Moving Beyond Aircraft Health Monitoring Mythology: The Journey Toward Obtaining Tangible Life-Cycle Management ROI.

Stephen R. Hall (halls@celeris.ca)

Celeris Aerospace Canada Inc., 880 Taylor Creek Drive, Orléans, Ontario, K1C 1T1, CANADA

Abstract

In their most rudimentary form, Aircraft Structural Health Monitoring Programs commenced some 50+ years ago following the de Havilland Comet accidents. The rapid developments in micro-processor, computer, sensor and data acquisition technologies that started in the late 1980's and continue through today have revolutionized the nature and the quantities of data that are being acquired. Yet despite all the technological advances, many Aircraft Structural Health Monitoring Programs are still failing to realize their full potential and provide ongoing tangible safety and economic life-cycle management benefits. This paper looks at some of the reasons why the full potential of these programs has still to be realized. Practical steps are proposed to translate the "myth" of what is promised to the reality of robust, viable structural health management capabilities upon which ongoing short and long-term life-cycle management decisions of fixed and rotary wing fleets can be confidently based.

Keywords: Aircraft Structural Health Monitoring, SHM, HUMS, Aircraft Life-Cycle Management, LCM, ASIP, Strategic Structural Health Management Plan, SSHMP.

Introduction

Amongst other things, the de Havilland Comet accidents in the 1950's [1] demonstrated the significance of metal fatigue and the importance of understanding the actual load spectra to which operational aircraft structures were subjected. Consequently, programs were initiated to measure the vertical accelerations (Nz_{cg}) experienced at an aircraft's centre-of-gravity. The reasons for selecting Nz_{cg} were primarily due to the limitations of the recording and sensor technology at the time (rudimentary electro-mechanical recorders) and the potential location of many of the structural problems that were liable to be encountered (the centre-wing and centre-fuselage). The maximum loads reacted by the centre-wing and centre-fuselage invariably necessitated the development of substantial structures whose stiffness allowed the structure in this area to be assumed rigid relative to the aircraft's centre-of-gravity. Therefore, relatively simple transfer functions that allowed Nz_{cg} to be related to the stresses at critical locations within the centre-fuselage/centre-wing area could be derived. This allowed stress spectra for the analytical and experimental validation of the cyclic integrity of the structure to be developed. While it was appreciated that the reduced stiffness/structural flexibility of outer-wing and empennage structures negated the rigid-body assumptions, thereby limiting the application of this technique to the centre-wing/fuselage, at the time this was not a concern as this was where most of the problems were being encountered. Over time, as wing and centre fuselage design benefitted from the data that was collected, other structural areas such as the outer wing and empennage started to become the critical life items. Unfortunately, determining the stress spectra in locations where flexibility effects could no longer be ignored required the application of more complex techniques.

As the understanding of the cyclic loading environment grew, a number of incidents demonstrated that a considerable amount of inherent scatter was associated with the prediction of fatigue and crack propagation in aircraft structures. Of particular concern was the very real possibility that cyclic damage might not be detected prior to the occurrence of catastrophic operational structural failures. This resulted in the regulatory approach to cyclic structural design and airworthiness criteria evolving from a safe-life approach (components can be replaced before any damage develops), through a fail-safe (redundant load paths) and damage tolerance approach (structure can function with damage/cracks that can be detected prior to becoming critical) to the multi-site, multi-element and Widespread Fatigue Damage (WFD) approach that is embodied in recent regulations [2].

While operational aircraft health monitoring programs were primarily pioneered by the fixed wing aircraft community, the rapidly reducing costs and increased capability of the technologies cited, plus the potential life-cycle management benefits of such programs led to their adaptation/adoption by the rotary wing (Health and Usage Monitoring Systems – HUMS) and engine communities. Additionally, ongoing usage monitoring programs for the purposes of ongoing training/feedback/fleet management (Flight Operations Quality Assurance – FOQA) and accident/incident investigation (Flight Data Monitoring – FDM) came into being as related but separate disciplines.

Concurrently, as structural design and airworthiness criteria were evolving, so too were other technologies that could potentially provide even greater insight into aircraft operational loads. Advances in and affordable computers, microprocessors, numerical simulation/analysis techniques, digital signal acquisition/processing and sensor technology made it possible to extensively instrument aircraft and acquire more data than it had ever been imagined was possible. Combining all these technologies resulted in a "tsunami" of data which it was assumed would provide considerable insight into actual aircraft usage and allow for the optimized management of critical components. The underlying mantra was that "more data is always beneficial" and that if the data could be measured "why not acquire it anyway, just in case …?"

Unfortunately, even though monitoring programs would intuitively be expected to realize numerous safety, operational and aircraft maintenance benefits, this has often not proved to be the case. Many organizations have ended up investing a considerable amount of resources in programs which have either failed to realize their full potential [3], [4] and/or to detect "unanticipated" or "unexpected" structural problems in a timely manner. Such failures can have catastrophic economic and operational consequences and, in some instances, lead to loss of life [5]. Consequently, the tangible benefits of implementing these types of programs are being increasingly questioned. For new aircraft acquisitions, this generally translates into health monitoring programs being relegated to "optional extras" which are rarely exercised due to cost or other factors. In some instances, even when a health monitoring capability exists on an aircraft fleet, the data is just stockpiled due to the organization having no idea as to its significance or what to do with it or because the reliability and utility of the recorded data itself is in doubt [6].

This paper attempts to provide some insights as to why, despite the potential of these programs to provide greater insight and economic and safety benefits, many aircraft health management programs fail to provide few if any tangible Returns-on-Investment (ROI); particularly with regard to individual or aircraft fleet life-cycle management. Based on the insights provided, the paper proposes ways in which tangible ROI can be obtained, thereby translating the implementation of successful programs from myth to reality!

Background Information

Before discussing some of the reasons why aircraft health monitoring programs can fail to generate tangible LCM/ROI, it is appropriate to provide some background discussion to facilitate a better understanding of the rationale for the subsequent comments/conclusions that are presented in this paper.

First, it is important to realise what is actually being monitored and the associated implications/limitations related to the information gathered. While ensuring that the instrumentation itself is gathering valid data is by no means trivial, this paper will primarily focus on the importance of understanding the structural implications of the data that is being acquired and analysed.

The static loading capability of a structure defines the point at which the structure will fail when it is subjected to an increasing load. To avoid in-service failures, factors of safety are applied to these loads. Hence there is the concept of Design Limit Load (the maximum load an aircraft or critical component is anticipated to see once in its lifetime). Design Limit Load (DLL) occurrence should not cause any permanent damage or deformation to a structure. Design Ultimate Load (DUL) is 1.5 x DLL and is the load which the aircraft is designed to sustain without catastrophic failure although in all likelihood the structure will incur, permanent structural damage/deformation.

Any flaw or crack in a structure that will degrade the ultimate strength of a structure. The degraded strength of the structure in the presence of a crack or flaw is termed its residual strength. Current airworthiness requirements, such as Reference [7] and [8], require that inspection and maintenance procedures be implemented to ensure that the residual strength of a structure (i.e.: its static strength in the presence of a crack or flaw) will not fall below the required DUL capacity of the structure prior to the detection of any existing crack or flaw. It is this type of requirement that often dictates the structural limits of an aircraft's flight envelope.

The occurrence of high loads and damage as a result of harsh or unusual usage is generally something that can readily be detected (felt) by aircrew, is fairly evident in any recorded data (such as might be captured by a FOQA program) and/or may be visible to the naked eye. As such load exceedances may have immediate airworthiness implications (e.g.: the need to implement an overstress inspection/maintenance procedure before the aircraft flies again) one important aspect of any aircraft monitoring program is the implementation of an effective Harsh or Unusual Usage Response Program (HUURP) [9]. This ensures that critical overload occurrences are rapidly detected and addressed in a timely manner.

Conversely, cyclic loading capability is a cumulative process where the finite cyclic loading capacity (life) of a structure is consumed over time. While the occurrences of individual high load cycles are detrimental to the cyclic life of a structure, so too are the many smaller load cycles to which a structure is subjected over its operational lifetime. In other words, the cyclic life of a structure can be consumed by a few large load cycles, multiple small load cycles or any combination thereof. Unlike, large load excursions, the immediate or cumulative impact of smaller load cycles may not be that intuitive to deduce or easy to detect. Continuous variations in the frequency and magnitude of load cycles due to such factors as mission type, prevailing weather, the season of operation, the terrain over which the aircraft is operating (e.g.: mountainous, flat plains or coastal), the point in the sky at which they occurred (e.g.: airspeed, altitude, orientation) and, last but by no means least, pilot input all need to be considered. For aircraft fleets that consistently fly the same missions within their original design intent at higher altitudes using a variety of crews, it is possible to develop a reasonable fleet spectrum over time. This can be used to derive or modify fleet-wide inspection and maintenance intervals so that they are tailored to actual usage without becoming overly conservative or onerous. It is also marginally viable, although not recommended, to maintain/inspect components based on rudimentary tracking parameters such as flight hours, or cycles (pressurization, take-off and landings etc.). However, for aircraft such as those which operate in special mission, predominantly low-level, roles such as aerial firefighting, Search and Rescue, interdiction, geophysical survey, crop spraying etc. the concept of using a fleet-wide spectrum for inspection and maintenance is fraught with danger and often results in catastrophic, unforseen, operational failures [5] [10].

For aircraft operating within their design intent, health monitoring programs are generally implemented to confirm their usage, extend inspection/maintenance intervals and/or extend the operational lives of the aircraft (the design assumptions are generally, although not always, conservative). Alternately, for aircraft operating outside their design intent in special mission roles, a monitoring program based on Individual Aircraft Tracking (IAT) is invariably required. IAT programs serve to confirm the adequacy of the original design inspection/maintenance intervals for the role in which the aircraft are actually operating thereby facilitating the ongoing structural health evaluation of individual fleet aircraft on an ongoing basis. Some operators seek to reduce the scope of an IAT program by attempting to base fleet inspections/maintenance cycles on the "worst usage of an individual aircraft". The assumption is that this will provide an inspection/maintenance regime which is overly conservative. However this may not actually be the case if the missions in which the aircraft are used result in the presence of crack-retardation at critical structural locations [11].

Organizations can fall into the trap of thinking that the aircraft they are using need little or any monitoring as they were designed using "conservative load criteria". Invariably it turns out that "conservative load criteria" relate to experimental and/or analytical validation of the aircraft to sustain high static loads (often far greater than those anticipated in service). The implicit assumption is that the validation of a conservative (high) static load capability automatically ensures/validates a conservative cyclic load capability. Unfortunately, as has been demonstrated repeatedly throughout aviation history, nothing could be further from the truth [1] [10].

Second, in recent years the ongoing monitoring of aircraft structures has been complicated by the introduction of components manufactured from composite materials. While composite components offer many weight and performance advantages, they also introduce a number of new monitoring challenges. These are primarily related to the capability of a designer to design both the underlying material characteristics (lay-up) and geometry of composite components. Consequently, for reasons discussed in Reference [12], sensor and monitoring technologies that are more than adequate for metal structures may no longer be adequate for composite structures. Failure to account for these differences can lead to an inadequate number of channels being assigned to accurately characterize the response of a composite structure and the data that is obtained being totally misinterpreted.

Finally, as alluded to earlier, different groups within the larger aerospace community have introduced a plethora of aircraft health monitoring initiatives to ostensibly address their "specific and unique needs". A list of the more common types of program that are commonly used at the time of writing is provided in Table 1.

Table 1: Different Health Monitoring Programs

Acronym	Definition
HUMS	Health and Usage Monitoring System
SHM	Structural Health Monitoring
OLM	Operational Loads Monitoring
LE/SS	Loads Environment Stress (or Strain) Survey
IAT	Individual Aircraft Tracking
FOQA	Flight Operations Quality Assurance
MFOQA	Military Flight Operations Quality Assurance
IVHM	Integrated Vehicle Health Monitoring
CBM	Condition Based Monitoring (Maintenance)
FDM	Flight Data Monitoring

A recent on-line survey commissioned by Celeris Aerospace [3] probed participants involved in the programs listed in Table 1 about their understanding of the purpose of these programs. Even though the number of respondents was relatively small, the experience and calibre of the respondents (based on responses to contextual questions) provided both insightful and sobering responses.

Over 85% of the respondents indicated that they found the different program terminology to be confusing with over 54% indicating that they considered the programs to have essentially the same or similar objectives. Furthermore, around 22% of the respondents felt that there was **significant** overlap/duplication between the different programs, with a further 58% indicating that they considered there to be **some** overlap/duplication between the programs. Only forty-five percent (45%) of the participants indicated that the program objectives were well defined and communicated. This probably goes someway to explaining why only 37% of the respondents considered that the benefits they expected from the programs listed in Table 1 were actually realized!

Given the apparent confusion amongst program practitioners and users with regards to the function and purpose of the programs themselves, it is not surprising that aircraft health monitoring programs often fail to yield tangible LCM/ROI benefits.

Aircraft Health Monitoring Context

Traditionally, aircraft health monitoring programs have often been considered to be primarily engineering programs with possible life-cycle management benefits. Consequently, the majority of program resources are invariably focussed on ensuring that the engineering aspects of the program are addressed. Unfortunately, merely regarding these programs as solely engineering endeavours significantly increases the risk that the programs will either fail or fall significantly short of producing any tangible LCM/ROI. Unless these programs are placed in the context of overall fleet LCM objectives from the outset, it will be almost impossible to realise tangible LCM/ROI. This concept is illustrated schematically in Figure 1.

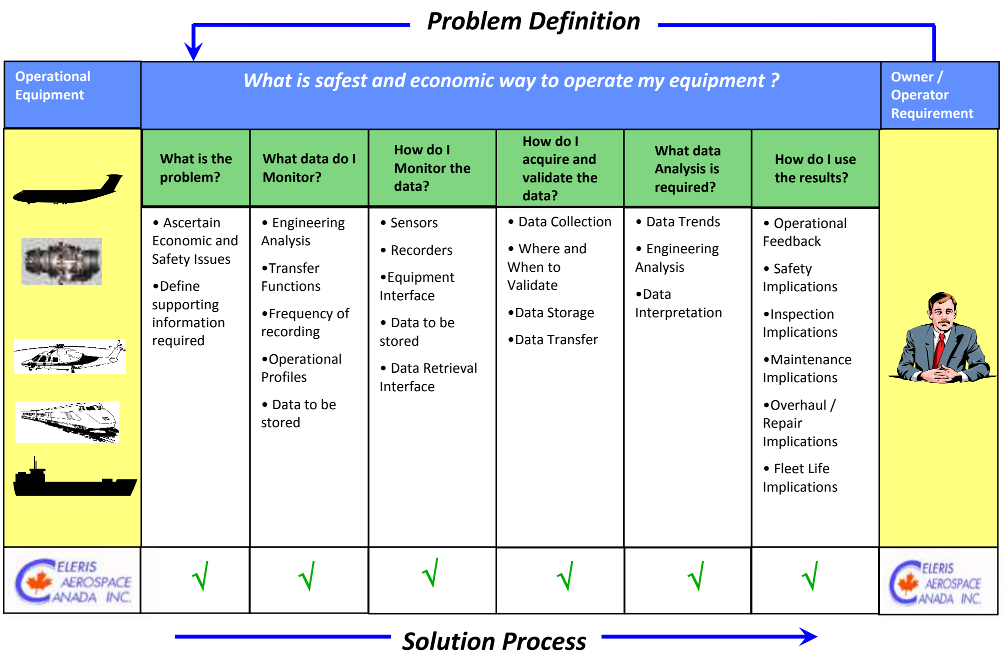

Figure 1: Aircraft Health Monitoring Program Context

As illustrated in Figure 1, to realize tangible LCM/ROI the scope/definition of any aircraft health monitoring program should at all times be focused on fleet life-cycle management requirements; not the other way around. While the engineering, instrumentation, data-management and analysis aspects are important tools that facilitate the achievement of fleet life-cycle management objectives, they are just tools and should never become the focus of the program. From a program perspective a proposed requirement for any engineering, instrumentation, data-management or analysis capability that cannot be directly related-back to a life-cycle management objective should be critically evaluated and in most cases discarded.

In reality, during their original or extended life-cycle, the usage of aviation assets will often be modified to meet changing operational needs and/or address new or modified roles. Different roles can result in the life of critical structures being consumed at an accelerated rate and/or the emergence of new critical structural locations not previously encountered. Consequently, aircraft health monitoring programs invariably require modification. Celeris Aerospace's approach to managing changing LCM requirements is defined and implemented through the use of the Strategic Structural Health Management Plan (SSHMP) Methodology. An SSHMP provides an integrated cradle (or current status) to grave approach for ensuring that aircraft meet their LCM goals, even though those goals may change over time. An important aspect, of this methodology, is the use of an appropriately tailored aircraft health monitoring plan that is compliant with aircraft fleet LCM objectives. This process is dynamic and ongoing in nature and requires periodic reviews of both the SSHMP and the monitoring system (typically on an annual basis). Further information related to SSHMPs can be found in Reference [13].

The Myths

Detailed accounts of many of the practical issues that need to be addressed to implement a well-structured aircraft health monitoring program can be found in References [14] and [15]. The remainder of this paper will focus on some of the commonly held myths about aircraft structural health monitoring programs that can contribute to their failure to generate tangible LCM/ROI.

While intuitively it would be anticipated that there are many life-cycle management and engineering benefits to be derived from monitoring an aircraft on an ongoing basis, what is often not fully appreciated is the need to develop and implement a structured, integrated and systematic approach to the program from the outset. This is critical if any of the potential benefits with these programs are to be realised. Aside of the logistical benefits of implementing such an approach, there is also a need to ensure that the multi-disciplinary and sometimes conflicting requirements of the different facets of a program (e.g.: sensors/instrumentation, data acquisition, validation, analysis and management, system airworthiness, LCM/Reliability Centred Maintenance (RCM)) are addressed and reconciled with overall fleet life-cycle management objectives; as defined by an SSHMP. Unfortunately, in an effort to “acquire data as soon as possible”; critical steps associated with the implementation of a systematic approach are often overlooked. This invariably has disastrous consequences with regards to both the integrity/relevance of the data that is obtained and the realization of tangible LCM/ROI benefits. Some of the more common rationalizations/misunderstandings (myths) associated with not implementing a systematic approach are:

1. **Myth # 1 - “Programs developed for the same aircraft type by other organizations can be applied directly to our aircraft (“There is no need to re-invent the wheel”).**

   It is important to realize that the loading experienced by identical aircraft is a function of both geometry/configuration and role/mission. Two identically configured aircraft flying in different roles/missions will often experience different loading. This can introduce significant variations in the cyclic lives of identical critical components. Harbingers of significant loading variations between identically configured aircraft include the acceleration/deceleration of life consumption in known critical structural locations and/or the introduction of new critical structural locations altogether! This type of problem will most frequently manifest itself in multi-purpose fleets of identically configured aircraft. For example consider the H-60 Helicopter. This helicopter has been sold around the world and in North America is used by armies, navies and air forces. A number of agencies have implemented Health and Usage Monitoring System (HUMS) programs on this platform with some success [16]. However, the way a coast guard utilises the aircraft tends to be significantly different from the way an army utilises the aircraft. Therefore, it is unlikely that the identical configuration of an aircraft health monitoring system developed to monitor an army H-60 fleet application would be well suited to monitor the coast guard H-60 fleet application and vice-versa. While, by virtue of being installed on sensibly similarly configured aircraft, both systems may contain a number of common elements, the final configuration used for each fleet application may need to differ substantially to support the different usage and life-cycle management objectives of each fleet.

2. **Myth # 2 - “We have a limited budget. Let’s just focus on getting some equipment on the aircraft so we can show people data”.**

   One of the most important aspects of collecting data is ensuring that in addition to capturing data of interest, data that will provide a context for its correct interpretation is also captured. A simple example that can illustrate the need for context is that of capturing the vertical acceleration experienced by an aircraft within its centre-of-gravity limits (Nz_{cg}).

Nz_{cg} is a parameter that has been collected for many years due to it being relatively straightforward to establish transfer functions relating this parameter to stresses at critical structural locations within the centre-wing and centre-fuselage (where it can be assumed the structure is sensibly rigid relative to the c.g.). However, to perform a fatigue or crack-growth analysis at any given critical location, it is necessary to derive a stress spectrum based on the recorded values of Nz_{cg}. A given value of Nz_{cg} may correspond to multiple values of stress at a given critical structural location depending on how it was configured and what it was doing at the time of the Nz_{cg} occurrence (i.e.: there is not always a one-to-one relationship between stress and Nz_{cg}). For example, to determine the stress at a given structural location from Nz_{cg}, it is necessary to know the weight of the aircraft at the time the given value of Nz_{cg} occurred. Hence weight becomes a context parameter that allows the analyst to determine the stress at a critical structural location for a given value of Nz_{cg}.

If an aircraft was only subject to level flight conditions, a combination of stress and weight would be adequate for developing a spectrum for critical locations in either wing. However an aircraft needs to turn and, in completing a coordinated turn, also needs to roll. This results in the lift on the wing inside the turning circle being reduced, while the wing on the outside of the turning circle being increased, i.e.: the wing loading is no longer symmetric between the left and right wing. Consequently, to interpret the structural significance of Nz_{cg} at a critical wing location it becomes necessary to know not only the weight but the lateral attitude of the aircraft at the time the value of Nz_{cg} was recorded. This process can be continued ad-infinitum. For example if the coordinated turn takes place while the aircraft on its final approach circuit prior to landing, the flaps may be deployed at one of three or four different angles. This too will impact the stress at critical structural locations and hence some way of acquiring the flap position when the value of Nz_{cg} was recorded will be required.

In a practical situation, such a process would not be continued ad-infinitum as a more cost-effective solution would be to locally measure, as opposed to globally infer, the value of the stress at the critical location with the use of a dedicated sensor such as a strain gauge; even though this may introduce some additional challenges. However, the cited example serves to illustrate that if recorded data is to be used as a means of obtaining tangible LCM/ROI, it is important that it can be related to aircraft life-cycle management objectives and correctly and meaningfully interpreted as a result of appropriate contextual data being gathered. Merely just collecting data without a prior understanding of its potential usage and significance will significantly increase the risk that the tangible LCM/ROI will not be obtained from a health monitoring program.

3. Myth # 3 - "Data acquisition hardware and software is cheap; let's monitor everything".

The implicit underlying assumption in the "let's monitor everything" approach is that when there is little or no understanding of what to measure; measure everything as "more data is always better". This approach is also assumed to provide a "fail-safe" approach as any data that is not well understood can always be reviewed at a later date when a better understanding of the loads that are actually being experienced is obtained. There are several potential pitfalls associated with this type of approach. These include, but are not necessarily limited to:

a. Assuming that a correct understanding of an aircraft's operational loads environment and its interpretation can be obtained by merely collecting more data.

If little or no effort has been made to understand how critical structure will function in the anticipated loads environment, it is going to be difficult to both validate and interpret the data that is obtained;

b. Considering that all the potential data to which there is access is adequate to characterize the structural response of critical structural locations. Once again, if there is little to no understanding as to how the structure might respond, how can the adequacy of the parameters selected to truly characterize the operational loads environment be evaluated?; and

c. Obscuring significant data trends in the large quantities of data obtained and assuming that data correlation alone can be used to identify and extract structurally significant causal factors (See discussion on Myth # 10).

From a cost perspective, over the life-cycle of an aircraft structural health monitoring program the major costs are not related to the data acquisition equipment, or supporting software, but rather the ongoing acquisition, validation, interpretation, dissemination and management of the data itself [13]. Careful consideration should be given to eliminating the recording of any data which cannot be directly related to aircraft fleet life-cycle management goals.

4. Myth # 4 – "The aim of any aircraft health monitoring program is to support a "Condition-Based Maintenance".

There are two parts to Condition Based Maintenance (CBM) (On-Condition Maintenance) programs; the engineering analysis and the operational/maintenance analysis. Unfortunately, many CBM programs implicitly assume that all the benefits and associated cost savings identified during the engineering analysis can be realized. Based on this assumption, projected cost-savings are generated which invariably prove to be overly optimistic. This is due to a failure to evaluate the potential savings identified by an engineering analysis against practical operational and maintenance constraints. In fact, in some circumstances, if all the potential savings identified in the engineering analysis were implemented carte-blanche, it is possible that the overall cost of operating and maintaining the aircraft could increase as opposed to decrease!

The engineering analysis of CBM involves recording the actual operational loads environment in which the aircraft operates. This data is used to generate the input spectra for fatigue and damage tolerance analysis/tests which determine critical component safe lives (where damage tolerance approaches are not viable) or required component inspection and maintenance intervals. For aircraft operating either in their design-intent roles, or roles that are less severe than assumed during the original aircraft design, critical component replacement or inspection and maintenance intervals can often be extended and life-extension contemplated. Conversely, in instances where the aircraft are operating in more severe loads environments than anticipated during their original design, critical component replacement or inspection and maintenance intervals may need to be reduced and/or the aircraft retired/replaced prematurely.

Where the engineering analysis indicates that there is potential for increased critical component replacement intervals, inspection and maintenance intervals or component life extension, it is important that the viability of such an approach includes a review of both the engineering maintenance/operational constraints. For example, there is a limited amount of hanger space and tooling that can be allocated to the inspection, maintenance and overhaul of an aircraft fleet. Hence it may prove to be impractical to implement customized schedules based on actual usage for every aircraft in the fleet. Similarly, while the engineering analysis may indicate that the life of a critical component can be extended several hundred hours, the maintenance schedule and associated level of structural tear-down may be such that unless the component can be extended by one complete maintenance cycle, it would not be economically viable to take advantage of the potential additional life derived from the engineering analysis. This is a result of the major cost associated with the maintenance or overhaul being the tear-down costs as opposed to the cost of the component itself (analogous to replacing a low-cost transmission shaft seal on the gear box of a car while the gear box is removed, even though the existing seal still has some useful remaining life). Other cost factors also come into play before a true assessment of whether it is economically viable to realise the potential inspection/maintenance interval or life extension predicted from the engineering analysis. For example, an engineering analysis may indicate the potential to extend the inspection/maintenance interval of a component by several hundred hours. However, in so doing, the required inspection/maintenance activity would need to be undertaken in an operational inspection/maintenance environment as opposed to a depot level environment. If the inspection requires specialized equipment and specialized operator training, the costs associated with equipping each operational unit with the specialized equipment and trained operators needs to be compared to the cost of continuing to undertake the inspection at depot level where only one piece of specialized equipment and one or two trained operators may be required.

While moving towards Condition Based Maintenance certainly has merit, it is important to remember that the engineering analysis only identifies potential LCM benefits. Before these benefits can be realized and tangible LCM/ROI obtained, their implementation costs have to be evaluated in the light of practical operational and maintenance constraints. This is illustrated in Figure 2 which attempts to define the contribution towards achieving tangible LCM/ROI made by the different monitoring programs defined in Table 1.

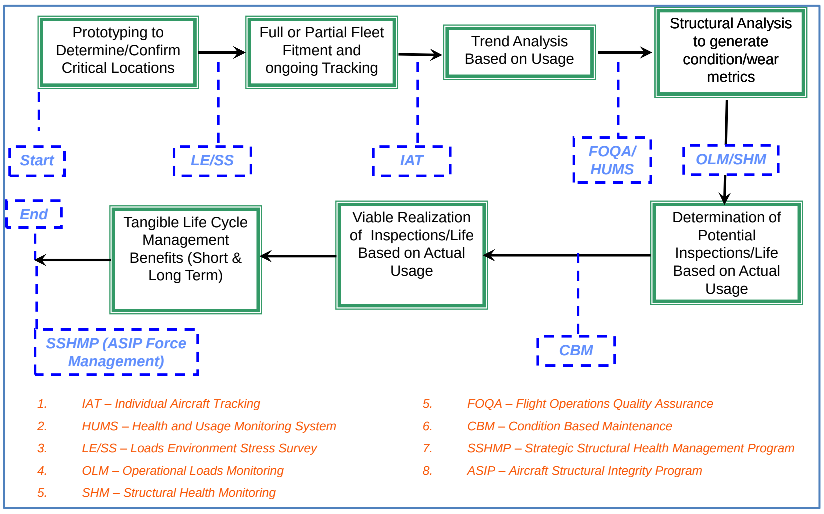

Figure 2: Monitoring Programs Contribution to Tangible LCM/ROI Benefits

5. Myth # 5 – "Operators have no need to understand aircraft health monitoring programs; they just need to follow OEM/Regulatory guidelines and direction".

As a result of many operational loads monitoring programs apparently contributing little or any tangible LCM/ROI, from an operational standpoint they tend to be regarded as "engineering or airworthiness" programs with little operational utility or benefit. While broadly being aware of their intent operational staff often consider them to be yet another program that "has to be managed by scarce and often overworked front-line resources". Consequently, at the operational level, there is limited incentive to go much beyond providing the minimal amount of required data/support with little or no operational resources being allocated to truly understanding the significance of the data that is collected. In some instances so little operational benefit has been derived from aircraft health monitoring programs that the data being recorded is either being stockpiled without examination or "optional" health monitoring systems are not installed on aircraft fleets.

There are three reasons why operators need to have more than a cursory in-house understanding of the purposes of a structural health monitoring program. These are:

- <u>Enhanced Safety:</u>

  Health monitoring systems and sensors can alert operators to potential safety problems providing the significance of the data being received is understood. Little or no understanding of the data at the operational level can result in harbingers of an impending problem not being fully appreciated with disastrous consequences [5] [17]. A front-line operational understanding of the significance of the data being presented is essential as it maximizes the possibility of timely intervention, thereby minimizing the chances of catastrophic operational failures.

- <u>Understanding the source of damage:</u>

  Aircraft structural health monitoring programs can assist operators to understand which missions/roles, or segments of missions/roles, are the most structurally damaging [10]. Once this has been established, there are a number of tangible LCM/ROI benefits that can be realized. The first benefit comes from evaluating whether modifications to the way the mission/role is flown can alleviate the incremental structural damage that is sustained. In some instances, relatively minor modifications to the way a mission/role is flown can significantly reduce the incremental structural damage sustained during a mission thereby reducing component life consumption.

  Accounting for different mission severities is often addressed through the development of Mission Severity Factors (MSF). Based on measured data and its subsequent analysis, the damage sustained for different missions is computed and compared to a baseline mission (typically the design-intent mission as that is what all the OEM inspection, maintenance and overhaul limits are based on). The inspection, maintenance and overhaul intervals are then scheduled based on Equivalent Baseline Hours (EBH) where EBH are equal to the actual hours flown multiplied by the MSF. While this approach can provide a reasonable initial front-line fleet management tool, it needs to be applied with caution as the data obtained can easily be misapplied or misinterpreted [18]. Once again this illustrates the need for more than a cursory understanding of health monitoring programs at the operational level.

- <u>Implementing proactive fleet management to maximize fleet operational life.</u>

  Once the most damaging missions/roles have been identified, it becomes feasible to implement pro-active fleet management techniques such as fleet rotation to optimize overall fleet operational life. For example, where practicable rotating fleet aircraft in and out of the more severe missions/roles for finite periods of time will result in the attainment of a greater operational fleet life than will dedicating one aircraft to the more severe missions/roles until its life has been expended and then replacing it with another fleet aircraft.

6. Myth # 6 - Advances in hardware and software technology will ensure program success.

While advances in hardware and software have certainly improved the capability and insights provided by health monitoring programs over the years, the availability of advanced technologies do not negate the need to understand the technical and operational environments in which the technologies will be used. Although advanced technologies can be useful in providing insight that facilitates a better understanding of well-defined and structured problems, they tend to do little with regards to providing insight into poorly defined problems. One disturbing trend that has become evident in recent years is the tendency to focus research and effort on developing or refining solutions for areas that are not the weakest link in the chain.

A good example of where interesting and challenging research is being conducted in an area that is not necessarily the weakest link in the chain can be found in the rotary wing environment. For rotary wing aircraft the most critical components are invariably the rotating components, shafts and gears associated with the gear-boxes and transmission components. While some success has been obtained through detecting changes in vibratory signature patterns (i.e.: base-lining a signature for specific operating conditions and then periodically reviewing the signature to see if there are any significant deviations), relatively little success has been obtained with regard to quantifying the ongoing health of critical mechanical rotating components. From an ongoing health monitoring perspective the ongoing health evaluation of rotating components is very challenging. They are subject to many high-cycle vibratory inputs and operate in environments where the durability of sensors that might be installed for direct monitoring is often tenuous (e.g.: oil being cycled over critical gears and shafts in a transmission unit). Consequently, parametric and/or indirect monitoring techniques often have to be adopted. This often results in signals related to critical components proving to be extremely challenging to detect and synthesize. In an attempt to address this situation, the rotary wing community has developed Condition Indicators (CIs) which attempt to determine the health of critical components through indirect (inferred) measurements. CI formulations can be relatively straightforward or complex equations comprised of readily available flight parameters. While the absolute value of a CI itself may be difficult to physically evaluate, some success has been obtained through the trend monitoring of these values, i.e.: providing the CIs stay within a certain range(s), the integrity of the components they are monitoring are deemed to be acceptable. Conversely, trends that stray outside acceptable limits may be indicative of a problem which requires some immediate intervention. Depending on the component being monitored the intervention may range from a simple inspection through to the removal and replacement of a complete transmission component (e.g.: a gearbox). Original Equipment Manufacturers (OEMs), and to some extent regulatory agencies, devote a lot of time to developing and validating appropriate CI formulations and ranges for critical components These formulations, which are invariably considered proprietary, are often incorporated into an on-board HUMS box. As with any health monitoring indicator, success is measured by the predictive reliability and consistency of the indicator with particular emphasis being placed on minimizing the number of false positives.

While the concept of CIs is theoretically appealing for a variety of technical and practical reasons, they are not as successful as one might hope in predicting impending problems and failures. In an attempt to improve the reliability of CIs, programs are being aggressively pursued both in Europe (Advanced Anomaly Detection (AAD) [19] [20]), and North America (Normalized Power (NP)). Essentially, both these techniques seek to refine the analytical techniques used to trend CI data. While there may be merit to pursuing such an approach, currently it is not the trending of the CIs that is the weakest-link in the process, but rather the inherent instability of the underlying CIs themselves. Until more robust methods of improving the reliability of the CIs themselves, or more reliable alternate health monitoring indicators are developed, there is relatively little to be gained from advanced trending of the underlying data which is inherently unstable and/or unreliable. Unless the reliability of the underlying data can be addressed, merely refining data trending capabilities only serves to create a false-sense of security that is not particularly helpful[17] [21] [22].

7. Myth # 7 - Prognostic technologies will improve the reliability and accuracy of health monitoring programs".

In recent years, much emphasis has been placed on moving from reactive (deterministic) health monitoring (i.e.: detecting when a critical component is close to failure or has failed) to prognostic (proactive) health monitoring (i.e.: predicting when failure of a critical component will occur so that timely intervention can be scheduled prior to component failure). While the introduction of any reliable prognostic capability is beneficial, an understanding of the bases on which the proactive predictions are based is essential if results that yield tangible LCM/ROI are to be obtained.

It is important to appreciate that although this technology is termed prognostic; there is still a deterministic element to it. For example, prognostic analysis assumes that the current state (health) of a structure and how it will be used in the future is known. This translates into fusing results from Non-Destructive Inspections (NDI), purpose built sensors and recorded health monitoring data to define how aircraft have been used in the past and to predict critical component(s) responses and damage accumulation for the future. An implicit, but significant assumption that is encapsulated in such an approach is that the usage experienced in the past will adequately characterise future usage. This might not always be the case for two reasons. First current usage might not be representative of future flying [18]. Second the usage might change over time, even though the aircraft is nominally kept in the same role [9]. To ensure the accuracy of prognostic predictions constant vigilance of actual usage is required to ensure it is consistent with the anticipated usage on which the prognostic predictions have been based.

8. Myth # 8 – The adoption of "physics-based" models will improve the accuracy of health monitoring analysis.

There is a lot of inherent scatter associated with fatigue and durability analysis. This scatter can be introduced as a result of the quality/processing of the materials data itself (e.g.: inclusions at the micro-mechanical level), manufacturing flaws (e.g.: the so-called rogue flaw) at the sub-component or component level and/or an unrealistic assessment/characterisation of the usage spectra assumed during both test and analysis. Traditionally, this scatter has been dealt with through the development of semi-empirical analytical methods that are calibrated through extensive test programs [11]. In conservatively trying to accommodate such scatter, it is recognized that such an approach can result in the premature inspection or removal of critical components.

In an attempt to better optimize the design and analysis process, significant effort has been expended in recent years on developing “physics-based models” to replace many of the previously developed semi-empirical models. The basic premise behind this effort is that by accurately modelling the damage process at the micro-mechanical level and simulating its propagation into the macro-mechanical level a more accurate analysis of the response of critical structural components can be obtained. Based on this premise, it is hypothesized that the conservatism required by semi-empirical methods can be reduced or eliminated.

From the perspective of improving the overall understanding of the damage propagation process such research is commendable and should be encouraged, as it can be used to refine the semi-empirical models previously mentioned. However, from the perspective of improving the tangible LCM/ROI that is realized from aircraft monitoring programs, practical constraints may limit the benefits that can be realized by such endeavours. For example suppose a micro-mechanical model that is able to successfully model crack-growth propagation at the grain boundary level is developed. To provide an accurate assessment, it is necessary to interface such models at a large number of locations in the macro-structure models so that the predicted damage can be propagated appropriately and predict failure. When considering whether such an approach is warranted, it is necessary to quantify the LCM/ROI benefits that will be gained by adopting a physics based modelling approach as opposed to a semi-empirical approach. The physics-based modelling approach will probably be more costly to implement as:

- The analytical and experimentally derived parameters needed to generate an accurate model are liable to require far more sophisticated equipment and operator expertize than do those used for semi-empirical models;

- The time required to generate run and validate the analytical models themselves, especially in the context of a complete structure, will probably be significantly greater than would be required for a semi-empirical analysis; and

- Once the micro models have been used to generate the initial crack propagation models, they are going to be interfaced with macro models (e.g.: Finite Element Models) which are generally using smeared (average) material and thickness properties, which may well negate the “additional accuracy” obtained from the micro models.

The purpose of the preceding discussion is not to argue the relative merits of a physics-based versus semi-empirical modelling one way or another, but rather to suggest that whatever approach is adopted needs to be placed in the context of the resulting LCM/ROI benefits. For instance, if the adoption of a physics-based approach resulted in an improvement in achievable life limits or inspection/maintenance intervals (see discussion on Myth # 4) of 2% but at an increased infrastructure/support cost to implement the method of 30%, it may well not be economically viable. Conversely, if the physics based approach resulted in an improvement in realisable potential life limits or inspection/maintenance intervals of 30% for an increased infrastructure/support cost to implement the method of 30% it may prove to be economically viable. Each potential application needs to be evaluated on a case-by-case basis.

9. Myth # 9 – Data validation can be accompanied solely through statistical analysis.

When processing large quantities of data, statistical techniques are often used to identify outlier data on the assumption that any data isolated using these tools is either erroneous or invalid. The general approach is to remove this data from any data-set that will be used for the purposes of spectrum development. However there is another possible explanation for the existence of outlier data which, if overlooked, can result in valid data being discarded. This explanation relates to the implicit weighting of combined data-sets versus individual data-sets and can be illustrated using a simple example. Consider a fleet of twenty aircraft consistently gathering 40 hours of data per month, twelve months of the year with the data being downloaded on a monthly basis. The amount of accumulated data is summarized in Table 2.

Table 2: Data Accumulated by a Fleet of Aircraft Over Time

Number of aircraft in fleet =	20
Hours flown/month =	40
Annual Hours per aircraft =	480

Year	**Fleet Hours**		**Percentage contribution of Individual aircraft data**	
	Annual	**Accumulated**	**Monthly**	**Annually**
1	9,600	9,600	0.417%	5.000%
2	9,600	19,200	0.208%	2.500%
3	9,600	28,800	0.139%	1.667%
4	9,600	38,400	0.104%	1.250%
5	9,600	48,000	0.083%	1.000%
6	9,600	57,600	0.069%	0.833%
7	9,600	67,200	0.060%	0.714%
8	9,600	76,800	0.052%	0.625%
9	9,600	86,400	0.046%	0.556%
10	9,600	96,000	0.042%	0.500%

A common approach to identifying outliers is to compare the trends in a given role/mission for an individual aircraft data-set against the cumulative data for the fleet. Even though in the example provided the quantity of data gathered by an individual aircraft on a monthly basis is constant, its percentage contribution to the overall fleet accumulated data on both a monthly and annual basis diminishes with time as illustrated in Table 2. Consequently, the influence of the cumulated fleet data-set relative to the individual data-sets becomes increasingly dominant (weighted) over time. Over time, any comparative trend analysis will automatically reject any data that differs from the accumulated fleet trends. If there is a rational explanation regarding why data has been identified as anomalous; e.g.: a faulty recorder or inadequate EMI/EMC shielding such rejection may be appropriate. However, if there is no apparent physical explanation as to why the data is anomalous, it could be that the data is accurate and what is happening is that either overall fleet usage is gradually changing over time or one or more aircraft are being used differently (even though their nominal role/mission remains unchanged). In such instances data rejection is inappropriate and hence methods that are capable of distinguishing between erroneous data and changes in aircraft historic usage need to be available. Once changes in usage have been identified, data preceding the changes needs to be frozen as historic data and the modified/changed mission data accumulated in a separate data-set. Any subsequent durability and damage tolerance analysis needs to be based on the historic data plus the data being accumulated in the modified/changed mission or role. Constant vigilance of the data and a good understanding of its significance are critical if the inadvertent rejection of perfectly valid data is to be avoided; particularly as this data may be a harbinger of pending critical structural issues. Maintaining such vigilance in a fleet monitoring environment can prove to be challenging particularly as the large quantities of data involved often result in the automated rejection of anomalous/outlier data. It is imperative that any potential outlier/anomalous data is isolated and investigated in a timely manner until the causal factors for the anomalies are identified.

Warnings that data is being inadvertently rejected due to changing usage include, but are not necessarily limited to:

- Increasing quantities of previously validated data being rejected; and

- Constant rejection of data from one or more aircraft within a fleet, even though physical reason(s) for the data rejection (e.g.: faulty sensor, faulty recorder, extraneous EMI/EMC emissions) cannot be isolated.

10. Myth 10 - "Data management and analysis can easily be done using spreadsheets; we will deal with the data once the engineering is sorted out".

When contemplating the implementation of a health monitoring program, many organizations first implement a prototype program on one or two aircraft and analyse the data with a desktop spreadsheet program using engineers who are intimately involved with the program. This prototype program is often coupled with flight-playback capabilities that have been either supplied with the recording system or independently acquired flight play-back facilities that can read exported data from the recorder in a commonly used output format (e.g.: CSV, XLS, TXT). Individual flights are examined in detail by knowledgeable individuals and interesting or significant trends isolated by interactively comparing the engineering parameters with the flight playback visuals. Unfortunately, based on this type of analysis and/or "demonstration of capability", the data management aspects of a health monitoring program are assumed to be addressed and little thought is given with regards to how data from a full-blown fleet implementation will be acquired, validated, transferred, analysed, interpreted, stored and managed. While there is nothing wrong with prototyping a concept, it should be appreciated that the data management and analysis aspects of a fleet implementation are far from trivial and failure to consider them concurrently with the engineering aspects from the outset of the program can jeopardize the program generating any tangible LCM/ROI.

The reasons for this are primarily attributable to a data volume/management effect. There is a big difference between obtaining and analysing data-sets from one or two aircraft on a flight-by-flight basis and getting data from two or three flights a day, of varying lengths from a number of geographically dispersed aircraft fleet(s) [23].

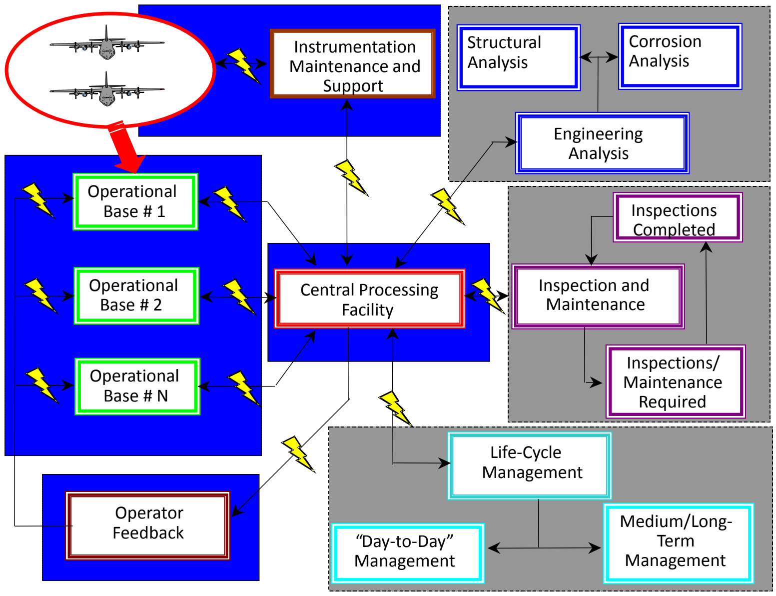

Figure 3: Overview of Typical Fleet Data Acquisition Process

Figure 3 provides an overview of a fairly typical health monitoring scenario. One or more aircraft types can be located at multiple, geographically dispersed, operational units where the data is downloaded and transmitted to a central processing facility. The downloading/transfer will invariably be overseen by operational maintenance personnel, with varying levels of computer skills for whom this task is a secondary duty. Therefore the data acquisition, initial validation and identification of any potential immediate safety of flight/airworthiness issues need to be as autonomous as possible (see Section on Background Information). Based on the data obtained, timely feedback related to usage and the health of the monitoring systems need to be provided to the operational units. Depending on the nature of the usage feedback and system status information required follow-up actions may need to be initiated and tracked until they are resolved. From an analysis, inspection and LCM perspective, the raw data itself needs to be accessed, processed and analysed to facilitate the timely management of the aviation assets in a safe and economic manner. In many cases, meaningful interpretation of data from the health monitoring system requires it to be integrated with data obtained from other sources such as aircraft log-books, aircraft configuration management systems (defining which critical components are associated with which aircraft tail number at any given point in time as typically a number of components can be exchanged between aircraft), and inspection/maintenance systems to name but a few. Similarly the output from the subsequent analysis of the data may need to be interfaced to other systems such as those used for Reliability Centred Maintenance [6].

From an analysis perspective the increased data volume also has an impact. While during the prototyping phase it was possible to look at each data-set in detail and interactively view high load manoeuvres through flight-playback software, the volume of data gathered from even one or two aircraft can frequently result in this no longer being viable; there is simply just not enough time to go through every data-set in detail. Therefore, automated methods of isolating significant events that merit further detailed analysis have to be developed. These methods have to be capable of detecting events that require immediate intervention (e.g.: HUURP) or trends that indicate potentially detrimental changes in usage. This requires the development of simple and complex criteria, often termed rules, which can be used to rapidly process large quantities of data and isolate significant or anomalous events. While some of these rules are well defined, others are more "fuzzy" in nature and capabilities that allow analysts to explore the impact of combinations of parameter ranges rather than exact values often need to be developed.

To address these data volume issues, a number of organizations have explored the application of Data-Mining/Big Data techniques to aircraft health monitoring data. Although under the auspices of these two disciplines, there are a number of different techniques that can be employed, the basic premise is to interrogate large quantities of data and “extract” correlations between parameters which can then be used to isolate significant events. While this is certainly a legitimate approach to isolating data requiring further investigation, it is important to remember that although correlation is a necessary condition for establishing causality it is not a sufficient condition, i.e.: because combinations of different parameters can be correlated, it does not necessarily imply that they have physical significance. This again illustrates that while engineering understanding and assessment of critical structural locations can be facilitated by Data-Mining/Big Data techniques, the application of Data-Mining/Big Data techniques does not negate the need for engineering understanding and assessment.

A final consideration regarding the data volume effect that needs to be kept in mind is that although the data management and engineering aspects of an operational health monitoring program are often considered separately, there is considerable interaction between these two aspects. For example the methods implemented to acquire and manage the data can often impact the ease with which the data can be analysed and manipulated and vice-versa. This is why the probability of obtaining tangible LCM/ROI from health monitoring programs is increased significantly if the data management and analysis aspects of the program are considered concurrently with the engineering aspects of the program right from the program outset. The adoption of a systematic and structured approach similar to that encapsulated in an SSHMP [13], provides a framework that can be used to accomplish the concurrent development of both the data management and engineering aspects of an operational health monitoring program.

While discussions in this paper only begin to scratch the surface of the data management and analysis issues related to operational aircraft health monitoring, they will hopefully serve to convince the reader that for most operational programs, tools developed during the system prototyping stage (e.g.: using spreadsheets and other desktop software) will be inadequate for fleet implementation purposes. Far more robust tools featuring SQL compliant relational data-bases and rules engines with well-defined data-exchange interfaces will be required. The fact that the scope of the data management and analysis issues are often not fully appreciated or addressed at the outset of a program may go some way to explaining why that in a recent structural health management survey [3], over 63% of the participants considered the volume of data obtained and its associated management to be problematic.

Myth Busters!

The myths that have been outlined in the previous section indicate that while in-depth technical issues will always arise and need to be addressed during the implementation of any aircraft health monitoring program, the failure of many operational health monitoring programs to realize tangible LCM/ROI can be attributed to a lack of development and implementation of a systematic and structured LCM-centred approach. Additionally, any systematic and structured approach that is implemented will require multi-disciplinary input which often results in the need for trade-offs between conflicting requirements. If tangible LCM/ROI benefits are to be achieved in such an environment, it is imperative that the relatives merit(s) be evaluated in the context of overall LCM program objectives as opposed to merely from the perspective of reducing the inconvenience/costs associated with one specific discipline alone. For example, deciding to sample the data at an inordinately high sample rate rather than understand the sample rates that are actually required might well reduce some of the up-front engineering development and time-to-implementation costs. However, this decision will significantly, impact the ongoing data acquisition, validation, transfer, analysis and management costs. Over the life of the program these latter costs will exceed any of the initial engineering/implementation savings that were realized.

Providing a systematic and structured approach within the overall context of LCM is implemented, there are a number of ways that an effective program can be developed. As previously noted, Celeris Aerospace's approach to this process has led to the development of the Strategic Structural Health Management Planning (SSHMP) methodology [13].

An SSHMP is a dynamic process which needs to be periodically updated (typically once a year) to ensure it is effectively managing actual usage and remains consistent with fleet LCM goals. To provide some idea of the types of issues that need to be addressed in an SSHMP Figure 4 depicts the operational considerations that need to be addressed to define critical locations, acquire and validate appropriate operational and contextual data (Note: this by no means depicts the full scope of an SSHMP, it is merely one facet that has been provided for illustrative purposes).

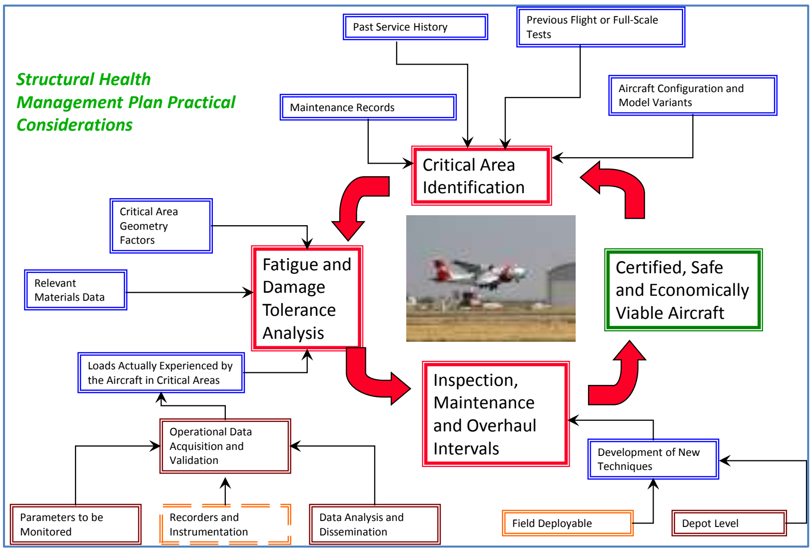

Figure 4: SSHMP Practical Operational Considerations

Generally an SSHMP is implemented in the following phased manner [13]:

1. **Phase 1 - Development and confirmation of a Strategic Structural Health Management Plan:** A well-structured SSHMP is capable of being applied to one or multiple fleets. Regardless of the aircraft type, there is a lot of infrastructure that is common to all programs. By properly exploiting this commonality significant cost-savings can be realized by operators of multiple fleets. The SSHMP itself:

    a. Provides a plan with a sliding window (typically five years) that looks at the implications of LCM requirements in the context of the current operational environment of each fleet;

b. Prioritizes fleet implementation and equipment acquisition in accordance with LCM requirements and any associated budgetary or operational constraints. Wherever practicable an SSHMP should seek to leverage and enhance existing infrastructure; and

c. Provides a framework for structural health management that will ensure regulatory agency compliance/acceptance by involving regulatory agencies from the outset.

2. **Phase 2 – Definition of Development and Implementation Plans for each fleet:** Even though infrastructure commonality is exploited as much as possible, customization for individual fleets cannot be avoided as a result of different aircraft (fixed or rotary wing) performance, capabilities and roles. Individual Fleet plans address fleet-specific issues such as sensor locations, data acquisition rates etc. The adaptation to individual fleets will require detailed review/analysis of such items as Inspection/Maintenance records, Service Bulletins, available engineering data, airworthiness directives etc. as illustrated in Figure 4. The fleet implementation plan will also define the criteria to be used for calibrating any installed monitoring system and for characterizing the anticipated loading that can be expected from individual roles. This latter activity is essential for the meaningful development of data reasonableness and error validation criteria that can be used to ensure the ongoing integrity and validity of measured fleet data.

3. **Phase 3 - Structural Health Monitoring System selection, installation, verification and training.** Once the requirements have been placed in the context of fleet LCM objectives (i.e.: Phase 1 and 2), it is appropriate to initiate the acquisition of the required system components such as the hardware, software and training elements needed to ensure a successful system deployment. It is interesting to note, that in their zeal to get "something on an aircraft and gather data as soon as possible (see Myth # 2)" many organizations class Phases 1 and 2 as "nice-to-have but not necessary" and either attempt to undertake them half-heartedly or even skip them altogether. Experience suggests that organizations which adopt such an approach invariably derive little to no LCM/ROI from their operational monitoring programs.

   Based on the Fleet Implementation Plan derived in Phase 2, it will be necessary to address the following issues in Phase 3:

   a. Selecting the aircraft monitoring hardware that best meets the fleet implementation program requirements via an appropriate RFP process (some trade-offs will invariably be required). In cases where the aircraft are being retrofitted with a monitoring system, the timing of the installation may also need to be considered, particularly with regard to evaluating if it is viable to install and calibrate the system during schedule maintenance or whether additional down-time, together with the associated loss of operational availability, will be required;

   b. Integrating existing software or acquiring new software as seamlessly as possible will facilitate the timely and accurate processing of the data. This involves far more than merely extracting the data from a recorder using manufacturer's proprietary software. Consideration also needs to be given to acquiring software that will:

i. Undertake the initial data acquisition validation and transfer to a central repository;

ii. Provide initial data processing that is capable of generating rapid and timely alerts for ground, maintenance and engineering personnel with respect to potential data errors, anomalies and harsh or unusual usage;

iii. Store the data and facilitate meaningful and prompt data retrieval;

iv. Undertake detailed analysis of critical structural areas based on the recorded data and evaluate the implications with regard to the ongoing structural integrity of the fleet;

v. Track and revise inspection/maintenance intervals as required; and

vi. Evaluate the life-cycle management implications of current and projected usage.

c. Training anticipated users of the system to ensure that the data is correctly acquired, appropriately responded to and correctly interpreted in a timely manner. Failure to provide such training can have disastrous consequences [17].

4. **Phase 4 - Ongoing data acquisition, validation and analysis**. Phase 4 is the phase in which the day-to-day monitoring activities actually take place. As illustrated in Figure 5 the acquisition, validation and transfer of the data typically requires the integration of a number of different technologies which if not addressed systematically can result in the loss of significant amounts of data, thereby increasing the risk that little to no tangible LCM/ROI will be realized.

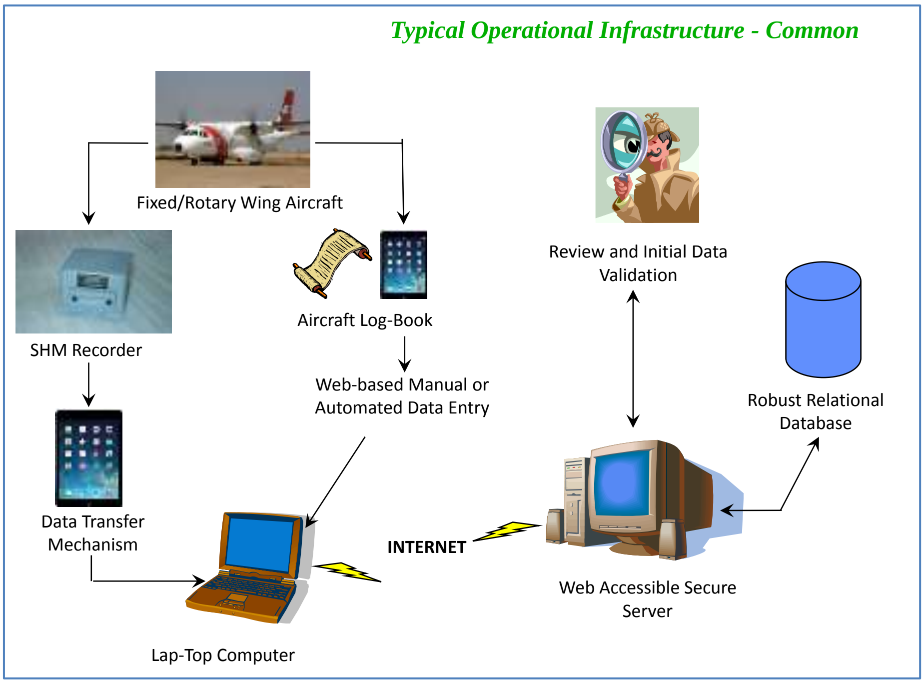

Figure 5: Typical Operational Data Acquisition and Validation Infrastructure

5. **Phase 5 - Ongoing review of usage and structural data to assess operational and life-cycle management implications:** Due to either limited understanding of the significance of the data acquired and/or budgetary restrictions, some organizations stockpile the data and only undertake periodic (sometimes around 5 years or more) detailed analysis of the data. This is usually a recipe for disaster as when the data is analyzed it is possible to find there are large quantities of invalid data (due to faulty sensors etc.) and/or manoeuvres or roles that prove to be severe. Had the data been analyzed on a more regular basis, corrective action could have been implemented to reduce or eliminate the causes of invalid data and mitigate the accumulated damage resulting from severe manoeuvres or roles. As a rule, any purported economy that may be derived from infrequent analysis of relevant data that is gathered generally turns out to be a false economy.

   Assuming that an effective operational health monitoring program has been implemented in a systematic manner it is recommended that the data that is gathered needs to be reviewed/analysed on a regular basis which should not be in excess of one year. The scope of the review/analysis should consider, but not necessarily be limited to:

   a. <u>Analysis of usage trends:</u> Is the aircraft being operated in a way that is consistent with past usage or is the usage and associated load spectrum at critical locations changing (even though the aircraft might be operating in nominally the same role)?

   b. <u>Evaluation/isolation of structurally damaging manoeuvres:</u> What are the roles/ manoeuvres that are causing the most damage at critical structural locations? Are these roles and manoeuvres necessary and, if so can their execution be modified to decrease the severity of the damage that is incurred? If for operational reasons the severe roles/manoeuvres cannot be eliminated or modified, can the damage accumulated by the fleet be better managed through fleet rotation strategies etc. (see Myth # 5)?

   c. <u>Addressing the life-cycle management implications of current and projected usage:</u> How will any change in usage that is detected impact anticipated fleet retirement dates? Based on current usage will major component replacement be required to ensure the aircraft meet their anticipated fleet retirement dates? How cost effective will any proposed modification prove to be? For example the merits of spending significant money on replacing a centre wing box to extend the life of a wing ten or twenty years needs to be considered carefully in the context of the potential replacement times/costs of other critical components (e.g.: the empennage). If the life consumption of other critical components result in the aircraft no longer being viable to repair/ maintain before the cost-benefits of the centre-wing box replacement can be fully realized, the LCM benefits of undertaking the centre wing box replacement may well be negated.

   d. <u>Evaluating the SSHMP to ensure it truly reflects current and projected usage:</u> As previously noted an SSHMP needs to reflect current LCM goals and as such is a dynamic document which reflects changes in fleet usage/role and/or changing corporate requirements. As the plan evolves over time it can also be used to assess/manage the viability of fleet life-extension initiatives and/or viable fleet replacement options.

Assessing the Effectiveness of Aircraft Health Monitoring Programs

Having outlined the need for a systematic and structured approach to Aircraft Health Monitoring programs within the context of overall fleet LCM requirements, the question arises as to how the effectiveness of current or proposed programs can be evaluated? While there are a number of ways that could be used to assess whether an Aircraft Health Monitoring program will achieve tangible LCM/ROI, experience has shown that developing an objective appraisal of current or proposed programs using the following questions can be beneficial:

1) **Is there a Strategic Structural Health Management Plan (or something similar) in place?**

   a) If so how current is it?

   b) If not, how is the structural health of fleet aircraft being managed?

2) **What are the top three structural issues for each fleet that is being operated?**

   a) If they are known; how are they going to be addressed in a timely manner?

   b) If they are not known; why are they not known?

3) **For each Structural Health Monitoring system that is being installed on operational fleets:**

   a) Can each of the selected parameters be related to specific fleet LCM objectives?

   b) What detailed plans exist to acquire, validate and use the data that is collected?

4) **What provisions, if any, exist to address emergency Airworthiness Directives(ADs) that are issued as a result of the use of similar aircraft by other operators (i.e.: how quickly can you establish that their problems might not be yours?)?**

   a) How long could an Aircraft-on-Ground (AOG) situation from a problem with similar or identically configured aircraft, which may not necessarily be applicable your fleet as a result of different usage, be sustained?

5) **How are the current short and long-term life cycles of your fleets managed?**

   a) Do you know whether the aircraft are being used as intended?

   b) How confident are you that you will meet your required fleet retirement dates?

   c) How do you intend to validate the cost-benefit of proposed structural life overhauls/modifications/upgrades?

   d) What basis will you use to define the requirements for any replacement aircraft?

If you or your organization are struggling to generate robust and satisfactory answers to these questions then the effectiveness of any existing or proposed program should be critically re-evaluated as there is a significant risk that minimal LCM/ROI will be obtained.

A Case for Program Convergence?

As previously noted in Table 1, there is a plethora of aircraft health monitoring initiatives to ostensibly address "unique needs". There is significant amount of overlap between many of these programs, the data from which contributes towards the overall LCM process to a greater or lesser degree (Figure 3). This was something that was recognized by many of the participants in a recent survey on aircraft structural health management/monitoring programs [3]. Unfortunately, these programs not only differ in name but also in the monitoring systems that are required to be installed on the aircraft, many of which draw upon the same data sources. Consequently, it is possible for aircraft to be equipped with both a Flight Data Recorder (FDR) and a Health and Usage Monitoring System (HUMS) which are largely recording the same data. While many of the individual program types (FDR, SHM, FOQA, HUMS etc.) have attempted to standardize data formats and sample rates etc. within a program type/system, there have been few if any attempts to standardize data formats and sample rates etc. across program type/system. Although attempting to standardize data formats and sample rates across program type/system will require some initial effort, it is considered that the downstream benefits of establishing a unified data structure and format could be significant. These include, but are not necessarily limited to:

1) Less confusion regarding the purpose of the data is gathered and its limitations, as it would now be placed in the context of the LCM process;

2) The development of more cost-effective monitoring programs;

3) Less burden on operational resources;

4) Simplification of software data acquisition and validation systems;

5) Reduced hardware installation and maintenance costs (one system addressing multiple requirements).

Applicability to Non-Aerospace Domains

While the examples presented in this paper have addressed aircraft health monitoring programs, the principles upon which the concepts have been developed follow an object oriented methodology. As illustrated in Figure 6, at the highest level this means that many of the concepts developed can be readily adapted to non-aerospace applications.

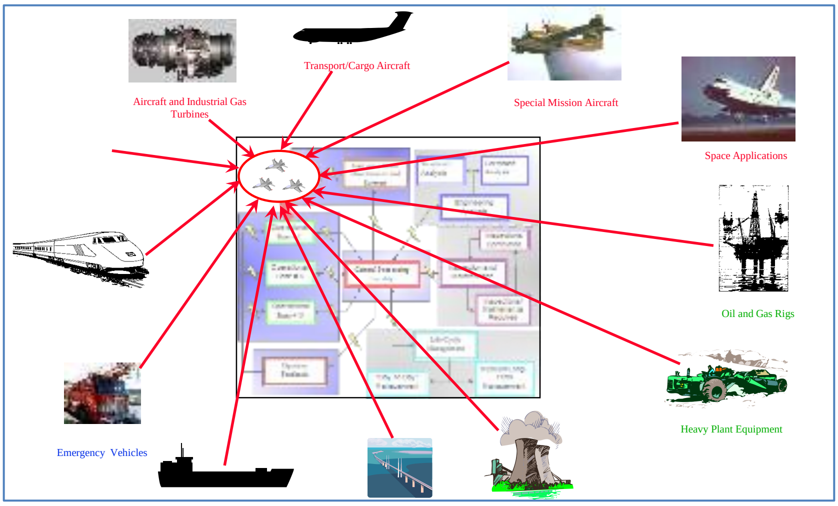

Figure 6: Leveraging the SSHMP Methodology Across Multiple Applications

Conclusions

Aircraft Health Monitoring programs have existed in one form or another for over 60 years. Since the mid-1990's rapid advances in micro-processor, computer, sensor and data acquisition technologies have revolutionized the quantity and quality of the data that can be acquired. These advances have birthed a number of different monitoring program types which are ostensibly used for different purposes.

Despite all the technological advances cited, many programs fail to realize their full potential and generate tangible life-cycle management returns-on-investment. This paper has attempted to provide some insight into why such programs fail to reach their full-potential by identifying a number of common misunderstandings (myths) related to their function and implementation. Solutions as to how this situation might be best turned around have been proposed. Key factors to achieving tangible LCM/ROI from these programs are considered to include, but not necessarily be limited to:

1) Adopting a concurrent structured and systematic approach to the data-management and engineering aspects of these programs right from the program outset based on corporate life-cycle management objectives. One such approach is encapsulated in the Strategic Structural Health Management Plan (SSHMP) developed by Celeris Aerospace.

2) Ensuring that the goals and objectives of any structural health monitoring programs are clearly defined and communicated to all potential stakeholders involved in the acquisition and/or utilization of the data; and

3) Instituting frequent and ongoing reviews of the data that is collected to facilitate rapid and timely responses to harsh or unusual usage and changes in utilization that will detrimentally impact the operational lives and ongoing safety of the aircraft being monitored.

Questions that can be used to realistically evaluate the chances of achieving tangible LCM/ROI from current or proposed programs have been suggested. Additionally, a case for consolidating many of the existing different program types (e.g.: SHM, FOQA, HUMS etc.) to facilitate the more cost effective implementation of these programs has been outlined.

Finally, while the examples provided in this paper have related to aircraft health monitoring programs, the object oriented nature of the concepts presented means that many of them can be readily adapted to non-aerospace applications.

Acknowledgements

The majority of the concepts outlined in this paper have been developed as a result of the author having been privileged to collaborate with many talented individuals working in the field of aircraft health monitoring over a number of years. In particular, the author is indebted to the insight into many of the data management concepts provided by Mr. John Miner of Celeris Aerospace.

References

1. *"Report of the Court of Inquiry into the Accidents to Comet G-ALYP on 10th January, 1954 and Comet G-ALYY on 8th April 1954"*, British Ministry of Transport and Civil Aviation, Her Majesty's Stationary Office, February 1955. Available from: http://lessonslearned.faa.gov/Comet1/G-ALYP_Report.pdf .

2. Hall, Stephen R. and Vellacott, Mark, *"Safe and Economic Management of Widespread Fatigue Damage (WFD) Using Prognostic/Diagnostic Health and Usage Monitoring"*, Proceedings of 5th DSTO International Conference on Health and Usage Monitoring, Melbourne, Australia, 20-21 March 2007.

3. Hall, Stephen R., *"Analysis of the Data Obtained from a Survey on Aircraft Structural Health Management/Monitoring Programs"*, Celeris Aerospace Canada Inc. Technical Report CAC/TR/14-001, 16 June 2014.

4. Hall, Stephen R., *"The Data Delusion: Do we really need more data?"* Presentation to the Ottawa Chapter Society of Reliability Engineers, 29th October 2014.

5. Hall, Stephen R., *"Lessons Learned from Recent North Sea Super-Puma, Incidents"*, Presented at the Bristow HUMS/HFDM Workshop, HAI Heli-Expo, February 2014.

6. Hall, Stephen R., *"The Reliability of Structural Reliability Metrics for Fixed and Rotary Wing Aircraft"*, Presentation to the Ottawa Chapter Society of Reliability Engineers, 26th February 2014.

7. *"Damage Tolerance and Fatigue Evaluation of Structure"*, U.S. FAA Advisory Circular, AC 25.571-1D, 13 January 2011. Available from: http://www.faa.gov/documentLibrary/media/Advisory_Circular/AC%2025.571-1D%20.pdf .

8. *"Fatigue, Fail-Safe and Damage Tolerance Evaluation of Metallic Structure for Normal, Utility, Aerobatic and Commuter Category Airplanes"* , U.S. FAA Advisory Circular AC 23-13A, 29 September 2005, Available from: http://rgl.faa.gov/Regulatory_and_Guidance_Library/rgAdvisoryCircular.nsf/list/AC%2023-13A/$FILE/ac23-13A.pdf

9. Hall, Stephen R. and Stubbs, Norman, *"Ensuring the Continuing Airworthiness of Aerial Firefighting Aircraft"*, Presentation to the California Department of Forestry, Sacramento, California, 25 April 2006.

10. Hall, S.R., *"Consolidation and Analysis of Loading Data in Firefighting Operations: Analysis of Existing Data and Definition of Preliminary Airtanker and Lead Aircraft Spectra"*, FAA Final Report DOT/FAA/AR-05/35, October, 2005, available from: http://www.tc.faa.gov/its/worldpac/techrpt/ar05-35.pdf

11. Broek, D. 1989, "*The Practical Use of Fracture Mechanics"*, A. A. Dordrecht, The Netherlands, Kluwer Academic Publishers (ISBN 0-7923-0223-0).

12. Hall, Stephen R., *"A Practical Method for Assessing the Validity of Strain Gauge Data Obtained from Advanced Composite Structures"*, Proceedings of the Fifth International Conference on Composite Structures (ICCS 5), Paisley, Scotland, 1989.

13. Hall, Stephen R., *"Strategic Structural Health Management Programs: An Essential Tool for Safe and Cost-Effective Life-Cycle Management"*. Celeris Aerospace Webinar, 2014.

14. Reed, S. C. and Holford, D. M., *"Guidance for Aircraft Operational Loads Measurement Programmes"*, MASAAG Paper 109, 31 May 2007.

15. Hall, Stephen R., *"Practical Considerations Associated with the Development and Implementation of Aircraft Structural Health Management Programs"*, Celeris Aerospace Canada Inc. Webinar, 2011.

16. Hasty, Joshua C., Speaks, Kenneth W. and Kennedy, Joshua S., *"Comparison of HUMS Benefits – A Readiness Approach"*, Presented at American Helicopter Society Forum, Grapevine, Texas, 27-29 May 2009.

17. *"Report on the accident to Aerospatiale (Eurocopter) AS332 L2 Super Puma, registration G-REDL 11 nm NE of Peterhead, Scotland on 1 April 2009"*, UK Air Accidents Investigation Branch, Department of Transport, 24 November 2011, Available from: http://www.aaib.gov.uk/cms_resources.cfm?file=/2-2011%20G-REDL.pdf .

18. Hall, Stephen, R., *"Celeris Aerospace Webinar Series: The Use and Abuse of Aircraft Structural Severity Factors"*, Celeris Aerospace Technical Memorandum, CAC/TM/13-001, February 2013.

19. *"Intelligent Management of Helicopter Vibration Health Monitoring Data"*, U.K. CAA Paper 2011/01, May 2012. Available from: http://www.caa.co.uk/docs/33/2011_01RFS.pdf .

20. *"The Application of Advanced Anomaly Detection to Tail Rotor HUMS Data"*, U.K. CAA Paper 2012/01, December 2012. Available from: http://www.caa.co.uk/docs/33/PAPER%20201201%20AAD%20HUMS%20DATA%20310113.pdf .

21. Hall, Stephen R., *"Observations Pertaining to Recent EC225/AS332 Structural Failures and Their Resolution."*, Celeris Aerospace Technical Memorandum CAC/TM/13-002, 30 April 2013.

22. Hall, Stephen R., *"Observations Related to the Proposed EC225 Return-to-Flight Methodology"*, Celeris Aerospace Technical Memorandum CAC/TM/13-005, 23 July 2013.

23. Hall, Stephen R., *"Knee-Deep In Data with Nowhere to Go? (Restoring Sanity to Structural Health Monitoring Programs)"*, Proceedings of the FAA/DoD/NASA 2000 Aging Aircraft Conference, St. Louis Missouri.