Integrated Software Platform for Diagnostics and Prognostics with Air Vehicle HUMS

Carl S. Byington[1], Romano Patrick[1], Matthew J. Smith[1], George J. Vachtsevanos[2]

1. Impact Technologies, LLC., 200 Canal View Blvd., Rochester, NY 14623, U.S.A

2. Dept. of Electrical and Computer Engr., Georgia Institute of Technology, 777 Atlantic Dr. NW, Atlanta, GA 30332-0250, U.S.A.

Abstract

Although typical Health and Usage Monitoring Systems (HUMS) intend to support a transition from scheduled part replacements to performing maintenance upon evidence of need, they generally exhibit a limited ability to diagnose component faults early and accurately in complex systems such as a helicopter drive train. Consequently, the traditional approach to implementing Condition Based Maintenance (CBM) programs is slow, requires substantial amounts of human supervision, and ultimately lacks prognostics. A team of developers from Impact Technologies, the U.S. Army Research Laboratory and the Georgia Institute of Technology, with support from the U.S. Army, have been working over the past 3 years to develop a methodology to improve the performance of U.S. Army helicopter CBM systems, and facilitate transition from scheduled maintenance to implementation of prognostics. This paper presents an integrated diagnostic enhancement and prognostic architecture and the corresponding software suite, and discusses how a hybrid and systematic approach to sensing, data processing, fault feature extraction, fault diagnosis, and parallel health-based and usage-based failure prognosis can be used to improve the performance of HUMS and CBM activities. The paper also provides an application example, and discusses validation steps for the new approach to diagnostic enhancement and prognostics.

Introduction

In recent years, the U.S. Army has witnessed various helicopter component failures that are currently driving the need for improved health monitoring and fault prediction that will be implemented under the broader initiative for Condition Based Maintenance (CBM) of the U.S. Department of Defense, known as CBM+. Since some of the most effective methods for diagnosing faults in components involve the analysis of sensor signals and the derivation of condition indicators (Smith et al., 2009, Watson et al., 2007), existing hardware/software health monitoring systems and digital source collectors (DSC), such as VMEP/MSPU (Branhof et al., 2005) and IMD-HUMS/IVHMS (Dora et al., 2004), collect vibration and other pertinent flight regime data and attempt to detect a fault condition using condition or health indicators derived via data processing algorithms. However, there is room for improvement in two key general aspects: (1) developing supporting technologies to enhance fault detection performance and capabilities (Byington et al., 2007; Byington et al., 2008a), and (2) predicting the remaining life or proper maintenance times for worn or failing components with sufficient warning time.

To specifically address the improvement of fault detection and the implementation of failure prediction methods, the U.S. Army Research Laboratory, Impact Technologies, LLC, and the Georgia Institute of Technology worked collaboratively for 3 years to develop, test and evaluate modular software components that provide enhancements to diagnostic systems already in service and add failure prognosis capabilities for critical Army aircraft components. The work is part of the now completed "Air Vehicle Diagnostic and Prognostic Improvement Program" (AVDPIP), whose ultimate goal was to complement existing Army Digital Source Collector (DSC) systems in support of early warning of impending failure conditions and optimization of aircraft repair, maintenance and overhaul practices.

Although typical Health and Usage Monitoring Systems (HUMS) support the transition from scheduled part replacements to performing maintenance upon evidence of need, they generally exhibit a limited ability to diagnose component faults early and accurately in complex systems such as a helicopter drive train. Consequently, the traditional approach to performing Condition Based Maintenance (CBM) requires extensive human supervision (including manual, case-by-case data analysis and results verification) and ultimately shuns prognostic activities.

The software modules developed under AVDPIP use a methodology and corresponding software architecture to integrate the operations of sensor data validation and pre-processing, advanced fault feature extraction, and different supporting techniques for effective fault diagnosis to overcome the aforementioned typical challenges, and provide implementation of reliable prognostics. The development contains generic software components and algorithms that build upon model based and data driven methodologies that are applicable to a variety of components in complex systems such as those found in a helicopter drive train. This paper presents the software components and capabilities, and illustrates their use with the the case study of a helicopter drive train bearing, integrating technologies presented in a preceding paper by Smith et al. (2009).

Application Example: H-60 Helicopter Bearing

As reported in Smith et al. (2009), a representative example of the CBM challenges discussed earlier is given by a bearing inside the oil cooling subsystem of a family of helicopters in service for the U.S. Army. The transition to performing maintenance upon need (condition based) instead of basing decisions on flight hours is of interest due to the criticality of the component and the relatively high incidence of replacements and faults. Enhanced diagnostic algorithms capable of detecting a fault in its early stages of development, and the use of a bearing life prognostic system would both clearly support the intended transition.

The oil cooler is a component of the H-60 tail rotor drive train assembly whose primary function is to cool the helicopter transmission lubricant while transmitting power to the tail rotor drive shaft through the oil cooler shaft. Bearings with heavily contaminated grease and exhibiting corrosion on the raceways, balls and cage surfaces have been found in the field in multiple instances (Baker et al., 2007). Flight hour limits have been imposed and later updated for the bearing, but there are

questions about the adequacy of the limits and the potential impact on maintenance scheduling (Smith et al., 2009). Hence, an effective diagnostic and prognostic system in support of life extension would offer substantial benefits.

The fault modes of the oil cooler bearing have been shown to be detectable through vibration analysis (Keller et al., 2005). Vibration data is collected from accelerometers mounted on the oil cooler housing. When bearing damage has initiated on rolling surfaces, specific signal frequencies associated with the location of the defect are excited. The amplitude and time duration of the defect frequency are expected to be indications of defect severity (Smith et al., 2009). Various features or condition indicators, both in the time and frequency domain, serve as a means for detecting these faults, in line with standard vibration based bearing health monitoring.

The AVDPIP team used different techniques to analyze the vibration and corresponding features prior to performing diagnostics and prognostics of a bearing's health. For example, Figure 1 illustrates on the comparative performance of traditional techniques vs. use of shock pulse amplification software with an active band selection (ABS) algorithm vs. use of feature fusion techniques, as reported by Smith et al. (2009). More details about these techniques are listed further below. Also, an anomaly detection methodology is used to detect as early as possible and with minimum false alarms an incipient fault (Zhang et al., 2008).

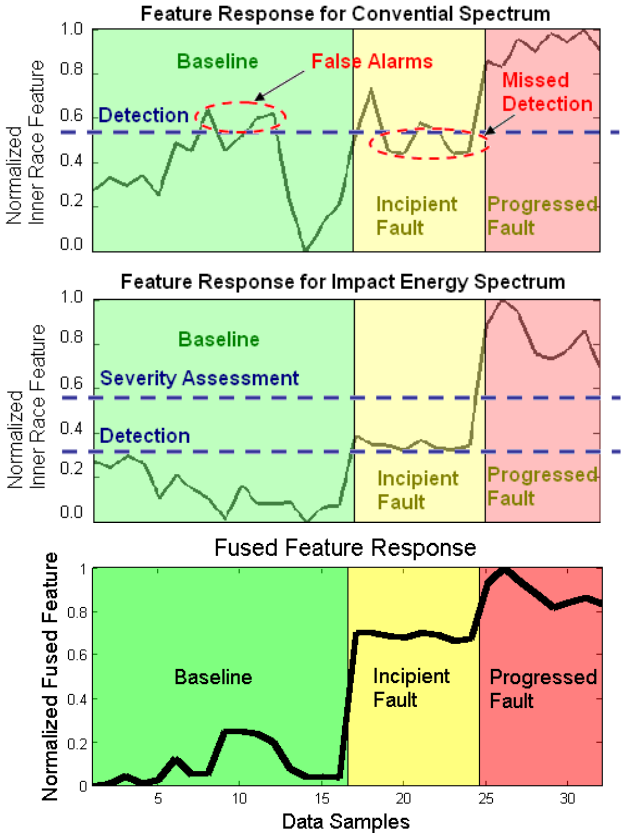


Figure 1. Comparison of bearing vibration processing results with different detection techniques (from Smith et al., 2009)

The AVDPIP Architecture

CBM programs commonly attempt to detect a fault condition in a component whose condition is under monitoring. Condition monitoring uses techniques to collect vibration or other sensory data, as well as other pertinent flight regime information, to derive condition or health indicators and interpret or classify them via data processing algorithms. However, as suggested previously, it is not uncommon to find that early detection, fault diagnostic accuracy, and prognostic capability are inadequate in CBM systems. Causes of these limitations, which ultimately lead to an underrepresentation of prognostics in fielded CBM systems, include (Zakrajsek et al., 2006): (i) the sensitivity of sensors and condition indicators to signal noise, specific fault modes, and variations in environmental and operating conditions; (ii) the performance of diagnostic processes, which attempt to make a health assessment using condition indicators that, in many instances, are chosen empirically and without full understanding of their fleet-wide behavior; (iii) the uncertainties in damage or wear progression as well as in future usage, and the corresponding difficulty in implementing a reliable degradation tracking algorithm that reliably captures these uncertainties; (iv) the inherent risk of relying on an algorithm or prognostic system that attempts to predict how much longer a component can operate even if it is expected to fail; and (v) the lack of verified fault case studies with sufficient representative data as well as diagnostic and prognostic algorithm validation.

It is possible to increase the performance, accuracy and detection capabilities of typical CBM systems by utilizing adequate data preprocessing techniques, advanced condition indicator evaluation, and detection and diagnostic enhancement algorithms. Furthermore, it is also possible to safely implement prognostic health assessments with the integral use of (a) the aforementioned enhanced diagnostic processes, (b) appropriate component degradation models, (c) uncertainty representation algorithms, (d) usage, loading and operating conditions data, and (e) the use of available calibration and validation data sets possibly supplemented with seeded fault tests.

To improve the performance of CBM diagnostic processes and facilitate reliable implementation of prognostics, the AVDPIP team has developed a methodology and a set of software components that are capable of addressing the five challenges listed earlier (i through v). We refer to this methodology as the “AVDPIP architecture” for the integrated operation of fleet data analysis, diagnostic enhancement, and safe implementation of prognostics.

The AVDPIP architecture was developed with the objective of maintaining modularity to allow for extensibility to a wide variety of systems and components, but still uses a systematic approach to integrate sensing, data processing, fault feature extraction, fault diagnosis, and failure prognosis, as illustrated in Figure 2.

Signal processing and CI/feature extraction

Data is retrieved from an available repository such as preexisting databases, a digital source collector (DSC), a HUMS system, a data acquisition system, etc. The AVDPIP architecture allows for two types of data to be used: (1) “raw” data from sensors (i.e., unprocessed signals), or (2) condition indicators (features) preprocessed

from the raw data by preexisting (e.g., traditional) methods and systems. In the case that raw data is available for processing, advanced feature extraction techniques can be used to derive condition indicators with increased accuracy and more desirable behaviors. A variety of techniques can support this task, including signal noise removal (or data de-noising, active noise cancelation, etc.), sensor validation (to remove corrupted or invalid signals), advanced demodulation and filtering, and analysis band selection optimization, among others. On the other hand, if preexisting condition indicators are used, the architecture can proceed to directly use them in analysis and post processing tasks as described below.

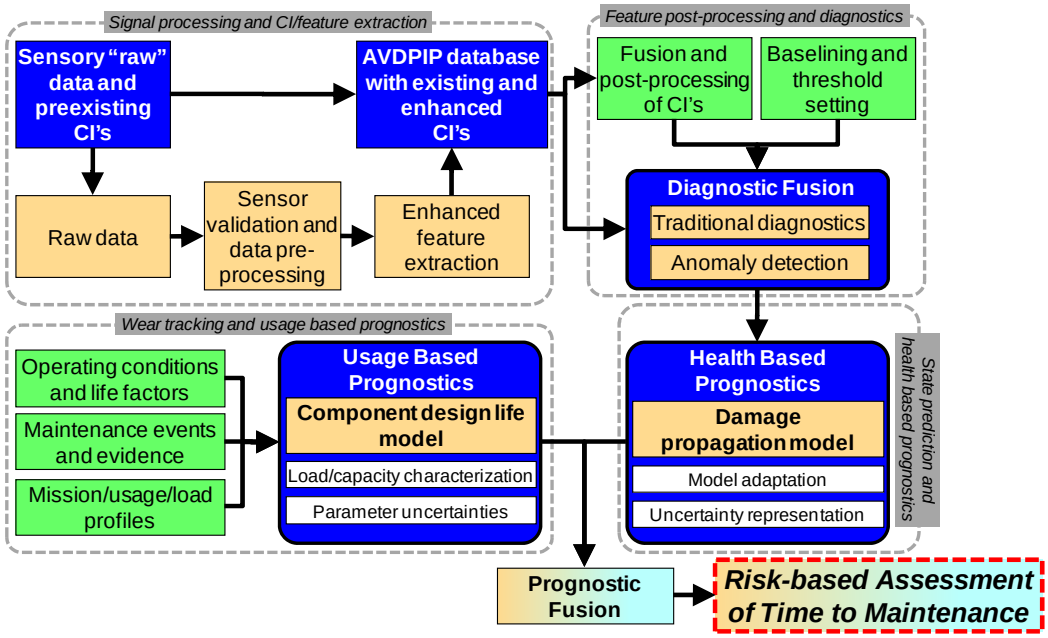


Figure 2. Block diagram of the AVDPIP architecture.

Feature or condition indicator post-processing

After features or condition indicators are generated, the AVDPIP architecture calls for their processing and analysis to perform: (1) anomaly or fault detection, (2) fault isolation and/or mode classification (source and category), and (3) fault identification (severity assessment), as necessary. Data processing and signal conditioning techniques can be used to provide diagnostic evidence for performing these three tasks. Some relevant supporting techniques include feature normalization and fusion, statistical analysis and optimal threshold setting, and the use of anomaly detection and diagnostic algorithms. Details about each of these varying techniques are discussed in this paper.

Diagnostics and prognostics

The processes listed above provide a means to improve the ability of diagnostic activities to detect early-stage faults and mitigate the adverse effects introduced by signal noise, difference in behaviors due to multiple fault modes, and variations in environmental and operating conditions (such as loads, speeds, flight regimes, etc.). Addressing these weaknesses is a key prerequisite for implementing effective prognostic algorithms. However, even if these techniques are proven

successful at improving fault detectability and diagnostic performance, the realization of accurate prognostic health assessment is still challenging, and remains as an activity worth of care to minimize potential risks. To reduce diagnostic and prognostic performance risks, the AVDPIP architecture combines two types of diagnostic assessment with two different approaches to prognostics. The operation, benefits and risks of these four techniques are as described in Table 2.

Table 1. Complementary diagnostic and prognostic strategies of the AVDPIP program

	Technique	Operation	Advantages	Limitations and risks
Diagnostics	**Traditional Diagnostics**	Compares feature or CI values against pre-established thresholds indicative of hazardous conditions; the thresholds are derived from statistical studies of fleet wide behaviors and known cases of faults	Useful for monitoring fault modes that are known to be severe, frequent and "testable" on the basis of a Failure Modes, Effects, and Criticality Analysis (FMECA)	Useful only for a finite number of fault modes with well identified and sufficiently understood behaviors. Good correlation between specific condition indicators and each fault mode must be sufficiently well established; monitoring is applicable to predefined and highly repeatable operating conditions, or else the risk for false alarms increases considerably
	Anomaly Detection Diagnostics	Uses a model of the system under consideration and the observation of an "innovation" or "discrepancy" between the actual plant output and the model output, for all possible operating conditions, to detect an unanticipated fault	Provides greater "coverage" in terms of the number of fault modes that can be detected, isolated, and identified. Can in effect detect a fault as an "anomalous" or "extraneous" behavior in the system that should be given attention	Requires an accurate model of system behaviors and operating modes. Performance limited by model accuracy; fault mode coverage limited by model breadth or complexity. Requires a set of baseline (reference) conditions representative of healthy operating conditions that are difficult to define and bound adequately in the context of potential, not-fully-understood system anomalies. Inadequate "baselining" can lead to poor detection performance
Prognostics	**Usage Based Prognostics**	Combines usage monitoring (load/stress tracking) with a wear or life model to estimate the rate at which a component degrades, accumulates damage or "consumes" its design life. A prognosis is based on remaining design life at any given time	Can track wear or degradation in parts and, by using planning triggers dependent upon usage of assets, can also offer support for replacement logistics that are "smarter" than scheduled maintenance programs. Reduces risk of life limiting approach when usage is intense	Does not take into account unanticipated faults. Most applicable for components that degrade as intended by design, although it can use "life factors" or adjustments for well-established and well-understood conditions leading to modified rates of degradation (outside of nominal). Requires an estimate of future system usage to provide an accurate prognosis or else a "worst-case-scenario" must be assumed for future system usage, potentially leading to excessively conservative results
	Health Based Prognostics	Combines health monitoring (damage assessment via diagnostics) with a damage progression model to estimate the rate at which *a faulted component* continues to degrade	Capable of tracking a degraded condition (once a fault or anomaly is detected) and providing early warnings for components in need of maintenance due to unanticipated faults or the presence of damage. It can potentially support life extension of degraded components because *prognosis estimates future damage progression in a faulty component*	Works only as a follow up to the detection of a fault or anomaly (i.e., requires a positive diagnosis of a fault or anomaly), which immediately implies that the *prognosis is operating over a damaged component that is at risk of failing*. Because planning actions are triggered depending on system condition, there is the risk that an error in the characterization of future damage progression can lead to (a) maintenance actions need to be rescheduled if the prognosis changes, or (b) a prognosis with a long lead time can lead to a false sense of safety

Usage Based Prognostics

The usage-based prognostics approach incorporates reliability data, life usage models and varying degrees of measured or proxy data to forecast durability of a component. The forecast is based on actual usage when such is known and a suitable representation (load and condition data series) is available. Incipient fault detection may not be available due to sensor or fault mode coverage limitations, but on the other hand, the usage based prognostic can make durability assessments even if no fault is detected. At the heart of the usage based prognostic algorithms of AVDPIP, we find the use of component life models that are used to determine the durability of vehicle components taking into account usage patterns and parameter uncertainties. The approach is illustrated in Figure 3. An example of use of this approach with our example case of helicopter bearings is illustrated in Figure 4. It should be noted that this algorithm is expected to be useful for anticipated types of faults. For example, in the case of a bearing, the model will be able to characterize specific degradation modes and parameters, including fatigue wear, effects of variable loading, effect of certain operating conditions in the form of life factors (typical in bearing lifing), etc. However, certain types of faults cannot be characterized by the usage-based approach proposed. For example, faults that would make the usage-based prognosis inapplicable include manufacturing defects or anomalies, improper installation, inadequate maintenance or unaccounted-for operating conditions.

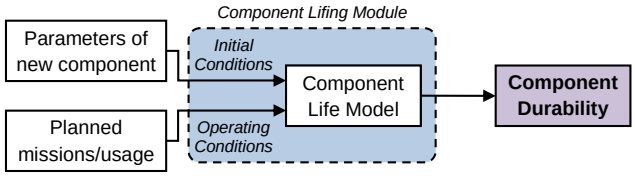


Figure 3. AVDPIP approach to model-based component lifing

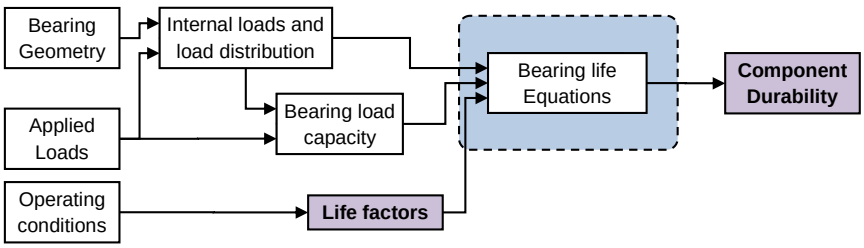


Figure 4. Approach to usage-based bearing lifing used by AVDPIP

Health Based Prognostics

The health-based prognostics approach involves utilizing the assessed health or diagnostic fault classifier output to predict evolution of a component fault. Feature trending or physics-of-failure based prediction may then be used. Incipient fault detection and diagnostic isolation is absolutely necessary, and thus the health based prognostic system cannot operate until a fault is detected. The AVDPIP program uses Particle Filters to perform feature trend predictions (Zhang et al., 2008). Particle filtering is an application of Bayesian state estimation that calculates an a posteriori

probability density function (PDF) of a state of a system based on a priori observations or measurements. If the calculation of the future state of the system is extended in multiple steps with the use of a model, the particle filtering algorithm can perform long term predictions. In this case, the system observations are initially used to build a PDF of the "present" or "current" system condition, as illustrated conceptually in Figure 5.

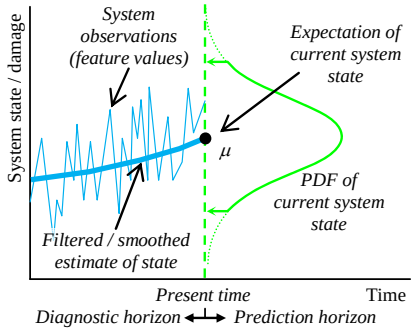


Figure 5. Determination of the state of a system as a PDF based on feature values

This PDF is then sampled into "particles" representative of potential system states with individual weights. Using the model, the prognostic algorithm simulates the progression of the weights in time to do a prediction of possible future system states, as illustrated in Figure 6.

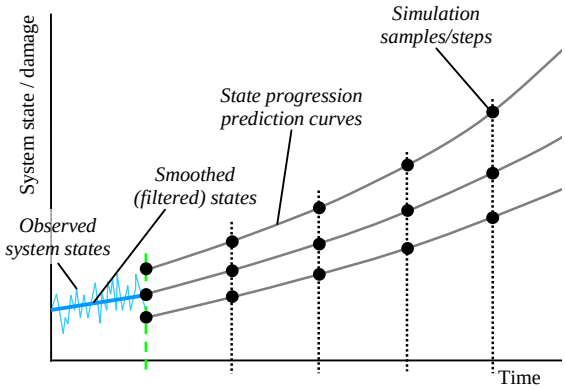


Figure 6. System state prediction and progression curves

Just as with the initial state, future states of the system can be represented by PDFs. Once the progression of the system state has been determined, the algorithm can be used to predict the time required for the system to reach a condition of interest, such as a need for maintenance. The condition predicted is represented by a "prediction threshold" line. Because there is uncertainty in the future system states (as represented by the different state progression curves), there is also uncertainty in the predicted time to reach the threshold. This uncertainty in time is represented also by a PDF, referred to as the "time-to-threshold" (TTT) PDF. The definition of prognostic confidence is tied to how the area of the TTT PDF is divided. To determine the minimum time remaining to reach the prediction threshold, called the "just-in-time" point, a confidence specification is required. Figure 7 illustrates how a 95% prediction confidence is used to determine the just-in-time point. The AVDPIP software suite implements the processes and algorithms described above.

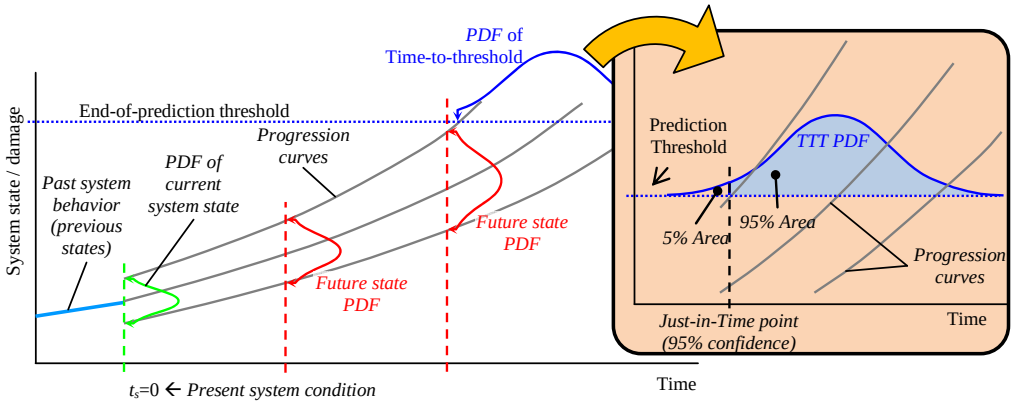


Figure 7. Determination of the prediction time to reach a prognostic threshold with a given prognostic confidence (the inlay box provides an example using 95% confidence)

The AVDPIP Software Suite

The AVDPIP software suite uses the AVDPIP architecture to integrate data pre- and post-processing with diagnostic and prognostic operations. The major functional software modules are described in Table 3.

Table 2. Functional software modules of the AVDPIP software suite.

Module name	Primary purpose	Description
TEDS™: Toolkit for Enhanced Diagnostics	Data validation and feature extraction	Sensor validation, raw sensor data enhancement and preprocessing, and extraction of features (condition indicators) from vibration signals
IMDx™: Integrated Mechanical Diagnostics Database Export Tool	Feature exporting	Exporting of condition indicators from Army HUMS databases into AVDPIP databases
NBATS™: Normalization, Baselining and Thresholding Software	Statistical analysis	Analysis of feature value distributions, normalization of data, and determination of optimal detection thresholds
Fusion Module	Sensor and Feature Fusion	Fusion of sensory data to decrease effects of noise and increase detectability of faults, and fusion of feature data into “meta-features” to improve classification and trending operations
MPUGS™: Mission Profile Usage Generation and Simulation	Usage profile definition	Definition/creation of usage / loading / mission profiles driving wear and degradation rates in operational components
PIP™: Prognostics Improvement Program	Diagnosis and prognosis	Analysis and determination of the progressing condition of a component using data from TEDS, IMDx, NBATS, Fusion module and MPUGS, as well as different diagnostic and prognostic algorithms.

TEDS

The TEDS tool performs advanced data processing and signal conditioning techniques aimed at providing enhanced diagnostic evidence for improved air vehicle diagnostics and prognostics. The tool utilizes such modules as:

- The ImpactEnergy™ shock pulse amplification software, which uses a multi-step signal processing routine prior to feature extraction that increases the visibility of shock-pulse events indicative of specific bearing faults, thus uncovering frequency spectrum peaks that are otherwise hidden in the broadband spectrum, and allowing for detection of faults in their incipient state (Sheldon et al., 2007; Byington et al., 2006a)
- The ABS™ (Automated Band Selection) software, which maximizes fault detection capabilities by using techniques to identify the best regions of the broadband spectrum to perform more effective fault frequency demodulation (potential resonances) for bearing vibration feature/CI extraction
- The FirstCheck™ software for real-time sensor validation, which uses time and frequency domain methods to detect various potential sensor faults including faulty connections, loose accelerometer mounts, and damaged accelerometer elements

TEDS allows a user to control how the analysis algorithms will be applied to raw vibration data files to extract feature values. It uses a robust database structure allowing to catalog data from different data sources (aircraft platforms), vehicle components, and sensors. Detailed meta-data (descriptive information) and date/time information is included in the catalogs. Feature extraction can be performed on data subsets selected by the user and saved the options for a given extraction saved into "analysis" files that allow for repeated, comparable extractions over different data sets. Analysis options include the selection of feature extraction algorithms, sensor validation, and vibration data filters among other advanced processing options. Additionally, TEDS is prepared to perform near real-time analysis on data that is streamed continually, opening the possibility to realize condition monitoring of operating assets.

IMDx

The Integrated Mechanical Diagnostics Database Exporting Tool (IMDx) is designed to extract condition indicators from Army databases to allow diagnostic/prognostic analyses with the PIP software (see below), but using condition indicators that have already been calculated with existing systems.

NBATS™ and Fusion modules

These tools provide feature pre-processing and post-processing algorithms to enhance the performance of diagnostic and prognostic analysis on series of feature values for groups or individual aircraft. The tools allow:

- Feature data normalization and rescaling to allow for more effective analysis of data across fleets, operating conditions and flight regimes.
- Definition and analysis of baseline distributions of feature values for different fleets, aircraft, operating conditions, flight regimes, etc.
- Design of optimal detection and prediction thresholds to enhance diagnostic/prognostic performance.
- Fusion at the sensor level, combining vibration signals of multiple sensors to decrease the effects of random noise and increase visibility of subtle signs of a fault (Byington et al., 2008b)

- Fusion of feature series into meta-features that maximize fault class separation, increase fault detection confidence, and exhibit improved monotonicity and correlation with fault conditions.

MPUGS

MPUGS is a tool for generating/assembling mission or usage profiles for Army helicopters. The tool can generate a set of loads on components for a usage pattern that an analyst specifies. The load sets can be representative of known past operation of a given aircraft (i.e., cataloguing missions already flown and the corresponding operating conditions experienced by the asset) or sets of expected/planned missions that aircraft are expected to undergo. The definition of expected/planned missions can be used to perform prognostic analyses and what-if-scenario simulations accurately. MPUGS allows an analyst to store load profiles into files that are readable and ready for use by the prognostic algorithms (PIP tool).

PIP

The Prognostics Improvement Program software application is used to perform diagnostic and prognostic assessments of the health of a given aircraft component based on feature values extracted from sensory data and on usage and maintenance information as specified by a user. The feature values are obtained from the AVDPIP database, which contains feature data as generated by the TEDS, IMDx, NBATS and Fusion modules. The usage information can be made available to PIP using the MPUGS tool. The user can also specify events that have a potential impact in the life of a component that is to be analyzed (for example, maintenace operations, inspection results, operating environment, etc.), and the software is capable of integrating this information to update the component life calculations.

Similarly to the TEDS tool, PIP is intended to potentially perform near real-time processing with feature data to monitor assets during operation. PIP is capable of using the two different diagnostic algorithms described earlier to perform feature value analyses. The results of the two algorithms can be fused using different techniques according to the user specifications: an anomaly detector (Zhang et al., 2008) or a more traditional threshold-based detection process. The anomaly detection technique compares a predefined baseline distribution of feature values (fixed) with a distribution of “current” feature values (changing as more data is processed). The baseline distribution can be defined by the user directly in PIP or determined using NBATS. False alarm rate, probability of detection, and other diagnostic settings can be configured for a given analysis. Two thresholds are utilized in diagnostic operations: a warning threshold (or “yellow” condition) and an alarm threshold (or “red” condition). The traditional thresholding diagnostic technique uses predefined warning and alarm thresholds (for the “yellow” and “red” conditions, respectively), which can be specified by the user or potentially generated by an NBATS analysis.

PIP allows for the use of both of usage-based and health-based prognostics. The usage-based approach to prognostics uses three types of loading profiles to define how an asset is used:

- Past usage: These are the missions or load profiles that a component is assumed or known to have undergone from the start of a prognostic analysis up to the "current" processing time. The "current" time is continually updated as the analysis progresses and more data is processed. This time horizon represents loading that the component has already experienced.
- Future loading (immediate): this mission is used by the prognostic algorithms only once, starting at the "current" time. This time horizon represents loading that the asset will undergo imminently.
- Future loading (extended): this mission is run for as many times as needed until the prognostic simulation reaches a critical or failure condition, starting right after the immediate future mission. Multiple uses of the extended future loading are needed to allow the prognostic algorithms to determine the time at which a component fails or needs maintenance.

The health-based prognostic algorithm is run by PIP only after a fault is detected by the fusion of the Anomaly Detection and Traditional Thresholding diagnostic algorithms. The health-based prognostic extends the trend of a feature progression being analyzed by using a degradation progression model and predicts times of needed maintenance actions based on when the predicted feature trend reaches an "end-of-life" condition as specified by the user.

When PIP performs an analysis, it returns diagnostic and prognostic results. For diagnostics, PIP displays a series of results as illustrated in Figure 8. The results include feature value graphs and normalized PDF plots of system conditions, as well as a series of system health indicators in the form of traffic lights displaying either of a green, yellow or red condition.

The health indicator traffic lights provide the final result for the diagnostic and prognostic processes, after the fusion of available evidence and following the settings specified by the user or analyst. As feature data is processed, diagnostic and prognostic results are updated. Figure 9 shows an example of fault detection by the anomaly detection algorithm.

The prognostics results screen of PIP displays usage-based and health-based prognosis assessments separately, as illustrated in Figure 10. The 'PDF of Time to Maintenance' histogram shows the normalized probability density of the expected time (in hours) at which point maintenance action will be required. The histogram provides a more detailed view of the information conveyed by the intersection of the predicted damage progression curves (lower, expected, and upper bounds) with the maintenance threshold (horizontal red line).

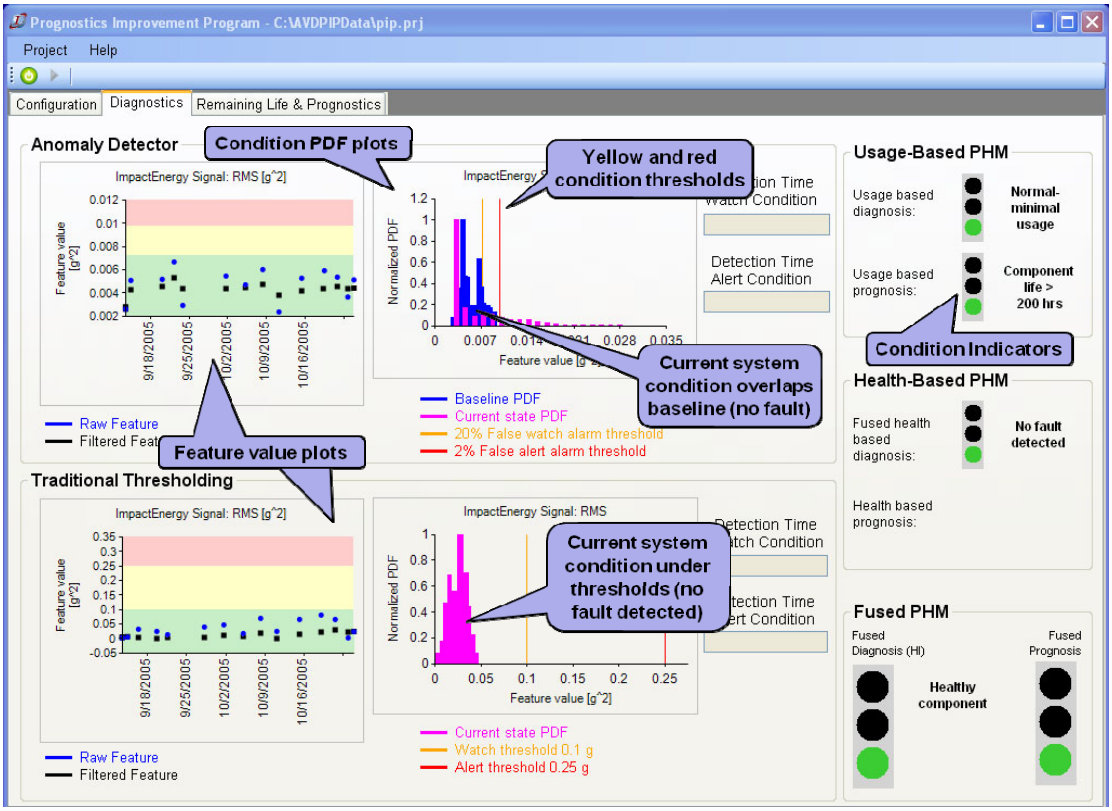


Figure 8. PIP diagnostics results screen (before fault detection)



Figure 9. PIP diagnostics results screen (after fault detection)

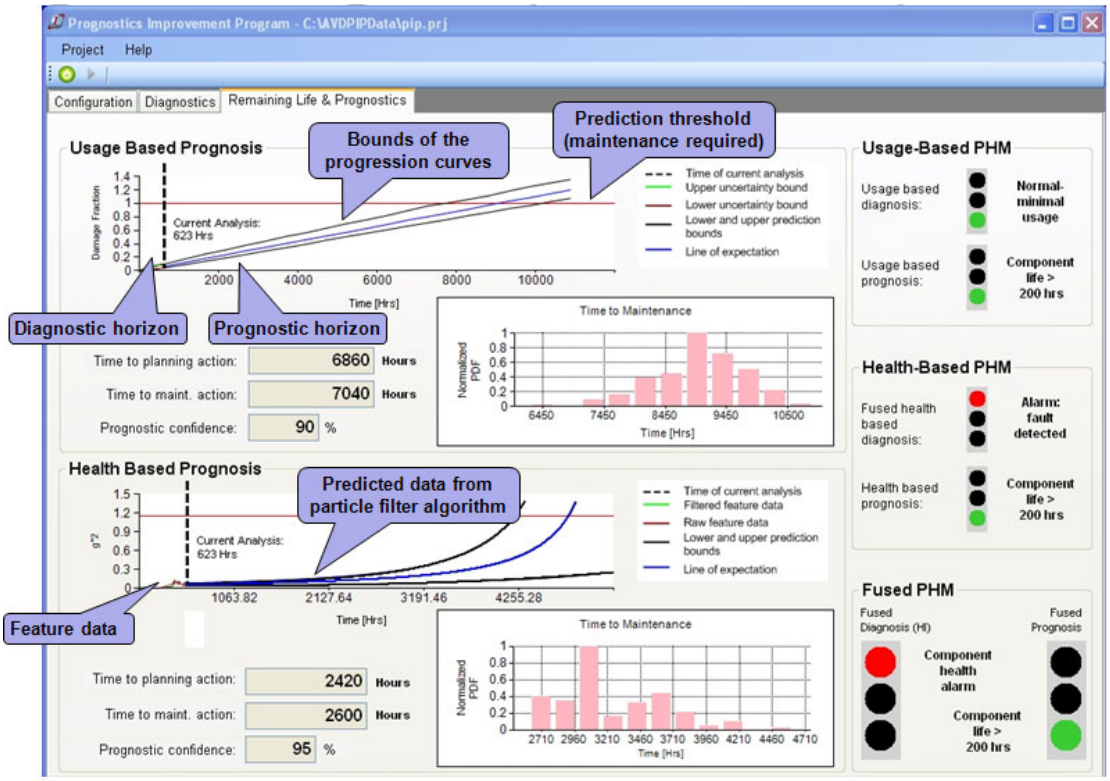


Figure 10. PIP prognostics results screen

Prognostic Validation

Because appropriate validation of fault detection techniques is an ever important aspect of engine and drive train monitoring technologies, (Byington et al., 2006b) fleet data and seeded fault tests are being used to demonstrate and validate the ability of the diagnostic enhancement and prognostic software modules to detect component faults early and accurately, as well as predict the rate of wear or damage progression. Early fault detection and failure/wear prediction methods are used to determine safe times for performing maintenance actions (planning and servicing) on an aircraft component. The example platform providing validation data is a bearing used by the oil cooling subsystem of the H-60 series of helicopters deployed by the U.S. Army (Smith et al., 2009). The software modules allow for performing a thorough analysis of the durability and behavior of failing oil cooler bearings, as well as adjusting and comparing the performance of a set of diagnostic enhancement and prognostic methods for realizing a reliable monitoring system. Case studies are being performed to compare diagnostic and prognostic results to ground truth data sets and known cases of fault in the helicopter fleet. Prognostic validation of these cases is planned separately for the usage based and health based prognostic algorithms. The following approach has been proposed for each.

Health based prognostics validation

To evaluate the accuracy of a health based prognostic assessment, it is necessary to utilize components that have experienced a fault and remain in operation so that the fault progresses. This is necessary because the health based prognosis

algorithm initiates operation once a fault is detected in its early stages, and proceeds to make predictions for a time when the fault will reach a more severe condition of interest with progressed damage. Hence, the following procedure can be used for validation of the predictions of times to reach the progressed damage condition:

1) The health based prognostic algorithm activates upon initial detection of a fault by the diagnostic system, which corresponds to the crossing of a yellow condition threshold.
2) The health based prognostic algorithm is configured to continually predict or estimate the operational time it will take for the system to reach the red condition threshold (time of required maintenance action).
3) As the system continues to operate, the time prognosis is automatically updated by the prognosis algorithm, so that the prognosis performance curve can be generated, as illustrated in Figure 11.
4) Once the system reaches the red condition, the damaged component is retired. The performance of the prognostic algorithm can be evaluated on the resulting prognosis performance curves: just-in-time line and expectation line.

Clearly, this approach is applicable to both of seeded fault testing and known cases of faults with fleet data.

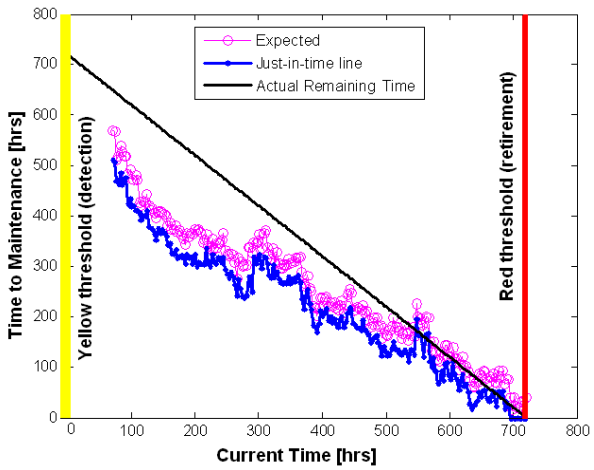


Figure 11. Proposed path to health based prognostics validation

Usage based prognostics validation

Because usage based prognosis focuses on assessing durability of components with regards to their expected design life (even if with the use of life modification factors for certain potentially harsh operating conditions), to evaluate the accuracy of a usage based prognostic assessment it is necessary to utilize components that have not experienced an unexpected or uncharacterized fault. Furthermore, the components must typically remain operational: (1) under known conditions so that the component life models can make a valid assessment of component degradation for the corresponding operating conditions, and (2) until the end of the design life of the component is reached. Generally speaking, this last requirement may be needed because the end of the design life of a component might be the only “verifiable” condition or event in the internal degradation process (operational wear) of a

component, unless partial degradation of the component can be accurately quantified at any point during its operational lifetime. For example, when estimating the operational life of a bearing, it is not generally possible to reliably measure the amount of damage it has accumulated at some point in time as a percentage of its total design life. If one intends to compare the durability of a bearing against a calculated life (design life estimate), the bearing must necessarily operate until it completely consumes its life, or else a measurement of the bearing durability will not be available for comparison with the calculated durability.

Unfortunately, running a system component until its end of life is reached is generally not feasible in operational systems and equipment. For example, we cannot keep operating a bearing in a helicopter until the bearing fails for purposes of validating a usage-based bearing-life prognosis algorithm. Nevertheless, this run-to-failure scenario may be attained in certain test platforms. For example, one may use an appropriate test rig that can support run-to-failure tests of bearings.

Hence, under these conditions and for components in real world service (as opposed to testing), it is generally more difficult to validate the performance of a usage based prognostic assessment of a component's end-of-life than the performance of health-based prognostics. Nonetheless, using a combination of real world service time followed by run-to-failure testing, it may be possible to compare durability calculations to actual component lengths-of-service. The following procedure is thus proposed for validation of usage based component life predictions:

1) A component (e.g., a bearing) is retired from service in operational (non-test) equipment before its design life is consumed (e.g., service time limit is reached); there may be no obvious damage or wear present in the component.
2) No unexpected or out-of-design damage or wear is detectable in the retired component. Furthermore, the approximate loading / stressing / mission / usage profile experienced by the component throughout its service life is known.
3) The component is installed in suitable test equipment and run until failure (end of design life) with a predetermined, known loading profile (controlled usage).
4) Using necessary adjustments, the in-service and in-test times are added to determine the actual total durability/life of the component.
5) Using necessary adjustments, the "predicted" or modeled durability of the component is calculated as follows:
   a. Use the in-service loads and the in-service time with the component life model to calculate the percent of life consumed during the in-service period.
   b. Use the in-test loads to calculate the time needed to consume the life of the component that was not consumed during the in-service time.
   c. The addition of this calculated time with the in-service time provides the "predicted" component durability.
6) The performance of the usage based prognostic algorithm can be evaluated by comparing the "predicted" component durability (step 5.c) to the actual component life (step 4).

Conclusion

This paper describes an integrated software architecture developed to support helicopter drive train component diagnostics and prognostics but generally applicable to a wide variety of complex systems. The predecessor to this paper is a paper by Smith et al., 2009, where individual processing techniques are described. In this follow-up paper the integration of those individual techniques into a modular, highly capable and functional software architecture is described. Implementation of the software with such side ranging functional capabilities is an important undertaking not often seen in the PHM/CBM arena. The effort demonstrates that it is possible to build comprehensive tools for enhanced, state-of-the-art diagnostics and prognostics, even for such complex systems as a helicopter drive train. Validation of the different comprising analysis tools is still ongoing, as more cases and data are needed from the field. Nevertheless, the paper also proposes how validation may continue as more cases and data become available. The AVDPIP program has now finalized, but the team is pursuing gradual implementation of its developments for the U.S. Army, to support the Army's efforts to improve diagnostic performance and implement prognostic capabilities for aircraft component CBM.

Acknowledgments

This work has been partially supported with a cooperative agreement by the U.S. Army Research Laboratory under contract number W911NF-07-2-0075. The authors would like to thank ARL technical points of contact, including Dr. Romeo del Rosario, Mr. Kwok Tom, and Dr. Canh Ly, as well as Army AED personnel and contractor representatives from organizations supporting the Army Utility (Black Hawk) Program Manager, including Mr. Daniel Wade, Mr. Daniel Suggs, Mr. Don Estes, Mr. Carlos Rivera, Dr. Jon Keller, Mr. Jeremy Branning and Dr. Robert Vaughan. This work has also benefitted greatly from consultations with other Army Research Laboratory and NASA Glenn researchers, including: Dr. Timothy Krantz, Mr. Brian Dykas, Mr. Harry Decker, Mr. David Lewicki, Dr. Hiralal Khatri, and Mr. Ken Ranney.

References

Baker, C., Marble, S., Morton, B.P., and Smith, B.J. "Failure Modes and Prognostic Techniques for H-60 Tail Rotor Drive System Bearings". IEEEAC paper #1122, IEEE, 2007.

Branhof, R.W., Grabill, P., Grant, L., and Keller, J.A. "Application of Automated Rotor Smoothing Using Continuous Vibration Measurements". American Helicopter Society 61st annual forum, Grapevine, Texas, June 1–3, 2005.

Byington, Orsagh, Kallappa, Sheldon, DeChristopher, Amin, and Hines, "Recent Case Studies in Bearing Fault Detection and Prognosis," IEEE Aerospace Conference, Big Sky, MT, March 2006a.

Byington, Orsagh, Sheldon, Kallappa, DeChristopher, Amin, "Verification and Validation of Incipient Fault Detection Techniques for Engines and Drivetrains," 60th Meeting of the Society for MFPT, April 2006b.

Byington, C. S., Smith, M. J., and Kalgren, P. W., "Improving Diagnostic Isolation in Drivetrain Components: A Swashplate Bearing Example," 5th DSTO International Conference on Health and Usage Monitoring (HUMS), Melbourne, Australia, March 20-22, 2007.

Byington, C.S., Watson, M.J., Amin, S.R., and Begin, M., "False Alarm Mitigation of Vibration Diagnostic Systems," IEEE Aerospace 2008, Big Sky, MT, March 1-8, 2008a.

Byington, C.S., Watson, M.J., Lee, H., and Hollins, M., "Sensor-Level Fusion to Enhance Health and Usage Monitoring Systems," AHS 64th Annual Forum and Technology Display, Montreal, Canada, April 29 – May 1, 2008b.

Dora, R., Wright, J., Hess, R., and Boydstun, B. "Utility of the IMD HUMS in an Operational Setting on the UH-60L Blackhawk". American Helicopter Society 60th annual forum, Baltimore, Maryland, May 7–10, 2004.

Keller, J.A., Branhof, R., Dunaway, D., and Grabill, P. "Examples of Condition Based Maintenance with the Vibration Management Enhancement Program". American Helicopter Society 61st Annual Forum, Grapevine, Texas, June 1 – 3, 2005.

Sheldon, J. S., Klenke, C., and Byington, C. S., "Detection of Incipient Bearing Faults in a Helicopter Gas Turbine Engine," AHS 63rd Annual Forum and Technology Display, Virginia Beach, VA, May 1-3, 2007.

Smith, M.J., Zhang, B., Patrick, R., Byington, C.S., Vachtsevanos, G.J., Del Rosario, R., Wade, D.R., and Suggs, D.T.; "Combination of Fusion and Preprocessing Techniques to Enhance Air Vehicle HUMS". Sixth DSTO International Conference on Health & Usage Monitoring, 2009.

Watson, M. J., Sheldon, J. S., Amin, S. R., Lee, H., Byington, C. S., and Begin, M., "A Comprehensive High Frequency Vibration Monitoring System for Incipient Fault Detection and Isolation of Gears, Bearings and Shafts/Couplings in Turbine Engines and Accessories," Proceedings of ASME Turbo Expo 2007: Power for Land, Sea and Air, Montreal, Canada, May 14-17, 2007.

Zakrajsek, J.J., Dempsey, P.J., et al. "Rotorcraft Health Management Issues and Challenges". NASA report TM—2006-214022. February, 2006.

Zhang, B., Sconyers, C., Byington, C.S., Patrick, R., Orchard, M.E., and Vachtsevanos, G.J. "Anomaly Detection: A Robust Approach to Detection of Unanticipated Faults". International Conference on Prognostics and Health Management, Denver, Colorado, October 6-9, 2008.