

Cost/Risk Optimisation (C/RO)

John Woodhouse

1 Introduction

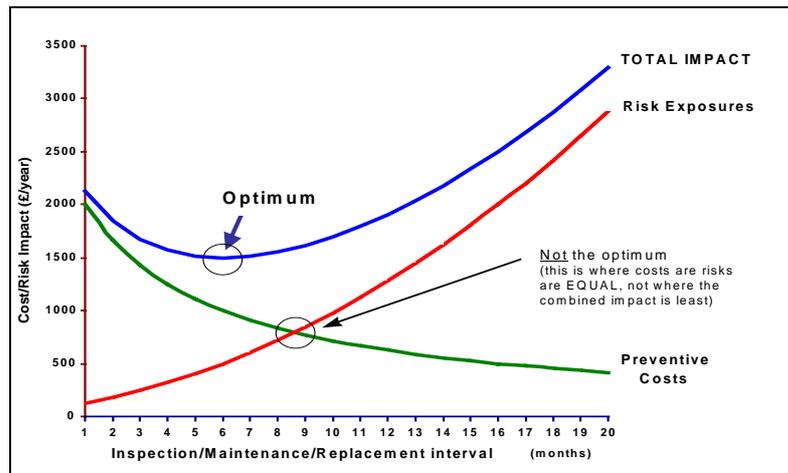
Most engineering, maintenance and operating decisions involve some aspect of cost/risk trade-off. Such decisions range from evaluating a proposed design change, determining the optimal maintenance or inspection interval, when to replace an ageing asset, or which and how many spares to hold. The decisions involve deliberate expenditure in order to achieve some hoped-for reliability, performance or other benefit. We may know the costs involved, but it is often difficult to quantify the potential impact of reduced risks, improved efficiency or safety, or longer equipment life. Not only are the benefits difficult to quantify, but the objectives often conflict with each other (we could clean the heat exchanger more often to achieve better performance, but the cleaning may damage the tubes and shorten their life). Finding the optimal strategy is difficult, therefore, but the wrong maintenance interval will result in excessive costs, risks or losses.

The European collaboration project “MACRO”, has developed a structured set of procedures to make sure that the right questions are asked, and sophisticated “what if?” analysis tools to calculate the optimum combinations of equipment reliability, performance, lifespan and cost. Specifically designed to be used where hard data is poor, and engineering judgement or range-estimates provide the main raw material, these Cost/Risk Optimisation techniques are the acknowledged ‘best practice’ when justifying design, maintenance, condition monitoring, replacement or spares decisions. The following ‘mini-guide’ outlines the underlying concepts of the approach, with illustrations of their application to a variety of decisions.

2 What is “Optimisation”?

The first concept that needs clarifying is the meaning of “optimum”. The word is often used very loosely in phrases such as “the optimum maintenance strategy” or “the optimum performance”. In areas where there are *conflicting interests*, such as pressures to reduce costs at the same time as the desire to increase reliability/performance/safety, an optimum represents some sort of compromise. It is clearly impossible to achieve the component ideals - zero costs at the same time as total (100%) reliability/safety etc. Reliability costs money, or, to put it the other way around, to spend less money we must choose what *not* to do or achieve. The resulting and inevitable trade-off can be drawn graphically (see figure 1), but we must be careful with the labelling.

2.1.1 Optimum is defined as minimal Total Business Impact



Many such diagrams show the bottom of the Total Impact curve neatly aligned above the cross-over point of the conflicting components, giving rise to confusion as to where and what is the true optimum. The Total Impact is the sum of the costs and risks etc. When this sum is at a minimum, we have defined the optimum combination of the components: the best value mixture of costs incurred, residual risks, performance losses etc. Crossover points do not signify the optimum; they merely show where the components are equal (i.e. the risks or losses have the same value as the amounts being spent to control them). The concepts of ‘balancing costs and risk’ or finding a ‘breakeven point’ are dangerous, therefore, because they imply this equivalence point as a target, rather than focus on the best value-for-money *combination*.

3 Why this is difficult to find

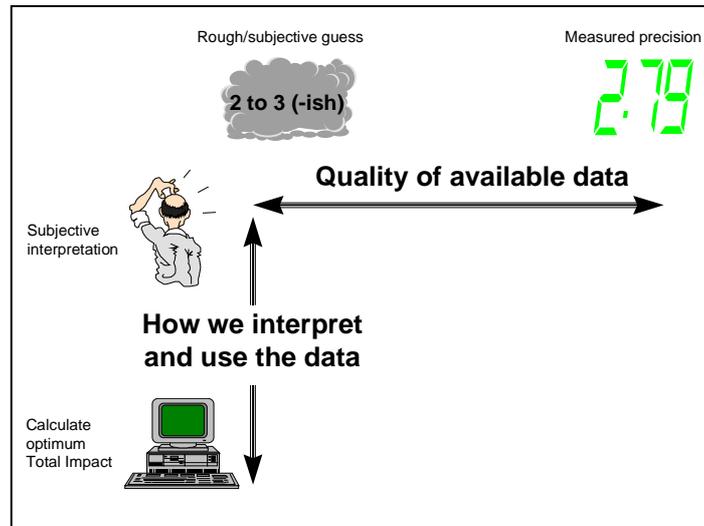
If we knew exactly what the risks were, and what they are worth, we could calculate the optimum amount of risk to take, and costs to incur. Similarly, we could make better (more optimal) decisions if we knew the value of improved performance, longer life, greater safety or quality. But the risks are difficult to quantify and many of the factors are “intangible”; i.e. it is extremely difficult to place an economic value on them. The first barrier, therefore, to cost/risk optimisation is the

- **lack of relevant hard data.**

This is not the only constraint. Whether or not there is suitable information, the complexity of the interactions is also a barrier. Reliability is a complex subject: the effects of one failure mode upon the probabilities of suffering other forms of failure involve nasty mathematics. These relationships have been known for a long time (over 20 years) but, especially in the absence of useful data, they have been limited in usefulness to academic or special case studies. So, whatever the state of information, the additional problem is

- **how we would use the data if it were available.**

3.1.1 “What data?” versus “How would we use it?”

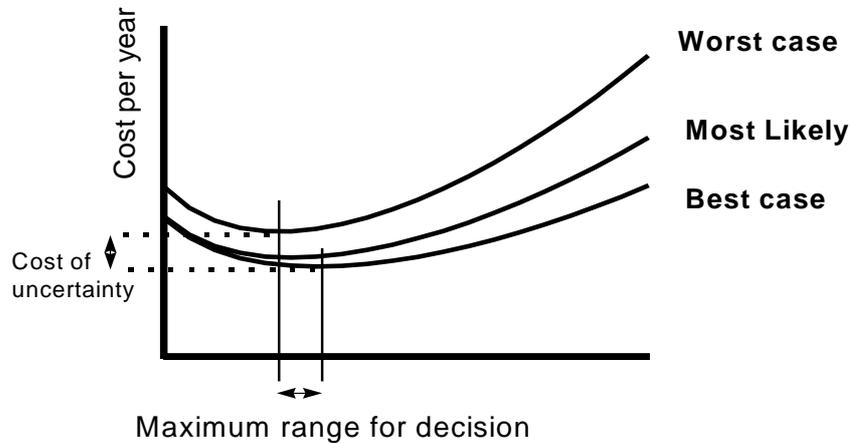


These problems appear to be linked (if we do not have suitable data, how can we improve the usage mechanisms?) but have, in fact, quite separate effects. The traditional reaction to poor data and subjective decision-making is to a) start collecting more/better data and b) hope that it will somehow tell us what to do. *This approach does not work.* Without knowing how we would use it, how do we know what data is worth collecting in the first place? Even if we were lucky enough to guess correctly on the data that is needed, how (and when) would we know that we had collected enough? What is “enough”, and is it physically/economically possible to collect it? Without a clear idea of how it will be used, and the sensitivity to data inaccuracy, it is impossible to say what data is needed, and to what precision. The first challenge is therefore the understanding of what information is required for specific decisions, and how it should be used. This issue can be addressed by designing and using templates and checklists; to **make sure that the right questions are asked in the first place.**

Even if hard data is not available, there is a considerable volume of knowledge in the operators, maintainers and engineers. This can be obtained in the form of range estimates or “worst case” and “best case” extremes of opinion. With a range of possible interpretation, we can see if the decision is affected – whether we need to dig deeper, and at what cost. This is achievable if we have the means rapidly to calculate the Total Impact for different assumptions. We must adopt a “What if?” approach to the problem: try the optimistic extreme and the pessimistic – does the data uncertainty have a significant effect?

The calculations require specialist software tools – the maths are too hard for mental arithmetic or even spreadsheets. Given their availability, however, even rough or range estimates can be explored for their effect. Sensitivity testing reveals which pieces of information are important, and which have little or no effect upon the relevant decision. Even with rough guesses, we can find the ‘envelope’ in which the optimum strategy must lie. In a surprising proportion of cases, this reveals that the decision is robust over the full range of data interpretation (i.e. the range estimates are enough to arrive at a firm and provable conclusion).

3.1.2 Using range estimates to locate optimum strategy



3.2 Example: pump overhauls

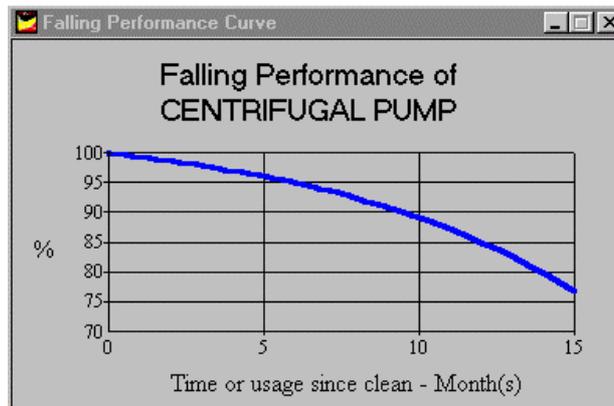
If the performance of a pump deteriorates as its impeller becomes fouled, and the reduced capacity is having an effect upon production or process efficiency, then there must be an optimum time to clean the impeller. To determine the best maintenance strategy, we need to know how the performance falls with time or use, the economic effect of the losses (perhaps the pump has to operate for longer to deliver the required volumes, or maybe the drive motor draws more electricity to compensate). We also need the cost of cleaning (including any operational downtime to do it). Some of this information may be known if there is some operational experience, but otherwise it must be range-estimated and explored for sensitivity:

3.2.1 Data estimates:

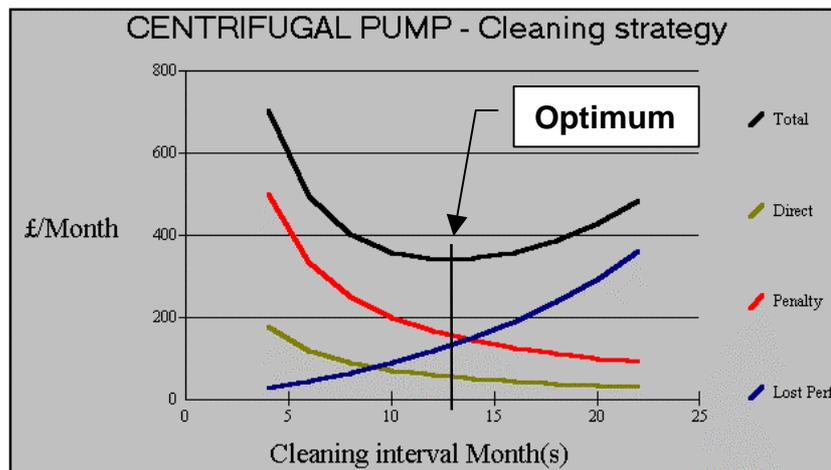
- By 6 months of operation, pump performance is 5-10% down, and this is likely to accelerate if left further.
- 10% lost performance is worth £10-30/day in extra energy/production impact or extended operating costs.
- The costs of cleaning or overhaul are £6-800 in labour and materials, and 2-3 hours downtime to swap over to an alternative pump.

3.2.2 Calculating the impact

The first step involves ‘fitting’ a performance curve to the examples given:

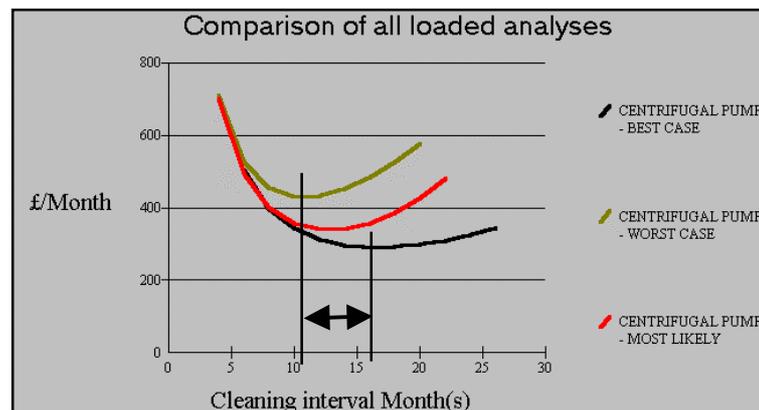


Then, a series of calculations can show the Total Impact of performance losses, cleaning costs and equipment downtime for various maintenance intervals:



3.2.3 Sensitivity testing

The “worst case” and “best case” interpretations combine the extremes of all the range-estimates. They show that the cleaning interval must be between 11 and 16 months. No interpretation of the problem could justify more, or less, frequent cleaning:



3.2.4 Extensions to include reliability characteristics

The optimisation can be extended to include the reliability of the pump, with a variety of failure modes. The complexity (which can be handled by modern software tools) lies in the *interaction* between failure risks. Historically, reliability studies have been obliged to simplify their assumptions to the point of impracticality – assuming just one mode of failure, only randomness, or all repairs “reset the clock” to “as new” condition. Real life is much more complicated: maintenance-induced failures (such as misalignments, faulty work or materials) influence the rates of subsequent deterioration. What might seem “random” in one view (e.g. lubrication failures of young or old pumps) is certainly not random in another context (e.g. time since last checking of the oil level).

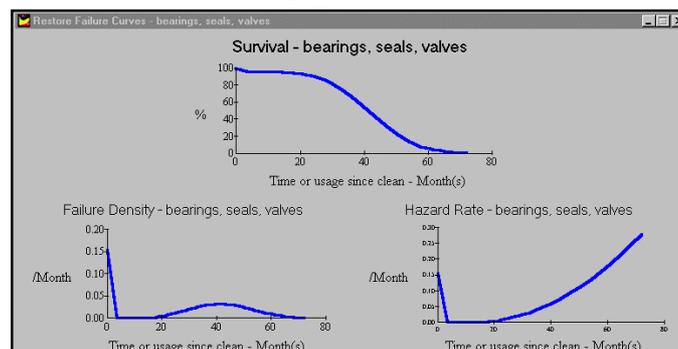
Maintenance options are nearly always faced with combinations of (interacting) failure modes. To make sense of the navigation, therefore, a disciplined structure is vital. This has been developed as part of the MACRO programme: it reveals how *cumulative* effects are much more useful than estimates of specific probabilities. We can make estimates of “how long things will last” much more easily than “the chance of an xyz failure is”. The cumulative information is called the Survival Curve and the following is a typical description of a complex mixture of failure modes:

With respect to **time since last overhaul**;

- Typically 5-10% will need repeat work (maintenance errors etc) in the first month
- Most (80-0%) survive the first year or two without failure
- Not many (less than 20%) would last longer than 5 years without some sort of failure.

Computer software can fit a curve to this Survival information, and calculate the pattern of risks that would be necessary for these symptoms to be achieved. In fact there are two further forms of this reliability data. The Hazard Rate is the “conditional” chance of failure, assuming the equipment has survived so far. The Failure Density quantifies how many would fail at different time points (i.e. a combination of how many reach each point, and the risks they face). It is the Hazard Rate that is needed for decisions about how long to leave the equipment in service (and risk failure), or deliberately maintain/replace it instead. APT-MAINTENANCE software is currently the only tool available to perform this analysis comprehensively:

3.2.5 Different views of reliability patterns:



4 Other cost/risk optimisation areas

The same essential process applies to a wide range of decisions and the MACRO project has developed seven sets of procedures, training courses and analysis tools to cope with the variety of “what if?” investigations that are necessary. These are

- Project cost/benefit and risk evaluation: 1-off investments, change proposals, modifications or procedural changes.
- Asset replacement and Life Cycle Costing: repair versus replace options, life extension projects, alternative cost/performance designs etc.
- Planned maintenance strategy: preventive versus on-failure, preventive versus predictive, optimal maintenance intervals, impact of different designs, maintenance procedures, quality etc.
- Inspection, testing and condition monitoring: optimal inspection or testing intervals, condition reaction points, alternative monitoring methods.
- Shutdowns and work grouping: optimal combinations of work content and timing, opportunities and alignment, shutdown intervals
- Spares and materials management: stock-holding, purchasing and supply options, spares ‘pooling’, centralised versus distributed warehousing, min/max and reorder quantities.

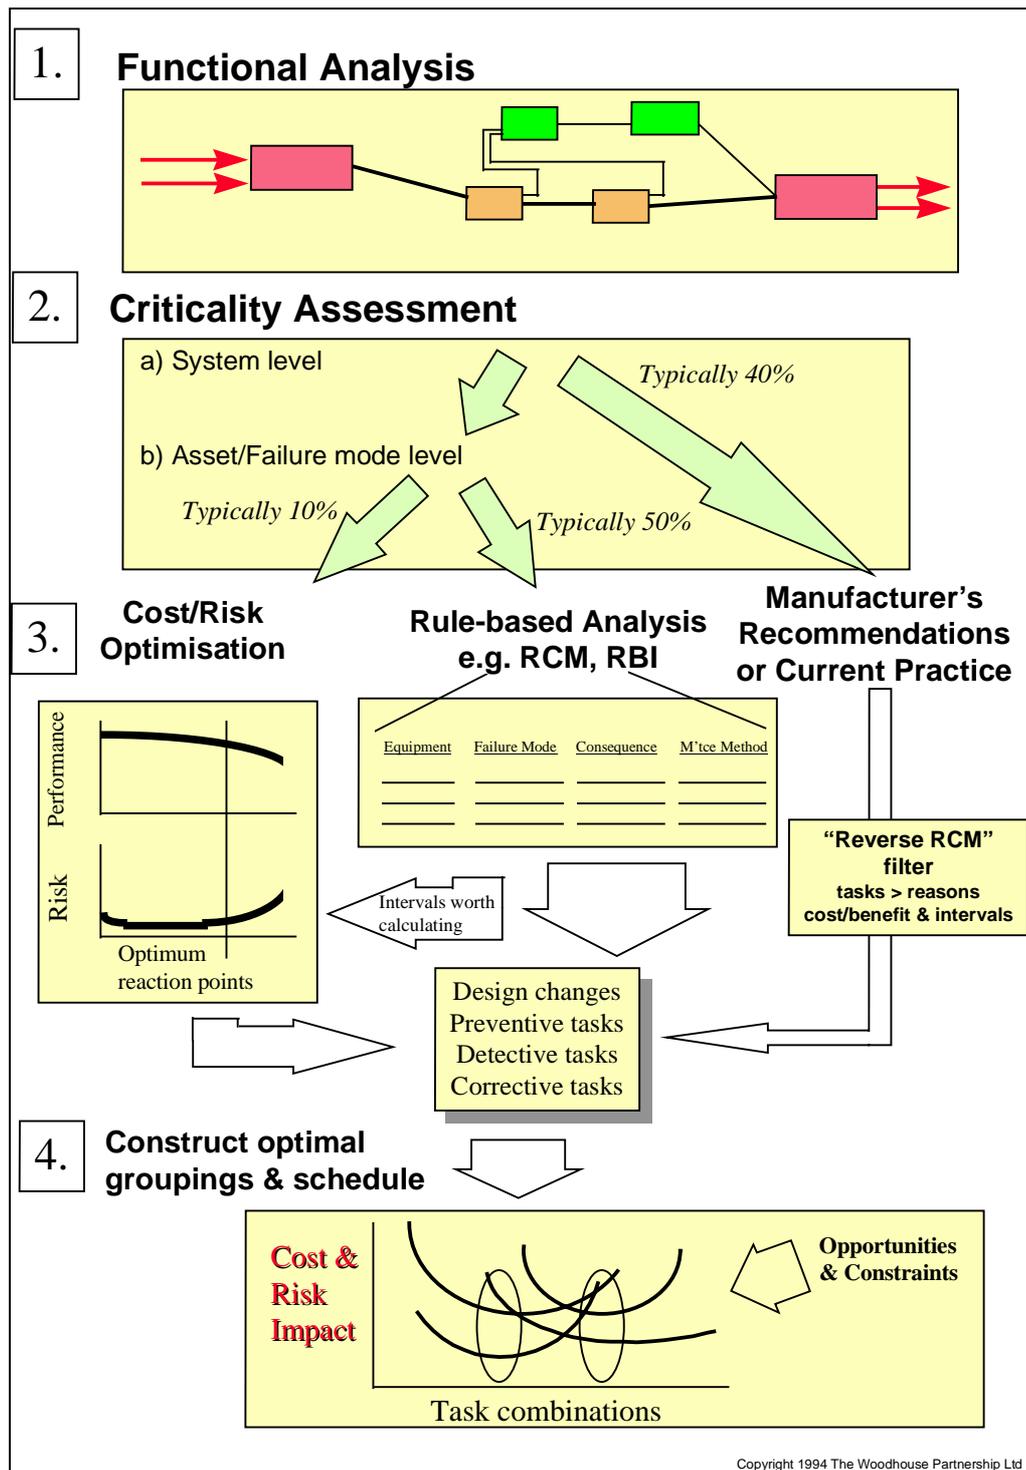
5 Where C/RO fits into reviews of maintenance strategy

The conclusion being reached by an increasing number of organisations (and the MACRO project) is that we need a mixture of methods to determine what maintenance is worth doing and when. No single formula yet on offer was found to be suited to different industries, or even to different processes, plant types or departments within the same company. The amount of analysis effort, and the payback of such analysis, is clearly dependent upon the importance of arriving at the correct strategy. Criticality prioritising of systems, equipment and failure modes is vital to avoid ‘analysis paralysis’ and loss of direction. The overall flowchart that has emerged from the MACRO work is a three-tier analysis. Dependent upon *process* or *functional* criticality, different levels of effort should be applied:

- A. At the top end, perhaps 5-10% of the most vital corporate functions, quantitative risk and performance analysis is warranted. This is where Cost/Risk Optimisation methods should be used straight away – the quick-win, high impact strategies, where error has significant costs or risks attached.
- B. For the next 40-60% of ‘core business’ activities, template and rule-based methods (such as RCM and RBI) are more appropriate, particularly if supplemented with some economic analysis or justification of the resulting strategies. Cost/Risk Optimisation methods supply this final step – putting the £-sign into the equation (setting appropriate intervals, or screening proposed design changes etc.)

- C. At the lower levels of criticality, not even the simple FMEA study is worthwhile - a cruder but quicker 'filter' is required. "Reverse-RCM" or "Review of Existing Maintenance" (REM) provide such a simple sanity check: to verify that there is a) a valid reason for doing the job and b) the cost & interval are reasonable in relation to the risks/consequences if it were not performed.

5.1.1 Strategy methods depending upon Operational Criticality



6 Case studies and examples

The following is a sample of real-life application of Cost/Risk Optimisation methods in different industries:

- N.Sea oil industry: 400 change proposals screened in 6-week period. £2.5million of unjustifiable expenditure avoided.
- Petrochemical refining: entire investment budget of largest refinery in the world evaluated on project-by-project basis. Average evaluation time reduced from 8 hours to 30 minutes *with higher confidence and auditability in the resulting conclusions.*
- Water services utility: maintenance intervals and safety testing strategy for borehole pumping and chlorination plant. Pump maintenance and electricity consumption optimised, saving over \$500,000/year.
- Electrical distribution company: high voltage protection equipment showed justifiable reduction in testing by 50%.
- Process manufacturing company: corrosion monitoring of critical vessels and pipework needed to be doubled, with net risk reduction worth hundreds of thousands per year.
- Industry materials stockist: 60% reduction in slow-moving inventory *without impact on risk exposures.*

7 Suggested further reading

MACRO project papers in wide variety of conferences (note: the MACRO component procedures are being published as a CD-ROM and on the World Wide Web during 1999: monitor the websites of www.twpl.co.uk (training) and www.aptools.co.uk (analysis tools) to learn more.

MAINTENANCE & ASSET MANAGEMENT Journal (magazine of the Institute of Asset Management)

“Managing Industrial Risk”, John Woodhouse, Publ. Chapman & Hall 1993

American Petroleum Institute Recommended Practice RP580 “Risk-Based Inspection”.