NASA/TM—2010-216358

An Integrated Architecture for On-Board Aircraft Engine Performance Trend Monitoring and Gas Path Fault Diagnostics

Donald L. Simon
Glenn Research Center, Cleveland, Ohio

May 2010

NASA STI Program . . . in Profile

Since its founding, NASA has been dedicated to the advancement of aeronautics and space science. The NASA Scientific and Technical Information (STI) program plays a key part in helping NASA maintain this important role.

The NASA STI Program operates under the auspices of the Agency Chief Information Officer. It collects, organizes, provides for archiving, and disseminates NASA's STI. The NASA STI program provides access to the NASA Aeronautics and Space Database and its public interface, the NASA Technical Reports Server, thus providing one of the largest collections of aeronautical and space science STI in the world. Results are published in both non-NASA channels and by NASA in the NASA STI Report Series, which includes the following report types:

- TECHNICAL PUBLICATION. Reports of completed research or a major significant phase of research that present the results of NASA programs and include extensive data or theoretical analysis. Includes compilations of significant scientific and technical data and information deemed to be of continuing reference value. NASA counterpart of peer-reviewed formal professional papers but has less stringent limitations on manuscript length and extent of graphic presentations.

- TECHNICAL MEMORANDUM. Scientific and technical findings that are preliminary or of specialized interest, e.g., quick release reports, working papers, and bibliographies that contain minimal annotation. Does not contain extensive analysis.

- CONTRACTOR REPORT. Scientific and technical findings by NASA-sponsored contractors and grantees.

- CONFERENCE PUBLICATION. Collected papers from scientific and technical conferences, symposia, seminars, or other meetings sponsored or cosponsored by NASA.

- SPECIAL PUBLICATION. Scientific, technical, or historical information from NASA programs, projects, and missions, often concerned with subjects having substantial public interest.

- TECHNICAL TRANSLATION. English-language translations of foreign scientific and technical material pertinent to NASA's mission.

Specialized services also include creating custom thesauri, building customized databases, organizing and publishing research results.

For more information about the NASA STI program, see the following:

- Access the NASA STI program home page at *http://www.sti.nasa.gov*

- E-mail your question via the Internet to *help@sti.nasa.gov*

- Fax your question to the NASA STI Help Desk at 443–757–5803

- Telephone the NASA STI Help Desk at 443–757–5802

- Write to:
  NASA Center for AeroSpace Information (CASI)
  7115 Standard Drive
  Hanover, MD 21076–1320

NASA/TM—2010-216358

An Integrated Architecture for On-Board Aircraft
Engine Performance Trend Monitoring and
Gas Path Fault Diagnostics

Donald L. Simon
Glenn Research Center, Cleveland, Ohio

National Aeronautics and
Space Administration

Glenn Research Center
Cleveland, Ohio 44135

May 2010

An Integrated Architecture for On-Board Aircraft Engine Performance Trend Monitoring and Gas Path Fault Diagnostics

Donald L. Simon
National Aeronautics and Space Administration
Glenn Research Center
Cleveland, Ohio 44135

Abstract

Aircraft engine performance trend monitoring and gas path fault diagnostics are closely related technologies that assist operators in managing the health of their gas turbine engine assets. Trend monitoring is the process of monitoring the gradual performance change that an aircraft engine will naturally incur over time due to turbomachinery deterioration, while gas path diagnostics is the process of detecting and isolating the occurrence of any faults impacting engine flow-path performance. Today, performance trend monitoring and gas path fault diagnostic functions are performed by a combination of on-board and off-board strategies. On-board engine control computers contain logic that monitors for anomalous engine operation in real-time. Off-board ground stations are used to conduct fleet-wide engine trend monitoring and fault diagnostics based on data collected from each engine each flight. Continuing advances in avionics are enabling the migration of portions of the ground-based functionality on-board, giving rise to more sophisticated on-board engine health management capabilities. This paper reviews the conventional engine performance trend monitoring and gas path fault diagnostic architecture commonly applied today, and presents a proposed enhanced on-board architecture for future applications. The enhanced architecture gains real-time access to an expanded quantity of engine parameters, and provides advanced on-board model-based estimation capabilities. The benefits of the enhanced architecture include the real-time continuous monitoring of engine health, the early diagnosis of fault conditions, and the estimation of unmeasured engine performance parameters. A future vision to advance the enhanced architecture is also presented and discussed.

Nomenclature

H	Module performance deterioration influence coefficient matrix
Nf	Fan speed
P	Health parameter covariance matrix
PLA	Power lever angle
R	Sensor measurement covariance matrix
V^{*}	Health parameter vector to tuner vector linear transformation matrix
h	Health parameter vector
u	Actuator commands
v	Sensor measurement noise
x	State vector
y	Measured output vector
z	Unmeasured output vector
Δ	Denotes change from baseline
$\Delta\Delta$	Denotes gradient in change from baseline

Superscripts, subscripts, and diacritical marks

$\dagger$	Pseudo-inverse (superscript)
T	Transpose (superscript)
cmd	Commanded value (subscript)
i	Sensor residual index (subscript)
$\wedge$	Estimated value (diacritical mark)

Introduction

Aircraft engines are highly complex systems consisting of static and rotating components, along with associated subsystems, controls, and accessories. They are required to provide reliable power generation over thousands of flight cycles while being subjected to a broad range of operating loads and conditions, including extreme temperature environments. Over repeated flight cycles the life of many engine parts will be consumed, and engine malfunctions may occur. Engine Health Management (EHM) systems play a critical role in assisting aircraft operators in managing the safety, reliability, availability, and affordability of their gas turbine engine assets (Refs. 1 to 4). The functionality provided by an EHM system includes deterioration trend monitoring, life usage tracking, diagnosis of engine faults, and providing recommended inspection and maintenance actions. The individual subsystems that comprise an EHM system vary by application, but generally consist of lubrication system health management, structural vibration monitoring, structural life usage tracking, and gas path health management. This paper will specifically focus on the gas path health management aspects of an EHM system.

The goal of propulsion gas path health management (GPHM) is to reliably assess and manage the health of gas turbine engine flow-path components and controls accessories (Refs. 5 and 6). It is performed by relating observed changes in sensed engine variables to internal performance-related changes within the gas turbine engine cycle. Monitoring and processing engine sensed measurements enables the detection and isolation of problems, ultimately enabling corrective action to be taken (Ref. 7). Examples of the types of events that can be addressed by GPHM include gradual component deterioration, sensor faults, actuator faults, and turbomachinery damage. Because GPHM provides a system level assessment of engine health, it is a cornerstone technology of an overall engine health management system.

GPHM primarily relies upon the gas path sensors installed on the engine for control purposes. Historically, GPHM came into prevalence concurrent with the introduction of digital engine controls and avionics. Prior to this, gas path performance trending was conducted manually based on cockpit gauge readings hand-recorded by pilots. The introduction of digital avionics and controls revolutionized this process by providing access to additional sensed measurements, along with automated data acquisition and processing capabilities. This enabled the inclusion of on-board engine diagnostic functionality, as well as ground-based fleet-wide EHM functionality in computer ground stations, which processes engine data acquired in-flight. Operators now rely on this combined on-board and off-board functionality to manage the health of their engine assets.

Continuing advancements in avionics are enabling more sophisticated on-board GPHM functionality. Towards that vision, this paper presents a proposed enhanced GPHM architecture that incorporates on-board model-based diagnostic and performance estimation functionality. The enhanced architecture gains real-time access to an expanded quantity of engine parameters and provides real-time continuous monitoring of engine health, the early diagnosis of fault conditions, and the estimation of unmeasured engine performance parameters that can be directly applied for health management and controls applications. These benefits, along with identified technology challenges, are discussed in this paper. Additionally, the research and development being conducted by NASA to make the enhanced GPHM architecture a reality is presented.

The remaining sections of this paper are organized as follows. First, the conventional aircraft engine GPHM architecture commonly applied today is presented and discussed in more detail. This is provided to describe the current state-of-the-art and to serve as an introduction to the engine performance trend monitoring and diagnostics process. Next, the proposed enhanced GPHM architecture is presented, along with ongoing NASA research and development efforts focused on maturing that architecture. The paper concludes with a discussion of practical issues for implementing the technology, remaining technology challenges, and a summary.

Conventional GPHM Architecture

A representative example of a conventional GPHM architecture is shown in Figure 1. As shown, this architecture includes both on-board and ground-based functionality. On-board fault detection and isolation (FDI) logic continuously monitors control sensors and actuators. Additionally, automated data acquisition logic is applied to collect in-flight engine measurements. These engine measurements, along
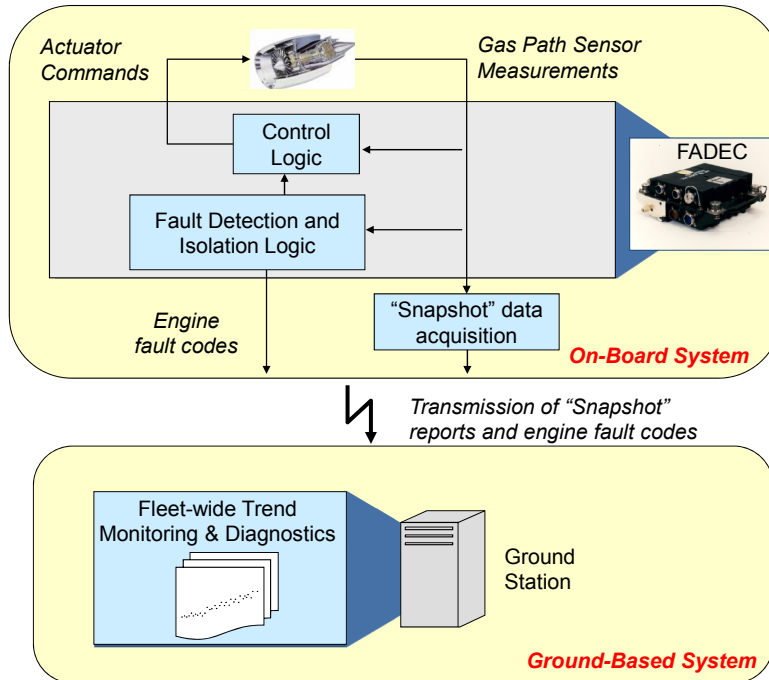
Figure 1.—Conventional gas path health management architecture.

with any identified engine fault conditions identified by the on-board FDI logic, are transmitted off-board to fleet-wide ground stations that apply additional monitoring functionality. The elements of this architecture are discussed in more detail in the following subsections.

Gas Path Sensor Measurements

Most of the sensors used to conduct GPHM are installed on the engine for control purposes. The fact that these sensors also enable GPHM is an added benefit. While some engine designs do include additional sensors specifically for gas path monitoring purposes, their inclusion comes at an expense. Additional sensors add weight, cost, and complexity, and reduce system reliability. Therefore, any additional sensors must provide sufficient benefit to "buy their way on-board." The types of gas path sensed measurements available typically include rotor speeds, pressures, temperatures, and actuator positions. The number of sensors available to conduct GPHM will vary by application, but sensor suites ranging from four to eight sensed measurements are common.

On-Board Fault Detection and Isolation Logic

Conventional on-board GPHM capabilities include FDI logic residing in the engine control computer. The FDI logic operates at an update rate consistent with the engine control computer (approximately 50 Hz). Conventional on-board FDI logic includes checks on the range and rate of change of sensed measurements and built-in-test (BIT) logic to monitor sensor and actuator electronics. Digital engine control systems commonly employ a dual-redundant architecture design to provide increased reliability. This includes physically redundant control sensors with cross-checks between each of the physically redundant sensor channels. In the event of a disagreement, simple analytical models are used as a third vote to help isolate the faulty sensor channel. If a fault is detected by the on-board FDI logic, a fault code that uniquely identifies the fault type is generated. The fault mitigation response then taken by the system is dependent on the fault type. For minor faults that do not compromise the operation of the system, the fault code is simply archived for post-flight indication to maintenance and operations personnel. The operator must comply with dispatchability regulations that dictate the number of future flights, if any, during which the vehicle may continue to operate prior to addressing the fault code through corrective

maintenance. Critical faults, which require an immediate response, are annunciated through cockpit alerts and warnings. Some automated accommodation is also incorporated within the control system. For example, if a failure is detected in a sensor channel, the system can automatically switch to the physically redundant backup channel. The control system may also include revisionary backup modes that can be entered upon the detection of certain fault types. This allows the engine to continue operating, albeit in a reduced performance mode, in the presence of faults.

On-Board Snapshot Data Acquisition and Data Transfer to Ground Station

During flight, the on-board data acquisition system collects and archives engine flight data for subsequent download to a ground station computer. The quantity of data archived varies by application, but is usually limited to "snapshot" engine reports, or measurements, taken at a limited number of engine operating points each flight (e.g., takeoff, climb and cruise). The vehicle to ground station data transfer may occur wirelessly (either in-flight, or while on the ground at an airline gate), or through the manual transfer of mass storage media. Transmission bandwidth limitations, transmission costs, and data archival costs currently restrict the amount of data collected and downloaded from each engine each flight. However, as data transmission and archival capabilities continue to advance these restrictions are expected to lessen.

Ground Station Fleet-Wide Trend Monitoring and Diagnostics

Off-board GPHM functions are performed within EHM computer systems, commonly referred to as ground stations. These ground stations provide fleet-wide engine health management functionality, including performance trend monitoring and gas path diagnostics (Ref. 8). This is accomplished by processing the snapshot engine data acquired in-flight. Performance trending enables operators to monitor the gradual performance deterioration that each engine will naturally undergo over its lifetime of use (Refs. 9 and 10). This deterioration is caused by factors such as fouling, corrosion, and erosion of turbomachinery, increased clearances, and seal leakage. Operators utilize performance deterioration trend information to assist them in making maintenance and overhaul decisions regarding their engine assets. The gas path diagnostic functionality conducted within the ground station is complementary to that conducted by the on-board FDI logic. In general, the ground station is capable of detecting smaller incipient faults, including those faults that evolve over several flights, while the on-board system provides the real-time diagnosis of larger magnitude faults. Reference 2 provides a top-level description of the ground-based fleet-wide trend monitoring and diagnostic process. A representation of this process is shown in Figure 2.
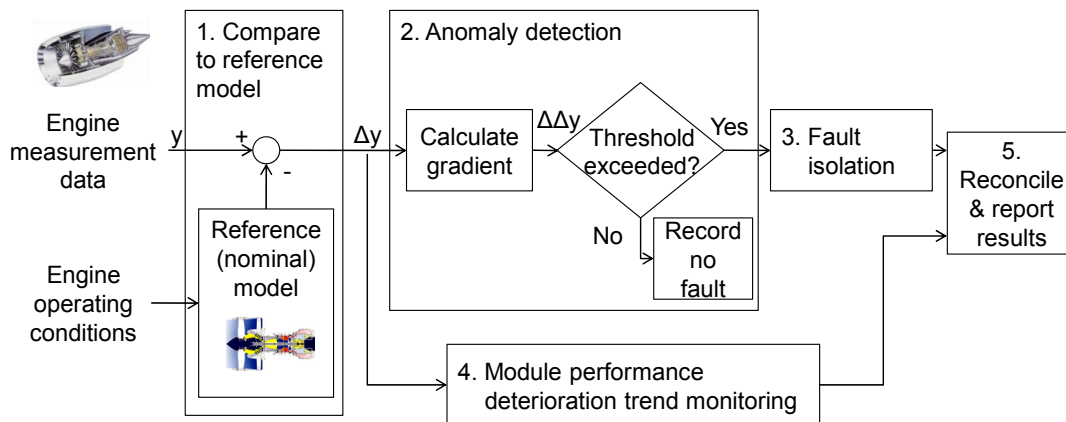

Figure 2.—Ground station performance trend monitoring and gas path diagnostic process.

The individual steps of this process are as follows:

Step 1—Compare to reference model: Incoming data is provided in the form of a snapshot measurement vector, *y*. Typically, this data has been corrected, or normalized, to a standard operating condition to help reduce the effects of operating condition variance from flight to flight (Ref. 11). The incoming data is compared against a reference model (run at the same operating condition and power setting as the measured engine data). Typically, this model is a fleet average engine model representing average engine performance across a fleet of aircraft. The resulting residuals between the measurement data and the reference model are referred to as delta-measurements, or Δ*y*'s.

Step 2—Anomaly detection: Anomaly detection logic is applied to detect any unanticipated rapid shifts in the observed Δ*y* measurements. This logic provides the capability to discriminate between gradual performance changes due to deterioration, which is not a fault, and rapid performance shifts that may be indicative of a fault (Ref. 2). An illustration of gradual versus rapid performance shifts is shown in Figure 3. Here time history data is shown for a single engine sensor residual, Δy_i. To improve the trend monitoring process it is typical to include some form of smoothing or filtering of the data to reduce the effects of measurement noise (Refs. 12 and 13). The blue points in the figure represent the Δy_i residual time history, and the solid red line represents its moving average. Anomaly detection is performed by referencing the engine's current Δ*y* against an archived pre-trend shift Δ*y* vector from the same engine. The change, or gradient, in this vector forms a ΔΔ*y*, or delta-delta measurement vector. If the ΔΔ*y* vector exceeds a defined threshold, a rapid trend shift event is detected and the logic attempts to isolate the cause (step 3). Otherwise the logic proceeds in performing a module deterioration assessment (step 4).

   A noteworthy aspect of this process that warrants emphasis is the coupling between gradient calculation and anomaly detection functionality. The distribution in engine-to-engine performance variation across a fleet of engines can be significant. Due to build tolerances and the range of deterioration levels engines will incur over their lifetime of use, no two engines exhibit identical performance. If detection thresholds were applied to Δ*y* (instead of ΔΔ*y*) they would have to be set suitably large to maintain an acceptable false alarm rate, that in turn would limit the system's ability to detect smaller magnitude faults. By leveraging recent past engine performance information, the detection approach can be tailored to the current performance level of the engine. This allows the application of tighter anomaly detection thresholds and, consequently, the detection of smaller magnitude incipient faults. This is a key difference between the on-board versus ground-based FDI logic applied in the conventional GPHM architecture. The on-board FDI logic is in general a static design, whereas the ground-based FDI logic adapts to the current performance level of the engine. Overall, the on-board and off-board logic are complementary. The on-board FDI logic detects larger magnitude abrupt faults that require more immediate attention, whereas the ground-based FDI logic enables the early detection of smaller magnitude incipient faults before they escalate to more severe consequences.
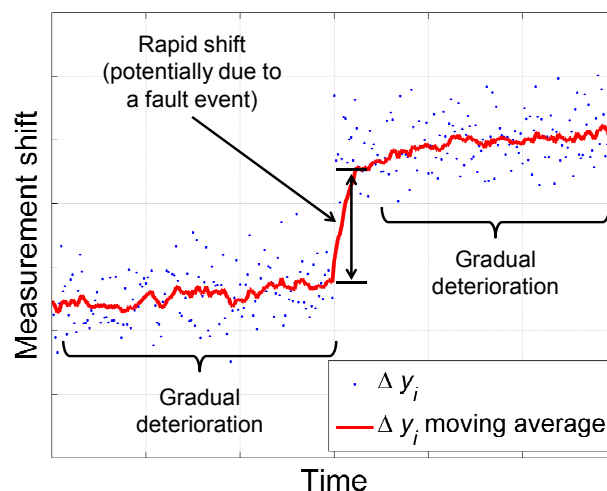

Figure 3.—Gradual versus rapid performance shifts.

Step 3—Fault isolation: Upon anomaly detection, the ΔΔ*y* vector will reflect the signature of the event. The ΔΔ*y* vector is then processed to isolate the root cause of the trend shift. A variety of methods have been applied for gas path fault isolation (Ref. 14). Most of these methods are based upon the single fault assumption. The simultaneous occurrence of multiple faults is usually not considered, which is reasonable considering the high reliability of modern aircraft gas turbine engines. A single fault isolator will individually evaluate how closely each candidate fault case matches the observed ΔΔ*y* measurement vector. It is common to incorporate additional system knowledge to help improve the accuracy of the diagnostic assessment. This includes component fault rates, fleet-wide trends, engine operating cycle counts, and recent maintenance actions. Upon completion, the single fault isolator will produce a diagnostic assessment isolating the cause for the observed trend shift. This diagnostic assessment may be a ranked list of plausible faults (including confidence levels of each), or isolation to an ambiguity group (a group of faults known to produce similar ΔΔ*y* signatures). The diagnostic assessment may also include the capability to classify non-fault events known to cause rapid performance trend shifts such as compressor water washes.

Step 4—Module performance deterioration trend monitoring: If a rapid trend shift event is not detected, the logic proceeds in performing a module deterioration assessment. This assessment is commonly referred to as gas path analysis. The objective of this assessment is to estimate and trend the engine's level of performance deterioration over its lifetime of use. Several parameters are trended including fuel consumption rate, turbine temperature margin, and module health parameters such as efficiencies and flow capacities. The health parameters are not sensed directly, but are instead estimated using model-based estimation techniques. To illustrate this, consider the following linear steady-state representation of the Δ*y* measurement process

$$\Delta y = H \Delta h + v \tag{1}$$

where Δ*h* is a vector of unmeasured module health parameters that impact engine performance, *H* is an influence coefficient matrix relating health parameters to Δ*y* measurements, and *v* represents random measurement error. The matrix *H* is typically derived by linearizing a physics-based aerothermal engine model, or cycle deck. The estimation objective is to estimate Δ*h* given Δ*y*. This typically poses an underdetermined estimation problem as there are usually fewer sensors than health parameters. A common estimation approach is to apply least squares maximum *a posteriori* (MAP) estimation (Refs. 5 and 15) given as

$$\Delta \hat{h} = \left(P^{-1} + H^T R^{-1} H \right)^{-1} H^T R^{-1} \cdot \Delta y \tag{2}$$

where *P* is the covariance matrix of the health parameters, and *R* is the covariance matrix of the measurement vector. By applying this process, operators can estimate and monitor $\Delta \hat{h}$ levels and trends to assist them in making engine maintenance and overhaul decisions.

Step 5—Reconcile and report results: The final step in the process is to reconcile and report the outputs from the fault isolator and the module deterioration assessment. Reconciliation is done to increase confidence in the provided diagnostic assessment and to resolve any conflicts prior to relaying the information to the end user.

Enhanced GPHM Architecture

The previously described conventional GPHM architecture has several notable shortcomings. First, it conducts ground-based FDI based on the post-flight processing of a small number of engine snapshot measurements collected each flight. Often, this processing does not happen until several days after the flight has occurred causing significant diagnostics latency. Additionally, engine performance deterioration estimates produced within the conventional architecture are only available in ground-station. They are not available for on-board controls and health management applications, and thus are not utilized to their full potential. The enhanced GPHM architecture design addresses these limitations by migrating portions of

the conventional ground-based functionality on-board. This is done to enable the real-time continuous monitoring of engine health, the earlier diagnosis of faults, and the estimation of unmeasured engine performance and health parameters that can be directly applied for on-board controls and health management applications.

The enhanced architecture, like the previously presented conventional architecture, contains both on-board and ground-based functionality. In fact, many of the components that make up the two architectures are identical. The primary difference is the inclusion of real-time on-board model-based GPHM functionality within the enhanced architecture. The inclusion of on-board models is an emerging practice in the aircraft engine industry. These models have the ability to self-tune to match the current performance level of the engine as it degrades over time (Refs. 16 to 18). Current applications of on-board self-tuning engine models are largely for sensor validation. The future vision is to expand the utility of such models for use in engine diagnostics and to provide more accurate estimation of unmeasured engine parameters to support advanced model-based controls and life usage monitoring applications. The proposed enhanced GPHM architecture that leverages on-board model-based technology is shown in Figure 4. The enhanced functionality included in the architecture is outlined by dashed boxes.

Enhanced On-board Model-Based GPHM Architecture

A more detailed view of the on-board portion of the enhanced GPHM architecture is shown in Figure 5. It includes many of the same components as the conventional architecture including control logic, conventional FDI logic, and snapshot data acquisition capabilities. Added enhancements include two on-board real-time aerothermal engine models, enhanced fault detection and isolation logic, and real-time performance estimation capabilities. Pivotal to the operation of this architecture are the two engine models that operate in parallel. The first model, referred to as the real-time adaptive performance model (RTAPM), continuously self-tunes internal tuning parameters that allow it to match the current sensed performance level of the physical engine. The tuning parameters are proxies for the engine's health, or level of deterioration. Through this tuning parameter adjustment, the model is able to track engine performance while the engine degrades with use. The second model, referred to as the performance baseline model (PBM), serves as a baseline of recent past engine performance. Residuals between the engine and PBM outputs are monitored for anomaly detection purposes. The PBM also adapts to account for engine performance deterioration by accepting tuning parameter inputs calculated within the RTAPM. However, updates of the tuning parameters used in the PBM only occur on a periodic basis (e.g., after each flight). This periodic update is necessary to ensure that the PBM doesn't immediately absorb the effect of any rapid trend shifts, thus compromising anomaly detection capability. The on-board elements of the enhanced GPHM architecture are discussed in more detail in the subsections that follow.

Control Logic

The control logic included in the enhanced architecture is identical to the control logic included in the conventional architecture. Figure 5 provides an example in the form of a fan speed (Nf) controller, but the overall architecture is applicable for any controller design. In this example, power level angle (PLA) commands are converted to a commanded fan speed (Nf_{cmd}). The error between actual and commanded fan speed is fed into the control logic that in turn generates a set of actuator commands, u, provided as engine inputs. Also shown in the control logic block is the conventional FDI logic. This is the same on-board FDI logic included in the conventional GPHM architecture. It includes BIT and cross-channel check logic applied for sensor and actuator validation purposes. This proven technology has been flight certified and plays a critical role in ensuring the safe, reliable operation of today's aircraft propulsion systems. Instead of completely replacing this technology with enhanced FDI logic, it is maintained and relied upon to make flight critical real-time FDI assessments. Conversely, the enhanced FDI logic, to be described later, will enable the diagnosis of smaller magnitude faults to be mitigated through post-flight maintenance.
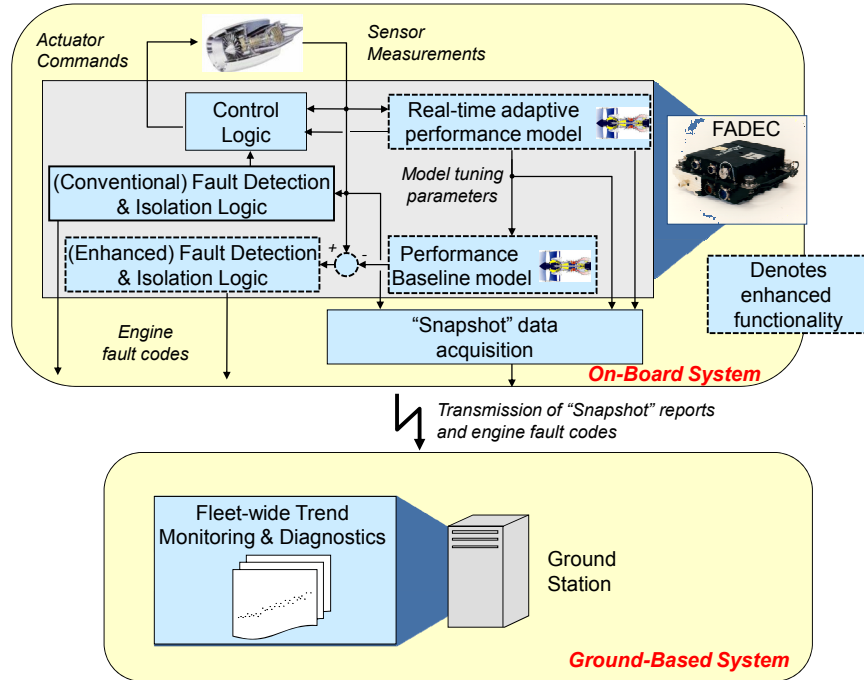
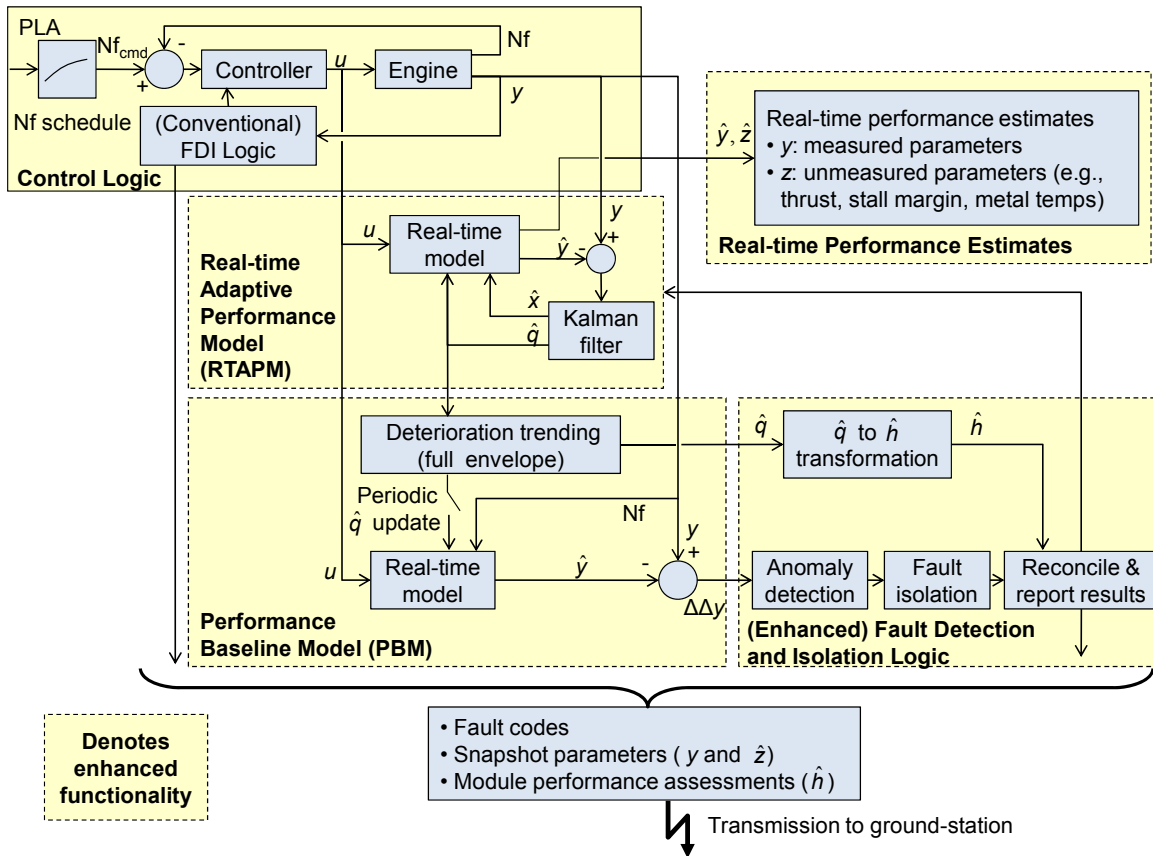Figure 4.—Proposed enhanced gas path health management architecture.

Figure 5.—Proposed enhanced on-board model-based GPHM architecture.

Real-Time Adaptive Performance Model (RTAPM)

The purpose of the RTAPM is to provide a continuous assessment of engine performance. It consists of a real-time physics-based aerothermal engine model and an associated tracking filter that automatically tunes the model to track the performance of the physical engine based on sensed feedback measurements. Kalman filters are typically applied to implement the tracking filter. The tracking filter produces real-time estimates of the engine's state variables, $\hat{x}$, and tuning parameters, $\hat{q}$. The tuning parameters reflect the effects of performance deterioration on engine outputs. The $\hat{x}$ and $\hat{q}$ variables along with actuator commands, u, are provided as inputs to the real-time model, allowing the model-produced estimates to "track" the outputs of the engine. This includes estimates of measured engine outputs, $\hat{y}$, and unmeasured, or auxiliary, engine outputs denoted as $\hat{z}$. The auxiliary outputs are user specified and may include model calculated parameters such as thrust and stall margin. In the upcoming real-time performance estimates sub-section, a methodology to help improve the estimation accuracy of the model will be discussed.

Performance Baseline Model (PBM)

The PBM provides a baseline of recent past engine performance used for trend monitoring purposes. It incorporates a version of the same real-time model used in implementing the RTAPM. The PBM accepts continuous inputs consisting of actuator commands, u, and a trim parameter representative of engine power such as fan speed. The PBM is trimmed to the power reference parameter to prevent the model from inadvertently diverging from the engine. In addition to the continuous inputs, the model also automatically receives periodically updated tuning parameter estimates, $\hat{q}$. This periodic update of the tuning parameters allows the PBM to adapt to normal performance deterioration that occurs gradually over time, but prevents the model from immediately adapting to rapid performance shifts caused by faults. The concept of periodically updating the tuning parameters used within an on-board model applied for diagnostics purposes was first introduced in Reference 19. The frequency of the periodic tuning parameter update is user-specified, but is typically on the order of once per flight.

To illustrate the benefit of updating the PBM tuning parameters on a periodic basis, a simulated fault example is given. Here, the dynamic response of the PBM is compared to that of the RTAPM. First, consider the results shown in Figure 6 which are based on the RTAPM. Shown in this figure are time history plots of: (a) the sensed and model produced estimate of a single sensed output (top plot); (b) the corresponding residual between these two parameters (middle plot); and (c) the Kalman filter estimated tuning parameters (bottom plot). In this example the engine experiences a high pressure compressor fault at time = 10 sec. Since the Kalman filter provides continuous adjustments to the tuning parameters the model estimated output is able to closely track the sensed engine output and the post-fault residuals between the model and engine outputs are minimal. Because of this, the RTAPM alone is not ideal for anomaly detection purposes. Conversely, consider Figure 7 which shows the PBM response to the same fault event. Here the PBM tuning parameters are held constant (bottom plot). In this example the tuner values have been established based on the mean pre-fault Kalman filter estimates shown in Figure 6. When the fault occurs, the change in the residual between the PBM output and the engine output is clearly discernable (middle plot). The vector of residuals generated by the PBM is information that can be directly monitored by on-board logic for anomaly detection and fault isolation purposes.

(Enhanced) Fault Detection and Isolation Logic

By design, the on-board enhanced FDI logic closely emulates the FDI logic applied in the ground station (see Fig. 2). The residuals between the PBM outputs and the sensed engine outputs are analogous to the $\Delta\Delta y$ vector used in the ground station diagnostic process. They reflect shifts in current engine performance relative to recent past engine performance, and can be monitored directly to detect performance trend shift events. If a performance shift is detected, these same residuals can be processed within a single fault isolation process to determine the most likely root cause for the event. Two key differences between the on-board and ground-based solutions are processing latency and available data quantity. The ground-based approach performs post-flight processing of a small number of engine snapshot measurements collected each flight while the on-board approach performs real-time processing of engine data collected at a much higher sample rate. Additionally, the on-board approach will have
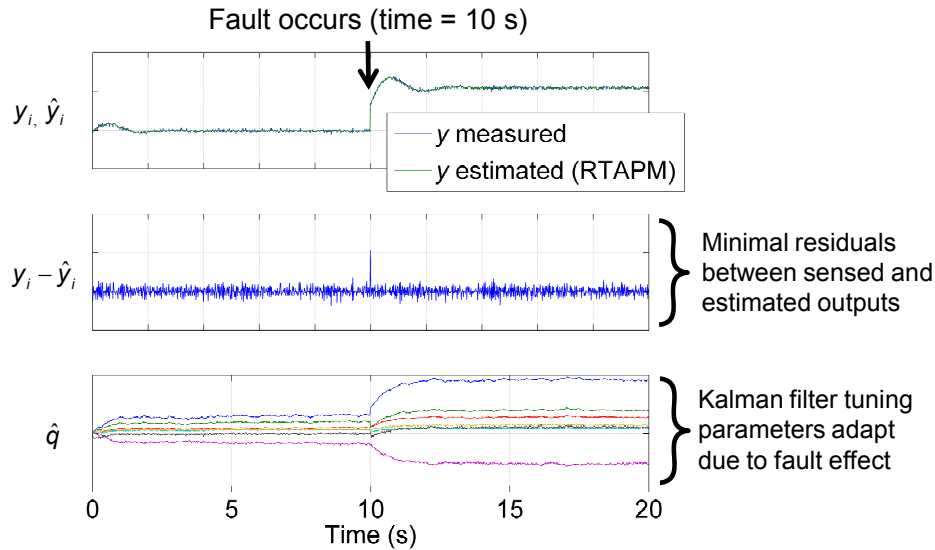
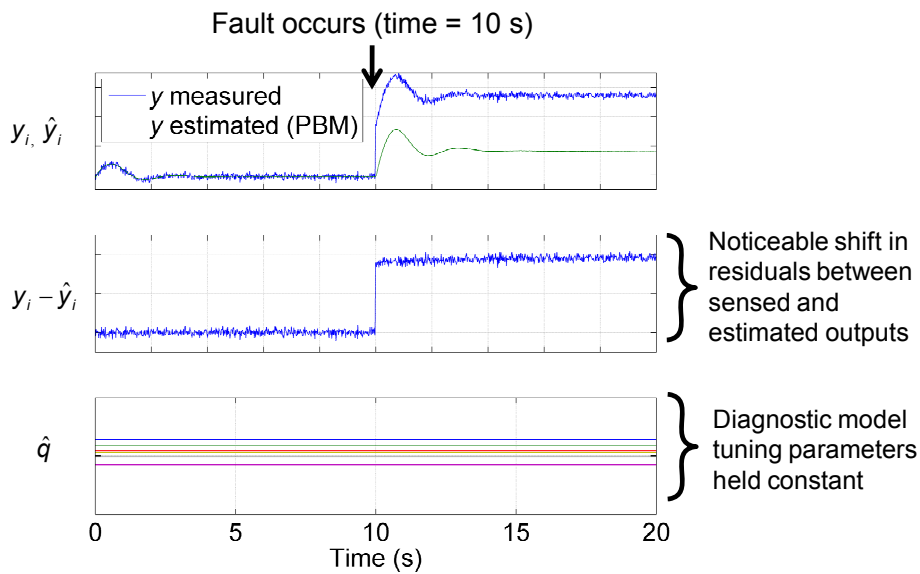Figure 6.—Real-time adaptive performance model (RTAPM) response to an abrupt fault.

Figure 7.—Performance baseline model (PBM) response to an abrupt fault.

access to significantly more engine parameters (i.e., all engine control computer parameters) than what can practically be downloaded to a ground station due to data telemetry and archival restrictions. These advantages enable the on-board FDI solution to produce diagnostic assessments with improved accuracy, and reduced diagnostic latency.

Real-Time Performance Estimates

An additional benefit provided by the RTAPM is the generation of real-time estimates of measured and unmeasured engine outputs. This information can be used for a variety of on-board applications including sensor validation, model-based controls, deterioration trend monitoring, and component life usage calculations. The accuracy of these estimates is dependent on a number of factors including the accuracy of the physics-based model, the available sensor suite (type and accuracy), and the ability of the tracking filter to capture engine performance deterioration effects.

As previously mentioned, engine performance deterioration is generally described in terms of an unmeasurable vector of health parameters, h. Previously, Equation (2) introduced the least squares maximum *a posteriori* (MAP) estimation approach commonly applied in ground stations for module deterioration assessment purposes. The MAP estimator includes *a priori* information regarding health parameter covariance, P. This enables the estimator to produce estimates of all health parameters, albeit *biased estimates*, in the underdetermined case where there are fewer sensors than health parameters. However, the MAP estimator, unlike a Kalman filter, is not a recursive estimator and does not take advantage of past measurements to enhance its estimate at the current time step. Furthermore, the MAP estimator only considers the static steady-state relationship between system states and measured outputs—it does not consider system dynamics. While the MAP estimator is ideal for processing steady-state snapshot engine measurements in ground stations, it is not suitable for the on-board processing of real-time dynamic engine measurements. Through Kalman filter-based estimation techniques, h can be estimated on-board in real-time, given that there are at least as many sensors as parameters to be estimated. However, in an aircraft engine the number of sensors available is typically less than the number of health parameters, presenting an underdetermined estimation problem. A common approach to address this shortcoming is to estimate a subset of the health parameters, referred to as model tuning parameters, q. The q vector serves as a proxy for capturing full-order engine deterioration effects as reflected by the h vector. Applying a reduced-order q vector will enable the Kalman filter to tune the on-board model to track measured (sensed) engine outputs. However, the Kalman filter produced estimate of unmeasured engine outputs may be biased due to the fact that the impact of the entire vector of health parameters will not be accurately represented within the model (Ref. 20).

In a departure from the conventional technique of selecting a subset of h to serve as q, an innovative methodology was developed that constructs a tuning parameter vector that is a linear combination of all health parameters and of appropriate dimension to enable Kalman filter estimation (Refs. 20 and 21). The relationship between h and q is defined by a linear transformation matrix, V^*, as shown in Equation (3).

$$q = V^* h \tag{3}$$

The selection of V^* is performed such that the Kalman filter's theoretical mean squared estimation error in the parameters of interest is minimized. For a complete derivation see Reference 20. Given the $\hat{q}$ estimates produced by the Kalman filter, an approximation of the health parameter vector, $\hat{h}$, can be obtained as

$$\hat{h} = V^{*\dagger} q \tag{4}$$

where $V^{*\dagger}$ is the pseudo-inverse of V^*. These health parameter estimates are used on-board to help reconcile diagnosed fault conditions. They can also serve to augment or verify the health parameters estimated and trended within the ground station applying MAP estimation. Since these on-board estimates are updated continuously, they will be able to reflect engine performance trend changes earlier than the ground-based module deterioration assessment conducted using snapshot engine measurements alone.

Figure 8 provides an illustration of the improvement in estimation accuracy that can be gained by applying the enhanced tuner selection approach (bottom plot) versus the conventional approach of defining tuning parameters as a subset of health parameters (top plot). In each plot the red line denotes actual net thrust and the cyan line denotes the Kalman filter estimate of thrust. Applying the enhanced tuner selection approach helps to significantly improve estimation accuracy by reducing the bias and the variance in the estimate.
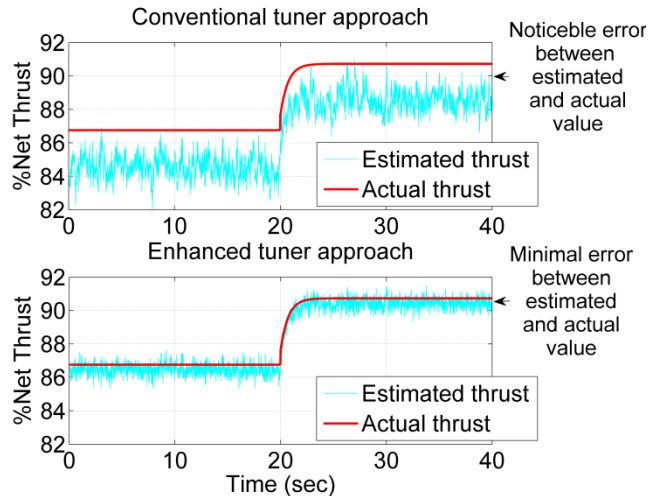
Figure 8.—Thrust estimation accuracy comparison of conventional versus enhanced estimation approaches.

On-Board Snapshot Data Acquisition and Data Transfer to Ground Station

The enhanced GPHM architecture, like the conventional architecture, contains both on-board and ground-based functionality. As such, the ability to perform on-board data acquisition and data transfer to the ground station is required. As with the conventional architecture, the data transfer can occur either in flight, or on the ground at an airline gate. The acquired and transferred data will consist of sensed measurements and engine fault codes. Additionally, the auxiliary parameter estimates, health parameter estimates, and information on any faults detected by the enhanced on-board FDI logic will also be acquired and transferred to the ground station.

Ground Station Fleet-Wide Trend Monitoring and Diagnostics

Although the proposed enhanced architecture holds promise for improving GPHM capabilities, it does not eliminate the need for EHM ground stations or the functionality that they provide. Manual human interrogation of the diagnostic assessments produced by the GPHM system will always be necessary to ensure that the recommended mitigation actions (e.g., inspections, repairs, overhauls, etc.) are indeed warranted.

Furthermore, the on-board and ground-based elements of the architecture are viewed as providing synergistic diagnostic functionality. On-board functionality provides continuous processing resulting in reduced diagnostic latency compared to ground-based processing of snapshot measurements. On-board continuous processing is also expected to be better suited for detecting intermittent fault behavior, which is difficult to detect relying on snapshot measurements alone. Conversely, the conventional ground-based approach is better suited for the detection of faults that evolve over multiple flight cycles. The periodic updating of PBM tuning parameters applied on-board will eventually absorb the effect of any performance changes. If an update of the tuning parameters occurs before a fault evolves to a detectable magnitude, the fault may not be detected. In combination, the on-board and ground-based functionality provide operators information that can be combined to gain increased confidence in GPHM assessments. A notional illustration of the fault detection capability provided by on-board versus ground-based FDI logic is shown in Figure 9. The subplots on the left are representative of the once per flight $\Delta\Delta y$ snapshot engine measurements available in a ground-station, while the subplots on the right represent available on-board $\Delta\Delta y$ engine measurements sampled at a higher frequency (1000 samples per flight in this example). This figure plots nominal (green points) and faulty (red points) $\Delta\Delta y$ residual engine data observed in two sensors under the following four fault scenarios: a) large magnitude abrupt fault; b) small magnitude abrupt fault; c) intermittent fault; and d) evolving fault. The first three scenarios show 20 flights of nominal

pre-fault data and one flight of faulty data. The fourth scenario, evolving fault, shows 20 flights of pre-fault data and 10 flights of faulty data to illustrate how an evolving fault signature can be absorbed within the PBM update. Each scenario is further discussed below:

a) *Large magnitude abrupt fault*: In this scenario the faulty versus nominal $\Delta\Delta y$ points are clearly discernable in both the ground-based and on-board FDI systems. Both systems are expected to perform well in detecting these types of faults within one flight.

b) *Small magnitude abrupt fault*: Small magnitude faults are much more difficult to detect based upon a single measurement sample. A fault of the magnitude shown would need to persist for several flights in order for the ground-based FDI approach to detect it while maintaining an acceptable false alarm rate. Conversely, the fault is more clearly noticeable in the on-board system due to the increased number of samples available.

c) *Intermittent fault*: Here the intermittent nature of the fault will cause the system to randomly alternate between exhibiting nominal and anomalous behavior. The ground-based FDI system will only be able to detect these types of faults if the snapshot data acquisition occurs while the system is exhibiting the fault. Conversely, the on-board FDI system is much more likely to detect these types of faults due to its higher sample rate.

d) *Evolving fault*: This scenario reflects a fault that grows in magnitude over 10 flights. The ground-based FDI approach reflects $\Delta\Delta y$ residuals generated by referring current engine measurements against those collected prior to fault initiation. After several flights the fault becomes apparent within the ground-based FDI approach. Conversely, the fault is not apparent within the on-board FDI system due to the PBM update, once per flight in this example, which continuously absorbs the fault signature as it evolves over time.
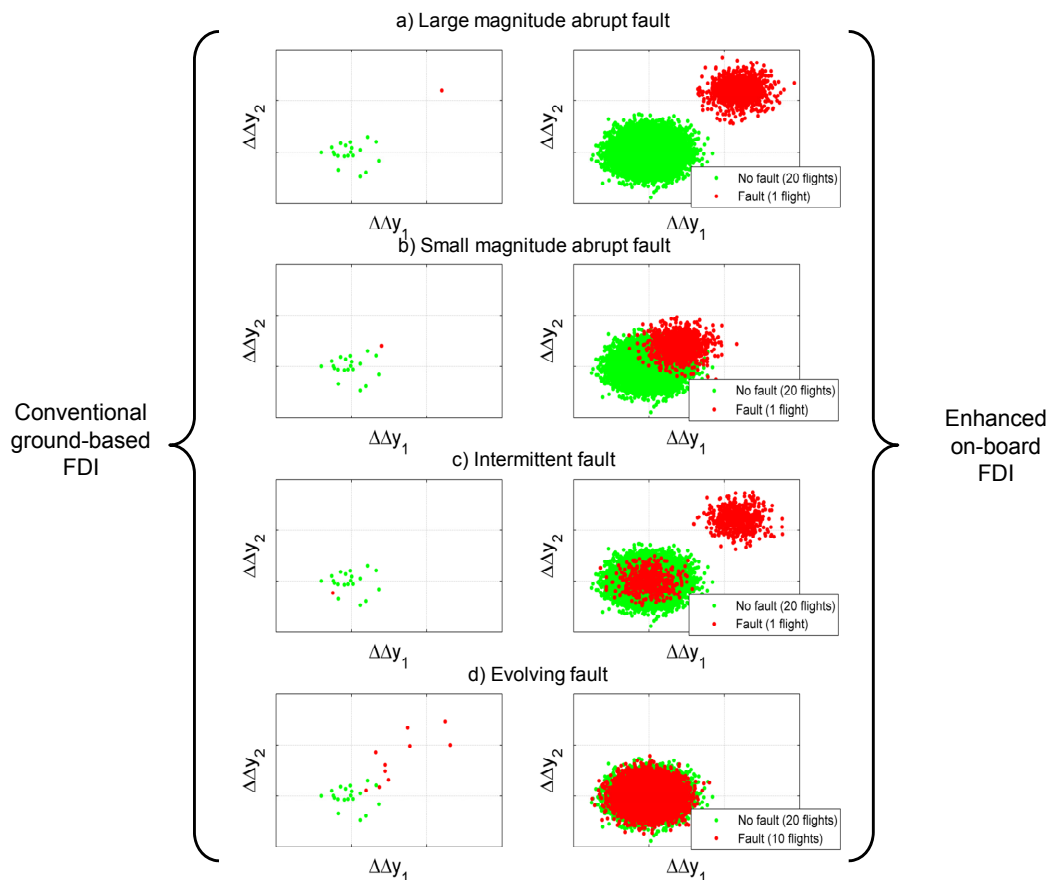

Figure 9.—Comparison of ground-based and on-board fault detection capabilities under different fault scenarios.

Future Work and Technology Challenges

There are several technology challenges that need to be addressed to bring the proposed enhanced GPHM architecture to fruition. These challenges and the plans to address them include:

- *On-board full operating envelope deterioration trending*: The tuning parameter inputs provided as inputs to the PBM are not expected to be uniform over the entire engine operating envelope. Therefore, a full-envelope approach for tuning parameter trending and archiving is necessary. This potentially could be accomplished by partitioning the operating envelope into regimes, and updating the average tuning parameter values within each regime over time. As the engine transitions throughout its operating envelope, the PBM would transition to the corresponding tuning parameter vector. Such an approach was taken in Reference 17, which describes the development of a hybrid (empirical and analytical) self-tuning on-board engine model. The empirical element of the model was trained and implemented by partitioning the model into different "cells" spanning the engine operating envelope. In addition to full-envelope trending, some means of full-envelope "resetting" of the tuning parameters is also necessary. Otherwise maintenance events that result in a performance trend shift (e.g., compressor water washes, replacing biased sensors, etc.) could potentially cause a false alarm.
- *Adaptive anomaly detection thresholds*: The on-board anomaly detection logic monitors $\Delta\Delta y$ residuals between the PBM and engine outputs, and declares an anomaly if these residuals exceed defined thresholds. The applied threshold levels will have a profound effect on the system's false alarm rate, and can render the system useless if not set properly. Scenarios that have the potential to cause large residuals include engine transients, noisy sensor measurements, operating in fringe operating regimes where the tuning parameters do not accurately reflect current engine performance, and long delays between the periodic PBM tuner updates. These factors should be thoroughly considered when defining the thresholds. An approach that may be beneficial is the application of *adaptive thresholds* (Refs. 22 and 23). The concept is that the detection thresholds are relaxed within operating scenarios known to produce large residuals, and tightened in other scenarios where the residuals are expected to be lower.
- *Quantify the diagnostic benefits of the enhanced architecture*: The overall cost versus benefit of the enhanced architecture needs to be demonstrated. While the enhanced architecture is expected to yield diagnostic improvements, it will add cost and complexity to the overall system design. NASA plans to perform a benchmarking exercise to quantify and compare the expected diagnostic performance of the on-board system relative to a conventional ground station gas path diagnostic approach. Additionally, NASA is creating and distributing publicly available gas path diagnostic benchmarking problems to the EHM community. These standard benchmark problems will allow others to apply and evaluate their GPHM approaches as well (Refs. 24 and 25).
- *Software certification*: The inclusion of additional on-board model-based software poses certification challenges beyond those typically faced by conventional designs. Of primary concern is ensuring that the software does not cause unintended consequences. To address this concern a functional hazard risk assessment must be conducted to identify all possible software failures, assess the consequences of each failure, and ensure that each failure will be appropriately mitigated if it occurs. Integral to this assessment is defining how the software impacts the system. Software used for control, real-time fault accommodation, or cockpit indications will require a higher level of certification than software used to generate post-flight maintenance or inspection advisories. These are factors that should be considered early in the design process.

Long Term Vision

A key justification for the inclusion of more sophisticated GPHM functionality on-board is the limited access to continuous full-flight engine sensor measurements off-board. Today, the cost of transmitting and archiving full-flight engine data is cost prohibitive. In the future, as data transmission and data archival costs are lessened, this will become less of a hurdle and ground station access to full-flight measurement data may become available. If this occurs, the question of where the engine GPHM diagnostic functionality resides should be reconsidered. If the same data is available on-board and off-

board, then the on-board system may only require the FDI functionality required to address faults that require an immediate on-board mitigation response.

Nevertheless, on-board self-tuning models are expected to continue to play a major role in future aircraft engine designs. Over time, as on-board model-based gas path diagnostics grows in use and credibility, the technology is expected to be applied for on-board fault mitigation functions. Additionally, perhaps the biggest technology benefit offered by the on-board models is the real-time performance estimation capabilities that they provide. It is not possible to directly measure engine parameters such as thrust or stall margin in flight. Today, engine controllers are designed to indirectly control these parameters based upon other sensor measurements. While functional, this current approach does not provide direct control of the main engine parameters of interest. Accurate on-board model-based estimates of thrust and stall-margin could revolutionize engine controller designs. This model-based control technology pull is only expected to increase in the future.

Summary

This paper reviewed the conventional aircraft engine gas path health management (GPHM) architecture commonly applied today, and also proposed an enhanced architecture that incorporates on-board model-based functionality. The inclusion of enhanced on-board GPHM functionality is becoming possible due to the increasing processing capabilities of modern avionics computers. Both the conventional and the enhanced architectures include on-board and ground-based functionality. By design, the two architectures follow a similar diagnostic process consisting of performance trend monitoring, anomaly detection, and fault isolation. A key difference is the fact that on-board systems are able to conduct the continuous real-time monitoring of engine sensor measurements, while ground-based systems are primarily relegated to the post-flight processing of snapshot engine measurements collected at a small number of operating points each flight. The enhanced model-based architecture allows on-board fault detection and isolation (FDI) logic to be tailored to the performance level of each individual engine, similar to the ground-based FDI logic. This promises to enable the earlier detection of gas path faults. In addition, the on-board model allows the real-time estimation of unmeasured engine parameters. This could prove to be highly enabling for model-based controls, module performance deterioration tracking, enhanced cockpit displays and annunciations, and usage-based lifing calculations. NASA is considering future work to quantify the diagnostic benefits provided by the enhanced architecture described in this report. This includes quantifying the reduction in diagnostic latency and the reduction in detectable fault magnitudes offered by the on-board approach.

References

1. Jaw, L.C., (2005), "Recent Advances in Aircraft Engine Health Management (EHM) Technologies and Recommendations for the Next Step," ASME Turbo Expo Conference, paper GT2005-68625, June 6–9, Reno, NV.
2. Volponi, A., Wood, B., (2005), "Engine Health Management for Aircraft Propulsion Systems," The Forum on Integrated System Health Engineering and Management (ISHEM) in Aerospace, November 7–10, Napa, CA.
3. Jaw, L.C., Mattingly, J.D., (2009), "Aircraft Engine Controls: Design, System Analysis, and Health Monitoring," *American Institute of Aeronautics and Astronautics*, Reston, VA.
4. Tumer, I.Y., and Bajwa, A., (1999), "A Survey of Aircraft Engine Health Monitoring Systems," 35th AIAA Joint Propulsion Conference and Exhibit, paper AIAA–99–2528, June 20–23, Los Angeles, CA.
5. Volponi, A.J., et al., (2003), "Gas Turbine Condition Monitoring and Fault Diagnostics," Von Karman Institute for Fluid Dynamics, Lecture Series 2003-01.
6. Doel, D.L., (1994), "TEMPER – A Gas Path Analysis Tool for Commercial Jet Engines," *Journal of Engineering for Gas Turbines and Power*, Vol. 116, No. 1, pp. 82–89.
7. Urban, L.A., (1974), "Parameter Selection for Multiple Fault Diagnostics of Gas Turbine Engines," ASME Gas Turbine Conference and Products Show, paper 74–GT–62, March 30-April 4, Zurich, Switzerland.
8. Society of Automotive Engineers E-32, (2005), "A Guide to the Development of a Ground Station for Engine Condition Monitoring," SAE Aerospace Information Report 4175A.

9. Sallee, G.P., (1978) "Performance Deterioration Based on Existing (Historical) Data – JT9D Jet Engine Diagnostics Program," NASA Contractor Report No. CR—135448, United Technologies Corporation, Pratt & Whitney Aircraft Group Report PWA-5512-21.

10. Wulf, R.H., (1980), "Engine Diagnostics Program CF6-50 Engine Performance Deterioration," NASA Contractor Report No. CR—159867.

11. Volponi, A.J., (1999), "Gas Turbine Parameter Corrections," *Journal of Engineering for Gas Turbines and Power*, Vol. 121, No. 4, pp. 613–621.

12. DePold, H.R., Gass, F.D., (1999), "The Application of Expert Systems and Neural Networks to Gas Turbine Prognostics and Diagnostics," *Journal of Engineering for Gas Turbines and Power*, Vol. 121, No. 4, pp. 607–612.

13. Ganguli, R., (2002), "Data Rectification and Detection of Trend Shifts in Jet Engine Gas Path Measurements Using Median Filters and Fuzzy Logic," *ASME Journal of Engineering for Gas Turbines and Power,* Vol. 124, No. 4, pp. 809–816.

14. Y.G. Li, (2002), *"*Performance-Analysis-Based Gas Turbine Diagnostics: A Review,*" Proceedings of the Institution of Mechanical Engineers, Part A: J. Power and Energy*, Vol. 216, pp. 363–377.

15. Doel, D., (1994), "An Assessment of Weighted-Least-Squares-Based Gas Path Analysis," *ASME Journal of Engineering for Gas Turbines and Power*, Vol. 116, No. 2, pp. 366–373.

16. Luppold, R.H. Roman, J.R., Gallops, G.W., Kerr, L.J., (1989), "Estimating In-Flight Engine Performance Variations Using Kalman Filter Concepts," AIAA 25th Joint Propulsion Conference, paper AIAA–89–2584, July 10–14, Monterey, CA.

17. Volponi, A., (2008), "Enhanced Self-Tuning On-board Real-Time Model (eSTORM) for Aircraft Engine Performance Health Tracking," NASA Contractor Report No. CR—2008-215272.

18. Kumar, A., Viassolo, D., (2008), "Model-Based Fault Tolerant Control," NASA/CR—2008-215273.

19. Kobayashi, T, and Simon, D.L., (2007), "Integration of On-Line and Off-Line Diagnostic Algorithms for Aircraft Engine Health Management," ASME Turbo Expo Conference, paper GT2007-27518, May 14–17, Montreal, Canada.

20. Simon, D.L., Garg, S., (2009), "Optimal Tuner Selection for Kalman-Filter Based Aircraft Engine Performance Estimation," *Journal of Engineering for Gas Turbines and Power*, Vol. 132, No. 3, /031601-1.

21. Litt, J.S., (2008), "An Optimal Orthogonal Decomposition Method for Kalman Filter-Based Turbofan Engine Thrust Estimation," *Journal of Engineering for Gas Turbines and Power*, Vol. 130, No. 1, /011601-1.

22. Chen, J., and Patton, R.J., (1999), "Robust Model-Based Fault Diagnosis for Dynamic Systems," Kluwer Academic Publishers. London, England.

23. Yedavalli, R.K., Li, W., (2007), "Aircraft Engine Fault Detection Using Dynamic/Adaptive Threshold Approach," ASME Turbo Expo, paper GT2007-28180, May 14–17, Montreal, Canada.

24. Simon, D.L., Bird, J., Davison, C., Volponi, A., Iverson, R.E., (2008), "Benchmarking Gas Path Diagnostic Methods: A Public Approach," ASME Turbo Expo Conference, paper GT2008-51360, June 9–13, Berlin, Germany.

25. Simon, D.L., (2010), "Propulsion Diagnostic Method Evaluation Strategy (ProDiMES) User's Guide," NASA/TM—2010-215840.

Appendix—Glossary

BIT	Built-in-test
EHM	Engine health management
FDI	Fault detection and isolation
GPHM	Gas path health management
Hz	Hertz
MAP	Maximum *a posteriori*
PBM	Performance baseline model
RTAPM	Real-time adaptive performance model
s	Seconds

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED *(From - To)*
05-01-2010	Technical Memorandum	

4. TITLE AND SUBTITLE	5a. CONTRACT NUMBER
An Integrated Architecture for On-Board Aircraft Engine Performance Trend Monitoring and Gas Path Fault Diagnostics	
	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S) Simon, Donald, L.	5d. PROJECT NUMBER
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER WBS 645846.02.07.03.12

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)	8. PERFORMING ORGANIZATION REPORT NUMBER
National Aeronautics and Space Administration John H. Glenn Research Center at Lewis Field Cleveland, Ohio 44135-3191	E-17294

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSORING/MONITOR'S ACRONYM(S)
National Aeronautics and Space Administration Washington, DC 20546-0001	NASA
	11. SPONSORING/MONITORING REPORT NUMBER NASA/TM-2010-216358

12. DISTRIBUTION/AVAILABILITY STATEMENT
Unclassified-Unlimited
Subject Category: 07
Available electronically at http://gltrs.grc.nasa.gov
This publication is available from the NASA Center for AeroSpace Information, 443-757-5802

13. SUPPLEMENTARY NOTES

14. ABSTRACT
Aircraft engine performance trend monitoring and gas path fault diagnostics are closely related technologies that assist operators in managing the health of their gas turbine engine assets. Trend monitoring is the process of monitoring the gradual performance change that an aircraft engine will naturally incur over time due to turbomachinery deterioration, while gas path diagnostics is the process of detecting and isolating the occurrence of any faults impacting engine flow-path performance. Today, performance trend monitoring and gas path fault diagnostic functions are performed by a combination of on-board and off-board strategies. On-board engine control computers contain logic that monitors for anomalous engine operation in real-time. Off-board ground stations are used to conduct fleet-wide engine trend monitoring and fault diagnostics based on data collected from each engine each flight. Continuing advances in avionics are enabling the migration of portions of the ground-based functionality on-board, giving rise to more sophisticated on-board engine health management capabilities. This paper reviews the conventional engine performance trend monitoring and gas path fault diagnostic architecture commonly applied today, and presents a proposed enhanced on-board architecture for future applications. The enhanced architecture gains real-time access to an expanded quantity of engine parameters, and provides advanced on-board model-based estimation capabilities. The benefits of the enhanced architecture include the real-time continuous monitoring of engine health, the early diagnosis of fault conditions, and the estimation of unmeasured engine performance parameters. A future vision to advance the enhanced architecture is also presented and discussed.

15. SUBJECT TERMS
Systems health monitoring; Gas turbine engines; Fault detection; Kalman filters

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON STI Help Desk (email:help@sti.nasa.gov)
a. REPORT U	b. ABSTRACT U	c. THIS PAGE U	UU	23	19b. TELEPHONE NUMBER *(include area code)* 443-757-5802