SPE 112051

Subsea Condition Monitoring: A Path to Increased Availability and Increased Recovery

John D. Friedemann, Anil Varma, Piero Bonissone, and Naresh Iyer, General Electric

This paper was prepared for presentation at the 2008 SPE Intelligent Energy Conference and Exhibition held in Amsterdam, The Netherlands, 25–27 February 2008.

Abstract

The application of subsea technology to reservoir development has been accepted as a valid production solution despite the known effects on facility availability and on ultimate recovery. The primary causes of the reduced availability for subsea tree-based solutions are directly related to the differences between the facility designs. A topsides drilling unit provides simple and direct access to the well for purposes of well logging and other data acquisition, and well interventions to replace equipment.

Other benefits of surface tree solutions are related to ease of well testing and chemical treatments. These in turn simplify the hardware maintenance because the facilities can be inexpensively directly accessed. In essence, the surface access solution gives increased well recovery because maintenance and diagnostic investigations are based on existing infrastructure (a permanently mounted drilling rig).

The subsea industry shares a common feature with other, more traditional, industries like Energy and Rail: the occurrence of unplanned outages results in increased costs and reduced revenues in each of these industries; a condition-based planning approach is the more desirable solution. Rail and Energy industries have successfully evolved to solutions based on the use of contractual frameworks aimed at guaranteeing equipment availability and efficiency.

This paper will examine the applicability of condition-monitoring technologies to subsea infrastructure in the oil production industry. The paper will present relevant case studies from the Energy and Rail transport industries. It will describe data and information handling which are required to enable the full development of a qualified predictive maintenance tool for the oil industry, with the view of taking the oilfield towards a paradigm of a contractually managed asset with the goals of improving service life and availability.

Introduction

The Subsea Domain

From a modern process control viewpoint, the subsea production system has typically consisted of a relatively simple network of piping and instrumentation designed to gather information from individual wells. Over time, this has expanded in complexity as systems began to be more complex. This can be seen in the comparison between Figures 1 and 2. Figure 1 represents a typical subsea manifold as constructed in the early nineties. These often involved relatively simple trees configured with one or two pressure and temperature sensors intended for but useless ass reservoir condition monitoring devices.

These early subsea facilities were typical of the period. Surface facilities were constructed for subsurface monitoring via direct intervention. At the time, typical topsides facilities have

- Drilling rigs for wireline logging and downhole sampling
- Test separators for rate determination, sampling and well diagnostics
- Convenient process access for observing process behavior and diagnostics

The control and monitoring of a topside system was therefore heavily dependent upon direct human access using relatively simple observation means. This became a problem as subsea systems became more complex. Access for legally mandated reservoir surveillance became difficult and a challenge arose. Ultimate recoveries for subsea field developments were found to be significantly lower; not because the production system was in anyway inefficient, but because the topsides based surveillance methods were, themselves, outdated.

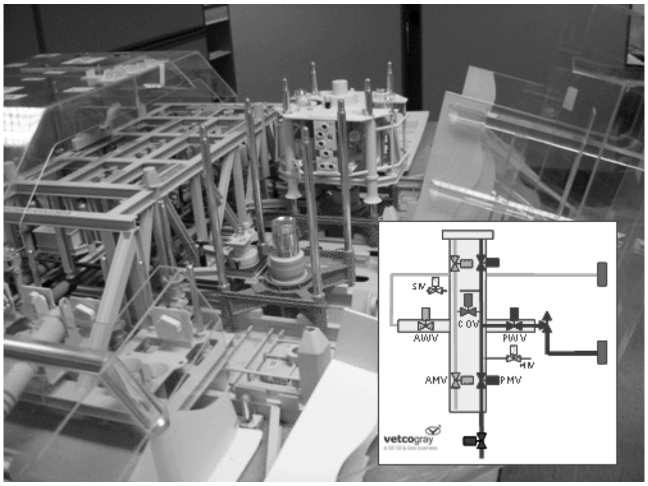

Figure 1: Early, simple subsea manifold.

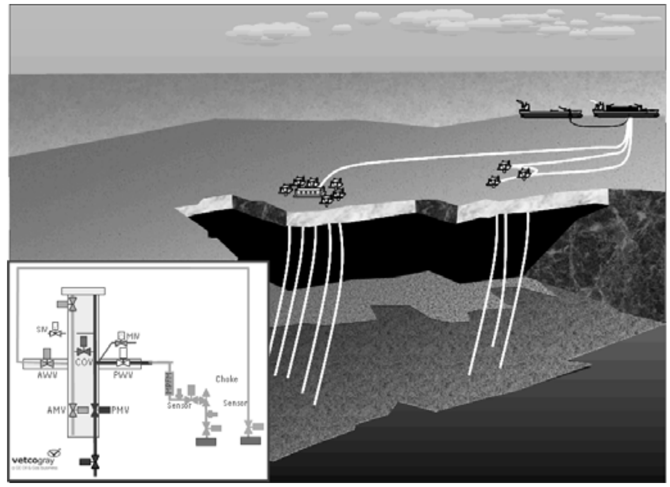

Figure 2: Typical subsea tree instrumentation of a modern subsea field

A simple comparison between figures 1 and 2 shows the level of instrumentation being utilized for condition monitoring of the field itself. Unfortunately the equipment itself needs equipment monitoring. The type of information being collected is illustrated in Table 1, which also includes an indication of the suitability of the information gathering technique for automation (0 being unsuitable for automation and 3 being already partially automated). One outstanding issue is the fact that systems for collating and interpreting all of these data are not well defined despite the fact that they can be quite time consuming to execute.

Table 1: Overview of traditional well-related condition monitoring activities.

Condition Monitoring Need		**Suitability for Automation**	**Condition Monitoring Need**
Template/Manifolds	Erosion / Corrosion	3	Template/Manifolds
	Gas Lift Control	2.5	
Well Flow Behavior	Test Separator	1	Well Flow Behavior
	Multiphase Flowmeter	2	
	Virtual Flowmeter	3	
Choke Behavior	Test Separator Interpretation	0	Choke Behavior
	Flowmeter Interpretation	2	
	Virtual Flowmeter	1	
Tubing Behavior	Wireline/Coiled Tubing/Tractor Surveys	0	Tubing Behavior
	SCRAMS Monitoring	2	
	Virtual Flowmeter	2	

Current market challenges

The installation of subsea processing stations poses a significant challenge for reservoir monitoring and control. In essence the facility functions as a data filter. Data acquisition will have to include facilities for online and real-time remote monitoring. The exact same phenomenon is experienced with the installation of downhole pumps and even subsea pumps.

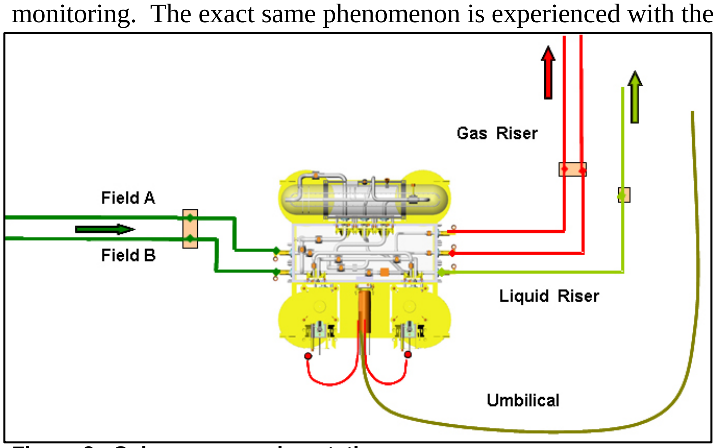

Figure 3: Subsea processing station

In the case of a separation facility like the one illustrated in figure 3, the complexity of the installation poses an issue in that the reliability of the installation is dependent upon the interplay of:

- Well-to-well fluid variations and potential effects on separation and other fluid chemistry related phenomena
- Methodologies for interpreting injector to producer interplay and communication
- Well-to-well sand production affecting chokes, flowmeters, and pumps
- Rate stability influencing the separator level control
- Separator efficiency causing pump and riser flow instabilities
- Pump design and robustness

All of this cause a challenge for interpretation when the interpretation must be made at a system level and effects must be inferred from available data. The challenge is to identify technologies to ensure the income from the customer's investment.

How and where can condition-monitoring technologies help?

The day-to-day job of the reservoir production engineer is, in fact, condition monitoring applied in a remote monitoring environment. Because of the limitations in the available data acquisition technologies, the application of model-based interpretation techniques is predominant. This works quite well for the simpler activities such as:

- In-flow diagnostics (water, gas breakthrough, loss of productivity, pressure decline, etc)
- Equipment monitoring (change in choke capacity, sensor failures, etc)
- Flow instability
- Other behaviors with known cause-effect relationships

This provides a basis for the application of simple principles of condition monitoring. But this would ignore the additional potential within condition monitoring technologies. The tools employ statistical analysis methods, which enable the observation of indirect causes-effect relationships such as:

- Producer-injector relationships by statistical analysis and artificial intelligence
- Failure identification and cause interpretation by reviewing historical data and comparison
- Maintenance identification and prediction based on failure data sets

In other words, condition monitoring serves as means to the ultimate end of improving system-level metrics like reduced costs, increased availability, improved recovery, planned maintenance, better efficiency and others. A broader framework is necessary that will synthesize information produced using condition-monitoring to further enable optimal decision-making across the lifecycle and management of the entire system, including operations and maintenance. The suite of technologies that fall under this framework are commonly recognized as Prognostics and Health Management or PHM. We identify the primary elements of this suite in the next section and present two case studies of real-world applications of PHM with the goal of improving system-level metrics.

Prognostics and Health Management

The Subsea industry shares a common feature with other, more traditional, industries like Energy and Rail: the occurrence of unplanned outages results in increased costs and reduced revenues in each of these industries. Also, factors like shift and variability in operating condition, operational modes and infrastructure-aging lead to newer modes of failure, thereby contributing to increased variability in how assets fail, or how long they last. Maintaining high reliability amidst such challenges often requires very conservative maintenance and operational policies, which come at a high cost and support burden. Condition-based strategies for maintenance and operations attempt to address and manage this burden by driving policies based on accurate and personalized assessment of equipment health.

Successful execution of condition-based strategies requires a careful tradeoff between proactive and reactive actions. Reactive strategies apply when there is confidence that a failure has occurred, but often at a cost of operational disruption, secondary damage costs and potential safety incidents. Proactive strategies can avoid unanticipated downtimes, but with the potential of unneeded and premature actions in maintenance and operations. We propose that the accurate and timely information about the health of the asset components is the best foundation for informed and condition-based decision-making.

Generally speaking, PHM technologies also comprise of hardware solutions, including various kinds of sensors to capture state information of the various components, communication devices that allow these sensors to transmit the state information to condition-monitoring stations as well as basic IT-infrastructure consisting of computing, control and visualization consoles. However, the primary focus of our paper will be software PHM techniques that enable intelligent reasoning by analyses of operating data.

PHM is primarily composed of two suites of technologies: health assessment and health management. These two suites and the manner in which they interact are illustrated more clearly in Figure 4. The modules within health assessment are shown in the left dotted box in Figure 4 while those within health management are enclosed in the right dotted box.

There are many ways to distinguish between these two sets of modules: for example, health assessment modules perform analysis at a much smaller, unit level (subsystems, components), while health management modules consider information across a fleet of assets or collection of components seen as an integrated, interacting system of systems. Health assessment involves classical bottom-up analysis of a device whereby it is assumed that the state of the device can be explained or inferred in terms of the states of its components. Health management, on the other hand, is more top-down synthesis of the information produced by the health assessment pieces to come up with strategic policies by means of which the entire logistics of sustaining the system can be managed optimally. A health assessment module could indicate the remaining useful life for a certain system component while a health management module takes a input the RUL estimates of many components of the system and recommends strategies by which the overall system performance in terms of operational yield, efficiency, repair scheduling, spare planning, revenues and costs can be optimized. Such strategies involve taking appropriate actions and making decisions in the domains of operations, maintenance, supply-chain etc. in the presence of constraints acting from outside the system. Thus, typical capabilities in health assessment library include diagnostics, prognostics, and anomaly detection while those in health management include decision-making, multi-objective optimization and scheduling. The circular arrows in the middle of the figure indicate the dependence between the two: information produced by the health assessment modules is used by the health management modules for synthesis; conversely, actions prescribed by the health management modules will impact health assessment and is information that the modules must use.

Looking into the boxes in Figure 4, health assessment proceeds by analysis of operational data collected for the system by multiple sensors monitoring various state variables of the system. Using field-deployed, remote sensors, the operational data is collected and preprocessed (segmented, filtered, validated, etc.). Then it is summarized by a set of significant features that provide a more compact, yet informative representation of the data. Anomaly-detection modules further analyze the features to identify presence of operational anomalies. If any anomaly is detected, the module additionally determines the time when the anomaly is first noticed and the possible source of anomaly (usually a coarse identification at the systems/subsystem level). This information allows a diagnostic module to focus on a given subsystem, analyze key variables associated with such subsystem, and look for the presence of a failure, possibly by using a pre-existing library of signatures associated with failure modes. The result of diagnosis is a ranked list of possible failure modes present in the system at the time of examination. A prognostics module will update a deterioration index for the platform (sub-) system, and compute the expected Remaining Useful Life (RUL) using an appropriate wear trajectory. The failure time and mode determine the inflection point in such curve and the steepness in deterioration, respectively. A prerequisite to leverage this RUL estimation is to have a tight confidence interval, such that this information is actionable and can be used for health management.

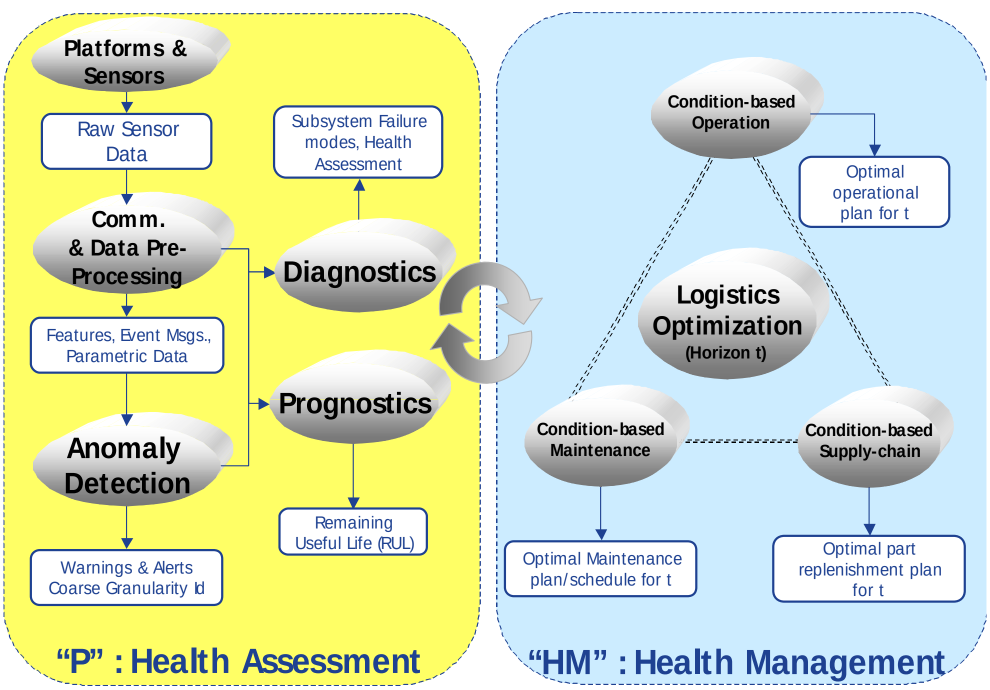

Figure 4: PHM as composed of health assessment ("P) and health management ("HM")

Health Management involves synthesis of the information created by the health assessment modules to explore, evaluate and propose appropriate responsive and strategic actions and decisions along operational, maintenance, supply-chain and other such logistics related areas. For example, information about remaining useful life for various operation-critical components along with knowledge regarding current operational requirements and constraints on the maintenance and supply-chain results in a plan optimization problem, where appropriate actions in operations, maintenance and supply-chain have to explored to find the best plan that optimizes overall system performance and meets requirements and constraints. Other logistics problems like repair scheduling, operational reconfiguration, spare allocation, customer notification also come under the umbrella of health management.

In summary, health assessment provides the key knowledge required to improve system performance and metrics that are typically crucial from the customer's perspective. However, once such knowledge has been generated, the problem of applying this knowledge to the devising of appropriate system-level strategies is a non-trivial one and requires the application of health management technologies like optimization, scheduling and decision support.

Case Studies from Rail and Energy Industries

The subsea industry shares a common feature with other, more traditional, industries like Energy and Rail: the occurrence of unplanned outages results in increased costs and reduced revenues in each of these industries; a condition-based planning approach is the more desirable solution. Rail and Energy industries have successfully have successfully made use of PHM technologies and evolved to solutions based on the use of contractual frameworks aimed at guaranteeing equipment availability and efficiency. In this section, we briefly describe successful case studies demonstrating the application PHM technologies to the Rail industry and the Energy industry with the overall goal of improving customer-centric metrics like increased availability and efficiency. For Rail locomotives, we describe a deployed PHM system Expert-On-Alert™ (EOA™) that is used to provide customer-centric services for management of a fleet of locomotives; EOA™ currently monitors over 8000 locomotives. From the energy domain, we present a successful case study of a real-world deployed product, KN3, that demonstrates the use of health management techniques to optimize the operation of a coal-fired boiler so as to simultaneously minimize emissions and maximize efficiency for a given load demand.

Expert-On-Alert™: A RM&D system for Rail Locomotives

GE Transportation, headquartered in Erie, PA initiated a remote monitoring and diagnostics program in 1997 to proactively monitor its customers' locomotives. The service is called EOA (Expert on Alert). Following the general trend in the equipment manufacturing industry, GE has been entering long time service contracts with purchasers of rail locomotives instead of selling individual spares. Locomotives are complex electromechanical systems equipped with the capability to monitor their state and generate fault messages. A typical GE rail locomotive has around 200,000 parts with more than 24 microprocessor controllers. These locomotives typically operate in extreme environmental conditions, logging over 100000 miles per year on average. On average, each locomotive has 3-4 scheduled shop-visits per year and about 4-5 un-scheduled shop-visits in addition. Since diagnostics and repairs are time consuming and complex tasks, it is essential that the number of shop visits and the time spent on the shop floor by a locomotive is minimized for a long-term service contract to be profitable. Therefore, making the service contracts profitable requires that the deployed Remote Monitoring and Diagnostics (RM&D) systems identify problems occurring on the equipment while in operation so a) the repair can be scheduled best keeping with the severity of the problem and b) the complete set of problems is identified so the time in the repair shop is utilized at not merely fixing one problem but releasing an overall healthy machine

These RMD systems are based on the information that is produced from the on-board control systems. The control-system generated symptom data is useful for field technicians in order to detect, diagnose, and fix equipment problems. Locomotive fault logs are accumulated on-board the locomotive and are periodically uploaded to a database for access in case a diagnostic need arises. Highly skilled field engineers at General Electric Transportation Systems have acquired expert knowledge over time that enables accurate diagnosis of locomotive problems from an examination of the fault log. While this provides positive evidence for the diagnostic significance of fault logs, the volume of logged data makes it impossible to rely on human examination alone for reliable and consistent identification of locomotive problems on many hundred locomotives on a daily basis.

The aim of the EOA service is to automate a significant portion of the diagnosis process to improve productivity and accuracy, and also move the service paradigm from being reactive to being proactive. The complete EOA system consists of several components: a RM&D center (which operates out of the company's Erie, Pa., headquarters), a locomotive on-board data collection and multi-mode communication module, a network backbone for communication, and the IT infrastructure to host different applications.

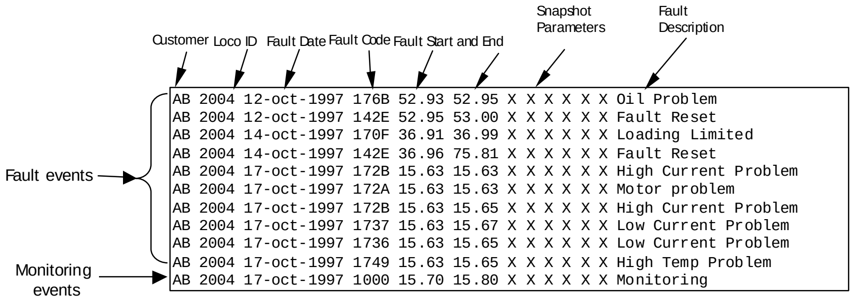

Figure 5: Fault Log Format

Operational data and fault logs are collected while the locomotive is on track in service. Figure 1 shows an example of fault log and operational data collected from the locomotive. The operational data contain various sensor readings and states of the locomotive. The data are then sent to the RM&D center via satellite or other communication modes depending on the coverage. The data are stored in the RM&D database and automatically processed and analyzed in the backend with the different application modules in the system. Based on the knowledge learned or encoded in the system, the system automatically creates cases when there are some problems with the locomotive based on the learned knowledge. Most of the time, the data indicates a healthy locomotive. In this situation, there are no cases created. This process is automatically done by the system without the diagnosis engineering involvement in sifting though the huge amount of fault log in order to determine whether a locomotive has any problems. When there are issues with a locomotive, alerts and diagnosis outputs are generated and presented to the RM&D engineers. The RM&D engineers review these alerts and diagnosis outputs generated from the system, and a final notification is sent to the field engineers and customer.

Description of Application

Figure 6 shows an overview of the EOA system. There are several modules in the system. At the top left there is the data collection and communication infrastructure. This data is fed to the reasoning engines. The recommendation generated by the engines is sent to the case management and presentation interface. The recommendation is delivered to the GE engineers and customer. Finally, feedback is collected and knowledge in the reasoning engines is updated. These modules are each described in this section. The whole system is a web-based system such that users with the right privilege can log in the system from any computer within the firewall.

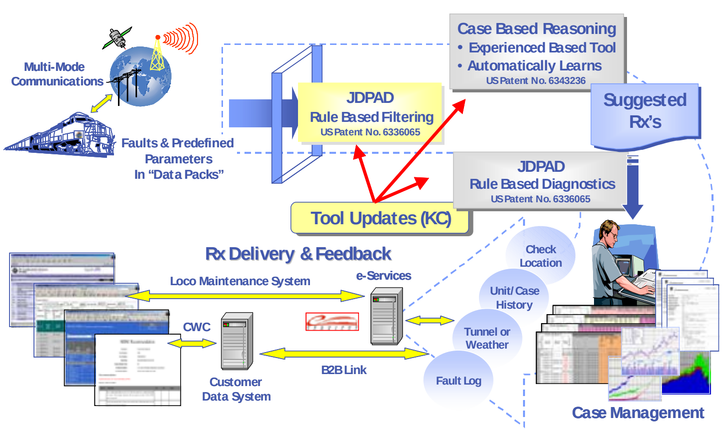

Figure 6: System Overview

Data collection and communication: This is application is hosted in the LOCOCOMM™, GE's integrated on-board computer and communication management system that serves as the basis for GE's current information-based services and systems. LOCOCOMM™ has extensive data acquisition capability and hosts a variety of different GE or third-party applications on its industrial standard Wintel platform. The data collection and communication module for RM&D is also hosted here along with other applications for asset tracking (PinPoint™), fuel management (Smart Fueling™), video and audio recording (LocoCAM™).

Reasoning engine: Figure 7 shows the functional modules of the main engine, which comprises data filtering, Case-Based Reasoning (CBR) engine, and Java Data Pack Anomaly Detection (JDPAD) rule engine. Once the data is collected to the RM&D data repository, a rule based filtering process is applied to filter out some nuisance faults that have no failure and diagnosis information. The filtered data are passed to two reasoners: the rule engine and the CBR engine. Both reasoners produce recommendations (Rx) independently. The system also tracks the performance the reasoners and outputs their prior performance for the reference of the remote diagnosis engineers. As we have discussed, the operational data stream flows to RM&D in real time, the reasoners are triggered periodically unless critical faults are found in the data log.

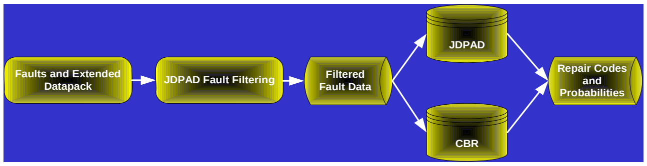

Figure 7: Functional Flow

Case management and presentation: The cases are automatically created from the reasoners along with an Rx. The new cases are entered into a workflow and are presented to the RM&D diagnosis engineers. Diagnosis engineers who logged into the system can pick a case from the workflow and work on the case. Priorities are assigned to each case based on urgency so that the urgent ones are competed first. Once an engineer starts to work on a case, the locomotive's basic information is presented, such as customer, model, configuration, recent history, and so on. The engineer can go deeper into the data and recommendations with the screen shown in figure 8.

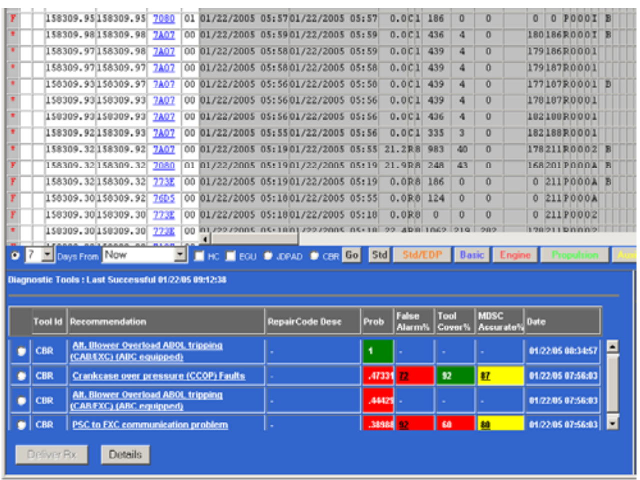

Figure 8: Data and recommendation screen for diagnosis engineers

Recommendation delivery and feedback: The engineers review the cases with Rx automatically generated from the system. A prior history of accuracy is checked to make a final decision. Other conditions are taken into consideration, such as weather and location, in order to come to a very accurate diagnosis. Once the diagnosis engineers confirm that there are problems with the locomotive, recommendations are sent to both railroad customers and service shops. In this way, the customers may act upon the recommendations based on their criticality to determine when to pull the locomotive off track to a shop or perform a run-through maintenance if the shop visit is not necessary. This information is also critical for the customers to proactively adjust their planning and optimize asset utilization. On the other hand, the shop also knows what are the problems for an incoming locomotive, which can greatly reduce time to repair, parts availability and work order planning.

Knowledge update: A key component for an intelligent system is continuous learning. In EOA, the system knowledge is continually updated based on feedback. The CBR engine is updated automatically with new golden cases (Cases with verified feedback) are put into the case base. When feedback is received from customer and service shop, high quality cases are marked and save into the case base. The JDPAD rule engine is updated periodically through a rigorous process and tracking of performance metrics. Within the ICARUS system, there is a web-based rule management and editing module that helps knowledge engineers to easily perform any update of the rules.

EOA's success really highlights the fact that intelligent use of analytics and data can reveal a significant amount of information about the condition of a complex system, without the upfront need to put dedicated PHM sensors in place.

KN3: A System for Operational Optimization of Coal-fired Boilers

Problem Description and Solution Architecture

Kn3 is a recently unveiled software platform that exemplifies GE's Ecomagination efforts. With this platform we can predict the performance of a power plant, which can be used to simultaneously reduce emissions, improve efficiency and increase availability. We approach plant power management as a Multi-Criteria Decision Making problem.

We use an ensemble of neural networks (NN) as the predictive model to map a vector of operational set points to their expected effects in terms of heat rate – inverse proportional to efficiency – and NOx emissions. Then, we rely on multi-objective evolutionary algorithms (MOEA) to generate the Pareto set for the plant at a given load. This set contains all non-dominated set points in HR. Finally, we use an automated decision-making that selects the point on the Pareto that best represent the desired tradeoff between fuel costs and emission credits (Subbu et al. 2006, Subbu et al. 2007). The predictive models are adaptive, and continually update themselves to reflect with high fidelity the gradually changing underlying system dynamics. The integrated approach, embedded in a real-time plant optimization and control software environment has been deployed to dynamically optimize emissions and efficiency while simultaneously meeting load demands and other operational constraints in a complex real-world power plant. The architecture for this process is shown in Figure 9.

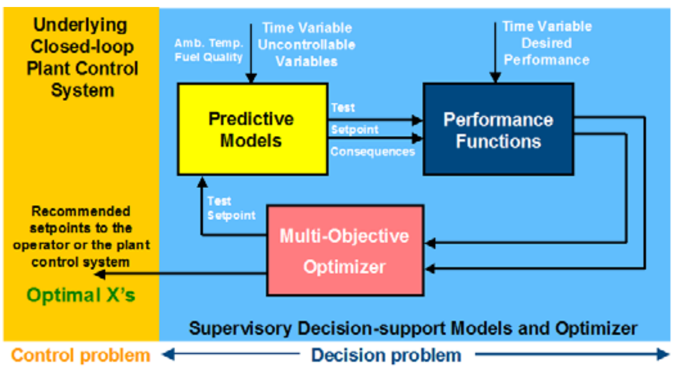

Figure 9: Architecture of model-predictive multi-objective optimization

Nonlinear neural-networks are used to represent mappings between the inputs space of control variables (X's) and time variable ambient uncontrollable variables and the various outputs (objectives and constraints) of interest. The evolutionary multi-objective optimizer generates test inputs and receives as feedback the corresponding output performance metrics after transformation by suitable objective (performance) functions. The multi-objective optimizer uses this feedback to generate and identify the Pareto-optimal set of input-output vector-tuples that satisfy operational constraints. A decision function is superimposed on this Pareto-optimal set of input-output vector-tuples to identify a deployable input-output vector, which is then dispatched to the underlying plant control system, or recommended to the operator for execution.

Key challenges

Beside the obvious problem of dealing with the conflicting goals of reducing emissions while improving efficiency (higher temperature combustions are more efficient but generate larger amount of emissions), the most challenging aspect of this problem is the management of uncertainty embedded in the data and the models themselves.
A prerequisite for model-based optimization is to be able to use accurate and reliable models over the operational range of interest to evaluate potential solutions. Unfortunately models do not always provide such accurate evaluations.

Uncertainty derived by violating model assumptions. In the case of physics-based models, usually they are built using a set of (simplifying) assumptions, such as use of lump, time-invariant parameters, linearization around operational points, etc. The model's accuracy depends on the degree to which such assumptions are satisfied when the model is used to evaluate a point in the solution space. If the transients are too slow, we might need to use more complex models with distributed parameters and partial differential equations. If system degradation and wear are important, then we use parameters that vary over time. If the point under evaluation is far from any operational point, then we should employ a more accurate non-linear model structure. If all design assumptions are satisfied then the model will provide accurate evaluations (module measurement errors and system disturbances).

Uncertainty derived by using functional approximations. In the case of empirical, data-driven models (such as the NN's), and the model's accuracy depends on how close the point under evaluation is from the training data. Several researchers have studied the error bounds of NN (Barron 1994) and special types NN, such as Radial Basic Functions (Townsend and Tarassenko 1997). To reduce this kind of uncertainty, we use of a committee of predictive models (NN's) and an intelligent fusion of their predictions (Xue et al. 2006), which aggregates the results of multiple predictive models based on local accuracy measures of these models in the neighborhood of the probe point for which we want to make a prediction. Fusing the outputs from an ensemble of models in an effective way can often boost overall model accuracy. This fusion method may be applied to develop highly accurate predictive models. The locally weighted fusion method boosts the predictive performance by 20~40% over the baseline single model approach for the various prediction targets. In (Bonissone et al. 2008) we further refine the performance by using CART algorithms to pre-compile a segmentation of the input space for each model. Using this approach we improve the predictive performance by 34~48% over the same baseline. Relative to these approaches, fusion strategies that apply averaging or globally weighting only produce a 2~6% performance boost over the baseline. A prerequisite for a successful fusion is to create a strong diversity of the models to be fused (Kuncheva and Whitaker 2003).

Uncertainty derived by model extrapolation errors. Such predictive models can tolerate moderate extrapolations, i.e. they can stretch to make predictions in inputs spaces that are proximal to regions of training data. However, their extrapolative prediction becomes inaccurate as this extrapolation is extended to spaces that are far from the training data. A solution to this problem is to restrict the search to areas close to the regions of the available historical data. Over time, deploying new set points that are close to the available historical data will push the envelope of the historical data, and lead to improved, model-based prediction capabilities. An efficient method to enable such a restricted search is to scan the historical data for operating points that were deployed when the ambient conditions were "close" to the current operating conditions. For instance, if the current load demand is 350MW and the ambient temperature is 70° F, then it would be appropriate to scan the historical database for set points that were deployed when operating conditions were close (within specifiable bounds of the current load demand and current ambient temperature), and use these set points as seed points to initiate a restricted search.

A Pareto-optimal front that jointly minimizes NOx and Heat Rate (inversely related to efficiency) for a 400MW target load demand in a 400MW power plant is shown in Figure 6. In this figure, the circles show the range of historical operating points from a NOx—Heat Rate perspective. The stars and inter-connecting line show the optimized Pareto frontier in the NOx—Heat Rate space. Each point not on this frontier is a sub-optimal operating point—the goal being the operation of the plant or process at a Pareto optimal point at all times. The act of moving the system operation from the interior of the decision space to the Pareto frontier results in a large operational savings opportunity.

Experimental Results

The multi-objective optimizer in conjunction with the predictive models and the decision function solve a decision problem as a function of time. Control of the transition of the plant or process state to achieve the recommendation is delegated to the underlying plant control system. In a supervisory mode of deployment, a recommendation is transmitted to an expert human operator who then programs the recommendations in the plant control system, while in an automated mode of deployment; the recommendations are directly transmitted to the plant control system. Such use cases necessitate the use of automated down-selection to a solution from the Pareto frontier, for execution. This down- selection is part of the multi-objective decision-making step. The Pareto frontier in NOx—Heat Rate space identified from the multi-objective search is clipped by the systematic application of profit-based and operational-need constraints for each of NOx and Heat Rate. Next, a solution from this reduced frontier that is closest in inputs space to the current plant state is selected and transmitted to the plant control system. Such an approach minimizes the state deviations while achieving Pareto-optimal operation.

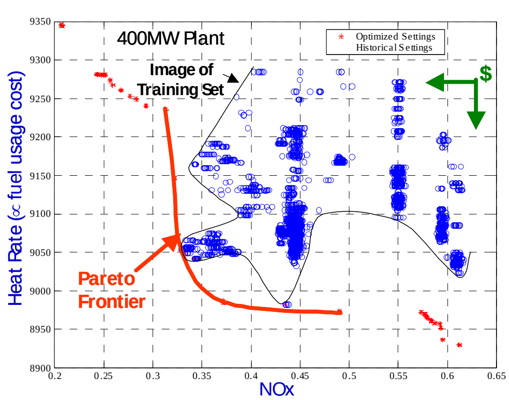

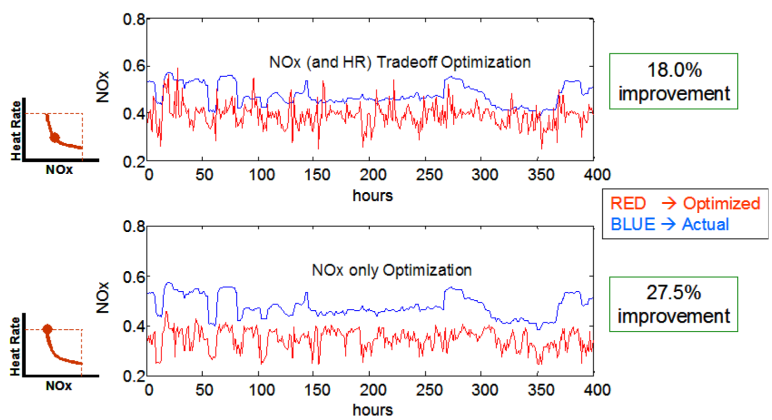

Figure 10 Plant operation represented in HR and NOx space. Figure 11 Tradeoffs in NOx and Heat Rate and NOx only optimization.

Figure 11 shows the performance gains achieved in NOx emissions using this decision-making approach. When a decision-making function is used which simultaneously considers a tradeoff Pareto point at each instant, roughly 18% improvement in NOx emissions may be achieved (upper figure half). However, if the optimization favors a NOx minimization that satisfies a given Heat Rate constraint, more significant NOx emissions improvement is possible (lower figure half). Similarly, 1-2%

improvements in Heat Rate are possible. In a typical power plant setting, such savings in NOx and Heat Rate are very significant and could lead to operational savings of hundreds of thousands to millions of dollars per year. The decision-making approach further highlights the inherent flexibility of Pareto frontier techniques whereby the entire efficient set of solutions is first identified without regard to situation specific down-selection, and later a flexible decision function is superimposed to identify a deployable input set (or set point).

This software platform has already been installed in five coal-fired power plants, resulting in significant operational and environmental savings. KN3 exemplifies is an example of evolutionary based optimization for power plants. For a more comprehensive description of industrial applications of evolutionary algorithms, the interested reader is referred to (Bonissone et al. 2006).

Generalizing lessons learned

Prognostics and Health Management requires both state awareness and the ability to adjust operations based on prognostic and diagnostic input. Expert on Alert has shown that a tight coupling between sensing degradation and proactive maintenance has increased utilization and decreased downtime. Road failure rates (unexpected downtime) have been reduced by over 50% and utilization of the fleet has increased by 2%. Both metrics have had tremendous financial impact in the operation of a rail fleet and network. Similarly, with the KN3 example, using plant data to optimize operations impacts utilization and minimizes emissions shows significant potential for efficiency of plant operation.

The key lessons from these case studies can be summarized as

a) New sensors are not always necessary for PHM. Significant impact can be achieved from existing data and a tight integration between sensing and health management.
b) Data-driven approaches can have an impact on complex domains since they can produce ‘reduced order’ causal models that are sufficient for making PHM decisions.
c) The investment into a remote monitoring and diagnostic solution and infrastructure creates the foundation for extensive data monitoring and collection. This usually has value and application far beyond the originally intended scope.
d) Learning and adapting the PHM logic with minimal human effort is an important capability for keeping the reasoning vital over long periods of time.

Instantiating the PHM Framework for Subsea

In deploying PHM solutions to take the entire subsea system towards a path of increased availability and recovery, one can view the oilfield as composed of multiple subsystems as shown in Figure 8. The multiple subsystems in this decomposition (*Downhole, Subsea, Surface, and Field*) imply varying requirements on the individual subsystems to improve overall system performance. For example, one path to *improving recovery* of the entire system can be tied to *better data collection* downhole, *erosion detection* at the subsea level, *improved debottlenecking* at surface level and *computing optimized recovery strategies* at the level of the field. Such a hierarchical view of the field also permits a gradual deployment and maturation of tactical capabilities at the subsystem level (“P”), while creating the information needed to eventually evolve towards providing strategic solutions that optimize metrics at the system-level (“HM”).

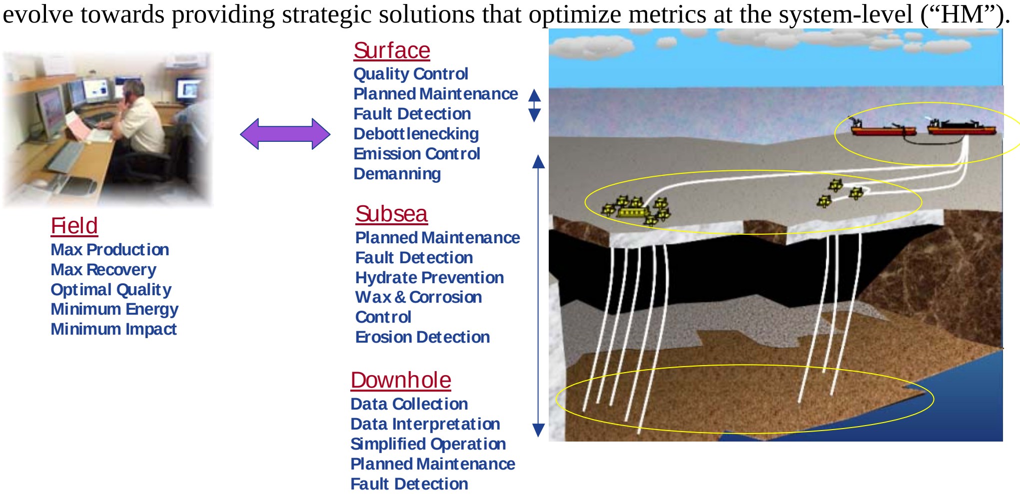

Figure 8: Elements of Integrated Subsea Production

Summary

The Subsea industry shares a common feature with other, more traditional, industries like Energy and Rail: the occurrence of unplanned outages results in increased costs and reduced revenues in each of these industries. From a modern process control viewpoint, the subsea production system has typically consisted of a relatively simple network of piping and instrumentation designed to gather information from individual wells. Over time, this has expanded in complexity as systems began to be more complex. In addition, factors like limited access to legally mandated reservoir surveillance, outdated surveillance methods pose significant challenges for interpretation when the interpretation must be made at a system level and effects must be inferred from available data. Also, factors like shift and variability in operating condition, operational modes and infrastructure-aging lead to newer modes of failure, thereby contributing to increased variability in how assets fail, or how long they last. Maintaining high reliability amidst such challenges often requires very conservative maintenance and operational policies, which come at a high cost and support burden. Condition-based strategies for maintenance and operations attempt to address and manage this burden by driving policies based on accurate and personalized assessment of equipment health. Condition-monitoring technologies generate information that can be synthesized using PHM technologies that further enable the generation, evaluation and optimization of condition-based strategies. The full development of such a capability for the oil industry will allow taking the oilfield towards a paradigm of a contractually managed asset with the goals of improving service life, availability and other system-level metrics.

We described two case-studies as examples of real-world deployed systems showing the application of PHM technologies for system-level management and optimization. The Expert-On-Alert case-study describes a deployed system that is currently used to monitor the condition of over 8000 Rail locomotives to assess the health of their subsystems in trying to alert the customer of potential events that can cause downtimes, thereby allowing planning of maintenance activities. The KN3 system demonstrates the use of PHM technologies to manage and optimize multiple system-level metrics like efficiency and emissions for a coal-fired boiler, while trying to meet the desired operational demand.

We presented a framework for PHM as a suite of technologies that perform two primary functions: health assessment and health management. In trying to apply the PHM framework to the oilfield, we presented a view of the oilfield that decomposes it into multiple subsystems that have varying requirements from the perspective of improving system-level metrics. We tied these requirements to appropriate PHM components from the framework using a matrix. This matrix outlines a path to develop and deploy tactical, subsystem-level solutions as well as strategic system-level solutions, either of which can be offered using contractual frameworks depending upon the planning horizon of interest to the customer. It also illustrates a gradual deployment and maturation of tactical capabilities at the subsystem level, while creating the information needed to eventually evolve towards providing strategic solutions that optimize metrics at the system-level.

References

1 Barron A. 1994. “Approximation and estimation bounds for artificial neural networks”, *Machine Learning* 14(1):115-133.

2 Bonissone P., Subbu R., Eklund N., Kiehl T. 2006. “Evolutionary Algorithms + Domain Knowledge = Real-world Evolutionary Computation,” *IEEE Transactions on Evolutionary Computation*, 10(3): 256-280

3 Subbu R., Bonissone P., Eklund N., Yan W., Iyer N., Xue F., Shah R. 2006. “Management of Complex Dynamic Systems based on Model-Predictive Multi-objective Optimization”, *IEEE CIMSA 2006*, La Coruña, Spain, July 12-14

4 Bonissone P., Xue F., Subbu R., 2008. “Fast Meta-models for Local Fusion of Multiple Predictive Models” to appear in *Applied Soft Computing Journal*.

5 Kuncheva L., Whitaker C.J. 2003. Measures of Diversity in Classifier Ensembles. *Machine Learning*, 51:181–207,

6 Subbu R., Bonissone P., Bollapragada S., Chalermkraivuth K., Eklund N., Iyer N., Shah R., Xue F., Yan W. 2007. “A review of two industrial deployments of multi-criteria decision-making systems at General Electric”, *1st IEEE Symposium on Computational Intelligence in Multi-Criteria Decision-Making (MCDM 2007)*, pp. 136-145, Honolulu, HI, USA, April 2,

7 Townsend N.W. and Tarassenko L. 1997. “Estimations of Error Bounds For RBF Networks”, *Proc. Artificial Neural Networks*, 7-9 July 1997, Conf. Publication No. 440, IEE.

8 Xue F., Subbu R., Bonissone P. 2006. “Locally Weighted Fusion of Multiple Predictive Models”, *IEEE International Joint Conference on Neural Networks (IJCNN’06)*, Vancouver, BC, Canada, July 16 –21