DOT/FAA/TC-16/29

Federal Aviation Administration
William J. Hughes Technical Center
Aviation Research Division
Atlantic City International Airport
New Jersey 08405

Interfaces for Sensory Prognostics and Management Systems

October 2016

Final Report

U.S. Department of Transportation
Federal Aviation Administration

NOTICE

Technical Report Documentation Page

1. Report No. DOT/FAA/TC-16/29	2. Government Accession No.	3. Recipient's Catalog No.
4. Title and Subtitle INTERFACES FOR SENSORY PROGNOSTICS AND MANAGEMENT SYSTEMS		5. Report Date October 2016
		6. Performing Organization Code
7. Author(s) Raj Bharadwaj and Darryl Busch		8. Performing Organization Report No.
9. Performing Organization Name and Address Honeywell Aerospace – Advanced Technology 1985 Douglas Drive N, Golden Valley, MN 5542		10. Work Unit No. (TRAIS)
		11. Contract or Grant No. DTFACT-12-C-00010
12. Sponsoring Agency Name and Address U.S. Department of Transportation Federal Aviation Administration William J. Hughes Technical Center Aviation Research Division Atlantic City International Airport, NJ 08405		13. Type of Report and Period Covered Final Report June 2012 – January 2016
		14. Sponsoring Agency Code

15. Supplementary Notes

The Federal Aviation Administration Aviation William J. Hughes Research Division COR was Michael Walz.

16. Abstract

Honeywell’s Sensory Prognostics and Management System (SPMS) is an ensemble of sensors, data capture trigger conditions, signal processing algorithms for feature extraction, prognostic reasoning to establish fault hypotheses, and maintenance reasoners to recommend an action. SPMS is envisioned to provide a look-ahead prognostic feature to the baseline diagnostic function that exists in today’s baseline aircraft. Within this program, the authors demonstrated prognostics use cases covering three aircraft subsystems: propulsion engines, auxiliary power units (APUs), and valves. The developed SPMS is shown to be effective at improving APU and valve prognostics. The engine SPMS shows promising results for a limited number of available cases. In addition, it is also shown that to implement successful SPMS for selected failure modes, standardizations are needed for collecting the sensor data during operation; recording on-wing field observations and maintenance actions; and recording the remove and repair findings. This studied approach provides the ability to trace health/degradation of a given asset using sensor data and correlates the improvements in health to appropriate on-wing repairs and replacements of the faulty parts. In addition, the collected data provide a means to evaluate prognostics accuracy and repair effectiveness.

17. Key Words Diagnostics, Prognostics, Health management, Sensor, Auxiliary power unit, engine, valve, Sensory prognostics and management system	18. Distribution Statement This document is available to the U.S. public through the National Technical Information Service (NTIS), Springfield, Virginia 22161. This document is also available from the Federal Aviation Administration William J. Hughes Technical Center at actlibrary.tc.faa.gov.

19. Security Classif. (of this report) Unclassified	20. Security Classif. (of this page) Unclassified	21. No. of Pages 179	22. Price

Form DOT F 1700.7 (8-72) Reproduction of completed page authorized

ACKNOWLEDGEMENTS

The authors would like to thank Michael Walz of the Software and Electronics Section, ANG-E271 FAA for his technical expertise and direction throughout the course of the program. The authors would also like to acknowledge numerous contributions from Honeywell staff supporting the work under this program. In particular, the authors wish to acknowledge Onder Uluyol, of the CBM team, for his technical discussions on valves, and Howie Wiebold for his many contributions. Further, the authors would like to thank Hal Voges for his technical feedback and support, Danny Turner for helping with all contract-related issues, and Savit Boyd for her patience and support on program management.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

LIST OF ACRONYMS

ACMS	Aircraft Condition Monitoring System
APU	Auxiliary power unit
ARP	Aerospace recommended practices
BIT	Built-in test
CA-CBM	Core Architecture-Condition-Based Maintenance
CBM	Condition-based maintenance
CI	Condition indicator
DP	Differential pressure
ECS	Environmental control system
EGT	Exhaust gas temperature
ECU	Engine control unit
EHM	Engine health monitoring
ETOPS	Extended time on partial systems
FADEC	Full Authority Digital Electronic Control
FC	Fault condition
FCU	Fuel control unit
FDAMS	Flight Data Acquisition and Management System
FF	Fuel flow
HI	Health indicator
HMA	Hydro-mechanical actuator
IBIT	Initiated built-in test
IGV	Inlet guide vane
IRM	Integrated reference model
IVHM	Integrated Vehicle Health Management
LDI	Loadable data image
LOWESS	Locally Weighted Scatterplot Smoothing
LRU	Line replaceable unit
MEA	More electric aircraft
MEL	Minimum equipment list
MES	Main engine start
MR	Maintenance report
MSA	Maintenance service agreement
MTBUR	Mean time between unscheduled removals
OEM	Original equipment manufacturer
OI	Operational interruption
PH	Flight phase
PIPS	Product In-Service Performance System
PTM	Performance trend monitoring
PTMD	Predictive Trend Monitoring and Diagnostics
PV	Prognostic vector
RMS	Root mean square
RTL	Ready to load
RUL	Remaining useful life
SCV	Starter control valve

SPMS	Sensory Prognostics and Management System
TAT	Total air temperature
TSB	Time series buffers
WF	Fuel flow sensor

EXECUTIVE SUMMARY

Honeywell's Sensory Prognostics and Management System (SPMS) is an ensemble of sensors, data capture trigger conditions, signal processing algorithms for feature extraction, prognostic reasoning to establish fault hypotheses, and maintenance reasoners to recommend an action. SPMS is envisioned to provide a look-ahead prognostic feature to the baseline diagnostic function that exists in today's baseline aircraft.

Within this program, four use cases were selected covering three aircraft subsystems: propulsion engines, auxiliary power units (APUs), and valves. These use cases are:

- Prediction of an impending fuel control unit (FCU) failure on a propulsion engine by monitoring engine behavior during startup and trending over many starts.
- Predictive trending of APUs using multivariate Bayesian techniques designed to track and disambiguate multiple forms of deterioration/failure throughout the lifecycle of the APU.
- Maintenance reasoning to provide a time-on-wing extension for an APU near the end of life. The maintenance action consists of replacing the APU fuel nozzles, which can become clogged, resulting in uneven burn and excessive heating (and is one of the causes of ongoing deterioration of the APU). Clogged nozzles are not the primary cause, so the extension is expected to be modest.
- Prediction of impending failure of a valve.

The first three use cases were chosen as potential enhancements to Honeywell's existing commercial services providing predictive trend monitoring, diagnostics, and prognostics for the propulsion engines and APUs. The fourth use case, valves, explores the relatively new environmental control system domain.

Honeywell's run-time Core Architecture-Condition-Based Maintenance framework was adopted as a means to systematically evaluate choices for SPMS applications. In this framework, any SPMS application is expressed as a sequence of five interfaces through which information is exchanged. The first year enumerated the choices available for each of these interfaces. The focus of the second year was to evaluate these choices and document the authors' understanding. The efforts in year three focused on completing the SPMS development, evaluating the performance, and recommending data standardization.

The APU use cases used historical data collected by Honeywell's Predictive Trend Monitoring and Diagnostics (PTMD) service for monitoring and trending the health of APUs in the field. PTMD has been in use for approximately 13 years and monitors over 1900 APUs of various models. The data consist of snapshots of APU parameters and built-in test results at several key points of a flight cycle, including start-up of the APU, and when the APU is used to power main engine start (MES). Maintenance data describing APU removals from wing and subsequent shop findings were also used. The propulsion engine use cases used historical data spanning 5 years beginning in 2002 that were collected by a data collection system Honeywell temporarily installed on 35 regional jets. The data consist of 182 parameters collected at 1 Hz or faster from various sensors installed on the engines, APU, flight management system, navigation system, bleed system, and landing system. Corresponding maintenance actions; parts replaced and repaired; and field service engineer observations were also used. The valve use case used simulated data because no historical data were available.

The significant findings from the project include the following:

- Hot-section deterioration is the reason for 59% of the APU removals because of failure; the deterioration is readily tracked by trending exhaust gas temperature (EGT) measured during MES.
- The multivariate Bayesian approach for monitoring APUs has promise but proved too difficult to solve on this project. The techniques of dynamic Bayesian networks and particle filters were explored. The main challenge is solving the hybrid problem of discrete faults and continuous degradations.
- Ten historical examples of on-wing nozzle changes for APUs were analyzed; no evidence was found that they provided any beneficial change in the degrading trend or extended on-wing time near the end of the APUs lifecycle.
- Fuel controller problems manifest during engine start, eventually leading to a “hot start.” The rate of change of engine speed (N2 derivative) and the rate of change of EGT derivative provide trendable indicators of this fault progression.
- Trending and diagnosing the fuel controller fault has some remaining challenges. The engine controller is designed to handle some degree of fuel valve and pressure regulation loss, which can mask the fuel controller fault. In addition, comparison of the fuel controller faults across multiple engines showed that there is no fixed threshold for the trendable indicators. Instead, thresholds may need to be defined with respect to a healthy baseline for each engine position.
- For the valve use case, capturing 1) max downstream manifold pressure and 2) minimum delta pressure across the valve provide sufficient information for generating a trendable condition indicator.
- The APU and engine analysis would have benefited from maintenance records for any on-wing service or repair of the engines and APUs, which were not available for this study. The on-wing maintenance records would have provided ground truth to help explain certain anomalous changes in trends (e.g., sudden improvements in EGT) that sometimes happen during APU lifecycles. Without this information, the authors were left with either accommodating the anomalies as normal APU behavior (resulting in dispersion in predictions) or attributing the behavior to a hypothesized failure mode and on-wing repair, but without much justification.
- To address this lack of on-wing data, the authors proposed machine-readable interfaces for recording on-wing and shop maintenance actions for APUs and propulsion engines.

1. INTRODUCTION

Honeywell's Sensory Prognostics and Management System (SPMS) technology is an open, scalable, and reusable approach for performing advanced diagnostics and prognostics on aircraft subsystems and other types of assets. The objective of this SPMS project is to:

- Formalize the interfaces for standardization of an open-architecture SPMS system.
- Identify any certification or regulatory issues that could hinder the deployment or reduce the benefit of an open-architecture SPMS system.
- Develop and demonstrate advanced diagnostic and prognostic features for inclusion in the architecture.
- Demonstrate prognostic reasoning for condition-based maintenance (CBM).
- Identify data and interfaces for standardization of prognostics development.

The SPMS study is divided into nine major tasks:

1. Establish program plan: Define metrics to be used in the SPMS study.
2. Identify high potential SPMS applications: Identify three major subsystems that would benefit from SPMS technology.
3. Define Core Architecture-Condition-Based Maintenance (CA-CBM) interface requirements: This task summarizes the SPMS interfaces for the selected subsystems. It also captures the high-level design of the prognostics and the maintenance reasoner that are used/developed in this program in subsequent tasks.
4. Define reasoner requirements: The designed reasoner implementation is completed during this task.
5. Design SPMS configurations: This task captures the reasoner test plan.
6. Test and evaluation of SPMS system: Formal tests and tabulation of SPMS results.
7. Technology demonstrations: The results of the SPMS study are summarized in this task.
8. Develop SPMS standards and roadmap: This task captures the roadmap for standardization of the SPMS technologies.
9. Analysis reporting of results and final deliverable.

The remainder of the report is organized as follows. The target systems chosen for this program and a means to measure the effectiveness of SPMS are described in sections 2 and 3, respectively. The rationale for selecting the target subsystems and available interface choices for designing an SPMS are described in sections 4 and 5, respectively. Section 6 describes the test plan and includes a detailed description of the data used for evaluating various SPMS design choices. The evaluation approach and results for the three SPMS subsystem applications are described in sections 7 and 8. Recommendations for data collection and standardization are covered in section 9.

2. TARGET AIRCRAFT AND SUBSYSTEMS

In this section, baseline aircraft and target systems for the SPMS study are discussed. A range of commercial and business jet aircraft were selected to address the applicability and efficacy of the SPMS application. The target platforms include single-aisle aircraft, such as the Airbus A319, A320, and Boeing B737, and regional jets, such as the Falcon 900, RJ85, and Gulfstream G150. The notations GXXX and A3XX are used to denote the family of Gulfstream and Airbus aircraft, respectively. The range of aircraft represents augmented aircraft that form the baseline of the SPMS study. The augmented aircraft bring together the richness of flight and simulation data that are used to evaluate the SPMS system.

The target aircraft subsystems are engines, auxiliary power units (APUs), and the mechanical perimeter, including the poppet valve used many places in the environmental control system (ECS). Data for demonstrating SPMS effectiveness come from these three subsystems installed on a variety of single-aisle and business aircraft, as shown in table 1.

Table 1. SPMS target systems to baseline aircraft mapping

Subsystems	A319/A320/B737	Falcon900	RJ85	GXXX
Engine TFE731/LF507		x (TFE731)	x (LF507)	
APU 131-9X	x			
Valve	x			x

The selection of these target systems is based on a combination of available data, aviation safety drivers, and maintenance drivers. The potential application of SPMS to these target subsystems is summarized in table 2. In the subsequent sections, the authors provide detailed descriptions of these target systems; evaluate potential SPMS applications with respect to cost and benefits; and define specific accuracy metrics that the authors plan to achieve to demonstrate the effectiveness of this SPMS program.

Table 2. SPMS subsystems, data sources, and metrics

Subsystem	Failure Modes	Safety/Maintenance Impact Examples	Prognostic Metrics	Data Source to Demonstrate Metrics
Engine	Fuel controller	In-flight shutdown, aircraft on ground, engine over temperature alarm	30 flights before in-flight shutdown/BIT 90% detection with 2–3 flight prognostic window	LF507 data reports
APU	Load compressor, fuel nozzles, inlet fouling	Time on wing extension: water-wash/nozzle replacements, borescope inspections	30 flights before auto shutdown	Historical fleet data
Valve	Valve regulation failures	Valve failure	90% detection within 100 ms	Honeywell simulation data

BIT = built-in test

3. TASK 1: SPMS EFFECTIVENESS METRICS

This section defines the metrics that are used to evaluate the SPMS. A successful SPMS includes advanced aircraft subsystem prognostics and a reasoner that determines the health of the asset with a very high level of confidence while working within the computation and communication constraints of current day aircraft.

A number of diagnostic and prognostic metrics have been reported in the literature (e.g., references 1–3) for well-circumscribed algorithms that apply to specific aircraft subsystems. Though this provides general guidelines for measuring the accuracy of SPMS, the practical realization is determined by factors such as computation, interface, and architecture metrics.

Existing metrics for evaluating fault detection, fault isolation, and prognostics schemes are directly applicable to SPMS. These metrics are described in a generic context here. Though these provide a good set of guidelines, practical and economic factors quickly narrow this list. These metrics measure accuracy of inferences; latency in making inferences; sensitivity to different fault and degradation conditions; and effectiveness of mitigation actions. The metrics discussions also narrow them down to specific aircraft systems with maintenance and safety impact.

There also are metrics, such as computation and interface metrics, unrelated to accuracy. These metrics arise from the fact that SPMS needs to support health management functions at the aircraft level, derive an overall aircraft health by analyzing evidence from various subsystems, and effectively communicate this information to the maintainer. The plan for this study is to evaluate the authors' approach with respect to these metrics, either by analysis or simulation studies.

The SPMS accuracy metrics define the quality of the system in terms of accuracy, sensitivity, confidence, and timeliness of the diagnostics and prognostics solution.

The accuracy metrics include:

- Diagnostic accuracy: This is defined as the accuracy of diagnostic conclusions by the SPMS. This is accomplished by comparing the final SPMS conclusion with the truth for faults scenarios. The diagnostics accuracy captures the following submetrics:

  – Probability of false alarm: The probability of signaling detection when no faults are present.
  – Probability of true detection: The probability of signaling detection of the correct/true fault or an ambiguity group that includes the true fault, given the presence of a fault.
  – Prognostic accuracy.
  – Prediction horizon: Time to failure from first detection of fault precursors.
  – Prediction confidence: Variability of prediction.

Table 3 shows the accuracy metrics. Note that some of the goals are application-dependent. This is especially true in the cases in which the application has economic benefit. The threshold in these cases is generally selected by the operator to maximize benefits given the operational and cost constraints. For this program, the thresholds were picked in consultation with the FAA.

Table 3. Accuracy goals

Category	Metrics	Goal
Diagnostics	Probability of false alarms Probability of true detects Probability of false detects Probability of isolation	5% 95% Based on application >85%
Prognostics	Prediction specificity Prediction horizon Prediction confidence	>80% Based on application >90%

SPMS can greatly alleviate diagnostics and prognostics accuracy problems by generating more meaningful condition indicators (CIs). Metrics for measuring the effectiveness of CIs vary. A commonly used metric is the early warning time window provided by a CI when compared with an existing fault code generated by a built-in test (BIT). The reduction in the number of troubleshooting steps starting from the CI, when compared with an existing fault code, is another commonly used metric. Often, the elimination of specialized ground-test equipment in the troubleshooting step outweighs the reduction in steps. With SPMS, the accuracy of CIs is included as a metric. However, the reduction in number of troubleshooting steps was not evaluated.

Specific accuracy metrics in this program are based on two key drivers: increasing time on wing and preventing unscheduled maintenance. Consistent with this, table 2 identifies aircraft target subsystems to demonstrate SPMS metrics. A combination of historical fleet data, simulation models, and test cell/bench data are used to derive SPMS performance metrics.

4. TASK 2: IDENTIFICATION OF HIGH-POTENTIAL APPLICATIONS

The application of SPMS and potential impact of SPMS on the three selected subsystems are summarized in table 4. The thought process for selecting these applications is described in the remainder of this section.

Table 4. High-potential applications of SPMS technology and target metrics

SPMS Technology	Metric	Current Metric (Without SPMS)	Target Metric With SPMS Technology	Achievable Target on This Program
Propulsion engine equipped with a FADEC communicating with an aircraft FDAMS recorder				
CI Generation	Probability of predicting an impending fuel control unit (line replaceable unit) problem five cycles before an uncommanded engine shutdown	0%	90%	75%
APU–Aircraft configured to send APU performance report to a ground station				
Predictive Trending	Accuracy of predicting APU removal based on deterioration	0%	90%	75%
Maintenance Reasoning	Temporary APU performance extension by a fuel nozzle remove and replace action	0%	15 cycles	10 cycles
Valve controlled by a digital controller				
Onboard CI generation	Probability of predicting an impending valve failure five flight cycles in advance	0%	90%	50%

FADEC = Full Authority Digital Electronic Control; FDAMS = Flight Data Acquisition and Management System

4.1 ENGINE

The baseline aircraft for engines is a business and regional jet, such as the Falcon 50, Falcon 900, Lear Jets, or Avro RJ 85. The SPMS application requirements for propulsion engines installed on these business and regional jets are relatively uniform.

Most of the baseline business jet aircrafts are equipped with two rear-mounted engines, whereas the regional jets are equipped with four on-wing engines. Figure 1 shows the Honeywell TFE731 and LF 507, which are from a family of geared turbofan engines commonly used on business and regional jet aircraft.

(a) (b)

Figure 1. The (a) Honeywell TFE 731 powers several business aircraft and the (b) Honeywell LF 507 powers the Avro RJ-85 regional jet

As the first high-bypass, fuel-efficient business aviation engine, the TFE731 engine has logged over 51 million operational hours and features a long history of proven reliability. The LF 507 is a geared turbofan engine produced by AlliedSignal and then Honeywell Aerospace. It was first certified as Lycoming ALF 502. The LF 507 is an improved higher-thrust engine used on the Avro regional jet (RJ-85).

4.1.1 SPMS Design

Dispatch availability is one of the key drivers for business aircraft. Most engine manufacturers offer a maintenance service agreement (MSA) that guarantees a minimum availability for aircraft customers. This study includes engines under a maintenance service plan. In this plan, the customers pay for the uptime guarantees.

Unscheduled engine removals negatively impact aircraft availability and cause operational interruptions (OIs). Two factors contribute to an OI: 1) no start rate, in which the engine fails to start and, therefore, the aircraft is unavailable, and 2) uncommanded auto shutdowns while the aircraft is on the ground or in flight. Though an in-flight engine shutdown presents a safety hazard, OI result in economic penalties, both of which can be impacted by using SPMS.

Data obtained from the reliability engineering group help to understand OI causes:

- High-time: This refers to overall deterioration of the engine. The extent of this deterioration is large enough to make the operations uneconomical.
- Core faults: The engine core includes the power plant, fans, and gearbox. Failures such as seal leaks, blade erosion, and combustor liner burn-through may cause an OI.
- Engine line replaceable unit (LRU): Current LRU fault detection is achieved through the use of BITs. Unfortunately, these BITs detect a failed LRU rather than a failing LRU and, therefore, may cause an OI.

SPMS can reduce OI by providing a prognostic window within which a maintenance action can be performed. Five SPMS technologies applicable to engines are described below:

- Performance trend monitoring (PTM): Thermodynamic efficiency of an engine can be calculated based on internal temperatures, pressure, load, and inlet conditions. Trending this efficiency over time is called performance trending; as the engine deteriorates, this trend decreases. By setting an appropriate lower-bound threshold, predictive algorithms can calculate estimated time-to-failure and improve maintenance planning.
- Mechanical systems monitoring: Mechanical systems refer to the bearings and gearbox. A gearbox is an integral part of engines that power the business aircraft, because these engines contain a low-pressure section to drive a high-bypass fan.
- Gas path analysis: Gas path analysis can improve fault isolation by detecting predefined patterns of abnormal behavior based on temperature, speed, pressure, and inlet conditions. Patterns could be developed and mapped to specific faults in the gas path components, such as compressor erosion, bleed leaks, combustor liners, and turbine wear.
- Usage-based lifing: Applicable for life-limited parts (blade tips and hubs), usage-based lifing measures the actual stress incurred by these components rather than simple cycle counting. This enables an extension between component removals and differential repairs without compromising engine safety.
- Onboard CI generation: Transients such as engine startup, acceleration, no-load to full-load, and shutdown provide rich time-series data for tracking the response of several LRUs, such as fuel controller, oil coolers, surge valves, sensors, and guide vanes. CIs consolidate this transient response to provide an indication of a failing LRU. These CIs are then downloaded to a ground station for further analysis.

4.1.2 Cost-Benefit Analysis

Unlike general transport, and similar to a military aircraft, business jets tend to operate in bursts. That is, a business aircraft may perform six consecutive flights over a 2-day period and then not be used, for example, for the next 2 months. Contrast this with a large transport aircraft that operates more or less on regular duty cycles. In addition, unlike military aircraft, business aircraft tend to not operate in severe weather. Therefore, the combination of a more predictable operating envelope and lower-duty cycles influences the cost benefit of SPMS applications.

Because most business jet engines do not accumulate usage hours as rapidly as those installed on a large transport aircraft, this study's reliability experts indicate a less favorable benefit for usage-based lifing for engines installed on business aircraft than for engines installed on a transport aircraft, such as the A320 or B737.
PTM is commonly offered by engine manufacturers to access general health of the engine. Data from this study's customer support engineers indicate a gradual increase in the accuracy of PTM and a steady decrease in OI caused by high-time. If the PTM was 100% accurate, OI caused by high-time would be zero.

An interesting metric tracked by engine manufacturers is the ratio of OI caused by engine core and OI caused by non-engine core. Figure 2 shows this OI ratio broken down by non-engine core LRUs. For example, the engine control unit (ECU) causes approximately 2.5 more OI than those caused by the engine core. Similarly, starters cause more than twice the OI than the engine core. The fuel control unit (FCU) causes 1.5 times more OI than the engine core.

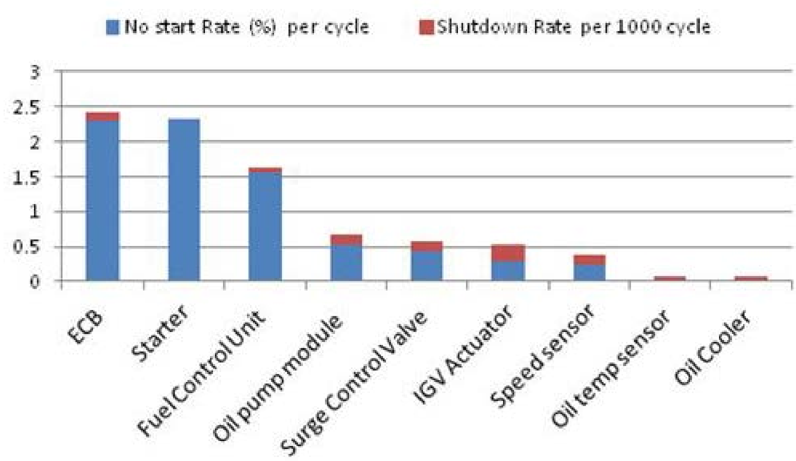

Figure 2. LRU-induced OI events per unit OI caused by engine core

Another consideration is that OI caused by a fault in the engine core takes more time to repair and is more expensive. Given the MSA, the length of the delay is more disruptive than the cost of the repair because most owners of business aircraft demand aircraft availability.

It is for this reason that the authors believe reducing LRU-induced OI is a favorable SPMS application for engines installed on business aircraft, given their sparse operational characteristics and prevalence of trend monitoring provided by most engine manufacturers. In addition, the review of LF 507 data also shows that detection of the onset of the FCU faults impacts the on-wing shutdown and no-start performance of the LF 507 engine.

4.1.3 SPMS Metrics

This section combines the outcome from the cost-benefit analysis described in section 4.1.2 with historic data to which the authors have access. The goal is to define an end-state metric for the SPMS technologies applied to a propulsion engine. The down selection process is shown in figure 3.

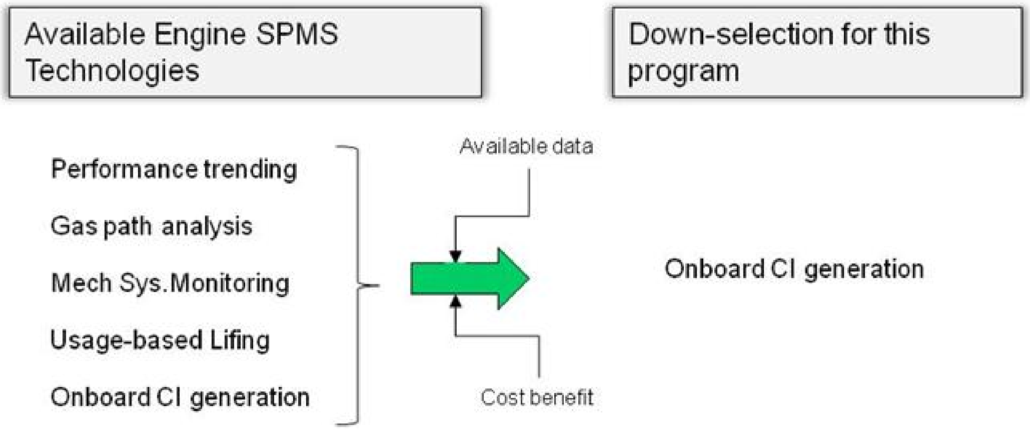

Figure 3. Down selection of SPMS technologies for propulsion engine installed on a business aircraft

Current LRU fault detection is achieved using BITs. These BITs implement simple threshold checks (i.e., hard faults) without taking a systems perspective of the propulsion system. As a result, when the BIT fails, it is too late and manifests as an engine shutdown or loss of power control at the propulsion-function level. The SPMS application planned to demonstrate in this program is “Onboard CI generation” for specific engine LRU.

To evaluate the effectiveness of SPMS, a term called “trendable CI” is defined. A trendable CI provides the maintainer an early indication before the corresponding BIT fails. For an LRU, this study’s experts believe the ideal prognostic window is approximately five engine cycles. There is some rationale for this window. The operational burst for a business aircraft is noted earlier. For the aircraft installed with the Honeywell TFE 731 engines, this operational burst tends to be 5–6 flight legs. Therefore, the authors believe a prognostic window of five engine cycles is a good threshold.

Based on the reliability data, the following engine LRUs are good targets for onboard CI generation:

- Starter: Generating CI requires access to electric current and voltage. Often these sensors are not available on smaller engines.
- ECU connectors: Loose connectors are one of the primary failure modes that drive ECU-related OI. Detecting this problem using reflectometry methods would necessitate installing additional circuitry within the connectors. The authors do not have access to these data for use in this program.
- FCU: This includes the fuel regulators, actuator, and position sensors. The authors have access to these data and, therefore, it makes a good SPMS demonstration candidate.

Table 5 summarizes the SPMS metrics for the propulsion engine that the authors plan to demonstrate in this program.

Table 5. SPMS technology metrics for propulsion engine

SPMS Technology	Metric	Current Metric (Without SPMS)	Target Metric With SPMS Technology	Achievable Target on this Program
CI Generation	Probability of predicting an impending FCU problem 5 cycles before an uncommanded engine shutdown	0%	90%	75%

4.2 APU

The baseline aircraft for APUs is mid-size transport planes such as the Airbus 320, 319 and the Boeing 737. The SPMS application requirements for an APU would be relatively uniform across these aircraft types.

4.2.1 General Description

Figure 4 shows the mounting of Honeywell’s 131-9A APU within the A320 aircraft tail cone. The main functions of the APU are:

- Provide bleed air for main engine start (MES).
- Provide bleed air for the ECS when the main engines are off.
- Provide electric power for the airplane.

Figure 4. Position of the APU within the A320 aircraft

Most APU models use a single-spool architecture wherein the power plant (the assembly that generates power by combusting fuel) and a load section are mounted on the same shaft. The load section is essentially a compressor that provides high-pressure bleed air for MES or the ECS. The power plant of the APU has an air intake, compressor, combustor, and turbine to extract energy from the expanding combustion gases. Typically, the bleed air exiting the load compressor is in the 50–60 psia pressure range and 180°–200 C temperature range. The APU is typically operated to provide a constant bleed flow and pressure. Demands imposed during the MES and ECS mode are automatically met by adjusting the fuel burn in the power plant/adjusting the guide vanes at the inlet of the load compressor.

4.2.2 SPMS Design

The two primary manufacturers of the APU, Honeywell and Hamilton Sundstrand, continue to refine the devices and are preparing many new technologies aimed at boosting reliability and “on-tail” time while decreasing emissions and fuel burn. Often, a life limit number (such as 12,000 hrs) is placed on the rotating components based on the probability of a high-energy component failure of the “hub” during operation. However, within this life, an APU is typically removed, refurbished, and returned to service multiple times because of blade erosion from normal wear/its failure to provide the above listed basic functions.

When an APU does fail, the main engines can often be started by a Ground Power Control Unit (also known as a ground cart) and, in some cases, the minimum equipment list (MEL) provides a means to permit dispatch. However, the first occurrence of the APU failure invariably leads to a delay and, in some cases, even a cancellation. To minimize such OIs, APU manufacturers track unscheduled removals. The mean time between unscheduled removals (MTBUR) is an important metric tracked by various airlines and is defined as:

$$\text{MTBUR} = \frac{\text{Total operating hours in a quarter}}{\text{Total number of unscheduled removals in the quarter}} \tag{1}$$

Several economic, competitive, and environmental reasons continue to drive an increase in MTBUR. SPMS can positively impact the APU MTBUR. The impact is first explained qualitatively and later one SPMS technology is picked to demonstrate the quantitative impact.

The primary reasons for removing an APU, based on the data obtained from this study's reliability engineering group, are shown in figure 5. An APU can fail to start the main engines because the APU may fail to start or it may not provide bleed air at the right pressure. Occasionally, the seals (which separate the oil and air flow) can leak, causing the contaminated bleed air to infiltrate the cabin or manifest as smoke on the tail pipe. In some cases, the exhaust gas temperature (EGT)—the amount of fuel that is not converted to useful work—may exceed the economic and environmental constraints, thereby forcing an auto shutdown.

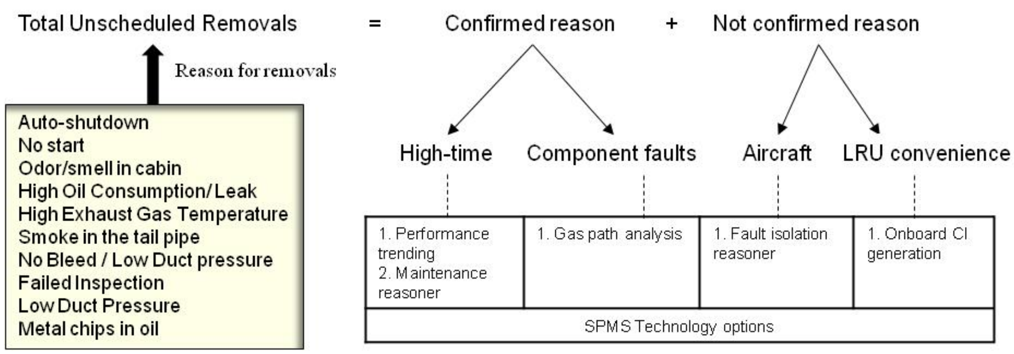

Figure 5. APU maintenance drivers and applicable SPMS technologies

APU removal from an aircraft can be confirmed by the presence of an actual component fault or general deterioration. Some removals are unjustified. For example, an aircraft fault can be confused with an APU failure/induced by an installation error. However, for an airline, both of these factors cause OIs, negatively impact the MTBUR, and provide the motivation for seeking advanced technologies (such as SPMS, better material, data management, etc.).

The root causes for APU removal are described below:

- High-time: An APU operates on the ground and ingests air that is not as clean as the air at altitude. This can lead to accelerated wear on turbine blades and rotating parts. For example, in hot climates with lots of pollution, particularly India and the Middle East, sulfur can build up on the blades, making them more prone to corrosion. High-time refers to general aging of the APU and overall deterioration, at which time it becomes inefficient to operate or causes an unsafe condition within the machine. Both of these conditions manifest as auto-shutdowns.
- Component faults: An APU is a small gas turbine engine and, therefore, inherits many of the mechanical complexities encountered in a large propulsion engine. High temperatures experienced by the turbine blades make them susceptible to high cycle fatigue (cracking and breakage) and debris that passes through the inlet screen can erode compressor tips and seals that keep the oil pathways separated from the air flow and lead to a leak. Though there is no requirement for a periodic hot-section inspection for the APUs considered in this project, airlines may detect internal damage during a random check. Seal leaks often manifest as odor in the cabin or smoke in the tail pipe.
- Aircraft: An APU interfaces with the aircraft bleed system through a load or starter control valve (SCV). It is generally believed by this study's reliability experts that SCV tends to stick because of condensation (note that bleed air from the APU is not conditioned). Often, this intermittent behavior is confused as an APU problem and one of the leading causes of unnecessary APU removal. Problems in the SCV actuation mechanism (solenoid) and duct leakages are two of the common aircraft problems that lead to unnecessary APU removal.
- LRU convenience: An LRU supports the power plant and load section within the APU. Some of these LRUs are designed to be removed and replaced without removing the APU from the tail cone. Common LRUs include the starter motor, oil cooler, sensors, and digital controller. Sometimes they operate intermittently and, therefore, faulty behavior often confuses the maintainer. Under extreme pressures to clear a BIT fault code, the maintainer may choose to remove the entire APU rather than spend time troubleshooting with existing tools. Such removals are unnecessary and negatively impact the MTBUR metric.

SPMS can impact the APU MTBUR metric by 1) increasing the numerator, which translates to keeping the APU longer before two consecutive removals, and/or 2) decreasing the denominator by detecting the underlying cause and planning a removal. Five such technologies listed in figure 5 are described below:

1. Performance trending: Thermodynamic efficiency of an APU can be readily calculated based on internal temperatures, pressure, load, and inlet conditions. Trending this efficiency over time is called performance trending; as the APU deteriorates, this trend decreases. By setting an appropriate lower-bound threshold, predictive algorithms can calculate estimated time-to-failure and improve maintenance planning.
2. Maintenance reasoner: Though overall APU deterioration is irreversible, the rate of deterioration can be slowed with the appropriate replacement of certain LRUs. A maintenance reasoner may provide a ranked list of possible actions based on the deterioration profile to provide additional operating hours and defer the APU removal to meet higher-priority organizational needs.
3. Gas path analysis: Gas path analysis can improve fault isolation by detecting predefined patterns of abnormal behavior based on temperature, speed, pressure, and inlet conditions. Patterns could be developed and mapped to seal wear, fuel pump failure, compressor erosion, and bleed leaks. Lack of available sensor measurements can greatly reduce the effectiveness of the application with respect to fault isolation.
4. Fault isolation reasoner: The APU, bleed duct, SCV, heat exchangers, and valves in the ECS are causally connected. Failure in one component cascades as failures in this connected system. A system-level diagnostic reasoner can leverage the observed symptoms and account for the propagation of fault effects across this connected system.
5. Onboard CI generation: Transients, such as APU startup, acceleration, no-load to full-load, and shutdown, provide rich time-series data for tracking the response of several LRU such as fuel controllers, oil coolers, surge valves, sensors, and guide vanes. CIs consolidate this transient response to provide an indication of a failing LRU. These CIs are then downloaded to a ground station for further analysis.

4.2.3 Cost-Benefit Analysis

An understanding of how APUs are serviced and maintained is needed to analyze the cost benefit of applying SPMS to APU. Almost all APUs installed on this study’s baseline airplanes have an MSA that transfers the maintenance cost risk and responsibilities to the APU manufacturer. The APU manufacturer, however, is paid based on the operating hours recorded by the asset.

An unscheduled removal typically causes an operational interrupt for the operator. Such delays can be avoided either by changing out the APU or by informing the ground crew about an appropriate manual action (if the APU is not on the MEL). In this case, a prediction horizon of 15 days provided by SPMS is desirable to avoid such interruptions. In some cases, SPMS needs to extend this horizon by up to 5 days to accommodate remote locations and outliers in the supply-chain logistics faced by the APU manufacturer.

Once a failing APU is removed and replaced and the potential delay is avoided, there is little benefit to isolate the cause of the failing APU for the airline operator. However, some preliminary indication of the root cause is beneficial to scope repair at the authorized shop. Based on discussions with Honeywell's field service engineers, the authors have discerned that SPMS needs to provide fault isolation at the following component level to benefit from the downstream repair processes:

- 1st stage turbine wheel
- Power plant aft bearing
- Load section seals
- Power plant (turbine) seals
- Load section shroud/impeller
- Load compressor bearing
- Compressor housing

A high degree of accuracy allows the downstream repair technician to skip troubleshooting steps and fix the APU within a shorter duration. A misclassification is not critical because the airline operator removes the APU regardless of the underlying fault, and the cost incurred by the MSA provider will not change. An isolation accuracy of 80% is generally acceptable to initiate permanent changes to the repair processes.

Because APUs have a non-negligible false removal rate, the second favorable application of SPMS is to minimize false removals. As described in section 4.2.2, the causes of this false removal rate are confusing an aircraft problem as an APU problem and the inability to isolate the root cause LRU in the APU. Figure 6 shows the interaction of an APU with other aircraft subsystems.

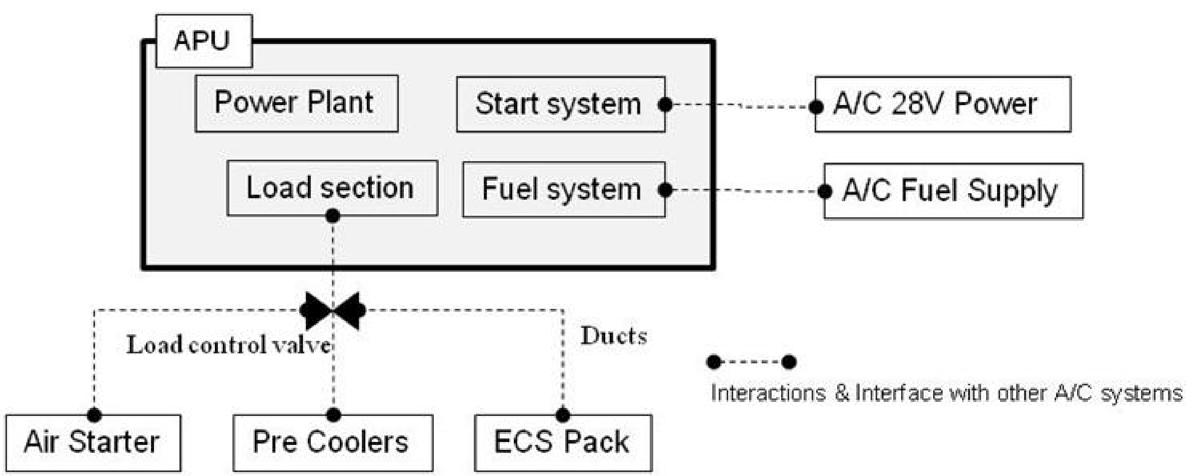

Figure 6. Interaction of an APU with aircraft A/C systems

SPMS needs to provide a clear isolation by reasoning the fault propagation between the following components to minimize the APU false removal rate:

- APU load section
- Aircraft 28 V power supply
- Air SCV
- Bleed air pre-cooler valves
- ECS pack
- Bleed duct leakages
- Aircraft fuel delivery

In general, an accuracy of 80% (correct 4 out of 5) is acceptable to initiate changes in the airline maintenance manuals.

An inability to recognize an LRU fault primarily arises from the lack of descriptive symptoms recorded by single variable threshold exceedance in typical BITs. The many-to-many symptom-fault mapping results in the technician taking the convenient approach: remove the APU. Furthermore, from an operator's point of view, the delay has already occurred. This combined with the MSA makes removing the APU more convenient for an operator than troubleshooting the problem. SPMS can greatly alleviate this problem by generating more meaningful CIs by analyzing multivariate time-series signals for the following APU LRUs:

- Oil cooler
- Load section delta-P sensor
- FCU
- Electronic control unit
- Starter motor
- Load control valve

The cost side of SPMS involves:

- Installing additional sensors primarily for health monitoring. This may result in adding wires and an increased power need.
- Collecting and recording more data from existing sensors. This may cascade as additional communication, computation, and storage needs.
- Development of first principles models to understand propagation effects/estimate hidden states in the APU.
- Development of health-monitoring algorithms for pattern recognition, reasoning, and feature extraction.
- Addition of processing power on the aircraft and specialized ground-support equipment for data transfer.
- Cost of changing maintenance manuals and service bulletins.

Figure 7 lists the desired parameters to develop various SPMS technologies for APUs. The cost benefit was primarily determined by assigning cost attributes to generate the required measurement/parameter. As shown in table 6, the cost of adding a sensor is 100 times more than using an existing sensor and 10 times more than collecting high-frequency data from an existing sensor. Consequently, the benefits from adding a new sensor would need to be 100-times more to justify the corresponding SPMS technology.

In summary, SPMS applications that can keep the APU on the airplane longer are favored from the manufacturer's point of view.

		SPMS Technologies				
		Increase planned removal			Decrease false removals	<-- Benefits
	Measurement/Parameter	PTM	MR	GPA	FIR	CI
Inlet	Ambient temperature TAT	x	x	x		x
	Inlet pressure P2	x	x	x		x
	Electrical load, GLA	x	x	x	x	
APU	Inlet guide valve position IGV	x	x	x	x	x
	IGV command signal					x
	Surge valve position SV					x
	SV command signal					x
	Bleed flow, WB	x	x	x	x	
	Parasitic load	x				
	Plenum delta-P		x	x		
	Exhaust gas temperature EGT	x	x	x		
	Air flow			x		
	Fuel flow, Wf			x		x
	Fuel flow demand signal					x
	Fuel temperature					x
	Core spool speed, N2			x		x
	Compressor discharge pressure			x		
	Bleed pressure, PB	x	x	x	x	
	Oil debris sensor			x		
	Oil Pressure			x		x
	Oil Level			x		x
	Oil Temperatures (3)			x		x
	Load section bearing vibration (1)					x
	Power plant bearing vibration (2)					x
Aircraft System	Start Control Valve SCV position				x	
	SCV solenoid current				x	
	Starter current				x	
	ECS Pack BIT				x	
	Precooler Temperature				x	
	Precooler Pressure				x	
	Bleed air regulator position					
	Bleed air regulator solenoid current				x	

Color Coding

- sensed and available
- sensed, need time series data
- additional sensor
- derived using models

Acronyms

- PTM: Performance Trending
- MR: Maintenance Reasoning
- GPA: Gas Path Analysis
- FIR: Fault Isolation Reasoner
- CI: Onboard Condition Indicators

x indicates requirements

Figure 7. Desired parameters to support SPMS of an APU

Table 6. Relative weights for requirements imposed for applying SPMS to an APU

	Black text	1x
	Blue text	10x
	Green text	25x
	Red text	100x

4.2.4 SPMS Metrics for APUs

This section combines the outcome from cost-benefit analysis described in section 4.2.3 with historic data to which this study has access. The goal is to define an end-state metric for the SPMS technologies applied to an APU. The down selection process is shown in figure 8.

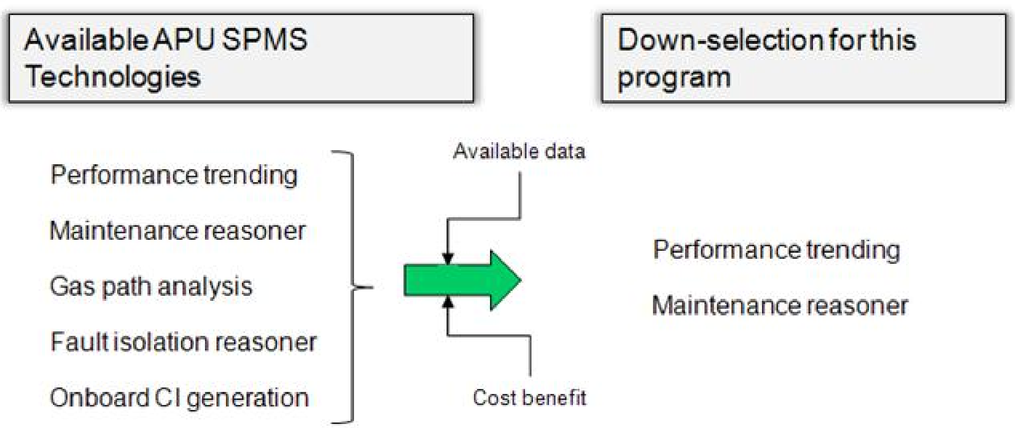

Figure 8. Down selection of SPMS technologies for APU

This is completed in four steps:

1. Define the metric.
2. Define a baseline value for the metric based on historic data.
3. Define a “to-be” value for the metric based on interviewing customer service engineers and leveraging their knowledge of potential business impact.
4. Based on the data in hand, modify the “to-be” metric to a metric that can achieve and demonstrate within the boundaries of this project.

The authors prefer to use the “operating cycle” as a unit of measure for several APU metrics. An APU operating cycle is defined as a sequence of: power-up → startup → idle → load →

shutdown → power down. Because the number of operating cycles per day depends on the airline, the authors prefer to use this unit of measure rather than calendar time.

4.2.4.1 Predictive Trending

Metric description: Over a 3-month interval, assume there were *N* unscheduled APU removals caused by "high-time." If predictive trending SPMS technology was developed for the entire APU fleet—such that *M* out of these *N* APU provide an early warning, with >80% probability at 30 cycles before the APU was actually removed—the accuracy metric α_{PTM} is defined as:

$$\alpha_{PTM} = \frac{M}{N} \times 100 \tag{2}$$

Suppose there were *P* APU that were falsely removed by the SPMS trending technology. The false alarm metric would be defined as:

$$\beta_{PTM} = \frac{P}{N} \times 100 \tag{3}$$

Current (baseline) metric value:

$$\alpha_{PTM} = 0\% \tag{4}$$

End-state for the metric based on discussions with Honeywell business leaders:

$$\alpha_{PTM} = 90\%, \beta_{PTM} < 5\% \tag{5}$$

Achievable target for the metric: Assuming a worldwide fleet reporting and the authors' access to data, the authors set:

$$\alpha_{PTM}^{achievable} = 75\%, \beta_{PTM}^{achievable} \leq 10\% \tag{6}$$

4.2.4.2 Maintenance Reasoning

Metric description: The maintenance reasoner SPMS technology can recommend a fuel nozzle replacement action as means to defer APU removal once its efficiency falls to near threshold and thereby extend life on wing. The metric the authors want to evaluate is the minimum deferral window, τ_{MR}, by applying the SPMS technology across the entire fleet.

Current (baseline) metric value:

$$\tau_{MR} = 0 \tag{7}$$

End-state for the metric based on discussions with Honeywell business leaders:

$$\tau_{MR} \geq 15 \text{ cycles} \quad (8)$$

Achievable target for the metric: Assuming a worldwide fleet reporting and the authors' access to data, the authors set:

$$\tau_{MR}^{achievable} = 10 \text{ cycles} \quad (9)$$

Table 7 summarizes the SPMS metrics for the APU.

Table 7. SPMS technology metrics for an APU

SPMS technology	Metric	Current metric (without SPMS)	Target metric with SPMS technology	Achievable target on this program
Predictive Trending	Accuracy of predicting APU removal based on deterioration	0%	90%	75%
Maintenance Reasoning	Defer APU removal by a fuel nozzle remove and replace action	0%	15 cycles	10 cycles

4.3 VALVE

The ECS of an aircraft provides air supply, cabin air conditioning, thermal control, cabin pressure control, and ventilation systems. Air is first compressed to higher pressure and temperature and then conditioned, a process in which excess moisture is removed and the air is heated or cooled to maintain a comfortable environment. The SPMS application requirements for valves will be relatively uniform across the business jet and medium-transport aircraft types. The most common type of valve is the regulation valve.

4.3.1 General Description

Figure 9 shows the three functional elements of a regulation valve system. It consists of a solenoid, a regulator that provides the reference pressure, and the valve that restricts the flow path using a plate or poppet. Despite their varied applications, the basic principles of a pneumatic regulating valve made by Honeywell or used on Honeywell engines remain the same. Specifically, a regulator provides the reference pressure; diaphragms and pistons convert pressure into force, which is opposed by a spring; and the valve opens and closes to close the force balance.

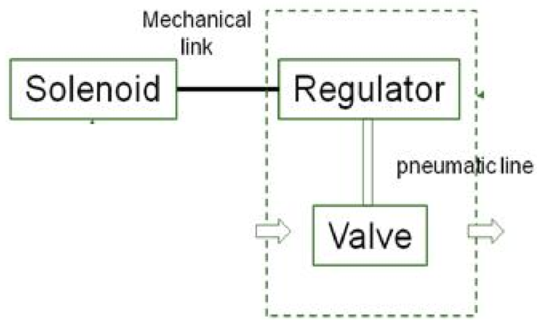

Figure 9. Functional elements of every inline regulation valve

An inline valve uses the "pressure tap" on the flow line to provide the regulation. The downstream pressure is connected under the diaphragm in such a way that the valve will automatically move to whatever position it needs to satisfy the force balance. This last statement is important for designing prognostics and health management for regulation valves. Measurement of the valve position is inconsequential for health monitoring because, by design, the valve position will "float" to meet the regulatory demands, signifying a healthy valve. Internal leaks open undesirable pathways inside the valve; this leads to a loss of regulation and eventually a complete loss of the valve function (i.e., it fails to open or close).

4.3.2 SPMS Design

The valve system is generally designed to meet the overall system reliability. Both the operators and the original equipment manufacturers (OEMs) track unscheduled failures by using the MTBUR metrics. The MTBUR goal for major assemblies is set at 16,000 hours. The goal helps the operators to schedule ECS component removal to coincide with major maintenance intervals pertaining to the aircraft.

The SPMS effort on the valve-driven system focuses on monitoring and detection technologies that address unscheduled removal of valve systems. A review of the reliability data reveals, for the ECS systems example, that the MTBUR is driven by the failure rate of valves. The latest valve systems integrate all of the cabin pressure control functions into one compact valve unit. It combines the three major system components—controls, sensors, and actuators—into a single valve that preserves the functionality while also minimizing direct maintenance costs for the operator through a reduction in components.

Figure 10 shows the key drivers for the valve.

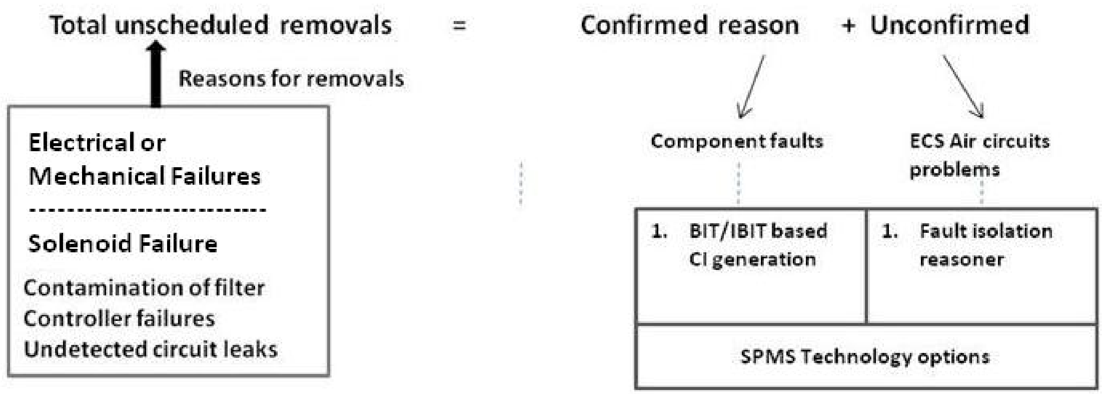

Figure 10. Poppet valve maintenance drivers and application of SPMS technologies

When a valve is removed from the aircraft, the reason for this action can be confirmed by the presence of an actual component fault or general deterioration. Some removals may have no justifiable reason and are classified as unconfirmed.

The root causes for valve removal include:

- Mechanical faults: The valves tend to fail open, closed, or chatter or become sluggish, thereby triggering ECS BITs and fault messages. The filters get clogged because of contaminants.
- Electrical faults: The solenoid mainly fails because of thermal and mechanical degradation.

The fielded systems do not perform any usage monitoring. However, valves and control assemblies can wear out faster in the presence of environmental contaminants like sand and particulates in the atmosphere.

SPMS can impact the MTBUR metric by 1) increasing the numerator, which translates to keeping the valve or control assemblies longer before two consecutive removals, and/or 2) decreasing the denominator by detecting the underlying cause and planning a removal. Two such technologies are described below as applicable to the ECS system:

1. Performance trending: Thermodynamic efficiency of an ECS system can be estimated based on internal/external temperatures, pressure, load estimates, and inlet conditions. The viability of system-level performance estimation and trending for valves was explored in the SPMS framework.

2. BITs/initiated built-in tests (IBITs)—IBIT-based onboard CI generation: BITs and IBITs can be used to generate the CI to aid the reasoning processes within the SPMS framework. The BIT or IBIT generation is done on-aircraft. The SPMS reasoners can be on-aircraft or on the ground.

4.3.3 Cost-Benefit Analysis

The current trend for the aircraft and ECS valves is to migrate toward a more electric system for increased reliability and efficiency. Over the last decade, most of the ECS valves are electrically driven; therefore, the areas of focus for the SPMS program are engine valves, wing anti-ice valves, and valve drive/controller monitoring. SPMS will cover the solenoid-driven poppet valves. Covering these valves with SPMS technologies also enables mitigation of long-term technology risk with electrical valves.

4.3.4 SPMS Metrics

This section combines the outcome from cost-benefit analysis described in the previous section with historic data to which the authors have access. The goal is to define an end-state metric for the SPMS technologies applied to valves.

Current LRU fault detection is achieved using BITs that are generally computed in the controller. The BIT implements simple threshold checks (i.e., hard faults) for controller and sensor faults without taking a systems perspective for the ECS or aircraft. As a result, when the BIT fails, it is too late and results in shutdown rather than a warning of impending failures. The SPMS application planned to be demonstrated in this program is “Onboard CI generation” for specific motor valve controller assembly faults by enhancing the BIT/IBIT from the valve. Table 8 shows the SPMS metrics for valves.

Table 8. SPMS technology metrics for ECS system

SPMS technology	Metric	Current metric (without SPMS)	Target metric with SPMS technology	Achievable target on this program
Onboard CI generation for valves	Probability of predicting an impending valve failure 5 flight cycles in advance of full functional failure	0	90%	50%

5. TASK 3: SYSTEM SPECIFICATIONS AND INTERFACE DESIGN

The SPMS architecture is based on Honeywell’s CA-CBM. Details of CA-CBM are described in Felke et al. [4]. Within the CA-CBM framework, an SPMS application implements one or more of the following functions:

- Measure
- Extract
- Interpret
- Act
- Interact

The target systems to be demonstrated in this program cover four of these functions. Each function has a well-defined set of interfaces in the CA-CBM framework. The authors first describe the five functions followed by a description of the interfaces. Next, choices available for the three target SPMS applications are listed. Table 9 shows the mapping of the interfaces to CA-CBM functions.

Table 9. CA-CBM functions

CA-CBM Functions					
		Measure	Extract	Interpret	Act
CA-CBM Functions (section 5.1.2)	Sensor interface	I			
	Data interface	O	I		
	Monitor interface		O	I	
	Health interface			O	I
	Action interface				O

I = Input, O = Output

5.1 SPMS ARCHITECTURE

5.1.1 Functional Blocks of SPMS

Within the CA-CBM architecture, a set of 5 functions define a run-time SPMS application. These 5 functions are shown in figure 11 and are described in this section.

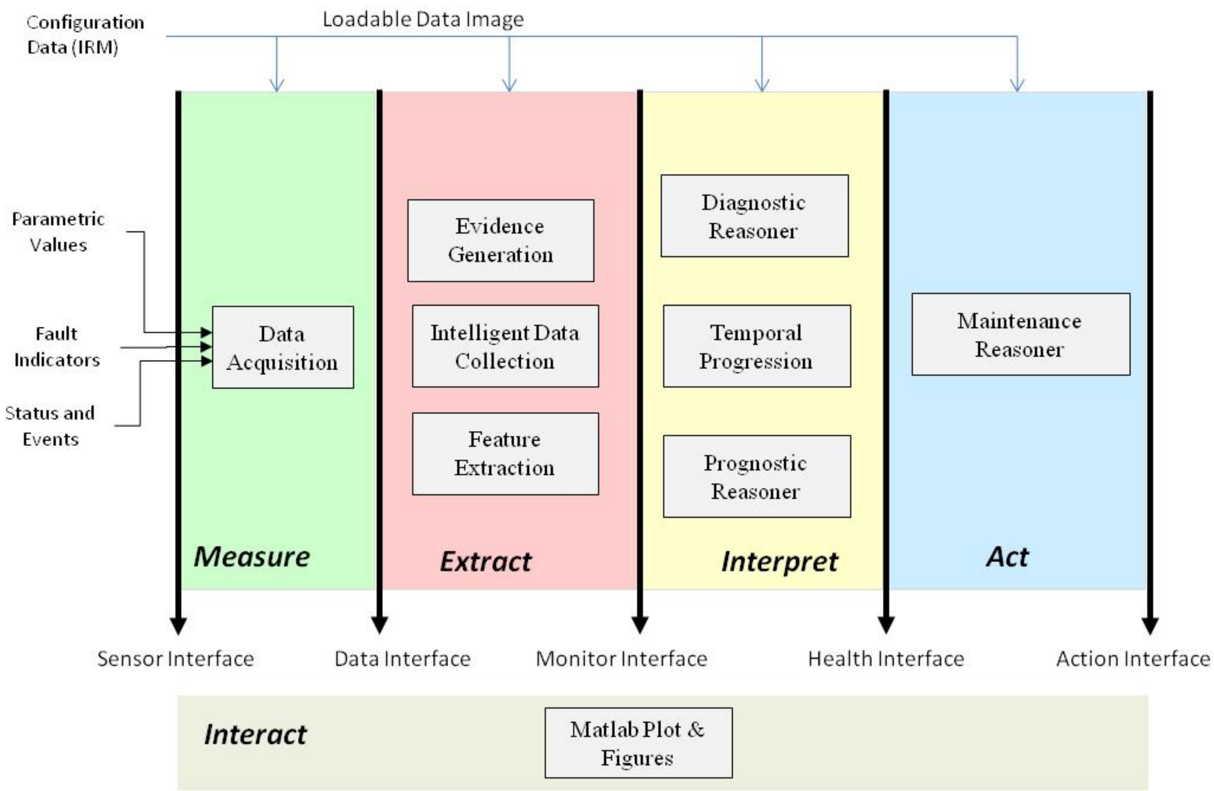

Figure 11. CA-CBM run-time modules to be exercised for demonstrating SPMS application for the three target subsystems

Table 10 summarizes the functional blocks the authors plan to exercise and enumerate for the three target systems scoped for this program. The five functions are described in the remainder of this section.

Table 10. SPMS functions exercised for the three target subsystems scoped in this program

	Target subsystems		
CA-CBM functions	Engines	APU	ECS
Measure	x	x	x
Extract	x	x	x
Interpret	x	x	
Act		x	
Interact	CA-CBM interact is not exercised in this program. Instead, interact is used for MATLAB® plots and display for the developed functionality		

The run-time SPMS is partitioned into five major functional areas that operate together as described below:

1. Measure: All functionality associated with sensing, measuring, reading of data from data bus messages, receipt of data from separately hosted portions of the CBM system, receipt of aircraft downlinks, and receipt of data from other systems.

2. Extract: All functionality associated with extracting the evidential information from the data provided by the measure function. This area also includes advanced algorithms to predict values of non-measured parameters based on the values of received data.

3. Interpret: All functionality associated with estimating the current and future health of the assets being monitored and determining the functional implications of that assessment. These functions convert evidence into actionable conclusions.

4. Act: All functionality associated with determining the required operational and maintenance actions that are appropriate to address the health status of the assets. This assessment identifies actions that are required now and those that will be required in the near future.

5. Interact with user: All functionalities for interacting with the user through a user interface. Interactions include displaying health data and status (text and graphics); displaying documents; controlling indicator lights and gauges; and receiving commands and data inputs from the user.

Each functional area consists of reusable software functions that can be adapted to support SPMS applications across a broad range of equipment types, operating environments, and maintenance practices. This is achieved using a data-driven approach. This means that rather than implementing the desired functionality in the software source code (i.e., hard coding), the functionality is specified in data structures that are referenced and interpreted by the runtime code, which cause it to perform the appropriate functions. This “reference data structure” that contains this configuration data is called the integrated reference model (IRM). Information contained in the IRM is provided to run-time SPMS applications using a loadable data image (LDI). This allows a complete reuse of the run-time software and specific aircraft-specific requirements are met by supplying it an appropriate LDI.

Reusable software modules that support run-time functions of SPMS are shown in figure 11.

The CBM core system run time functions are:

- Measure function: The sensor interface describes the input to this SPMS function and the data interface defines the outputs produced by this function. This function has one software module:
  - Data acquisition that receives incoming signal and decodes structured variables into a usable form.
- Extract function: Inputs to this SPMS function are defined by the data interface and the monitor interface defines the outputs generated by this function. Internally, the function consists of three software modules:
  - Intelligent data collection: Comprises data logging, trending, asset history tracking, and event-specific data capture functionality needed by most programs. Time series buffers (TSB) recorded by this function provide inputs to downstream computation or offline engineering analysis.
  - Feature extraction: Performs a sequence of analytic functions against streaming data, snapshot data, trending data, historic data, or event data to retain features of interest.
  - Evidence generation: This module applies an appropriate threshold function to determine if a CI is outside its acceptable range to indicate the presence of a not-normal condition and, therefore, evidence for health monitoring.
- Interpret function inputs and outputs to this SPMS function: These modules are defined by the monitor interface and health interface. This function consists of two software modules:
  - Diagnostic reasoner: Identifies likely fault causes that may be occurring in the aircraft based on all observed evidence. The output is a structure called a fault condition (FC), which asserts one or more hypotheses regarding prevailing faults in the system.
  - Temporal progression: Calculates the predictive behavior of health states based on future evolution of health indicators (HIs). This progression is expressed as a probability-and-time piecewise curve called a prognostic vector (PV).
  - Prognostic reasoner: Identifies likely fault causes that may be occurring in the aircraft based on all observed evidence including the prognostics in the form of a vector or temporal progression. The output is an FC.

- Act function within SPMS consists of the following two software modules: Inputs to this function are defined by the health interface; the outputs are described by the action interface:

  - Maintenance reasoner: Current fault indications and diagnostic conclusions are converted into required maintenance actions through a corrective action mapping.

- Interact function within SPMS consists of modules that communicate results from other CA-CBM functions to a human user/display device that interfaces with a pilot and maintainer crew. The exact output interface depends on the device displaying the information:

  - MATLAB® plots and figures: In this program, this module receives data from other CA-CBM functions and displays the information as MATLAB plots and figures.

5.1.2 SPMS Interfaces

Within the CA-CBM architecture, a set of five interfaces define an SPMS application:

- Sensor interface
- Data interface
- Monitor interface
- Health interface
- Action interface

Table 11 summarizes the interfaces the authors plan to exercise and enumerate for the three target systems scoped for this program. The five interfaces are described in the remainder of this section.

Table 11. SPMS interfaces exercised for the three target subsystems scoped in this program

	Target subsystems		
SPMS interfaces	Engines	APU	Valve
Sensor interface	x	x	x
Data interface	x	x	x
Monitor interface	x	x	x
Health interface	x	x	
Action interface		x	

The sensor interface defines how input values enter the SPMS application. The provider of this value could be a sensor or a digital bus. If the provider is a sensor, then the data acquisition module must convert an analog signal to an appropriate digital quantity and apply the necessary scaling to convert the quantity to an engineering value. If the provider is a digital bus, then the data acquisition module must decode the input byte stream (message) based on a predefined protocol and retain the engineering value contained in the message.

Typical choices for signal interface are:

- Analog sensors (e.g., pressure, temperature, flow, position).
- Digital bus communicating using the ARINC 429 protocol.
- Digital bus communicating using the ARINC 624 protocol.

In either case, the data acquisition module produces "sampled engineering values." Because this data forms the foundational element for SPMS applications, the data interface enforces critical requirements to ensure that downstream software modules can differentiate between usable and non-usable data. Furthermore, when the data are tagged as usable, they contain some common attributes. Data interface requirements are listed below:

- Sensor name
- Sampling interval–the time between updates to the value
- Minimum range value–the lower limit for the value
- Maximum range value–the upper limit for the value
- Engineering units

The monitor interface defines the outputs generated from the EXTRACT functions (intelligent data collection module, evidence generation module, and feature extraction module) within an SPMS application:

- TSB: This is the simplest form of output that the monitor interface needs to handle. The concept of "simple" is used with respect to its information content. A TSB is simply a collection of predefined sensor values collected when the aircraft is present in a particular operating mode. The mode itself is calculated by evaluating whether a specific trigger condition is present or absent. A TSB has minimal information content; it just signifies that these data may be interesting to analyze to support SPMS applications. The monitor interface allows a TSB to express this information by specifying the number of variables being buffered and the condition that triggered the collection of these data.
- CI: The summary based on sensor values. A CI provides more information than a TSB that the monitor interface needs to handle. A CI can be calculated by retaining predefined features from the sensor signal/applying an appropriate operator such as minimum, maximum, average, etc. A CI is more than a numeric value that provides a summary; it must also specify an acceptable region within which the CI should be considered normal.

The monitor interface allows a CI to express this information either by specifying an upper or lower threshold for the magnitude of the CI.

- HI: An HI provides indicting or exonerating evidence toward a failure mode set. That is, when the HI is indicting, it asserts the fact that one of the failure modes may be present in the aircraft. Conversely, when the HI is exonerating, it asserts the fact that none of the failure modes are present in the aircraft. Therefore, an HI provides the highest form of information from the extract function. The monitor interface allows an HI to express this information using either a detection probability (probability that HI is indicting when a specific failure mode is present) and/or false alarm probability (probability that the HI is indicting when none of the failure modes is present).

In the context of the SPMS, monitors represent deviations from nominal patterns. There are three preferred choices of "distance measures" available to SPMS to capture this deviation and, therefore, generate monitors. These are:

1. Euclidian distance measure: This is the simplest form of generating monitors. The values of the selected features provided by the data interface are recorded and subtracted from corresponding nominal feature values; each of the terms are squared; and an overall sum is calculated. If this sum is greater than a threshold, a monitor is generated. However, in closed loop systems, such as engines, the controller often compensates for incipient problems, making the calculated features correlated and small changes unobservable. Therefore, the Euclidian distance measure for generating monitors is not a good choice.
2. Mahalanobis distance measure: With this measure, monitors are generated based on transforming the original features to an Eigen space and calculating distance in this space. This transformation explicitly accounts for correlations and the corresponding statistics reflect deviations from a baseline correlation. Because the Mahalanobis distance measure uses the covariance information, it is better suited for observing feature changes in closed-loop controlled systems.
3. Max distance measure: This is the maximum of the absolute difference between every feature and its baseline value. Also called the L_∞ norm, this tends to emphasize the deviation of "single" features from its baseline value. Based on the physics of target systems, the authors recognize that these features are correlated and, therefore, picking "one" defeats the purpose of multivariate analysis. Therefore, using max distance to generate monitors is not a good choice.

The health interface defines the outputs generated from the INTERPRET functions (diagnostic reasoner module and temporal progression module) within an SPMS application:

- FC: An FC asserts the current diagnostic start of a target subsystem. The health interface allows an FC to express this information by specifying an ambiguity group of failure modes that may be occurring within the target system, a likelihood number that measures the strength of this failure mode occurring relative to a healthy subsystem. The interface also allows an FC to express future evidence (trending information) that it is expected given the current health of the subsystem.
- PV: A PV asserts the behavior of the target system in future time. The health interface allows a PV to express this information as a series of pairwise time and likelihood values. The likelihood here refers to the relative likelihood of the failure mode occurring in the target subsystem relative to a healthy subsystem. Though an FC expresses the current health of a target subsystem, the PV allows the health to be expressed in some future time.

The action interface defines the outputs generated from the ACT function (e.g., the output of the maintenance reasoner module) within an SPMS application:

- Maintenance report (MR): The MR provides an actionable conclusion for the target subsystem. The action interface allows an MR to express a prioritized list of tasks that a maintainer may perform based on inputs received from the INTERPRET function. The interface also allows an MR to express a time window within which this action needs to be performed and the system-level impact of not performing these tasks.

5.2 INTERFACE SPECIFICATIONS FOR THE TARGET SYSTEMS

Section 2 described the CA-CBM framework; it is comprised of five functions, each of which has well-defined interfaces for exchanging data and information. These functions can be executed on-aircraft (at the component-level, subsystem-level, or a central maintenance computer) or within a maintainer station or repair depot. Information exchange within these interfaces can take place using wires/in a wireless environment. The choices made regarding these functions and the interface depends on the SPMS application that needs to be built.

In a previous task (SPMS Task 2), three target subsystems (see table 12) were defined. Based on maintenance cost drivers and available data, an SPMS application for these subsystems was defined.

Table 12. SPMS applications planned for demo in this program

Subsystem	SPMS technology	Metric
Engine	CI Generation	Probability of predicting an impending FCU problem five cycles before an uncommanded engine shutdown
APU	Predictive Trending	Accuracy of predicting APU removal based on deterioration
	Maintenance Reasoning	Defer APU removal by a fuel nozzle remove and replace action
ECS	CI generation for ECS valve	Probability of predicting an impending trim-air valve failure five flight cycles in advance of full functional failure of ECS pack

5.2.1 Engines

The primary function of a propulsion engine is to propel the aircraft. In addition, the engines are the primary source of bleed air and electric power within the aircraft. Large propulsion engines are difficult to start with battery power and are, instead, started using a starter system, which can be a small electric motor or air-turbine starter (pneumatic). Irrespective of the source that provides the initial torque, the engine must sustain itself by burning appropriate fuel and maintaining the right combustion conditions throughout the flight.

Startup of an engine is a complex operation and includes several milestones:

1. Engaging the starter (electric motor/air-turbine starter)
2. Providing the initial torque
3. Energizing the igniter
4. Introducing fuel at the right time
5. Auto-combustion
6. Continuously metering fuel to the combustion chamber to follow an acceleration schedule until the engine reaches its idling speed

Engine startup is a critical phase for making dispatch decisions. An engine that fails to start definitely results in delays (engine manufacturers recommend a minimum cooling time between successive starts) and, in some cases (e.g., multiple failed starts), can ground the airplane. The cost-benefit analysis indicates that a good SPMS application can detect impending failures that could avoid such expensive disruptions.

The cost-benefit analysis indicated that disruptions caused by engine no-start events are expensive. Therefore, the SPMS application for engines is focused to provide a probability of an

impending FCU problem five cycles before an uncommanded engine shutdown value either during a ground start or an in-flight engine start.
Most modern large civil aircraft use an Aircraft Condition Monitoring System (ACMS) to acquire the data for engine health monitoring (EHM). This system captures three types of reports:

1. Snapshot reports: The sensor data are captured and collected in a small report during takeoff, during climb, and a few times when the airplane is in cruise.
2. Summary reports: Usually produced at the end of the flight; they collect key statistics such as maximum and minimum values during the entire flight.
3. Exceedance reports: Triggered by unusual engine conditions. These reports contain time-series data of key parameters.

The software mechanism for generating all of these reports is well-established and parameterized. Most engine OEMs provide offline configuration tools that allow an end user to select from a pre-existing list of reports. Furthermore, the sensors installed on the engine are multi-purpose; they are used to control the engine, provide indications of engine operation to the pilot, and for EHM.

To support engine SPMS (including the startup example focused on in this project), a new type of report needs to be defined: a relaxation of the exceedance report. Essentially, this report records time-series data of key engine parameters whenever the engine enters a specific operating mode rather than an exceedance condition. Table 13 lists the sensors that should be part of a startup time series report to support the engine SPMS application.

Table 13. SPMS sensor interface for engines

SPMS “sensor interface” for engines	Sensor rate and frame duration	Justification
EGT sensor	≥10 Hz time-series data. 60 seconds data frame. One frame every engine cycle. Can tolerate 2% loss of data due to drop out. Frame start: engine command to start	Engine controller works with the data at 50 Hz. Startup algorithms are designed to handle limited data dropouts. Preliminary tests with engine data show a 10 Hz or greater rate is expected to be sufficient.
Engine shaft speed (the power generator core)		
Engine fuel flow sensor		
FADEC fuel command signal		
Inlet temperature sensor		
Inlet pressure sensor		
Inlet guide vane position sensor		

FADEC = Full Authority Digital Electronic Control

The prognostic algorithm in SPMS is looking for specific patterns in this time-series data described mathematically using features. This summarization of time-series data into features can be considered a modification of the existing summary reports. However, the operations would need to be expanded from min/max/average to more complex mathematical functions such as Eigen values, singular values, or wavelet coefficients. Table 14 lists the features that should be part of a startup summary report to support the engine SPMS application.

Table 14. SPMS monitor interface for engines

SPMS "monitor interface" for engines	Data format
Starter deviation distance measure	Four numbers from each engine start. Non-dimension numbers, each with an acceptance region when no maintenance action is needed
Igniter deviation distance measure	
Fuel controller deviation distance measure	
Engine anomaly distance measure	

A departure of these features, called monitors, from a baseline value can be mapped to an incipient FC. This comparison may involve "more than," "less than," or "greater than" checks, such as statistical hypothesis testing—in which a confidence interval is based on a probability density function—or fuzzy threshold checking using a sigmoid or triangular membership function. Nevertheless, the result of this comparison operator is a new summary value for the engine cycle. Table 14lists these monitors to support the engine SPMS application.

The interfaces needed to support the engine SPMS are summarized in table 15. The remainder of the section describes the selection rationale, advantages, and disadvantages for each of the choices.

Table 15. Summary of interfaces for engines SPMS

Interface	Details	Choice for the required CA-CBM module
Sensor	≥10 Hz time-series data from sensors listed in table 14. One frame every engine start: Begin: engine commanded to start End: engine reaches idle	Engine FADEC
Data	Batch processing (non real time) to generate fixed number of features, once per engine start (see table 16)	Engine FADEC ACMS Ground computer
Monitors	Four CIs along with an acceptance range, once per engine start (see table 14)	Engine FADEC ACMS Ground-based computer
Health	Generate ambiguity group of engine components that may need maintenance. Once per engine cycle	Ground-based computer Central Maintenance Computer

FADEC = Full Authority Digital Electronic Control

Table 16. Engine SPMS interface options for the measure function

Where the "measure" function executes	Advantages	Disadvantages
In the FADEC	SPMS function is well-contained and could be provided as a standard upgrade process. Does not expose the intellectual property in the measure algorithms	The measure function may need a higher certification level. FADEC may need additional software partitioning to support SPMS
In the ACMS	Minimal changes to the FADEC and minimizes certification costs. Modification to the measure function (increasing/decreasing sampling frequency) does not alter the software criticality level	Increases the bandwidth on existing data buses by 2 orders of magnitude
In a ground-based computer	Propagating advances in feature extraction algorithms is easy	Download cost may be limited by speed and connectivity

During engine startup, the sequence is controlled based on external conditions, internal constraints, and proprietary control laws. Though all of the data needed for the SPMS are often available to the engine Full Authority Digital Electronic Control (FADEC), transmitting and sharing these "raw" time-series data represents a sensitive issue for all engine OEM. This represents a sensitive issue because raw sensor data can reveal internal engine states that are competitive and, in some cases, export controlled. Therefore, engine OEM are inclined to provide a mere "success/failure" status rather than detailed sensor data.

Conversely, monitors (like the ones listed in table 16) provide summarized data for prognostic decision making without revealing sensitive data. However, generating these monitors requires temporary storage of time-series data and executing algorithms that are more complex than simple threshold exceedance calculation. Because a FADEC is certified at level-A, extending its functions brings in additional software verification cost. Engine OEMs are willing to make these additions if there is sufficient economic payback.

Though SPMS uses data from sensors that are needed for the basic control of an engine, it does not interfere with the control laws. Almost all FADECs maintain a rolling buffer of high-frequency sensor data as supplemental reports when an engine exceedance event is recorded. Furthermore, because most modern engines have built-in EHM modules, SPMS merely requires creating new reports like the ones summarized in this section. Though the necessary software and hardware infrastructure exists in most modern FADECs, the updated functions would require stricter software assurance rigor. It is the authors' observation that making these changes to support SPMS does not significantly alter the software assurance rigor from what it is today. Engine OEMs are willing to make these additions if there is sufficient economic payback.

The primary benefit of providing a prognostic indicator is the avoidance of an operational disruption. Though there may be clear incentive for an airline operator to minimize airplane-on-ground time, this payback needs to propagate to the engine OEM for the necessary changes to be made.

5.2.1.1 Sensor Interface

The baseline aircraft consists of a turbine engine controlled by a FADEC. In the baseline aircraft, the FADEC interfaces with all engine-installed sensors and communicates to an ACMS through dedicated twisted pairs using the ARINC 429 protocol.

To define the SPMS interface criterion, it is important to understand a typical[1] startup sequence. This helps to identify the types of data that are needed and the rate at which they are needed so that meaningful numerical analysis can be performed by SPMS:

1. The starter on phase occurs when the starter is turned on and begins to rotate the engine. The maximum engine speed gradient occurs when the engine speed has its highest rate of change during startup, usually a few seconds after the starter is switched on.
2. Light-off phase occurs when ignition successfully completes and the combustor is able to sustain combustion. This causes a sudden temperature rise (controlled explosion) in the combustion chamber and the temperature of the gas exiting the engine increases sharply. A maximum temperature gradient follows the light-off.
3. Acceleration-to-self-sustain-speed occurs as the engine reaches and accelerates through a speed regime where the airflow through the front stages of the compressor becomes unstable and overcomes a rotating stall to push more air through the combustion chamber. This flow could be regulated by appropriately positioning the guide-vane based on ambient temperature (e.g., pushing large amounts of cold could kill the combustion process and, therefore, its flow needs to be modulated). This causes a temporary drop in the EGT as the air rushes through the chamber and exits the engine.

4. After auto-combustion is sustained, the engine FADEC starts to schedule fuel to maintain a specific speed schedule[2]. Close-loop control algorithms adjust fuel valve position to achieve smooth speed acceleration without violating temperature constraints.
5. When the engine reaches a predefined speed, the starter disengages to avoid damage. For an electric starter, the current is cut out. For a pneumatic air turbine starter, the bleed supply valve is commanded to close.

[1] Typical sea-level start. The spectrum includes cold start, hot start, and high-altitude starts.
[2] The selection of the control schedule depends on hot/cold/altitude/start.

6. The engine continues to accelerate as the rate of metered fuel increases and the combustion process provides a continuous stream of hot gases for expansion at the turbine and extracts mechanical energy. The EGT rises until the engine reaches its idling speed.
7. When the engine reaches its idling speed, the rate of fuel metering is set to zero and the EGT stabilizes at a steady value after a few seconds of lag.
8. The FADEC monitors the dwell time of speed, fuel, and EGT and sends a signal to the cockpit indicating a successful engine start.

The engine shaft speed, abbreviated as N2, is the primary indicator of an accelerating engine. Though the fuel flow sensor (WF) is the primary indicator that fuel is flowing into the combustion chamber, the actual combustion process can only be monitored through the measurement of the EGT.

Therefore, interfacing with the N2, EGT, and WF sensors is a primary requirement for SPMS to generate the necessary CIs that map to failing components, which may cause an engine no-start FC. To calculate various milestones numerically, this interface must provide instantaneous and synchronized values for all three of these sensors during the entire startup period.

It was also noted that numeric values of, for example, peak temperatures and fuel speed depend on the ambient temperature as it rushes through the combustion chamber during the acceleration-to-self-sustain-speed phase. Therefore, SPMS needs to interface with appropriate sensors that provide this inlet conditions to minimize this incorrect recommendation. Inlet conditions include temperature, pressure, and guide-vane positions that control the volume of air. Furthermore, the command signal sent by the FADEC to the fuel-metering value indicates how the acceleration process is being controlled.

There are no loads on the engine during that startup mode. Therefore, SPMS needs no interfaces to measure electrical loads or bleed loads. The SPMS sensor interface is summarized in table 14.

All sensors listed in table 14 are needed for controlling the engine during its startup phase and, therefore, data from these sensors are digitized and available to the engine FADEC. This is the primary advantage of this interface: no new sensors are needed.

Because the baseline aircraft consists of an engine with a FADEC, SPMS needs to merely access and retrieve the value. This is known as the data interface, and the choice for this is described section 5.2.1.2.

5.2.1.2 Data Interface

The following observations can readily be made based on the description of engine start:

- The time when certain milestones are achieved is important. Achieving or not achieving these milestones at predefined times provides valuable indicators for SPMS. Therefore, the interface must provide time-synchronized values. Otherwise, SPMS would incur some error.
- The definition of a milestone varies. It can be as simple as a sensor value crossing a threshold or looking at drops or maximum values. This implies that the SPMS interface must provide enough samples to derive numerically robust features.
- Algorithm development using historical data [5] indicates that a shift of as little as ~3 seconds in the timing of milestones differentiates a normal engine from a degraded engine. To numerically conduct these comparisons with 95% confidence, the gradual degradation in the sensor's signal-to-noise ratio as the engine ages needs to be taken into account. Based on earlier work [6], a sampling frequency requirement of no less than 10 Hz on the interface is recommended for providing these data to the SPMS.
- The startup of the engine typically begins when the engine is commanded to start and ends when the FADEC sends a signal to the cockpit. The SPMS interface must be able to recognize this "start" and "end" for the data collection process.
- As the cold air rushes through the combustion chamber, numeric values of temperature peak depending on ambient conditions. Therefore, SPMS needs to interface with sensors that can provide this inlet condition information.

To calculate various milestones numerically, this interface must provide time-series data for all three sensors during the entire startup period (i.e., the data collection needs to start when the engine is commanded to start and the FADEC indicates that the engine is in its idling state).

Because the baseline aircraft consists of an engine with a FADEC already generating data, SPMS needs to merely access and retrieve the value. This is the primary advantage of this interface.

However, data from these sensors are not always available for recording. Often, engine OEMs provide "startup summary" reports that summarize predefined milestones during the startup phase. Therefore, to support SPMS, the existing interface between the engine FADEC and ACMS needs to be expanded to transmit and record time-series data during the entire startup phase.

To continue the analysis, an understanding of how SPMS uses the data provided by this interface is needed.
Consider an engine startup phase that lasts for 60 seconds. At 16 Hz sampling, the SPMS interface provides seven time-series vectors with $60 \times 16 = 960$ samples. That is, the "measure" function within SPMS provides 960 values. Within the "extract" CA-CBM function, a series of signal processing algorithms are applied to retain "interesting patterns." As an example, consider

the algorithm described by Kyusung et al. [5] that defines the six milestones in table 17 and retains specific parameters at each of these milestones.

Table 17. SPMS data interface for engines

Milestone	Parameter
Peak N2	Time, peak N2 value
Peak EGT	Time, peak EGT value, N2 value, WF value
Fuel enable	Time, N2 value, WF value, EGT value
Light off	Time, N2 value, WF value, EGT value
Peak EGT-Dot	Time, peak EGT value, N2 value, WF value
Idle	Time, N2 value, WF value, EGT value
Total = 22 features (or numbers) every engine start	

The extract function within SPMS provides exactly 22 data points for every engine start. The location for executing the "extract function" generating 22 features from the raw time-series sensor data can be in the 1) FADEC, 2) ACMS, or 3) in a ground-based computer. The advantages and disadvantages for each of these locations are listed in table 16.

The next module in SPMS is to generate evidence from the data provided by the measure function via the data interface defined in tables 14 and 17.

5.2.1.3 Monitor Interface

In the context of the engine SPMS, monitors represent deviations from nominal patterns. Figure 12 shows a monitor example that captures the deviation of the typical engine start (blue line) and the start profile that had issues with light-off (red line).

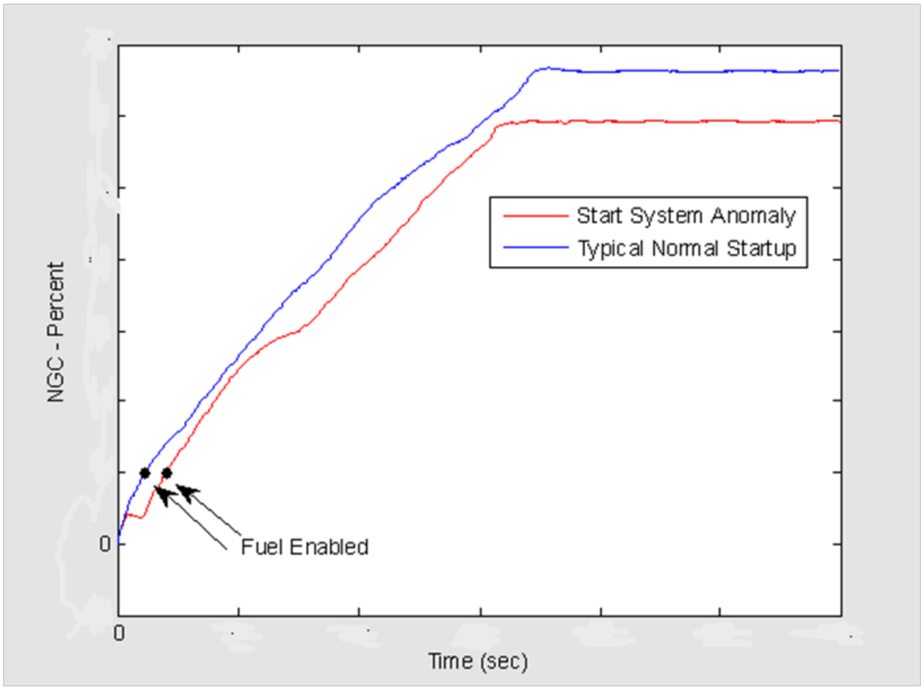

Figure 12. Example of an engine startup monitor

Mahalanobis distance is used for generating the engine startup monitor. This transformation explicitly accounts for correlations and the corresponding statistics reflect deviations from a baseline correlation. From the understanding of how engine starts and the features are calculated, the Mahalanobis distance measure is best suited for generating engine startup monitors from SPMS.

To support the engine SPMS function, irrespective of the actual mathematical method used for calculating the Mahalanobis distance, the monitor interface needs to provide one single number *CI* for every engine start and an acceptable region, such that $CI_L \leq CI \leq CL_U$ indicates that the underlying component must be considered healthy.

The numeric calculation for generating monitors happens in the "evidence generation" module within the CA-CBM framework. Clearly, the calculation of these monitors requires a detailed understanding of the engine startup process. Furthermore, establishing an acceptance range (CI_L, CI_U) requires subject matter expertise. There are three options for where this calculation can be executed. The advantages and disadvantages for each option are listed in table 18.

Table 18. Engine extract function options for SPMS

Where the "extract" function executes	Advantage	Disadvantage
In the FADEC	SPMS function is well-contained and could be provided as a standard upgrade process. Generates a new SPMS-rich report from the engine.	FADEC needs more computational power to do additional calculation. Furthermore, because the baseline for calculating deviation keeps changing, this would require periodic software modifications.
In the ACMS	80% reduction in the size of a download report per engine start. Reduction from raw data numbers to 22 features data point numbers per engine start.	Usually, computation power to do additional calculations is available; however, because the baseline for calculating deviation keeps changing, this would require periodic software modifications.
In a ground-based computer	One time, remote software updates.	There are no serious disadvantages.

5.2.1.4 Health Interface

Because the SPMS application is focused on detecting failures that may prevent an engine startup or cause an in-flight shutdown caused by fuel metering components, the health interface within SPMS needs to clearly list the engine LRU that is most likely going to cause this engine failure event. This determination is made by the diagnostic reasoner module within the Act CA-CBM function.

Calculations needed for the diagnostic reasoner module can be conducted on a central maintenance computer or ground-based computer. Table 19 lists the advantages and disadvantages for these choices.

Table 19. Engine interpret function options for SPMS

Where the "diagnostic reasoning" module executes	Advantage	Disadvantage
On the central maintenance computer	Can use data from multiple engines to reduce false alarms. Over time, SPMS output can be used to make dispatch decisions regarding the MEL.	Cost of developing onboard software and the rigor it entails.
In a ground-based computer	Deployed as a Web service and makes the output available whenever needed and to whomever needs it.	SPMS remains as a maintainer advisory system .

5.2.2 APU

An APU is a small turbine engine. Ambient air entering the APU is split at the plenum; part of the ambient air enters the main compressor and the remainder enters the load compressor. The flow of air through the load compressor is regulated by inlet guide vanes (IGVs). High-pressure air from the main compressor enters the combustor, where fuel is introduced using a series of annular nozzles. The air/fuel mixture is burned continuously and smoothly in the combustor. Hot combustion gases are expanded in the turbine, which drives the shaft. Exhaust air from the turbine is expelled through the aircraft exhaust. The useful shaft work is expended by the load compressor to provide bleed air and available at the generator to meet the electrical needs of the aircraft.

There are significant incentives to keep an APU operating until such time as it can no longer be operated. From the cost-benefit analysis, it was found that an SPMS application for an APU needs to identify conditions under which APU removal can be deferred by a fuel nozzle remove-replace maintenance action.

Turbine blades with protective coating are designed to withstand the high temperatures of the combustion gases. However, this protective coating degrades over time, leading to permanent damage. As the turbine blade tips erode, their ability to extract mechanical energy from the hot combustion gases decreases. Consequently, the gas that exits the turbine remains hot. This "leftover energy" is sensed by the EGT sensors. Therefore, an increasing EGT is a necessary indicator of hot-section deterioration. The cost-benefit analysis indicates that SPMS needs to do more than just detect this deterioration; it must also provide recommendations for changing the fuel nozzle, which could defer the APU removal, whenever appropriate.

Most modern APU use an ACMS to acquire the data for APU health monitoring. Similar to EHM, the ACMS captures three types of reports:

1. Snapshot
2. Summary
3. Exceedance

Like the engine, the software mechanism for generating all of these reports is well-established and parameterized. Most APU OEMs provide offline configuration tools that allow end users to select from a pre-existing list of reports.

To support APU SPMS (including the maintenance example focused on in this project), new reports do not need to be defined; instead, expand the "snapshot report" to include two power settings: one when the APU experiences the maximum load at MES, and the other at ready to load (RTL).

The interfaces needed to support the above-described engine SPMS are summarized in table 20. The remainder of the section describes the selection rationale and advantages and disadvantages for each of the choices available.

Table 20. Summary of interfaces for APU SPMS

Interface	Details	Choice for the required CA-CBM module
Sensor	10 Hz data of sensors listed in table 21	APU FADEC
Data	RTL snapshot: Values averaged for 1 second when APU transitions to RTL mode MES snapshot: Values averaged for 1 second when APU transitions to MES mode	APU FADEC
Monitor	Batch processing (non real-time) to generate the following evidence: • EGT margin • Bleed pressure margin	Ground-based computer
Health	Batch processing (non real-time) to generate EGT margin deterioration pattern	Ground-based computer, Central maintenance computer
Action	Recommend nozzle replacement to extend APU time-on-wing	Ground-based computer

Table 21. SPMS sensor interface for APU

SPMS sensor interface for APU	
EGT sensor	≥ 10 Hz time-series data
Inlet temperature sensor	
Inlet pressure sensor	
IGV position sensor	
Bleed pressure sensor	
Bleed flow sensor	
Generator load (current) sensor	
APU WF	
APU shaft speed sensor (optional)	

Though SPMS uses data from sensors that are needed for basic control of an engine, it does not interfere with the control laws. Almost all APU FADEC provides the two snapshot reports. Therefore, no updates are needed to create these reports.

The primary benefit of providing a prognostic indicator is the time-on-wing extension it can provide. Extending this time on-wing provides an opportunity window for an operator and APU OEM to plan this removal. As long as the payback is greater than the cost of downloading the report periodically, realizing the APU SPMS application presents no major challenges.

To provide this information, SPMS needs to interface with other sensors in the APU. This interface is described in section 5.2.2.1.

5.2.2.1 Sensor Interface

The EGT is the primary indicator of hot-section deterioration. This EGT interface is a primary requirement for SPMS. In addition to hot-section deterioration, inlet conditions cause the EGT sensor to increase (or decrease). Under these conditions, recommendation of a fuel nozzle change or an APU removal would be inappropriate. Therefore, SPMS needs to interface with appropriate sensors that provide these inlet conditions to "minimize" this incorrect recommendation. Inlet conditions include temperature, pressure, and guide-vane positions that control the volume of air.

There are two primary loads on the APU: 1) the generator, which provides electrical power when the aircraft is on the ground and 2) the load compressor, which provides high-pressure bleed air to start the propulsion engine. If these loads increase/decrease, the APU controller adjusts the fuel flow (FF) to meet the demand and maintain constant speed. Burning more/less fuel directly impacts the "energy of the exhaust gas" and, therefore, the EGT. To account for these factors, SPMS needs to interface with sensor that 1) measure the additional fuel that is being combusted and 2) measure the magnitude of the load to eliminate the corresponding effect on the EGT.

Generator demand can be inferred by a current sensor, whereas the bleed flow and pressure sensor provide demand imposed by the load compressor. The WF that measures the fuel being delivered to the APU combustion chamber provides information regarding how the APU is meeting these demands. If the APU is designed to operate at constant speed, the interface to the speed pickup sensor is optional for SPMS. These SPMS data interfaces are summarized in table 21.

All of the sensors listed in table 21 are needed for controlling the APU operations and, therefore, data from these sensors are digitized and available to the APU digital controller. Because the baseline aircraft consists of an APU with a FADEC, SPMS needs to merely access and retrieve the value. This is the primary advantage of this sensor interface.

Though hot-section deterioration occurs over several hundred operating hours, an APU can cause an operational disruption because of bearings, leaks, stuck values, frozen sensors, etc. For SPMS to address these problems, the sensor interface needs to include:

- Replacing the oil filter delta-pressure switch (on-off) with a differential pressure (DP_ sensor. With this sensor and an appropriate interface that can record delta-pressure values at 15-minute intervals over the entire APU operating mode, SPMS can trend the pressure loss across the filter and avoid disruptions caused by hot oil and low oil pressure.
- Replacing the oil quantity switch with a DP sensor. With this sensor and an appropriate interface that can record delta-pressure values at once when the APU enters the RTL mode, SPMS can trend the oil loss and avoid disruptions caused by hot oil and low oil pressure and quantity.
- Start motor current sensor. With this sensor and an appropriate interface that can record the current value at 50 Hz during the startup mode of the APU, SPMS can analyze these profiles and avoid disruptions by APU auto shutdowns and no-starts.

The rationale for selecting a “point in the APU cycle” when SPMS needs to retrieve these data is described in section 5.2.2.2.

5.2.2.2 Data Interface

During its normal[3] operations, an APU goes through a series of modes, which follow:

1. Startup: The APU is commanded to start and uses the aircraft's 28 V power supply or built-in batteries to accelerate the turbine from 0% speed to an auto-combustion speed. After combustion is sustained, the FADEC accelerates the APU to reach its target 100% speed.
2. RTL: The APU is idling and waiting for loads.
3. Electrical power: In this mode, the APU is producing useful electrical power by rotating the generator shaft through a gearbox.
4. MES: The APU is providing high-pressure bleed air to the air-turbine starter, which in turn is providing torque to start the main propulsion engine.
5. ECS: APU provides both pneumatic and electrical power. The bleed air is used for environment control rather than MES.
6. Shutdown: The FADEC decelerates the APU from 100% speed to 0% speed.

An APU cycle is defined as one transition from the startup to the shutdown mode. Because the FADEC is actively engaged in every APU operating mode, data from the sensors listed in table 21 are continuously available to it and, therefore, an appropriate time when SPMS can retrieve the data needs to be chosen.

Hot-section deterioration occurs slowly. Though an APU cycle may last for more than 1 hour (clock time) and the above-listed sensor data are available to the FADEC throughout this period (number of samples = clock time divided by the controller update rate), the change in the hot-section deterioration within a cycle is small. Therefore, a summary of the sensor data for each cycle is sufficient for SPMS to detect long-term hot-section deterioration and make its maintenance recommendations. Furthermore, this summary must be clearly timestamped so that SPMS can trend the deterioration chronologically. Therefore, SPMS needs to interface with a DATE sensor that can provide the APU cycle number and duration of the cycle.

[3] Abnormal mode includes auto shutdown or uncommanded shutdown.

The MES and RTL modes impose, respectively, the maximum and minimum load for the APU. These modes, therefore, provide operating conditions at which maximum and minimum hot-section deterioration occurs. Therefore, a summary generated during the MES mode and one generated during the RTL mode provide a statistically bounded observation for long-term trending. A simple average of samples over a predefined time interval provides a robust summary value from a signal-to-noise point of view. Because SPMS needs to capture the data at the peak load condition, the time period for calculating this average should not be more than 1 second:

- 1-second average after APU enters RTL mode
- 1-second average after APU enters MES mode

Because APU OEM provides a standard performance report that calculates these averages when the APU transitions to the MES and the RTL modes, this is a primary advantage of this interface. In addition, these reports also record the APU cycle number that helps to index these reports for long-term trending that SPMS needs to achieve its objectives.

The primary disadvantage of this SPMS interface manifests in the following two limitations:

1. Snapshot limitation: A summary report provides a single observation point for the entire APU cycle. Therefore, SPMS is not capable of detecting faults, such as oil leaks, fuel leaks, stuck valves, etc., that arise within one APU cycle.
2. Sensor limitation: SPMS relies on sensors primarily used for controlling the APU. Failures related to lubrication oil and bearings do not manifest in these sensors and cannot be detected by SPMS.

5.2.2.3 Monitor Interface

The basic requirement is that this monitor must be reflective of the hot-section deterioration and must trend monotonically as the magnitude of the deterioration increases. When the hot-section degrades, its ability to extract energy from the expanding gas decreases. This makes the EGT hotter than its baseline design value. This departure from the baseline design value is called EGT margin. Furthermore, because part of this energy is provided to the load compressor to provide bleed air, the flow and pressure of this bleed air can also decrease because of hot-section degradation.

Bleed air flow depends on the opening of the downstream bleed valve. With the selected sensor and the data interface, the position of this valve is unknown. Therefore, using the bleed-pressure margin as a monitor for hot-section deterioration is not very effective. Conversely, the EGT margin is a good monitor indicator for hot-section deterioration. However, to ensure that the degradation in the load compressor is not adding to the EGT margin and causing SPMS to wrongly conclude about hot-section deterioration, SPMS needs bleed-pressure margin to eliminate this incorrect or false alarm. In summary, though EGT margin is a primary monitor for hot-section deterioration, SPMS needs the bleed-pressure margin to do its reasoning and increase its confidence in its conclusions.

5.2.2.4 Health Interface

The next interface considered is the health interface. Given that the SPMS monitor interface provides EGT and bleed-pressure margins, an understanding of the primary factors that can cause these monitors to appear is needed; that is, from the probabilistic point of view, a list of failure modes that can trigger these monitors is needed. The top-four factors are fuel nozzles, inlet fouling, 1st stage turbine stations, and compressor erosion.

Calculations needed for the progression trending module can be completed on a central maintenance computer or ground-based computer. Table 22 lists the advantages and disadvantages of these choices.

Table 22. APU interpret function options for SPMS

Where the "diagnostic reasoning" module executes	Advantage	Disadvantage
On the central maintenance computer	Over time, SPMS output can be used to make dispatch decisions regarding MEL.	Progression trending of an APU needs several hundred hours of operational data. Maintaining this historical data on an aircraft may need additional memory and disk space. This also entails the additional efforts and rigor needed to develop onboard software.
On a ground-based computer	Deployed as a Web service and makes the output available whenever needed and to whoever needs it.	SPMS remains a maintainer advisory system.

5.2.2.5 Action Interface

After SPMS has isolated the underlying factor that best explains the observed monitors or symptoms, the Act interface needs to provide recommendations. To define this interface, one needs to understand activities a maintainer can do. If the problem is turbine stator vanes, the APU needs to be sent back to the OEM. The same action is applicable if the compressor has eroded. However, if the problem is fuel nozzles, a maintainer can replace them without sending the APU back to the OEM. If the root problem is inlet fouling, a maintainer can clean the inlet without sending the APU back to the OEM.

The cost-benefit analysis indicates that the optimal use of SPMS is to recommend a fuel nozzle replacement whenever applicable and, therefore, extend the APU time-on-wing before sending it to the repair shop. With this objective, the SPMS “action interface” needs to provide the following three outcomes.

1. Replace fuel nozzles
2. Inlet cleaning
3. Send to repair shop

5.2.3 Valve

The main subsystems of the valve were covered in detail in the SPMS Task 2 section (see section 4). From the cost-benefit analysis, it was identified that an SPMS application for valves would generate “Condition Indicator” and “Condition Trending.” The ”Condition Indicator” disambiguates between valve faults and clogged filter faults based on CI generation using additional sensors, whereas ”Condition Trending” identifies the deterioration trend to enable valve replacement or filter cleanup in a scheduled maintenance.

Most modern aircraft use an ACMS to acquire the data for valves. This captures BIT and initiated-BIT results.

The software mechanism used for generating all BIT reports is well-established and parameterized. The sensors installed in the valve support both the control and health management requirements. However, unlike the APU and engine, the OEM seldom provides offline tools that allow an end-user to generate new reports. This adds further challenges to the task of adding new reports.

To support valve SPMS, new reports may need to be defined, such as the valve turn-on snapshot. These reports do not have the same criticality as the control function and can be computed in a separate, low-priority partition on the controller. It should be recognized that the additional report may be difficult to generate if the valve SPMS requester is not the OEM valve.

The interfaces needed to support the above-described SPMS are summarized in table 23, and the remainder of the section describes the selection rationale and advantages and disadvantages of each of the available choices.

Table 23. Summary of interfaces for ECS SPMS

Interface	Details	Choice for the required CA-CBM module
Sensor	Real-time (streaming) of sensors listed in table 24. Begin: valve commanded to open End: valve commanded to close	Controller
Data	Batch processing (non real-time) to generate fixed number of features (see table 25)	Controller Aircraft ACMS Ground-based computer
Monitors	One monitor per valve	Controller Aircraft ACMS Ground-based computer
Health	Batch processing (non real time) to generate fault ambiguity group	Ground-based computer Central maintenance computer

Table 24. SPMS sensor interface for ECS

SPMS "Sensor Interface" for ECS
Manifold pressure and temperature **Valve open/close sensor*** **Differential/delta pressure*** **Flow measuring orifice*** **Valve position sensor (hall effect)***
Valve open/close sensor* **Valve differential pressure sensor*** **Inlet pressure sensor** **Temperature sensor** **Pressure sensor (calculate valve position)** **Micro switch to detect fully closed valve***

*Set of sensors not deployed on all valve systems

Table 25. SPMS data interface for ECS

SPMS "data interface" for ECS	Sampling precision	Reaction time	Sample rate (Hz)
Servo-valves: 10-second buffer at 20/50 Hz after an open command was issued. Feature calculated: rise time and steady state settling valve.	12-bit sampling	10 ms	50
Solenoid-controlled valves: 1 Hz data 1 second before and 5 seconds after command issued. Feature calculated: rise time and steady state settling valve.	Single precision for all calculated features	10 ms	1

5.2.3.1 Sensor Interface

There are two types of electrically driven valves used in the ECS system: the servo or rotary valves and the solenoid valves (also referred as the torque motor). One or more combinations of sensors on the valve monitor the valve position (or on/off state), valve actuation current, fluid pressure, and differential pressure and provide feedback to the controller. Often, these valves use some kind of inlet filter that keeps the particulate out. Filter clogging is a common failure mode (for valves with a filter). In addition to filter clogging, valve system failure modes include the spring failure in solenoid-type valves, actuator response faults in rotary valves, sensor failure, and insufficient hysteresis in the valve controller which results in valve chatter (eventually valve failures due to cycling). Deployed systems use BITs and initiated BITs to detect controller and valve failures. The incorporation of valve performance trends also enhances the ability to predict, isolate, and repair while maintaining the operator's schedule.

For SPMS to generate appropriate CIs for the filter, it needs to interface with a differential pressure sensor across the filter and valve systems. Additional sensors for SPMS are listed in table 24. The trending allows for the establishment of the rate of performance degradation, which can be used as a prognostics indicator.

The baseline aircraft has electrically operated valves with a digital controller. The valve controller provides an ARINC 429 connection to a central ACMS. All listed sensors in table 24 are available to this central ACMS node for download. The new sensors marked with "*" in table 24 need to be connected to the existing valve controller, which can sample and average the measurement during switch on and high-load regions during the valve operation cycle.

5.2.3.2 Data Interface

To better understand the performance degradation of the actuator/valve systems, the authors propose data collection capturing the valve dynamic response once per flight at max load. For the rotary valve, this includes high-rate samples of the valve response. For the solenoid valve, this would include a snapshot of the valve response. The data interface is shown in figure 13.

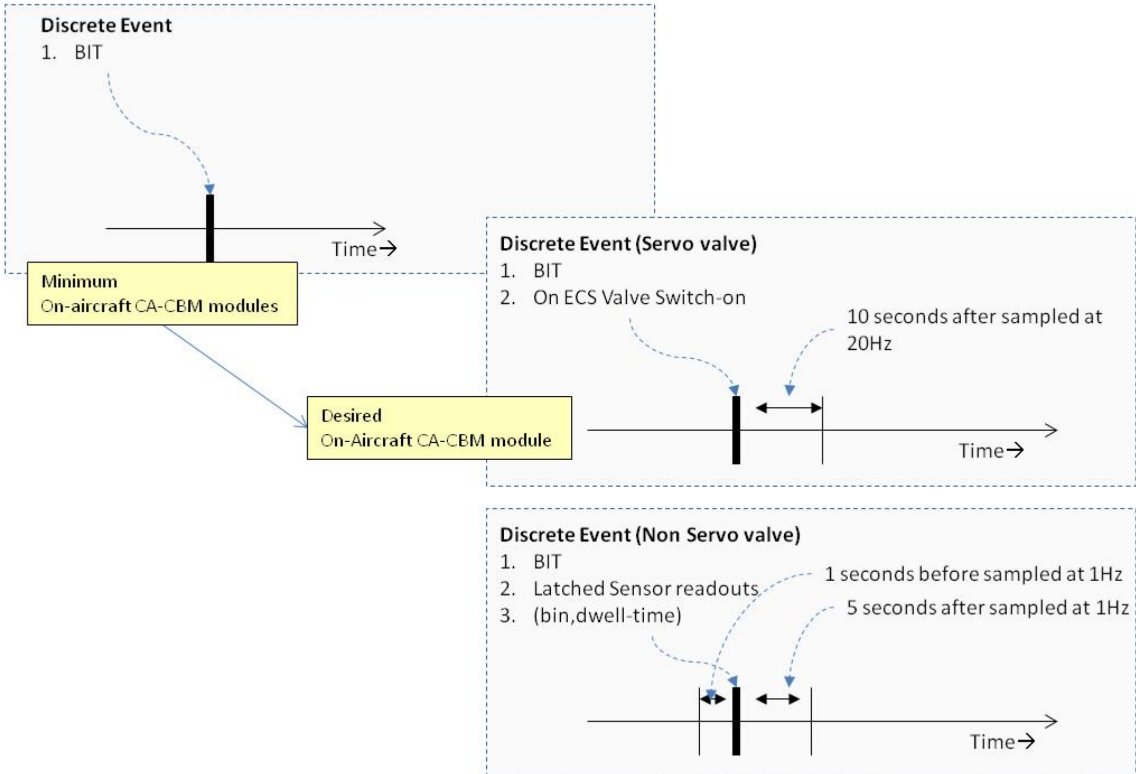

Figure 13. Data collection and evidence generation to support SPMS

Because these sensors directly interface with the valve controller, executing the intelligent data collection algorithms within the controller would provide access to timestamped sensor readings for subsequent SPMS calculations. The time delay from the sensor to the sampler should not exceed 10 milliseconds. This would allow the diagnostics algorithm to use the data from multiple sensors to be used in the same calculation frame.

Valve deterioration happens slowly. The change in the valve deterioration within an operation cycle is small. When valves degrade, they becomes sluggish, which affects how fast the valve opens when commanded to do so. This "rise time" is an important feature of this dynamic response of a valve and a good indicator of sluggishness. The second failure stiction manifests as the valve either not opening fully or not closing fully. It could get stuck at some intermediate position. Knowing this final position of the valve is a good indicator of a sticking valve. It is these two features that SPMS must capture to render its conclusion about valve health. This data interface is listed in table 25.

The extract function within SPMS provides exactly two data points for every valve transient: the rise time and the steady-state settling value. The location for executing the "measure function" that generates these two features from the raw time-series sensor data can be 1) the valve controller, 2) the ACMS, or 3) a ground-based computer. The advantages and disadvantages of each are discussed in table 26.

Table 26. Valve SPMS interface options

Where the "measure" function executes	Advantage	Disadvantage
In the valve controller	SPMS function is well-contained and could be provided as a standard upgrade process. Does not expose the intellectual property in the measure algorithms.	The measure function may need higher certification level. Controller may need additional software partitioning to support SPMS.
In the ACMS	Minimal changes to the controller. Modifications to the measure function (increasing/decreasing sampling frequency) do not alter the software criticality level.	Increases the bandwidth on existing data buses by 2 orders of magnitude.
In a ground-based computer	Propagating advances in feature extraction algorithms is easy.	Download cost may be limited by speed and connectivity.

5.2.3.3 Monitor Interface

In the context of the engine SPMS, monitors represent deviations from nominal patterns. Figure 14 shows a monitor example that captures the deviation of the typical valve opening start (blue line) and the profile as the valve degrades.

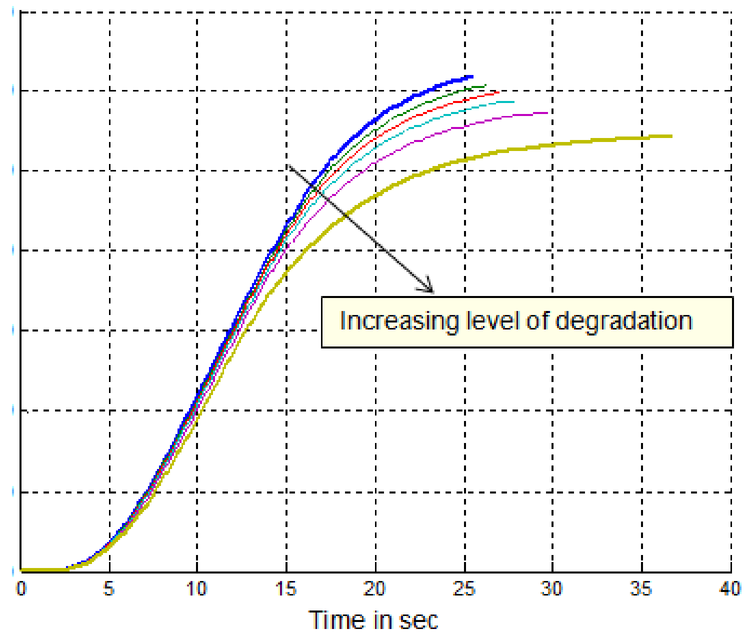

Figure 14. Example of a valve response curve (gold line shows a deteriorating valve)

To support the valve SPMS function, a distance measure is used to calculate the monitors. The monitor interface needs to support the:

- Valve rise-time distance measure.
- Valve steady state bias measure.

The numeric calculation for generating monitors happens in the "evidence generation" module within the CA-CBM framework. There are two choices regarding where this calculation can take place and the advantages and disadvantages of each are listed in table 27.

Table 27. ECS extract function options for SPMS

Where the "extract" function executes	Advantages	Disadvantages
In the ECS Controller	SPMS function is well-contained and could be provided as a standard upgrade process. Generates a new SPMS-rich report from engine.	ECS controller needs more computation power to do additional calculations. Furthermore, because the baseline for calculating deviation keeps changing, this would require periodic software modifications.
In the ACMS	Significant reduction in the size of a download report per valve transient.	Usually computation power to do additional calculation is available. However, because the baseline for calculating deviation keeps changing, this would require periodic software modifications.
In a ground-based computer	One-time remote software updates.	There are no serious disadvantages.

5.2.3.4 Health Interface

Earlier in the program, it was planned that the valve SPMS application would be focused on disambiguates between valve faults and clogged filters. If the underlying problem is related to a filter, then the health interface within SPMS needs to clearly list the valve that is faulty and the state of the filter. This determination could be made by the "diagnostic reasoner" module within the "Act" CA-CBM function.

Calculations needed for the diagnostic reasoner module can be done on a central maintenance computer or ground-based computer. Table 28 lists the advantages and disadvantages for these choices.

Table 28. ECS interpret function options for SPMS

Where the "diagnostic reasoning" module executes	Advantages	Disadvantages
On the central maintenance computer	Can use data from multiple engines to reduce false alarms. Over time, SPMS output can be used to make dispatch decisions regarding MEL.	Cost of developing onboard software and the rigor it entails.
In a ground-based computer	Deployed as a Web service and makes the output available whenever needed and to whoever needs it.	SPMS remains as a maintainer advisory system.

However, because only simulation data were used for the SPMS evaluation, the exercising of the health interface performed on this program is not extensive; this is due to the limitations of the simulator to add realistic noise and disturbance.

5.3 CERTIFICATION ISSUES

The systems approach has to be used to identify all potential SPMS certification issues. In this section, potential certification issues with target SPMS systems are discussed. A summary of the discussions from target systems in section 3 is included here for completeness. The certification challenge is best understood by synthesizing the SPMS functionality and implementation starting with the detailed analysis of required sensing, data acquisition, data communication, and computation along with system impact and any unintended consequences.

For both engines and valves, SPMS only aids the maintainer by providing enhanced diagnostics. However, for the APU, it also aids dispatch. SPMS's extension of time on wing for APUs does not have any certification issues for non-extended time on partial systems (ETOPS) operation. However, for ETOPS operation, APU time extension on wing needs to demonstrate that the probability of failed starts for APU is at or below the acceptable level. For use with ETOPS operation, the SPMS application needs to be certified to a level of APU BIT.

SPMS application for engines and APU works with the sensors installed for normal control operations. Most modern aircraft use an ACMS to acquire the data for engine, APU, and ECS valve health monitoring. There aremechanism and infrastructure to interface with the controllers to get the existing reports/BITs on the ACMS enabled aircraft. The engine and APU health monitoring captures three types of reports: snapshot reports, summary reports usually produced at the end of the flight (which collect key statistics), and exceedance reports. To support engine and APU SPMS, the reports need to be expanded. Though the necessary software and hardware infrastructure exists in most modern FADECs, the "updated functions" would require stricter software assurance rigor. Therefore, it is the authors' observation that making these changes to support SPMS does not significantly alter the software assurance rigor from what it is today. However, making these updates is not trivial. Therefore, engine and APU OEMs are only willing to make these additions if there is sufficient economic payback for them. Because new and detailed reports contain sensor data that can reveal internal engine states (making them competition sensitive), the economic payback has to overcome this barrier.

SPMS application for valves includes the addition of new sensors and generation of summary reports by the valve controllers. The authors have identified the addition of differential pressure and valve position sensors to support SPMS application. As many of these types of sensors are already certified to be used for airborne use, no sensor hardware certification issues are foreseen. The additional sensors have to be sampled by the respective valve controller. By design, the sensor data collection and summary report generation tasks are of lower priority than the controller control tasks. This minimizes the certification challenge associated with the sensor sampling and summary generation. With controllers that do not have time and space partitions, any additional code to trigger, collect, or send the data would need to be certified to the criticality level of the controller data collection functions.

From the ACMS to the ground-based computer, the SPMS uses the certified data downloading options available on the aircraft; therefore, there are no additional certification requirements/issues associated with data transmission and downloading.

The prognostics reasoner can be implemented as ground software; therefore, no certification issues are identified. In addition to the issues identified above, the use of a CA-CBM framework allows the separation of reasoner model (data) from reasoner calculations (software). This allows for the certification and maintenance of the model separately from the software. This approach saves certification costs in lower-criticality reasoning systems.

6. TASK 5: DESIGN SPMS TEST PLAN

Table 4 lists the possible impacts SPMS can make for the four use cases. SPMS is a modular system and described functionally using a series of five interfaces (described in section 5.1.2). Section 5.2 enumerated the choices for the three target subsystems. In this section, the procedure planned to evaluate some of the choices for these interfaces is described. These procedures are influenced by availability of data and, therefore, the ability to make a statistically viable comparison. The definitions of these procedures and the outcome used to make these comparisons were the primary goals of “Test Plan Activity/Task 5.”

The five SPMS interfaces are listed in table 29. An “x” in the cell indicates the test plan shall evaluate available choices and compare the outcome of specific metrics.

Table 29. SPMS target system functions and interfaces to be tested using the test plan

Test plan quantifies the impact of these choices on the accuracy and false alarm performance of SPMS		Target subsystems. An "x" indicates the coverage of the test plan on this program		
SPMS interfaces	Choices to be made	Engines	APU	Valve
Sensor interface	Which sensor to include (e.g., flow vs. temperature vs. speed)?			x
Data interface	How and when to acquire data from the sensor (e.g., sample rate, frame size, trigger conditions)?	x		x
Monitor interface	What processing needs to happen before recording the signal?	x	x	x
Health interface	What faults are distinguishable?	x	x	x
Action interface	What can be recommended to extend time-on-wing?		x	

6.1 ENGINE SPMS TEST PLAN

This study includes engines under a maintenance service plan. In this plan, the customers pay for the uptime guarantees. Unscheduled engine removals negatively impact aircraft availability and cause OIs. One such factor is related to a malfunctioning FCU.

For the engines subsystem, reliability data and dialogues with Honeywell's customer support resulted in one SPMS application, which is related to providing an actionable indicator of an impending FCU failure that could cause a potential engine no-start or in-flight engine shutdown. Because the FCU is relatively easy to swap out and does not need any special qualification tests after it is replaced, experts within Honeywell believe a prediction horizon of five cycles (five flights) is sufficient to drive SPMS value.

Beginning in 2002, Honeywell installed an elementary data collection system on 35 identical regional jets. Because no specific SPMS application was targetted, the system was set up to collect 182 parameters from various sensors installed on the four engines, APU, flight management system, navigation system, bleed system, and landing system. The list of 182 parameters was set based on physical memory locations imposed by the existing Flight Data Acquisition and Management System (FDAMS). Data collection was maximized to fill all 256 words that the ARINC 717 encoding would allow.

Actual maintenance procedures performed and parts replaced and repaired are documented and archived in relational databases maintained at three primary repair locations. Field service engineer observations are recorded as free-form text entries and archived together with repair records. Because the SPMS is targeted at the engine FCU, all of the available data cannot be used; for example, data surrounding an engine bird strike event is not within this study's defined SPMS scope.

Figure 15 lists a subset of field service observations that indicate the cause for removal of these engines in this time period. The figure only shows the subset of events in which the removal was initiated by the pilot or field maintainer after noticing something unacceptable with the FCU.

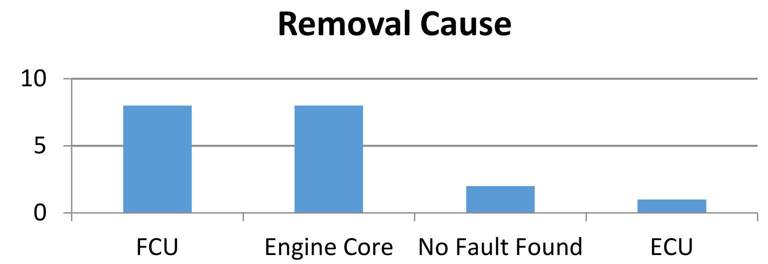

Figure 15. Number of cases for SPMS evaluation

To develop an SPMS application for the FCUs, an interface with sensors besides those installed on the engine is needed (e.g., the SPMS sensor interface includes flight phases [PHs] calculated from the flight management system). Furthermore, the rate at which the SPMS needs to analyze the sensor data is more than once per engine cycle. That is, the SPMS needs to specify a new data interface. Finally, the data need to be analyzed to generate an indicator that a maintainer can understand and, therefore, react to. That is, the engines-targeted SPMS needs to specify the monitor and health interface. A plan for evaluating these interfaces is described in section 6.1.1.

6.1.1 Sensor Interface

The baseline aircraft consists of a turbine engine controlled by a FADEC. In the baseline aircraft, the FADEC interfaces with all engine-installed sensors and communicates to a central ACMS through dedicated twisted pairs using the ARINC 429 protocol.

On a multiple-engine aircraft, these sensors are available for individual engines. Consider cases when a malfunctioning FCU may cause a slow start for the engine. A slow start may also be caused by extremely low lubrication oil temperature. If this were true (i.e., lube oil is cold soaked), then all engines on a multiple-engine aircraft would experience similar slow starts. Therefore, if SPMS for engine "A" had access to sensor data from engine "B" on the same aircraft, then it would not generate a false indicator and implicate the FCU.

The test plan for the sensor interface involves evaluating whether sensors from engine A can help the SPMS for engine B (see table 30).

Table 30. Sensor interface for engine-FCU targeted SPMS

SPMS "sensor" interface choices	Test plan	Outcome
EGT sensor Engine shaft speed (the power generator core) Engine WF FADEC fuel command signal Inlet temperature sensor Inlet pressure sensor IGV position sensor	Pick and choose sensors from other engines on the same aircraft	Two "no fault found" cases. The outcome counts the number of these cases in which SPMS successfully declares "FCU is not a problem."

6.1.2 Data Interface

The test plan for the data interface (see table 31) is focused on evaluating the sampling choice (i.e., the rate at which the sensors are sampled and the collected data is retained and the impact of this choice on the detection accuracy of FCU-related problems). Specifically, the focus is on the following regimes: engine startup and takeoff.

Table 31. Data interface options to be tested for the engine FCU-targeted SPMS application

SPMS "data" interface choices	Test plan	Outcome
Sampling frequency of 1, 2, and 4 Hz 60-second data collection during engine start 30-second data collection during takeoff (engine max power)	Study the sensitivity of sampling at rate and the regime then data is collected	There are eight cases of FCU removals. How does the detection metric (# of successful detections divided by eight) change with these choices?

6.1.3 Monitor Interface

The focus of the SPMS application is FCUs. The approach as described in this report is to analyze relatively high-frequency, time-series data collected during predefined phases of engine operations. The pattern of given time-series data is compared against a baseline pattern. The monitor measures the dissimilarity between the baseline pattern and given pattern.

To clarify, let *X(t)*denote the timeseries data collected from one of the sensors listed in the previous section. Let $\overline{X}(t)$ denote the baseline pattern for the same sensor. The monitor measures the dissimilarity between these two patterns. That is:

$$\text{Engine Monitor } \propto (X(t) - \overline{X}(t)) \tag{10}$$

SPMS monitors for the engine are:

The starter monitor is the deviation of the engine speed v/s time data before the engine reaches its auto combustion point:

$$m_{SM}(n) \tag{11}$$

The ignitor monitor is the deviation of the engine speed v/s time data, when fuel is introduced into the combustion chamber for the first time:

$$m_{IG}(n) \tag{12}$$

The fuel controller monitor is the deviation of the engine speed v/s time data, when the fuel controller is actively trying to attain its set point and meet the fuel demands:

$$m_{FCU}(n) \tag{13}$$

The idle speed monitor is the deviation of the engine speed after it reaches its idling speed:

$$m_{core}(n) \tag{14}$$

Here, (n) denotes that the monitors are calculated at the end of the n^{th} engine cycle.

The choice here depends on how the deviation is calculated. There are two options for calculating these: 1) the simple Euclidian distance, in which two signals are compared sample by sample, and 2) the Mahalanobis distance, in which observations are made of the distance of more than one signal by transforming them to a new signal.

The choice can impact the numeric value of the monitor and, therefore, placement of the decision boundary determining whether the SPMS declares the data indicates an impending FCU fault. The test plan for this interface is summarized in table 32.

Table 32. Monitor interface options to be tested for the engine FCU-targeted SPMS application

SPMS monitor interfaces choices	Test plan	Outcome
Profile deviation monitors : 1) speed v/s time profile before the engines reaches its fuel is introduced, 2) engine speed v/s time profile when fuel is introduced into the combustion chamber point, 3) speed v/s time data when fuel controller is governing, 4) engine speed deviation after it reaches its idling point Methods for calculating deviation: 1) Euclidian, 2) Mahalanobis	Study the sensitivity and contribution of various metrics detecting FCU problems. Robustness of using two methods for calculating profile deviation.	Coverage of FCU-related engine removals.

6.1.4 Health Interface

Given deviation from baseline pattern monitors, the health interface provides a ranked list of probable root causes that represent the internal health state of the engines. The intent is to help the ground maintainer and engine OEM to get ready and avoid any operational disruption. The outputs provided by the SPMS monitor interface provide evidence for several engine components.

For example, a deviation in the engine speed v/s time curve before auto-combustion can be due to a starter problem, low bleed from the APU, and/or increased resistance from the engine turbine bearings. The last problem could in turn be caused by low lube temperature and increased oil viscosity. Deviations during the light off phase could be due to igniters or incorrect fuel atomization. Deviations at engine idle could be caused by incorrect fuel metering or loss of engine performance due to the erosion of rotating components, such as turbine blades.

The health interface provides a probability (a number between 0 and 1) associated with the following four causes:

$$\{FM\} = \begin{Bmatrix} \text{FCU} \\ \text{Starter Motor (SM)} \\ \text{Igniter} \\ \text{Engine Core} \end{Bmatrix} \qquad (15)$$

The interface provides:

$$P(\{FM_j\}) = \begin{Bmatrix} > 0 \text{ implies failure mode is present} \\ = 0 \text{ otherwise} \end{Bmatrix} \qquad (16)$$

Here the notation $P(.)$ indicates the probability and $\{FM\}$ denotes the set consisting of the four failure modes listed. Because these conditions can exist independently, there is no requirement that the probabilities add to 1.

Because the SPMS is targeted at the engine FCU, the interest is in $P(FCU)$. Therefore, the SPMS health interface should provide an actionable indicator of an impending FCU failure that could cause a potential engine no-start or inflight engine shutdown. The rest are important because $P(FM \neq FCU) > 0$ indicates an ambiguity group and does not provide a clear direction to the maintainer. For example, assume that SPMS provided $P(FCU) = 0.85$ and $P(SM) = 0.72$.. Though mathematically the FCU has a higher probability, the maintainer is confused and may replace the starter motor instead of the FCU, and the problem may persist.

The primary metric that measures the effectiveness of the health interface and, to a large extent, engine-FCU targeted SPMS, is called fault isolation. There are several formulas for fault isolation. The following was used for this study:

$$I_F = \frac{P(FCU)}{\Sigma_j P(FM_j)} \tag{17}$$

If $P(FCU) = 1$ and $P(FM \neq FCU) = 0$, then $I_F = 1$, indicating perfect fault isolation. Because there are $N = 19$ cases, the average over all N cases is taken.

The metric fails if all $P(FM) = 0$. That is, none of the failure modes is detected. Therefore, the secondary metric is defined as a detection.

Rationale: Over a 3-month interval, assume there were R engine removal incidents. If the health interface of engine-targeted SPMS generated an output such that $P(FCU) \geq 0$ at least five cycles before the removal date, this study claims that the SPMS indicated the FCU removal. The number "5 cycles" was set after discussions with Honeywell's field service engineers. This organization within Honeywell interfaces with airline operators and, based on the feedback they receive, they indicate a 5-cycle early warning is needed to respond.

If SPMS provides $P(FCU) \geq 0$ at least five cycles before the removal date, this study determines this a successful detection. Let M be the count of such successful detections.

Detection accuracy is defined to evaluate the SPMS health interface. The accuracy metric α_{ENG} is defined as:

$$\alpha_{ENG} = \frac{M}{R} \times 100 \tag{18}$$

Checking the four monitors (m_{SM}, m_{IG}, m_{FCU}, m_{core}) at the end of the n^{th} usage cycle against four individual thresholds may not provide the 5-cycle prediction horizon. Therefore, the authors believe that changes and improvements need to be made to get better results. The technology the authors want to develop is multivariate feature extraction. That is, looking at all four monitors together. This combination may be a linear weighting or Boolean combination formula. Algorithmically, the plan involves testing the following two options for calculating the combination function:

1. Linear weighting: The following provides an example linear weighting function that the authors plan to test in this program:

$$P(FCU) = a_1 m_{SM}(n) + a_2 m_{FCU}(n) + a_3 m_{IG}(n) + a_4 m_{core}(n) \tag{19}$$

2. Boolean rule: The following provides an example combination function the authors plan to test in this program:

$$P(FCU) = AND(m_{FCU}(n) > \theta_{FCU}, m_{SM}(n) \le \theta_{SM}) \tag{20}$$

In either case, the linear weighting coefficients $a_1, \ldots, a_4$ and the thresholds $\theta_{SM}, \theta_{FCU}$ need to be calculated That is, the authors will develop the health interface for the engine-targeted SPMS application and test it.

The test plan for this interface together with the outcome is summarized in table 33.

Table 33. Interface test plan

SPMS health interfaces choices	Test plan	Outcome
Coverage elements: 1) FCU, 2) starter, 3) ignitor, 4) engine core Combination function for monitors: 1) weighted average with a threshold, 2) and/or logic	Study the ability to discriminate between various components (overlap in the decision space) based on the combination function.	Detection accuracy fault isolation: In how many cases can the FCU be pinpointed among other components?

6.2 APU SPMS TEST PLAN

In the past 10–13 years, Honeywell's PTMD application has been monitoring over 1900 APUs. This PTMD program has archived ~11 years of APU operational data. Actual maintenance procedures performed and parts replaced and repaired are documented and archived in relational databases maintained at three primary repair locations. Field service engineer observations are recorded as freeform text entries and archived together with repair records. By cross-correlating operational data with the APU maintenance logs, the APU historic data can be segmented as a sequence of on-aircraft and repair-shop lifecycle events. The authors' intent is to use this lifecycle data spanning over 1800 APUs to exercise the SPMS interfaces and evaluate their effectiveness.

The primary function of an APU is to provide bleed air for starting the main propulsion engines. It also drives the generator and cabin air when the aircraft is on the ground. Therefore, an APU is removed from service by an airline operator when it fails to provide these basic functions. Internally, the compressor that provides the bleed air is driven by a power plant (gas turbine engine). Being a small gas-turbine engine, an APU is subject to the same deterioration phenomena as a larger propulsion engine; namely, the turbine blades erode over time and reduce the cycle efficiency. This lowers the energy available to drive the bleed air compressor. The second leading cause is the APU's inability to start because of oil/fuel or other non-rotating component failures.

Table 34 illustrates the data provided by the PTMD program cross-linked with the maintenance records that provides the "shop finding." The table shows two cases (APU serial numbers 102 and 104), both of which were removed because of a performance problem; the repair shop confirmed the root cause as the first-stage turbine stators. The table also shows serial number 106, which was removed because of an auto shutdown issue; the repair shop confirmed an oil cooler seal problem as the cause.

Table 34. Historical data available for APU-targeted SPMS interface testing

APU serial number	Usage hours	Date	EGT margin	Bleed pressure margin
102	8,253	12/2/2008	102	6.71
	8,257	12/3/2008	103	6.63
	82,59	12/4/2008	95	6.87
	…	…	…	…
	APU removed because of low performance; first stage turbine damage confirmed. Date: 1/4/2010.			
104	4,567	8/12/2007	121	7.23
	4,568	8/12/2007	113	6.71
	4,571	8/13/2007	118	6.86
	…	…	…	…
	APU removed because of low performance; first stage turbine damage confirmed. Date: 6/13/2008.			
106	3,331	8/12/2007	56	2.89
	3,334	8/12/2007	42	3.13
	33,336	8/13/2007	48	3.16
	…	…	…	…
	APU removed because of shut down; oil cooler seal leak confirmed. Date: 9/20/2007.			
Total such cases available N = 1200				

The PTMD is an ongoing program and, therefore, the number of cases keeps growing. At the time of this writing, 1200 such cases were accumulated; all of these cases are available for SPMS testing. The notation N was used to denote the count of these cases. In this case, N = 1200.

6.2.1 Sensor and Data Interfaces

Though these historical data provide a rich source of successive lifecycle events, they also impose a few constraints on the interfaces that can be tested. The sensor and data interfaces are particularly affected. For example, the sensor interface that describes the sensors from which SPMS needs to collect data is affected. Because this is an existing aircraft-installed APU, the authors' test plan cannot modify it; that is, neither the sensor locations nor the underlying modality they are sensing can be changed.

The same discussion holds for the data interface that defines the frequency of sampling and how often the data are collected from the sensor. The historic data that is planned to be used on this program use standard reports from the APU. These reports have predefined trigger conditions, the sampling rate, and the digitization rules (number of bits to be used to save the data). Therefore, the test plan cannot modify these.

The test plans for these two interfaces are trivial and summarized in table 35.

Table 35. Sensor and interface test plan

SPMS interfaces	Test plan	Outcome
Sensor interface	This interface cannot be exercised on historical data.	N/A
Data interface	This interface cannot be exercised on historical data.	N/A

6.2.2 Monitor Interface

The monitor interface consists of two performance CIs called EGT margin and bleed-pressure margin. EGT margin and bleed-pressure margin measure the ability of the APU to deliver bleed air without exceeding internal temperature and surge limits. When the APU is new, the margins are very high, indicating that the APU is operating far away from its internal limits. As the APU ages, margins decrease indicating that the APU is operating closer to its internal limits. Therefore, a threshold θ_{EGT} and θ_{PB} can be assigned, such that when the current EGT margin and bleed pressure margin are less than these thresholds, the SPMS will recommend a removal.

Note that, as shown in table 34, not all APUs are removed because of performance issues. That is, the monitor interface that provides only the EGT and bleed pressure margins would not be able to detect these removals. This is called the "coverage metrics;" of the 1200 cases available to the authors, "coverage metrics" is the number of cases that cannot be detected by the SPMS monitor interface. Let this number be *R*. Then, the coverage metric is defined as:

$$\gamma_{PTM} = \frac{R}{N} \times 100 \tag{21}$$

The desired result is $R \rightarrow N$ such that the the SPMS monitor interface is sufficient to cover all of the APU removed in the past 11 years. This is clearly not the case. The test plan is to establish this number.

The PTMD application calculates these margin values after the APU completes one full start-stop cycle. These calculations are performed using a software module executing on a ground-based computer. Because the PTMD is a production/live-customer application, the software that calculates the EGT and bleed pressure margins cannot be altered. However, the following can be done: 1) add additional sensors as monitors, 2) assume that some input values are unavailable, or 3) introduce random noise in the inputs. That is, this sensitivity analysis can be performed and the impact of these error sources on the coverage metrics can be documented.

The test plan for this interface, together with the outcome, is summarized in table 36.

Table 36. SPMS monitor interface test plan

SPMS monitor interfaces	Test plan	Outcome
Monitor choices: 1) EGT_{margin}, 2) Bleed pressure margin, 3) Start time, 4) IGV position	Add random noise and study the impact on detection metrics. Sensitivity analysis of adding new monitors.	Detection accuracy: How many APUs removed because of performance issues can be detected by adding these monitors?

6.2.3 Health Interface

Given EGT and bleed-pressure margins, the health interface provides a ranked list of probable root causes that represent the internal health state of the APU. The intent is to help the ground maintainer and APU OEM be prepared to avoid any operational disruption. This interface provides a probability (a number between 0 and 1) associated with the following four causes:

$$\{FM\} = \left\{ \begin{array}{c} \text{Fuel Nozzles} \\ \text{Inlet Fouling} \\ \text{1st Stage Turbine Stators} \\ \text{Compressor Erosion} \end{array} \right\} \tag{22}$$

Mathematically, this is:

$$P(\{FM\}) = \left\{ \begin{array}{c} > 0 \text{ if } EGT \text{ margin} < \theta_{EGT} \text{ or } PB \text{ margin} < \theta_{PB} \\ = 0 \text{ otherwise} \end{array} \right\} \tag{23}$$

Here, the notation $P(.)$ indicates the probability and $\{FM\}$ denotes the set consisting of the four failure modes listed in equation 22.

Because these conditions can exist independently, there is no requirement that the probabilities add up to 1.

Rationale: Over a 3-month interval, assume there were *R* unscheduled APU removals because of performance issues. If the health interface of APU-targeted SPMS generated an output such that one of the health states attained a probability of ≥80% at least 30 cycles before the removal date, this study claims that the SPMS indicated the APU removal. The number "30 cycles" was set after discussions with Honeywell's field service engineers. This organization within Honeywell interfaced with airline operators and, based on the feedback they received, indicated that a 30-cycle early warning is needed to respond. That is, in the 30-cycle prediction horizon provided by 30 cycles (this typically corresponds to 10–12 calendar days), the airlines can react appropriately and ensure that a spare APU is available. With fewer than 30 cycles, the response may be delayed because of supply chain and logistics delays, making SPMS less effective.

If SPMS provides a probability of ≥80% at least 30 cycles before the removal date, this is called a successful detection. Let *M* be the count of such successful detections.

Two metrics are defined to evaluate the SPMS health interface: accuracy (true positives) and false positives. The accuracy metric α_{PTM} is defined as:

$$\alpha_{PTM} = \frac{M}{R} \times 100 \tag{24}$$

Suppose there were *P* APUs that were falsely removed by the SPMS trending technology. Then, the false alarm metric is defined as:

$$\beta_{PTM} = \frac{P}{R} \times 100 \tag{25}$$

Next, the test for calculating the detection accuracy metrics for the APU-targeted SPMS health interface is described; namely, the values for α_{PTM} and β_{PTM} using historical data.

Clearly checking the EGT and bleed pressure margins at the end of the n^{th} usage cycle against the threshold θ_{EGT} and θ_{PB} would not provide the 30-cycle prediction horizon. Therefore, the authors believe that changes and improvements need to be made to get better results. The technology the authors want to develop should be based on the rate of change (slope) of the margins rather than current value. Furthermore, θ_{EGT} and θ_{PB} may not be a single number and one may have to use a fuzzy threshold. That is, trending and assignment algorithms (mathematically, they are called decision boundaries) need to be developed to test the health interface.

Algorithmically, this plan involves testing the following two options for aggregating the monitors:

1. Univariate (EGT_{margin} only) trending and looking for a threshold crossing.
2. Multivariate Bayesian inferencing: The decision to remove an APU because of performance is a random variable that depends on several inputs, such as 1) EGT_{margin}, 2) bleed-pressure margin, 3) start time, and 4) IGV position.

For the decision boundaries, the plan involves testing the following two options:

1. Constant threshold: Discussions with Honeywell subject matter experts indicate:

$$\theta_{EGT} = 15C \text{ and } \theta_{PB} = 0.5 \text{ psia} \tag{26}$$

2. Threshold that is derived using a training and testing data set picked randomly from the APU population.

The test plan for this interface, together with the outcome, is summarized in table 37.

Table 37. Test plan for SPMS health interfaces

SPMS health interfaces choices	Test plan	Outcome
Aggregation function: 1) univariate trending, 2) Bayesian multivariate inference. Decision threshold: 1) fixed by subject matter experts, 2) empirically derived from the APU population.	Evaluate options for single variable trending. Evaluate options for decision boundaries.	Detection accuracy: How many APUs removed because of performance issues can be detected? False positives: How many times SPMS determines that the threshold is crossed and the APU needs to be removed, but the repair shop shows no performance issues.

6.2.4 Action Interface

After the SPMS has isolated the underlying factor that best explains the observed monitors or symptoms, the action interface provides recommendations. The cost-benefit analysis indicates the APU-targeted SPMS action interface needs to recommend a fuel nozzle replacement whenever applicable and, therefore, extend the APU time-on-wing before sending it to the repair shop.

Rationale: The APU EGT margin trend line is a measure of APU aging and, therefore, it is a characteristic slope as shown in figure 16. When the APU is first installed on the aircraft, the margin remains more or less constant. This is called the “break-in” phase. As the APU ages, the trend line starts its downward slope. The slope increases over time, which is called the deterioration phase. At some point, the trend line crosses the zero line and the APU is removed and sent to the repair shop for maintenance. The cycle repeats after the APU is repaired and reinstalled.

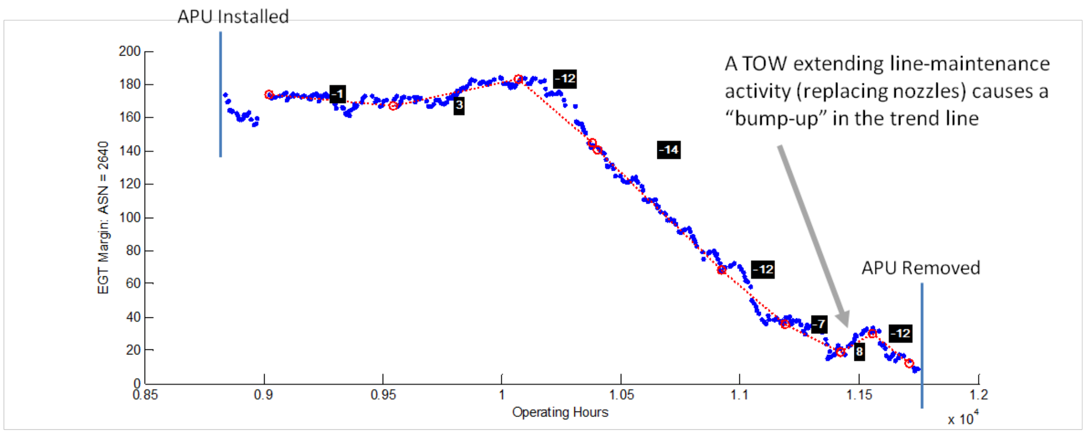

Figure 16. SPMS-generated EGT_{margin} trend

Beginning at approximately 10,200 hours until 11,400 hours, the APU is in the deterioration phase. At approximately 11,400 hours, the slope of the trend line increases temporarily because of a fuel nozzle change maintenance action. Clearly, this temporary maintenance action effectively slowed the deterioration rate. However, the “bump up” was temporary and lasted for approximately 100 hours, after which the APU deterioration continued and it was eventually removed and sent to the Honeywell repair shop. Interviews with Honeywell field service engineers indicate that the “bump-up” was due to a fuel nozzle replacement. That is, for this APU, the field service engineer combined their knowledge of the APU deterioration with the SPMS output (from the health interface) to infer that a fuel nozzle change under these conditions should “buy some time” for the operator.

The primary objective of the APU-targeted SPMS action is to "algorithmically capture this knowledge," and the test plan evaluates how well it was captured. This was not a simple task. For example, other sensor data may need to be taken into account, such as APU start times, inlet temperature, and bleed valve positions, in addition to the EGT margin. Developing this multivariate pattern recognition algorithm (which the authors call "maintenance reasoned") is a significant part of testing the SPMS action interface. This involves 1) characterizing the pattern using historical data and 2) defining a similarity measure that recommends the fuel nozzle maintenance action based on closeness of the prevailing APU health as provided by the SPMS health interface. Two candidate options the authors plan to evaluate are:

1. Qualitative trend analysis: In this case, the down-up-down template pattern is approximated as a series of qualitative shapes. Similarity measures the recurrence of this shape sequence.
2. Match filter: In this case, the down-up-down pattern is treated as a time-series vector. A given time-series vector, *x(t)*, matches a template pattern, *p(t)*, when the dot product of the two vectors is 1. Similarity measures the cosine of the angle between *x(t)* and *p(t).*

Next, a metric is defined to evaluate this SPMS interface based on discussions with Honeywell's field service and business leaders who manage the APU maintenance service plans.

Metric description: Over a 3-month interval, assume there were *F* APU removals, such that within 50 cycles prior to the removal some or all fuel nozzles were replaced by the line maintainer. The metric the authors want to track is the minimum deferral window, τ_{MR}, from applying the SPMS technology across the entire fleet. Clearly, the authors want to establish this bound with a >90% confidence level, such that the SPMS can increase its role in maintenance-related decision making. The authors clearly want $\tau_{MR} \gg 0$ so that there is enough financial incentive to make this nozzle change.

The current (baseline) metric value is $\tau_{MR} = 0$. The end-state for the metric based on discussions with internal Honeywell business leaders is $\tau_{MR} \geq 15$ cycles.

This study's plan is to use historic data to develop the necessary rules that mathematically characterize the "bump" or down-up-down patters in the EGT margin trend line. Because the fuel nozzle replacement action is a line maintenance action and not recorded as repair logs, the authors have to rely on field service engineer interviews to annotate these events. This implies that only a subject of N = 2000 cases is available for testing this SPMS interface. The procedure for identifying this subset, $F \ll N$, is described below and shown in figure 17:

- For each of the available N = 2000 cases, the authors identified the "down-up-down" trend line that signifies a possible fuel nozzle replacement action. This yielded 44 cases.
- If this pattern occurs "close" to the APU removal guidelines Honeywell provides to the airline operators, this increases the likelihood that a fuel nozzle replacement was completed.
- The APU serial number and approximate date-time are cross-checked with "squawks" recorded by the field service engineer. The net result is eight cases that are available for testing this action interface of the APU-targeted SPMS application.

These steps are depicted in figure 17. Because the PTMD is an ongoing program, the authors expect these test cases to increase over time. To date, eight cases have been identified, but this may grow during the course of this program.

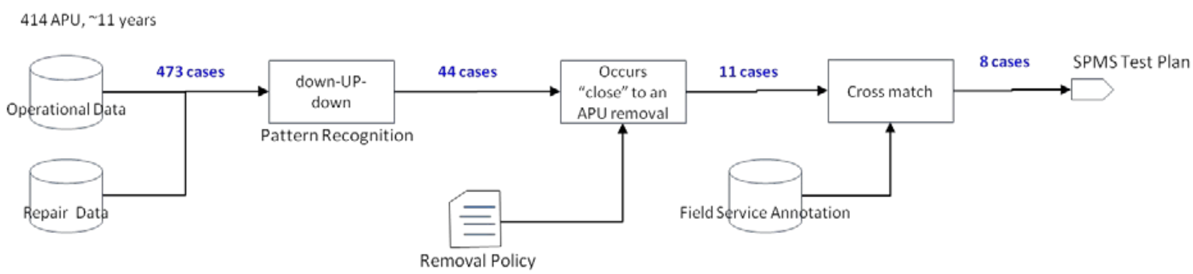

Figure 17. Steps for identifying test cases for the APU-targeted SPMS application

The test plan for this interface, together with the outcome, is summarized in table 38.

Table 38. Test plan for Action interface

SPMS interfaces	Test plan	Outcome
Action interface	Evaluate options for multivariate pattern recognition	On-wing extension time (in APU cycles)

6.3 VALVE SPMS TEST PLAN

The mechanical perimeter includes multiple valves that are tailored to meet the cabin and avionics heating and cooling loads, deicing, and cabin pressurization needs. In general, the valves are becoming more electrically driven due to an emphasis on fuel economy and more electric aircraft (MEA) architectures. Talks with Honeywell experts and data from field reliability indicate that the valve and filter issues are key maintenance drivers. However, the valve seats wear faster than the electrical components of the valve. Therefore, the SPMS valve efforts were refocused to trending the valve seat wear. The associated target is shown in table 39.

Table 39. SPMS value expected results

SPMS application	Problem/target system	Potential impact of SPMS	Expected results
Onboard CI generation for valve	Probability of predicting an impending valve failure	Identify the rate of degradation and predict the removal due to valve seat issues	Identify the valve seat failure 10 flights in advance

However, Honeywell's engineering teams develop very detailed physics-based models to simulate the valve wear for the poppet valves. These models are calibrated using well-designed test-bench experiments. One such model is the ECS system for the Airbus A350 (see figure 18), which is used to evaluate the ECS-targeted SPMS design.

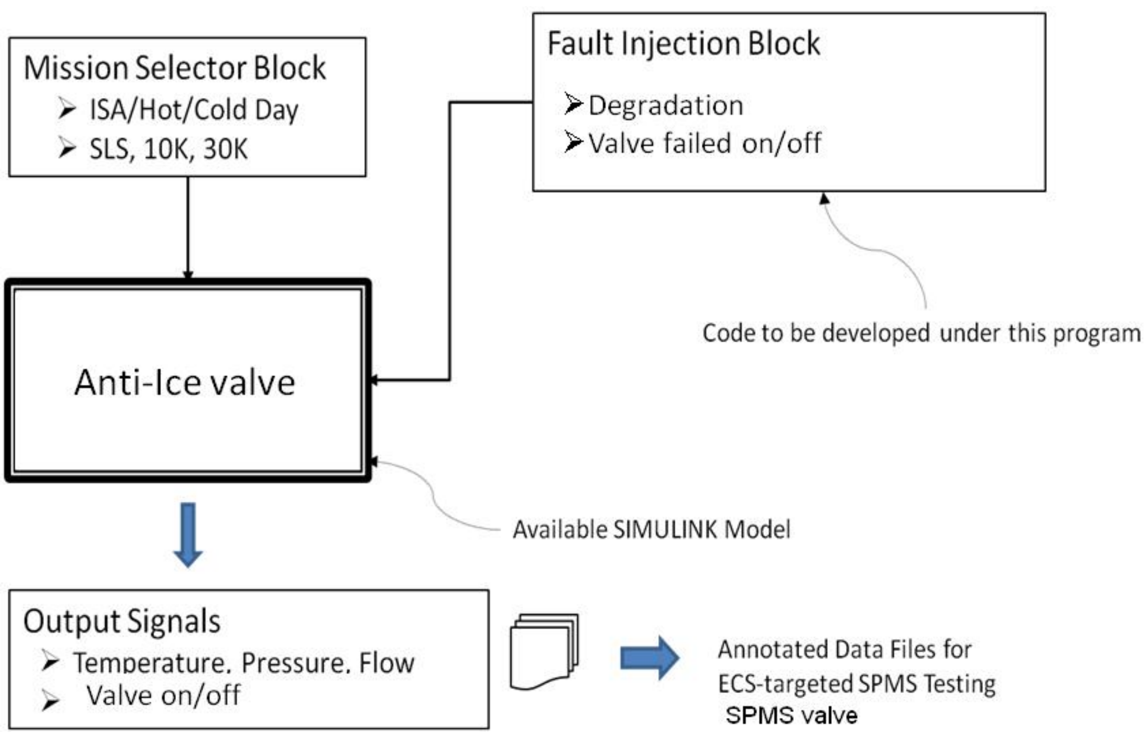

Figure 18. A "fault injection block" added to the existing Simulink model of poppet valve degradation to generate the necessary data for testing SPMS interfaces

Because the objective is to estimate the effectiveness of SPMS in detecting different levels of degradation that can be used to provide advance warning, an unacceptable level for the valve response needs to be defined. Valve responses slower than this level are tagged as unaccepted and, therefore, the valve needs to be removed or replaced or the filter needs to be cleaned. Let

δ_V and δ_F represent different levels of valve and filter degradation, respectively. The valve charactistics will be developed by using the data sets in a design of experments (see table 40).

Table 40. Design of experiments to generate test cases from the A350 ECS Simulink model

Fault injection	Severity	Data sets
Valve degradation based on duty cycles	Normal, three degradation profiles for δ_V	4
Stuck open	On/Off	2
Stuck closed	On/Off	2
Total: 17 cases		

Table 41 summarizes the ECS-targeted SPMS interfaces that the authors plan to test and the outcomes from these tests.

Table 41. ECS-trim valve-targeted SPMS test plan

SPMS interfaces	Test plan	Outcome†	Additional details
Sensor interface	Pick and choose signals available from the Simulink model	FC, DA, FI	See section 6.3.1
Data interface	Choose the rate at which data are recorded from the Simulink model	FC, DA, FI	See section 6.3.2
Monitor interface	Evaluate options for time-series analysis	FC, DA, FI	See section 6.3.3
Health interface	Develop and test options	FC, DA, FI	See section 6.3.4
Action interface	N/A	N/A	N/A

FC = fault coverage, DA = detection accuracy, FI = fault isolation

6.3.1 Sensor Interface

The baseline aircraft has ECS electrically operated valves with a digital controller. Unfortunately, the model available for this study can only simulate pressure, flow, and thermal aspects of the valve. Therefore, the fault injection is restricted to mechanical failures and their impact on valves.

Several "signals" are available from the model. This, in turn, defines the choices that can be exercised for the ECS-targeted SPMS sensor interface. The choices available for SPMS testing are:

- Manifold air pressure, temperature, and flow
- Valve positions
- Differential/delta pressure across all valves within the ECS

The test plan (see table 42) for the sensor interface involves evaluating the minimum number of sensors the ECS-targeted SPMS needs to interface and its impact on overall accuracy of the CIs.

Table 42. SPMS sensor interface test plan

SPMS interfaces	Test plan	Outcome
Sensor interface	Pick and choose signals from the Simulink model to measure the value of adding a sensor at that location in an actual aircraft.	Accuracy of the HI for isolating malfunctioning valve from the clogged filter.

6.3.2 Data Interface

To calculate the necessary CIs, the data interface needs to provide time-series data for all the sensors. The rationale for collecting this data when the valve is operational was also described. The test plan for this data interface is focused on evaluating this choice. That is, 1) the time duration during which the authors need to collect the signals, and 2) the number of samples needed (sampling frequency) that are best for the valve-targeted SPMS application (see table 43).

Table 43. Data interface options to be tested for the ECS valve targeted SPMS application

Trigger event buffer	Sampling frequency
$\pm 10, \pm 20, \pm 30$ seconds before and after the commanded change from the controller.	Study the sensitivity of sampling at 5,10, 20, and 50 Hz

6.3.3 Monitor Interface

Valves are the focus of the SPMS application for ECS. The approach is to analyze relatively high-frequency time-series data collected before and after the valves receive command signals from the controller; that is, the transient response of the system. The pattern of given time-series data is compared against a baseline pattern. The monitor measures the dissimilarity between the baseline pattern and the given pattern.

To make things clear, let $X(t)$ denote the timeseries data collected from one of the sensors listed in the previous section. Let $\overline{X}(t)$ denote the baseline pattern for the same sensor. The monitor measures the dissimilarity between these two patterns. In other words:

$$\text{Valve Monitor } \propto (X(t) - \overline{X}(t)) \tag{27}$$

There are several choices for this.

1. Root mean square (RMS) deviation: The difference between the baseline response and the current response is compared at each sampling interval and added to generate a mean difference between the two signals. This becomes the valve monitor:

$$m_{RMS} \leftrightarrow \frac{1}{T}\sum_{t=1}^{T}(X(t)-\overline{X}(t))^2 \qquad (28)$$

2. The response can be approximated using a first order response to the input command signal, *u*(*t*). This model has two parameters: gain, *K*, and time constant, τ. In this case, $\overline{X}(t)$ is given as:

$$\overline{X}(t) = u(t) * K * (1-e^{-t/\tau}) \qquad (29)$$

The monitor is defined as the model-based residual:

$$m_{MB} \leftrightarrow \frac{1}{T}\sum_{t=1}^{T}(X(t)-u(t)*K*(1-e^{-t/\tau}))^2 \qquad (30)$$

3. The response of the valve, $X(t)$, to the given input command, $u(t)$, can be modeled as a first order response. In this case, the observed data are used to estimate the first order gain, $\widehat{K}$, and the time constant, $\hat{\tau}$. These estimated model parameters become the monitors:

$$m_{par} \leftrightarrow \min_{K,\tau}\sum_{t=1}^{T}(X(t)-u(t)*K*(1-e^{-t/\tau}))^2 \qquad (31)$$

Let m_{RMS}, m_{MB}, and m_{par} denote the three available choices. These notations are used in the next section.

Furthermore, if during a flight the valve experiences *C* command signals, then several choices are available for the monitor interface:

1. Min/Max/Average: Calculate the monitors from all *C* trigger conditions and take a minimum/maximum/average of them.
2. Max loading: Only calculate monitors for the trigger that experiences the maximum command signal.
3. Averaging time series and calculating the monitors based on the average response time-series profile.

The test plan for the montor interface and expected outcome is provided in table 44.

Table 44. Monitor interface options to be tested for the ECS targeted SPMS application

SPMS interfaces	Test plan	Outcome
Monitor interface	Options for summarizing the time-series response to 1–2 numbers	Coverage of valve faults from all the test cases

6.3.4 Health Interface

Given "deviation" from baseline pattern monitors, the health interface provides a ranked list of probable root causes that represent the internal health state of the valve. The intent is to help the ground maintainer be prepared and avoid any operational disruption.

The health interface provides a probability (a number between 0 and 1). The interface provides:

$$P(\{FM_j\}) = \left\{ \begin{matrix} > 0 \text{ implies specific failure mode is present} \\ = 0 \text{ otherwise} \end{matrix} \right\} \tag{32}$$

Here, the notation $P(.)$ indicates the probability and $\{FM\}$ denotes the set consisting of the four failure modes listed. Because these conditions can exist independently, there is no requirement that the probabilities add up to 1.

Rationale: Of the N cases listed in table 40, two scenarios correspond to a no-fault case. If the health interface of valve-targeted SPMS generated an output such that $P(V) \geq 0$ for M cases, the authors claim that the SPMS indicated a valve removal.

Define detection accuracy to evaluate the SPMS health interface. The accuracy metric α_{ECS} is defined as:

$$\alpha_{ECS} = \frac{M}{N-2} \times 100 \tag{33}$$

The authors plan to keep the choices for the health interface simple. Let m_{RMS}, m_{MB}, and m_{par} be the three monitors available from the monitor interface. The technology the authors want to test is multivariate feature extraction; that is, looking at all three monitors together. This combination may be a linear weighting or Boolean combination formula. Algorithmically, the plan involves testing the following two options for calculating the combination function:

1. Linear weighting: The following provides an example linear weighting function that the authors plan to test on this program:

$$P(V) = \sum_i a_i \times m_i \tag{34}$$

2. Boolean rule: The following provides an example and combination function that the authors plan to test on this program:

$$P(V) = AND(m_i(n) > \theta_i), i = 1, 2, \ldots \quad (35)$$

In either case, the linear weighting coefficients, a_i, and the thresholds, θ_i, need to be calculated. That is, the authors will develop the health interface for the valve-targeted SPMS application and test it. The test plan for this interface, together with the outcome, is summarized in table 45.

Table 45. Health interface test plan

SPMS interfaces	Test plan	Outcome
Health interface	Evaluate options for monitor combination rules (e.g., linear weighting, Boolean rules)	Accuracy of the HI detecting an impending malfunctioning of the valve because of seat leakage

7. TASK 4: REASONER REQUIREMENTS

7.1 INTRODUCTION

The objective of this task is to design and develop the SPMS reasoners for the SPMS applications the authors identified for the three target subsystems: APU, engines, and ECS. The designed reasoner is used to evaluate the interface selections listed in the Task 3 report and quantify the performance metrics specified in the Task 5 report.

7.2 ENGINE SPMS

Unscheduled engine removals negatively impact aircraft availability and cause OIs. The engine SPMS targets OIs because of a malfunctioning FCU.

Beginning in 2002, Honeywell installed an elementary data collection system on 35 identical regional jets. Because no specific SPMS application was selected, the system was set up to collect 182 parameters from various sensors installed on the four engines, APU, flight management system, navigation system, bleed system, and landing system. The list of 182 parameters was set based on physical memory locations imposed by the existing FDAMS. Data collection was maximized to fill all 256 words that the ARINC 717 encoding would allow.

Actual associated maintenance procedures performed and parts replaced and repaired are documented and archived in relational databases at repair locations. Field service engineer observations are recorded as freeform text entries and archived together with the repair records. Note that, based on the authors' experience, the field observation date (which is in the form of a text note) was not accurately recorded. The actual dates tended to be within a range of ±2 days from the actual event. Table 46 lists a subset of field service observations that may be caused by a malfunctioning FCU.

The regional airline data, along with the associated FCU replacement observations, are used to design, develop, and test the engine SPMS.

Table 46. Historical data available for engine FCU-targeted SPMS interface testing

	5	Engine will not start	FCU replaced
	9	Engine went sub-idle and over-temped	FCU replaced
	10	Engine slow light off	FCU replaced
	15	Intermittent engine on fire alarm due to fuel problems	FCU replaced
	17	Engine hung (twice)	FCU replaced
	17	Engine idling high at 55% N2	FCU replaced
	19	Over-speed temperature and engine shutdown	FCU replaced
	29	Engine starting hot	FCU replaced Fuel leak

7.2.1 Design of Engine SPMS

An analysis of table 46 shows that six out of eight cases resulted in engine startup issues. This is not surprising considering that engine startup is a key stressor. Furthermore, the startup also amplifies control issues and provides opportunity to observe response-related features. Therefore, analysis of the startup profile is essential for the design of engine SPMS functionality. The fuel-controller problems can lead to poor regulation, which results in too much or too little fuel being supplied to the engine. Too little fuel is generally caused by obstruction in the fuel lines due to, for example, debris accumulation, damage, ice buildup, or improper valve operation that eventually results in engine flameout or hung start. Too much fuel in the engine due to valve and regulation faults results in higher EGT. As the fault progresses, EGT increases, which eventually exceeds the safe limit of an aircraft; it experiences a “hot start.” It is caused by too much fuel entering the combustion chamber or insufficient turbine rpm. If an engine has a hot start, appropriate maintenance and an inspection are required. Figure 19 shows the startup profile for one engine. Key milestones and parameters of the engine start, such as starter on, light off (after engine reaches stoichiometry), and idle speed are also shown in the figure. The acceleration of the engine post light on calculated by calculating the derivative of the engine speed and gradient of the EGT also serve as a measure for controller effectiveness. Detailed analysis of the startup measurements are performed to identify the observable features for the engine SPMS.

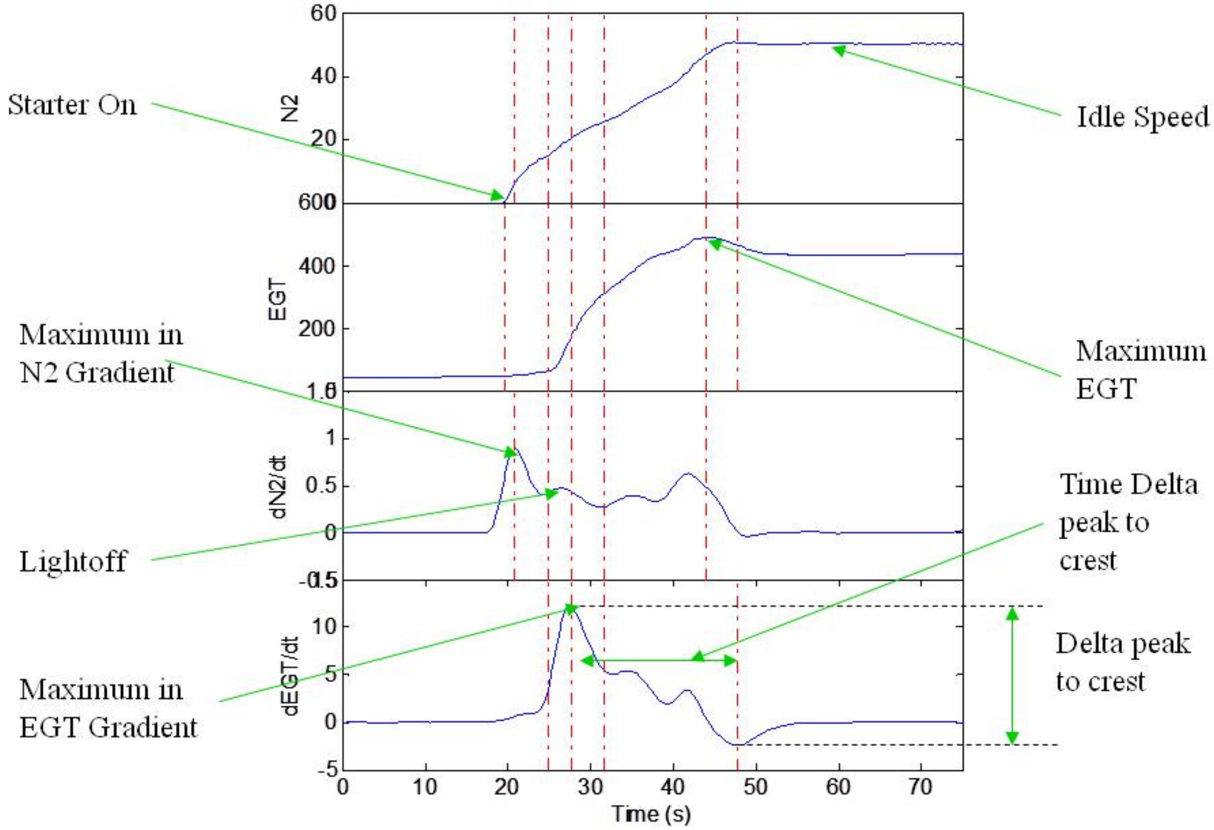

Figure 19. Startup profile—engine core speed, N2, EGT, and derivatives

7.2.2 Preliminary Analysis Engine SPMS

There are 182 parameters recorded in the regional airline records. A preliminary analysis of the field data used in the interface definition task above identified the list of sensors to be used for design. These data fields (see table 47) form inputs for three algorithms that are used to generate the FCU CIs and HIs at startup, takeoff, and engine roll down. The snapshots detect temperature, speed excursions, and peaks along with profile slopes and time (see figure 20).

Table 47. Engine SPMS signals recorded per flight

Signal Name	Description	Sample rate (per sec)	Relationship with FCU faults
DateTime	Time recorded per sample or at the beginning of record	4/1 per set	Needed for association and timeline
EGT	Engine EGT for all four engines	4	Correlated
TAT	Total Air Temperature (ambient)	4	Used for correction
N2	Engine core speed (rpm) for all engines	4	Correlated
N1	Fan speed	4	Not used
PLA	Power lever angle or engine thrust setting	4	Not used
PALT	Pressure altitude	4	Used for flight vs. ground determination
PH	Flight phase from ACMS	4	Redundant info
Engine Position	Identify the engine position for the collected data	1 per set	Engine position dependent
BeginHours	Begin hours of engine operation	1 per set	
EndHours	End hours of engine operation	1 per set	
BeginCycles	Begin engine cycles (starts)	1 per set	
EndCycles	End engine cycles (starts)	1 per set	
Flight Track	Latitude, longitude	1	

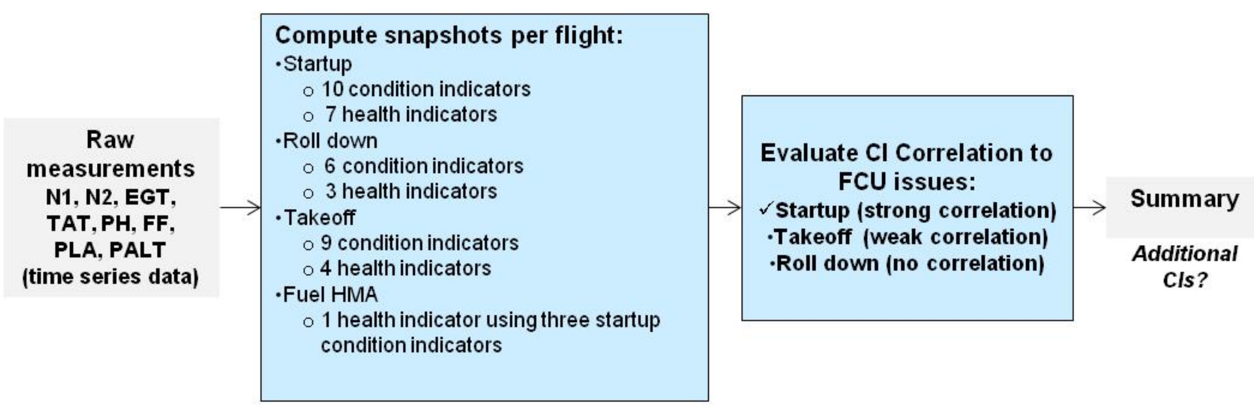

Figure 20. Steps used to screen engine FCU-related CIs

The CI algorithms were used to analyze and compute snapshots per flight for corresponding to engine startup, aircraft takeoff, and engine roll down (see figure 20). The CI with a threshold is also used to generate corresponding HI. The analysis also included results from the engine hydro-mechanical actuator (HMA) monitor. Figure 21 shows results for 230 flights for relevant takeoff monitors and HMA heath indicators. The x-axis of the plots shows "flight index," which indexes the flight from removal, with "0" corresponding to first flight post-FCU repair/removal (marked on the plots with a cyan line). The negative flight indices show flight before repair, whereas the positive indices show flight post-repair. In this case, engine 4, labeled as "E4" in all of the plots, has FCU issues that are eventually repaired. The other FCUs are treated as healthy and the CI spread from them is used to determine the natural variability of the CI.

The screening results show that the startup CIs are correlated to the FCU condition. The specifics of the results are subsequently discussed. The roll-down CIs do not show any correlation to the FCU health. The takeoff monitors in figure 21(a) show that both the "take off N2" (engine speed at takeoff) and "peak EGT" are correlated to the engine condition. However, the peak EGT from engine 4 is within the spread of the CIs from the other engines. The E4 engine speed at takeoff shows a spread that is beyond the nominal CI from other engines. However, it has a region of overlap and is not clearly distinguishable. In summary, the engine takeoff CIs show correspondence but weak discrimination capability (i.e., overlapping nominal and faulty CI distribution). The HMA HIs are driven by startup CIs such as high idle speed, fast startup, and hot start. In addition, figure 21(b) shows that the HMA HIs do not show a good discrimination capability for this data set. The results show that the high idle HI (at engine startup) is triggered often and shows some correlation to the fault. The fast hot start and fuel HMA HI (logical combination of high idle and fast hot start) are not triggered consistently. Fuel HMA HI is not detecting the FCU faults consistently. Because the HMA HIs are driven by startup CIs, better CI threshold may improve performance.

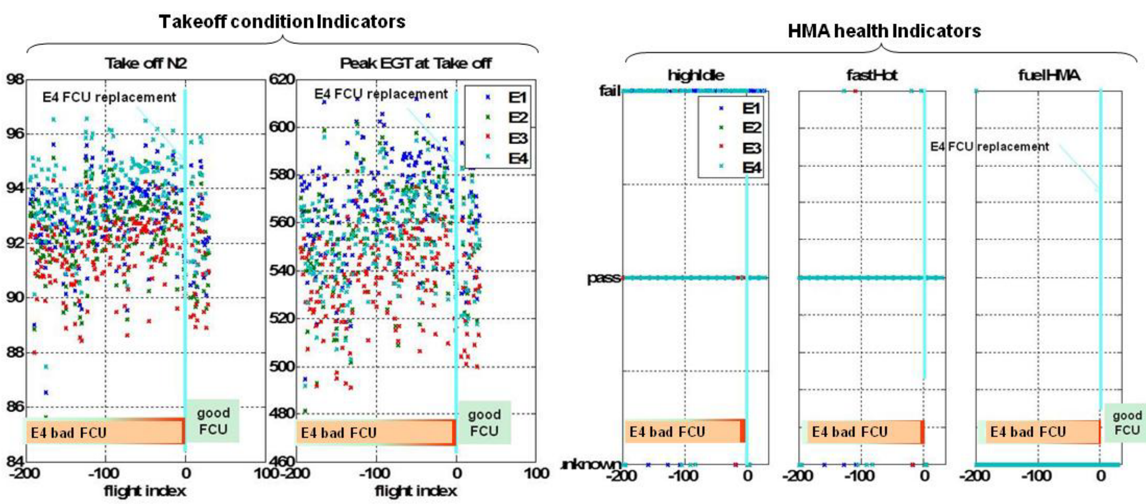

(a) (b)

Figure 21. Takeoff and fuel hydro mechanical actuator super monitor response

Table 48 shows a signal-to-FCU fault correlation evaluation summary condensed back to the raw signal measurements. Based on this analysis, the startup CIs are selected for further development.

Table 48. Engine SPMS signals used

	Description	Relationship with FCU faults and comments
	Engine EGT for all four engines	Correlated
	Total air temperature (ambient)	Used for correction
	Engine core speed (rpm) for all engines	Correlated
	Fan speed	Not correlated
	FF (lb/hour) for all engines	Does not measure FF
	Power lever angle or engine thrust setting	Not correlated
	Pressure altitude	Used to determine flight vs. ground runs
	Flight phase from ACMS	Redundant info

Figure 22 shows the steps for engine SPMS algorithm development/testing. The screening tests shown in figure 20 represent step "0" of algorithm development. The algorithm development and option selections are discussed below:

1. Align the airline data with the FCU replacement field service records. FCU replacement records typically do not identify the engine position or serial number. Furthermore, date of replacement was found to be approximate. This study is using a ±2-day alignment window.
2. This is a data-labeling step. The idea is to distribute the data into two classes:

    a. Before the replacement: The CI from before replacements should show the presence of underlying fault. This assumes that the fault is observable through the respective CIs. There is an additional assumption regarding the slow progressive nature of faults. This implies that the observed CI shall indicate increasing severity or a prognostic trend for the fault.
    b. After the replacement: The CIs after the replacement should show the absence of fault.

3. In this step, the raw sensor values are read from the input file. The N2 and EGT are corrected using the ambient air temperature to create the total air temperature (TAT). The EGT and N2 profile are used to identify the engine startup. Startups, failed start, and hung start are identified. In addition, the algorithms identify ground runs and flights for each profile. This step is critical for downstream analysis and algorithm development.
4. Filter out profiles with non-auto/failed starts and ground runs.
5. The signals are filtered and the smoothed rate d/dt is computed for the corrected N2 and EGT.
6. Compute CIs/features associated with startup.

7. Identify the CI trends that are impacted by FCU fault.
8. Apply the decision logic to test and generate the FCU alarms. Measure the effectiveness of the engine SPMS.

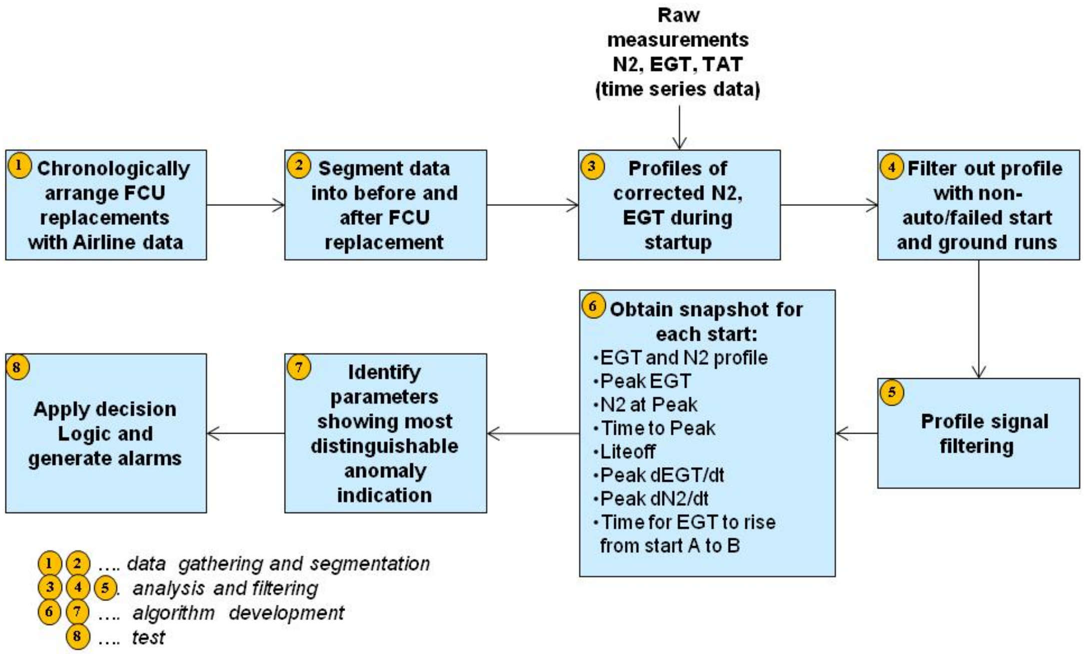

Figure 22. Engine SPMS FCU algorithm development

Startup profiles at the end of step 5 are plotted and analyzed. Figure 23 shows the N2 and EGT profile for 230 consecutive flights from the same engine. The blue lines show 30 flights post-FCU replacement. The red lines show the profile 30 flights before FCU replacement. Cyan lines represent 200 flights before and 30 flights after the FCU replacement. The engine no-starts and startups corresponding to ground runs are not plotted. The startups are lined when N2 exceeds zero. Figure 23(a) shows the N2 at zero due to filtering. The fuel in the turbine is introduced after it has reached 25% of the idle speed. The following observations can be made based on the analysis shown in figures 23(a) and (b):

- Ambient condition correction: For N2, ambient condition correction (using TAT) is working. The corrections for TAT do not line up all the profiles.
- Filtering: Filtering of N2 profiles introduces a steeper gradient (at the start of the profile). This gradient is ignored because it is before the turbine light on and, therefore, is not affected by the FCU condition.
- Alignment of EGT profiles: The EGT profiles do not temporally line up. This is because of variability in the starting conditions of the turbine.
- N2 profile: Based on the FCU condition (between blue and red/cyan), the N2 profile looks different. However, the N2 profile shows large variability even with the

healthy/repaired FCU. This could be there because of insufficient ambient-conditions correction. The N2 rate (dN2/dt) for prior to repair is higher than post-repair.

- EGT profile: The EGT profile before the repair show a faster rise time. This is clearly visible in the EGT rate (dEGT/dt) plot. The rate of EGT rise is much higher prior to FCU repair.
- EGT and N2 profiles (and the rates) do not show a trend that decreases/increases with each flight; instead, the profile shows that healthy and faulty engines can be identified prior to the on aircraft BIT and removal.

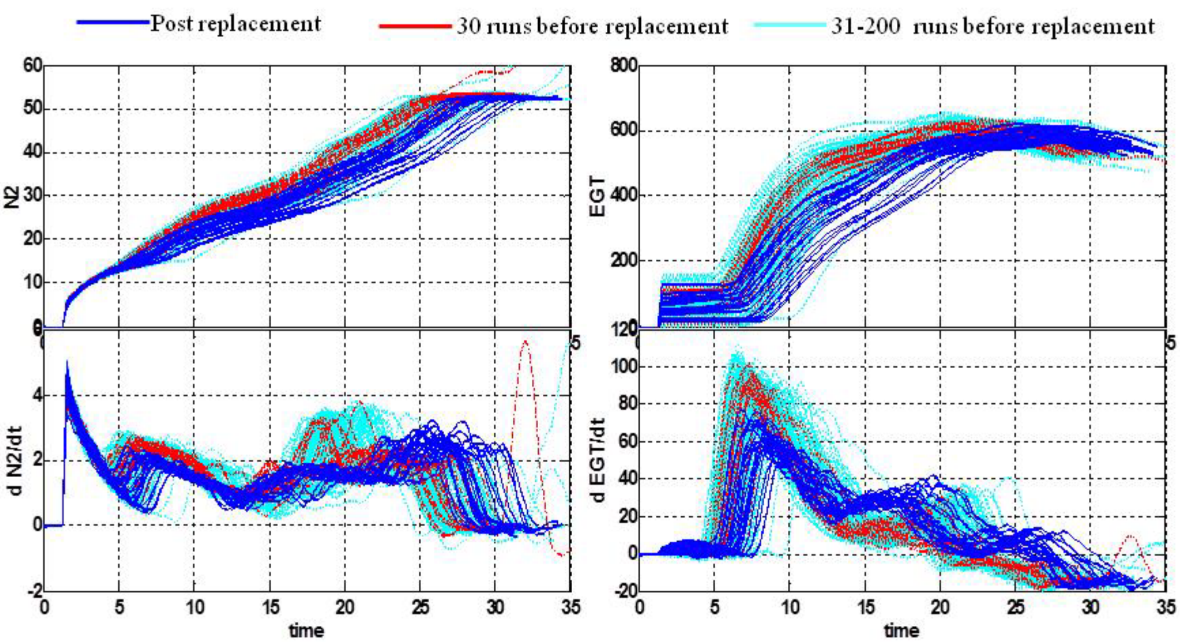

(a) (b)

Figure 23. Start profile for 200 flights before and 30 flights after the FCU replacement for engine 4

Next, the analysis of the entire engine CIs index from the FCU removal event was computed. Results that show some of the promising indicators are shown in figures 24 and 25. The results presented are for 200 flights before FCU removal.

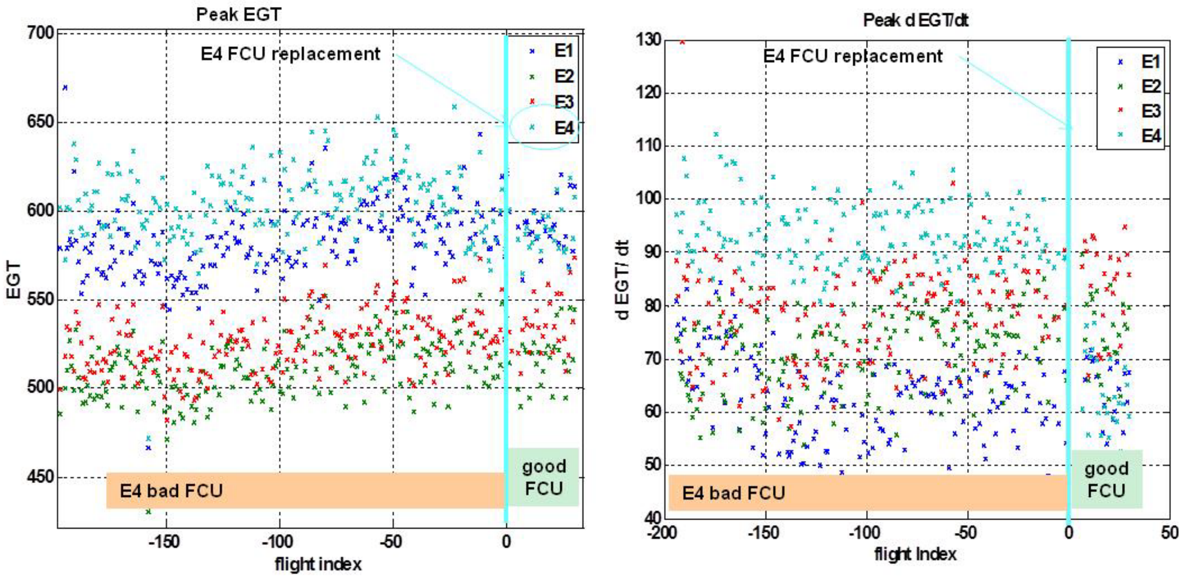

(a) (b)

Figure 24. EGT-based CIs using (a) peak EGT and (b) peak dEGT/dt

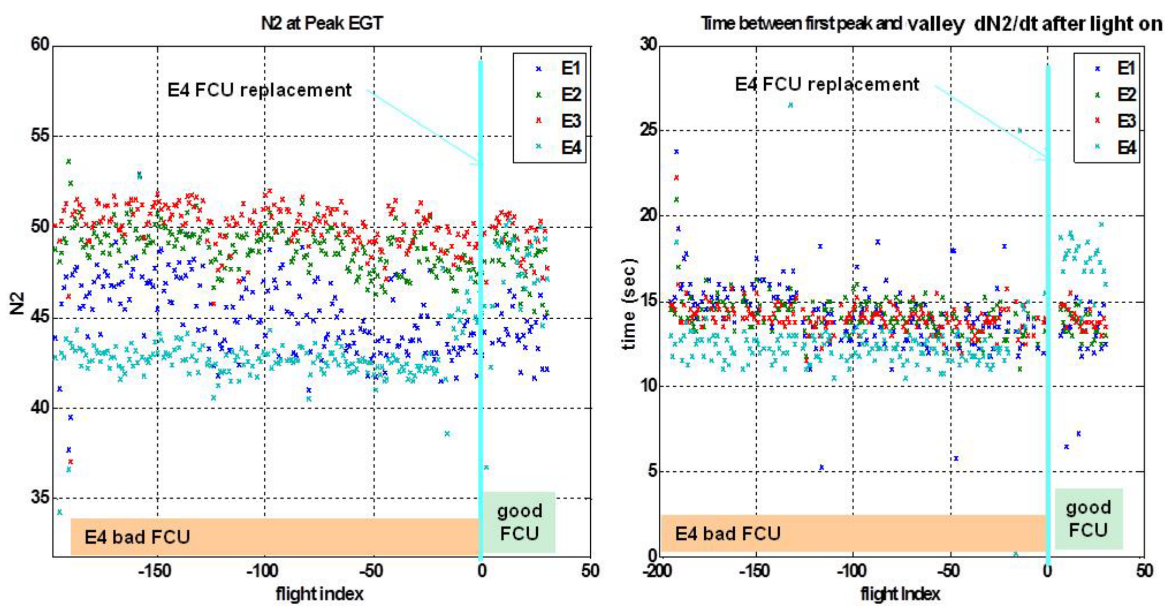

(a) (b)

Figure 25. CIs based on N2 and EGT interplay:
(a) N2 at peak EGT and (b) time between first peak and valley of dN2/dt after light on

Figure 26 shows the analysis for 1000 flights. Column 2 of figure 26 shows the histogram bin of the filtered CI. For this analysis, a 10-sample moving average filter was used. It can be seen that the individual CI are not capable of predicting the FCU removal. However, analyzing all four together, a truth table or weighted indicator that distinguishes before and after the FCU removal can be determined. This exercise is made difficult due to an incomplete record of field symptoms evident through a comparison of figures 26–29.

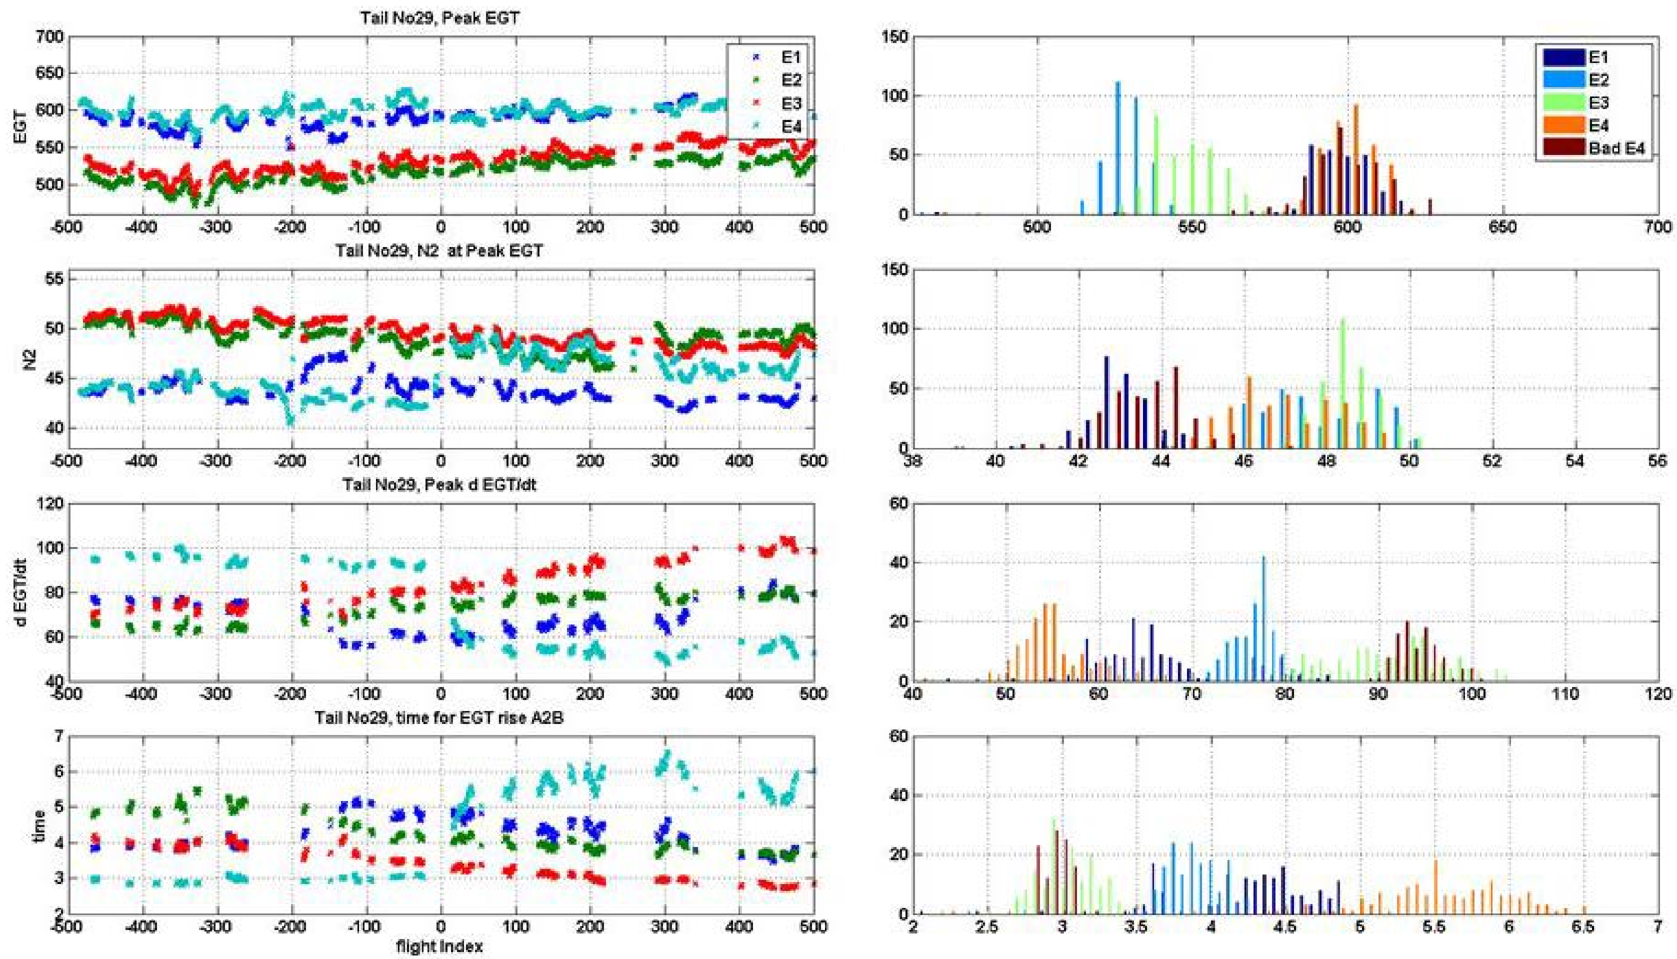

Figure 26. Filtered indicator with histogram (column 2) for tail 29

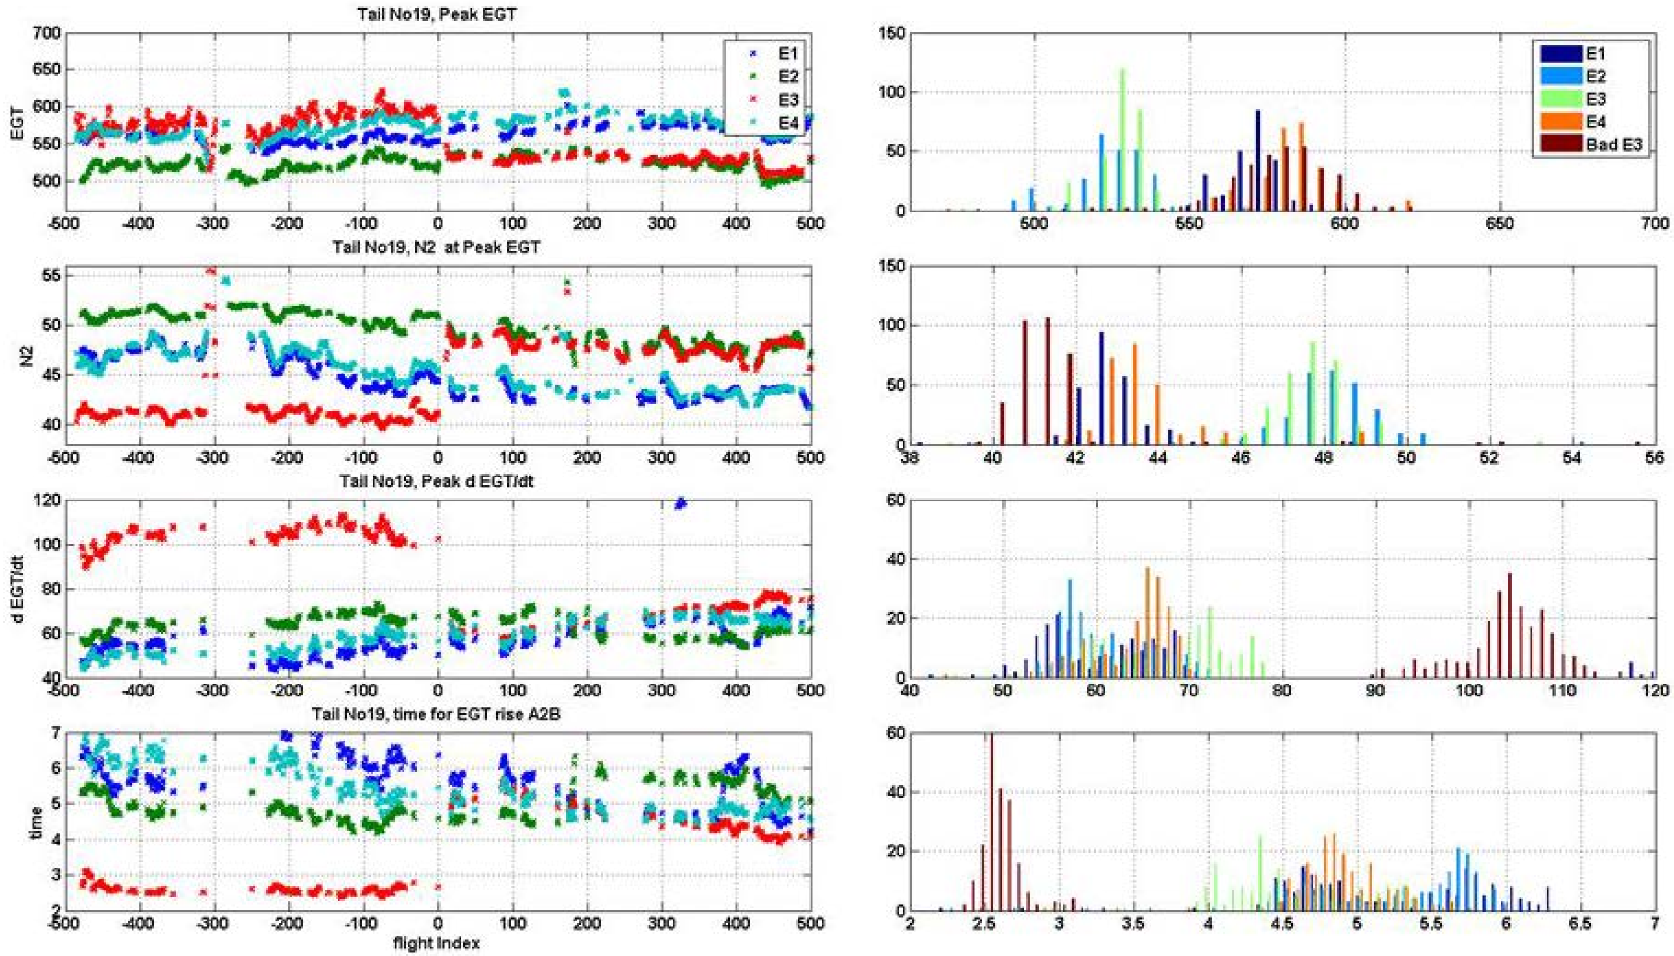

Figure 27. Engine 3 running hot, FCU replaced tail 19

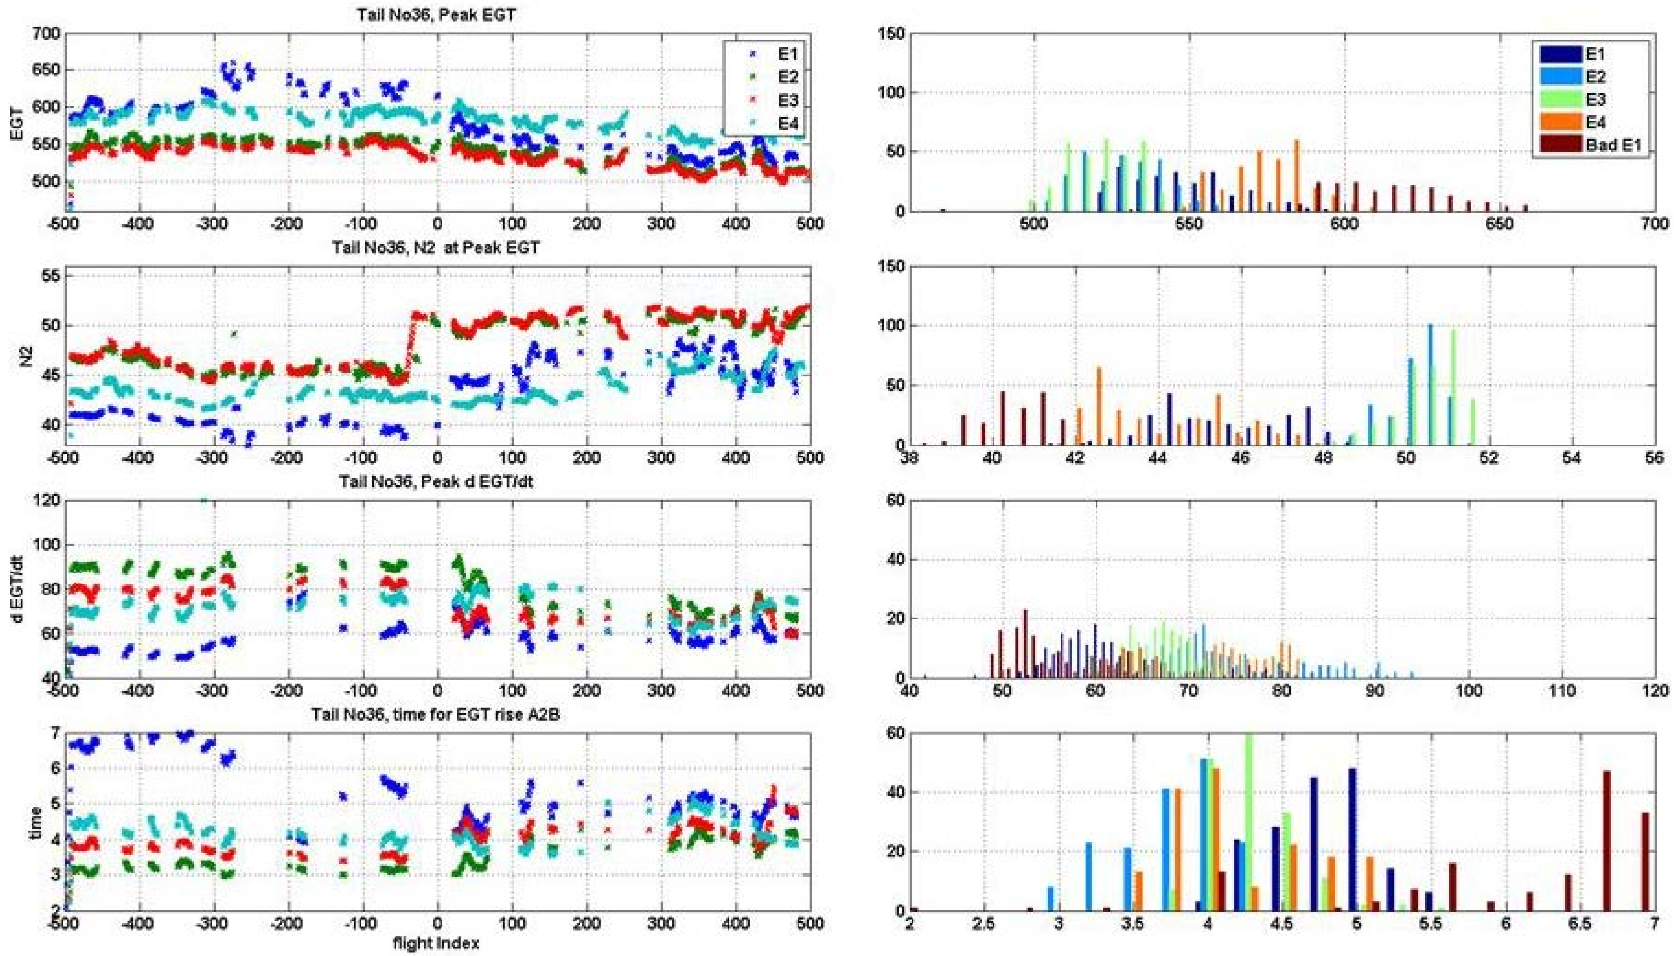

Figure 28. Engine 1 running hot, ECU replaced, filtered CIs with histogram (column 2) for tail 36

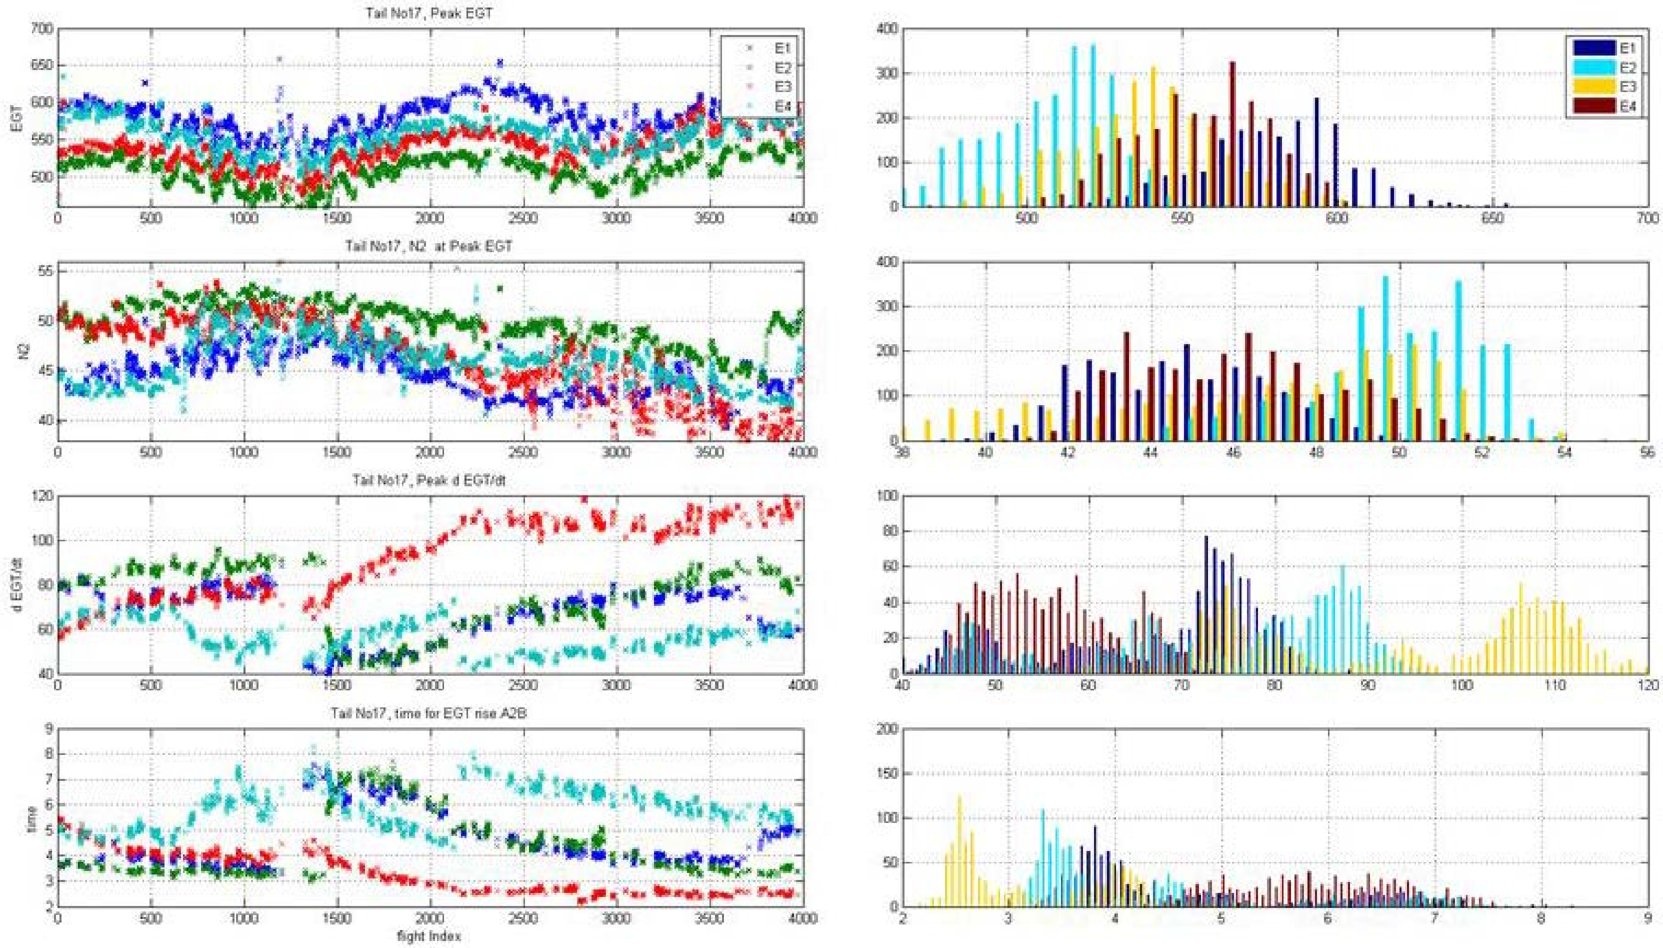

Figure 29. Engine CIs for 4000 flights (tail 17) show the spread of the individual CIs (includes multiple engine removals)

7.3 APU SPMS

7.3.1 Predictive Trending of EGT_{margin}

An APU is a small turbine engine that provides auxiliary power. With their protective coating, the turbine blades of the APU are designed to withstand the high temperatures of the combustion gases. However, this coating degrades over time, leading to permanent damage. As the turbine blade tips erode, their ability to extract mechanical energy from the hot combustion gases decreases. The fuel controller increases FF to meet demands and the "left over energy" is sensed by the EGT probes called EGT sensors. Excessive EGT is an indicator of hot section deterioration.

EGT varies for reasons other than deterioration, including environment (ambient air temperature and pressure) and load on the APU (generator and load compressor). Having other sources of variation is undesirable for a CI. The derived quantity EGT_{margin} is a more salient indicator of hot section deterioration than EGT. EGT_{margin} uses the measured EGT at a particular time, MES, when the load demands on the APU are generally highest. EGT_{margin} uses model-based corrections to remove variations attributable to environmental and load conditions. The corrected EGT value, $EGT_{corrected}$, represents what the EGT value would have been if it were measured under fixed reference conditions. The fixed reference conditions are chosen to be indicative of the APU meeting its requirements under relatively challenging conditions.

The EGT_{margin} is the difference between $EGT_{corrected}$ and a threshold called EGT_{design}:

$$EGT_{margin} = EGT_{design} \quad EGT_{corrected} \qquad (36)$$

A healthy APU has an $EGT_{margin} > 0$ and trends downward as the hot section slowly deteriorates with use. $EGT_{margin} < 0$ represents an APU that is undesirable to operate further. Nevertheless, operators can, and do, sometimes continue operating the APU beyond the threshold.

The threshold EGT_{design} can address several issues:

- Setting EGT_{design} to a lower corrected EGT can reduce the cost to repair the deterioration and return the APU to service. Hot-section temperature is both an indicator of deterioration and a cause of further deterioration. Operating at high temperatures causes severe deterioration throughout the hot section, which is costly to rebuild. The initial deterioration is often at the blade tips. Setting EGT_{design} to a lower corrected EGT can limit deterioration to the blade tips, which are much less costly to refurbish.
- Setting EGT_{design} to higher amounts to keep APUs on-wing longer despite deterioration and, therefore, fewer remove and replace actions.
- A more deteriorated APU (higher EGT_{design}) has poorer fuel efficiency and emissions.
- Eventually, deterioration affects functionality. The APU controller monitors the EGT and prevents it from exceeding an upper limit, where significant damage would occur. It does this by modulating the IGV closed to reduce the bleed air load and also reduce fuel flow. The limit is referred to as the "IGV Trim Limit." An APU that has deteriorated to this

point is not meeting its functional role for MES. EGT_{design} is set below the IGV Trim Limit. For the APU model selected for SPMS, the limit varies with ambient temperature approximately 1-to-1 and IGV Trim Limit = 1220 F when ambient temperature = 100 F.

- The APU controller performs autoshutdown of the APU if the EGT exceeds the "Overtemperature Shutdown Limit." The Overtemperature Shutdown Limit is higher than the IGV Trim Limit, so the EGT_{design} is set well below this point.

The following parameters are used to compute $EGT_{corrected}$ and EGT_{margin}. All are acquired during MES:

- EGT (degrees C or F): For the APU model chosen for SPMS, healthy EGT values are in the range 995–1095 F, corresponding to an EGT_{margin} of 200–1000 F. The EGT_{design} might be 1195 F or 1220 F, depending on the operator.
- T2, inlet air temperature (degrees C or F): This is the ambient air temperature at the inlet of the APU and is arguably the most important parameter for computing $EGT_{corrected}$ from EGT. Experience shows T2 extremes can differ by over 100 F due to season, location, time of day, and/or weather condition. The T2 correction is performed by table lookup but is approximately a 1.5-to-1 ratio for EGT increases with T2. Therefore, variations in T2 could induce 150 F variations in EGT, which would obscure trends in EGT_{margin} if left uncorrected.
- P2, inlet air pressure (psi): This is the ambient air pressure at the inlet of the APU. Lower P2 results in higher EGT.
- L, generator load (kW or % of full load): This is the load on the electrical generator driven by the APU. Higher electrical load requires more power from the APU, more fuel, and, therefore, higher EGT.
- IGV, IGV position (degrees): This is the measured angular position of the IGVs, which the controller uses to modulate bleed air production. For the APU used for SPMS, 90 degrees corresponds to fully open, and 0 degrees corresponds to closed. Normally, the controller holds the IGV fully open during MES to provide maximum bleed air for starting. When the APU is severely deteriorated, the APU controller modulates IGV partially closed to keep EGT at or below the IGV Trim Limit.

7.3.1.1 Slow Smooth EGT Trend

Figure 30 shows an example of the parameters (blue graphs in each subplot) listed in section 7.3.1 for an APU operating in the field over a complete lifecycle (a lifecycle is the period from installation until removal). The parameters were recorded once per flight. The figure also shows the computed quantities of $EGT_{corrected}$ (green graph in upper left subplot) and EGT_{margin} (green graph in lower left subplot). The red lines in figure 30 show smoothed quantities.

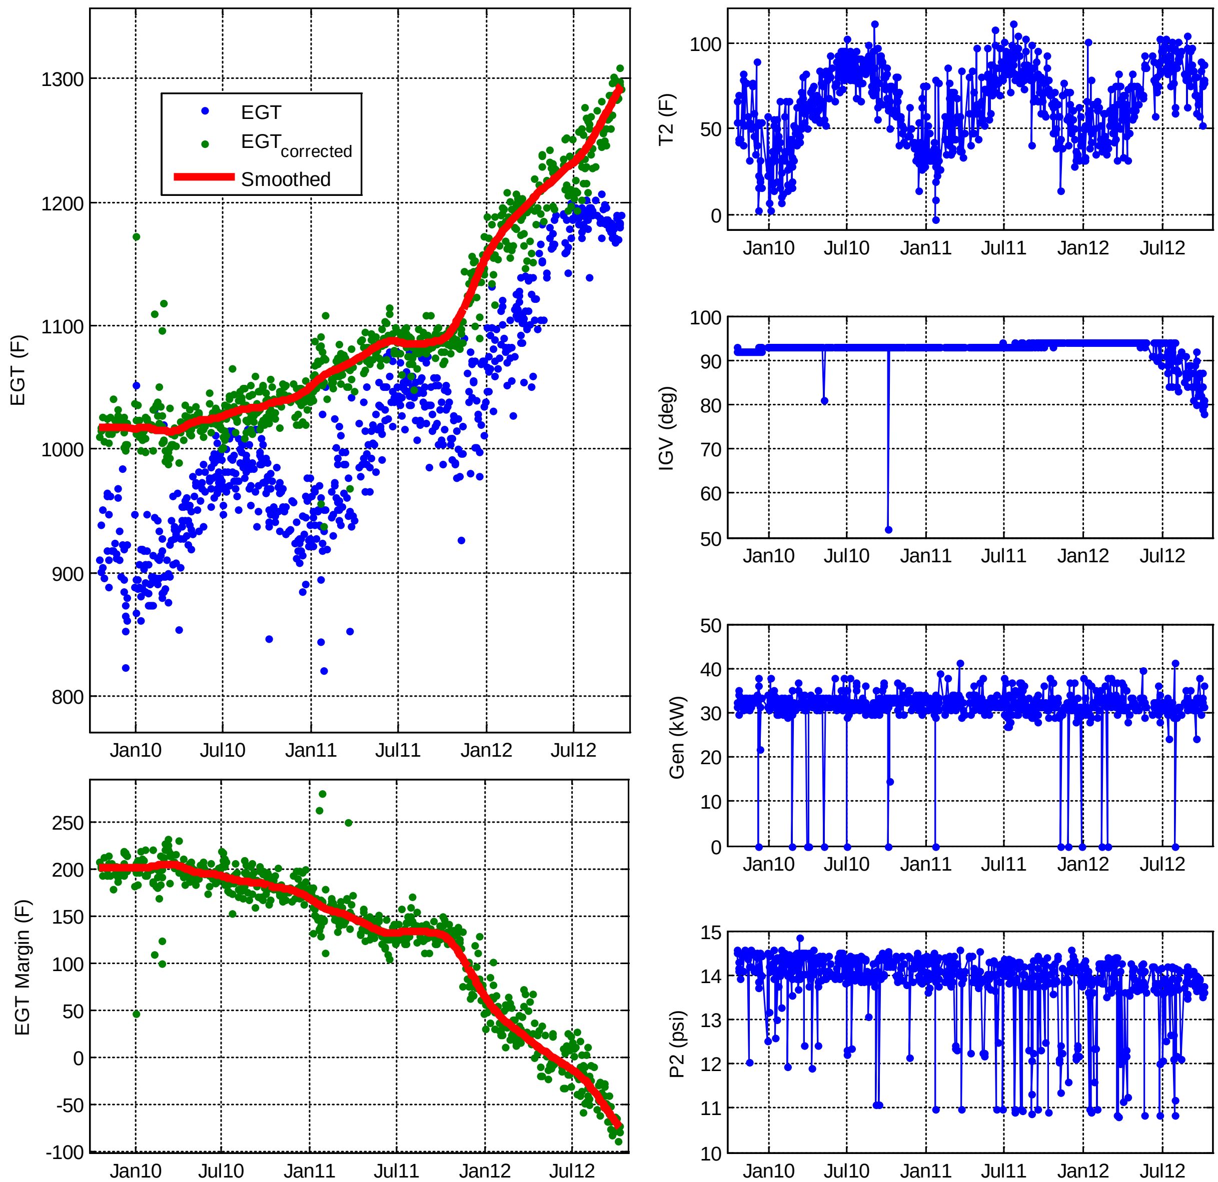

Figure 30. Example lifecycle of EGT trending parameters

For the APU lifecycle depicted in figure 30, the EGT_{margin} started relatively flat at +200 F. It began a slow downward trend approximately April 2010, transitioned to a steeper trend in October 2011, and crossed through 0 F in May 2012. The zero crossing is when Honeywell would recommend removing the APU for servicing. The need for IGV modulation began near this date. The operator nevertheless kept the APU in service another 5.5 months and eventually removed it in October 2012 because of autoshutdown and no-start issues (according to the maintenance shop report). By this time, the EGT_{margin} had reached -75 F. The maintenance shop report described the APU as having medium to heavy hot-gas erosion of the turbine wheels.

The EGT_{margin} graph in figure 30 shows that this technique enables the operator to monitor the graceful degradation of the APU hot section over months and years.

For the APU lifecycle shown in figure 30, the ambient temperature T2 exhibits high frequency variation, due to location, time of day, and weather, superimposed on more regular seasonal cyclic variation. This pattern of T2 is common. The effect of the seasonal T2 variation on EGT is visible in the EGT graph, which shows seasonal variation superimposed on an upward trend. The $EGT_{corrected}$ graph shows that the T2 correction removes the seasonal variation from EGT quite well.

The IGV is essentially constant at 93 degrees for most of the lifecycle and then decreases near the end when the controller modulates IGV to keep EGT below the IGV Trim Limit. The effect of the controller modulation on EGT is visible in the EGT graph: EGT stops its upward trend and becomes less noisy. If the controller had not modulated the IGV, the EGT would have continued increasing along a trend like $EGT_{corrected}$. When the controller is modulating IGV, noise in the EGT decreases because the EGT becomes a controlled parameter. Residual noise in EGT is partly due to the fact that the IGV Trim Limit is a function of T2 and increases approximately 1-to-1 with T2.

The generator load, L, is near 33 kW for this lifecycle, with a few downward excursions. Close inspection shows that many of the excursions to 0 kW coincide with downward spikes in the EGT graph, suggesting that the downward spikes are not simply data dropouts. A downward spike from 30 kW to 0 kW should produce a 45 F downward spike in EGT based on the tables used to correct EGT for L.

The ambient pressure P2 hovers at approximately 14 psi for the entire lifecycle, with downward spikes to 11 psi. These downward spikes are only expected to produce 11 F upward spikes in EGT, which is in the noise and not visible in the EGT graph.

The correction parameters fix the offset in EGT, not just noise. The values of T2 and L in this lifecycle are generally lower and less challenging than the fixed reference conditions of 100 F and 65 kW used for computing $EGT_{corrected}$ for this model of APU. Recall the fixed reference conditions are chosen such that EGT_{margin} is indicative of the APU being able to meet its requirements under relatively challenging conditions. The resulting difference between $EGT_{corrected}$ and EGT is approximately 120 F in the winter and 50 F in the summer.

The smoothed lines in figure 30 were computed using a common technique called "robust LOWESS" (locally weighted scatterplot smoothing) that is built into MATLAB. The term "robust" refers to the algorithm's ability to discard outliers so that they do not skew the solution; the seven outliers in the EGT_{margin} graph in figure 30 show that this is important. The choice of LOWESS over other possible robust smoothing/filtering techniques was simply a matter of convenience.

7.3.1.2 Sudden Drops in EGT_{margin}

Not all lifecycles trended as gracefully as figure 30. A minority of lifecycles had a sudden drop in EGT_{margin} somewhere in the trend. Figure 31 provides an example. In August 2012, the EGT_{margin} made a sudden drop from +90 F to -20°F between successive flights. Even so, the operator was able to continue using the APU until December 2012, approximately 4 months after the sudden drop. The operator removed the APU because of oil smell in the cabin. The report from the repair shop indicated finding coking of the gearbox vent line and fan/load compressor witness (consistent with the oil smell). The report also indicated finding high-time wear-out (consistent with negative EGT_{margin}). Unfortunately, the report from the repair shop for this particular lifecycle did not provide any findings that would explain the sudden drop in EGT_{margin}. Other reports have sometimes reported turbine blade shift or cracked liner, which might cause a sudden drop.

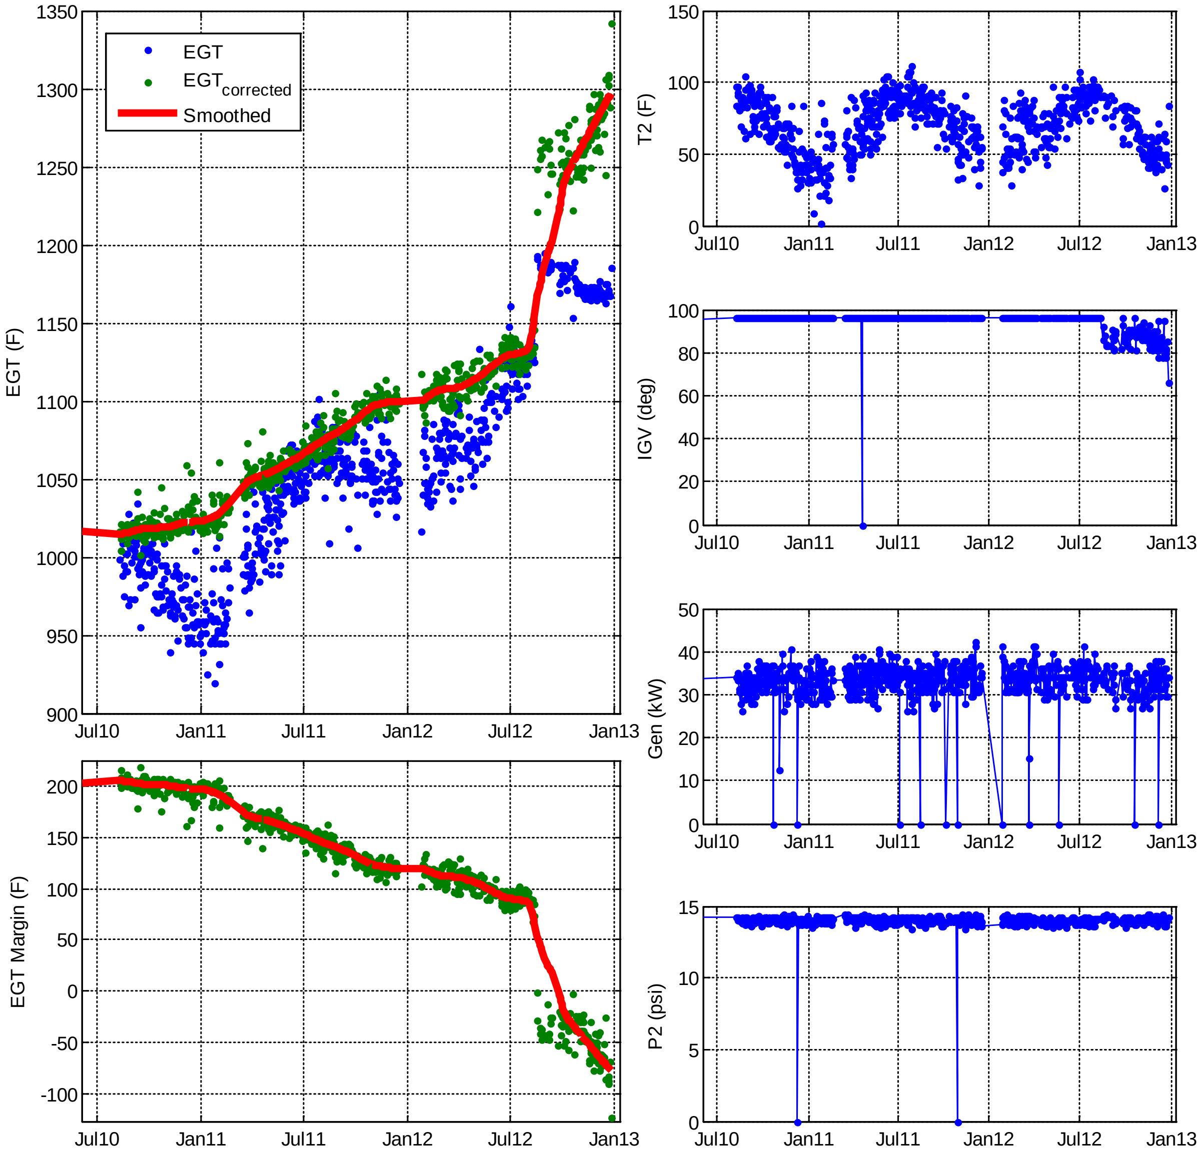

Figure 31. APU exhibits a sudden drop in EGT_{margin}

7.3.1.3 Bi-Nominal Distribution

There is evidence in the historical data that there is at least one additional as-yet undetermined correction parameter that could improve the computed values of $EGT_{corrected}$ and EGT_{margin}. Figure 32 provides an example lifecycle that highlights this issue. The graphs of EGT_{margin} and

$EGT_{corrected}$ in figure 32 show a bi-modal distribution, with a separation of approximately 50 F between the two modes. Other lifecycles collected by Honeywell exhibit this bi-modal distribution, though not the majority. The separation is consistently 50 F. The cause of the bi-modal distribution is unknown. There is nothing unusual in the correction parameters T2, IGV, L, or P2 that could have produced this effect and other parameters in the recorded data were checked and none were found that is correlated with the bi-modal distribution. The unknown parameter appears to be discrete and binary, because this could produce a bi-modal distribution (e.g., EGT increases by +50 F when the parameter is true).

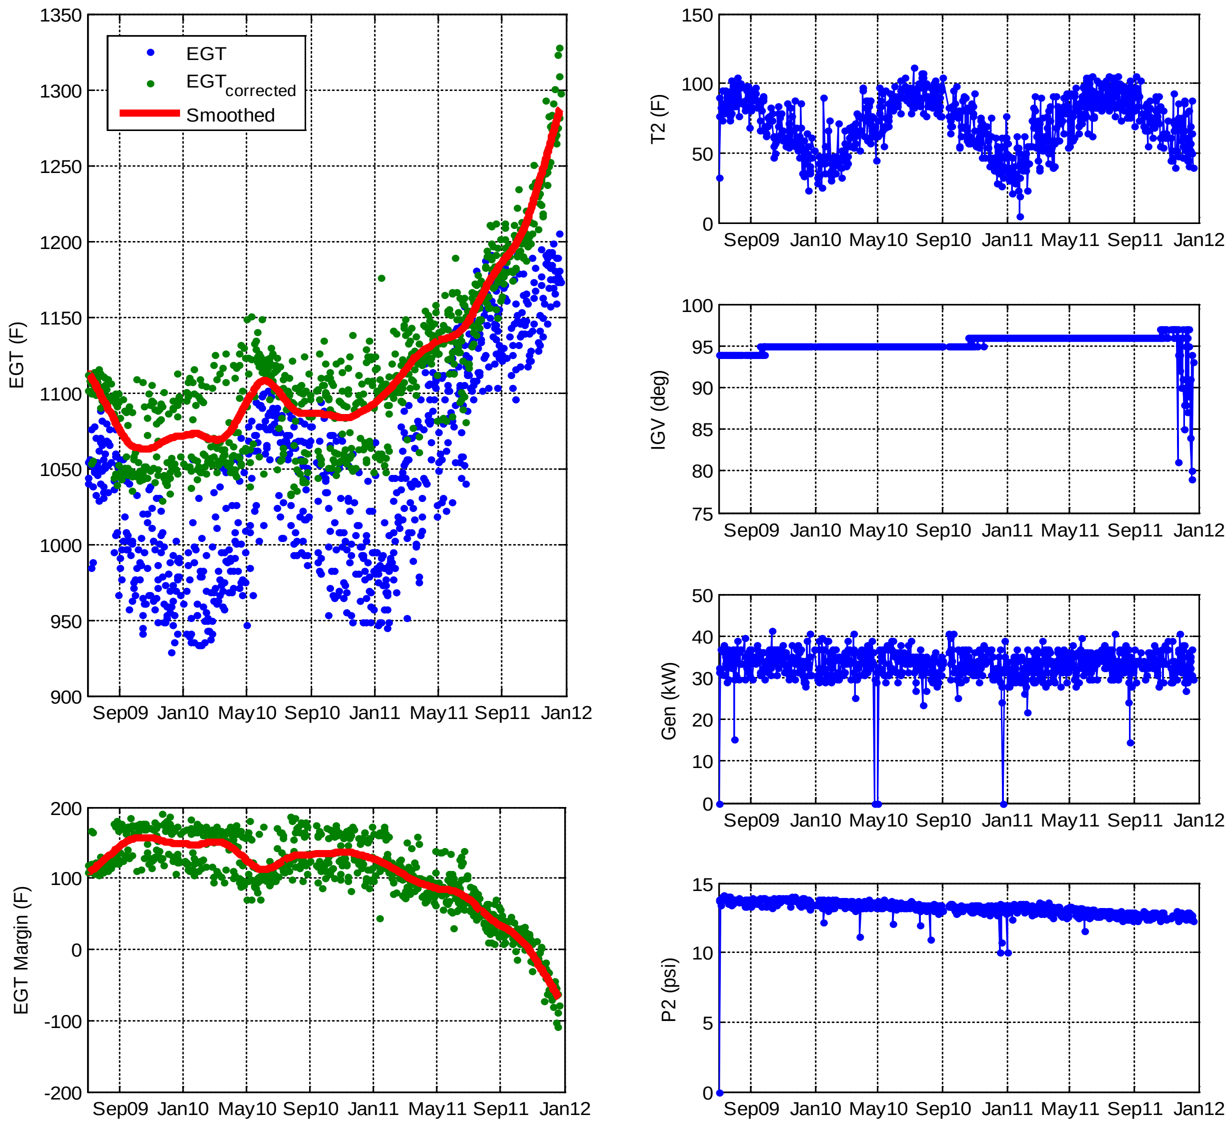

Figure 32. APU exhibits a bi-modal distribution for EGT_{margin} and $EGT_{corrected}$

7.3.1.4 Failure Probabilities

The repair shop records produced at the end of 300 APU lifecycles were analyzed. The repair shop records contain information including:

- Why each APU was pulled from service (e.g., autoshutdown, no-start, time-based removal, order in cabin).
- What fault was found in the APU (e.g., hot section erosion, seal leakage, no fault found, APU ok).
- Identifiers to associate repair shop data with the EGT, T2, etc. measurements acquired in the field.

The repair shop records allow the authors to characterize how often trending the EGT_{margin} successfully tracks hot-section health and predicts the need for repair. Examples of lifecycles considered successful include:

- The EGT_{margin} begins at some high value at the start of the lifecycle, gracefully trends downward over months or years, crosses through zero, the APU is removed from service due to problems consistent with hot-section deterioration (e.g., autoshutdown), and the shop finds there was indeed significant hot section deterioration.
- The EGT trends gracefully downward to zero and the operator removes the APU based on $\mathrm{EGT}_{\mathrm{margin}}$. The shop finds there was indeed significant hot-section deterioration.
- The EGT drops suddenly to/through zero, but the operator continues to use the APU with $EGT_{margins}$ below zero for a month or more. The shop finds there was indeed significant hot-section deterioration.
- The EGT_{margin} trends gracefully through zero and the operator continues to operate the APU far longer than recommended. Eventually, some other fault causes the APU to be removed. The shop finds there was hot-section deterioration and the other fault.

Examples considered unsuccessful include:

- The APU is removed from service because of some fault other than hot-section deterioration while the EGT_{margin} is still acceptable.
- The EGT_{margin} drops suddenly through zero and the operator needs to remove the APU very soon afterward. This is a case in which EGT_{margin} lacked predictive power.

Table 49 shows the results of the analysis. The left-most portion shows the successful lifecycles. There were 121 lifecycles considered successful, most of which exhibited a graceful trend.. A few had a sudden drop but still provided predictive value.

The middle portion of table 49 shows the unsuccessful lifecycles. Most are where the APU was pulled from service due to some fault that is not detected by EGT_{margin}. A few are where EGT_{margin} lacked predictive power because it dropped suddenly and the APU was unusable.

Table 49. Lifecycle statistics

Lifecycles ending in failure trendable by EGT	
EGT trends to failure	97
Assumed (gaps in data)	14
But 0<EGT<60	2
Yes, but also other fault	2
Yes, but sudden	3
By sudden drop	3
Total:	121

Other lifecycles ending in failure	
Failure not seen w/ EGT	59
Assumed not (gaps in data)	14
External contamination	3
EGT drops too sudden	1
Inspection detected	3
Other fault primary	5
Total:	85

Lifecycles ending without failure	
Time-based early removal	49
Lease end-OK	23
Convenience-OK	9
N/A-No data	6
N/A-Service bulletin	7
No fault found	6
Precautionary early removal	3
Request early removal	2
Time-based, turbine shifted	2
N/A	1
N/A-Faulty or not?	2
N/A-Inspection early removal	2
N/A-Maintenance error	2
Total:	114

The left and center sections of table 49 show that of the 206 (=121 + 85) lifecycles that end in failure, 59% are successfully tracked by EGT_{margin} trending and reasoning.

The right section of table 49 shows 114 of the 320 lifecycles did not end in failure or degradation. Typically, these are because of time-based removal and/or end of lease. Presumably, the 59% success rate observed in the failure lifecycles would apply to these APUs left in service.

7.3.1.5 Requirements Summary

Based on the discussion in the preceding sections, the following are required for an EGT trending reasoned:

- EGT should be acquired and recorded at a fixed condition that represents a challenging requirement for the APU (e.g., during MES). The use of a fixed condition lessens the fluctuations that are caused by conditions and need for corrections.
- The reasoner should compute a continuous-valued EGT_{margin} that is an indicator of how close the EGT is to hitting limits during challenging conditions.
- The noise in EGT_{margin} necessitates smoothing or filtering by the reasoner.
- The smoothing or filtering should not obscure sudden drops in EGT_{margin}.
- Because the trend in EGT_{margin} is so slow, it is sufficient to acquire EGT only once per flight. Less frequent acquisition is tolerable, especially during the flat part of the trend. However, too infrequent acquisition would negatively impact smoothing.
- T2 should be acquired and recorded for each EGT. The reasoner should use T2 to remove variations and offsets in EGT that are attributable to T2. The T2 correction can be 150 F on a cold winter day for this APU.
- Because T2 is expected to vary on timescales of minutes or hours, T2 need not be acquired at the exact time as EGT.
- IGV should be acquired and recorded for each EGT. The prognostic reasoner should use IGV to remove variations and offsets in EGT that are attributable to IGV. IGV correction is especially important near the end of the life of the APU when the controller modulates IGV to keep EGT within limits. The IGV correction can reach 80 F for this APU.
- IGV must be acquired at the same time as the EGT, because the IGV setting can change rapidly. Otherwise, IGV modulation could lead to misleading results.
- Electrical generator load L should be acquired and recorded for each EGT. The prognostic reasoner should use L to remove variations and offsets in EGT that are attributable to L. The L correction can reach 90 F for this APU.
- Ambient pressure P2 may be acquired and recorded for each EGT. The prognostic reasoner should use P2 to remove variations and offsets in EGT that are attributable to P2. The P2 correction is relatively small for this APU.
- Intermittent gaps in the data may be tolerable, even spanning a few months, especially during the flat part of the trend.

7.3.2 Maintenance Reasoning: Fuel Nozzle Remove and Replace Action

7.3.2.1 Fuel Nozzle Background

Dirty or clogged fuel nozzles are a causal contributor to hot-section deterioration because of the uneven combustion flame they produce. Fuel nozzle streaking or torching can deteriorate the first stage turbine nozzle vanes over time (the first turbine nozzle vanes are stator vanes at the start of the hot section before the turbine wheel). Deteriorated first stage turbine nozzle vanes reduce the efficiency of the turbine, forcing the controller to increase FF to obtain the same output power. This produces higher temperatures in the hot section, which causes deterioration throughout the hot section over time. The deteriorated hot section has reduced efficiency, necessitating increased FF, producing higher hot-section temperatures, and causing even further deterioration.

The elevated hot-section temperatures are detectable through elevated EGTs and reduced EGT_{margin}. However, dirty or clogged fuel nozzles do not directly impact EGT_{margin} (or EGT). Instead, they impact EGT_{margin} (and EGT) indirectly through the accumulated deterioration they cause. Notionally, if an experiment can be performed by alternating between clean and dirty nozzles and measuring the EGT_{margin}, the EGT_{margin} would not be expected to jump up and down between clean and dirty nozzles. Instead, the rate of change of the EGT_{margin} would be expected to flatten and steepen with clean and dirty nozzles, respectively.

A remove and replace action of dirty or clogged fuel nozzles can correct one of the causal contributors to further hot-section deterioration. If the old fuel nozzles were dirty or clogged, then the action can slow the rate of ongoing deterioration and thereby flatten what was a downward trend in EGT_{margin}. There are a few caveats:

- The benefits of a fuel nozzle remove and replace action can only be obtained when the old nozzles were actually dirty/clogged. There is currently no BIT to detect dirty/clogged nozzles before removal.
- Even if the old nozzles were dirty/clogged, the action cannot undo hot-section deterioration that already exists. A fuel nozzle remove and replace action can flatten a downward trend in EGT_{margin}, but EGT_{margin} is not expected to suddenly improve.
- Any existing deterioration to the hot section (including first stage nozzle vanes) is a cause for further deterioration. If the contribution from existing deterioration is already large, then there is less benefit to changing the fuel nozzles. More benefit is expected when the APU has few hours on-wing (for example, <4000 hours' time since repair) and EGT_{margin} is still high.

Maintainers perform fuel nozzle remove and replace field actions on a scheduled basis to help prolong the on-wing time of the APU. For example, fuel nozzles may be changed at an on-wing time of 2000 hours.

Maintainers can also perform a fuel nozzle remove and replace action on a condition-basis in an attempt to flatten a downward trend in EGT_{margin}.

7.4 VALVES SPMS

7.4.1 Basic Principles

Figure 33 shows the typical components of a spring-loaded regulation valve such as those used for engine anti-ice and bleed applications Spring-loaded regulators operate strictly on force balance. If closing forces exceed opening forces, the poppet is against the seat, and the excess forces are reacted against the seat. If opening forces exceed closing forces, the poppet is against the open stop, which reacts to the excess forces. If closing forces and opening forces are equal, the poppet/diaphragm assembly is in static balance (positioned somewhere in mid-stroke).

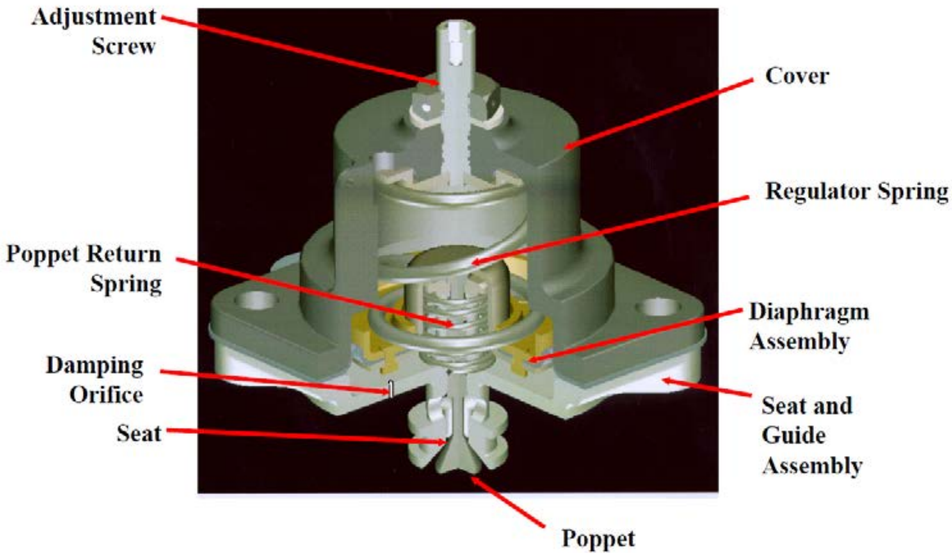

Figure 33. Spring-loaded poppet regulation valve

In an inline valve, the downstream pressure is connected under the diaphragm in such a way that the valve automatically moves to whatever position results in the pressure being at the valve set point. That is, it uses the "pressure tap" on the flow line to provide the regulation. This last statement is important for designing an SPMS for regulation valves. Measurement of the valve position is inconsequential for health monitoring; by design, the valve position "floats" to meet the regulatory demands, signifying a healthy valve. This study used a simulation-based approach to evaluate SPMS options for monitoring an inline regulation valve. The setup of this simulation is described in section 7.4.2.

7.4.2 Simulation Setup

The simulation setup was implemented using Matlab/Simulink.

7.4.2.1 Fleet/Operator Definition

This study's fleet consisted of three categories of airline operators. These categories were characterized as:

- Airplanes operating in a dirty environment.
- Short haul airplanes with a quick turnaround between flights.
- Typical long haul airplanes.

Within each category, an operator can have several airplanes, which results in several regulating valves.

7.4.2.2 Usage Definition (Aircraft Flight Profiles)

This specifies the flight patterns for the airplane, which results in the valve experiencing varying duty cycles. The study included the following flight profiles:

- Short cold day flight–icing at lower altitudes.
- Long cold day flight–icing at lower altitudes.
- Short standard flight and day–no icing.
- Long standard flight and day–no icing.
- Short cold day flight with hold–icing at lower altitudes.
- Long cold day flight with hold–icing at lower altitudes.
- Short warm day flight with icing on descent.
- Medium length warm day flight with hold and icing on descent.
- Warm day flight with no icing.
- Warm day flight with icing on descent.

Valve wear was characterized using the following parameters:

- Valve set-point calibration.
- Valve hysteresis: This is the difference between an upward and a downward movement of the valve.
- The valve drifts over life. A drift can happen either in the calibration/hysteresis, both of which lead to loss of regulation. This parameter specifies a temporal pattern for this change.
- The HI generation module: This module allowed the authors to evaluate the SPMS interfaces—specifically, the following three SPMS interfaces:

  – Sensor interface: Options include upstream pressure, downstream pressure, and manifold temperature.
  – Data interface: Options include sampling frequency. This interface was not evaluated in this experiment. The sampling frequency was fixed at 1 Hz.
  – Monitor interface: Options include retaining 1) average, 2) maximum manifold pressure, and 3) minimum manifold pressure.

At the heart of the simulation was the valve model. Implemented in MATLAB/Simulink®, this describes a quasi-steady state response of the valve to various inlet pressures. The response depends on various valve parameters such as hysteresis and calibration errors.

7.4.3 Data Generation

Data are generated from the simulation setup using a Monte Carlo experiment, as follows:

- Each experiment is initiated by selecting a valve, which picks the wear profile randomly and sets the calibration error, hysteresis, and its evolution over the life of the valve.Each trial of the Monte Carlo experiment consists of picking an operator (short haul, long haul, dirty environment) and a usage profile (altitude, cold day, hot day).
- The experiment is stopped when the valve is unable to regulate the downstream pressure, signaling a "hard fault" and a fault message reported to the onboard maintenance system. The experiment can also be stopped after a predefined number of calendar days have elapsed.

Figure 34 shows the time trace of the outputs generated from one such experiment. The experiment shows the data after approximately 210 calendar days. The inlet pressure, P_{in}, downstream pressure, P_{dwn}, and inlet temperature, T_{in}, are shown in three subplots. In this case, the valve is installed on a short haul aircraft and designed to regulate the pressure at approximately 30psi ± 3 psi. That is, $P_{reg} = 30$ psi. As shown by the *x*-axis, the valve accumulated approximately 130 hours of operations within the 210 days of operations. The valve usage to flight usage ratio is 0.6(≈ 130/210).

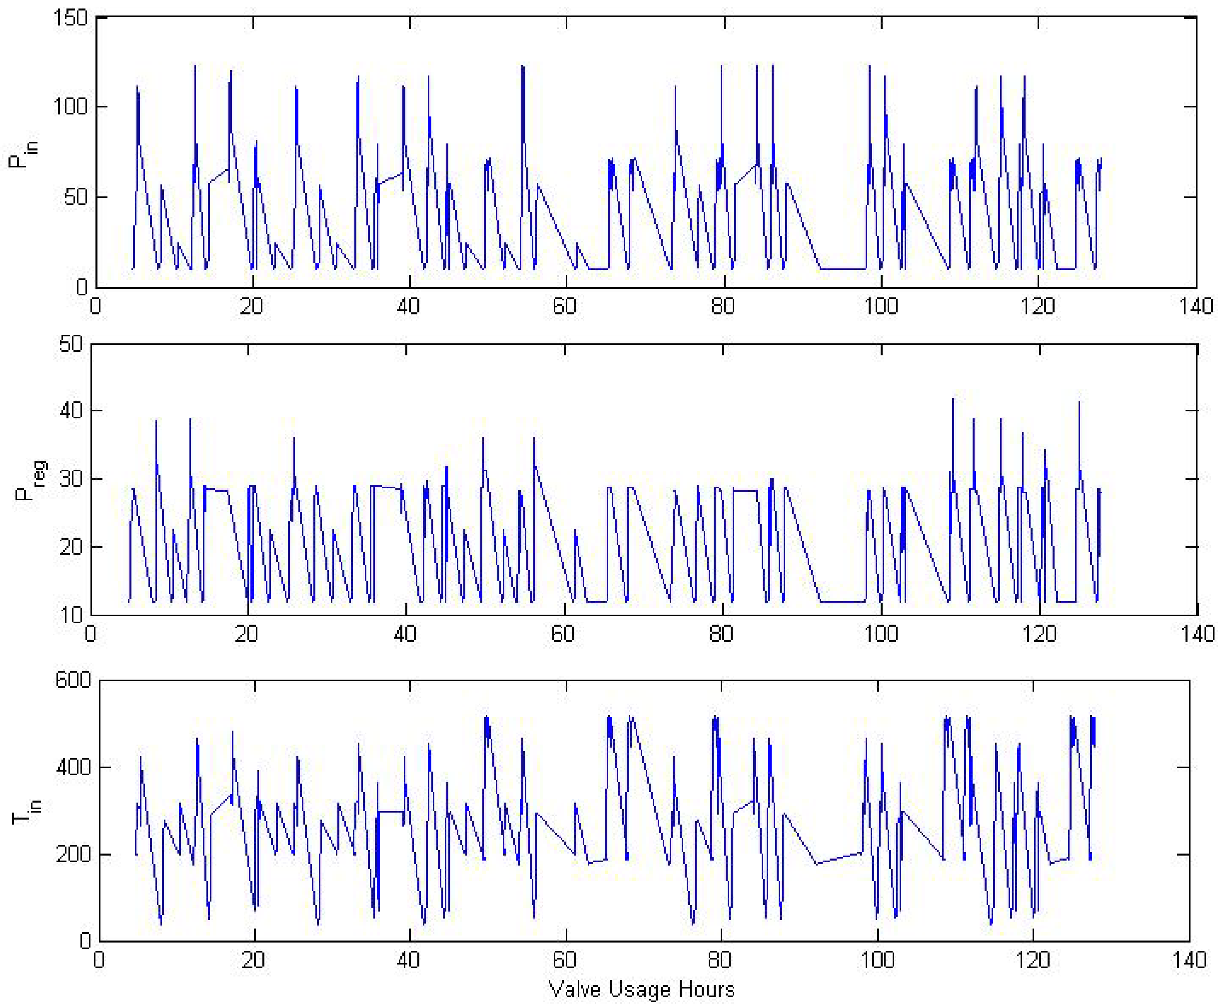

Figure 34. Time trace data generated from the valve simulation setup

7.4.4 SPMS Interface Evaluations

7.4.4.1 Sensor Interface

Measurements available from the simulator for each experiment are listed below:

- Valve serial number: This is a constant. Because a family of valves is simulated, this provides a reference to track data from each valve.
- Valve on/off command: Binary 0/1 value indicating if the valve was activated.
- Valve downstream pressure, P_{dwn}, in psig.
- Valve inlet pressure, P_{in}, in psig.
- Valve inlet temperature, T_{in}, in F.
- Aircraft altitude, F_{alt}, in ft.
- Ambient pressure, P_{amb}, in psia.
- Ambient temperature, T_{amb}, in F.

- Cumulative flight time, t_{flt}, in hours.
- Valve cumulative operating time, t_{vlv}, in hours. The cumulative time when the valve was commanded to open and actively regulate the pressure. By definition, $t_{vlv} \leq t_{flt}$.

Because the ambient temperature characterizes the flight altitude, this measurement was not included in the SPMS sensor interface. Furthermore, valve cumulative flight time can be readily derived from valve on/off times and the flight times. This assumes that the valve installation did not change, which is an unreasonable assumption to make. Therefore, this sensor was excluded.

The SPMS sensor interface for inline regulation valves is summarized in table 50.

Table 50. SPMS sensor interface for inline regulation valves

SPMS sensor interface
Valve on/off command: Binary 0/1 value indicating if the valve was activated. Valve downstream pressure, P_{dwn}, in psig. Valve inlet pressure, P_{in}, in psig. Valve inlet temperature, T_{in}, in F. Ambient temperature, T_{amb}, in F.

7.4.4.2 Data Interface

The bottom subplot of figure 35 shows the transient response of the valve when the inlet pressure (see top subplot) rises from 15 psig to approximately 70 psig. High variability in operating conditions and the associated valve response makes a transient-based approach highly difficult to consistently capture and analyze. Therefore, a robust data interface for SPMS must reduce the monitoring window by applying some form of data filtering and retain only the relevant part of the sensor measurements. The SPMS data interface specifies this for the inline regulation valves.

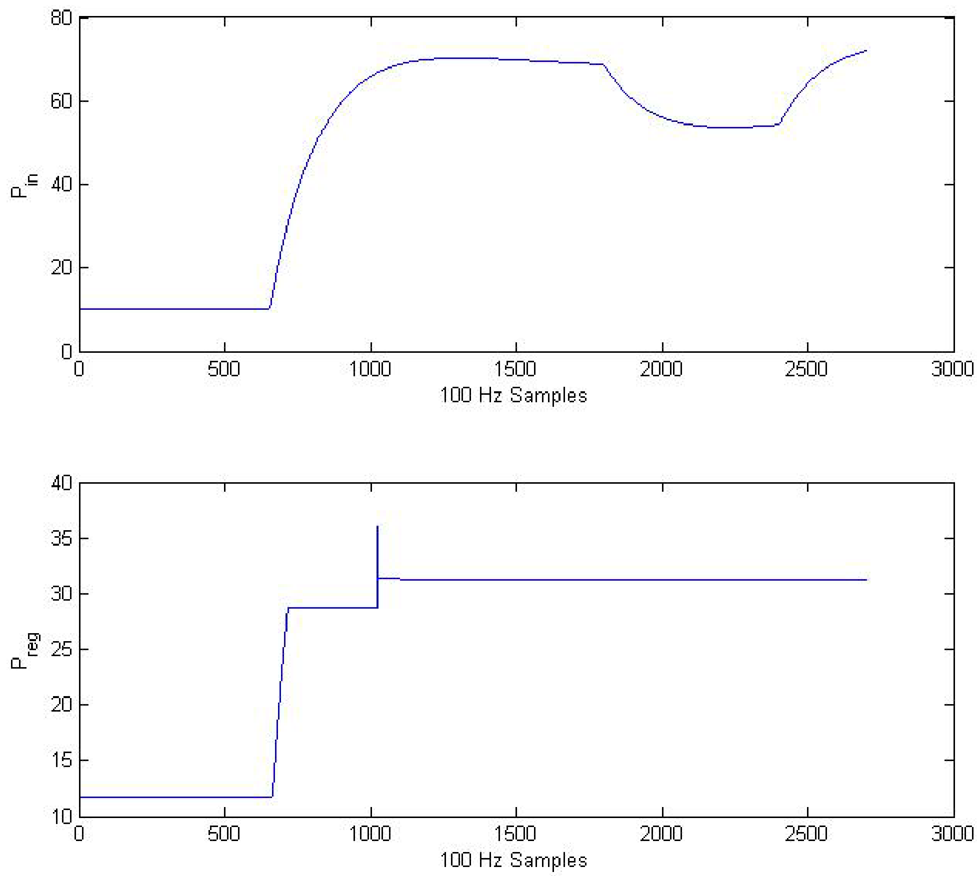

Figure 35. Transient response of the valve

Though a valve operates intermittently over the course of each flight, discussions with subject matter experts indicate the amount of wear is greater at pressures around its calibration point; the valve is supposed to hover around this point gently but often oscillates, causing wear on the seals and valve seat. Therefore, the SPMS data interface is defined to include regulation pressure, P_{dwn}, samples taken at a relatively low frequency. The SPMS data interface is summarized below:

- Capture trigger logic:
  - Valve is commanded to be ON.
  - If the downstream pressure is not indicating a "stuck closed fault."
  - Inlet pressure is above the regulator set point, $P_{in} > P_{reg}$.
- Delta pressure between the valve inlet and downstream is greater than a certain amount.
- That is, $P_{in} - P_{dwn} \geq \theta$. Typically, θ = 10% of P_{reg}.
- Data frame.
- Variables: P_{in}, P_{dwn}, T_{in} at 1-minute intervals.

7.4.4.3 Monitor Interface

The SPMS data interface as described in section 7.4.4.2 provides a series of P_{in}, P_{dwn}, T_{in} sensor values. In any given flight, if the valve is open for more than n minutes, the data interface can provide, at most, n such data frames. The SPMS monitor interface choices evaluated are:

- Maximum downstream pressure. That is, $\max(P_{dwn})$.
- Minimum delta pressure across the vale. That is, $\min(P_{in} - P_{dwn})$.

The traces of the CIs generated are shown in figure 36.

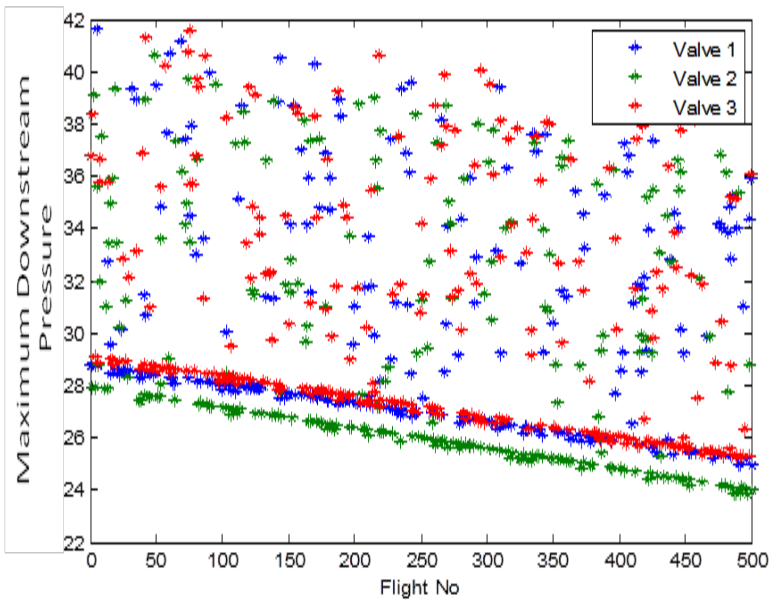

Figure 36. Monitor interface choices for generating CIs

It is important to note that despite the high variability in the data, both $\max(P_{dwn})$ and $\min(P_{in} - P_{dwn})$ can be trended and changes in the slope indicate a degrading valve or gradual loss of the regulation function, which would eventually warrant a removal action.

7.4.4.4 Health Interface

The health interface provides a yellow or red indication for removing the valve because of internal wear and eventual loss of the regulation function. The indicator is predictive in nature and, if no maintenance is performed, the valve issues a fault message to the onboard maintenance system. The yellow and red indicator threshold is critical. If max(P_{dwn}) is chosen, then the health interface needs to establish a "lower bound," below which an indicator is generated. Conversely, if min(P_{in} – P_{dwn}) is chosen, the health interface needs to establish an "upper bound," above which an indicator is generated:

- Lower bound on max(P_{dwn}): Trigger a notification when max(P_{dwn}) is less than this threshold.
- Upper bound on min(P_{in} – P_{dwn}): Trigger a notification when max(P_{dwn}) is greater than this threshold.

A complete evaluation of the above-listed health interfaces was not done because of difficulty in simulating the disturbance conditions. The evaluation needs to establish the upper and lower bounds using a large valve population. Furthermore, the authors believe a combination of the two options can provide a more robust solution.

8. TASK 6: TEST RESULTS

8.1 ENGINE SPMS TEST RESULTS

SPMS targets the engine's fuel control system to provide an actionable indicator (see table 51) of an impending FCU failure that could cause a potential engine no-start or in-flight engine shutdown. The FCU is relatively easy to swap out and does not need any special qualification tests after it is replaced.

Table 51. Engine tests

SPMS application	Problem/target system	Potential impact of SPMS
Actionable Indicator	Probability of predicting an impending FCU problem within the propulsion engine.	Five cycles before an uncommanded engine shutdown.

8.1.1 Data

Beginning in 2002, Honeywell installed an elementary data collection system on 35 identical regional jets. Because no specific SPMS application was selected, the system was set up to collect 182 parameters from various sensors installed on the four engines, APU, flight management system, navigation system, bleed system, and landing system. The list of 182 parameters was set based on physical memory locations imposed by the existing FDAMS. Data collection was maximized to fill all 256 words that the ARINC 717 encoding would allow.

Actual maintenance procedures performed and parts replaced and repaired are documented and archived in relational databases maintained at three primary repair locations. Field service engineer observations are recorded as freeform text entries and archived together with repair records. Because the SPMS is targeted at the engine FCU, all of the available data cannot be used; for example, data surrounding an engine bird-strike event is not within this study's defined SPMS scope.

Table 46 lists a subset of field service observations that may be caused by a malfunctioning FCU. To evaluate the SPMS application, the shop finding needs to be known. For example, consider the event on March 9, 2003 for Tail #25. To link this field observation with a malfunctioning FCU, this field observation needs to be correlated with the repair finding. Table 46 does not list this information because the authors only identified the source when they could get this information, and they do not have it at this time. Through analysis of the airline data it was observed that the shop findings have not been consistently recorded. For testing the engine SPMS, the cases in table 52 were used.

Table 52. Historical data available for engine FCU-targeted SPMS interface testing

AC tail #/engine	Field observation	Shop finding
27	Engine starting slow	FCU replaced
33	Engine starting hot	FCU replaced
29	Engine starting hot (over-temperature)	FCU replaced fuel leak
19	Engine starting hot, over-speed temperature, and engine shutdown	FCU replaced
Total = 4 selected from total of 19 cases		

Table 53 summarizes the engine FCU-targeted SPMS interfaces that were tested.

Table 53. Engine FCU-targeted SPMS test plan

SPMS interfaces	Test plan	Outcome
Sensor interface	Pick-and-choose sensors from other engines on a multi-engine aircraft	Selected list of sensors
Data interface	Sampling error analysis	No significant effect
Monitor interface	Evaluate options for time series analysis	Evaluated the monitors, included new startup monitors (CIs)
Health interface	Develop and test options	Studied the relationship of CI patterns to the HI conclusion
Action interface	N/A	N/A

8.1.2 Sensor Interface

The test plan for the sensor interface involves evaluating whether sensors from engine A can help the SPMS for engine B.

Table 54 shows a signal-to-FCU fault correlation evaluation summary condensed back to the raw signal measurements. Based on this analysis, the startup CIs were selected for further development.

Table 54. Engine SPMS signals used

Signal Name	Description	Used on engine SPMS	Relationship with FCU faults and comments
EGT	Engine EGT for all four engines	✓	Correlated
TAT	Total air temperature (ambient)	✓	Used for correction
N2	Engine core speed (rpm) for all engines	✓	Correlated
N1	Fan speed	x	Not correlated
FF	FF (lb/hour) for all engines	x	Does not measure FF
PLA	Power lever angle or engine thrust setting	x	Not correlated
PALT	Pressure altitude	✓	Used to determine flight vs. ground runs
PH	Flight phase from ACMS	x	Redundant info

8.1.3 Data Interface

To calculate the necessary CIs, the authors specified that the data interface would provide time-series data for all of the sensors. The authors also described the rationale for collecting this data during the entire startup period (see table 55). The test plan for this data interface is focused on evaluating this choice; that is, which of the data collected during engine startup or acceleration is best for the FCU-targeted SPMS application. The tests set with multiple sampling rates were carried out. It should be noted that the EGT and N2 measure used in the CI computation is filtered; therefore, down sampling by 4 does not impact the CI. Moreover, for modern aircraft engines, the down sampling by 4 is not expected to have a significant impact on memory requirements for recorders and engine controllers.

Table 55. Data interface options to be tested for the engine-FCU targeted SPMS application

Flight phase	Sampling frequency	SPMS test summary
Study the impact on detection metrics when data is collected during engine startup.	Study the sensitivity of sampling at 1, 2, and 4 Hz data for 60 seconds.	Summarized with the health interfaces.
Study the impact on detection metrics when data is collected during engine acceleration.	Study the sensitivity of sampling at 1, 2, and 4 Hz data for 30 seconds.	Not tested separately. Indirectly tested as part of test above.

8.1.4 Monitor Interface

The focus of the SPMS application is FCUs. The approach, as described in this report, is to analyze relative time-series data collected during predefined phases of engine operations. The pattern of given time-series data is compared against a baseline pattern. The monitor measures the dissimilarity between the baseline pattern and the given pattern. There are two sets of CIs that have been tested. Figure 23 shows the N2 and EGT profile for 230 consecutive flights from the same engine. The blue lines show 30 flights post-FCU replacement. The red lines show the profile 30 flights before FCU replacement. Cyan lines mark the flights 200 to 30 before the FCU replacement. The engine no-starts and startups corresponding to ground runs are not plotted. The startups are lined when N2 exceeds zero. Figure 23(a) shows the N2 at zero due to filtering. The fuel in the turbine is introduced after it has reached 25% of the idle speed. Because the EGT profiles do not line up, the new CIs compute the relative distance between two points on the profile (e.g., time4EGTa2b and timeMinN2P2L).

Next, the analysis of the entire engine CI index from the FCU removal event was computed. The results showing some of the promising indicators are shown in figure 23. The results shown are for 230 flights before and after the incident. Figure 26 shows the analysis for 1000 flights. Column 2 of figure 26 shows the histogram bin of the filtered CI. For this analysis, a 10-sample moving average filter was used. It can be clearly seen that the individual CIs are not capable of predicting the FCU removal. This exercise is made difficult due to incomplete records pertaining to repair actions (figures 26–29).

In summary, it is important to point out two limitations that were encountered:

1. Because of incomplete records of the repair actions, it is hard to identify the false alarm rates, especially in the case in which the CI patterns match a known fault pattern but do not have any corresponding repair actions.
2. The interaction between the fuel controllers with the underlying FC produces a non-trendable CI. This happens because the fuel controller continues to compensate for the fault in the actuator until it can no longer produce sufficient actuation, leading to eventual failure. Because of the compensation by the controller, it is difficult to observe the underlying failure mode growth without analyzing the controller commands and responses.

Both of these limitations make it very difficult to develop a robust approach to prognostics development. Therefore, identifying and separating the faulty assets from the health asset were the focus of the authors' efforts. A cluster analysis of all CIs was the start of the authors' renewed focus.

Figure 37 shows the engine start sequence from four engines. The top subplot shows the peak EGT (corrected) during the engine start. The bottom subplot shows the corresponding start time, idle speed, engine speed (N2) at peak EGT, light off, and time for EGT to peak. The data set includes the fast start, slow start, over temp, and health starts sequenced together. Samples from this data set have been divided into training and testing sets for the purpose of developing clustering and inferencing systems.

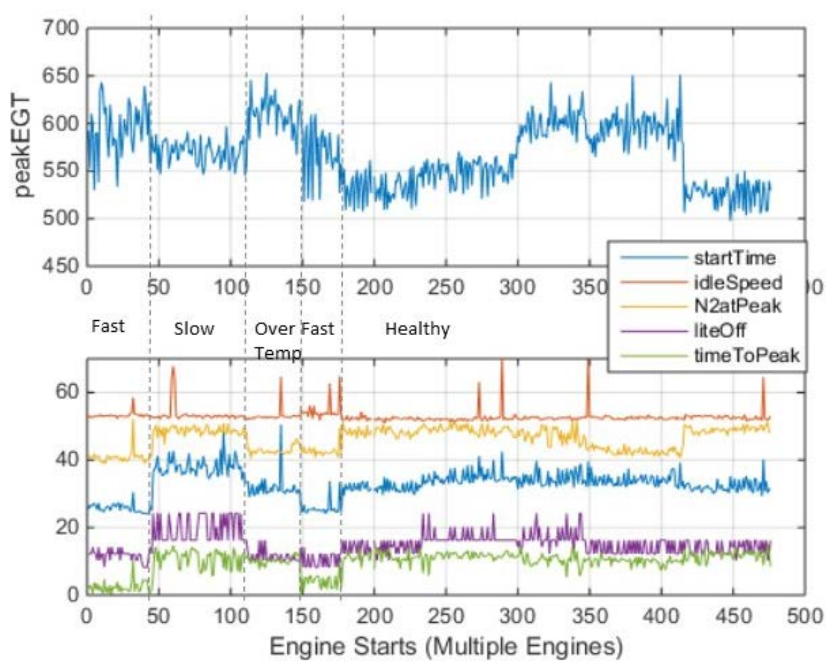

Figure 37. CI plotted sequence per engine start

To better understand the relationship between the CIs and their separability, multiple clustering approaches were explored, such as subtractive fuzzy clustering and fuzzy k-means clustering. Both of these approaches yield a clearly separable class for the health and fast start data sets. However, the other fault clusters were only marginally separable. Figures 38–40 show the scatter plot of CIs vs. the start time. The scatter plots include a continuous distribution plot next to the x-axis and y-axis that shows the distribution of the respective CIs. The scatter plots, together with the continuous histograms, give a clear picture of which of the CIs can be used to detect the underlying failure mode associated with the FCU faults. The analysis of the CIs shows that:

- Fast starts can be detected using the start times (see figure 38). Figure 38 shows that the histogram of start time on the x-axis is clearly separable (to the left of) histograms for all other health classes, namely healthy, fast, slow, and over temperature.
- A slow start can be detected by tracking the light-off time (see figure 39) in the y-axis. The histogram mean for light-off is higher compared to other health classes.
- Over temperature can be detected by looking at peak temperatures (see figure 40). In this case, the peak EGT is a differentiator. It can be seen that this is not a very strong indicator as the “over temp” and “slow” start cases do not have much of a separation (y-axis).

It is also clear that none of these clusters has clear linearly separable boundaries; therefore, the use of fuzzy and clustering approaches is justified.

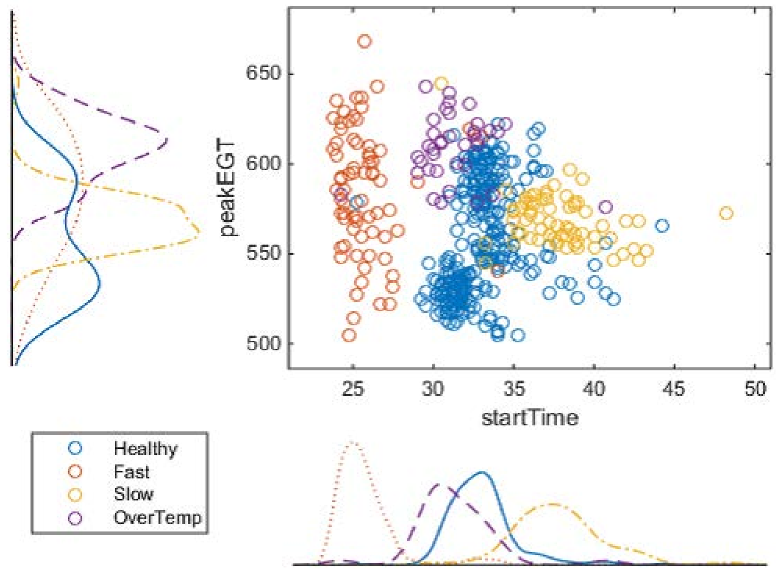

Figure 38. Scatter plot of peak EGT vs. idle speed (training data from engine position 3, 1)

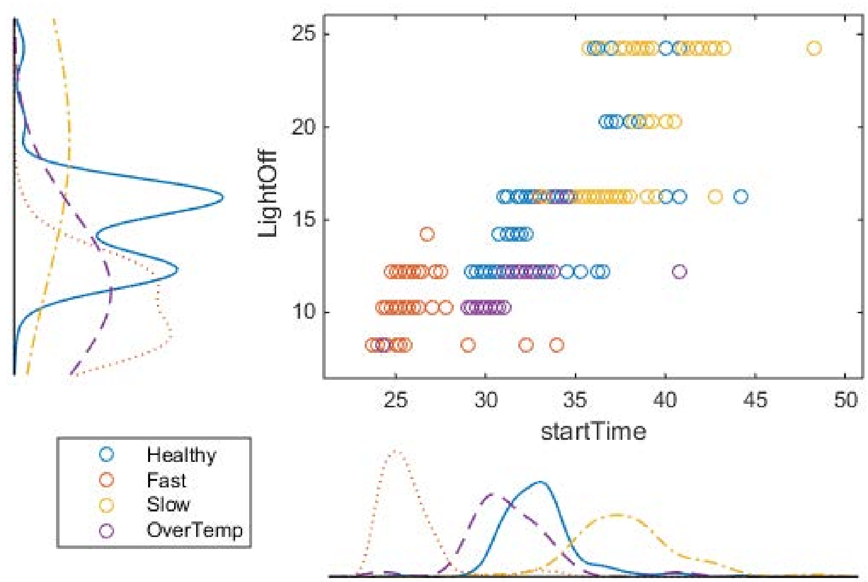

Figure 39. Scatter plot of light-off vs. idle speed (training data from engine position 3, 1)

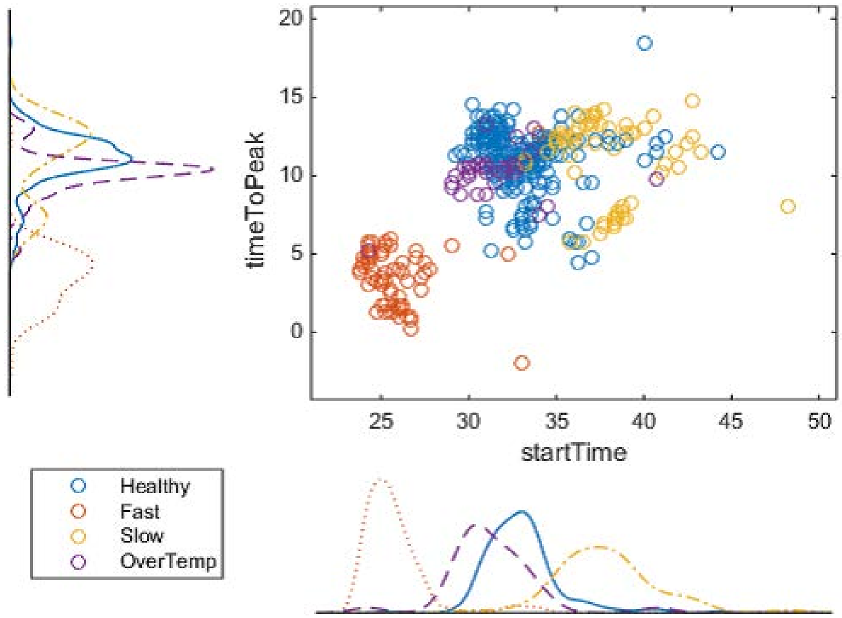

(a)

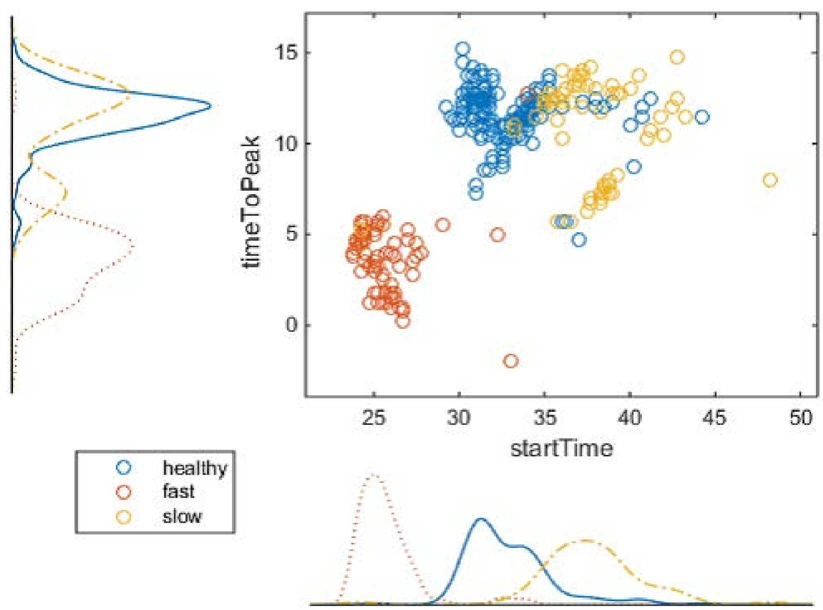

(b)

Figure 40. Scatter plot of time for the EGT to peak during engine startup sequence vs. idle speed: (a) training data from engine position 3, 1; and (b) training data from engine 3 only show more separation

After reviewing these results with the experts (see table 56), it was decided to use engine position to improve the classifiers. Because two of the three cases were for engine position 3, it was decided to train a new classifier using the engine positionspecific data. This is more practical because engine start sequence impacts whether they are started using APU bleed or engine bleed and, therefore, have impact on start times. Figure 40(b) shows the scatter plot for the training set for engine position 3.

Table 56. Monitor interface options to be tested for the engine-FCU targeted SPMS application

SPMS interfaces	Test results	Outcome
Monitor interface	Options for time-series data analysis showed that the CIs were not trendableand fault cases were limited. Therefore, Markovian approaches were not tried. All CI metrics were computed using Euclidean distance measures.	100% coverage for four cases of faults and healthy turbines.

The startup CIs are in two groups: the startup CIs developed before the start of the program and the startup CIs implemented during this program (listed as “New Startup CIs;” see table 57). Additional/new startup CIs are based on first-order differential components for the N2 and EGT. After analyzing the clusters together and in light of the limited number of cases, the authors decided to develop HI using a limited set of CIs, as shown in table 57.

Table 57. Engine startup CIs

Label	CI Name	CI Group	Description	Used for HI
x1	startTime	Startup CI	Time it took the turbine to start up from roll-up to reaching idle speed	✓
x2	idleSpeed	Startup CI	N2 speed after engine start when engine has attained stoichiometry	✓
x3	peakEGTC	Startup CI	Peak EGT at engine startup	✓
x4	N2atPeak	Startup CI	Engine speed at peak EGT	✓
x5	liteOff	Startup CI	Time the ignition is turned off and now the fire burning in the hot section supplied by air from the compressor	✓
x6	timeToPeak	Startup CI	Time to get to peak EGT	✓
x7	timeMaxEgtP2L	New Startup CI	Time from start of profile to reach max dEGT/dt	✗
x8	maxEgtP2L	New Startup CI	Magnitude of max dEGT/dt	✗
x9	timeMinEgtP2L	New Startup CI	Time to get to min dEGT/dt	✗
x10	delTimeEgtPC	New Startup CI	Time between first peak and crest of dEGT/dt	✗
x11	deltaEgtPC	New Startup CI	Magnitude between peak and crest of dEGT/dt	✗
x12	timeMaxN2P2L	New Startup CI	Time for max dN2/dt in the window of max to min dEGT/dt	✗
x13	maxN2P2L	New Startup CI	Max dN2/dt in the window of max to min dEGT/dt	✗
x14	timeMinN2P2L	New Startup CI	Time for minimum dN2/dt in the window of peak to trough of dEGT/dt	✗
x15	minN2P2L	New Startup CI	Minimum N2 peak to peak to trough of dEGT/dt	✗
x16	time4EGTa2b	New Startup CI	Time for EGT to rise from A to B above ambient temperature (A + 300 degrees). This is a pseudo indicator for rate of EGT rise.	✗

8.1.5 Health Interface

Given the "deviation" from baseline pattern monitors, the health interface provides a ranked list of probable root causes that represent the internal health state of the engines. The intent is to help the ground maintainer and engine OEM to get ready and avoid any operational disruption. The outputs provided by the SPMS monitor interface provide evidence for several engine components.

For example, a deviation in the engine speed v/s time curve before auto-combustion can be due to a starter problem, low bleed from the APU, and/or increased resistance from the engine turbine bearings. The last problem could in turn be caused by low lube temperature and increased oil viscosity. Deviations during the light-off phase could be due to igniters or incorrect fuel atomization. Deviations at engine idle could be caused by incorrect fuel metering or loss of engine performance due to erosion of the rotating components, such as the turbine blades.

The health interface provides a probability (a number between 0 and 1) associated with the following four causes.

To develop the HI from the CIs discussed in the previous section, the fuzzy inference system and a classification tree method of supervised learning were used. Figure 41 shows the fuzzy rules learned by the inferencing system. The rules represent a 6×5 matrix (six inputs by five fuzzy clusters per input) that can be used to fully describe the training set outputs.

Figure 41. Engine position-specific fuzzy inference rules for engine position 3

Figure 42 shows the full classification tree learned using the same training set. The classification tree shown in figure 42 detected FCU health associated with effects (0 = healthy, 1 = fast start, 2 = slow start).

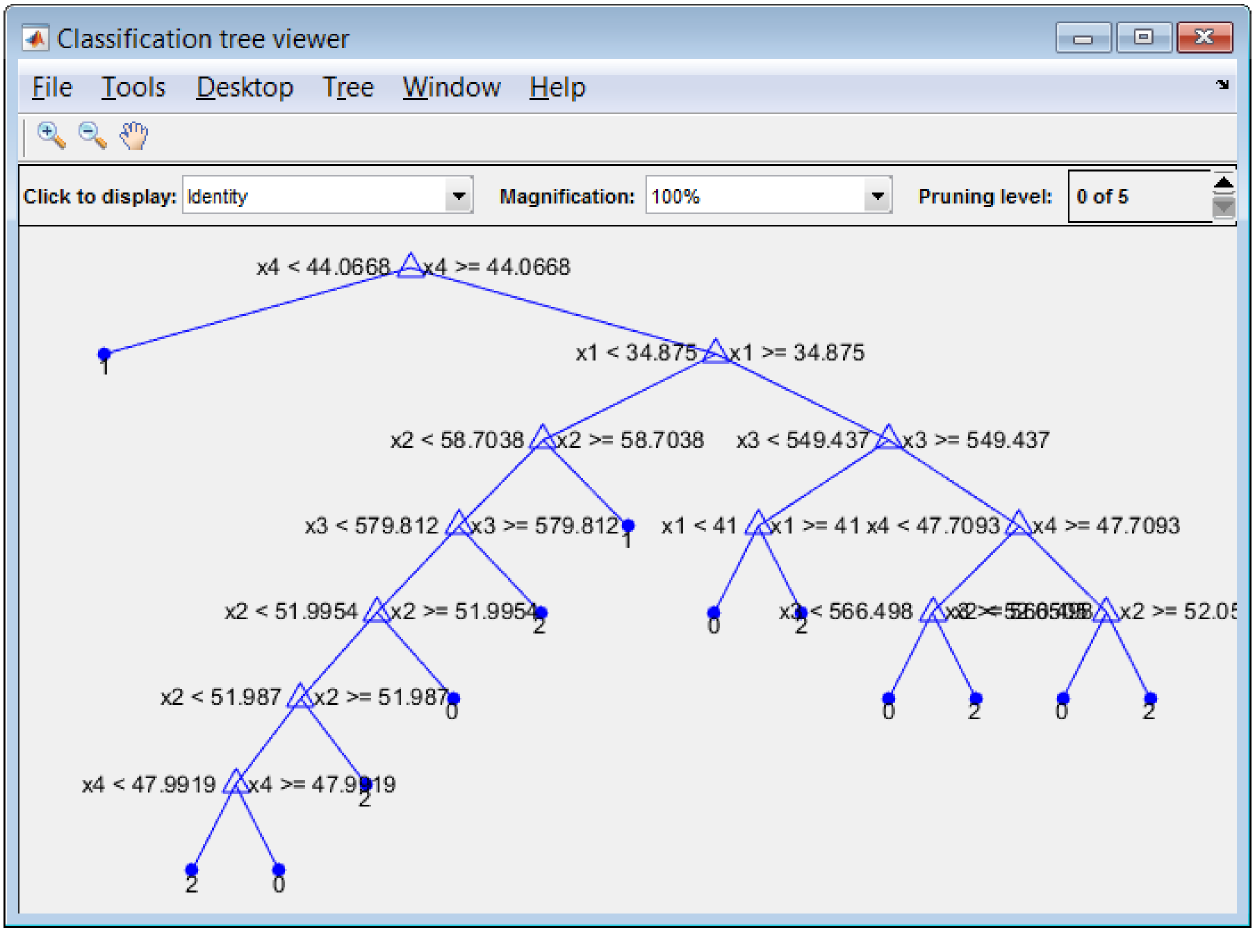

Figure 42. Classification tree for engine-specific classifier (engine position 3) to detect FCU health associated with effects (0 = healthy, 1 = fast start, 2 = slow start)

The developed HIs using fuzzy clustering and the classification tree are very sensitive to input errors. Therefore, it is essential to perform simple range checks using the guidance from the scatter histograms to filter starts. In addition, because the CIs are not clearly separable, the de-bounce approach using a finite impulse response filter and median filter are used on the generated HIs to validate the approach. Figure 43 shows the validation run from the four data sets discussed earlier. It can be clearly seen that the classification tree yields better results. Figure 43(b) shows the validation run for engine position 3. The results show a clear improvement when the inferencing (fuzzy and classification) is performed on an engine-position basis.

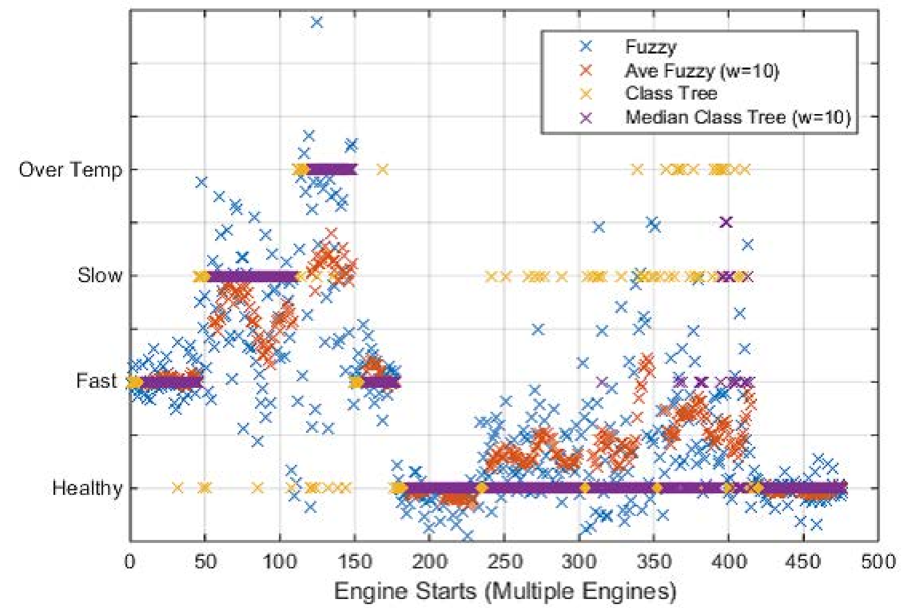

(a)

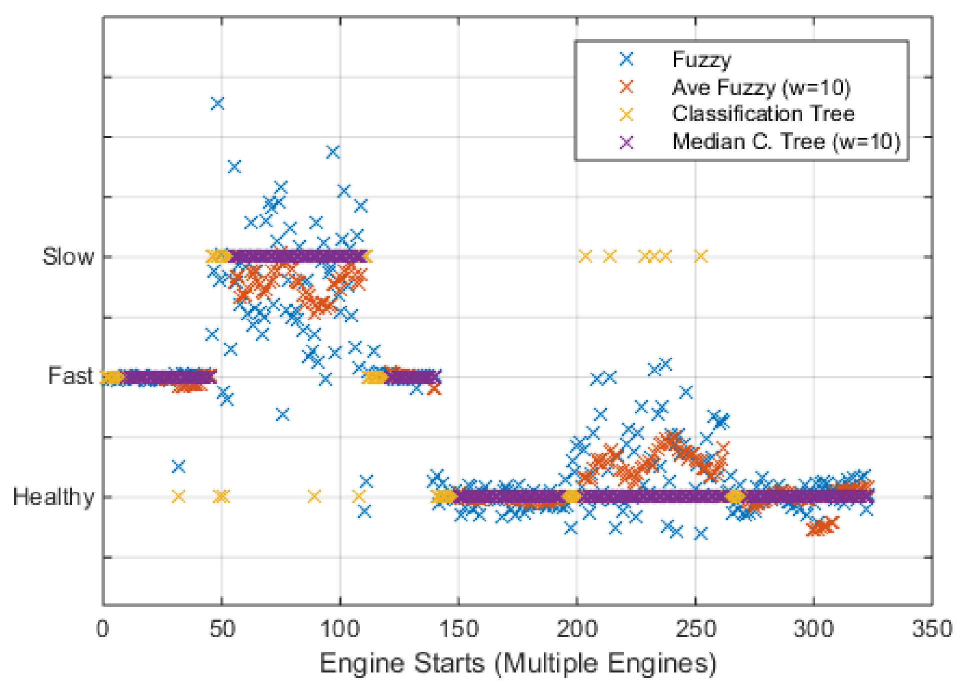

(b)

Figure 43. Validation results using fuzzy inference system and classification tree: (a) validation from engine positions 1, 3, and 4; and (b) validation from engine 3 only shows less false alarms ((b) also includes a squashing filter to eliminate bouncing between two classes)

Note that, because of the limited number of test cases, both of these approaches can yield overfitted inference models that cannot be fully tested using this data set. Figure 44 shows the HI for 26,000 engine 3 starts. It can be clearly seen that there are many cases that correspond to fast and slow health states. It can also be seen that sometimes these conditions persist for up to 4000 engine starts (see figure 44). The subplot on figure 45 shows the zoomed up HI bouncing between healthy and fast start for 4000 engine starts. Figure 45 shows the corresponding CIs which are consistent with the CIs for fast start (see figure 37). Figure 46 shows the scatter plot for CI light-off versus the engine start time for the same data set. It can be clearly seen that slow starts and healthy starts are more separable than fast starts and healthy starts, as indicated by overlap on both the x-axis and y-axis histograms. Cluster labeling on the validation is achieved using the trained classifier with median-floor filter. Based on these spot checks, the following conclusions can be drawn:

- HIs for the engine FCU are working well.
- HIs can distinguish between fast and slow starts.
- HIs work better when developed for respective engine position.
- The HIs are predictive (n-class prediction: healthy, fast, slow, and over temp start) but because of the limited number of cases and the nature of the CIs, the HIs do not provide prognostic (time to failure) information.
- Finally, the engine SPMS (as designed and tested) can only be used as a maintainer aid to guide them if a failure mode is present/absent in the engine FCU. Therefore, the engine SPMS does not have any certification impact.

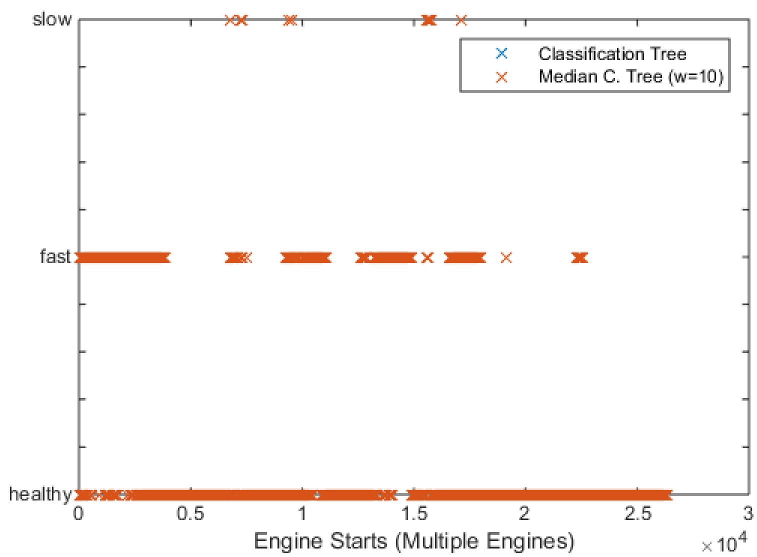

Figure 44. HI run on 26,000 engine 3 starts from five tail numbers

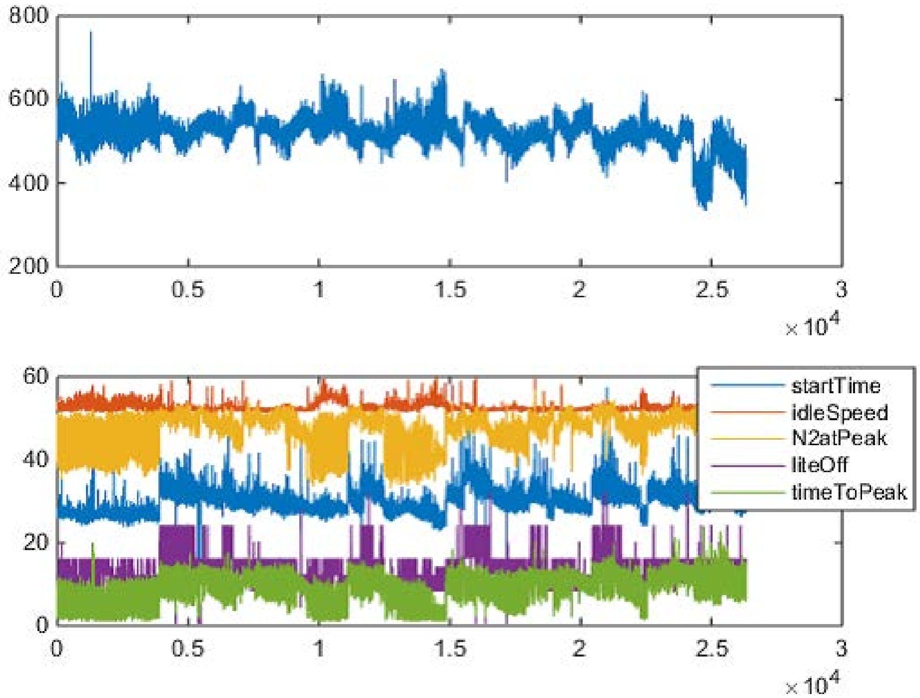

Figure 45. CI corresponding to 0–4000 startups (HI shown in figure 13)

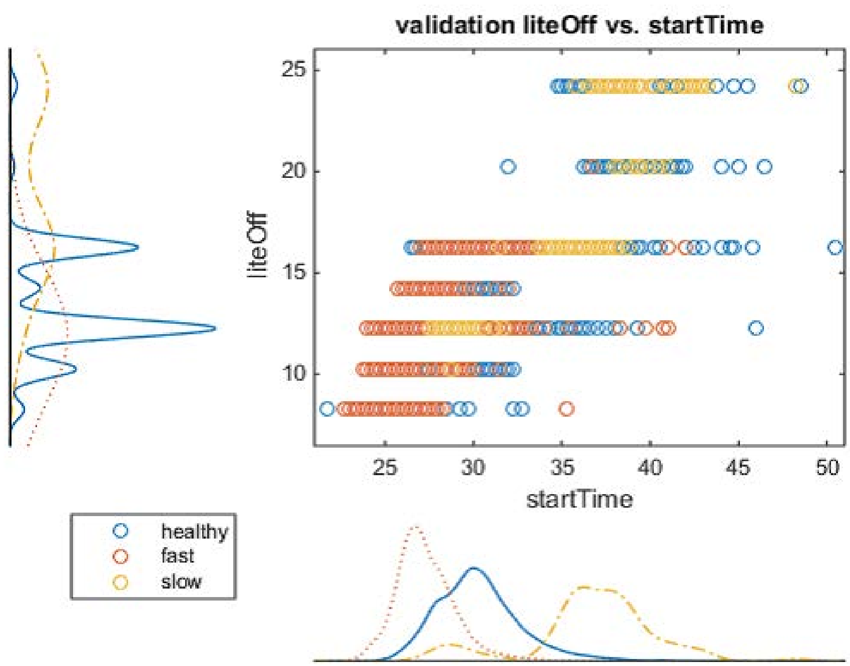

Figure 46. CI scatter light-off vs. start time for engine position 3 (the histograms show considerable overlap between healthy and fast starts)

The test results for this are summarized, for a limited number of labeled test cases, in table 58.

Table 58. Test results for engine health interfaces

SPMS interfaces	Test results	Outcome (for classification tree with median filtering)
Health interface	Evaluate options for monitor combination rules	Detection accuracy = 100% for the four cases for all engines (100% for three cases for engine position 3) Fault isolation = 100% for the four cases False alarm rate < 10%

8.2 APU SPMS TEST RESULTS

8.2.1 APU monitor interfaces

The test plan for this interface is summarized in table 59. The test plan addresses two questions: 1) whether using additional monitors, beyond EGT_{margin}, would provide a significant improvement for monitoring performance, and 2) the impact of noise on monitoring.

Table 59. Test plan for APU monitor interfaces

SPMS monitor interfaces	Test plan	Outcome
Monitor choices: • EGT_{margin} • Bleed pressure margin • APU start time • IGV position	Sensitivity analysis of adding new monitors. Add random noise and study the impact on detection metrics	Detection accuracy. How many APUs removed due to performance issues can be detected by adding these monitors?

This test evaluates the potential benefits that additional monitors, beyond EGT_{margin}, could provide for monitoring and trending APU health. This study's assumption going into this test is that EGT_{margin} is the most important monitor for the APU and that the other monitors would provide relatively modest benefits. The reason is that hot-section deterioration is the dominant wear-out pattern for APUs and EGT_{margin} is a good indicator for that deterioration. If the additional monitor is another HI for hot-section deterioration, then it has the disadvantage of only providing corroborative evidence for what is already a good monitor. If the additional monitor is an indicator of some other failure mode, then it has the disadvantage that the other failure mode is of lesser importance than hot-section deterioration because it occurs less often, does not prevent operations, or cannot be repaired on-wing.

8.2.1.1 IGV Monitor

A monitor of the IGV position during MES can potentially provide additional health information beyond that provided by EGT_{margin}. The available information can be understood based on the IGV theory of operation.

8.2.1.1.1 IGV Theory of Operation

For a nominally healthy APU during MES, the APU controller commands the IGV to its fully open position so that the APU can supply maximal bleed air to start the main engine. This is one of the most stressful times for the APU hot section and the APU controller continuously reads the EGT sensor to ensure that the APU does not overheat. Excessive heat leads to very rapid deterioration and could cause catastrophic damage. If the EGT reaches an "IGV trim limit" threshold, then the APU controller "trims" the IGV position toward the closed position, thereby reducing the load on the APU to keep the EGT from exceeding the threshold. A trimmed IGV constitutes a reduction in functionality for the APU. The trimmed position of the IGV tends to vary with inlet temperature, because EGT varies with inlet temperature. In addition, the IGV trim limit threshold varies slightly with inlet temperature because it is acting as a proxy to limit the internal hot-section temperature. In summary, when the hot section is healthy or mildly deteriorated, the IGV is expected to be at the fully open position (approximately 90°) during MES for each flight and the EGT is generally flat or trending upwards, with fluctuations due to external conditions. When the hot section is severely deteriorated, the IGV tends to be at a trimmed position that varies from flight to flight with inlet temperature and the EGT tends to be at the IGV trim limit threshold.

In the lifecycle of the APU, once the APU switches over from fully open IGV to trimmed IGV, the trimmed IGV position becomes the key parameter for trending hot-section deterioration. EGT becomes less important because the controller limits EGT to the IGV trim limit threshold. As a rule of thumb, the start of IGV trimming occurs approximately when EGT_{margin} is crossing zero. This is by no means exact, as IGV trimming is most directly tied to EGT, which depends on the inlet temperature, inlet pressure, and generator load. The rule of thumb merely conveys that IGV becomes the key parameter near the end of the lifecycle.

The computed EGT_{margin} already takes the IGV position into account, so a separate IGV monitor for the sole purpose of measuring hot-section deterioration would not add new information. Nevertheless, IGV monitoring can provide value for other purposes, as described below.

8.2.1.1.1.1 Quantifying Degraded Bleed Air Functionality

Whenever the APU controller trims IGV to limit excessive EGT, it negatively impacts the APU's function of supplying bleed air during MES. Therefore, the trimmed IGV position could be considered a proxy for how much functionality has been lost and it serves as part of the explanation for how that functionality was lost.

8.2.1.1.1.2 IGV Position Noise

When the IGV is in the fully open position, the IGV monitor can check for noise in the measured position, which could be caused by sensor looseness, sensor electrical noise, or slop in the actuation mechanism. In the lifecycles examined for this study, the fully open IGV position was noisy in 57% of the lifecycles, with noise levels of 1 or 2 degrees standard deviation. This fault does not seem to be of concern to operators. No APUs in the data were removed because of IGV position noise. It does not appear that noisy IGVs were being repaired on-wing or between lifecycles (i.e., noisy IGVs rarely transitioned to noise-free).

8.2.1.1.1.3 Anomalously Low IGV

The IGV monitor can check that the IGV does indeed go to the fully open position during MES, except during trimming situations. The monitor can distinguish between trimming and non-trimming situations by whether the EGT is very near the threshold. Only the non-trimming situations with low IGV are considered anomalous. In the lifecycles examined for this study, anomalously low IGV occurred only sporadically, amounting to less than 0.5% of the flights. Having a low IGV is inconsistent with the understanding of the MES conditions, for which the data were supposed to have been acquired, and the anomalously low IGV points often resulted in outlier values for the EGT_{margin}. Therefore, this monitor could be used as a means for rejecting outliers and thereby produce a more sensible EGT_{margin} trend.

The recorded data also includes several BIT results and discrete signals that are relevant to IGV monitoring.

8.2.1.1.1.4 IGV Actuator BIT

This BIT reports whether there is some problem with the actuator that moves the IGV. When anomalous behavior is observed in the IGV data, this BIT can sometimes provide an explanation.

8.2.1.1.1.5 IGV Position vs. Command BIT

This BIT reports whether there is a discrepancy between the commanded position for the IGV and the actual position. The problem could be in the actuation mechanism or the IGV position sensor. When anomalous behavior is observed in the IGV data, this BIT can sometimes provide an explanation. In addition, this BIT can be used to clean out questionable data prior to trending EGT_{margin}, because a fail result from this BIT means it is ambiguous whether the recorded IGV value in the data is accurate.

8.2.1.1.1.6 MES Mode Discrete

The MES mode discrete reports whether the APU is being commanded to MES mode when the data point is acquired. The APU is capable of collecting data in non-MES regimes, and these non-MES data are mixed in with the MES data. In the lifecycles examined for this study, a small percentage of the points are non-MES. When anomalous behavior is observed in the IGV data or EGT_{margin}, this discrete can sometimes provide an explanation. In addition, this discrete can be used to clean out questionable data prior to trending EGT_{margin}, because it is generally preferable to trend only data that was acquired under consistent conditions.

8.2.1.2 APU Start Time Monitor for Hot-Section Deterioration

The recorded data from each flight include parameters acquired when the APU itself was started (i.e., conditions other than MES). One of these parameters is “APU start time,” which is the time in seconds it took to start the APU. For the APUs evaluated on SPMS, the healthy APU start time is nominally 45–55 seconds. As the hot section deteriorates, the APU start times increase mildly and have fluctuations to much higher values. APU start time fluctuations as high as 155 seconds were observed in the lifecycles evaluated for SPMS.

The top subplot of figure 47 shows an example of the recorded APU start times acquired over an APU lifecycle. Early in the lifecycle when the hot section is healthy, as evidenced by EGT_{margin} = 200°F in the middle subplot, this APU has a healthy baseline start time of approximately 47 seconds and relatively small fluctuations. By the end of the lifecycle, the hot section has deteriorated to $EGT_{margin} < 0$ and the APU start time has fluctuations as high as 120 seconds.

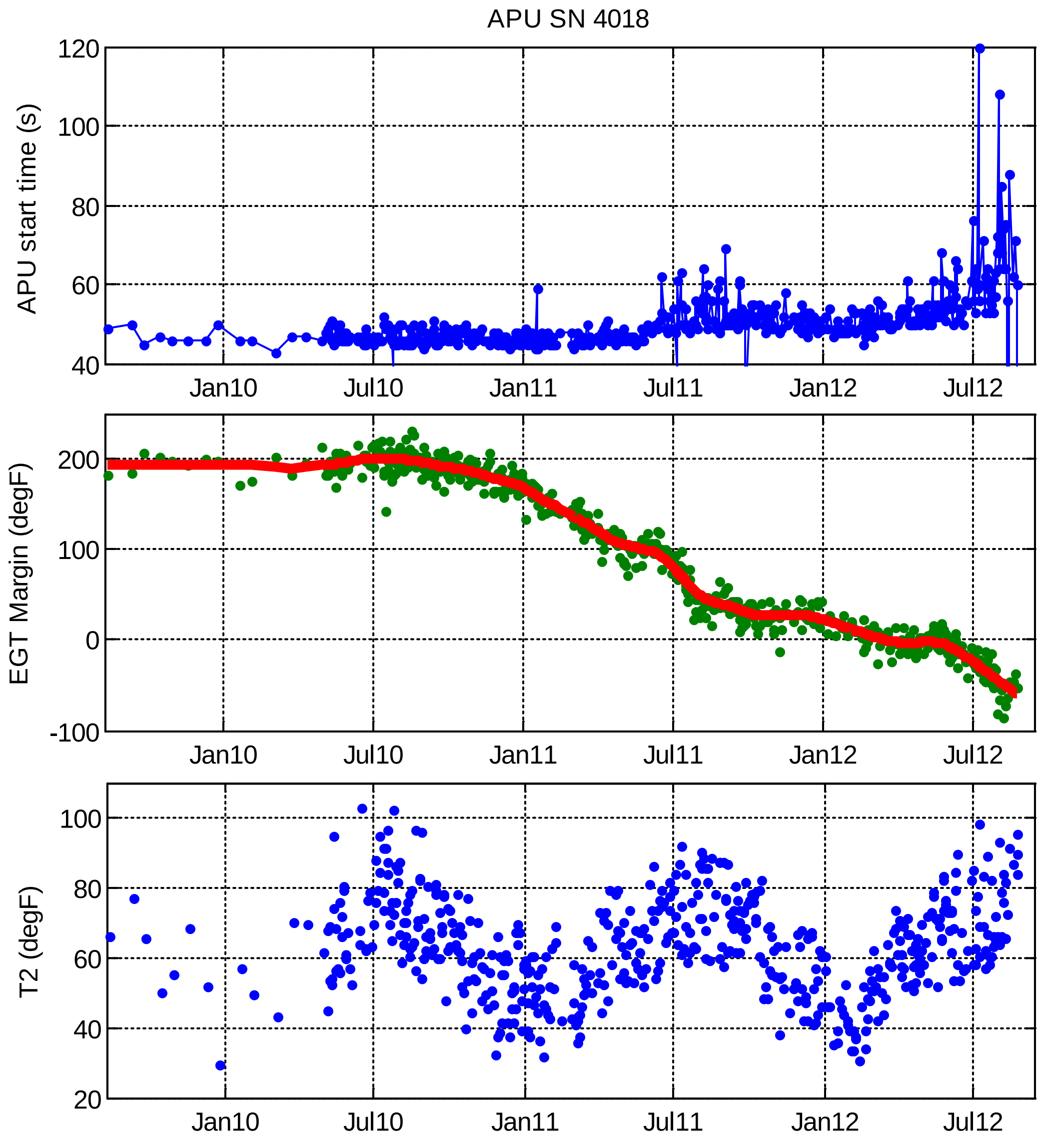

Figure 47. Example APU start time trend over a lifecycle

This APU has a healthy baseline APU start time of 47 seconds (top subplot from 2008–January 2010). The APU start time fluctuates to higher times as the hot section deteriorates (top subplot in September 2011 and April 2012), where the deterioration is evidenced by the EGT_{margin} (middle subplot). Warmer inlet temperatures T2 (bottom subplot) in July 2010 and April 2011 produce more fluctuations, whereas colder T2 in January 2011 brings the APU start times to near healthy baseline levels.

Figure 47 shows that the APU start time is also dependent on the inlet temperature T2 (shown in the bottom subplot in figure 47), with higher inlet temperatures leading to fluctuations and higher start times. During the summer of 2011, when the inlet temperatures T2 was warm and EGT_{margin} ≈ 30°F, the APU start times where slightly elevated and fluctuating. APU start times returned to near baseline during the following winter, when T2 is cooler, and the APU start times grew again in the summer of 2012.

Figure 48 shows a similar lifecycle, but where the inlet temperatures T2 were approximately 20°F hotter during summer than that of figure 47, reaching over 110°F on multiple occasions. For this case, the APU start times showed a visible increase during the summer of 2011, even though EGT_{margin} ≈ 170°F at that time was relatively healthy.

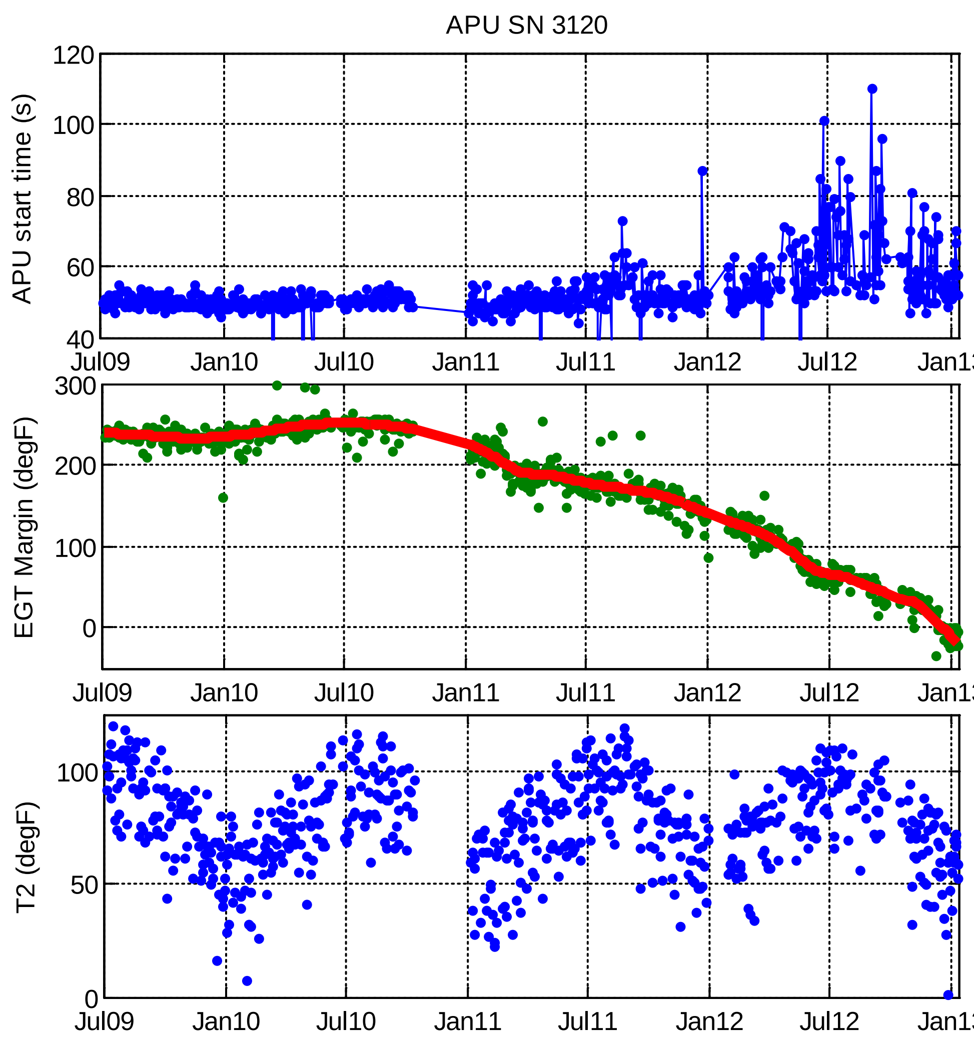

Figure 48. Example APU start times for a case with hot summers

Therefore, the challenges with using APU start time as an HI for hot-section deterioration are as follows:

- The deterioration is most visible as frequent upward fluctuations of start time to varying height. Even when deterioration is bad, some start times remain near the baseline. Fluctuations are generally harder to trend than a shift in the mean.
- The effect is one-sided with respect to a healthy baseline start time and, therefore, non-linear. APU start times increase relative to the healthy baseline due to deterioration or temperature, but the start times do not decrease below the healthy baseline when conditions are good.
- Cool inlet temperature T2 can squash the start time fluctuations and, therefore, make a deteriorated hot section look healthy. Therefore, an APU start time monitor that receives healthy-looking start times obtained during a cold period would be correct to say the results are inconclusive. A monitor that produces inconclusive results for the duration of each winter season would be harder to trend.
- Hot inlet temperature T2 can cause increased start time fluctuations and, therefore, make a relatively healthy hot section look unhealthy. If the behavior is reproducible across APUs, then a monitor may be able to factor this out using T2 corrections.

  The APU start time is also sensitive to faults in the starter motor, which are fairly common.

These challenges make APU start time a less effective monitor for hot-section deterioration than EGT_{margin}.

The authors developed a simple APU start time monitor that attempts to estimate the envelope of the fluctuations and also rejects isolated outliers. The detection threshold for the monitor was set to 60 seconds. The authors evaluated the monitor on the PTMD data and checked the end of the lifecycle for the ending EGT_{margin} and ending start time monitor value. As expected, the start time monitor often corroborated the EGT_{margin} results, but it depended on the degree of deterioration and on T2. For lifecycles ending in severe deterioration ($EGT_{margin} < -100°F$), the start time monitor exceeded 60 seconds at the end of 87% of the lifecycles. The remaining 13% of the lifecycles happened to end during a cold period. For lifecycles ending at the more mild deterioration ($0 < EGT_{margin} < 25°F$), where PTMD would be warning of a need for upcoming removal, the start time monitor exceeded 60 seconds at the end of 40% of the lifecycles.

Though the statistics are only modest for APU start time's ability to corroborate EGT_{margin} on hot-section deterioration, APU start time is an independent measurement and, therefore, has the potential to detect deterioration when EGT_{margin} somehow misses it; two such cases are available from the PTMD data, both caused by a faulty EGT sensor. These cases are shown in figures 49 and 50. The EGT_{margin} computation goes awry because of a faulty sensor and actually trends to higher values. The high APU start times could have alerted the operator that APU was deteriorating.

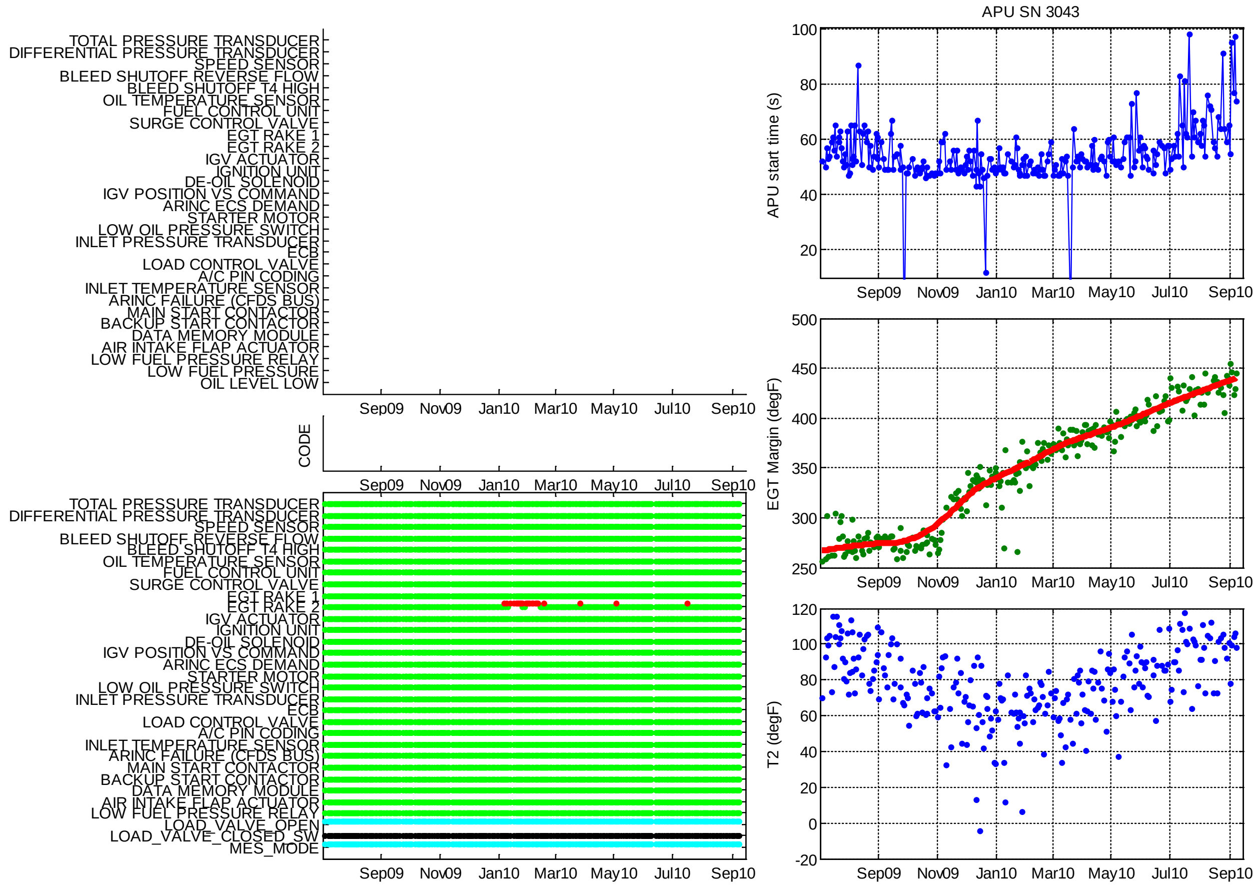

Figure 49. Case in which start times detect hot-section deterioration but EGT did not

The elevated APU start times (upper right) reveal that the APU was deteriorating, and shop findings confirm hot section deterioration. The EGT_{margin} computation (middle right subplot) went awry and increased to implausible values. The EGT problem appears to be caused by the EGT rake sensor, as indicated by startup BIT (lower left subplot). No shutdown codes (middle left subplot) or shutdown BITs (upper left subplot) were recorded for this lifecycle.

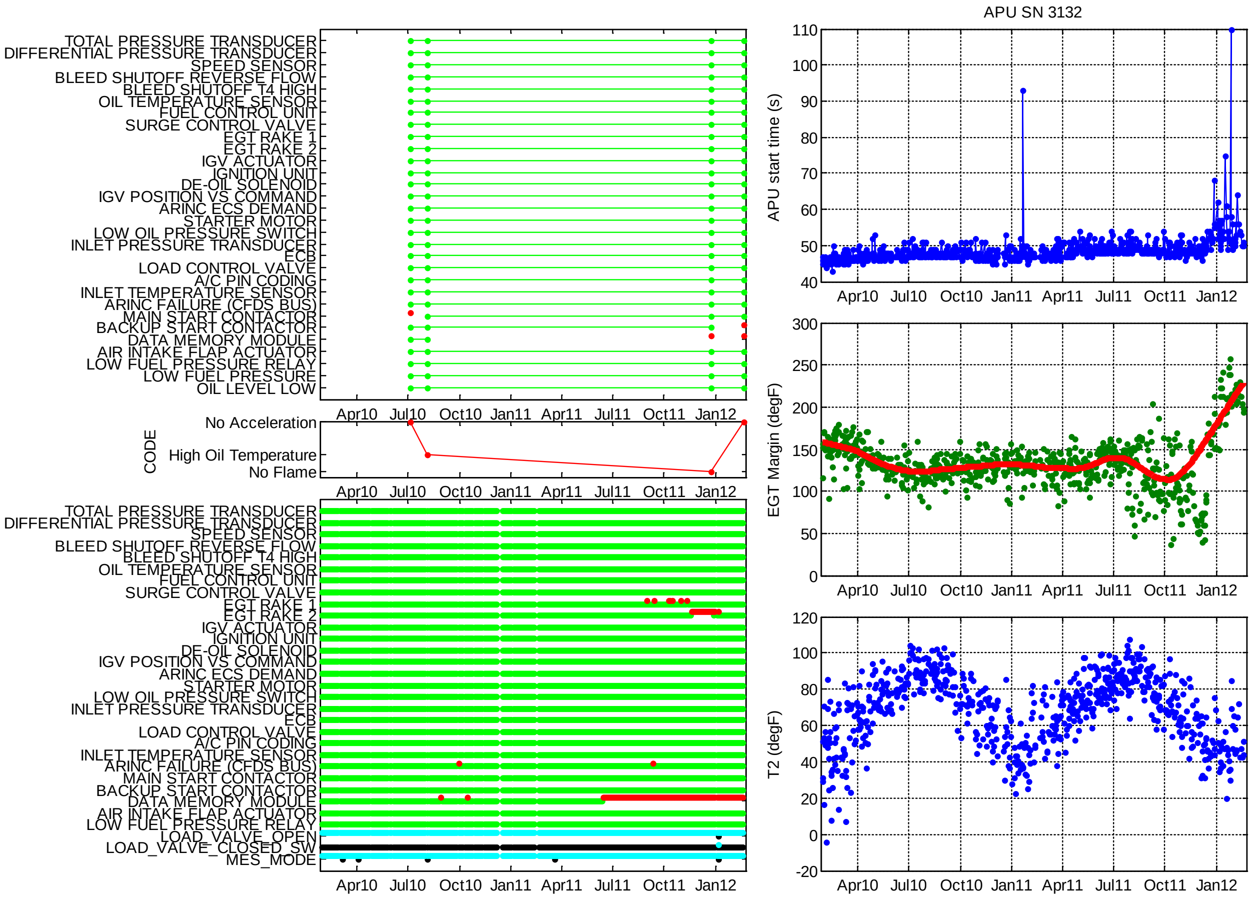

Figure 50. Another case in which start time reveals deterioration, but EGT has gone awry

APU start times (upper right subplot) are fluctuating in January 2012, indicating that the hot section is deteriorating. At the same time, EGT_{margin} has gone awry, becoming erratic and increasing implausibly. Shop findings describe that a turbine blade broke off and caused damage to downstream components. The startup BIT (lower left subplot) suggests the EGT rake sensor was damaged. The shutdown codes (middle left subplot) and shutdown BIT (upper left subplot) do not provide evidence related to the issue in this case

8.2.1.3 APU Start Time Monitor for Starter Motor Faults

APU start time is also sensitive to impending failure of the APU starter motor. Figure 51 shows an exemplary case. The APU start time trends upward for approximately four months in this case, until late April 2010, when the starter motor causes an autoshutdown. An on-wing corrective action is performed and the APU start time immediately drops to the healthy baseline.

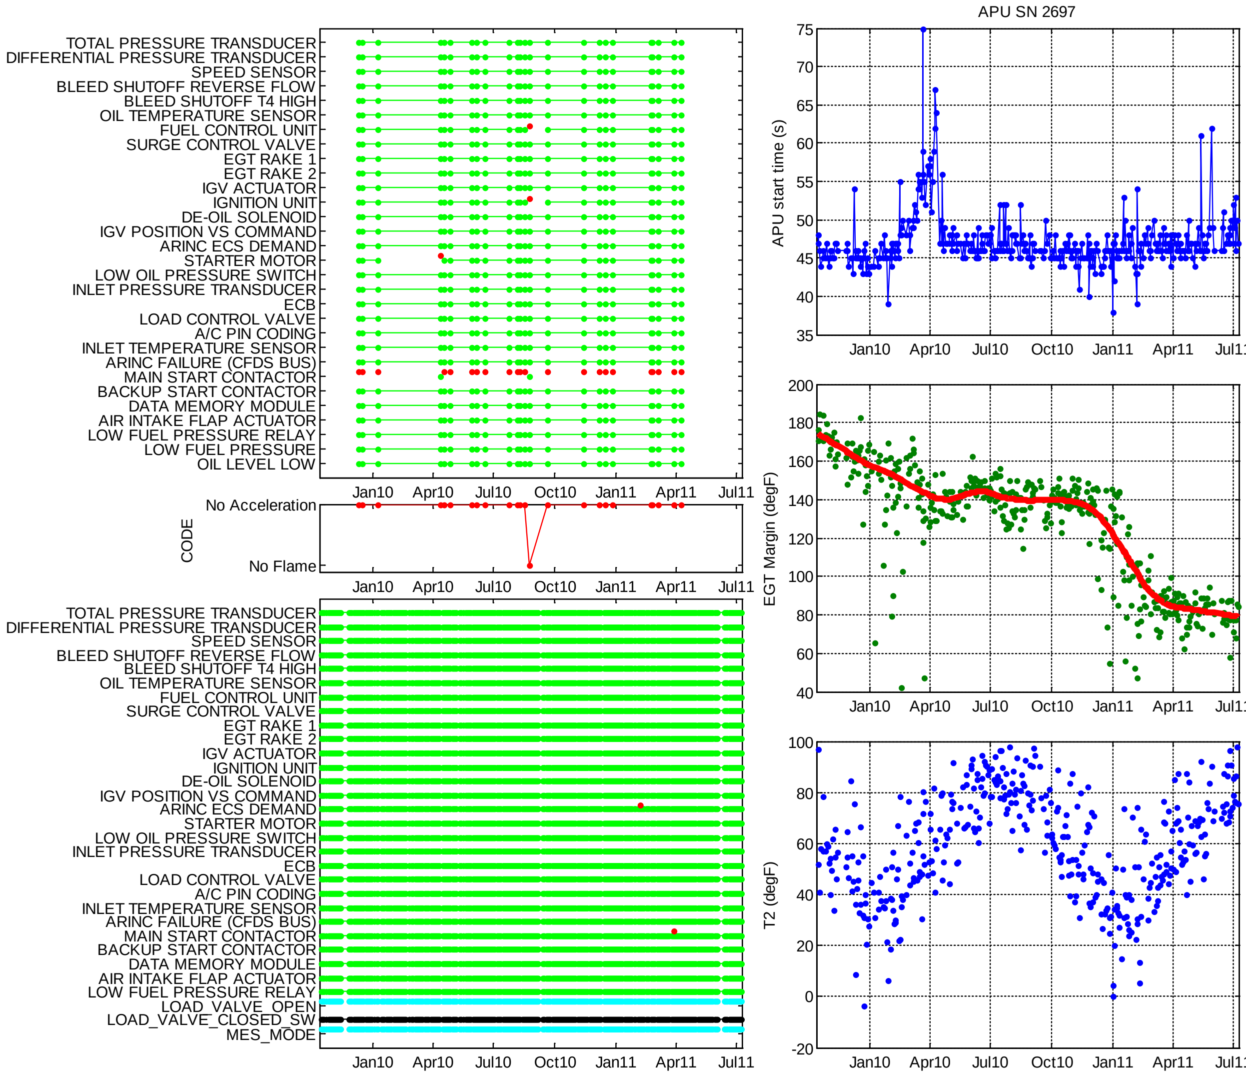

Figure 51. Lifecycle in which APU start time trending detects a fault corrected on-wing

The upper right subplot shows the APU start time slowly increases for months around April 2010 to as high as 75 seconds, until it suddenly drops back to normal, suggesting a corrective action took place on-wing. The corrective action coincides with a shutdown event because of No Acceleration (middle left subplot) and firing of the starter motor BIT (upper left subplot). The starter motor is an LRU. The EGT_{margin} (middle right subplot) and inlet temperature T2 (lower right subplot) did not contribute to the elevated start time in April 2010. The startup BIT (lower left subplot) did not detect an issue.

A failed starter motor is correctable on-wing (i.e., without removing the APU and sending it back to Honeywell), so it is less compelling than hot-section deterioration. However, it would be useful to give operators an advanced warning so they could plan the maintenance/avoid an autoshutdown. Analysis of the PTMD data shows that starter motor failure and on-wing correction happens during approximately 11% of the lifecycles.

Only a subset of cases of impending starter motor failure were slowly evolving, as in figure 51. Other cases are better described as a sudden jump followed by autoshutdown a few flights later. Figure 52 shows one of these quicker cases. The APU start time makes a sudden jump upwards and stays high for four flights until an autoshutdown on the fifth flight. A monitor that uses APU start time to warn of impending starter motor failure would want to distinguish between the slowly evolving and quickly evolving trends to judge time to failure.

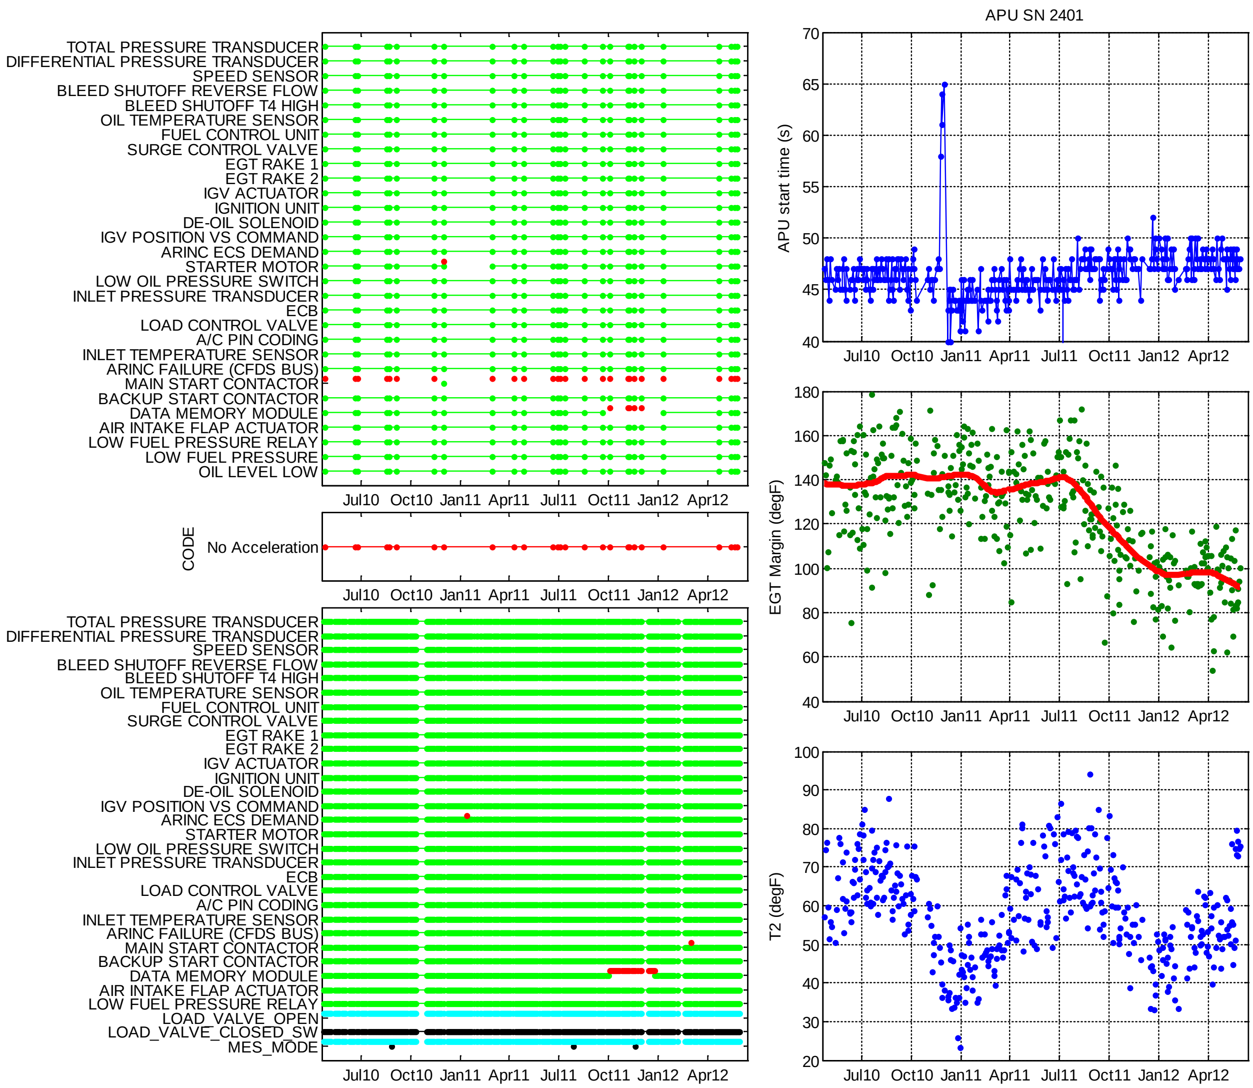

Figure 52. Another starter motor failure detected by APU start time trending

In this case, only four flights of elevated APU start times (upper right subplot) were observed until a shutdown event occurred because of No Acceleration (middle left subplot) and starter motor BIT (upper left subplot), followed by an on-wing corrective action and the APU start times returning to normal.

8.2.2 APU Health Interface

SPMS application is focused on detecting failures that may prevent an APU startup or cause a shutdown caused by fuel metering components. The health interface within SPMS needs to clearly list the APU LRU that is most likely going to cause this APU failure event. This determination is made by the diagnostic reasoner module within the Act CA-CBM function. The test plan and expected outcome for this are summarized in table 60.

Table 60. Test plan for APU health interface

SPMS health interfaces choices	Test plan	Outcome
Aggregation function: 1) univariate trending, and 2) Bayesian multivariate inference Decision threshold: 1) fixed by subject matter experts, and 2) empirically derived from the APU population	Evaluate options for single variable trending. Evaluate options for decision boundaries.	Detection accuracy: How many APUs removed because of performance issues can be detected? False positives: How many times SPMS determines that the threshold is crossed and the APU needs to be removed, but the repair shop shows no performance issues?

8.2.2.1 Bayesian Multivariate Inference

This test could not be performed per the plan because of unsolved challenges in developing Bayesian inference for the APU. The original vision was that the authors could prototype a thoroughly principled Bayesian multivariate inference reasoning and trending approach that would use the parametric observations (e.g., EGT, T2, P2, IGV, start time, etc.) and discrete observations (e.g., startup BIT, shutdown BIT, etc.) as functions of time to make sound inferences about the hidden discrete and parametric health state of the APU. The tests would show that the Bayesian multivariate approach was more accurate than the baseline approach of univariate trending.

The inferences from the Bayesian multivariate inference reasoning and trending approach were envisioned to include:

- Determine if any discrete fault is present in the APU. Examples of discrete faults include:
  - Intermittent EGT sensor faults that sometimes manifest as a bimodal distribution in EGT and sometimes manifest in EGT RAKE BIT.
  - Noisy IGV.
  - Degraded starter motor; manifests as elevated APU start times.

- If there is ambiguity as to which fault is present, infer an ambiguity group of faults with their corresponding probabilities. Also consider the possibility that the evidence is a false alarm and provide a no-fault probability.
- Use all available information to estimate the current, past, and future trend in hot-section deterioration (a parametric fault) including error bounds. Available information includes:

  - The parameters used to compute EGT_{margin}: EGT, P2, IGV, P2, generator load.
  - Other indicators such as APU start times.

- The potential of some faults to impact these parameters, such as the impact of a sensor fault on the EGT measurement or the impact of a starter motor fault on the start times.

This study's tests explored several Bayesian approaches including dynamic Bayesian networks and particle filters. However, there were challenges that could not be overcome within the scope of SPMS and given the suite of available Bayesian tools. These challenges include:

- The hybrid problem of combining discrete and continuous known and hidden variables (e.g., discrete faults and parametric hot-section degradation) in the time domain is still an area of active research. Dynamic Bayesian networks can solve certain hybrid problems but not the general problem in which the hidden states include discrete and continuous variables, which is the case for the APU. Particle filters can solve more general problems but are challenged by low-probability discrete variables, such as faults.
- Some distributions for continuous variables are non-Gaussian, including:

  - EGT_{margin} spends much of the lifecycle near its initial healthy value and less time at lower values as it degrades downward. EGT_{margin} is unlikely to spend time above its initial healthy value.
  - The time derivative of EGT_{margin} is frequently near zero and much more likely to be negative than positive.
  - IGV spends much of the lifecycle at the fully open position and spends less time at lower values during IGV trimming. It spends no time above fully open.
  - APU start time spends much of the time near its healthy baseline value. It fluctuates upward away from the baseline for hot-section degradation, but does not fluctuate downward.

- Some relationships between parametric variables are non-linear:

  - EGT_{margin} has a slightly non-linear relationship on T2, IGV, P2, and generator load.

- The strength of the correction relationship of EGT_{margin} on T2 could vary from APU to APU. For most APUs, the default correction relationship worked acceptably well. For some APUs, the correction removed T2-related variations more completely than when the default was scaled up or down.
- APU start time has a one-sided and therefore non-linear relationship on T2 and hot-section degradation. T2/hot-section degradation can cause APU start time to fluctuate upwards from its healthy baseline, but it does not do so downward.

The authors believe these challenges might be solvable with further research. This could entail moving from the tools used on SPMS for this study to a more current tool designed for this type of problem. It might also entail finding the right compromises between a full Bayesian treatment and a treatment that has tractable solutions.

8.2.2.2 EGT Decision Threshold

The threshold relevant to EGT_{margin} trending is called EGT_{design}. It is the upper limit for the corrected EGT, whereas EGT_{margin} is the difference between $EGT_{corrected}$ and EGT_{design}:

$$EGT_{margin} = EGT_{design} - EGT_{corrected} \tag{37}$$

The threshold is set so that a healthy APU has $EGT_{margin} > 0$ and trends downwards as the hot section slowly deteriorates with use. $EGT_{margin} < 0$ represents an APU that is severely deteriorated. Nevertheless, operators can, and sometimes do, continue operating the APU beyond the threshold.

The threshold EGT_{design} can reflect several concerns. These concerns are discussed in section 7.3.1. Nothing was seen in the analysis of the PTMD data that would lead to a suggestion to incurease or decrease the value of EGT_{design}. In general, the shop findings did report that hot-section deterioration is present in APUs with low values of EGT_{margin}. In no case did the shop find that a hot section looked healthy when the EGT_{margin} was low. In addition, EGT_{margin} generally coincides with the onset of loss of functionality due to IGV trimming. Therefore, the threshold is at least coarsely correct and there is no reason to believe that the APUs are being removed from wing too soon.

The shop findings used for this study did not have quantitative assessments of the deterioration of the APU. In addition, cost-of-repair data were not available.

8.2.3 APU Action Interface

The original test plan for the APU action interface was to evaluate by how much a deteriorating APU's life can be extended by an on-wing replacement of the APU fuel nozzles. The authors' understanding was that field service engineers sometimes perform a fuel nozzle replacement on a deteriorating APU to extend its life. However, after analyzing the available data on fuel nozzle replacements, there is no evidence that replacing fuel nozzles is an effective tactic to extend time on-wing near the end of the APU's lifecycle. The authors do not think a reasoner should recommend a nozzle replacement action.

Honeywell has informally collected 10 reports of nozzle remove and replace actions performed in the field. The authors express it this way because the reports are based on notes taken during regular status meetings with customers. The reports are not necessarily exhaustive regarding all of the nozzle replacements that were performed in the field. In addition, only an approximate date of the nozzle remove and replace action based on the date of the customer meeting is available. The date of the nozzle replacement could have been in the reported month or the month before.

Table 61 summarizes the results of the authors' analysis of the nozzle remove and replace cases. For each case, the EGT_{margin} graph was inspected to determine whether the trend appears to flatten as a result of the nozzle change. Figure 53 shows graphs of the EGT_{margin} trends. The four shaded rows in table 61 represent cases in which the EGT_{margin} trend does appear to flatten. The five rows that are not shaded are cases in which EGT does not flatten; in some cases, it worsened. For the cases in which the EGT_{margin} does flatten, the duration of the flattening is between 1 and 5 months. Note that the cases that exhibit flattening could be coincidental; the EGT_{margin} sometimes exhibits flattening in other examined lifecycles, but no information is available of nozzle replacements being performed for those lifecycles.

Table 61. Summary of nozzle remove and replace actions

Reported month of nozzle change	EGT_{margin} in reported month (degrees F)	Observed EGT_{margin} behavior
April 2012	25	Flattens one month before reported month. Flat for four months.
September 2010	60	Flat before. Sharp 60° F drop one month after reported month. Flat after drop.
March 2012	0	Slope worse after reported month.
January 2012	60	Slope about the same after reported month.
August 2012	150	Slope about the same or worse after reported month.
April 2012	0	Maybe flattens at reported month. Flat for two months.
November 2012	20	Flattens three months before reported month. Bumps up for five months.
October 2012	-40	Maybe bumps up or flattens. Noisy.
February 2013	0	Flattens two months before reported month. Rapid 40° F drop at reported month. Trends down after reported month (not much data).

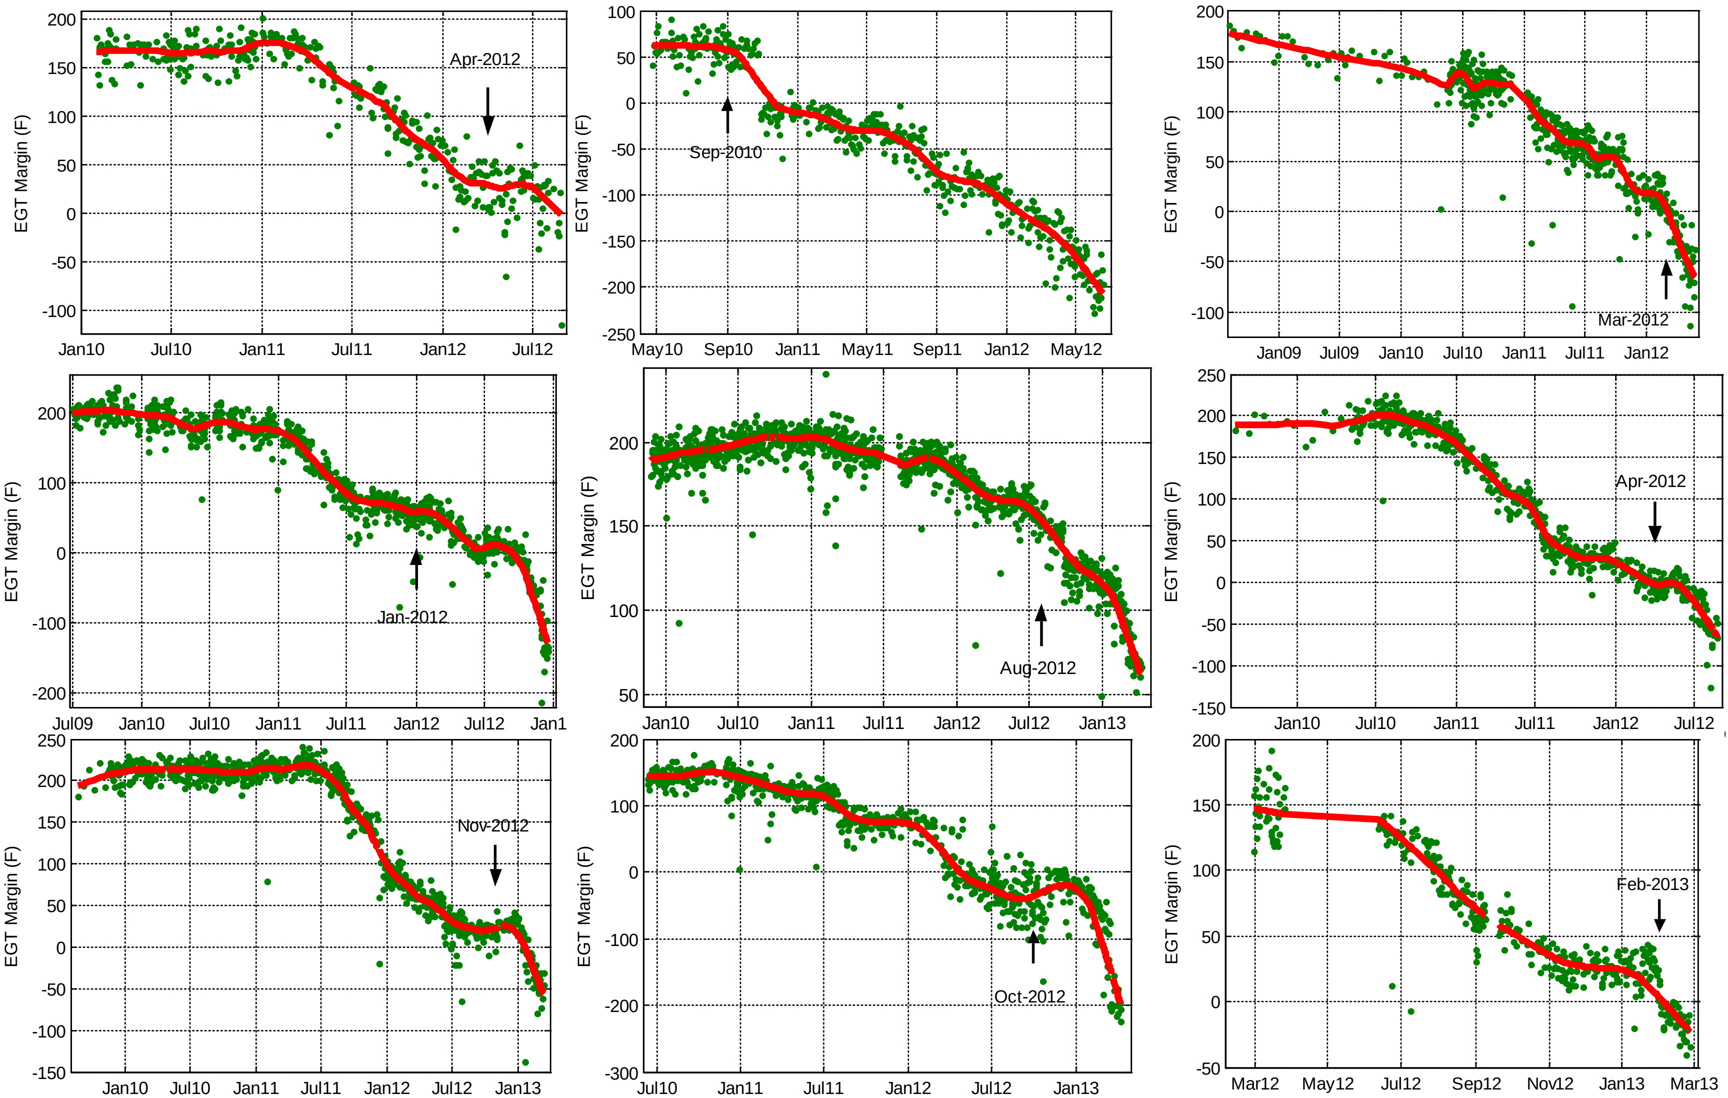

Figure 53. ***EGT$_{margin}$*** **trends for the nozzle change cases**

8.3 VALVE SPMS TEST RESULTS: EVALUATIONS

The ECS is a collection of multiple components tailored to meet the cabin and avionics heating and cooling loads. In general, ECS components are becoming more electrically driven due to an emphasis on fuel economy and MEA architectures. Talking with Honeywell experts and data from field reliably indicate that valve and filter issues are key maintenance drivers for the ECS. Because most of these components are relatively easy to maintain, remove, and replace, experts within Honeywell believe isolation to the right component is sufficient to drive SPMS value. The isolation problem is made challenging due to the lack of sensors on the aircraft.

Despite their varied designs, the basic principles of a pneumatic regulating valve remain the same. Consequently, despite the variety of distinct failure modes that can cause internal leakages, the net effect is the same: loss of regulation.

Simulation study shows capturing downstream manifold and inlet pressure provides sufficient information for generating a trendable HI for mechanical pneumatic problems.

Monitoring logic entails counting the number of pressure measurements in predefined bins. For each flight, calculate a weighted average over these bins; the minimum delta pressure, $P_{in} - P_{dwn}$, across all these bins; and the maximum across all these bins.

The three CIs are trended after a post-processing smoothing step (low-pass filter, min-hold or max-hold) is applied.

The test objective is to establish the correctness and robustness of these trends.

8.3.1.1 Sensor Interface

Measurements available from the simulator for each experiment are listed below:

- Valve serial number: This is a constant. Because a family of valves is simulated, this provides a reference to track data from each valve.
- Valve on/off command: Binary 0/1 value indicating whether the valve was activated.
- Valve downstream pressure P_{dwn} in psig.
- Valve inlet pressure P_{in} in psig.
- Valve inlet temperature T_{in} in F.
- Aircraft altitude F_{alt} in ft.
- Ambient pressure P_{amb} in psia.
- Ambient temperature T_{amb} in F.
- Cumulative flight time t_{flt} in hours.
- Valve cumulative operating time t_{vlv} in hours. The cumulative time when the valve was commanded to open and actively regulate the pressure. By definition: $t_{vlv} \leq t_{flt}$.

Because the ambient temperature characterizes the flight altitude, the authors decided not to include this measurement in the SPMS sensor interface. Furthermore, valve cumulative flight time can be readily derived from valve on/off times and the flight times. This assumes that the valve installation did not change, which is an unreasonable assumption to make. Therefore, the authors decided to drop this sensor. The valve on/off command is a binary 0/1 value indicating if:

- Valve downstream pressure P_{dwn} in psig.
- Valve inlet pressure P_{in} in psig.
- Valve inlet temperature T_{in} in F.
- Ambient temperature T_{amb} in F.

The SPMS sensor interface is summarized in table 62.

Table 62. Sensor interface for inline regulation valve

<table>
<tr><th>Sensor</th><th>Sensor interface</th><th>Trigger</th></tr>
<tr><td>Serial number</td><td>System data input at setup</td><td>-</td></tr>
<tr><td>Model number</td><td>System data input at setup</td><td>-</td></tr>
<tr><td>Ref designator</td><td>System data input at setup</td><td>-</td></tr>
<tr><td>Aircraft tail</td><td>System data input at setup</td><td>-</td></tr>
<tr><td>ON</td><td>Solenoid current or valve on/off command binary value indicating if the valve was being actively commanded by the controller</td><td rowspan="4">Valve is ON and inlet pressure is above the regulator set point</td></tr>
<tr><td>Downstream pressure</td><td>Valve downstream pressure, P_4, in psig</td></tr>
<tr><td>Inlet pressure</td><td>Valve inlet pressure, P_1, in psig</td></tr>
<tr><td>Inlet temperature</td><td>Valve inlet temperature, T_{in}, in F</td></tr>
</table>

8.3.1.2 Data Interface

The bottom subplot of figure 54 shows the transient response of the valve when the inlet pressure (top subplot) rises from 15 psig to approximately 70 psig. High variability in operating conditions and the associated valve response make a transient-based approach highly difficult to consistently capture and analyze. A robust data interface for SPMS reduces the monitoring window by applying data filtering and retaining only the relevant part of the sensor measurements. The SPMS data interface specifies this for the inline regulation valves.

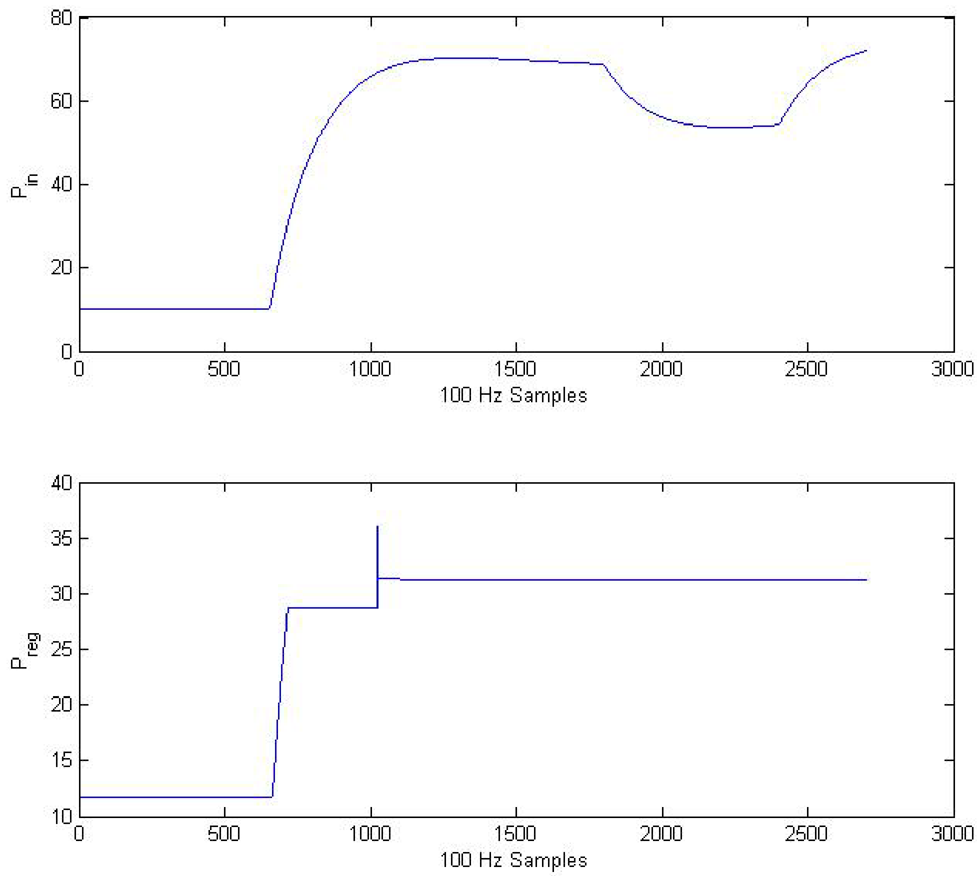

Figure 54. Transient response of the valve

Though a valve operates intermittently over each flight, discussions with subject matter experts indicate the amount of wear is greater at pressures around its calibration point; the valve is supposed to hover around this point but often oscillates, causing wear on the seals and valve seat. Therefore, the SPMS data interface is defined to include regulation pressure P_{dwn} samples taken at relatively low frequency. The SPMS data interface is summarized below:

- Capture trigger logic:
  - Valve is commanded to be ON.
  - If the downstream pressure is not indicating a "stuck closed fault."

    Inlet pressure is above the regulator set point, $P_{in} > P_{reg}$.
- Delta pressure between the valve inlet and downstream is greater than a certain amount, $P_{in} - P_{dwn} \geq \theta$. Typically: θ = 10% of P_{reg}.
- Data frame variables: P_{in}, P_{dwn}, T_{in} at 1-minute interval.

8.3.1.3 Monitor Interface

The SPMS data interface provides a series of P_{in}, P_{dwn}, T_{in} sensor values. In any given flight, if the valve is open for more than n minutes, the data interface can provide, at most, n such data frames. There are many choices to summarize the valve behavior from this flight. The SPMS monitor interface choices evaluated are discussed earlier. The traces of the CIs generated from these two choices (see figure 55) are:

- Maximum downstream pressure, max(P_{dwn}).
- Minimum delta pressure across the vale. That is, min(P_{in} P_{dwn}).

It is important to note that, despite the high variability in the data, both max(P_{dwn})and min(P_{in} P_{dwn}) can be trended and changes in the slope indicate a degrading valve or gradual loss of the regulation function that will eventually warrant a removal action.

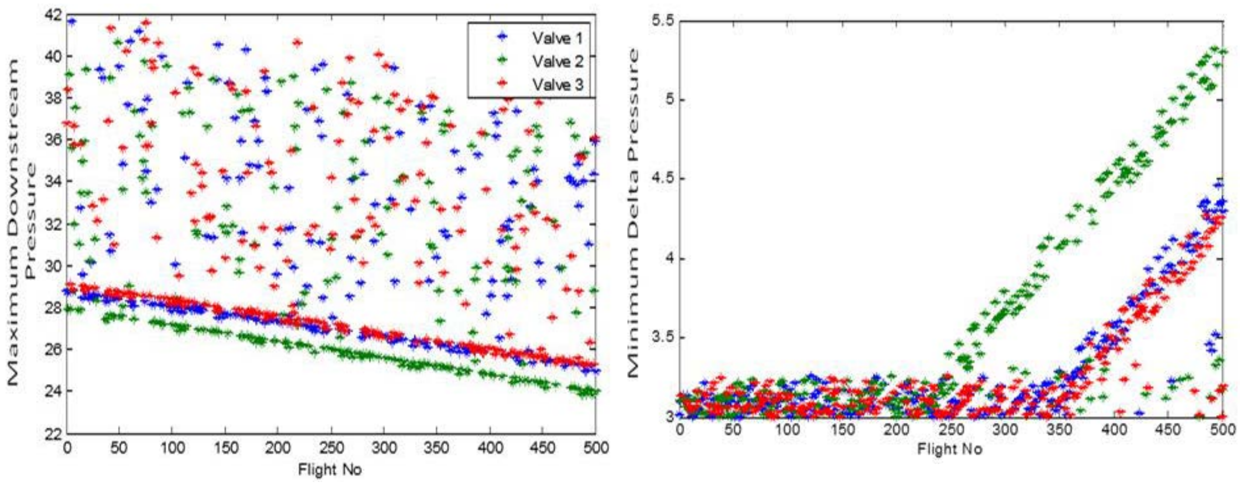

Figure 55. Two monitor interface choices for generating CIs

8.3.1.4 Health Interface

The health interface provides a yellow or red indication for removing the valve because of internal wear and eventual loss of the regulation function. The indicator is predictive in nature and, if no maintenance is performed, the valve will issue a fault message to the onboard maintenance system. The yellow and red indicator thresholds are critical. If max(P_{dwn}) is chosen, then the health interface needs to establish a "lower bound" below which an indicator is generated. Conversely, if min(P_{in} P_{dwn}) is chosen, then the health interface needs to establish an "upper bound" above which an indicator is generated:

- Lower bound on max(P_{dwn}): Triggers a notification when max(P_{dwn}) is less than this threshold.
- Upper bound on min(P_{in} P_{dwn}): Triggers a notification when max(P_{dwn}) is greater than this threshold.

The evaluation needs to establish the upper and lower bounds (see figure 56) using a large valve population. Furthermore, the authors believe a combination of the two options provides a more robust solution (see figure 57).

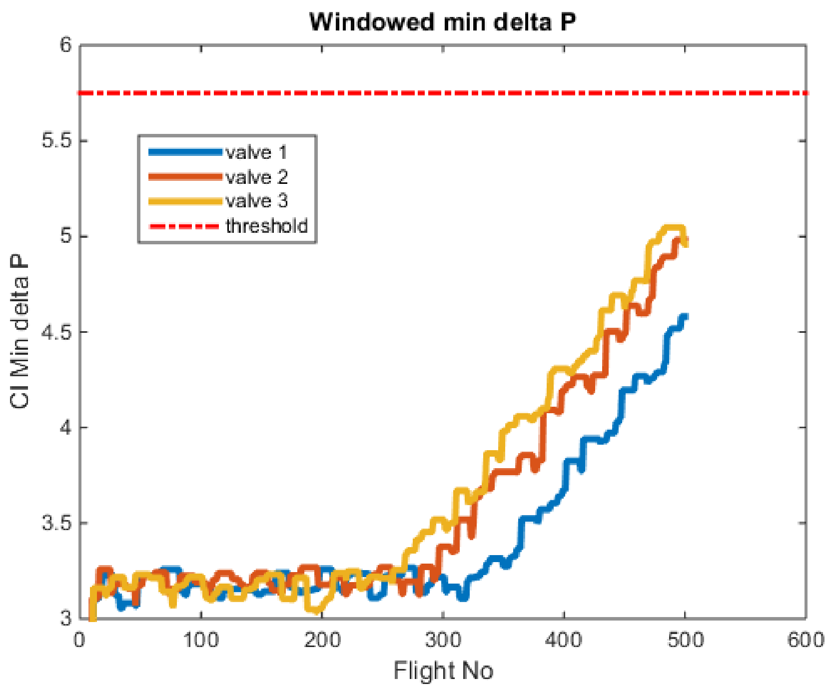

Figure 56. Upper bound on minimum pressure drop using 10-flight running window

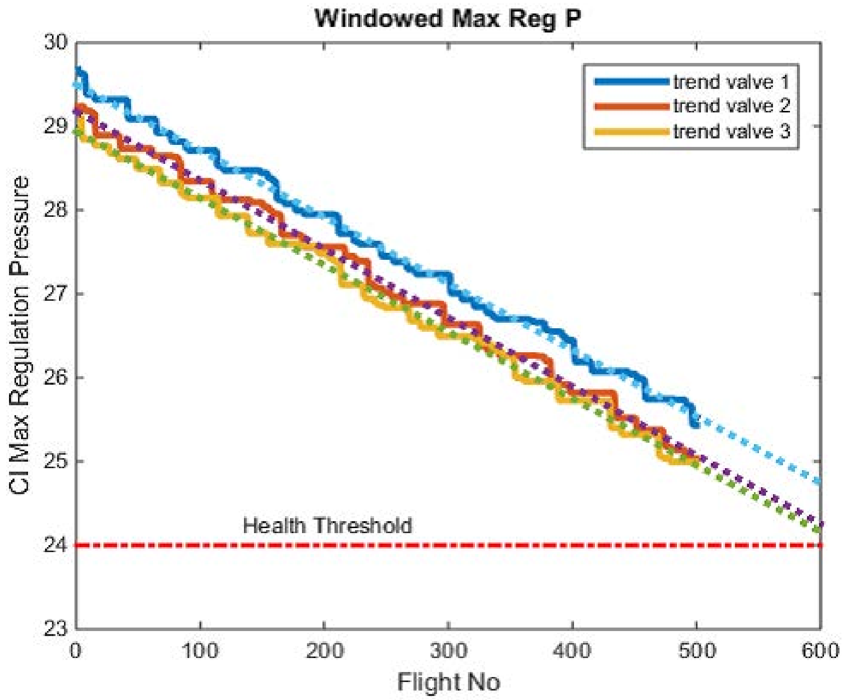

Figure 57. Maximum regulation pressure over 10 flights with linear trends

8.3.1.5 Impact of Interface Parameters on System Performance

This section reviews the impact of windowing and down sampling on system performance. First, to understand the impact of windowing, window sizes 5, 10, 20, and 30 samples were used to compute the trends for max regulation pressure and minimum delta P. Figure 58 shows the impact of window size on the minimum delta P HI trend. It shows that window size 10 and above is good for this indicator. However, the indicator shows a staircase shape as the windowing increases. The window size does not impact the maximum regulation pressure HI trend significantly (see figure 59). Figure 60 shows the impact of down-sampling on both CIs. The CIs become noisy only after significant down sampling (to 0.05 Hz). The valve SPMS used simulation data that could explain why the CIs behaved gracefully with down sampling and windowing.

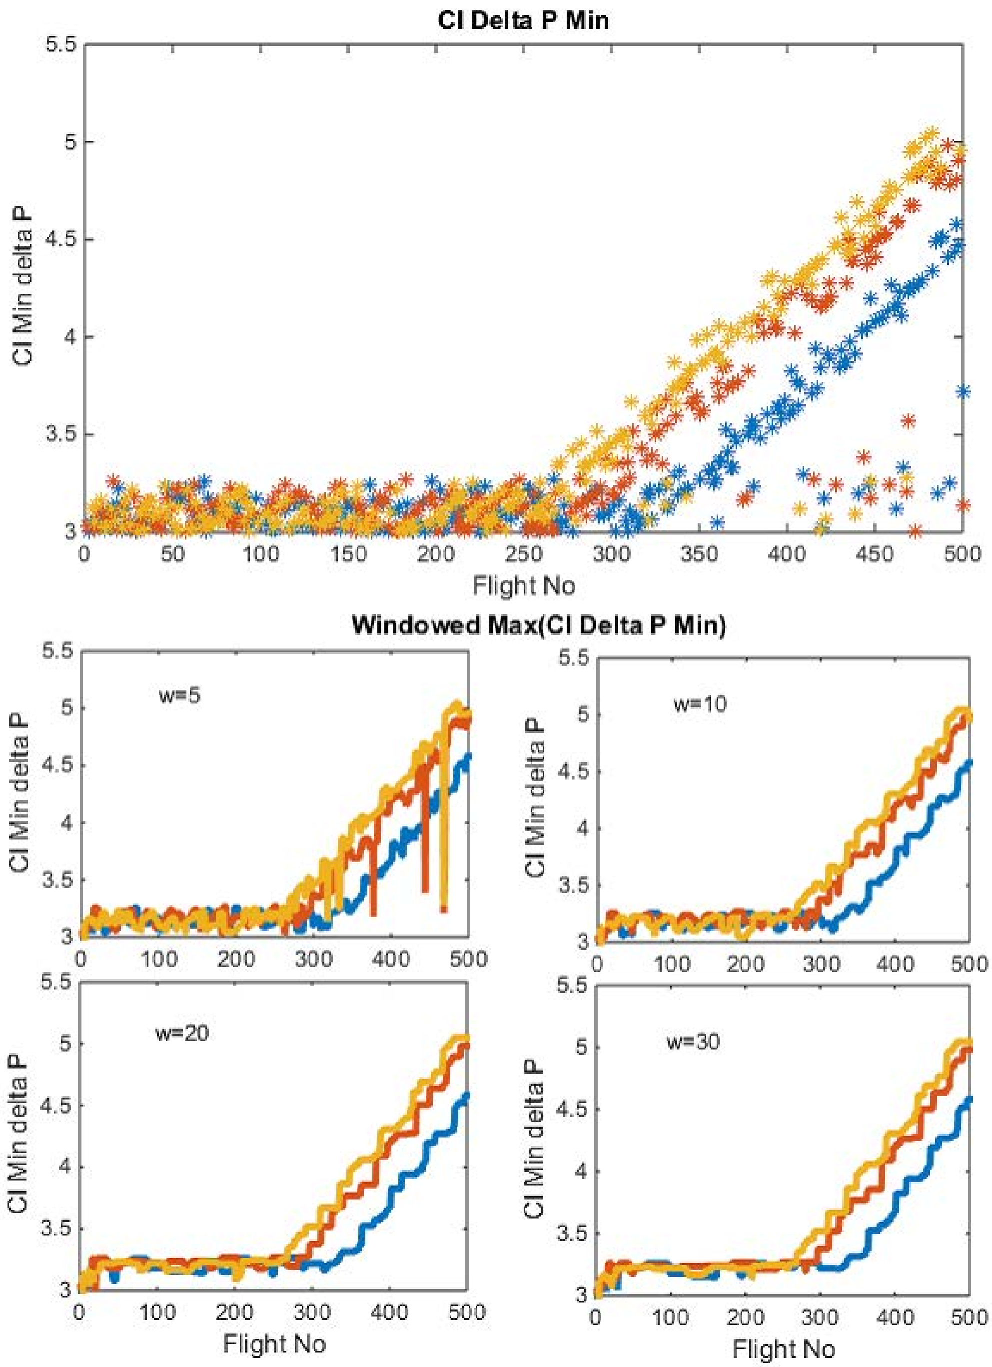

Figure 58. Impact of window size on the minimum delta P HI trend

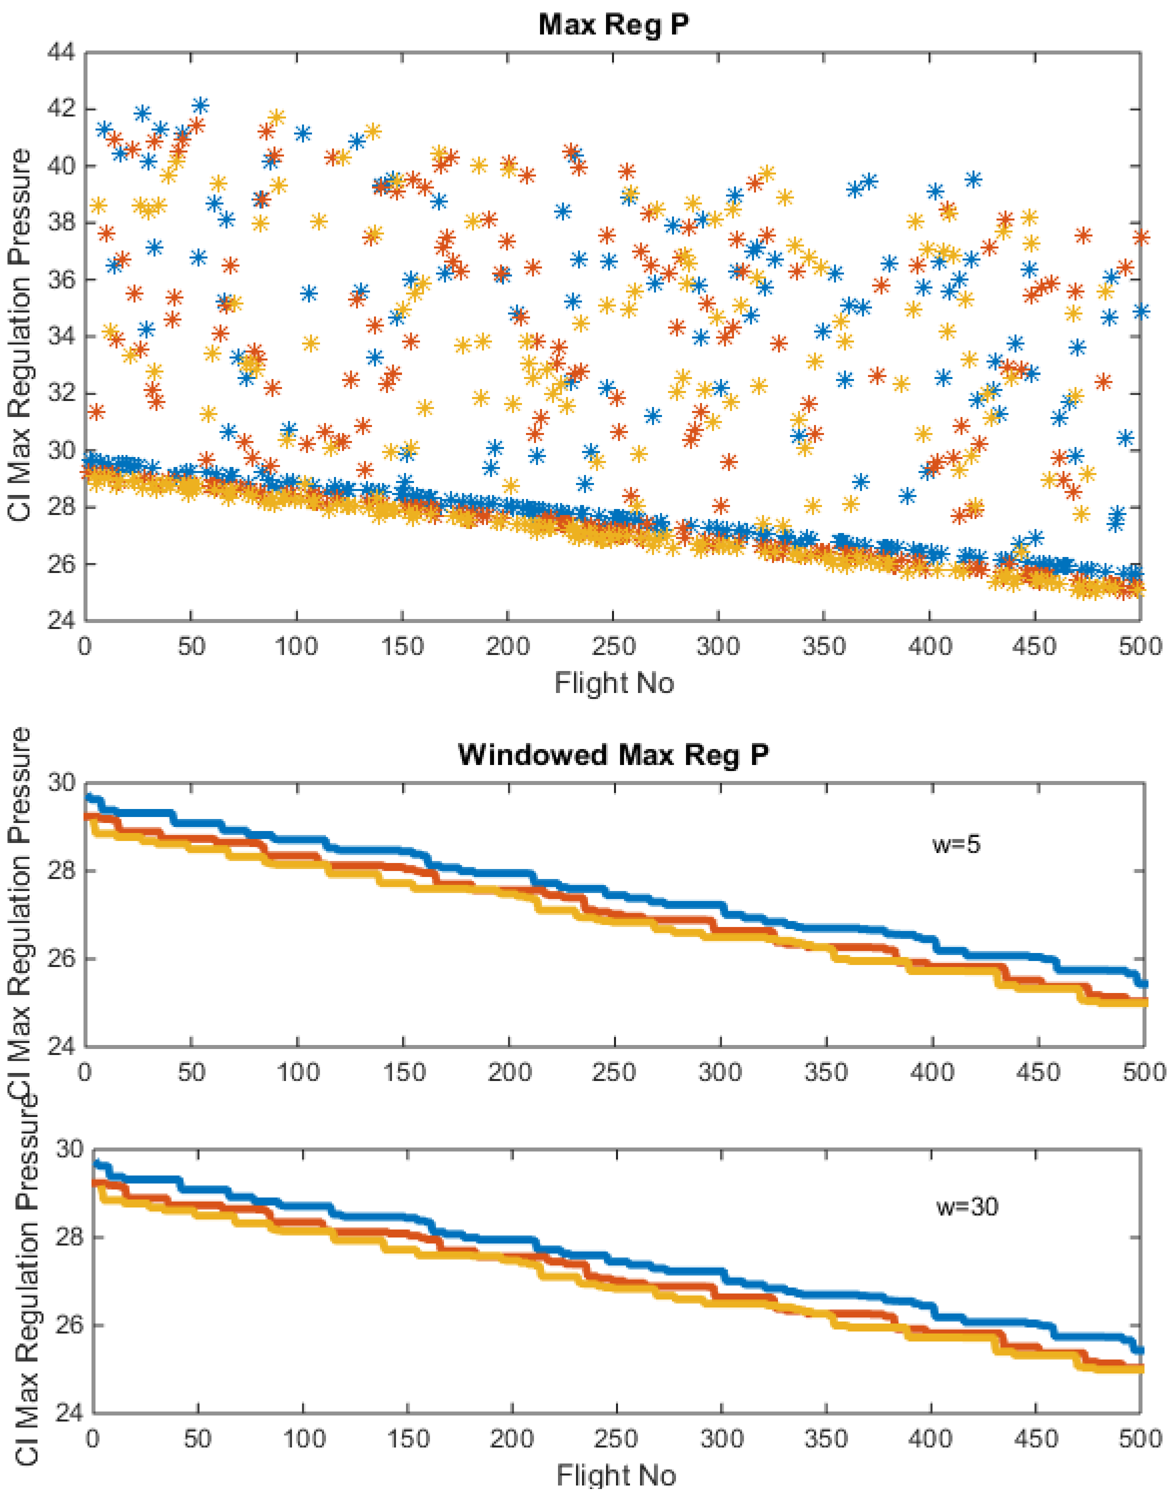

Figure 59. Impact of window size on the maximum regulation pressure HI trend

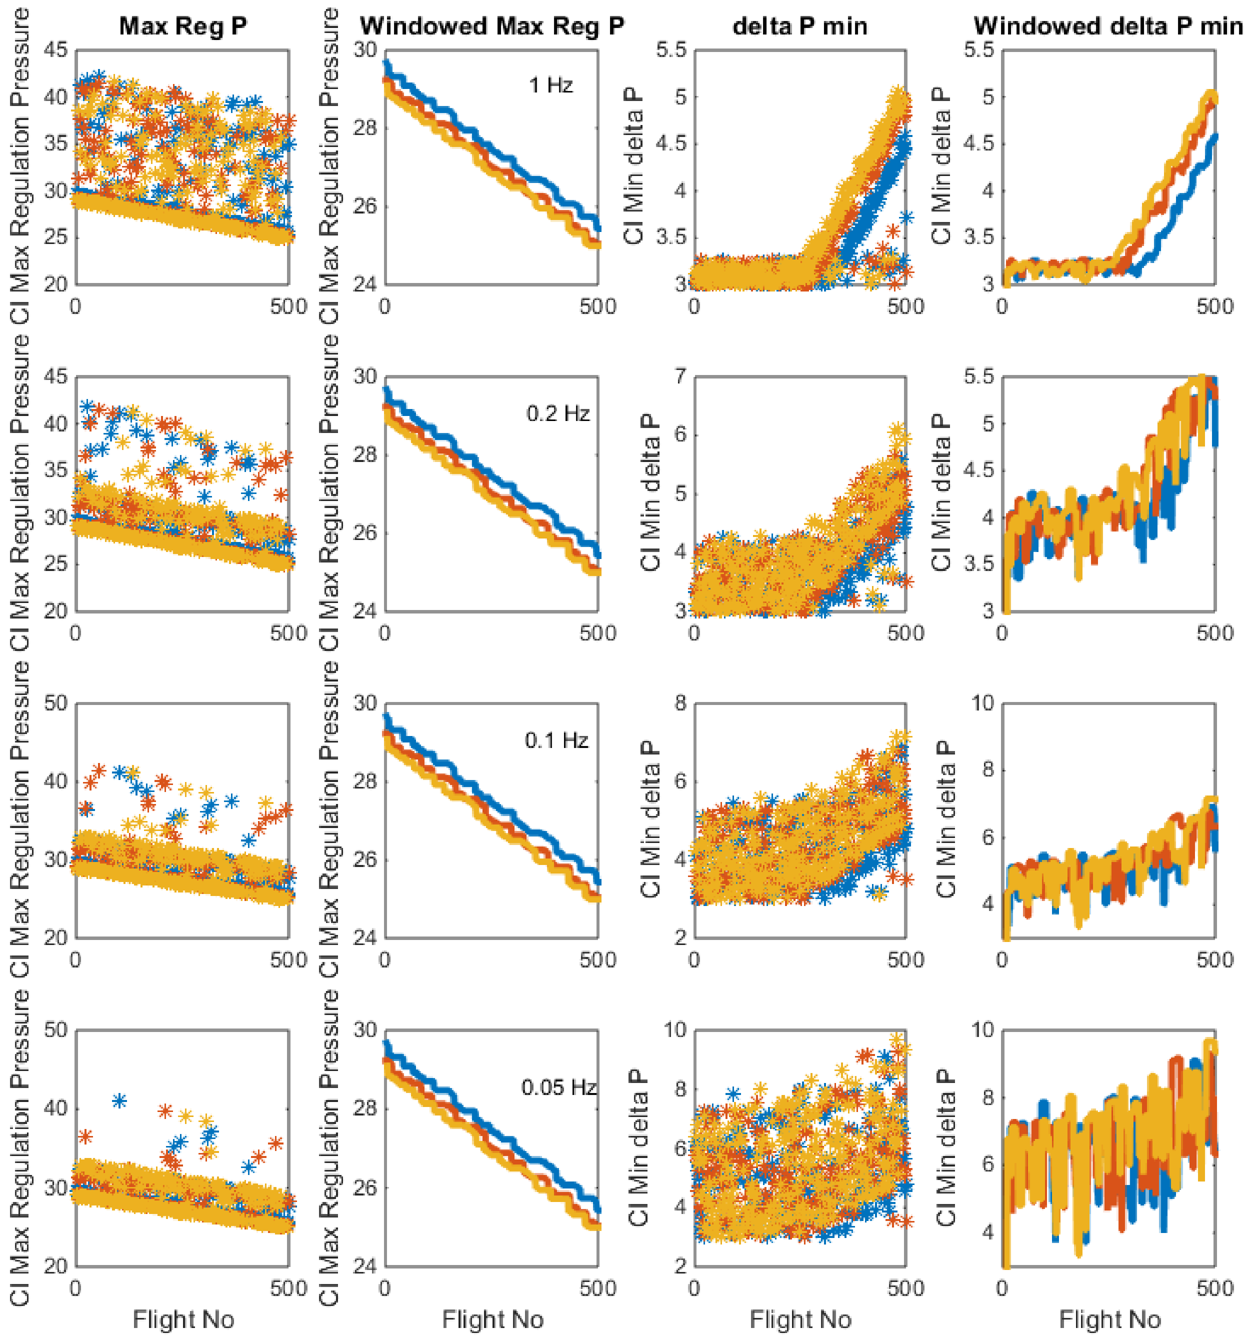

Figure 60. Impact of down sampling on HI trends

9. TASK 8: STANDARDIZATION OF INTERFACES

This report provides a methodology for specifying component/subsystem/system design characteristics which facilitates the SPMS development for monitoring, detection, and management of degraded states, enhanced diagnostics, and prognostic health assessments of aerospace systems. At the high-level, the SPMS systems embody many of the features currently under consideration by SAE International under the technical committee HM-1, Integrated

Vehicle Health Management (IVHM). The HM-1 committee is developing multiple aerospace recommended practices (ARPs) to help mitigate existing barriers to the successful implementation of IVHM technology in the aerospace sector. One of the documents under development, ARP-6268, “Design & Online Communication Standards for Health-Ready Components,” addresses requirements for health-ready components. The health-ready components are expected to incorporate the following features to varying degrees:

- Sensor mechanisms to continuously monitor critical system functions and environmental parameters.
- Raw sensor data acquisition and data processing that produce health-state parameters or indicators.
- State detection and health assessment functions that can synthesize health-state information to define system degradation severity levels (e.g., normal, warning, critical).
- Prognostic algorithms to predict remaining useful life (RUL), also known as “performance life remaining” of the degraded system.
- Enhanced diagnostic functions that correlate health-state data, including BIT results across system functions or subsystems.
- Communications interface module to transmit raw and health-state data to the data acquisition subsystem, data recorder, or vehicle control system.[4]

These features are also common to the development of SPMS systems. In addition to the expressed features, it is essential to gather system lifecycle (i.e., design, operation, and repair) data for SPMS development. Section 9.1 elaborates on the data collection for the target systems selected for this SPMS study.

9.1 SPMS DATA COLLECTION GUIDELINES

This section captures SPMS architecture and interface designs for three selected systems: the engines, APU, and valves. To implement SPMS for selected failure modes of the target systems, the sensor data collection during operation needs to be standardized; on-wing field observations and maintenance actions need to be recorded; and the remove and repair findings need to be recorded. Figure 61 shows the lifecycle data sets required for the development of SPMS applications. This approach provides the ability to trace health/degradation of a given asset using sensor data and correlates the improvements in health to appropriate on-wing repairs and replacements of the faulty parts. The collected data provide means to evaluate prognostics accuracy and repair effectiveness.

[4] SAE ARP6268, Design & Online Communication Standards for Health-Ready Components (currently in draft stage; expected to be released for committee review in spring 2016).

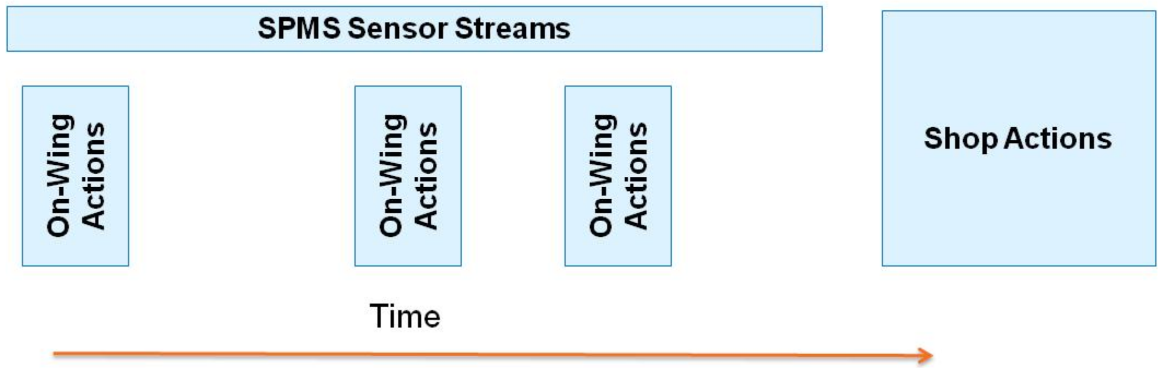

Figure 61. Lifecycle data sets required for the development of SPMS applications

The SPMS sensor stream is a collection of sensor data from the aircraft. The data streams can contain time-series or snapshot data. The choice of time-series versus snapshot is based on maturity; in the early stages of an SPMS, the organization may choose to collect large amounts of time-series data for subsequent analysis and refinement of their techniques. Once an SPMS is mature, the organization may choose to collect only snapshots of the needed data. In sections 9.1.1–9.1.3, the developed engine SPMS uses the time-series data, whereas the APU and valve SPMSs use the snapshot data as sensor streams.

In summary, when developing new SPMS applications, it is advisable to use the time-series data initially. When the SPMS target is clear, a snapshot or event data can be used as a sensor stream to build the SPMS application.

9.1.1 Engines

To implement SPMS for selected failure modes associated with the aircraft engine fuel controller failures, sensor data collection; on-wing observations and maintenance actions recordings; and shop remove and repair recordings need to be standardized. The sensor data collection requirements for engine SPMS are shown in table 47. These sensor signals need to be captured at prescribed rates per engine start including both ground runs and flights. Table 63 defines the associated header information needed with each record set collected using table 47 specifications. The collected header data can also be useful for reliability and product improvement programs. The run time data/signals are used by the engine SPMS to generate CIs. The CIs are associated with the field observations and findings from both the on-wing maintenance actions and shop maintenance. Table 64 captures the interface information associated with both on-wing actions and shop actions. Table 63 shows the information summaries with maintenance events, and table 65 lists the repair actions. The information captured in tables 64 and 65 associates the CI trends with the appropriate failure modes.

Table 63. Information recorded per download

Fields	Description
Date time	Time for the record
Model number	Engine model number
Serial number	Serial number
Aircraft tail number	Aircraft trail number
Engine position	Position of the engine when all of the engine records are not in the same file
Record ID	Record ID for the sampled data set
Engine hours	Engine hours (end engine hours – beginning hours)
Engine cycles	Engine cycles (end engine cycles – beginning cycles)

Table 64. Information recorded for on-wing engine actions

Information tag	Description
Model number	Engine model number
Serial number	Serial number
Tail number	Aircraft trail number
Engine position	Engine position
Installation date	Engine install date
Date time	Repair date
Hours	Hours
Cycles	Cycles
Parts removed	List of parts/assemblies removed
Part serial numbers	Removed parts (LRU) serial numbers
Reason for repair	Reason for repair for LRU
Actions	Actions performed
Did this work	Yes/No, additional text inputs

Table 65. Recording of the engine shop repair events

Repair Information	Description
Model Number	Engine model number
Serial Number	Serial number
Equipment condition	Details in repair findings below
Installation date	Installation date
Repair date	Repair date
Hours	Hours (if available)
Cycle	Cycled (if available)
Repair findings (condition of components)	Seals
	Cooling
	Components removed
	HMA
	Controller
	Igniter
	Valves

9.1.2 APUs

Based on Honeywell's experience with PTMD and on this study, the authors consider the following types of data to be of use in an SPMS for APUs:

- MES: These data are automatically generated by the APU from sensor streams and control signals and characterize the APU's performance and behavior while it is being used to start the main engines (nominally once per flight, before each flight). MES is of interest because it is where the load on the APU is typically highest. The data are used to trend deterioration of the hot section and, therefore, the APU's ability to meet the load demands of MES without overheating.
- APU startup: These APU-generated data characterize the APU's behavior while the APU is being started (nominally once per flight, before each flight). The startup data can provide evidence of hot-section deterioration or impending starter failure.
- Summary of each use: These APU-generated data summarize each use of the APU from startup to shutdown at some point after MES. The data include the results of BITs and the cumulative usage in hours and cycles.
- Autoshutdown and failed startup events: These APU-generated data characterize the state of the APU during an autoshutdown or a failed start event, in which the APU controller shuts down the APU because of a critical fault and/or to protect the APU. Such events can occur sporadically throughout an APU lifecycle, not just the end of a lifecycle. The data include the fault code for the event (e.g., no flame, no acceleration, overtemperature) and results of BITs.

- On-wing maintenance: These operator-generated data describe any maintenance or corrective actions performed by the operator while the APU remained on-wing; the aircraft operator's maintenance team could perform the actions without needing to send the APU back to the manufacturer. Such events can occur sporadically, or not at all, during an APU lifecycle. These can include removal and replacement of a failed LRU on the APU (e.g., replacing a failed starter motor), a water-wash of the gas path to maintain performance, etc.
- APU removal: These operator-generated data include the operator's reason for removal and any supporting observations. Example reasons include failed starts, high EGTs, odor in cabin, and scheduled removal.
- Shop maintenance: These data are generated by the APU manufacturer's maintenance shop for each APU returned from the field. The data consist of health observations and any corrective or regular maintenance actions needed to return the APU to healthy working order before returning it to service. The observations and actions address the specific problems (if any) that led to removal and any other problems found during the inspection of the APU. Shop maintenance marks the end of an APU lifecycle; the next lifecycle starts when the APU is installed on another aircraft.

Figure 62 shows these data for the APU. The Honeywell APU used for this project records startup, MES, and summary data in a combined report, with the report for each flight recorded to non-volatile memory after the APU is commanded to shutdown. For cases in which the APU fails to start or performs an autoshutdown, the APU records the associated data to shutdown or failed start report and there is no combined startup-MES-summary report. For this project, the APU removal data and shop maintenance data was accessed via Honeywell's Product In-Service Performance System (PIPS) database. The PIPS database contains an abridged version of each shop maintenance record; the full shop records were not readily available for this study. No on-wing maintenance records were available to use for this study, as represented in figure 62 (by the on-wing MR being grayed out). Tables 66–70 list the fields in each report in figure 62, with the exception of the on-wing MR.

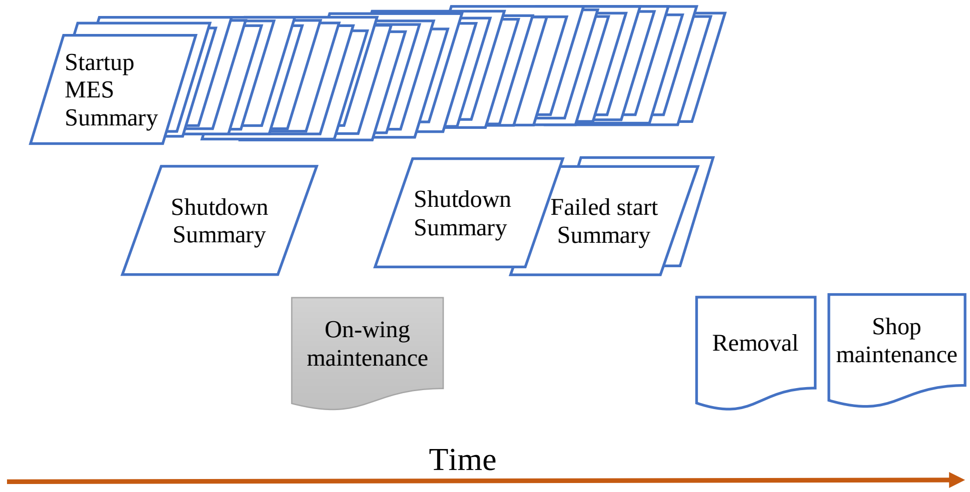

Figure 62. Data associated with an APU lifecycle (the rhombus shapes represent the many reports that the APU generates automatically from sensor data, nominally one report per flight; the "document" shapes represent human entry)

Table 66. Combined report contents for APU startup, MES, and summary

Parameter	Group	Units	Description
APU model	Summary	N/A	Model name of the APU. Used to identify the APU that produced the report.
APU serial number	Summary	N/A	Serial number of the APU. Used to identify the APU that produced the report.
Report date and time	Summary	Date	Date number representing date and time when the APU received the start command.
Cumulative operating hours	Summary	Hours	Cumulative operating hours since the APU was new. A measure of cumulative usage.
Cumulative start cycles	Summary	Cycles	Cumulative start-stop cycles since the APU was new. A measure of cumulative usage.
BIT results	Summary	N/A	Pass/fail BIT results spanning startup to shutdown. A fail result effectively latches for purposes of the final result.
TAT	Summary	Degrees C	Total air temperature, as measured by other sensors on the aircraft.
APU start time	Startup	Seconds	Time it took for the APU to startup.
Inlet temperature, startup	Startup	Degrees C	Inlet air temperature when the APU was started. Air temperature can impact start times and EGT.
EGT peak during APU start	Startup	Degrees C	The maximum EGT reached during startup.
Speed at EGT peak	Startup	% rpm	The shaft speed of the APU at the time of the EGT peak.
EGT	MES	Degrees C	The temperature of the APU exhaust gases during MES. An indicator of hot-section health, after correcting for external factors.
Inlet temperature, MES	MES	Degrees C	Inlet air temperature during MES. Used to compute a “corrected EGT” for trending.
Generator load	MES	% full load	The generator load on the APU during MES. Used to compute a “corrected EGT” for trending.
Inlet pressure	MES	psia	Inlet air pressure during MES. Used to compute a “corrected EGT” for trending.
IGV position	MES	Degrees	Angle of the IGVs of the load compressor during MES. Used to compute a “corrected EGT” for trending.
LC discharge pressure	MES	psia	Load compressor discharge pressure during MES. An indicator of load compressor health, after correcting for external factors.

Table 67. Report contents for APU autoshutdown and failed startup

Parameter	Group	Units	Description
APU model	Summary	N/A	Model name of the APU. Used to identify the APU that produced the report.
APU serial number	Summary	N/A	Serial number of the APU. Used to identify the APU that produced the report.
Report date and time	Summary	Date	Date number representing date and time when the APU received the start command.
Cumulative operating hours	Summary	Hours	Cumulative operating hours since the APU was new. A measure of cumulative usage.
Cumulative start cycles	Summary	Cycles	Cumulative start-stop cycles since the APU was new. A measure of cumulative usage.
BIT results	Summary	N/A	PASS/FAIL result of each BIT for the period spanning startup to shutdown. A FAIL result effectively latches for purposes of the recorded result.
TAT	Summary	Degrees C	Total air temperature, as measured by other sensors on the aircraft.
Altitude	Autoshutdown	Feet	Altitude above sea level at the time of the autoshutdown decision.
APU speed	Autoshutdown	% full speed	The shaft speed of the APU at the time of the autoshutdown decision.
Exhaust gas temp	Autoshutdown	Degrees C	The temperature of the APU exhaust gases at the time of the autoshutdown decision.
Inlet temp, MES	Autoshutdown	Degrees C	Inlet air temperature during MES at the time of the autoshutdown decision.
Inlet pressure	Autoshutdown	psia	Inlet air pressure at the time of the autoshutdown decision.
IGV position	Autoshutdown	Degrees C	Angle of the IGVs of the load compressor at the time of the autoshutdown decision.
LC discharge pressure	Autoshutdown	psia	Load compressor discharge pressure at the time of the autoshutdown decision.
Fault code	Autoshutdown	N/A	Numeric fault code indicating why the APU controller switched to autoshutdown.

Table 68. Fields in an APU removal report

Field	Description
APU model	APU model name. Used to identify the APU for the report
APU serial number	APU serial number. Used to identify the APU for the report
A/C model	The model and variant of the aircraft from which the APU was removed (e.g., Airbus A320-232)
A/C serial number	The serial number of the aircraft from which the APU was removed
A/C registration	The registration code of the aircraft from which the APU was removed
Operator organization	The operator of the aircraft (e.g., the name of the airline)
Facility	The facility where the APU removal was performed
Removal date and time	The date and time the APU was removed
Symptom(s) /reason(s)	A fixed-field containing maintainer's reason(s) for removing the APU, such as one or more symptoms of malfunction or scheduled removal. It is particularly important for symptoms that are not automatically measured and included in the APU summary, startup, MES, or shutdown reports. An example is provided below in the form of check boxes
Symptom comment	A free-text field for the maintainer to provide any further details on the reason(s) why the APU was removed
Inducing cause	A fixed-field for reporting any external causes that caused a problem in the APU. Examples may include: • None or N/A • Aircraft system • Assembly error • Foreign object damage • Handling damage • Maintenance damage • Maintenance error • Shop error

Table 69. Fields in the APU repair shop report

Field	Description
APU model	APU model name. Used to identify the APU for the report.
APU serial number	APU serial number. Used to identify the APU for the report.
Repair facility	Identifies the repair facility (e.g., Honeywell Repair and Overhaul, Phoenix, AZ).
Arrival date and time	The date and time that the APU arrived at the repair facility.
TSN	Cumulative operating hours since the APU was new.
TSR	Cumulative operating hours since the APU was last repaired.
TSO	Cumulative operating hours since the APU was last overhauled.
TSI	Cumulative operating hours since the APU was last inspected.
CSN	Cumulative startup-shutdown cycles since the APU was new.
CSR	Cumulative startup-shutdown cycles since the APU was last repaired.
CSO	Cumulative startup-shutdown cycles since the APU was last overhauled.
CSI	Cumulative startup-shutdown cycles since the APU was last inspected.
Findings summary	A free-text field that summarizes the health issues found in the APU, particularly issues related to the symptoms reported by the operator that led to APU removal; also, any other health issues discovered independently by the repair facility. The findings summary should span initial observations made by receiving, borescope inspections, and teardown findings.
Detailed findings by part	Detailed good and bad findings for each subassembly/part of the APU, broken out by subassembly/part. The detailed findings should contain computer-readable information that captures the important aspects of the findings, supplemented by human-readable information. The details can include: • The part name/number. • A fixed-field check list of potential health issues and evidence of issues, with those that are actually present checked off. • A numeric or enumerated quantification of each health issue or symptom, if applicable. • A free-text description of the health issues and evidence. • A photograph of the part showing any health issues and symptoms. The subassemblies of the APU use on this project are listed after this table.
Detailed actions by part	Detailed listing of the actions performed, broken out by subassembly/part on the APU. The actions could take the form of a check list that includes whether the part was used as-is, reworked, or replaced. The subassemblies of the APU are listed after this table.

Table 70. Information recorded for on-wing APU actions

Field	Description
APU model	APU model name. Used to identify the APU for the report.
APU serial number	APU serial number. Used to identify the APU for the report.
A/C model	The model and variant of the aircraft on which the APU was repaired (e.g., Airbus A320-232).
A/C serial number	The serial number of the aircraft on which the APU was repaired.
A/C registration	The registration code of the aircraft on which the APU was repaired.
Operator organization	The operator of the aircraft (e.g., the name of the airline).
Facility	The facility where the APU repair was performed.
Repair date and time	The date and time the APU was repaired.
Symptom(s)/reason(s)	A fixed-field containing maintainer's reasons for repairing the APU.
Symptom comment	A free-text field for the maintainer to provide any further details on the reasons why the APU was repaired.
Hours	APU hours.
Cycles	APU cycles.
Parts removed	List of parts/assemblies removed.
Part serial numbers	Removed parts (LRU) serial numbers.

Figure 63 is an example set of check boxes for completing the "Symptoms(s)/reason(s)" field in table 68. These are in use by Honeywell. Note that the symptoms/reasons are not necessarily a mutually exclusive set. Also note that the example provides a simple alphabetical listing; an alternative could be a hierarchical representation.

□ A/C LEASE RETURN
□ AUTOSHUTDOWN
□ AUTOSHUTDOWN - HIGH EGT
□ AUTOSHUTDOWN -HIGH OIL TEMP
□ AUTOSHUTDOWN - LOW OIL PRESSURE
□ AUTOSHUTDOWN - ON SPEED
□ AUTOSHUTDOWN - ON START
□ CONVENIENCE
□ DAMAGED
□ EGT FLUCTUATES
□ ENGINEERING REQUEST
□ FAILED BORESCOPE INSPECTION
□ FAILED TEST/INSPECTION
□ FIRE WARNING
□ FIRE, FLAMES FROM TAILPIPE
□ FLUCTUATING DUCT PRESS
□ FLUCTUATING SPEED
□ FUEL LEAK
□ HIGH EGT
□ HIGH OIL CONSUMPTION
□ HOT SECTION INSPECTION
□ HOT START
□ HUNG START
□ INOPERATIVE
□ INTERNAL FAILURE
□ LOAN RETURN
□ LOW AIR OUTPUT
□ LOW DUCT PRESSURE
□ LOW FUEL FLOW
□ LOW OIL PRESSURE
□ LOW OIL QUANTITY
□ METAL IN OIL/FILTER/CHIP DETECTOR
□ MODIFICATION/SERVICE BULLETIN
□ NO BLEED AIR
□ NO START
□ NOISY
□ ODOR IN CABIN
□ OIL IN TAILPIPE
□ OIL LEAK
□ OIL TEMPERATURE HIGH
□ OTHER (DESCRIBE IN REMARKS)
□ OVERHAUL
□ OVERTEMP
□ POOR PERFORMANCE
□ PRECAUTIONARY REMOVAL
□ PRECAUTIONARY TROUBLESHOOTING
□ PREDICTIVE TREND MONITORING
□ SCHEDULED MAINTENANCE
□ SEIZED
□ SLOW START
□ SMOKE FROM EXHAUST
□ SMOKE IN CABIN
□ SPARKS OUT TAILPIPE
□ STALLS
□ SURGING
□ TIME LIMIT REACHED
□ VIBRATION
□ WON'T CARRY LOAD (ELECTRICAL)

Figure 63. Example set of check boxes for completing the "Symptoms/Reasons" field

The following is the subassembly list for the APU used by Honeywell on this project, as used in the "Detailed Findings by Part" and "Detailed Findings by Part" fields in table 69. The parts within each subassembly are omitted for brevity. Subassemblies and parts would depend on the APU model:

- Cooling fan
- Engine assembly
- Power section assembly
- Gearbox assembly
- Actuator, IGV
- Surge control valve
- Cooling valve
- Thermocouple
- Ignition lead
- Low oil press switch
- Oil level transmitter switch
- Oil pump assembly
- Indicator
- Fuel control assembly

- Nozzle and shroud assembly
- Starter motor assembly
- Wiring harness
- Pressure relief valve assembly
- Shutoff valve
- Load compressor

9.1.3 Valves

Despite their varied designs, the basic principles of a pneumatic regulating valve remain the same. Consequently, despite the variety of distinct failure modes that can cause internal leakages, the net effect is the same: loss of regulation. To develop a cost-effective SPMS solution, the focus needs to be on this functional trending without adding the complexity of detecting or isolating specific failure modes or faulty valve parts.

Capturing 1) downstream manifold pressure, P_4, 2) inlet pressure, P_1, and 3) solenoid current provides sufficient information for generating a trendable HI for pneumatic problems. Valve cumulative flight time can be readily derived from valve on/off times and flight times. The sensor interface is summarized in table62 along with the data collection trigger.

The CI trends on filtered CI are calculated using data collected using the specifications in table 62. The CI trend is then associated with the appropriate repair records to SPMS trends (see table 71).

Table 71. Remove and replace record gathered per valve removal

Repair information	Description
Serial number	-
Model number	-
Reference designator	-
Aircraft tail	-
Valve identifier position	-
Reason for maintenance/complaints/gripes	-
Findings	Repair findings
Actions	Actions performed
Did this work?	Yes/No, additional text inputs
Installation date	-
Repair date	-
Time since repair	-
Cycle since repair	-

9.2 STANDARD RECOMMENDATIONS

Honeywell is a long-time member of the SAE HM-1 committee and actively contributes toward the development of various recommended practices. The charter of HM-1 is: “The SAE HM-1 Integrated Vehicle Health Management committee serves as a forum to gather, develop, record and publish expert information in the discipline of aerospace Integrated Vehicle Health Management.”

The list of standards development/revision activities within SAE HM-1 follows in table 72.

Table 72. SAE HM-1 documents

Project	Title
AIR6212	The Potential Usage of Health Monitoring Techniques and Technologies on Aircraft Operations During or After Active Volcanic Events
AIR6334	A Power Usage Metric for Rotorcraft Power Train Transmissions
ARD6888	Functional Specification of Miniature Connectors for Health Monitoring Purposes
ARP6268	Design & Online Communication Standards for Health-Ready Components
ARP6290	Guidelines for the Development of Architectures for Integrated Vehicle Health Management Systems
ARP6407	Integrated Vehicle Health Management Design Guidelines
ARP6803	HM-1 Cornerstone Document
ARP6883	Guidelines for Writing IVHM Requirements for Aerospace Systems
ARP6887	Verification & Validation of Integrated Vehicle Health Management Systems and Software
AS5391A	Health and Usage Monitoring System Accelerometer Interface Specification

An inspection of table 72 shows that there is no existing standard or activity that covers recording of on-wing and shop maintenance activity. The authors’ experience with SPMS has shown that knowledge of the maintenance activities is crucial for building and refining a successful diagnostics and prognostics system. Knowledge of on-wing activities sometimes explains otherwise anomalous changes in trends. Knowledge of shop findings and on-wing repairs provide ground truth regarding what faults are present, which allows an SPMS to reinforce or correct its inferences through learning. Accordingly, to fill this gap, the authors recommend development of a new standard within SAE HM-1, to be named “Recording of On-Wing and Shop Maintenance to Support Diagnostics and Prognostics.” The work done on engines, APUs, and valves under SPMS will serve as example use cases for the standard.

9.3 TECHNOLOGY TRANSITION PLAN AND ROADMAPS

The current SPMS improvements for APU are being explored in detail by Honeywell for enhancements for APU prognostics trend monitoring offerings. As a result of this program and other internal research efforts, Honeywell Advanced Technology is working internally on programs to enhance PTMD with field and shop data findings.

The authors propose a two-pronged approach for further development and sustainment both within Honeywell and in the greater IVHM community:

1. Pursue further development of prognostics with better maintenance data capture.
2. Pursue the data collection standardization through a new standard within SAE HM-1.

10. CONCLUSIONS AND RECOMMENDATIONS

This section captures SPMS conclusions and recommendations for three selected systems:the engines, APU, and valves. The engine, APU, and valve SPMS can improve significantly by capturing the complete lifecycle data, as recommended in section 9. Additional improvements in performance can be achieved by exploring better correction factors for ambient conditions for engine and APU SPMS. Finally, the use of Bayesian network learning is recommended for trend improvements and time-series anomaly detection using complete lifecycle data to fully realize the SPMS prognostics capability.

10.1 ENGINES

The authors ran the startup monitors through 180,972 flight and ground run files. The engine startup monitors detect no start, slow start, hung start, high EGT, multi-start, hot start, and unusual dwell. Using engine startup monitor with nominal threshold for known cases of FCU replacement resulted in an alarm in 7% of the flights. This was reduced through threshold tune-up. However, clear trends consistent with alarm were observed without any associated repair records. Therefore, the engine SPMS development suffered because of the lack of adequate total lifecycle data, including on-wing maintenance actions and repairs. The data capture recommendations have been summarized in table 61. The lack of on-wing repair records impacts the SPMS effectiveness. The classification effectiveness for FCU failure classifiers increased for engine position specific classifiers.

10.2 APU

Hot-section health deterioration reduces efficiency. More fuel and poorer energy extraction lead to hotter EGT. It was observed that some lifecycles have sudden degradations. Step-drop suggests sudden damage in the APU.

The APU SPMS development also suffered for the lack of adequate total lifecycle data, including on-wing maintenance actions and repairs. The date of change of nozzle is not obvious from the trend nor recorded accurately. APU data do not show nozzle change as a discrete jump in the EGT trend. The effects of APU nozzle change are weak and discernable from noise factors. Therefore, the life on-wing hypothesis extension using nozzle replacement cannot be proven by the APU trend data.

The EGT trend is accurate in predicting the APU RUL. Bayesian multivariate inference, using Bayesian networks and particle filters, was tried for improving the prognostics trend. The trend did not improve trend accuracy significantly; this could be because of the effect of inlet air temperature, which has to be normalized. Inlet temperature in some cases was found to show unusual bimodal noise characteristics. This unexplained nonlinear behavior contributes to the performance mismatch.

10.3 VALVES

The authors started the project focusing on many valves in the ECS system. However, the ECS field data was not available. The authors explored the use of ECS models to generate data for SPMS development. The complete ECS systems model is complex, slow, and does not have valve degradation built in. Therefore, the poppet valves were selected for which the degradation models were already available. The pneumatic regulating valve, such as the poppet valve, fails because of internal leakages. The net effect is the same: loss of regulation. The results show that trendable HIs can be developed using downstream manifold pressure and inlet pressure.

10.4 FUTURE RESEARCH

There are three main barriers for the application of SPMS methodology to realize the benefits of prognostics in aircraft operations, maintenance, and reliability: a lack of lifecycle data, a lack of in-depth knowledge of the sensor noise characteristics under anomalous conditions, and a lack of machine learning methods for anomaly detection in unequal length time series. Additional targeted research is needed to provide prognostics coverage for uncovered domains such as avionics, new electrical systems, and composite structures. Future SPMS research needs to be focused in these areas:

- Lifecycle data: Capture the complete lifecycle data, including sensor data, on-wing maintenance, and shop findings for SPMS target systems (see section 9).
- Expert knowledge: Obtain expert knowledge to explain the non-linear, correlated noise patterns and better normalize sensor data from different operating conditions.
- Anomaly precursors identification for unequal time-series data: Though the machine learning literature is full of approaches to deal with time series data, these approaches fall short when comparing and identifying precursors for unequal length time series. More research needs to be completed to apply and improve the methods in literature to the SPMS problems in aviation.
- Extend SPMS to cover new systems.

11. REFERENCES

1. Byington, C.S., Watson, M., Kalgren, P., Safa-Bakhsh, R., and Forsyth, G.F. (Ed.). (2003). Proceedings from the Third International Conference on Health and Usage Monitoring - HUMS2003: *Metrics Evaluation and Tool Development for Health and Usage Monitoring System Technology*. Melbourne, Australia.

2. Kurtoglu, T., Mengshoel, O., and Pol, S. (2008). Proceedings from the International Conference on Prognostics and Health Management: *A Framework for Systematic Benchmarking of Monitoring and Diagnostic Systems*. Denver.
3. Saxena, A. et al. (2008). Proceedings from the International Conference on Prognostics and Health Management: *Metrics for Evaluating Performance of Prognostic Techniques*. Denver.
4. The American Institute of Aeronautics and Astronautics (2010). *Architectures for Integrated Vehicle Health Management* (AIAA-2010-3433). Felke, T., Hadden, G.D., Miller, D., and Mylaraswamy, D.
5. The American Society of Mechanical Engineers (2005). *Fault Diagnosis and Prognosis for Fuel Supply System in Gas Turbine Engines*. Kim, K., Uluyol, O., and Ball, C.
6. The International Society of Airbreathing Engines (2011). *Low-cost embedded scouts for engine health monitoring*. Mylaraswamy, D., Parthasarathy, G., Kim, K., and Uluyol, O.