Lessons Learned in Implementing a Practical Aircraft System Health Management (ASHM) System

Joshua Koelle[1], Matthew Smith[2], Peter Sulcs[3], Greg Kacprzynski[4] and Rhonda Walthall[5]

[1,2,3,4]*Impact Technologies, a Sikorsky Innovations Company, Rochester, NY, 14623, USA*

joshua.koelle@impact-tek.com
matthew.smith@impact-tek.com
peter.sulcs@impact-tek.com
greg.kacprzynski@impact-tek.com

[5]*UTC Aerospace Systems, San Diego, CA, 92123, USA*

rhonda.walthall@utas.utc.com

ABSTRACT

Aircraft System Health Management (ASHM) is a web application used for Boeing 787® and Airbus A320® and A380® aircraft system monitoring by airlines and field engineers worldwide which also serves all existing Aircraft Condition Monitoring Function (ACMF) reports, Flight Deck Effects (FDE) records and aircraft metadata to UTC engineering teams to aid in efficient aftermarket support. This enables creation, testing and fielding of off-board diagnostics and prognostics modules of varying levels of sophistication, that convert this abundance of existing data into actionable and timely knowledge about a/c fleet health. ASHM encourages and promotes cross functional collaboration allowing those with the most subject matter expertise within the enterprise to access the field data they need to observe operational performance and to create, test and field modules that can actively diagnose and warn field service professionals of problems when and potentially before they arise. A practical case study related to monitoring of the novel Boeing 787® electromechanically driven distributed aircraft environmental systems is presented. This use case motivates a discussion of pragmatic lessons learned in the fielding of diagnostic and prognostics solutions.

1. INTRODUCTION

Modern commercial aircraft contain computerized maintenance systems that have replaced the dials, indicators, switches, and diagnostic read-outs of prior generations of aircraft. These systems in addition to performing on-board diagnostic functions also record high value parametric data that can be used for system and component health tracking, fleet data studies, and prognostics. By observing changes in component performance or recognizing abnormal response behaviors it is possible to observe incipient fault conditions before they grow into significant problems that are recognized by the on-board Built-in Test (BIT) checks which in the worst case may cause a delay or cancellation of service. Looking at data trends across a fleet of aircraft can identify outlier behaviors that are indicative of degraded health, determine the effect of usage factors on component life, and optimize maintenance practice. Prognostics goes beyond fault assessment to project remaining useful life, allowing advanced scheduling of maintenance procedures, proactive replacement part allocation, and enhanced fleet deployment decisions based upon the estimated progression of component life usage.

This aircraft data and diagnostic information is of high value to multiple groups who have interest in the current and future health states of these critical systems. First, the airline operators of these aircraft are looking for a change in aircraft health that can be used to optimize maintenance and prevent the occurrence of delays and cancellations. Second, the field service engineers who support the airlines benefit from information that allows them to better support maintenance troubleshooting and logistics. Finally, the system engineers who support each of the aircraft systems can obtain critical information to characterize product issues and develop enhancements to diagnostic capability. To serve each of these important user groups, UTC Aerospace Systems has developed and continues to improve and expand the Aircraft System Health Management (ASHM) Tool.

ASHM takes in ACMF reports for selected subsystems and components of each supported aircraft platform, parses and

Joshua Koelle et al. This is an open-access article distributed under the terms of the Creative Commons Attribution 3.0 United States License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

processes the reported parameters against thresholds, computes estimated or expected values for some key parameters, and serves the report data and the processed results as part of a fleet view available to airline, maintenance, and engineering users. The application allows for the creation, integration, and execution of custom analytic modules that extract enhanced diagnostic and prognostic information from the raw report data. This information is also made available for visualization, trending, and alerting.

In this paper, the sub-systems of the Boeing 787 will be used as a case study of how this high value parametric data and derived diagnostic and prognostic information can be used to enhance commercial aircraft maintenance practice. The Boeing 787 is outfitted with a modern computerized maintenance system that records key data for each of these aircraft systems. The available on-board data is transformed by ASHM into actionable component and system fleet information to guide fleet troubleshooting, opportunistic maintenance, and logistics. In addition to the Boeing 787, the ASHM software is also being used to support the Airbus A320 and A380 platforms.

2. AIRCRAFT DATA SOURCES

The ASHM software tool collects, organizes, and stores aircraft data from two different sources. The first class of data is comprised of system status flags that report anomalous behavior or degraded performance as obtained from on-board BIT checks and diagnostic functions. These events which are commonly referred to as Flight Deck Effects (FDEs) are captured along with linked maintenance messages that capture the symptoms of the observed condition. The second class of data is recorded by the Airplane Condition Monitoring Function (ACMF) and is specifically targeted at long term analysis of aircraft health and usage (Ramohalli 1992).

The ACMF report data is of particular interest to the ASHM fleet monitoring software. These reports capture aircraft parametric data based upon triggering criteria and a format that have been established by aircraft system domain experts. The specific content in each report can be targeted at periods of operation sensor signals that are of particular interest for a given component or fault mode. Some reports trigger based upon entry into a given operating mode that is appropriate for system performance characterization or detection of anomalies. Other reports are triggered by the occurrence of specific events that are noteworthy, abnormal, or perturb the system in a way that makes performance or fault conditions more observable. Each report collects an assortment of sensor data, state information and contextual metadata that is of interest to the system or component monitored by the report. This data includes typical operating parameters such as: temperature, pressure, position, and speed. The reports in some cases record operational state values or calculated values that have been derived from the raw parametric data. Often, the aircraft system data is reported along with aircraft level data such as altitude, air temperature, and Mach number that provide important context about the operating conditions at the time of acquisition. The data may be acquired as a single snapshot in time, a set of statistical metrics such as mean or peak value, or as a time series history.

3. ASHM SOFTWARE ARCHITECTURE

ASHM is an enterprise application that monitors systems on multiple commercial aircraft platforms using ACMF reports, automated parameter alerts and notifications and diagnostic and prognostic custom analytic modules. The overall software architecture is depicted in Figure 1. ACMF reports and other information sources are generated by the on-board maintenance system. The reports are offloaded and automatically processed by ASHM using a fully automated workflow. The ASHM application parses and stores incoming reports to generate alerts and execute advanced algorithms that were developed or adapted using existing system models and data. The event driven architecture powers the real-time web portal where airlines, maintenance support and engineers can analyze reports and be notified in advance of potential issues. The portal also displays high level dashboards for easy information consumption and drill-downs, graphs and report viewers for detailed analysis.

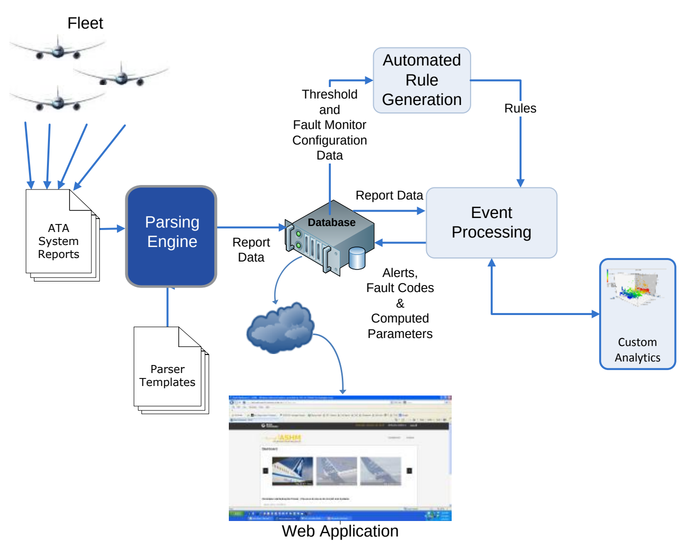

Figure 1. ASHM Software Architecture

An extract, transform, load (ETL) engine is utilized in ASHM. As different report types arrive they are bucketed by type into directories for further processing. Each subsystem has one or more report types, and as ASHM grows to process more systems of each aircraft platform, and adds more aircraft platforms, the number of report types will grow accordingly. In the first parsing step for each report, the report is parsed at a high level to determine the specific type of report, specifically the subsystem, report type, version that report type, and based on that information a determination is made and action taken to move that report to the proper staging directory. If the report is not recognized it is placed in a separate bucket. Then, the Data Transformation agent reads each report as one record of input, and parses the parameters from that report. It stores each parameter as part a unique record for that report in the application database.

The complex event processing portion of the ASHM software is shown in Figure 2. Data flows through RulePoint by report type, originating with a SQL Source that acquires parameter instance data from the ASHM database and pushes it into a RulePoint® Topic. A Rule references one or more topics and may use data from those topics to 1) determine anomalous conditions, e.g. value out of range, 2) compute new values based on those parameters, 3) send those computed values or detected conditions to a Responder that is responsible for storing new data back to the same ASHM database. For the ASHM project automatic alert rule generation based on thresholds defined in the database is employed. This custom tool uses a Java API Adapter to 1) connect to the development RulePoint instance, 2) remove all previously generated (as opposed to hand entered) rules, and 3) generate a new set of rules based on those thresholds.

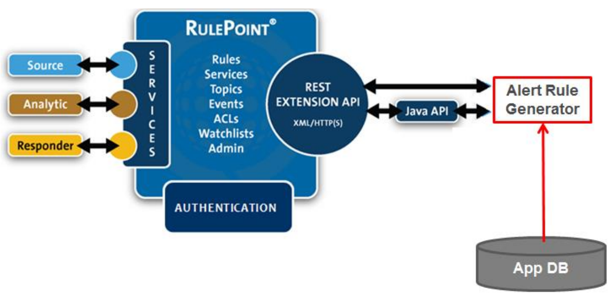

Figure 2. Detail View of Event Processing Software

The ASHM application checks for out of range "alert" conditions on selected incoming report parameters, looking for warning or alarm conditions that are higher or lower than expected under normal operating conditions. Each "alertable" parameter has its own set of thresholds defined in the database for low and high warning and alarms. There are also mechanisms in place to define two additional criteria which are when the thresholds are to be ignored, say when some (the same or another) parameter's value meets a certain conditional relationship with a fixed value, e.g. <= some value, = some value, or >= some value. The parameter alert rules store parameter out of range conditions back to the database, where they are used to display those anomalous conditions to the end user in the web application. In addition to simple thresholds, the ASHM application can invoke a custom analytic that performs calculations on the report data. This functionality is used to run diagnostic and prognostic algorithms that extract refined system health information from the raw parametric data. The process for developing, vetting, and integrating these custom analytic modules is described in detail in the following section.

4. CUSTOM ANALYTIC DEVELOPMENT PROCESS

The development and use of custom analytic modules enables health monitoring capability that extends far beyond what is possible using only parameter alerting and trending of raw ACMF data. These modules use ACMF report data to extract subtle fault indicators that give advanced notice of impending failures at severity levels below that captured by the on-board diagnostics. The overall development process for the ASHM custom analytic modules is shown in Figure 3.

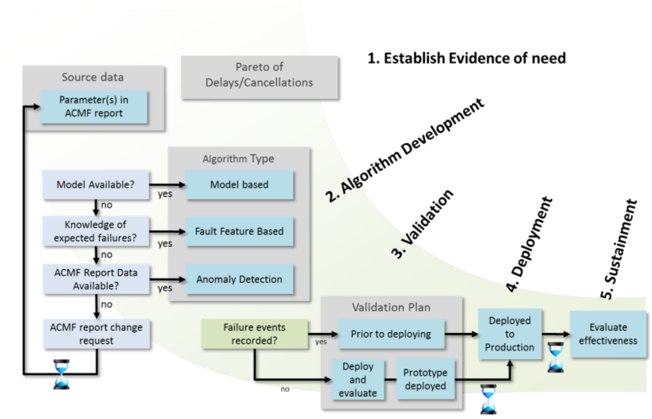

Figure 3. Custom Analytic Development Process

The first step in the development process is the establishment of need. When deciding where to allocate effort to create new diagnostic and prognostic capability, the first priority is given to components and fault modes that cause flight cancellations and delays. These service interruptions (SIs) have significant negative effects for all interested parties from the airlines, field service support, and product engineers. The primary objective of custom analytic development and ASHM in general is the minimization of SIs throughout the fleet. The leading causes of SIs are established by analyzing fleet reliability data compiled by the worldwide network of field service personnel. This data is summarized using Pareto chart analysis and the most significant contributors to SIs are given the highest priority when developing new custom analytics. After SI data, priority is given to the components and systems that are creating the highest load on aircraft maintainers. This is quantified by analyzing LRU removal data compiled by field service personnel to determine which components are involved in the highest number of maintenance procedures. While the removal data aligns with the SI records in many situations, there are some cases where the component criticality or the complexity/time-intensiveness of a given maintenance procedure may affect priority for a given component or system. Lastly, there is the issue of component logistics. Replacement of large, expensive LRUs like Auxiliary Power Units (APUs) requires careful planning due to limited product availability and transportation time. Therefore these systems are prime candidate for development of PHM technologies that provide an assessment of component health, and forecast future maintenance needs.

The second step in the custom analytic development process is creation of the diagnostic or prognostic algorithm. In this step the PHM design experts collaborate with the subject matter experts to determine the available raw materials to support algorithm development, and how this information aligns with known system issues. When possible, model based approaches to diagnostics and prognostics are preferred. The raw ACMF report data represents only one snapshot in time of system operation. It can be difficult to know if a change in a parameter is due to fault or the normal variations that occur over the operating range of the system. The use of system and component models that simulate normal system response can be used to reduce the effect of typical system variations on health assessment performance. Modern engineering design makes extensive use of simulation during the product development phase. For many of the systems the component design models have been readily accessible for implementation into multiple custom analytic modules. However, a suitable model is not available in every case, and in some cases the data required to exercise the model is not present in the related ACMF reports. In those cases, an approach is used that acts directly on the raw report data. If there is a good understanding of the failure modes, either via documented field or laboratory failure data or an understanding of the underlying physics, a feature based approach can be developed to extract known fault indicators from the parametric data. If there is limited understanding of the nature of failure, or if a generalized anomaly detection capability is desired, an approach that is based purely on the known range of healthy data is used.

Regardless of the approach, the module must be designed so that it has sufficient robustness to accommodate variations in the report data that are not associated with system performance. The ASHM software package is independent of the on-board aircraft software and therefore must accommodate changes in report format and content as they occur. The nature of the report data may be affected by one or more component operating states, and these must be observed and tracked to ensure proper operation. The software must also be prepared to recognize on-board data collection irregularities and screen out the affected values so that spurious fault correlations are not reported.

If the current ACMF data is deemed insufficient for enhanced system monitoring, the team turns its focus to establishing how the report can be modified to obtain the highest value condition information. The reports were designed to provide the most important system condition information as understood at the time of implementation. On a new aircraft like the Boeing 787®, this means that these decisions were made based upon a theoretical understanding of the system or using the available test data. An examination of field data can identify opportunities to improve the available system configuration information based upon actual usage.

After the diagnostic or prognostic method has been created, it is subjected to a series of validation tests to ensure that it provides the desired system heath information with an acceptable level of performance. Generally the performance of diagnostic methods is evaluated by determining the correctness of fault detection results, and the accuracy of fault severity assessment metrics. Of particular interest are the rates at which two diagnostic results occur: false alarms, or when the diagnostic system detects a fault but the condition of the system is not significantly degraded, and missed detection when the diagnostic system does not indicate a fault when one is known to be present. To support this activity, again the field SI and removal data is used to obtain ground truth information about the health state of the fleet. Documented failures, particularly those with a conclusive root cause assessment are extremely valuable in establishing system response at degraded health states. The fleet service history not belonging to known fault cases can be used to establish baseline system performance. Laboratory test results can be used to supplement field data experience, particularly in cases where practical fault experience is limited. The validation of prognostic approaches is a far more complex topic and has a more significant set of input requirements (Byington, Roemer, Kalgren, & Vachtsevanos, 2005). Given the long timescale of component life, it is generally impractical to complete significant validation prior to algorithm deployment. However, it is of critical importance to establish the accuracy and uncertainty bounds of the models used to assess and predict system health progression.

Validation is difficult in cases where the fault mode does not result in a condition that requires documented maintenance actions. For example, the air flow pathways in environmental control systems or air management systems may become contaminated by foreign material that may become lodged in components such as heat exchangers (Najjar, Hare, D'Orlando, Leaper 2013). This condition is problematic and requires a cleaning operation, but may not result in a removal or SI. Knowing when these events occur is essential to establishing health state ground truth for the related fault modes. This example illustrates the importance of communication between the operators and field support engineering in developing and validating effective PHM methods.

This document does not contain any export controlled technical data.

Validation may be performed at various points throughout the development cycle. It is preferable to conduct significant validation prior to deployment of the analytic to ASHM. This is more realistic for mature platforms that have been in service for a significant amount of time. For example, the ASHM team has created a custom analytic module to evaluate the progression of Airbus A320® APU health. There is a wealth of data for over a decade of operation for a very large fleet of aircraft, and significant validation was possible prior to analytic deployment. By comparison the Boeing 787® is a relatively new platform, with significantly less observed health progression and field maintenance issues throughout the fleet. While is it possible to confirm proper basic functionality and baseline response, it may not be possible to completely validate algorithm response to field failures prior to deployment. In these cases, it may be necessary to deploy a version of the algorithm for validation against new fleet data as it arrives.

When a custom analytic reaches a state that is mature enough for implementation in ASHM, it enters the production rollout phase. The algorithms as created by the PHM and system domain experts are translated into production software modules. The engineering and software teams work together to define a set of verification tests that evaluate all relevant logical paths within the software. These tests ensure that the production implementation matches exactly the approach that was validated during engineering sandbox development. The input and output data streams are established and any relevant contextual information is integrated as meta-data that is cataloged in the production database and made available to down-stream processes that consume the custom analytic output. Finally, the module output is integrated into the downstream ASHM processes that will serve this enhanced system condition information to the users. This includes configuration of custom data visualizations, data plotting and trending, and definition of parameter alerts including warning and alarms.

Upon successful deployment of a custom analytic module, it enters the sustainment phase. The output of the module will be regularly inspected and compared to the documented SIs, field failures and maintenance actions. The report data will be monitored for format changes or other updates that require reconfiguration of the event processing configuration or custom analytic software. If new classes of failure or degradation are observed, or the understanding of a given fault mode changes, an alternative version of the custom analytic can be created and evaluated in the engineering sandbox as a software upgrade candidate.

5. GENERATION OF ACTIONABLE INFORMATION

ASHM aids the user in extracting actionable information for short term proactive fleet support from the raw data sources. It does this by raising visibility of event reports (system reports that are only generated when an anomaly is encountered), automatically interpreting various error codes generated by the monitored equipment and triggering ASHM alerts based on threshold exceedances on reported parameters.

Custom analytics provide the means of generating more sophisticated health indicators from the raw data. These health indicators provide actionable information in the following ways: diagnostic and or prognostics indicators augment the raw report within ASHM's report viewer; alerts based on computed parameters are displayed alongside the ones based on raw reported parameters; and finally computed parameters can be included in the ASHM graphs pages to be monitored and observed visually for trends or anomalous behavior.

The final way that ASHM provides actionable short term information is by allowing the user to compare health indicator parameters by aircraft across the whole fleet thus focusing attention on the aircraft with health indicators that are abnormal compared to the fleet.

ASHM and the associated Data Analytics Tool also provide actionable information for a different audience with a longer term interest: the Engineering teams responsible for supporting fielded systems as well as new product design. ASHM aggregates operational field data that is invaluable in terms of closing the loop between Engineering and Field Support. It provides unprecedented visibility into how the systems that Engineering designs are operating in the field. This allows a closed loop refinement of all the assumptions made at design time, to both improve the current product offering and enable better design assumptions for the next generation of new products.

The goal of any deployed custom analytic is to provide the user of ASHM with actionable information. The following case study is a simple example of extracting invaluable information available by implementing a straight forward custom analytic.

6. CASE STUDY – ELECTROMECHANICAL SUBSYSTEMS

The transition from engine driven hydraulic subsystems to distributed electric motor driven subsystems has led to ACMF reports that capture characteristics from a wide array of systems despite being targeted to one component. A smart motor controller may now be responsible for 3 or more different tasks throughout the duration of a single flight. These tasks range from air management, to motor start procedures. The case study presented here will examine the complexities of the Boeing 787® motor controller system and its effect on the interpretation of ACMF data.

In previous generation civilian aircraft, the main engines generally provide four main sources of auxiliary power: electrical, pneumatic, hydraulic, and mechanical. The electrical system supplies power for avionics equipment, lighting, and in-flight entertainment. The pneumatic bleed air system supplies power for cabin pressurization and wing anti-ice systems. The hydraulic system provides power for flight control systems and auxiliary systems, and the mechanical power is used within the engine for oil and fuel management.

The electrical subsystems incorporated on next generation electric aircraft combine some of these power systems into one, through the use of generators and smart power distribution systems. This reduces system hardware complexity, resulting in weight reduction and efficiency improvements. A secondary result of this change is an increase in electrical system complexity. There are multiple electrical power distribution systems and new FAA certified software accompanied with these systems (Wheeler & Bozhko 2014).

The Common Motor Start Controller (CMSC) ACMF reports contain parametric data on component power draw, fluid temperatures, power frequency, and active mode information. Among these reports, data is provided in two different formats; time series data and snapshot data. The time series data provides prognostic systems with valuable information regarding how a system reacts to applied power. With this data it is possible to measure spool up time for an engine, or when a starter generator is drawing maximum power. ACMF reports which contain time series data typically target vital aircraft procedures, like Main Engine start or APU start. The snapshot data, in contrast, allows maintainers to track and trend parameters and features over extended periods of time. ACMF reports, which contain snapshot data, capture data at key points throughout each portion of a flight, from Taxi to Landing. The snapshot data is generally divided into two capture methods, peak data value and average data value. Both of these data points are useful in application to component fault detection. While the ACMF data provides vital information about the functions of these subsystems, it is equally as important to understand the operational modes and connectivity of the CMSC subsystem. Without this understanding, the data loses its diagnostic and prognostic value. Figure 4 is a simplified representation of the CMSC component assignments. In this figure, each CMSC is connected to two unique components, Cabin Air Compressors (CAC) and Generators (G), which then provide power to critical systems.

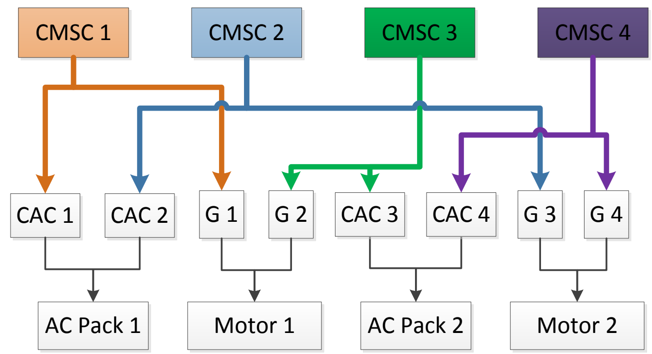

Figure 4: Representation of CMSC Component Assignments

An important part of establishing component connectivity is the operation mode flag. Each CMSC reports an operational flag associated with each data point, whether single point or time series. That operational flag indicates the CMSC's active mode, which ties to one of the Boeing 787's redundant system components. These operational flags have a significant bearing on the development of custom analytics for this system.

In the time series ACMF data, the operational mode flag can be used in determining how long the CMSC was driving a specific component and, in combination with the parametric data, provides information about the response of the controller and motor pair. The process involves filtering the data based on the operational mode flag. The time series ACMF reports contain data for a defined period of time after a trigger has occurred. This results in the ACMF report containing data from a CMSC driving multiple components over the course of that time series. When a CMSC is responsible for driving components of vastly different power requirements, it becomes clear that it is vital to filter the ACMF reports by the operational mode prior to applying statistical analyses to the data. Development of analytics which harness this time data can assist maintainers and engineers in analyzing and categorizing the effect of routine procedures on key aircraft components. Figure 5 illustrates the operational mode filtering utilized to select a targeted subset of data.

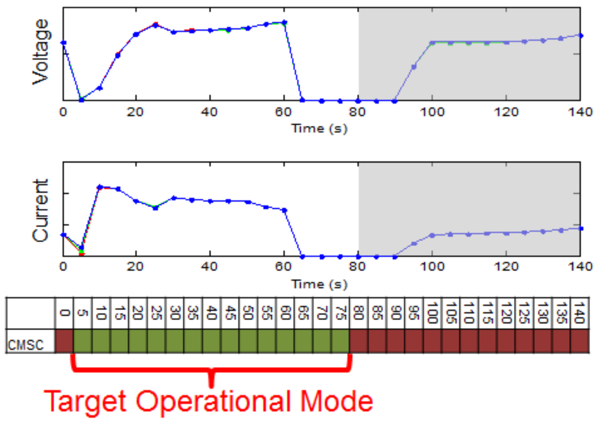

Figure 5: Filtering of Data Using Operational Mode Flag

By contrast, the snapshot ACMF reports contain peak or average data from the CMSC at specific points in each flight leg. In this data, the operational mode flag is useful in plotting component trends over time. For example a motor which draws more power each flight, while performing the same duty cycle, might indicate that the motor or driven component is experiencing degradation. Pairing this information with removals, allows for the development of fault detection thresholds which can be alerted on in ASHM, notifying the maintainers that maintenance action should be taken to avoid future Service Interruptions.

It is not enough, however, to know that a CMSC is driving a specific type of component. For redundancy and distributed system operation, multiple instances of many components exist on the aircraft. A complex logic dictates which of these components is active, and which motor controller is driving each component. A calendar rotation is applied to certain redundant components, thus distributing the wear among each. In a second common arrangement, redundant components share the load on a common task. When one fails in this arrangement, the healthy component is responsible for providing the power to compensate for the failed component, usually at a reduced level of performance. Another common practice is for components to operate in a master slave arrangement where one of the components is present simply to provide a backup should the primary system fail.

The goal for these analytics is not only to identify issues in the motor, but issues in the intermediate components driven directly by the CMSC. The smart power distribution system is responsible for dynamically assigning tasks to each CMSC based on priority and system availability. This information can be useful in fault isolation. For example, if an issue is present with a starter generator, and that issue is prevalent when driven by multiple motor controllers, we now have evidence that the issue exists within the starter generator and not within the power generation system. An ACMF report is designed to capture data during a specific routine during a flight leg. Understanding which routine the ACMF report is designed to monitor provides important information regarding which CMSCs will be providing parametric data within that report. Combining this information with the CMSC operational mode flag will result in the isolation of data which was captured when a specific CMSC was driving a specific component. More importantly, this prevents the ASHM system from producing false alarms on irrelevant data. This method of data fusion was prevalent throughout the development of analytics for the CMSC subsystem on the Boeing 787. After filtering the data, simple statistical analyses were performed on the data to provide parametric data and binary flags. An example of these statistical analyses on a filtered data set, is shown in Figure 6. The parametric data derived from the ACMF reports resulted in the application of the following statistical features:

- Statistical Electrical Power Calculations
- Amount of time the CMSC powered a specific component
- Amount of time required for a component to reach a speed increment
- Oil Temperature Trending: Also useful for trending engine temperature

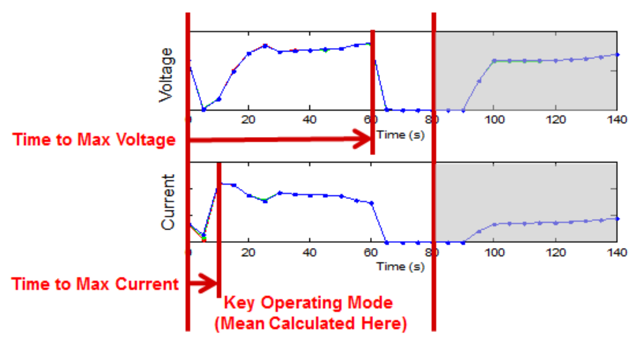

Figure 6: Time Series Plots of Parametric Data from Target Components

Enumerated flags were also generated from this filtered report content. The focus was on certain key operating routines/conditions which were highlighted for investigation. Utilizing the filtered data, enumerated flags were incorporated to track the frequency at which these key operating conditions were performed on the aircraft. The example shown in Figure 7: Sample Data Used for Development of Enumerated FlagsFigure 7 highlights a flag which would identify when a specific component was operating at High, Limited, or Low power. The assumption was made that high power operation exposed the component to high load and thus more demanding operation.

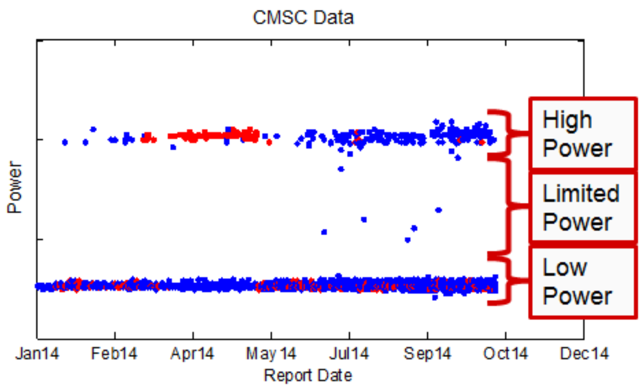

Figure 7: Sample Data Used for Development of Enumerated Flags

Other detection logic was put in place to recognize maintenance events and environmental conditions that result in accelerated system life usage. These flags appear as calculated parameters and can be plotted in ASHM. This gives maintainers and engineers the ability to correlate observed usage patterns with field failures.

7. CONCLUSION

The ASHM software tool provides enhanced health monitoring capability for the commercial aircraft fleet. The development and implementation of custom analytic modules unlocks the full potential of the recorded ACMF data for diagnostics and prognostics. By employing a multi-step analytic development and validation strategy, software development is accelerated while ensuring the quality and accuracy of the actionable condition information provided to the fleet stakeholders. Care must be taken to create robust algorithms that recognize irregularities in the report data, selectively filter applicable data, and ignore any potentially spurious or errant output. Effective validation requires communication between all interested parties to ensure that the high value system health ground truth information is documented and included in the technology assessment.

A case study for the Boeing 787® Common Motor Start Controller subsystem illustrates how trending and alerting on raw data alone is not enough for effective aircraft system monitoring. A complete understanding of system connectivity and operation states is required. The analytics developed for the CMSC subsystem follow three basic steps. First, they filter and down select the data. Each analytic is designed to target a specific system component. This filtering is achieved through the use of the operational mode flags, system connectivity information, and ACMF report information. Second, statistical power features, oil monitoring, temperature monitoring, and speed monitoring parameters are calculated from the data. These calculated features provide a summary of the target component during the report time period. Third, enumerated flags are generated from reported parameters or calculated features. These flags act to communicate relevant events to the operator such as maintenance procedures, abnormal or damaging environmental conditions and differentiate these noteworthy events from standard operation.

These calculated features can be observed and plotted over time in ASHM to provide insight into fleet trends or individual aircraft trends. Maintainers and engineers can then assess this data to indict specific LRUs and then proactively plan maintenance avoiding service interruptions.

The ASHM software tool is a key enabling technology for condition based maintenance of commercial aircraft, and provides the capability needed to reduce the rate of service interruptions and improve field service logistics operations.

NOMENCLATURE

ACMF	Aircraft Condition Monitoring Function
API	Application Program Interface
APU	Auxiliary Power Unit
ASHM	Aircraft System Health Management
BIT	Built In Test
CAC	Cabin Air Compressor
CMSC	Common Motor Start Controller
FDE	Flight Deck Effect
G	Generator
LRU	Line-replaceable Unit
PHM	Prognostics and Health Management
SI	Service Interruption
SQL	Structured Query Language
UTC	United Technologies Corporation

REFERENCES

Ramohalli, G., “The Honeywell On-Board Diagnostic & Maintenance System for the Boeing 777,” IEEE/AIAA 11th Digital Avionics Systems Conference, 1992.

Peters, G., Mosher, M., and Carlson, C., “Bringing the Aircraft System Health Management to Life Using the Informatica Tool Suite,” Annual Conference of the Prognostics and Health Management Society, 2013.

Byington, C., Roemer, M., Kalgren, P., and Vachtsevanos, G., “Verification and Validation of Diagnostic/Prognostic Algorithms,” 59th Annual Meeting of the Society for Machinery Failure Prevention Technology (MFPT), 2005.

Najjar, N., Hare, J., D'Orlando, P., Leaper, G. et al., "Heat Exchanger Fouling Diagnosis for an Aircraft Air-Conditioning System," SAE Technical Paper 2013-01-2250, 2013, doi:10.4271/2013-01-2250.

Wheeler, P., Bozhko, S. (2014) “The More Electric Aircraft: Technolgy and Challenges.” *IEEE Electrification Magazine*, December 2014 7. doi:10.1109/MELE. 2014.2360720