WHITE PAPER

Predict, Prescribe, Profit: Creating a World that Doesn't Break Down

Robert Golightly, Senior Manager, Product Marketing,
Aspen Technology, Inc.

aspentech

Creating a World That Doesn't Break Down

The objective of the management within any process industry is to minimize the stock level and increase the availability time, production and quality rate.

Reducing unplanned downtime and increasing asset utilization are the single biggest opportunities for financial improvement in production operations.

It's a $20 billion USD problem for the process industries. Millions are being spent on new maintenance solutions, but is that even addressing the problem?

> "If I had an hour to solve a problem and my life depended on it, I would use the first 55 minutes determining the proper question to ask, for once I know the proper question, I could solve the problem in less than five minutes."
> Albert Einstein

The Goal to Reduce Maintenance Costs

Over the past 50 years, maintenance practices have evolved in attempts to better serve the reliability and availability needs of manufacturing organizations. Strategies include:

- Run-to-failure: Maintenance only takes place when breakdowns occur.
- Calendar: Inspections and service occur on a fixed periodic basis in an attempt to be forewarned of potential issues.
- Usage-based: Attention is applied based on usage – as measured by run-time hours or cycles, for example – and is seen as a better indicator of potential failure.
- Condition-based: Routine inspections occur based on changes in monitored conditions. Some form of standardized monitoring, either automatically and continuously or manually at discrete intervals, aims to detect symptoms that herald harbingers of impending failures.
- Reliability-centered maintenance (RCM): An "umbrella" process with goals to establish operational risk, assess safe minimum levels of maintenance, and develop all the procedures and strategies to assure equipment performs "the way it should" without incident. It is often heralded as "the ubiquitous maintenance solution" for its ability to drive cost-effectiveness, increase machine uptime and improve management of the level of risk that the organization is undertaking.

A Flawed Strategy

Though often overlooked, there are two problems with all of these maintenance strategies. First, they are successively more complicated paths to the same outcome: the timing of when to inspect and service. Second, each one has a dominant focus on normal wear and tear as the root cause of failures. A failure mode effects analysis (FMEA) showed that approximately 80 percent of potential failures that could occur to the asset (the overall system) are not wear-based. That covers all degradation and failure in mechanical equipment that is process-induced; operating equipment outside of safety and design limits. Those circumstances include pump cavitation, liquid carry-over into compressors, dirty feedstock, pumps running dry, and so on. The implication is that most maintenance strategies are only looking for around 20 percent of the causes of failure.

In spite of these maintenance initiatives, equipment continues to break down, causing process interruptions, crises and sometimes loss of life. Effective, earlier detection of issues would eliminate those failures. This is where big data emerges to solve reliability problems that only a few years ago were unsolvable.

The Dirty Little Secret of RCM

The current hype is about reliability-centered maintenance (RCM). RCM differs in that it is more of a framework than a specific maintenance strategy, allowing the use of any of the existing maintenance strategies when appropriate. The upfront work required, however, can be daunting.

RCM relies on risk assessment and ranking for each and every piece of equipment, determined by the reliability probability number (RPN). The RPN proposes to offer the priority order for managing equipment. However, RPN is a simplistic calculation – the multiplication of three ordinals ascribed to represent the following:

- Severity – a rating of the severity and potential effect failure
- Occurrence – a rating of the likelihood the failure will occur
- Detection – a rating of the likelihood the problem can be detected before it reaches the end user/ customer.

Each of these ratings is determined by the subjective opinion of the participants. They are multiplied (not added) as equally weighted (should they be?) and therefore statistically guarantee a wildly skewed distribution1. After all of that work, you are left with an assessment that was true at that single point in time. There is no process to determine priorities on an ongoing basis.

[1] See Donald J Wheeler's article, Problems with Risk Priority Numbers: Avoiding more numerical jabberwocky at **http://www.qualitydigest.com/inside/quality-insider-column/problems-risk-priority-numbers.ht**ml

Preventative Maintenance for Process-Induced Downtime?

Traditional condition-monitoring technologies (originally called machinery monitoring) essentially examined the machine through specific machine sensors such as vibration probes, accelerometers, speed indicators, current draw, and so on. These are typically excellent at late stage detection of issues, preventing catastrophic, dangerous breakdowns and safety events. However, they usually only examine the machine by detecting normal wear and tear and, for the most part, ignore the manufacturing process in which the machine plays its role.

> "A useful prognostics solution is implemented when there is sound knowledge of the failure mechanisms that are likely to cause the degradations leading to eventual failures in the system."
>
> Peter Reynolds, Senior Analyst, ARC Advisory Group

Since most degradation and damage is borne of excursions outside of appropriate design and safe operating limits, new technologies are looking further afield than traditional reliability and maintenance products. The machine and the process are integral to each other: You cannot separate the machine from the process, and vice versa. We can let machine learning do the heavy lifting of determining the operating conditions and patterns that have a deleterious impact on the asset by capturing the patterns of process operation and merging them with failure information. The result is a comprehensive picture of the impact of the process on the asset.

An Ounce of Prevention...

The main problem is that we are not detecting problems before damage occurs. Both operations and maintenance departments agree that operating equipment outside of safety and design limits causes reliability issues. It's an insidious problem in that by the time we discover it, the damage is done.

Another facet of this relates to process conditions. Upstream conditions affect the monitored equipment. Machine learning can identify the changes in upstream process behavior that can have deleterious effects on downstream equipment integrity. Such damaging process patterns include liquid entrainment, low-temperature deviations causing gases to drop below a dew point before a compressor, cavitation in pumps, and so forth. Casting a wider "net" around equipment assures the capture of the damage-causing process-related issues that evade traditional maintenance techniques.

Boeing, a leader in the aircraft industry where RCM began, has stated that up to 85 percent of all equipment failures happen on a time-random basis, irrespective of inspection or service frequency. "Opening the box to see if there are filings in the oil pan," or waiting until the check engine light comes on, are not working.

The recent emergence of machine learning, big data, and analytics has created the opportunity to more precisely look at the data sets across process variables and asset health.

Adding voice to the argument for revisiting current practices, a large automation and reliability equipment vendor proclaims that 63 percent of all maintenance is unnecessary – and most cause more damage than if left alone. A major refiner validated this by reporting that one of its five tank gauging systems has far fewer faults and failures than the other four – the one where no planned maintenance takes place.

It's Not Just a Maintenance Problem

We are operating plants in ways not envisioned during the original design. We run closer to asset constraints, and we drive processes harder than ever. Profitability demands these processes operate as close to key limits as possible. Process excursions can quickly put an asset at an undesirable operating point where damage or excessive wear and tear to the asset occurs.

Understanding the impact on the asset and the process will mitigate maintenance decisions. To understand the impact requires a new generation of analytical capabilities which provide deeper insights into the asset, the process and the interaction between the two. Operators need predictive solutions that tell them of impending trouble, with prescriptive guidance in the software guiding them away from trouble. Those two factors will determine the potential pool of viable solution providers, as accomplishing this requires deep domain and process expertise, and the ability to reach into any and all design, production and maintenance systems for data.

> One [faulty] assumption associated with maintenance theory is that there is a fundamental cause-and-effect relationship between scheduled maintenance and operating reliability
>
> RCM RELIABILITY CENTERED MAINTENANCE GUIDE

How Do Things Break Down?

If maintenance alone doesn't provide a complete answer to the downtime and availability problems, what does? We need to look at this considering design, operations and maintenance. Today these three workflows (design, operate, maintain) are not understood from a cost trade-off perspective. These functions don't share knowledge and maintenance is not done based on need.

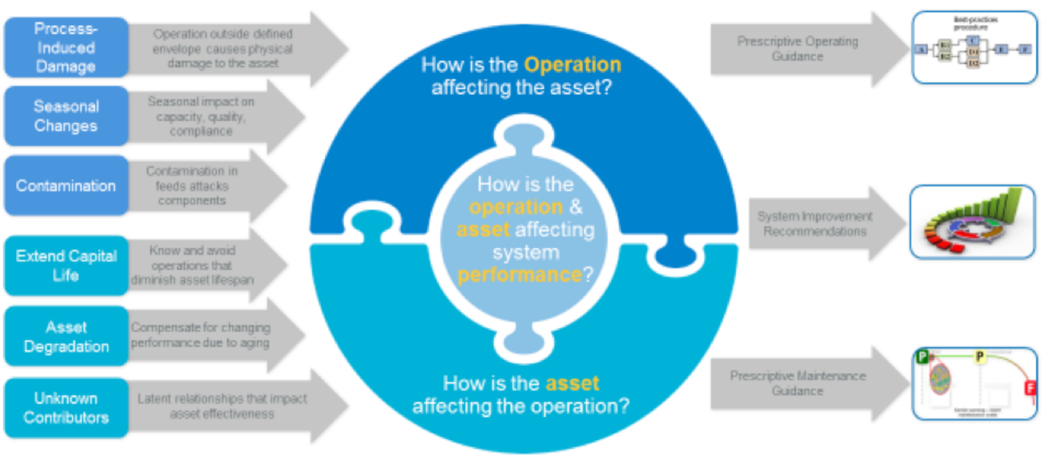

Figure 1: A New Approach to Asset Performance Management

Solving this will require better predictive and prescriptive process diagnostics to help operators respond quickly and effectively. We need to understand the reliability issues inherent in the assets, and we need predictive solutions that can prevent asset damage by guiding operators to keep the process away from areas that degrade the asset.

A New Approach to Asset Performance Management (APM)

Aspen Asset Performance Management (APM) combines 35 years of design experience and 25 years of process optimization experience with reliability analysis, machine learning and maintenance management to provide a comprehensive solution to improving asset effectiveness. Using asset-specific analytics, we can provide early detection and guidance to help operators keep the process away from conditions that are detrimental to the asset. Our reliability modeling provides the information needed to make economically optimal decisions and to focus on the most vulnerable assets or failure modes. Machine learning agents scour real-time data streams to predict situations before they cause losses or asset damage.

A recent publication by the ARC Research Group made the point that APM 2.0 incorporates new analytics and data from control systems and asset management applications which provide new opportunities to optimize availability and operational performance. In that document, they also provided a summary view of how companies are measuring the performance, which is depicted in Figure 2.

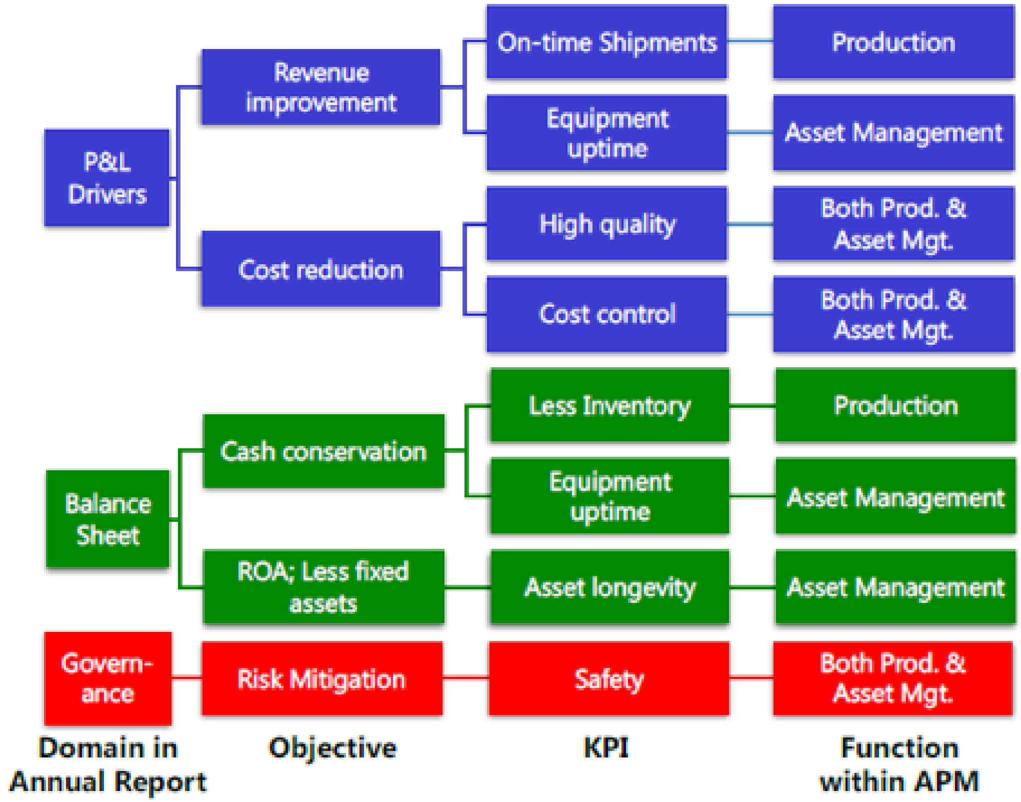

Figure 2: ARC Advisory Group: Key Performance Indicators

It's easy to see how intertwined these decisions are and the cross-functional nature of the decision making across production and asset management. The situation calls for a new level of collaboration between operations and maintenance. ARC articulated the point nicely in their white paper on APM 2.0.

APM 2.0

While the first generation of APM solutions provided value, they were focused exclusively on the maintenance side of the problem. As we have learned, since a significant amount of downtime is process-induced, new maintenance programs have not had the desired impact on asset performance. They lacked that consolidated view of the maintenance, asset and process perspectives.

APM 2.0 strategies include sharing of information from other systems – such as manufacturing execution systems (MES) and enterprise asset management (EAM) systems – to deliver a comprehensive view of production and asset performance. That amalgam of data supports an improved understanding of risks and the ability to balance operational constraints to improve return on assets (ROA).

APM 2.0 Collaboration

With a good APM strategy, operations and maintenance groups become more collaborative, exchanging information to manage critical issues and operational constraints while improving overall operating performance. Combining the information from the traditionally separate operations and maintenance solutions improves the effectiveness of both areas, and offers new opportunities for managing risk and optimizing performance.

ARC Advisory Group

Aspen Asset Performance Management™

With advanced analytics, we can now provide operators with an ensemble of models that provide a holistic view of the process and the asset. Aspen APM combines asset analytics, reliability modeling and machine learning to provide comprehensive analysis, understanding and guidance.

Principles of data analytics and data science enable the reliability strategy, and likely need some explanation. The term analytics is used to mean the techniques used for understanding known data, whereas data science pertains to predictions of unknown data. An essential core element of leading solutions is the use of machine learning as a fundamental analytical and data science technique. Machine learning is a computer science subfield that has evolved from the study of pattern recognition using algorithms that learn and make predictions from data. It is now the dominant predictive analytics technology in all IT fields.

However, machine learning does little by itself; it is complex, difficult to master, and takes enormous efforts to develop truly workable and scalable solutions that do not require squadrons of data scientists. Additionally, within manufacturing, it is not just about the numbers. Proper application on manufacturing assets requires domain-specific knowledge of the chemical processes, the mechanical assets, maintenance practices and so on. “Industrialized” machine learning requires nuance and inside knowledge of interpreting and massaging complex, problematic sensor and maintenance event data to manage the health of industrial equipment.

Deeper Insights via Asset Analytics

Enriched analytics are at the root of APM 2.0. Asset-specific analytics provide information that was starkly absent in the first-generation APM solutions. AspenTech’s approach is quite different. Using ensemble modeling, we can construct applications that provide multiple perspectives

APM 2.0 incorporates the advanced analytics that will predict issues and prescribe operator actions. A function of the analytics will be to discern the process operating conditions that result in (latent) damage and diminished reliability. For example, in early research with customers, we identified problems with operators unable to determine what is going on inside some distillation columns. The columns sometimes go into undesirable states of flooding or weeping. In the flooding example, as the name implies, the liquid levels in the column become too high and affect the efficiency of separation. A somewhat related condition is weeping; the vapor pressure in the column is too low to prevent the liquid flowing across the trays from “weeping” down through the holes in the trays. Existing process information does not help operators see these conditions coming. Recovery can take many hours, costing tens of thousands of dollars. Operators need new insights into the internal conditions to avoid process-induced damage and downtime.

“AspenTech’s new Asset Analytics contains a unique set of modeling and data science-based technologies. Utilizing the additional process insight available from this promising new software solution brings with it the potential to operate closer to the true flooding limit on this tower. For a world scale olefins unit, this would be worth millions of dollars per year.” — LyondellBasell

> "[…] entirely new and more affordable manufacturing analytics methods and solutions have emerged, and they are now reaching market maturity as part of Industry 4.0. These solutions—which provide easier access to data from multiple data sources, along with advanced modeling algorithms and easy-to-use visualization approaches—could finally give manufacturers new ways to control and optimize all processes throughout their entire operations."
>
> McKinsey & Company

Not Just Any Analytics

Industry 4.0 began as a German initiated high-tech strategy to promote smart manufacturing concepts. It includes cyber-physical systems monitoring physical processes and making decentralized decisions. This concept is based on a marriage between the asset and the analytics that are responsible for monitoring and managing its health and effectiveness. Analytics of this form are chartered to create a multi-faceted view of the asset that enables fact-based decisions that consider a broader set of trade-offs.

One of the challenges in producing these solutions is the sheer number of asset types to cover. While there are thousands of engineering simulation models in existence, we need an efficient way to transform them. As shown in Figure 3, that transformation consists of taking a powerful, generalized model that is capable of being utilized for many different use cases and converting it into a model that is packaged and tuned to a very specific purpose. That transformation strips away unnecessary complexity and results in a model that delivers detailed information without requiring an army of engineers to maintain and calibrate.

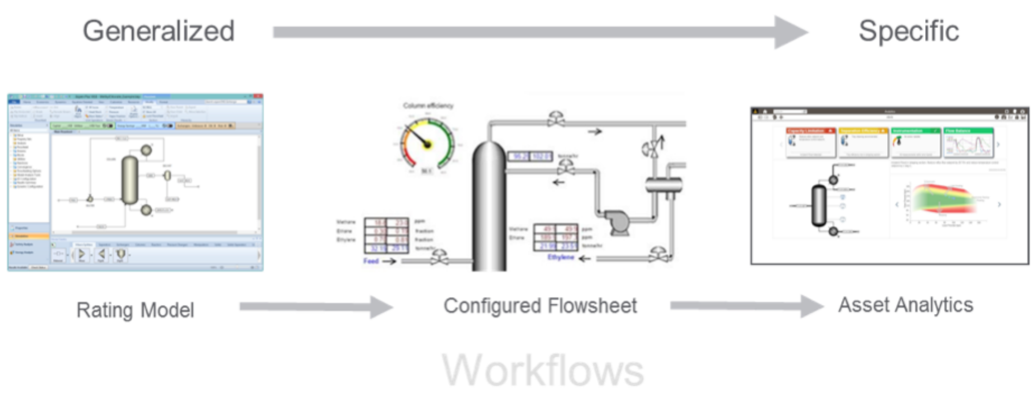

Figure 3: Transformation of a Generalized Rating Model

We have executed this transformation for certain distillation tower configurations. We have learned that it will require multiple models in these asset-specific applications where each model is responsible for a particular facet of information. As you'll see, we have successfully blended engineering simulation models with pattern recognition and other predictive technologies that work in concert to deliver a look-ahead view of production. That said, there is a need to scale this process and create tools to aid in these workflows. This is where domain expertise and breadth of experience is critical.

Asset Analytics for Distillation

Figure 4 is a screen shot of the operator interface for our new Asset Analytics for Distillation. There are several models running in the background to produce this view. Let's examine these underlying models and how they work together to build a comprehensive view.

The image at the lower left in Figure 4 is the typical process graphic that operators see on the distributed control system (DCS) console. It provides the data on process settings. To the right of this is a tailored Aspen Plus model of the hydraulics of the column. This model is used to tell operators when they are moving in a direction likely to result in flooding or weeping in the column.

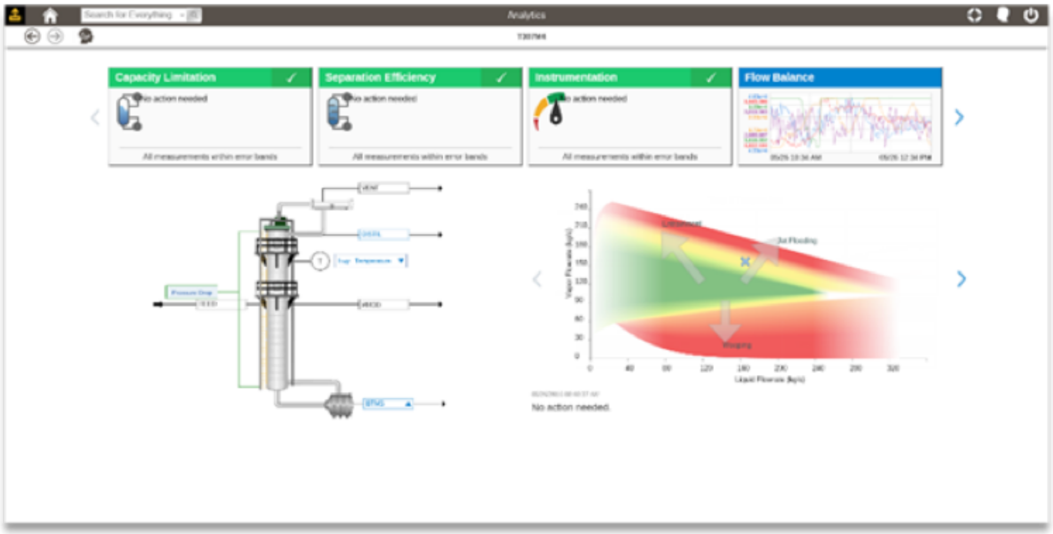

Figure 4: Aspen Asset Analytics™: Operator Interface

The next component is using machine learning to monitor real-time data streams, looking for patterns in the data that indicate potential problems. When the signature of a known failure mode is detected, the operator is alerted.

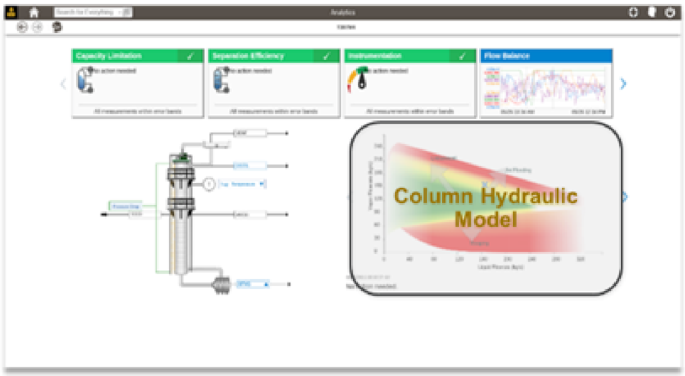

Figure 5: Column Hydraulic Model in Aspen Asset Analytics™

The column analytics application includes models that look for problems in the plant instrumentation. As shown in Figure 6, inferential models in the application can detect when instrument values don't make sense and might indicate a calibration or instrument failure.

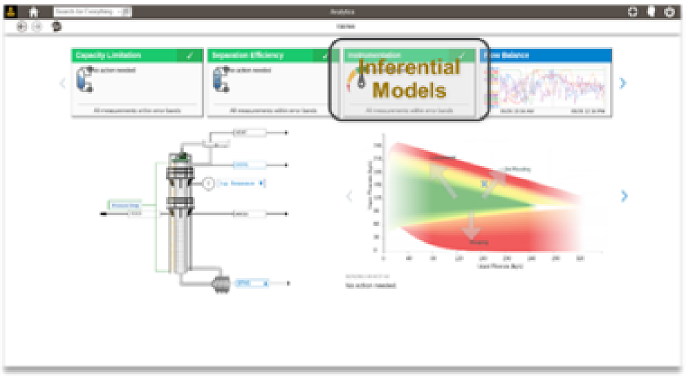

Figure 6: Inferential Models in Aspen Asset Analytics™

With this comprehensive information, we can provide an interpretation of the situation to the operators. Figure 7 depicts the summary when no abnormal conditions exist.

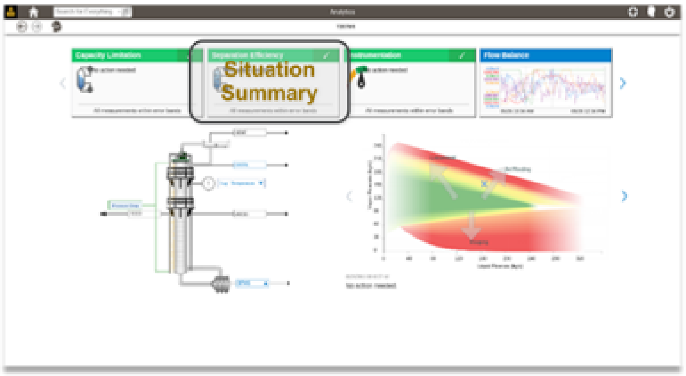

Figure 7: Situation Summary in Aspen Asset Analytics™

Figure 8 illustrates the view of the operator interface when a process problem has been detected. In the upper left corner, we have the status panel. It provides a summary of the process and asset condition. In this case it tells the operators that they are in a capacity limited state due to low separation efficiency in the column. The lower right area of Figure 8 now shows a probability plot for the predicted event. This area contains recommended actions for the operator to take to resolve a situation. The block above shows the recommended process setpoints for the operator to change to bring it back to a stable point.

Figure 8: Operator Interface Detects Process Problem

Asset Analytics are tailored to a specific asset type, with preconfigured and self-calibrating models. In the case of the hydraulic model, we have removed unnecessary degrees of freedom in the model and set parameter values appropriately. The result is a fit-for-purpose model without unnecessary complexity, without a maintenance burden, and without the need for an expert in simulation.

Pattern Matching

"I could see that coming from a mile away." It's an expression we hear frequently. The fact is that there are (sometimes subtle) "signatures" in the data that portend trouble. Pattern matching is like having a handwriting expert looking for those signatures in the real-time data, and even the predictive data streams.

Pattern matching is different than machine learning where there is no a priori model. In pattern matching, the user visually describes the relationships of interest and the algorithms go happily searching through data looking for those signatures.

Pattern matching is being successfully applied to monitoring real-time data for problems, and in data mining of process history to find operating periods of interest. It is also useful in preparing for the development of prescriptive applications by capturing unstructured information to create context for the process history. This context will be used in future endeavors to mine process history to encapsulate best practices in procedural automation that guides the human interaction with the process automation.

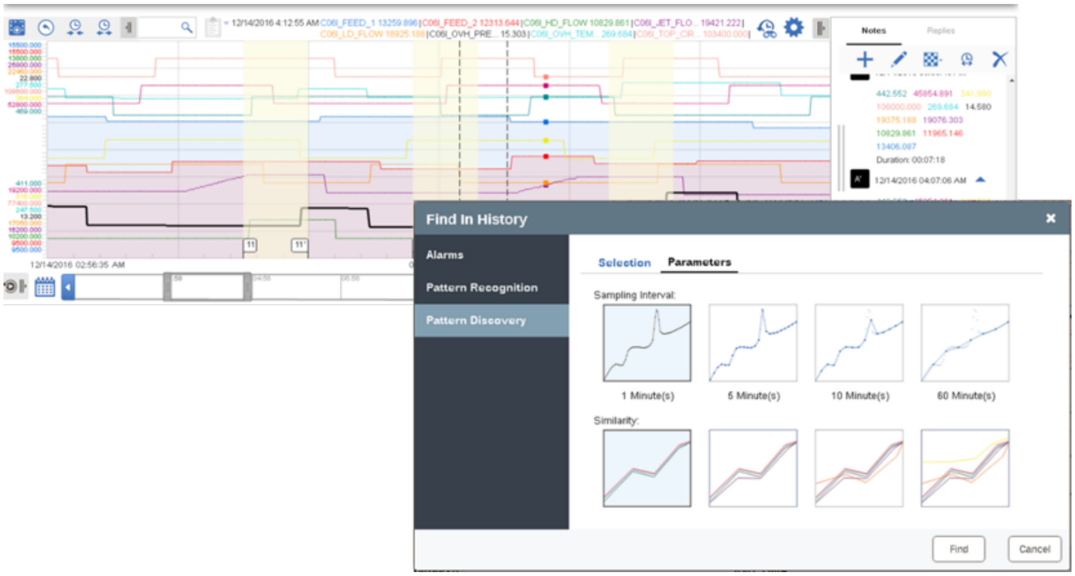

Figure 9: Pattern Matching Used in Data Mining of Process History

Root Cause Analytic

Let's assume we are incurring a disruptive event. Subject to cascading alarms, profitability trade-offs, and dozens or even hundreds of process inputs to manage, operators are often at a loss to determine the most effective course of action. With so many possible process handles, operators are not always fully aware of how and when their control actions will impact the process. While standard operating procedures provide useful information, they do not always inform operators about the cause and effect relationships.

There is no fixed frequency of these types of disturbances and the subsequent analysis. They can be the result of changes in feed composition, equipment issues, control problems or a host of other reasons.

The operator typically does not have a good understanding of the process fundamentals, and relies on training and experience. The operator will first try to resolve the problem before escalating it to the supervisor or process engineer, and their resolution attempts may exacerbate the problem. When finally escalated, it may be too late to avoid a shutdown or serious incident.

The process engineer will need to work fast to avoid further escalation, and there may not be time for a detailed analysis. The root cause may not be understood so treating the symptom may further exacerbate the problem. While the engineer will talk to the operators and specialists and then attempt to quickly prioritize the many possible causes and responses, engineers are often new recruits lacking understanding of the cause and effect relationships.

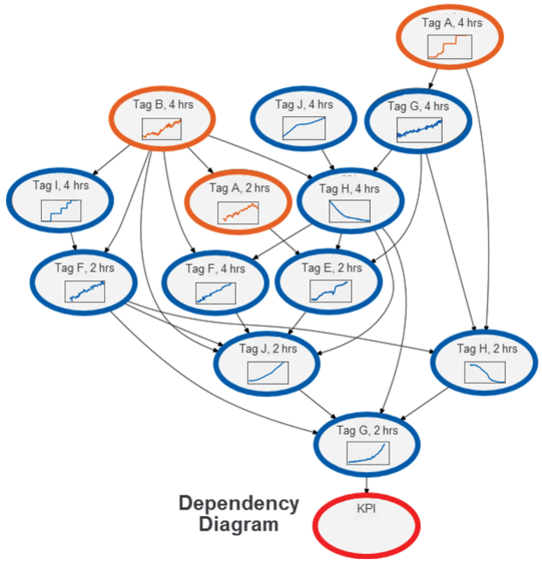

Figure 10: RCA Dependency Diagram

Machine Learning

Earlier techniques for condition-based monitoring relied upon specific detailed mathematical and statistical models for explicit equipment types. Here, intense engineering and mathematical skills formalize relationships between variables in the form of mathematical equations. Previous efforts to deploy machine learning often resulted in a cacophony of alarms, many of which were false. Operators quickly learned to tune them out.

Newer machine learning technologies take a different approach. The patterns of behavior do not come from rules and equations. Rather, the behavioral models are derived automatically from the available data without formal rules-based programming. If the data are available, machine learning can use data mining to derive patterns of behavior from archived sensor and maintenance event information, and thereby predict future behavior.

Machine learning approaches can encompass nuance and detail that eludes model-based approaches, and can rapidly learn on previously unseen equipment. Traditionally, model-based approaches have been applied to rotating equipment such as pumps, compressors, turbines, motors, and drives with success. The newer machine learning approaches fare equally well and – besides being easier and faster to deploy – often exceed the capabilities of models.

In an operating company, board operators noticed a data drift but had no DCS or secondary alarms. They noted that the differential pressure was increasing on a reactor.

When it got worse the operators escalated the problem to the operations engineer, but the engineer's analysis was too late to avoid a shutdown.

They could have avoided the incident with an analytic tool that provided a decision tree linking increased pressure with plugging in the reactor.

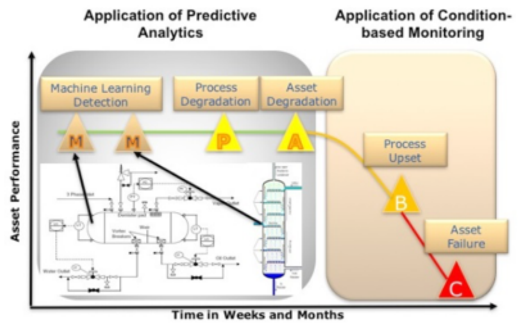

Figure 11: From ARC Advisory Group

MACHINE LEARNING

We can let machine learning do the heavy lifting of determining the operating conditions and patterns that have a deleterious impact on the asset by capturing the patterns of process operation and merging them with failure information.

In addition, previously built models are not a limitation for novel machine learning approaches. Consequently, the technology is applied to static equipment where mathematical models have not been available, such as mobile transportation equipment, and is emerging into process equipment such as exchanges, chillers, and incinerators. It can even be used in areas of piping that tend to plus and to look for corrosion and breakthrough in heat exchangers,

The technology is self-learning and adaptive which implies that the same core machine learning competence can serve many markets including oil and gas, heavy chemicals, pharmaceuticals, pulp and paper, water/waste water, mining and even discrete industries and semiconductors.

Machine learning technologies in oil and gas and heavy chemicals have delivered weeks' notice of emergent compressor issues such as liquid carry-over, dirty feedstock affecting seals and bearings, repeat valve problems, and similar issues for centrifugal and electric submersible pumps. Similar results have occurred in pulp and paper and autoclave incidents. In transportation, such technology has eliminated catastrophic engine failure in locomotives and detected failure issues in the six main systems of the heavy-haul trucks used in mining.

Figure 12: Software agents monitor real-time data for precursor patterns

For those interested in the details, you can read about our patented technology here. (**link to patent grant**).

Reliability Modeling

Aspen APM uses reliability analysis to create currency, or measure of greatest need. It does this by not only calculating a probability matrix, but by looking at flows and other measures that enable accurate estimates of the cost of lost production along with reliability factors.

When new concerns arise from machine and process analytics, the Monte Carlo simulations can be performed to determine how priorities and potential losses are stacked. This creates a scalable way of managing large numbers of events without increasing the chaos in decision making. Codifying governance directives for safety and protection of assets will be a key enabler of scalable yet agile enterprises.

Avoiding catastrophes is highly correlated with equipment reliability. If a machine does not degrade and break it will not cause the issues. In contrast to reliability, availability is a measure of the time the equipment is usable. Therefore, reliability and availability are different but related. Two factors influence availability:

1. Reliability [a function of mean time between failures (MTBF)]
2. Maintainability [a function of mean time to repair (MTTR)]

Generally, a machine with high reliability imparts high availability. As a cautionary note, however, measuring MTBF as an indicator and absolute metric is ambiguous. As explained in the next section, this is due to the fact that the causes of failure are not statistically deterministic. Improving reliability reduces risks — safety, environmental and economic.

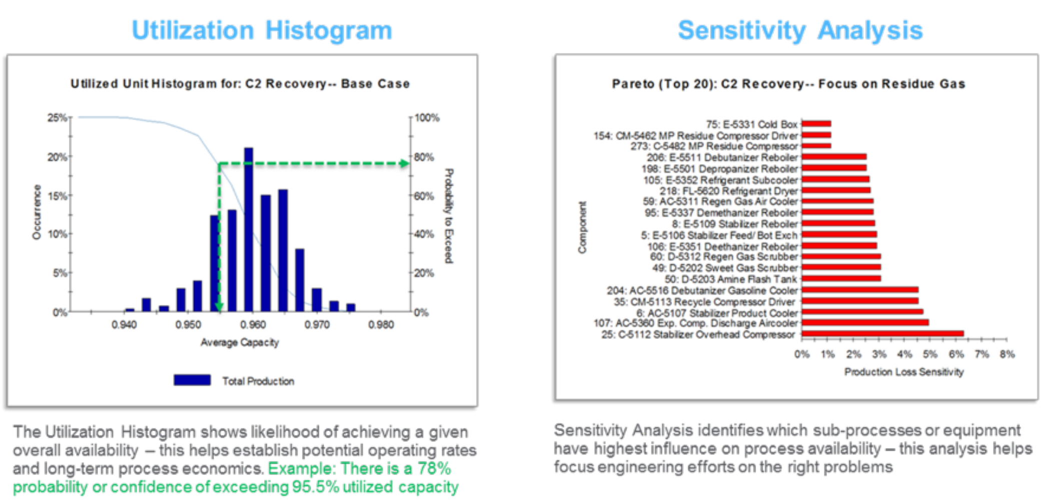

Figure 13: Analysis Artifacts from Aspen Fidelis Reliability Study

Reliability factors should be modeled and incorporated into maintenance and operations decisions. With reliability modeling we identify all causes of lost production or revenue and quantify each contributor. More importantly, the financial impact of the various solution scenarios can be quantified. With that information, we can make fact and monetary-based decisions.

The Fidelis Reliability software is able to dynamically solve the most complex, integrated process systems. It takes into account local operations and maintenance practices, redundancy, unit turn-up capacity, buffering, alternate flow paths, seasonal demands, weather and product shipping logistics in order to provide an accurate prediction of future performance.

Armed with the reliability analysis, and the impact analysis, we can move from emotional to fact-based decisions. Reliability modeling helps us understand and quantify the risks and downside costs, organizing those elements according to lost profit, not just lost time or lost production.

Few, if any, organizations have the resources to address everything at once, so to keep things moving we need a "currency" that normalizes the current basket of problems. Here's where we have taken the concepts of RCM and made them a dynamic process rather than a point-in-time analysis.

Bringing It All Together

Predictive analytics will go a long way towards reducing downtime, but we are seldom in situations where we are managing a single disruption. There are dozens of reliability, process and asset issues in play at any given moment in time. One of the problems with RCM is that it is a static assessment. When new conditions arise, you must go back through the analysis process from the beginning. That delays decisions when time is big money.

What is needed is dynamic assessment. As new warnings come from the machine and asset analytics, it must be evaluated alongside all of the other active conditions to set priorities and allocate resources. We usually can't address everything at once, so a system is needed to address problems according to the level of risk.

With Aspen APM, each new alarm triggers a recalculation of the risk profiles, guaranteeing that the most current financial and risk probability assessment is used to determine service priorities.

Making It Happen for You

In no uncertain terms, deploying analytics to successfully improve safety and reliability requires several elements: the right goals, culture, focus, people, technology, data strategy, software/hardware, and urgency.

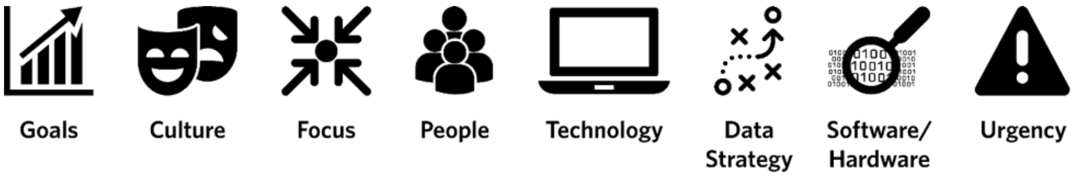

Goals: What is the organization trying to achieve? This often gets lost in the tension and enthusiasm when deploying the latest glorified technology. Not a new phenomenon, years ago I witnessed my employer's IT department proposing initiatives for Business Warehouse and Business Intelligence deployments and my response was, "That's an answer, what's the question?" In other words, do not follow a solution looking for a problem. Rather, understand what specific business problem needs to be solved and then look for the appropriate mechanism to solve it.

Culture: It's a new world, driven by data, but culture dictates if and how something can be achieved. It determines how an organization assesses performance, allocates resources, and administers actions. Culture propels intrinsic motivation. Leading organizations now use big data analytics to reach past "what was," to achieve "what can be." Has the organization seized a cultural desire for change?

Focus: Analytics work must align with important business goals. The organization develops an analytical aptitude; data-driven rather than opinion-driven – data do not have opinions. Focus permits the understanding of the differences between lagging and leading indicators and how to act.

People: Who are the business stakeholders and users? What experience and skills do the organization's people need? Impulsive stakeholders hire highly skilled data scientists, who usually have little domain-specific knowledge and are motivated by technology rather than business outcome. The choice of people and the choice of technology, strategy and solution are highly interdependent. An understanding of use cases is critical. How will the mechanics of the solution process enable the people to predict an outcome, prescribe the potential changes, and implement changes appropriately?

Technology: No matter how good the technology, if the organization requires large work process changes, it is unlikely to be extensively deployed or "move the business" in any significant way. The leading machine learning tools widely surpass the best of yesterday's technologies in the analytics fields. Again, the difference will be the time, skills and experienced needed to master the tools and technologies that attempt a solution.

Data Strategy: The analytics program must align with the fundamental business goals including performance, quality and net product output. It begins with what the organization is trying to achieve – and avoids the "answer looking for a problem" and "all dashboard, no improvement" syndromes. Understand business requirements, identify data requirements, and carefully select the appropriate tools and products that directly align with the analytics job at hand. Understanding data context is necessary for analytical tools to provide good solutions – you cannot just stir the data and apply algorithms and get meaningful solutions. Don't just copy the bank. The bank has far different data and is not using machine learning to solve the same problem you will in the same way. Do the collected data correlate symptoms and cause? Do the data types you gathered match the data technology? For example, supervised machine learning aligns patterns with events to enable the detection of recurrences. Unsupervised learning finds "random" patterns that might be meaningful. Does the data strategy support both? Does it need to?

Software: The tools and products chosen must support the problem to be solved, the available skill sets, the delivery time in which results and action can be taken following event detection, and the repeatability needed for useful deployments. Is the solution a project development kit, where each application is a new ground-up development? Or do you have an out-of-the-box product delivering fast, repeatable applications that require skills the organization already has? Use an Impact and Effectiveness Matrix to assess the products in the market and determine how well they can fit your organization's solution and people needs – e.g., skill set required versus solution efficacy.

Hardware: Modern analytical solutions require beefed up hardware beginning with "serious" Windows servers, and advancing to Linux, Hadoop clusters and in-memory computing hardware. Does the organization have an appetite for allocating such hardware resources? What's needed to get started? Can your chosen solution scale?

Urgency: Does the organization "walk the walk" or does it just "talk the talk?" Leading companies have C-level sponsored initiatives for analytical solutions, aligned with specific corporate objectives; for example, to improve operational effectiveness, reduce unplanned outages, and improve environmental and safety issues caused by equipment behavior. If this is a pressing priority, then a budget is appropriated and made available. If not, a different implication can be drawn.

AspenTech is a leading supplier of software that optimizes process manufacturing — for energy, chemicals, engineering and construction, and other industries that manufacture and produce products from a chemical process. With integrated aspenONE® solutions, process manufacturers can implement best practices for optimizing their engineering, manufacturing, and supply chain operations. As a result, AspenTech customers are better able to increase capacity, improve margins, reduce costs, and become more energy efficient. To see how the world's leading process manufacturers rely on AspenTech to achieve their operational excellence goals, visit **www.aspentech.com**.

Worldwide Headquarters
Aspen Technology, Inc.
20 Crosby Drive | Bedford, MA 01730 | United States
phone: +1-781-221-6400 | fax: +1-781-221-6410 | **info@aspentech.com**

Regional Headquarters
Houston, TX | United States
phone: +1-281-584-1000

São Paulo | Brazil
phone: +55-11-3443-6261

Reading | United Kingdom
phone: +44-(0)-1189-226400

Singapore | Republic of Singapore
phone: +65-6395-3900

Manama | Bahrain
phone: +973-13606-400

For a complete list of offices, please visit **www.aspentech.com/locations**