Kevin P. Logan and Capt. John G. Walker (USN Ret.)

Embedded Diagnostics for LSD Machinery Control System Upgrade

ABSTRACT

As part of the U.S. Navy's LSD 41/49 Class Mid-Life Upgrade program, diagnostic software agents are being embedded into the modernized Machinery Control System (MCS) to support sailors in keeping ship equipment running reliably within the context of reduced total ownership costs. Software agents will automate the essential analytical work necessary to implement an effective Condition-Based Maintenance (CBM) program at minimum cost to the Navy. Intelligent software agents will deliver real-time diagnostics and predictive analytics for various mission-critical systems, such as the main propulsion diesel engines, electric generators, etc. It is envisioned that software agents will become as common place on future shipboard networks as the anti-virus software running on most PCs today.

The technological innovations described in this paper deliver new capabilities for automated diagnosis of equipment faults, early warning of equipment health problems, automatic generation of maintenance work orders, and timely delivery of equipment health information to both the crews and shore-side support staff. Class-wide LSD deployment will occur across the four-year LSD mid-life upgrade program.

INTRODUCTION

Navy fleet maintenance has recently transformed into "Operations-Focused Maintenance", striving to improve operational readiness and availability while providing an enhanced surge capability. In the past, surface fleet maintenance needs were met through a two-year cycle. Ships were deployed for six months and then spent the next 18 months primarily in a maintenance and training sequence. A ship would then be ready for another deployment. Although the six-month deployment schedule still remains, ships now must enter the basic training phase earlier, perhaps before maintenance has been completed. They must also be maintained in a higher status so that they can surge if required. Also, due to uncertain operational demands, ships may be deployed for longer or shorter periods than in past tradition. The primary challenge now is establishing the management processes to identify and complete necessary maintenance activities while meeting surge and other unpredictable operational demands.

Navy requirements for Operations-Focused Maintenance include (Brooks 2008):

- Engineering for reduced maintenance,
- Extending time-between-overhauls,
- Reducing time in depot maintenance,
- Optimizing continuous maintenance through effective distance support, and
- Improving onboard equipment health monitoring with CBM technologies.

All of these factors can significantly reduce maintenance costs.

An urgent need to engineer for reduced maintenance becomes readily apparent when considering two major "maintenance mega-trends" taking place; 1) new technology has dramatically increased the complexity of many ship systems, and 2) there is a continuing decrease in the skilled maintenance labor pool, caused partly by new systems complexity (i.e. new skill set requirements not being met), workforce attrition, and manning reductions. These mega-trends will continue into the foreseeable future.

Complex systems on naval ships are difficult to cost effectively maintain (Dean, Reina, and Bao 2008, US OMB 2002). In general, the three types of maintenance applied to any equipment, irregardless of its degree of complexity, include corrective, preventative, and predictive maintenance. Of these three, predictive maintenance is impacted the most by complexity issues, as it requires expert-level system assessments to identify incipient and often very subtle problems such that maintenance action can be taken to avoid failures. Increased complexity of ship systems, particularly the automation, have made equipment failure detection more difficult, more time consuming, and more technically sophisticated.

The second maintenance mega-trend involving the declining maintenance workforce is supported by a recent study that found a 4.1% decrease in DOD and a 9.8% decrease in civilian maintenance workforces during the 1997-2001 timeframe (Clifford, Callendar, and O'Meara 2003). The study also concluded that, based on the average years of service, the DOD maintenance workforce is becoming younger and less experienced. The declining workforce has been accompanied by a significant increase in deferred maintenance actions across the Atlantic Fleet during the 1995-2000 timeframe, a significant portion of which was attributed to maintenance training and manpower issues (Yardley, et al. 2006). The situation is likely to become more acute as even more complex, high-tech ship systems are introduced into the fleet.

While technology may be contributing to these problems, it can also offer potential solutions, particularly with regard to outfitting "Intelligent" ships with intelligent systems. The levels of machinery automation in future ship designs will continue to increase, providing massive amounts of data for equipment health monitoring, by some estimates, on the order of 100,000 data points (Logan Jan/Feb2007). 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. Simply put, 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 unexpected machinery failures will likely draw a larger percentage of available onboard human resources.

Machinery performance monitoring is an area where the immediate exploitation of software agent technology can yield substantial economic benefits (Logan Dec2007). Software agents can continuously monitor machinery, identify impending failures, and predict remaining useful equipment life. Agents can perform the tedious, repetitive, time-consuming, or analytically complex tasks on behalf of sailors who may not have the time, skills, or motivation to perform these tasks themselves. Software agents can serve as expert assistants in equipment health monitoring, no matter how complex the machinery process. Agents can automatically monitor and analyze hundreds of thousands of machinery plant sensors, just as the anti-virus software on most PCs analyzes hundred of thousands of virus signatures unobtrusively in the background.

Software agents can be empowered with computer representations of human knowledge, allowing them to perform information processing tasks exceeding the sailors' capabilities. Once constructed, the agents become a valuable resource that can be distributed when and where needed to enhance operations and performance. Their intelligent processing functions will allow ship maintenance managers to leverage valuable diagnostic knowledge across a ship fleet to enhance readiness while proactively minimizing maintenance costs. Software agents can serve as *maintenance workforce productivity multipliers*, allowing sailor-agent teams to achieve higher levels of ship readiness at reduced maintenance costs.

DIAGNOSTIC SOFTWARE AGENTS FOR FLEET MODERNIZATION

CBM and reliability-centered maintenance (/RCM) strategies are widely used today by an ever growing number of organizations, including the Navy. The main objectives of CBM/RCM are to optimally manage the risk of equipment failure, while maximizing equipment reliability and minimizing maintenance costs. Meeting these objectives is becoming increasingly challenging due to the maintenance mega-trends previously discussed. Increasingly complex automation and voluminous data collection will demand advanced operator/analyst skill sets to implement effective CBM programs, yet fewer resources will be available (people and money). This is exactly where intelligent software agents can help.

Software agents can be rapidly integrated into existing shipboard computer networks, adding significant capability and value to automation systems. The agents can serve as *virtual* maintenance team members whose job is to oversee these advanced systems. This capability will become increasingly important with increases in system complexity, as maintainer training costs will become unrealistically high. (One cannot help to draw from the analogy of automobiles, which are now far too complex for the typical owner to troubleshoot and maintain).

Intelligent software agents offer many advantages as embedded components of *intelligent* ships. Human intelligence can be cloned and then replicated by creating distributable knowledgebases, for example, equipment diagnostic or troubleshooting knowledge. Human-like reasoning can be emulated with advanced inferencing algorithms, such as artificial neural networks. By coupling the knowledgebases and the inferencing techniques, the software agents can leverage *corporate knowledge* assets to automate complex and tedious tasks for sailors. Corporate knowledge assets are derived from the best possible sources, which typically include OEM technical information and Failure Mode and Effects Analysis (FMEA) performed by Subject Matter Experts. Figure 1 illustrates the consolidation of the diagnostic information into a knowledgebase, which is subsequently transformed into a neural network reasoner. The neural network becomes the embedded intelligence of the diagnostic and prognostic software agents.

The Navy can benefit greatly from the agents' ability to enhance the performance of inexperienced or overloaded ships force, as well as shore-side maintenance support staff.

Agent Intelligence Creation

The "knowledgebase-centric" software agents are characterized by three types of intelligence: 1) a reasoning mechanism, 2) diagnostic knowledge, and 3) equipment performance models. Each of these will now be addressed.

Diagnostic Reasoning with Artificial Neural Networks

The fielded software agents employ artificial neural networks as their reasoning mechanism. Neural networks are well-proven pattern recognition devices, tolerant of noisy or incomplete input data, handle uncertainty well, are capable of real-time pattern recognition, and learn classification tasks from training data.

Neural networks are modeled after the neuron processing elements occurring in biological nervous systems. Neurons (processing units) receive input signals from other neurons, perform a weighted summation of those inputs and generate an output signal while performing a nonlinear transformation. An error correction procedure of some type is typically employed for network learning via weight adjustments (Rumelhart, Hinton, and Williams 1986). Neural network-based diagnostics are more robust than rule-based systems because when one or more input values are missing, the network is still able to make a similarity pattern match from the training data that it has learned.

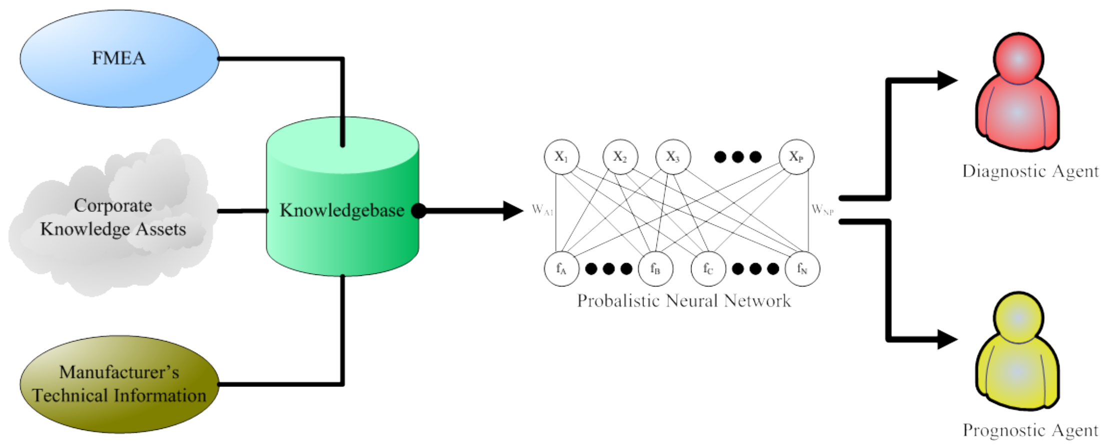

Figure 1 –Software Agent Intelligence Creation

The critical factor in deploying neural networks is having good training data. In the case of diagnostics, this is data that relates equipment faults to specific measurable symptoms. A practical approach to generating this data is to conduct a FMEA on the equipment targeted for diagnostic coverage. The FMEA documents potential equipment failures, expected symptoms, and the probable causes for the failures. It is typically performed by subject matter experts, such as the Navy's In Service Engineering Agents. Other possible diagnostic information sources include manufacturer's technical documentation, equipment procedures for operation, maintenance and troubleshooting, in-house system expertise, and historical data. The results of the FMEA represent valuable diagnostic knowledge and can be rapidly incorporated into a knowledgebase and, in turn, a neural network. Knowledge gaps or inaccuracies in the FMEA are dampened by the use of *probabilistic neural network* techniques.

While many neural models exist, the Probabilistic Neural Network (PNN) was chosen for embedded agent intelligence, primarily due to the following advantages:

- Fault classification probabilities are directly output from the PNN, generated by its nonlinear decision surfaces.
- PNN is capable of handling situations in which one or more input variables are missing or corrupted using valid statistical treatment of uncertainty.
- PNN can be rapidly deployed using existing experiential and empirical knowledge and can be readily updated as new knowledge is acquired.
- PNN training is effectively instantaneous, unlike other neural methods that require hundreds of thousands of learning iterations.

A PNN is used to classify symptom patterns according to the faults that may have generated the alarm conditions. A single PNN can be pre-trained to learn the associations between a large number of faults and their corresponding symptom patterns. Once trained, the PNN is connected to the machinery control system to perform real-time diagnostics. Automated prognostics are a direct extension of diagnostics, coupled with an automated trending analysis function (Logan ASNE ISS 2007, Logan 2003).

Many neural networks perform statistical computations on patterns contained in a training data set. These internal statistics are then used to classify new patterns presented as inputs to the trained network. The classification problem can be posed as an example of Bayesian classification in which it is desired to categorize a set of inputs (symptoms). In the context of the diagnostic and prognostic applications, the classification categories represent the different machinery fault conditions.

PNN uses a Bayesian decision strategy to classify unknown symptom patterns in a way that minimizes expected risk. As an example, consider only two possible fault categories, θ_A and θ_B. Based on a set of measurement symptoms:

$$X^T = [X_1, X_2, ..., X_P],$$

the task is to decide whether the fault (or "state of nature") is either θ_A or θ_B. The Bayes decision rule for this classification task, denoted d(X), can be expressed as:

$$\text{d(X)} = \theta_A \text{ if } P_A^{'} L_A P_A(X) > P_B^{'} L_B P_B(X)$$

$$\text{d(X)} = \theta_B \text{ if } P_A^{'} L_A P_A(X) < P_B^{'} L_B P_B(X) \qquad (1)$$

where:

$P_A^{'}, P_B^{'}$ = a-priori probability of faults θ_A and θ_B, respectively,

L_A = loss associated with decision d(X) = θ_B, when the fault = θ_A

L_B = loss associated with decision d(X) = θ_A, when the fault = θ_B,

P_A (X), P_B (X) = probability density functions for θ_A and θ_B, respectively.

The structure of the Bayesian decision model is powerful and very useful if the a priori probabilities and loss functions are available for use in the model. It is worth examining the component factors of the decision model, as useful statements about equipment reliability and readiness can be derived from it.

The a-priori probabilities of the two faults, $P_A^{'}$ and $P_B^{'}$, are difficult to accurately determine based on typical Navy maintenance information management practices. The recommended approach is to derive component failure distributions from historical maintenance records, assuming comprehensive records are kept. Beyond this, OEM failure statistics, such as Mean-Time-Between-Failure, can be used for a-priori fault probability estimations. In either case, the work effort in deriving these estimates can be substantial. Lacking any useful information, the effects of varying a-priori's can be removed from the decision analysis by assuming that all faults are equally likely.

The effects of making the wrong fault call, which in practice could also equate to making no call when a fault was actually present, are modeled though the loss functions in (1), i.e. L_A and L_B. Recall that L_A is the loss associated with making the wrong the decision, d(X) = θ_B, when the fault was actually θ_A. Loss values can be defined on any scale, whether monetary or in criticality metrics; however, as with the a-priori's, the work effort involved in derivation of the loss functions could be substantial, as expert subjective analysis is typically required.

Bayesian Fault Classification

The PNN's function is to classify or assign an input symptom pattern to a fault category. For most practical pattern classification problems, the Bayes optimal decision surfaces in this multi-dimensional "pattern space" are non-linear, making classification assignment difficult. For the two category example involving the decision rule from (1) above, the decision boundary between the two fault categories, θ_A and θ_B, is given by:

$$P_A\,\text{(X)} = \text{C}\ P_B\,\text{(X)}, \qquad (2)$$

$$\text{where C} = \frac{P_B^{'} L_B}{P_A^{'} L_A}.$$

The decision surface represented by (2) can be highly nonlinear, since there are few restrictions on the density functions, P_A (X) and P_B (X). In practice, the decision surfaces for diagnostic problems involve a large number of fault categories.

The key to using the PNN for diagnostics is estimating the class probability density functions

(PDFs) so that the Bayes decision rules can be implemented. Non-parametric methods of PDF estimation have been long proposed by many researchers and used with good accuracy (e.g. Parzan 1962, Murthy 1966). PDF estimation embedded in the fielded software agents use a Gaussian PDF estimator (Logan 2003). The PNN is readily trained from the results of the FMEA on the machinery plant.

Embedded Diagnostic Knowledge for LSD

Diagnostic knowledgebases were developed by performing a FMEA to identify all common or likely machinery faults. The targeted equipment for diagnostic coverage for the LSD class ships included four Colt Pielstick PC2.5 Main Propulsion Diesel Engines (MPDEs), four Fairbanks Morse OP38ND8-1/8 Ship Service Diesel Engines (SSDEs), and four KATO Ship Service Diesel Generators (SSDGs). In performing the FMEA, each fault was characterized by its measurable symptoms in the plant, as monitored by the MCS sensor instrumentation. Figure 2 shows the five salient steps in organizing and constructing a diagnostic knowledgebase.

USS Gunston Hall

In addition to the MCS signal inputs, performance data manually recorded from a cylinder combustion analyzer is also input for diagnostic processing. A common set of cylinder combustion-related faults were developed and applied individually to all 112 cylinders comprising the four 16-cylinder MPDEs and the four 12-cylinder SSDEs.

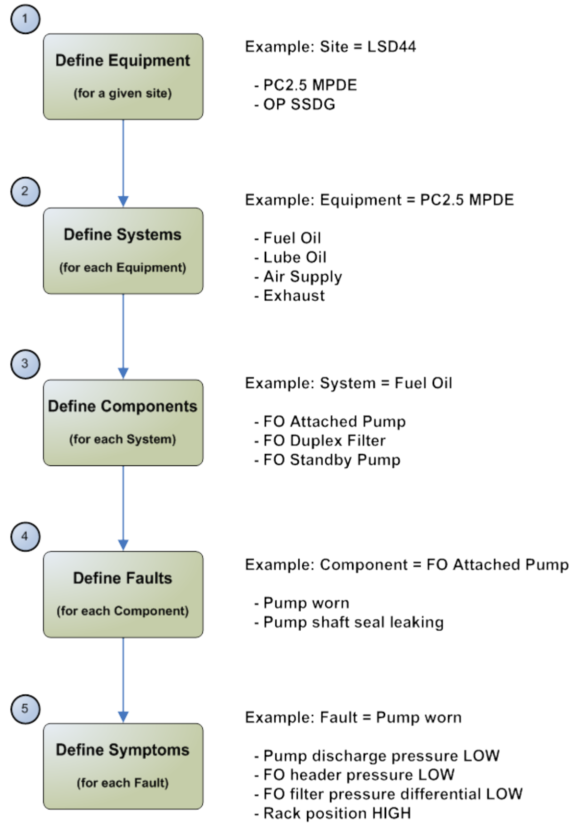

Figure 2 - Steps in Building Diagnostic Knowledgebase

In total, approximately 4400 individual engine fault diagnostics are performed by software agents aboard each LSD vessel.

The *USS Gunston Hall* (LSD 44) is the first vessel of the class to receive diagnostic software agents as part of her Mid-Life Upgrade. Installation was completed in July 2009. The *USS Germantown* will be the second installation during FY09, followed by ten additional vessels.

Diagnostic Agent Information Processing

Multi-agent techniques have been successfully applied to various distributed problem solving, information fusion, computing applications. In addition to localized data processing and computation, advantages include scalability, since once an agent's standard data processing, collaboration, and communication functions have been defined, additional agents can be deployed as needed when new equipment is added to an existing configuration.

The case for distributed machine intelligence and decentralized ship system architectures was presented in (Drew and Scheidt 2004), which proposed intelligent software agents as a generic deployment model to control and autonomously reconfigure all major ship systems, particularly under battle damage situations. As increases in processing power and advances in control theory make modern controllers more capable, the expected complexity of ships systems will likely increase exponentially in next generation ships. Future ship systems may be comprised of tens of thousands of connected components. Autonomous control using intelligent software agents appears to be one of the few options available for dealing with an enormous number of data elements and possible equipment states. Software agents used to implement distributed machine intelligence will allow collaborative control during reconfiguration situations. Diagnostics must be performed within the embedded agent intelligence on each device. Health status information can then be propagated over the control network to facilitate system level, real-time, dynamic control. The LSD modernization represents the first-generation deployment of such agents in the fleet.

Figure 3 depicts an overview of sensor data flow and diagnostic processing for a generic diagnostic agent (Logan 2007). This architecture can be applied across a wide range of device complexity, as 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 the neural network techniques previously described are amenable to chip-level deployment. The design provides for standard interfaces for reporting device health status to other integrated systems as part of an open-systems architecture. For LSD Mid-Life, the architecture illustrated in figure 3 was deployed for health monitoring of the MPDEs, SSDEs, and SSDGs.

Referring to figure 3, real-time data is input to the diagnostic module and used for both model estimation and residual generation. For LSD, sensor inputs are acquired directly via two data communications interfaces; one with the MCS and a second interface to the Shipboard Automated Maintenance Management (SAMM) system. The MCS interface provides strictly real-time data from the machinery plant. The SAMM interface supplies data manually collected via Personal Digital Assistants (PDAs) for sensor readings not currently connected to the MCS.

Agent – Machinery Control System Interface

The LSD MCS consists of multiple Programmable Logic Controllers (PLCs) forming the Advanced Engineering Console System (AECS) network. Roughly 4000 data signals are transmitted across the AECS LAN amongst the various PLCs. The MCS software acts as an intermediate “data broker” to all external data consumers, such as the software agents. The MCS software developed by NSWCCD Philadelphia shields data consumer applications from signal changes at the PLC level that may occur over the life of the ship. The LSD diagnostic software agents acquire approximately 450 MCS signals in real-time.

The LSD network architecture shown in figure 4 allows the software agents to distribute diagnostic/prognostic results to MCS workstations located throughout the AECS LAN. Sailors have the capability to monitor the equipment’s health from any workstation location.

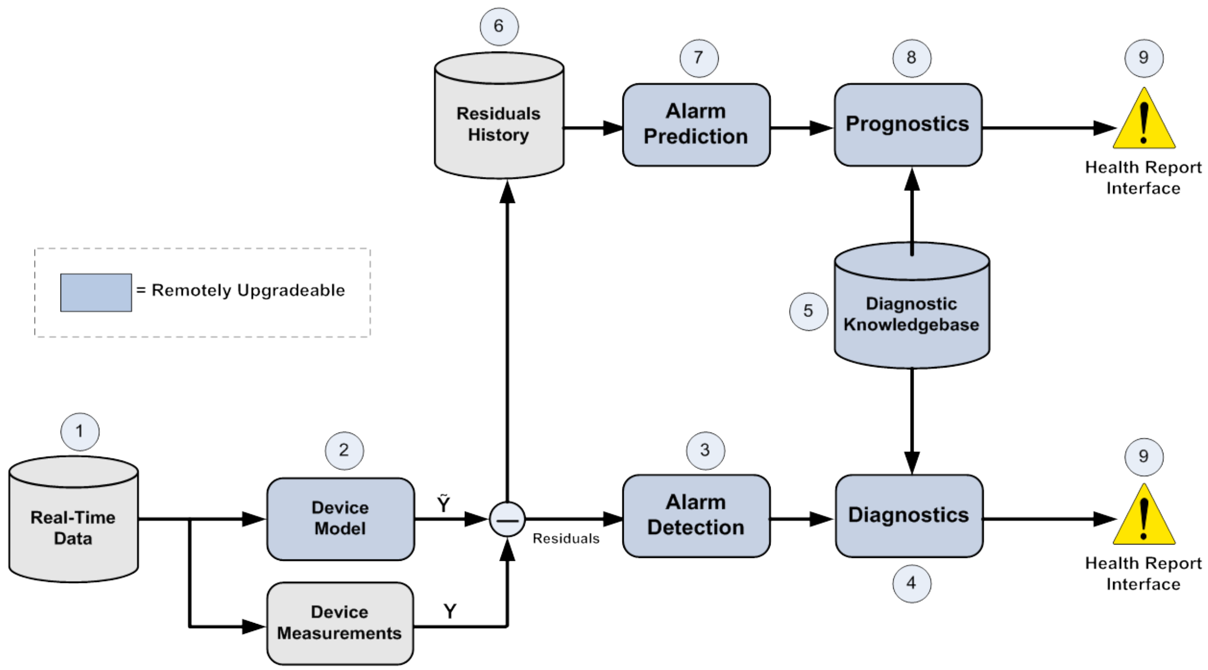

Figure 3 - Diagnostic/Prognostic Software Agent Data Processing

Shipboard Automated Maintenance Management System Interface

The Shipboard Automated Maintenance Management (SAMM) system, developed by the Military Sealift Command, has been adopted by the Navy for use aboard the LSD and other vessel classes as a preventive maintenance tool. SAMM is an automated system for documenting onboard maintenance and equipment configuration control. SAMM integrates various software applications, including maintenance scheduling, machinery history record keeping, machinery vibration monitoring, lube oil analysis, electronic watchkeeping data collection, and diesel engine combustion analysis. SAMM also provides fleet-wide maintenance data management through periodic transmissions of individual shipboard SAMM databases to a centralized, fleet database ashore.

Referring back to figure 3, the software agents employ *model-based diagnostics*. 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. These residuals are used in all diagnostics and prognostics, as opposed to raw

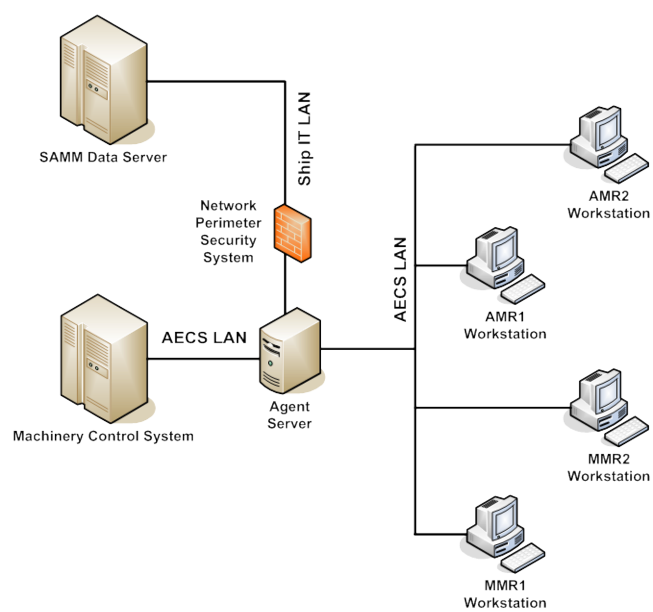

Figure 4 – Software Agent Network Architecture

sensor measurements. For LSD, initial baseline models were developed from test data provided by the engine OEM. A more complete modeling effort was then performed using actual machinery data collected during sea trials.

An *alarm detection* process detects anomalous conditions for the device and generates the symptoms used for agent diagnostics. The distribution of residuals for healthy devices is statistically quantified. 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. All abnormal conditions are recorded to a database along with the timestamp of occurrence.

The *diagnostic reasoner* implements the PNN previously described. The PNN algorithms are applied to associate detected residual alarm conditions with known fault conditions. This technique is fast, memory efficient, performs near real-time, and produces Bayesian probability estimates based on the similarity between stored diagnostic knowledge and detected alarm conditions.

The *diagnostic knowledgebase* maintains the essential associations between fault conditions and expected alarm conditions. This is the knowledge acquired from the FMEA process previously discussed. The knowledgebase can be continually enhanced over time and the agents can use upgraded knowledgebases without reprogramming.

Fault predictions (prognostics) are based on the *residuals history*. This function manages recording and maintenance of the historical data store. System configuration settings control 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.

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 determines if and when a device alarm condition will occur within the prediction horizon.

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 pre-trained PNN as diagnostics, but outputs predicted device faults based on detected trends in its residuals. Trends are also used to determine remaining time until predicted alarm occurrence. While the diagnostics function is more concerned with restoration of service following a fault, the prognostic function addresses fault prevention and maintenance cost avoidance.

The *health report interface* depicted in figure 3 is a standard mechanism for passing diagnostic and prognostic information to other software. For LSD, the software agents report all equipment failure alerts to the SAMM system, where maintenance Work Requests are automatically generated to both the sailors and, through distance support, to the support staff ashore. This process is described more fully below. Figure 5 shows the contents of a typical prognostic alert message sent from the agent to the SAMM system.

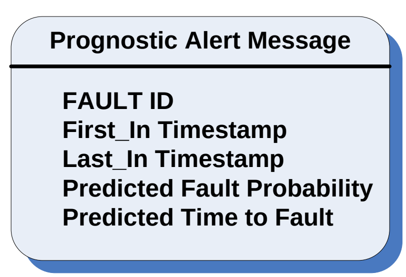

Figure 5 – Prognostic Agent Alert Message

Automatic Generation of Maintenance Work Orders

In an effort to streamline essential business processes related to reducing total ownership costs, the diagnostic and prognostic alerts generated by the LSD software agents are fed into the SAMM maintenance system for automatic conversion into Work Requests. These Work Requests are subsequently transmitted ashore for ISEA review. Approved Work Requests are converted into Work Orders via SAMM processing for follow-up action by the crew, as depicted in figure 6 below, Hence, the combination of software agents and SAMM automate the process of converting large amounts of raw engine data into *actionable information*. The Work Request/Work Order mechanism directs the crew and/or shoreside maintenance engineers to take follow-up action on detected machinery faults.

BUSINESS CASE ANALYSIS (BCA) AND PROJECTED COST SAVINGS

A BCA was performed to estimate the Return on Investment (ROI) and payback period for diagnostic software agent technology insertions into the LSD 41/49 Class' Mid-Life program (NAVSEA/PMS470 2007). The BCA assessed only MPDE maintenance cost data that included for Organizational, Intermediate, and Depot level activities for a two year period. The objective was to determine the cost avoidance potential of software agents for MPDE maintenance. *(SSDG maintenance cost avoidance was not considered in the BCA and represents additional potential savings to the Navy.)*

This study examined approximately ten-thousand Navy casualty data records and modeled this data through the software agent's diagnostic features to determine casualty avoidance data if the agents had been installed. Casualty avoidance data was then converted to avoided MPDE maintenance cost. From this analysis, the projected cost avoidance from deploying software agents is approximately 6%. ROI determination was therefore based on an assumed 6% cost avoidance in annual LSD 41/49 MPDE maintenance.

Based on the estimated 6% maintenance cost avoidance and the projected cost of software agent deployment, the ROI for diagnostic software agents on the LSD 41/49 Classes was estimated at $13 to $1, with a projected payback period of 13 months. The projected payback was considered conservative since the BCA excluded potential savings for SSDG maintenance.

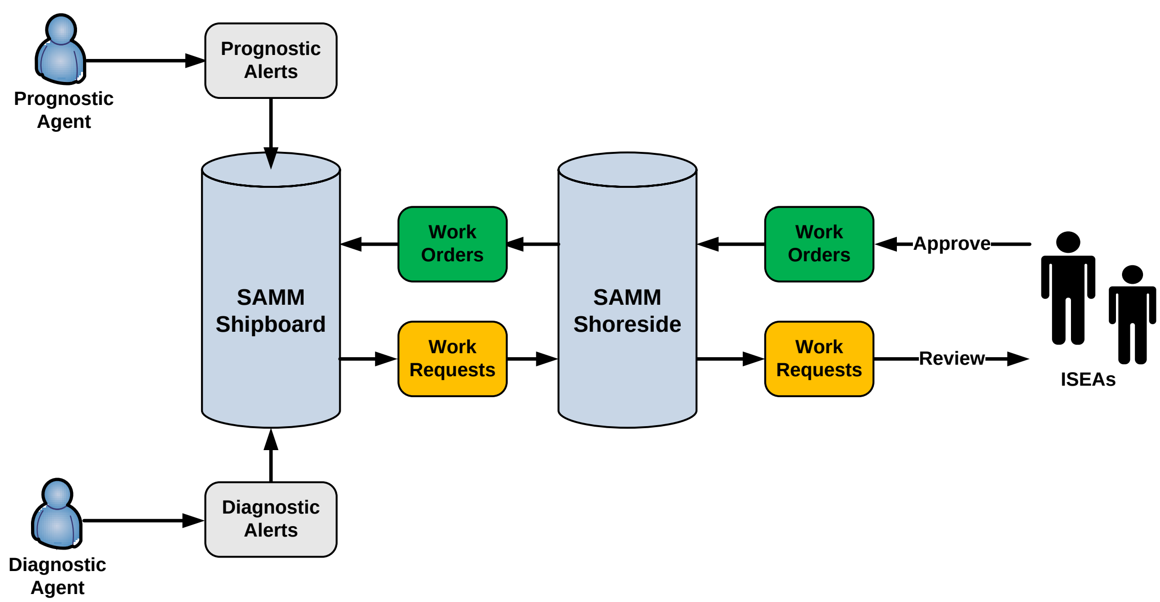

Figure 6 – Automatic Generation of Work Requests for Enhanced Distance Support

The BCA concluded that the installation of diagnostic software agents on all LSD 41/49 Class MPDEs would provide distinct condition based monitoring and assessment advantages. Each ship outfitted with real-time, onboard diagnostic/prognostic software agents can help sailors and maintainers correct minor operational deficiencies before they cascade into major engine casualties, thereby reducing the overall Operation & Support costs. The average annual maintenance costs on the MPDEs have consistently increased over the last few years. Installation of diagnostic software agents will provide today's shipboard engineers with immediate notification of impending MPDE casualties, a series of suggested maintenance and/or corrective actions to avoid these casualties, and the means to accurately document actual MPDE operating parameters prior to and during a casualty, should one occur. If these impending casualty notifications are promptly investigated and rectified, the number of catastrophic failures resulting from a series of minor cascading problems should decrease and the annual maintenance costs should be significantly reduced.

CONCLUSION

New diagnostic technology is being inserted into the LSD Mid-Life AECS/MCS upgrade in the form of software agents. These first-generation software agents will maintain a continuous health watch over mission-critical equipment, such as the main propulsion engines, electric generators, and related auxiliary systems. The agents implement expert-level diagnostic intelligence for over 4400 unique faults, provide early warning of equipment health problems, and data management services for life-cycle maintenance management. Through an interface to a computerized maintenance system (SAMM) developed by the Military Sealift Command and now used aboard LSD vessels, software agents trigger the automatic generation of maintenance Work Requests and Work Orders, while providing immediate delivery of equipment health information to the crew.

The underlying technology behind the software agents has been field-proven for approximately ten years aboard other naval and commercial ships, as well as in various stationary electric power generation applications. Artificial neural networks, notably Probabilistic Neural Networks, form the basis of intelligent agent reasoning by encoding expert diagnostic knowledge into concise software modules for real-time, onboard implementation on COTS computer systems.

The primary challenge of "Operations-Focused Maintenance" is establishing the management processes to identify and complete necessary maintenance activities while simultaneously meeting surge and other unpredictable operational demands. By automating the essential analytical work necessary to implement an effective CBM program, software agents can support the urgent need to keep ship equipment running reliably in the context of reduced shipboard maintenance capabilities and strengthened distance support.

Major *maintenance mega-trends* taking place, such as increasing ship systems complexity and a decreasing skilled maintenance labor pool, will continue into the foreseeable future. Software agents can serve as expert assistants in equipment health monitoring, no matter how complex the machinery processes. Their intelligent processing functions will allow ship maintenance managers to leverage valuable diagnostic knowledge across a ship fleet to enhance readiness while proactively minimizing maintenance costs. The agents can be distributed when and where needed to enhance operational performance. They represent valuable *maintenance workforce productivity multipliers*, as sailor-agent teams can achieve higher levels of ship readiness at reduced maintenance costs. An independent business case analysis estimated their return-on-investment of over 13-to-1 in reduced LSD maintenance cost.

The levels of machinery automation being installed during vessel upgrades, as well as that planned for future ship designs, will continue to increase, providing massive amounts of data for equipment health monitoring. Sailors will be hard pressed to transform voluminous amounts of raw data into useful information without the help of software agents similar to those described in this paper. Software agent technology represents a potential solution for building truly intelligent ships of the future that can automatically assess and report their equipment health.

REFERENCES

"Interview with RADM Jeff Brooks, Director of Maintenance for US Fleet Forces Command", ASNE Naval Engineers Journal: Special Issue on Fleet Maintenance, Vol. 120, No. 2, 2008.

Clifford, R.S., Callendar, J., and O'Meara, N., "Maintenance Workforce Characteristics", Logistics Management Institute Report No. LG101T3, Logistics Management Institute, McLean, VA, 2003.

Dean, A.W., Reina, J.J., and Bao, H.P., "Identification of Supplementary Metrics to Sustain Fleet Readiness from a Maintenance Perspective", ASNE Naval Engineers Journal: Special Issue on Fleet Maintenance, Vol. 120, No. 2, 2008.

Drew, K., and Scheidt, D., "Distributed Machine Intelligence for Automated Survivability", Engineering the Total Ship 2004 Proceedings, ASNE, Gaithersburg, MD, March 2004.

Logan, K.P., **"**Intelligent Diagnostic Requirements of Future All-Electric Ship Integrated Power System", IEEE Trans. On Industry Applications, Vol. 43, Issue No. 1, Jan-Feb 2007.

Logan, K.P., "Reducing Maintenance Costs through Agent-based Predictive Analytics", Intelligent Ships Symposium 2007, ASNE.

Logan, K.P., "Prognostic Software Agents for Machinery Health Monitoring", 2003 IEEE Aerospace Conference, Vol. 7, Big Sky, Montana, March 2003.

Murthy, V.K., "Nonparametric Estimation of Multivariate Densities with Applications", Multivariate Analysis, Krishnaiah (Ed.), Academic Press, New York, 1966.

LSD 41/49 Classes Main Propulsion Diesel Engine (MPDE) Remote Machinery Monitoring System Capability Assessment BCA Study Rev 3, NAVSEA PMS470, July 2007.

Economic Report of the President, U.S. Office of Management and Budget, 2002.

Parzan, E., "On Estimation of a Probability Density Function and Mode", Annals of Mathematical Statistics, 33, 1962.

Rumelhart, D.E., Hinton, G.E., and Williams, R.J. Learning internal representations by error propagation. *Parallel Distributed Processing – Explorations in the Microstructure of Cognition*, Vol. 1, MIT Press, 1986.

Yardley, R.J., Raman, R., Riposo, J., Chiesa, J., and Schank, J.F., "Impact of the Fleet Response Plan on Surface Combatant Maintenance", The RAND Corp., Santa Monica, CA, 2006.

ACKNOWLEDGMENTS

The authors gratefully acknowledge PMS 470 (NAVSEA 21) for their continuing support of the work described in this paper. In addition, the contributions of Mr. John L. Pelosi of the Machinery Enterprise Information Systems Branch, NAVSEA-Philadelphia are sincerely appreciated.

Kevin Logan *is President/Founder of MACSEA Ltd., a 27-year old company specializing in machinery diagnostic and predictive analytic systems. Mr. Logan holds a BA in Mathematics, attended graduate studies in Operations Research, and has performed decades of applied research in vessel/machinery performance analysis, machinery diagnostics, and software agent technology. His most recent research includes intelligent software agents, artificial neural networks for automated shipboard learning systems, and automated machinery model synthesis in 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.

Capt. John Walker *recently founded Thor Solutions LLC, a Service-Disabled Veteran-Owned Small Business specializing in engineering services. Capt. Walker's 26 years of US Naval service encompassed a variety of sea and shore billets including command of two Navy ships and as Major Program Manager for Mine, Auxiliary, Amphibious, and Command Ships (PMS 470). Capt Walker holds a BA in Management Information Systems from the University of Oklahoma, and an MMAS in Middle Eastern Studies from the U.S. Army Command and General Staff College.*