Kevin P. Logan

Software Agent Technology as a Maintenance Workforce Multiplier

ABSTRACT

The ability to predict failure of mission-critical equipment at an early stage can substantially reduce maintenance costs for the Navy. Unplanned down-time can have severe consequences on safety, readiness, and operational and support costs. Condition-Based Maintenance (CBM) strategies have been mandated by the Navy; however, the practical implementation of an effective predictive maintenance program requires diligence in data gathering, data analysis, and data conversion to maintenance decision information. While CBM has been well-proven as a cost-effective maintenance strategy, the Navy infrastructure may not yet have adapted to this new way of doing business. To make matters worse, the introduction of more complex heavily instrumented ship control and weapon systems aboard new vessels presents additional maintenance challenges. Valuable equipment performance information can be "mined" from the thousands of sensor measurements acquired by these modern systems; however, the extraction process can involve substantial human resources, encompassing a broad range of special analytical and computer skills, which are often lacking and/or expensive to develop and maintain throughout the organization. Simply put, there will be more data than people to analyze it.

The exploitation of software agent technology for equipment health monitoring is rapidly becoming the only viable solution for converting voluminous raw machinery plant data into maintenance decision information. Software agents can automatically monitor, troubleshoot, and predict failures in complex machinery processes in support of drastic manning reductions on future ships, with a backdrop of more complex machinery systems and orders of magnitude more data to monitor. This approach has become commonplace in other contexts, considering that the average personal computer has several such software agents monitoring Internet security, virus detection, hard drive performance, etc.

Software agents have been deployed to continuously monitor the health of main propulsion diesel engines on various Navy ships (MSC). The agents "live" on the shipboard data networks and perform comprehensive predictive analytics of engine performance to aid in maintenance management decisions as part of a CBM strategy. This paper provides a detailed description of diagnostic software agent analytics and presents a recent case history of engine faults called out by the agents that were subsequently verified during a shipyard engine overhaul. The value of software agents for effective equipment health monitoring on minimally-manned ships and as rational maintenance decision aids for both onboard and distance support are also discussed.

Introduction

A primary goal for introducing new technology into shipbuilding and operations is minimizing life cycle cost. Among the major cost factors for ship operations are manning and maintenance. Design strategies focused on reducing crew size involve more extensive automation of machinery monitoring functions, primarily through increased sensor instrumentation. Other design strategies, particularly for naval vessels, focus on increasing survivability through decentralized, distributed systems, resulting in more complex machinery plants, with redundant systems. Such designs also generate increased requirements for monitoring and control automation.

Following either of these strategies creates maintenance-related operational challenges. More complex machinery plants generate requirements for more extensive equipment monitoring, as well as more comprehensive knowledge and skills on the part of operations and maintenance crews. Crew training requirements will grow with system complexity. There will be more machinery data to monitor, with fewer people having less time to analyze it. Yet with reduced manning, the importance of keeping a constant vigilance in machinery performance assessment will never be greater, as the attendance to machinery failures will draw a larger percentage of available onboard human resources.

Reliability-Centered Maintenance (RCM) is a process for analyzing and establishing maintenance strategies for complex systems to determine system functions, equipment failure modes and causes, the impact of functional failures, and optimal strategies for managing potential failures, including predictive maintenance. RCM is part of overall process that manages the risk of losses associated with equipment failures through an effective maintenance program. Resources are allocated to equipment maintenance based on the risk impact of failures. In this context, maintenance is one of the many opportunities to improve equipment and system reliability. RCM also aids in identifying premature equipment failures through condition monitoring of machinery health, that is, through Condition-Based Maintenance (CBM).

A CBM maintenance strategy typically involves performing maintenance only when there is objective evidence of need, while ensuring safety, equipment reliability, and reduction of total ownership cost. The fundamental goal is to optimize availability (readiness) while reducing maintenance and manning requirements. Proper application of CBM can reduce operating and support costs by providing a basis for maintenance decisions that focus scarce resources on that maintenance most needed to ensure safety and mission readiness. However, the implementation of an effective CBM program is not without cost. In order to be effective, equipment health monitoring and analytical tasks must be diligently and regularly performed. This can generate significant workload requirements on the part of the both shoreside and shipboard engineering personnel.

Recent advances in intelligent software agent technology provide embedded analytical capabilities, such as failure trend analysis and enhanced prognostic and diagnostic techniques, to automate the bulk of the work necessary to continuously monitor machinery health. Software agents can autonomously perform complex information processing tasks to identify impending failures and accurately predict remaining useful equipment life. Software agents can be deployed to automatically monitor and analyze hundreds of thousands of data points, while being integrated into existing shipboard automation system environments.

Software agents can clone human intelligence in varying degrees, perform human-like reasoning, and interact with maintenance engineers. They can perform tedious, repetitive, time-consuming, and analytically complex tasks more accurately and reliably than people. They can serve as expert assistants in monitoring, troubleshooting, and predicting failures in complex machinery processes in support of manning reductions on future ships. Imparting intelligent processing functions into software agents will allow the Navy to leverage valuable equipment OEM and organizational knowledge across a geographically distributed ship fleet. Agents can be distributed when and where needed to enhance fleet operational efficiency, platform and crew performance, and mission readiness. Agent intelligence can also be upgraded remotely throughout the platform lifecycle. The human-agent team can provide higher levels of platform readiness/reliability at far less cost than that of the equivalent human resource required to perform the same work. Hence, software agents can be viewed as *maintenance workforce multipliers.*

Figure 1 illustrates the division of labor between software agents and engineering crews for typical equipment health monitoring tasks. The six main CBM processes include data acquisition, equipment performance analysis, condition assessment, fault diagnosis and isolation, problem verification, and maintenance/repair action. As shown here, agents can automate the majority of these processes, reducing the crew's role to fault verification and corrective action. Furthermore, these crew activities only become necessary after a problem has been identified by an agent, resulting in very significant savings in manpower required for CBM implementation.

Intelligent Software Agents for Condition-Based Maintenance

Intelligent software agents will play an increasingly important role in monitoring, controlling, and troubleshooting complex machinery processes aboard future ships. A key benefit of software agents is their ability to automatically perform complex information processing tasks without being constantly controlled by people. Software agents can assist the crew in complex decision-making and other knowledge processing tasks related to CBM requirements. The role of prognostic software agents will grow as higher levels of plant automation and complexity raise the cost of continuous machinery monitoring and CBM beyond what the Navy can afford. New agent technologies can be deployed to automatically monitor and analyze hundreds of thousands of data points, while being integrated into existing shipboard automation system environments.

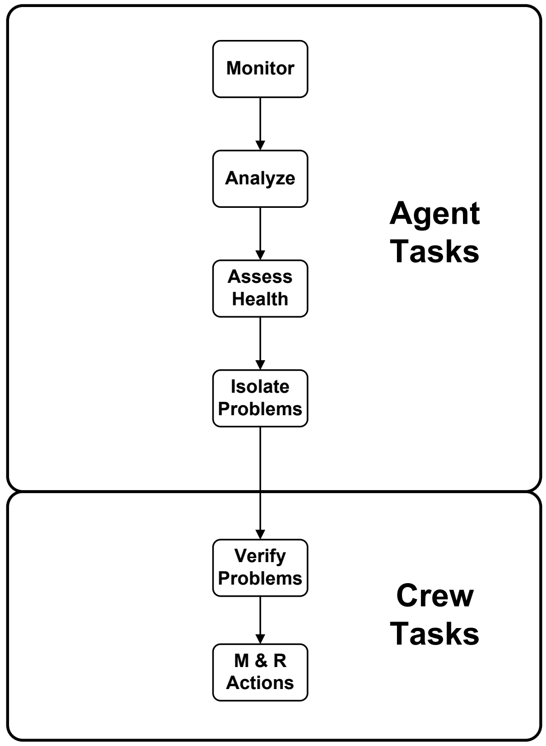

Figure 1 – Division of Labor for CBM Tasks

The main functions of software agents are:

Clone human intelligence – Rare or valuable human intelligence embedded into agents can be used by people who may be less experienced about a particular application (e.g. integrated power system diagnostics).

Perform human-like reasoning - Software agents are empowered with computer representations of human knowledge, allowing them to perform information processing tasks on behalf of their human counterparts. Agents can perform tedious, repetitive, time-consuming, or analytically complex tasks on behalf of people who may not have the time or requisite skills to perform these tasks themselves. The ability to impart intelligent processing functions into artificially constructed agents allows the Navy to leverage valuable knowledge across a geographically distributed work environment, such as a ship fleet. Knowledgebases can be constructed as combined repositories of experiential and technical diagnostic knowledge from OEMs, subject matter experts, experienced operators, etc. Coupled with such knowledgebases, software agents become a valuable resource that can be distributed when and where needed to enhance operations and performance.

Interact with human clients – Agents are designed to interact with people, rather than to replace people. Instead, they are a valuable extension of the human client. Agents tirelessly and autonomously perform their work in the background. This allows the crew to address higher-level problems within their work environment. Through an agent's direct interaction capabilities, they inform the crew of equipment status and report important operational events. The agents can also be called upon demand to perform specific tasks when a crewmember needs them done. Software agents can work autonomously, as the crew performs other work in parallel. The human-agent team can provide higher levels of productivity at practically the same cost as that of just the human resource alone. By applying software agents as *workforce productivity multipliers*, the Navy will be able to leverage its intellectual assets for maximum effectiveness by deploying agents throughout the fleet.

Commercial-Off-The-Shelf (COTS) agent-based systems for diesel and gas turbine engine health monitoring have been in operation aboard various naval vessels for several years. The software agents are "knowledge-centric", which means that each agent is linked to a specific diagnostic knowledgebase. The knowledgebase defines specific equipment failure modes and their related measurable effects through existing machinery plant sensors. Multiple knowledgebases can be created, with each pertaining to a separate machinery plant, specific system within a plant, or even an individual piece of equipment.

The agents use artificial neural networks for diagnostic reasoning of machinery faults. The neural networks learn to associate patterns of alarm conditions with the machinery faults entered into knowledgebases. Once the fault-symptom associations have been learned, an agent uses this knowledge to perform real-time diagnostics and prognostics. The agent assesses current alarm conditions from real-time sensor inputs. It then recalls from its neural network memory those faults having symptom patterns most closely matching detected alarms.

A separate neural network is created for each knowledgebase. The agent uses the neural network associated with the knowledgebase to which it is attached. Several different agents can share the same neural network. For example, separate diagnostic and prognostic agents may be created that both attach to the same knowledgebase. Both of these agents will use the same neural network for their monitoring and analysis tasks, one for diagnostics and the other for prognostics, as depicted in figure 2 below.

Anatomy of a Diagnostic Software Agent

Figure 3 depicts a conceptual view of sensor data flow and diagnostic processing for a generic diagnostic agent. The circled numbers next the various components shown in figure 3 are discussed below. As indicated, various software components are remotely upgradeable. This architecture can be applied across a range of device complexity. For example, at the simplest level, it can be applied for sensor diagnostics/prognostics. While being deployed primarily for complex devices, such as diesel and gas turbine engines, the underlying diagnostic technology is applicable to any electro-mechanical device and is extensible to large-scale systems. The design also accommodates the implementation of component-level intelligence, as it provides standard interfaces for reporting device health status.

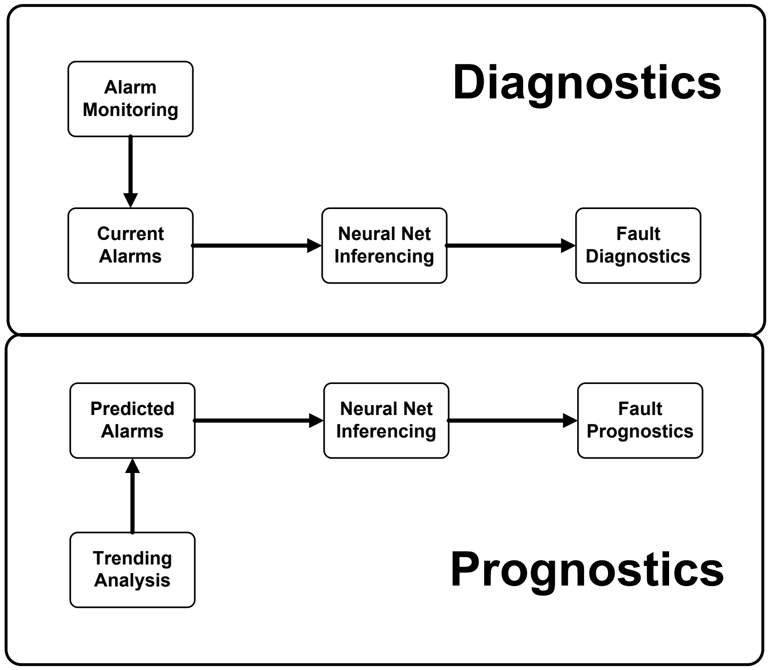

Figure 2 – Knowledge-centric Intelligence for Diagnostics and Prognostics

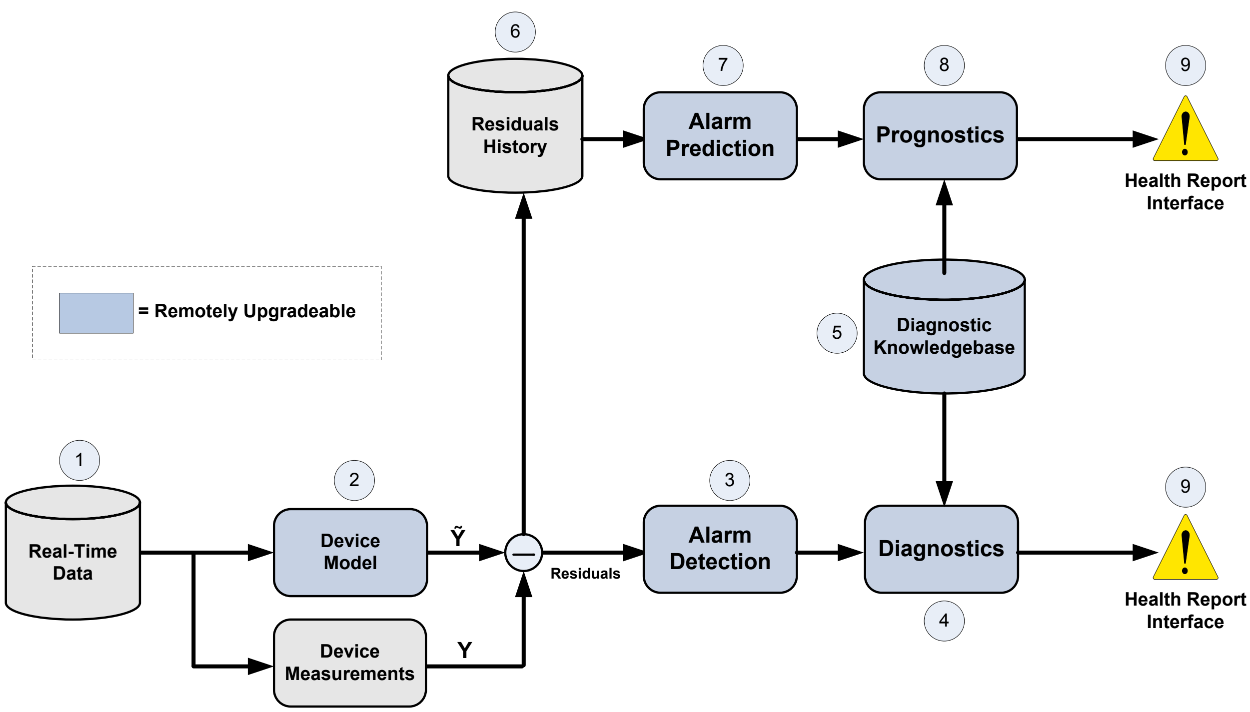

Figure 3 –Diagnostic/Prognostic Software Agent Data Processing

1 ***Real-time data*** is input to the diagnostic module and is used for both model estimation and residual generation. Sensor inputs can be acquired directly through internal electrical interfaces of the diagnostic module or be obtained via data communications interfaces with existing plant automation (e.g. Ethernet, wireless, etc.)

2 ***Model-based diagnostics*** rely on some type of model for the device under diagnosis. For key diagnostic or performance parameters, the *device model* is used to derive expected device outputs ($\hat{Y}$) from other salient device measurements. The estimated output is then compared to actual device measurements (Y). The difference between expected and measured outputs forms a residual which is key to detecting anomalous device behavior.

3 The ***alarm detection*** process detects anomalous conditions for the device. The distribution of residuals for healthy devices is statistically quantified. If the device model accurately reflects the device's behavior, the residuals can be represented by a zero-mean Gaussian process with known variance. Instantaneous residual values exceeding statistically derived confidence regions are classified as anomalous and, along with other similarly classified alarms associated with other device parameters, are input to a diagnostic reasoner for interpretation.

4 The ***diagnostic reasoner*** performs pattern recognition based on internal representations of diagnostic knowledge acquired from pre-training with the diagnostic knowledgebase. Advanced pattern recognition, neural network algorithms are applied to associate detected residual alarm conditions with known fault conditions. This technique is fast, memory efficient, capable of real-time performance, and produces Bayesian probability estimates based on the quality of match between stored diagnostic knowledge and detected alarm conditions.

5 The ***diagnostic knowledgebase*** maintains the essential associations between fault conditions and expected alarm conditions. This knowledge can be acquired from a variety of sources, such as device manufacturer or other experts, failure experiments on the actual device, computer simulation experiments, historical customer trouble call or maintenance records, etc. This knowledge is typically derived through a failure mode and effects analysis on the device.

6 Fault predictions (prognostics) are based on the ***residuals history***. This function manages recording and maintenance of the historical data store. System configuration settings are used to control history length/memory and these are dictated by the prediction horizon of the prognostic (i.e. how far ahead one wishes to predict faults). Prognostic schemes have been developed to predict across multiple prediction horizons

(e.g. short, medium, and long-term prediction) using varying time resolutions of residual histories.

7 ***Alarm prediction*** involves quantifying the trends in device residuals over time and using detected trends to predict future alarm conditions. This function analyzes the residuals history through statistical trending techniques to determine if any significant trends are occurring. A residuals trend indicates a discrepancy between the device's actual behavior and its model estimate. These trends are early indicators of an anomaly, either in the device itself or one of its sensors. This function uses the same attributes of the residuals statistical distribution as the alarm detection function to determine if and when a device alarm condition will occur within the prediction horizon.

8 The ***prognostics*** function performs similarly to the diagnostics function, but inputs predicted alarms instead of current alarms. It relies on the same diagnostic knowledgebase and pattern recognition function as diagnostics, but outputs predicted device faults based on detected trends in its residuals. As subsequently discussed below, trends can also be used to determine remaining time until predicted alarm occurrence. This important information can be relayed to maintenance decision makers in advance of an equipment failure to avoid disruptions of service, thus improving reliability, mission readiness, and platform availability. While the diagnostics function is more concerned with restoration of service, the prognostic function addresses avoidance of loss of service.

9 The ***health report interface*** is a standard mechanism for passing diagnostic and prognostic information to other software, such as graphical user interfaces, as well as other systems such as maintenance management systems, etc. This interface can also be configured to replicate machinery health information ashore via standard Internet connections for various distance support applications.

The agent software is designed to allow remote upgrades to embedded device diagnostic intelligence throughout the ship's life cycle. The blue-shaded components of figure 3 (circles 2, 3, 4, 5, 7, 8) are specifically designed for remote upgrading. To avoid product obsolescence, it is imperative that embedded diagnostic knowledge be current and as complete and accurate as possible. Diagnostic knowledge management is a critical supporting technology and updating onboard intelligence with new experiential knowledge accumulated over time is required throughout the ship's life cycle. Remote upgrading will also minimize long-term technical support/service costs.

Model-Based Equipment Performance Analysis

Machinery performance assessment is accomplished by deriving baseline performance models for the salient parameters across the machinery operating range. Sensor measurements are then compared to the "healthy state" baseline model and performance deviations are computed. During anomaly prediction, these deviations are trended to determine if they will exceed statistical limits for the machinery process. Predicted deviations falling beyond established operating thresholds are considered as anomalous behavior and are flagged as predicted alarms. Figure 4 shows an example of a typical baseline performance relationship. Measurements of charge air pressure are plotted against engine load. A regression curve has been fitted through the raw data to derive an equation expressing the mathematical relationship between the two measurements. The equation represents a baseline model and can be used to estimate expected values of charge air pressure for any engine load.

Assuming the baseline relationship represents healthy equipment behavior, measured machinery performance can be expected to follow the baseline. If a machinery problem develops, its behavior will no longer follow the baseline and an anomaly between measured versus estimated values will occur. These discrepancies can readily be detected by monitoring the computed deviation between measured and estimated performance parameters. Software agents automatically compute, record, and analyze deviation values.

The statistical distribution of deviation values are typically assumed to be Gaussian and are characterized by their mean and standard deviation parameters. The standard deviation is often multiplied by some factor to establish statistical thresholds defining a region within which the machinery performance deviations should normally vary during "healthy" operation. By comparing machinery performance deviations to their statistical thresholds, anomalies can readily be detected and predicted. Deviation values that exceed established threshold limits are declared as alarm conditions by the agents.

Early Failure Detection via Automatic Predictive Analytics

Prognostic agents automate the highly technical analytical work that typical ship crews do not have the time or skills to perform. This work is essential, however, to achieving the benefits of condition-based maintenance. These agents perform historical data archiving, model-based performance assessment, automated trending analysis, alarm prediction, fault prediction, and prognostic event logging. Prognostic software agents predict future machinery faults and determine when maintenance should be carried out. By predicting machinery problems before they occur, unexpected breakdowns can be avoided. In the absence of significant trends, equipment overhaul periods may be rationally extended, thereby eliminating unnecessary maintenance work. The ability to predict future maintenance requirements leads to improved maintenance planning and cost management. Maintenance and repair decisions can be tied to actual plant operating conditions based on the severity of degrading trends and predicted plant problems.

During anomaly prediction, historical deviations are trended to determine if they will exceed statistical limits in the future. Predicted deviations falling beyond established operating thresholds are considered as anomalous behavior and are flagged as predicted alarms. The software agents automatically compute, record, and analyze deviation values as the basis for prognostics.

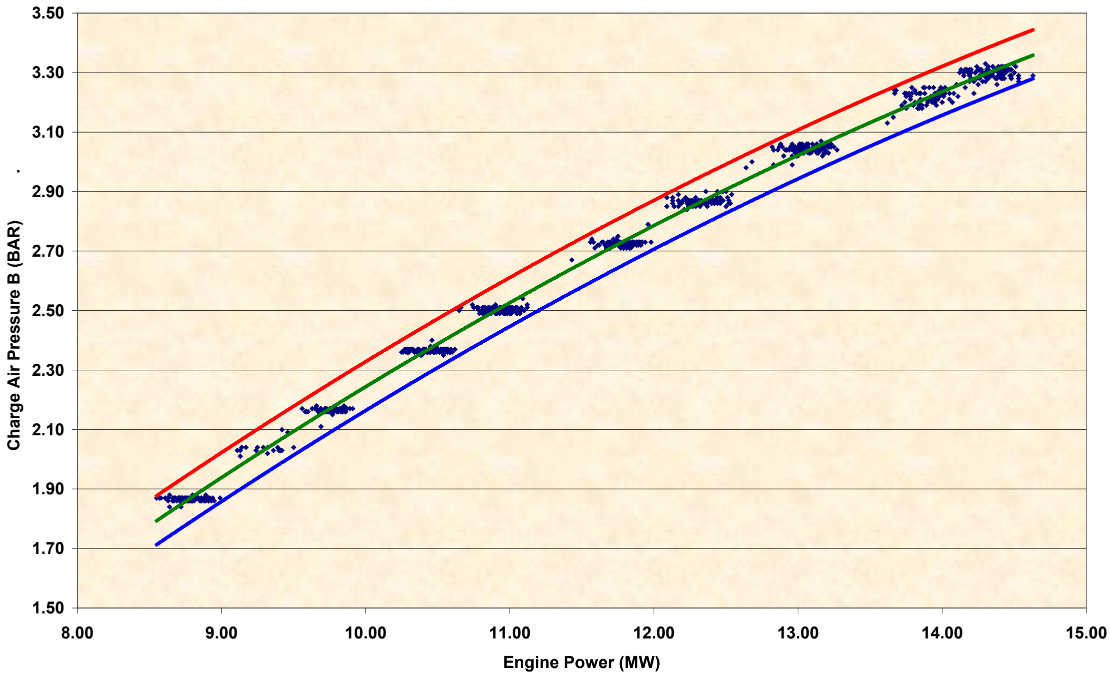

Figure 4 – Typical Baseline Performance Relationship

The exact point in time that an alarm is predicted to occur correlates directly with the predicted time to failure of the related machinery component. The prognostic agent will report the estimated time-to-failure associated with each predicted machinery fault. This estimate is derived from the time available to run the equipment until the first predicted alarm associated with a given fault. The varying severity of different deviation trends may result in some alarms predicted to occur before others. By reporting the earliest predicted alarm, the user is given the maximum amount of time to take corrective action prior to actual equipment failure.

The prediction process is illustrated in figure 5. The ***Maximum History Length*** and the ***Prediction Horizon*** are prognostic agent attributes. All alarm conditions predicted in this manner are automatically assessed by the prognostic agent's neural network.

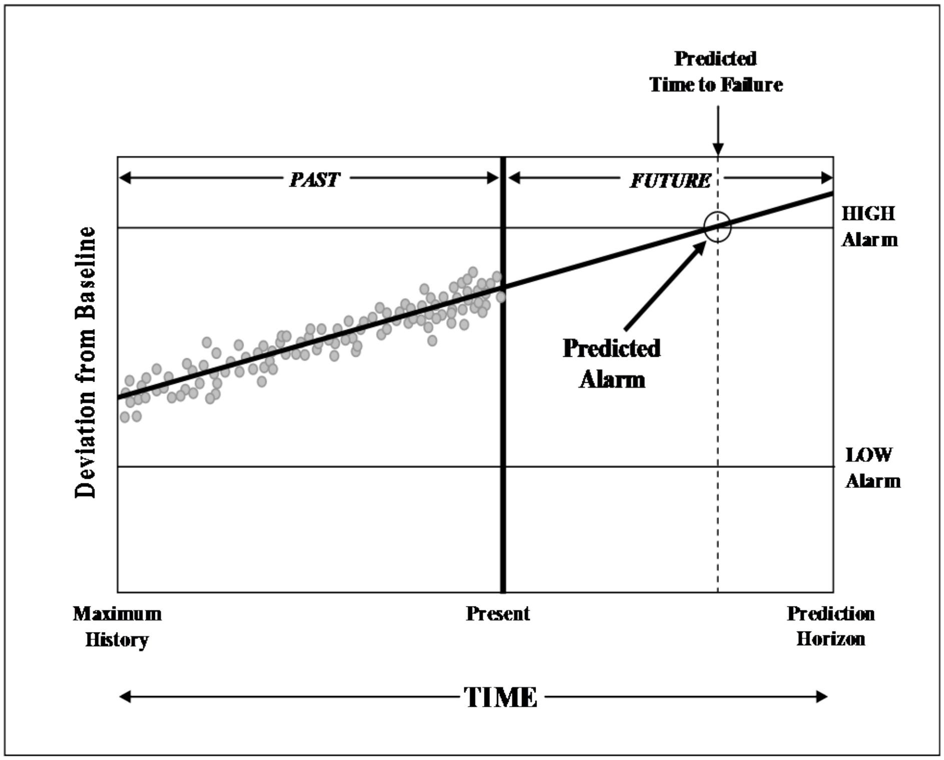

Figure 5 – Predicting Future Alarm Conditions from Deviation Trends

Case History: Diesel Engine Turbocharger Fouling

Diagnostic and prognostic agents have been deployed on a number of Navy ships as part of MACSEA's DEXTER™ system. In late August 2006, the USNS Big Horn was upgraded with the latest version of the agent software. Diagnostic event data retrieved from the vessel covering the September 2006 time period revealed several high probability diagnostic calls relating to the port main diesel engine (PME) turbochargers. In fact, turbocharger fouling was called out with 100% probability based on multiple symptoms detected on both banks of the engine.

Diagnostic Agents on USNS Big Horn

Subsequent discussions with the Chief Engineer aboard indicated that the engines had been frequently operated at low loads as mission requirements dictated. This practice is allowed by the manufacturer's technical documentation but is not considered the preferred method of operation. Diesel engines are designed to operate most efficiently under loaded conditions (typically 80-90% of maximum power rating). Low loading causes low combustion temperatures, incomplete combustion and results in increased depositing of material on exhaust train components. When an engine must be operated at low loads for extended time periods, the manufacturer recommends periodically running the engine at 70% maximum continuous range (MCR) or greater for 15 to 20 minutes every 10 hours. This action is intended to prevent the accumulation of unburned fuel or oil in the cylinders. The Chief Engineer indicated that he was following this recommended practice during low power operations by routinely increasing engine power about every 4 hours to burn out carbon deposits. He also indicated that no obvious signs of system degradation were observed throughout the deployment and that the turbochargers were scheduled for regular overhaul during the ship's upcoming shipyard period in December 2006.

During the shipyard engine repair work, it was noted that significant carbon buildup and fouling had occurred throughout the port main engine and its components, including the air coolers and turbochargers. This verified that the diagnostic alert calls made by the software agents were indeed correct.

Figures 6(a) through 6(d) show the main symptomatic conditions of turbocharger fouling detected by the software agents. The figures compare these performance parameters from before the overhaul (September 2006) to after the overhaul (December 2006). Turbocharger performance is noticeably improved subsequent to the overhaul.

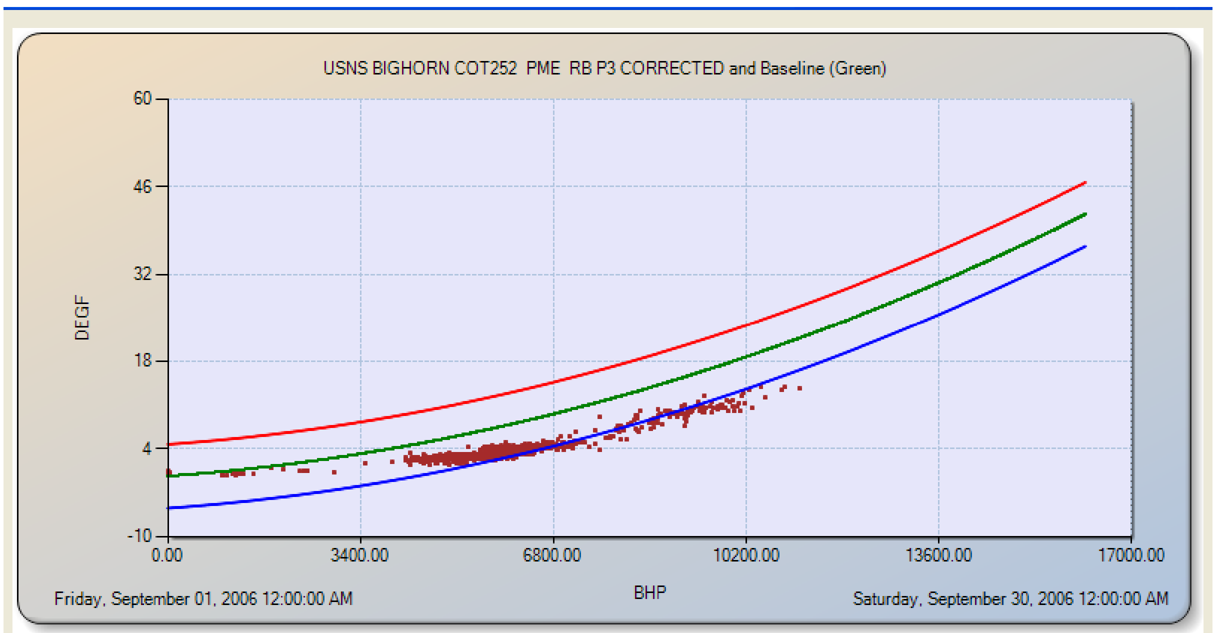

Figure 5(a) – PME Right Bank Charge Air Pressure – BEFORE OVERHAUL

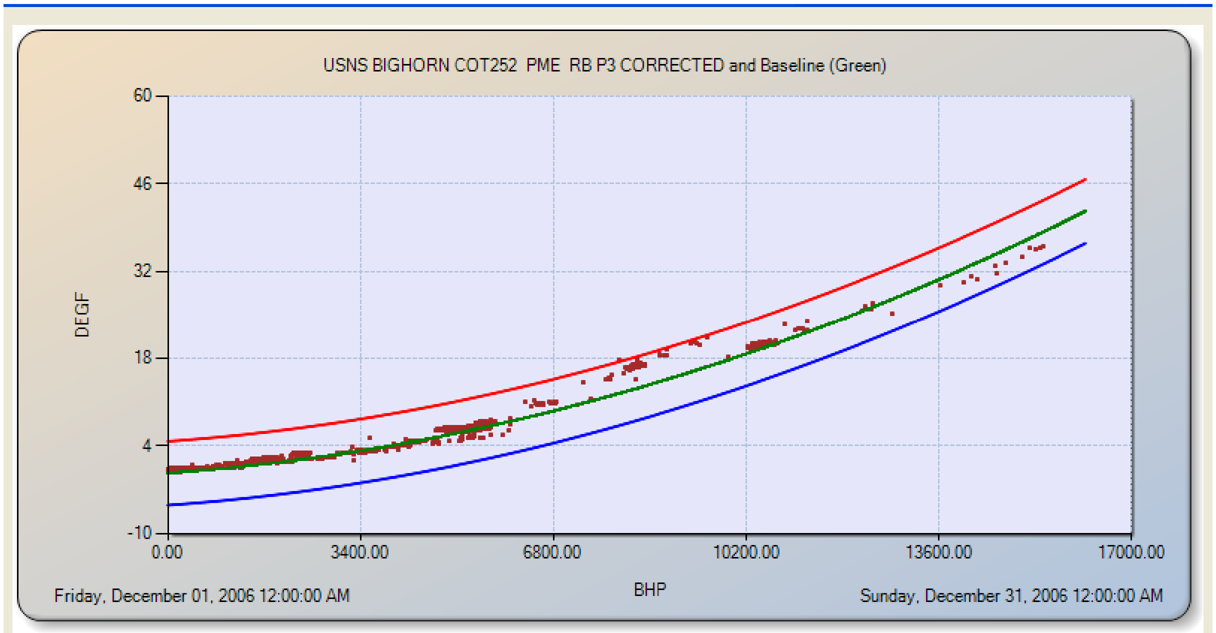

Figure 5(b) – PME Right Bank Charge Air Pressure – AFTER OVERHAUL

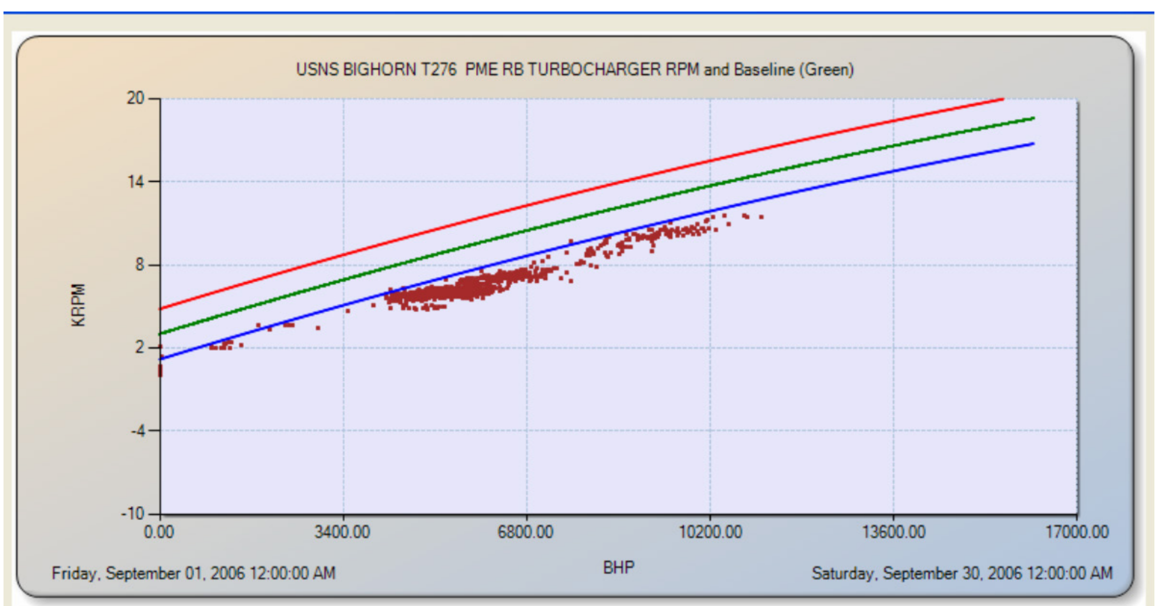

Figure 5 (c) – PME Right Bank Turbocharger Speed – BEFORE OVERHAUL

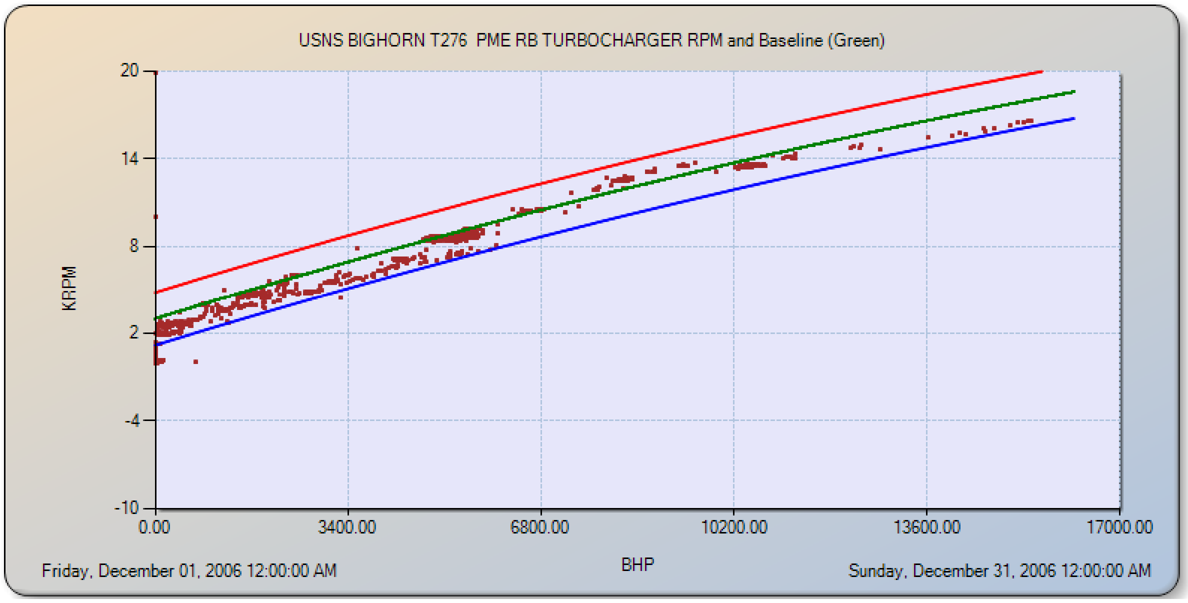

Figure 5(d) – PME Right Bank Turbocharger Speed – AFTER OVERHAUL

Agents Value for Reducing Maintenance Costs

Software agents automate the data reduction and analysis work required to implement a fleet-wide Condition-Based Maintenance strategy. A key aspect of this involves monitoring for changes in machinery operating performance over time. This follows a model-based diagnostic approach, whereby baseline models of engine performance are continuously compared to actual performance. When actual engine performance deviates from the model (expected) behavior beyond reasonable statistical limits, software agents issue alerts to the crew, reporting the most likely machinery faults based on their embedded diagnostic intelligence.

The analytics provided by the agents are fairly complex and would require excessive time and special engineering/statistical skills for the crew to perform. These analytical tasks typically are not performed on a consistent and reliable basis, yet in order to realize the benefits of CBM, this work must be accomplished. Software agents automate this work for the crew and provide actionable information to head off problems, not just raw data that sometimes is difficult to assimilate and correlate, particularly when large amounts of machinery data are being acquired.

Software agents can provide ship-owners with these key benefits:

- **consistent and reliable equipment health monitoring**
  - The analysis procedures can be as analytically complex as necessary
  - Expert diagnostic knowledge embedded into agents can be distributed across the fleet for a consistent, repeatable, scientific analysis of engine health, no matter what crew is aboard

- Because it's all automatic, there's no additional crew workload
- Historical diagnostic alert information is recorded for further review when it's time to make informed maintenance decisions
- Agents work 24/7, reliably and continuously

- **transformation of voluminous raw data into actionable information**
  - Agents are an analytical tool to be used by the crew
  - Hundreds or thousands of data points can be monitored
  - Agents correlate, assimilate, and make sense out of even very subtle performance changes that often remain hidden in the volumes of data
  - Agents will direct the crew to detected problems that can then be investigated in more detail
  - Agents can help avoid large, expensive problems by alerting the crew at the earliest possible stage using predictive analytics

Conclusions and Recommendations

The levels of machinery automation in future ship designs will continue to increase, providing massive amounts of data for health monitoring of ship systems. The crews aboard future minimum-manned ships will be hard pressed to transform this raw data into information that supports an effective condition-based maintenance program. There will be more machinery data to monitor and fewer people with less time to analyze it. Yet with reduced manning, the importance of keeping a constant vigilance in machinery performance assessment will never be greater, as the attendance to machinery failures will draw a larger percentage of available onboard human resources. Software agents can continuously and automatically monitor machinery, identify impending failures, and predict its remaining useful life (or equivalently time to failure). Software agents can be empowered with computer representations of human knowledge, allowing them to perform information processing tasks on behalf of their human counterparts. The ability to impart intelligent processing functions into software agents will allow ship maintenance managers to leverage valuable diagnostic knowledge across a ship fleet to enhance readiness while proactively minimizing maintenance costs. Once constructed, the agents become a valuable resource that can be distributed when and where needed to enhance operations and performance. Software agents will become *workforce productivity multipliers*, as the human-agent team will provide higher levels of productivity at practically the same cost as that of just the human resource alone.

Diagnostic and prognostic software agents have been deployed to continuously monitor the health of main propulsion diesel engines on various Navy ships (MSC). The agents "live" on the shipboard data networks and perform comprehensive predictive analytics of engine performance to aid in maintenance management decisions as part of a CBM strategy. Diagnostic software agents are delivering value to the Navy, as described herein, by identifying equipment faults at the earliest possible stage, where remedial action is the least expensive. A recent case history of a diesel engine turbocharger fault called out by software agents, verified during a subsequent shipyard engine overhaul, demonstrates that software agent technology has value for effective equipment health monitoring on minimally-manned ships and as rational maintenance decision aids for both onboard and distance support.

The Navy has recently initiated a project to implement software agent technology aboard the LSD class ships for diagnostics and prognostics of the main propulsion diesel engines and ship service generators. The initial installation is scheduled as part of the mid-life upgrade of the first ship of the class in July 2008.

To date, shipboard diagnostic software agents have been applied primarily to main propulsion diesel and gas turbine engines, as well as their supporting auxiliary subsystems. The implications for expanding agent application to other equipment are far reaching, ranging from health monitoring of very simple to very complex devices. A recommended strategy is to begin at the simple end of the spectrum and proceed to develop a comprehensive, integrated, multi-agent environment. Sensor diagnostics, for example, is the essential enabling technology for all shipboard control devices that require reliable data to synthesize decisions and actions. Yet, this area has traditionally been neglected for cost reasons and sensor diagnostic systems aboard Navy ships are severely lacking (e.g. as compared to aerospace systems). If a critical sensor fails, there is typically no estimate available of the measurement that the failed sensor should have provided. In the context of future all-electric Navy ships, this technology shortfall represents a significant barrier to implementing advanced system functionality that will support drastic manning reductions, integrated power system automation, and intelligent system reconfiguration under failure and/or battle damage conditions. None of this higher-level functionality can occur without accurate, reliable sensor diagnostic algorithms that can not only detected a sensor failure, but can also provide an analytical estimate of the lost sensor measurement to higher-level control systems.

Once a reliable foundation of sensor systems has been established, agent technology can then be expanded to more complex devices, such as the integrated power system (IPS). Reduced manning initiatives will place high value on agent-embedded, upgradeable diagnostic knowledge within such complex devices. These systems should be self-diagnosing and provide actionable information to the crews to mitigate failures at the earliest possible time. Figure 7 highlights some of the major categories of IPS components for which diagnostic software agent development should be considered.

The Navy seeks both affordability and reliability for the current, and more importantly, the next generation ships. Software agent technology supports new diagnostic engineering paradigms for dealing with complexity of ship system designs, as well as to manage and extract maximum value from diagnostic knowledge. Previous diagnostic practices and tools applicable to past generations of ships, relying on people and paper-based troubleshooting, are inadequate for future diagnostic technology delivery, particularly for advanced IPS designs. Integrating diagnostic knowledge via software agents into the ship systems will facilitate maintenance cost containment over the ship's lifecycle.

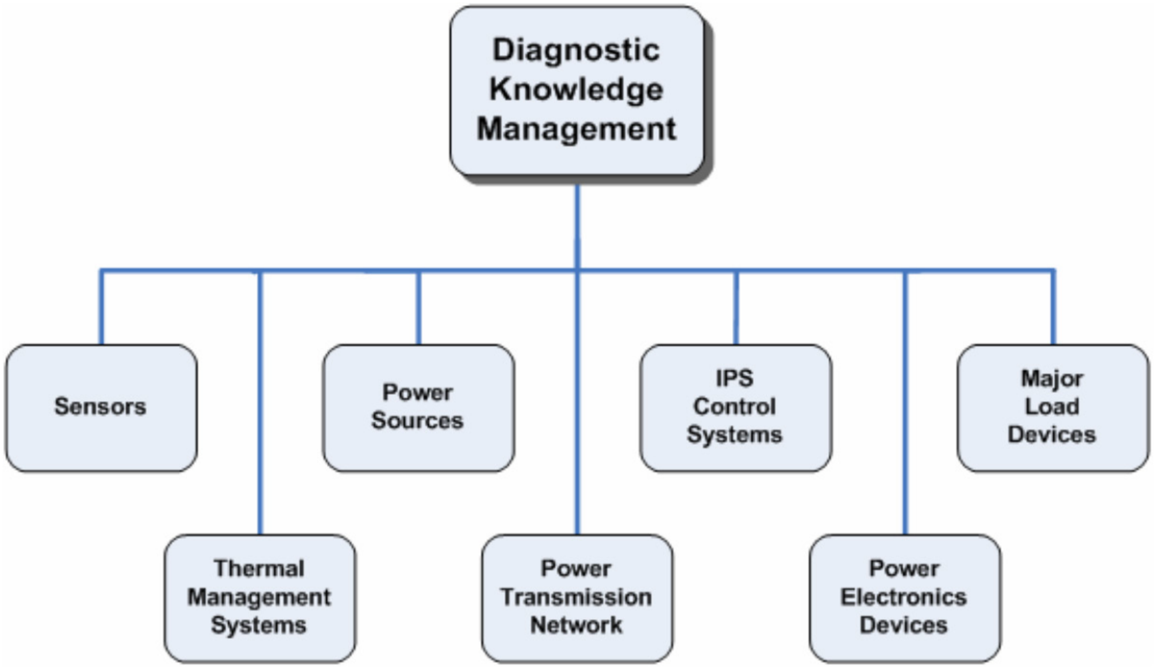

Figure 7 – Target IPS Components for Diagnostic Software Agent Application

Author Biography

Kevin P. Logan is President/Founder of MACSEA Ltd., a 25-year old company that manufactures machinery diagnostic systems. He has a BA in Mathematics from Eastern Connecticut State University and performed graduate studies in Operations Research at the University of New Haven. Mr. Logan has performed applied research in vessel performance analysis, machinery diagnostics, and software agent technology throughout his 27 year career in the maritime industry. His most recent research includes intelligent software agents, neural networks for automated shipboard learning systems, and automated machinery model synthesis for model-based reasoning systems. He has published over 16 technical papers and is a member of ASNE, IEEE, SNAME, MFPT, and the International Neural Network Society.