Modeling Propagation of Gas Path Damage

Kai Goebel*
RIACS, NASA Ames Research Center
MS 269-4, Moffett Field, CA 94035
650-604-4204
goebel@email.arc.nasa.gov

Hai Qiu, Neil Eklund, Weizhong Yan
GE Global Research
One Research Circle, Niskayuna, NY 12309
{qiu, eklund, yan} @ research.ge.com

Abstract— This paper describes how damage propagation can be tracked and modeled for a range of fault modes in some modules of commercial high bypass aircraft engines. To that end, response surfaces of all sensors are generated via a thermo-dynamical simulation model for the engine (cycle deck) as a function of variations of flow and efficiency of the modules of interest. These surfaces are normalized and superimposed. Next, sensor readings are matched to those surfaces and – using an optimization approach – the corresponding flow and efficiency pair is found that best explains the sensor data. This flow and efficiency pair is then compared to previous pairs and the direction of the change as well as the rate of change is determined. The whole trajectory is then projected into the time domain. An extrapolation of the curve to the limit (which is established via operational margins) yields the remaining life. In a backward mode, the extrapolated curve is discretized and estimated future flow and efficiency pairs are retrieved. These pairs are then input to the cycle deck to produce future anticipated sensor readings as well as confirmatory trips of operational margins. Changes of the future sensor readings with real readings are used to adjust the remaining life calculations. The method is demonstrated on time series of historical engine faults.

TABLE OF CONTENTS

1-4244-0525-4/07/$20.00 ©2007 IEEE
IEEEAC paper #1032, Version 3, Updated December 26, 2006
* Author conducted work while he was with GE Global Research

INTRODUCTION

At the heart of model-based prognostics is the ability to properly model the propagation of damage. To that end, considerable effort is currently being placed into understanding the mechanisms that influence the propagation at the materials level. This requires an in depth understanding of the local conditions the particular component is exposed to. For example, for spall propagation in bearings, it is required to have knowledge about the local load, speed, and temperature conditions at the site of the damage, e.g., at the outer race (or ball or cage). In addition, it is required to have knowledge about the geometry and local material properties at the suspected damage site. This information is then used to derive the stresses the component is expected to experience, typically using a finite element approach. The potential benefit of this – arguably somewhat tedious – process is the promise of accurate prediction of when the bearing will fail. For a different fault mode, the process has to be repeated. Because of the cost and effort involved, this method is reserved for a set of components that, if left undetected and without remaining life information, might experience catastrophic failure that transcends the entire system and causes system failure. However, there is a large set of components that will not benefit from this approach, either because the number of system failures caused for individual fault modes is low or because typically diagnostics can catch the fault before it causes system failure. In case of the small number of system failure per fault mode, the number of failures for the sum of all fault modes at a particular component might still be large. This might justify the need for prognostic techniques that can employ techniques that are not as labor-intensive and cost-intensive, perhaps at the expense of some accuracy. A similar argument can be made for the case where diagnostics detects the fault but causes wasting in-service life, or worse, unscheduled maintenance, the latter typically causing significant interruption of service.

It is therefore desirable to be able to increase coverage of prognostics for a range of fault mode. To that end, the techniques would ideally utilize existing models and sensor data.

As mentioned earlier, the key to prognostics is the ability to model the propagation of faults. Some information exists from diagnostics about how faults manifest themselves in sensor signatures. However, little is known about how faults propagate and what effect this propagation has on the sensor signature. Similarly, little is known about how one would go about modeling the propagation of particular faults in existing tools such as a cycle deck.

This paper addresses this issue by learning the signature of faults in the sensor data and mapping them back into parameters that can be changed in cycle decks. Initial results indicate that particular faults have preferred directions in the health related parameter space. By extrapolating the propagation in this parameter space and by mapping the extrapolation into the time domain, remaining life information can be derived.

The following section briefly discusses prognostics, along by a review of health parameter modeling. This is followed by a description of how to accomplish prognostics in the health parameter space. A results section shows examples cases. The paper concludes with a discussion and remarks for future research.

PROGNOSTICS

Predicting life is not straightforward because, ordinarily, remaining life is conditional on future usage conditions such as load and speed, among others. Finding solutions to the prognostics problem is a very active research area. It has the promise of allowing users to avoid unscheduled maintenance and to increase equipment usage. In addition, it might potentially improve the safe operation of the equipment.

Prognostics is closely linked with diagnostics. In the absence of any evidence of damage or faulted condition, prognostics reverts to statistical estimation of fleet-wide life. It is more common to employ prognostics in the presence of an indication of abnormal wear, faults, or other non-normal situation. It is therefore critical to provide accurate and quick diagnostics to allow prognostics to operate.

Diagnostics

As mentioned earlier, condition-based prognostics is reliant on diagnostics. It is assumed that the latter will provide a trigger point for the prognostic algorithms. That is, no prognostic estimates are calculated before diagnostics has detected a fault condition. In the absence of abnormal conditions – or fault conditions – the best estimates for remaining component life are often fleet wide statistics expressed by Weibull curves or other suitable mechanism. Condition-based systems depend on reliable fault diagnostics to initiate the prognostic algorithms. If diagnostics recognizes the start point of damage too late, the damage propagation models will always lag reality and keep underestimating the damage. If prognostic algorithms are kicked off when there is no real damage, the benefit of true remaining life estimate is erased. It is therefore important to optimize the diagnostic capability to attain optimal prognostics. In this paper we will assume the presence of an accurate fault detection algorithm.

Prognostics Definition

Prognostics is here defined as the estimation of remaining useful component life. The remaining useful life (RUL) estimates are in units of time or cycles. The time estimate typically has associated uncertainty that is described as a probability density curve around the actual estimate. Operators can choose a confidence that allows them to incorporate a risk level into their decision making. Typically, the confidence level on RUL estimates increases as the prediction horizon decreases.

Remaining life estimates provide indispensable information for operation of modern complex equipment. They provide decision making aids that allow operators to change operational characteristics (such as load) which in turn may prolong the life of the component. It also allows planners to account for upcoming maintenance and set in motion a logistics process that supports a smooth transition from faulted equipment to fully functioning. Examples of these types of equipment are aircraft engines (both military and commercial), medical equipment, power plants, etc.

A common approach to prognostics is to employ a model of damage propagation contingent on future use. Such a model is often times based on detailed materials knowledge and makes use of finite element modeling. Another approach is to take advantage of time series data where equipment behavior has been tracked via sensor measurements during the normal operation all the way to equipment failure. When a reasonably-sized set of these observations exists, pattern recognition algorithms can be employed to recognize these trends and predict remaining life (albeit, often times under the assumption of near-constant future load conditions). However, often times, run-to failure data are not available because, when the observed system is complex and expensive and safety is critical (e.g., in aircraft engines), faults are repaired before they lead to system failure. This deprives the data driven approach of critical information.

Damage Modeling

Tracking and predicting the progression of damage in a turbo machinery has some roots in the work of Kurosaki et al. (Kurosaki et al., 2004). They identify the efficiency and the flow rate deviation of the compressor and the turbine based on operational data for fault detection. Chatterjee and Litt (Chatterjee and Litt, 2003) investigate engine degradation by exploring flow and efficiencies.

Our proposed process is broken down into an off-line training process and an on-line monitoring process.

Off-line process

During the off-line process, the sensor response surfaces are acquired from the cycle deck as a function of flow and efficiency for specific modules (see Figure 1). Specifically, for each module in the gas path (HPC, HPT, and LPT), the efficiencies and flows are incrementally changed and the cycle deck is then run under reference cruise conditions. Some resulting HPC module response surfaces for the core speed sensor (N2), the exhaust gas temperature (EGT), the compressor inlet pressure (ps3), and the compressor inlet temperature (T3) are shown in Figure 1 –Figure 4, respectively. The range of the flow and efficiency is the same in all figures. Response surfaces for the other modules are generated by the same process.

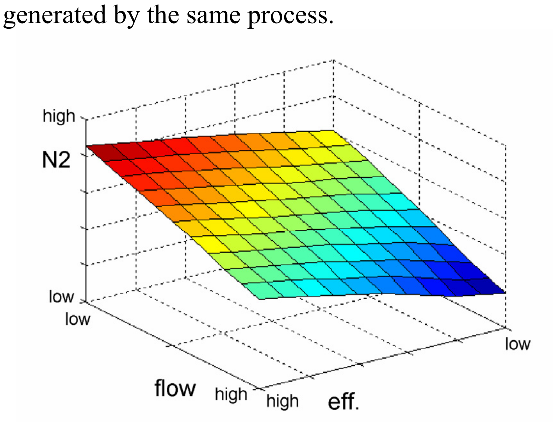

Figure 1: Response surface of N2

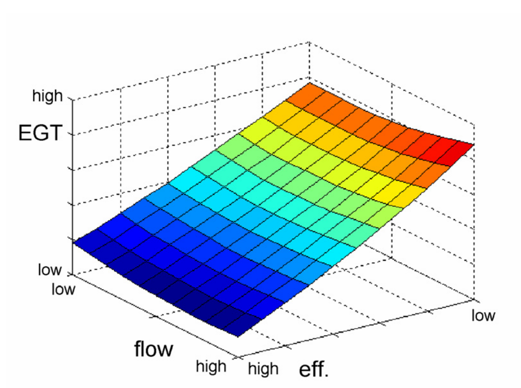

Figure 2: Response surface of EGT

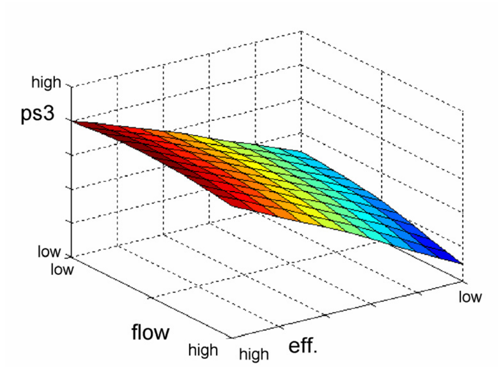

Figure 3: Response surface of ps3

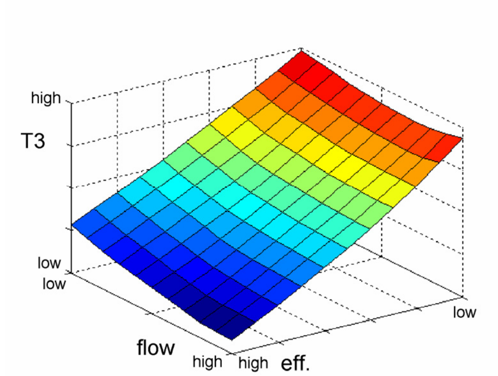

Figure 4: Response surface of T3

The response surfaces for all gas path modules are then overlaid and normalized for use with the online process. There, a plurality of sensor measurements is fitted to the response such that a best match is established. The corresponding flow f and efficiency η pair is assumed to be the parameter set for this flight.

Online Process

The online process can be divided in several steps:

1.) Based on the information from diagnostic module, a set of response surfaces associated with the faulty component is selected. Sensor measurements are normalized and their best match with the existing surfaces is determined. This can be phrased as an optimization problem where the objective function is written as,

$$\min\left(\sum \mathrm{w}_i (dist_i)^2\right), \mathrm{i} \in \{\mathrm{N}_2, ps_3, T_3, EGT, wf_{36}\}$$

where the $dist_i$ are the distances from the specific measurements to the respective response surface. This step is illustrated in Figure 5 where a vector of measurements is fitted to the respective response surfaces. After sensor measurements collected at different flight cycles have found their corresponding best matching pairings on the same $\eta - f$ space, a $\eta - f$ trajectory can be generated by connecting all those best matches.

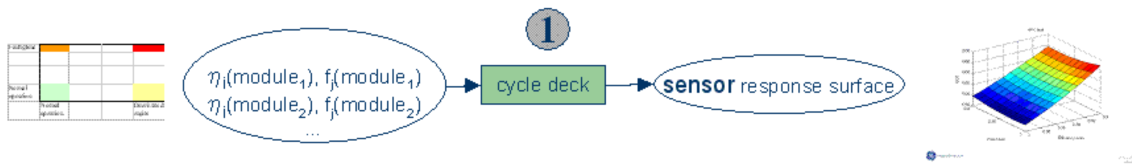

Figure 6 Steps for off-line process

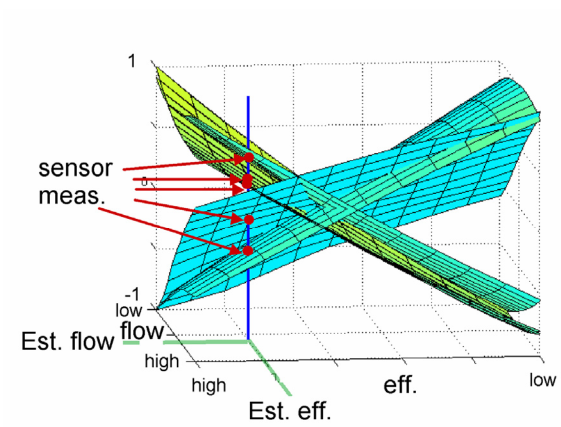

Figure 5: Process of finding best match between sensor vector and response surfaces

2.) Next, the best matching pairing η, f is compared with previous η, f pairings. In the presence of a fault and depending on the fault mode, the direction of the trajectory in the η, f space is determined. This is a key insight because different fault modes may be manifested by different progressions in this space. An example for the directional information is shown in Figure 13 that shows the fault of a static structure in the gas path.

3.) As a next step, the rate of change in the η, f space is observed and recorded for the trajectory. A critical aspect of prognostics is the ability to establish operability limits. This is a threshold beyond which the equipment cannot be operated. Often times, this is the point where failure is assumed to occur (or failure will occur for some proportion of equipment given a particular risk level). Failure can be expressed in many different ways. For the purpose of this study, we consider zero operational margin (such as stall margin) as the failure threshold.

4.) Among the margins considered, some are directly measurable, such as core speed limits and upper EGT thresholds. Others are "virtual" margins established through simulation in the cycle deck. A normalized margin is used to quantify the health of engine modules. The underlying premise is that if one engine with certain η and f pairing violates either one of operational margins under any possible operational conditions, such as hot day take off, maximum climb, or cruise, its health index would be zero. Otherwise, whichever minimal margin would be its current health index. Figure 7 shows a normalized minimal margin surface of the HPC under faulty condition. The pink contour line presents the boundary of zero margin, which means any η and f pairings beyond that line would have a minimal margin less than zero, i.e., the health index is zero.

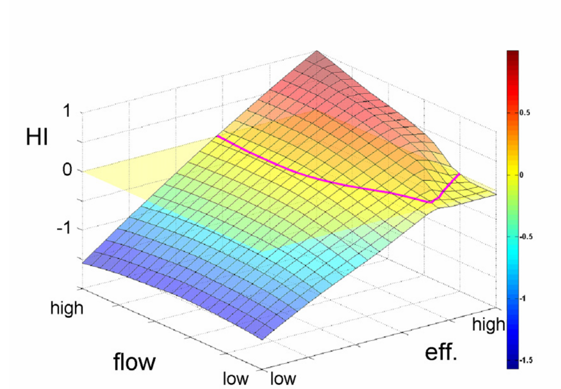

Figure 7 - Minimal margin surface of HPC fault

The trajectory is then projected into the time domain (Figure 8).

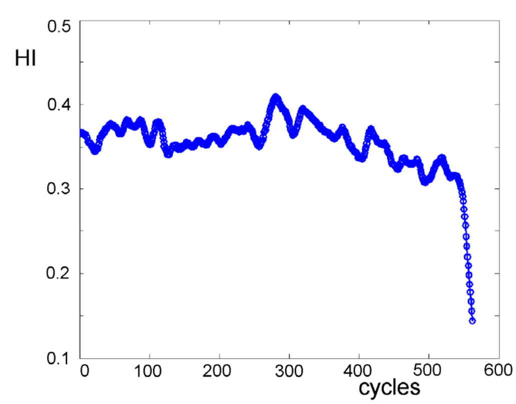

Figure 8 – Trajectory projected into time domain

5.) A suitable function is used for regression and extrapolation of the curve to known health limits. Figure 9 illustrates that concept where the circles represent observed data, the dashed line is extrapolated. The intersect of the extrapolated line with zero health gives the remaining life. Usually, the prognostics process stops with this step

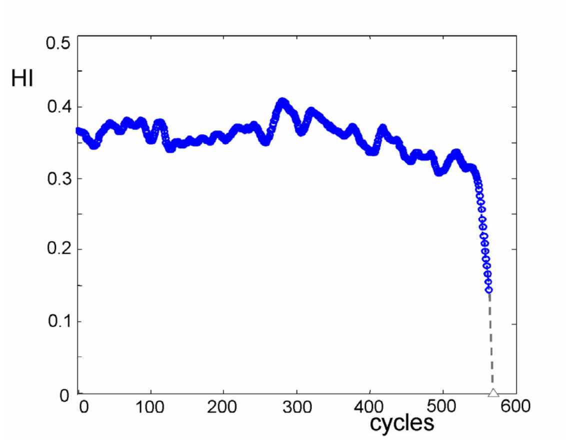

Figure 9: Extrapolation of projected health parameters.

However, there is additional information that can be gleaned from the calculated value.

6.) Specifically, in a backward chaining move, the extrapolated η, f pairings are discretized. Next, thay can be used as input to the cycle deck. The resulting expected sensor outputs are compared with the real sensor output to assess prediction accuracy. The remaining life estimates are further confirmed with operational events in this forward mode. In addition, a distribution of η, f readings can be input to the cycle deck, which will produce operational events at different times. This distribution of module failure times can be used as an uncertainty estimation tool.

Figure 12 gives an overview of the online process.

RESULTS

The process described has been employed on real engine data from a high bypass commercial jet engine. The resulting trajectory for a static gas path structure fault can be seen in Figure 10. The start point is the point when a diagnostic tool has indicated the presence of a fault (the presence of a high fidelity diagnostic tool is assumed for the purpose of this study). The end points reflect when maintenance was performed. Prior to fault initiation, the movements in the η, f-space were very inert (with some random changes of small magnitude), essentially milling around a stagnant operating point, only moving slightly with normal wear over thousands of cycles. Soon after the fault starts, the engine exhibits distinct changes in the η, f space. The resulting trajectory has been smoothed to accentuate the directional properties.

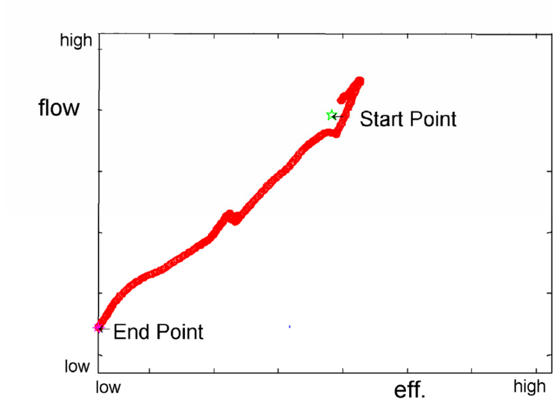

Figure 10: Directional information in η, f space for fault of static structure in gas path

Although the diagnostic tool is assumed to give accurate information about which module exhibits the fault, it is possible to confirm that information with the process described herein. Specifically, one can compute a residual error of fit for the surface mapping tasks (step 1 of the online process). Figure 11 shows the residual of surface fits applied to the different module HPT and LPT. In this case, the problem was a LPT problem, which is confirmed by the lower LPT surface fit residual.

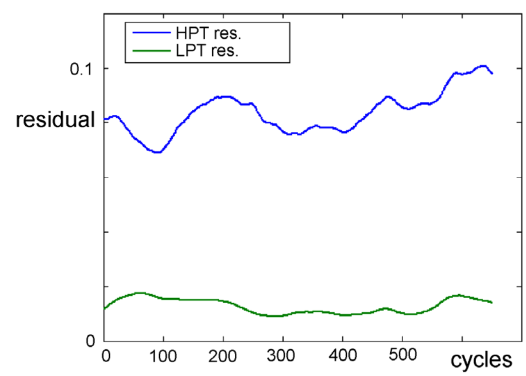

Figure 11 – Residual of surface fit for different modules

Figure 13 illustrates nicely that the directional properties of the trajectories in the η, f space are retained for different engines with the same fault mode. Only the starting point and ending points are dissimilar. The starting points are determined by engine-to-engine variation and priori deterioration of the engine.

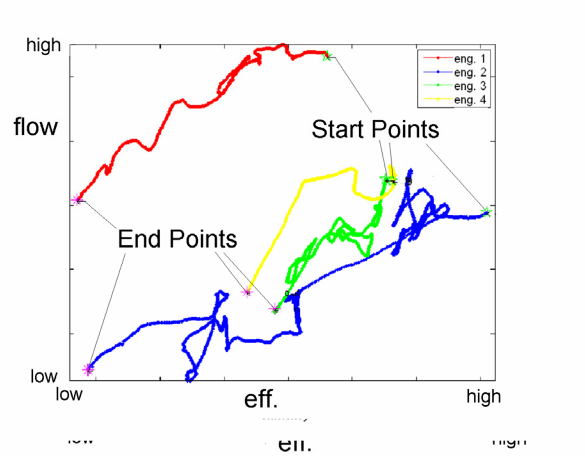

Figure 13 – Trajectories for similar faults in different engines

While this was not the primary objective of this project, it is interesting to note that other fault modes may be represented by different trajectories in the η, f space, which reveal different processes of damage propagation. Figure 14 gives an example for a different fault mode where the predominant effect is an increase in flow capacity with comparatively little loss in efficiency.

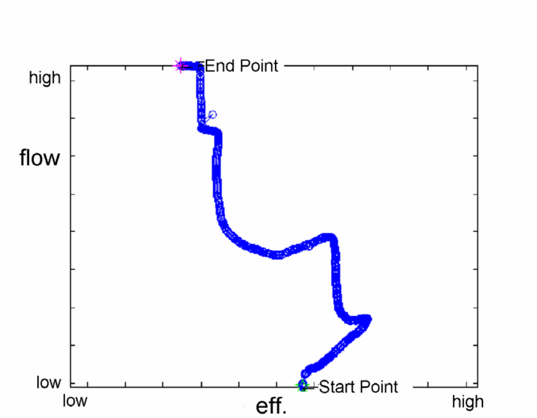

Figure 14 – Trajectory for different fault

SUMMARY & CONCLUSIONS

This paper describes how one can use cycle deck models in conjunction with real engine data to estimate damage propagation for the purpose of estimating remaining life of a component or a subsystem. This approach trades off accuracy with coverage: instead of describing the damage propagation at the materials level, it addresses it at the module level. In other words, it does not attempt to model a subset of faults with a physics of failure approach. Instead, it observes the effect of the damage propagation and extrapolates the damage to an operational limit. To that end, the response surfaces of all sensors are generated first as a function of variations of flow and efficiency of the modules. These surfaces are normalized and superimposed. Next, the sensor readings are matched to that surface and – using an optimization approach, the corresponding flow and efficiency pair is found that best explains the sensor data. Next, the flow and efficiency pair is mapped into a safety margin space which all kind of flight conditions are considered. The margin trajectory is then compared to previous cycle and the direction of the change as well as the rate of change is determined. Finally, the whole trajectory is projected into the time domain. An extrapolation of the curve to the limit (which is established via operational margins) yields the remaining life. In a backward mode, the extrapolated curve is discretized, and flow and estimated future efficiency pairs are retrieved. These pairs are then input to the cycle deck to produce future anticipated sensor readings as well as confirmatory trips of operational margins. Changes of the future sensor readings with real readings are used to adjust the remaining life calculations. In addition, the variation of the flow and efficiencies quantifies the uncertainty of the remaining life output.

Future work will tackle the validation of this approach which is an issue in general for prognostics but even more so here where failure data are missing. In addition, uncertainty management needs to be addressed in a more rigorous manner.

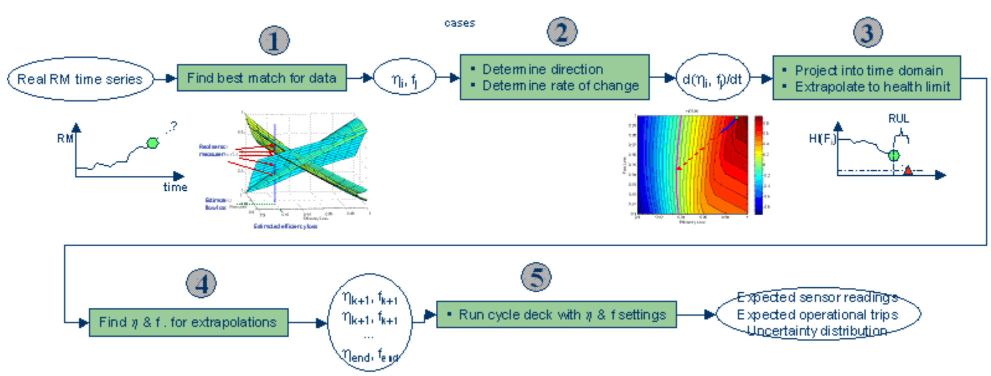

Figure 12: On-line Process

REFERENCES

S. Chatterjee and J. Litt, Online Model Parameter Estimation of Jet Engine Degradation for Autonomous Propulsion Control, Technical Memorandum NASA TM-2003-212608, 2003.

N. Aretakis, K. Mathioudakis, A. Stamatis, Nonlinear Engine Component Fault Diagnosis From a Limited Number of Measurements Using a Combinatorial Approach, Journal of Engineering for Gas Turbines and Power, JULY 2003, Vol. 125, pp. 642-650, 2003.

M. Kurosaki, T. Morioka, K. Ebina, M. Maruyama, T. Yasuda, M. Endoh, Fault Detection and Identification in an IM270 Gas Turbine Using Measurements for Engine Control, Journal of Engineering for Gas Turbines and Power, Oct. 2004, Vol. 126 , pp. 726-732, 2004.

BIOGRAPHY

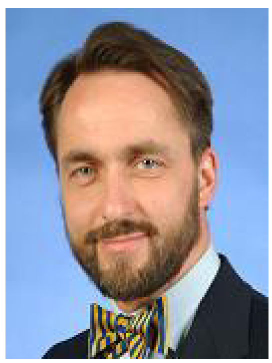

Kai Goebel *received the degree of Diplom-Ingenieur from the Technische Universität München, Germany in 1990. He received the M.S. and Ph.D. from the University of California at Berkeley in 1993 and 1996, respectively.*

Dr. Goebel joined RIACS in 2006. He is currently leading the Prognostics Center of Excellence at NASA Ames. Prior to that, he was with General Electric's Corporate Research and Development facility in Schenectady, NY from 1997 to 2006. He has carried out applied research in the areas of artificial intelligence, soft computing, and information fusion. His research interest lies in advancing these techniques for real time monitoring, diagnostics, and prognostics. He has fielded numerous applications for aircraft engines, transportation systems, medical systems, and manufacturing systems. He has published more than 75 technical papers in these areas and he holds 6 patents.

Dr. Goebel has been an adjunct professor of the CS Department at Rensselaer Polytechnic Institute (RPI), Troy, NY, since 1998 where he taught classes in Soft Computing and Applied Intelligent Reasoning Systems.

Hai Qiu *is a Research Scientist in the GE Global Research at Niskayuna, New York. Prior to join GE in 2005, he was a Research Assistant Professor in the Department of Mechanical, Industrial and Nuclear Engineering of the University of Cincinnati and served as the Lead Researcher of the NSF Industrial/University Cooperative Research Center for Intelligent Maintenance Systems (IMS). He obtained his Bachelor and PhD degree in mechanical engineering from the Xi'an Jiaotong University in 1995 and 1999 respectively. He has conducted a wide variety of research projects in the fields of prognostics and intelligence maintenance systems, funded by the NSF and industry. His current research areas include intelligent diagnostics and prognostics, advanced signal processing, and applied artificial intelligence.*

Neil Eklund *received B.S. in 1991, two M.S. degrees in 1998, and a Ph. D. in 2002, all at the Rensselaer Polytechnic Institute.*

Dr. Eklund was a research scientist at the Lighting Research Center from 1993 to 1999. He was in the network planning department of PSINet from 1999 to 2002, before joining General Electric Global Research in Niskayuna, NY in 2002. He has worked on a wide variety of research projects, including early detection of cataract using intraocular photoluminescence, multiobjective bond portfolio optimization, and on-wing fault detection and accommodation in gas turbine aircraft engines. His current research interests involve developing hybrid soft/hard computing approaches for real-world problems, particularly real time monitoring, diagnostics, and prognostics.

Dr. Eklund is an adjunct professor in the Engineering/CS department of the Graduate College of Union University, Schenectady, NY, since 2005 where he teaches classes in Computational Intelligence and Machine Learning.

Weizhong Yan *is a Research Engineer in the Computing and Decision Sciences of GE Global Research Center. His research interests include pattern recognition and classification, data analysis and modeling, and soft computing. His specialties include application of soft computing technologies to monitoring and diagnosis of gas turbine engines and other mechanical systems. Dr. Yan is an Adjunct Professor of the Mechanical Engineering Department at Rensselaer Polytechnic Institute where he teaches Control and Modeling*