Research Article

Assessment of an Anomaly Detector for Jet Engine Health Monitoring

Sebastien Borguet and Olivier Léonard

Turbomachinery Group, University of Liège, Sart-Tilman Campus B52/3, 4000 Liège, Belgium

Correspondence should be addressed to Sebastien Borguet, s.borguet@ulg.ac.be

Received 7 July 2010; Accepted 11 January 2011

Academic Editor: Robert Hendricks

Copyright © 2011 S. Borguet and O. Léonard. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

The goal of module performance analysis is to reliably assess the health of the main components of an aircraft engine. A predictive maintenance strategy can leverage this information to increase operability and safety as well as to reduce costs. Degradation undergone by an engine can be divided into gradual deterioration and accidental events. Kalman filters have proven very efficient at tracking progressive deterioration but are poor performers in the face of abrupt events. Adaptive estimation is considered as an appropriate solution to this deficiency. This paper reports the evaluation of the detection capability of an adaptive diagnosis tool on the basis of simulated scenarios that may be encountered during the operation of a commercial turbofan engine. The diagnosis tool combines a Kalman filter and a secondary system that monitors the residuals. This auxiliary component implements a generalised likelihood ratio test in order to detect abrupt events.

1. Introduction

Predictive maintenance aims at scheduling overhaul actions on the basis of the actual level of engine deterioration. The benefits are improved operability and safety as well as reduced life cycle costs. Generating reliable information about the health condition of the gas turbine is therefore a requisite and has been the subject of intensive research in the community.

In this paper, *module performance analysis* is considered. Its purpose is to detect, isolate, and quantify the changes in engine module performance, described by so-called health parameters, on the basis of measurements collected along the gas path of the engine [1]. Typically, the health parameters are correcting factors on the efficiency and the flow capacity of the modules (fan, lpc, hpc, hpt, lpt) while the measurements are intercomponent temperatures, pressures, and shaft speeds.

Figure 1 sketches a typical degradation profile of fan efficiency versus engine usage time. As far as time scale is considered, alterations in engine health can be split into two groups. On one hand, *gradual deterioration* (due to erosion, corrosion, or fouling for instance) occurs during normal

operation of the engine and affects all major components at the same time. On the other hand, *accidental events*, caused for instance by foreign object damage (FOD) or hot restarts, impact one (at most two) component(s) at a time and occur infrequently. As depicted in Figure 1, occurrence of such an abrupt fault often results in an unscheduled maintenance action, and, therefore, these events should be detected and addressed in a timely manner.

Among the numerous techniques that have been investigated to monitor the performance of an engine, see [2] for a detailed survey, the popular Kalman filter [3] has received special attention. Initially devised as a minimum variance estimator of the states of a dynamic system, the Kalman filter is seen here as a recursive algorithm for the identification of the health parameters. The Kalman filter has proven its capability to track gradual deterioration such as engine wear with good accuracy. Indeed, the Kalman filter embeds a transition model that describes a “relatively slow” and smooth evolution of the health parameters. On the other hand, a sluggish Kalman filter response is observed in the face of rapid variations of the engine condition, leading to a long delay in recognising such short-time-scale events.

One way to tackle this problem is to reconsider it in the realm of *adaptive estimation* [4]. Basically, the idea consists of an online tweaking of the bandwidth of the filter in order to improve its behaviour with respect to rapid degradation. Willsky and Jones [5] have proposed an appealing solution that relies on a modified transition model for the system states in order to account for possible “jumps” in these states. This technique is the backbone of an adaptive diagnosis tool developed by the authors in [6]. The resulting algorithm combines a Kalman filter, which relies on the assumption of a smooth variation of the health parameters, and a secondary system that monitors the residuals. This auxiliary component implements a *generalised likelihood ratio test* in order to detect short-time-scale events. As a result, the adaptive algorithm provides not only the same performance as the standard Kalman filter under normal operation (long-time-scale deterioration), but also reduced detection delay of accidental events.

In the present article, an extensive assessment of the detection capability of the adaptive diagnosis tool is reported. Simulated scenarios are representative of degradation profiles that can be encountered on a commercial aircraft engine. The library of abrupt events encompasses module faults, system faults, and instrumentation faults as well as unreported maintenance actions such as compressor water-washes. The performance of the adaptive tool is evaluated in terms of false alarm and missed detection rates, and detection delay.

2. Description of the Method

The scope of this section is to provide the mathematical background of the adaptive algorithm. The generic diagnosis tool relying on a Kalman filter is briefly presented. Then, based on a model that can handle abrupt variations in the health parameters, the auxiliary component that performs the detection of the “jumps” is introduced. Finally, the integration of the adaptive component with the Kalman filter is described.

2.1. Simulation Model. One of the key components in module performance analysis is a model of the jet engine. Considering steady-state operation of the gas turbine, these simulation tools are generally nonlinear aerothermodynamic models based on mass, energy, and momentum conservation laws applied to the engine flow path. Equation (1) represents such an engine model, where k is a discrete time index, $\mathbf{u}_k$ are the parameters defining the operating point of the engine (e.g., fuel flow, altitude, and Mach number), $\mathbf{w}_k$ are the health parameters, and $\mathbf{y}_k$ are the gas path measurements.

$$\mathbf{y}_k = \mathcal{G}(\mathbf{u}_k, \mathbf{w}_k). \quad (1)$$

Module performance analysis is a relative approach in the sense that it assesses the changes in engine performance relative to some reference level. Accordingly, the quantity of interest is the difference between the actual engine health condition and a reference one. In the recursive approach that is used here, this reference value is represented by a so-called prior value $\hat{\mathbf{w}}_k^-$ which designates a value of the health

FIGURE 1: Typical fan efficiency evolution showing gradual and abrupt changes.

parameters that is available before the measurements $\mathbf{y}_k$ are observed. Applying a Taylor series expansion around this prior value to the function $\mathcal{G}(\mathbf{u}_k, \mathbf{w}_k)$ and truncating to the first order, (1) becomes

$$\mathbf{y}_k = \mathcal{G}(\mathbf{u}_k, \hat{\mathbf{w}}_k^-) + \mathbf{G}_k(\mathbf{w}_k - \hat{\mathbf{w}}_k^-), \quad (2)$$

where the influence coefficient matrix $\mathbf{G}_k$ is evaluated at the prior value $\hat{\mathbf{w}}_k^-$

$$\mathbf{G}_k \triangleq \left. \frac{\partial \mathcal{G}(\mathbf{u}_k, \mathbf{w}_k)}{\partial \mathbf{w}_k} \right|_{\mathbf{w}_k = \hat{\mathbf{w}}_k^-}. \quad (3)$$

A random vector $\boldsymbol{\epsilon}_k \in \mathcal{N}(\mathbf{0}, \mathbf{R}_y)$ is added to the deterministic linearised engine model (2) to account for sensor inaccuracies and modelling errors. Equation (4) is therefore termed the statistical model.

$$\mathbf{y}_k = \mathcal{G}(\mathbf{u}_k, \hat{\mathbf{w}}_k^-) + \mathbf{G}_k(\mathbf{w}_k - \hat{\mathbf{w}}_k^-) + \boldsymbol{\epsilon}_k, \quad (4)$$

$$\Leftrightarrow \mathbf{r}_k = \mathbf{G}_k(\mathbf{w}_k - \hat{\mathbf{w}}_k^-) + \boldsymbol{\epsilon}_k, \quad (5)$$

where $\mathbf{r}_k$ are the *a priori* residuals defined as

$$\mathbf{r}_k \triangleq \mathbf{y}_k - \mathcal{G}(\mathbf{u}_k, \hat{\mathbf{w}}_k^-). \quad (6)$$

2.2. Kalman Filter-Based Diagnostics. The tenet behind estimation of the health parameters relies on (5). The residuals are a measure of the discrepancy between the actual measurements taken on the engine and the value predicted with the performance model. The purpose of the estimator is to adjust the health parameters so as to cancel the residuals on average.

The estimator used in this study is the celebrated Kalman filter, a recursive algorithm first developed for the estimation of the state variables of a dynamic system. In that framework, both the measurements and the health parameters are considered as Gaussian random vectors, which are completely described by their respective mean value and covariance matrix. At time index k , the update rule for the mean value of the health parameters is a linear function of the measurements through

$$\hat{\mathbf{w}}_k = \hat{\mathbf{w}}_k^- + \mathbf{K}_k(\mathbf{y}_k - \mathcal{G}(\mathbf{u}_k, \hat{\mathbf{w}}_k^-)). \quad (7)$$

The first term on the right-hand side of (7) is the *a priori* estimate of the health parameters, obtained from past data

- (1) $\hat{\mathbf{w}}_k^- = \hat{\mathbf{w}}_{k-1}$ and $\mathbf{P}_{\mathbf{w},k}^- = \mathbf{P}_{\mathbf{w},k-1} + \mathbf{Q}_k$
- (2) acquire $\mathbf{u}_k$ and $\mathbf{y}_k$
- (3) $\mathbf{r}_k = \mathbf{y}_k - \mathcal{G}_k(\mathbf{u}_k, \hat{\mathbf{w}}_k^-)$
- (4) compute Jacobian matrix, $\mathbf{G}_k$, as per(3)
- (5) $\mathbf{P}_{\mathbf{y},k} = (\mathbf{G}_k \mathbf{P}_{\mathbf{w},k}^- \mathbf{G}_k^T + \mathbf{R}_y)$
- (6) $\mathbf{K}_k = \mathbf{P}_{\mathbf{w},k}^- \mathbf{G}_k^T \mathbf{P}_{\mathbf{y},k}^{-1}$
- (7) $\hat{\mathbf{w}}_k = \hat{\mathbf{w}}_k^- + \mathbf{K}_k \mathbf{r}_k$ and $\mathbf{P}_{\mathbf{w},k} = (\mathbf{I} - \mathbf{K}_k \mathbf{G}_k) \mathbf{P}_{\mathbf{w},k}^-$

ALGORITHM 1: Basic steps of the extended Kalman filter.

up to time index $k - 1$. To generate these *a priori* values, a model describing the temporal evolution of the parameters must be supplied. Generally, little information is available about the way the engine degrades, which motivates the choice of a random walk model

$$\mathbf{w}_k = \mathbf{w}_{k-1} + \mathbf{v}_k, \quad (8)$$

where the random vector $\mathbf{v}_k \in \mathcal{N}(\mathbf{0}, \mathbf{Q}_k)$ is the so-called process noise that provides some adaptability to track a time-evolving fault.

The second term on the right-hand side is a correction term that accounts for the information contained in the latest data sample $\mathbf{y}_k$. This corrective action is proportional to the *a priori* residuals defined in (6). The gain matrix $\mathbf{K}_k$ is selected so as to minimise the *a posteriori* covariance matrix of the health parameters defined as

$$\mathbf{P}_{\mathbf{w},k} \triangleq E\{(\mathbf{w}_k - \hat{\mathbf{w}}_k)(\mathbf{w}_k - \hat{\mathbf{w}}_k)^T\}, \quad (9)$$

where E is the mathematical expectation operator.

Algorithm 1 summarises, in a pseudocode style, the basic processing steps of the extended Kalman filter. This algorithm has a predictor-corrector structure and involves only basic linear algebra operations. On line 1, prediction of the prior values of the health parameter distribution are made through the transition model (8). Then the data are acquired and used for building the *a priori* residuals (lines 2 and 3). The Jacobian matrix is assessed on line 4 and subsequently used in the computation of the covariance matrix of the residuals $\mathbf{P}_{\mathbf{y},k}$ (line 5) and of the Kalman gain $\mathbf{K}_k$ (line 6). Finally, the *a posteriori* distribution is assessed at the corrector step (line 7). Loosely speaking, the Kalman gain controls the contribution of the residuals to the *a posteriori* estimate. If the prior uncertainty in the parameters (represented by matrix $\mathbf{P}_{\mathbf{w},k}^-$) is low with respect to the uncertainty in the residuals (represented by matrix $\mathbf{P}_{\mathbf{y},k}$), then the Kalman gain is small and the residuals do not contribute much to the estimate. In the opposite case, the Kalman gain is large and the *a posteriori* estimate relies more on the residuals. To complete the picture, the block diagram in Figure 2 shows the closed-loop, predictor-corrector structure. The interested reader is directed to [7] for a comprehensive derivation and complementary details.

2.3. Incorporating Adaptability in the Diagnosis Tool. To improve the tracking abilities of abrupt events without

FIGURE 2: Health parameter update mechanism using an extended Kalman filter.

sacrificing the reliability of the estimation of long-time-scale deterioration, adaptive estimation is considered. The approach is based on the assumption that abrupt events may occur, but that they occur *infrequently*. This means that the transition model (8) is valid most of the time, except in the case of anomalies. This assumption seems a reasonable one for the intended application of engine trend monitoring, as an accidental event does not occur—fortunately—on each flight.

Inspired by the work of Willsky and Jones [5], the core of the adaptive algorithm consists of a Kalman filter, which relies on the assumption of a smooth variation of the engine condition. A secondary system that monitors the residuals of the Kalman filter complements the design. It implements a *generalised likelihood ratio* (GLR) test in order to detect rapid events. The milestones of the technique are reported below, the interested reader can find a detailed development in [6].

The root of the adaptive algorithm is an enhanced transition model of the health parameters that accounts for possible abrupt events

$$\mathbf{w}_k = \mathbf{w}_{k-1} + \mathbf{v}_k + \Delta_{\mathbf{w}} \delta_{\tau,k}, \quad (10)$$

where

- (i) $\Delta_{\mathbf{w}}$ is a vector modelling the jump,
- (ii) τ is a positive integer that represents its time of occurrence, and
- (iii) $\delta_{i,j}$ is the Kronecker delta operator.

Note that $\Delta_{\mathbf{w}}$ and τ are regarded here as *unknown* parameters and not as random variables, which means that no prior distribution is attached to them.

The strategy of adaptive estimation comes from viewing the new state-space model (10) and (4) according to two different hypotheses:

- (i) H_0 : no jump up to now ($\tau > k$),
- (ii) H_1 : a jump has already occurred ($\tau \leq k$).

Under assumption H_0 (no jump), the Kalman filter provides an optimal estimation of the health parameters in the least-squares sense. Under assumption H_1 , the residuals $\mathbf{r}_k$ can be expressed as a function of the jump characteristics τ and $\Delta_{\mathbf{w}}$. Given the first-order approximation on the measurement equation, the residuals $\mathbf{r}_k$ can be split into two terms

$$\mathbf{r}_k = \mathbf{r}_{k,H_0} + \mathbf{H}_{k,\tau} \Delta_{\mathbf{w}}, \quad (11)$$

where $\mathbf{r}_{k,H_0}$ are the residuals in the no-jump case, distributed as $\mathcal{N}(\mathbf{0}, \mathbf{P}_{y,k})$, and the second term represents the influence of a jump $\Delta_{\mathbf{w}}$ that has occurred at time τ on the residuals at time k . Matrix $\mathbf{H}_{k,\tau}$ is computable from the enhanced state-space model (10) and (4) and the equations of the Kalman filter in Algorithm 1, see [6] for further details.

2.4. The GLRT as an Event Detector. In order to determine which hypothesis between H_0 and H_1 is true, a GLR test (see [8]) is applied. In short, it is a statistical test in which a ratio is computed between the maximum probability of a result under two different hypotheses, so that a decision can be made between them based on the value of this ratio.

As highlighted in [9], the original implementation of the GLRT algorithm involves storage and computational resources that grow linearly over time. To keep the problem tractable, the jump detection is therefore restricted to a sliding window of width M . Provided the window is sufficiently wide to ensure detection of all major events, this is a reasonable approximation.

Essentially, the procedure consists of a first step that computes the maximum likelihood estimates $\hat{\tau}$ and $\hat{\Delta}_{\mathbf{w}}$ from the residuals $\mathbf{r}_{k-M}, \dots, \mathbf{r}_k$ assuming H_1 is true. These values are then substituted into the likelihood ratio test for H_1 versus H_0 . All probability densities being Gaussian, the log-likelihood ratio becomes

$$l_{k,\tau} = \mathbf{d}_{k,\tau}^T \mathbf{C}_{k,\tau}^{-1} \mathbf{d}_{k,\tau}, \quad k - M < \tau \leq k, \quad (12)$$

where matrix $\mathbf{C}_{k,\tau}$ is deterministic and *does not depend* on the data while vector $\mathbf{d}_{k,\tau}$ is a linear combination of the residuals

$$\begin{aligned} \mathbf{C}_{k,\tau} &= \sum_{j=\tau}^k \mathbf{H}_{j,\tau}^T \mathbf{P}_{y,j}^{-1} \mathbf{H}_{j,\tau}, \\ \mathbf{d}_{k,\tau} &= \sum_{j=\tau}^k \mathbf{H}_{j,\tau}^T \mathbf{P}_{y,j}^{-1} \mathbf{r}_j. \end{aligned} \quad (13)$$

These two equations show that the likelihood ratio (12) actually implements a matched filter, that is, a correlation test between the variations in the residuals and the signature of a jump, represented by $\mathbf{H}_{k,\tau}$.

The value $\hat{\tau}$ that maximises $l_{k,\tau}$ represents the most likely time at which a jump occurred during the last M time steps. The decision rule to choose between H_0 and H_1 is

$$\begin{aligned} &H_1, \\ &l_{k,\hat{\tau}} \geq \eta, \\ &H_0. \end{aligned} \quad (14)$$

FIGURE 3: Integration of the GLR detector with the Kalman filter.

There is a direct relation between the threshold η and the probability of false alarm P_F in jump detection through

$$P_F = \int_{\eta}^{\infty} p(l = L | H_0) dL, \quad (15)$$

where $p(l = L | H_0)$ is the probability density of $l_{k,\tau}$ conditioned on H_0 which is a central chi-squared density with n degrees of freedom, see [8] for a proof. Practically, (15) is inverted numerically (e.g., with the function `chi2inv` in Matlab) to obtain the threshold η from a prescribed probability of false alarm.

2.5. Implementation of the Adaptive Algorithm. Two parameters are available to tune the GLRT system: first, the threshold η (or equivalently the probability of false alarm P_F) in the hypothesis testing (14) as discussed previously, and, second, the width M of the sliding window. The selection of M is dictated by a tradeoff between accurate and fast detection of the events. The former implies to choose M large enough while the latter advocates a small-sized buffer.

Figure 3 depicts the integration of the adaptive component, which comprises all the elements in the dashed box, with the Kalman filter. For the sake of clarity, only the most relevant data streams are sketched in the diagram. Briefly explained, the adaptive component works as follows.

- (i) The “GLR” box updates the quantities $\mathbf{d}_{k,\tau}$ and $\mathbf{C}_{k,\tau}$ in the M -sized buffer. The likelihood ratio is then assessed through (12) and its maximum value $l_{k,\hat{\tau}}$ is searched for.
- (ii) This value is compared to the threshold η in order to determine whether a jump has occurred.
- (iii) In case hypothesis H_1 is true, a flag is issued that would subsequently trigger a fault isolation logic.

From a computational standpoint, it is worth noting the two following points. First, recursive relations can be derived for the quantities involved in the GLRT, see for instance [6]. Second, the aforementioned operations performed in the adaptive component are roughly equivalent to $M + 1$ runs of a linear Kalman filter (i.e., no call to the nonlinear engine

FIGURE 4: Turbofan layout with station numbering and health parameter locations.

model) per time step. So the increase in computational load is directly proportional to the size of the buffer. Nonetheless, the most demanding part of the whole adaptive algorithm lies in the evaluation of the Jacobian matrix G_k and the prediction of the measurements. This requires, in the case of centred-differences, $2n + 1$ calls to the nonlinear engine model. Consequently, the overhead in CPU time needed by the adaptive algorithm is rather limited (about 15%) for common window widths of $M \approx 10 - 20$.

3. Application of the Method

3.1. Engine Layout. The application used as a test case is a high bypass ratio, mixed-flow turbofan. The engine performance model has been developed in the frame of the OBIDICOTE project (a Brite/Euram project for on-board identification, diagnosis, and control of turbofan engine) and is detailed in [10]. A schematic of the engine is sketched in Figure 4 where the location of the health parameters and the station numbering are also indicated.

A total of 12 health parameters are considered to simulate engine deterioration. 10 of them are usual correction factors that determine the change in efficiency and flow capacity of each turbomachinery component with respect to a reference condition, see [11]. The last two represent deviations relative to the nominal schedule of the variable geometry devices, namely, variable stator vanes and blow-off valves. They model either a fault on the sensed actuator position or a fault on the actuator itself (e.g., mechanical failure). No health parameters are attached to the combustor because its deterioration does not cause significant changes in the engine performance, see [12, 13].

The sensor suite selected to perform the engine diagnostics is similar to the instrumentation available onboard contemporary turbofan engines and is detailed in Table 1 where the nominal accuracy (uncertainty is three times the standard deviation σ) of each sensor is also quoted. The table is complemented with the sensors used to define the operating conditions of the engine.

The engine model has no built-in control system. As a consequence, the engine is run at a prescribed fuel flow and

TABLE 1: Gas-path and operating condition instrumentation.

Label	Description	Uncertainty
P13	Fan outlet total pressure	± 100 Pa
T26	lpc outlet total temperature	± 2 K
P3	hpc outlet total pressure	± 5000 Pa
T3	hpc outlet total temperature	± 2 K
N_{lp}	Fan speed	± 6 rpm
N_{hp}	Core speed	± 12 rpm
T49	lpt inlet total temperature	± 8 K
W_f	Fuel flow	± 5 g/s
P2	Fan inlet total pressure	± 100 Pa
T2	Fan inlet total temperature	± 2 K

TABLE 2: Operating envelope.

Name	Units	Lower bound	Upper bound
Fuel flow	kg/s	0.350	0.360
Altitude	m	10,058	11,278
Mach	—	0.78	0.82
ΔT_{ISA}	K	-10	+10

the variable geometry devices are set according to an open-loop schedule in the simulations.

3.2. Definition of the Scenarios. Inspired from [14], a routine dedicated to the generation of so-called scenarios has been written for the present study. Each scenario consists of a database of the sensed engine parameters given in Table 1 simulated along unique operating history and engine deterioration profile. The length of a scenario is arbitrarily set to 3000 flights.

The virtual data collection is performed during the cruise part of the flight. All operating points are randomly selected in the envelope defined in Table 2. The degradation profile is composed of a gradual wear of the engine modules plus a single event picked from a predefined library. The abrupt event is superimposed to the gradual deterioration profile at a flight drawn at random between flight 1450 and flight 1550. Complementary details about each category are given in the next two sections.

3.2.1. Gradual Deterioration. Progressive deterioration, representing engine wear, is modeled by altering the efficiency SEi and flow SWiR correcting factors of the engine modules (fan, lpc, hpc, hpt, lpt). Despite the decrease in performance it causes, gradual deterioration is not regarded as a faulty condition in this report, but rather as a normal mechanism induced by engine usage. This point of view is shared by several authors, see for example [14, 15].

An average profile was created by means of a linear + exponential law fitted to historical data available in the open literature [16, 17]. Moreover, random variations are added to the initial and final values of the 10 health parameters as well as to the shape parameters of the linear + exponential fit. These modifications account for various effects,

such as engine-to-engine manufacturing variations and more/ less rapid and severe deterioration of each module.

3.2.2. Library of Abrupt Events. The events are picked up from the library summarised in Table 3. Most of the 20 classes impact only one entity which can be a component (FC's 1–5), a system (FC's 6–7) or a sensor (FC's 8–17). Unreported maintenance actions (FC's 18–20) impact one or several modules.

Turbomachinery *module faults* involve alterations in both the efficiency and flow correcting factors. During the generation of a scenario, two related quantities, namely the fault magnitude f_m and the fault ratio f_r , are used. They are uniformly distributed in the intervals quoted in Table 3. The fault magnitude f_m is defined as the Euclidian norm of the efficiency and flow variations. The fault ratio f_r , also termed coupling factor in the following, is defined as the ratio between the change in flow capacity and the change in efficiency

$$\begin{cases} f_m \stackrel{\text{def}}{=} \sqrt{\Delta \text{SEi}^2 + \Delta \text{SWiR}^2} \\ f_r \stackrel{\text{def}}{=} \frac{\Delta \text{SWiR}}{\Delta \text{SEi}} \end{cases} \Rightarrow \begin{cases} \Delta \text{SEi} = \frac{f_m}{\sqrt{1 + f_r^2}} \\ \Delta \text{SWiR} = f_r \cdot \Delta \text{SEi} \end{cases} \quad (16)$$

Once randomly drawn, f_m and f_r are converted back to deviations in the health parameters according to (16). For compressors, ΔSWiR and ΔSEi have the same sign while for turbines the coupling factor can be either positive or negative.

The vbv and vsv *system faults* are implemented as true off-schedule deviations. The uniformly distributed magnitude for these fault types is reported here as some kind of severity index, for the sake of simplicity. A unit value corresponds to a small modification with respect to the nominal setting (e.g., only a slight mistuning of the vsv) while a value of 5 hints at a deep malfunction (e.g., fully open vbv).

Instrumentation faults are modeled as biased readings from one sensor, either in the flow path or for the operating conditions. The magnitude, expressed in Table 3 in units of sensor standard deviation σ , and the sign of the bias are both randomly selected.

Finally, the last three events in Table 3 represent *maintenance actions* that might have been unreported by the personnel. Unlike the previous types, these three events lead to improvement in the engine performance. Nonetheless, they will also result in a shift in the sensed engine parameters that the algorithm should detect. Water-washes (event no. 18) are assumed to bring the performance of all compressor devices (fan, lpc, and hpc) back to their initial level. This is a somewhat idealised situation, as part of the deterioration is not recoverable with a simple water-wash [12]. Event no. 19, named hpt service, consists of the replacement of the hpt module and is supposed to lead to a restoration of the hpt performance to its initial level too. Event no. 20, named lpt service, is the counterpart of no. 19 for the lpt.

TABLE 3: List of considered abrupt events—FC stands for fault code.

FC	Description	Magnitude f_m	SW/SE ratio f_r
1	Fan fault	1 to 5%	+(0.5 to 2.0)
2	lpc fault	1 to 5%	+(0.5 to 2.0)
3	hpc fault	1 to 5%	+(0.5 to 2.0)
4	hpt fault	1 to 5%	$\pm(0.5 \text{ to } 2.0)$
5	lpt fault	1 to 5%	$\pm(0.5 \text{ to } 2.0)$
6	vbv fault	1 to 5	/
7	vsv fault	1 to 5	/
8	P13 bias	$\pm 1 \text{ to } 5\sigma$	/
9	T26 bias	$\pm 1 \text{ to } 5\sigma$	/
10	P3 bias	$\pm 1 \text{ to } 5\sigma$	/
11	T3 bias	$\pm 1 \text{ to } 5\sigma$	/
12	N_p bias	$\pm 1 \text{ to } 5\sigma$	/
13	N_{hp} bias	$\pm 1 \text{ to } 5\sigma$	/
14	T49 bias	$\pm 1 \text{ to } 5\sigma$	/
15	W_f bias	$\pm 1 \text{ to } 5\sigma$	/
16	P2 bias	$\pm 1 \text{ to } 5\sigma$	/
17	T2 bias	$\pm 1 \text{ to } 5\sigma$	/
18	Water-wash	Restore fan, lpc, and hpc perfs.	
19	hpt service	Restore hpt perfs.	
20	lpt service	Restore lpt perfs.	

3.2.3. Snapshot Generator. The adaptive diagnosis tool analyses data collected once per flight at cruise conditions. As explained in [14], the data acquisition system records measurements over a window of time and saves the averaged values for later exploitation. In an attempt to mimic this on-board archival of engine data, the snapshots are generated in the following way.

- (1) Select a random operating condition from the distribution specified in Table 2 and read the deterioration relative to the current flight.
- (2) Run the engine model to generate 25 data samples for these inputs, this number corresponds to a recording window of 2.5 seconds at a sample frequency of 10 Hz, which are common values (see [18]).
- (3) Add Gaussian noise, whose magnitude is specified in Table 1, to the noise-free simulated measurements. In case of a sensor fault, also add the bias to the faulty sensor.
- (4) Average the readings and store them in the database. This last step provides a first decrease in the noise level of the measurements.

3.3. Selected Metrics. The present work is focused on the detection part of the diagnosis problem. Two metrics, recommended in [14], have been selected appropriately to assess the detection capability of the adaptive diagnosis tool. For the sake of completeness, the metrics are briefly introduced below. The reader may recall that gradual deterioration is not considered here as an event due to its continuous nature. An event is any of the types specified in Table 3.

TABLE 4: Sample detection decision matrix (DDM).

		Predicted state	
		Event	No event
True state	Event	True Positives	False Negatives
	No event	False Positives	True Negatives

The first metric is the *detection decision matrix* (DDM). As shown in Table 4, it is a square matrix of dimension two. The elements on the main diagonal reflect correct predictions (the predicted and true states are the same). False negatives are cases where an event is not detected by the algorithm. For this reason, they are also called *missed detections*. False positives are cases where the algorithm detects a nonexistent event. This situation is known as a *false alarm*.

As by-products of this matrix, the true positive rate (TPR) and false positive rate (FPR) are given by

$$\text{TPR} = \frac{\text{TP}}{\text{TP} + \text{FN}}, \quad \text{FPR} = \frac{\text{FP}}{\text{FP} + \text{TN}}. \quad (17)$$

The second metric is the *detection delay*, which is defined as the time required to detect an event after its initiation. It is desirable to have a minimum detection delay in order to take the corrective maintenance actions as soon as possible.

3.4. Results. In the present study, 100 scenarios have been generated for each of the 20 events listed in Table 3. Such a number allows a fair coverage of the fault pattern for module faults (both in magnitude and coupling factor), as well as for system and sensor faults. Another batch of 2000 no-event scenarios (which are composed of gradual deterioration only according to our convention) has also been processed to assess in a meaningful way the DDM and its related metrics.

The adaptive diagnosis tool is configured to estimate only the efficiency SE_i and flow SW_iR correcting factors of the turbomachinery components. Indeed, the main purpose of the Kalman filter is to track the gradual deterioration of the engine. Neither the systems nor the sensors are expected to exhibit this long-time-scale degradation trend. The tuning parameters of the anomaly detector are set to the following values: a sliding window of width $M = 15$ flights and a probability of false alarm $P_F = 10^{-5}$. As will be shown below, these settings were found to lead to satisfactory results for the scenarios under investigation.

Figure 5 depicts the tracking of pure engine wear (i.e., no-event scenario), with the adaptive diagnosis tool. All subplots but the bottom-right one show the true and estimated health parameters, expressed in terms of percentage of deviation with respect to a reference value that is equal to one for all parameters. On the abscissa, time is expressed in terms of flights. It can be seen that the identified values are in good agreement with the true ones, especially for the lpc, hpc and hpt. The estimation error is slightly larger for the health parameters of the fan and the lpt and especially for their efficiencies SE₁₂ and SE₄₉. It is worth noting that both the lpc and the lpt have much slower deterioration rates

TABLE 5: Detection decision matrix.

		Predicted state		
		Event	No event	
True state	Event	1672	328	TPR = 83.6%
	No event	0	2000	FPR = 0.0%

than the other three components. This observation will be recalled later. The subplot in the bottom-right corner shows the output signal from the event detector which remains at zero (i.e., no event detected) during the whole scenario.

As far as event detection is concerned, the global performance of the adaptive tool is summarised in the detection decision matrix given in Table 5. First, it can be seen that the false positive rate is equal to zero, which means that *no false alarm* was issued over the 2000 no-event scenarios processed in this study. As a consequence, when a detection flag is raised, it can be concluded with a high confidence that an anomaly indeed occurred on the engine. Looking at the other row of the DDM, it can be seen that the number of missed detections amounts to 328 out of 2000 faulty scenarios. This translates in an encouraging value of 83.6% for the true positive rate. The next step is to analyse the performance for each type of event separately.

Table 6 provides an overview of the detection capability for the different events. The third column gives the percentage of detected events for the said type. The mean detection delay is reported in the fourth column. This average value is computed from the detection delay, as defined in Section 3.3, of the detected cases. Finally, the fifth column gives the so-called “span”, defined as the difference between the maximum and minimum detection delays for a specific event. This quantity provides an adequate measure of the variability in the detection delay.

It can be seen that the PCD reaches 100% for all module faults (FC’s 1–5), meaning that for the range of magnitudes and coupling factors considered in this study, any abrupt deterioration of a component will be captured by the algorithm. The small-valued MDD’s hint at a rapid detection of the fault, actually one or two flights after its initiation for all modules but the lpc. For this component, the span is quite large—almost equal to the window length. Lpc faults of small magnitude are hence the most difficult to detect. On the contrary, hpt faults are the easiest to detect, whatever their structure (magnitude and coupling factor), this is confirmed by the related MDD and span.

The difference in the ability to detect the various component faults can be explained by taking a look at Figure 6. The graph shows the relative sensitivity of the sensor set with respect to a one-percent change in each of the health parameters. It can readily be seen that the hpt efficiency factor SE₄₁ has the largest impact on the measurements while the lpc flow capacity factor SW₂R has the lowest one.

Detection of system faults (FC’s 6–7) is successful over their whole range of severity index. This is confirmed by the PCD’s of 100% and the low values of MDD and span. Unlike vsv faults, vbv faults show a detection delay that depends on its intensity.

FIGURE 5: Tracking of engine wear with the adaptive diagnosis tool—in the legends, “est” stands for estimated.

The detection performance of gas-path instrumentation faults (FC's 8–14) looks worse. Indeed, the PCD ranges between 55% and 79% depending on the sensor. The MDD's rise up to 5–6 flights, but at the same time, the span is rather large so that the detection delay of gas-path sensor faults is highly dependent on its magnitude.

Timely detection of operating condition instrumentation faults (FC's 15–17) is partly better than that for gas-path sensor faults. Detection of a biased W_f sensor is the most effective, with PCD, MDD, and span values comparable to what is obtained for a fan fault for instance. Malfunction of the T_2 sensor is also quite efficiently caught by the adaptive diagnosis tool. The three detection metrics for this fault type are of the same order as for a vbv fault, which means that the detection delay is dependent on the magnitude of the bias. The situation for P_2 is nearly the same as for the gas-path sensors.

Finally, the behaviour of the event detector with respect to the unreported maintenance actions can be easily

FIGURE 6: Relative sensitivity of the measurements with respect to the health parameters.

explained from the very nature of these events. Indeed, a water-wash implies a simultaneous modification in the condition of the fan, lpc, and hpc. It can, therefore, be considered as a combined anomaly on these three modules.

TABLE 6: Detection performance for the different events.

FC	Description	PCD	MDD	Span
1	Fan fault	100	1.23	1
2	lpc fault	100	1.83	11
3	hpc fault	100	1.13	1
4	hpt fault	100	1.00	0
5	lpt fault	100	1.04	1
6	vbv fault	100	1.73	5
7	vsv fault	100	1.00	0
8	$P13$ bias	55	5.95	11
9	$T26$ bias	78	5.41	13
10	$P3$ bias	79	6.15	13
11	$T3$ bias	75	5.69	13
12	N_{lp} bias	58	5.09	13
13	N_{hp} bias	60	5.48	12
14	$T49$ bias	57	6.17	12
15	W_f bias	100	1.29	1
16	$P2$ bias	69	4.19	13
17	$T2$ bias	100	2.02	7
18	Water-wash	100	1.00	0
19	hpt service	100	1.00	0
20	lpt service	41	7.24	11

PCD: percent correctly detected, MDD: mean detection delay.

As abrupt faults on compressors are very efficiently detected by the algorithm, so is a water-wash. Event no. 19 can be thought of as an anomaly on the hpt. As a consequence, an hpt service is as detectable as an hpt fault. A similar reasoning applied to event no. 20 would lead to detection performance of an lpt service as good as that for an lpt fault. Looking at the last line of Table 6, the metrics are, however, worse than expected, with a PCD of 41% and quite high values for the MDD and span. These bad numbers can be explained by the fact that the lpt deteriorates quite slowly, as has already been pointed out when analysing Figure 5. In most cases, the magnitude of the event caused by an lpt replacement is so low that either it is not detected, or it is detected after several flights.

To conclude the analysis of the results, it is interesting to come back to instrumentation faults. As previously mentioned, the detection rate wavers between 55% and 79%. This means that some sensor faults are not detected. Table 7 reports the minimum level of sensor bias (negative ones in the central column, positive ones in the right column) that was successfully detected by the algorithm. The values quoted in the table are normalised with the respective standard deviation of the sensor and are obtained from the processing of the batch of scenarios previously defined.

First, it can be seen that the bounds are almost symmetric for positive and negative biases. The bounds for $P13$, N_{lp} , N_{hp} , and $T49$ are on the order of three standard deviations, which seems a logical value given the assumption of a Gaussian measurement noise. The bounds for $T26$, $P3$ and $T3$ are on the order of two standard deviations. Considering

TABLE 7: Minimum levels of bias (expressed in number of standard deviations σ) for sensor fault detection.

Sensor	Negative bias	Positive bias
$P13$	-2.72	2.95
$T26$	-1.82	1.91
$P3$	-2.10	2.26
$T3$	-2.27	2.00
N_{lp}	-3.22	2.93
N_{hp}	-2.60	2.70
$T49$	-2.90	2.81
W_f	-1.23	1.15
$P2$	-2.48	2.05
$T2$	-1.03	1.06

the instrumentation related to the operating point of the engine, the bounds for W_f and $T2$ are in the vicinity of one standard deviation. As all instances of these faults were effectively detected, the values quoted in Table 7 are actually the minimum values for these biases in the batch of simulated scenarios.

4. Discussion

The analysis of the results has illustrated the good performance of the adaptive diagnosis tool as far as fault detection is concerned. It should nonetheless be realised that the algorithm has processed simulated data, which are always “better looking” than true operational data. In the remainder of this section, some ideas that may lead to complementary work are discussed.

A first idea is to perform a parametric study to assess the influence of the two tuning parameters of the anomaly detector, namely, the window length M and the probability of false alarm P_F , on the performances of the adaptive algorithm, both at the global level, through the detection decision matrix, and at the local level, through the percent correctly detected and the mean detection delay.

A second axis for further development of the anomaly detector would be to embed some robustness against outliers in the data. So far, the detection flag is issued as soon as the likelihood ratio exceeds some threshold value. In the presence of “spiky” data samples, this could cause false alarms. The addition of fault persistency logic when a fault is detected could be a convenient solution.

The adaptive diagnosis tool evaluated in this paper performs a double task: estimation of gradual deterioration and anomaly detection. Once an event is detected, the next piece of valuable information to generate is a localisation of the event. The authors have presented, in another contribution [19], a sparse estimation tool dedicated to fault isolation. A third possibility for future work is, therefore, the combination of the adaptive Kalman filter used in this paper with the sparse estimation tool to offer a complete solution for performance monitoring of jet engines.

5. Conclusions

In this contribution, the detection capability of an adaptive algorithm for engine health monitoring has been assessed. The diagnosis tool combines a Kalman filter, which provides an accurate estimation of the health condition for long-time-scale deterioration (such as engine wear), and an adaptive component which monitors the residuals and looks for abrupt changes in the health condition. The adaptive component relies on a generalised likelihood ratio test to detect rapid variations in the engine condition.

The performance of the adaptive algorithm has been evaluated in terms of detection decision matrix, detection delay, and complementary metrics from the processing of a large number of degradation scenarios that may occur during the operational life of a commercial turbofan. Each scenario combines gradual deterioration and an abrupt event picked from a library including component, system, and instrumentation faults as well as unreported maintenance actions.

Nomenclature

Abbreviations

$\hat{\cdot}$:	Estimated value
$\cdot^-$:	Prior value
FC:	Fault code
GLRT:	Generalised likelihood ratio test
hpc:	High pressure compressor
hpt:	High pressure turbine
lpc:	Low pressure compressor
lpt:	Low pressure turbine
vbv:	Variable bleed valves behind the lpc
vsv:	Variable stator vanes on the hpc
$\mathcal{N}(\mathbf{m}, \mathbf{R})$:	A Gaussian probability density function with mean $\mathbf{m}$ and covariance matrix $\mathbf{R}$.

Scalars

k :	Discrete time index
l :	Likelihood ratio
m :	Number of monitored gas path variables
n :	Number of health parameters
N :	Rotational speed
P_i :	Total pressure at station i
P_F :	Probability of false alarm
SE $_i$:	Efficiency factor of the component whose entry is located at station i
SW $_i$:	Flow capacity factor of the component whose entry is located at station i
T_i :	Total temperature at station i
η :	Detection threshold
σ :	Standard deviation of a measurement
τ :	Time of occurrence of the abrupt event.

Vectors and Matrices

$\mathbf{G}$:	Influence coefficient matrix
$\mathbf{K}$:	Gain matrix of the Kalman filter

$\mathbf{P}$:	Covariance matrix
$\mathbf{Q}$:	Covariance matrix of the process noise
$\mathbf{R}_y$:	Covariance matrix of the measurement noise
$\mathbf{r}$:	Vector of residuals
$\mathbf{u}$:	Vector of control parameters
$\mathbf{w}$:	Vector of health parameters
$\mathbf{y}$:	Vector of monitored gas path variables
$\Delta_{\mathbf{w}}$:	Abrupt variation of the health parameters
ϵ :	Random vector of measurement noise
$\mathbf{v}$:	Random vector of process noise.

Acknowledgment

The authors wish to thank Mr. Donald L. Simon, with the NASA Glenn Research Center, for his input in defining the scenarios.

References

- [1] A. J. Volponi, "Foundation of gas path analysis (part i and ii)," in *von Karman Institute Lecture Series, number 01 in Gas Turbine Condition Monitoring and Fault Diagnosis*, 2003.
- [2] L. Marinai, D. Probert, and R. Singh, "Prospects for aero gas-turbine diagnostics: a review," *Applied Energy*, vol. 79, no. 1, pp. 109–126, 2004.
- [3] R. E. Kalman, "A new approach to linear filtering and prediction problems," *Journal of Basic Engineering*, vol. 82, pp. 35–44, 1960.
- [4] R. K. Mehra, "Approaches to adaptive filtering," *IEEE Transactions on Automatic Control*, vol. 17, no. 5, pp. 693–698, 1972.
- [5] A. Willsky and H. Jones, "A generalized likelihood ratio approach to state estimation in linear systems subject to abrupt changes," in *Proceedings of the 13th IEEE Conference on Decision and Control*, pp. 846–853, Phoenix, Ariz, USA, 1974.
- [6] S. Borguet and O. Léonard, "A generalized likelihood ratio test for adaptive gas turbine performance monitoring," *Journal of Engineering for Gas Turbines and Power*, vol. 131, no. 1, Article ID 011601, 2009.
- [7] S. M. Kay, *Fundamentals of Statistical Signal Processing*, Prentice Hall, London, UK, 1993.
- [8] H. L. van Trees, *Detection, Estimation and Modulation Theory*, John Wiley & Sons, New York, NY, USA, 1968.
- [9] A. S. Willsky and H. L. Jones, "A generalized likelihood ratio approach to the detection and estimation of jumps in linear systems," *IEEE Transactions on Automatic Control*, vol. 21, no. 1, pp. 108–112, 1976.
- [10] A. Stamatis, K. Mathioudakis, J. Ruiz, and B. Curnock, "Real-time engine model implementation for adaptive control and performance monitoring of large civil turbofans," in *ASME Turbo Expo*, New-Orleans, La, USA, 2001, 2001-GT-0362.
- [11] K. Mathioudakis, Ph. Kamboukos, and A. Stamatis, "Turbofan performance deterioration tracking using nonlinear models and optimization techniques," *Journal of Turbomachinery*, vol. 124, no. 4, pp. 580–587, 2002.
- [12] I. S. Diakunchak, "Performance deterioration in industrial gas turbines," *Journal of Engineering for Gas Turbines and Power*, vol. 114, no. 2, pp. 161–168, 1992.
- [13] S. O. T. Ogaji, S. Sampath, R. Singh, and S. D. Probert, "Parameter selection for diagnosing a gas-turbine's performance-deterioration," *Applied Energy*, vol. 73, no. 1, pp. 25–46, 2002.

- [14] D. L. Simon, J. Bird, C. Davison, A. Volponi, and R. E. Iverson, "Benchmarking gas path diagnostic methods: a public approach," in *Proceedings of the ASME Turbo Expo*, vol. 2, pp. 325–336, Berlin, Germany, 2008, GT2008-51360.
- [15] L. C. Jaw, "Recent advancements in aircraft engine health management (EHM) technologies and recommendations for the next step," in *Proceedings of the ASME Turbo Expo*, vol. 1, pp. 683–695, 2005, GT2005-68625.
- [16] G. P. Sallee, "Performance deterioration based on existing (historical) data—JT9D jet engine diagnostics program," Contractor Report CR-135448, NASA, 1978.
- [17] O. Sasahara, "JT9D engine/module performance deterioration results from back to back testing," in *Proceedings of the 7th International Symposium on Air Breathing Engines*, pp. 528–535, Beijing, China, September 1985.
- [18] D. Simon and D. L. Simon, "Aircraft turbofan engine health estimation using constrained Kalman filtering," in *ASME Turbo Expo*, pp. 485–492, June 2003, GT2003-38584.
- [19] S. Borguet and O. Léonard, "A sparse estimation approach to fault isolation," *Journal of Engineering for Gas Turbines and Power*, vol. 132, no. 2, Article ID 021601, 2010.