2010 NDIA GROUND VEHICLE SYSTEMS ENGINEERING AND TECHNOLOGY SYMPOSIUM
VEHICLE ELECTRONICS & ARCHITECTURE (VEA) MINI-SYMPOSIUM
AUGUST 17-19 DEARBORN, MICHIGAN

PROGNOSTIC ALGORITHMS IN CONDITION BASED MAINTENANCE

Gregory Harrison, PhD
Sreerupa Das, PhD
Michael Bodkin
Stefan Herzog
Richard Hall
Lockheed Martin Global Training and Logistics
Orlando, FL

ABSTRACT

Machinery health management is becoming increasingly important and the diagnosis of failures based on machinery condition has been analyzed in-depth in the last few decades, and is relatively well understood. However, prognostic evaluation of faults in a machine is a harder task that involves predicting impending faults in the system and determining remaining useful life of the machinery. A survey of algorithms, and a detailed description of a hybrid CBM prognostic techniques being investigated for use in ground vehicle systems will be presented. The system incorporates a number of techniques to process and analyze the current condition of a ground vehicle, and to generate a prognosis for each subsystem in the vehicle. The discussion will describe a means of testing, verifying and iteratively improving prognostic capabilities throughout the lifecycle of the platform.

INTRODUCTION

Vehicles generally have some form of diagnostic indicators and the ability to perform diagnostic tests and sense the condition of the vehicle. Producing a viable prediction from these condition indicators (CI) is the next step for vehicle and fleet integration. Although Prognostics has been studied for a very long time in varied disciplines, its real-time use in ground vehicles is just starting. It is generally accepted that an electronics-based means of prognostics that relies on computers will be the main form of the future vehicle prognostics system, but some other types of prognostics technologies will be discussed, and perhaps they could also play a part in the overall prognostics mechanisms. The Condition Based Maintenance+ (CBM+) research serves as guidance into the overall operation and layout of the information system infrastructure. Ideally, the prognostics sections to be employed will also form an open-standards type of arrangement that will allow modules from different researchers and producers to be easily integrated in a system that produces a prognostic metric. Defining the prognostic metric is the interesting thing, where it can be used and updated by many different organizations. The Remaining Useful Life (RUL) metric or prognostic mechanism may be the result, with suitable modifications and additional surrounding technologies.

PROGNOSTICS

Correctly predicting the future is a profitable ability. In the case of predicting the RUL of a vehicle component, it can lead to cost savings by waiting until maintenance is needed instead of performing maintenance actions on a set interval. Correctly predicting RUL also can help prevent situations where the component stops working at the wrong time.

History

This subsection deals with the formal mathematical use of prognostics, since there have been so many other uses through the millennia. Predicting the future seems to have always been a skill in demand. The science of statistics provides a mathematical basis for predicting the future. Statistics are all that is really known about anything, so basing predictions on statistics must be a good start. This is the basis behind the RUL form of prognostication: if you know how long a certain component statistically lasts in a given environment, then you can figure how much time it has left to operate, given that you know how much time it has already operated.

One form of non-data-based prognostication is the use of a canary [1]. A canary, also sometimes known as a fuse, is generally a piece of hardware that is meant to fail before the main system fails, thus providing an indication of what may be coming in the future. It got its name from coal mining warning systems where they carried a canary bird into the mine with them, and if the canary died, then the miners knew that there were dangerous gases present, and they needed to leave immediately. Canary has then become the word to describe this sort of early warning system. Another form of canary could be considered to be the indicator that tire manufacturers sometimes put into their tires that shows when only a little tread life remains by providing a comparison plateau that is reached by the remaining tread as it wears. If the tread height reaches the plateau, then it is nearly time to replace the tire. Another form of canary is a deliberately-weakened electronic circuit on an integrated circuit chip. If the circuit becomes inoperable, then it is an indication that the entire chip may soon become inoperable. Hopefully the canary gives enough time to attain a safe state.

Canaries are effective and provide a definite physical indication that there is a problem and that time is running out for the health of the system. Unfortunately they only give off one warning, and until the canary dies, there is ostensibly no indication that the system remaining useful life is just about running out. A Prognostics Health Management (PHM) system can use canaries as long as they can be somehow connected into the PHM system. A canary will allow all the mathematical predictions to be immediately reset to the ground truth of the system health status, and can carry on to predict RUL from there.

Similar to a canary is a failure precursor, where an impending failure is predicted by conditions detected in a data stream or information about the system. It may take the form of a metallic shell around a component glowing red hot, or a signal being sensed from a vibration transducer has become twice as powerful.

More subtle shifts in operating condition of a machine can be sensed with technologies like Fuzzy Adaptive Resonance Technology (Fuzzy ART) Neural Networks [2]. For the case of a vibration signal, the frequency response is split into bins of an arbitrary resolution, and applied to the neural network, where a representative spectrum is learned for each of the various operating modes that are expected to be encountered. After these have been learned, and then may have had been learned on the actual machine or vehicle of interest thus making them even more sensitive, the system is switched into monitor mode, and if it senses a deviation from the automatically-learned envelope of operating spectral characteristics, it will provide an indication therewith. This technology has been expanded to work in a direct prognostics application, and for many different types of signals, not just vibration. Either way, having a precursor is a great thing, and similar to a canary it allows the operator to take immediate action to safeguard the machinery.

Another type of PHM technology that has attained a good level of use it that of remaining-useful-life estimation. The simplest form of remaining useful life calculation is that such as is found on many automobiles for the RUL of the oil. The system assumes that the oil needs to be changed every 3000 miles. The car computer also has access to the distance that had been driven since the last oil change. It can reason that if 2000 miles has already been driven since the last oil change, that therefore the oil needs to be changed in 1000 miles. If I have not been really stressing the car, and it has already been broken in, then I might wait longer, or if I have been driving cross country or through a dusty environment, I might reason to change it sooner than the blind-reasoned car software decision has told me to. Having such a drivertronic reasoner is a good improvement over the blind reasoner, and probably saves me money, or at least eases up on having to get the car in for an oil change the minute the indicator says that there are zero miles remaining before an oil change is needed. The example systems that are described later in the paper illustrate things that can be done to improve the reasoning capabilities of the RUL-type reasoner by using some of the concepts of the failure precursor-type and the canary–type prognostication system.

Statistics

The medical and insurance industries rely heavily on the use of statistics for prognostication. Extensive testing and studies are conducted to find the improvement rates from different diseases, given particular treatments. Insurance agents work with actuarial tables to try to predict how long someone may live, in order to determine how much to charge them for insurance. Both of these studies are completely statistical. They are not focused on a particular individual. For medical applications, a doctor could refer to these studies and probabilities, and suggest a course of treatment, or maintenance for the individual, knowing the individual's particular circumstances. That is what the new systems for embedding in vehicles will provide: a continuous doctor's perception of the state of the entire vehicle, at least for the parts that are sensorized. Of course the doctronics needs to have familiarity with the various fields and components, and these components are provided by human experts, and integrated in, perhaps licensed in, for online use.

To know what a baseline system is expected to perform like requires statistics. Many vehicles need to be studied, with test equipment that measures the performance characteristics, and determines the expected actuarial tables, or predicted lifespan, of each of the pieces of equipment being monitored.

A common approach to learning the baseline of a system is to apply Bayesian Learning techniques to learn the probability of failure of the life of the system. If the engineering information indicates a causality situation, such as when the wheel hub degrades, it increases the probability of the axle failing. A Bayesian Network can be created to fuse the engineering data with the historical data to increase the accuracy of the prognostics even in the absence of significant number of instances of the event in the historical data. When the engineering data is missing, a hidden Bayesian Network can be trained to learn the probability of failing. The limitation is it will not learn the above event unless there are a statistically significant number of occurrences of it in the historical data.

Algorithms

In order to operate a vehicle reliably, it is necessary to be aware of the time left before its functionality is impaired. Determining Remaining Useful Life (RUL) of a machine or a vehicle is one important aspect of Prognosis. This is especially important in mission critical tasks, such as in a vehicle attempting to undertake a mission in a battlefield. An unanticipated vehicle downtime could lead to catastrophic failure, lost lives, or a lost mission. RUL of a vehicle helps quantify how much time is left until the vehicle becomes dysfunctional due to failure of one or more of its components. Awareness of RUL ahead of time not only helps determine a vehicle's reliability in undertaking an operation, it also helps prepare for maintenance of the vehicle, chart future operations based on the remaining life of available vehicles, and help plan logistics.

Given a new component, it is possible to define its estimated life in an ideal condition from the manufacturer's specification, e.g., a battery may run out of life out after 100 days of operation or a gear may break a tooth after traveling 100 miles. The estimated life can also be determined statistically through measuring a sampling of units to determine how long they have historically operated, on average. The units for RUL may be defined in terms of time, distance or other unit as is appropriate for the component whose RUL is being determined. This initial form of RUL prognosis may be represented in a graph as shown in figure 1, where lifetime is defined in terms of time, instead of usage or distance.

In figure 1, Life Credits (LC) are used as a gauge of useful remaining life. Life credits are used up at one single rate that has been determined in some way, and in this case is represented as a slope, m_x in the LC curve having the value:

$$m_x = -\frac{ILC}{t_{EOUL}}$$

where ILC is the initial amount of life credits, say 1000, and t_{EOUL} is the time it has been found that is the expected end of life for that component being monitored. The slope is negative since it is meant to decrease the remaining life credits over time or usage.

The graph generally always starts with the maximum LC, shown as Initial LC (ILC) after the machinery is installed or serviced. As the machinery is operated, its LCs keep decreasing, terminating at the End of Useful Life (EOUL). Hence in an RUL plot, the slope of the line is always monotonically decreasing with a constant slope.

The *LC(t)* line in figure 1 is shown as extending below the zero LC axis. This may be used to show increased urgency in the need to replace a certain component, if that component has not been replaced or serviced before the predicted EOUL as it should have been.

This simple case of one straight line at a given slope is used in situations such as the remaining miles able to be driven in a vehicle based upon the miles-per-gallon usage of the car and the amount of fuel remaining. In that case, the units of the abscissa of the graph of figure 1 would change from a unit of time to now represent miles.

Prediction of the RUL using the system of figure 1 is such that: given the time now, t_{now}, and the amount of life credits remaining, LC_{now}, the remaining useful life time period, T_{RUL}, can be determined as:

$$T_{RUL} = \frac{LC_{now}}{m_x}$$

which is the predicted time remaining until the end of useful life.

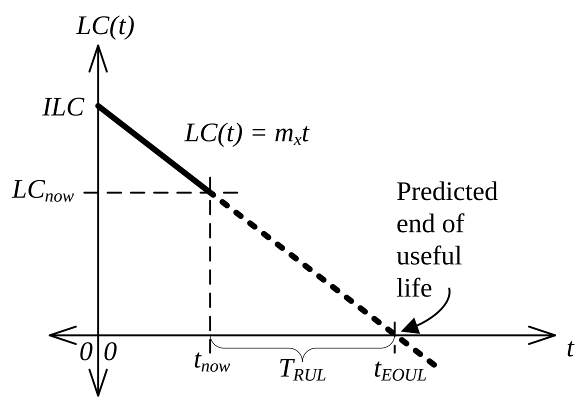

Figure 1: RUL plot in the most simplistic case. LC is Life Credits, ILC is Initial Life Credits, T_{RUL} is the Remaining useful life in terms of time, and t_{EOUL} is end of useful life.

Modifications to this RUL model can likely provide more accurate predictions of the future. These modifications take at least three forms:

1. Real-time LC offsets and slope changes
2. Predicted LC offsets and slope changes

Real-time LC Offsets and Slope Changes

Something may happen that causes the remaining LC to decrease instantaneously. For instance, a canary sensor may trigger, indicating that the RUL has dropped to a low amount or a shock detector could detect that a powerful shock was received, such as by driving over a deep pothole. In the canary sensor case, the life credits can be modeled to drop **to** a certain amount, while in the shock detector case, the LC drops **by** a certain amount. In the case of a shock absorber, severe shock impact is a reoccurring discrete event. A circuit breaker may be considered as canary event where the circuit breaker may be reset to indicate future trips. In the circuit breaker case, the root cause of the circuit breaker trip should be determined and new life credits added to the system through maintenance.

The canary algorithm can be visualized as a binary bit that is off until the canary dies. The death of the canary can be modeled as a Dirac delta function, $\delta(t)$, which is a single event in time, as seen in figure 2a. Mathematically, this function has infinite height and infinitely narrow width, integrating to a value of one. The nomenclature $\delta(t-t_0)$ indicates that the Dirac delta function is offset in time so that at time t_0, the quantity $t-t_0$ has a value of zero and makes the function fire at that instant.

The USMC Embedded Platform Logistics System (EPLS) ground vehicle support apparatus that is used to monitor and collect vehicular system/subsystem mission critical data from various sensors on a military ground vehicle, represents most of the data using events. These events can be viewed as Dirac delta functions. They indicate that a certain event occurred at a certain time. Perhaps a pump started running. Then an event would be sent along the internal vehicle data bus, that pump_x started running at time t_0, but subsequent to that, no more information about pump_x is sent until it turns off.

The canary event can be used to trigger an instantaneous drop of the remaining life credit to a low value, indicative of the prediction that the component is very near the end of its life. This is shown in figure 2d, where the life credits drops immediately all the way down to the Canary Predicted Life Level (CPLL), and continues to decline through usage from that new level. The slope of decline after the canary dies can also be a different slope, as indicated by the new slope callout of $m_{x,2}$ instead of $m_{x,1}$ in the figure. Thus, for a canary event,

$$LC(t) \leftarrow CPLL$$

and

$$m_x \leftarrow m_{x,2}$$

depending on which of the canary-induced changes to the LC update algorithm are being used.

To process the information about the canary event as just a change in slope instead, the event needs to be stretched into a step function as seen in figure 2b. This step function is also called the Heaviside function [3], and can be created mathematically by integrating the Dirac delta function, yielding a function, *u(t)*, that has the following characteristics:

$$u(t) = \begin{cases} 0, & t < 0 \\ 1, & t \geq 0 \end{cases}$$

The offset, $t-t_0$, makes the function shift in time to fire at t_0. Converting the vehicle-monitoring canary event to a step function allows its use in the calculation and continuation of the new slope of the RUL graph.

In figure 2c), it can be seen that the slope changes after the canary fires. The initial slope can be represented as $m_{x,1}$, and the post-canary-firing slope as $m_{x,2}$, where it can be modeled as:

$$m_{x,2} = m_{x,1} + \Delta m_{canary}$$

Thus the RUL curve attains a continuous added slope of a certain amount, and the slope is required to be negative, or $\Delta \mathrm{m}_{\mathrm{canary}} < 0$.

Another form of instantaneous change in life credits can be produced by the response to a sensed, potentially damaging event, such as a shock from driving over a pothole. This produces a life credit offset, which may be proportional to the sensor reading for the shock, for instance. In a life credit offset, the RUL would be dropped a certain amount from where it currently was, as seen in figure 3. The slope of life credit decline may be modeled to change after each impact also.

In reality, during the lifetime of the component, it hardly ever operates in one operating condition or regime. A vehicle may have to operate under very low or very high temperatures, drive very slowly or very fast (on a variety of terrains or situations), hence impacting the RUL of the component. Using historical data, it is possible to estimate the changing remaining life when operated in different operating regimes. The changing RULs during the component's lifetime may be represented by changing slopes on an RUL plot. For example, figure 4 shows of a graph of altering RUL over six operating regimes.

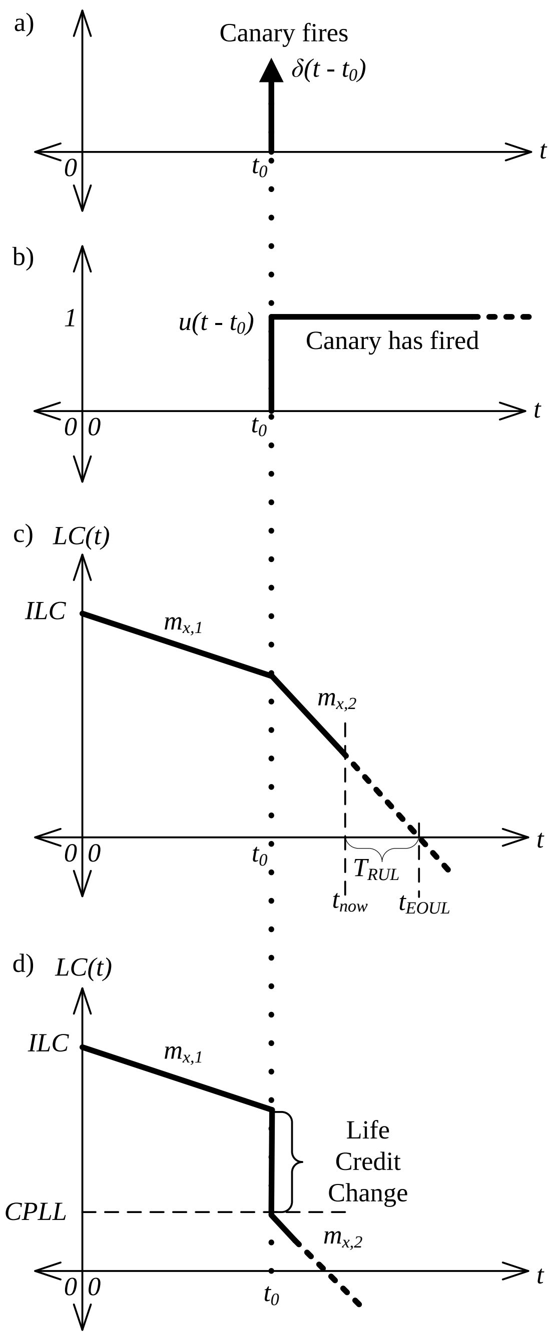

Figure 2: Information from a canary is translated into a modification of the prediction of remaining useful life. CPLL is Canary-Predicted Life Level. Note: a canary generally only activates once.

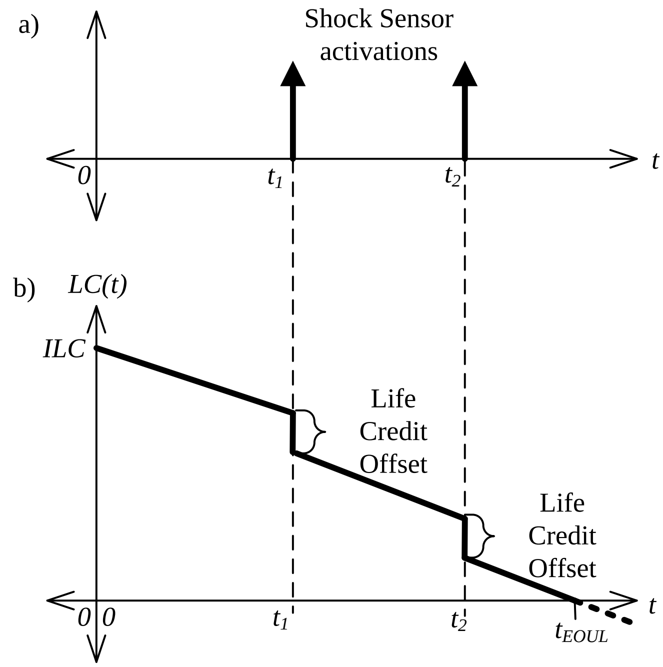

Figure 3: A Life Credit Offset is different from the canary-based modification in that it just subtracts a bit of life, but doesn't reduce the remaining life to a set amount.

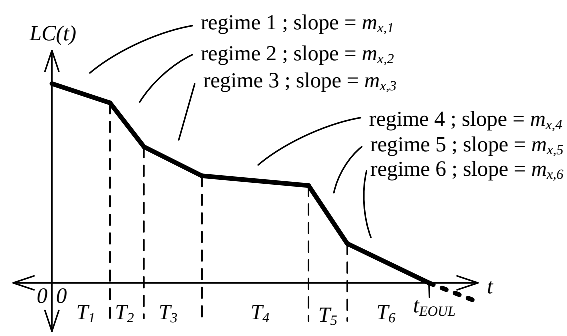

Figure 4: This RUL prognostics plot with six operating regimes shows the complete life cycle of the component.

From the historical data the following can be obtained:

a) The regimes in which the vehicle (or component) has operated (e.g., regime1, regime2, regime3,….)
b) time period it has been in each regime (e.g., T_1, T_2, T_3, …, T_6) and
c) the total life time of the vehicle (or component) before it failed.

Predicted LC Offsets and Slope Changes

The system operates in the t_{now} time. The future consists of a prediction, but knowing the state of the system as represented by the remaining Life Credits, and the planned operating regimes, and expected system shocks, the RUL may be predicted for a given mission plan. All the components have a set of slopes that are appropriate for a set of given regimes. If a mission is known in advance as to the expected external environment regimes that will be encountered, then these regimes and their durations can be presented to the RUL model, and run over time to determine when and if components are expected to fail during the mission. Pre-emptive maintenance may be performed ahead of the mission, if possible, or mission plans can be changed based upon the foreseen capabilities of the fleet of vehicles available for the commander’s use.

Rolling up RUL to entire system

The RUL of an entire system, such as a vehicle, is based on the RUL of its components. Failure of one component may be sufficient to fail the system. Since the interactions and interdependencies of the components may not be very obvious, it is helpful to use statistical methods, such as Bayesian networks, to provide reasonable estimates of RUL based on the data available [4][5].

As the individual RULs are rolled up to a system level, an accurate picture of the readiness of a vehicle or a fleet of vehicles can be obtained. This picture may then be used to plan for an upcoming mission.

OFF-PLATFORM ANALYSIS

The prognostics analysis does not have to be performed only on the platform of interest, but may be performed at a computation center that is located off-platform. A prognostics analysis center may be located anywhere that may be communicated to by the vehicle, directly or indirectly. This center produces analysis of the vehicle health on an on-going basis. Usage rates for components will be needed to be sent back to the off-platform site to allow mission planning based on remaining life values. Another type of off-platform analysis center creates the algorithms that detect the failure precursors and create new memories for the neural networks and new LC usage slopes.

The Off-System CBM+ components, as seen in figure 5, also provide a data warehouse, a data mining and analysis capability, and a capability to update the condition monitoring algorithms and advanced prognostic health indicators that will run on the On-System hardware. The term CBM+ indicates a more enterprise-level CBM system, with on-platform, at-platform, and off-platform components. The CBM+ off-platform communications could also be linked to a National-level Strategic Data warehouse [6].

Specific application tools such as Enterprise Logistics IT, prognostic algorithms, data mining and the like would comprise the Off-Platform or Off-System Architecture layer. Data mining and Business Intelligence Tools would identify trends, and improve and refine Maintenance, Diagnostic, and Prognostic capabilities.

Processing at the Off-system level is intended to provide a processing and data storage capability. Data is moved from the On-system components to the Off system environment using established data communication networks Data compression algorithms would be used to decrease the network loading required to move large data sets.

Normal operating equipment will not need to have data sent back to the Off-System location, but when anomalies are detected, or a piece of equipment is determined to be failing but the condition was not indicated, then the set of data recorded for that piece of equipment in its degraded or failing state would be sent to the Off-System site for analysis to develop an algorithm that would detect the failure in the future. This updated algorithm, perhaps consisting of neural network parameters, slopes, or other controls for the On-System hardware, would then be distributed to all On-System monitors for the affected equipment model.

The Off-System architecture requires the use of carefully designed database schemas in order to coordinate large datasets of recorded machinery operation, enabling the precise application of the analysis and condition indication algorithms.

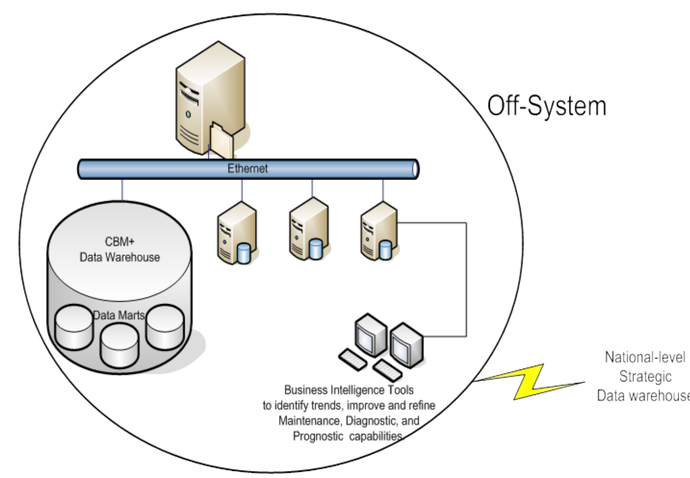

Figure 5: Off-System CBM+ Architecture.

POTENTIAL USES OF ANALYSIS

Now that technology is available for creating intelligent platforms, collecting system data, analyzing and maintaining the data one simple question remains “so what?” How does access to this information affect material changes in the operation of a specific platform? There are three areas where significant impacts can occur: at the operator, the fleet manager, and potentially the original equipment manufacturer.

A question an operator may ask himself may be: "If I keep driving my vehicle will I be able to complete my mission?" The system discussed above will allow for real time alerts to be presented to the vehicle driver or commander in real time. As canaries "die" and LC's (the CPLL) are impacted, messages can be sent to the operator and presented as an alert. Unlike current oil RUL indication this canary type prognostication tool provides a better understanding of the current vehicle/component health based on the inputs discussed previously.

Operational planners, or fleet managers also require the information created through the use of this system. Having a CBM+ enabled architecture allows for the collection of data from multiple platforms: for example, a Battalion or Brigade's worth of tactical wheeled vehicles. The S3, or operations officer, can identify vehicles which are more fit or in a better condition to meet upcoming mission requirements, based on the inputs obtained from individual platforms within his span of control. Similarly, logisticians can begin to utilize this data for understanding when or if to order certain repair parts. By doing this logisticians can provide the right repair parts to the field, without having to pre-position parts based on demand data from previous deployments.

Finally, there is an engineering/OEM aspect to this intelligent system as well. Over time, failure modes can be isolated, as well as the associated inputs leading up to the component failure. This will allow for the study of the reasons behind the failures and provide data as to the cause, be it engineering design, unintended use of the system, or operator training, for instance. All these questions affect overall performance and LC of the system.

REFERENCES

[1] M.G. Pecht, "Prognostics and Health Management of Electronics", Wiley, Hoboken, NJ, 2008.

[2] G.A. Harrison and F.J. Taylor, "Gas Turbine Vibration Analysis with Fuzzy ART Neural Network" Proceedings of the 1999 International Joint Conference on Neural Networks, Washington DC. 1999.

[3] Interactive Mathematics, Laplace transforms -1a. The Unit Step Function (Heaviside Function), web page: http://www.intmath.com/Laplace-transformation/1a_Unit-step-functions-definition.php viewed July 18, 2010.

[4] S. Das, G.A. Harrison, M.A.Bodkin, S. Herzog, R.Hall, "CBM+ for Helicopter Health Management" European Rotorcraft Forum, 2010.

[5] B. Saha, Kai Goebel, Scott Poll2 and Jon Christophersen, "A Bayesian Framework for Remaining Useful Life Estimation" Association for the Advancement of Artificial Intelligence Fall Symposium, 2007.

[6] Harrison, G.A, Bodkin, M.A., Das, S., Worden, E.W., Hall, R.M, & Herzog, S.,"CBM+ System for Ship Condition Monitoring," Proceedings of the 14th International Ship Control Systems Symposium, Ottawa, Canada. 2009.