Advanced Diagnostics and Prognostics for Gas Turbine Engine Risk Assessment

Michael J. Roemer and Gregory J. Kacprzynski
Impact Technologies, LLC
125 Tech Park Drive
Rochester, NY 14623
716-424-1990
mike.roemer@impact-tek.com

Abstract- Real-time, integrated health monitoring of gas turbine engines that can detect, classify, and predict developing engine faults is critical to reducing operating and maintenance costs while optimizing the life of critical engine components. Statistical-based anomaly detection algorithms, fault pattern recognition techniques and advanced probabilistic models for diagnosing structural, performance and vibration related faults and degradation can now be developed for real-time monitoring environments. Integration and implementation of these advanced technologies presents a great opportunity to significantly enhance current engine health monitoring capabilities and risk management practices.

This paper describes some novel diagnostic and prognostic technologies for dedicated, real-time sensor analysis, performance anomaly detection and diagnosis, vibration fault detection, and component prognostics. The technologies have been developed for gas turbine engine health monitoring and prediction applications which includes an array of intelligent algorithms for assessing the total 'health' of an engine, both mechanically and thermodynamically.

TABLE OF CONTENTS

1. INTRODUCTION

There is a great opportunity for military jet engines to become more reliable, operationally available and economically maintained through the use of enhanced diagnostic and prognostic strategies such as those presented in this paper. The development and integration of enhanced diagnostic and prognostic algorithms that can predict, within a specified confidence bound, time-to-failure of critical engine components can provide many benefits including:

- Improved safety associated with operating and maintaining gas turbine engines.
- Reduced overall life cycle costs of engines from installation to retirement.
- Ability to optimize maintenance intervals for specific engines or fleets of engines and prioritization of tasks to be performed during the planned maintenance events.
- Increased up-time/availability of all engines within a fleet.
- Provides engineering justification for scheduling maintenance actions with corresponding economic benefits clearly identifiable.

The development of enhanced diagnostic and prognostic strategies built upon existing engine condition monitoring platforms can allow for continuous monitoring and prediction of component failure rates and degraded engine performance. A block diagram illustrating how prognostic technologies can be integrated within existing diagnostic system architectures is shown in Figure 1.

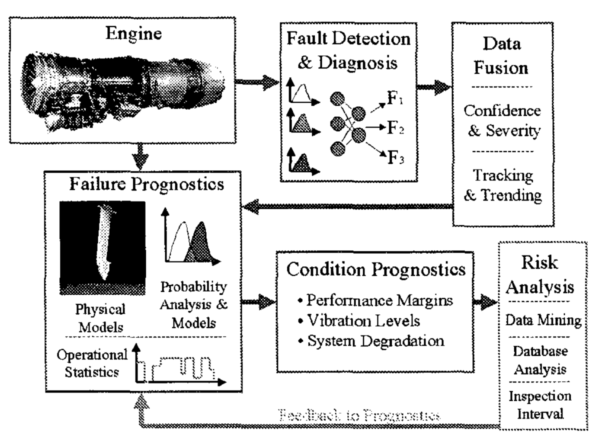

Figure 1 Integration of Prognostics within Diagnostic Framework

From Figure 1, the integration of prognostic technologies within existing diagnostic systems begins with validated sensor information currently measured on the engine being feed directly into the diagnostic algorithms for fault detection/isolation and classification. The ability of an enhanced diagnostic system to fuse information from multiple diagnostic sources together to provide a more confident diagnosis is emphasized along with a system's ability to estimate confidence and severity levels associated with a particular diagnosis. In a parallel mode, the validated sensor data and real-time current/past diagnostic information is utilized by the prognostic modules to predict future time-to-failure, failure rates and/or degraded engine condition (i.e. vibration alarm limits, performance margins, etc.). The prognostic modules will utilize physics-based, stochastic models taking into account randomness in operation profiles, extreme operating events and component forcing. In addition, the diagnostic results will be combined with past history information to train real-time algorithms (such as a neural networks or real-time probabilistic models) to continuously update the projections on remaining life. The specific approaches and algorithms for determining these component prognostic results are described in this paper.

Once predictions of time-to-failure or degraded condition are determined with associated confidence bounds, the prognostic failure distribution projections can be used in a risk-based analysis to optimize the time for performing specific maintenance tasks. A process which examines the expected value between performing maintenance on an engine or component at the next opportunity (therefore reducing risk but at a cost of doing the maintenance) versus delaying maintenance action (potential continued increased risk but delaying maintenance cost) can be used for this purpose. The difference in risk between the two maintenance or operating scenarios and associated consequential and fixed costs can then be used to optimize the maintenance intervals or alter operational plans.

2. Technical Approach

In order to evolve purely diagnostic health monitoring systems to those that are capable of more robust diagnostics and failure prognostics, a probabilistic framework is advantageous. Certainly, a prognostic system output that only reports a specific time-to-failure without having any confidence (uncertainty) bound associated with the prediction would be unwise. This is true for simple prognostic approaches that only utilize historical reliability data (such as Weibull distributions) to the more advanced prognostic modeling approaches that take design parameter and operating condition uncertainties into account.

The prognostic modeling approach implemented in this paper takes advantage of the directly sensed parameters, fused and diagnostic EHM system results, as well as inspection and historical reliability data to provide critical inputs for producing accurate failure predictions. As shown in Figure 2, this process begins with a comprehensive evaluation of the sensor data including multiple techniques that are fused together for identifying incipient sensor failure modes. Once the possibility of a sensor malfunction has been minimized, multiple diagnostic algorithm outputs are combined with diagnostic fusion techniques to enhance the fault identification capabilities of the engine health monitoring system. An example of diagnostic fusion might be intelligently processing information from vibration data, performance parameters, and oil related measurements to detect bearing faults with higher certainty. Next, in order to obtain the prognostic projections on component failures or unsatisfactory engine operation, different levels of prognostic strategies can be utilized. Approaches to be considered in this paper include probabilistic, physics-based models that take design parameter and operating condition uncertainties into account.

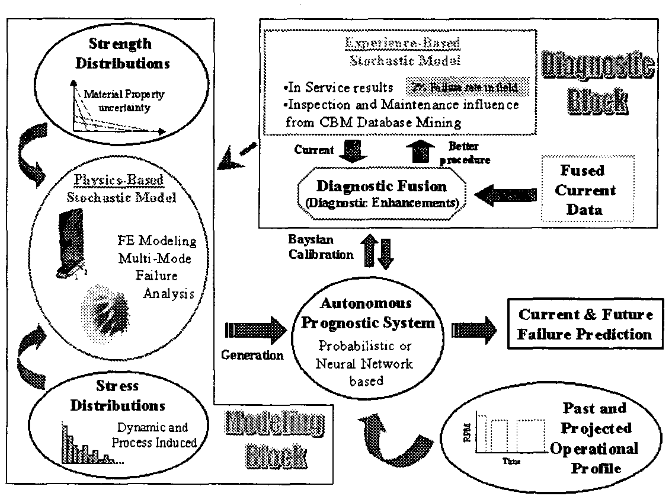

Figure 2 Prognostic Enhancements to Diagnostic Systems

3. Enhancements to Engine Health Monitoring and Diagnostic Systems

A few technologies for improving the robustness of existing engine health monitoring (EHM) and diagnostic systems will be discussed next which include; fused sensor validation and recovery, probabilistic anomaly detection and diagnosis and stochastic vibration fault classification. These real-time diagnostic enhancements are capable of accounting for uncertainties from engine transient conditions, random measurement fluctuations and modeling errors associated with model-based diagnostic procedures. Besides providing more robust diagnostic results, the techniques yield direct confidence and severity levels associated with particular diagnoses that can be directly utilized by prognostic algorithms which are also probabilistic in nature.

Fused Sensor Validation and Diagnostics

A necessary front-end to all engine health monitoring systems must insure the integrity of the measured parameters. Sensor problems such as ground loop faults, sensor drift or electrical noise can often appear as the onset of a performance or vibration fault and must be isolated and detected properly. The sensor analysis enhancements described herein must validate the integrity of sensor signals with multiple techniques that isolate particular failure modes, therefore also providing a level of fault diagnosis for the sensor system itself.

The sensor validation and diagnostic process is performed using multiple and collaborative techniques that offer advantages for isolating and detecting specific sensor failure modes (Figure 3). Some available techniques that have been implemented with success include; trained neural networks, fuzzy logic analysis, signal auto and cross correlation, and high-pass filtering.

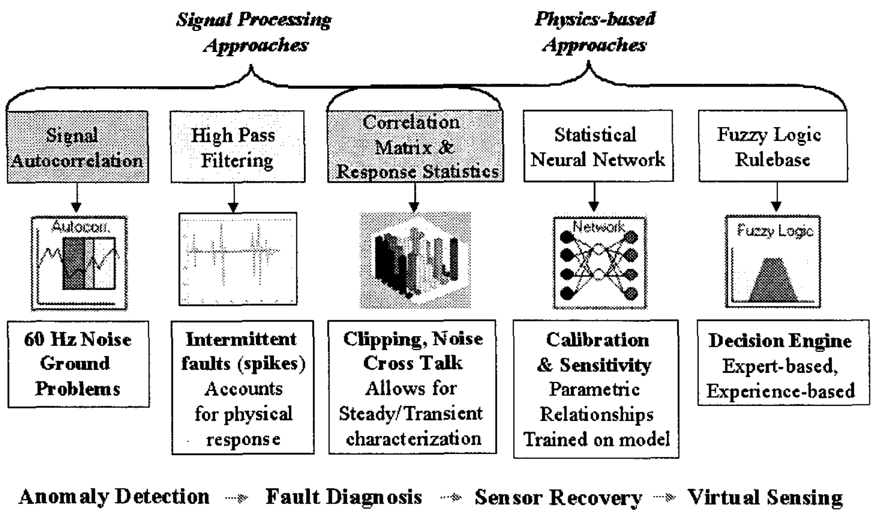

Figure 3 Sensor Validation and Fault Diagnostics

The sensor validation and diagnostic techniques fall in two categories; signal processing-based and physics-based. The signal processing techniques accomplish their tasks (detecting and diagnosing sensor anomalies) independent of a monitored system's characteristics. Conversely, the physics-based techniques accomplish more advanced tasks (sensor recovery and virtually sensing parameters) because they are developed from a-priori knowledge of the system characteristics. Signal auto-correlation and high-pass filtering are two signal processing techniques used for identifying grounding and intermittent spikes at high levels of sensitivity. The correlation matrix tracks the degree of co-linearity between signals in real-time and can detect clipping, noise and multiplexer cross-talk. This technique bridges signal processing and physics-based approaches because system response characteristics may be used to aid in fault identification in highly dynamic situations.

The more advanced detector schemes are rooted in the engine's physical characteristics. The fuzzy logic based sensor analysis continuously assesses the "normal" bands associated with each sensor signal at the current operating condition. When a signal goes outside these bands, while others remain within, an anomaly is detected associated with those specific sensors. The neural network operates by comparing the physical relationships between signals as determined from a gas path model of the engine's performance parameters. The neural network has the ability to recovery sensors that have failed or "virtually" sense critical diagnostic parameters that is not monitored. All of these parallel algorithms may be combined in a data fusion process that determines the final confidence levels that a particular sensor has either failed or has suspect operation. [3].

Figure 4 shows an example of some engine test cell results where the fuel flow sensor began "spiking" at full throttle. Both the high pass filter and autocorrelation technique were able to autonomously detecting this sensor fault.

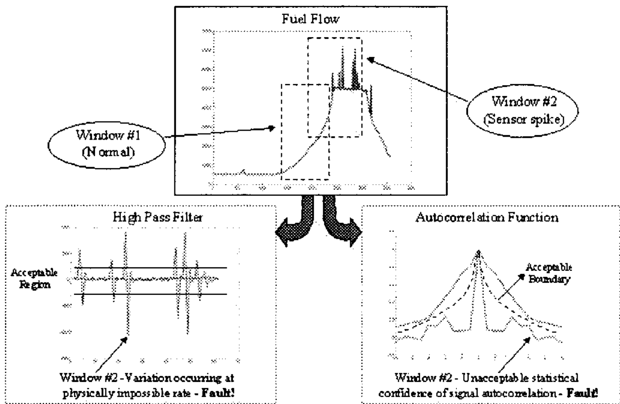

Figure 4 Detection of Spikes in Fuel Flow Sensor

Statistical Engine Baseline Models for Robust Fault Detection

A critical aspect of most model-based diagnostic or prognostic processes is the physics-based engineering analysis and/or computer simulation of the engine and associated baseline data that is used to develop the "baseline signature" model. The analysis model will be used to model the engine's "normal" behavior, and any statistical deviation from that model will used to detect engine faults. Baseline data is typically analyzed for each engine over the entire envelope of operating conditions. The corrected parametric curves are then developed as the "engine signature" models. Subsequent engine performance assessments depends on evaluating the distribution shifts associated with the normal deviations from these baseline parametric models. This concept is illustrated at the top of Figure 5.

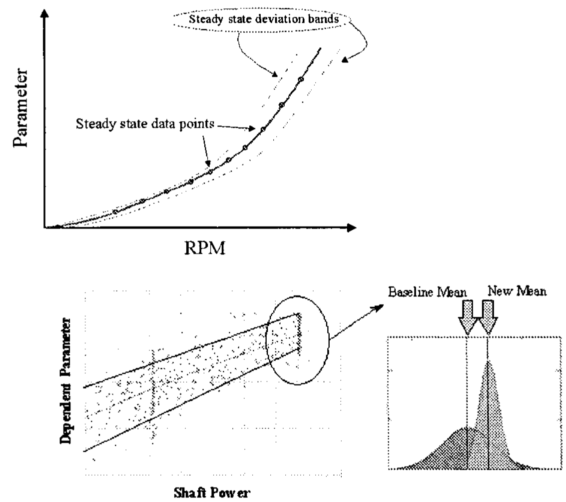

Figure 5 Engine Signature Model and Statistical Anomaly Detection

Engine performance anomaly detection algorithms are designed to statistically detect the manner in which performance parameters are shifting over time and to analyze the confidence intervals associated with engine performance degradation issues.

The algorithms statistically trend and analyze how the measured parameters compare against the engine baseline "signature" model. A sample parametric plot is shown in the bottom left of Figure 5. The bottom right of Figure 5 illustrates the normal deviation from a baseline engine signature model as the PDF or Probability Density Function identified with "baseline mean" under a specific operating condition. When a sufficient amount of new data at this condition has been collected, its PDF (identified by "new mean") may be statistically compared with the baseline PDF to determine the confidence that the mean value has shifted. As an example, we may be 75% confident that the mean has shifted by more than 3%. Continuously tracking significant parameters such as flows, vibration fault frequency amplitudes, etc. in this statistical manner can yield robust and pertinent diagnostic information regarding degraded engine operation.

Fused Fault Pattern Recognition and Probabilistic Engine Diagnostics

Utilizing collaborative probabilistic [2] and pattern recognition techniques [3] to diagnose particular fault error patterns with associated confidence and severity levels can significantly enhance current engine diagnostics. The probabilistic fault identification process utilizes known information on how measured parameters degrade over time and compares them with the current parameter distribution shifts (calculating the degree of overlap between the known fault joint probability density function and currently measured joint PDF) to identify potential fault scenarios.

Pattern recognition of a fault pattern has been successfully performed using an unsupervised neural network (Kohonen Map) for clustering fault types in series with a standard back propagation neural network to classify specific fault patterns in 2-D space [4]. The diagnostic predictions made by these two techniques are then fused together using a Bayesian inference algorithm that either increases or decreases the confidence of a particular diagnosis based on the outputs agreement or disagreements. A discussion and example of the probabilistic fault diagnostic method is provided in this paper.

As mentioned, the probabilistic fault identification process utilizes known information on how measured parameters degrade over time and compares them with the current parameter distribution to track and identify degraded fault conditions. The amount of each parameter's shift from the expected baseline operation at a particular operating condition is continuously analyzed along with the associated statistical confidence level. The shifted distributions are compared against the baseline parameter's probability distribution for all known faults. Comparing the measured distributions against all the known fault distributions yields the confidence of a known performance degradation problem.

A stochastic interpretation of the measured and known parameter distributions can provide a powerful means of calculating multiple fault contributions to a current engine operating condition and envisioning the most likely evolution of the fault. A manipulation of the First Order Reliability Method (FORM) can be used to gage a fault's evolution in state vector space where the faults and the current condition are not known with complete certainty. The conceptual framework for structural reliability and more specifically the FORM algorithm is provided by classical reliability theory [5].

In essence, the probabilistic fault diagnostic process involves assigned PDF's to performance error patterns associated to known faults in N-dimensional space. Similarly, the current error exists as a PDF in the space as well. The probability that the current condition (C), may be attributed to a given fault (F) is determined by their joint probability density function. If C and F can be assumed to be normally distributed, the probability of association (Pa) can be found using:

$$p_a = 2\Phi\left(-\frac{\overline{F-C}}{\sqrt{\sigma_f^2+\sigma_c^2}}\right) = 2\Phi(-\beta) \qquad (1)$$

where:

$\overline{F}, \overline{C}$ = the mean of the distributions F and C respectively
σ_f, σ_c = the standard deviation of the F and C distributions

The function $\Phi(\ )$ is the standard normal cumulative distribution. The notation β is the reliability index. If F and C are not normal or lognormal variables then the probability of failure can be determined only by using a computer algorithm as described in [4].

Two different fault measures have been developed from this stochastic technique [1]. First, a cumulative sensitivity index defined by the global non-dimensional variation from the initial state, at time 0, to the final state, at time t (over the time interval [0,t]). Next, an evolutionary sensitivity index defined by the local non-dimensional variation from an intermediary state, at time ti, to another intermediary state, at time ti+1, over the time interval [ti,ti+1]. Both of these indices are given below.

Cumulative Sensitivity Index:

$$C_{0,t} = -\frac{\beta_t - \beta_0}{\beta_0} \qquad (2)$$

Evolutionary Sensitivity Index:

$$E_{ti,ti+1} = -\frac{\beta_{ti+1} - \beta_{ti}}{\beta_{ti}} \qquad (3)$$

where β is the reliability index related to the Euclidean distance between the current conditional distribution and a given fault distribution.

Figure 6 illustrates the fault detection and classification process in a two-dimensional parameter space. Starting at the origin, (representing initial normal engine operation) the measured parameter distributions begin to shift as some type of performance degradation begins to occur. After several missions when the anomaly level is reached an anomaly detection warning is issued. As shown in Figure 6, the current measured PDF moves to point T1 then to T2 as time progresses, ultimately approaching the 4% fault. The described situation in Figure 6 can be handled with mathematical accuracy by using the proposed probabilistic fault diagnostic calculations previously described.

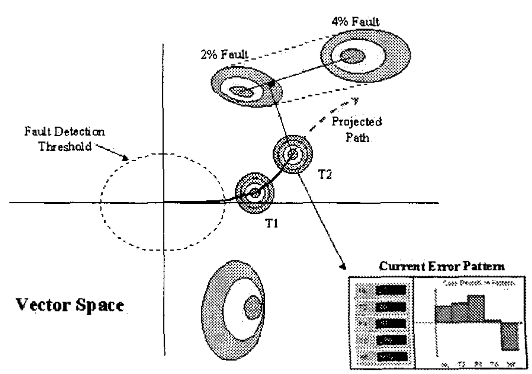

Figure 6 Probabilistic Fault Diagnosis Process

Figures 7 and 8 provide a further example of the evolution of a performance error pattern. Three different performance faults at a 2% level were identified using a gas path model of an engine. In this example, and shown in Figure 7, the PDF of the current error pattern initially evolves towards a 2% HP Compressor efficiency fault. This is also shown in Figure 8 from the fact that from T=0 to T=3 the Euclidean distance between the current PDF and the HPC fault gets smaller. However, as time goes by, the current condition evolves towards, and eventually past, the HPC fault. From Figure 8, at T=5 the current PDF is closest to the 2% HP Turbine Efficiency fault. Figure 7 illustrates that the engine's degradation has indeed evolved beyond association with the HP Compressor fault to high association with the HP Turbine fault. In this example, the final position rests at 9.98% association with the HPC fault and 22.8% association with the HPT fault.

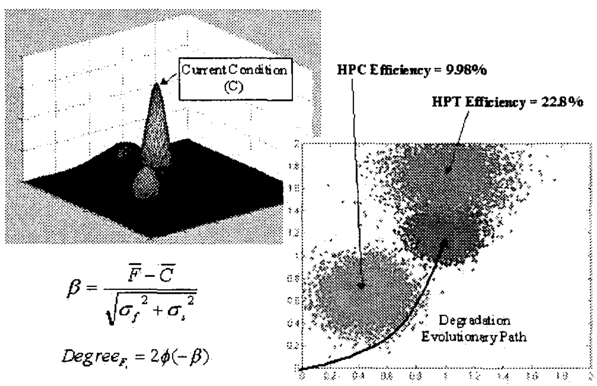

Figure 7 Degree of Fault Association

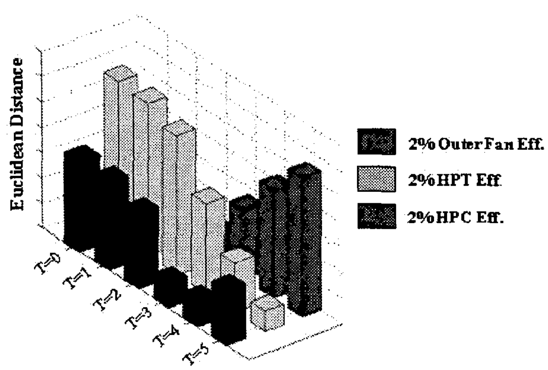

Figure 8 Fault Evolution

Enhanced Vibration Fault Detection and Diagnostics

Real-time assessment of mechanical faults (bearing, rotordynamic, structural, etc.) utilizing vibration signatures collected from accelerometers at specified locations on the machine can be enhanced using neural-fuzzy diagnostic techniques. Domain knowledge associated with particular vibration fault frequencies, fixed frequency ranges, per-rev excitations, structural resonance's, etc. are extracted from the vibration spectrums and used to develop a knowledge base from which the fuzzy logic membership functions and rulebases are developed. This feature extraction process is illustrated in Figure 9. Non-vibration related data such as the performance parameters, oil analysis data, etc. can also be integrated (knowledge fusion) into the fuzzy expert system to collaborate on a particular diagnosis.

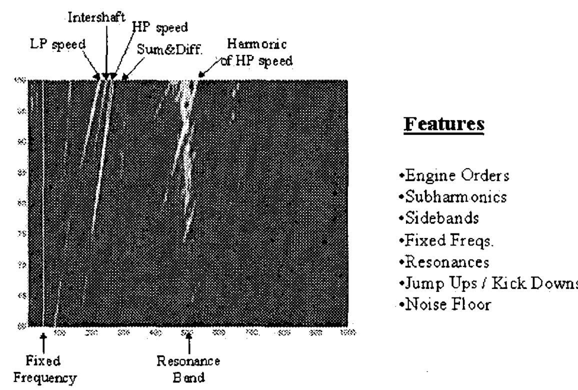

Figure 9 Feature Extraction from Vibration Spectrums

Most current vibration anomaly detection schemes rely solely on amplitude levels. However, the vibrational analysis scheme presented here investigates amplitude levels, critical spectral features and utilizes a shape-based statistical analysis of the tracked order coupled with a intelligent rulebase to detect <u>and</u> diagnose mechanical faults. The combination of these 3 techniques allows for robust and more sensitive diagnostic capability. Figure 10 shows the average shape of a HP shaft tracked order and +/- 2 standard deviations determined from testing of multiple engines. The bold line in Figure 10 shows a simulated tracked order of an engine with a different structural resonance. In this example, a simple amplitude level band on the racked order would not detect a problem, however, a statistical analysis of the <u>shape</u> of the tracked order is able to detect a fault.

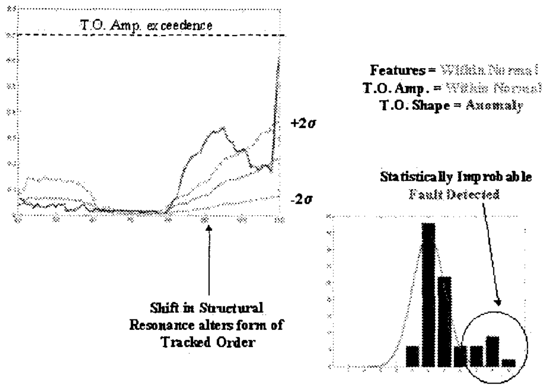

Figure 10 – Shape-based Vibration Diagnostics

Database Analysis

Utilizing probabilistic and artificial intelligence methods to record and trend critical component life usage, instrumentation problems, as well as vibration and performance faults over the life of the machine is an important feature of an advanced diagnostic and prognostic system. A database would at a minimum consist of error pattern trend charts, life accumulation charts, as well as selected data and anomoly detection logs. The error pattern trend chart shown in Figure 11 is a "to-date" snap-shot of how performace or vibration faults have manifested themselves over the life of the machine. Similarly, the life accumulation chart would show how the remaining life for a given critical compent has reduced over time.

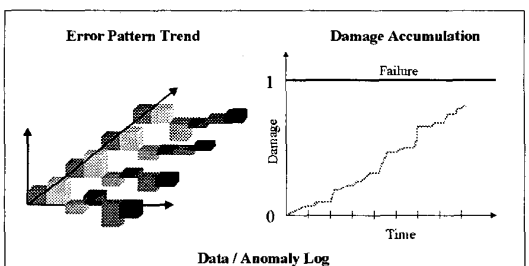

Date/Time	A	B	C	D	E	F	G	H	Comment
2/3/99	716.3	12.5	26.6	1838	261	32.233	432.23	12024	none
2/4/99	723.8363	18.3855	31.51332	1844.999	261.6869	38.22378	436.5645	12025.31	none
2/5/99	723.8631	23.79259	38.12661	1853.82	268.5282	43.6136	442.8088	12028.18	none
2/6/99	727.6233	29.11339	41.62234	1861.33	272.1812	44.55892	449.7705	12034.95	HPC eff loss
2/7/99	731.6712	32.75117	49.15439	1864.455	274.4742	50.91318	453.9047	12039.72	HPC eff loss
2/8/99	734.9802	36.413	49.88262	1872.291	281.821	56.5552	459.0287	12045.59	HPC eff loss
2/9/99	735.1179	38.93704	54.35586	1876.883	283.9431	56.7189	466.632	12046.32	HPC eff loss

Figure 11 Error Pattern Trending

4. Component Prognostic Modeling

A physics-based stochastic model is a technically comprehensive modeling approach that has been traditionally used for component failure mode prognostics. It can be used to evaluate the distribution of remaining useful component life as a function of uncertainties in component strength/stress or condition for a particular fault. The results from such a model can then be used to create a neural network or probabilistic-based autonomous system for real-time failure prognostic predictions. Other information used as input to the prognostic model includes diagnostic results, current condition assessment data and operational profile predictions. This knowledge-rich diagnostic information is generated from multi-sensory data fusion combined with in-field experience and maintenance information obtained from data mining processes.

A prognostic model must have ability to predict or forecast the future condition of a component and/or system of components given the past and current information. The realm of prognostics is sometimes divided into failure and condition prognostics. Failure prognostics often refers to the continuous accumulation of damage and/or life on components or systems of components, with or without the presence of any identified faults. Components governed by mechanical wear and failure often fit into this category (i.e. prediction of crack initiation without the presence of a fault detected). In contrast, condition prognostics is most often associated with a fault being diagnosed prior to a vibration or performance related limit being exceeded. A detected fault must be isolated and assessed for severity so that the remaining useful life can be determined. This useful life is defined by the operating time between detection and an unacceptable level of degradation.

Failure Prognostics using Physics-Based Models

A physics-based stochastic model for failure prognosis typically incorporates mechanical (finite element) or thermodynamic (through-flow model) deterministic models as their basis. The probabilistic procedure for addressing inherent modeling uncertainties must be built into these models using statistical distributions of the parameters that most directly effect the component life limiting factors. Some of these factors include the material properties, dynamic forcing, and process variability. The distribution on the current remaining life in a component life prediction may be determined by calculating all possible combinations of these life-limiting factors in a stochastic process given past maintenance and operating conditions. Operating hours can be statistically analyzed, trended and projected into the future to provide the prognosis of remaining life. More advanced stochastic models that represent failure mode uncertainties, projected operational parameters, and rare/random events can be used to help predict failure mode propagation.

This analytical model should be calibrated using in-service data to clearly reflect the root cause of the in-service failure mode experiences. In the case where a finite element model can be used (gearing, blading, impellers, or rotors for example) crack initiation regions should agree with any in-field experience and inspection data. More empirically analyzed components such as bearings should have clearly identified relationships between diagnosed fault severity and life consumption.

Example of a Stochastic Physics-Based Model

Sophisticated fracture mechanics and damage accumulation analysis have shown that accelerated crack nucleation and micro-crack formation in components can occur due to start-ups and shutdowns, transient load swings, higher than expected intermittent loads, or defective component materials. More commonly, normal wear causes configuration changes (loose fit of assembled parts, work hardened surfaces, and reduced structural section areas) that contribute to increased or unexpected dynamic loading conditions. High cycle dynamic and transmission loads cause micro-crack incubation and formation at material grain boundaries in stress concentrated regions (especially between hardened surfaces and softer subsurface material interfaces, and at acute changes in component material geometry). The majority of crack growth evolves in a sub-critical propagation process of crack tip blunting, unstable crack formation, and crack elongation. As super-critical loading in the cracked material region is approached, growth accelerates resulting in material dislocation and detachment. Sub-critical crack evolution is highly dependent on a component's material, geometry, loading conditions, and the particulars of the unique component crack growth cycle. This kind of failure-mode knowledge is often times overlooked in determining the potential usefulness of a particular prognostic or diagnostic algorithm.

The available time to take corrective or compensatory actions during specific periods of micro-crack incubation, formation, and sub-critical propagation in the material of a faulted component must be considered. Based on this understanding, either of two beneficial actions could be taken: a corrective one to perform maintenance to repair or replace the part, or a compensatory one to reduce system operational loads to extend the life of the faulted part. The informed decision exists only if the diagnostic/prognostic system has the ability to detect that the fault exists, isolate it to the specific component, and assess its severity in a timely manner.

A stochastic physics-based model of a turbine blade will be used to describe the modeling approach described above and is shown below in Figure 12. Although each component prognostic modeling procedure is different based on the failure modes being predicted, a process that utilizes the raw, database and processed diagnostic data through a physical-based model is still applicable.

The factors and associated level of uncertainty that most directly effect the remaining useful life on a component must be identified in this physical model. One of these factors specific to a turbine blade is the steady stress at the critical locations in the root region. The uncertainty associated with this steady stress due to variations in the operating environment, temperatures, manufacturing tolerances, etc must be accounted for with a statistical distribution computed by testing and field service data. Another factor critical in predicting turbine blade life is the dynamic stress as a function of the uncertainty in harmonic excitation and response characteristics. The strength or resistance capability of the material must also be considered as a function of the uncertainty in material properties.

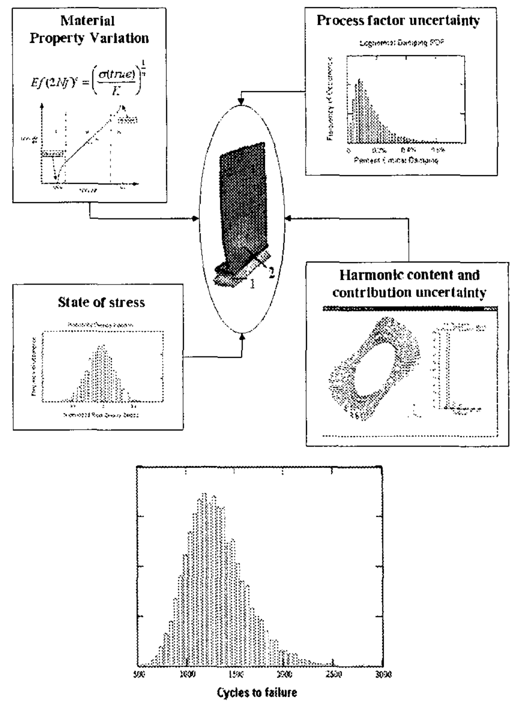

Figure 12 Physics-Based Stochastic Model

In an effort to describe the process of prognostic modeling for a specific component failure mode (only one aspect of this turbine blade model) the LCF fatigue life for the root location that experiences stress cycling in excess of the material's yield strength will be described in the equation (4) below [6].

$$Nf1_L = \frac{1}{2}\cdot\sigma_L(true)^{\left[\frac{1}{(n-c)}\right]}\cdot K^{\left[\frac{1}{(n-c)}\right]}\cdot Ef^{\left(\frac{-1}{c}\right)} \qquad (4)$$

All of the parameters involved in calculating LCF life have levels of uncertainty associated with them and are therefore given as probability distributions that may or may not be Gaussian. The distributions are combined using a Monte-Carlo simulation. The Monte-Carlo simulation is an automatic process that randomly selects thousands of different values from each of the life-limiting factor distributions. Over the entire simulation, the randomly chosen values are combined to generate the distribution of a parameter that may have been very difficult or impossible to calculate in a strict analytical sense. The result of the simulation is also shown in Figure 12.

The damage due the low cycle fatigue may be given by a non-linear damage accumulation rule proposed by Gary Halford at NASA Langley [7]:

$$Damage = \left(\frac{n_1}{Nf1_L} \right)^{r_1} \qquad (5)$$

The complete turbine blade prognostic model must further account for the other failure modes at other critical locations on the blade, however that is outside the illustrative scope of this paper. The net result, however, is the path that's been taken to determine the current component life consumption as shown in Figure 13.

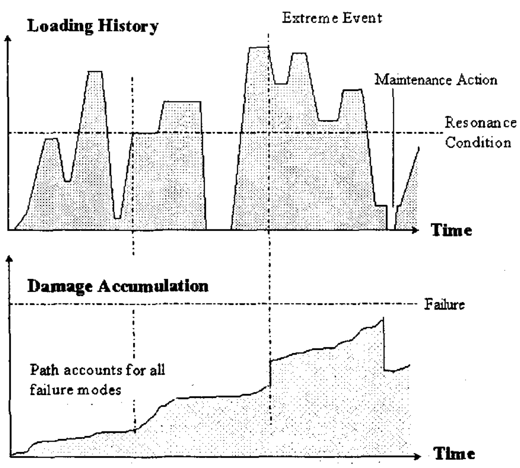

Figure 13 Damage Accumulation and Projected Remaining Useful Life

When statistics on the past operating profile of the machine are tracked, projected future operating conditions and maintenance actions can be estimated and utilized by the prognostic model in order to forecast the remaining life in the blade.

5. Conclusions

An integrated set of turbomachinery health monitoring, diagnostic and prognostic technologies have been presented, that when implemented will offer significant potential for reducing current turbomachinery Life Cycle Costs (LCC). These technologies can be implemented across the entire spectrum of turbomachines from mid-sized pumps to land-based gas and steam turbines as well as aircraft engines. Implementation of these technologies is advantageous in nearly eliminating sensor problems, improving maintenance decision effectiveness by providing early warning of incipient performance and vibration faults and gauging remaining life and predicting future usage associated with critical components.

6. References

[1] Bowerman and O'Connel, Forecasting and Time Series, Duxbury Press, 1993.

[2] M.J. Roemer, and B. Atkinson, "*Real-Time Engine Health Monitoring and Diagnostics for Gas Turbine Engines*," Proceedings of the International Gas Turbine & Aeroengine Congress and Exhibition, Orlando, 1997

[3] M.J. Roemer, and D.M. Ghiocel, "*A Probabilistic Approach to the Diagnosis of Gas Turbine Engine Faults*" 53rd Machinery Prevention Technologies (MFPT) Conference, Virginia Beach, VA, April 1999

[4] T. Kohonen, Self Organizing and Associative Memory, New York, Springer-Verlag, 1987

[5] P. Bjerager, "*On Computational Methods for Structural Reliability Analysis*", Proceedings of the International Workshop on Structural System Reliability, Boulder, CO, 1988

[6] G.J. Kacprzynski, and M.J. Roemer, "*Outage Interval Optimization of a Large Fossil Unit Using Stochastic Lifing Techniques for Blade Failure Risk Evaluation*", Proceedings of the IJPGC, Baltimore, MD, 1998

[7] G. Halford, "Cumulative Fatigue Damage Modeling – Crack Nucleation and Early Growth" First International Conference on Fatigue Damage, Sept, 1996

Dr. Michael J. Roemer is the Director of Engineering at Impact Technologies in Rochester, NY and Adjunct Professor of Mechanical Engineering at the Rochester Institute of Technology. He has over 12 years experience in the development of real-time, state-of-the-art, automatic monitoring and diagnostic systems for large turbomachinery including large steam turbines and gas turbine engines. He has developed several diagnostic and prognostic capabilities for turbomachines utilizing probabilistic methods that are directly linked to maintenance planning and operation of turbomachines. He is also currently the Chairman of the Machinery Failure Prevention Technology (MFPT) Society a Division of the Vibration Institute.

Mr. Gregory J. Kacprzynski is a Project Manager at Impact Technologies with 5-yrs. experience in the development and testing of diagnostic/prognostic R&D systems for compressors, pumps, transmissions, gas and steam turbines. He has been involved in developing real-time, intelligent health monitoring systems for gas turbine engines for on-wing and test cell applications. Greg has published papers and developed technologies in the areas of turbomachinery health monitoring, prognostics, and maintenance optimization.