SAE International

SURFACE VEHICLE/ AEROSPACE RECOMMENDED PRACTICE

SAE JA6097 MAY2013

Issued 2013-05

Using a System Reliability Model to Optimize Maintenance Costs
A Best Practices Guide

RATIONALE

Complex repairable systems that consist of many individual components can be designed to be quite reliable, at least when they're still new. However, as such a system continues to be used and repairs are made to address failures that occur, the interval between failures gets shorter and shorter. Eventually, that complex system becomes a collection of parts, each with a different amount of operating time and differing reliabilities. Maintenance intervals that may have aligned with one another when the system was new will eventually become misaligned, making maintenance more complicated, or at least sub-optimal and more expensive. It is not uncommon to have a repaired system returned to service only to be shut down shortly thereafter because of another failure.

The classical Reliability Centered Maintenance (RCM) process does a good job of focusing attention on the actual reliability and failure modes that is (or will be) exhibited by a particular system, and helps a system designer, implementer, or operator develop an effective maintenance strategy for that system. However, the RCM decision logic only looks at the individual components and doesn't address how to reconcile different maintenance intervals for each component across the entire system. Furthermore, RCM focuses more on preventive maintenance (lubrication and failure finding tasks) and doesn't really address corrective maintenance. So when a complex system is taken out service for maintenance, there is no guidance as to what other optional maintenance should be performed to increase the time until the next failure and/or reduce long-term operating costs.

This "Best Practices Guide" describes a proven approach for objectively determining what other maintenance should be performed when a system is being repaired to improve system reliability and reduce long-term operating costs.

FOREWORD

The classical RCM analysis process is designed to look at each individual component of a complex system (such as an aircraft, military tank, ship, radar system, factory, etc.), determine how each component can fail, identify the consequences of each failure mode, and decide which maintenance strategy best meets the needs of each component in order to maximize the reliability of the system. The RCM maintenance strategy options are: impose a scheduled discard event, running a component to failure, initiating failure-finding tasks, implementing a preventive maintenance program, or a one-time change (redesigning a part if its failure would pose a safety or environmental risk and such a failure cannot be mitigated through other means, operational change, etc.). At the end of this process, the analyst will have a structured, thorough, maintenance plan that will account for all parts of the system being evaluated.

TO PLACE A DOCUMENT ORDER: **Tel: 877-606-7323 (inside USA and Canada)**
Tel: +1 724-776-4970 (outside USA)
Fax: 724-776-0790
Email: CustomerService@sae.org
SAE WEB ADDRESS: http://www.sae.org

However, classical RCM (as defined in the current RCM Standards) doesn't take into account situations where maintenance on one piece of equipment may disturb or drive maintenance on one or more other pieces of equipment. It also doesn't tell the analyst, when the RCM analysis specifies maintenance intervals for various portions of the system that are all different from each other, how to reconcile those differences in such a way that will keep maintenance costs, down-time, or any other system-wide maintenance optimization, to a minimum. Furthermore, it will not address how to best perform ***corrective*** maintenance – i.e., how to decide what other optional tasks should be performed while the system is down for maintenance.

Undoubtedly, there will be differences between different types of systems, users, operating locations, etc., and they all would have some kind of criteria to measure the cost-effectiveness of any maintenance decisions that are contemplated. The real issue is how to calculate the cost effectiveness in all those various situations. The costs themselves may be easy enough to compile, but the "effectiveness" part of the equation is usually a bit more difficult to determine.

This guide is an attempt to document and explain one general method for determining the cost-effectiveness of various maintenance options that are possible when a complex, heavily integrated system is taken off-line for maintenance. The techniques described herein have been developed for, and successfully applied to, aircraft and aircraft engines. They allow a comparison of a wide variety of considerations to be made whereby the cost and the "effectiveness" are determined computationally (vice notionally), and gives system maintainers a more unambiguous recommendation for the maintenance that should be performed that will result in the lowest practical long-term maintenance cost.

Since the techniques described herein have been successfully applied to aircraft and aircraft engines, the examples below will frequently refer to these applications. However, in the course of the following discussion, this guide will try to provide enough explanation of the issues involved so that the reader will be able to identify and apply the same techniques to their own situations and address any similar requirements and constraints that pertain to them.

TABLE OF CONTENTS

1. SCOPE

SAE JA6097 (“Using a System Reliability Model to Optimize Maintenance”) shows how to determine which maintenance to perform on a system when that system requires corrective maintenance to achieve the lowest long-term operating cost. While this document may focus on applications to Jet Engines and Aircraft, this methodology could be applied to nearly any type of system. However, it would be most effective for systems that are tightly integrated, where a failure in any part of the system causes the entire system to go off-line, and the process of accessing a failed component can require additional maintenance on other unrelated components.

2. REFERENCES

2.1 Applicable Documents

The following publications form a part of this specification to the extent specified herein. Unless otherwise indicated, the latest issue of SAE publications shall apply.

2.1.1 SAE Publications

Available from SAE International, 400 Commonwealth Drive, Warrendale, PA 15096-0001, Tel: 877-606-7323 (inside USA and Canada) or 724-776-4970 (outside USA), www.sae.org.

SAE JA1011 Evaluation Criteria for Reliability-Centered Maintenance (RCM) Processes

SAE JA1012 A Guide to the Reliability-Centered Maintenance (RCM) Standard

2.2 Related Publications

The following publications are provided for information purposes only and are not a required part of this SAE Technical Report.

F. Stanley Nowlan and Howard F. Heap, “Reliability-Centered Maintenance”, Department of Defense, Washington, D.C., 1978, OASD Report Number AD-A066-579

2.2.1 US Government Publications

Condition Based Maintenance Plus DoD Guidebook, May 2008

Condition Based Maintenance Plus (CBM+) for Materiel Maintenance, DoDI 4151.22, December 2, 2007 (http://www.acq.osd.mil/log/mpp/cbm+/CBM_DoD_Guidebook_May08.pdf)

Guidelines for the Naval Aviation Reliability-Centered Maintenance Process, NAVAIR 00-25-403, U.S. Naval Air Systems Command

Planned Maintenance System: Development of Maintenance Requirement Cards, Maintenance Index Pages, and Associated Documentation, MIL-P-24534 (U.S. Naval Sea Systems Command)

Procedures for Performing a Failure Mode, Effects and Criticality Analysis, MIL-STD-1629A, Notice 2, 1984, Department of Defense, Washington, D.C. (NOTE: canceled without replacement, August 1998.)

Reliability Centered Maintenance for Aircraft, Engines and Equipment, MIL-STD-1843, United States Air Force (NOTE: canceled without replacement, August 1995.)

Reliability-Centered Maintenance Requirements for Naval Aircraft, Weapons Systems and Support Equipment, MIL-HDBK-2173(AS), U.S. Naval Air Systems Command (NOTE: canceled without replacement, August 2001.)

2.2.2 Books

Anderson, Ronald T. and Neri, Lewis, "Reliability-Centered Maintenance: Management and Engineering Methods," Elsevier Applied Science, London and New York, 1990

Andrews, J.D. and Moss, T.R., "Reliability and Risk Assessment," Longman, Harlow, Essex (UK), 1993

Blanchard, B.S., D. Verma and Peterson, E.L., "Maintainability: A Key to Effective Serviceability and Maintenance Management," John Wiley and Sons, New York, 1995

Cox, S.J. and Tait, N.R.S., "Reliability, Safety and Risk Management," Butterworth Heinemann, Oxford, 1991

Jones, Richard B., "Risk-Based Management: A Reliability-Centered Approach," Gulf Publishing Co., Houston, TX, 1995

Moubray, John, "Reliability-Centered Maintenance", 1997

Rausand, Marvin and Høyland, Arnljot, "System Reliability Theory, Models, Statistical Methods, and Applications", 2004

Smith, Anthony M., "Reliability Centered Maintenance," McGraw-Hill, New York. 1993

Zwingelstein, G., "Reliability Centered Maintenance, a practical guide for implementation," Hermès, Paris. 1996

2.3 Other Publications

"Dependability Management - Part 3-11: Application guide - Reliability centred maintenance," International Electrotechnical Commission, Geneva, Doc. No. 56/651/FDIS

MSG-3, "Maintenance Program Development Document," Air Transport Association, Washington, D.C. Revision 2, 1993

NES 45 - "Naval Engineering Standard 45, Requirements for the Application of Reliability-Centred Maintenance Techniques to HM Ships, Royal Fleet Auxiliaries and other Naval Auxiliary Vessels" (Restricted-Commercial)

3. DEFINITIONS

3.1 PROBABILISTIC

Of, based on, or affected by probability, randomness, or chance (American Heritage Dictionary). For the purposes of this document, a probabilistic failure is an event where the system (or a component thereof) becomes unserviceable, and the underlying probability of that type of failure mechanism occurring at a given point in time can be represented by a mathematical function, such as an Exponential, Weibull or LogNormal distribution. See also Stochastic (below).

3.2 STOCHASTIC

Of or pertaining to a process involving a randomly determined sequence of observations each of which is considered as a sample of one element from a probability distribution (Random House Dictionary).

3.3 SUNSHINE COSTS

Costs incurred to correct conditions discovered only when a part is exposed and a fault with that part becomes visible to, and detected by, the technician.

3.4 WORKSCOPE

The scope (amount or extent) of work you plan to do during a maintenance event.

4. ACRONYMS

4.1 ATOW

Average Time On-Wing. This is the historical time between aircraft engine removals.

4.2 CBM

Condition-Based Maintenance. CBM is an approach to maintenance whereby maintenance is only performed on evidence of need identified through direct or indirect monitoring. CBM requires specific knowledge of an asset's condition at any given time in its operating life such that the maintenance action can be planned with sufficient lead time to minimize the cost and operational impact of the occurrence of a failure. CBM differs from "on-condition" maintenance in that under CBM, knowledge of asset condition at any given time provides an understanding of how much time is available before the required maintenance must be performed.

4.3 CPFH/CPEFH

Cost Per (Engine) Flying Hour. The cost of maintenance (parts and labor, not fuel) divided by time the system operates from the time it is repaired until it fails again, or the estimated amount of time it could be expected to operate given a certain level of repair.

4.4 ETOW

Estimated Time On-Wing. This is a future projection (estimate) of aircraft engine life expectancy (time between engine removals). It is the "expected value" (E(t)) term derived by integrating the overall system reliability function (R(t)).

4.5 LCF

Low Cycle Fatigue.

4.6 LLP

Life-Limited Part. Engine parts that have a limited allowable service life.

4.7 OCM

On-Condition Maintenance. Scheduled maintenance performed to detect potential failures. This usually consists of inspections, tests, or measurements to determine if an item is or is expected to remain in serviceable condition until the next scheduled maintenance event.

4.8 PM

Preventive Maintenance. Maintenance tasks such as inspections, adjustments, tests, as well as lubrication and other servicing actions performed to extend the life of a component or system.

4.9 RBD

Reliability Block Diagram. RBD models are a way of graphically representing the reliability relationships of various components in a system, including the potential for single point failures as well as redundancies.

4.10 RCM

Reliability-Centered Maintenance. RCM is a specific process used to identify the policies which must be implemented to manage the failure modes which could cause the functional failure of any physical asset in a given operational context. (See SAE JA1011).

4.11 TSO/TSN

Time Since Overhaul or Time Since New. Different measures of service life accumulation.

5. BACKGROUND

As we have all been reminded many times over the last few years, aircraft operators (specifically, commercial airlines and cargo delivery service providers) are under intense pressure to reduce operating costs. Rising fuel prices, increased competition spurred by de-regulation, and aging aircraft fleets requiring increased levels of maintenance are all cutting into the airlines' profit margins. Several airlines have even been driven into bankruptcy as a result.

The single largest controllable expense for aircraft fleet owners or operators is maintenance. In fact, the aircraft engines alone typically account for about 50% of all aviation fleet maintenance costs. Since other costs, such as fuel, are driven by external market conditions and are not really controllable, those aircraft fleet owners have typically looked at how they perform maintenance as a means to reduce costs.

Historically, maintenance decisions have been made based only on short-term needs, not based on the long-term financial impact. One of the reasons is that the maintainers have had no objective method for determining what maintenance is the most cost-effective whenever something breaks; it's all based on 'experience' or 'gut instinct', neither of which is very accurate and can sometimes be totally wrong. In addition, aircraft engine maintenance is complicated by an abundance of rules, regulations, technical manual requirements, the complex design of the engines, a wide variety of configurations, high asset value, high cost impact when it's unavailable, all of the possible decisions to be made when engines are brought in for repair, and whether it is a scheduled or unscheduled maintenance event. Any method that could objectively determine what maintenance to perform, with the goal of lowering the long-term operating cost and increasing system availability, needs to address all of these issues.

6. LIMITATIONS OF 'CLASSICAL' RELIABILITY-CENTERED MAINTENANCE

Recall the steps of the traditional Reliability-Centered Maintenance (RCM) system analysis process (shown below). Basically, an analyst looks at all of the components of a system, identifies all the ways that each component can fail, determines the consequences of each failure, then decides (using the RCM decision logic) what maintenance strategy will best address that failure mode.

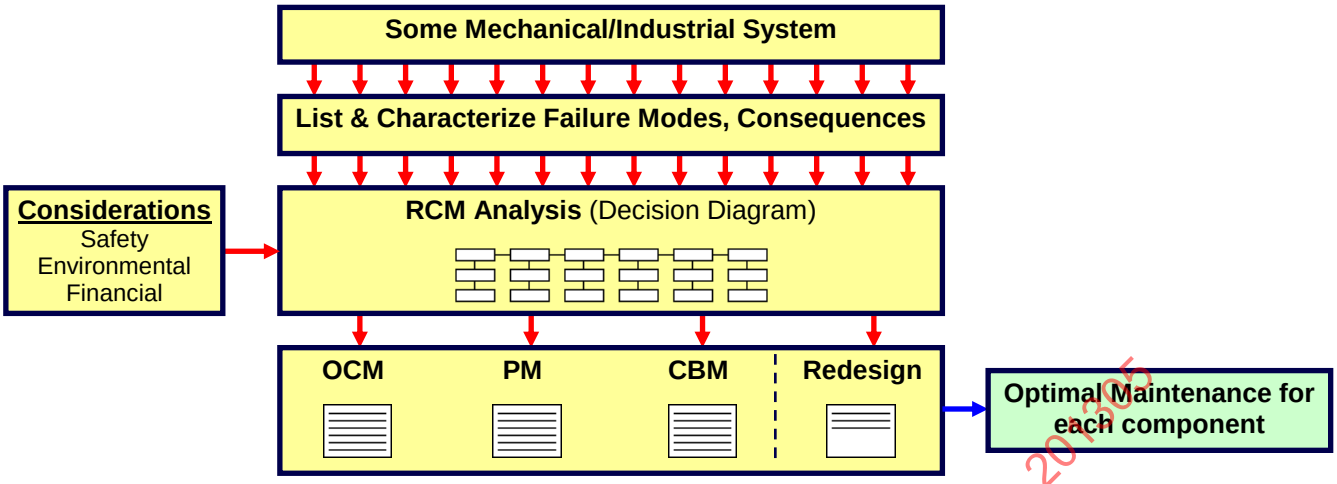

FIGURE 1

The traditional RCM analysis results in the optimal maintenance strategy for each individual component of a system, regardless of how complex that system may be. That, in turn, usually means many components will have slightly different timetables for maintenance. This may be perfectly suitable (or at least acceptable) for systems like a factory where most pieces of equipment are accessible without disturbing anything else and can be repaired independently of other equipment. Often, production processes are specifically designed so they can tolerate one or more pieces of equipment being inoperable for short periods of time and still cause no significant disruption to the overall production flow.

But what happens for systems where maintenance decisions on one portion of the system have direct ramifications on maintenance requirements for the rest of the overall system, or where the failure of a single component will cause the entire system to go off-line? As mentioned above, prime examples of systems like this include aircraft and aircraft engines. In addition, other examples may also include (at least for some kinds of maintenance actions) locomotives and/or railway cars, tanks and other heavy military vehicles, construction equipment (such as bulldozers, graders, etc.), ships, some types of chemical plants, nuclear reactors, and other complex systems. In these situations, the classical RCM maintenance approach is incomplete, and other techniques are needed to address this problem.

The systems mentioned above all share some common characteristics. The first is a high cost to bring those systems down for maintenance. Second, for aircraft and aircraft engines as well as the other examples cited above, a failure of one portion of the system will often leave the entire system inoperable. Also, performing maintenance on any portion of these systems often requires taking the entire system out of service. The maintenance may involve large complex equipment and many personnel to disassemble the areas that require attention and the disassembly may take a considerable amount of time to perform. As a result, the cost impact of a failure in the system can be significant.

Another common characteristic of these types of systems is that, given how tightly integrated and/or closely packaged some of the components are, getting access to the failed component(s) can require extensive disassembly of other uninvolved portions of the system. Exposing those non-failed portions of the system often creates considerable opportunities to find additional substandard conditions and the possibility that additional maintenance (primarily on those portions that are exposed) will be required in order to comply with government safety regulations (especially true for aircraft and possibly for nuclear reactors), the manufacturer's requirements, or operational needs. When they're discovered, the additional maintenance needed to address the other non-conformances then leads to increased costs over and above what was originally anticipated.

All of the factors described above mean that whenever maintenance is performed on systems like these, anyone interested in keeping costs down needs to make each maintenance event count and get the most 'bang-for-the-buck' when that maintenance is performed. In fact, it also means it can be financially advantageous to perform extra maintenance if it will reduce the chances of future failures.

In order to ensure the maintenance is as cost-effective as possible, the system owners and/or maintainers need a structured methodology to facilitate the decision-making process that can objectively account for the reliability of key components by using all applicable failure distributions to determine the effectiveness of different workscopes. It needs to weigh the cost of performing the different repairs against the benefits (at the system level) to identify the workscope that yields in the lowest long-term operating cost. This approach goes beyond the classical RCM analysis process by optimizing the maintenance for all components in the system.

So the question then becomes, how can system maintenance be optimized on complex integrated systems like the ones described above? First, let's take a specific example and dissect the system to identify the major subsystems, issues, and constraints. Then, once we understand the complexity of the problem, we can identify how to approach the solution.

7. MAINTENANCE REQUIREMENTS AND CONSTRAINTS – AIRCRAFT ENGINES

NOTE: Since the process described in this guide had been successfully developed and applied to aircraft engines, the following paragraphs will focus on that type of system, just for the sake of discussion. This does not mean that the methodology isn't suitable for, or cannot be applied to, other systems. Quite the contrary; it only means that the application of this technique to aircraft engine maintenance is far better understood at this time.

Readers already familiar with aircraft and/or aircraft engines may not need to refer to the following, although some of the terminology discussed below will be used later in this document.

7.1 Basic Parts of an Aircraft Engine

A typical jet engine consist of several "modules", such as one or more compressor(s), a diffuser, a combustion section, one or more turbine(s), and one or more external gearboxes. The compressor(s) and turbine(s) can consist of one, two, or even three major rotating assemblies, and they are enclosed within what are usually referred to as "cases". These rotating assemblies usually ride on roller or ball bearings mounted in heavy structural supports (often called "frames") at different locations throughout the engine. The modules, frames, cases, and major rotating assemblies constitute what are generally referred-to as the engine "core".

Also, there are so-called Life-Limited Parts, or LLPs – these are typically individual components of the primary rotating assemblies where the manufacturer is concerned about the effects of dynamic stresses (in the form of Low-Cycle Fatigue or LCF), heat, mechanical wear, deterioration through rework, etc., and must impose a maximum service life to ensure those parts nearly never experience an in-service failure. Jet engines also include a number of external accessories associated with the oil, fuel, pneumatic, hydraulic, and electric systems, etc., but they will not be discussed here since they generally don't drive engine removals and can be replaced without involving many other components (if any).

It is normal for the major engine subassemblies (modules) and critical components (LLPs) to be tracked in terms of the total amount of accrued operating time, cycles, etc. The amount of "Life" accumulated by these various parts of the engine drives maintenance decisions. For example, LLPs are not allowed to operate beyond the life limits set by the OEMs, and compliance with those limits is required by the FAA (or other applicable regulatory agencies around the world) – see below for an extended discussion of this topic.

Modules and other parts may also have time limits where inspections are required at fixed intervals. Those limits may be "hard" or "soft" (meaning there may be some flexibility as to when the maintenance must be performed), and those limits may be imposed by the OEM or the owner/operator.

7.2 Aircraft Engine Maintenance 101

The ideal situation from the standpoint of maintenance planning would be that all aircraft engine removals would be "scheduled". In other words, the operator and maintainer would be able to forecast all removals well in advance so that the necessary parts, manpower, equipment, and facilities could be budgeted ahead of time and the removal timed to be the least disruptive to the aircraft fleet's operational schedule. It would also mean that decisions on what maintenance should be done on an engine could be made only once for all engines, and all the maintainer would have to do is look up the 'rules' for what maintenance would be needed based on the specific situation for that particular engine removal.

Even though we don't live in such an ideal world, there are several types of maintenance actions that can be handled as scheduled events. The primary ones are for the expiration of life or inspection (time) limits (applied to LLPs and modules, respectively) which are based on the current flying (utilization) rate and the life used so far. Some maintenance requirements are strictly calendar-driven (e.g., ones that pertain to corrosion), so they're independent of the utilization rate. Also, some types of engines have such predictable performance degradation that the end of their useful life (in terms of engine performance) can be forecasted as well.

Engines are also subject to unscheduled removals for "probabilistic" (or stochastic) failures. Oftentimes, these may be characterized as wear-out, random, or infant mortality failures. Unscheduled engine removals caused by these types of failures can dominate fleet maintenance costs and will derail almost any scheduled maintenance plans.

Examples of some "probabilistic" failures include:

- Wear-out – Erosion from sand ingestion, or tread wear on tires. Both of these are directly related to the length of time the component has been in service, where the deterioration accumulates over time. While this implies you could establish a hard time limit to remove all components before they experience a wear-out failure, that would be an incorrect conclusion. The probabilistic aspect of wear-out behavior means you can never be 100% sure a component will always reach a particular limit, even if you set the limit very, very, low. Since asset owners also want to extract as much useful service life out of their equipment as possible, they would want to set these limits as high as possible and maximize their return on investment. As a result, this implies an assumption of a certain degree of risk that some unscheduled maintenance events will occur due to wear-out failures.

- Random events – A flat tire from a road hazard, or foreign object ingestion in a jet engine. Neither of these types of events have any relationship to how long a component has been in service. This is typically what most people think of when they refer to "probabilistic" failures.

- Infant mortality – Defects introduced during manufacture or rework, such as residual swarf remaining in a gearbox leading to bearing or gear failure, or faulty belts in a tire. Failures in this category can often benefit from some form of testing prior to being returned to service to capture failures before the asset reaches the customer, such as 'burn-in' testing to check for many types of electronics failures.

As implied above, these "probabilistic" failures can't be forecasted, but some of them can be anticipated with the aid of various types of monitoring equipment. Examples include vibration sensors, real-time oil debris detectors, and other specialized signal processing devices. The failures can be anticipated as long as there is sufficient time between when an impending failure can be detected and the time when the failure would actually occur. However, a significant limitation in this strategy is that it only works if the failure is something that can be detected accurately, consistently, and with enough advance notice that the maintainer can proactively prepare for the engine removal.

Once removed, these engines require workscoping – determining what work needs to be performed based on the reason for removal, the life remaining for the various modules and components of the engine, the equipment manufacturer's maintenance requirements and recommendations, and any applicable special safety inspections, repairs, and/or modifications and the associated requirements for each. Each module may have its own workscope, where each different workscope would entail different tasks & associated costs. Some workscopes will drive the replacement of certain parts, which then entails the cost of replacing those parts.

A typical aircraft engine can't be disassembled at any arbitrary location. In other words, the engine must be disassembled in a predetermined sequence, possibly resulting in many portions of the engine being exposed that are unrelated to the intended maintenance task and have no pre-identified defects. A simple analogy would be performing a brake job on a car with drum brakes – you couldn't just repair the brakes without doing some disassembly first: you'd first have to remove the hub cap, unscrew the lug nuts, remove wheel and tire assembly, and then remove the brake drum. The disassembly would pretty much have to be done in that sequence every time before the brake components would be exposed for maintenance.

The components on the exterior of a turbine engine are, for the most part, accessible without disturbing many (if any) other components. The engines have various types of components for fuel delivery (pumps, filters, fuel metering, flow meters, etc.), oil system operation (pressure & scavenge pumps, filters, breather valves, etc.), anti-icing, ignition, electrical power generation, hydraulic pumps, starters, pneumatic controls, and so forth. On nearly all modern engines, these components are replaceable without removing the engine from the aircraft or disassembling the engine core.

7.3 Contractual and Regulatory Requirements for Maintenance

From time to time, unusual conditions and an increased potential for dangerous failures to occur are identified. When that happens, the various aviation regulatory agencies around the world require aircraft operators and maintainers to undertake special inspections, repairs, or modifications to (1) identify the potential failures, (2) correct problems introduced through faults in manufacturing, repair, handling, or operation, and/or (3) address deficiencies in the design of critical parts (either through replacement or modification). Circumstances that lead to these types of special inspections or repairs can arise when a previously unrecognized failure mode has occurred, a manufacturing defect had been found in a group or type of part, or a flaw was found in the way certain parts were repaired.

For US civil aviation, these special inspections and mandatory maintenance orders are commonly known as Airworthiness Directives (ADs) or Service Bulletins (SBs), depending on their severity. The names given by other aviation regulatory agencies may be different, but the general principles of ADs and SBs would still apply. They can be required immediately (i.e., aircraft are grounded until the AD or SB is completed), at the first maintenance opportunity (whenever the aircraft, engine, or applicable major system component is removed from service), or at the convenience of the operator/maintainer (for low priority tasks). As such, the decision to perform maintenance for other reasons can also trigger a requirement to comply with one or more ADs or SBs, which would make an otherwise optional maintenance task much more expensive.

Another element in the engine maintenance decision-making process is that most aircraft owners/operators default to the engine manufacturer's technical requirements when it comes to maintenance. While airlines are allowed to establish their own engineering authority and override the manufacturer's requirements under certain circumstances, they are becoming less willing to do so because of the difficulty in establishing the justification to satisfy aviation regulatory agencies (such as the Federal Aviation Administration (FAA), European Aviation Safety Agency (EASA), etc.), their unwillingness to assume the liability risks by making engineering decisions that may contradict the manufacturer, or simply because they don't want the expense of having a large engineering staff to support that capability. The manufacturer's maintenance requirements are typically driven by how long certain portions of the engine have been in service since its last overhaul, what's been exposed in the course of performing other maintenance, how many times a component has been refurbished in the past, etc. Those requirements can frequently take the form of "*whenever part X is exposed, it must be removed, cleaned, inspected, and refurbished, unless it has been refurbished more than N times already, whereupon part X must be replaced*", or some variation thereof. Such a requirement would mean that any discretionary maintenance that exposes '*part X*' would need to be carefully considered prior to undertaking that maintenance task due to the extended costs (which might be unrelated to the primary maintenance task) that could be incurred.

8. MAINTENANCE REQUIREMENTS AND CONSTRAINTS – NON-AIRCRAFT SYSTEMS OR EQUIPMENT

While the maintenance requirements for aircraft engines were discussed in-depth above, there are other systems that have at least some similarities when it comes to maintenance. Railroad Locomotives, Military Tanks, bulldozers, and ship engines all may require some very heavy equipment to disassemble them. Certain industrial processes (like those in nuclear reactors, petrochemical plants, etc.) require the entire system (or at least a large portion of it) to be taken off-line whenever certain types of maintenance are performed. The common thread among these examples is that maintenance events can be expensive, and some future failures may be averted by performing some preemptive repairs when the opportunities arise.

For example, some types of chemical plants have complex handling systems that are tightly integrated from start to finish, making any kind of system shutdown an expensive and time-consuming undertaking. Their design may also mean that an in-service failure could be catastrophic, such as in the following example:

> A consumer products company had a manufacturing plant where they operated a Nylon injection mold production line. The injection molds were supplied with liquefied Nylon that was pumped through the center tube of an extensive network of double-walled stainless steel pipes that ran throughout the plant. The Nylon was kept liquefied though the use of molten Sodium flowing through the annular passage in the double-walled pipes. Sometime in the late 1950's or early 60's, through some kind of system breakdown (the exact nature of which is not clear, but it doesn't really matter for the purposes of this example), the Sodium was not kept hot enough, resulting in both the Sodium and the Nylon cooling and solidifying in the pipes. Since neither the Sodium nor the Nylon could be re-liquefied in the pipes once they cooled, this resulted in the total loss of that portion of the plant's Nylon distribution system. While the cost of that failure is hard to assess nearly 50 years after-the-fact, it certainly had to be significant considering that half of their facility was permanently shut-down following that incident.

Whether the failure described above was the result of a single point failure, or a combination of failures, RCM analysis should have been able to identify the critical component(s) and where the system was most vulnerable. Most likely it would have shown the circumstances and potential for a single point failure. Presumably, this would have resulted in a system redesign to incorporate higher reliability components or some kind of redundancy, such as the use of a backup heating system for the Sodium, along with some sort of alarm or monitoring to ensure that component failures could be identified before major damage was done to the rest of the system.

But more to the point, as far as this best practices guide is concerned, this is an example of an industrial system whose components are tightly integrated, and the failure of any one of them could have rendered the entire system inoperable. It is also an example of a case where shutting the system down to repair a failed component would be complicated, expensive, and time consuming. When the system was shut down for maintenance, it would have been a prime opportunity to perform additional maintenance on components other than just the one(s) that had failed or were about to fail, and some form of cost optimization would help to identify which other areas of the system should be addressed while the system was down for maintenance.

While we could engage in an extensive discussion of the relative merits of different factors to include when considering the extended costs (and resulting benefits to be derived thereof) for various workscope options, the reality is that it would be impossible to account for every conceivable cost and benefit for every possible situation and type of system this approach might be applied to. Those factors may be applicable and suitable in some situations and not for others. Therefore, we will leave it to those implementing this approach to decide for themselves whether those other factors should be included or not, and how to account for them.

9. SELECTING AN OPTIMUM WORKSCOPE

NOTE: This section will continue to refer to the aircraft engine example discussed earlier.

After the obvious question of "*What's wrong with it?*", the second question the technician must answer when an aircraft engine is brought in for repair is: "*At this time, what other work should be performed beyond what's needed to repair the primary failure?*" Unless a lot of diagnosis or investigation is required, the decision of what needs to be done to fix the primary problem is usually pretty straightforward. However, determining what else should be done while the engine is in the shop is usually much more nebulous. This is sometimes referred to as the "Maintainer's Dilemma".

There is typically no clear guidance for determining the maintenance workscope, and it's often based only on policy, common practice, or 'gut instinct'. And there is no "Bang-for-the-Buck" analysis behind any of these that takes into account all of the variables involved. As a result, whatever work is done is usually **not** the most cost-effective. When not enough maintenance is done, areas that could be addressed are not touched making the engine more likely to fail in the near future, resulting in the engine being returned to the shop for yet another repair too soon. Conversely, if too much work is done, areas are refurbished that don't appreciably improve the engine's overall reliability, resulting in time and money being expended for maintenance that really wasn't needed and increasing the cost of that shop visit. Somewhere in between these two extremes, there is an "optimum" workscope.

Determining an optimum workscope starts by defining what is meant by the "goodness" of the maintenance that is performed. In other words, it would be the "effectiveness" portion of a cost-effectiveness evaluation. Since the purpose of maintenance is to help keep a system operational and to correct problems when they occur, effectiveness could be expressed as how long the system would be expected to stay operational before it would fail again, or how quickly the system could be made operational again (as a measure of system availability). For the purposes of this discussion, let effectiveness be measured as the life expectancy for an aircraft engine given the maintenance performed, otherwise referred to as the "Expected Time On-Wing" (ETOW).

Recall that the definition of something being an "optimum" involves balancing the influence of competing variables, such as finding a minimum or maximum point on a curve, there needs to be another parameter to compare against the basis – a "y" to compare to the "x". In this context it would be a term that establishes the basis against which "goodness" can be measured. The basis for establishing an optimum will vary from one industry to another, one customer to another, and possibly one maintenance event to another. Typically, it will take the form of the lowest long-term cost per operating hour (or cycle, calendar day, mile, units produced, etc.), lowest total down-time, highest availability, or something similar. Whatever the situation, the asset owner, operator, and/or maintainer would have to decide which definition best suits their particular needs.

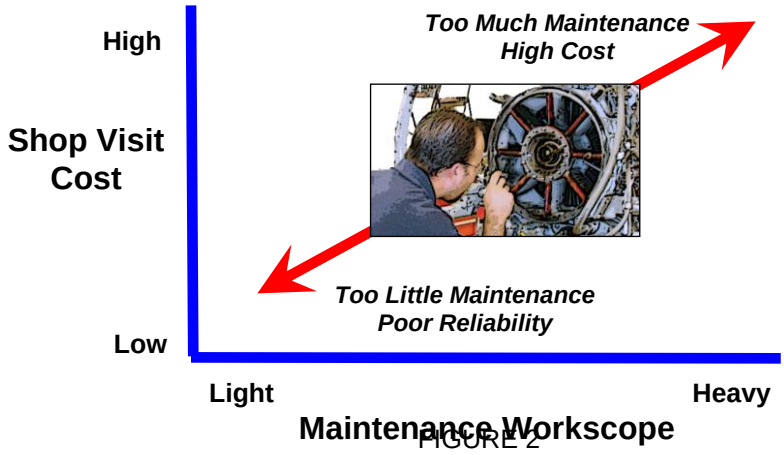

FIGURE 2

A term that is commonly used in the aviation industry is also the most appropriate basis for measuring an optimum aircraft engine workscope – the "Cost Per Flying Hour" (CPFH). This would be calculated by dividing the expected cost of a repair by the life expectancy (ETOW) of the engine once the repair was performed.

If values for both ETOW and CPFH can be determined for each possible maintenance workscope, then an optimum workscope can be selected that would yield the lowest overall CPFH. Let's look at how ETOW might be estimated.

9.1 Constructing a System Reliability Model

The next step is to identify and characterize (mathematically) the dominant failure modes, if it hadn't already been done as part of an RCM analysis. First, this would mean collecting failure data and performing Weibull analysis on as many of the observed failure modes of the system as possible. This also means that operating time (or cycle) data for all components in the system needs to be established. If this information is not readily available, there are techniques for approximating that data (which won't be discussed here). Keep in mind that it is often unnecessary to characterize every possible failure mode for each and every piece of equipment in the system – the 80-20 rule is often enough. In other words, characterizing the top 20% of the failure modes that cause 80% of the maintenance events should be sufficient in most situations. Most of the remainder will probably constitute 'rare event' failures for which there will be very few data points (due to a low rate of occurrence) and may be treated as a "miscellaneous" failure category.

This needs to be followed by an assessment of other like or similar equipment in order to collect or generate suspension data for use in the Weibull analysis. If the failure distribution data is available from the equipment manufacturer, that may be sufficient, but keep in mind that such a failure distribution would likely be an overall industry average, rather than one that reflects the specific operating conditions present in your particular situation.

Once the characteristics for all of the dominant failure modes have been established, an overall system reliability model must then be developed. At a minimum, the model must include all of the dominant failure modes, along with a means of accounting for all remaining maintenance events (such as a 'miscellaneous' category), so that the system reliability function provides as much of an accurate representation of the system's failure behavior as possible. The model must also be able to mathematically tie together the operating times of each applicable piece of equipment that the failure modes apply to. As different repairs are contemplated, those operating times can be adjusted for the components whose reliability has been restored as part of the repair, as well as accounting for components returned to service time-continued (i.e., where their reliability is not fundamentally different than when they were removed from service). Once built, this reliability model will be used to compute the system's conditional probability of failure once it has been repaired.

This bears repeating – the system model needs to account for the differences in operating times (since each component was new or their reliability was restored through refurbishment) so that the conditional probability of failure for each component can be computed individually. If the system reliability model includes failure distributions for components that are not or cannot be time-tracked, then the conditional probability of failure for components that have accrued time in service cannot be determined. This capability is critical to being able to compute the system's overall conditional probability of failure for every combination of workscope that's being considered for the system.

There are many commercial off-the-shelf software tools available that can aid in the construction of a mathematical system reliability model. One approach would be to use a software package to build a Reliability Block Diagram (RBD) model. Using the software, it is possible to build an RBD model for nearly any kind of system you might want to represent. The software can even be used to verify and validate the failure mode interdependencies so that the proper mathematical relationships can be established. Each failure mode (or possibly a group of failure modes) would be represented by a single block, where the 'life used' for each block would be tied to the operating time of the applicable piece of equipment that the failure mode applies to.

An RBD model is a valuable tool to (1) clarify the interrelationships (from a reliability standpoint) between components in the system, (2) ensure the system model reasonably reflects the overall reliability behavior of the real-world system, and (3) identify which components of the system need to be time tracked so the system's future reliability can be estimated.

Once the system model has been validated, constructing the mathematical formula from the RBD model simply becomes an exercise in decomposing the RBD representation into its individual pieces and relating the component failure distributions accordingly (there are a number of Reliability Analysis references that can guide you through this process).

Once an overall system reliability model has been built, where it knows which modules and components are actually installed (and the operating times for each of them), an overall system reliability function can be defined. If the system can be represented as a collection of independent failure modes with no redundancy, then the system reliability function would be the product of the individual reliability functions for each failure mode (a.k.a. a "series" reliability model) which would be represented as follows:

$$R_s(t) = \prod_{i}^{n} R_i(t) \qquad \text{(Eq. 1)}$$

where $\boldsymbol{R_i}$ is the reliability function for $\boldsymbol{i}^{th}$ component. If the system reliability function needs to take into account the various operating times of individual components (as is usually the case for any system that includes a mixture of new and used components), the system reliability function would become:

$$R_s(t) = \prod_{i}^{n} R_i(t \mid t_i) \qquad \text{(Eq. 2)}$$

where $\boldsymbol{t_i}$ is the time accrued by the $\boldsymbol{i}^{th}$ component as of the current maintenance event, unless the $\boldsymbol{i}^{th}$ component will have its reliability restored as part of the planned workscope (either because (a) that component has failed, (b) there is another requirement to refurbish that component, or (c) refurbishment is being considered as an option for reliability or financial reasons), in which case $\boldsymbol{R_i}$ degenerates to the unconditional form. The $\boldsymbol{R_i}$ functions are commonly expressed as Weibull, Lognormal, Normal, or Exponential failure distributions, although other mathematical functions could be used as well.

More complex systems with redundancy or interdependencies between components and/or failure modes will entail a more complex system reliability function than the ones shown above. A number of good technical texts have been published that delve into the mathematics of these more complex systems and should be referred to when developing the underlying mathematics for this type of analysis.

9.2 Expected System Life After Repair

Once the overall system reliability function is defined, it then becomes possible to determine the expected service life after repair for a specified workscope (denoted by the term $\boldsymbol{E(t)}$) by integrating this system reliability function using the conditional reliability for each of the major components installed in the system. By methodically considering all of the different combinations of which components to re-use, repair, or replace, it is possible to determine the impact to the overall system reliability after repair for each combination. Invariably, some workscopes will result in a much higher expected life compared to other workscopes.

The general formula for determining the expected life of a system $E_s(t)$ is given by:

$$E_s(t) = \int_0^\infty R_s(t)dt \quad \text{(Eq. 3)}$$

where the function $R_s(t)$ is the system reliability function, as described above. In general practice, it's unlikely that the $R_s(t)$ function will exist in a form that could be solved analytically, but a numerical integration routine to calculate the value of $E_s(t)$ should not be too difficult to construct. The value of $E_s(t)$ then represents the estimated amount of time that the system can be expected to function (on average) given the repairs that comprise a particular anticipated workscope.

9.3 Estimated Cost of Each Workscope

Just as there will be differing life expectancies for each of the different possible workscopes considered above, there will be different expected costs as well. Some workscopes will drive the replacement of parts that may not have failed, just because of OEM-defined maintenance policies or due to requirements established by an applicable regulatory agency (such as the FAA), resulting in an additional cost. As the extent of maintenance increases, so do the chances of incurring additional costs to correct previously unknown problems that must be addressed once discovered, sometimes referred-to as "Sunshine Costs". In other words, "Sunshine Costs" are costs incurred to correct conditions discovered only when a part is exposed and a fault with that part becomes visible to, and detected by, the technician. Replacing parts that haven't yet failed also means that some of the useful life in those discarded parts will be wasted; in some cases, that residual value should also be included in the cost calculations.

In addition to the cost of the replacement parts and the associated labor (anticipated or not), there are the costs associated with depreciation over the expected time in service, and "risk" costs (that address the costs that might be incurred if a repair is not performed) that may need to be included in the cost calculations.

10. COST OPTIMIZATION EXAMPLE #1 – REPAIR OF AN AUTOMOBILE ENGINE

NOTE: The following example is intended to merely illustrate the concepts of maintenance optimization. Some of the specific details shown below on how to perform the calculations are not done using the most rigorous mathematical methods but are used to convey the basic principles involved.

Scenario: A car with about 55,000 miles is brought to a repair shop to have its failed water pump replaced. The car is equipped with an in-line 4-cylinder overhead cam engine. In order to replace the pump, a significant portion of the engine around the pump has to be disassembled. The repair technician then asks if he should replace the timing belt as well, given that it is a known failure item and it would involve only a small amount of additional effort to reach the timing belt.

This example is exactly the kind of question that maintenance (or in this case, cost) optimization is designed to answer. The basics of the cost optimization decision process would work like this:

It would probably be a valid assumption that timing belt failures follow a wear-out pattern (using a Weibull failure distribution, the shape parameter "β" would be somewhat greater than 1). Also, let's assume the typical life expectancy of a timing belt (comparable to the Weibull characteristic life, η) is ~90,000 miles. If the belt were to fail in service, let's assume it could cost anywhere from ~$1,000 to $3,000 in collateral damage to the engine (from pistons hitting valves that are stuck in the open position, etc.). The cost of a new belt is ~$90 (or ~$1 per 1000 miles of service life), and the cost of making a separate visit to replace only the belt would be ~$350 for labor + the cost of the belt itself.

The impact of replacing the timing belt now would be that it avoids a future shop visit (time & inconvenience), it reduces the chance of an in-service failure, but it does result in throwing away >⅓ of the useful life of the belt.

The costs to defer the belt replacement would add-up as follows:

- The cost of replacing the water pump ≈ $500
- Risk cost of in-service failure = $1-3k x F(t + dt | t)

If we assume the conditional probability of failure is ≈ 5%, this would equate to ≈ $100 Risk Cost.

On the next shop visit, the costs would be:

- Separate visit to replace the belt ≈ $350
- Cost to replace the belt at next visit = $90
- Some sunshine risk cost = (Probability of finding an unexpected problem) x (cost to repair such a problem)

If we assume the chance of finding an unexpected problem is ≈ 5%, and the average cost to correct such a problem in this area is ≈ $500, this would equate to a sunshine risk cost of ≈ $25

The total expected cost of the option to defer replacement of the timing belt is approximately $600 (the cost of replacing the water pump plus the cost associated with the risk that the belt could fail sooner while it is still in service). If we assume the expected life until the subsequent shop visit to replace the belt is 35,000 miles, this equates to a long-term cost of $17.14 / 1000 miles. Even if it were based on the expected life after the belt was replaced, then it would be $1065 (the cost of this shop visit plus the next) divided by 35,000 + 90,000 miles or about $8.52 / 1000 miles.

The costs associated with replacing the timing belt now would be:

- The cost of replacing the water pump ≈ $500
- Residual value of old belt ≈ $35
- Cost to replace the belt now = $90
- Additional labor to replace belt now ≈ $50
- Some sunshine risk cost = (Probability of finding an unexpected problem) x (cost to repair such a problem)

If we assume the chance of finding an unexpected problem is ≈ 5%, and the average cost to correct such a problem in this area is ≈ $500, this would equate to a sunshine risk cost of ≈ $25

Therefore, the total expected cost of the option to replace the timing belt now would be about $700. With an expected life of 90,000 miles, this yields a long-term operating cost of $7.78 / 1000 miles – which is the lowest cost option in the long run.

The decision is ultimately based on a comparison between the expected cost of a minimal repair divided by the expected life for that workscope vs. the cost of replacing the timing belt divided by the expected life after replacing the belt. This would be expressed as follows:

$$\frac{E_{Cost}(Minimal\ Repair)}{E_{Life}(Minimal\ Repair)} \quad vs. \quad \frac{E_{Cost}(Belt\ Replacement)}{E_{Life}(Belt\ Replacement)} \qquad \text{(Eq. 4)}$$

The term that yields the lower value (on the left or the right side of the expression above) would be the optimum choice between these 2 options.

11. COST OPTIMIZATION EXAMPLE #2 – REPAIR OF AN AIRCRAFT ENGINE

Now it's not hard to understand that the process is much more involved for complex systems like an aircraft or aircraft engine, but the basic elements are the same, it just needs to be more methodical and structured since there are so many more options to consider. It also benefits from a graphical representation of all the different combinations of possible workscopes and the corresponding long-term costs.

Recall the Maintainer's Dilemma which shows shop visit cost vs. workscope. The problem is that this chart doesn't identify an optimum workscope, so it can't answer the maintainer's question – *What other repairs should be done at this time?*