Embedded Prognostics Health Monitoring

Dr. Frank L. Greitzer
Chief Scientist
Pacific Northwest National Laboratory
frank.greitzer@pnl.gov

Ronald A. Pawlowski
Senior Research Engineer
Pacific Northwest National Laboratory
ron.pawlowski@pnl.gov

KEYWORDS

Prognostics; engine health monitoring; embedded diagnostics and prognostics; predictive maintenance; logistics.

ABSTRACT

The Pacific Northwest National Laboratory (PNNL) has conducted R&D for the US Army on prognostics health monitoring (PHM). The main focus of the work was to demonstrate the feasibility of developing an onboard PHM system for the gas turbine engine used on the M1 Abrams tank. Research was performed on methods for real time, onboard prognostics/engine life expectancy forecasting, and a prototype system was designed, developed, and installed on several test tanks. The purpose of this presentation is to review the approaches to PHM employed in this research, to provide an overview of the PHM prototype and results obtained to date with data collected in the field, and to describe current work aimed at developing a PHM capability for diesel engines.

INTRODUCTION

Prognostics is the process of predicting the future state of a system. Prognostics systems comprise sensors, a data acquisition system, and microprocessor-based software to perform sensor fusion, analysis, and reporting/interpreting of results with little or no human intervention, in real time or near real time. The US Army intends to make extensive use of prognostics technology on weapons platforms, support vehicles and even munitions.

With technical assistance from Pacific Northwest National Laboratory (PNNL), the US Army Logistics Integration Agency (USALIA) has been assessing the state-of-the-art in prognostics technology, extending the capabilities of the technology, and examining implementation approaches to provide maximum benefit to the Army from prognostics data. PNNL has investigated prognostics applications for several US Army vehicles, performed research on methods for real time, onboard prognostics/engine life expectancy forecasting, and developed a feasibility demonstration prototype onboard Prognostics Health Monitoring (PHM) system for the gas turbine engine used on the M1 Abrams tank. The purpose of this paper is to review the approaches to PHM employed in this research, to provide an overview of the PHM prototype and results obtained to date with data collected in the field, and to describe current work aimed at developing a PHM capability for diesel engines.

APPROACH TO PROGNOSTICS

An engineering and cost/benefit approach is used to guide PHM development. Basic steps that are addressed for a given platform in deciding the scope of PHM are:

- Identify major maintenance/life-cycle cost drivers for the platform
- Assess root-causes and engineering/physics of failure that underlie component failures
- Examine onboard existing sources of sensor data that apply to the component failures
- Identify possible additional sensors to improve real time diagnostics and prognostics
- Assess feasibility and cost/benefit of equipment modifications/upgrades to support PHM.

COST/BENEFIT APPROACH

It will take many years and substantial investments to implement fully the Army's vision for self-reporting platforms with embedded diagnostics and prognostics capabilities. This also requires a fundamental change in logistics system architectures that support proactive business practices. The Army Transformation, with its concomitant business process changes and technology insertion, is in part driven by a cost/benefit approach to tackle the question of which platforms get how much capability in what time frame [1]. The Army's cost/benefit analysis (CBA) approach considers development, procurement and operating and maintenance costs and both monetary and non-monetary benefits [2]. Non-monetary benefits include creating tactical advantage in combat situations, improving the readiness of combat forces, and reducing the amount of logistics support required on the battlefield. Monetary benefits include near-term reductions in operations and support costs and reductions in the life-cycle costs of owning and operating a platform. Near-term cost reductions might result from reduction of collateral damage from catastrophic failure of components, reduction in the number of maintenance actions that replace good parts as the result of misdiagnosis, and reduced customer wait time resulting from anticipatory maintenance actions that ensure that parts, tools and personnel are in the pipeline at the earliest possible time. Life cycle cost reductions might result from reduced overhaul costs, extended time between overhauls, and extended life of components.

The Pacific Northwest National Laboratory (PNNL) has performed several preliminary CBA assessments for the USALIA. The results are being used to guide development of a prognostics implementation strategy. Early work has focused on drive-train components, since they tend to be major cost drivers and the technology for embedded prognostics is maturing rapidly. The results of these efforts are described in [1] and summarized below.

- Turbine Engine Prognostics for the M1 Abrams Tank. The analysis indicated that development and deployment of an engine prognostics system with about a dozen auxiliary sensors (thermodynamic and vibration sensors installed via a wiring harness) would result in a benefit-to-cost ratio of about 11:1. This estimate is in general agreement with outcomes of similar analyses. For example, a CBA conducted for an expert systems-based engine monitoring system for helicopter gas turbine engines resulted in a ROI of about 12:1 [3].

- Drive Train Prognostics for the Interim Armored Vehicle (IAV). A preliminary analysis indicated that onsite fluid analysis would yield a positive return on investment of at least 6:1, and this may be enhanced slightly by developing advanced prognostics software for the IAV diesel engine and incorporating it into the engine's support system.
- Prognostics for Legacy Army Diesel-Powered Vehicles. Although a cost-benefit analysis was not conducted for the US Army legacy diesel-powered fleet, the results of a qualitative assessment indicate that the potential economic benefits of applying best commercial diagnostics and prognostics practices to the US Army's legacy diesel fleet are enormous.

A comparison of military versus commercial diesel engine life for different applications reveals a tenfold increase in service life of commercial equipment versus military equipment. The commercial advantage is partly due to the difference in duty cycles between commercial and military applications. Other factors are the design choices and engineering compromises that must be made when developing a vehicle for combat or combat support. Often the desire to maximize available horsepower while minimizing weight and vehicle size results in a design with less longevity than would be experienced without such constraints.

For legacy military vehicles, opportunities to insert embedded diagnostic and prognostic technology are limited to infrequent equipment modifications, vehicle upgrades, or recapitalization initiatives; resource and schedule constraints further reduce the scope of these upgrades, and the opportunity to leverage these investments to include prognostics. CBAs are needed to prioritize proposed embedded diagnostics and prognostics functionality so that options may be selected to maximize the economic benefits and impacts on sustainability, deployability, readiness, and reliability.

ENGINEERING APPROACH

An onboard PHM system comprises (a) a set of sensors mounted to the engine, (b) electronics for collecting and processing sensor signals, and (c) microprocessor(s) to process and analyze the sensor data and perform statistical prediction analyses. Installed onboard the vehicle, the system senses and records engine operational status and performs diagnosis and prognosis of the engine's condition. To design a PHM system on a particular set of equipment, it is necessary to develop requirements that target high-value system components and specify sensor requirements to support prognostics.

In the analysis conducted for the M1 Abrams tank AGT1500 gas turbine engine, several different deployment options were considered. For example, different collections of sensors were evaluated based on the value of the data provided by the candidate sensors (toward diagnosing and predicting failures) and the cost and ease of installation of the sensors. This analysis led to the development of an initial prototype system called TEDANN (Turbine Engine Diagnostics using Artificial Neural Networks) [4], which was later expanded to incorporate prognostics and a larger number of potential failures/faults—the new version was called REDI-PRO (Real time Engine Diagnostics-Prognostics).

REDI-PRO receives input from 38 sensors mounted on the AGT1500 engine. Of these sensors, 25 are factory installed for engine control and basic diagnostics performed by the engine control unit. The other thirteen sensors—retrofitted to the engine using a wiring harness—include pressure sensors, temperature sensors, and vibration sensors located at strategic points on the engine to provide more detailed thermodynamic picture of the engine's state. The REDI-PRO system architecture includes real time data acquisition, sensor validation, digital signal processing of vibration waveforms, engine health analyses using artificial neural networks and rule-based algorithms, and prognostics analyses that perform engine health forecasting. While the ANNs and rule-based diagnostic algorithms are specific to the AGT1500 engine, the methods and architecture employed in REDI-PRO have general applicability, as do the life-expectancy forecasting methods employed for prognostic analysis.

Figure 1 illustrates the processing flow. The sensor values are first analyzed to determine if they are valid, i.e., within expected operating ranges. If any sensor is determined to be faulty, a modeled value is substituted for the observed value. Artificial neural networks and a set of rules are used to model the sensor values and support sensor validation. Following the sensor validation, a diagnostic module processes the sensor data using rule-based and ANN-based analyses. Rule-based analyses check to see if one or a few sensor values exceed thresholds or fail to follow thermodynamic relationships. ANN-based analyses provide diagnoses of complex faults requiring parallel analysis of a large number of sensors. An unsupervised, self-organizing ANN classifies engine operations into states, such as low-idle, tactical idle, full power, etc. Other supervised, feed-forward ANNs perform engine modeling and pattern recognition to diagnose specific faults and conditions. Such model-based diagnoses are output as parameters that are analyzed in REDI-PRO's prognostic module.

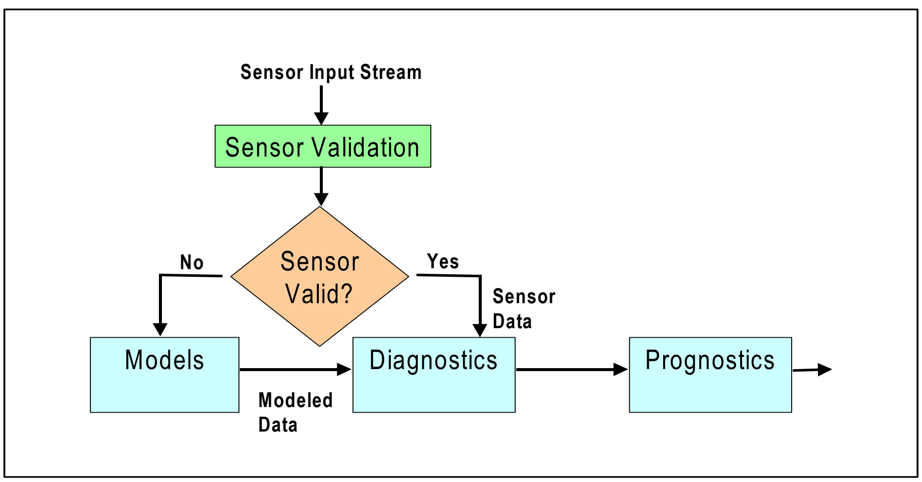

Figure 1. REDI-PRO Analysis Overview.

Table 1 indicates some of the forty-two engine conditions that are monitored by REDI-PRO.

Table 1. Representative List of Engine Conditions Monitored by REDI-PRO

Engine Condition	Analysis Approach/Method
Overtemperature	Simple rule
Coking of bearings	Simple rule
Low and high-pressure compressor efficiency	Supervised feedforward ANN
Overboard air leakage	Supervised feedforward ANN
Bypass gas leakage	Supervised feedforward ANN
Shaft imbalances	Model-based feature extraction from vibration signals
Bearing defects	Model-based feature extraction from vibration signals

REDI-PRO has undergone a limited amount of field-testing. While no rigorous cost savings data have been collected to enable an assessment of the assumptions made in the initial CBA, the development costs have been in line with expectations. If deployed, REDI-PRO could become a component of upgraded onboard systems (such as a new engine controller). Incorporation of this functionality into onboard systems would reduce the costs of deployment and yield an even more favorable ROI. At the present time, the possible integration of the REDI-PRO onboard prognostics capability into the M1 Abrams tank is very much uncertain. Because of plans to replace the AGT1500 engine with a new engine (LV100), the expense of upgrading the old engine is an obstacle, despite the fact that the old engines are expected to remain in active status for at least another ten years. Therefore, the path forward for integrating prognostics into the M1 Abrams platform appears to be most promising for the new engine. Nevertheless, major obstacles still must be overcome in convincing decision makers that embedded prognostics will significantly impact operational and life-cycle costs.

APPROACH TO PREDICTION

The approach to prediction is based in part on research conducted at PNNL under an internally funded Laboratory Directed Research and Development project called Life Extension Analysis and Prognostics (LEAP). The LEAP project ([5], [6]) focused on analytic methods to enhance the quality of predictions by prognostics systems.

To predict a *failure*—the inability or at least serious degradation of the platform to perform its intended function—three things typically must be known: the system's current degree of fault as quantified by a Figure of Merit (FOM); a theory about the progression of the fault, so as to postulate the system's degree of fault at a particular point in time in the future; and the level of the fault, as quantified by the FOM, that will produce a *failure* of the platform. The specification of these factors is typically done through engineering/analytical studies such as Failure Modes and Effects Analysis (FMEA). These analyses and expert judgments yield descriptions of how the system fails, what faults can be measured given available sensors, and the values expected for the sensors when these failures occur. Except in unusually simple cases, the FOM determined for a fault, failure, or system condition is a function of a combination of sensor values (i.e., through sensor fusion), rather than a single sensor. As a result, it is typically *not* sufficient to monitor and trend individual sensor values, independently, to perform diagnostics and prognostics.

As data are collected, regression models are applied to the data to determine trends in FOMs. These FOMs are compared, in real time, to previously established metric failure limits. The point of predicted failure is calculated as the intersection of these lines (see Figure 2). Uncertainty intervals (dashed lines surrounding the trend lines) also may be derived to incorporate uncertainty estimates into the prediction. In Figure 2, predicted time of failure is indicated by time t_2. The method estimates failure occurrence unlikely before time t_1 and normal operations unlikely to extend beyond time t_3.

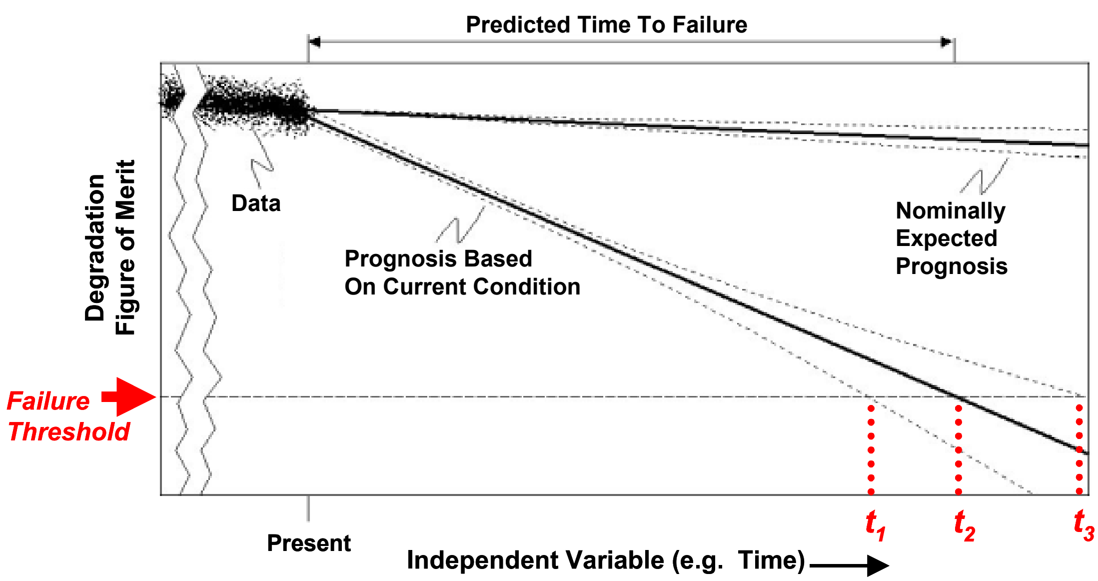

Figure 2. Regression line intersecting failure threshold indicates predicted time to failure.

Because FOMs are mathematical functions of one or more measured variables, normal variation in the FOMs and underlying sensor values must be distinguished from values that indicate degradation. An inevitable trade-off exists between using a large set of data to reduce the consequence of noisy sensor values and inherent system variability and using a smaller set of data to be responsive to changing system characteristics that may occur when a system health problem begins to manifest itself. Use of large amounts of data spanning a long window of data acquisition tends to yield more stable, less variable predictions. However, it also may yield a prediction that is less sensitive to recent changes. Use of a smaller data set spanning the most recent operating history tends to produce predictions with more variability but more sensitivity to current operating conditions (see regression line labeled "Prognosis Based on Current Condition" in Figure 2). The goal of the prognostic analyses is to choose among varying window sizes to maximize the system's adaptability to change while maintaining an acceptable amount of predictive variability/uncertainty. A statistical technique called LEAP-Frog regression [5] was developed to accomplish this goal.

Using a relatively small set of time windows, the analysis makes a prediction at the *current time* for the value (and uncertainty intervals) of the FOM at a *future time*. The method begins with a regression analysis using large time window for data acquisition, which will likely give the best estimate of the regression fit, if the platform is at a constant rate of change of health (maybe steady with a slow rate of degradation). This prediction and an uncertainty distribution about the estimates are tested to see if the prediction is reasonably compatible with the most recent data points. If so, then the regression is used. If not, then the size of the regression window is reduced and the analysis is repeated. This method continues until it yields a small enough window that is compatible with the most recent data points. As a result, this method can detect if the most recent data points indicate a change from the long-term regression (as would be the case if the platform had a change in FOM). In the end, the method uses the longest regression window that does not result in evidence (based on the most recent records) that refutes the assumption of a good linear fit; and then uses this window to predict the future state of the system or its remaining useful life.

RECENT FINDINGS FROM PROGNOSTIC ANALYSES

The REDI-PRO prototype system has been installed on two US Army National Guard tanks for several years, and in 1999 it was installed on some US Army tanks at Yuma Proving Ground for several months. During the data collection period at Yuma Proving Ground there was one occasion when data were recorded prior to and during a failure (the failure occurred as a result of an extreme air filter clog). In deriving the FOM for the air filter clog, the REDI-PRO analysis computes the pressure differential on both sides of the air filter, which is customary. However, in addition to computing the delta-pressure estimate, the analysis also takes variables into account that tend to affect the estimate—in particular, the RPM and temperature—statistically removing their effects, and focusing the prognostics trend analysis on the residual. As reported in [5], the REDI-PRO prognostics analysis was shown to detect a change in trend sufficient to predict the occurrence several operating hours ahead of the detection of an air-filter clog by on-site maintenance staff.

More recently, REDI-PRO data have been analyzed from onboard vibration sensors on M1 Abrams tanks at the US Army National Guard, Yakima Training Center. Vibration analysis methods for fault diagnosis and prognosis on the Abrams tank AGT1500 turbine engine were developed to meet the following unique requirements and constraints:

- This engine's bearing failures are more frequent and more serious than gearbox failures.
- Independent measurements of shaft speed were not available to the vibration analysis.
- Localizing shaft imbalances and bearing defects will improve prognostics capability.
- Cost and deployability considerations limited the number of vibration sensors to install.
- It is not practical to store and download large amounts of raw vibration waveform data.

These factors led to the choice of vibration analysis methods that are quite different from the methods that are typically used in other applications, such as Cepstrum analysis and wavelet techniques [7]. Algorithms had to be developed to provide the broadest possible capability with only a few vibration sensors. The new approach was to develop a parametric model of the vibration waveform based on a postulated submodel for each type of defect. For example, the characteristic waveform excited by an outer race defect was postulated to be a repeating series of damped oscillations. This series repeats at the characteristic ball pass frequency of the outer race that is well described in the literature [8]. Since the load on a defect in the outer race does not vary with the rotation angle of the shaft, the amplitude of the postulated waveform was assumed not to vary with rotation angle (such a variation, if present, would result in sidebands around the defect peak in a frequency spectrum). However, the modeled amplitude was postulated to vary with the rotation speed of the bearing. Models for three types of defects – specifically, outer race defects, inner race defects and rolling element defects - were developed for each bearing or set of bearings. In addition, models of shaft imbalance or misalignment defects were also included. Due to the predominant history of these kinds of defects within the AGT1500, and to limit the size of the model, no other types of defects were included in the analysis. It should be noted that few of these models are independent; in particular, most of the defect models depend on rotational speed of a particular shaft.

Models of the vibration waveforms produced by each of the defects were developed and superimposed to obtain an overall vibration model of the engine. The resulting model was superimposed and compared to the actual waveform. Using a model fitting procedure, the global difference (i.e. Chi-squared) between the model and the actual data is minimized by varying the model parameters. The resulting model parameters represent diagnostic figures of merit such as imbalance amplitudes and defect amplitudes. The implementation of the model-fitting algorithm on a desktop computer environment is ongoing. Early results limited to the detection of shaft imbalances in test waveforms indicate that the procedure has promise. However, the bearing defect equations and their associated partial derivative equations are significantly more complex than the analogous equations for shaft imbalance. Therefore, a significant amount of analysis and programming remains before the method can be shown to detect bearing defects.

DIESEL ENGINE APPLICATIONS

Because of the age and large number of diesel engines on military vehicles, there is interest in developing embedded diagnostic and prognostic capabilities for these engines. Despite significant differences in the physics and operation of diesel versus gas turbine engines, many of the general diagnostic methods used by REDI-PRO for gas turbine applications remain pertinent, such as methods for rule-based and model-based fault severity assessment. Furthermore, the prediction methods as embodied in the LEAP-frog methods also remain generally applicable. Clearly, however, there are differences in the specific types of faults of interest and methods for assessing the severity of those faults.

Four basic categories of faults are of interest in diesel engines: compression problems (faulty piston rings, bad valves or valve seats); injector or fuel system problems; turbocharger problems; and oil condition problems. Faults in the first three of these fault categories can result in reduced engine power. They can also result in excessive smoke, which has undesirable environmental as well as combat stealth consequences. Decreases in fuel economy are also of concern. Finally, problems with compression or with the fuel system may result in difficulties in starting the engine. Oil condition is a concern with respect to the cost of routing maintenance, since significant savings in maintenance costs can be obtained if the oil is maintained on condition rather than in accordance with a preventive maintenance schedule. These faults can also be interdependent, in that the presence of one can increase the rate at which another will progress. For example, bad piston rings result in excessive blowby, thus allowing more combustion products into the crankcase to contaminate the oil. Working in the other direction, bad rings can allow excessive passage of oil from the crankcase to the combustion chamber, thus resulting in the formation of coke deposits, which interfere with valve seats.

For diesel engines more than gas turbines, it is desirable to perform diagnostics and prognostics with a minimum of additional sensors installed on the engine. Sensor retrofits are costly in terms of labor and downtime, and since the per-unit cost of diesel engines is much less than that of gas turbines, the benefit-to-cost ratio is reduced for diesel engines. Therefore, we are investigating ways to use existing sensors and sensor wiring in novel ways to perform diesel engine diagnostics and prognostics.

CONCLUSIONS

This research demonstrates the feasibility of applying prognostic health monitoring systems to gas turbine engines—and, it is hoped, through continued research and reductions in developmental costs, to less expensive diesel and reciprocating engines. Onboard, self-reporting prognostic systems will enable maintenance personnel to be dispatched in anticipation of failures with the right equipment, parts, and tools; thereby minimizing equipment down time and customer wait time.

REFERENCES

[1] Greitzer, F. L., Hostick, C. J., Rhoads, R. E., and Keeney, M. Determining *how* to do prognostics, and then determining *what* to do with it. *AUTOTESTCON 2001*, Valley Forge, Pennsylvania, August 20-23, 2001.

[2] CEAC. Economic Analysis Manual. U.S. Army Cost and Economic Analysis Center. Department of the Army, July 1995.

[3] Collinge, K, and K. Schoff. TEXMAS- An expert system of gas turbine engine diagnosis and more. SAE Technical Paper 871737. Aerospace Technology Conference, Long Beach CA Oct. 5-8, 1987.

[4] Greitzer, F. L., L. J. Kangas, K. M. Terrones, M. A. Maynard, B. W. Wilson, R. A. Pawlowski, D. R. Sisk, and N. B. Brown. Gas Turbine Engine Health Monitoring and Prognostics. SOLE '99 Symposium, August 31- September 2, 1999, Las Vegas, Nevada.

[5] Greitzer, F. L. and T.A. Ferryman. Predicting Remaining Life of Mechanical Systems. ASNE Intelligent Ship Symposium IV, April 2-3, 2001, Philadelphia, Pennsylvania.

[6] Greitzer, F. L., E. J. Stahlman, T. A. Ferryman, B. W. Wilson, L. J. Kangas, and D. R. Sisk. Development of a Framework for Predicting Life of Mechanical Systems: Life Extension Analysis and Prognostics (LEAP). SOLE '99 Symposium, August 31-September 2, 1999, Las Vegas, Nevada.

[7] Grabill, P., J. Berry, L. Grant and J. Porter. Automated Helicopter Vibration Diagnostics for the US Army and National Guard. American Helicopter Society 57th Annual Forum, May 9-11, 2001, Washington, DC.

[8] Widner, R.R. and W.E. Littman, Bearing Damage Analysis. National Bureau of Standarsds special publication No. 423, April 1976.

This research was performed by PNNL for the US Army Logistics Integration Agency under a Related Services Agreement with the U.S. Department of Energy, Contract DE-AC06-76RLO 1830.