

MAINTENANCE MANAGEMENT - A New Paradigm -

By John Moubray, Managing Director, Aladon Ltd

Abstract:

This paper attempts to summarise fifteen of the most important areas of change which have occurred in the field of physical asset management over the past fifteen years.

The paper is based on the experience gained by the author and his associates from the application of RCM and its more modern derivative, RCM2, on more than 500 industrial sites in 27 countries over a period of 10 years.

On the basis of this work, it has become evident that any organisation that wishes to achieve rapid, substantial and lasting improvements in maintenance effectiveness – in other words, in physical asset performance – must ensure that everyone associated with the operation and maintenance of these assets profoundly understands and internalises both the nature and significance of these changes.

PARADIGM:
Pattern or model
Oxford Dictionary

MAXIM:
Piece of wisdom or rule of conduct expressed in a sentence
Oxford Dictionary

Introduction

The subject of change dominates nearly everything currently being written about management. All disciplines are being exhorted to adapt to changes in organisation design, in technology, in leadership skills, in communication – in fact, in virtually every aspect of working life.

Perhaps nowhere is this felt more broadly and deeply than in the field of physical asset management.

A striking feature of this phenomenon is the number of changes which have occurred simultaneously. Some have occurred at a strategic – almost philosophical – level, while others are more tactical – or technical – in nature. Even more striking is the extent of the changes. Not only do they involve radical changes of direction (some diametrically opposed to the way things have been done in the past), but a few ask us to come to terms with entirely new concepts.


This paper identifies fifteen key areas of change. Each of them *on its own* is sufficiently far-reaching to merit a great deal of attention in most organisations. Together they amount to a whole new paradigm. Accommodating this paradigm shift means that for most of us, the management of physical assets is going to become a monumental exercise in change management over the next few years.

Each of the changes on its own is also sufficient to form the subject of one – if not several – books (think of all the books available on the subject of predictive maintenance alone), so a short paper like this cannot hope to explore all the changes in detail. In fact, it goes to the opposite extreme by reducing each area of change to two maxims followed by a short explanation. In each case, one maxim attempts to summarise the way things used to be, while the other summarises the way things are – or should be – now.

There is bound to be an element of oversimplification in any attempt to summarise any *one* issue in one or two sentences, let alone *fifteen* issues. However, such a summary does fulfil two purposes:

- it provides a quick overview of what the changes are
- it provides a basis for comparing the different decision support tools and management philosophies that claim to provide a basis for action (RCM, FMECA, MSG3, HAZOP, TPM, RCA, RBI, RCM2 and others).

This paper only summarises the fifteen areas of change. Time constraints mean that use of this summary to compare different maintenance tools will have to be the subject of a later paper.



Maxim 1

OLD Maintenance is about preserving physical assets

Most people become engineers because they feel at least some affinity for things, be they mechanical, electrical or structural. This leads us to derive pleasure from assets in good condition, but feel offended by assets in poor condition.

These reflexes have always been at the heart of the concept of preventive maintenance. They have given rise to concepts like "asset care", which as the name implies, seeks to care for assets *per se*. They have also led maintenance strategists to believe that maintenance is all about preserving the inherent reliability or built-in capability of any asset.

In fact, this is not so.

As we gain a deeper understanding of the role of assets in business, we begin to appreciate the significance of the fact that any physical asset is put into service because someone wants it to do something. So it follows that when we maintain an asset, *the state which we wish to preserve must be one in which it continues to do whatever its users want it to do*. This in turn implies that we have to focus our attention on maintaining what each asset *does* rather than on what it *is*.

Clearly, before we can do this, we must gain a crystal clear understanding of the functions of each asset together with the associated performance standards

For instance, Figure 1 shows a pump with a nominal capacity of 400 litres/minute pumping water into a tank from which it is being drawn at a rate of 300 litres/minute. In this case, the primary function of the pump is "to supply water to the tank at not less than 300 litres/minute". Any main-

NEW Maintenance is about preserving the functions of assets

tenance programme for the pump should try to ensure that its performance does not drop below 300 litres/minute. (Note that in seeking to ensure that the tank does not run dry, the maintenance program does *not* try to ensure that the pump continues "to be able to supply water to the tank at not less than 400 litres/minute".)




FIGURE 1

However, if exactly the same pump is moved to a different tank where the offtake is 350 litres/minute, the primary function is changed accordingly and so the maintenance programme must now change to accommodate the higher performance expectation.

Functions and performance expectations not only cover output. They also concern issues such as product quality, customer service, economy and efficiency of operation, control, containment, comfort, protection, compliance with environmental regulations, structural integrity and even the physical appearance of the asset.

Maxim 2

OLD Routine maintenance is about preventing failures

A detailed analysis of an average industrial undertaking is likely to yield between five and ten thousand possible failure modes. Each of these failures affects the organisation in some way, but in each case, the effects are different. They may affect operations. They may also affect product quality, customer service, safety or the environment. They will all take time and cost money to repair.

It is these consequences which most strongly influence the extent to which we try to prevent each failure. If a failure mode has serious consequences, we are likely to go to great lengths to try to prevent it. If it has little or no effect, then we may decide to undertake no preventive action.

In other words, the consequences of failures are far more important than their technical characteristics.


For example, one failure which could affect the pump shown in Figure 1 is "bearing seizes due to normal wear and tear". Assuming that it takes 4 hours to replace a failed bearing, and that the unanticipated failure of the bearing only comes to the attention of the operators when the level in the tank drops to the low level switch, the tank only

NEW Routine maintenance is about avoiding, reducing or eliminating the consequences of failures

contains a 2.5 hour supply of water, so it would run dry and remain empty for 1.5 hours while the bearing is repaired.

One condition-based task which could apply to this bearing is to monitor vibration levels using a vibration analyser. If incipient failure is detected, the first priority of the operators would be to *fill the tank* before the bearing seizes, thus giving themselves 5 hours to do a 4 hour job. This in turn enables them to avoid the consequences of an empty tank (and also avoid possible secondary damage to the pump). *The task does not "save" the bearing* – that is doomed whatever happens.

This example demonstrates that the main reason for doing any kind of proactive maintenance is to avoid, reduce or eliminate the *consequences* of failure. A formal review of failure consequences focuses attention on maintenance tasks that have most effect on the performance of the organisation, and diverts energy away from those which have little or no effect. This helps ensure that whatever is spent on maintenance is spent where it will do the most good.



Maxim 3

OLD

The primary objective of the maintenance function is to optimise plant availability at minimum cost

Downtime has always affected the productive capability of physical assets by reducing output, increasing operating costs and interfering with customer service. By the 1960's and 1970's, this was already a major concern in the mining, manufacturing and transport sectors. In manufacturing, the effects of downtime are being aggravated by the world-wide move towards just-in-time systems, where reduced stocks of work-in-progress mean that quite small breakdowns are now much more likely to stop a whole plant. In recent times, the growth of mechanisation and automation means that reliability and availability have now also become key issues in sectors as diverse as health care, data processing, telecommunications and building management.

The cost of maintenance has also been rising at a steady pace over the past few decades, in absolute terms and as a proportion of total expenditure. In some industries, it is now the second highest or even the highest element of total costs. So in only 40 years maintenance has moved from nowhere to the top of the league as a cost control priority.

The importance of these two aspects of asset management means that many maintenance managers still tend to view them as the only significant objectives of maintenance.

However, this is no longer the case, because the maintenance function now has a wide range of additional objectives. These are summarised in the following paragraphs.

Greater automation means that more and more failures affect our ability to achieve and sustain satisfactory quality standards. This applies as much to *standards of service*

NEW

Maintenance affects all aspects of business effectiveness and risk – safety, environmental integrity, energy efficiency, product quality and customer service, not just plant availability and cost

as it does to *product quality*. For instance equipment failures affect climate control in buildings and the punctuality of transport networks as much as they interfere with the consistent achievement of specified tolerances in manufacturing.

Another result of growing automation is the rising number of failures which have serious safety or environmental consequences, at a time when standards in these areas are rising fast. Many parts of the world are reaching the point where organisations either conform to society's *safety and environmental expectations*, or they get shut down. This adds an order of magnitude to our dependence on the integrity of our physical assets – one which goes beyond cost and becomes a simple matter of organisational survival.

At the same time as our dependence on physical assets is growing, so too is their *cost – to operate and to own*. To secure the maximum return on the investment which they represent, they must be kept working efficiently for as long as their users want them to.

These developments mean that maintenance now plays an increasingly central role in preserving *all* aspects of the physical, financial and competitive health of the organization. This in turn means that maintenance professionals owe it to themselves and their employers to equip themselves with the tools needed to address these issues continuously, proactively and directly, rather than deal with them on an ad hoc basis when time permits

Maxim 4

OLD

Most equipment becomes more likely to fail as it gets older

For decades, conventional wisdom suggested that the best way to optimise the performance of physical assets was to overhaul or replace them at fixed intervals. This was based on the premise that there is a direct relationship between the amount of time (or number of cycles) equipment spends in service and the likelihood that it will fail, as shown in Figure 2. This suggests that most items can be expected to operate reliably for a period "X", and then wear out.

Classical thinking held that X could be determined from historical records about equipment failure, enabling users to take preventive action shortly before the item is due to fail in future. This predictable relationship between age and failure relationship is indeed true for some failure modes. It tends to be found where equipment comes into direct contact with the product. Examples include pump impellers, furnace refractories, valve seats, crusher liners, screw conveyors, machine tooling and so on. Age-related failures are also often associated with fatigue and corrosion.

NEW

Most failures are *not* more likely to occur as equipment gets older




FIGURE 2

However, equipment in general is much more complex than it was even fifteen years ago. This has led to startling changes in the patterns of equipment failure, as shown in Figure 3. The graphs show conditional probability of failure against operating age for a wide variety of electrical and mechanical items.

Pattern A is the well-known bathtub curve, and pattern B is the same as Figure 2. Pattern C shows slowly increasing probability of failure with no specific wear-out age. Pattern D shows low failure probability to begin with then a rapid increase to a constant level, while pattern E shows





FIGURE 3

a constant probability of failure at all ages. Pattern F starts with high infant mortality and drops eventually to a constant or very slowly increasing failure probability.

Studies on civil aircraft showed that 4% of the items conform to pattern A, 2% to B, 5% to C, 7% to D, 14% to E and no fewer than 68% to pattern F. (The distribution of these patterns in aircraft is not necessarily the same as in industry, but as equipment grows more complex, more and more items conform to patterns E and F.)

These findings contradict the belief that there is always a connection between reliability and operating age – the belief which led to the idea that the more often an item is overhauled, the less likely it is to fail. In practice, this is hardly ever true. Unless there is a dominant age-related failure mode, fixed interval overhauls or replacements do little or nothing to improve the reliability of complex items.

Most maintenance professionals are aware of these findings, and coming to terms with the reality of randomness after decades in the bathtub. However, the fact that the bathtub curve still features in so many texts on maintenance is

testimony to the almost mystical faith that some people still place in the relationship between age and failure. In practice, this faith has two serious drawbacks, as follows:

- it leads to the belief that if we don't have any hard evidence at all about the existence of an age-related failure-mode, it is wise to overhaul the item anyway from time to time "just-in-case" such a failure mode does exist. This ignores the fact that overhauls are extraordinarily invasive undertakings that massively upset stable systems. As such, they are highly likely to induce infant mortality, and so cause the very failures which they seek to prevent. This is illustrated in Figure 4.




FIGURE 4

- at a more philosophical level, bathtub believers convince themselves that it is more conservative (in other words, safer) to assume that everything has a life – and hence to overhaul equipment on the basis of an assumed life – than to assume it could fail at random. After implementing overhaul programs based on this assumption, they then assume that no failures should occur between overhauls, and that any which do occur cannot be attributed to maintenance "because we overhauled it last week/last year/whenever". The possibility that the overhaul itself may be the cause of the failure is usually completely lost on such people. More seriously, they simply refuse to accept the most important conclusion associated with maxim 4, which is summarised below:

*In the absence of any evidence to the contrary, it is more conservative to develop maintenance strategies which assume that any failure can occur at **any** time (in other words, at random), and not to assume that it will only occur after some fixed amount of time in service.*

Maxim 5

OLD

Comprehensive data about failure rates must be available before it is possible to develop a really successful maintenance program

A surprising number of people believe that effective maintenance policies can only be formulated on the basis of extensive historical information about failure. Thousands of manual and computerised technical history recording systems have been installed around the world on the basis of this belief. It has also led to great emphasis being placed on the failure patterns discussed in the previous section of this paper. Yet from the maintenance viewpoint, these patterns are fraught with practical difficulties, conundrums and contradictions. Some of these are summarised below:


NEW

Decisions about the management of equipment failures will nearly always have to be made with inadequate hard data about failure rates

Sample size and evolution:

Large industrial processes usually possess only one or two assets of any one type. They tend to be brought into operation in series rather than simultaneously. This means that sample sizes tend to be too small for statistical procedures to carry much conviction. For new assets with high levels of leading-edge technology they are always too small.

These assets are also in a continuous state of evolution and modification, partly in response to new operational requirements and partly in an attempt to eliminate failures



which either have serious consequences or which cost too much to prevent. This means that the amount of time any asset spends in any one configuration is relatively short.

So actuarial procedures are not much use in these situations because the database is very small and constantly changing. (The main exception is undertakings using large numbers of identical items in an almost identical manner.)

Complexity:

The sheer number and diversity of assets present in most industrial undertakings means that it is simply not possible to develop a complete analytical description of the reliability characteristics of an entire undertaking – or even any major asset within the undertaking.

This is complicated by the fact that many functional failures are caused not by two or three but by two or three dozen failure modes. As a result, while it may be fairly easy to chart the incidence of the functional failures, it is a major statistical undertaking to isolate and describe the failure pattern which applies to each of the failure modes. This alone makes sensible actuarial analysis almost impossible.

Reporting failure:

Further complications arise due to differences in reporting policy from one organisation to another. For example, an item may be removed on one site because it is failing while on another site it is removed because it has failed.

Similar differences are caused by different performance expectations. A functional failure is defined as the inability of an item to meet a desired standard of performance. These standards can of course differ for the same asset if the operating context is different, so what we mean by failed will also differ. For instance, the pump shown in Figure 1 has failed if it is unable to deliver 300 litres per minute in one context and 350 litres per minute in another.

These examples show that what is a failure in one organisation - or even one part of an organisation - might not be a failure in another. This can result in two quite different sets of failure data for two apparently identical items.

The ultimate contradiction:

An issue which bedevils the whole question of technical history is the fact that if we are collecting data about failures, it must be because we are not preventing them. The implications of this are summed up most succinctly by Resnikoff (1978) in the following statement:

"The acquisition of the information thought to be most needed by maintenance policy designers – information about critical failures – is in principle unacceptable and is evidence of the failure of the maintenance program. This is because critical failures entail potential (in some cases, certain) loss of life, but there is no rate of loss of life which is acceptable to (any) organisation as the price of

failure information to be used for designing a maintenance policy. Thus the maintenance policy designer is faced with the problem of creating a maintenance system for which the expected loss of life will be less than one over the planned operational lifetime of the asset. This means that, both in practice and in principle, the policy must be designed without using experiential data which will arise from the failures which the policy is meant to avoid."

Despite the best efforts of the maintenance policy designer, if a critical failure actually does occur, Nowlan and Heap (1978) go on to make the following comments about the role of actuarial analysis:

"The development of an age-reliability relationship, as expressed by a curve representing the conditional probability of failure, requires a considerable amount of data. When the failure is one which has serious consequences, this body of data will not exist, since preventive measures must of necessity be taken after the first failure. Thus actuarial analysis cannot be used to establish the age limits of greatest concern – those necessary to protect operating safety."


This brings us to the ultimate contradiction concerning the prevention of failures with serious consequences and historical information about such failures: that successful preventive maintenance entails preventing the collection of the historical data which we think we need in order to decide what preventive maintenance we ought to be doing.

This contradiction applies in reverse at the other end of the scale of consequences. Failures with minor consequences tend to be allowed to occur precisely because they do not matter very much. As a result, large quantities of historical data are available concerning these failures, which means that there will be sufficient material for accurate actuarial analyses. These may even reveal some age limits. However, because the failures don't matter much, it is highly unlikely that the resulting fixed interval maintenance tasks will be cost effective. So while the actuarial analysis of this information may be precise, it is also likely to be a waste of time.

Conclusion:

Perhaps the most important conclusion to arise from the above comments is that maintenance professionals should turn their attention away from counting failures (in the hope that an elegantly constructed scorecard will tell us how to play the game in the future), towards anticipating or preventing failures which matter.

So to be truly effective, we simply have to get comfortable with the idea of uncertainty, and deploy strategies which enable us to deal with it confidently. We also need to recognise that if the consequences of too much uncertainty cannot be tolerated, then we must change the consequences. In extreme cases of uncertainty, the only way to do so may be to abandon the process concerned.



Maxim 6

OLD

There are three basic types of maintenance:

- predictive
- preventive
- corrective

Most of what has been written to date on the general subject of maintenance strategy refers to three – and only three – types of maintenance: predictive, preventive and corrective.

Predictive (or condition-based) tasks entail checking if something is failing. Preventive maintenance usually means overhauling items or replacing components at fixed intervals. Corrective maintenance means fixing things either when they are found to be failing or when they have failed.

However, there is a whole family of maintenance tasks which falls into none of the above categories.

For example, when we periodically activate a fire alarm, we are not checking if it is failing. We are also certainly not overhauling or replacing it, and nor are we repairing it.

We are simply checking if it still works.

Tasks designed to check whether something still works are known as "functional checks" or "failure-finding tasks". (In order to rhyme with the other three families of tasks, the author and his colleagues also call them "detective" tasks because they are used to detect if something has failed.)

Detective maintenance or failure-finding applies only to hidden or unrevealed failures, and hidden failures in turn only affect protective devices.

If one applies scientific maintenance strategy formulation techniques to modern, complex industrial systems, it is not unusual to find that up to 40% of failure modes fall into the hidden category. Furthermore, 80% of these failure modes usually require failure-finding, so up to *one third of the tasks generated by scientific maintenance strategy development programs – such as RCM2 – are detective tasks.*

NEW

There are four basic types of maintenance:

- predictive
- preventive
- corrective
- detective

On the other hand, the same analytical techniques reveal that it is not unusual for condition monitoring to be technically feasible for no more than 20% of failure modes, and worth the investment in less than half these cases. (This is not meant to imply that condition monitoring should be not be used – where it is good it is very, very good – but that we must also remember to develop appropriate strategies for managing the other 90% of our failure modes.)

A rather more troubling finding is that most traditionally derived maintenance programs provide for fewer than one third of protective devices to receive any attention at all (and then usually at inappropriate intervals). The people who operate and maintain the plant covered by these traditional programs are aware that another third of these devices exist but pay them no attention, while it is not unusual to find that no-one even knows that the final third exist. This lack of awareness and attention means that most of the protective devices in industry – our last line of protection when things go wrong – are maintained poorly or not at all.

This situation is completely untenable.

If industry is serious about safety and environmental integrity, then the whole question of detective maintenance – failure-finding – needs to be given top priority as a matter of urgency. As more and more maintenance professionals become aware of the importance of this neglected area of maintenance, it is likely to become a bigger maintenance strategy issue in the next decade than predictive maintenance has been in the last ten years.

Maxim 7

OLD

The frequency of condition-based maintenance tasks should be based on the frequency of the failure and/or the criticality of the item

When people are discussing the frequency of predictive (or condition-based) maintenance tasks, one often hears either – sometimes both – of the following statements:

- it doesn't fail so often, so we don't need to check it so often
- we need to check more critical plant more often than less critical plant.

In both cases, the speakers are wrong.

The frequency of predictive maintenance tasks has nothing to do with the frequency of failure and nothing to do with the criticality of the item. The frequency of any form of condition-based maintenance is based on the fact that most failures do not occur instantaneously, and that it is often possible to detect the fact that the failure is occurring during the final stages of deterioration.

NEW

The frequency of condition-based maintenance tasks should be based on the failure development period (also known as the "lead time to failure" or "P-F interval")

Figure 5 shows this general process. It is called the P-F curve, because it shows how a failure starts and deteriorates to the point at which it can be detected (the potential failure point "P"). Thereafter, if it is not detected and suitable action taken, it continues to deteriorate – usually at an accelerating rate – until it reaches the point of functional failure ("F").

The amount of time (or the number of stress cycles) which elapse between the point where a potential failure occurs and the point where it deteriorates into a functional failure is known as the P-F interval, as shown in Figure 6.

The P-F interval governs the frequency with which the predictive task must be done. The checking interval must be significantly less than the P-F interval if we wish to detect the potential failure before it becomes a functional failure.






FIGURE 5

The P-F interval can be measured in any units relating to exposure to stress (running time, units of output, stop-start cycles, etc), but it is most often measured in terms of elapsed time. For different failure modes, the P-F interval can vary from fractions of a second to several decades.

FIGURE 6

The amount of time needed to respond to any potential failures which are discovered also influences condition-based task intervals. In general, these responses consist of any or all of the following:



- take action to avoid the consequences of the failure
- plan corrective action so that it can be done without disrupting production and/or other maintenance activities

- organise the resources needed to rectify the failure.

The amount of time needed for these responses also varies, from a matter of hours (say until the end of an operating cycle or the end of a shift), minutes (to clear people from a building which is falling down) or even seconds (to shut down a machine or process which is running out of control) to weeks or even months (say until a major shutdown).

Unless there is a good reason to do otherwise, it is usually sufficient to select a checking interval equal to half the P-F interval. This ensures that the task will detect the potential failure before the functional failure occurs, while providing a net interval of at least half the P-F interval to do something about it. However, it is sometimes necessary to select a checking interval which is some other fraction of the P-F interval. For instance, Figure 7 shows how a P-F interval of 9 months and a checking interval of 1 month give a nett P-F interval of 8 months.

If the P-F interval is too short for it to be practical to check for the potential failure, or if the nett P-F interval is too short for any sensible action to be taken once a potential failure is discovered, then the condition-based task is not appropriate for the failure mode under consideration.




FIGURE 7

Maxim 8

OLD
If both are technically appropriate, fixed interval overhauls/replacements are usually both cheaper and more effective than condition-based maintenance

NEW
If both are technically appropriate, condition-based maintenance is nearly always both cheaper and more effective than fixed interval overhauls/replacements throughout the life of the asset

The new maxim 8 is now well understood by most maintenance professionals, and the change is really only recorded here for the sake of completeness. However, there still remain a small number of people who subscribe to the old maxim, so it is worth summarising briefly why the new maxim is valid. Perhaps the best way to do so is by means of an example.

Most countries nowadays specify a minimum legal tread depth for tyres (usually about 2 mm). Tyres which are worn below this depth must either be replaced or retreaded. In practice, truck tyres – especially tyres on similar vehicles in a single fleet working the same routes – show a close relationship between age and the onset of failure due to normal wear. Retreading restores nearly all the original resistance to normal wear, so the tyres could be scheduled for restoration after they have covered a set distance. This means that all the tyres in the fleet would be retreaded after they had covered a specified mileage, whether or not they needed it.





FIGURE 8

Figure 8 shows hypothetical failure data for such a fleet which show that most of the tyres last between 30 000 miles and 50 000 miles. If a fixed-interval retreading policy designed to prevent all failures due to normal wear is adopted on the basis of this information, all the tyres will be retreaded at 30 000 miles. However, this policy also means that many tyres would be retreaded long before it was really necessary. In some cases, tyres which could have lasted as long as 50 000 miles would be retreaded at 30 000 miles, so they could lose up to 20 000 miles of useful life.



On the other hand, it is possible to define a potential failure condition for tyres related to tread depth. Checking tread depth is quick and easy, so it is a simple matter to check the tyres (say) every 2000 miles and to arrange for them to be retreaded only when they need it. This would enable the fleet operator to get an average of 40 000 miles out of his tyres without endangering his drivers, instead of the 30 000 miles which he gets if he does the scheduled restoration task described above - an increase in useful tyre life of 33%. So in this case a predictive task is much more cost-effective than scheduled restoration.

This example demonstrates that predictive maintenance tasks should be considered first, for the following reasons:

- they can nearly always be performed without moving the asset and usually while it is in operation, so they seldom interfere with production. They are also easy to organise.
- they identify specific potential failure conditions so corrective action can be clearly defined before work starts. This reduces the amount of repair work to be done, and enables it to be done more quickly.
- by identifying equipment on the point of potential failure, they enable it to realise most of its useful life. The number of removals for potential failures is only slightly larger than that which would result from functional failures, so total repair costs and spares requirements are minimised.

Maxim 9

OLD

Serious incidents/catastrophic accidents which involve multiple equipment failures are usually the result of "bad luck" or "acts of God", and are hence unmanageable

In the past, major industrial accidents tended to be seen as just another part of the overall risk of doing business. It was felt that it was just too expensive (if not impossible) to analyse industrial systems in enough detail to be able to manage the risks with any credibility.

In more recent times, reliability professionals have developed powerful tools (such as *probabilistic* or *quantitative risk assessment*) to assess the cumulative probabilities of failure and associated overall levels of risk inherent in complex systems.

However, one limitation of these techniques, especially when applied to protected systems, has been a tendency to regard the probability of failure of the protected function and protective device as fixed. This leads to the belief that the only way to change the probability of multiple failures associated with such systems is to change the hardware (in other words, to modify the system), perhaps by adding more protection or by replacing existing components with ones which are thought to be more reliable.

In fact, it is now apparent that it is possible to vary both the probability of failure of a protected function and (espe-

NEW

To a considerable extent, the likelihood of a multiple failure is a manageable variable, especially in protected systems

cially) the downtime of the protective device by adopting suitable maintenance and operating policies. As a result, it is also possible to reduce the probability of multiple failures to almost any desired level within reason by adopting such policies. (Zero is of course an unattainable ideal.)

The probability which is considered acceptable for any multiple failure depends on its consequences. Sometimes levels of acceptability are specified by regulatory authorities, but in the vast majority of cases the assessment has to be made by the users of the asset. Since these consequences vary hugely from system to system, what is deemed to be acceptable varies equally widely. This means that there no universal standards of risk that can be applied to all systems of a particular type (at least, not yet).

But *someone* has to make a decision as to what level of risk is acceptable before it is possible to decide what must be done to design, operate and maintain protected systems. (In fact, merely persuading senior people to accept that this is a manageable variable *which they must therefore manage* is currently one of the biggest challenges facing maintenance professionals.)

Maxim 10

OLD

The quickest and surest way to improve the performance of an existing "unreliable" asset is to upgrade the design


As we learn more and more about what must be done to maintain our assets successfully, we learn just how many maintenance problems could have been avoided or eliminated on the drawing board. This is leading to a long overdue recognition that equipment designers should consider not only what must be done to create new equipment that works, but also what must be done to keep it working.

However, this realisation shows a sometimes alarming

NEW

It is nearly always more cost-effective to try to improve the performance of an unreliable asset by improving the way it is operated and maintained, and only to review the design if this cannot deliver the required performance

tendency to be applied inappropriately to the management of existing assets. A small but vocal minority of people seem to believe that the best way to deal with reliability problems is to go straight back to the drawing board, without stopping to ask whether improved maintenance practices may not in fact be the best solution to the problem.



In practice, maintenance should be considered before redesign for three reasons.

- most modifications take from 6 months to 3 years from conception to commissioning, depending on the cost and complexity of the new design. On the other hand the maintenance person who is on duty today has to maintain the equipment as it exists today, not what should be there or what might be there some time in the future. So today's realities must be dealt with before tomorrow's redesigns
- most organisations face many more apparently desirable design improvement opportunities than are physically or economically feasible. Trying first to extract the desired performance from assets as they are currently configured does much to help develop rational priorities for these projects, especially because it separates those that are essential from those that are merely desirable.

- there is no automatic guarantee that a new design will actually solve the problem. The scrapyards of the world are littered with modifications that "didn't quite work" – silent testimony that second guessing the original designers often becomes an expensive exercise in futility.

However, this is not meant to imply that we should never redesign existing assets. Occasions often arise where the desired performance of an asset exceeds its inherent reliability, in which case no amount of maintenance can deliver the desired performance. In such cases "better" maintenance cannot solve the problem, so we need to look beyond maintenance for the solutions. Options include modifying the asset, changing operating procedures, or simply lowering our expectations and deciding to live with the problem.

Maxim 11

OLD Generic maintenance policies can be developed for most types of physical asset

The belief that generic maintenance policies can and should be applied to most types of assets lies at the heart of nearly all traditional maintenance programs. For instance, how often does one hear people say things like "the maintenance policy which we apply to all our pumps is X" or "we have a type Y calibration policy for all our instruments"?

However, the correct use of techniques like RCM2 to develop maintenance programs quickly shows why the inappropriate use of generic maintenance is one of the main reasons why so many traditional programs do not achieve their full potential. The main reasons are as follows:

- *functions*: the narrative accompanying Figure 1 on page 2 explains how a pump can have one performance expectation in one location, and a different expectation in another location. Different performance standards of this type inevitably call for different standards of maintenance. (This is especially true where otherwise identical machines are used to produce products which have widely differing quality standards.)
- *failure modes*: when otherwise identical equipment is used in even slightly different locations (an area of high humidity, an unusually dusty environment) or to perform slightly different tasks (cutting a harder than usual metal, operating at a higher temperature, pumping a more abrasive or a more acidic liquid), the possible failure modes vary drastically. This in turn means that failure management strategies need to vary accordingly
- *failure consequences*: different failure consequences also demand different maintenance strategies. This is illustrated by three otherwise identical pumps in Figure 9. Pump A stands alone, so if it fails, operations will be affected sooner or later. As a result the users and/or maintainers of Pump A are likely to make some effort *to anticipate or prevent its failure*. (How hard they try will be governed both by the effect on operations and by the severity and

NEW Generic policies should only be applied to identical assets whose operating context, functions and desired standards of performance are also identical

frequency of the failures of the pump.)

However, if pump B fails, the operators simply switch to pump C, so the only consequence of the failure of pump B is that it must be repaired. As a result, it is probable that the operators of B would at least consider letting it *run to failure* (especially if the failure of B does not cause significant secondary damage.)

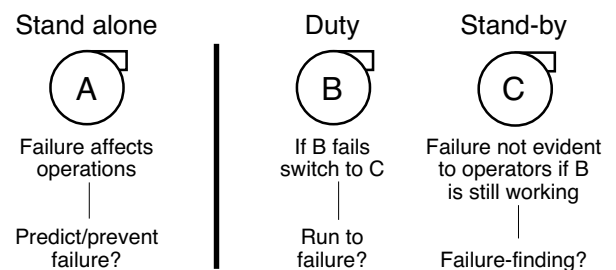



FIGURE 9

On the other hand, if pump C fails while pump B is still working (for instance if someone cannibalises a part from C), it is likely that the operators will not even know that C has failed unless or until B also fails. To guard against this possibility, a sensible maintenance strategy might be *to run C from time to time to find out whether it has failed*.

This example shows how three identical assets can have three totally different maintenance policies, because the failure consequences are different in each case.

- *maintenance tasks*: different organisations – or even different parts of the same organisation – seldom employ people with identical skillsets. This means that people working on one asset may prefer to use one type of proactive technology (say high-tech condition monitoring) to anticipate failure, while another group working on an



identical asset may be more comfortable using another (say a combination of performance monitoring and the human senses). It is surprising how often this does not matter, as long as the techniques chosen are cost-effective. In fact, many maintenance organisations are starting to realize that there is often more to be gained from ensuring that the people doing the work are comfortable with what they are doing than it is to compel everyone to do the same

thing. (The validity of different tasks is also affected by the operating context of each asset. For instance, think how background noise levels affect checks for noise.)

All of this means that special care must be taken to ensure that the operating context, functions and desired standards of performance are all virtually identical before applying a maintenance policy designed for one asset to another.

Maxim 12

OLD

Maintenance policies should be formulated by managers and maintenance schedules drawn up by suitably qualified specialists or external contractors (a top-down approach)

The traditional maintenance planning department typified the old maxim in action. A key responsibility of this department was usually to compile maintenance schedules for all the equipment in the plant. Maintenance planners often devoted immense amounts of time and energy to this exercise (the author knows – he was one once). However, more often than not, their schedules died when they reached the shop floor. This happened for two main reasons:

- *technical validity*: the planners who wrote the schedules were usually out of touch with the equipment (if they had ever been in touch to start with). As a result, they often had a less than adequate understanding of the functions, the failure modes and effects and the failure consequences of the assets for which their schedules were being written. This meant that the schedules were usually generic in nature, so people who were supposed to do them often saw them as being incorrect if not totally irrelevant
- *ownership*: people on the shop floor (supervisors and craftsmen) tended to view the schedules as unwelcome paperwork which appeared from some ivory tower and disappeared after it was signed off. Many of them learned that it was more comfortable just to sign off the schedules and send them back than it was to attempt to do them. (This led to inflated schedule completion rates which at least kept the planners happy.) The main reason for the lack of interest was undoubtedly sheer lack of ownership.

NEW

Maintenance policies should be formulated by the people closest to the assets. The role of management is to provide the tools to help them make the right decisions, and to ensure that the decisions are sensible and defensible

The only way around the problems of technical invalidity and lack of ownership is to involve shop floor people directly in the maintenance strategy formulation process. This is because they are the ones who really understand how the equipment works, what goes wrong with it, how much each failure matters and what must be done to fix it.

The best way to access their knowledge on a systematic basis is to arrange for them to participate formally in a series of meetings. However, it is essential to ensure that these meetings do not just become another bunch of inconclusive talkfests. This can be done by arranging for the participants to be trained professionally in the use of RCM2, and to provide them with skilled guidance in the application of this technique.

Done correctly, this not only produces schedules with a much higher degree of technical validity than anything that has gone before, but it also produces an exceptionally high level of ownership of the final results.

(A word of caution at this stage: It is wise to steer clear of the temptation to use external contractors to formulate maintenance strategies. An outsider's sheer ignorance of almost all the issues discussed in connection with maxims 1 through 11 insofar as they affect *your* plant means that all you are likely to get is a set of elegantly completed forms that amount to little or nothing. Using such people to develop maintenance programs is to wander into the hazy – and dangerous – region where delegation becomes abdication.)

Maxim 13

OLD


The maintenance department on its own can develop a successful, lasting maintenance programme

Maxim 12 above reminds us of the need to involve shop floor people as well as managers in the maintenance strategy development process. Maxim 13 concerns what is often a much more difficult challenge in many organizations – the almost impenetrable divide between the maintenance and production functions.

NEW

A successful, lasting maintenance program can only be developed by maintainers and users of the assets working together

In fact, as the very first maxim in this series makes clear, maintenance is all about ensuring that assets continue to function to standards of performance required by the users. In nearly every situation, the "users" are the production or operations functions. This means that modern maintenance strategy formulation starts by asking the



users what they want, with a view to setting up asset management programs whose sole objective is to ensure that the users get what they want. Clearly, for this to be possible, the users must be prepared to specify exactly what they require. (If they do not bother to state the performance they require from each asset with adequate precision, then of course they cannot hold maintenance responsible for delivering that performance.) Both users and maintainers must also take care at this stage to satisfy themselves that the asset is capable of delivering the required performance to begin with.

In addition to spelling out what they want the asset to do, operators also have a vital contribution to make to the rest of the strategy formulation process.

By participating in a suitably focused FMEA, they learn a great deal about failure modes caused by human error, and hence what they must do to stop breaking their machines. They also play a key role in evaluating failure

consequences (evidence of failure, acceptable levels of risk, effect on output and product quality), and they have invaluable personal experience of many of the most common warnings of failure (especially those detected by the human senses). Finally, involvement in this process helps users to understand much more clearly why they sometimes need to release machines for maintenance, and also why operators need to be asked to carry out certain maintenance tasks.

In short, from a purely technical point of view, it is rapidly becoming apparent that it is virtually impossible to set up a viable, lasting maintenance program in most industrial undertakings without involving the users of the assets. (This focus on the user – or customer – is of course the essence of TQM.) If their involvement can be secured at all stages in the process, that notorious barrier rapidly starts to disappear and the two departments start to function, often for the first time ever, as a genuine team.

Maxim 14

OLD

Equipment manufacturers are in the best position to develop maintenance programs for new physical assets

A universal feature of traditional asset procurement is the insistence that the equipment manufacturer should provide a maintenance program as part of the supply contract for new equipment. Apart from any thing else, this implies that manufacturers know everything that needs to be known to draw up suitable maintenance programs.

In fact, manufacturers are usually at best no better informed than traditional maintenance planners about the operating context of the equipment, desired standards of performance, context-specific failure modes and effects, failure consequences and the skills of the user's operators and maintainers. More often the manufacturers know nothing at all about these issues. As a result, schedules compiled by manufacturers are nearly always generic, with all the drawbacks discussed under maxim 11.

Equipment manufacturers also have other agendas when specifying maintenance programs (not least of which is to sell spares). What is more, they are either committing the users' resources to doing the maintenance (in which case they don't have to pay for it, so they have little interest in minimizing it) or they may even be bidding to do the maintenance themselves (in which case they have a vested interest in doing as much as possible).

This combination of extraneous commercial agendas and ignorance about the operating context means that maintenance programs specified by manufacturers tend to embody a high level of over-maintenance (sometimes ludicrously so) coupled with massive over-provisioning of spares. Most maintenance professionals are aware of this problem. However, despite our awareness, most of us persist in demanding that manufacturers provide these programs, and then go on to accept that they must be followed in order for warranties to remain valid (and so bind ourselves contractually to doing the work, at least for the duration of the warranty period).

NEW

Equipment manufacturers can only play a limited (but still important) role in developing maintenance programs for new assets


None of this is meant to suggest that manufacturers mislead us deliberately when they put together their recommendations. In fact, they usually do their best in the context of their own business objectives and with the information at their disposal. If anyone is at fault, it is really us – the users – for making unreasonable requests of organisations which are not in the best position to fulfil them.

A small but growing number of users solve this problem by adopting a completely different approach to the development of maintenance programs for new assets. This entails asking the manufacturer to supply experienced field technicians to work alongside the people who will eventually operate and maintain the equipment, to use RCM2 to develop programs which are satisfactory to both parties.

When adopting this approach, issues such as warranties, copyrights, languages which the participants should be able to speak fluently, technical support, confidentiality, and so on should be handled at the request for proposal/contracting stage, so that everyone knows what to expect of each other.

Note the suggestion to use field technicians rather than designers (designers are often surprisingly reluctant to admit that their designs can fail, which reduces their ability to help develop a sensible failure-management program). The field technicians should of course have unrestricted access to specialist support to help them answer difficult questions.

In this way, the *user* gains access to the most useful information that the manufacturer can provide, while still developing a maintenance program which is most directly suited to the context in which the equipment will actually be used. The *manufacturer* may lose a little in up-front sales of spares and maintenance, but will definitely gain all the long-term benefits associated with improved equipment performance, lower through-life costs and a much better understanding of the real needs of his customer. A classic win-win situation.



Maxim 15

OLD

It is possible to find a quick, one-shot solution to all our maintenance effectiveness problems

If one takes a moment to review the breadth and depth of the paradigm shift implicit in the foregoing paragraphs, it soon becomes apparent just how far most organisations have to move in order to adopt the new maxims. It simply cannot happen overnight.

Nonetheless, RCM2 enables most users to put into practice most of the changes described in this paper in less than a year, and to recoup the associated investment in a matter of months (if not weeks). However, so great is the obsession which many businesses have with quick results that even this is just not quick enough. Financial, regulatory and competitive pressures all conspire to make people want lasting change *now*. As a result, people tend to fall into the last and often saddest trap of all – the quest for shortcuts.

Unfortunately, in the experience of the author, this quest is invariably counterproductive. Firstly, the development of the "shortcut" itself takes time – time which is spent reinventing perfectly round wheels instead of getting on

NEW

Maintenance problems are best solved in two stages: (1) change the way people think (2) get them to apply their changed changed thought processes to technical/process problems – one step at a time

with the job of improving asset performance. Secondly, shortcuts nearly always end in sub-optimal solutions – so much so that they often result in little or no change at all.

In fact, people who seek an effective, enduring maintenance program which has universal support should not lose sight of the fact that improvement is a journey, not a destination (the essence of the Kaizen philosophy). In the field of asset management, this means that we should turn away from the search for a silver bullet which will blow away all our problems in an instant – the 1 x 100% solution. Success is far more likely to be assured if we think in terms of a silver shotgun shell, and set out to blow away our problems one pellet (or one failure mode) at a time – the 1000 x 0.1% approach which virtually guarantees 100% success to those with the patience to try it.

Thank you – and good hunting!

References:

Moubray J M (1991) "*Reliability-centred Maintenance*". Butterworth-Heinemann, Oxford

Nowlan F S and Heap H (1978) "*Reliability-centered Maintenance*". National Technical Information Service, US Department of Commerce, Springfield, Virginia

Resnikoff H L (1978) "*Mathematical Aspects of Reliability-centered Maintenance*". Dolby Access Press, Los Altos, California