



PLANT ENGINEERING AND MAINTENANCE
A CLIFFORD/ELLIOT PUBLICATION
Volume 23, Issue 6

The

Reliability HANDBOOK

From downtime
to uptime
— in no time!

John D. Campbell, editor

Global Leader, Physical Asset Management

PricewaterhouseCoopers LLP.

The evolution of reliability

Take stock of your operation

Is RCM the right tool for you?

The problem of uncertainty

Optimizing time based
maintenance

Optimizing condition
based maintenance

Reliability HANDBOOK

CONTENTS

INTRODUCTION

5 Reliability: past, present, future

by John D. Campbell

Establishing the historical and theoretical framework of RCM

CHAPTER ONE

7 The evolution of reliability

by Andrew K.S. Jardine

How RCM developed as a viable maintenance approach

CHAPTER TWO

9 Take stock of your operation

by Leonard Middleton and Ben Stevens

Measuring and benchmarking your plant's reliability

- Benefits of benchmarking
- What to measure?
- Relating maintenance performance measures to the business objectives
- Excess capacity (cost-constrained)
- Capacity-constrained businesses
- Compliance with requirements
- How well are you performing?
- Findings and sharing results
- External sources of data
- Researching secondary data
- General benchmarking considerations
- Internal benchmarking
- Industry-wide benchmarking
- Benchmarking with comparable industries

CHAPTER THREE

23 Is RCM the right tool for you?

by Jim V. Picknell

Determining your reliability needs

- The 7-step RCM process
- The RCM "product"
- What can RCM achieve?
- What does it take to do RCM?
- Can you afford it?
- Reasons for failure of RCM
- "Flavours" of RCM
- Capability-driven RCM
- How do you decide?
- RCM decision checklist

CHAPTER FOUR

39 The problem of uncertainty

by Murray Wiseman

What to do when your reliability plans aren't looking so reliable

- The four basic functions
- Summary
- Typical distributions
- An example
- Real life considerations — the data problem
- Censored data or suspensions

CHAPTER FIVE

49 Optimizing time based maintenance

by Andrew K.S. Jardine

Tools for devising a replacement system for your critical components

- Enhancing reliability through preventive replacement
- Block replacement policies
- Statement of problem
- Result
- Age-based replacement policies
- When to use block replacement over age replacement
- Setting time based maintenance policies

CHAPTER SIX

57 Optimizing condition based maintenance

by Murray Wiseman

Getting the most out of your equipment before repair time

- Step 1: Data preparation
- Events and inspections data
- Sample inspection data
- Cross graphs
- Cleaning up the data
- Data transformations
- Step 2: Building the proportional hazards model
- Step 3: Testing the PHM
- Step 4: The transition probability model
- Discussion of transition probability
- Step 5: The optimal decision
- Step 6 Sensitivity Analysis
- Conclusion

APPENDIX

71 Searching the Web for reliability information

by Paul Challen

Looking for useful Internet sites? Here's where to start

- Reliability Analysis Center
- Book and print material sources
- Professional organizations
- General information

A word from PEM

Our tradition of publishing annual, fact-filled *PEM* handbooks continues with this issue, *The Reliability Handbook*. As soon as the ink was dry on last year's *MRO Handbook*, we went back to John Campbell and his team of experts at PricewaterhouseCoopers to see if they could provide our readers with current and to-the-point information on the subject of reliability-based maintenance, along with the tips they'd need to put this information into action. Well, the team at PWC came through with flying colours, and what you'll see on the next 72 pages represents the cutting edge of reliability research and implementation techniques from a firm that's one of the world's leading providers of this kind of information to plant professionals around the world. All of us at *PEM* hope you enjoy it, and use it well! — **Paul Challen**

ABOUT THE CONTRIBUTORS:

JOHN D. CAMPBELL is a partner in PricewaterhouseCoopers and director of the firm's maintenance management consulting practice. Specializing in maintenance and materials management, he has more than 20 years of worldwide experience in the assessment/implementation of strategy, management and systems for maintenance, materials and physical asset life-cycle functions. He wrote the book *Uptime: Strategies for Excellence in Maintenance Management* (1995), and is co-author of *Planning and Control of Maintenance Systems: Modeling and Analysis* (1999). You can reach him at 416-941-8448, or by e-mail at john.d.campbell@ca.pwcglobal.com.

ANDREW JARDINE is a professor in the Department of Mechanical and Industrial Engineering at the University of Toronto and principal investigator in the Department's Condition-Based Maintenance Laboratory where the EXAKT software has been developed. He also serves as a senior associate consultant in the International Center of Excellence in Maintenance Management of PricewaterhouseCoopers. Dr. Jardine wrote the book *Maintenance, Replacement and Reliability*, first published in 1973 and now in its sixth printing. You can reach him at 416-869-1130 ext. 2475, or by e-mail at andrew.k.jardine@ca.pwcglobal.com.

JAMES PICKNELL is a principal in PricewaterhouseCoopers Maintenance Management Consulting Center of Excellence. He has more than twenty-one years of engineering and maintenance experience including international consulting in plant and facility maintenance management, strategy development and implementation, reliability engineering, spares inventories, life cycle costing and analysis, benchmarking for best practices, maintenance process redesign and implementation of Computerized Maintenance Management Systems (CMMS). You can reach him at 416-941-8360 or by e-mail at James.V.Picknell@ca.pwcglobal.com.

LEN MIDDLETON is a principal consultant in PricewaterhouseCoopers' Physical Asset Management consulting practice. He has more than twenty years of professional experience in a variety of industries, including an independent practice of providing project management and engineering services. Project experience includes a variety of technical projects in existing manufacturing sites and green-field sites, and projects involving bringing new products into an existing operating plant. You can reach him at 416 941-8383, ext. 62893, or by e-mail at leonard.g.middleton@ca.pwcglobal.com

BEN STEVENS is a managing associate in PricewaterhouseCoopers' Maintenance Management Consulting Centre of Excellence. He has more than thirty years of experience including the past twelve dedicated to the marketing, sales, development, justification and implementation of Computerized Maintenance Management Systems. His prior experience includes the development, manufacture and implementation of production monitoring systems, executive-level management of maintenance, finance, administration functions, and management of re-engineering efforts for a major Canadian bank. You can reach him at 416-941-8383 or by e-mail at ben.stevens@ca.pwcglobal.com.

MURRAY WISEMAN is a principal consultant with PricewaterhouseCoopers, and has been in the maintenance field for more than 18 years. He has been a maintenance engineer in an aluminum smelting operation, and maintenance superintendent at a large brewery. He also founded a commercial oil analysis laboratory where he developed a Web-enabled Failure Modes and Effects Criticality Analysis system, incorporating an expert system and links to two failure rate/ mode distribution databases at the Reliability Analysis Center. You can contact him at 416-815-5170 or by e-mail at murray.z.wiseman@ca.pwcglobal.com.



PLANT ENGINEERING AND MAINTENANCE
A CLIFFORD/ELLIOT PUBLICATION
Volume 23, Issue 6 December 1999

EDITOR
Paul Challen
pc@pem-mag.com
PRESIDENT/PUBLISHER
George F.W. Clifford
PUBLISHER
Joanna Malivoire

EDITORIAL DIRECTOR
Jackie Roth
ASSOCIATE EDITORS
Todd Phillips
Nathan Mallet

PRODUCTION/OPERATIONS EDITOR
David Berger, P.Eng. (Alta.)

CONTRIBUTING EDITORS
Wilfred List
Ken Bannister

ART DIRECTION
Ian Phillips

NATIONAL SALES MANAGER
Joanna Malivoire

DISTRICT SALES MANAGER
Julie Clifford

DISTRICT SALES MANAGER
Alistair Orr

MARKETING SERVICES COORDINATOR
Corina Horsley

PRODUCTION MANAGER
Christine Zulawski

EDITORIAL PRODUCTION COORDINATOR
Nicole Diemert

ASSISTANT ART DIRECTOR
Julie Bertoia

CIRCULATION MANAGER
Janice Armbrust

GENERAL MANAGER
Kent Milford



CLIFFORD/ELLIOT

MRO Handbook is published by Clifford/Elliott Ltd., 209-3228 South Service Road., Burlington, Ontario, L7N 3H8. Telephone (905) 634-2100. Fax 1-800-268-7977. Canada Post - Canadian Publications Mail Product Sales Agreement 112534. International Standard Serial Number (ISSN) 0710-362X. Plant Engineering and Maintenance assumes no responsibility for the validity of the claims in items reported.

*Goods & Services Tax Registration Number R101006989.



Reliability: past, present, future

Establishing the historical and theoretical framework of RCM

by John D. Campbell

Not all that long ago, equipment design and production cycles created an environment in which equipment maintenance was far less important than run-to-failure operation. Today, however, condition monitoring and the emergence of Reliability Centred Maintenance have changed the rules of the game.

At the dawn of the new millennium, it is fitting that this edition of *PEM's* annual handbook should be a discussion on reliability management. A half a millennium ago, Galileo figured that the earth revolved around the sun. A thousand years ago, efficient wheels were made to revolve around axles on chariots. Four thousand years ago, the loom was the latest engineering marvel. A little closer to the present day, maintenance management has evolved tremendously over the past century.

The maintenance function wasn't even contemplated by early equipment designers, probably because of the uncomplicated and robust nature of the machinery. But as we've moved to built-in obsolescence, we have seen a progression from preventive and planned maintenance after WWII, to condition monitoring, computerization and life cycle management in the 1990s. Today, evolving equipment characteristics are dictating maintenance practices, with predominant tactics changing from run-to-failure, to prevention and now to prediction. We've come a long way!

Reliability management is often misunderstood. Reliability is very specific — it is the process of managing the interval between failures. If availability is a measure of the equipment uptime, or conversely the duration of downtime, reliability can be thought of as a mea-

sure of the frequency of downtime.

Let's look at an example. In case 1, your injection moulding machine is down for a 24-hour repair job in the middle of what was supposed to be a solid five-day run. Its availability is therefore 80 percent (120hr-24hr/120hr). The reliability of the machine is 96 hours (96hr/1 failure).

In case 2, your machine is down 24 times for one hour each time. Its availability is still 80 percent (120hr-24x1hr/120hr), but its reliability is only four hours (96/24 failures)!

The two measures are, however, closely related:

$$\text{Availability} = \frac{\text{Reliability}}{(\text{Reliability} + \text{Maintainability})}$$

where maintainability is the mean time to repair.

One of the most robust approaches for managing reliability is adopting Reliability Centred Maintenance. The recently-issued SAE standard for RCM is a good place to start to see if you are ready to embrace this methodology. Although SAE defines RCM as a "logical technical process...to achieve design reliability," it requires the input of everyone associated with the equipment to make the resulting maintenance program work. Even at that, we are still dealing with a lot of uncertainty.

Dealing with uncertainty in equip-

ment performance is, in some ways, like dealing with people. If, for the past three generations, your forefathers lived to the ripe old age of 95, then there is a strong likelihood you will live into your 90s, too. But there are a lot of random events that will happen between now and then. Despite this randomness, we can use statistics to help us know what tasks to do (and not to do) to maximize our life span, and when to do them. Do exercise three times a week and don't smoke — and by analogy, do vibration monitoring monthly and don't rebuild yearly.

When we put cost into this equation, we start to enter the realm of optimization. For this, there are several modeling techniques that have proven quite useful in balancing run-to-failure, time based replacement and condition based maintenance. Our objective is to minimize costs and maximize availability and reliability.

Our Physical Asset Management team at PricewaterhouseCoopers is pleased to provide this introduction to reliability management and maintenance optimization. If you are interested in a broader discussion on these topics, be sure to read our new book *Maintenance Excellence: Optimizing Equipment Life Cycle Decisions*, published by Marcel Dekker, New York (Spring 2000). ■

The evolution of reliability

How RCM developed as a viable maintenance approach

by Andrew K.S. Jardine

Reliability mathematics and engineering really developed during the Second World War, through the design, development and use of missiles. During this period, the concept that a chain was as strong as its weakest link was clearly not applicable to systems that only functioned correctly if a number of sub-systems must first function. This resulted in “the product law for series systems,” which demonstrates that a highly reliable system requires very highly reliable sub-systems.

To illustrate this, consider a system that consists of three sub-systems (A, B, and C) that must work in series for the complete system to function, as in Figure 1.

In the system, where subsystems have reliability values of 97 percent, 95 and 98 percent, respectively, then the complete system has a reliability of 90 percent. (This is obtained from $0.97 \times 0.95 \times 0.98$). This is not 95 percent, which would be the case under the case under the weakest-link theory.

Clearly, for real complex systems where the design configuration is dramatically more complex than Figure 1, the calculation of system reliability is more complex. Towards the end of the 1940s, efforts to improve the reliability of systems focused on better engineering design, stronger materials, and harder and smoother wearing surfaces.

For example, General Motors extended the useful life of traction motors used in locomotives from 250,000 miles to one million miles by the use of better insulation, high temperature testing, and improved tapered-spherical roller bearings.

In the 1950s, there was extensive development of reliability mathematics by statisticians, with the U.S. Department



Figure 1: Series system reliability

of Defense coordinating the reliability analysis of electronic systems by establishing AGREE (Advisory Group on Reliability of Electronic Equipment).

The 1960s saw great interest in aerospace applications, and this resulted in the development of reliability block diagrams, such as Figure 1, for the analysis of systems. Keen interest in system reliability continued into the 1970s, primarily fuelled by the nuclear industry. Also in the 1960s, there grew an interest in system reliability in the civil aviation industry, out of which resulted Reliability Centred Maintenance (RCM), which is covered in Chapter 3 of this handbook.

Subsequent to the early success of RCM in the aviation industry RCM has become the predominant methodology within enterprises (military, industrial etc.) in the 1970s, 80s and 90s to establish reliable plant operations.

The next millennium will see RCM continuing to play a significant role in establishing maintenance programs, but a new feature will be the focus by maintenance professionals of closely examining the plans that result from an RCM analysis, and using procedures which enable these plans to be optimized.

Other sections (Chapter Five, “Optimizing time-based maintenance”, and Chapter Six, “Optimizing condition based maintenance”) of this handbook address procedures already available to assist in this thrust of maintenance optimization. Both of them cover tools, policies, analysis methods, and data preparation techniques to further the implementation of an RCM strategy.

References

Reliability Engineering and Risk Assessment, Henley and Komamoto, Prentice-Hall, 1981. ■

Take stock of your operation

Measuring and benchmarking your plant's reliability

by Leonard G. Middleton and Ben Stevens

Does your business operate using “cash-box accounting”? That’s where you count your money at the beginning and end of each month. If you have more when month’s end rolls around, then that’s good. If not, well you’ll try to do better next month — or at least until the money runs out. No truly successful large business operates using this model, yet maintenance is often organized and performed without proper measures to determine its impact on the business’s success. The use of performance measurement is rapidly increasing in maintenance departments around the world. This stems from a very simple understanding: You can’t manage what you don’t measure. Performance measurement is therefore a core element of maintenance management. The methodologies for capturing performance are of vital importance, since unreliable measurements lead to unreliable conclusions, which lead to faulty actions.

The benefits of benchmarking

Benchmarking reinforces positive behaviour and resource commitment. It enables faster progress towards goals by providing experience, through the example of the participant and improved “buy-in.” By using external standards of performance, the organization can stay competitive by reducing the risk of being surpassed by competitors’ performance, and by customers’ requirements.

Most importantly, benchmarking meets the information needs of stakeholders and management by focusing on critical, value-adding processes, while achieving an integrated view of the business. Benchmarking is a process that makes demands of all its participants, but it is one of the best ways to find practices that work well within comparable circumstances.

What should you measure?

There is no mystique about performance




Figure 2: Inputs, outputs, and measures of process effectiveness

measurement — the trick is in how to use the results to achieve the actions that are needed. This requires a number of

conditions to be in place — including consistent and reliable data, high quality analysis, a clear and persuasive presenta-




Figure 3: **Maintenance optimising — where to start**

tion of the information, and a receptive work environment (see Figure 2).

In accordance with the theme that maintenance optimization is targeted at the company's executive management and the boardroom, it is vital that the results are shown as a reflection of the basic business equation: Maintenance is a business process turning in-

puts into usable outputs. Figure 3 shows the three major elements of this equation — the inputs, the outputs and the conversion process within examples of performance measures.

Input measures are the resources we allocate to the maintenance process. Output measures are the outcomes of the maintenance process. Measures of

maintenance process effectiveness determine where improvements should be made. Specific measures for reliability include MTBF (reliability), MTTR (maintenance serviceability).

Relating performance measures to business objectives

Three basic business operational scenarios impact the focus and strategies of maintenance. They are:

- 1. Excess operations capacity;
- 2. Constrained by operations capacity;
- 3. Focus on compliance to service, quality or regulatory requirements.

Benchmarking measures must reflect the predominant business operational scenario at the time. While some measures are common to all three scenarios, the maintenance focus must align with the current focus of the organisation. The operational scenario may change due to changes in the economy. For example, a strong economy or a reduction in interest rates could cause an increase in demand of building materials. As a consequence, building materials industries (like gypsum wallboard and brick making) will shift from scenario 1 (excess capacity) to scenario 2 (constrained capacity).

Most of the inputs in Figure 1 are

quite familiar to the maintenance department and readily measured — like manpower, materials, equipment and contractors. There are also inputs that are more intangible, more difficult to measure accurately — such as experience, techniques, teamwork, work history — yet each can have a very significant impact on results.

Likewise some of the outputs are easily recognised and equally easily measured; others are more difficult to measure effectively. As with the inputs, some are intangible, such as the contribution to team spirit that comes from completing a difficult task on schedule. Measures of attendance and absenteeism are very inexact substitutes for these intangibles, and overall indicators of maintenance performance are much too broad to resolve them. Notwithstanding the contribution of these intangibles to the overall maintenance performance mosaic, the focus of this handbook will be on the tangible measurements.

The process of converting the maintenance inputs into the required outputs is the core of the maintenance manager's job — yet rarely is the absolute conversion rate of much interest in itself. Converting manpower hours

consumed into reliability, for example, probably makes little or no sense — until it can be used as a measure of comparison, through time or with another similar division or company. Similarly, the average consumption of materials per work order carries little significance — unless it is seen that, say, press A consumes twice as much repair material as press B for the same production throughput. Indeed, one simple way of reducing the consumption of materials per work order is to split the jobs and therefore increase the number of work orders.

Thus the focus must be on the comparative standing of a company or division, or the improvement in the maintenance effectiveness from

one year to the next. These comparisons highlight another outstanding value of maintenance measurement — namely its use in regular comparisons of progress towards specific goals and targets. This process of benchmarking — through time, with other divisions, or other companies — is increasingly being used by senior management as a key indicator of good maintenance management, and frequently discloses surprising discrepancies in performance. A recent benchmarking exercise (see Figure 4) turned up some interesting data from the pulp industry.

A quick glance at the results shows some significant discrepancies — not only in the overall cost structure, but

	Average US\$	Company X US\$
Maintenance costs per ton output	78	98
Maintenance costs per unit equipment	8900	12700
Maintenance costs as % of asset value	2.2	2.5
Maintenance management costs as % of total maintenance costs	11.7	14.2
Contractor costs as % of total maintenance costs	20	4
Materials costs as % of total maintenance costs	45	49
Total number of work orders per year	6600	7100

Figure 4: Maintenance costs in pulp industry benchmarking survey




Figure 5: Measuring the gap

also in the way “company X” does business — a heavier management structure and far less use of outside contractors, for example. What is clear from this high level benchmarking is that to preserve company X’s competitiveness in the marketplace, something needs to be done. Exactly what, though, requires more detailed analysis.

As one may guess, the number of potential performance measures far exceeds the ability (or the will) of the maintenance manager to collect, analyse and act on the data. An important part, therefore, of any program to implement performance measurement is the thorough understanding of the few, key performance drivers. Maximum

called “production constrained,” and is likely to achieve maximum payoff from focussing on maximizing outputs through reliability, availability and maintainability of the assets.

Excess capacity (cost-constrained) businesses

Typical maintenance financial measures could include:

- 1. Maintenance budget versus expenditures (i.e. predictability of costs);
- 2. Maintenance expenditures cash flow (i.e. impact on ability to pay for expenditures);
- 3. Maintenance expenditures, relative to output (e.g. maintenance costs per production unit);

An important part of any program to implement performance measurement is the thorough understanding of the key performance drivers — with prime emphasis given to the indicators that show progress in the areas that need the most help within a company’s maintenance operations.

leverage should always take top priority, that is, you must first identify the indicators that show results and progress in those areas that have the most critical need for improvement. As a place to start, consider Figure 3 on page 10.

If the business would be able to sell more products or services if their prices were lowered, then the business is said to be “cost-constrained.” Under these circumstances, the maximum payoff is likely to come from concentrating on controlling inputs — i.e. labour, materials, contractor costs, and overheads. If the business can profitably sell all it produces, then it is

- 4. Derivative measures of expenditures describing the actual activities the expenditures are used for: emergency repair, condition based monitoring, corrective planned work, shutdown work, process efficiency, cycle time percentages, or production rate percentages (not absolute production rate).

Capacity-constrained businesses

Typical maintenance productivity or output measures include:

- 1. Overall equipment effectiveness and each of a piece of equipment’s components of availability, production rate, and quality rate, as primary measures;

- 2. MTBF and MTTR, as secondary measures used to analyse problems with respect to availability.

RCM analysis (see Chapter 3) should include all operations bottlenecks and critical equipment (allowing for impact of system or equipment redundancy, or parallel processes), in evaluating tactics.

Compliance with requirements

The critical success factors of the organization often depend heavily on compliance with a set of requirements determined by regulatory agencies or by the customer base. Compliance may apply to the operation, such as effluent monitoring equipment. Regulated utilities, pharmaceutical and health care products, or “prestige” products are examples of businesses whose margins and nature are such that compliance is their most critical operational aspect. Financial measures and output measures remain important, though secondary in focus.

Typical maintenance measures within this operating scenario include:

- 1. Quality rate;
- 2. Availability (e.g. regulatory compliance requirement);
- 3. Equipment or system precision or repeatability.

RCM analysis should, therefore, focus on the critical aspects of the compliance, as required by the critical outside stakeholders.

How well are you performing?

The 10 maintenance process items of Figure 5 (above) result from the benchmarking exercise, and a subsequent analysis. They provide the broad targets and the ability to measure progress towards goal achievement.

Findings and sharing results

“Best practices” are identified by benchmarking organizations, who take into account constraints that may exist. An implementation plan is developed to make the “best practices” part of the benchmarking organization’s maintenance process (see Figure 6). In the spirit of benchmarking, the results of the study are shared with the participants. They receive a report detailing the finding of the benchmarking study, without recommendations specific to the benchmarking organization. This report is critical, as it is the essential value they receive, for their effort expended.

External sources of data

Legal means of getting additional data for performance comparison require




Figure 6: Acting on the results of a benchmarking study

researching for secondary data (described below) and benchmarking. Benchmarking may be one or more of the following:

- 1. internal benchmarking;
- 2. benchmarking within the industry;
- 3. benchmarking with comparable organisations not in the industry.

After obtaining it, a critical issue with external data is how comparable it is. Financial data is particularly difficult to compare, because both financial accounting (including GAAP restrictions) and management cost accounting practices vary according to the organisation's objectives.

Questions you should ask when analyzing accounting data include: Are MRO materials and outside services included in the maintenance budget or purchasing budget? Are replacement-in-kind projects and turnarounds included in the maintenance budget or in the capital budget? In calculating equipment availability, does it consider all downtime, or just unscheduled downtime? The answers will depend upon what the organization is trying to achieve with the calculation.

Researching secondary data

Secondary data is information collected for other purposes. Typically these

could be government-generated reports, industry association reports, or annual reports of publicly-traded companies. The data has a number of limitations. It is likely to be quantitative with little information available to understand the context. It may not be current. Comparable performance measures may not be available because their underlying data were not collected.

General considerations in benchmarking

Benchmarking is the sharing of similar information among benchmarking participants. Benchmarking can be comprehensive, covering the entire business organisation or focused at a particular process or set of measures (see Figure 7).

Benchmarking can take a number of forms. Specific, though usually qualitative information can be acquired through short (15 to 30 minute) telephone surveys. Detailed information can be obtained through comprehensive questionnaires, but participation then becomes an issue. A focused benchmarking study can address specific interests, but the content has to be broad enough to obtain participation (see Figure 8).

Internal benchmarking

The principal difficulty in benchmarking is getting the critical data needed for comparison. Internal benchmarking addresses this issue by comparing data from other organisations and divisions within the company. Information can be freely exchanged since it does not provide an advantage to a competitor. Data exchanged through internal benchmarking programs is typically quantitative because qualitative data requires considerably more analysis.

The limitation (of internal benchmarking) is that knowledge of one's performance relative to the external companies remains unknown. Without outside comparisons, a company is unable to determine whether it is performing at its highest possible capability.

Industry-wide benchmarking

In some industries, there is sharing of data through a third party. It may be




Figure 7: Benchmarking comparators

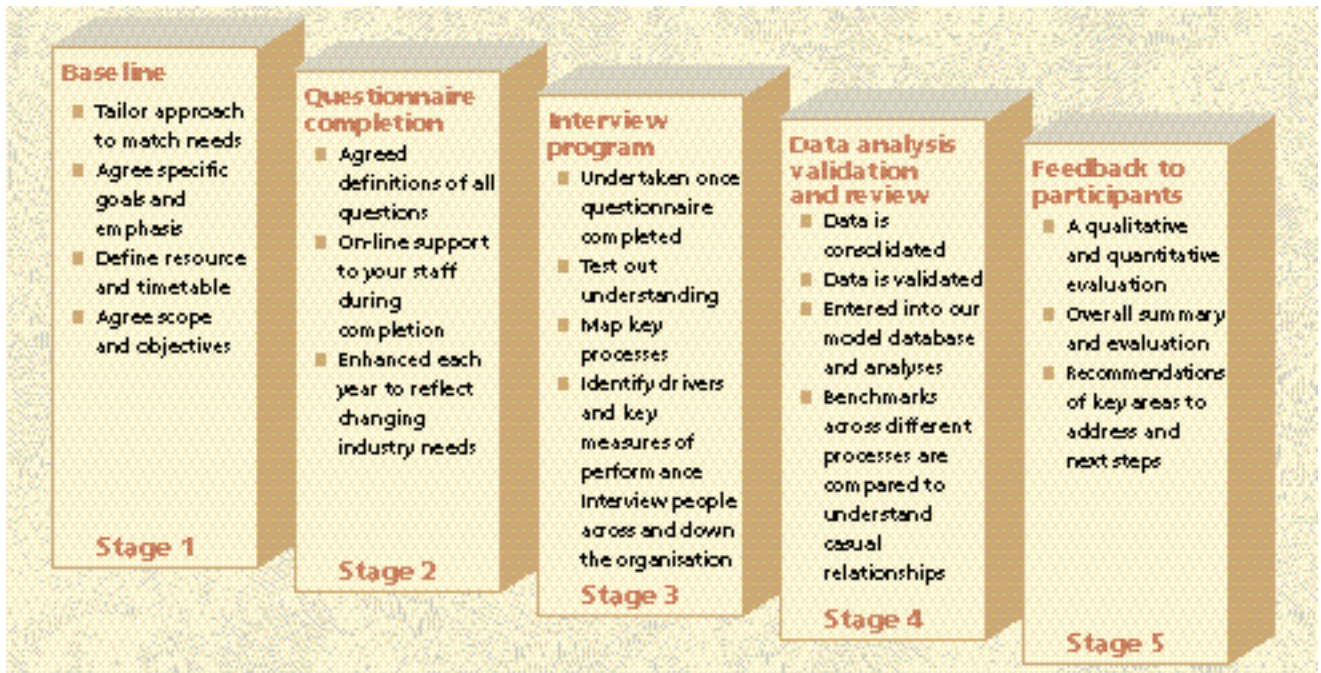


Figure 8: Stages of benchmarking

through an industry association or through an outside party that has developed industry expertise and wants to remain a focus point of the industry (e.g. the PWC Global Forestry benchmarking report). This third party ensures that the data remains confidential and is not directly attributed to any participating organisation. The organisation would also develop the questionnaire, although input from the participants would help provide direction to ensure the results have the highest utility to the participants. It is sometimes difficult to get participation of all the critical parties, as some companies view their operations as a strategic advantage over their competition. If they are indeed better than their competitors and have nothing to learn from them, that would be true — but there is always something to learn.

Benchmarking with comparable industries

Where specific measures are desired by an organization, it is possible to perform a focused benchmarking study. Using an outside third party to maintain confidentiality (that is, who belongs to what data), an organization can measure and compare specific parts of its process with the best practices revealed by the exercise. As it is often difficult to get information from direct competitors, it is usually possible to get information from other industries with similar process issues and constraints. A spin-off benefit of benchmarking with comparable industries may be the discovery of new usable ideas that may not be common practice in one's own industry.

For example, a client in the oil and gas refining industry wished to bench-

mark electrical and instrumentation maintenance management. The participant selection criteria for other organizations were:

- 1. Electrical and instrumentation maintenance is critical to reliable operations;
- 2. Production is a continuous process requiring operating 24 hours a day, 7 days a week;
- 3. Ramifications of unscheduled downtime are severe and there is considerable effort and focus by maintenance to avoid downtime; and
- 4. Maintenance is pro-active.

The list of possible participants that met most or all of the requirements, included chemical sites, electrical power generation sites, waste water treatment plants, steel mills, as well as other oil and gas refining sites.

Maintenance is an essential part of the overall business of an organization, and it must therefore take its lead from the objectives and direction of that organization. Maintenance cannot operate in isolation. The continuous improvement loop that is key to improvement in maintenance, must be driven by and mesh with the corporation's own planning, execution and feedback cycle.

Disconnects frequently occur because of the failure of maintenance to correlate from the corporate to the department level — for example if the company places a moratorium on new capital expenditures, then this must be fed into the maintenance department's equipment maintenance and replace-




Figure 9: The maintenance continuous improvement loop

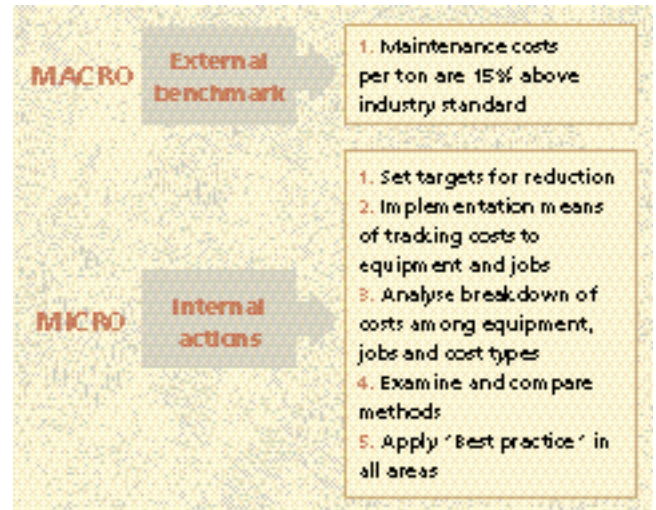


Figure 10: Relating macro measurements with micro tasks

ment strategy. Likewise, if the corporate mission is to produce the highest possible quality product, then this is probably not in synch with a maintenance department's cost minimisation target. This type of disconnect frequently crops up inside the maintenance department itself; if the maintenance department's mission is to be the best performing one in the business, then a strategy which excludes condition-based maintenance and reliability is unlikely to achieve the results (see Figure 9). Similarly, if the strategy statement calls for a 10 percent increase in reliability, then reliable and consistent data must be available to make the comparisons.

Conflicting priorities for the maintenance manager

In modern industry, all maintenance departments face the same dilemma — which of the many priorities is at the top of the list? (And dare one add “this week”?) Should the organization minimize maintenance costs — or maximize production throughput? Does it minimise downtime — or concentrate on customer satisfaction? Should it spend short-term money on a reliability program to reduce long-term costs?

Corporate priorities are set by the senior executive and ratified by the board of directors. These priorities should then flow down to all parts of the organization. The maintenance manager's task is to adopt those priorities, and convert them into the corresponding maintenance priorities, strategies and tactics which will achieve the results; then track them and improve on them.

Figure 10 (above) shows an example of how the corporate priorities can flow down through the maintenance priorities and strategies to the maintenance tactics that control the everyday work of the maintenance department. Hence if the corporate priority is to maximise product sales, then this can legitimately be converted into maintenance priorities that focus on maximizing throughput and therefore equipment reliability. In turn, the maintenance strategies will also reflect this, and could include (for example) implementing a formal reliability enhancement program supported by condition-based monitoring. Out of these strategies, the daily, weekly and monthly tactics flow — providing the lists of individual tasks which then become the jobs that will appear on the work orders from the EAM or CMMS. The use of the work order as the “prompt” to ensure that the inspections get done is widespread; where organizations frequently fail is to ensure that the follow-up analysis and reporting is completed on a regular

and timely basis. The most effective method of doing this is to set them up as weekly work order tasks which are then subject to the same performance tracking as the preventive and repair work orders.

In seeking ways to improve performance, a maintenance manager is confronted with many, seemingly conflicting alternatives. Many review techniques are available to establish where organizations stand in relation to industry standards or best maintenance practice. The best techniques are those which will also indicate the pay-off to be derived from improvement and therefore the priorities. The review techniques tend to be split into the macro (covering the full maintenance department and its relation to the business) and the micro approaches (with the focus on a specific piece of equipment or a single aspect of the maintenance function).

The leading techniques are:

1. Maintenance effectiveness review: This covers the overall effectiveness of the maintenance function and its relationship with the organization's business strategies. These can be conducted internally or externally, and typically cover areas such as:

- Maintenance strategy and communication;
- Maintenance organization;
- Human resources and employee empowerment;
- Use of maintenance tactics;
- Use of reliability engineering and reliability-based approaches to equipment;
- Equipment performance monitoring and improvement;
- Information technology and management systems;
- Use and effectiveness of planning and scheduling;
- Materials management in support of maintenance operations.

2. External benchmark: This draws parallels with other organizations to establish the organizations standing relative to industry standards. Confidentiality is a key factor here, and results are typically presented as a range of performance indicators and the target organization's ranking within that range. Some of the topics covered in benchmarking will overlap with the maintenance effectiveness review; and additional topics include:

- Nature of business operations
- Current maintenance strategies and practices
- Planning and scheduling
- Inventory and stores management practices
- Budgeting and costing
- Maintenance performance and measurement
- Use of CMMS and other IS tools
- Maintenance process re-engineering

3. Internal comparisons: These will measure a similar set of parameters as the external benchmark, but will be drawn from different departments or plants. As such, they are generally less expensive to undertake and, provided the data is consistent, can illustrate differences in the maintenance practices among similar plants. These differences then become the basis for shared experiences and the subsequent adoption of best practices drawn from these experiences.

4. Best practices review: Looks at the process and operating standards of the maintenance department and compares them against the best in the industry.

This is generally the starting point for a maintenance process upgrade program, and will focus on areas such as:

- Preventive maintenance;
- Inventory and purchasing;
- Maintenance workflow;
- Operations involvement;

- Predictive maintenance;
- Reliability based maintenance;
- Total productive maintenance;
- Financial optimisation;
- Continuous improvement.

5. Overall Equipment Effectiveness (OEE): A measure of a plant’s overall operating effectiveness after deducting losses due to scheduled and unscheduled downtime, equipment performance and quality. In each case, the sub-components have been

defined meticulously, to provide one of the few reasonably objective and widely-used indicators of equipment performance. In looking at the following summary of the results from one company (see Figure 11), remember that the category numbers are multiplied through the calculation to derive the final result. Thus, Company Y which achieves 90 percent or higher in each category (which look like pretty good numbers), will only have an

OEE of 74 percent. This means that by increasing the OEE to, say, 95 percent, Company Y can increase its production by 28 percent with minimal capital expenditure $(95-74)/74 = 28\%$. Doing this in three plants prevents the fourth from being built.

	Target	Company Y
Availability	97%	90%
x		
Utilization rate	97%	92%
x		
Process efficiency	97%	95%
x		
Quality	99%	94%
=		
Overall equipment effectiveness	90%	74%

Figure 11: Overall equipment effectiveness

These, then, are some of the high level indicators that serve to provide management with an overall comparison of the effectiveness and comparative standing of the maintenance department. They are very useful for highlighting the key issues at the executive level, but require more detailed evaluation to generate specific actions. They also typically will require senior management support and corporate funding.

Fortunately, there are many measures that can (and should) be implemented within the maintenance department which do not require external approval or corporate funding. These are important to maintainers, as they can be used to stimulate a climate of improvement and progress. Some of the many indicators at the micro level are:

- 1. Benefits realization assessment following the purchase and implementation of a system (EAM) or equipment against the planned results or initial cost-justification;
- 2. Machine reliability analysis/failure rates: targeted at individual machine or production lines;
- 3. Labour effectiveness review: measuring the allocation of manpower to jobs or categories of jobs compared to last year;
- 4. Analyses of materials usage, equipment availability, utilization, productivity, losses, costs, etc.

All of these indicators give useful information about the maintenance business and how well tasks are being performed. The effective maintenance manager will need to be able to select those that most directly contribute to the achievement of the maintenance department’s goals as well as the overall business goals. ■

Is RCM the right tool for you?

Determining your reliability needs

by Jim V. Picknell

In this chapter, we define Reliability Centered Maintenance (RCM) as a “logical, technical process for determining the appropriate maintenance task requirements to achieve design system reliability under specified operating conditions and in the specified operating environment.”

The recently issued SAE Standard, JA1011, “Evaluation Criteria for Reliability-Centered Maintenance (RCM) Processes” outlines a set of criteria with which any process must comply to be called RCM. While this new standard is intended for use in determining if a process qualifies as RCM, it does not specify the process itself. The standard does present seven questions that the process must answer. This chapter describes that process and several variations. You can check the SAE standard for a comprehensive understanding of the complete RCM criteria.

We cannot achieve reliability greater than that designed into systems by their designers. Each component has its own unique combination of failure modes, with their own failure rates. Each combination of components is unique and failures in one component may well lead to the failure of others. Each “system” operates in a unique environment consisting of location, altitude, depth, atmosphere, pressure, temperature, humidity, salinity, exposure to process fluids or products, speed, acceleration, etc. Each of these factors can influence failure modes making some more dominant than others. For example, a level switch in a lube oil tank will suffer less from corrosion than the same switch in a salt water tank. And an aircraft operating in a temperate maritime environment is likely to suffer more from corrosion than one operating in an arid desert.

Technical manuals often recommend a maintenance program for equipment and systems. They sometimes take ac-

count of different operating environments to the extent that they can. For example an automobile manual will specify different lubricants and anti-freeze densities that vary with ambient operating temperature. But they don’t often specify different maintenance actions based on driving style — say, aggressive vs. defensive — or based on use of the vehicle — like a taxi or fleet versus weekly drives to church or to visit grandchildren. In an industrial setting, manuals are not often tailored to your particular operating environment. Your instrument air com-

pressor installed at a sub-arctic location may have the same manual and dew point specifications as one installed in a humid tropical climate. RCM is a method for looking out for your own destiny with respect to fleet, facility and plant maintenance.

The 7-step RCM process

RCM has seven basic steps:

- 1. Identify the equipment / system to be analyzed;
- 2. Determine its functions;
- 3. Determine what constitutes a fail-




Figure 12: RCM process overview




Figure 13: What's at stake when it fails?

ure of those functions;

- 4. Identify the failure modes that cause those functional failures;
- 5. Identify the impacts or effects of those failures' occurrence;
- 6. Use RCM logic to select appropriate maintenance tactics; and
- 7. Document your final maintenance program and refine it as you gather operating experience.

These seven steps are intended to answer the seven questions posed in the new SAE standard. Figure 12 (page 23) depicts the entire RCM process.

In the first step, the RCM practitioner must decide what to analyze. It is the most critical items that require the most attention. There are many possible criteria and a few are suggested here:

- Personnel safety
- Environmental compliance
- Production capacity
- Production quality
- Production cost (including maintenance costs) and;
- Public image.



Figure 14: How complex? Which functions are most easily defined?

When a failure occurs in any system, equipment or device it may have varying degrees of impact on each of these criteria, from "no impact" through "increased risk" and "minor impact" to "major impact". Each of these criteria and impacts can be weighted. Items with the highest combined impact over all criteria should be analyzed first.

The functions of each system are what it does — in either an active or passive mode. Active functions are usually the obvious ones for which we name our equipment. For example a motor control center is used to control the operation of various motors. Some systems also have less obvious secondary or even protective functions. A chemical process loop and a furnace both have a secondary function of containment and may also have protective functions provided by thermal insulating or chemical corrosion resistance properties.

It is important to note that some systems do not perform their active role until some other event occurs, as in safety systems. This passive state makes failures in these systems difficult to spot until it's too late. Each function also has a set of operating limits. These parameters define "normal" operation of the function. When the system operates outside these "normal" parameters, it is considered to have failed. Defining functional failures follows from these limits. We can experience our systems failing high, low, on, off, open, closed, breached, drifting, unsteady, stuck, etc.

Functions are often more easily determined for parts of an assembly than


for the entire assembly. There are two approaches to determining functions that dictate the way the analysis will proceed. One alternative is to look at equipment functions. To think of all the failure modes it is necessary to imagine everything that can go wrong with this fairly high level of assembly. This approach is good for determining the major failure modes but can miss some less obvious. An alternative is to look at "part" functions. This is done by dividing the equipment into assemblies and parts much the same way as if you took the equipment apart. Each part has its own functions and failure modes.

A failure mode is "how" the system fails to perform its function. A cylinder may be stuck in one position because of a lack of lubrication by the hydraulic fluid in use. The functional failure is the failure to stroke or provide linear motion but the failure mode is the loss of lubricant properties of the hydraulic fluid. Of course there are many possible causes for this sort of failure that we must consider in determining the correct maintenance action to take to avoid the failure and its consequences. Causes may include: use of the wrong fluid, the absence of fluid due to leakage, dirt in the fluid, corrosion of the surfaces due to moisture in the fluid, etc. Each of these can be addressed by checking, changing or conditioning the fluid. These are maintenance interventions.

Not all failures are equal. The consequences of failure are its effects on the rest of the system, plant and operating environment in which it is taking place. The failure of the cylinder above may cause excessive effluent flow to a river if it is actuating a sluice valve or weir in a treatment plant — that is, severe impact. The effects may also be as minor as failing to release a "dead-man" brake on a forklift truck that is going to be used for a day of stacking pallets in a warehouse — that is, relatively minor impact. In one case the impact is on the environment and in another it may be only a maintenance nuisance.

By knowing the consequences of each failure we can determine if the failure is worthy of prevention, efforts to predict it, some sort of periodic intervention to avoid it altogether, redesign to eliminate it or no action.

Figure 15 (page 26) graphically depicts the RCM logic. RCM logic helps us to classify failures as being either hidden or not and as having safety or environmental, production or maintenance impacts. To simplify the classical RCM logic diagram we have shown the failure classifying questions nearer the end of






Figure 15: RCM decision logic

the logic tree. Investigations of failure modes reveal that most failures of complex systems made up of mechanical, electrical and hydraulic components will fail in some sort of random fashion – they are not predictable with any degree of confidence.

Many of these failures will however be detectable before they have reached a point where the functional failure can be deemed to have taken place. For example, failure of a booster pump to refill a reservoir that is in use to provide operating head to a municipal water system may not cause a loss of system functionality. It can be detected before we lose municipal water pressure, however, if we are watching for it – that is the essence of condition monitoring. We look for the failure that has already happened but hasn't progressed to the point of degrading system functionality. By finding these failures in this "early" failed state we can avoid the consequences on overall functional perfor-

mance. Since most failures are random in nature RCM logic first asks if it is possible to detect them in time to avoid loss of the function of the system. If the answer is yes then the need for a "condition monitoring" task is the result.

To avoid the functional failure event we must monitor often enough that we are confident in detecting the deterioration with sufficient time to act before the function is lost. For example, in the case of our booster pump we may want to check for its correct operation once a day if we know that it takes a day to repair it and two days for the reservoir to drain. That provides at least 24 hours between detection of a pump failure and its restoration to service without loss of the reservoir system's function. Optimization of these decisions is discussed further in Chapter 7.

If the failure is not detectable in sufficient time to avoid functional failure then the logic asks if it is possible to repair the item failure mode to reduce failure rate.

Some failures are quite predictable even if they can't be detected early enough. For example we can safely predict brake wear, belt wear, tire wear, erosion, etc. These failures may be difficult to detect through condition monitoring in time to avoid functional failure, or they may be so predictable that monitoring for the obvious is not warranted. Why shut down equipment to monitor for belt wear monthly if you know with confidence that it is not likely to appear for two years? You could monitor every year but in some cases it may be more logical to simply replace the belts without checking for their condition every two years. Yes there is a risk that failure occurs early and there is also a risk that the belts will be fine and you are replacing belts that are not worn. These decisions are discussed further in Chapter 6.

If it is not practical to replace components or to restore "as new" condition through some sort of usage or time based action then it may be possible to replace the entire equipment. Usually this makes sense if the loss of function is very critical because this implies an expensive sparing policy to support this approach. Perhaps the cost of lost production in the downtime associated with part replacement is too expensive and the cost of entire replacement spare equipment is less. Again, this sort of decision is discussed further in Chapter 6.

In the case of hidden failure modes that are common in safety or protective systems it may not be possible to monitor for deterioration because the system is normally inactive. If the failure mode is random it may not make sense to replace the component on some timed basis because you could be replacing it with another like component that fails immediately upon installation.

You simply can't tell. In these cases RCM logic asks us to explore functional failure finding tests. These are tests that we can perform that may cause the device to become active, demonstrating the presence or absence of correct functionality. If such a test is not possible you should re-design the component or system to eliminate the hidden failure.

In the case of failures that are not hidden and you can't predict with sufficient time to avoid functional failure and you can't prevent failure through usage or time based replacements you can either re-design or accept the failure and its consequences. In the case of safety or environmental consequences you should re-design. In the case of production related consequences you may chose to redesign or run to failure de-

pending on the economics associated with the consequences. If there are no production consequences but there are maintenance costs to consider you make a similar choice. In these cases the decision is based on economics — that is, the cost of redesign vs. the cost of accepting failure consequences (like lost production, repair costs, overtime, etc.).

Task frequency is often difficult to determine with confidence. Chapters 6 and 7 discuss this problem in detail, but for the purpose of this chapter, it is sufficient to recognize that failure history is a prime determinant. You should recognize that failures won't happen exactly when you predict, so you have to allow some leeway. Recognize also that the information you are using to base your decision upon may be faulty or incomplete. To simplify the next step, which entails grouping similar tasks, it makes sense to pre-determine a number of acceptable frequencies such as daily, weekly, every shift, quarterly, annually, units produced, distances traveled or number of operating cycles, etc. Select those that are closest to the frequencies your maintenance and operating history tells you make the most sense.

After having run the failure modes through the above logic, the practition-

er must then consolidate the tasks into a maintenance plan for the system. This is the final “product” of RCM. When this has been produced the maintainer and operator must continually strive to improve the product. Task frequencies that are originally selected may be overly conservative or too long.

If you experience too many failures that you think you should be preventing then you are probably not performing your proactive maintenance interventions frequently enough. If you never see any of what used to be common failures that have little consequence or your preventive costs are higher than your costs when you did no preventive maintenance then it is possible that you are maintaining the item too frequently. This is where optimization techniques come in. (See Chapters 6 and 7 for a complete discussion.)

The RCM “product”

The output of RCM is a maintenance plan. That document contains consolidated listings with descriptions of the condition monitoring, time or usage based intervention and failure finding tasks, the re-design decisions and the run-to-failure decisions. This document is not a “plan” in the true sense.

It does not contain typical maintenance planning information like task duration, tools and test equipment, parts and materials requirements, trades requirements and a detailed sequence of steps.

In a complex system there may be thousands of tasks identified. To get a “feel” for the size of the output consider a typical process plant that carries spares for only 50 percent or so of its components. Each of those may have several failure modes. That plant probably carries some 15,000 to 20,000 individual part numbers (stock keeping units) in its inventory. That means that there may be some 40,000 parts with one or more failure modes and task decisions.

Fortunately, there are a limited number of condition monitoring techniques available to us and these will cover much of the output task list, because most of failure modes are random. These tasks can be grouped by technique (e.g. vibration analysis) and by location (like the machine room) and by sub-location on a route. It may be possible to group tens and even hundreds of individual failure modes this way so as to reduce the number of output tasks for detailed maintenance planning. It is necessary to watch the

frequencies at which the grouped tasks were specified.

Time- or usage-based tasks are also easy to group together. All the replacement or refurbishment tasks for a single piece of equipment may be grouped by task frequency into a single overhaul task. Similarly, multiple overhauls in a single area of a plant may be grouped into a single shutdown plan.

Another way to group the outputs is by

mately 30 years. It began with studies of airliner failures in the 1960s, to reduce the amount of maintenance work required for what was then the new generation of larger wide-bodied aircraft. As aircraft grew larger and had more parts and therefore more things to go wrong, it was evident that maintenance requirements would similarly grow and eat into flying time which was needed to generate revenue. In the extreme, safe-

craft safety performance has been improving dramatically.

Outside the aircraft industry, RCM has also been used successfully. Military projects often mandate the use of RCM because it allows the end users to experience the sort of highly reliable equipment performance that the airlines experience. The author participated in a shipbuilding project where total maintenance workload on the ship's crew was reduced by almost 50 percent from that experienced on a similarly sized class of ships. At the same time the ship's availability for service was improved from 60 to 70 percent through the reduced requirement for downtime for maintenance intervention.

The mining industry typically finds itself operating in remote locations that are far from sources of parts and materials and replacement labour. Consequently miners want high reliability and availability of equipment — minimum downtime and maximum productivity from the equipment. RCM has been helpful in improving availability for fleets of haul trucks and other equipment while reducing maintenance costs for parts and labour and planned maintenance downtime.

RCM has also been successful in

The success of the airline industry was a highly-visible endorsement of the success of RCM, and showed the world the benefits of an almost entirely proactive maintenance approach.

who does them. Tasks assigned to operators are often done using the senses of touch, sight, smell or sound. These are often grouped logically into daily or shift checklists or inspection rounds checklists.

In the end you should have a complete listing that tells you what maintenance must do, and when. The planner has the job of determining the details of what is needed to execute the work.

What can RCM achieve?

RCM has been around for approxi-

ty could have been very expensive to achieve and could have made flying uneconomical. The success of the airline industry in increasing flying hours, drastically improving its safety record and showing the rest of the world that an almost entirely proactive maintenance approach is possible, all attest to the success of RCM.

New aircraft that had their maintenance determined using RCM required fewer maintenance man-hours per flight hour. Since the 1960s, air-



Figure 16: Aircraft are more reliable now than several years ago, due in part to RCM.

chemical plants, oil refineries, gas plants, remote compressor and pumping stations, mineral refining and smelting, steel, aluminum, pulp and paper mills, tissue converting operations, food and beverage processing and breweries. Anywhere that high reliability and availability is important is a potential application site for RCM.

What does it take to do RCM?

RCM is not a household word (or acronym). It must be learned and practised to attain proficiency and to gain the benefits that can be achieved. Implementation of RCM entails:

- Selecting a willing practitioner team;
- Training them in RCM;
- Teaching other “stakeholders” in the plant operation and maintenance what RCM is and what it can achieve for them,
- Selecting a pilot project to improve upon the team’s proficiency while demonstrating success and
- A roll-out of the process to other areas of the plant.

One key to success in RCM is the demonstration of success. Before the analysis begins, the RCM team should determine the plant baseline measures for reliability and availability as well as proactive maintenance program coverage and compliance. These measures will be used later in comparisons of what has been changed and the success it is achieving.

The team must be multi-disciplinary, and able to draw upon specialist knowledge when it’s needed. It requires knowledge of the day-to-day operations of the plant and equipment, along with detailed knowledge of the equipment itself. This dictates at least one operator and one maintainer. Knowledge of planning and scheduling and overall maintenance operations and capabilities is also needed to ensure that the tasks are truly doable in the plant environment, and senior level operations and maintenance representation is also

needed. Finally, detailed equipment design knowledge is important to the team. This knowledge requirement generates the need for an engineer or senior technician / technologist from maintenance or production, usually with a strong background in either the mechanical or electrical discipline.

The team now numbers five, and experience shows that this is optimum. Too many people will slow progress, and too few means that a lot of time is spent in seeking answers to the many questions that inevitably arise.

Initially, the team will need help to get started. Training can take from a week to a month depending on the approach used. It is usually followed up with the pilot project. The pilot is part of the training that is used to produce a real product.

Training for the team should take about a week. Training of other stakeholders can take as little as a couple of hours to a day or two depending on their degree of interest and “need to know”.

The pilot project time can vary widely depending on the complexity of the equipment or system selected for analysis. A good guideline is to allow for a month of pilot analysis work to ensure the team knows RCM well and is comfortable in using it. Each failure mode can take

about half an hour of analysis time. Using the process plant example from before a very thorough analysis of all systems entailing at least 40,000 items (many with more than one failure mode) would entail over 20,000 man-hours (that’s nearly 10 man-years for an entire plant). When you divide that by five team members you can expect the analysis effort to take up to two years for an entire plant of that size.

Can you afford it?

So, what can you expect to pay for training, software, consulting support, and your staff’s time? You can see from the example that a large process plant will require a lot of effort to analyze. That effort comes at a price. Ten man-years at an average of, say, \$70,000 per person tallies up to \$700,000 for your staff time alone. The training for the team and others will require a couple of weeks from a third-party expert. The expert should also be retained for the duration of the pilot project — and that’s another month.

Several software tools exist that step you through the RCM process and store your answers and results as you produce them. Some of the software can be bought for only a few thousand dollars for a single user license. Some of it comes with the training in RCM and some of it comes as part of large computerized maintenance management systems. Prices for these high-end systems that include RCM are typically in the hundreds of thousands of dollars.

To assist in determining task frequencies it will be necessary to have an understanding of your plant failure histories or to be able to interrogate databases of failure rates. Your plant failure history should be available to you already through your maintenance management system. You may require help in building queries and running reports and that may require time from your programming staff. External reliability databases are available although



Figure 17: RCM can be a lot of work!

not always easily located. Access to them may require a user or license fee.

RCM has experienced a great deal of success and widespread acceptance in some industries where safety and high reliability have been drivers. It has also failed in many other attempts.

Reasons for Failure of RCM

There are many reasons for failure, including, but not necessarily limited to:

- Lack of management support and leadership;
- Lack of “vision” of the end result of the RCM program;
- No clearly stated reason(s) for doing RCM (i.e. it becomes another “program of the month”);
- Lack of the right resources to man the effort especially in “lean manufacturing” environments;
- A clash of RCM’s proactive underpinnings with a traditional and highly reactive plant culture;
- Giving up before it’s complete;
- Continued errors in the process and results that don’t stand up to practical “sanity checks” by dirt-under-the-fingernails maintainers. The wrong team composition or lack of understanding contribute to this one;
- Lack of available information on the equipment/systems under analysis. In reality this need not be a significant hurdle, but it often stops people cold;
- Disappointment occurs when the RCM-generated tasks appear to be the same as those already in the PM program that has been in use for some time. Criticism arises that it’s a big exercise which is merely proving what you are already doing;
- Lack of measurable success early in the RCM program. This is usually because a starting set of measures wasn’t taken, a goal did not exist and no ongoing measurements are taken;
- Results don’t happen quickly enough. The impacts of doing the right type of PM often don’t happen immediately and it takes time for results to show up — typically 12 to 18 months;
- There is no compelling reason to maintain the momentum or even start the program;
- The program runs out of funding; or
- The organization lacks the ability to implement the results of the RCM analysis (e.g. no functional work order system that can trigger PM work orders on a pre-determined basis).

There are a number of solutions to this problem. One which often works is the use of an outside facilitator or consultant. A knowledgeable facilitator can help get the client through the process

and help to maintain momentum. Also a few shortcut methods have been developed to help companies get beyond these problems.

“Flavours” of RCM

In one methodological “flavour” of RCM, logic is used to test the validity of an existing PM program. A drawback is that this approach fails to recognize what you are already missing in your PM program. For example, if your current PM program makes extensive use of vibration analysis and thermographic analysis but nothing else, it may very adequately address failure modes that result in vibrations or heat. But it will miss other failure modes that manifest themselves in cracks, reduction in thickness, wear, lubricant property degradation, wear metal deposition, surface finish or dimensional deterioration, etc. Clearly this program does not cover all possibilities.

In another flavour of RCM, criticality is used to weed out failure modes from ever being analyzed. Typically, the failure modes being ignored either arise in parts of equipment that is deemed to be non-critical or the failure modes and their effects are deemed to be non-critical and the RCM logic is not applied. In these cases the program is disregarding failures because they do not exceed some hurdle rate. The savings arise because analysis effort is reduced. When criticality is applied to the failure modes themselves there is relatively little risk of causing a critical problem. The disadvantage of this approach is that you spend most of the effort and cost getting to a decision to do nothing — remember that you are at step five of seven here. Therefore, relatively little is saved.

When a criticality hurdle rate is applied to equipment, the decisions to reduce analysis can be made before most of the analysis is done (at step one). Some but not necessarily all of the equipment failure modes will be known intuitively to those performing the criticality analysis but they will not be documented at this point. Little effort is expended and considerable program costs can be saved. This method has appeal to companies that are limiting their budgets for these proactive efforts.

This latter flavour of RCM may be entirely acceptable if the consequences of possible failure are known with confidence and are acceptable in production, maintenance, cost, environmental and human terms. For example, many failures have relatively little consequence other than the loss of production and manufacturing process disruption,

along with their associated costs.

Purists may argue that doing anything less than full RCM is irresponsible because without the full analysis you are potentially ignoring real and critical failures, even if inadvertently. This is of course true and it is this very concern that prompted the SAE to develop JA-1011.

Unfortunately, it is also true that many companies suffer from one or many of the reasons for failure described, and possibly others. Without the force of law, RCM standards such as SAE JA-1011 and Nowland and Heap and others are mere guidelines that may or may not be followed, depending on the decisions made at each company. These decisions are often made by those who, although senior, lack a comprehensive RCM knowledge.

Knowledgeable practitioners or enthusiastic supporters of being proactive may recognize that their particular plant suffers from one or more of these potential causes of failure. We should do what we can to mitigate potential consequences and avert risk. As responsible engineers and maintainers we may find ourselves in a position that demands we start slowly and build up to full RCM.

Simply reviewing an existing PM program using RCM logic will accomplish very little. Reviewing only critical equipment does more and does it where it counts the most. Reviewing critical equipment first and then moving down the criticality scale does more again and eventually achieves the full objectives of RCM.

If performing RCM is simply too much to expect realistically in your company, then an alternative approach may be the best you can do and it may be much more achievable. This will also reduce risk by at least some amount and is better than doing nothing at all.

Capability driven RCM

If RCM logic progresses from equipment to failure modes and then through decision logic to a result, can the opposite process flow not address many of the failures even though they may not be clearly identified? Why not reverse the logic, start with the solutions (of which there are a finite number) and look for good places to apply those solutions?

It's worth looking at the following questions as well:

- What's wrong with using existing condition monitoring techniques and extending their use to other pieces of equipment? If you can do vibration analysis on some equipment why not do it elsewhere?
- What's wrong with looking specifically

for wear out type failures and simply deciding to do time based replacements? If you can identify major wear out problem areas why not use this technique?

■ What's wrong with simply operating standby-by (redundant) equipment to ensure that it works when needed thus performing a "failure finding task"?

The answer to all three of these questions is that you do run the risk of over-maintaining some items, you may miss some failure modes and their

maintenance actions due to the lack of rigor and you will miss re-design opportunities. It does however take advantage of your capabilities to perform PM and holds true to RCM principles as a way of building up to full RCM. We call this Capability driven RCM or CD-RCM.

A maintenance practitioner may take steps based on the above which will help to mitigate the consequences of failures. Those steps, when success

is demonstrated, may result in the maintainer gaining sufficient influence to extend his proactive approach to include RCM analysis. This approach is not intended to avoid, or as a shortcut for, RCM but is a preliminary step that will provide positive results that are not inconsistent with RCM and its objectives.

The CD-RCM approach that will accomplish this is:

■ Ensure that your PM Work Order system actually works — that is, that PM work orders can be triggered automatically, the work orders get issued, the work orders get done as scheduled. (If this is not in place, you should stop reading now. You need help beyond the scope of this article);

■ Identify your equipment / asset inventory (this is part of the first step in RCM),

■ Identify the available conditioning monitoring techniques that may be used (which is probably limited by your plant capabilities),

■ Determine the types of failure modes that each of these techniques can reveal;

■ Identify the equipment on which these failure modes are dominant;

■ Decide on appropriate frequencies to perform these monitoring tasks and implement them in your PM work order system;

■ Identify the equipment which has dominant wear-out failure modes;

■ Schedule regular replacement of those wearing components and others that are disturbed in the replacement as time based maintenance using your PM work order system;

■ Identify all your standby equipment and safety systems (alarms, shut-down systems, stand-by redundant equipment, back-up systems, etc.). These are systems and equipment that are normally inactive until some other event occurs to cause their usage to be triggered,

■ Determine appropriate tests to exercise this equipment on a periodic basis so that those failures that can be detected are revealed and then implement them in your PM work order system.

■ Examine failures that are experienced after the maintenance program is put in place to determine the root-cause of the failures so that appropriate action may be taken to eliminate those causes or their consequences.

The result of applying this CD-RCM approach may look like:

■ Extensive use of CBM techniques like: vibration analysis, lubricant / oil analysis, thermographic analysis, visual inspections and some non-destructive testing.

■ Limited use of time based replace-

ments and overhauls.

- In plants having a great deal of redundancy, extensive “swinging” of operating equipment from A to B and back, possibly combined with equalization of running hours.

- Extensive testing of safety systems.

- The systematic capture of information about failures that occur and analysis of that information and the failures themselves to determine root-causes of the failures so that they are eliminated.

While this approach does not achieve the results that full RCM analysis will accomplish, it is founded on RCM principles and will move the organization in the direction of being more proactive. CD-RCM is intended to build upon early success with proven methods targeted where they make sense so that credibility is gained and the likelihood of implementing full RCM is enhanced.

How do you decide?

You can see that RCM is a lot of work and can be expensive to perform. There are alternatives that are less rigorous. You may be faced with the challenge of justifying the costs associated with a full RCM program and be unable to say what sort of savings will arise with any confidence. This is a tough situation to be faced with and each situation will have its own peculiarities and twists and personalities to deal with.

RCM is the most thorough and complete approach you can take to determine the right proactive maintenance approaches to use in achieving high system reliability. It is expensive and time consuming — the results, although impressive, can take time to accomplish. Often this time is sufficient that RCM does not exceed the hurdle rates often called for in the modern world of business investments.

Simply reviewing your existing PM program with an RCM approach is not really an option for a responsible manager — it simply risks missing too much that may be critical and safety or environmentally significant.

Streamlined (or “Lite”) RCM may be appropriate for industrial environments where criticality is recognized and used to guide the analysis efforts using the limited or time resources that can be made available. This will achieve the desired RCM results on a smaller but well targeted sub-set of the failure modes on the critical equipment and systems.

Where RCM investment is not an option, the final alternative is to build up to it using CD-RCM which adds a bit of logic to the old approach of get-

ting a new technology and applying it everywhere. In CD-RCM we take stock of what we can do now, make sure we are applying that as widely as possible and demonstrate success by complying with the new program. After success is demonstrated you can expand the program using that success as “evidence” that it works and produces the desired results. Eventually, RCM can be used to ensure that the program is complete.

RCM decision checklist

You must answer a number of questions and evaluate several alternatives to determine if RCM is right for you. These questions are posed throughout this chapter and are summarized here for quick reference.

- 1. Can your plant or operation sell everything it can produce? If the answer is yes, then high reliability is important and RCM should be considered, and you can skip to question five. If the an-

swer is no then you need to focus on cost cutting measures.

■ 2. Do you experience unacceptable safety or environmental performance? If yes, then RCM is probably for you — skip to question five.

■ 3. Do you already have an extensive preventive maintenance program in place? If the answer is yes you may benefit from RCM if its costs are unacceptably high. If the answer is no you may

■ safety or environmental problems;

■ an expensive and low performing PM program, or;

■ no significant PM program and high overall maintenance costs.

■ 5. RCM is for you. You need to ensure that your organization is ready for it. Do you already experience a “controlled” maintenance environment where most work is predictable and planned and where you can confident-

When full RCM investment is not an option, there are options, like “RCM Lite” or Capability-driven RCM (CD-RCM) that might do the trick for your operation in the short run.

still benefit from RCM if your maintenance costs are high compared with others in your business.

■ 4. Are your maintenance costs high relative to others in your business? If yes then RCM is for you — proceed to question

■ 5. If not, then you probably won’t benefit from RCM, and you can stop here.

At this point you have one or several of the following:

■ a need for high reliability;

ly expect that planned work, like PM and PdM, will get done when scheduled? If yes, you pass the very basic test of readiness — your maintenance environment is under control — and can proceed to question 6. RCM won’t work well if you can’t do it in a controlled environment. If not, then you need help beyond what RCM alone can do for you. Get help in getting your maintenance activities under control first, before going any further.

You need RCM and your organiza-

tion is under control already — you are ready for it. Now you need to get the OK to go ahead with it. If you can simply approve it then go for it. Otherwise:

■ 6. Can you get senior management support for the investment of time and cost in RCM training and piloting and rollout? If yes, then you are ready for RCM and it’s ready for you — so stop here, your decision is made. If not then you need to consider the alternatives to full RCM — proceed to question 7.

■ 7. Can you get senior management support for the investment of time and cost in RCM “Lite” training and piloting? This investment will require about one month of your team’s time (5 persons) plus a consultant for the month. If yes then you should consider the RCM “Lite” method to demonstrate success before attempting to roll RCM out across the entire organization. You can stop here — your decision is made. If not then you are faced with the challenge of proving yourself to your senior management and demonstrating success with a less thorough approach that requires little up front investment and uses existing capabilities. Your remaining alternative here is CD-RCM and a gradual build up of success and credibility to expand on it. ■

The problem of uncertainty

What to do when your reliability plans aren't looking so reliable

by Murray Wiseman

When faced with uncertainty, our instinctive, human reaction is often indecision and distress. We would all prefer that the timing and outcome of our actions be known with certainty. Put another way, we would like all problems and their solutions to be *deterministic*. Problems where the timing and outcome of an action depend on chance are said to be probabilistic or *stochastic*. In maintenance, however, we cannot reject the latter mode, since uncertainty is unavoidable. Rather, our goal is to quantify the uncertainties associated with significant maintenance decisions to reveal the course of action most likely to attain our objective. The methods described in this chapter will not only help you deal with uncertainty but may, we hope, persuade you to treat it as an ally, rather than a foe.

How many of us have been told since childhood that “failure is the mother of success” and that “a fall in the pit is a gain in the wit”. Nowhere is this folk wisdom more valued than in a maintenance department employing the tools of reliability engineering. In such an enlightened environment, failures — an impersonal fact of life — can be leveraged and converted to valuable knowledge followed up with productive action. To achieve this lofty but attainable goal in our own operations we require a sound quantitative approach to maintenance uncertainty. So, let’s begin our ascent on solid ground — the easily conceptualized relative frequency histogram of past failures.

The four basic functions

In this section we discover the Relative Frequency Histogram and the four basic functions: 1) the Probability Density Function; 2) the Cumulative Distribution Function; 3) the Reliability




Figure 18: Monthly failure ages for 48 in-service items

Function; and 4) the Hazard Function.

Assume that a population of 48 items purchased and placed in service at the beginning of the year all fail by November.

List the failures in order of their failure ages as in Figure 18. Group them in convenient time segments, in this case by month. Plot the number of failures in each time segment as in Figure 19 (on page 40). The high bars in the centre of Figure 19, represent the most “popular” (or most probable) failure times. By adding the number of failures occurring prior to

April, that is, 14, and dividing that by the total population of items, 48, we may estimate that the cumulative probability of the item failing in the first quarter of the year is 14/48. The probability that all of the items will fail before November is 48/48 or 1.

By transforming the numbers of failures into probabilities in this way, the relative frequency histogram may be converted into a mathematically more useful form called the probability density function (PDF). To do so, the data are replotted such that the area under the curve represents the




Figure 19: Histogram of failure ages

cumulative probability of failure as shown in Figure 20. (How the PDF plot is calculated from the data and drawn will be discussed more thoroughly in Chapter 5).

The total area under curve of the probability density function $f(t)$ is 1, because sooner or later the item will fail. The probability of the component failing at or before time t is equal to the area under the curve between time 0 and time t . That area is $F(t)$, the cumulative distribution function (CDF). It follows, therefore, that the

remaining (shaded) area is the probability that the component will survive to time t , and is known as the reliability function, $R(t)$. $R(t)$ can, itself, be plotted against time. If one does so, the mean time to failure (MTTF) is the area below the Reliability curve or $\int R(t)dt$ [ref. 1]. From the reliability, $R(t)$ and the probability density function, $f(t)$ we derive the fourth useful function, the failure rate or hazard function, $h(t) = f(t)/R(t)$ which can be represented graphically as in Figure 21. The hazard function is the instantaneous probability of failure at a given time t .

Summary

In just a few short paragraphs we've learned the four key functions in reliability engineering: the PDF or probability density function $f(t)$; the CDF or cumulative distribution function $F(t)$; the reliability function $R(t)$; and the hazard function $h(t)$. Knowing any one of these, we can derive the other three. Armed with these fundamental statistical concepts, we go fourth to do battle with the randomness of failures occurring throughout our plant. Even though failures are random events, we shall, nonetheless, discover how to de-

termine the best times to perform preventive maintenance and the best long run maintenance policies. Having "confidently" (that is to say with a known and acceptable level of confidence) estimated the PDF (for example) we shall call upon it or its sister functions to construct optimization models. Models describe typical maintenance situations by representing them as mathematical equations. That makes it convenient if we wish to optimize the model. The objective of optimization is very often to achieve the lowest long run, overall, average cost of maintaining our production equipment. We discuss and build models in Chapters 5 and 6.

Typical distributions

In the previous section we defined the four key functions which we may apply to our data once we have somehow transformed it into a probability distribution. That prerequisite step of converting or fitting the data is the subject of this section.

How does one find the appropriate failure PDF for a real component or system? There are two different approaches to this problem:

- 1. Curve-fit the failure data obtained

from extensive life testing, or

■ 2. Hypothesize it to be a certain parameterized function whose parameters may be estimated via statistical sampling techniques, and conducting numerous statistical confidence tests. We will adopt the latter approach.

Fortunately, we discover from past failure observations, that the probability density functions, (hence their derived reliability, cumulative distribution, and hazard functions) of real data usually fit one of a number of mathematical formulas, whose characteristics are already familiar to reliability engineers. These known distributions include the exponential, Weibull, Lognormal, and Normal distributions. They can be fully described if one can estimate the their parameters.

For example the Weibull CDF is:

$$F(t) = 1 - e^{-\left(\frac{t}{\eta}\right)^\beta}, \text{ where } t \geq 0$$

The parameters β and η can be estimated from the data using the methods to be described. Through one of the common distributions, once we will have confidently estimated their parameters, we may conveniently process our failure and replacement data. We do so by manipulating the statistical functions




Figure 20: Probability density, cumulative probability and reliability functions

we learned in the previous section to: a) understand our problem; and b) forecast failures and analyze risk to make better maintenance decisions. Those decisions will impact upon the times we choose to replace, repair, or overhaul machinery as well as help us optimize many other maintenance decisions.

The trick is threefold: 1) to collect good data; 2) to choose the appropriate function to represent our own situation, then to estimate the function parameters, and finally; 3) to evaluate the

level of confidence we may have in the resulting model. Modern software makes this process easy and fun. Much more than a toy for engineers, though, reliability software gives us the ability to communicate and share with our management, the common goal of business — to devise and select procedures and policies which minimize cost and risk while maintaining and even increasing product quality and throughput.

The four failure rate functions or hazard functions corresponding to the




Figure 21: Hazard function curves for the common failure distributions

four probability density functions (exponential, Weibull, lognormal, and normal) are shown in Figure 21.

Of the four, the observed data most frequently approximates the Weibull distribution. That's fortunate — and understated by Waloddi Weibull himself who said while delivering his hallmark paper in 1951 that it "...may sometimes render good service". The initial reaction to his paper varied from disbelief to outright rejection. Eventually, as great ideas sink into fertile ground and sprout life, the U.S. Air Force recognized and funded Weibull's research for the next 24 years until 1975. Today, Weibull Analysis is the leading method in the world for fitting life data [ref. 3].

The objective of this chapter is to

bring such decision-making methodologies to the attention of the maintenance professional and to show that they are worthwhile and rewarding endeavours for managing his or her company's physical assets.

An example

Here is a look ahead using an example illustrating how one may extract meaningful information from failure data. Assume that we have determined (using the methods to be discussed later in this chapter and those in chapter 7, that an electrical component has the exponential cumulative distribution function, $F(t) = 1 - e^{-\lambda t}$, where $\lambda = .0000004 \text{ hr}^{-1}$.

We then have to answer the following:

a) What is the probability that one of these

parts fails before 15,000 hours of use?

b) How long do we have to wait to expect 1 percent failures.

$$\text{a) } F(t) = 1 - e^{-\lambda t} = 1 - e^{-.0000004 \times 15000} = 0.001 = .6\%$$

Rearranging

$$\text{b) } t = -\ln(1-F(t))/\lambda = \ln(1-.01)/.0000004 = 25,126 \text{ hr.}$$

c) What would be the mean time to failure (MTTF)?

$$\text{MTTF} = \int_0^\infty R(t)dt = \int_0^\infty e^{-\lambda t} dt = 1/\lambda = 250,000 \text{ hr}$$

d) What would be the median time to failure (the time when half the population will have failed)?

$$F(T_{50}) = 0.5 = 1 - e^{-\lambda T_{50}}$$

$$T_{50} = \ln 2/\lambda = 0.693/.0000004 = 1,732,868 \text{ hr.}$$

This is the kind of information we can expect to get by examining our data using the reliability engineering principals embodied in user friendly software. Read on to discover how.

Real life considerations — the data problem

Ironically, we often let the minimal

data required for reliability management slip through our fingers as we relentlessly pursue the ever elusive control over our plant and production equipments' maintenance costs. Data management is the first step towards successful physical asset management. Good data embodies rich experience from which one may learn — and thereby improve upon — one's current maintenance manage-

continuously, their own maintenance and replacement data with renewed appreciation of its high value to their organization's bottom line.

Without doubt, the first step in any forward looking activity is to get good information. In importance, this step outweighs the subsequent analysis steps which may seem trivial by comparison. The history of humanity is one of progress built on its experi-

tive to world class industry best practices, is the extent to which their data is fed back to guide their maintenance decisions, tactics, and policies.

Here are some examples of decisions based on reliability data management:

■ A maintenance planner notes three in-service failures of a component during a three-month period. The superintendent asks, to help plan for adequate available labour, "How many failures will we have in the next quarter?"

■ To order spare parts and schedule maintenance labour, how many gear-boxes will be returned to the depot for overhaul for each failure mode in the next year?

■ An effluent treatment system requires a regulatory shutdown overhaul whenever the contaminant level exceeds a toxic limit for more than 60 seconds in a month. What level and frequency of maintenance is required to avoid production interruptions of this kind?

■ After a design modification to eliminate a failure mode, how many units must be tested and for how long to verify that the old failure mode has been eliminated, or significantly improved with 90 percent confidence.

■ A haul truck fleet of transmissions is routinely overhauled at 12,000 hours as

Ironically, we often let the minimal data required for reliability management to slip through our fingers, as we're busy pursuing the ever-elusive control over the cost of maintenance in our plants and processes.

ment process. An essential role of upper level managers is to place ample computer resources and scientific methodologies into the hands of trained maintenance professionals who collect, filter, and process data for the expressed purpose of guiding their decisions. The intention of this chapter is to inspire dedicated tradesmen, planners, engineers, and managers to gather, earnestly and

ences. Yet there have been countless moments when opportunity was squandered by neglecting to collect and process readily available data. Today many maintenance departments, unfortunately, fall into that category. That is why one of the important measurements used by PricewaterhouseCoopers' Centre of Excellence in Maintenance Management to benchmark companies rela-

stipulated by the manufacturer. A number of failures occur before the overhaul. By how much should the overhaul be advanced or retarded to reduce average operating costs?

■ The cost in lost production is four times that of the cost of a preventive replacement of a worn component. What is the optimal replacement frequency?

■ Fluctuating values of iron and lead from quarterly oil analysis of 35 haul

the system level, and where warranted, even at the component level. The life data for a given component or system comprise the records of preventive replacement or failure ages. When a tradesmen replaces a component, say a hydraulic pump, one of several identical ones on a complex machine whose availability is critical to the company's operation, he or she should indicate which specific pump failed. Further-

One of the unavoidable problems of managing this kind of data is that at the time of our observations and analysis, not all of the units will have failed. We know the actual ages of the “unfailed” units, but we don’t know their failure ages.

truck transmissions along with the failure times of the 35 units over the past three years are all available in the database. What is the optimal preventive replacement time, given the unit's age today and the latest lab results for iron and lead? (This problem will be examined in Chapter 6.)

Therefore it is, without doubt, worth our while to implement procedures to obtain and record life data at

more he or she should also specify how it failed (the failure mode), for example “leaking” or “insufficient pressure or volume.” Given that the hours of equipment operation are known, the lifetimes of individual critical components can then be calculated and tracked by software. That information will become a part of the company's valuable intellectual asset — the reliability database.

Censored data or suspensions

An unavoidable problem in data analysis is, that at the time of our observation and analysis, not all the units will have failed. We know the age of the currently “unfailed” units, and we know that they are still in the unfailed state, but we do not (obviously) know their failure ages. Some units may have been replaced preventively. In those cases, too, we do not know their age at failure. These units are said to be suspended or right censored. While not statistically ideal, we can still make use of this data since we know that the units lasted at least this long. Good reliability software such as Winsmith for Windows, REL-CODE, and EXAKT described in Chapters 5 and 6, can handle suspended data properly.

References:

- 1. *An Introduction to Reliability and Maintainability Engineering*, Charles E. Ebeling, ISBN 0-07-018852-1, 1997.
- 2. *Systems Reliability and Risk Analysis*, Ernst G. Frankel.
- 3. *The New Weibull Handbook* 2nd edition, Robert B. Abernethy.
- 4. www.barringer1.com. ■

Optimizing time based maintenance

Tools for devising a replacement system for your critical components

by Andrew K.S. Jardine

The goal of this chapter is to introduce tools that can be used to derive optimal maintenance and replacement decisions. Particular attention is placed on establishing the optimal replacement time for critical components (also known as line-replaceable units, or LRUs) within a system.

We'll also take a look at age and block replacement strategies on the LRU level. This enables the software package RelCode to be introduced as a tool that can be used to assist maintenance managers optimize their LRU maintenance decisions. We will also take a look at the optimizing criteria of cost minimization, availability maximization, and safety requirements.

Enhancing reliability through preventive replacement

The reliability of equipment can be enhanced by preventively replacing critical components within the equipment at appropriate times. Just what the best time is depends on the overall objective, such as cost minimization or availability maximization. While the best preventive replacement time may be the same for both cost minimization and availability maximization, this is not necessarily the case. Data first needs to be obtained and analyzed before it is possible to identify the best preventive replacement time. In this chapter several optimizing procedures will be presented that can be used with ease to establish optimal preventive replacement times for critical components.

Block replacement policies

The block replacement policy is sometimes termed the group or constant interval policy since preventive replacement occurs at fixed intervals of time with failure replacements occurring whenever necessary. The policy is illustrated in Figure 22 where C_p and C_f are the total costs associated with preventive and failure replacement respectively. t_p is the fixed interval between preventive replacements. In the figure, you can see that for the first cycle there is no failure, while there are two in the second cycle and none in the third or fourth. As the interval between preventive replacements is increased there will be more failures occurring between the preventive replacements

and the optimization is to obtain the best balance between the investment in preventive replacements and the consequences of failure. This conflicting case is illustrated in Figure 23 for a cost minimization criterion where $C(t_p)$ is the total cost per week associated with a policy of preventively replacing the component at fixed intervals of length t_p , with failure replacements occurring whenever necessary. The equation of the total cost curve is provided in several text books including the one by Duffuaa, Raouff and Campbell [see ref. 1]. The following problem is solved using the software package RelCode, which incorporates the cost model, to establish the best preventive replacement interval.




Figure 22: Block replacement policy




Figure 23: Block policy: optimal replacement time




Figure 24: RelCode output: block replacement




Figure 25: Age-based replacement policy

Statement of problem

Bearing failure in the blower used in diesel engines has been determined as occurring according to a Weibull distribution with a mean life of 10,000 km and a standard deviation of 4,500 km. Failure in service of the bearing is expensive and, in total, a failure replacement is 30

times as expensive as a preventive replacement. We need to determine the optimal preventive replacement interval (or block policy) to minimize total cost per kilometer. What is the expected cost saving associated with the optimal policy over a run-to-failure replacement policy?

Given that the cost of a failure is

\$2000.00, what is the cost per km associated with the optimal policy?

Result

Figure 24 shows a screen capture from RelCode from which it is seen that the optimal preventive replacement time is 4,140 kilometers. In addition the Figure provides much additional information that could be valuable to the maintenance planner. For example:

- The cost saving compared to a run-to-failure policy is: \$ 0.1035/km (55.11%)
- The cost per kilometer associated with the best policy is: \$ 0.0843/km

Age-based replacement policies

The age-based policy is one where the preventive replacement time is dependent upon the age of the component. If a failure replacement occurs then the time clock is reset to zero, which differs from the block replacement policy. Figure 25 illustrates an age-based policy where we see that there is no failure in the first cycle. After the first failure the clock is set to zero, and the component reaches its planned preventive replacement age, t_p . After this second preventive replacement, the component again survives to the planned preventive replacement age.

The conflicting cost consequences associated with this policy are identical to that depicted on Figure 23 except that the x-axis measures the actual age (or utilization) of the item, rather than a fixed time interval. The following problem is solved using the software package RelCode, which incorporates the cost model, to establish the best preventive replacement age.

Statement of problem

A sugar refinery centrifuge is a complex machine composed of many parts and subject to sudden failure. A particular component, the plough-setting blade, is considered to be a candidate for preventive replacement. The policy to be considered is the age-based policy with preventive replacements occurring when the setting blade reaches a specified age. What is the optimal policy to so that total cost per hour is minimized?

To solve the problem the following data have been acquired:

- 1. The labour and material cost associated with a preventive or failure replacement is \$2000;
- 2. The value of production losses associated with a preventive replacement is \$1000 while for a failure replacement it is \$7000;
- 3. The failure distribution of the setting blade can be described adequately by a Weibull distribution with a mean




Figure 26: RelCode output: age-based replacement

life of 152 hours and a standard deviation of 30 hours.

Result

Figure 26 shows a screen capture from RelCode where the optimal preventive replacement age of the centrifuge is 112 hours. Additional key information is also provided in the figure that can

be used by the maintenance planner. For example, we see that the optimal policy costs 45.13 percent of that associated with a run-to-failure policy, and therefore it is clear that preventive replacement is a very worthwhile maintenance tactic. Also, the total cost curve is fairly flat in the region 90 hours to 125 hours, thus providing

flexibility to the maintenance scheduler on when to plan preventive replacements.

When to use block replacement over age replacement

Age replacement may seem to be more attractive than block replacement since a recently installed component is never replaced on a preventive basis. The component always is allowed to remain in service until its scheduled preventive replacement age.

To implement an age-based replacement policy, however, requires that a record is kept of the current age of the component and that if a failure occurs then the expected preventive replacement time is changed. Clearly, for expensive components it will be economically justifiable to monitor the age of an item, and to accept rescheduling of the item's change-out time. For inexpensive components it may be appropriate to adopt the easily implementable block policy, knowing that at some of the regularly performed preventive replacements the working unit being preventively replaced may be quite new due to being installed recently at the time of a failure replacement. This compromise is obviated by modern enterprise asset management




Figure 27: Optimal replacement age: risk based maintenance

software which can conveniently track the component working ages having been advised of the failure and replacements events.

Setting time based maintenance policies

Safety constraints: So far, our discus-

sion has assumed that the objective was to establish the best time to replace a component preventively, such that total cost was minimized.

If the goal is to ensure that the probability of failure before the next preventive replacement does not exceed a particular value, such as five

percent, then the time to schedule a preventive replacement can be obtained from the failure distribution as illustrated in Figure 27. That is, we simply need to identify on the x-axis the time that corresponds to a value of five percent on the y-axis.

Cost minimization and availability maximization: In the above discussion, the objective was cost minimization. Availability maximization simply requires that in the models, the total costs of preventive and failure replacement are replaced by the total downtime associated with a preventive replacement and the total downtime associated with a failure replacement. Minimization of the total downtime is then equivalent to maximizing availability. Readers who wish to work through a variety of problems using the RelCode software may do so by obtaining a demonstration version of the software from the author [ref. 2].

References:

- 1. Duffuaa S.O., Raouff A., Campbell J.D., *Planning and Control of Maintenance Systems*, Wiley 1998.
- 2. Readers can be contact the author via e-mail at andrew.k.jardine@ca.pwcglobal.com. ■

Optimizing condition based maintenance

Getting the most out of your equipment before repair time

by Murray Wiseman

Condition Based Maintenance (CBM) is an obviously good idea. It stems from the logical assumption that preventive repair or replacements of machinery and their components will be optimally timed if they were to occur just prior to the onset of failure. Our objective is to obtain the maximum useful life from each physical asset before taking it out of service for preventive repair.

Translating this idea into an effective monitoring program is impeded by two difficulties.

The first is: how to select from among the multitude of monitoring parameters, those which are most likely to indicate the machine's state of health? And the second is: how to interpret and quantify the influence of the measurements on the remaining useful life (RUL) of the machinery? We'll address both of these problems in this chapter.

The essential questions posed when implementing a CBM program are:

- 1. Why monitor?
- 2. What equipment components to monitor?
- 3. How (what parameters) to monitor?
- 4. When (how often) to monitor?
- 5. How to interpret and act upon the results of condition monitoring?

Reliability Centred Maintenance (RCM) as described in Chapter 3 assists us in answering questions 1 and 2. Additional optimizing methods are required to handle questions 3, 4, and 5. We learned in Chapters 4 and 5 that the way to approach these types of problems is to

build a model describing the maintenance costs and reliability of an item. In Chapter 5 we dealt uniquely with situations in which the lifetimes of components were considered independent random variables, meaning that no other information, other than equipment age, was to be used in scheduling preventive maintenance. CBM introduces new information, called covariates, which influence the probability of failure at time t . Consequently the models of Chapter 5 will be extended to include the impact of these measured covariates (for example, the parts per million of iron in an oil sample, the amplitude of vibration at 2 x rpm, etc.) on the remaining useful life of machinery or their components. The extended modelling method we introduce in this chapter, which takes measured data into account, is known as Proportional Hazards Modelling (PHM).

Since D.R. Cox's [ref. 1] 1972 pioneering paper on the subject of PHM, the vast majority of reported uses of Proportional Hazards Modelling have been for the analysis of survival data in the medical field. Since 1985 there have

been an increasing number of references which include applications to marine gas turbines, motor generator sets, nuclear reactors, aircraft engines, and disk brakes on high-speed trains. In 1995 A.K.S. Jardine and V. Makis at the University of Toronto initiated the CBM Consortium Lab [ref. 2] whose mission was to develop general-purpose software for proportional hazards models analysis. The software was designed to be integrated into the operation of a plant's maintenance information system for the purpose of optimizing its CBM activities. The result in 1997 was a program called EXAKT (produced by Oliver Interactive) which, at the time of this writing, is in its second version and rapidly earning attention as a CBM optimizing methodology. The example and associated graphs and calculations given in this chapter have been worked using the EXAKT program.

Readers who wish to work through the examples using the program may do so by obtaining a demonstration version from the author [ref 3].

We, as maintenance engineers, planners, and managers try to perform Con-

dition Based Maintenance (CBM) by collecting data which we feel is related to the state of health of the equipment or component. These condition indicators (or covariates, as they are called in PHM) may take various forms. They may be continuous, such as operational temperature, or feed rate of raw materials. They may be discreet such as vibration or oil analysis measurements. They may be arithmetic combinations or transformations of the measured data such as rates of change of measurements, rolling averages, and ratios. Since we seldom possess a deep understanding of underlying failure mechanisms, the choices for condition indicators are endless. Without a systematic means of discrimination and rejection of superfluous and non-influential




Figure 28: The actual transition is A-B-C-D and not A-B-D

data, CBM can be far less useful as a maintenance decision tactic than should otherwise be the case. Proportional hazards modelling is an effective approach to the problem of information overload because it distills a large set of basic historical condition and failure data into an optimal decision rec-

ommendation founded upon the equipment's current state of health.

In this chapter we discover the proportional hazards modelling process by describing each step through the use of examples. Statistical testing of various hypotheses along the way is an integral part of the process. That will help us avoid the

Figure 29 Haul truck transmissions inspection and event data

IDENT	DATE	W_AGE	HN	P	EVENT	IRON	LEAD	CALCIUM	MAG
HT-66	12/30/93	0	1	4	B	0	0	5000	0
HT-66	1/1/94	33	1	0	*	2	0	3759	0
HT-66	1/17/94	398	1	0	*	13	1	3822	0
HT-66	2/14/94	1028	1	0	*	11	0	3504	0
HT-66	2/14/94	1028	1	1	OC	0	0	5000	0
HT-66	3/14/94	1674	1	0	*	10	0	4603	0
HT-66	4/12/94	2600	1	0	*	14	2	5067	0
HT-66	4/12/94	2600	1	1	OC	0	0	5000	0
HT-66	5/9/94	2927	1	0	*	7	1	4619	2
HT-66	6/4/94	3522	1	0	*	14	0	4784	2
HT-66	6/4/94	3522	1	1	OC	0	0	5000	0
HT-66	7/4/94	4177	1	0	*	13	0	4517	1
HT-66	8/2/94	4786	1	0	*	9	1	4062	3
HT-66	8/2/94	4786	1	1	OC	0	0	5000	0
HT-66	8/29/94	5392	1	0	*	8	0	4562	3
HT-66	9/26/94	6030	1	0	*	9	3	4409	3
HT-66	9/26/94	6030	1	1	OC	0	0	5000	0
HT-66	10/24/94	6693	1	0	*	14	2	5895	0
HT-66	11/21/94	7319	1	0	*	11	1	4827	0
HT-66	11/21/94	7319	1	1	OC	0	0	5000	0
HT-66	12/19/94	7902	1	0	*	5	0	5313	0
HT-66	1/16/95	8474	1	0	*	8	4	5138	0
HT-66	1/16/95	8474	1	1	OC	0	0	5000	0
HT-66	2/13/95	9108	1	0	*	8	2	5039	3
HT-66	3/13/95	9732	1	0	*	18	5	4050	5
HT-66	3/13/95	9732	1	1	OC	0	0	5000	0
HT-66	4/10/95	10320	1	0	*	25	3	5576	3
HT-66	4/23/95	10524	1	2	EF	25	3	5576	3
HT-66	4/24/95	10524	2	4	B	0	0	5000	0
HT-66	5/8/95	10886	2	0	*	12	0	4584	5
HT-66	5/8/95	10886	2	1	OC	0	0	5000	0
HT-66	6/5/95	11457	2	0	*	14	1	5218	5
HT-66	7/3/95	12011	2	0	*	8	0	4955	7
HT-66	7/3/95	12011	2	1	OC	0	0	5000	0
HT-66	8/1/95	12670	2	0	*	10	1	4830	7
HT-66	8/28/95	13215	2	0	*	12	0	4287	9
HT-66	8/28/95	13215	2	1	OC	0	0	5000	0
HT-66	9/25/95	13834	2	0	*	10	2	5523	10
HT-66	10/23/95	14441	2	0	*	10	1	5413	7
HT-66	10/23/95	14441	2	1	OC	0	0	5000	0
HT-66	11/20/95	14915	2	0	*	8	0	6605	6
HT-66	12/18/95	15523	2	0	*	6	0	6542	5
HT-66	12/18/95	15523	2	1	OC	0	0	5000	0
HT-66	1/15/96	16037	2	0	*	4	0	5377	5
HT-66	2/12/96	16694	2	0	*	11	0	5441	5
HT-66	2/12/96	16694	2	1	OC	0	0	5000	0
HT-66	3/12/96	17134	2	0	*	4	0	5040	0
HT-66	4/9/96	17760	2	0	*	7	1	5349	4
HT-66	4/9/96	17760	2	1	OC	0	0	5000	0
HT-66	5/6/96	18377	2	0	*	9	0	3170	14

trap of blindly following a method without adequate verification of the assumptions and the appropriateness of the model to the situation and to the data.

We divide the problem of optimizing condition based maintenance into six steps:

- 1. Studying and preparing the data.
- 2. Estimating the parameters of the Proportional Hazards Model.
- 3. Testing how “good” the PHM model is.
- 4. Building the transition probability model.
- 5. Making the optimal decision for lowest long run maintenance cost.
- 6. Sensitivity analysis.

Step 1: Data preparation

No matter what tools or computer pro-

grams are available, the modeller should always “look” at the data in several ways [see ref. 4]. For example, many data sets can have the same mean and standard deviation and still be very different — and that can be of critical significance. Maintenance modellers must be involved in their applications and understand the context.

Too little attention has been paid to data collection, notwithstanding the elaborate and powerful computerized maintenance management systems (CMMS) and enterprise asset management (EAM) systems in growing use. Change management strategies promoting education and pride of ownership and the alteration of behaviours and attitudes regarding the recognition of the value of data will inspire its meticulous

collection by tradesmen when they remove and replace failed components. In the new industrial world of Total Productive Maintenance (TPM), they may be considered the true custodians of the data and the models derived from them.

Events and inspections data

The data required are of two types — events data and inspection data. Three types of events data, at a minimum, are required to define a component’s lifetime. They are:

- 1. The Beginning (B) of the component life or the time of installation.
- 2. The Ending by Failure (EF).
- 3. The Ending by Suspension (ES) due to a preventive replacement.

Additional events should be included in the model if they are known to di-

Figure 29 Haul truck transmissions inspection and event data *continued*

IDENT	DATE	W AGE	HN	P	EVENT	IRON	LEAD	CALCIUM	MAG
HT-66	5/10/96	18378	2	0	*	10	1	4385	12
HT-66	6/3/96	18914	2	0	*	11	0	3586	21
HT-66	6/3/96	18914	2	1	OC	0	0	5000	0
HT-66	7/1/96	19338	2	0	*	8	0	4885	5
HT-66	7/29/96	19876	2	0	*	16	0	4430	6
HT-66	7/29/96	19876	2	1	OC	0	0	5000	0
HT-66	8/26/96	20425	2	0	*	6	0	5249	2
HT-66	9/24/96	21034	2	0	*	10	0	4932	2
HT-66	9/24/96	21034	2	1	OC	0	0	5000	0
HT-66	10/21/96	21626	2	0	*	5	0	5070	2
HT-66	11/18/96	22266	2	0	*	3	0	5538	0
HT-66	11/18/96	22266	2	1	OC	0	0	5000	0
HT-66	12/9/96	22706	2	3	ES	3	0	5538	0
HT-66	12/10/96	22706	3	4	B	0	0	5000	0
HT-66	12/16/96	22862	3	0	*	4	0	4962	4
HT-66	1/13/97	23499	3	0	*	9	0	5593	2
HT-66	1/13/97	23499	3	1	OC	0	0	5000	0
HT-66	2/11/97	24084	3	0	*	11	0	5361	8
HT-66	3/9/97	24491	3	0	*	5	0	4916	100
HT-66	3/9/97	24491	3	1	OC	0	0	5000	0
HT-66	4/7/97	25053	3	0	*	8	1	4321	83
HT-66	5/5/97	25666	3	0	*	11	0	4316	100
HT-66	5/5/97	25666	3	1	OC	0	0	5000	0
HT-66	6/2/97	26289	3	0	*	6	0	5013	21
HT-66	6/29/97	26884	3	0	*	10	0	5293	24
HT-66	6/29/97	26884	3	1	OC	0	0	5000	0
HT-66	7/28/97	27519	3	0	*	7	0	4933	15
HT-66	8/25/97	28157	3	0	*	8	0	5648	56
HT-66	8/25/97	28157	3	1	OC	0	0	5000	0
HT-66	9/22/97	28784	3	0	*	6	0	4642	47
HT-66	10/20/97	29379	3	0	*	3	2	5826	27
HT-66	10/20/97	29379	3	1	OC	0	0	5000	0
HT-66	11/17/97	29921	3	0	*	5	0	5522	24
HT-66	12/15/97	30507	3	0	*	5	1	5294	58
HT-66	12/15/97	30507	3	1	OC	0	0	5000	0
HT-66	1/12/98	31133	3	0	*	1	0	6124	16
HT-66	2/9/98	31724	3	0	*	3	0	5685	20
HT-66	2/9/98	31724	3	1	OC	0	0	5000	0
HT-66	3/9/98	32335	3	0	*ES	2	0	5000	8
HT-67	1/19/94	0	1	4	B	0	0	5000	0
HT-67	1/20/94	3	1	0	*	1	0	3860	0
HT-67	2/17/94	657	1	0	*	15	0	3482	0
HT-67	3/17/94	1299	1	0	*	17	0	3717	0
HT-67	3/17/94	1299	1	1	OC	0	0	5000	0
HT-67	4/14/94	1922	1	0	*	7	0	3680	8
HT-67	5/12/94	2516	1	0	*	18	2	3655	9
HT-67	5/12/94	2516	1	1	OC	0	0	5000	0
HT-67	6/9/94	3129	1	0	*	9	0	4609	4
HT-67	7/6/94	3680	1	0	*	10	0	4526	3
HT-67	7/6/94	3680	1	1	OC	0	0	5000	0

rectly influence the measured data. One such event type can be an oil change (designated by “OC” in Figure 29). One should “tell” the model that at each oil change, some covariates such as the wear metals are expected to be reset to zero. This additional intelligence will preclude the model from being “fooled” by periodic decreases in wear metals. This is illustrated in Figure 28.

Periodic tightening, alignment, balancing, or re-calibration of machinery may have similar effects on measured values (such as vibration readings) and should be accounted for in the model.

Sample inspection data

Figure 29 (beginning on page 58) displays a partial data set from a fleet of four haul truck transmissions. The complete

data set is available as a MSAccess database file from the author [see ref. 3]. Such data is necessary for building a proportional hazard model (and ultimately an optimal decision policy). The inspections, in this case are oil analysis results, and are designated by an asterisk in the “Event” column. The data comprise the entire history of each unit identified by the designations HT-66, HT-67, HT-76, and HT-77, between December 1993 and February 1998. The event and inspection data are displayed chronologically by equipment number.

Cross graphs

The data of Figure 29 must be “understood” by the modeller. The data preparation phase includes activities whereby the modeller may become fa-

miliar with the data using a number of software graphical tools. The cross graph (Figure 30, page 61) is very convenient for graphical statistical analysis. For example, it readily shows possible correlation between diagnostic variables. Correlation between two variables becomes evident when the points are clustered around a straight line. If the points are randomly scattered as they are in Figure 30 (which is a plot of Lead vs. Iron) one can easily see that there is no correlation between the two covariates. Should correlation be evident, this will be useful knowledge in subsequent modelling steps.

Cleaning up the data

The technical term used by statisticians when the data contains inappropriate

Figure 29 Haul truck transmissions inspection and event data *continued*

IDENT	DATE	W_AGE	HN	P	EVENT	IRON	LEAD	CALCIUM	MAG
HT-67	8/4/94	4331	1	0	*	9	2	4701	2
HT-67	9/1/94	4977	1	0	*	22	0	4639	2
HT-67	9/1/94	4977	1	1	OC	0	0	5000	0
HT-67	9/29/94	5597	1	0	*	15	1	4574	3
HT-67	10/27/94	6196	1	0	*	15	0	5555	0
HT-67	10/27/94	6196	1	1	OC	0	0	5000	0
HT-67	11/24/94	6760	1	0	*	19	2	4536	2
HT-67	12/22/94	7378	1	0	*	25	5	4279	2
HT-67	12/22/94	7378	1	1	OC	0	0	5000	0
HT-67	12/29/94	7523	1	2	EF	25	5	4279	2
HT-67	12/29/94	7523	2	4	B	0	0	5000	0
HT-67	1/19/95	7982	2	0	*	7	1	3284	57
HT-67	2/16/95	8623	2	0	*	5	2	3077	51
HT-67	2/16/95	8623	2	1	OC	0	0	5000	0
HT-67	3/16/95	9243	2	0	*	6	2	4985	11
HT-67	4/13/95	9866	2	0	*	4	1	5068	9
HT-67	4/13/95	9866	2	1	OC	0	0	5000	0
HT-67	5/11/95	10507	2	0	*	4	2	4386	8
HT-67	6/8/95	11107	2	0	*	3	3	4501	7
HT-67	6/8/95	11107	2	1	OC	0	0	5000	0
HT-67	7/7/95	11716	2	0	*	6	0	4862	7
HT-67	7/26/95	12082	2	0	*	10	0	2375	5
HT-67	8/3/95	12230	2	0	*	5	0	5435	7
HT-67	8/3/95	12230	2	1	OC	0	0	5000	0
HT-67	8/31/95	12846	2	0	*	7	1	5216	28
HT-67	9/28/95	13476	2	0	*	9	0	4708	34
HT-67	9/28/95	13476	2	1	OC	0	0	5000	0
HT-67	10/26/95	14099	2	0	*	6	0	5114	12
HT-67	11/22/95	14723	2	0	*	5	0	5684	26
HT-67	11/22/95	14723	2	1	OC	0	0	5000	0
HT-67	12/21/95	15325	2	0	*	2	2	5306	100
HT-67	1/18/96	15815	2	0	*	6	0	5058	87
HT-67	1/18/96	15815	2	1	OC	0	0	5000	0
HT-67	2/15/96	16370	2	0	*	8	0	4928	18
HT-67	3/11/96	16932	2	0	*	6	0	5230	17
HT-67	3/11/96	16932	2	1	OC	0	0	5000	0
HT-67	4/11/96	17532	2	0	*	0	0	5838	1
HT-67	5/9/96	18153	2	0	*	9	1	5389	6
HT-67	5/9/96	18153	2	1	OC	0	0	5000	0
HT-67	6/5/96	18751	2	0	*	12	0	5435	2
HT-67	7/4/96	19277	2	0	*	8	1	4104	8
HT-67	7/4/96	19277	2	1	OC	0	0	5000	0
HT-67	7/18/96	19575	2	3	ES	8	1	4104	8
HT-67	7/19/96	19575	3	4	B	0	0	5000	0
HT-67	8/1/96	19897	3	0	*	12	0	4133	1
HT-67	8/29/96	20387	3	0	*	7	0	5008	4
HT-67	8/29/96	20387	3	1	OC	0	0	5000	0
HT-67	9/26/96	21032	3	0	*	10	0	4996	7
HT-67	10/24/96	21660	3	0	*	10	0	4545	31
HT-67	10/24/96	21660	3	1	OC	0	0	5000	0

and misleading events and values is “dirty”. Dirty data must be cleaned up before synthesizing it into a model to be used for future decisions and policy. That would include verifying the validity of outliers in the inspection data set.

Data transformations

The modeller needs to have at his fingertips, not only the actual data, but also any combinations (transformations) of that data which he feels may be influential covariates.

One obvious transformed data field of interest is the lubricating oil’s age, which is usually not directly available in the database, but can be calculated knowing the dates of the oil change (OC) events. In the current example one would expect the “wear metals”




Figure 30: Cross graph of lead vs iron ppm

iron or lead to be low just after an oil change and increase linearly as the oil ages and as particles accumulate and stay resident in the oil circulating system. A cross graph (Figure 31, page 62)

of oil age and iron negates this theory except for very low ages of the lubricating oil. Hence the modeller may conclude that oil age, in this system, is not a significant covariate.

Figure 29 Haul truck transmissions inspection and event data *continued*

IDENT	DATE	W AGE	HN	P	EVENT	IRON	LEAD	CALCIUM	MAG
HT-67	11/21/96	22309	3	0	*	11	0	5034	11
HT-67	12/19/96	22874	3	0	*	9	0	4883	15
HT-67	12/19/96	22874	3	1	OC	0	0	5000	0
HT-67	1/14/97	23458	3	0	*	7	1	5742	8
HT-67	2/12/97	24126	3	0	*	8	0	4935	9
HT-67	2/12/97	24126	3	1	OC	0	0	5000	0
HT-67	3/15/97	24706	3	0	*	4	0	5205	6
HT-67	4/10/97	25079	3	0	*	3	0	5550	7
HT-67	4/10/97	25079	3	1	OC	0	0	5000	0
HT-67	5/8/97	25642	3	0	*	7	2	4393	18
HT-67	6/4/97	26198	3	0	*	6	4	4744	17
HT-67	6/4/97	26198	3	1	OC	0	0	5000	0
HT-67	7/2/97	26782	3	0	*	8	4	4182	10
HT-67	7/31/97	27415	3	0	*	7	5	5562	6
HT-67	7/31/97	27415	3	1	OC	0	0	5000	0
HT-67	8/28/97	27954	3	0	*	2	0	4520	15
HT-67	9/24/97	28591	3	0	*	2	0	5290	9
HT-67	9/24/97	28591	3	1	OC	0	0	5000	0
HT-67	10/23/97	29222	3	0	*	5	0	4989	12
HT-67	11/20/97	29847	3	0	*	5	0	4440	14
HT-67	11/20/97	29847	3	1	OC	0	0	5000	0
HT-67	12/18/97	30381	3	0	*	16	0	5008	44
HT-67	1/15/98	30954	3	0	*	2	0	5484	14
HT-67	1/15/98	30954	3	1	OC	0	0	5000	0
HT-67	2/10/98	31544	3	0	*	2	0	5612	17
HT-67	3/12/98	32214	3	1	*ES	2	0	5612	17
HT-77	3/14/95	0	1	4	B	0	0	5000	0
HT-77	3/15/95	15	1	0	*	2	0	4222	1
HT-77	4/22/95	871	1	0	*	8	3	3031	2
HT-77	5/20/95	1530	1	0	*	8	2	3474	3
HT-77	5/20/95	1530	1	1	OC	0	0	5000	0
HT-77	6/17/95	2147	1	0	*	7	1	4633	6
HT-77	7/15/95	2779	1	0	*	10	3	4923	7
HT-77	7/15/95	2779	1	1	OC	0	0	5000	0
HT-77	8/12/95	3419	1	0	*	8	0	5287	6
HT-77	9/9/95	4052	1	0	*	77	2	5021	4
HT-77	9/9/95	4052	1	1	OC	0	0	5000	0
HT-77	9/21/95	4274	1	2	EF	77	2	5021	4
HT-77	9/22/95	4274	2	4	B	0	0	5000	0
HT-77	9/23/95	4352	2	0	*	4	0	5037	9
HT-77	9/23/95	4352	2	1	OC	0	0	5000	0
HT-77	10/7/95	4631	2	0	*	10	0	4768	7
HT-77	11/4/95	5285	2	0	*	39	0	5225	8
HT-77	11/4/95	5285	2	1	OC	0	0	5000	0
HT-77	12/2/95	5919	2	0	*	20	1	6117	6
HT-77	12/30/95	6552	2	0	*	18	0	5901	6
HT-77	12/30/95	6552	2	1	OC	0	0	5000	0
HT-77	1/17/96	6946	2	2	EF	18	0	5901	6
HT-77	1/18/96	6946	3	4	B	0	0	5000	0
HT-77	1/27/96	7178	3	0	*	11	0	2903	0

Step 2: Building the Proportional Hazards Model

This step is performed entirely by the software. The parameters of the PHM equation are estimated.

Equation 1:

$$h(t) = \frac{\beta}{\eta} \left(\frac{t}{\eta} \right)^{\beta-1} e^{\gamma_1 Z_1(t) + \gamma_2 Z_2(t) + \dots + \gamma_n Z_n(t)}$$

Examining equation 1 we see that it extends the Weibull hazard function described in Chapter 4 and applied in Chapter 5. The new part factors in (as an exponential expression) the covariates $Z_i(t)$ which are the set of measured CBM data items, for example, the parts per million of iron or other wear metals present in the oil sample. The covariate parameters γ_i specify the relative “influence” that each covariate has on the hazard (or failure




Figure 31: Iron vs oil age

rate) function. A very low value for γ_i would tend to indicate that the corresponding covariate is of little influence and not worth measuring. The software to test whether each covariate is insignificant uses a statistical test. In fact, a variety of statistical tests within the software (Maximum Likelihood Estimates, Wald, Chi Square,

Cox-generalized residuals, Kolmogorov-Smirnov) provide systematic criteria for testing various hypotheses concerning model confidence and significance.

The software’s algorithms fit the proportional hazards model to the data providing estimates not only of the shape parameter β and scale parameter

Figure 29 Haul truck transmissions inspection and event data *continued*

IDENT	DATE	W_AGE	HN	P	EVENT	IRON	LEAD	CALCIUM	MAG
HT-77	2/25/96	7830	3	0	*	6	1	3540	0
HT-77	2/25/96	7830	3	1	OC	0	0	5000	0
HT-77	3/23/96	8032	3	0	*	0	0	3660	0
HT-77	4/20/96	8717	3	0	*	1	0	3885	0
HT-77	4/20/96	8717	3	1	OC	0	0	5000	0
HT-77	5/18/96	9230	3	0	*	8	0	4237	4
HT-77	6/15/96	9852	3	0	*	10	0	4994	2
HT-77	6/15/96	9852	3	1	OC	0	0	5000	0
HT-77	7/13/96	10418	3	0	*	6	0	4570	6
HT-77	8/8/96	10952	3	0	*	3	0	4901	53
HT-77	8/8/96	10952	3	1	OC	0	0	5000	0
HT-77	9/7/96	11600	3	0	*	5	1	5470	20
HT-77	10/5/96	12175	3	0	*	3	0	4998	19
HT-77	10/5/96	12175	3	1	OC	0	0	5000	0
HT-77	11/2/96	12807	3	0	*	3	0	4588	5
HT-77	11/30/96	13422	3	0	*	2	0	5547	1
HT-77	11/30/96	13422	3	1	OC	0	0	5000	0
HT-77	12/28/96	14015	3	0	*	2	2	5929	0
HT-77	1/25/97	14624	3	0	*	2	0	5351	1
HT-77	1/25/97	14624	3	1	OC	0	0	5000	0
HT-77	2/22/97	15190	3	0	*	3	0	4822	6
HT-77	3/22/97	15768	3	0	*	2	0	4937	4
HT-77	3/22/97	15768	3	1	OC	0	0	5000	0
HT-77	4/19/97	16417	3	0	*	3	0	5420	4
HT-77	6/14/97	17516	3	0	*	1	0	4996	62
HT-77	7/12/97	18217	3	0	*	4	0	4633	59
HT-77	7/12/97	18217	3	1	OC	0	0	5000	0
HT-77	8/10/97	18816	3	0	*	3	0	5679	20
HT-77	8/25/97	19146	3	3	ES	3	0	5679	20
HT-77	8/26/97	19146	4	4	B	0	0	5000	0
HT-77	9/6/97	19424	4	0	*	24	0	4939	7
HT-77	9/6/97	19424	4	1	OC	0	0	5000	0
HT-77	10/4/97	20038	4	0	*	26	3	3588	10
HT-77	11/1/97	20662	4	0	*	26	3	4977	11
HT-77	11/1/97	20662	4	1	OC	0	0	5000	0
HT-77	11/29/97	21259	4	0	*	8	0	4430	16
HT-77	12/27/97	21931	4	0	*	8	0	5496	16
HT-77	12/27/97	21931	4	1	OC	0	0	5000	0
HT-77	1/24/98	22561	4	0	*	3	0	5414	16
HT-77	2/21/98	22917	4	0	*	12	2	4571	9
HT-77	2/21/98	22917	4	1	*ES	12	2	4571	9
HT-79	4/15/95	0	1	4	B	0	0	5000	0
HT-79	6/23/95	1307	1	0	*	13	9	3784	4
HT-79	7/21/95	1958	1	0	*	19	10	3177	5
HT-79	7/21/95	1958	1	1	OC	0	0	5000	0
HT-79	8/18/95	2463	1	0	*	15	2	4906	7
HT-79	9/15/95	3106	1	0	*	17	3	4898	6
HT-79	9/15/95	3106	1	1	OC	0	0	5000	0
HT-79	10/13/95	3725	1	0	*	13	0	5293	8

η as was the case in the Weibull examples of Chapter 5, but also estimates of each covariate parameter γ_i .

Step 3: testing the PHM

Let's review our objectives. We are searching for a decision mechanism, which will, over the long run, result in the lowest total cost of maintenance. Recall that the total cost of maintenance includes the cost of failure repairs (which typically include a variety of additional costs such as the cost of lost production). Hence it is essential that we are confident that the model adequately and realistically reflects our system or component's failure characteristics.

Residual Analysis is a procedure that tells us how well the PHM Model fits the data. The method of Cox-generalized residuals is applied to test the

model fit. The method mathematically generates numbers known as "residuals." The residuals are then examined graphically.

There are a variety of types of residual plots used to evaluate how well the model "fits". One of them is the "Residuals In Order of Appearance" graph, Figure 32 (page 64). This graphical method plots the residuals in the same order as the histories that appear in Figure 29. The average residual value must equal 1. So the random scatter of points around the horizontal line $y=1$ is expected if the model fits the data well. Note that the residuals obtained from censored values are always above the line $y=1$. (Censored data was discussed in Chapter 4.) To help in examining residuals, the appropriate upper and lower limits are included on the graph.

If the PHM fits the data well, at least 90 percent of the residuals are expected within these limits. Varieties of additional graphical and mathematical tools are called upon in step-wise fashion to help the modeller develop, test, compare, and gain confidence in his or her model.

Step 4: The transition probability model

At this juncture in the CBM optimization process the PHM will have been developed and tested by the modeller using the software tools described in the preceding sections. He or she is presumably satisfied and confident that the model fits the cleaned data well. For each of the covariates in the model, the modeller must now define ranges of values or states, for example

Figure 29 Haul truck transmissions inspection and event data *continued*

IDENT	DATE	W AGE	HN	P	EVENT	IRON	LEAD	CALCIUM	MAG
HT-79	11/10/95	4456	1	0	*	20	0	5513	8
HT-79	11/10/95	4456	1	1	OC	0	0	5000	0
HT-79	12/9/95	4772	1	0	*	1	0	7175	4
HT-79	1/5/96	5951	1	0	*	9	0	6814	6
HT-79	1/5/96	5951	1	1	OC	0	0	5000	0
HT-79	2/2/96	6102	1	0	*	5	0	5046	4
HT-79	3/1/96	6614	1	0	*	6	0	5924	3
HT-79	3/1/96	6614	1	1	OC	0	0	5000	0
HT-79	3/29/96	7257	1	0	*	7	0	5032	5
HT-79	4/26/96	7881	1	0	*	8	0	5826	3
HT-79	4/26/96	7881	1	1	OC	0	0	5000	0
HT-79	5/24/96	8515	1	0	*	8	0	4881	4
HT-79	6/19/96	9055	1	0	*	10	0	4817	5
HT-79	6/19/96	9055	1	1	OC	0	0	5000	0
HT-79	6/25/96	9468	1	2	EF	10	0	4817	5
HT-79	6/25/96	9468	2	4	B	0	0	5000	0
HT-79	7/19/96	9629	2	0	*	13	1	4710	2
HT-79	8/16/96	10244	2	0	*	14	4	4652	2
HT-79	8/16/96	10244	2	1	OC	0	0	5000	0
HT-79	9/13/96	10887	2	0	*	7	0	5214	3
HT-79	10/11/96	11437	2	0	*	9	0	5593	0
HT-79	11/7/96	12042	2	0	*	7	0	5714	0
HT-79	12/6/96	12691	2	0	*	7	0	5787	2
HT-79	12/6/96	12691	2	1	OC	0	0	5000	0
HT-79	1/3/97	13104	2	0	*	3	0	4790	3
HT-79	1/31/97	13680	2	0	*	4	0	5144	3
HT-79	1/31/97	13680	2	1	OC	0	0	5000	0
HT-79	2/28/97	14291	2	0	*	4	0	4985	4
HT-79	3/28/97	14821	2	0	*	3	0	4944	4
HT-79	3/28/97	14821	2	1	OC	0	0	5000	0
HT-79	4/24/97	15417	2	0	*	3	0	5088	6
HT-79	5/23/97	16045	2	0	*	6	0	5921	7
HT-79	5/23/97	16045	2	1	OC	0	0	5000	0
HT-79	6/20/97	16459	2	0	*	5	2	5629	4
HT-79	7/18/97	16969	2	0	*	6	8	5090	6
HT-79	7/18/97	16969	2	1	OC	0	0	5000	0
HT-79	7/29/97	17653	2	2	EF	6	8	5090	6
HT-79	7/30/97	17653	3	4	B	0	0	5000	0
HT-79	9/13/97	18157	3	0	*	20	4	5602	5
HT-79	9/13/97	18157	3	1	OC	0	0	5000	0
HT-79	10/10/97	18758	3	0	*	9	1	5221	10
HT-79	11/6/97	19353	3	0	*	11	2	5545	10
HT-79	11/6/97	19353	3	1	OC	0	0	5000	0
HT-79	12/5/97	20029	3	0	*	5	4	3915	16
HT-79	1/2/98	20627	3	0	*	5	6	5834	17
HT-79	1/2/98	20627	3	1	OC	0	0	5000	0
HT-79	1/30/98	21271	3	0	*	1	0	6699	16
HT-79	2/26/98	21688	3	0	*	3	3	4718	15
HT-79	2/27/98	21688	3	1	*ES	3	3	4718	15

low, medium, high or normal, marginal, critical levels. These covariate bands are set by the modeller by gut feel, plant experience, eyeballing the data, or using some statistical method such as placing the boundary at two standard deviations above the mean to define the medium level and at three standard deviations to define the high range of values.

Discussion of transition probability

Once the modeller has established the covariate bands, the software calculates the Markov Chain Model Transition Probability Matrix (Figure 33), which is a table that shows the probabilities of going from one state (for example, light contamination, medium contamination, heavy contamination) to another, between inspection intervals. Or more formally stated “The table provides a quantitative estimate of the probability that the equipment will be found in a particular state at the next inspection, given its state today.”

It means that given the present value range of a variable we can predict (based on history) with a certain probability, its value range at the next inspection moment. For example, if the last inspection showed an iron level less than 13 parts per million (ppm) we may, in this example, predict that the next record will be between 13 and 26 ppm with a probability of 12 percent and more than 26 ppm, with probability 1.6 percent. Probabilities such as the 12 percent and 1.6 percent are called transition probabilities.




Figure 32: The transition probability model

IRON	0 to 13.013	13.013 to 26.949	Above 25.949
0 to 13.013	0.864448	0.11918	0.0163723
13.013 to 25.949	0.139931	0.657121	0.202947
Above 25.949	0	0	1

Figure 33: The Markov chain model transition probability matrix

Step 5: The optimal decision

In this step we “tell” the model (consisting thus far of the PHM and the Transition Probability models) the respective costs of a preventive and failure instigated repair or replacement.

The replacement decision graph, Figure 34 (page 66), is the culmination of the entire modelling exercise to date. It combines the results of the proportional hazards model, the transition probability model, and the cost function to display the best decision policy regarding the component or system in question.

The ordinate is the composite co-

variate, Z — a weighted sum of those covariates having been statistically determined to influence the probability of failure. Each covariate measured at the most recent inspection will have contributed its value to Z. That contribution will have been weighted according to its degree of influence on the risk of failure in the next inspection interval.

The advantage is evident. On a single graph one has obtained the distilled information upon which to base a replacement decision. The alternative would have been to examine trend graphs of dozens of condition parameters and “guess” at whether to repair or replace




Figure 34: The optimal replacement decision graph

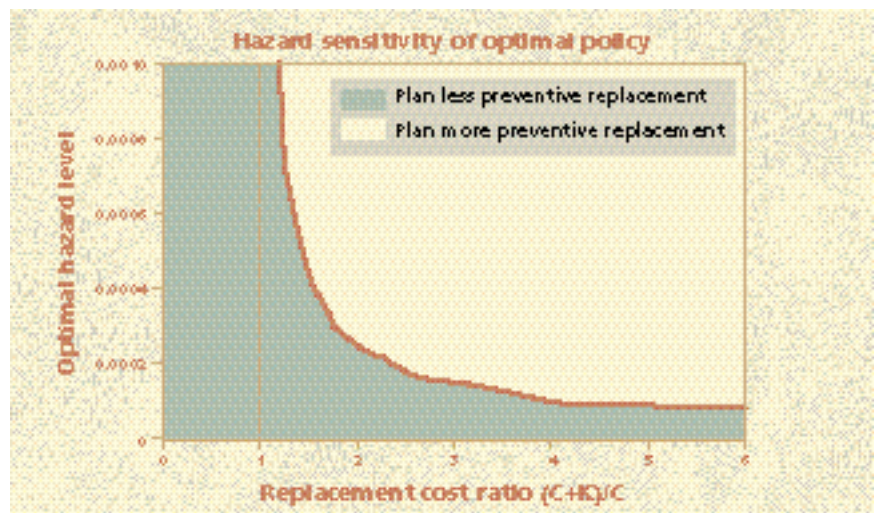


Figure 35: Sensitivity of optimal policy to cost ratio

the component immediately or wait a little longer. By accepting the recommendation of the Optimal Replacement Decision Graph, one acts according to the best known policy to minimize the long run maintenance cost.

Step 6: Sensitivity analysis

There is one more step to complete. How do we know that the Optimum Replacement Decision Graph truly constitutes the best policy in the light of our plant's ever-changing operating situation? Are the assumptions we used still valid, and if not, what will be the effect of those changes? Is our decision still optimal? These questions are addressed by sensitivity analysis.

The assumption we made in building the cost function model centered on the relative costs of a planned replacement versus those of a replacement forced by a sudden failure. That cost ratio may have changed. Our accounting methods may not currently provide us with the precise costs of re-

pair and therefore we had to estimate them when building the cost function. In either case we have doubts about whether the policy dictated by the Optimal Replacement Decision Graph is well founded given the uncertainty of the costs upon which it was calculated.

The purpose of the sensitivity analysis is to allay such fears when they are not warranted and to direct the modeller to expend some effort to obtain a more precise estimate of maintenance costs where they are needed.

Figure 35, the Hazard Sensitivity of Optimal Policy graph, shows us the relationship between the optimal hazard or risk level and the cost ratio. If the cost ratio were low, less than 3, then optimal hazard level would increase exponentially. Hence we need to track costs very closely in order to assure ourselves of the benefits calculated by the model.

Conclusion

As acute competition of the global market touches an industry, visibility of each

aspect of the production process, and in particular that of equipment reliability, will become an urgent business requirement. Integration and fluidity of the supply chain across multiple partnering businesses connected electronically will force maintenance, availability, and reliability information to be mission critical.

User friendly, yet sophisticated, software will empower maintenance professionals to respond with agility to the incessant fluctu-

model, an unambiguous yet appropriate optimal decision must be made and executed quickly. The stage has been set, and maintenance players, in various states of readiness, must meet entirely new challenges as the curtain rises on a rapidly transforming business culture. The management of this change will necessarily revolve around the meticulous collection and analysis of data.

Existing maintenance information

The stage has been set, and maintenance players, in various states of readiness, must meet entirely new challenges as the curtain rises on a rapidly transforming business culture.

ating demand of unprecedented market forces. Mathematical statistical models such as those developed with the help of the software tools discussed in this chapter, along with expert systems founded upon the principles of reliability centered maintenance, will be the “watchdogs” operating silently within a company’s computerized maintenance management system. Their function — to continually monitor incoming condition and age data. When condition data triggers an alert with respect to the optimal maintenance policy

management database systems are underused and inadequately populated mainly because maintenance tradesmen and employees are not yet convinced that there is a relationship between accurately recorded component lifetime data and their own effectiveness to keep the physical assets of their organization functioning. It is the author’s hope that the methods described here will assist maintenance personnel in their decision-making tasks as they progress towards ultimate plant reliability at lowest cost.

References:

- 1. *Statistical Methods in Reliability Theory and Practice*, Brian D. Bunday, Ellis Horwood Limited, 1991.
- 2. www.mie.utoronto.ca/labs/cbm
- 3. murray.z.wiseman@ca.pwcglobal.com
- 4. *Applied Reliability*, Paul A. Tobias, David C. Trinidade, 1995 Van Nostrand Reinhold.
- 5. *The New Weibull Handbook*, Robert E. Abernethy, 2nd ed Dr. Robert B. Abernethy SAE TA 169 A35 1996X C.1 Engi.
- 6. “Applications of maintenance optimization models: a review and analysis”, Rommert Dekker, *Reliability Engineering and System Safety* 51, 1996, 229-240 Elsevier Science Limited.
- 7. “On the application of mathematical models in maintenance,” Philip A. Scarf, *European Journal of Operational Research*, 1997 Elsevier Science.
- 8. “On the impact of optimization models in maintenance decision making: the state of the art,” Rommert Dekker, Philip A. Scarf, *Reliability Engineering and System Safety*, 1998 Elsevier Science Limited.
- 9. *Mine Planning and Equipment Selection*, A.A. Balkema, 1994, Proceedings of the Third International Symposium on Mine Planning and Equipment Selection, Istanbul/Turkey. ■

Searching the Web for reliability information

Looking for useful Internet sites? Here's where to start

by Paul Challen

In the preceding pages of *The Reliability Handbook*, we've seen a full discussion about ways in which maintenance professionals can increase uptime in their plants. Part of this discussion hinges on the need to use technology to garner the kind of information needed to implement reliability decisions. Along with the various software packages and databases our authors have recommended, there are a myriad of other resources available to people searching for reliability information — and one of the most useful places to look is on the Internet.

The following is a by-no-means-comprehensive listing of some of the reliability resources you'll find on the Net. We encourage readers to explore these sites and to follow the many reliability links emanating from them.

Reliability Analysis Center

Located at <http://rac.iitri.org>, this site bills itself as the “center of reliability and maintainability excellence for over 30 years.” The RAC is operated by the IIT Research Institute in the U.S., and provides information to industry via data bases, methodology handbooks, state-of-the-art technology reviews, training courses and consulting services. Its mission is to provide technical expertise and information in the engineering disciplines of reliability, maintainability, supportability and quality and to facilitate their cost-effective implementation throughout all phases of the product or system life cycle.

The RAC web site contains a number of useful information resources, including a bibliographic database of books, standards, journal articles, symposium papers, and other documents on reliability, maintainability, quality, and supportability. It also maintains a calendar of upcoming events throughout the industry, information about links to other databases, a “data sharing consortium” with information on non electronic parts reliability data, failure mode distributions, and electrostatic discharge susceptibility data.

The RAC has also compiled two important lists, one of frequently used acronyms in the reliability world, and another (at <http://rac.iitri.org/cgi/rac/sites?0>) of navigable links to other Internet sites on related topics.

Book and print material sources

For those reliability professionals who are searching for good old-fashioned

print material on the subject, there is a comprehensive general list of books located at <http://www.quality.org/Bookstore>. For books on reliability that you can order on-line through the Amazon site (www.amazon.com) on the subject of reliability, you can try: <http://www.quality.org/Bookstore/Reliability.htm>.

There is also a complete bibliography of U.S. government reliability documents at <http://www.incose.org/lib/sebib5.html> that covers a wide range of technical topics.

Professional organizations

The RAC site also contains a huge list of relevant professional organizations, at <http://rac.iitri.org/cgi-rac/sites?00013>.

One of the most useful of these, from a reliability standpoint, is the Society for Maintenance & Reliability Professionals (SMRP) site at <http://www.smrp.org>. This organization is an independent, non-profit society devoted to practi-



<http://www.world5000.com>.

tioners in the maintenance and reliability fields.

The society is, according to its mission statement, “dedicated to excellence in maintenance and reliability in all types of manufacturing and service organizations, and to promote maintenance excellence worldwide.” It also contains links to other organizations on reliability and maintenance issues.

The Society of Reliability Engineers (SRE) web site at <http://www.sre.org> is another helpful resource on reliability. Of particular interest will be the society’s newsletter “Lambda Notes,” and its directory of reliability utilities.

For a global perspective, you can look



<http://www.reliability-magazine.com>.

at the web site of the World Reliability Organization at <http://www.world5000.com>. The organization’s site describes their plan to implement an International Reliability Index, in which each country will be rated on a number of indexes for an overall total reliability score.

General information

Reliability Magazine, hailed as “the first trade journal dedicated specifically to the predictive maintenance industry, root cause failure analysis, reliability centered maintenance and CMMS,” has a site at <http://www.reliability-magazine.com>.

The University of Tennessee’s

Maintenance & Reliability Center (MRC) is a newly-mounted site that uses research and cutting-edge technology to help member companies reduce losses caused by equipment downtime. You can visit the MRC at <http://www.engr.utk.edu/mrc>.

The Centre for the Management of Industrial Reliability and Cost Effectiveness, located at the University of Exeter in England, promotes national and international collaboration with industry and other academic and research institutions in the the reliability and maintenance areas. Their web site is at <http://www.ex.ac.uk/mirce>.

The Equipment Reliability Institute (ERI) “links” page at <http://www.equipment-reliability.com/ERILinks.html> provides resources for finding standards organizations, technical societies, and sites that provide equipment and services that the ERI says “can help organizations achieve high reliability and durability.”

The Vibration Institute (at <http://www.vibinst.org/>) is a non-profit organization dedicated to the exchange of practical vibration information on machines and structures. The Institute’s activities also publishes *Vibrations* magazine, Proceedings of its Annual Meetings, and “Short Course Notes”. ■

COST SAVINGS IN

PHYSICAL ASSET MANAGEMENT

CONTRIBUTE DIRECTLY TO BOTTOM LINE PROFITS.



Our Physical Asset Management Group provides best practice and systems consulting services in all areas of maintenance including equipment, plant, fleet and facilities; any business whose bottom line performance can be improved by increasing the cost effectiveness of productive assets. We can help your company with:

- Strategic Cost Reduction
- Physical Asset Productivity
- Autonomous Maintenance Techniques
- Reliability Centered Maintenance
- Maintenance Benchmarking Studies
- Maintenance Diagnostic
- Enterprise Asset Management/Computerized Maintenance Management Systems

For more information please contact:

John D. Campbell

Global & Americas Leader

Toronto, Canada

Phone: (416) 941-8448

Fax: (416) 941-8419

Email: john.d.campbell@ca.pwcglobal.com



Join us. Together we can change the world.™