimensional traces displayed on the CRS can be quite difficult, particularly if during flight the aircrew did not mark the incident you are investigated. The data manipulation tools on the CRS are basic and it can be time consuming to extract the information required.

68. The data manipulation tools developed under FUMS have greater functionality. They will allow the user to design his preferred screens and reports that contain the important FDR information. The enhanced tools have the ability to very quickly translate the data sets onto a graph and correlate the outputs. The tools also permit the user to graphically construct Structured Query Language (SQL) queries on the flight data. By using simple mouse drag/drop actions, fields can be selected, compared and anomalies identified without any typing. In this way the users workload is reduced and complex SQL statements are created without the need to learn SQL. Importantly by using the enhanced data manipulation tools, the ability to write incident/investigation reports will be simplified as it is possible to directly export any resultant data tables and graphs into most word processing documents.

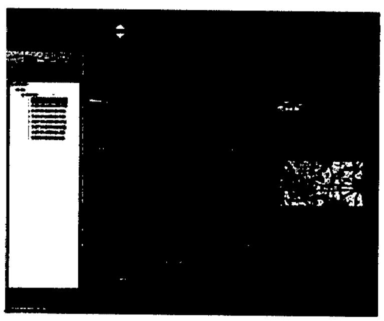

Figure 3: FUMS Flight Animation Screen Shot

69. To complement the data manipulation tools, a flight animation facility is being provided. A typical screen shot is shown in Figure 3. This is to aid the engineer to visualise what is happening to the aircraft at time of the incident. The flight display is correlated directly to the flight data and the display can be used to fast forward to the relevant point in the flight of interest. The animation tool allows multiple views of the aircraft flying, including views from inside the cockpit with real time instrument display and external views from any angle. Side and plan views are also provided and the animation can also be flown against satellite imagery and textured maps. If the FUMS tools are utilised, the fatigue damaging events can also be displayed. The flight animation tool also has a huge potential as a pilot training aid for reviewing sorties particularly if it is combined with the mission planning system and threat data.

Operational Data Recording Exercise

70. To comply with flight safety regulation each helicopter fleet has to conduct an Operational Data Recording Exercise (ODRE) every 5 years. The ODRE requires an aircraft to be fitted with strain gauges so that the structural loads experienced in flight can be measured directly. The measured loads are analysed to assess whether the assumptions made in the aircraft manufacturers fatigue calculations remain valid (i.e. the stresses experienced map the stresses predicted). To complement the strain gauges a full set of flight data is required. The instrumentation for the ODRE is quite extensive, and as the exercise typically last 12 months, therefore the cost of ODRE can be quite large. HUMS can help reduce the cost in two areas

71. The first is to reduce the cost of the ODRE installation by utilising HUMS FDR information to support the data gathered from the strain gauges. A second memory card receptacle will be installed in the Chinook for its ODRE to collect the FDR data. This is to simplify the conduct of the exercise without affecting normal crew operation. It is likely that for future ODRE's the data will be held in a partition of the main memory card.

72. The second means to reduce ODRE cost is to utilise the advanced usage monitoring techniques from FUMS. The techniques can map the actual aircraft usage against the design usage spectrum. Therefore eliminating the need to conduct ODRE every 5 years. The requirement can be satisfied by conducting a load survey on introduction into service, and then continuously monitoring the aircraft against the manufacturers design usage spectrum.

CONCLUSIONS

73. The HUMS programme in UK Military helicopters took a long time to become established and was subject to a number of technical difficulties, all have now been resolved and the lessons learnt from the Chinook programme are benefiting the follow on programmes. The follow-on programmes are now firmly underway.

74. It is still early days for Chinook HUMS but already sizeable benefits are being realised and expectations from the project definition study are broadly being met. Even with the system in its infancy, it is necessary to look forward and start to look at ways the system can be improved and the HUMS data further exploited.

75. The actual cash sum saved by HUMS will always be contentious and difficult to prove. It is however very evident that HUMS is saving engineering man-hours and improving the

operational availability of the aircraft. Similar benefits are eagerly anticipated from the Sea King, Puma and Lynx user communities.

ACKNOWLEDGMENTS

The author would like to acknowledge the support of Cdr James Gourlay, for his assistance in proof reading the paper, and his MoD colleagues in the HUMS IPT and AD/AHMG. The support to the HUMS programme from the team at Smiths Aerospace Electronic Systems Southampton is also acknowledged. Finally the dedication, enthusiasm and perseverance of the staff at RAF Odiham, particularly in the Health and Usage Centre, in making Chinook HUMS a success is also acknowledged and greatly appreciated.

REFERENCES

1. Draper, A. MoD HUMS: Today And Tomorrow, ERA Avionics 2000, London, Nov 2000.

2. Draper, A. An Update On HUMS In UK Military Helicopters, American Helicopter Society Forum 58, Montreal, June 2002.

3. Cook, J. *et al.* Critique Of The Chinook GenHUMS Ground Support System. ES(AIR)(G)/8976.1 AD MM, July 2001 - Unpublished MoD report.

4. Cokayne, I. Chinook HUMS In-Service Reliability Demonstration Final Report, HUMS/76/8/1/3, Nov 2002 - Unpublished MoD report.

5. Cook, J. *et al.* Metallurgical And HUMS Data Evaluation Of Failed Roller Bearing, TM02.216, AD AIM, Aug 2002. - Unpublished MoD report.

6. Larder, B. Project Definition Study for HUMS for MoD Helicopters, SHL1156, Dec 1994 - Unpublished MoD report.

7. Draper, A. Fatigue Usage Monitoring In UK Military Helicopters, HUMS 1999, Melbourne, Feb 1999.

8. Azzam, H. FUMS Software Functionality For Retrofit And Future Prognostics Health Management, 20 Nov 2002 - Unpublished MoD report.

A Study of Usage Variability and Failure Probability in a Military Helicopter

Lt Cdr Jon Milsom RN
Equipment Support (Air), Structures Support Group, Aircraft Structural Integrity 3 - Helicopters, RAF Wyton, Cambs, UK

Nick Toulas
Cranfield University, Cranfield, Beds, MK43 0AL,
Tel 44-(0)-1234-750111ext2371; Email n.toulas.1999@cranfield.ac.uk

Phil Irving
Cranfield University, Cranfield, Beds, MK43 0AL,
Tel 44-(0)-1234-754129; Email p.e.irving@cranfield.ac.uk

ABSTRACT

A study has been made of the loads variability for each manoeuvre and the variability of each manoeuvre during normal operation of a UK military helicopter. Three components were monitored. One of these was a location on the lift frame, the others were a main rotor blade linkage and a rotor pitch change link termed the 'spider arm'. Service loads were measured using strain gauged components and a data recording system. Usage was monitored by manual identification and recording of manoeuvres throughout the helicopter flight. It was found that usage variability is very great, with coefficients of variation in excess of 100% for the majority of manoeuvres. Using a deterministic approach to damage calculation it was found that taking worst case values of manoeuvre damage, usage measured in service was significantly more benign for all types of mission, than that assumed in design, allowing the possibility of future life extension past the nominal safe life on usage monitored aircraft. These conclusions are compared with results from Monte Carlo simulations involving entire distributions rather than worst case assumptions. Uncertainties and variability in individual manoeuvre damage are compared with variability in component life.

INTRODUCTION

Both helicopter health and helicopter usage monitoring technologies are being implemented in civil and military helicopters. While the benefits of health monitoring are very clear in terms of avoidance of potentially fatal accidents arising from unexpected component failure, the benefits and consequences of usage monitoring are more complex and in many ways less obvious. This is particularly true if usage monitoring is to be used as a basis for prognostics, prediction of the future life of the component. In this context issues of probability of particular damage levels occurring become of importance. These are related to both variability in materials properties as well as usage variability.

USAGE MONITORING - CURRENT TECHNIQUES

The manufacturer substantiated Safe Life of a helicopter component is derived from the fatigue characteristics of the component material, the S-N curve, the assumed usage of the helicopter in service and a knowledge of the loads experienced during manoeuvres. The assumed usage spectrum of the design in terms of type and frequency of manoeuvres per flying hour and the operating mass and altitude bands is generally agreed with the principal customer. This is termed the Design Spectrum. The helicopter is then designed to meet this requirement using loads derived from calculation and previous designs. When the prototype is constructed, it is instrumented to measure the actual loads and flown throughout the Design Spectrum. This exercise is termed the Flight Load Survey (FLS). The scope of the FLS varies between manufacturers from a single representative manoeuvre to a large number of repetitions at varying masses with the highest recorded load being assumed to occur in all instances for fatigue substantiation. The measured loads are compared with the calculated values and the life adjusted accordingly. This substantiated life is usually re-qualified with each subsequent production version of the design. The techniques do vary between manufacturers with varying safety factors applied at each part of the substantiation process. Some countries, such as the UK, mandate the fatigue techniques and factors to be applied [1].

There have been a number of investigations into practical methods of measuring both the usage and the corresponding loads of in-service helicopters. The measurement of usage is the simpler. This can be done as simply as recording the actual usage using an experienced non-flying operator to record the manoeuvres flown using a pen and paper or laptop computer. The recorded manoeuvres are normally the same as those detailed in the Design Spectrum but this can be expanded to include any new manoeuvres identified with the operator prior to the commencement of the exercise. This type of exercise is termed a Manual Data Recording Exercise (MDRE) by the UK Ministry of Defence (UK MoD) and is used as the basis of a Statement of Operating Intent and Usage (SOIU) document. This document is passed back to the manufacturer, (Design Authority - DA) by the UK MoD, who review the revised spectrum to ensure that the original design fatigue lives remain valid. This approach was also used by the Dutch Navy to measure the actual usage of their Lynx fleet and achieved a 10% extension in the airframe life of their fleet [2]. This type of data gathering exercise is quite cheap to instigate but does require non-flying aircrew to occupy a dedicated seat in the helicopter and accurately interpret manoeuvres during normal in-service flying operations. It is also necessary to perform the exercise over a 12-18 month period to ensure that sufficient data is gathered to produce a statistically valid average usage. From the data it is possible to provide the DA with a series of typical sorties types dependant upon the role in which the helicopter is being used. Depending upon the severity of usage in each role it may be possible to apply a penalty factor to certain components when flying more damaging types of sortie.

A more accurate type of usage measurement is to automatically record and interpret the manoeuvres flown on completion of the flight by downloading captured parameter data to recognition software that produces an output of the types and duration of manoeuvres performed during that sortie. This is generally termed Flight Condition Recognition (FCR). As its name implies, FCR involves the recognition of specified flight manoeuvres by analysis of characteristic parameters. The flight manoeuvres recognised can correspond to the manoeuvres used in the Design Spectrum or they can be expanded upon to reflect any number of frequently performed and hence recognisable manoeuvres.

A number of FCR systems have been developed. The US Army, who were keen exponents of the helicopter's potential, had commenced work on a concerted effort to acquire realistic service data for their operational helicopters in 1964 [3]. The US Army Helicopter Service Usage Program was a comprehensive recording exercise of the usage of 5 helicopter types in 4 operating environments. In addition an Operational Load Survey was conducted on an AH-1G attack helicopter in order to obtain detailed knowledge of the flight loads experienced in each manoeuvre. By 1975, this program had matured into the Structural Integrity Recording System (SIRS) which was fitted to small numbers of AH-1G and the later AH-1S variant. This system included load measurement of one component and an FCR system that drew data from 8 aircraft instrument parameters

The US Navy, having used counting accelerometers ('g' meters) on its fixed wing fleet to monitor usage for many years, opted to replace this with a Structural Data Recording System (SDRS) [4]. This microprocessor based multichannel system was essentially an FCR system with accompanying dedicated 3-axis 'g'-meter. The success of this system on the fixed wing fleet lead to its adaptation for use in the rotary wing fleet.

The SDRS identifies patterns in the combinations of a number of parameters from onboard instruments. The instruments are a combination of outputs from aircraft systems, transducers within the SDRS system and a number of fixed parameters input prior to take-off such as take-off fuel load and expendable weapon load to allow the gross weight to be determined. The parameters selected are the minimum number required to successfully distinguish between each manoeuvre. A certain manoeuvre can be identified when each parameter satisfies a set of conditions. These conditions were identified by analysis of traces coupled with a corresponding 'pilot -event log' to provide a prompt when a particular manoeuvre was being flown. The SDRS software is designed to highlight a manoeuvre that did not comply with an algorithm as an 'unrecognised' manoeuvre instead of allocating it to the nearest. This highlights any changes of usage and allows a new algorithm to be written. Similar parameter algorithm recognition techniques have been used on the AH-64 Enhanced Diagnostic System [5], and Project HeliSTAR on a Bell 412 [6].

An alternative method proposed by Polanco [7] uses a more mathematical approach that makes allowance for the different weighting in terms of the varying damage rate or cost for each manoeuvre in the spectrum to be taken into account through the use of 'amplification factors'. The method was developed primarily as a means of rapidly calculating the effect, in terms of Life Cycle Cost (LCC) and fatigue damage incurred, of modifying the usage spectrum of an aircraft or helicopter.

An extension of usage monitoring is to carry out the recording of both usage and loads to obtain a more accurate determination of the actual consumption of fatigue in service. This method is generally termed Flight Load Monitoring (FLM). The two principal FLM techniques use either direct load measurement or indirect load measurement that uses mathematical synthesis to identify component fatigue damage.

Direct load monitoring is the simplest method and most accurate method to measure the actual loads that occur during a manoeuvre. Loads acting on parts of the rotables or airframe structure are measured using a number of fixed strain gauges on an example airframe. The output is a strain history that can then be converted into a load for further fatigue evaluation. This technique is termed Operational Data Recording (ODR) by the UK MoD

An FLM system was developed by McDonnell Douglas Helicopters (formerly Hughes, now Boeing Helicopters) for the AH-64 helicopter [5]. This system used 9 strain gauges installed in relatively sheltered positions to measure strains at key points in the transmission and structure. Although results were favourable, it was not adopted by the US Army.

There are a number of techniques that have been developed to calculate or predict the load, or fatigue damage, that occur during manoeuvres. These are characterised by the use of a model based framework that utilises techniques such as regression analysis, holometrics or neural networks. The main purpose of indirect monitoring systems is to organise and develop useful relationships between the chosen parameters and fatigue damage. The indirect systems weigh and associate the parameters in such a way that fatigue damage attributable to a manoeuvre can be confidently calculated.

The principal advantage of indirect techniques is that they dispense with the need to fit costly and unreliable measurement equipment. However, some of the proposed systems are only theories and few of the systems that do exist [8], [9], [10] could not yet be described as mature and are still under development.

One advantage of neural networks is the potential to adjust the 'weighting' of the final synthesised output value due to the detection by the Helicopter Health and Usage Monitoring System (HHUMS), or the external input, such as a fixed value, of an event or situation that may not be distinguished by the usual input parameters. Thus faults due to rotor imbalance, blade delamination, moisture absorption etc can be fed into, and accounted for, by a suitably trained network. Smiths Aerospace has also investigated algorithms that can constantly compute an aircraft's AUM and C of G position [11].

The use of FCR to provide the actual usage of individual airframes reduces one of the levels of uncertainty in one of the 3 factors that go into the deterministic substantiation of the Safe Life of a component. The use of direct load measurement or indirect load synthesis reduces the level of uncertainty in a second factor. This improved knowledge may allow the reduction of some of the additional factors, such as 1.2 on loads and 1.5 on life presently applied to account for currently unknown load and usage scatter.

However, no method can take into account material performance scatter between individual components. The behaviour of such components cannot, presently, be predicted and so the current material Life and Safety factors used to generate a 'working' S-N curve will need to remain in place.

PROBABILISTIC METHODOLOGIES - PREVIOUS RESEARCH

The application of probabilistic methods was investigated by several researchers. Moon, Menon and Brandt [12] used the US Navy SDRS fitted to a fleet of 50 US Navy rotorcraft to explore the variation in usage distribution in the helicopter population. They found that usage variability in each manoeuvre was better modelled using a Weibull distribution. However, the fit was far from perfect. Variability in usage was large, with coefficients of variation ranging from 33% to 273%. These workers also present limited data on the variability of loads in each manoeuvre, with coefficients of variation ranging from 15% to 57%. They use this in conjunction with material fatigue data to perform Monte Carlo Simulation to derive the cumulative distribution of probability of failure of selected components.

Thompson and Adams [13] also performed a Monte Carlo simulation concerning three UH-60 main rotor components. They note that component strength was the overwhelming factor in the determination of component life and that the effects of spectrum and flight loads tended to average out over a lengthy period of simulation. In view of a six-nines reliability it is concluded that three nines are derived from component strength, one nine from usage and two nines from treatment of flight loads. Boorla and

Rotenberger [14] performed a flight loads variability analysis using a Bell OH-58C helicopter. Six pilots flew, 5 times each, a representative flight of 33 manoeuvres. Load variability appears to be normally distributed with a median coefficient of variability of approximately 13%. It was also noted that use of peak loads to characterize load variability may result in more scatter than is truly representative of inherent load variability. Zion [15] carried out a variability study on two BH Model 234 dynamic components indicating that current deterministic methods are conservative enough to exceed six nines reliability. The study was based on loads and strength variability but it is also noted that operational variability in the fleet would be beneficial for component reliability when load severity of the flight test aircraft is high relative to the fleet. However, if the flight test aircraft is flown with a load severity which is low relative to the fleet, the inclusion of operational variability will adversely affect component variability. Harris et al [16] applied a probabilistic methodology to calculate the retirement life of a critical helicopter dynamic component, demonstrating a potential gain of 30% in life compared to the conservative deterministic approach.

MEASUREMENT TECHNIQUES AND DATA ANALYSIS

USAGE MEASUREMENT

The original design usage spectrum of the Lynx Mk9 helicopter is shown in table 1, together with a modified spectrum which reflects the more limited range of manoeuvres recognized by the usage measurement technique. A few manoeuvres, notably cruise turns and control reversals either were not recognised or did not occur in the service measurement of manoeuvre usage. The percent time which these manoeuvres represented were re allocated proportionately among the remaining manoeuvres so that the total manoeuvre time still totalled 100% of a flight hour. Service measurement of manoeuvre occurrence was performed on a lap top computer by observers flying with the helicopter. The measurement records what manoeuvre the aircraft is carrying out and for how long. The measurement not only records what manoeuvre the aircraft is carrying out and for how long but also the mass of the aircraft at take-off, mass of freight carried, no of crew and passengers, any additional fuel taken on board and the condition of the landing. A total of 22 helicopters flying 72 hours (49 missions) was observed to produce the data analysed in this work. The manoeuvre usage spectra and the associated fatigue damage were analysed in terms of operation type and also were compared with the design spectrum. Four types of operation, namely training, anti- tank, personnel carrying, and a utility mission were investigated.

FLIGHT LOADS MEASUREMENT

The investigation was performed using 10 hours of flight load data gathered on a Westland Lynx Mk9 helicopter. Data were recorded as 55 channels of flight parameters, of which 7 were strain or bending moment data from specific components. In this research load data from the dogbone rotor linkage, the spider rotor component and a lift frame component (station 420A) were studied. The strain or bending moment - time data were converted to stress - time data. The stress-time history between a start and stop point for each manoeuvre was then rainflow cycle counted. The rainflow spectrum was then used together with the component constant amplitude S-N curve, in a proportional damage (Miner summation) calculation of the damage and damage rate for each occurrence of each manoeuvre type. Failure was defined as a damage sum of unity. The Goodman equation was used to account for mean stress effects. Constant amplitude (S-N) fatigue data for each component was supplied by the design authority. The mean S-N curve was reduced by applying factors on both stress and life. A factor of 1.2 on the load spectrum was applied. Having calculated the damage content of each manoeuvre, the damage rate was derived by dividing by the manoeuvre duration.

MONTE CARLO SIMULATIONS

In order to establish the influence of usage variability and of different patterns of usage on probability of achievement of service lives, a series of Monte Carlo simulations were performed. These were not full simulations in that variability in fatigue S-N data was not simulated. Fully factored S- N curves were used to calculate a damage distribution for each manoeuvre. This was repeatedly sampled, together with the measured usage distributions found in the different service types. The damage was accumulated, until failure was predicted at a damage sum of unity. The process was continued to yield a distribution of lives for the different usage types. In performing these calculations, both the damage produced by each manoeuvre and the usage spectrum was assumed constant over a 1 hour period of flight. The calculated distributions of service lives were compared with lives from an average usage for all mission types, and also with the distribution predicted for the design usage spectrum. Further simulations were performed to investigate the influence of service time for which the usage spectrum remains unchanged.

RESULTS

USAGE AND USAGE VARIABILITY

A comparison of the measured incidence of manoeuvres (or usage) compared with the original usage spectrum used at the design stage is shown in figure 1 for a selection of the most commonly occurring manoeuvres. Not surprisingly there are significant deviations from the original spectrum, with some manoeuvres having greater incidence than design, (e.g. forward flight at 0.6-0.8 VNE) and others less (e.g. Hover and forward flight at 0.8-0.9 VNE. The coefficients of variation of the usage (std dev/mean values) were very large for all manoeuvres in all types of operation and frequently exceeded 100%. This is shown in figure 1 as range bars around the mean levels.

Figures 2 and 3 show the differences in measured and design usage mean values for the most significant manoeuvres expressed as % time occupied by the manoeuvre in a flight hour. There are major changes in how the helicopter spends each flight hour, when design and measured usage values are compared. Forward flight, hover and rotor turning on the ground occupy almost 70% of the design spectrum and almost 90% of the measured spectrum. As these manoeuvres are largely undamaging these differences do not matter greatly from the viewpoint of component life.

Figures 4 and 5 show a similar plot but based on damages caused by the incidence of the manoeuvres shown in figures 2 and 3. The damage figure used is that caused by a single instance of the manoeuvre, with 50% probability of occurrence. Because the overall average damage per hour is reduced from that of the design spectrum by a factor of about 2.5 for the dogbone component, it is possible for manoeuvres to have a reduced incidence from the design spectrum, but a greater measured percent damage contribution. An example of this is the hover manoeuvre.

On the damage charts shown in figures 4 and 5, the majority (70-80%) of the damage in the dogbone component is caused by hover and sideways flight, with other manoeuvres making little contribution. Monitoring helicopter usage for this type of operation therefore reduces largely to monitoring the incidence of hover and sideways flight and perhaps one or two other manoeuvres. The contribution of other manoeuvres to accumulated damage is negligible.

This comparison is for the dogbone rotor component, other components such as the spider will have different manoeuvres which cause the most damage, and differing damage levels. This is shown in figures 6 and 7 where for the spider, measurements show that helicopter descent produces 74% of the total damage per hour.

Figures 8 and 9 show a comparison for selected manoeuvres of the percent time spent in each manoeuvre for different types of service operation, and of the percent damage contribution. Significant differences both in terms of percent time and percent damage were found between the different types of service. However the overall conclusions noted earlier still applied, in that in terms of time spent in manoeuvres, forward flight and hover manoeuvres were responsible for the majority of time in flight, but in terms of damage hover and sideways flight were responsible for the majority. These conclusions were generally irrespective of the operation type, with one or two exceptions.

LOAD AND DAMAGE VARIABILITY FOR EACH MANOEUVRE

The loads spectra measured for each manoeuvre, and the damage content derived from the load spectra also showed considerable variability. A typical example of the distribution of damage rates for the transition to hover manoeuvre for the component 420A is shown in figure 12 for 77 occurrences of the manoeuvre. Almost one third of the instances did not cause any damage, and the most common damage rate was small. The largest damage rate was over 10 times the smallest, again emphasising that the damage content of each manoeuvre also shows great variability. The level of damage and the type of manoeuvre causing the most damage also varied from component to component.

MONTE CARLO CALCULATIONS

During the Monte Carlo research, it was found that changing the damage content of individual manoeuvres each time the manoeuvre occurred, resulted in extremely narrow distributions of accumulated damage. The damage value tended towards an arithmetic mean value for the distribution for that manoeuvre. Monte Carlo analysis could frequently be simplified by using this mean figure.

Figure 13 shows the result of Monte Carlo simulations in which the extremely variable usage distributions, and the damage distributions which resulted from them, found in the different types of service operation were randomly sampled, and the damage accumulated to calculate distributions of service lives under particular types of usage. Material and damage for each manoeuvre were held constant in all simulations. All of the curves show small variability in life resulting from usage variation and manoeuvre damage variation alone. There are relatively small but distinct differences in life between the different types of usage, with the training spectrum being most severe, and the utility spectrum being most benign. For example at a cumulative probability of failure of 10^{-5} the training usage has a life of 6,500 flight hours, and the trooping spectrum a life of 7,500 hours. These are both significantly longer than the design spectrum life of under 5,000 hours at the same cumulative probability, and considerably more severe than the utility spectrum, which shows a life of over 12,000 flight hours.

These conclusions, regarding the comparison between different spectra, would not be significantly changed if damage in other components were considered. For example, applying a deterministic analysis (mean damage levels) to the dogbone component shows that the utility spectrum has a life of 1,455 flight hours and the trooping spectrum a life of 1,260 hours. Both are longer than the design spectrum life of less than 1040 hours. Also for the Spider component, all spectra have similar lives of around 17,600 flight hours, longer than the design spectrum life of less than 14,300 hours. These conclusions would not be significantly changed in a Monte Carlo analysis as figure 13 shows that the total variability in each usage case is small.

The role of period over which usage remains constant in determining life variability in a population of helicopters is shown in figure 14. This figure shows the results of Monte Carlo simulations on the dogbone component in a helicopter population, where the design usage was defined as the mean, and a coefficient of variation of 80% was simulated for each manoeuvre. These distributions were randomly sampled in the usual way. The flight period over which the manoeuvre proportions was held constant was systematically changed from every one hour of flight up to 500 hours. The extreme case was when the usage was held constant up to component failure. It was found that frequent usage change produces very narrow distributions with little variability, whereas changing the usage each 500 hours produces significant variability largely equivalent to keeping helicopter usage constant for the entire safe life.

DISCUSSION

Although the calculated component lives are similar to safe lives calculated by the design authority, the lives calculated in this analysis should not be compared directly with them. Design Authority life calculation procedures tend to apply different safety factors to damage calculations of different manoeuvres, depending on the level of information available about the manoeuvre.

As in this research a comparison between the severity of different types of usage was required, a common set of factors and procedure was used which could be applied across all manoeuvres. In all calculations, the S- N curve was fully factored and was fixed for each of the three components. Hence neither the component lives or the simulated distributions are real- however they do allow comparisons of the effects of different usage to be made.

The results demonstrate that as found in other recent work [17], [18] the damage content of different manoeuvres is extremely variable, as is the usage of different manoeuvres. However the damage content of different manoeuvres is relatively small requiring many repeated applications of the manoeuvre to cause failure. This means that variability in service life to achieve particular levels of accumulated damage will reduce with increasing service life. It was demonstrated for instance in Monte Carlo simulations that it made little difference to the calculated distribution, whether the variation in damage content of a manoeuvre was fully represented in the simulation input, or whether it was put in as a fixed quantity.

Hence the major loading variable influencing the distribution of lives is the usage type and its variability.

The damage calculations show both in terms of deterministic calculations and Monte Carlo based probabilistic analysis, that there are significant differences between the design spectrum and the actual one with the combined all usage spectrum having 50% greater lives than the design spectrum at a constant cumulative probability of failure of 10^{-5}. Operating the utility spectrum exclusively would result in lives almost double those of the combined spectrum and a factor of three greater than the design spectrum.

The variability in the calculated lives depends greatly on the conditions assumed in the Monte Carlo simulation for the service time interval over which the usage distribution is held constant. Long intervals produce increased variability, small intervals decrease the variability in an analogous way to the lack of variability produced by the changes in damage for an individual flight manoeuvre. To know which interval is appropriate to the service situation it is necessary to know the service time interval over which the manoeuvre usage spectrum remains essentially constant. This information is not available at the moment. Thompson and Adams [13] assumed that an aircraft changes assignments every 1000 hours, or about every three years.

The current data do show that substantial life extensions of this helicopter could be taken over the nominal design safe life, providing that there were no additional increase in likelihood of failure due to mechanisms other than fatigue, such as increased probability of corrosion and or mechanical damage, which was associated with the increased life.

CONCLUSIONS

(1) Measurements of service usage in the Lynx helicopter under 4 different types of service operation has shown that increases in life of up to 50% of the original design usage spectrum are possible.

(2) Although the usage spectra are dominated in time by forward flight, in terms of damage, both design and measured usage spectra for the dogbone rotor component show that hover and sideways flight cause over 80% of the damage. Other components such as the 'spider arm' have different manoeuvres that dominate the damage.

(3) Variability in calculated lives using Monte Carlo simulations is strongly dependent on the time interval over which the relative proportion of manoeuvres in the usage spectrum can be considered constant. There is currently a lack of service data on this issue.

ACKNOWLEDGMENTS

Thanks are due to Andy Stephenson of Agusta- Westland helicopters for provision of component S-N data, and together with Roger Cansdale (Qinetiq) for many helpful discussions on the problems of usage monitoring.

REFERENCES

1. UK Defence Standard 00-970. Leaflet 201.
2. Van Der Horst J. G., Melissant L., "*Lynx Life Extension activities in the Royal Netherlands Navy*", 25th European Rotorcraft Forum, Rome, Italy, Sep 1999
3. Merkley D., Reddick H. K., "*US Army Operational Flight Loads*", Proceedings of the 58th meeting of the Structures and Materials Panel, Sienna, Italy, AGARD-CP-375, Apr 1994
4. Moon Suresh and Hinger Carl, "*Helicopter dynamic components fatigue tracking system*", 18th ICAF symposium, Melburne, Australia, p 995-1010, May 1995
5. Harrington J., Chia D. Neff J., "*MDHC's Enhanced Diagnostic System, A Unique and Comprehensive Approach to Structural Monitoring*", National Technical Specialists Meeting on Advanced Rotorcraft Structures, Alexandria, VA, USA, Oct 1988
6. Basehore M. L., Dickson W., "*HUMS Load monitoring and Damage Tolerance: An Operational Evaluation*", Federal Aviation Authority, Nov 1998
7. Polanco Frank G., "*Usage spectrum perturbation effects on helicopter component fatigue damage and life-cycle costs*", Defence Science & Technology Organisation, http://www.dsto.defence.gov.au/corporate/reports/DSTO-RR-0187.pdf, Nov 2000
8. Azzam H., "*FUMS An Emerging Technology for Improved Safety, Reduced Costs and Increased Availability of Aircraft*", Farnborough International Technology Exploitation Conference, Sep 1998
9. Manry M., Hsieh C. H. Chandrasekaran H., "*Near-Optimal Flight Load Synthesis Using Neural Net Dept of Electrical Engineering*", Published on Internet only, University of Texas at Arlington,
10. Haas D.J., Milano J. Flitter L., "*Prediction of Helicopter Component Loads Using Neural Networks*", Structural Dynamics & Materials Conference, LaJolla, CA, Part 1 Proceedings of the 34th AIAA/ASME/ASCE/AHS/ASC Structures, Apr 1993
11. Wallace M., "*An Investigation into the Feasibility of helicopter AUM/C of G Calculations using HUMS/FDR Parameters*", Unpublished MoD (PE) Report, Dec 1997
12. Moon Suresh, Menon Dinesh, and Barndt Gene, "*Fatigue life reliability based on measured usage, flight loads and fatigue strength variations*", 52nd annual forum of the American Helicopter Society, Washington, p 366-389, Jun 1996
13. Thompson A.E. and Adams D.O., "*A computational method for the determination of structural reliability of helicopter dynamic components*", 46th annual forum of the American Helicopter Society, Washington, D.C., USA, p 859-873, May 1990
14. Boorla Raghupati and Rotenberger Kevin, "*Load variability of a two-bladed helicopter*", Journal of the American Helicopter Society, p 15-26, Jan 1997
15. Zion Lewis H., "*Safe life reliability: Evaluation of new statistical methods*", 47th annual forum of the American Helicopter Society, Phoenix, AZ, USA, p 257-272, May 1991
16. Harris W.D., Larchuck T., Zanoni E., and Zion L., "*Application of probabilistic methodology in the development of retirement lives of critical dynamic components in rotorcraft*", 55th annual forum of the American Helicopter Society, Montreal, Canada, p 1254-1267, May 1999
17. Irving P.E. and Hudson R.A., "*The contribution of health and usage monitoring systems to calculations of damage state and future life of helicopter components under safe life and damage tolerant designs*", RTO AVT specialists' meeting on "Exploitation of structural loads/health data for reduced life cycle costs", RTO MP-7, Brussels, Belgium, p 7-1 7-11, May 1998
18. Irving P.E., Strutt J.E., Hudson R.A., Allsop K., and Strathern M., "*The contribution of fatigue usage monitoring systems to life extension in safe life and damage tolerant designs*", CEAS forum on life extension, Royal Aeronautical Society, Cambridge, UK, Mar 1999

Table 1 Lynx Mk7 – Usage percentage per manoeuvre.

	Manoeuvre	Design		Measured Usage	
		A	B	Mean	COV
1	0 - 0.2 VNE	10.05	10.49	6.18	68%
2	0.2 - 0.4 VNE	5.43	5.67	3.34	68%
3	0.4 - 0.6 VNE	6.33	6.61	3.89	68%
4	0.6 - 0.8 VNE	12.66	13.21	37.53	51%
5	0.8 - 0.9 VNE	19.40	20.23	8.65	170%
6	0.9 - 1.0 VNE	0.90	0.95	3.06	336%
7	*VNO - 50kts (60deg)*	0.80			
8	*VNO - 20kts (45deg)*	0.64			
9	*VNO - 20kts (30deg)*	0.96			
10	0.9 VNE (30deg)	2.51	2.63	0.01	441%
11	1.0 VNE (20deg)	0.11	0.13	0.01	700%
12	*Control Reversal*	0.50			
13	Autorotation	0.50	0.53	0.26	315%
14	Hover	14.10	14.71	9.35	91%
15	Sideways	2.01	2.10	0.48	166%
16	Rearwards	1.01	1.06	0.01	320%
17	Spot Turns	1.86	1.95	1.75	98%
18	*Hover Control Reversals*	1.13			
19	Climb	5.53	5.77	0.84	180%
20	Descent	5.57	5.81	0.25	242%
21	*Take-off*	0.21			
22	Transition to Hover	0.70	0.74	0.49	135%
23	Transition from Hover	0.70	0.74	0.49	135%
24	Landing	0.35	0.38	0.06	259%
25	Rotor Turning on Ground	6.03	6.30	23.37	46%
	Total	100.0	100.0	100.0	

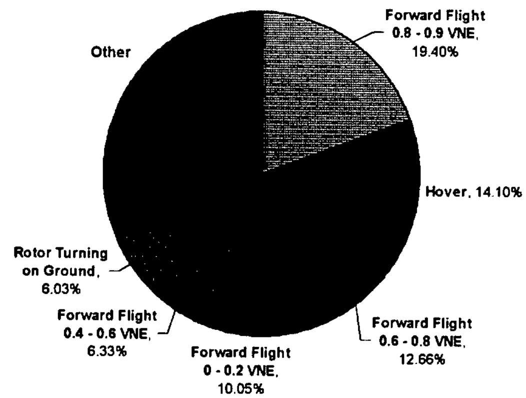

Figure. 2 Design usage percentage - Lynx Mk7.

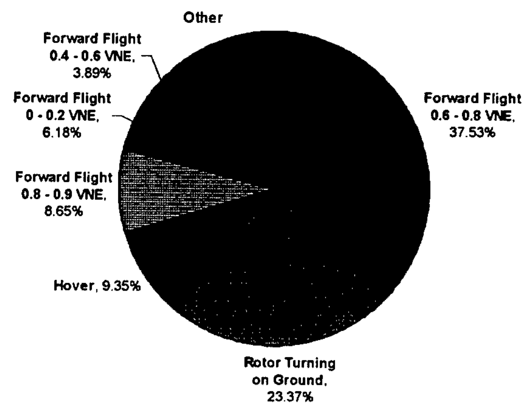

Figure. 3 Measured mean usage percentage – all mission types - Lynx Mk7.

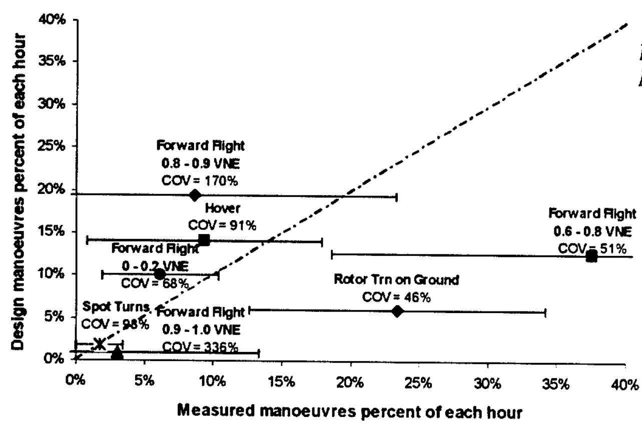

Figure. 1 Measured v Design usage spectrum. Data points are mean values, Bars represent one standard deviation.

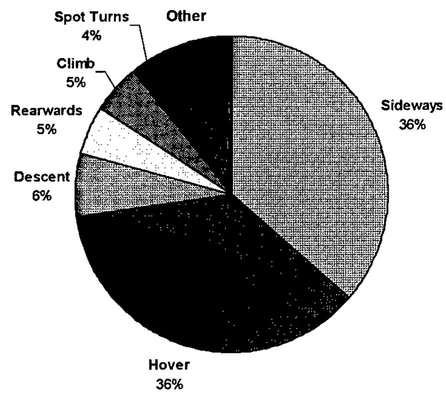

Figure. 4 Damage per hour contribution of each manoeuvre using Design spectrum – Dogbone component.

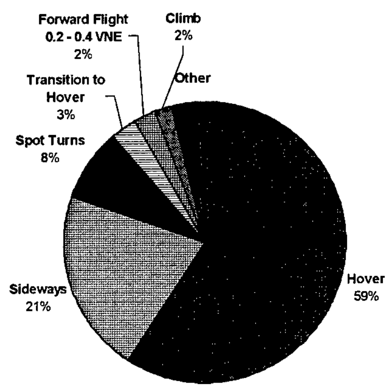

Figure. 5 Damage per hour contribution of each manoeuvre using measured usage spectrum – Dogbone component.

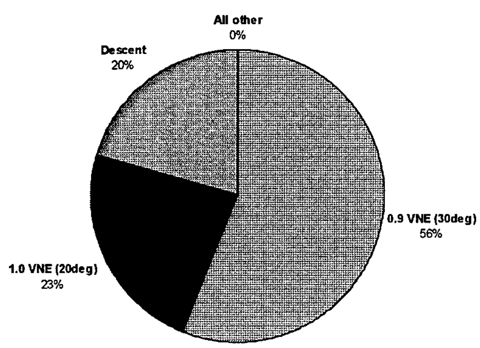

Figure. 6 Damage per hour contribution of each manoeuvre using modified Design spectrum – Spider component.

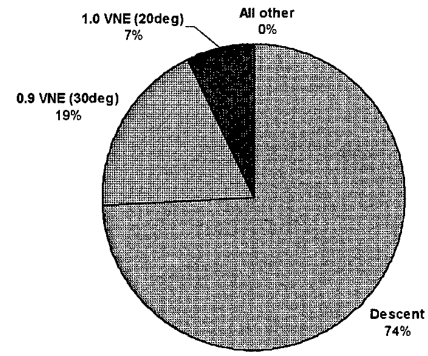

Figure. 7 Damage per hour contribution of each manoeuvre using measured usage spectrum – Spider component.

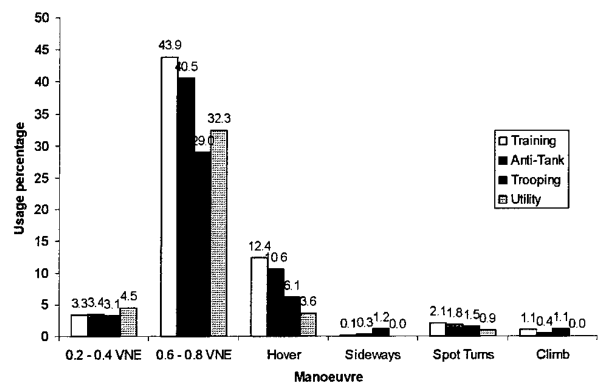

Figure. 8 Measured usage percentage per manoeuvre and mission type – major manoeuvres only.

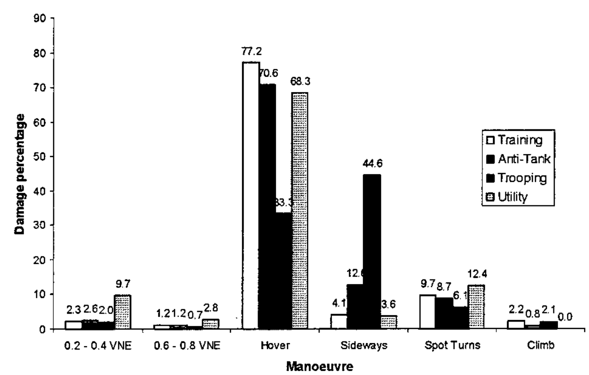

Figure. 9 Damage contribution percentage per manoeuvre and mission type – Dogbone component.

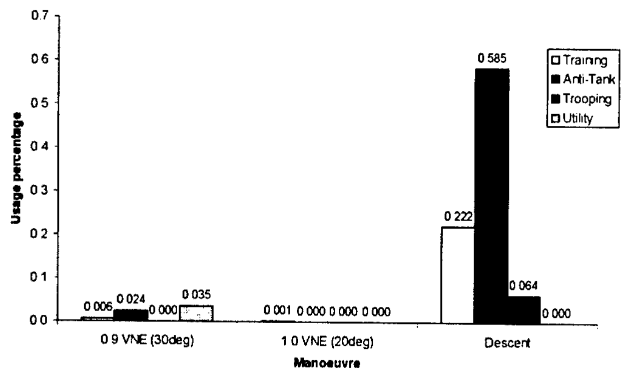

Figure. 10 Measured usage percentage per manoeuvre and mission type – minor manoeuvres only.

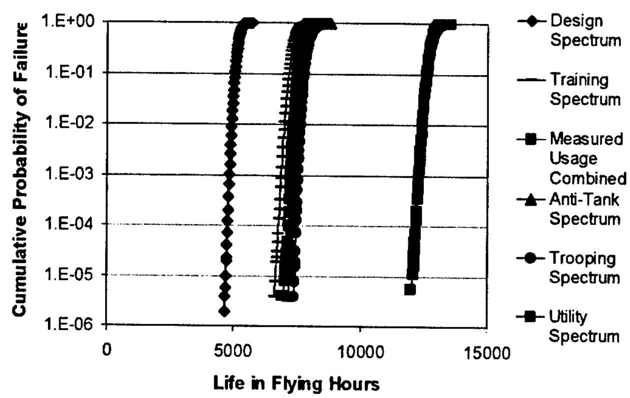

Figure. 13 Cumulative probability of failure v Life hours – (usage variability based on measured COV per manoeuvre) – Stn.420A component.

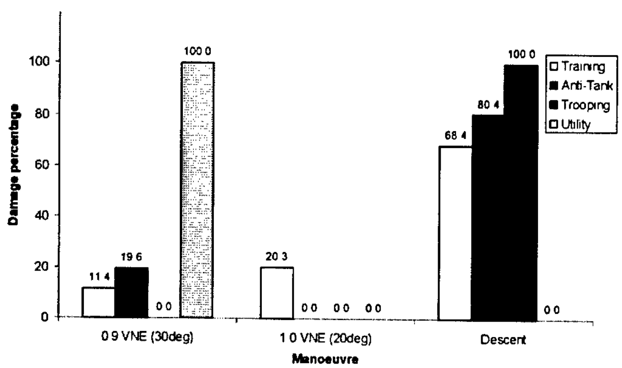

Figure. 11 Damage contribution percentage per manoeuvre and mission type – Spider – component.

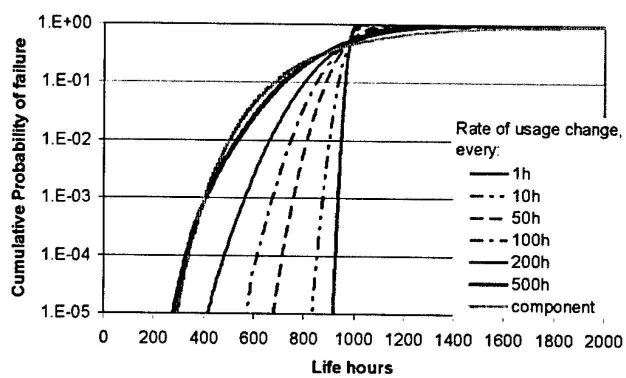

Figure. 14 Cumulative probability of failure v Life hours – (usage COV = 80%) – Dogbone component.

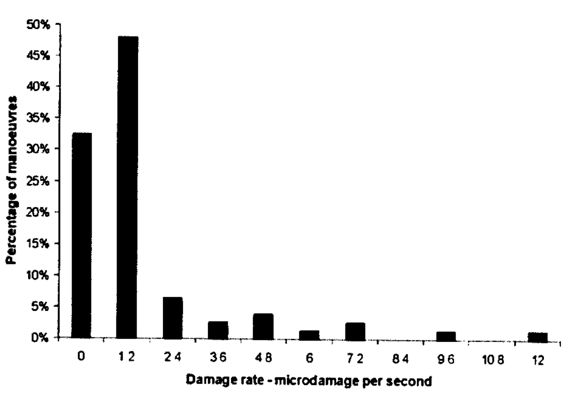

Figure. 12 Damage rate distribution – Transition to Hover manoeuvre – Stn.420A component.

Applying Rules For Isochronous Sampling Within Acquisition Cycles To Airborne Data Acquisition Systems

Richard Fielding
ACRA Control Ltd, Dublin, Ireland

Aaron Mc Nelis
ACRA Control Ltd, Dublin, Ireland

ABSTRACT

As the desire for large quantities of high accuracy and complex flight data increases (eg Airbus A380, JSF etc) traditional data gathering techniques will no longer be practical. The next generation of Airborne Data Acquisition systems must be far more flexible with respect to how data is gathered, transmitted, stored and distributed.

This paper examines two rules for data acquisition that have advantages for today's advanced Flight Test Instrumentation (FTI), Operational Loads Monitoring (OLM) and Health and Usage Monitoring (HUMS) systems where:

- Data is acquired from physically separate test equipment
- Deterministic (IRIG-106 (Ch. 4), ARINC 573 etc) and non-deterministic (ETHERNET, AFDX, JDANS etc) networks co-exist
- Data Acquisition Units (DAUs) from multiple vendors are required
- Signal lists and sampling rates change rapidly
- A time-coherent sampling strategy (even for smart sensors) is required

These rules may aid not only in the selection of the data acquisition equipment but also the definition of the sampling, transmission, storage and flight test data analysis strategies.

INTRODUCTION

There are many questions asked of Airborne Data Acquisition vendors today:

Can you supply a 1553, CAIS, NexGenBus, smart sensor controller? ✓

Have you a plan for FC, Firewire, FDDI, 1Gb Ethernet, 1553++, Loadnet? ✓

Have you thought about non-deterministic networks? ✓

Can you send packets of data to a certain Solid State Recorder (SSR)? ✓

Can you guarantee simultaneous sampling? ✓

Have you thought about packet structures and how they may be handled? ✓

Can I change my mind later about any of the above? ✓

Many Airborne Data Acquisition vendors are content to tick these boxes; and explain later the coherency issues and the problems of stale, skipped and lost packets. This paper argues that there exist some core axioms about which an Airborne Data Acquisition system must be designed. Obeying these axioms supplies a rigorous solution to these challenges rather than simply ticks in boxes.

This paper outlines some of these rules and illustrates some of the problems these rules solve.

THE GENERIC PROBLEM

Data is gathered from many sources:

- Analog to digital converters (strain, accel., video, synchros, and so on)
- Bus controllers (smart sensors, 1553, 429, CAIS and so on)
- Bus monitors (too many to mention)

Subsets of this data are being sent to many sinks

- A few to a cockpit display (VGA, 1553, 429 and so on)
- A few to a telemetry link (perhaps in packets and hence to a network)
- A lot to a recorder (e.g. SSR via FC-AE or 1Gb Ethernet or whatever)

How do we correlate, with respect to time, the data from the various sources? Furthermore how do we present the data to the various sinks in well-defined, succinct packets across networks with limited determinism?

Next, this paper takes an informal look at some old rules by which Airborne Data Acquisition programs were defined. This is followed by some new rules and a discussion of how some of the challenges outlined above are met.

THE OLD RULES FOR SOLVING THESE PROBLEMS

(i) TAG EVERYTHING - With time, stale, skipped and empty (at least)!
How else can data from the controller gathering data via MIL-STD-1553 be correlated with respect to the data gathered from the smart sensor belt and the myriad A/Ds about the system?

(ii) AVOID COMMERCIAL BUSSES - There be dragons!
There was a time when MIL-STD-1553 was a better choice than Ethernet, mainly because the latter did not exist, but even when it did, it was not much faster and had "determinism issues". With 1Gb Ethernet is this still the case?

There was a time when we could only dream of "Decomless" telemetry. Again mainly because the commercial world was not sending large packets of data via telephony in real-time? Is this still the case?

(iii) FORCE "SIMULTANEOUS SAMPLING MODES" - Whatever that is?
In those cases where parameters must be sampled "simultaneously" the Airborne Data Acquisition vendor must jump through some ill-defined hoops and support a "broadcast" sample command. We won't talk about what that means for parameters at different sampling rates or why we just don't do this for all parameters.

(iv) DEFINE DATA PACKETS IN DETAIL - Don't trust the Airborne Data Acquisition **vendor!**
How do we specify that all parameters be sent to the recorder and which subset to send to the telemetry link? One problem with old Airborne Data Acquisition systems is that a small change like adding a new parameter or changing the sampling rate of an existing parameter often meant a big change to the sampling sequence and hence time delays and so on.

THE NEW RULES FOR SOLVING THESE PROBLEMS

(i) Define an **acquisition cycle** time during which all parameters everywhere that are **potentially** of interest are sampled at least once.

(ii) Insist that all parameters everywhere be sampled at the start of the acquisition cycle and at even time-intervals thereafter.

These rules are deceptively tricky to understand and implement but are equally deceptively powerful once implemented completely. They describe an isochronous sampling system (Iso = same, chrono = time).

This goes beyond the mere synchronicity of a PCM stream or the type of "simultaneous sampling" boasted of by certain command-response busses. Before looking at the design elements of such a system let's first look at some situations where these rules may provide clarity.

COMMERCIAL NETWORKS - BEYOND MIL-STD-1553

Many ground stations today use networks to share telemetry data among multiple ground stations. As these networks get faster and the chip-sets associated with them get smaller it seems the next step may be to think of the Data Acquisition Systems (DAUs) from which the data was originally gathered as network nodes.

There are a few commercial busses under consideration by the avionics community: FDDI, Firewire, ATM, 1Gb Ethernet, and FC-AE to name but a few. It is also worth mentioning that considerable effort is being spent on faster, "enhanced", usually optic-fiber versions of MIL-STD-1553.

Deciding between the various options is not trivial. There are financial, mechanical and packet delivery time trade-offs. Once a network is chosen the learning curve is just beginning; for example:

- Fiber-channel (FC) is a commercial network standard
  - They want a very fast SCSI bus.
- FC-AE is an Avionics Environment group within FC
  - They want a very fast avionics bus (MIL-STD-1553+)
- NexGenBus is yet another group with an Airborne Data Acquisition focus
  - They want a very fast CAIS type bus

This paper does not advocate one bus over another. All the busses discussed above can operate comfortably in an environment of isochronous DAUs. In particular, it may be that in environments where more than one network is used then the DAUs *must* be isochronous.
This paper argues that, whatever network or flavor of network is chosen, if each DAU node is isochronous then at least the data collection or sampling is completely deterministic - even if the transfer of that data is not. Now the problem of determinism is purely on the receiver (ground station) side.

Figure 1 shows multiple DAUs operating isochronously (the mechanics of this are discussed later). Each is gathering data packets during each acquisition cycle. The network however transfers these packets in a way that first might appear as anathema to an Airborne Data Acquisition engineer - the DAUs can be read out of sequence and at varying intervals of time. All is not lost! Remember all the data is sampled isochronously within, for example, ±100ns. The large (super-set) packets going to the recorder can be sorted in time later. Also it may be that the even though the order in which the smaller (sub-set) packets are transmitted to the ground may change, the worst-case delay may be within some acceptable window, for example 300ms. In this case the design task on the ground becomes one of building a 300ms buffer on the ground - if a packet is received within 10ms; delay it by 290ms. For this to work each packet must have a time tag - axiomatic for anisochronous DAU network.

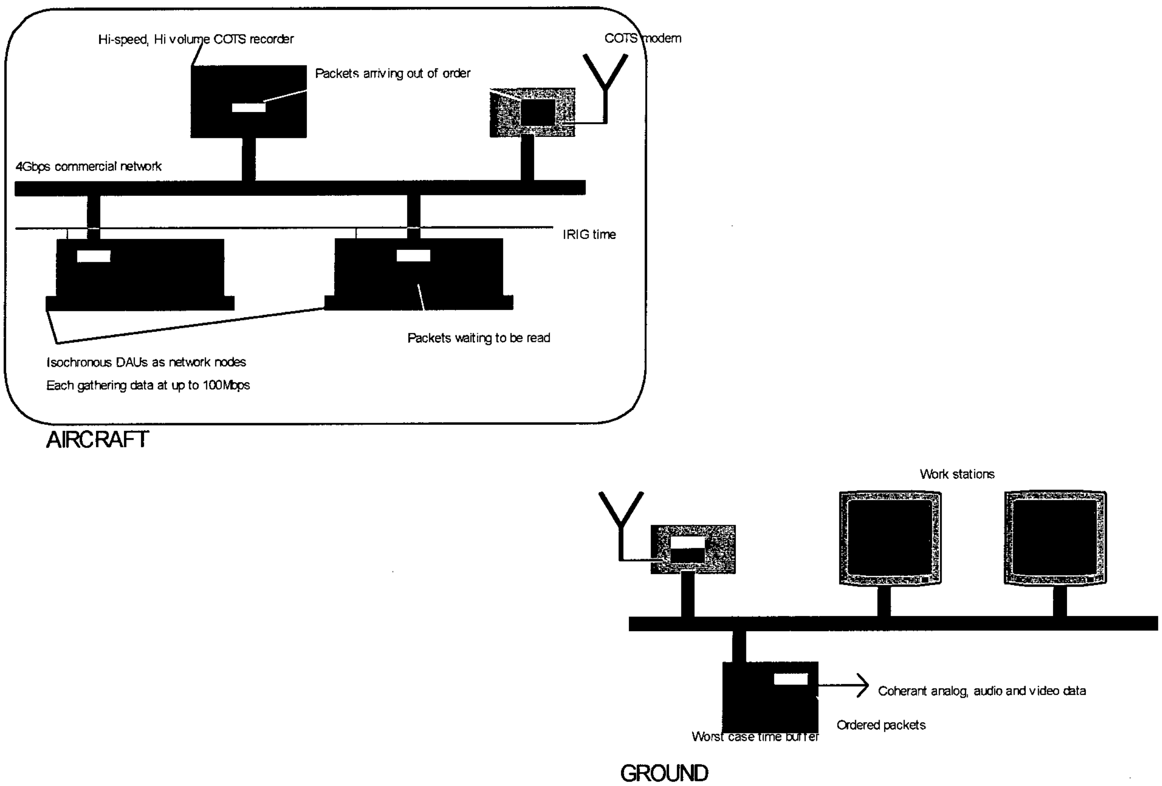

Figure 1 - Commercial networks – Beyond MIL-STD-1553

PACKET DEFINITION

It may not be obvious that isochronous systems have advantages with respect to packet definition. In particular once the two core axioms of isochronous operation are understood then significantly less communication is required when discussing topics such as:

- Sub-packets
  For example, a parameter sampled at 20kHz to an on-board recorder is also sent at a much slower rate to the ground and at an even slower rate to a cockpit meter. Which samples are sent where? Another request often made is - only send "interesting data" - if data is reduced, the structure of the new packet must be defined (or must it?).
- Time tagging
  Wouldn't it be great if one time tag tagged all parameters?
- Packet structure
  What information is needed in a packet?

In *Figure 2* there are two packets of data acquired during one acquisition cycle. One packet is going to a solid-state recorder and another, smaller one, is going somewhere else. The first element in each file is the packet identifier, the second is the packet's time tag. Each row after that contains samples of a particular signal.

In an isochronous acquisition system there is a lot of information in these packets.
The **packet identifier** points to a header packet that need only be sent occasionally. This packet contains information on the acquisition cycle time and signal names, ranges, units, delays and so on for each row.

The **time tag** (Tc0) is the precise time that all the samples in the first column of the packet were sampled. If the acquisition cycle length is Ta and there are only two samples in a row then the second sample was taken at exactly Tc0 + Ta/2 and so on.

Some observations:

(i) It may be worth considering having the packet identifier and time-tag incorporated into any file name associated with the packet, as it would make sorting easier.

(ii) One time tag tags everything - while this greatly reduces the tag information that must be transmitted it also means that even if the sampling rate changes or extra signals are added the engineers analyzing the data need not care.

(iii) Defining data reduction subsets becomes axiomatic. Each set must contain the first sample and all samples must be evenly spaced in time. For example see the third row of each packet.

(iv) In the first packet the parameter in row 5 is sampled at 50Hz and row 6 at 60Hz. Time correlation of these signals is straightforward. Remember the first sample of each row was taken at the same time.

(v) In the first packet the parameter in row 7 is also sampled at 60Hz, that means every sample in row 6 was taken at precisely the same time, as those in row 7.

P1234
2002 03 27 23 59 5999 9999

0001	0002	0003	0004	0005	0006	0007	0008	0009	000A
1001	1002								
2001	2002	2003	2004	2005	2006				
3001	3002	3003	3004	3005					
4001	4002	4003	4000	4005					

First (super-set) packet

P1234
2002 03 27 23 59 5999 9999

0001	0006			
1001	1002			
2001	2004			
3001	3002	3003	3004	3005
4001				

Second (sub-set) packet

Figure 2

BUS CONTROLLERS

Airborne Data Acquisition equipment is often the glue between high-speed busses used to transport all Airborne Data Acquisition data and slower sub-system busses such as MIL-STD-1553, ARINC-429, CAIS, smart sensor arrays or legacy 10-wire interfaces.

These sub-systems are typically command-response type architectures not designed for isochronous operation. However they can be adapted to co-exist in such an environment with immediate advantages with respect to time-tagging and coherency.

Figure 3 shows multiple bus controllers gathering data about their respective busses. For completeness data from an analog channel is being sampled along with data from an external analog multiplexer or scanner.

The analog signal is sampled at the start of the acquisition cycle and at equal intervals of time thereafter. The first sample is stored in address X of a current value table (CVT) the second in address Y and so on. So far so good - this would be expected from an isochronous channel. However in real life systems there must be an anti-aliasing filter and all filters have delays. So even though the A/D sampled the signal at the start of the acquisition cycle there is a fixed delay that must be noted.

The external multiplexer at first seems to violate the rule of sampling all parameters at the same time - this cannot be done with a multiplexer. However think of each channel of the multiplexer as being sampled after a fixed delay. Design the multiplexer controller to step through a sequence of channels at the start of the acquisition cycle and at equal intervals of time thereafter. The first sample from channel 1 is stored in CVT address A, the second sample from channel 1 in B and so on..

This concept (delays + sequences + equal sample intervals) is then extended to the case of the bus controllers. At the start of the acquisition cycle the MIL-STD-1553 controller requests data from a given remote terminal and sub-address and each data word is stored in a specific address in a CVT.

All this data from a myriad of analog channels and busses is stored coherently in a CVT, a snap-shot of which at the end of the acquisition cycle forms a super-set of all the data packets for that DAU. Even if the sampling sequence changes radically on one module, the other modules do not change, providing the acquisition cycle has remained the same.

With multiple data banks these packets can be stored for as long as it takes the network(s) to read the packets.

Furthermore, all these acquisition modules need not be in the same DAU providing all the DAUs are operating isochronously. It is very important to note that only the bus controllers were affected by the move to isochronicity - none of the remote terminals had to be redesigned. One hidden advantage of isochronous systems is that if they work once, they work always because everything that happens, happens always.

The next section discusses some of the design implications in designing an isochronous distributed data acquisition system.

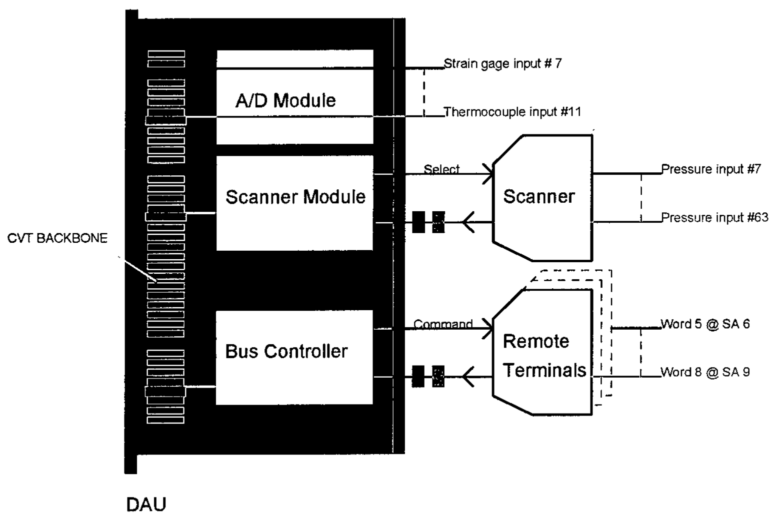

Figure 3 - Different modules in an isochronous environment

MAKING DAUS ISOCHRONOUS

At first glance it may appear that a simple broadcast at the start of an acquisition cycle to all DAUs is all that is required to make multiple DAUs operate isochronously. However even with 3ppm oscillators this would require a broadcast every 30ms to guarantee jitter of less than ±100ns. This may be an excessive overhead on many networks. By the way, 100ns jitter is not an unreasonable tolerance for oversampling digital filtering systems.

So an indication of the start of an acquisition cycle, and a regular metronome beat would be ideal. Also it may be desirable to have different sampling strategies or modes or formats. For example CAIS has program, verify and acquisition modes. Also, it may not be possible to always sample everything that might possibly be of interest in every stage of the flight - for this reason IRIG-106 Ch.4 supports format switching.

One solution to all three problems has been around the Airborne Data Acquisition world for decades. IRIG time can be connected to all DAUs with control function bits indicating the end/start of an acquisition cycle and the format to use during the next acquisition cycle.

The downside of this solution is, at worst, an additional single twisted pair looped to each DAU. The advantage however is that all DAUs are now gathering data coherently even across multiple airframes.

Remember that it is only bus controllers that are affected - not remote terminals. Also it is only network sources that are affected - not the sinks. In other words, the DAUs are affected - not the recorders. For example, a COTS SSR that supports an acceptable flavor of FC-AE (for example) need not be modified to gather data from multiple DAUs. However, each DAU must have enough buffer space to handle any lack of determinism in the network.

The final section looks at the various decisions that must be made in choosing an acquisition cycle length.

SOME THOUGHTS ON THE ACQUISITION CYCLE

Imagine an Airborne Data Acquisition system with multiple PCM streams, some CAIS equipment and a MIL-STD-1553 controller talking to a cockpit display. *(Figure 4)*

PCM streams are often defined using power-of-two rules such as 512 words per minor-frame, 128 minor-frames per major-frame and 8 major-frames per second.

Fifty cycles per second are often found in MIL-STD-1553 busses. With these two criteria alone, 500ms may be the optimum choice for acquisition cycle (500ms = 4 major-frames and 25 MIL-STD-1553 cycles). By the way, some smart sensor systems talk of EPOCHs - another word for acquisition cycle - that may also have to be factored into the choice of acquisition cycle length.

If 500ms is chosen then IRIG-G time (10ms/cycle) would be a better choice than IRIG-B (1s/cycle).

If the data packets are too big then the delay in gathering the data might be unacceptable (for example audio to ground). Also with a large file losing a small proportion of the file means losing a lot. Finally larger acquisition cycles mean that the recovery time from a power brown-out is longer.

On the other hand if the packets are too small then the protocol overheads (e.g. packet headers or file names) may be excessive, also sampling rates may be pushed too high.

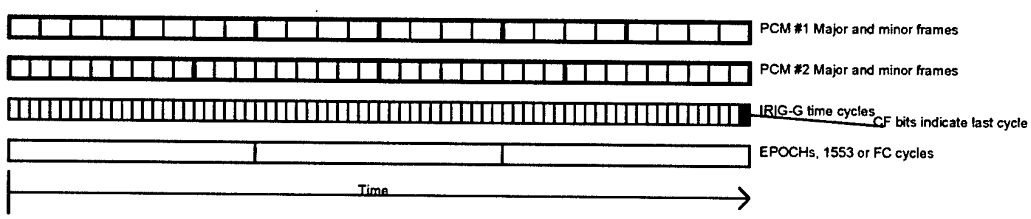

Figure 4 - An acquisition cycle and various sub-cycles

CONCLUSION

Forcing various elements of an Airborne Data Acquisition system to operate isochronously with respect to IRIG time requires some investment from each vendor, some training of program groups and may require a twisted pair to the controller of whatever network or bus is used to gather data.

However, when adhered to, these simple rules provide many advantages with respect to interoperability, network independence, future-proofing and packetization.

References

Fibre Channel: Replacement for MIL-STD-1553 & Next Generation Military Data Bus, Wilf Sullivan, Product Manager-Systems & Software, DY 4 Systems Inc.

A Sensor for Monitoring Corrosive Environments on Military Aircraft

P. N. Trathen and B. R. W. Hinton
Defence Science and Technology Organisation
506 Lorimer Street, Fishermans Bend, Vic, 3207

Abstract

A major part of the cost of maintaining aging military aircraft is repairing damage due to corrosion. Knowing when this corrosion occurs may help to optimise maintenance intervals and procedures. DSTO has designed a corrosion monitor, which has been fitted to RAAF P-3C maritime patrol aircraft. The monitor has been is use since 1998, and it has flown on 15 different aircraft. The monitor has a galvanic sensor that becomes active in corrosive environments, and records this data in internal memory. This paper discusses the development of the monitoring system, and the results obtained. The paper also analyses the relationship between corrosion activity, aircraft flight and the moisture and humidity levels at the airfield.

1. Introduction

Corrosion in military aircraft is responsible for a large proportion of maintenance time and expenditure. Most of the corrosion damage is revealed during periodic time-based inspections. A knowledge of when the corrosion occurs could be used to optimise maintenance intervals and procedures. A study [1] by the United States Air Force (USAF) in 1991, revealed that the cost of corrosion was $US718 million per year. Corrosion prevention measures accounted for 14% of the expenditure, while repair related costs made up 82%. The United States Navy (USN) spends more that $US1.5 billion annually on aircraft systems on repair and maintenance due to corrosion [2]. Of the aircraft in the USN fleet, 40% are over 20 years old.

With future condition based monitoring and management in mind, DSTO has been developing a monitor which will allow the severity of the corrosive environment within known corrosion prone areas in aircraft structure to be determined. The monitor will also indicate when the corrosive events occur, ie. what stage during or after the flight, and the locations and mission types that give rise to corrosive events. This paper describes the results of an initial program to determine this information from a monitor developed and built at DSTO. The system is currently operating in Royal Australian Air Force (RAAF) P-3C and F-111 aircraft, and approvals are being sought to install the unit in a Royal Australian Navy (RAN) Seahawk helicopter. The data discussed in this paper were obtained from a monitor located in P-3C Orion aircraft based at RAAF Edinburgh, South Australia.

2. Development of the Monitor

The DSTO monitor consists of a galvanic corrosion sensor, and supporting electronics. The sensor is made using printed circuit board technology with alternate tracks of tin and copper. These tracks are insulated from each other on the side of the circuit board that is exposed to the corrosive environment, and are connected beneath the circuit board to a current measuring device in the case. The galvanic current produced is proportional to the corrosivity of the environment. At pre-set intervals, the electronics power up and the galvanic corrosion current is measured and stored in computer memory. The monitor then enters a standby mode to conserve power. The monitor is battery powered and self-contained, with a battery life and data capacity of approximately 30 days, when readings are taken every 30 minutes. The monitor is based on a Motorola 6303 CMOS microprocessor and Military Specification grade components. The case (Figure 1) is constructed from aluminium alloy with dimensions of 130 x 130 x 85 mm and total weight is approximately 2 kilograms. The monitor measures corrosion current in units of micro-amps (μA). It has a linear range from zero to 25 μA. The resolution of the data is 0.098 μA. (The resolution is limited by an 8-bit processor). The data from the monitor are plotted as μA vs. time, and correlated with the flight activity of the aircraft. Data were also available from a data logger located on the flight line (the area where the aircraft are parked between flights), which in addition to galvanic corrosion sensors, contained temperature and humidity sensors. Data from this unit was used together with information from the Bureau of Meteorology automatic weather station located at RAAF Edinburgh.

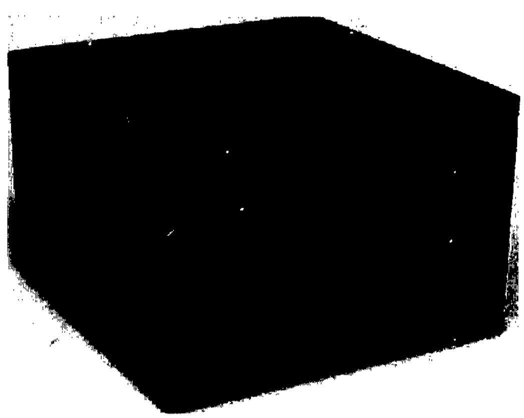

Figure 1. DSTO Corrosion Monitor

Before installation, the unit was assessed for the possible emission of electromagnetic radiation. Approval was also required from the RAAF Maritime Patrol Systems Program Office (MPSPO) in order to install the unit.

3. Installation of the Monitor

The P-3C Orion (Figure 2) is a 4 engine turbo-prop aircraft based on the Lockheed Electra airliner. It is used for long-range maritime patrol and surveillance, particularly anti-submarine and anti-shipping. The first of the current model aircraft were introduced into RAAF service in 1978. The recently released Defence White Paper [3] has speculated that they may be kept in service until 2040, giving a 'life of type' of 62 years. There are currently 19 aircraft in RAAF service, with an additional 3 recently procured (without surveillance sensors) for pilot training. A typical mission would see the aircraft operating at low levels over the sea for extended periods of time. It is quite common for the aircraft to return to base with visible salt build-up over large areas of the external surfaces. Corrosion has been an on-going problem for the P-3C since its introduction.

Figure 2. RAAF P-3C Orion

The monitor was located on one of the frames in the tail of the aircraft (Figure 3), behind the pressure bulkhead. The arrow shows the location. This is near the rudder torque-tube, which has been susceptible to corrosion. Access to the monitor was gained through the maintenance access panel used for servicing hydraulics. A Special Servicing Instruction was raised to manage the installation and removal of the monitor; a task performed by 92 Wing personnel at RAAF Edinburgh. Immediately after the monitor was installed in the aircraft, it was switched on and recorded data at 30 minute intervals. After each 30-day exposure, the unit was returned to DSTO to download the data and recharge the batteries. After a turn around period of 2-3 days, the unit was returned for the next exposure. Information on flights was sourced from the 92 Wing Maintenance Control Section.

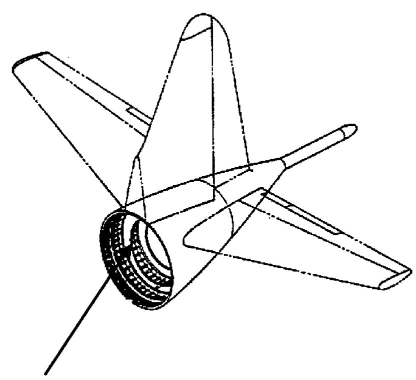

Figure 3. Location of monitor

4. Results

At the time of preparation of this paper, the monitor had been exposed in P-3C aircraft for a total of over 22,000 elapsed hours and had flown well over 230 flights. The monitor has been fitted to 15 different aircraft. During this time, the aircraft were in flight for 7.3% of the time and on the ground for 92.7% of time. Flights have been made to locations all over the world, however most of the flights have been within Australia. The majority of Australian flights were in the RAAF Edinburgh region. A total of 209 flights were analysed in detail for this paper.

Close analysis of the data obtained from the monitor revealed that they could be divided into 3 categories.

- Ground activity - output from the monitor that was clearly unrelated to any flight.
- Flight activity - output occurring during flight.
- Post flight activity - output that commenced within 30 minutes after the reported landing time.

Typical output data are shown as Figure 4. In this graph, the monitor output is shown as the solid line, the flight-line humidity is the dashed line, and the black bars represent the time the aircraft was airborne. The data show an example of activity over a period of 7 days, with the vertical gridlines at midnight. The time readings are elapsed hours from the start of the exposure period. Humidity readings over 100% are due to condensation settling on, and wetting the electronic sensor. There are three flights shown in this graph. The first two are transit flights, where the aircraft was flown at relatively high altitude for much of the flight. The third flight is a test flight, which was normally conducted at low altitude and consisted of a series of circuits in the vicinity of the airfield. From this graph, it can be seen that significant outputs occurred immediately after landing, also that these outputs tended to be higher if the flightline humidity was high. It can also be seen that not all flights produced an output.

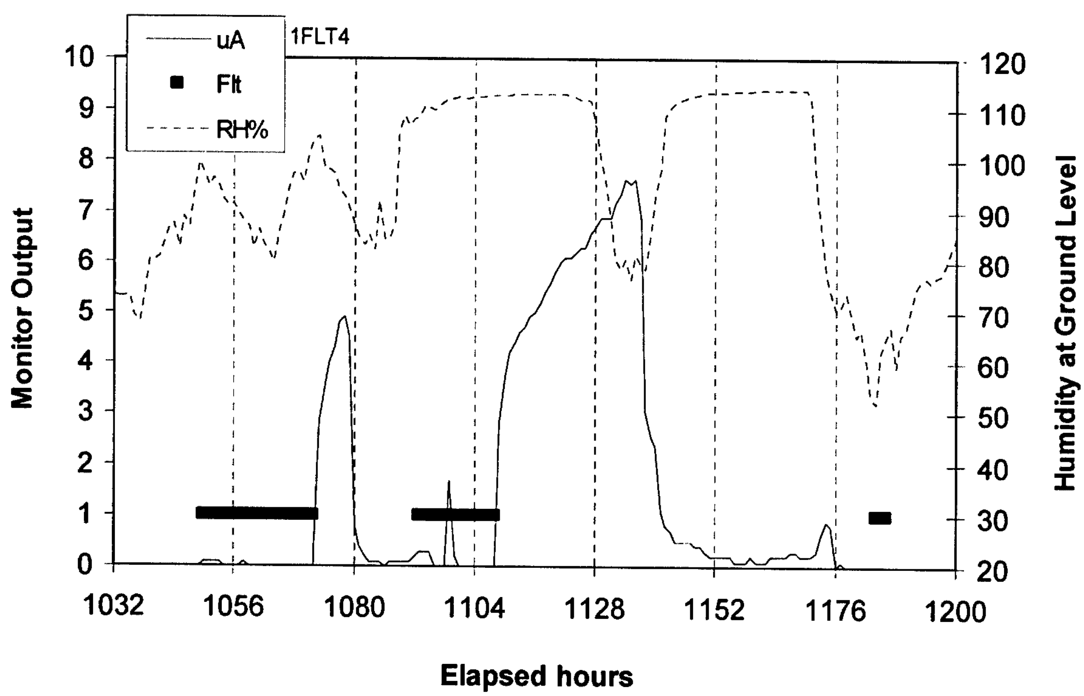

Figure 4. Typical output from the corrosion monitor

In order to analyse the output data in greater detail, they were characterised in the following way.

- Events - a period of continuous monitor output above zero.
- Duration - the length of time the event lasted, measured in hours
- Charge - a measurement of the total amount of charge passed between the sensor elements during a corrosion event. Through Faradays Law, this charge is proportional to the total or cumulative corrosion that occurs on the tin anode. The unit of charge used in this paper is millicoulomb (mC).

4.1 Ground Activity

It was clear from the data, that the occurrence of monitor output while the aircraft was on the ground was greater when the moisture levels were high. For the purposes of this report, three humidity zones have been defined. The zones are (1) 0 to 50%, (2) 50 to 80% and (3) 80 to 100%. The thresholds of 50% and 80 % were chosen as 50% is the approximate level where corrosion activity is believed to commence [4], and 80% is the level used to define a surface as being wet in the ISO standard [5].
Table 1 shows the percentage of corrosion events in each humidity zone from a 3 year period. Most of the corrosion events occurred in the high humidity zone, with only 9% occurring in the low zone. Also shown are the average activity duration and charge passed, which were greatest in the high humidity zone.

Table 1. Ground activity of the monitor in various humidity zones

Humidity Zone	Percentage of all corrosion events	Average event duration - hours	Average charge per event - mC
0 to 50	9	1.21	1.47
50 to 80	34	2.60	3.44
80 to 100	57	9.09	9.19

4.2 Flight Activity

For each flight, details of the take off and landing times, as well as the departure and arrival airfields were supplied by the RAAF. Figure 5 shows the percentage of flights having corrosion events, plotted against the flight duration in hours. There is no clear correlation of activity with flight time. No flights under 1 hour duration recorded any corrosion activity.

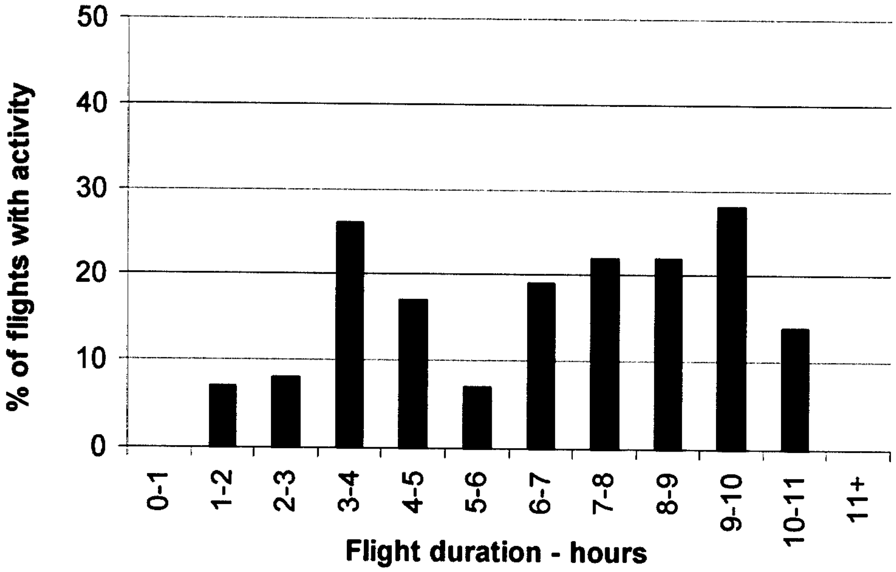

Figure 5. Corrosion activity for different flight durations

4.3 Post-Flight Activity

A significant proportion of flights produced activity from the monitor immediately after the aircraft landed. Edinburgh was chosen for more detailed analysis of this post-flight activity, as most flights landed there and detailed data was available for the ground conditions.

The data in Table 2 indicate that the post-flight monitor activity was related to the humidity at the destination airfield flight line. The post-flight activity was greatest at the higher humidity range (80 - 100), with approximately 45% of flights causing current to be recorded. At humidity less than 50%, only 10% of flights caused activity on the monitor. Both the duration of the corrosion activity, and the total cumulative charge, were highest when the humidity was in the high range.

Table 2. Post-flight activity for flights landing at RAAF Edinburgh, at various humidity levels

Relative Humidity %	% of flights with post-flight activity	Average event Duration - hours	Average charge per event - mC
0 to 50	10	2.83	6.64
50 to 80	19.7	4.5	6.39
80 to 100	45.2	9.82	41.49

The RAAF Edinburgh data were analysed to determine the effect of flight time. Figure 6 shows the percentage of flights having post-flight activity for a range of flight times, in the three humidity ranges at the airfield. No particular relationship was evident between flight time and post-flight activity, but once again, high ground humidity increased the chance of post-flight activity. Importantly, flights less than 1 hour recorded no post-flight activity.

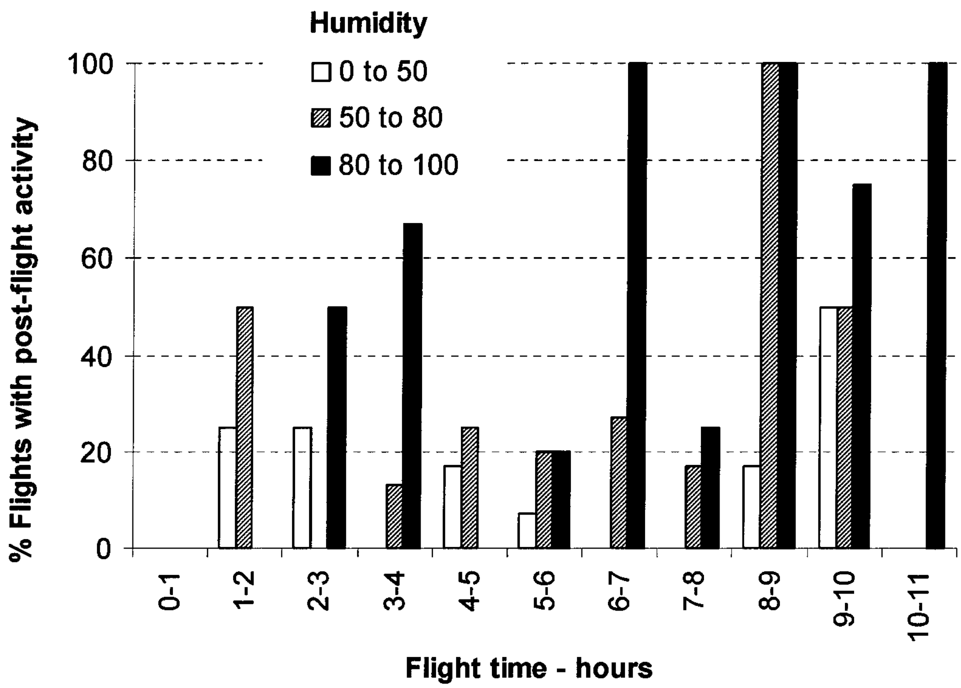

Figure 6. Proportion of flights with post-flight activity, at various humidity levels

5. Discussion

Previous work [6] using electrochemical noise and impedance sensors mounted in the wing bays of a Nimrod maritime patrol aircraft, found that changes of altitude and temperature lead to the formation of moisture on sensor surfaces and subsequent sensor activity. It was concluded that the activity was a function of operation of the aircraft and the prevailing environmental (atmospheric) conditions. The sensor activity was greater and more variable in the more exposed sections of the aircraft.

The USN [2] has used a galvanic sensor made from a gold-cadmium couple to investigate the corrosivity of the environments within aircraft in storage. It was found that even in the desert, the sensor was active, and that condensation from temperature excursions below the dewpoint were the likely cause of corrosion.

The reason for the higher post-flight activity at higher ambient humidity can be explained by the relationship between temperature, humidity and dew point. Air contains a certain amount of moisture vapour. At any given temperature, the amount of moisture vapour actually in the air, relative to the maximum amount it could hold (i.e. percentage saturation), is defined as the relative humidity. As the air temperature increases, its capacity to hold moisture increases. Similarly, if the air temperature decreases, the amount of moisture it could hold decreases. If the amount of moisture (measured in grams of moisture per kilogram of dry air) stays constant, then as the temperature of air decreases, its saturation level (humidity) will increase to the point where it is fully saturated. The temperature at which this occurs is the dew point. Cooling of surfaces below the dew point temperature will cause the water vapour to condense out of the air, forming a layer of moisture on the surface.

Figure 7 shows the relationship between relative humidity and dew point at an ambient temperature of 20°C, as calculated from psychrometric tables [7]. At 100% humidity, any temperature drop is sufficient to cool surfaces below the dew point. At lower humidity, a greater drop of surface temperature is required to cause condensation. Water will condense within an aircraft structure soon after landing if the temperature of the structure is below the dew point.

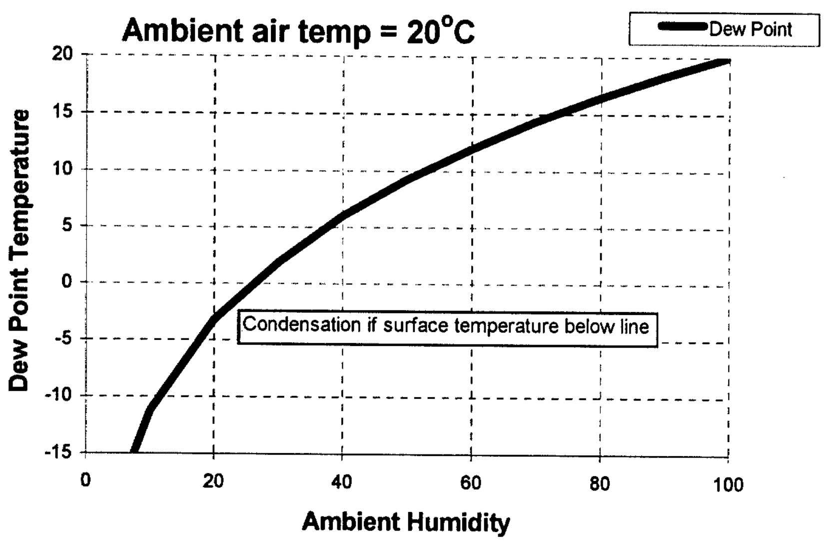

Figure 7. Relationship between humidity and dewpoint

The temperature that an aircraft structure will reach during flight depends on three main factors, (1) the ambient temperature, (2) the altitude, and (3) the aircraft speed. Guidelines exist for the estimation of this temperature [8]. Air temperature decreases with altitude, with the drop being approximately 18°C per 10,000 feet altitude. Aircraft therefore, cool significantly at cruising altitude. Against this must be contrasted the heating effects due to aerodynamic friction. The P-3C, being a propeller driven aircraft, is relatively slow and the frictional heating is limited to a few degrees °C. It is estimated that a P-3C would cool to a temperature of approximately -25°C during a high altitude transit flight.

The temperature of the airframe immediately after landing depends on the ambient temperature, and the speed and altitude profile of the flight just before landing. With the P-3C, this temperature could range from slightly above ambient, to perhaps 20 to 30°C below the ambient air temperature. As Figure 7 shows, at higher humidity levels (on the ground) less temperature drop is required to go below the dew point. This is reflected in the data shown at Table 2 and Figure 6 that show greater activity at high humidity. At low ambient humidity, the airframe temperature may have to drop 20 to 30C below ambient to produce condensation, and hence, monitor activity is less.

Most flights of a P-3C include a range of activities to maximise the training benefit. For example, a flight to search for illegal fishing vessels may also include time devoted to 'anti-submarine warfare' training. There are likely to be several changes in altitude and speed. Monitor activity during flight is likely to be due to changes in altitude causing changes in temperature of the aircraft structure, and therefore condensation of water on the monitor.

Knowing when corrosion is likely to occur would allow maintenance schedules to be tailored to suit aircraft operations. For example, aircraft operating from bases in the tropics during the wet season would be expected to experience much more condensation than aircraft flying in temperate areas. As the maintenance schedules are generally based on flying hours rather than location, knowledge of the corrosivity of actual operating conditions could enable the limited maintenance resources to be used most effectively.

6. Conclusions

- Corrosion activity was recorded by the monitor when; on the ground, during flight, and immediately after flight.
- Corrosion activity during flight was probably caused by condensation as a result of temperature change during altitude change.
- There was no clear correlation between the length of flight and monitor output, except that no activity was recorded for any flight less than 1 hour duration.
- Corrosion activity post-flight was greater for higher ambient humidities.
- Condensation within the airframe after flight was more likely at high ambient humidity, as the airframe temperature was more likely to be below the dew point temperature.
- Measuring the corrosivity of the actual operating environment could optimise maintenance intervals and procedures.

7. References

1. United States Air Force, Logistics Command, Corrosion Program Office Report -#F09603-89-C-3016 (1991)
2. Sensors for Monitoring Corrosion in Aircraft Systems, V. S. Argawala, 13th ICC
3. Defence Review 2000 - Our Future Defence Force, Department of Defence, June 2000
4. C. S. Laurent, "Corrosion Protection through the use of Desiccant Wheel Dehumidification" NACE Conference, Corrosion 89, Paper 319, New Orleans, April 1989
5. ISO 9223:1992 (E) Corrosion of metals and alloys - Corrosivity of atmospheres - Classification
6. Health and Usage Monitoring Systems - Corrosion Surveillance, D. Weetman, D. Eden, W. Cox, CAPCIS MARCH Ltd. 1993
7. PsyCalc, Psychrometric Calculator, Linric Company
8. Defence Standard 00-970/Issue 1, Design and Airworthiness Requirements for Service Aircraft, Volume 1 - Aeroplanes, Leaflet 101/3, Operation in Various Climatic Conditions - Temperature Limits for Design Purpose

Managing a Successful HUMS Operation

Andy J. Evans
Propulsion Department, Safety Regulation Group, UK Civil Aviation Authority

ABSTRACT©

Helicopters are potentially more vulnerable to catastrophic mechanical failures than fixed wing aircraft because of the number of single load path critical parts within the rotor and transmission systems and the reduced redundancy within their design. Since its introduction into the North Sea offshore helicopter fleet in 1991, the Vibration Health Monitoring (VHM) capability of Health and Usage Monitoring Systems (HUMS) has been shown to be an effective and practical means of reducing mechanical failures and most importantly those that prevent continued safe flight and landing. One major independent study has concluded that: 'HUMS was probably the most significant isolated safety improvement of the last decade'.

HUMS is not a conventional maintenance tool. VHM is primarily used as an early-warning system. While many defects exhibit 'classic' vibration characteristics, new defect indications are still common. So, although downloads are automatically screened, expert interpretation and/or physical investigation is usually necessary to detect incipient failures that would otherwise remain undetected.

Setting definite vibration rejection criteria is difficult, so most HUMS data needs careful interpretation. It is therefore be sensitive to normal human and organisational factors. To gain the full of HUMS benefits it is essential that an operator defines sound operating procedures, effectively a 'HUMS Management System'. This can fundamentally affect their maintenance philosophy and the maintenance programmes of each affected helicopter type.

This paper will identify HUMS Management System best practice.

INTRODUCTION

Helicopters are potentially more vulnerable to catastrophic mechanical failures than fixed wing aircraft because of the number of single load path critical parts within the rotor and transmission systems. Since its introduction into the North Sea offshore helicopter fleet in 1991, the VHM capability of HUMS has been shown to be a practical means of reducing the rate of hazardous and catastrophic failures that prevent continued safe flight and landing.

Before HUMS entered service an early study of past accidents by one operator (Gordon 1989) suggested that VHM would have been able to prevent 50% of airworthiness related accidents (i.e. those with a technical cause) both within their own twin-engined fleet worldwide and amongst other North Sea operators. This would have meant avoiding 17 accidents within 1.8 million flying hours, 7 of which were fatal accidents (including the world's worst civil helicopter accident in which 45 people died [AAIB 1988]). The same study also suggested that 60% of serious incidents could have been prevented by VHM. Another retrospective accident study conducted by the UK Ministry of Defence (MOD) concluded that HUMS would have cut their airworthiness accidents by 39% (Jeram-Croft 1994).

The predicted potential of VHM has been demonstrated by the initial retro-fitted HUMS that have since been termed 'first generation HUMS'. They were able to provide warnings for 69% of the failure types that occurred in the first 6 years of operation (McColl 1997), and were able to record vibration changes in a further 17% of cases (offering the potential for further improvements). This study showed that first generation HUMS warned successfully in 60% of the potentially catastrophic failure cases, with 80% a reasonable target with system improvements.

A study by a working group of the Helicopter Health Monitoring Advisory Group (HHMAG) showed that incidents of serious vibration occurring in-flight had reduced dramatically in the UK fleet after the introduction of HUMS (Evans 2002a).

HUMS regularly identifies mechanical problems. During 2000 and 2001 HUMS was the first indicator of a problem requiring significant maintenance action approximately once every 6000 flying hours (Evans 2002b). VHM has thus proven capable of filling most of the gap in existing established health monitoring techniques (e.g. MCDs, SOAP and oil pressure and temperature monitoring). Whilst these established techniques are effective in detecting wear debris and oil loss, VHM provides a means to detect those failures that result in misalignment or imbalance (such as fatigue failures). VHM is thus an important complement to traditional techniques.

A study by one helicopter constructor dispelled initial fears of false removals as it determined that even early in the history of HUMS, the main gear box false removals rate was as low as 1 per 147,000 flying hours (Kershner et al [1997]), with 5 of the 6 HUMS initiated MGB removals being due to confirmed anomalies.

The same study also showed that the only premature removals of deteriorating gearboxes (i.e. ones still judged serviceable at the subsequent strip) occurred prior to the end of 1994. Kershner et al (1997) attribute this to a HUMS diagnostic learning curve.

Some critics have suggested that spurious warnings would make HUMS impractical to use. While HUMS had a 'difficult' three-year introductory phase (HHMAG 1997), the proven safety benefit soon outweighed the perceived difficulties (Dobson 1997). These safety benefits are supplemented by economic, reliability and availability benefits (Dobson 1997 and Pouradier 2001).

Even during one of the early CAA sponsored operational trials of HUMS prototypes (CAA 1993b), less than 5% of acquisitions were corrupted, less than 2% were 'cautionary warnings' (all subsequently identified as wiring, connector or redatum errors) and there were no false drive train 'failure warnings'.

One major independent study (Hokstad et al 1999) concluded that: 'HUMS was probably the most significant isolated safety improvement of the last decade'. The same study stated that HUMS is 'expected to mature over the next decade and will probably contribute to a further risk improvement'. The CAA believes that these benefits will come from design-assessed helicopters (i.e. those meeting the latest certification requirements), which integrate HUMS into the helicopter's maintenance philosophy.

Worldwide HUMS experience now exceeds 2 million hours. HUMS is frequently required by oil companies as a contractual requirement. This is in response to their duty of care to their employees, recognising the contribution that helicopters make to the risk their staff are exposed to (Clark 1998).

In the UK, HUMS has been necessary on large helicopters certified since the design assessment concept was introduced by the CAA in 1985 (e.g. CAA 1992) and has been retrospectively made mandatory on older types by an Additional Airworthiness Directive (CAA 1999b), effectively giving two first generation HUMS airworthiness credit based on 9 years of successful service experience. Similar retrospective requirements are being proposed in Norway (CRHSNCS 2002).

HUMS is not a conventional maintenance tool. VHM should not be thought of as a measurement device, able to be calibrated to tight tolerances. Current VHM is primarily used to provide early-warning of impending failures. While many defects do exhibit 'classic' vibration characteristics, new defect indications are still common. So although HUMS data downloads can be automatically screened, warnings usually require expert interpretation and/or investigation to determine if the helicopter remains airworthy or to identify an appropriate maintenance action.

To gain the benefits of HUMS it is essential that the operator defines sound HUMS operating procedures, effectively establishing a 'HUMS Management System'. It is important to recognise that this will fundamentally affect both the maintenance programmes of each affected helicopter type and operator's maintenance philosophy.

The primary purpose of this paper is to identify 'best practice' for HUMS management. The structure of the paper is as follows:

1. How the term 'HUMS' is used in the context of this paper.
2. The framework of civil aviation maintenance in which HUMS is used.
3. The aspects that make HUMS unique.
4. HUMS Management System best practice.
5. How designers can make HUMS more effective in-service.
6. The role of the regulator with HUMS operation.

DEFINING HUMS

One pioneer in the field has reflected that the term HUMS 'never acquired a strong, tight meaning' (Stewart 2000). It is therefore worth identifying how the term will be used in this paper.

Joint JAA/FAA advisory material defines health monitoring as a means 'by which selected incipient failure or degradation can be determined' (FAA 1999]). The same source defined usage monitoring as a means 'by which selected aspects of service history can be determined'. In both cases these means include the 'equipment, techniques and/or procedures'. Hence the 'system' of HUMS is wider than just the installed hardware.

HUMS will be used in this paper to refer to any system capable of rotorcraft drive train health monitoring that also collects some form of usage data. This paper will concentrate on the management of health monitoring and in particular the warning of incipient failures. This is an aspect of particular interest to the CAA because of the proven safety benefit that HUMS delivers to areas of rotorcraft that have historically been responsible for the majority of serious airworthiness accidents (Vinall 1980, HARP 1984, Witham 1991 and Neubert 1997). It is for this reason that when the CAA made health monitoring mandatory on the majority of the large UK transport fleet (CAA 1999b) the expression used was the more specific term 'Health Monitoring Systems'.

The paper will specifically focus on the primary form of drive train health monitoring used to prevent failures: vibration health monitoring. VHM does not include the rotor track and balance function used to identify adjustments for vibration reduction. Some of the observations and issues identified are however also relevant to other forms of monitoring systems.

AN INTRODUCTION TO CIVIL AVIATION MAINTENANCE

In civil aviation, the operator of an aircraft used for commercial air transport must be approved. In the European JAA system, the operator is accountable for managing the maintenance of their aircraft and for having a system that ensures airworthiness. For European aircraft, maintenance can only be carried out by an organisation with a JAR-145 approval. Maintenance can be defined as 'any one or combination of overhaul, repair, inspection, replacement, modification or defect rectification of an aircraft/aircraft component' (JAA 2001a). After completion of any package of maintenance a 'certificate of release to service' is necessary before flight (JAA 2001a). In this a paper it will generally be assumed that the operator also conducts maintenance in-house, the simplest scenario. If line maintenance is sub-contracted, then the situation becomes slightly more complex due to the extra interfaces, however the same basic principles apply.

The basis for maintenance is a framework typically established by the 'instructions for continued airworthiness' issued by the appropriate design organisation (including for example the Maintenance Manual), Airworthiness Directives issued by the appropriate National Airworthiness Authorities and other regulations of these NAAs. In addition, the internal procedures of the approved maintenance organisation should also impose controls on the maintenance engineer so as to assure a consistently airworthy outcome. Wackers and Kørte (2001) observe that aircraft maintenance is thus a 'highly regulated space'.

WHY IS HUMS DIFFERENT?

The health monitoring aspects of HUMS are essentially means of maintenance inspection ('the examination of an aircraft/aircraft component to establish conformity with an approved standard' [JAA 2001a]). These inspections lead to other forms of maintenance when conformity (i.e. airworthiness) cannot be established.

For most conventional inspections, the criteria used to establish airworthiness are normally what could be described as 'hard criteria'. Hard criteria are those that can either be measured reliably with calibrated equipment (e.g. the pitch of a gear tooth) or are discrete conditions (e.g. 'cracked' or 'not cracked'). Exceeding hard criteria results in failure of the inspection.

One major advantage with HUMS is that it is not simply a post-flight inspection, but as Wakers & Kørte (2001) put it, HUMS 'constitutes an attempt to have the maintenance engineer's senses present' aboard the helicopter in-flight. However, as they note, 'vibrations do not respect component boundaries'. VHM data also varies widely dependent on flight conditions and torque, individual combinations of components and even between individual builds of the same components (Salzer 1994 and AAIB/N 2001). Most HUMS only acquire data in specific flight regimes ('recognised' by analysis of data being fed to the Flight Data Recorder). However there can still be large variations if the data is acquired at extremes within the acquisition envelope. It can also be difficult to correlate vibration levels recorded during test rig running with in-flight data because test rig mounting is so much stiffer (McFadden 1985). Trend analyses, and VHM thresholds based on the vibration history of that build (so called learned thresholds), are sensitive to maintenance actions that can change the datum vibration.

Hence it is difficult to establish clear and universal 'hard' vibration limits for in-flight acquired VHM data that can give guaranteed warning of failure without unnecessarily rejecting airworthy components.

For a new HUMS installation this means that the initial VHM thresholds will be subject to revision as data is gathered and analysed. The accuracy of these thresholds will improve with time and so modified and developed systems should achieve continual improvements. Also, those systems associated with an initial design assessment at certification may achieve improved initial thresholds. JAA/FAA advisory material on HUMS describes the concept of a Controlled Service Introduction (FAA 1999), which is a vehicle for validating thresholds. During a CSI the certifying aviation authority will closely monitor the system's performance. They must confirm that the HUMS has reached an acceptable level of maturity before closing the CSI. The CAA has supervised two such HUMS CSI programmes (both on systems required as part of a helicopter's certification basis). The only prospects to perhaps accelerate this process is the use of advanced data processing techniques, employing unsupervised or supervised learning, to model normality and detect deviations from the norm (e.g. CAA 1999c) during the CSI.

For HUMS retrofitted without the helicopter constructors assistance, in the few cases where any 'hard' vibration limits are available (e.g. based on established Maintenance Manual limits for external shafts [CAA 1993c]), they are usually only applicable during ground running and require an exercise in HUMS cross-calibration to allow for any differences from the traditional carry-out Ground Support Equipment (GSE). In such cases it is common to use HUMS to highlight if a vibration is getting close to the defined limit, prompting the use of the carry-out GSE (as called for in the Maintenance Manual [MM]) for a definitive inspection.

For retrofitted systems sound engineering judgement and fleet statistical data ~~usually~~ forms the basis for warning threshold levels (CAA 1993c and Nevins 2002). Even VHM integrated by the helicopter constructor rarely have genuine hard removal limits.

From the time of early VHM trials it was recognised that if a warning was found during 'first line' data review by a line engineer, a 'second line' analysis by a more experienced HUMS specialist would be required (CAA 1993a).

Routine manual review of all data is not practical as vast quantities can be collected. Typically 50-100 components will be monitored with perhaps 18 parameters generated for each (Kørte and Aven 2000). One operator generates over 100Mb each day from its fleet at Aberdeen alone (Dobson 2000a). A parameter is also likely to react to several types of defect (Pouradier 2001). Therefore in order to direct investigation effort to priority cases, many VHM criteria are what could be termed 'soft criteria'. Such vibration thresholds are designed to allow HUMS to automatically screen data, highlighting only vibrations that are deemed unusual and thus deserve closer attention. The setting of such thresholds is difficult because parameters can detect numerous failures modes, which will each have different vibration characteristics (e.g. AAIB 1998). VHM warnings require interpretation at a ground station by trained engineers in conjunction with fault-finding actions on the helicopter, a knowledge of the helicopter's maintenance history and access to fleet wide data (Evans 2002a).

VHM threshold exceedances are warnings not 'no-go' indications. A single exceedance is therefore not necessarily confirmation that a particular helicopter has an incipient failure.

To narrow the focus of subsequent inspections, HUMS should ideally identify the component affected, though it often only localises the defect to components that share the same shaft speed as the affected component. The general type of defect is however normally specified (e.g. bearing, gear tooth, shaft etc). If the tail rotor was identified, for example, detailed visual inspections and even a partial strip are possible (AAIB 1998). In the case of a suspected fault within a gearbox, other than the premature removal of the gearbox, the more limited options include checking the torque of external bolts, the application of wear and debris analysis (Barber 2002) or in rare cases boroscoping (Astridge 1984).

In some cases the vibrations may be ambiguous, inconsistent with established theories or dissimilar to historical failure experience. When this happens the primary form of inspection may initially be HUMS 'close monitoring' over subsequent flights. Close monitoring involves manually reviewing the trends of appropriate parameters after each HUMS download to establish if vibration levels are increasing and to localise the suspect component (if possible). Assessing trend graphs 'by eye' is still a quick and effective method (Dobson 2000b).

First generation HUMS suffered from a number of instrumentation faults and some algorithms have proved susceptible to 'spiking' (giving occasional false warnings). If inspections do not produce any physical evidence of a failure, often cautious observation of trends in subsequent downloads is the most sensible option (Dobson 2000b). The ideal trend is one that gradually increases over time with little scatter. Often scatter due to the varying acquisition environment makes trend identification difficult. Again subsequent close monitored flights can help resolve any uncertainty.

Some systems are actually configured to require further flights before all the necessary information is downloaded. For example one system with airborne VHM processing requires that after a warning the engineer selects appropriate parameters whose signal averages are to be added to the download file for the following flight.

All HUMS in-service in the North Sea depend on certain HUMS warnings being evaluated further by the HUMS developer's specialists. HUMS warnings are thus not the same as exceedance of a Maintenance Manual transmission torque limit, which can result in an immediate gearbox removal.

When analysing data, engineers must be aware that in some cases, though a crack may propagate over a considerable time period under low cycle fatigue (on rotor start up and shutdown), the vibration may only significantly increase when the crack reaches a size that it propagates under high cycle fatigue (i.e. due to shaft rotational speed). In some of these cases the potential warning interval may be very short.

Some soft VHM warning thresholds can become hard component rejection criteria over time as fleet experience is gained and maintainers are able to recognise similarities between past defects and current data. Even when a vibration level becomes a hard rejection criterion, a component would not be rejected without checking for the most powerful source of vibration signal possible, a loose accelerometer.

It is easy to present hard criteria in maintenance data. A measurement will be either within or outside the limit and a discrete condition will be determined to be present or not. If a component is found to be outside a hard limit (i.e. no longer airworthy), maintenance data will describe either an on-wing rectification process, or will require replacement, with off-wing rectification if possible.

For the soft criteria that produce HUMS warnings, the maintenance data must introduce a further series of inspections prior to the final determination of airworthiness. Some HUMS retro-fitted to existing helicopters come with diagnostic manuals that use flow-charts to prompt an appropriate sequence of checks in consultation with the helicopter's Flight Manual and Maintenance Manual. Other retrofitted systems leave the operator the task of identifying what tasks within the helicopter's existing maintenance data are appropriate for the investigation. There may be a need for the operators to introduce their own extra inspections. Such a situation puts an unnecessary maintenance development burden on operators. The most straightforward situation is however when the helicopter constructor integrates the fault-finding process into the standard maintenance data for the helicopter.

As so much HUMS data is acquired in-flight, HUMS also differs from normal maintenance tools in that the airborne element (and in particular the network of sensors positioned across the airframe) is routinely exposed to the harsh in-flight environment around the engines, transmission and rotors of a helicopter. This means the HUMS must itself be robust and data analysts must consider the increased possibility of corrupt data (for example caused by a loose accelerometer).

Yet HUMS is also different from a normal item of avionics. Firstly, as Jesse and Slssinger (1994) highlight its intended prime user 'is the line engineer' rather than the pilot and so 'it is the interface with the line engineer which is of greatest importance'. Secondly, for civil aircraft, Minimum Equipment Lists define the maximum rectification periods allowed when important systems become unserviceable. However HUMS functionality relates to being able to successfully download data for analysis. This depends on the serviceability of the airborne equipment, the functioning of the data transfer medium and the ground stations serviceability. Even then, data may not be downloaded successfully if the conditions necessary to acquire data in-flight were not achieved. This has been a common problem when helicopters make short flights, or if missions involve frequent manoeuvring rather than cruise, when airspeed limitations have been imposed because of airframe cracking or during night flying with optical blade trackers. It can be seen that even serviceability becomes a complex issue. Finally if HUMS is limited to providing warnings to maintenance engineers, then if it is unserviceable it can have no affect on the conduct of the flight.

HUMS BEST PRACTICE ISSUES FOR OPERATORS

If a helicopter operator is to be successful in their operation of HUMS, as well as procuring HUMS hardware, providing suitable facilities and provisioning the appropriate technical publications, they need to consider a number of key operational issues:

1) Threshold Changes

Changes to thresholds may be made to ensure new defects are captured, to improve warning times or to reduce false alarm rates. There is a danger that more importance will be given to the latter in the early stages of an operator's use of HUMS, before the HUMS has demonstrated that it can warn of otherwise undetected failures. This temptation must be resisted.

Reducing the false alarm rate simply by rasing thresholds may reduce the resources needed to clear the aircraft for flight but at the expense of reduced warning times and a risk of even disabling the system's ability to warn of some failure modes. While excessive false warnings drain maintenance resources and reduce confidence, they rarely result in unwarranted component rejections (Kershner 1997) and can often be quickly identified and dismissed (Dobson 1997). When changing thresholds the impact on both warning times and false alarm rates must be considered.

A common misconception after an incident is that the thresholds should have been lower as that would automatically improve the safety margin. This would be the case with hard criteria but fails to recognise the soft nature of HUMS warnings. The AAIB (1998) report a case were vibrations rose for 45 hours before exceeding a threshold, with a warning downloaded 5 hours later at the end of the flight, suggesting that the threshold used was 'set too high'. When improved threshold strategies (such as learnt thresholds) were retrospectively applied to this data the maximum practical warning time has been estimated to be 25 hours. Earlier examination of the vibration history may not have shown a sufficient trend to concern the engineer (especially as more helicopters will routinely be closer to the lowered threshold) and it would have been more difficult to correlate the HUMS warning with any physical symptoms on the helicopter. Hence the decision may have been to fly on and 'close monitor' the component, reducing the net improvement in warning from 20 hours to something lower. This case (considered further below) is significant in that despite a clear 50 hour trend and physical symptoms, the helicopter was actually flown on and a failure occurred in-flight. This is not to suggest that enhanced warning times are not important, they clearly are, but safety margin improvements are dependent on warnings being verifiable and/or sufficiently convincing.

The CAA views the setting of vibration thresholds as a design activity that should be controlled by either the helicopter constructor or the HUMS supplier. This does not however preclude the operator either having tailored thresholds issued by the design organisation or HUMS with an option of setting lower cautionary alert levels. Operators, with their intimate knowledge of HUMS performance and fleet trends, however can and must, participate in the complex process of adjusting thresholds (as suggested by Wackers and Kørte [2001] and discussed by Dobson [2000a]).

Changes to thresholds should be controlled in a formal manner (usually by the use of Service Bulletins that cover the issue of field loadable software - as discussed in CAA [2002b]). This should be irrespective of whether the software is being loaded onto an aircraft processor or a ground station. The same controlled process should be used for any other changes, be they changing the regime recognition envelopes for data acquisition or changing the algorithms.

With UK operators the CAA expect to be able to review changes periodically and be able to see a clear justification for each change recorded (CAA 1999b). Such changes should also be audited by the operator's quality system.

2) Download and Analysis Frequency

It is theoretically possible for an aircraft mounted HUMS processor to indicate to maintenance personnel that a warning has been generated, so that data is only downloaded and analysed when a warning occurs. However data is routinely downloaded from current HUMS. This has the advantages of automatically backing up the data and allowing fleet trending.

Assuming that warnings are only available after a download, it is clear that HUMS can only provide adequate warning of failure or degradation if the download periodicity is sufficiently frequent.

UK operators download HUMS data on a daily basis as a minimum (in North Sea operations the highest daily utilisation tends to be 5-7 hours). Even then the operators encourage downloads where possible after every shutdown. Frequent downloads allow the operators the maximum time to seek advice from the helicopter constructor or the HUMS supplier's vibration specialists if necessary. It also means that failures to acquire data (see below) become apparent as early as possible.

Westland (1993) give an example of an early 1980s vibration-sampling programme conducted ever 50 hours by the operator with carry-out GSE. One helicopter being monitored crashed after an input pinion failed. When the accumulated data was reviewed it became clear that the defect was detectable in the recording taken 103.6 hours before the accident. This emphasises that downloading the data is only part of the task.

To gain a safety benefit there must be a timely review of the data. Any warnings must be then be assessed and where appropriate investigated with further inspections. Only when the engineer is confident that the HUMS has not identified a non-conformity should the engineer sign for the completion of the check. In order to do this the engineer must also confirm that sufficient data has been acquired to make this determination (see below).

In some cases operators have included the HUMS download and analysis within their Daily check. Others have instituted a special check of HUMS data by a Technical Instruction or other similar procedure. The prime advantage of the latter is that authorisation can be limited to a smaller group of specially HUMS trained engineers. The operator's procedures must ensure that a helicopter cannot be released for service if the previous HUMS Daily check was not signed for.

3) Data Continuity

There are two significant aspects to data continuity. The first relates to the Maximum Permitted Period Between Successful Downloads (MPPBSD), an issue often linked with MELs. The second concerns access to a helicopter's past vibration history when analysing HUMS data.

There were two main drivers for the development of first generation HUMS. The first was the realisation that the helicopter safety record was not adequate and that health monitoring technology was close to maturing (HARP 1984 and Ingstad et al 1990). The second was a move by ICAO to recommend fitment of Flight Data Recorders (FDR) on helicopters over 2700kg (White 1990). This prompted decisions by the CAA and FAA to require such FDR and Cockpit Voice Recorders (CVR) for certain rotorcraft, encouraging the integration of HUMS and CVFDR functions (Barr and Brown 1991, Land and Daniels 1991). These integrated systems were initially ordered by operators partly to meet contractual requirements established by oil companies and partly as a voluntary initiative to improve safety of their operations. The HUMS functions were however only certified on a 'no hazard, no credit' basis, effectively accepting that they did not hazard the helicopter, but without validating their functionality.

While the CVFDR functions required for mandatory airworthiness requirements were recognised in MELs, it is now clear with hindsight that the 'no credit' health monitoring functions, because they were not legally required, were not adequately considered (despite the subject being identified by the CAA's pioneering operational health monitoring trial [CAA 1993a]).

In a case in 1997 a helicopter flew from a remote base for 430 hours (during a busy 10 week period of flying) with an accelerometer unserviceable (AAIB/N 2001). Tragically during this period a component that this accelerometer would have monitored began to crack (according to investigators at least 62 hours before it failed). The subsequent failure resulted in a rapid catastrophic chain of events and the loss of the helicopter and 12 lives. The smaller vibration increases sensed by other accelerometers convinced the accident investigators that HUMS would have provided adequate warning of this failure if the unserviceable accelerometer had been working. They concluded that 'this accident has shown that HUMS is capable of being an important tool in accident prevention' (AAIB/N 2001).

This fatal accident had a major impact on perceptions of the importance of HUMS serviceability and resulted in UK and Norwegian operators having a free exchange of views and re-evaluating their own internal procedures (Dobson 2000b).

The same accident was also one factor that prompted the CAA decision to introduce an AAD (CAA 1999b). In 1999 the AAD prompted UK operators to agree a common proposal for MPPBSDs. Previously operators had their own individual procedures. Initially assigning varying degrees of importance to each accelerometer, based on the criticality of the components monitored, was considered. However it was decided to propose a less complex scheme for retrofitted systems (Dobson 2000b), namely:

25 flying hours:	Health monitoring accelerometers if not being close monitored
10 flying hours:	Health monitoring accelerometers if in close monitor
100 flying hours:	Rotor track and balance (noting that carry-out equipment can be used as an alternative for RTB)

The tighter period for accelerometers being close monitored was adopted because it was considered that 25 hours was too long to be without data if a component had already generated an warning and was being reviewed at each download (Dobson 2000b).

This proposal was, for example, far more stringent than the '5% rule' applied at annual FDR serviceability checks (CAA 1999b). The CAA adopted this proposal, and the close monitoring concept, as the MPPBSD baseline for retrofitted HUMS, although on design assessed helicopters that failure analysis would usually be used as the basis of the MPPBSD (CAA 1992).

This scheme acknowledges that as HUMS is an airborne acquisition system that can only be partially tested on the ground, several flights may be necessary before the diagnostic process can be shown to be complete. For many HUMS aircraft hardware problems, practical avionics fault-finding can often only be conducted overnight, in a hangar's controlled environment (Dobson 2000b). In logistical terms if spares are not on-site, 25 flying hours is still a relatively short period considering that the remote locations helicopters tend to operate from are a logistical challenge to even the most capable express delivery service. As well as having helicopters based on North Sea rigs and remote Scottish islands, British helicopters routinely operate to these MPPBSDs in locations as far apart as Belize, the Falkland Islands, Kazakhstan, China and Brunei.

Prior to the AAD becoming effective, these periods were incorporated direct into the operator's MELs, with individual line items for every accelerometer. Over the first two years of auditing compliance with the AAD (and the associated advisory material [CAA 1999b]) it has become clear that this was an unduly complex implementation that put too much emphasis on the serviceability of aircraft components and insufficient emphasis on actually successfully downloading and then assessing the data. Future CAA approved MELs will simply have a one-line HUMS reference to the operator's 'HUMS Handbook' which must be produced for AAD compliance and approved by the CAA. This document will then define the MPPBSDs.

Once a MPPBSD is reached, maintenance test flights are necessary to gather data. In the UK this has been extremely rare.

A less complex aspect of data continuity is ensuring that when a diagnosis is being conducted, the maintenance engineer has access to sufficient historic data to help draw a conclusion. When helicopters are moved to another base temporarily, there should be provision for moving data with the helicopter (HHMAG 1997). If a helicopter moves frequently between bases, there may be a need to synchronise ground stations at each base. Jesse and Slssinger (1994) give two examples of operations affected by this. The history required need not be that extensive. As a minimum, access to 50 hours of data would be good practice, however as some degradations can be tracked over many hundreds of hours (Barber 2002), more data is preferable. As the value of such a history is reduced by components changes it is not sensible to be too prescriptive. When helicopters are moved between operators (on lease or when sold) a suitable HUMS history for the helicopter should also be transferred.

4) HUMS Serviceability and System Reliability

Operating within the MPPBSD is not sufficient. To give HUMS the opportunity to generate maximum warning times and detailed trends, the operator must strive to both increase the probability that the necessary flight conditions for data acquisition will occur (by flight planning or having the HUMS flight regime windows adjusted) and raise the reliability of the HUMS hardware.

HUMS should be considered within the operator's reliability programme. However such programmes are typically driven by metrics such as in-flight shutdowns, air turn backs, departure delays and high value unscheduled removals. To ensure HUMS is adequately considered, all failures to successfully download data and all HUMS system problems should be specifically recorded (for a good example of the latter see McKim [1997]). The causes should be identified and prioritised for corrective action.

This can help identify if the operator's flight profiles rarely achieve the standard data acquisition conditions. Often it is possible for the HUMS designers to respond by issuing tailored software that can increase the acquisitions success rate. This may be by acquiring data at a slightly lower speed or by shortening the time for the acquisition (reducing the quality of the signal average values, but ensuring that data is acquired more often).

The review of hardware reliability must consider each of the three HUMS elements: the airborne system, data transfer medium and ground station. Problems as simple as a bad batch of data transfer cards (Clark 1997) can dramatically reduce the download success rate. In order to increase the download success rate the operator should actively evaluate preventative maintenance options.

Any reliability review should also consider human factors as a possible cause for failures to successfully download data. Sometimes errors in configuration management mean that data will not be acquired or downloaded. It is not uncommon to find the download medium is disconnected before data transfer is complete or that following an incorrect sequence prevents a successful download. Once these failures have been spotted, corrective actions can be taken.

Although this activity is not directly mentioned in the CAA's guidance for its AAD, it is one of the 'general procedures' that at operator would be expected to have in place (CAA 1999b). Such efforts have in the past resulted in 'a dramatic improvement in reliability and system serviceability' (Clark 1997).

5) Structured Diagnostic Approach

In practice while the operator is accountable for determining if the helicopter is airworthy, the responsibility for this task is discharged by the licensed engineer who signs for the HUMS download analysis and/or the certificate of release to service. The added difficulty of dealing with the soft criteria that characterise HUMS must be balanced by a standard 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. Hopkins (1999) gives a good example of how the effectiveness of a warning system (in that case monitoring carbon monoxide levels in a coal mine) can be subverted by an inadequate approach to interpreting and reacting to the warning.

Some HUMS produce a report (or download log) that can include a simple warning message. Others use mimic representations of the helicopter and changes in colour to highlight components that have triggered a warning. After a warning the first step is usually to examine the downloaded data in more detail. One aim is to determine if the warning was triggered by a clear trend or a step change, another is to examine wether that parameter has been subject to excessive random scatter previously. The engineer will also need access to the helicopter's recent maintenance history, as that can explain some changes. As well as the parameter(s) that triggered the warning, it is usual to examine other associated parameters that may help clarify if a real degradation or mechanical failure is occurring. Some failures can also cause visible reactions on neighbouring accelerometer channels (AAIB/N 2001).

As HUMS warnings are generally soft inspection criteria, if it is believed that the warning could have been caused by something more than just routine scatter, further inspections on the helicopter are usually necessary. Some warnings can be identified as false alarms by careful fault-finding on the airborne HUMS element. Others will need more detailed visual and physical inspection of the suspect drive train components. Magnetic plug inspections are common after HUMS engine or gearbox warnings, though some health monitoring statistics have been skewed by the reason for removal being subsequently recorded as 'debris' (Dobson 1997). In certain cases ground running can gather useful VHM data for comparison. These need to be backed up by examining other sources of health monitoring such as magnetic chip detectors or spectrographic oil analysis programmes. Such wear and debris analysis techniques are complementary to HUMS and may provide the necessary insight (e.g. Barber 2002).

Completing further flights and close monitoring the subsequent downloads may be necessary to gather the information necessary to make a decision. The same approach may be taken initially with those warnings known to be erratic (Dobson 1997).

However it should not be assumed that just exceeding a threshold means that any failure is still some way off. If a clear trend has developed over a number of hours, see for example the case reported by AAIB (1998), there should be no need to gather further data, but the emphasis should be on detailed inspection around the affected component. In that incident a failure occurred in the very next flight, even though other helicopters in the fleet had previously operated close to the threshold without any problems.

Vaughan (1996) has conceived the concept of 'normalisation of deviancy'. This involves tolerating worsening service experience because the deterioration occurs in small increments that are individually deemed to be an acceptable extension of past experience. Tolerating such steadily deteriorating vibration data must be guarded against as warnings are 'particularly sensitive to this process of normalisation' (Hopkins 2001).

With rotors, because they are adjustable, particular care should be taken not to blindly follow RTB instructions and balance out developing rotor defects (see AAIB 1998).

If the maintenance engineer has any doubts over the data it is important to seek specialist support. All HUMS operators should have a HUMS coordinator appointed who can provide advice to line engineers (Dobson 1997). Since they will review data from across the fleet and see the operator's most challenging warnings, they are normally the most experienced HUMS analysts within the organisation. During the early stages of a CSI of a new HUMS there may be a need for daily multi-disciplinary meetings to resolve warnings and system problems (Dobson 1997).

The other sources of advice are the HUMS supplier and helicopter constructor. The operator remains accountable for the final decision on whether the helicopter remains airworthy and unless there is a clear need to remove a component, outside advice will usually be couched in liability deflecting terms. For example a report may state that 'the data does not correspond to any known past defect or any theoretical defect type'. There is of course always the possibility that a new failure mode indication is just about to be discovered. A good diagnostic report will however suggest what causes are unlikely, where future attention should be focused and how a known or theoretical defect may become apparent.

A just culture (Reason 1997) should exist within the company if a component is rejected that subsequently proves to be serviceable and also, as far as practical, contractually between an operator and a power-by-the-hour supplier for false rejections (Dobson 2000a). If a blame culture were to exist, after a few false rejections, there will develop a reluctance to reject components even if there is a clear HUMS indication.

If a false rejection occurs, or if a component does fail in service or if an incipient failure is safely spotted at a later stage (by HUMS or other means), the operator should reactively examine how improvements can be made in the diagnostic process.

A clear definition of the diagnostic process and how a helicopter is cleared for flight are key elements of CAA HUMS requirements (CAA 1999b), with the use of flowcharts a suggested means of defining the process. The CAA also requires that company diagnostic procedures are produced, though if the helicopter's standard maintenance data integrates HUMS in a detailed and comprehensible manner this can be cross referenced rather than duplicated.

6) Clear Responsibilities

The operator will need to establish clear definitions of who is responsible for which parts of the operation and management of HUMS, especially were any sub-contractors are used for specific HUMS support or line maintenance (CAA 1999b).

As already noted, the operator will need to appoint a HUMS coordinator (or equivalent) to manage much of the day-to-day support for HUMS. In small operators this may be one person acting part-time. The larger operators would normally have at least one person dedicated to HUMS and perhaps associated subjects (e.g. RTB, FDRs etc). Responsibilities for managing technical HUMS liaison with the HUMS suppliers, helicopter constructors, overhaul agencies and the NAA will normally rest with the HUMS coordinator. In the UK this includes submitting details of HUMS service experience to the CAA and the helicopter constructor / HUMS supplier (CAA 1992 and 1999b).

It should be clear who has responsibility for improving the serviceability of the HUMS equipment and improving the successful download rate. When maintenance action is carried out that could affect HUMS data, it should be clear who conducts the HUMS redatuming (necessary to avoid either false alarms or the corruption of learnt thresholds). It is good practice to include redatum actions in the stage sheets or work packs for the appropriate maintenance tasks.

Critically, it must be clear where individual responsibility rests for determining if the HUMS data confirms that the helicopter remains airworthy (and how that engineer is supported with additional diagnostic information and advice). Care must be taken to not to dilute responsibility by the impression that observing experts will spot an error (as Snook [2000] has noted). The HUMS coordinator and company management should however monitor the conclusions of significant diagnostic investigations and suspend flying if they ever believe an unacceptable risk remains.

In a case in 1995, an engineer rejected a tail rotor because of a combination of a HUMS warning and abnormal vertical play in the tail rotor shaft (AAIB 1998). The HUMS warning was on tail rotor gearbox output shaft once per rev vibration. In this installation the TGB accelerometer is aligned with the tail rotor lateral axis, making it sensitive to TGB and lateral tail rotor vibrations. It had triggered 5 flying hours earlier during the day's flying, and the vibration had been increasing over 50 hours.

The engineer had rejected the tail rotor because he suspected wear in the pitch change spider was the cause of the vibration and abnormal play. However the actual cause was cracking of the tail rotor flapping hinge retainer. A prior design failure mode analysis had identified this as a potentially catastrophic failure mode. The crack had propagated in low cycle fatigue over about 200 rotor start cycles, but only reached a stage were vibrations increased in-flight over the last 50 hours. The engineer had also followed the standard 3000 hour MM inspection on the tail rotor hub (coincidentally due at that time) without finding any defects, however this inspection did not specifically include the flapping hinge retainer. The accident investigators believe that the crack would have been difficult to spot during a zonal visual inspection because of difficulty getting good access to the area and as the crack only opened up at rotor start up. (AAIB 1998).

If the tail rotor had been removed and returned for an overhaul inspection, as the line engineer intended, this clear HUMS warning would have been a routine success story. The only fall out would have been the need to study ways of lengthening the warning.

However a subsequent shift supervisor decided to attempt to balance the tail rotor. As the tail rotor was successfully balanced within MM limits with carry-out GSE, it was decided to release the helicopter to service after a brief test flight and ground run. During the next revenue flight the hinge retainer finally fractured, causing intense vibration. The crew considered ditching, but for nearly an hour nursed the helicopter back to shore were a safe landing was made, and all 17 occupants disembarked uninjured (AAIB 1998).

Just like the 1997 accident mentioned earlier, this incident was a defining event in HUMS history. It convinced the CAA that there was more to HUMS than just installing the equipment. Furthermore people who doubted the ability of HUMS to detect catastrophic failures became believers.

In this incident the trend of the parameter that generated the warning was probably as clear and distinct a trend as it is possible to have. Hence an important issue is why, with that trend available and abnormal play, was the second decision made (over-ruling the original rejection of the tail rotor) releasing the helicopter for service after a balance. This can be compared to the 1986 decision to launch the *Challenger* Space Shuttle despite progressive deteriorations in-service (Vaughan 1996). Questions that could be asked include:

- Did the Shift Supervisor have sufficient HUMS training?
- Was the maintenance data clear and adequately detailed?
- Were company procedures clear and adequately detailed?
- How was the shift handover conducted?
- Was there a breakdown in communication?
- Was HUMS distrusted?
- Was there management pressure to release the helicopter?

Obviously this is a complex issue. The accident report does not explain why the original decision was over turned.

It is clear however that the response to HUMS warnings is as important as the warnings themselves and that human factors are important in the HUMS diagnostic process.

In order to have effective control over the use of HUMS, it should also be clear who within the operator approves HUMS procedures and policy. Ideally this should be the Quality Manager (or equivalent).

A particular senior manager should be responsibility for determining if their company's use of HUMS is effective and considering strategically how HUMS effectiveness can be optimised. This person should ideally be the Technical / Engineering Director, and should consider HUMS during its whole life cycle, from procurement, through commissioning into service. This may be by specific HUMS reviews or by considering HUMS as a regular item within existing management reviews. The CAA believes management understanding and commitment is vital to the effective use of HUMS (CAA 1992).

7) Training and Competency

Heather (1997) reports that one lesson from the initial introduction of HUMS is the need to 'emphasise training on all levels'. This needs to appropriately cover line maintenance personnel, the company's HUMS specialists and management. One recommendation from the CRHSNCS (2002) study was that 'requirements are made for training in the use of HUMS', and training forms an essential part of the CAA requirements (CAA 1999b).

When compiling HUMS training material there is some value in considering a system conceptually similar to the three levels used to categorise NDT personnel, a specialist activity with which HUMS has some similarities to.

It must be emphasised that there are no regulatory requirements to implement the sort of formal qualification systems used within the NDT industry, nor is this the only or the best option. However irrespective of the system used, staff must be adequately trained for their role.

At Level 1 an engineer could be trained (for example) to conduct routine line maintenance on the HUMS equipment and transfer data from the helicopter to the groundstation.

At Level 2, the engineer could be trained to review the data, assessing trends if necessary and control the diagnostic process. Someone with such training would be authorised to sign for a HUMS download check, so each operating base would need at least one person at the level (it is unlikely that this role could effectively be done remotely).

While the first two levels relate to line maintenance staff, Level 3 would be reserved for the HUMS coordinator and any other senior HUMS specialists. These people will have the greatest experience of HUMS diagnostics, supply detailed diagnostic support, provide technical supervision of the HUMS process, conduct fleet trends monitoring and be able to liaise effectively with outside organisations on HUMS matters.

One advantage of conceptualising training in this way is that it shows that (except in perhaps the smallest of operators) a single training standard is inappropriate. The best results come with such a pyramid approach, with more in-depth training and education for smaller number of Level 2 and 3 engineers. This approach also avoids diluting the experience base by allocating HUMS diagnostic tasks to a core of people who can become particularly experienced.

Because HUMS is a computer based avionic system that monitors mechanical drive train components, Level 2 and 3 personnel need to be comfortable with and have a good understanding of airframe, engine and avionics issues.

Finally as HUMS can become such an important part of an operator's maintenance system, as Heather (1997) suggests senior managers and other staff (e.g. quality) who may influence the HUMS activity also need appropriate familiarisation. It is important that all key personnel understand the capabilities and limitations of HUMS, the operator's system for using the technology and how HUMS responsibilities are divided.

Manufacturer's HUMS courses can provide in-depth education in VHM theory, the use of the equipment and its maintenance. Though one study commented that these courses tended to be 'too deep and too expensive for large numbers of line personnel' (HHMAG 1997). Most operators will usually run Level 1 and some Level 2 training in-house. All operators will need specific internal training on how they implement HUMS within their organisation.

As HUMS detectable failures will occur relatively frequent within a large fleet (Evans 2002b) and are useful for diagnostic training, such case studies are well suited for inclusion in a continuation-training programme (HHMAG 1997). The CAA's HHMAG meets twice each year in the UK and provides one vehicle for HUMS users to share such case studies. Continuation training should also cover other lessons learnt within the HUMS programme and changes in software. HUMS is also a useful subject for maintenance human factors training (see CAA 2002a). Both the high profile failures to successfully exploit HUMS (AAIB 1998 and AAIB/N 2001) provide many human factors lessons, as well as fitting the category of organisational accidents conceived by Reason (1997).

For a new operator of HUMS it is especially important to have a defined plan for how the necessary minimum level of HUMS knowledge and experience is to be gained before HUMS use commences. Established HUMS users will need a similar plan when introducing a new make of HUMS. One option may be to establish a support agreement with another suitably experienced organisation to provide a level of diagnostic support and hands-on instruction. However unless the arrangement extended to providing full line maintenance support (i.e. the maintenance organisation released the helicopter to service), any diagnostic support would only be advisory. Either way defining and understanding responsibilities is still important.

The CAA considers training to be fundamental to the successful operation of HUMS and consequently audit HUMS training plans and records (CAA 1999b).

8) Developing Insight Versus Developing The Technology

Wackers and Kørte (2001) warn of the phenomena of becoming fascinated by new technology. They suggest that it is possible for engineers who are not involved in day-to-day helicopter operations (such as the operator's HUMS Level 3 specialists and engineering management) to over-focus on the technology of HUMS and its long-term potential, at the expense of the short-term application and usability of HUMS.

While HUMS requires long-term commitment to continually optimise and improve the system, it has a great potential for reducing the accident rate now. Long-term development efforts must be balanced with short-term attention to two prime issues: the quality of downloads and the insightfulness of diagnoses.

The capability of a well-managed HUMS to prevent expensive accidents should be motivation enough to actively manage those two issues. It is clear that HUMS is one of the last barriers between incipient mechanical failures and an accident. It is vital to work to prevent cases were either no warning is given (e.g. AAIB/N 2001), or no suitable action is taken (AAIB 1998), when a detectable failure occurs. It is also likely that HUMS data and HUMS operation will be critically scrutinised by investigators after any airworthiness incident or accident.

9) Monitoring Performance and Making Improvements

Reason (1997) contends that 'it is often latent conditions created by maintenance lapses that either set the accident sequence in motion or thwart its recovery'.

Wackers and Kørte (2001) when considering one past accident (AAIB/N 2001), warn of 'drift towards a more vulnerable state'. The concept of drift (conceived by Snook [2000]) primarily addresses problems of poor coordination and organisational incompatibility. Reason (1997) has a complementary concept that addresses coordinated organisation-wide changes that increase risk. In his two-dimensional 'production-protection' space, a complacent organisation converts seemingly surplus defences into increased production, eventually moving unintentionally beyond the production-protection equilibrium towards catastrophe. This is perhaps a better model to at least partly explain the HUMS unserviceability aspect of the 1997 accident (AAIB/N 2001). In this case despite the initial objective of HUMS installation being to increase safety, over time it became accepted that accelerometer unserviceability was non-safety critical (Wackers and Kørte 2001).

Wackers and Kørte (2001) suggest there are few means to monitor for such drift or deterioration. However a solution may be the development of an effective safety culture (Reason 1997). Such a culture is made up of shared values and beliefs that help define 'how we do things around here'. It has been argued that 'safety performance will be greatly strengthened by the existence of a positive safety culture' (CAA 2001). Reason suggests that in order to have an effective safety culture, it is necessary to have reporting, just, flexible and learning sub-cultures.

The first two of these mean effective reporting of problems will occur in the just environment, allowing the organisation to learn (the fourth sub-culture) and act to reduce risk. In their study of high reliability organisations Weick et al (1999) note that 'maintenance departments... become central locations for organisational learning'. This is because 'they are front line observers in a position to give early warning of ways in which things could go wrong' (Hopkins 2001). This reporting and learning offers the prospect of continually improving common procedures in a systematic way that can avoid the inadvertent introduction of latent conditions elsewhere in the organisation, ensuring appropriate balance between production and protection.

Auditing is one means of proactively seeking out problems and initiating corrective action (consistent with reporting and learning sub-cultures). Hokstad et al (1999) however express concern over the focus of many audit programmes, noting 'too much emphasis is put on the processes' and warn of a 'lack of focus on the final quality of the product (i.e. well maintained, safe helicopters)'. When auditing something with the safety potential of HUMS it is vital that both aspects are considered.

One explanation for such limited auditing may be the origin of quality auditing in manufacturing industry. In mass-production, process consistency is an important goal. In more complex and safety critical arena, like aviation, where activities are more varied and less repeatable, it is unrealistic to expect that procedures will not evolve continually. This evolution aligns well with the fourth element of Reason's safety culture, the flexible sub-culture, which relates to how the organisation adapts to change.

Maintenance auditing must therefore not just seek out procedural non-compliances and inadequately defined and out-of-date procedures (although there will no doubt be examples to find). Audits must seek to identify any deterioration in the final product (i.e. deteriorations in airworthiness). The effectiveness of audits can be further improved by using a concept Reason (1997) proposed for regulatory audits, namely assessing the findings made to determine the underlying organisational conditions responsible for them. This idea adds a second dimension to audits, greatly increasing the insight that can be gained.

For the purposes of prioritisation, any safety related audit planning should also consider the hazards associated with the operation including rare but catastrophic hazards (Hopkins 1999). An understanding of the currently perceived hazards also allows an auditor to usefully seek out unrecognised hazards (Hopkins 2001). This approach may have helped prevent in the 1997 helicopter accident (AAIB/N 2001) and other accidents were defences have been eliminated over time.

Auditing must itself be of high quality. Appleton (1994) for example notes of auditing on the North Sea platform *Piper Alpha* prior to is destruction in 1988: 'there appeared to be sufficient effort put into safety auditing but it is evident that it was not of the right quality'.

In an attempt to improve maintenance auditing, increased emphasis has now been placed on product audits and auditor competence within JAA regulations (JAA 2001a). A further improvement will be introduced by NPA145-12 (JAA 2001b), which emphasises human factors and safety culture issues. The CAA has a general requirement for HUMS auditing (CAA 1999b). Both the central HUMS team and HUMS use within individual line stations should be audited.

Snook's concept of drift and Reason's production-protection space relate to changes in organisational practices that can detrimentally affect airworthiness. The previously discussed concept of the normalisation of deviance (Vaughan 1996) however can be applied directly to the definition of what is airworthy. Vaughan's concept explains how each worsening piece of service experience can be collectively accepted as a new norm, until the safety margin is completely eroded. This perhaps provides the other half of the explanation to the 1997 accident (AAIB/N 2001), because as longer periods of unserviceability became the norm, the priority of rectifying defective accelerometers decreased.

With HUMS the most dangerous form of normalisation of deviance is when increasingly more extreme vibrations are, over time, accepted as normal. This is difficult for non-HUMS specialists to detect.

In one initiative to supplement conventional maintenance auditing, Edwards (2002) proposes that 'supervisors or standards engineers should be routinely monitoring maintenance practices'. This proposal differs from conventional auditing because it recognises that some monitoring requires specialist knowledge of the activity being undertaken. The CAA also considers this in their requirements for NDT supervision (CAA 1999a), which places the responsibility for technical NDT supervision (including auditing and compliance verification) on a 'Nominated NDT Level 3'.

Similar measures will help detect if normalisation of deviance is occurring in HUMS diagnoses. A company's senior HUMS specialist can play a role in reviewing the diagnoses made at different line stations. Additionally some HUMS suppliers offer a programme of specialist HUMS audits as part of their standard support contracts while others conduct routine remote database reviews.

Employing a mix of safety reporting (including the previously mentioned reactive analysis of HUMS 'failures' and serviceability), conventional internal auditing and specialist compliance monitoring will help detect and prevent both human errors in the diagnostic process and the normalisation of deviance. These activities will all feed into management HUMS effectiveness reviews previously mentioned.

10) Bringing It All Together - A HUMS Management System

The previous nine issues all cover important aspects of HUMS operation. The final aspect is integrating their management.

Existing CAA requirements refer to having 'an effective health monitoring programme' (CAA 1992) and acceptable procedures (CAA 1999b).

The harmonised JAA/FAA guidance for future HUMS projects (FAA 1999) refers to the operator having a 'HUMS programme' (which should include were applicable a CSI implementation plan), where HUMS 'adds to, replaces or intervenes in industry accepted maintenance practices' (FAA 1999).

The means of meeting all these requirements and integrating the previous nine aspects together in an effective manner can perhaps be best envisaged as a 'HUMS Management System'.

In reality this is not a stand-alone system but a logical extension to the quality system required by the JAA for both maintenance organisations and operators (JAA 2001a and 2002) and the accident prevention and flight safety programme demanded of operators (JAA 2002).

The latter was introduced following it becoming an ICAO recommended practice and may itself be combined with the quality system. CAA (2001) describes how such a programme can be more effectively implemented as a Safety Management System.

SMS is both proactive and reactive, giving a means to anticipate and prevent or reduce the effect of risks.

There are essentially three basic prerequisites for a safety management system:

1. *A comprehensive approach to safety*
2. *An effective organisation for delivering safety*
3. *Systems to achieve safety oversight*

(CAA 2001)

There are two distinct functions within the SMS: the ability to assess risk and ability to properly address these risks in a timely manner (CAA 2001).

It can be seen that the SMS concept fits well with the discussion of previous HUMS issues.

Safety management systems, if not properly developed and implemented, can degenerate into nothing more than complex systems of paper' Hopkins (2001). The introduction of a SMS is however likely to be eased if there is a positive safety culture (as previously mentioned) within the organisation. However such a safety culture is likely to be more difficult to acquire than by simply establishing a management system (even before considering the hazard of subsequent drift [Snook 2000]). As Reason (1997) notes, a truly effective 'safety culture is something that is striven for but rarely attained'.

Critics may suggest that implementing the type of management system proposed for HUMS is unduly burdensome, however it is consistent with sound business practices when introducing any major change or managing a major investment. As a failure to fully exploit HUMS to prevent an accident is effectively what has been termed as an 'organisational accident', it must always be remembered that 'it only takes one organisational accident to put an end to all worries about the bottom line' (Reason 1997).

HOW THE DESIGNERS CAN ASSIST THE OPERATORS

HUMS is steadily spreading around the world as it becomes a de facto standard amongst customers. Many more operators will be faced with many of the challenges that North Sea operators rose to 10 years ago (Clark 1997). Designers (either at a HUMS supplier or the helicopter constructor) can have a major influence over the in-service effectiveness of HUMS.

In addition to the usual functional requirements (such as reliability, maintainability, usability etc) and long-term software development, product improvement and logistics, issues that they need to better address in future include:

1. Warning Quality: HUMS designers need to strive to further improve the quality of warnings (i.e. reducing false warnings, providing more convincing warnings of actual failures and increasing warning intervals). This may include the use of advanced data processing techniques to highlight abnormal data.
2. Threshold Development: An effective and ongoing process for timely optimisation of fleetwide thresholds (with the participation of the operators), with the ability to make local adaptations if a warning quality improvement can be justified.
3. Maintenance Data: HUMS for new helicopters should be accompanied by clear maintenance data (and training material) that fully integrates HUMS, including a structured diagnostic approach, into the helicopter's maintenance philosophy. The maintenance data for future retrofit systems should be designed for easy interface with the helicopter's maintenance data. Maintenance data must more explicitly embody HUMS service experience that today exists mainly in case study reports and failure databases.
4. Diagnostic Support: This needs to be timely and useful. In the 1995 incident the constructor's recommendation that the tail rotor should be replaced arrived while the helicopter was airborne on the flight it failed (AAIB 1998). One European organisation has, for example, located a diagnostic specialist on the Pacific Rim to better HUMS support local customers.

Helicopter constructors can make additional contributions to the effectiveness of HUMS:

5. Retrofit Involvement: A past study (HHMAG 1997) found a lack of clear commitment from constructors to third party HUMS. If constructors are to provide effective support to customers (and potential customers) of their existing types they need to be able to respond to customer demands for HUMS (Clark 1997). While many constructors now take a more positive attitude, the others will be at a commercial disadvantage if they do not put mechanisms in place to assist third party HUMS suppliers retrofit HUMS to their helicopters and help the in-service development of such systems.
6. Overhaul Instructions: In order to improve the diagnostic capability of HUMS, detailed and timely feedback from overhaul organisations is important to improve warning of failures and eliminate false removals (Clark 1997). Constructors should revise component overhaul instructions to ensure that adequate reporting requirements exist to give this feedback.
7. HUMS Friendly Design: Considering HUMS and the need to do HUMS prompted inspections when designing new helicopters will help make the use of HUMS simpler and cheaper. This may include specifically designing gearbox casings to give better accelerometer positioning, use of a suite of complementary health monitoring techniques, and even the widespread incorporation of borescope ports (Astridge 1984). Consideration of HUMS from the preliminary design phase would also reduce installation costs.
8. Using HUMS to Increase Design Insight: Operators have seen how HUMS can give them added insight into the condition of their helicopters. Dobson (1997) reports cases of careful balancing during ground runs with MM specified carry-out GSE that have been shown to be completely ineffective at cruise conditions by HUMS data. The constructors also need to make a mental shift to expect the insight that HUMS can give them, rather than treating in-flight acquired VHM data as 'alien data' that they cannot analyse.
9. One Fleet Concept: In order for HUMS to enter service with at least a minimum level of effectiveness HUMS needs to fly on a new type's flight test programmes from the very first flight. The best way for a new helicopter / HUMS combination to mature quickly (remember that first generation systems each achieved 200,000 flying hours in their first 5

years) is to make HUMS a standard fit as suggested by Clark (1997).

10. Maintenance Credits: In the early 1990's the HUMS community extensively debated the concept of using HUMS to relax the requirements for other maintenance actions (James 1994 and HHMAG 1997). Ten years on there is now harmonised JAA/FAA guidance on the certification of such credits, yet despite encouragement, few serious applications have made. If one of the new generation of helicopters currently in flight test are certified with an integrated HUMS as standard to compliment existing health monitoring (e.g. MCDs, SOAP etc), perhaps we shall see dramatic economic benefits that rival the proven safety benefit of HUMS.

THE ROLE OF THE REGULATORS

In their report into the 1997 accident the accident investigators characterised the national aviation authority involved as 'too passive' regarding HUMS (AAIB/N 2001). This is perhaps an example of what has been referred to as 'the regulators unhappy lot' (Reason 1997), in this case being criticised because of operational failings with a relatively new and technologically advanced system fitted voluntarily.

However today the safety benefit of HUMS is proven and even if HUMS is fitted voluntarily (e.g. as an oil company contractual requirement) the local regulator should still consider its operation as part of their evaluation of the operator's quality system and flight safety programme

When Hopkins (2001) observed that 'management systems can degenerate into nothing more than complex systems of paper', he also noted that 'they can be kicked into life by vigorous action by regulators'.

The CAA has recognised that HUMS as a technically feasible system that can have a major safety benefit when used to compensate for catastrophic and hazardous failure modes. The CAA has thus dedicated considerable effort building up the necessary specialist knowledge to regulate HUMS operations.

CONCLUSIONS

While HUMS make a significant contribution to the safe operation of rotorcraft, it can be vulnerable to human and organisation factors that reduce its effectiveness. Thus purchasing the equipment is only the first step, with many operational aspects needing to be successfully managed.

Implementing a HUMS Management System can considerably increase the probability that HUMS will be operated effectively and will thus be able to prevent accidents and incidents.

The effectiveness of HUMS is also dependent on the support given by the HUMS supplier and the helicopter constructor. With the right support and commitment from all parties there are further major advances can still be made in future HUMS safety and economic performance.

ABBREVIATIONS

AAD	Additional Airworthiness Directive
AAIB	Air Accidents Investigation Branch [UK]
AAIB/N	Air Accident Investigation Board - Norway
AC	Advisory Circular
ARB	Airworthiness Requirements Board [an industry group that advises the UK CAA]
CAA	Civil Aviation Authority [UK]
CAP	Civil Aviation Publication
COTS	Commercial Off The Shelf
CRHSNCS	Committee for the Review of Helicopter Safety on the Norwegian Continental Shelf [an ad hoc advisory committee of the Norwegian Ministry of Transport and Communications]
CSI	Controlled Service Introduction
CVR	Cockpit Voice Recorder
CVFDR	Cockpit Voice / Flight Data Recorder(s)
FAA	Federal Aviation Administration [US]
FDR	Fight Data Recorder
FLS	Field Loadable Software
GSE	Ground Support Equipment
HARP	Helicopter Airworthiness Requirements Panel [an ad hoc working group of the ARB]
HCF	High Cycle Fatigue
HHMAG	Helicopter Health Monitoring Advisory Group [a CAA coordinated industry group formed in 1985]
HMS	Health Monitoring System(s)
HUMS	Health and Usage Monitoring System(s)
HUMSMS	HUMS Management System
ICAO	International Civil Aviation Organisation [a United Nations body]
JAA	Joint Aviation Authorities [Europe]
LCF	Low Cycle Fatigue
MCD	Magnetic Chip Detector
MEL	Minimum Equipment List
MM	Maintenance Manual
MOD	Ministry of Defence
MPPBSD	Maximum Permitted Period Between Successful Downloads
NAA	National Aviation Authority [generic term]
NDT	Non Destructive Test
NPA	Notice of Proposed Amendment
PBH	Power-by-the-Hour
RTB	Rotor Track and Balance
SMS	Safety Management System
SOAP	Spectrographic Oil Analysis Programme
VHM	Vibration Health Monitoring

REFERENCES

AAIB "Report on the Accident to Boeing Vertol 234LR G-BWFC 2.5 miles East of Sumburgh, Shetland Isles on 6 November 1986", Aircraft Accident Report 2/98, HMSO, London, UK **1988.**

AAIB "Report on the Incident to Aerospatiale AS332L Super Puma G-PUMH over North Sea on 27 September 1995", Aircraft Incident Report 2/98, HMSO, London, UK **1998.** http://www.aaib.dft.gov.uk/formal/gpumh/gpumh.htm

AAIB/N "Report on Air Accident September 1997 in the Norwegian Sea Approx. 100 NM West North West of Brønnøysund, Involving Eurocopter AS332L1 Super Puma. LN-OPG, Operated by Helikopter Service AS", Report 47/2001, Air Accident Investigation Board - Norway, Oslo, Norway November **2001.** http://www.aaib-n.org/rapporter%202001/rapp%2047-01/index.htm

Appleton B. "Piper Alpha", in Kletz T. "Learning from Accidents", 2nd Edition, Butterworth-Heinmann, London, UK **1994.**

Astridge D.G. "Advanced Gearbox Health Monitoring Techniques", *The International Journal of Aviation Safety*, UK pp165-169 September **1984.**

Barber D. "British International S61 HUMS Case Study, Main Gearbox S/N 66-721 Ring Gear", Proceedings of the 35th Meeting of the Helicopter Health Monitoring Advisory Group, Gatwick, UK June **2002.**

Barr D.W. and Brown A.C. "Helicopter Combined Flight Data Recording and Health and Usage Monitoring Systems", Proceedings of the ISASI 1991 Forum, International Society of Air Safety Investigators, USA **1991.**

CAA "Special Conditions AS332L2/05 and /06", Civil Aviation Authority, Gatwick, UK November **1992.**

CAA "Helicopter Health Monitoring: Operational Trials Review", CAA Paper 93002, Civil Aviation Authority, London, UK, February **1993a.**

CAA "Helicopter Health Monitoring: Bristow Helicopters Ltd - Operational Trial of Helicopter Health Monitoring Techniques", CAA Paper 93003, Civil Aviation Authority, London, UK February **1993b**

CAA "Helicopter Health Monitoring: British International Ltd - Operational Trial of Sikorsky S61N HUMS Demonstrator", CAA Paper 93004, Civil Aviation Authority, London, UK, February **1993c.**

CAA "Airworthiness Notice 94: Personnel Certification for the Non-Destructive Testing of Aircraft, Engines, Components and Materials", Issue 7, CAP455, Civil Aviation Authority, Gatwick, UK, March **1999a.** http://www.caa.co.uk/publications/publicationdetails.asp?id=240

CAA "Acceptable Means of Compliance, Helicopter Health Monitoring CAA AAD 001-05-99", CAP693, Issue 1, Civil Aviation Authority, London, UK, May **1999b.** http://www.caa.co.uk/publications/publicationdetails.asp?id=241

CAA "Intelligent Management of Helicopter HUMS Data", CAA Paper 99006, Civil Aviation Authority, London, UK, September **1999c.**

CAA "Safety Management Systems for Commercial Air Transport Operations", CAP712, Civil Aviation Authority, Gatwick, UK, May **2001.** http://www.caa.co.uk/publications/publicationdetails.asp?id=236

CAA "Aviation Maintenance Human Factors (JAA JAR145)", CAP716, Civil Aviation Authority, Gatwick, UK, January **2002a.** http://www.caa.co.uk/publications/publicationdetails.asp?id=275

CAA "Airworthiness Notice 43: Aircraft Field Loadable Software", Issue 2, CAP455, Civil Aviation Authority, Gatwick, UK, October **2002b.** http://www.caa.co.uk/publications/publicationdetails.asp?id=240

Clark E. "Improving the Effectiveness of First Generation HUMS - A Customer View", Proceedings of the Rotor Transmission System Health Monitoring - Experience Update Seminar, S553, Institution of Mechanical Engineers, London, UK, November **1997.**

Clark E. "A Customer Perspective on Helicopter Safety in North Sea Operations", Proceedings of the Rotorcraft Flight Safety Conference, Royal Aeronautical Society, London, UK, November **1998.**

CRHSNCS "NOU 2002:17 Helicopter Safety on the Norwegian Continental Shelf - Review 2: Trends, Objectives, Risk Influencing Factors and Measures", Ministry of Transport and Communications, Oslo, Norway **2000.** http://odin.dep.no/sd/engelsk/publ/rapporter/028031-020002/index-dok000-b-n-a.html

Dobson I. "Transmission Health Monitoring Experience", Proceedings of the Rotor Transmission System Health Monitoring - Experience Update Seminar, S553, Institution of Mechanical Engineers, London, UK, November **1997.**

Dobson I. "IHUMS Data Collection and Analysis -Operators Perspective", Proceedings of the Monitoring and Intelligent Use of Data - 'Turning Water into Wine' Conference, Royal Aeronautical Society, London, UK, March **2000a.**

Dobson I. Personal communication with the author, June **2000b.**

Edwards C. "The Case for Compliance Monitoring in Maintenance Organisations: Process and Practice Monitoring" Presentation available at Royal Aeronautical Society Human Factors Group Web Site. Accessed: November **2002.** http://www.raes-hfg.com/eng0812.htm

Evans A.J. "Flight Deck Indication of Helicopter Health Monitoring Data - A Critique", *Proceedings of the Institution of Mechanical Engineers Part G - Journal of Aerospace Engineering*, Issue G5, UK November **2002a.**

Evans A.J. "UK HUMS Effectiveness Review: 2000 and 2001", Proceedings of the HUMS ITT Meeting, 58th American Helicopter

Society Forum, Montreal, American Helicopter Society, Alexandria, Virginia, USA June **2002b.**

FAA "Airworthiness Approval of Rotorcraft Health and Usage Monitoring Systems", Federal Aviation Administration, Washington DC, USA July **1999**.

http://www1.faa.gov/certification/aircraft/Rot_Pol_Hums.htm

Gordon A.C. "An Operator's Viewpoint", Proceedings of the Helicopter Airworthiness - A Post HARP Review Conference, Royal Aeronautical Society, London, UK, February, **1989.**

HARP "Review of Helicopter Airworthiness", CAP491, Civil Aviation Authority, London, UK June **1984.**

HHMAG " Helicopter Health Monitoring Advisory Group Working Group Reports: Maintenance Credit, Usage Monitoring and HUMS Implementation", Civil Aviation Authority, London, UK, February **1997.** Note: These Working Groups originally reported to the HHMAG in November 1992, January 1994 and November 1995 respectively.

Heather A. "Support of In-Service HUMS", Proceedings of the 27th Meeting of the Helicopter Health Monitoring Advisory Group, Gatwick, UK, November **1997.**

Hokstad P., Jersin E, Hansen G.K., Sneltvedt J and Sten T. "Helicopter Safety Study 2", Report STF38 A99423, SINTEF, Trondheim, Norway, December **1999.**

Hopkins A. "Managing Major Hazards", Allen and Unwin, St Leonards, NSW, Australia, **1999.**

Hopkins A. "Lessons From Longford: The Esso Gas Plant Explosion", CCH Australia, Sydney, NSW, Australia, **2001.**

JAA "JAR-145: Approved Maintenance Organisations", Amendment 4, Joint Aviation Authorities, Hoofddorp, The Netherlands, November **2001a.**

JAA "NPA145-12: Human Factors", Joint Aviation Authorities, Hoofddorp, The Netherlands, November **2001b.**

JAA "JAR-OPS 3: Commercial Air Transport (Rotorcraft)", Amendment 2, Joint Aviation Authorities, Hoofddorp, The Netherlands January **2002.**

James S.L. "HUMS - A Regulatory Viewpoint of its Roles, Benefits and the Implementation Challenges", Proceedings of the Health and Usage Monitoring of Rotorcraft Transmission Systems - Review of Service Experience Seminar, S247, Institution of Mechanical Engineers, London, UK, November **1994.**

Jeram-Croft M. "MOD Requirements and Justification for Advanced Health and Usage Monitoring Techniques for Military Helicopters", Proceedings of the Health and Usage Monitoring of Rotorcraft Transmission Systems - Review of Service Experience Seminar, S247, Institution of Mechanical Engineers, London, UK, November **1994.**

Jesse D. and Slssinger M "Achieving a Working HUMS from Design Philosophy", Proceedings of HUMS - Avionic Implementation Conference, Royal Aeronautical Society, London, UK, April **1994.**

Ingstad O., Rosness R., Sten T., Ullberg T., Rausand M., and Lydersen S. "Helicopter Safety Study", Report STF75 A90008, SINTEF, Trondheim, Norway, November **1990.**

Kershner S.D., Johnson J.B. and Gamauf M.D. "Sikorsky Support to Commercial Health and Usage Monitoring Systems - A Summary of Forty Months of Support", Proceedings of the 53rd American Helicopter Society Forum, American Helicopter Society, Virginia Beach, Virginia, USA May **1997.**

Kørte J. and Aven T. "Health Monitoring of Helicopter: A Framework for Decision Making Based on Health Information", in Cottam M.P., Harvey D.W., Pape R.P. and Tate J. (Eds.), "Foresight and Precaution: Proceedings of ESREL, SARS and SRA-Europe Annual Conference 15-17 May 2000", Balkema, Rotterdam, The Netherlands, pp933-940 **2000.**

Land J.E. and Daniels G. "Helicopter Flight Data Recorder and Health and Usage Monitoring Systems", Proceedings of the 16th Symposium on Aircraft Integrated Monitoring Systems, DLR, Braunschweig, Germany, September **1991.**

Nevins J.J. "Are Vibrations Dangerous?", *Flying Safety*, USA pp9-13 October **2002.**

Neubert C. and Mimnagh M. "Flight and Iron Bird Test Results of the US Navy SH-60 Helicopter Integrated Diagnostics System Program", Proceedings of the Rotor Transmission System Health Monitoring - Experience Update Seminar, S553, Institution of Mechanical Engineers, London, UK, November **1997.**

McColl J. "Overview of Transmission HUM Performance in UK North Sea Helicopter Operations", Proceedings of the Rotor Transmission System Health Monitoring - Experience Update Seminar, S553, Institution of Mechanical Engineers, London, UK, November **1997.**

McFadden P.D. "Analysis of the Vibration of the Input Bevel Pinion in RAN Wessex Helicopter Main Gearbox WAK143 Prior to Failure", Aero Propulsion Report 169, Defence Science and Technology Organisation, Melbourne, Australia **1985.**

McKim N. "Brunei Shell Petroleum IHUMS Management", Presentation to the European Helicopter Operator's Committee HUMS Working Group, Redhill, UK, April **1997.**

Pouradier J-M. "An Assessment of the Eurocopter Experience in HUMS Development and Support", Proceedings of the Aerospace Transmissions Technology into the 21st Centaury European Seminar, Marseille, Institution of Mechanical Engineers, London, UK, May **2001.**

Reason J.T. "Managing the Risks of Organisational Accidents", Ashgate Publishing, Aldershot, UK **1997.**

Salzer M.L.W. "The Variability of Vibration Monitoring Parameters from Helicopter Transmissions During Normal Operations", Proceedings of the Health and Usage Monitoring of Rotorcraft Transmission Systems - Review of Service Experience Seminar, S247, Institution of Mechanical Engineers, London, UK November **1994.**

Snook S. A. "Friendly Fire: The Accidental Shootdown of US Blackhawks over Northern Iraq", The Princeton University Press, Princeton, USA **2000.**

Stewart R.M. "Turning Water into Wine - Keynote Address", Proceedings of the Monitoring and Intelligent Use of Data - 'Turning Water into Wine' Conference, Royal Aeronautical Society, London, UK, March **2000.**

Vaughan D. "The Challenger Launch Decision: Risky Technology, Culture and Deviance at NASA", The University of Chicago Press, Chicago, USA **1996**

Vinall P.D. "Airworthiness of Helicopter Transmissions", Proceedings of Conference, Royal Aeronautical Society, London, UK, 6th February **1980**.

Wackers G. and Kørte J. "Drift and Vulnerability in a Complex System: Reliability of Condition Monitoring Systems in North Sea Offshore Helicopter Transport", in Dym C. L., and Winner L. (Eds.) "Mudd Design Workshop III: Social Dimensions of Engineering Design Proceedings a Workshop", Harvey Mudd College, Claremont, California, USA pp319-339 **2001**

Weick K., Sutcliffe K. and Obstfelf "Organising for High Reliability: Processes of Collective Mindfulness", *Research in Organisational Behaviour*, Volume 21 pp81-123 **1999**. Note: quoted in Hopkins (2001).

Westland "Vibration Health Monitoring: Improving Safety and Reducing Costs of Transmission Systems", Brochure B3273, Westland Helicopters Ltd, Yeovil, UK May **1993**

White E.R. (ed) "A Guide to Health Monitoring in Helicopters", Helicopter Health Monitoring Advisory Group, Civil Aviation Authority, London, UK September **1990**.

Witham J.E. "Rotorcraft Reliability and Maintainability - A CAA View of Future Trends", Proceedings of Rotorcraft Reliability and Maintainability - Future Design Requirements Conference, Royal Aeronautical Society, London, UK, March **1991**.

U.S. ARMY LEAD THE FLEET USAGE ANALYSIS

David White
Westar Corporation

Michael McFalls
U.S. Army Aviation & Missile Command

David Largess
COBRO (a Westar Company)

ABSTRACT

The purpose of the U.S. Army Lead the Fleet (LTF) program is to rapidly accumulate flight hours on selected U.S. Army helicopters to identify safety, Reliability, Availability, and Maintainability (RAM) issues before they occur during fleet operational usage. The assets of the LTF Program currently include one helicopter of each of the following model mission design series: UH-60A and UH-60L Black Hawk, CH-47D Chinook, AH-64A Apache, and AH-64D Longbow Apache. The LTF Program was originally conducted from 1986 to 1995 and re-instituted in April 2002. This paper will report the approach, status, and early results of the LTF usage analysis.

The U.S. Army Aviation Technical Test Center (ATTC) pilots fly aggressive mission profiles and rapidly accumulate flight hours to stimulate usage-related discrepancies prior to operational occurrences. Structural and system deficiencies are to be identified, addressed, and corrected prior to the need for costly fleet-wide groundings, restorations, modifications, or retrofits.

LTF provides aircraft usage information to correlate with discrepancies and establish meaningful usage-related safety and RAM trends. The amount of time each airframe and each dynamic component is exposed to damaging flight regimes is monitored and recorded. The basic parameters used to identify the flight regimes include gross weight, stores configuration, airspeed, altitude, roll angle, vertical acceleration, and ground-air-ground cycles. This paper discusses in detail the approach used to identify the helicopter flight regimes and usage intensity from data provided by the LTF instrumentation.

ATTC, with support from COBRO (a Westar Company), collects, verifies, processes, and archives LTF operational and maintenance data using the Unified RAM (UniRAM) Data Management System. The UniRAM database is transmitted to the LTF Data Analysis Team to process, evaluate, and analyze trends.

The Westar Corporation LTF Data Analysis Team is under the direction of the U.S. Army Aviation and Missile Command (AMCOM) Test and Evaluation Management Office (TEMO). The team structured the data analysis effort to evaluate LTF aircraft discrepancies and analyze the UniRAM database to establish safety and RAM trends that result from the rate and intensity of LTF usage.

A review of the previous FY86-95 LTF Program shows that LTF is able to identify problems that include:

- Airframe and dynamic component degradation.
- Engine, drive train, fuel systems, and hydraulics failures.
- Vibration and torsional stability related problems.
- Weapons effectiveness and avionics / electronics malfunctions.

This paper concentrates on the current LTF Program mechanical system and dynamic component structural issues.

BACKGROUND

The LTF Program Manager is Mr. Mike McFalls, who heads TEMO within the Aviation and Missile Research, Development, and Engineering Center (AMRDEC) of AMCOM at Redstone Arsenal, AL.

ATTC at Ft. Rucker, AL, with the support of Westar and COBRO, is responsible for all aspects of the LTF Program execution including aircraft operation; usage and maintenance data collection; and data verification, archival, and transmission. Westar analyzes data for the AMCOM TEMO in Huntsville, AL.

Each Program Management Office (PMO) with responsibility for the oversight of the following systems participates in the LTF Program:

- AH-64A/D Apache
- UH-60A/L/M Black Hawk
- CH-47D/F Chinook
- Aircraft Survivability Equipment (ASE)

The PMO is responsible for defining solutions and implementing fleet-wide fixes for adverse trends identified in the LTF Program.

The LTF Program findings are coordinated and shared with the following organizations:

- Aircraft System PMOs
- Aviation Engineering Directorate (AED)
- Integrated Materiel Management Center (IMMC)
- U.S. Army Safety Center (USASC)
- Army Materiel Systems Analysis Activity (AMSAA)

LTF APPROACH AND OBJECTIVE

The LTF Program approach is to fly controlled conditions and scenarios and accelerate up the "reliability curve" ahead of the fielded aircraft in order to:

- Identify helicopter system and component deficiencies and failure modes.
- Provide information to resolve safety, RAM, and logistics issues.
- Optimize system and component replacements and improvements in the ongoing U.S. Army Aviation recapitalization efforts.

The objectives of the LTF Program are to reduce operational and sustainment costs, improve system reliability, and increase safety across Army Aviation.

LTF FLIGHT OPERATIONS

LTF Flight Profiles: LTF pilots fly profiles that replicate the operational missions and exercises flown in the field. The LTF flight profiles include the following:

- Combat missions
- Maintenance test flights
- Internal / external loads
- Training flights
- Range operations
- Live fire operations

LTF pilots exercise all the aircraft systems:

- ASE
- Weapons
- Navigation equipment
- Night vision systems
- Communication equipment
- Auxiliary equipment (hoists, etc.)

LTF pilots operate the aircraft in representative field environments:

- High altitude
- Desert conditions
- Cold weather
- Aircraft are not hangared

LTF pilots fly engineering regimes to stress the aircraft and systems to detect problems before they occur in the field.

LTF Aircraft Plan: The aircraft planned for participation in the LTF Program are shown in Fig. 1. The aircraft currently in the LTF Program include the following:

- A/C with data bus:
  - AH-64A Apache
  - AH-64D Longbow Apache
- A/C without data bus:
  - UH-60A Black Hawk
  - UH-60L Black Hawk
  - CH-47D Chinook

Aircraft that will be added to the LTF as they become available include:

- A/C with data bus:
  - CH-47F (FY03)
  - UH-60M (FY04)
  - OH-58D (FY05)

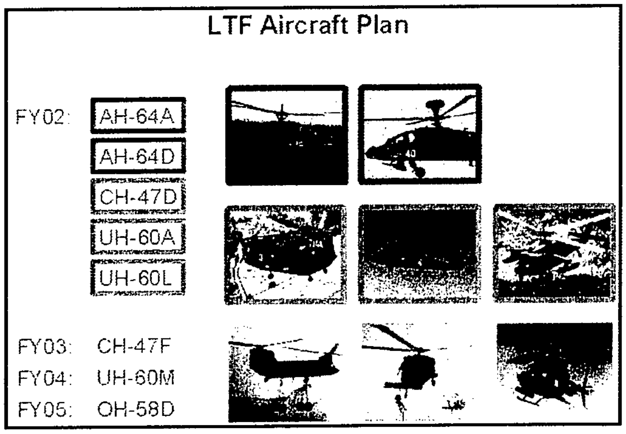

Fig. 1 ***LTF Aircraft Plan***

The remainder of this paper will concentrate on the LTF data collection and analysis for the AH-64A aircraft.

AH-64A LTF OPERATIONS TEMPO (OPTEMPO): The AH-64A planned LTF usage is 60 hours per month as shown in Fig. 2. This planned OPTEMPO is approximately 3.6 times the flight hour rate of the average U.S. Army AH-64A. The LTF execution plan is not linear because it includes scheduled phased maintenance at 250 flight hour intervals.

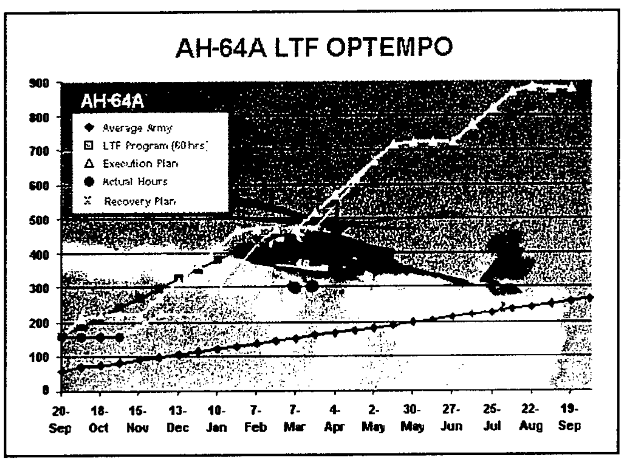

Fig. 2 ***AH-64A LTF OPTEMPO***

LTF DATA ACQUISITION, ANALYSIS, AND DISTRIBUTION

LTF Data and Action Flow: Fig. 3 shows the LTF usage and maintenance data and action flow. The AMCOM AMRDEC TEMO, who is responsible to a General Officer Steering Group for the overall success of the LTF Program, provides LTF Program Management.

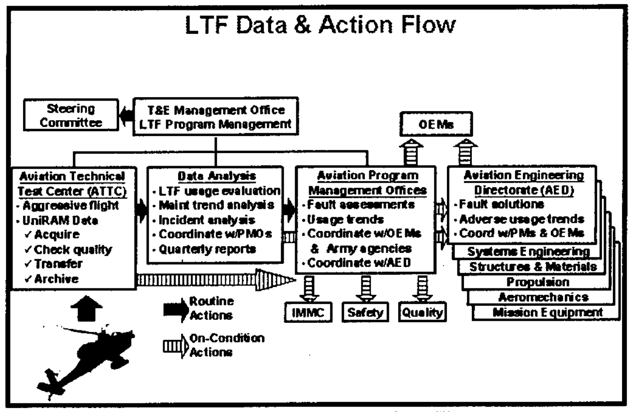

Fig. 3 ***LTF Data and Action Flow***

ATTC pilots fly the aggressive flight plan that includes the accelerated OPTEMPO, operational missions and flight profiles, and engineering regimes. ATTC, with support of COBRO, collects all LTF aircraft usage and maintenance data in the UniRAM database. COBRO acquires the data; checks the data quality; archives the database at ATTC, Ft. Rucker, AL; and replicates a copy to Westar in Huntsville, AL, daily. Unscheduled maintenance write-ups are scored as the events occur. A failure review board meets quarterly, validates the faults, and determines the final scoring of each event.

Westar evaluates the usage information to identify adverse maintenance and fault occurrence trends, and analyzes aircraft incidents. The resulting information is coordinated with the aircraft and aircraft equipment PMOs. LTF aircraft usage analysis information is reported to the PMOs quarterly.

The PMOs are ultimately responsible for resolution of adverse trends identified by the LTF data analysis. The PMOs coordinate with IMMC, USASC, and Quality on an on-condition basis to resolve LTF-identified issues. The PMOs also coordinate with the Systems Engineering Division of AED to resolve fault, failure, and adverse usage trends. If required, the PMO will work with the Original Equipment Manufacturer (OEM) to resolve aircraft problems and implement fleet-wide fixes.

AED Structures and Materials, Propulsion, Aeromechanics, and Mission Equipment Divisions participate in fault, failure, and adverse trend solutions as required to resolve aircraft problems.

UniRAM Data: UniRAM Data consists of aircraft operational usage and maintenance data, including:

- Pilot profile card
  - Aircraft configuration – fuel weight, stores, etc.
  - Profile flown – external fuel, external loads, etc.
  - Qualitative comments
- Enhanced Logbook Automation System (ELAS)
  - Operational information
  - Scheduled maintenance
  - Unscheduled maintenance
- Aircraft state and motion – from AMPOL

DataMARS bus monitor for bus aircraft or C-MIGITS GPS/INS system for non-bus aircraft.

- Ground-air-ground cycles
- Pitch and bank angle
- Airspeed
- Load factor
- Altitude
- Engine torque

LTF Bus Aircraft: AMPOL DataMARS: The AH-64A has a limited data bus with limited parameters available for usage analysis as shown in Fig. 4. Aircraft position, attitude, motion, and engine torque are used to identify the AH-64A usage in the engineering flight regimes and evaluate flight incidents. The AMPOL Data Monitoring, Analysis, and Recording System (DataMARS) monitors and records the data available from the AH-64A MIL-STD-1553 Data Bus.

Parameter	Aircraft Model		Parameter	Aircraft Model	
	AH-64A	AH-64D		AH-64A	AH-64D
Time	X	X	Rotor RPM		X
Indicated Airspeed	X	X	Pitch Rate	X	X
Pressure Altitude	X	X	Roll Rate	X	X
Radar Altitude	X	X	Yaw Rate	X	X
Heading	X	X	VX		X
Latitude	X	X	VY		X
Longitude	X	X	VZ		X
Pitch Angle	X	X	AX		X
Roll Angle	X	X	AY		X
Yaw Angle		X	AZ		X
Engine Torque	(2)	(2)	Lat Cyclic Position		X
Gas Producer Speed (NG)		(2)	Long Cyclic Position		X
Power Turbine Speed (NP)		(2)	Pedal Position		X
Turbine Gas Temp (TGT)		(2)	Collective Position		X

Fig. 4 ***AH-64A/D Data Bus Parameters***

LTF Data Analysis Process: The LTF data analysis process is shown in Fig. 5.

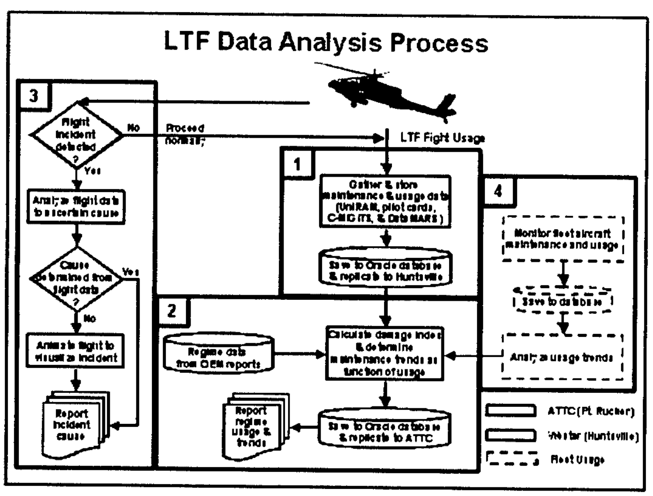

Fig. 5 ***LTF Data Analysis Process***

Block 1 of Fig. 5 illustrates the data acquisition, storage, and transmission functions that occur at ATTC. Block 2 summarizes the engineering analysis of aircraft usage performed by Westar to determine the aircraft usage intensity and identify adverse maintenance trends. The results of these usage analyses are saved to the UniRAM database and replicated back to ATTC. The results of the analysis are reported to the PMOs quarterly and are planned to be available on a near real-time basis at a password-protected web portal beginning in January 2003. Blocks 3 and 4 will be discussed later in this paper.

AH-64A Percent Time or Number of Events per Hour in Damaging Flight Regimes: A sample analysis of the time and number of events the AH-64A has spent in damaging regimes during LTF usage is shown in Fig. 6. The OEM determined the fatigue lives of AH-64A life-limited, flight-critical components when the aircraft was designed. These component lives were based upon a conservative assumption of the percentage of flight time and number of discrete events that each component could spend in damaging flight regimes. When the AH-64D was designed, the fatigue lives of parts that are common between the AH-64A and AH-64D were re-assessed based upon design assumptions of AH-64D separate fleet and training usage. Since the common parts are not tracked by AH-64A/D mission design series, the Army always assumes that the design fatigue life of a common part is based upon the usage that produces the shortest of the three possible fatigue lives: AH-64A usage, AH-64D fleet usage, or AH-64D training usage. Therefore, flight regime usage for any individual common part may be determined by any of the three usage scenarios. As shown in Fig. 6, the AH-64A and AH-64D have 100 discrete parts in common that may experience damage in any of 1,210 defined combinations of aircraft configurations and flight regimes. For simplicity, Fig. 6 groups the 1,210 flight regimes into 15 types. The sample results shown in Fig. 6 are based upon approximately 53 flight hours of AH-64A LTF usage.

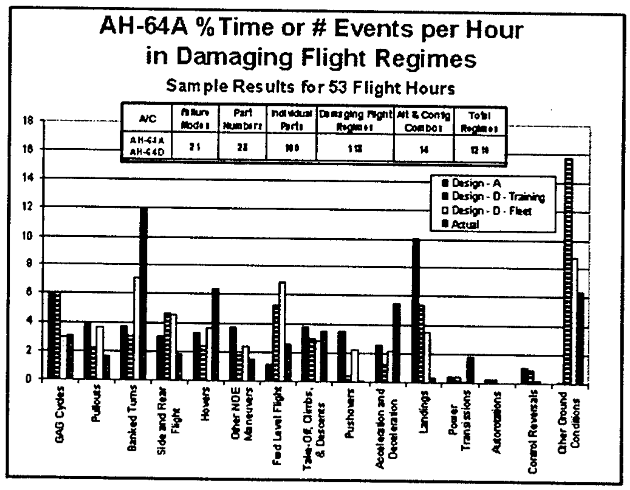

Fig. 6 ***AH-64A Regime Usage***

The black bars are an aggregate of the percent time or number of events per flight hour in the damaging regimes that the 21 AH-64A component failure modes were designed to experience. The black and gray bars are the aggregate of the AH-64D design training usage. The gray and white bars are the aggregate of the AH-64D design fleet usage, and the red bars are the aggregate of the actual LTF experience. Although a large percentage of LTF time was spent in the most damaging type regimes, such as banked turns, the time was spent in the lower aspects of the regime type. For example, most of the 53-hour sample of damaging LTF banked turns occurred at relatively low angles that produce only mild damage. The actual damage produced by the LTF flight regime experience can be seen in Fig. 7.

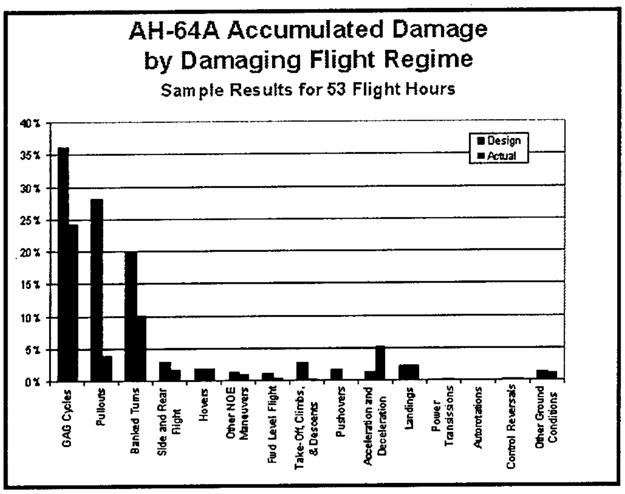

Fig. 7 ***AH-64A Accumulated Damage***

AH-64A Accumulated Damage by Damaging Flight Regime: Fig. 7 shows the aggregate of the damage experienced by the 21-component failure modes for each type of flight regime. This chart is a normalized comparison of the aggregate of the 28-component design damage with the actual damage experienced in the first 53 hours of LTF usage. Fig. 7 shows that most of the first 53 hours of AH-64A LTF damage occurred due to ground-air-ground cycles.

AH-64A Damage Rate Accumulation: The normalized cumulative AH-64A LTF damage for the aggregate of the 21 component failure modes is compared to the normalized design values in Fig. 8. The chart is the 53-hour accumulation of 35 flights and shows the overall damage intensity is about 50% of the design value for the aggregate of the 100 life-limited components.

AH-64A LTF Quality Deficiency Reports (QDR): The six initial AH-64A LTF QDRs are summarized in Fig. 9. The first two listed QDRs are related to the Pitch Change Link (PCL) assembly. The first was a Category I, or safety of flight, deficiency that involved a worn PCL rod end. The rod was on the verge of separating from the bearing, which could have resulted in loss of aircraft. The second listed QDR resulted from the discovery of excessive play between three PCL rod ends and their bearings. The PCLs are not fatigue life-limited components and are removed and replaced based upon inspection for excessive wear. However, the pitch housing is an adjacent life-limited component, and the PCLs and pitch housing are subject to similar loads and stresses, and therefore wear and damage, from the same type flight regimes.

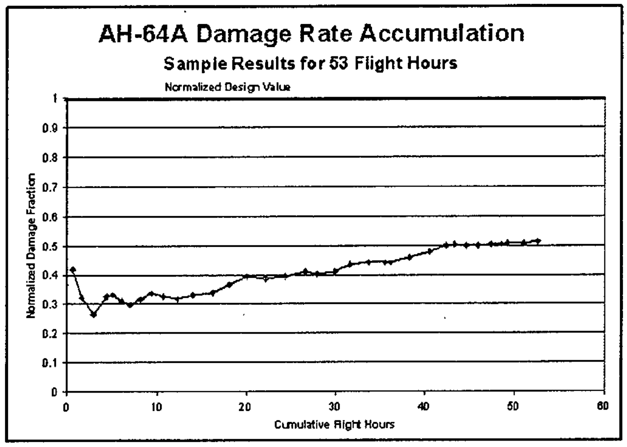

Fig. 8 ***AH-64A Damage Rate Accumulation***

Model	Category	Aircraft Component	QDR Report
AH-64A	I	Pitch Change Link Assembly	W31NWZ020012
AH-64A	II	(3) Pitch Change Link Bearings (Excessive Play)	W31NWZ020018
AH-64A	II	Engine #2 Nose Gearbox Failure	W31NWZ020988
AH-64A	II	Primary Hydraulic System Manifold	W31NWZ020008
AH-64A	II	Primary Hydraulic System Manifold	W31NWZ020013
AH-64A	II	Primary Hydraulic Axial Pistons Pump Leaking	W31NWZ020009

Fig. 9 ***AH-64A Initial Quality Deficiency Reports***

AH-64A Pitch Housing Damage Accumulation: The pitch housing is one of the 21 life-limited component failure modes used in determining the AH-64A damage index. The design and the first 53 flight hours of AH-64A LTF usage are shown in Fig. 10. The pitch housing damage for the flight in which excessive vibrations, due to PCL degradation, were experienced was tacked on to the 53 flight hours of damage as shown in Fig 10. It should be noted that the pitch housing damage rate for that flight exceeded the overall design damage rate by 130%. Although the data are not conclusive, it may be surmised that the intensity of the flight may have exacerbated the PCL wear.

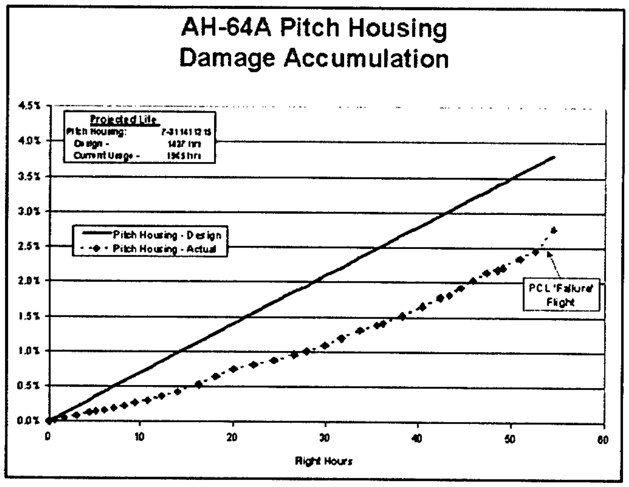

Fig. 10 ***AH-64A Pitch Housing Damage Accumulation***

AH-64A RH NGB Incident: Block 3 of Fig. 5 is a schematic of the LTF flight incident evaluation process. An example of such an incident occurred during an LTF flight at Ft. Rucker, AL, when a 'chip light' indicated a metal chip was detected in the Right Hand (RH) (or No. 1) Nose Gearbox (NGB). The aircrew performed the appropriate corrective action that resulted in the shutdown of the No. 1 engine and landing of the aircraft.

The time histories of the pertinent flight parameters for the NGB incident are shown in Fig. 11. Important incident events are noted and chronologically numbered.

A post-flight inspection of the RH NGB revealed that there were large metal fragments present in the gearbox. The NGB was held as a QDR exhibit. Teardown at the depot will document the nature of the failure.

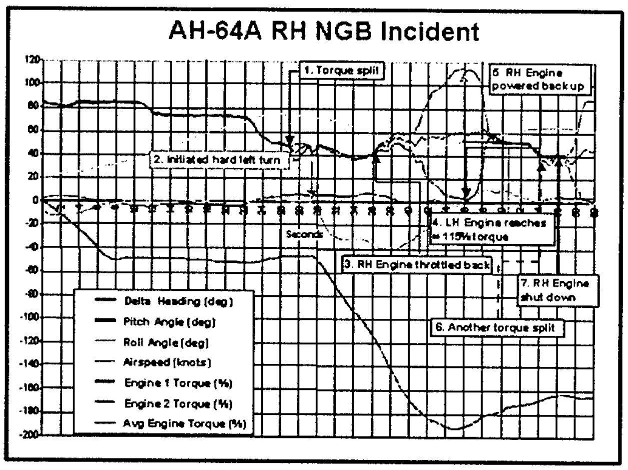

Fig. 11 ***AH-64A RH NGB Incident***

FLEET USAGE

Block 4 of Fig. 5 identifies an essential element in achieving the potential benefits of the LTF Program. Having a better understanding of actual fleet usage can enhance the value of LTF. Downloading and analyzing the AH-64D on-board Maintenance Data Recorder (MDR) data can provide fielded Apache aircraft usage information. These data will be compared to AH-64D LTF usage to ensure achievement of LTF goals. Also, AH-64D usage will be correlated with maintenance actions to determine the impact of usage intensity on maintenance actions. For example, unscheduled gearbox overhauls will be correlated with pertinent usage parameters, such as:

- Number of 'cold' starts
- Time spent at various engine torque levels
- Time spent in stressful aircraft maneuvers

AH-64A and AH-64D aircraft have common systems and components; therefore, knowledge obtained on the AH-64D will benefit the AH-64A.

SUMMARY

The LTF Program is a systems engineering approach to the evaluation and understanding of helicopter usage faults, failures, and trends. This objective is achieved by selectively accelerating U.S. Army helicopter usage to experience usage-related faults and failures prior to their occurring during fleet usage. The LTF Team accomplishes the data evaluation by collecting, archiving, transmitting, analyzing, and correlating helicopter maintenance, fault, and failure information with helicopter usage information to identify adverse usage trends. The LTF Program is also a systems approach to evaluating in-flight faults and failures by analyzing flight recorded parameter time histories to understand the nature and cause of the incident. In some cases, the analysis includes the assessment of the flight's geographic location and environment of the incident. When required, the flight is animated to understand aircraft motion and response at the time of the incident.

Abstracts of Papers Not Included in This Document

The following are the abstracts of papers accepted for the conference but where no formal paper is available or could be made available in time for inclusion in this document. These are listed in the order received.

vHUMS System Architecture, Adam Forsyth, VAMtec (GPS Online Ltd)

This presentation will cover the technical overview of issues facing a vHUMS implementation and possible technology architecture solutions.

vHUMS Applications in the United Kingdom, Keith Mowbray, VAMtec (Dytecna Ltd)

Case Study of vHUMS trials and issues being conducted in the UK as part of MoD Land Systems projects.

Engine Vibration Control: a step by step way to HUMS, François Cantegreil, SEMIA, France

Owing to the high cost of the equipment available on the market, the number of helicopters partially equipped with HUMS is only a few hundreds, that is a very small part of the worldwide helicopter fleets.

This situation holds reasons to be optimistic since:

1. The available target for a first installation of HUMS remains very important
2. The technology currently available enables to very markedly reduce the costs and weights of the proposed equipment
3. The regulation requires certain operators to imperatively equip themselves to be able to continue operating their helicopters

These three factors combine and the target aimed at by the components manufacturers then becomes very important, which enables to increase the quantities to produce, therefore still more reduce their costs.

SEMIA who holds a position as a leader in the field of the Engine Vibration Control which is one of the various modules of the HUMS, explains in its paper the development strategy followed by SEMIA: to enable an operator to get equipped progressively and for a reduced cost.

It presents also the progress of its product range which takes into account the necessary compatibility with the other modules of the HUMS.

Considering the Role of Health Monitoring in Fixed Wing Aircraft HUMS, L. Molent, DSTO Australia

The role of a Health and Usage Monitoring System (HUMS) for rotating and related components in helicopters airframes and engines is well defined and understood. In more recent times the term HUMS has been applied to fixed wing airframe systems. Generally these airframe HUMS consider the health (H) of an airframe through the fatigue monitoring or usage aspects (U) rather than a separate diagnostic capability monitoring health or structural integrity as found in most rotating component HUMS.

However with the development of so-called smart structures and miniaturisation some researchers are arguing that the "U" may be replaced by more direct "H", in other terms HuMS. In general it is envisaged that the health of the structure will be ensured through some "advanced" non-destructive inspection (NDI) techniques (using embedded sensors or the so-called smart structures etc) to detect damage, degradation and flaws before they cause structural failure. Whilst this may seem attractive, there are many significant through life structural integrity issues to be addressed before this can be viable. The purpose of this paper is to canvass issues regarding the application of HuMS to the through life management of fixed wing manned airframes. It specifically considers whether the role of prognostic usage monitoring should be replaced with direct health monitoring techniques.

Training Techniques and Methodologies for Helicopter Rotor Track and Balance, Mick Richmond, Australian Army Rotary Wing Aircraft Maintenance School

The Rotary Wing Aircraft Maintenance School (RAMS) is responsible for delivering all individual technical training to Army aviation tradespeople post initial employment trade training. This training includes all weapon system trade courses on UH-1H, S70A-9, BL206B and soon ARH Tiger.

In Feb 2002, training commenced in the Aircraft Advanced Technician Courses (ACAT). This course was designed to equip tradespeople with the fault diagnostic and management tools required prior to advancement to the supervisory level.

Previously, training of this type had not been conducted.

A significant portion of the course is in Rotor Track and Balance (RTB) and vibration analysis (VA) techniques. In the case of S70A-9 tradespeople, integration of this training into the Black Hawk Full Mission Flight Simulator (FMFS). The training educates tradespeople in the techniques that have been developed over the years that greatly assist in troubleshooting vibration related unserviceabilities. The formalising of this training was necessary to combat the dilution of experience at the operating units.

The presentation concentrates on these two particular areas. The structure of RTB and VA training conducted by RAMS, and the low cost data management tools that have been developed to reduce personnel workload such as associated with the routine vibration data collections from the 'on board' VA systems on the Australian Army CH47D.

Certification of Engine Health and Usage Monitoring Systems on Military Aircraft, SQNLDR Paul Parolo, ADF Directorate-General Technical Airworthiness

Australian military aircraft such as the Hawk 125 lead-In-Fighter and the C-130J transport come under the category of State Registered Aircraft and as such the Australian Defence Force (ADF) is the Certification Authority.

This paper includes current work on the Hawk and C-130J-30 as case studies.

HUMSSAVE 4 - An Upgraded Econometric Model for Cost/Benefit Studies of HUMS, Graham F Forsyth and Eric CJ Lee, DSTO Australia

One of the reasons to fit HUMS or conduct other Condition Monitoring processes is to save money overall, either by reduction in maintenance costs or by increased safety. For some years, the DSTO HUMSSAVE model has been widely used to model both the costs and returns from such installations.

The beta version of a new version of this modeling tool will be distributed at this conference. The new version widens the scope to include platforms other than helicopters and adds a number of indicators to help the user. This paper describes the design and use of this software.

Joint Advanced Health and Usage Monitoring System (JAHUMS) Advanced Concept Technology Demonstration (ACTD) Program, David Haas and Treven Baker, Naval Surface Warfare Centre and Aviation Applied Technology Directorate

The Joint Advanced Health and Usage Monitoring System (JAHUMS) Advanced Concept Technology Demonstration (ACTD) program is a joint U.S. Army/U.S. Navy Program. The JAHUMS ACTD supports a dynamic change in the maintenance philosophy of DOD helicopters and expands the HUMS paradigm in a number of areas. The objective is to demonstrate the utility of HUMS technologies in a military operational environment and to demonstrate and validate an open systems approach for technology insertion. Several technology modules developed by third party providers have been integrated into an open architecture system developed by Goodrich Aerospace for the U.S. Navy/Army Integrated Mechanical Diagnostic System. The presentation will describe the ACTD program and associated technologies and provide a summary of results to date and future plans.

A Planetary Gearbox Diagnostic Technique using Constrained Adaptive Lifting, Paul Samuel, Alfred Gessow Rotorcraft Centre

This paper presents a methodology for detecting and diagnosing gear faults in the planetary stage of a helicopter transmission. This diagnostic technique is based on the constrained adaptive lifting algorithm. The lifting scheme, developed by Wim Sweldens of Bell Labs, is a time domain, prediction-error realization of the wavelet transform that allows for greater flexibility in the construction of wavelet bases. Classic lifting analyses a given signal using wavelets derived from a single fundamental basis function. A number of researchers have proposed techniques for adding adaptively to the lifting scheme, allowing the transform to choose from a set of fundamental bases the basis that best fits the signal.

This characteristic is desirable for gear diagnostics as it allows the technique to tailor itself to a specific transmission by selecting a set of wavelets that best represent vibration signals obtained while the gearbox is operating under healthy-state conditions. However, constraints on certain basis characteristics are necessary to enhance the detection of local wave-shape changes caused by certain types of gear damage.

The proposed methodology analyses individual tooth-mesh waveforms from a healthy-state gearbox vibration signal that was generated using the synchronous signal-averaging algorithm. Each waveform is separated into analysis domains using zeros of its slope and curvature. The bases selected in each analysis domain are chosen to minimize the prediction error, and constrained to have the same-sign local slope and curvature as the original signal. The resulting set of bases is used to analyse future-state vibration signals and the lifting prediction error is inspected. The constraints allow the transform to effectively adapt to global amplitude changes, yielding small prediction errors. However, local wave-shape changes associated with certain types of gear damage are poorly adapted, causing a significant change in the prediction error.

The constrained adaptive lifting diagnostic algorithm is validated using data collected from the University of Maryland Transmission Test Rig and the results are discussed.

Analysis of Epicyclic Gearbox Vibration, David Blunt and David Forrester, DSTO Australia

Many aircraft transmissions use epicyclic gear trains, particularly helicopter main rotor gearboxes and propeller reduction gearboxes. As these gears form a non-redundant critical part of the drive to the main rotor, or propeller, it is important to have advanced techniques and tools to assess the condition of these components. One such tool is vibration analysis. However, epicyclic gear train vibrations are difficult to analyse. Not only are there multiple planet gears producing similar vibrations, but there are multiple and time-varying vibration transmission paths from the gear mesh points to any vibration transducer mounted on the gearbox housing. These factors combine to reduce the sensitivity of conventional fault detection algorithms when they are applied to epicyclic gears. This paper outlines the DSTO-developed technique for analysing epicyclic gear train vibration, based on an algorithm for separating the meshing vibrations from each planet. The results of applying this technique to seeded fault tests, using both DSTO and US Navy vibration data, are shown to significantly improve the detection of localised gear faults.

Use of Conductance Measurement For Battery Health Monitoring, Peter W. House, Ultra Electronics

The paper will outline the history of conductance measurement as applied to saturated lead acid batteries and its use for permanently monitoring batteries when in service.

Various system architectures will be presented showing how this may form part of a HUMS system on the vehicle and the potential expansion towards power management using data provided by the vehicle HUMS system.

High Precision and Bandwidth Efficient Data Recording Format, Dr. Balázs Bagó and Wolfgang Isermann, HEIM Systems

It's a common task for all data acquisition and recording system to find solution for capturing data the most accurately with the least needed bandwidth. In most test systems not only the data itself, but their accurate timing is very important. This article presents a recording medium independent data acquisition and recording technique with highly optimal bandwidth usage and very high absolute time accuracy.

First several possible data recording formats are presented which are optimized for capturing different class data sources - continuous, burst or message-oriented data types. The key features of a bandwidth optimized time stamping technique - based on the IRIG-107 data format - is outlined, which is used in the DATaRec recorder series.

The recording format needs low data processing during recording and also very effective for analyzing the recorded data with computer analysis tools. The method is shown, how the time synchronization is carried out to absolute time sources like GPS or IRIG time codes - along with the solution how the absolute time stamping is maintained even in cases when the time source is temporarily not receivable during a mission.

As an interesting and unique application of the presented technique shows how data reconstruction can be achieved exactly synchronized to a given absolute time - allowing the synchronization of several recorders during replay without any complicated synchronization mechanism between them.

Smart Structures and Materials: A DSTO Key Initiative, S. Galea, A. Wilson , C. Scala and A. Wong, DSTO Australia

The Australian Defence Force is increasingly facing escalating costs on through-life support for major platforms (ships, aircraft and land vehicles). The application of smart materials and structures technologies in platform management systems is seen as a very promising approach to reduce these costs and to potentially achieve significant enhancement of platform capability. A new DSTO Key Initiative, "Smart Materials and Structures", has been recently developed and funded to address these technologies. The Initiative will build on and grow the current activities within DSTO and promote collaboration with external Australian institutes and industry.

This paper will present an overview of the Initiative and the generic sensor and system issues inherent in the "whole-of-platform" and "whole-of-life" monitoring and management of major defence platforms. Examples for some particular elements of this will be drawn from current work in DSTO. Other presentations in the conference will cover the technical and scientific aspects of these in more detail.

RSL's AH64 T-HUMS - the HUMS that does the job! H. Silverman, RSL Electronics

RSL's T-HUMS (total health and usage management system) uses innovative diagnostics and prognostics technologies which are carried out by T-HUMS unique architecture. These enable real time health situation awareness of the entire helicopter including the rotors, gears, shafts, engines and other utility systems. The on-board / in-flight automatic data mining provides the pilots as well as flight-line technician explicit and simple decision support while minimizing data handling and data mining overheads, and makes T-HUMS the first market available true HUMS.

T-HUMS self-identifies the location of deterioration prior to the critical level land provides clear advisories how the aircraft should be flown to avoided unsafe condition or further damage while the flight line technician is directed what maintenance activity to be done without the need to analyze the data manually. T-HUMS which is a generic HUMS is currently being flown by the IAF on board the AH64 helicopters. T-HUMS already showed substantial benefits with respect to flight safety, availability and LCC through real life events. Further success has been demonstrated by T-HUMS derivatives for UAVs and jet engines.

Maintenance Analysis, Safety and Training (MAST) Program - A comprehensive application of digital flight data, Errol Farr, Smiths Aerospace

Civil and military helicopter operators are required to maintain higher operational readiness rates than ever before. They are also required to maintain comprehensive aircrew training and safety programs. The introduction of integrated onboard maintenance system technology helps to address these issues by integrating the functions of data collection and recording, health and usage monitoring, maintenance diagnosis, and fault prediction. The Maintenance Analysis, Safety and Training (MAST) program meets these objectives by implementing a Helicopter Fleet Management System to reduce Operations and Support (O&S) costs and enhance safety and training. The MAST system also responds to compelling operational needs identified by the US Army Aviation Modernization Plan, Army Aviation's current recapitalisation efforts and the requirement for a 90% Fully Mission Capable (FMC) fleet.

The MAST program for the U.S. Army 160th Special Operations Aviation Regiment (SOAR) will adapt commercial HUMS items and install an Aircraft Integrity Monitoring System (AIMS) on MH-47E Chinook helicopters. The AIMS will expand the currently installed Voice & Data Recorder (VADR®) system with the capability to perform embedded diagnostics including: automatic on-board rotor track and balance, usage monitoring, performance monitoring, exceedance detection, and vibration monitoring for helicopter airframes, power trains and rotor system dynamic components. The performance of the HUMS will be demonstrated and verified through a one-year operational evaluation on two helicopter installations. Digital flight data, which is readily available from the Voice & Data Recorder (VADR®), will be used to implement a Flight Operations and Quality Assurance (FOQA) Program and display Graphical Replay of flight data. During this evaluation, data will also be collected to assess the level of O&S cost savings.

The fully integrated system utilizes HUMS data, flight data in a Flight Operations Quality Assurance (FOQA) program, and Graphical Replay to provide benefits across the entire Special Operations warfighting spectrum. The integrated fleet management system will use the data collected via the VADR® and the AIMS for:

- Maintenance - Health and Usage Monitoring
- Safety - Flight Operations Quality Assurance
- Training - Graphical Replay

The presentation will describe the system capabilities and summarize the status of the program development.

Keeping Track of the Future, Mal Davies, Helitune Ltd

The tip of a rotor blade is accepted as the most accurate place to determine the track of a blade. Although this is fundamentally difficult to accurately achieve, Helitune developed the linescan camera to do just that. The first camera was put in production for the UK military in the late 1980's and has proved to be very successful.

Since then several other manufacturers have developed camera trackers. Although most provide accurate track pictures, generally not at the tip, all of these cameras suffer from struggling in particular light conditions. To help overcome this Helitune have developed the MkII linescan camera. By utilising the latest DSP and CCD technology this camera provides the platform for the future. This is the first step towards the Helitune fixed tracker and the more accurate "find edge" tracking system.

This presentation demonstrates what the issues are, how they are being overcome and explains the future direction of the Helitune tracker.

DISTRIBUTION LIST

Third International Conference on Health and Usage Monitoring - HUMS2003

Graham F Forsyth (editor)

AUSTRALIA

DEFENCE ORGANISATION

S&T Program

Chief Defence Scientist
FAS Science Policy
AS Science Corporate Management
Director General Science Policy Development
shared copy (Doc Data Sheet)
Counsellor Defence Science, London (Doc Data Sheet)
Counsellor Defence Science, Washington (Doc Data Sheet)
Scientific Adviser to MRDC Thailand (Doc Data Sheet)
Scientific Adviser Joint (Doc Data Sheet)
Navy Scientific Adviser (Doc Data Sheet and distribution list only)

Scientific Adviser - Army (Doc Data Sheet and distribution list only)

Air Force Scientific Adviser (Doc Data Sheet)
Director Trials (Doc Data Sheet)
Platforms Sciences Laboratory

Chief of Air Vehicles Division (Doc Data Sheet and Distribution List Only)
Research Leader (Doc Data Sheet and Distribution List Only)
Graham Forsyth 200 copies

DSTO Library and Archives

Library Fishermans Bend
Library Edinburgh 1 copy
Australian Archives

Capability Systems Staff

Director General Maritime Development (Doc Data Sheet only)
Director General Aerospace Development (Doc Data Sheet only)

Knowledge Staff

Director General Command, Control, Communications and Computers (DGC4)
(Doc Data Sheet only)

Navy

SO (SCIENCE), COMAUSNAVSURFGRP, NSW (Doc Data Sheet and distribution list only)

Intelligence Program

DGSTA Defence Intelligence Organisation
Manager, Information Centre, Defence Intelligence Organisation

Defence Libraries

Library Manager, DLS-Canberra
Library Manager, DLS - Sydney West (Doc Data Sheet Only)

UNIVERSITIES AND COLLEGES

Australian Defence Force Academy
  Library
  Head of Aerospace and Mechanical Engineering
Hargrave Library, Monash University (Doc Data Sheet only)
Librarian, Flinders University

OTHER ORGANISATIONS

National Library of Australia

OUTSIDE AUSTRALIA

INTERNATIONAL DEFENCE INFORMATION CENTRES

US Defense Technical Information Center, 2 copies
UK Defence Research Information Centre, 2 copies
Canada Defence Scientific Information Service, 1 copy
NZ Defence Information Centre, 1 copy

ABSTRACTING AND INFORMATION ORGANISATIONS

Engineering Societies Library, US
Materials Information, Cambridge Scientific Abstracts, US

INFORMATION EXCHANGE AGREEMENT PARTNERS

Acquisitions Unit, Science Reference and Information Service, UK

SPARES (5 copies)

Total number of copies: **224**

DEFENCE SCIENCE AND TECHNOLOGY ORGANISATION
DOCUMENT CONTROL DATA

1. PRIVACY MARKING/CAVEAT (OF DOCUMENT)

2. TITLE	3. SECURITY CLASSIFICATION (FOR UNCLASSIFIED REPORTS THAT ARE LIMITED RELEASE USE (L) NEXT TO DOCUMENT CLASSIFICATION)
Third International Conference on Health and Usage Monitoring - HUMS2003	Document (U) Title (U) Abstract (U)
4. AUTHOR(S) Graham F Forsyth (editor)	5. CORPORATE AUTHOR Platforms Sciences Laboratory 506 Lorimer St Fishermans Bend Victoria 3207 Australia

6a. DSTO NUMBER	6b. AR NUMBER	6c. TYPE OF REPORT	7. DOCUMENT DATE
DSTO-GD-0348	AR-012-516	General Document	February 2003

8. FILE NUMBER	9. TASK NUMBER	10. TASK SPONSOR	11. NO. OF PAGES	12. NO. OF REFERENCES
2002/54165/1	NAV 02/149	DGTA	137	Various

13. URL on the World Wide Web	14. RELEASE AUTHORITY
http://www.dsto.defence.gov.au/corporate/reports/DSTO-GD-0348.pdf	Chief, Air Vehicles Division

15. SECONDARY RELEASE STATEMENT OF THIS DOCUMENT

Approved for public release

OVERSEAS ENQUIRIES OUTSIDE STATED LIMITATIONS SHOULD BE REFERRED THROUGH DOCUMENT EXCHANGE, PO BOX 1500, EDINBURGH, SA 5111

16. DELIBERATE ANNOUNCEMENT

No Limitations

17. CITATION IN OTHER DOCUMENTS Yes

18. DEFTEST DESCRIPTORS

Health and Usage Monitoring Systems, Helicopter Maintenance, Aircraft Maintenance, Airworthiness, Condition Monitoring

19. ABSTRACT

This document includes formal papers for the Third International Conference on Health and Usage Monitoring, HUMS2003, which will be held in Melbourne in February 2003. The scope of papers covers a wide range of monitoring issues with focus on the application to helicopters, military aircraft and gas turbine engines.

This document includes only the formal papers made available for publication during December 2002. Other papers will be available from the conference CD-ROM which includes this document, other formal papers, presentations not available as formal papers and related software.

For a short period, all these will also be available from the conference website at URL: http://www.dsto.defence.gov.au/corporate/conferences/hums.