Wear Prognostic on Turbofan Engines

Jérôme Lacaille[1], Aurélie Gouby[1] and Olivier Piol[1]

[1]*Snecma, 77550 Moissy-Cramayel, France*

jerome.lacaille@snecma.fr
aurelie.gouby@snecma.fr
olivier.piol@snecma.fr

ABSTRACT

One of the most evident characteristic of wear for a turbofan engine is the exhaust gas temperature (EGT). It seems clear that this temperature increases when some carbon deposits on the turbine, when the compressor efficiency diminishes so the fuel flow should increase to produce the same amount of thrust, or even when some unbalance opens the spaces between the turbine and the casing. In any cases, an increase of the EGT should be analyzed because it is a wear symptom of the engine. It is mostly concluded by a water wash in the best case or a shop visit inspection and repair in the worst case. The engine manufacturer defines a schedule plan with its customer based on consumption of the EGT margin. This margin is the amount of available increase of the exhaust temperature before an inspection. Contractually, the engine is restored with a minimum EGT margin after each repair. Thus it is up to the manufacturer to understand how this margin is used to plan shop visits and to the company to estimate the current state of its engine. However the EGT measurement is subject to a lot of noise and the company regularly washes their engines to increase randomly the margin and their capabilities. In this article we present a simple, automatic and embeddable algorithmic method to transform the successive EGT measurements in a delay indicator computed after each flight giving the amount of available use time. One challenge is to take care of the random wash or repair executed by the user. Finally this indicator may be transmitted automatically with the other data broadcasted by the aircraft computer (ACMS/ACARS) and it is used by the manufacturer to prepare his shop logistic.

1. INTRODUCTION

Snecma is an engine manufacturer and produces turbofan for most of the short and medium range applications. For medium range specific fleets, aircrafts is doing around 4 flights per day whereas most of short range aircrafts are doing around 10 flights per day. During each flight, measurements are recorded on each engine. They are capitalized on the aircraft computer (ACMS) with other navigation observations such as the altitude, mach speed, etc. During each takeoff and a stable cruise phase a snapshot of measurement is broadcasted to the ground for monitoring.

Table 1. Short list of some snapshot measurements.

Name	Description
Index information	
AC_ID	*Aircraft ID*
ESN	*Engine Serial Number*
FL_DATE	*Flight Date*
Context information	
TAT	*External temperature*
ALT	*Altitude*
AIE	*Anti Ice Engine*
AIW	*Anti Ice Wings*
BLD	*Bleed valve position*
ISOV	*ECS Isolation Valve Position*
VBV	*Variable Bleed Valve Position*
VSV	*Variable Stator Vane Position*
HPTACC	*High Pressure Turbine Active Clearance Control*
LPTACC	*Low Pressure Turbine Active Clearance Control*
RACC	*Rotor Active Clearance Control*
ECS	*Environmental Control System*
TLA	*Thrust Lever Angle*
N1	*Fan Speed*
XM	*Mach Number*
Monitored measurements	
N2	*Core Speed*
FF	*Fuel Flow*
PS3	*Static pressure after compression*
T3	*Temperature after compresion*
EGT	*Exhaust Gas Temperature*

The ground operator follows each engine flight after flight. In fact there are three phases in the ground monitoring process as pictured on Figure 1 next page. The first one is a fast answer mainly based on FADEC error detections (the engine control computer); the second one is the trend monitoring we are interested in this article and which deals with successive flights of the same engine; and the last one is fleet monitoring (Lacaille & Come, 2011a) that compares engines to establish a prognostic on its internal state.

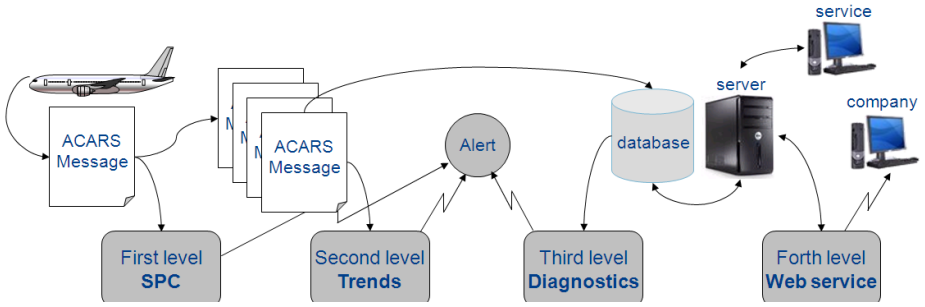

Figure 1. Three different diagnostic levels and a web service.

Trend monitoring algorithms look at successive snapshots of observations to help the fleet manager analyzing the wear trend of the engine. Some algorithms detect sudden changes of the engine behavior (Lacaille & Come, 2011b; Lacaille & Nya Djiki, 2009; Lacaille, 2009a) and classify the type of change (Bellas, Bouveyron, Cottrell, & Lacaille, 2012, 2013; Come, Cottrell, Verleysen, & Lacaille, 2010; Cottrell et al., 2009; Flandrois, Lacaille, Masse, & Ausloos, 2009; Lacaille, 2009b).

The EGT margin algorithm just looks at the exhaust temperature to predict a given drop and anticipate the need for a potential water wash or shop visit if the margin is really small.

2. MEASUREMENTS

The EGT margin is the difference between a maximum admissible value for the specific engine application and the observed temperature measured just before the exhaust nozzle. This value is given in °C and decreases progressively to zero. The maximum value corresponds to the certified maximal admissible temperature for an engine type. The value subtracted to this maximum threshold is an estimation of this maximal value for the current engine when measured at sea level with an external temperature equivalent to standard value plus 15°C during the most stressful moment of the takeoff. However, even if acquired with lot of care this measurement still depends on actual external conditions, engine thrust, aircraft speed. A normalization procedure is applied to suppress these last dependencies. This normalization is an analytical certified computation and a mathematical analysis (Lacaille, 2010) confirms the precision of this result.

For example, on Figure 2 the first plot shows the original measurement of an exhaust gas temperature during 300 successive flights. These measurements highly depend on external conditions like the flight mission: altitude, speed gross weight, localization (sea, land or desert) instant of acquisition and external conditions: weather, wind…

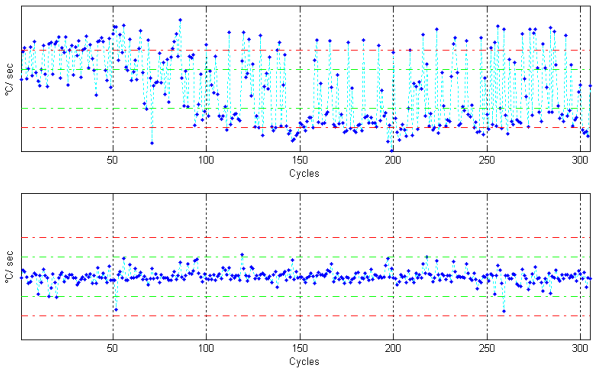

Figure 2. Normalization of the EGT measurement.

The two graphs have the same scale, they are centered on the EGT mean value; the top one is the original acquired measurements and the bottom one is the normalized data. Green dashed lined corresponds to 3σ bounds and red dashed lines to 6σ.

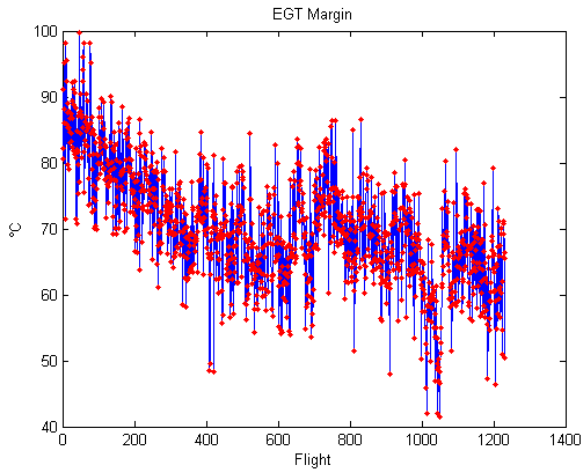

Figure 3. Plot of an exhaust gas temperature margin for an example engine.

A margin computation example is presented on Figure 3. It is based on normalized data. One observes that instead of continuously decreasing as we may have anticipated; the

signal is subject to random oscillations. To deal with this behavior we tried two approaches. The first one directly captures this phenomenon with a dynamic model of the temporal evolution. This algorithm is detailed in the next section. The second approach assumes that if the margin grows, it is because the airline decides an intervention on the engine: a water-wash for example. Then it may not be so important to deal with increasing margin and we just try to analyze the downward trends. This second solution is much simpler and may be efficient for the airline but not for the MRO to improve the shop logistic.

The next plot (Figure 4) presents a weekly smoothed version of the preceding signal. This signal clearly presents the oscillations of the margin over time. It seems that the increasing steps appear regularly but with different effects.

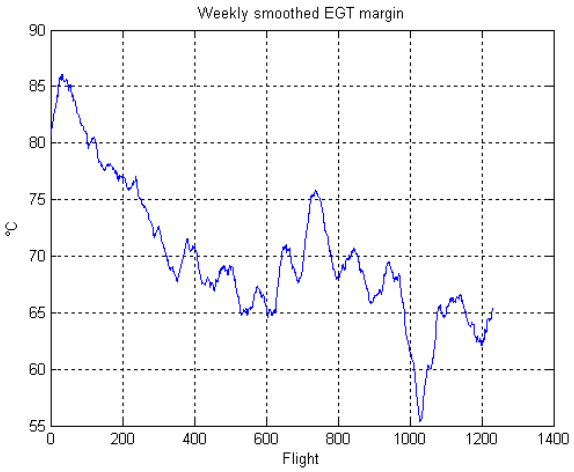

Figure 4. Moving average of the EGT margin measurement over a week.

The goal of the study is to estimate the probability to cross a minimum threshold before a given horizon h. In general this horizon corresponds to the necessary notice before an engine maintenance operation. This is the real need for maintenance operators and it is schematized on Figure 5 by a computation of the probability of detection (POD) after time $t+h$. We will also produce lower bound of the time left before crossing this threshold.

It is sometime easier to think about an estimate of the remaining useful life (RUL) and eventually to produce an indicator that corresponds to the probability that this RUL is less than a given delay. On Figure 5 we also introduce the probability of failure (POF) which is the probability that the remaining useful life is less than h.

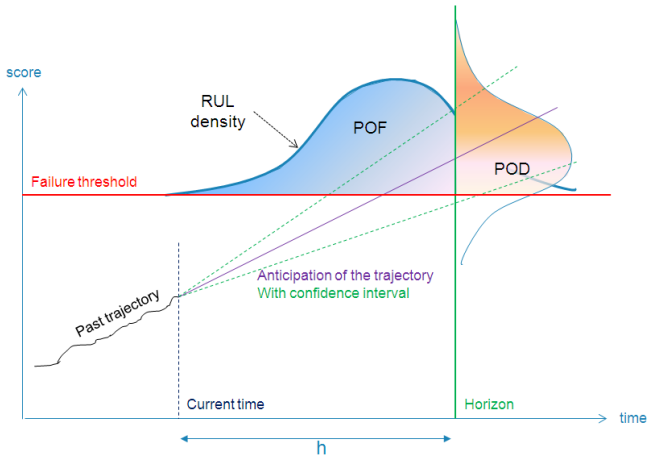

Figure 5. We introduce two output indicators: the probability of failure (POF) directly linked to the remaining useful life (RUL) and the probability of detection (POD) used in the maintenance logistic.

3. ANTICIPATION ALGORITHM

A nice method to anticipate a continuous process is to model its behavior with an autoregressive model and then filter the signal with a Bayesian update of the state. The standard way to use dynamic filter for anticipation purpose is to throw particles with a sampling scheme (Liu, 2001). However the particle filters or other equivalent Bayesian derived algorithms (Kalman, extended or unscented filters) need an input about the dynamic model to follow (An, Choi, & Kim, 2012), (Saxena, Goebel, Field, & Filter, 2012). In our case, the evolution of the EGT margin is unknown and probably depends on the company process. We need for example to find if it is possible to anticipate the availability of airport technicians for water washes, and it probably depends on the airline politics and its financial stress. Anyway we may assume some regularity and try to find a good autoregressive model: one that predicts this behavior depending on the past observations.

The mathematics to retrieve a good autoregressive (ARMA) model from observations is given in (Lacaille, 1998) with possible adaptation to non linear neural models. But most of today's algorithmic toolboxes are able to fit autoregressive models.

Hence it is possible to retrieve a hidden state (x_t) from observations (y_t) where

$$\begin{cases} x_t = Fx_{t-1} + e_t \\ y_t = Hx_{t-1} \end{cases} \quad (1)$$

with (e_t) a process noise, F the state transition matrix and H the observation matrix that define the dynamic system. The main problem is to find the good rank of the system (dimension of the transition matrix F), but a good guess is given using some information criterion like AIC (Akaike

information criterion) which was specifically built for this purpose). This model estimation is done after each new observation and for each different engine. The observation set taken into account for learning extends from six month to one year of observations (~1000 flights) in our study.

The second step, once the dynamic captured on the past observations, is to use a sampling algorithm to simulate particles (probable trajectories) and infer a probability to cross a threshold before a maximum time allowed by the company rules and the availability of a technical team to realize the maintenance operations.

Figure 6 presents the preceding curve where only the first 75% of the first flights are used for learning the dynamic model. Then a sampling algorithms simulate this model, weights are given to each trajectory according to the filter relevance computed on the first "observed" 75% points. The blue-to-green curves at the end of the graph shows the results of the simulations (100 particles in this case). In fact the dynamic is not sufficiently regular to establish a real prediction. We don't capture any regularity in time or in size, hence the probability density of the particles after three months (~300 flights) is almost a Gaussian noise.

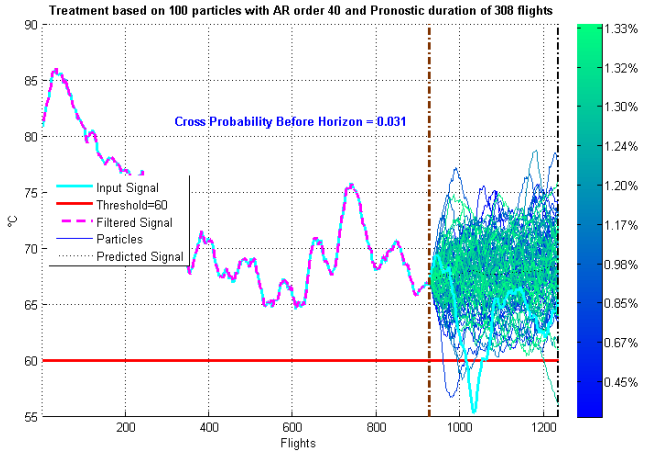

Figure 6. Estimation of the future of the margin at 75% of the available data. The blue-green color map shows the weight of the different particles. Green trajectories represent particles with higher weights.

However, three month anticipation is definitely too optimistic. Figure 7 is a zoom of the preceding graph near the anticipation point. One week corresponds roughly to 25 points and most of the repair procedures may be realized in two days (less than 10 flights) so even if this tool is not very efficient we should not be too hard with ourselves.

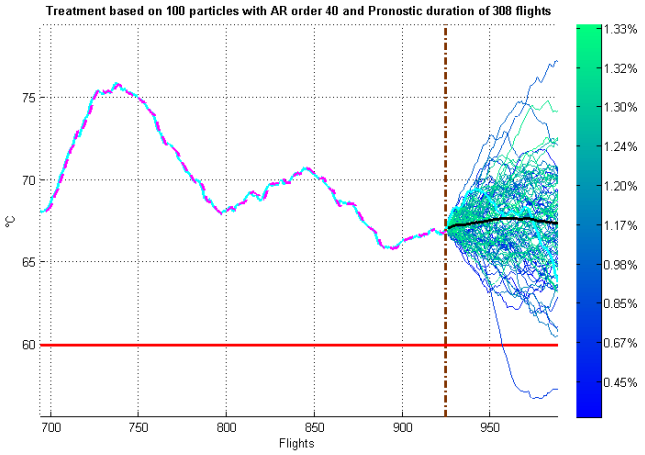

Figure 7. Zoom of the preceding estimation. The black curve presents the weighted mean of all trajectories.

4. SKIPPING THE AIRLINE INTERVENTIONS

In fact the underlying problem is much simpler if we don't bother with the random increases of the margin. We understand that those phases are completely unnatural because an engine cannot repairs by itself. Margin increases are the result of company operations and the widths of those phases are only computation artifacts due to smoothing and variation of the acquisition context.

We suppose that the statistic model of the step process $dX_t = X_{t+1} - X_t$ behind our observations is an independent process decomposable into two independent parts:

$$dX_t = U_t + Z_t G_t \qquad (2)$$

- a decreasing part $U_t \sim N(-u, \sigma_u)$ for example with a Gaussian distribution with a constant negative trend $-u$;
- and a step function representing the airline maintenance operation build from a product of a binomial distribution and a positive gap $Z_t G_t$
  - where $Z_t \sim B(p)$ with a small probability p to fire an maintenance operation;
  - and a $G_t \sim N(+g, \sigma_g)$ with also for example an Gaussian behavior and a gap of mean size g which is supposed positive and greater than normal trend u.

We are only interested in the decreasing part modeled by U_t, hence the best way is to find a set of instants t where $Z_t = 0$ The probability to have a maintenance operation is $p = P(Z_t = 1)$ is supposed small, so we will detect big gaps and neglect measurement around those instants.

Figure 8 presents the result of such a detector. The graph shows points localized around wide separated instants as we

imagined. The smoothing process and previous normalization procedure may induce some artificial thickness but as the number of other (decreasing) points is great enough for a model, so we may just ignore all detected instants without much loss and no risk to pervert the estimation of u and σ_u by mixing other distributions.

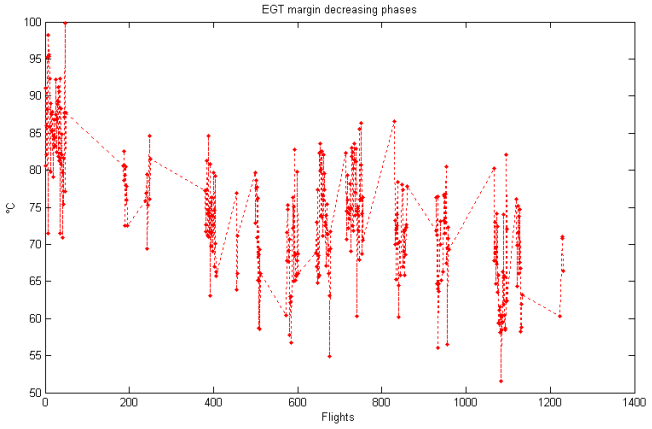

Figure 8. Search of the decreasing phases of the EGT margin. The observations where the EGT margin decreases are concentrated on small periods. The duration of those intervals is only a computation artifact due to the trend estimation.

Once all maintenance intervals suppressed from our dataset we concatenate the decreasing phases by just adding the necessary bias at each phase to ensure to obtain a continuous curve (Figure 9).

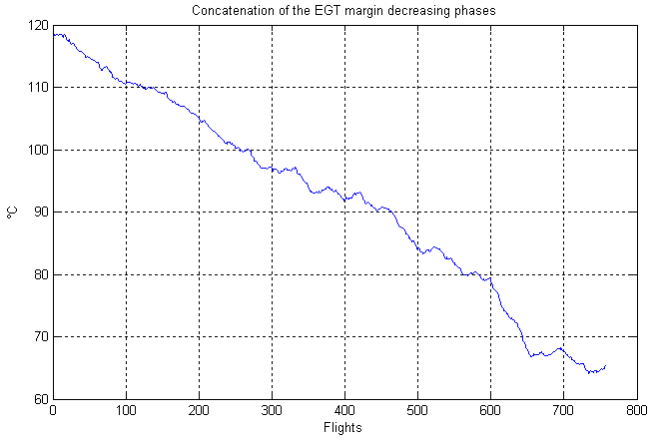

Figure 9. Concatenation of decreasing EGT margin phases.

Each small curve is added after the preceding, with corresponding bias to ensure continuity. In fact we begin with the last measurements (the last interval of data), then we concatenate to the left the preceding measurements' interval and so on. This is done by progressively adding from right to left the values of the variations $dx_t^- = x_{t-1} - x_t$. The last values of the last packet correspond to the real observations (higher values have no meaning, just the general trend is important.)

With no surprise the resulting curve is almost linear. The prediction shown on Figure 10 is a lot simpler. This time no real need to learn the dynamic of the signal: only the main trend is enough. However the algorithm used is still a particle filter because on younger engines like the one plotted on Figure 11 the behavior is not strictly linear but has a slow decrease of its trend.

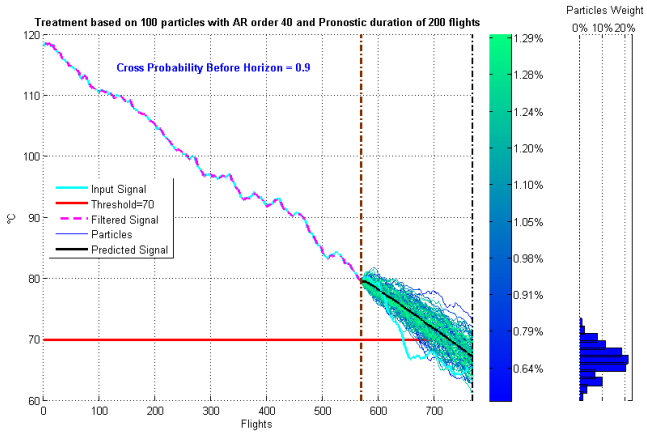

Figure 10. Prediction of the linear trend with a particle filter. The dynamic filter is just used here for presentation so we can observe the variance of the trend coefficient.

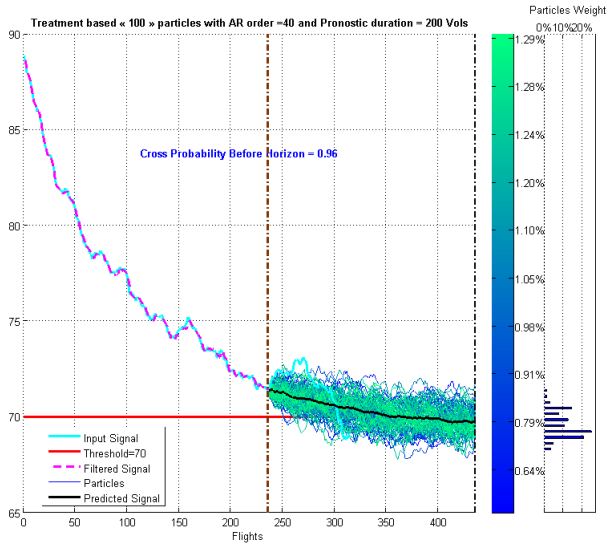

Figure 11. A younger engine with a slow decrease of its margin trend.

Eventually the final algorithm should not use a dynamic filter if a simple regression is sufficient to estimate the trend which value follows a Student law easy to estimate (Besse, 2003).

5. VALIDATION RESULTS

There are two required outputs of the algorithm. Those outputs are validated on an experiment set. A validation experiment is a selection of an engine at a given time with EGT margin computations for the 500 past flights for calibration and 200 next flights for confirmation. We used a set of $N = 150$ such experiments built on 50 different engines and 3 observation times per engine. We took wide intervals between each selected time instant in the engine's lives to get rid of local dependencies.

5.1. Estimation of the RUL for 10% of relative decrease of the margin

The requirement was to give an estimation of the RUL corresponding to a decrease of 10% of the last measure of the margin. This estimation should be given with a maximal lower bound ΔT_0 set at 95% of the distribution of the RUL estimation.

$$P\left(RUL_t^{10\%} \geq \Delta T_0\right) = 0.95 \quad (3)$$

where the $RUL_t^{10\%}$ is the delay after current time t before the EGT margin y_t cross a threshold that corresponds to a gap of 10% of its current value.

Figure 12 shows a distribution of the RUL on our set of experiments.

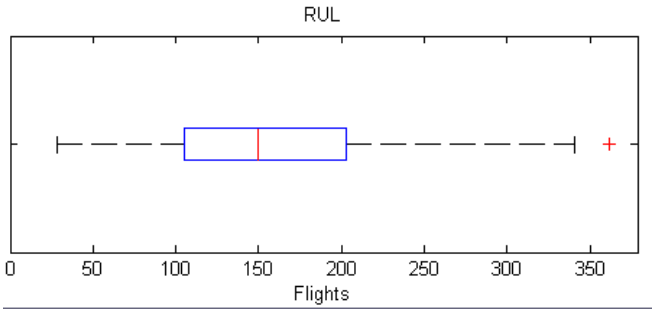

Figure 12. Distribution of the RUL computed on our experiment set.

This first output may be used as an alert indicator. Suppose we asked the client company to repair its engine at time $t + \Delta T_0$. Then the main mistake is to miss the threshold and wait too long. Our quality indicator is the proportion of misdetections ρ_{miss} on a set of N experiments. We will also look at the distributions of the delay error (or misdetection error) because it can modulate our result.

$$\rho_{miss} = \frac{1}{N}\sum_{i=1}^{N} 1_{\{RUL<\Delta T_0\}} \quad (4)$$

On a set of 150 experiments we obtain

$$\rho_{miss} = 22.1\% \pm 6.5\% \quad (5)$$

The 6.5% value after the proportion corresponds to a symmetric confidence interval computed by cross validation on the experiment set.

The next graph (Figure 13) shows a box plot of the misdetection errors.

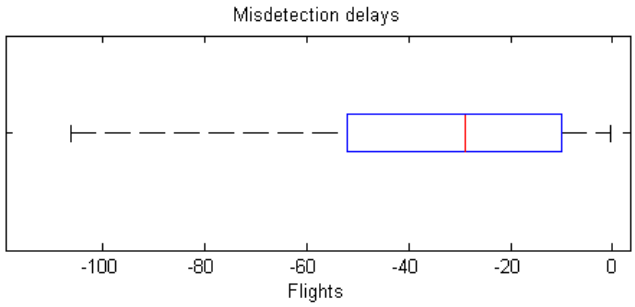

Figure 13. The delays of misdetections.

We observe that the values for those misdetections are around 30 flights and less than 50 with a 95% probability. This is less than two weeks of error for less than 25% of the estimations.

5.2. Probability to cross the 10% margin before one month (100 flights)

$$P_{100}^{10\%} = P(y_t - y_{t+100} \geq y_t * 0.1) \quad (6)$$

This second output is the probability that we cross a threshold before one month. The main goal of this algorithm is to space unnecessary interventions on the engine. The risk there is to prepare the shop when it was not really needed. Our quality indicator is a rate of unnecessary alerts ρ_{ua}, which should be as small as possible. Figure 14 presents a plot of the probability computation according to the real observed RUL.

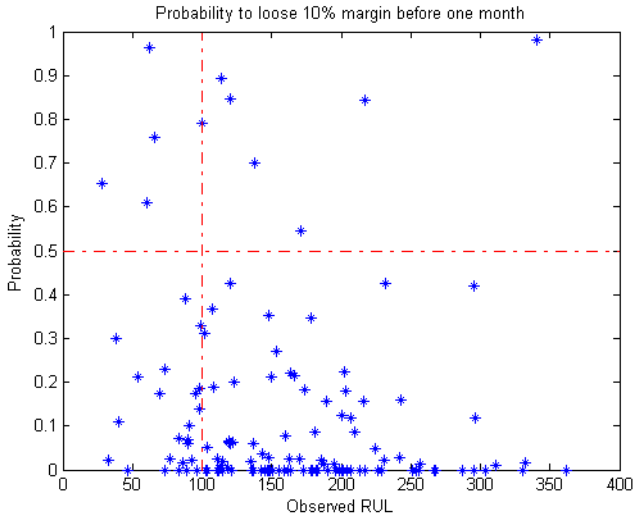

Figure 14. The probability to cross the threshold after a horizon of one month according to the real observed RUL.

The problem appears for observations on the upper right corner of this plot: when the probability is high and the RUL far away (after the month horizon). Arbitrary setting a probability threshold of 50% one may compute the rate of unnecessary alerts as

$$\rho_{ua} = \frac{1}{N}\sum_{i=1}^{N} 1_{\{(P>0.5)\&\ (RUL>100)\}} \quad (7)$$

This computation gives

$$\rho_{ua} = 6.5\% \pm 3.7\% \quad (8)$$

The 50% threshold seems to be a logical choice for decision purpose when one observes a probability of degradation and want to decide if an alert should be emitted.

6. CONCLUSION

We built a really simple algorithm able to predict with good efficiency the evolution of the EGT margin. The risks associated with both output uses (time or probability) are well mastered and not too big. In the first case the prediction error is of at most two weeks for a really small number of cases when the RUL estimation for 10% of margin decrease is between one to two months. On the other hand, the probability of threshold crossing before one month generates less than 10% early interventions.

The next step is to build a general decision rule based on both outputs which will help us to better master the risks.

NOMENCLATURE

ACARS	Aircraft Communications Addressing and Reporting System
ACMS	Aircraft Condition Monitoring System
AIC	Akaike Information Criterion
ARMA	Auto Regressive and Moving Average
LASSO	Least Absolute Shrinkage and Selection Operator
FADEC	Full Authority Digital Engine Control
MRO	Maintenance and Repair Overhaul
OSA-CBM	Open Systems Architecture for Condition-based Maintenance
PFA	Probability of False Alarm
POD	Probability Of Detection
POF	Probability Of Failure
RUL	Remaining Useful Life
SAMANTA	Snecma Algorithm Maturation And Test Application

REFERENCES

An, D., Choi, J., & Kim, N. H. (2012). A Comparison Study of Methods for Parameter Estimation in the Physics-based Prognostics. *PHM*. Minneapolis (MO).

Bellas, A., Bouveyron, C., Cottrell, M., & Lacaille, J. (2012). Robust Clustering of High-Dimensional Data. *ESANN* (pp. 25–27). Bruges (Bx).

Bellas, A., Bouveyron, C., Cottrell, M., & Lacaille, J. (2013). Model-based Clustering of High-dimensional Data Streams with Online Mixture of Probabilistic PCA. *ADAC*, 1–20.

Besse, P. (2003). *Pratique de la modélisation statistique*.

Come, E., Cottrell, M., Verleysen, M., & Lacaille, J. (2010). Self Organizing Star (SOS) for health monitoring. *ESANN*. Brugges (Bx).

Cottrell, M., Gaubert, P., Eloy, C., François, D., Hallaux, G., Lacaille, J., & Verleysen, M. (2009). Fault prediction in aircraft engines using Self- Organizing Maps. *WSOM*. Miami (FL).

Flandrois, X., Lacaille, J., Masse, J.-R., & Ausloos, A. (2009). Expertise Transfer and Automatic Failure Classification for the Engine Start Capability System. *AIAA InfoTech*.

Lacaille, J. (1998). Synchronization of multivariate sensors with an autoadaptive neural method. *Intelligent & Robotic Systems*, *21*(2), 155–165.

Lacaille, J. (2009a). An Automatic Sensor Fault Detection and Correction Algorithm. In American Institute of Aeronautics and Astronautics (AIAA) (Ed.), *Aviation Technology, Integration, and Operations Conference (ATIO)*. Hilton Head (SC).

Lacaille, J. (2009b). Standardized failure signature for a turbofan engine. *IEEE Aerospace conference* (p. 11/0505). Big Sky (MT).

Lacaille, J. (2010). Standardization of Data used for Monitoring an Aircraft Engine. US patent 2010076468A1

Lacaille, J., & Come, E. (2011a). Visual Mining and Statistics for a Turbofan Engine Fleet. *IEEE Aerospace Conference* (p. 11/0405). Big Sky (MT)

Lacaille, J., & Come, E. (2011b). Sudden change detection in turbofan engine behavior. *CM & MFPT*. Cardiff, UK: British Institute of Non-Destructive Testing.

Lacaille, J., & Nya Djiki, R. (2009). Model Based Actuator Control Loop Fault. *European Conference on Turbomachinery Fluid Dynamics and Thermodynamics*. Gratz, Austria.

Liu, J. S. (2001). *Monte carlo Strategies in Scientific Computing*. *Book* (p. 245). Springer.

Saxena, A., Goebel, K., Field, M., & Filter, E. K. (2012). Uncertainty Representation and Interpretation in Model-based Prognostics Algorithms based on Kalman Filter Estimation. *PHM*. Minneapolis (MO).

BIOGRAPHY

Jérôme Lacaille is a Safran emeritus expert which mission for Snecma is to help in the development of mathematic algorithms used for the engine health monitoring. Jérôme has a PhD in Mathematics on "Neural Computation" and a HDR (habilitation à diriger des recherches) for "Algorithms Industrialization" from the Ecole Normale Supérieure (France). Jérôme has held several positions including scientific consultant and professor. He has also co-founded the Miriad Technologies Company, entered the semiconductor business taking in charge the direction of the Innovation Department for Si Automation (Montpellier - France) and PDF Solutions (San Jose - CA). He developed specific mathematic algorithms that where integrated in industrial process. Over the course of his work, Jérôme has published several papers on integrating data analysis into industry infrastructure, including neural methodologies and stochastic modeling.

Aurélie Gouby is a Snecma engineer with a double skill in mathematics and informatics. She's working directly with the PHM team helping system engineers implementing algorithmic solution in the engine.

Olivier Piol is a Snecma engineer working on performance analysis for civil engines. He's working directly with the PHM team on gas path monitoring.