An Integrated Architecture for Aircraft Engine Performance Monitoring and Fault Diagnostics: Engine Test Results

Aidan W. Rinehart[1]
Vantage Partners LLC, Brook Park, OH, 44142

Donald L. Simon[2]
NASA Glenn Research Center, Cleveland, OH, 44135

This paper presents a model-based architecture for performance trend monitoring and gas path fault diagnostics designed for analyzing streaming transient aircraft engine measurement data. The technique analyzes residuals between sensed engine outputs and model predicted outputs for fault detection and isolation purposes. Diagnostic results from the application of the approach to test data acquired from an aircraft turbofan engine are presented. The approach is found to avoid false alarms when presented nominal fault-free data. Additionally, the approach is found to successfully detect and isolate gas path seeded-faults under steady-state operating scenarios although some fault misclassifications are noted during engine transients. Recommendations for follow-on maturation and evaluation of the technique are also presented.

Nomenclature

C-MAPSS40k	=	commercial modular aero-propulsion system simulation 40k lbf thrust
H	=	fault influence matrix
HPC	=	high pressure compressor
HPT	=	high pressure turbine
LPC	=	low pressure compressor
LPT	=	low pressure turbine
m	=	fault magnitude
$\hat{m}$	=	estimated fault magnitude
PBM	=	performance baseline model
RTSTM	=	real time self-tuning model
R	=	sensor measurement covariance matrix
u	=	command vector
VIPR	=	vehicle integrated propulsion research
$WSSEE$	=	weighted sum of squared estimation errors
$WSSR$	=	weighted sum of squared residuals
y	=	sensor vector
$\tilde{y}$	=	sensor residual vector
$\hat{y}$	=	estimated sensor vector
$\hat{\tilde{y}}$	=	estimated sensor residual vector

I. Introduction

CONVENTIONAL aircraft engine gas path diagnostic approaches are designed for ground-based post-flight processing of "snapshot" measurement data collected at a limited quantity of operating points each flight.[1,2,3] However, advances in onboard processing and flight data acquisition capabilities are providing access to increased quantities of flight data and enabling new diagnostic approaches. Analyzing full-flight streaming measurement data,

[1] Aerospace Engineer II, aidan.w.rinehart@nasa.gov
[2] Research Engineer, Intelligent Control and Autonomy Branch, donald.l.simon@nasa.gov

either onboard in real-time or post-flight, can help reduce diagnostic latency and improve overall engine safety and reliability. In response to this, several research efforts conducted within the aviation community focused on the analysis of streaming engine measurement data. Merrington *et al* applied analytical redundancy methods to process aircraft gas turbine engine transient measurement data[4], and Kerr *et al*[5], Dewallef *et al*[6], and Borguet *et al*[7] have each proposed Kalman filter-based approaches for the on-line processing of aircraft engine measurement data for diagnostic purposes. However, these approaches have primarily focused on either performance estimation or fault diagnostics as opposed to performing both tasks concurrently. Additionally, the approaches that have focused on joint performance estimation and fault diagnostics have only considered turbo machinery faults, not actuator or sensor faults. Recently, NASA has developed a model-based performance trend monitoring and gas path diagnostic architecture designed to process streaming full-flight aircraft engine measurement data.[8,9,10] This architecture provides the dual functionality of estimating and trending engine performance parameters and diagnosing the occurrence of gas path system faults and isolating the faults including turbo machinery, actuator or sensor faults. The fault isolation is completed under the assumption that only one gas path system fault will be present at a given time. Simulation studies have shown that this architecture holds promise for analyzing full-flight engine data.[9] This paper will present results from the application of this architecture on engine test data, including both nominal and faulty engine operating scenarios. The remainder of this paper proceeds as follows. First, the introduction of the architecture and the various components comprising it are explained. Next, the application of the architecture for processing engine test data is discussed and the associated results are presented. This is followed by a discussion of the results, planned future work, and conclusions.

II. Architecture Description

A block diagram schematic of the NASA-developed model-based performance trend monitoring and gas path fault diagnostic architecture is shown in Figure 1.[10] It contains three main components. These include: 1) a Real-Time Self Tuning Model (RTSTM); 2) a Performance Baseline Model (PBM); and 3) a Fault Diagnostics module. Each component is further discussed in the subsections below.

A. Real-Time Self Tuning Model

The RTSTM is a piecewise linear model that is self-tuned by a Kalman tracking-filter. The Kalman filter estimates model state variables and a set of model tuning parameters reflecting performance deterioration within the engine. Through these Kalman filter estimates, the RTSTM accounts for changes in engine dynamics and changes from the gradual performance deterioration that aircraft gas turbine engines naturally experience as they wear over time from standard use. The RTSTM was designed to provide accurate continuous real-time estimates of engine performance parameters over the lifetime of the engine.

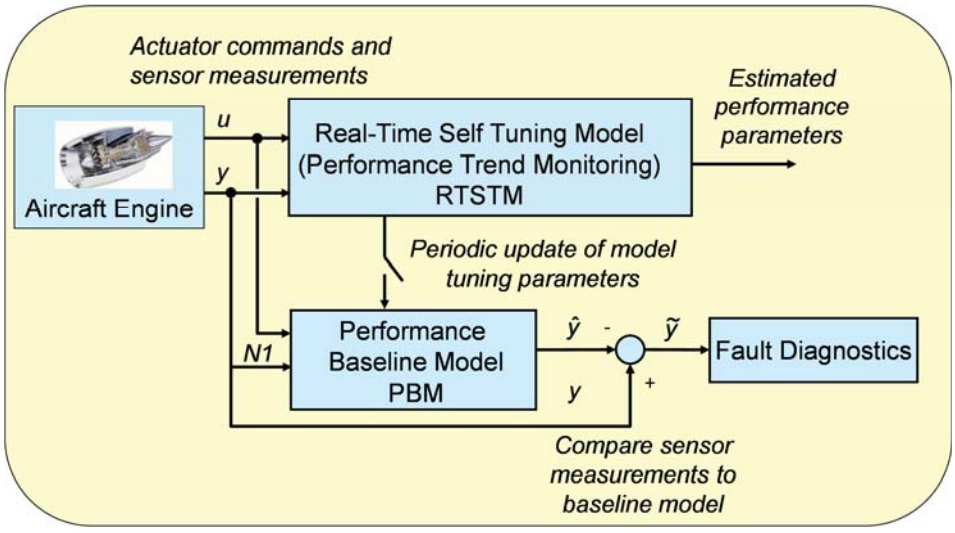

Figure 1. Performance Trend Monitoring and Gas Path Fault Diagnostic Architecture

B. Performance Baseline Model

The PBM is built from the same piecewise linear model as used in construction of the RTSTM. However, in the case of the PBM, no Kalman filter tracking input is applied. The PBM continuously receives actuator commands and sensed fan speed measurements as inputs. The PBM also receives tuning parameter updates periodically from the RTSTM. These periodic tuner updates enable the PBM to adapt to gradual long-term engine performance deterioration. However, unlike the RTSTM, which continuously receives updates via a Kalman filter, the periodic tuner updates provided to the PBM prevents engine performance changes, potentially due to a fault, from being instantaneously absorbed into the model. In this form, the PBM serves as a baseline for describing recent engine performance because it is updated at a slower rate than the RTSTM. By comparing sensed engine outputs to PBM estimated outputs, an abrupt performance change, which may be indicative of a fault, can be detected.

C. Fault Diagnostics

The overall fault diagnostics logic is a two-step process consisting of fault detection and then fault isolation. Fault detection is performed by monitoring a vector of measurement residuals, $\tilde{y}$, between sensed engine outputs, y, and the PBM estimated outputs, $\hat{y}$, defined by Eq. (1).

$$\tilde{y} = y - \hat{y} \tag{1}$$

If a gas path fault impacting engine performance occurs, the sensed engine and PBM outputs are expected to diverge resulting in an increase in the residuals. Fault detection is accomplished by calculating a weighted sum of squared residuals (*WSSR*) value based on the measurement residual vector and comparing this calculated value against a threshold value. The *WSSR* is calculated as

$$WSSR = \tilde{y}^T R^{-1} \tilde{y} \tag{2}$$

where R is the sensor measurement covariance matrix. Each time step, a new *WSSR* value is calculated based on the most recent sample of measurement data. Then the *WSSR* value is compared against a defined anomaly detection threshold. Logic for recognizing a fault initiates when the *WSSR* value exceeds the threshold. The fault recognition logic requires the high *WSSR* value to persist for a period to filter out anomalous spikes. After this persistence metric is met, then fault isolation logic is engaged to identify the root cause for the fault. For the first portion of a test, the *WSSR* calculation is not compared against the threshold. This is a *WSSR* blackout period to prevent engine start up dynamics from triggering a false alarm. The length of the engine start up dynamics defines the duration of the *WSSR* blackout period.

Fault isolation is accomplished through a linear least-squares estimation approach implemented with the assumption that at a given time only one fault may be present. This assumption is reasonable based on the high reliability of gas turbine engines and the low probability that any fault will be present. The isolation process compares the residual vector against a matrix of coefficients that were pre-calculated at different points throughout the engine operating envelope. This pre-calculated matrix is referred to as the Fault Influence Matrix, H, and it is a collection of single matrices assembled to allow interpolation calculations. The Fault Influence Matrix is populated with data collected while running the PBM under nominal and faulty operating scenarios. The resulting residuals between the nominal and the faulty PBM outputs, normalized by the magnitude of the fault, m, form a column of the H matrix. Each of the H matrix columns pertains to a specific fault at a particular operating condition. Equation (3) defines the calculation of a single element of the fault influence matrix, $H_{i,j}$, where i is the index of engine sensor in the residual vector and j is the index of fault type.

$$H_{i,j} = \frac{\tilde{y}_i}{m_j} \tag{3}$$

Upon fault detection, the observed measurement residual vector, $\tilde{y}$, is further processed to isolate the most likely root cause for the fault based on the type of faults contained in the H matrix. To do so an estimated magnitude for each possible fault is first calculated as follows

$$\hat{m}_j = \left(H_j^T R^{-1} H_j\right)^{-1} H_j^T R^{-1} \tilde{y} \tag{4}$$

These estimated fault magnitudes are used to produce an estimated vector of sensor residuals for each fault type as shown in Eq. (5)

$$\hat{\tilde{y}}_j = H_j \hat{m}_j \tag{5}$$

Finally, fault isolation is performed by identifying the fault type that produces the estimated sensor residual vector that most closely approximates the observed measurement residual vector in a weighted least squares sense. A vector of values pertaining to the weighted sum of squared estimation errors (*WSSEE*) is calculated using Eq. (6) for each potential fault type. The fault type yielding the smallest *WSSEE*, indicating the closest match to the provided fault types contained in *H*, is classified as the fault.

$$WSSEE_j = \left(\tilde{y} - \hat{\tilde{y}}_j\right)^T R^{-1} \left(\tilde{y} - \hat{\tilde{y}}_j\right) \tag{6}$$

III. Application

This section will cover the application of this performance trend monitoring and gas path fault diagnostic architecture using data obtained from the NASA Vehicle Integrated Propulsion Research (VIPR) engine tests.[12] The VIPR program is a series of ground-based engine tests conducted to mature aircraft engine health management technologies. To date, two VIPR tests (denoted as VIPR I and VIPR II) have been conducted, with a third test, VIPR III, scheduled to occur in 2015. These tests are ongoing at the NASA Armstrong Flight Research Center, formerly the Dryden Flight Research Center/Edwards Air Force Base, on a C-17 aircraft equipped with Pratt & Whitney F117 turbofan engines. The VIPR tests include baseline runs where the test engine is operating normally without faults as well as runs with faults applied that are non-damaging to the engines. These fault cases are created by operating the engine with mis-scheduled 14th stage bleed valve and station 2.5 bleed valve actuators. These faults provide a good test case because operating the actuators differently than their normal scheduled operation will result in observable gas path variations.

The diagnostic architecture applied for analyzing the available VIPR data was implemented using a NASA-developed 40,000 lbf thrust class engine simulation called Commercial Modular Aero-Propulsion System Simulation 40k (C-MAPSS40k).[11] Exercising the C-MAPSS40k simulation produced a data set that was reduced to the piecewise linear model used to construct the RTSTM and the PBM shown in Fig. 1. The models generated from C-MAPSS40k displayed notable differences in performance compared to the nominal steady-state performance data recorded from the normal operation of the F117 turbofan engines used in the VIPR tests. To account for these model to engine differences a method was developed and applied to modify, or re-trim, the piecewise linear models based on nominal steady-state performance data obtained in the VIPR tests.[10] This enabled better model to engine performance agreement and helped to improve overall diagnostic results. Table 1 contains the vector of eight gas path sensor measurements and Table 2 contains the four actuator commands. The two sensors in Table 1 denoted by an asterisk, T25 and P25, were only included during VIPR II testing.

The asterisk next to the BLD14 symbol in Table 2 indicates the 14th stage bleed valve was not one of the actuator commands included in the C-MAPSS40k model used to generate the piecewise linear models. As a result, the standard model-based method for calculating the column of the Fault Influence Matrix corresponding to the 14th stage bleed valve fault is not possible. Instead, 14th stage bleed valve measurement residual vectors were calculated based on actual VIPR data collected during the 14th stage bleed valve fault

Table 1. Gas Path Sensor Measurements

Symbol	Description
N1	Fan speed
N2	Core speed
T25*	Low pressure compressor exit temperature
T35	High pressure compressor exit temperature
T5	Low pressure turbine exit temperature
P25*	Low pressure compressor exit pressure
Ps3	High pressure compressor exit pressure
P5	Low pressure turbine exit pressure

*- Research sensors unique to VIPR II engine data

Table 2. Actuator Commands

Symbol	Description
Wf	Fuel flow
VSV	Variable stator vane
BLD25	Station 2.5 bleed valve
BLD14*	14th stage bleed valve

*- Valve not modeled in C-MAPSS40k

testing. Data collected at given steady-state power setting conditions were averaged and referenced against PBM estimated outputs to provide residuals analogous to those simulated for other fault types. This method allowed creation of Fault Influence Matrices that included the 14th stage bleed valve fault.

Prior to analyzing VIPR data with the performance trend monitoring and fault diagnostic architecture, Eq. (4) was used to construct a Fault Influence Matrices, and the *WSSR* anomaly detection thresholds were defined. The Fault Influence Matrices were constructed assuming eight possible single fault types, as identified in Table 3. While the only faults evaluated during the VIPR testing were a station 2.5 bleed fault and a 14th stage bleed valve fault, the six additional fault types were included as possible fault scenarios to provide a more representative and challenging fault isolation problem. The anomaly detection threshold and persistency values were established through a manual process of running non-faulted VIPR data through the architecture and selecting threshold and persistency values that did not produce false alarms. In conducting this process, the *WSSR* signal and the measurement residual vector, $\tilde{y}$, were filtered using a median filter to help reduce signal noise, and removed outliers in an effort to improve overall diagnostic results. The *WSSR* anomaly detection threshold and persistency values applied for analyzing the VIPR I and VIPR II data are not identical. This variation is because the test engine used in the VIPR II test contained two additional sensors, P25 and T25, that the VIPR I test engine did not. The increase in the number of sensors produced a larger *WSSR* while no fault was present. This required a slight increase in the anomaly detection threshold for VIPR II tests.

Table 3. Fault Influence Matrix Fault Types

Fault Index	Fault Types
1	Fan
2	Low Pressure Compressor
3	High Pressure Compressor
4	High Pressure Turbine
5	Low Pressure Turbine
6	Station 2.5 Bleed Valve
7	Variable Stator Vane
8	14th Stage Bleed Valve

IV. Results

The following is a sample of the results obtained from processing VIPR I and VIPR II data through the performance trend monitoring and fault diagnostic architecture. The illustration in Fig. 2 describes a typical engine test. As illustrated in Fig. 2, the tests were segmented into two parts, a steady-state portion followed by a transient power sweep portion. The steady-state portion involved stepping up and stepping down the engine fan speed. The steady-state testing held the engine fan speed at a constant value for a short period, and then the fan speed was changed to a new level. The transient portion included two transient operations. The first transient slowly ramped the fan speed up and down; whereas, the second transient was faster. All of the test cases presented below, except for the VIPR II baseline test, which did not include the transient power sweeps, followed the test format as depicted in Fig. 2. The test cases containing faults have the fault inserted for the entire test unless otherwise noted.

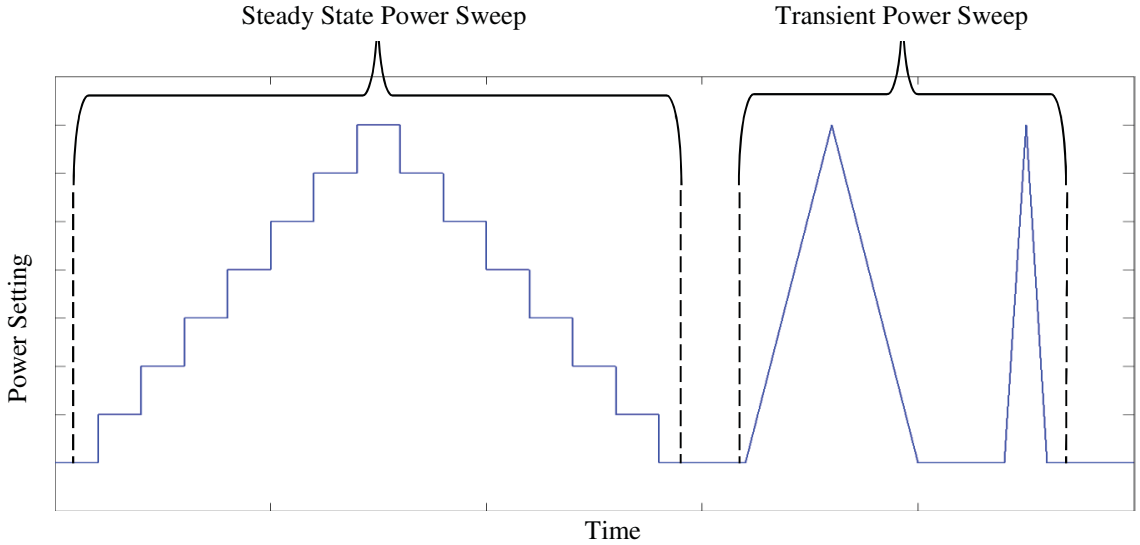

Figure 2. Representative VIPR Event Test Sequence

Results from the VIPR I and VIPR II baseline runs where the engine was operating nominally without any faults are shown in Figs. 3 and 4 respectively. In these figures, and all remaining figures of the document, parameter names and engineering units have been omitted due to the proprietary nature of the data. The top subplots of each figure show an unspecified gas path parameter plotted against time. For these subplots, the blue line represents the sensed engine measurement, the red line represents the PBM produced estimate, and the green line represents the RTSTM estimate. For both Figs. 3 and 4 the green line is not visible due to the close agreement between the RTSTM estimate and the sensed engine measurement. The sensed and PBM estimated values also match very well. No faults were present during the tests conducted to acquire the data displayed in Figs. 3 and 4, so the close agreement in the sensed and PBM estimated values was expected and desired. The middle subplots displays the *WSSR* values plotted against time in blue and the selected anomaly detection threshold are represented with a dash-red trace. For these two cases, the *WSSR* signal remained below the established anomaly detection threshold throughout the entire run. Since the *WSSR* signal never exceeded the threshold, the fault isolation logic never engaged and "No Fault" is reported over the duration of these test cases. The bottom subplots of Figs. 3 and 4 show the fault classification plotted against time, which in the baseline cases was "No Fault".

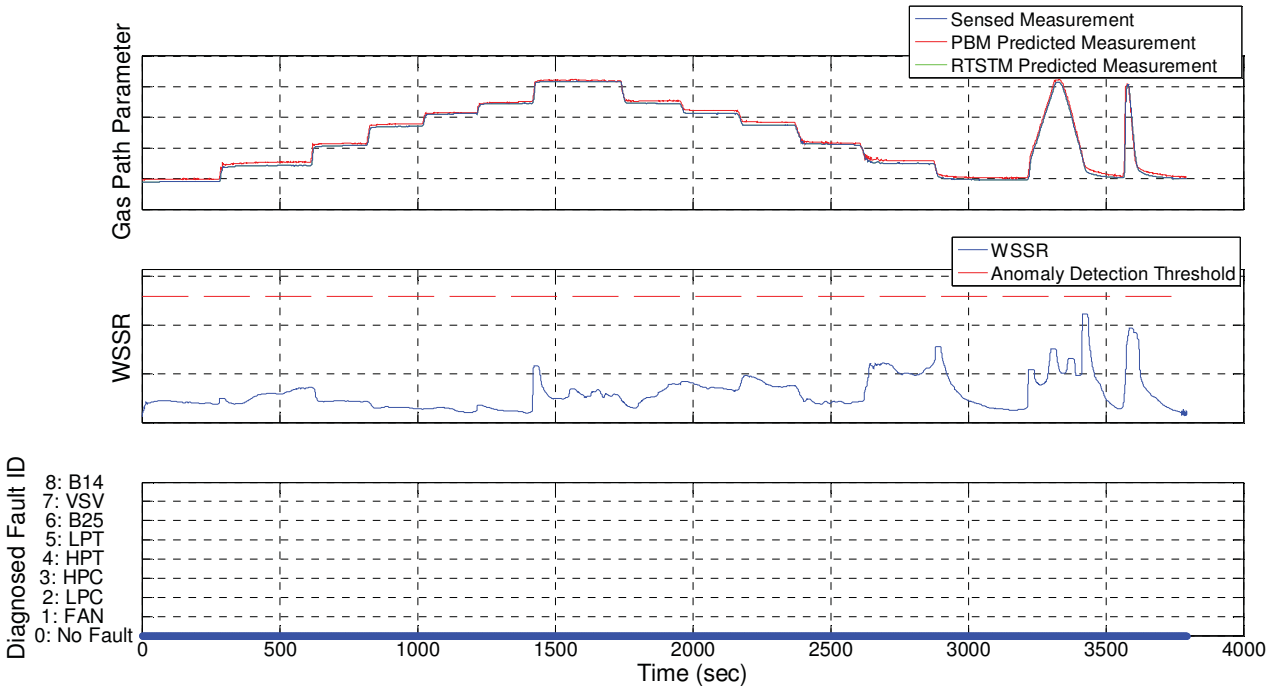

Figure 3. Baseline VIPR I Results

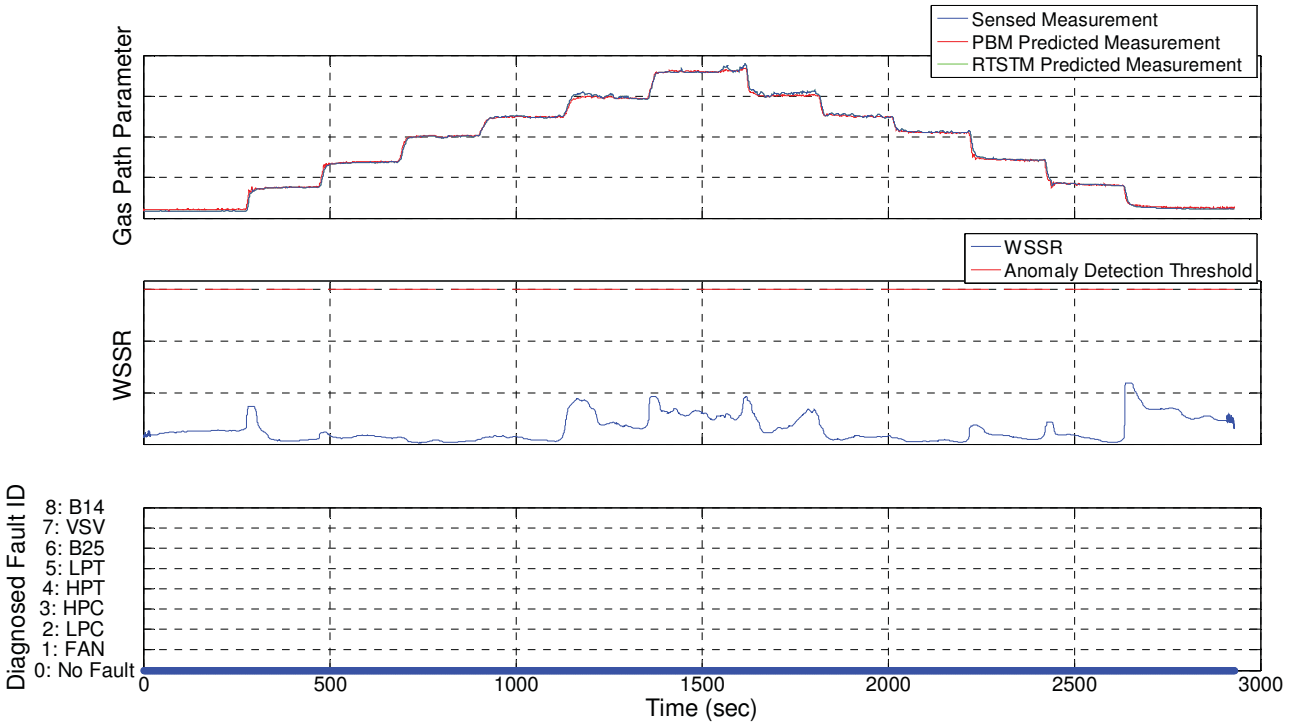

Figure 4. Baseline VIPR II Results

Figures 5 and 6 contain results from VIPR I and VIPR II test cases when the station 2.5 bleed valve was intentionally failed open for the entire test case. The gas path parameter depicted in the top subplots of Figs. 5 and 6 show very close agreement between the sensed measurement and the PBM estimate for the majority of this test. This is expected as the station 2.5 bleed valve is normally scheduled open during lower power settings and modulates closed as the engine increases in power setting. During lower power settings, the sensed engine parameter measurements and PBM produced estimates show close agreement, because the failed open valve position is the same as the normally scheduled position. Once the engine is operating at higher power settings, sensed engine measurements and PBM estimates begin to diverge and the presence of the fault becomes apparent. This is evident in the *WSSR* signal shown in the middle subplots of Figs. 5 and 6. In the middle subplots, the *WSSR* signal can be seen to increase as the engine increases to higher power settings and eventually surpassing the anomaly detection threshold. The fault isolation logic engages after the *WSSR* has exceeded the anomaly detection threshold and it remains above the threshold for a period that satisfies the persistency requirement. Comparing the middle subplots of Figs. 5 and 6, there is a noticeable difference in the duration that the *WSSR* signal exceeds the threshold. The VIPR II test in Fig. 6 shows a wider range of power settings in which an anomaly was detected compared against the VIPR I test in Fig. 5. In addition, the bottom subplot of Fig. 6 shows that faults are detected for both transients near the end of the test while in Fig. 5 a fault is only briefly detected during the second transient. The improved fault detection displayed in VIPR II data can be attributed to the additional P25 and T25 sensors added for this test. The faults that were misclassified can be attributed to dynamic modeling inaccuracies contained in the PBM. The misclassifications in Fig. 5 occurred during the second transient sweep. For the results illustrated in Fig. 6, misclassifications occurred in both the steady-state and transient portions of the test. The misclassification during the steady-state portion occurred while the engine speed was transitioning to a new power setting. Therefore, this error was not a steady-state operation misclassification. The other misclassifications in Fig. 6 were during the transient sweep portions of the test. However, both test results illustrated in Figs. 5 and 6 correctly classify the fault as a station 2.5 bleed valve fault during steady-state operation. As evident in the data from these experiments, accurate diagnostics during large rapid transient engine operation proved to be challenging for both VIPR I and VIPR II tests. The analysis shows faults were occasionally incorrectly classified during these tests; however, they were both correct while analyzing the steady-state data.

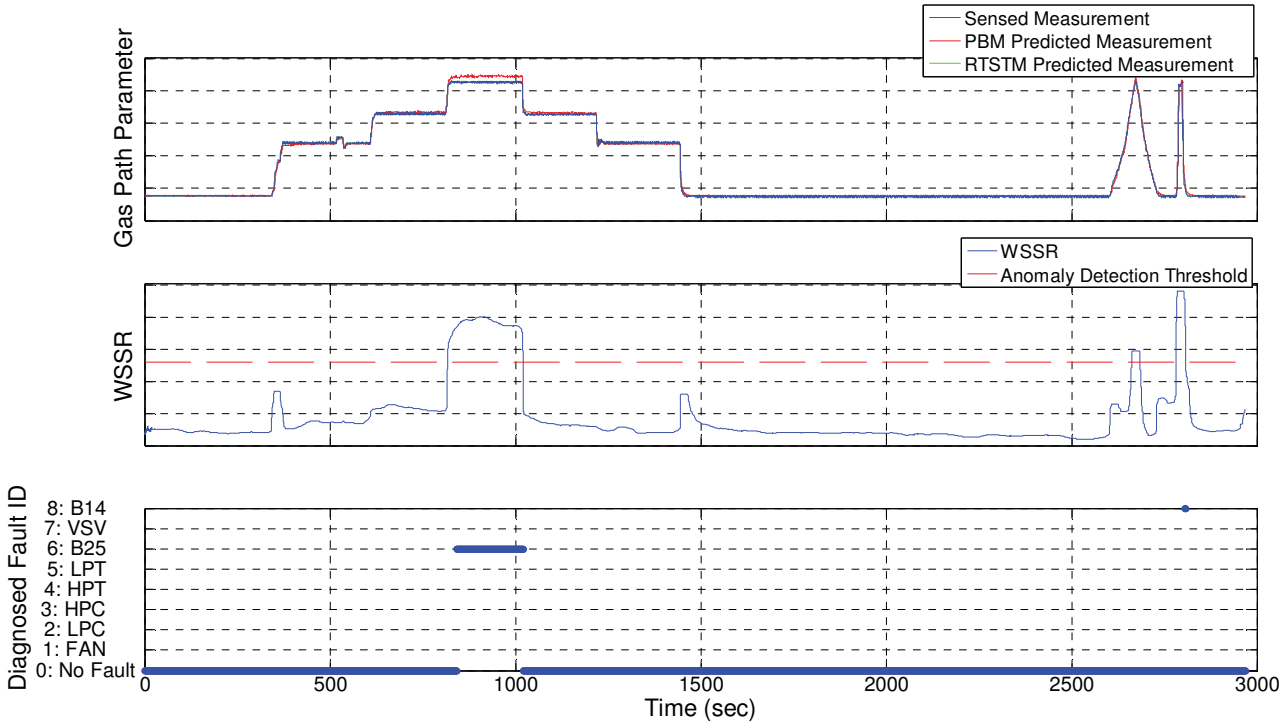

Figure 5. Station 2.5 Bleed Valve Fault VIPR I Results

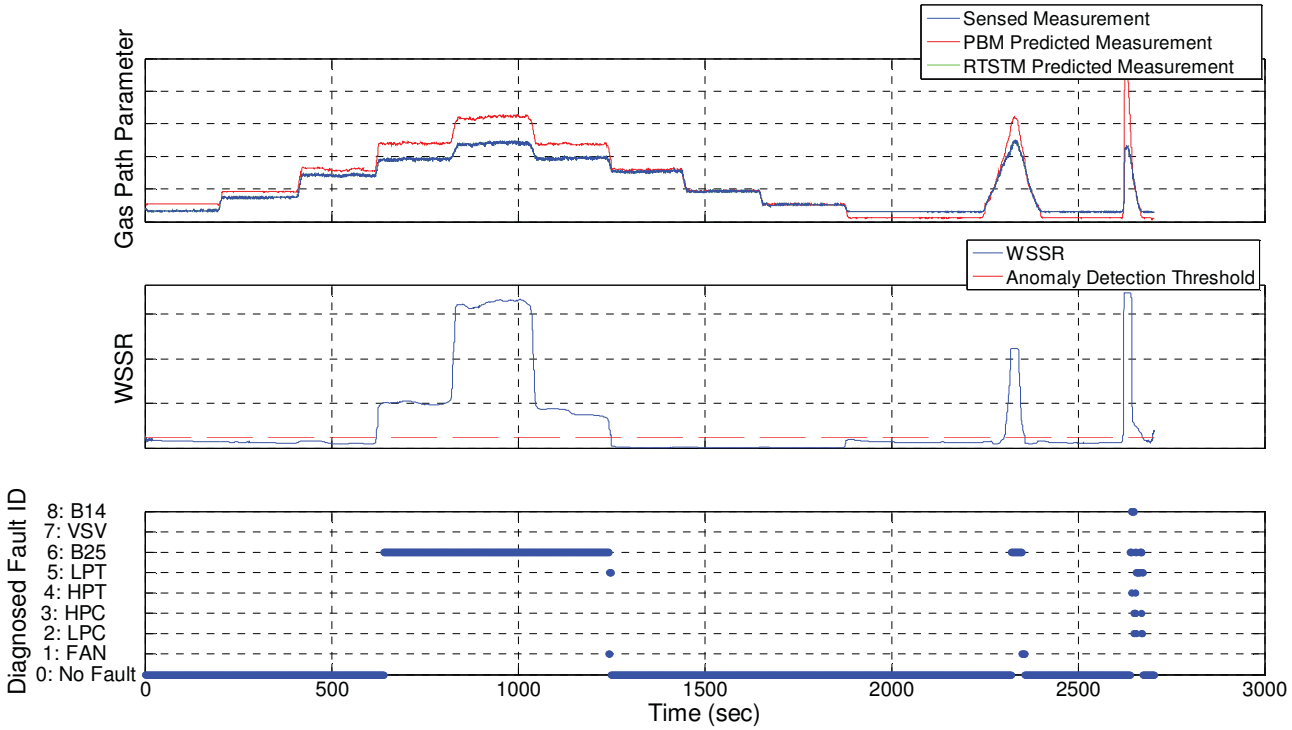

Figure 6. Station 2.5 Bleed Valve Fault VIPR II Results

Figures 7 and 8 contain VIPR I and VIPR II results for a 14th stage bleed valve fault case. The top subplots in Figs. 7 and 8 show divergence between the sensed measurement and the PBM estimate for most of the test case. In Fig. 7 the divergence is present from the beginning of the test until 2600 seconds when the fault is removed. The top subplot of Fig. 8 begins with the sensed measurement and PBM estimate in relative agreement for the first 300 seconds, then the fault was inserted. After fault insertion, the PBM estimate diverges from the sensed measurement. The middle subplots in Figs. 7 and 8 show that the calculated *WSSR* value has exceeded the anomaly threshold, which indicates a fault. Figure 7 shows the *WSSR* signal surpassing the threshold for the entire test until the fault is removed. Figure 8 illustrates that the *WSSR* signal was above the threshold once the fault was inserted around 300 seconds; however, the *WSSR* signal did drop below the threshold at 2150 seconds and at 2600 seconds. These drops in *WSSR* signal below the threshold occurred at low power settings. This point indicates that the sensed measurements and the PBM estimates show relatively close agreement, even though the fault is still present during this time. The bottom subplots of both VIPR I and II tests show accurate fault classification during most of the steady-state portions of the test. At the very beginning of the VIPR I test, "No Fault" is reported even though the middle subplot of Fig. 7 indicated an anomaly is present. This is due to the detection logic suppressing any faults for the first 40 seconds of data to insure engine start up dynamics do not trigger a false alarm. The bottom subplot of Fig. 7 shows a few instances where the fault is identified as something other than the 14th stage bleed valve fault. These are very brief instances, and they occurred while the engine was transitioning between power settings. The bottom subplot of Fig. 8 displays a few incorrect fault classifications when the *WSSR* signal dropped below the threshold and reported there was "No Fault" detected. In addition, this subplot illustrates a brief misclassification when reporting a LPC fault during an engine transient. The incorrect classifications show that there are areas for improvement with the method. However, despite these inaccuracies the method proves to be very accurate for steady-state engine operation.

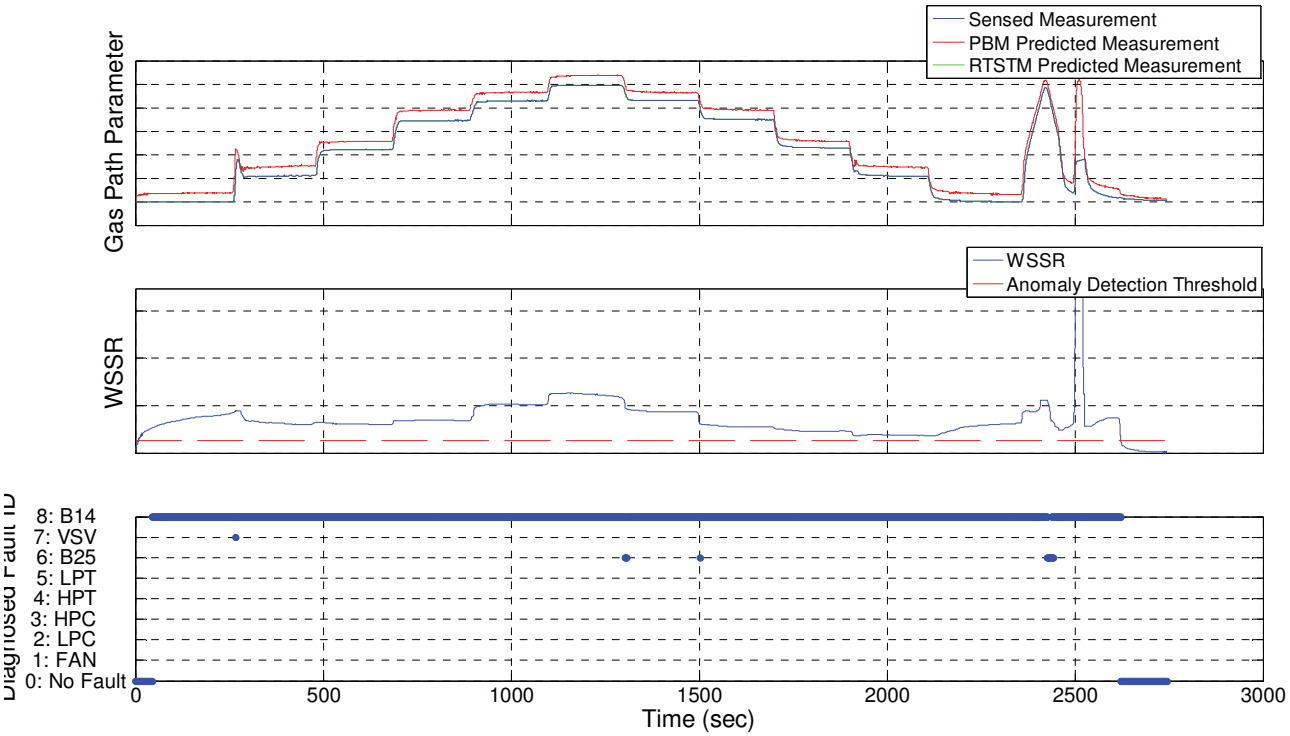

Figure 7. 14th Stage Bleed Valve Fault VIPR I Results

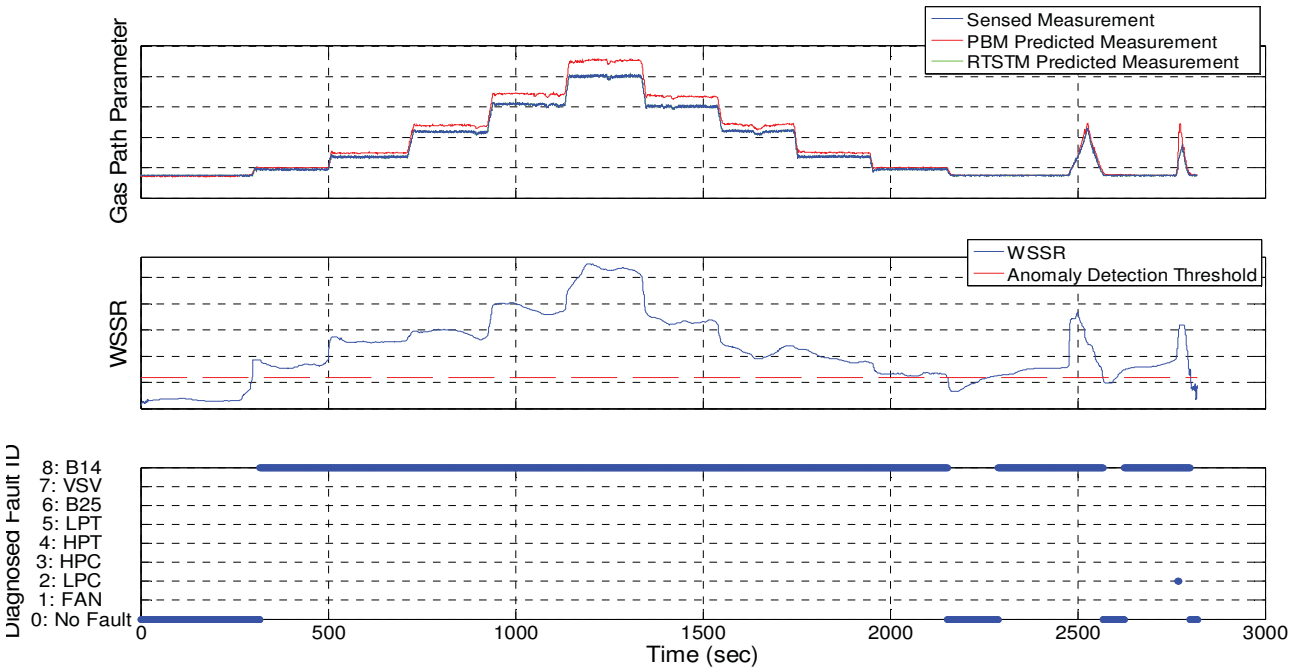

Figure 8. 14th Stage Bleed Valve Fault VIPR II Results

V. Discussion

Analysis of the VIPR engine test data marks the first time the model-based performance trend monitoring and gas path fault diagnostic architecture was applied for processing real engine data. Overall, the results are encouraging. The first step of fault diagnostics, fault detection, was shown to avoid false alarms when presented nominal data and to correctly detect anomalies when presented faulty data. Furthermore, the second step of the fault diagnostics, fault isolation, was found to correctly identify the fault type during steady-state engine operating conditions. However, this process did experience fault misclassifications during engine transients. This suggests that the PBM provides good steady-state agreement with the engine, but it exhibits some issues when working with dynamic transient behavior. More work is needed to investigate potential enhancements to help improve the PBM transient accuracy and the architecture diagnostic during transient performance. The ultimate goal would be to improve the architecture and model to a point where real-time processing of engine data could occur with confidence in the fault detection and isolation capabilities during transients. However, the first step is to implement this method as a ground based post-processing tool. In this scenario, the architecture could be easily modified to disengage during transients and only convey the fault most frequently diagnosed throughout the entire flight. The results shown above suggest that the architecture could successfully provide that type of analysis.

While the architecture is capable of combined performance trend monitoring and gas path fault diagnostics, this paper only reports on an assessment of the fault diagnostic results. The performance trend monitoring aspect of the architecture, provided by the RTSTM Kalman filter, was not able to be fully evaluated based on the available VIPR I and VIPR II data because the engines underwent little to no performance variations during these tests. However, the VIPR III test to be conducted in 2015 will intentionally degrade the engine through volcanic ash ingestion testing. As such, the VIPR III test data is expected to provide the opportunity to assess the performance deterioration trend monitoring functionality of the architecture.

VI. Conclusion

The model-based approach to gas path fault detection and isolation presented in this paper is a promising architecture for the processing of streaming engine sensor data. During steady-state operating periods, the architecture was able to avoid false alarms and was consistently able to correctly identify and classify the two bleed valve faults introduced in the VIPR test cases presented in this paper. The approach did experience fault misclassifications during engine transients. Inaccuracies between the model dynamics and the engine dynamics are believed to be the cause of these misclassifications, and further work is needed to focus on improving the model dynamic accuracy. The continuation of the VIPR test series includes an upcoming third test where engine performance will be degraded via volcanic ash ingestion testing. This test will provide data for evaluating the effectiveness of the performance trend monitoring estimation capability of the architecture.

Acknowledgments

This work was conducted under the NASA Aviation Safety Program, Vehicle Systems Safety Technologies Project.

References

[1]E-32 Aerospace Propulsion System Health Management, "A Guide to the Development of a Ground Station for Engine Condition Monitoring," *SAE Standard AIR4175 Revision A*, February 2005.

[2]Volponi, A., and Wood, B., "Engine Health Management for Aircraft Propulsion Systems," *First International Forum on Integrated System Health Engineering and Management (ISHEM) in Aerospace*, Napa CA, November 7-10, 2005.

[3]Doel, D. L., "TEMPER – A Gas Path Analysis Tool for Commercial Jet Engines," *Journal of Engineering for Gas Turbines and Power*, Vol. 116, No. 1, January 1994., pp. 82-89

[4]Merrington, G., Kwon, O. K., , Goodwin, G., and Carlsson, B., "Fault Detection and Diagnosis in Gas Turbines," *ASME Journal of Engineering for Gas Turbines and Power*, Vol. 113, No. 2, April 1991, pp. 276-282.

[5]Kerr, L. J., Nemec, T. S., and Gallops, G. W., "Real-Time Estimation of Gas Turbine Engine Damage Using a Control-Based Kalman Filter Algorithm," 91-GT-216, *International Gas Turbine and Aeroengine Congress and Exposition*, Orlando, FL, June 3-6, 1991.

[6]Dewallef, P., Léonard, O., and Mathioudakis, K., "On-Line Aircraft Engine Diagnostic Using a Soft Constrained Kalman Filter," *Proceedings of the ASME Turbo Expo 2004*, GT2004-53539, Vienna, Austria, June 14-17, 2004, pp. 585-594.

[7]Borguet, S., Dewallef, P., and Léonard, O., "On-Line Transient Engine Diagnostics in a Kalman Filtering Framework," *Proceedings of the ASME Turbo Expo 2005*, GT2005-68013, Reno NV, June 6-9, 2005, pp. 473-481.

[8]Simon, D. L., "An Integrated Architecture for Onboard Aircraft Engine Performance Trend Monitoring and Gas Path Fault Diagnostics," NASA TM-2010-216358, May 2010.

[9]Armstrong, J. B., and Simon, D. L., "Implementation of an Integrated On-Board Aircraft Engine Diagnostic Architecture," *47th AIAA Joint Propulsion Conference & Exhibit*, AIAA-2011-5859, San Diego CA, August 2011.

[10]Simon, D. L., Rinehart, A. W. "A Model-Based Anomaly Detection Approach for Analyzing Streaming Aircraft Engine Measurement Data," *Proceedings of the ASME Turbo Expo 2014,* ASME-GT2014-27172, Düsseldorf, Germany, June 16-19, 2014.

[11]May, R. D., Csank, J., Lavelle, T. M., Litt, J. S., and Guo, T. H., "A High-Fidelity Simulation of a Generic Commercial Aircraft Engine and Controller," *46th AIAA Joint Propulsion Conference and Exhibit*, AIAA-2010-6630, Nashville, TN, July 25-28, 2010.

[12]Hunter, G. W., Lekki, J. D., Simon, D. L., "Development and Testing of Propulsion Health Management," *Workshop on Integrated Vehicle Health Management and Aviation Safety*, Bangalore India, Jan. 9-12, 2012.