

Decision-support: technology and people in solving problems and making better Asset Management decisions

John Woodhouse, The Woodhouse Partnership Ltd
Presented to the ERTC conference, Berlin, March 2005

1 Introduction

Good decisions are at the heart of good management – but what is a ‘good’ decision? We certainly want to do the right things (be effective), and we want to do things right (be efficient). Of these two goals, success or failure most often rests on the first; choosing what to do, or what to spend, where and when (in other words, doing the right things, for the right reasons, at the right time). These decisions have a more profound effect on our results than efficiency improvements in *how* we do it. Yet it is still common to find our improvement efforts are directed only at greater efficiency (doing things quicker, better or cheaper) rather than challenging what it is that we do in the first place. If we focus too much on delivery efficiency, we run a significant risk *doing the wrong things 10% cheaper or quicker!*

The challenges of determining what is worth doing and when are significant. We don’t have all the data we would like, life and the future are both uncertain, competing influences are complex, there are short- and long-term conflicts in objectives or personal agendas, and stakeholders have incompatible expectations!

This paper looks at which methods or tools currently work best in which circumstances and, in particular, how we can cope with risk and uncertainty, data unavailability, the better use of ‘tacit knowledge’ and the incorporation of long-term consequences into short-term decisions.

2 A bit of background

Since the second world war, Deming, Juran and co. introduced quality management and statistical process control, formalising many of the concepts of fact (data) based problem-solving and decision-making. Kepner Tregoe and Edward de Bono have encouraged more logical organising of the ideas and options, giving rise to, among other tools, decision trees and dependency models. Some of these have been developed into problem-specific ‘rules’ to encourage greater decision consistency and thoroughness – such as Reliability Centred Maintenance for the selection of maintenance strategies, developed by the civil aviation sector in the 1970’s. In the 1990’s, the North Sea Oil and Gas sector developed an ISO standard¹ for Life Cycle Costing, the American Petroleum Institute published their Risk Based Inspection guidelines² and the Safety/Instrumentation world developed IEC61511 to help decision-making in levels of safety protection.

¹ ISO15663 obtainable from www.bsi-global.com

² API RP580/581

In the meantime, of course, computers have been increasingly useful – both in the easier storage and examination of data (relational databases, reporting and pattern-finding tools), and in the manipulations, calculations and simulations that enable “what if?” studies, cost/benefit appraisal and performance predictions (spreadsheets, modelling tools etc). In the specific area of Asset Management decision-making, the European MACRO³ project of the late 1990’s delivered an extraordinarily effective mix of structured quantification methods (‘how to ask the right questions’) and very flexible “what if?” calculator tools - in 42 areas of asset management decision-making. Since then, technology and management science have moved on even further, and this paper is a review of the methods, common sense and combined toolkit that is now available to reduce errors, truly optimise what we do and increase transparency in complex decisions.

3 Decision types & different approaches

There are now hundreds of clever analytical aids, methodologies, standards and In order to sort out the confusing language and overly optimistic claims of technical enthusiasts promoting their particular piece of the puzzle, I have clustered the different approaches to decision support into some simple families. Using a few examples of relevant or familiar tools, I will then discuss their strengths and weaknesses and ‘best fit’ roles within the Asset Manager’s decision toolbox.

Two main categories of decision-support aids need to be considered straight away. The aids help us to:

1. *detect, diagnose or characterise the problem,*
2. *choose, justify or optimally time/target the appropriate medicine*

The first category covers many condition monitoring, data collection, inspection, maintenance history, reporting, pattern-finding and root cause analysis tools. They aim to assist our decision-making by providing greater clarity about the nature of the ‘illness’ or opportunity to improve. This has two stages – the detection and the diagnosis. Detection aids comprise a wide range of monitoring, reporting and performance indicators, but they do all require pre-consideration of what symptoms represent a ‘problem’ – at what level to set the alarm bell. Furthermore, when faced with the inevitable conflicts between business priorities, improvements in one direction (e.g. production rates) may be associated with deterioration elsewhere (e.g. costs or risks). A ‘balancing’ mechanism is needed for the ‘scorecard’ if we are to be consistent in targeting the most important improvement opportunities.

Unfortunately the increasing ease of such data collection has, in many cases, resulting in more confusion than clarity – data overload rather than more intelligent, targeted discovery and diagnosis of the important issues. Technology certainly can assist, greatly, but there is a big danger of the ‘tail wagging the dog’!

³ See www.twpl.com - follow link to MACRO Navigator

The second category of decision support (evaluating solutions) is an even more complex one – there are many, confusing, methods to help choose between different actions, to evaluate their cost/benefit/risk impact, and to determine when, or how much intervention is appropriate. In some cases there are simple, common-sense aids to encourage greater consistency or more appropriate choices. For more complex trade-off’s or interactions, significant calculations, modelling or “what if?” assessments may be necessary. The following table (figure 1) provides a summary of the main groupings of requirements.

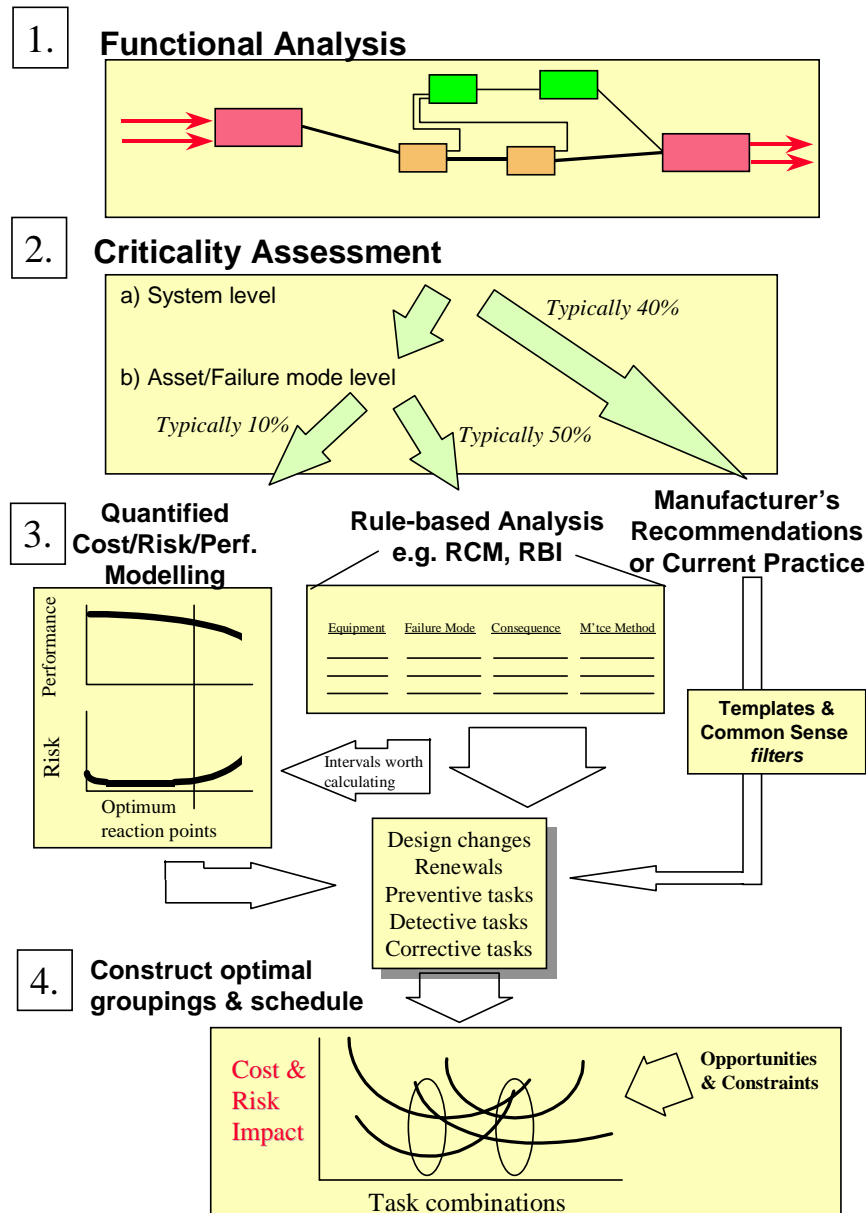
		<i>Increasing complexity of the decision being taken >></i>			
Criticality/size of the decision (appropriate sophistication & cost of method)		Simple Yes/No decisions	Option or scenario choices	Specific task evaluation & timing optimisation	Multiple tasks or systems optimisation
	Customised system/programme Simulation	n/a	5	n/a	5
	Quantified cost/benefit/risk Calculation	3		4	
	Rules, templates & decision trees		2		n/a
	Simple, structured common sense	1			n/a

Figure 1. Main blocks of decision complexity and criticality

Clearly the more complex and critical the decision, the more care and rigour is justified in evaluating options or optimising the appropriate actions. In operational practice, however, there are some natural groupings to the combination of decision type and most suitable technology or decision aid. These ‘best fit’ uses of different methods are illustrated in the numbers cells of Figure 1 above and will be discussed in more detail later. In the meantime, however, and from experience in hundreds of implementations, we can see an overall pattern emerge.

Around 5-10% of assets, equipments, projects and decisions are ‘super-critical’ and justify case-by-case quantified modelling, exploration and analysis. The next 30-50% of cases are too many for such individual and costly consideration, but are sufficiently important to justify an enforced rigour, discipline and cost/benefit/risk evaluation to minimise the errors of subjective judgement. The targeted application of RCM and RBI (to choose which type of risk control method is most appropriate) fit well into this category – they are sufficiently rigorous to achieve high confidence in the results, but they are not sophisticated enough to truly optimise what *combination* of actions and *how much* should be done (which are justifiable extra levels of consideration in the super-critical cases).

The remaining 40%+ of processes, equipment or projects are individually of low importance, but collectively still responsible for large amounts of budget, resource and impact. Case-by-case treatment of these decisions can only be justified if the method is extremely simple, rapid and cheap – so here we find sensible use of templates (sometime derived and ‘de-tuned’ from the higher-criticality cases) and simple procedures or value-for-money filters.



Copyright 1994 The Woodhouse Partnership Ltd

Figure 2. Analysis sophistication should be proportional to criticality

The essential message here is that a ‘mix-and-match’ approach is necessary to the various decisions and tools. Many organisations have tried to do systematic studies with a specific methodology such as RCM, RBI or, more recently, 6-Sigma only to find that, if there has been no selective focus onto areas where the method is most cost-effective, they reach

‘paralysis by analysis’ quite quickly. Each method has its place but the real art is in selective, targeted application!

4 Decisions involve trade-off, and we are not good at it!

In order to dig deeper now, we need to consider the underlying nature of many of the decisions we face. Again I am going to concentrate on the important ones – *what* is worth doing, *when* – rather than the fine-tuning aspects of *how* things should be done. In choosing what to do, there is always a compromise between the costs of the proposed action, and the reasons for doing it (or the consequences of **not** doing it). Sometimes this trade-off is simple – we can make the \$10,000 modification and achieve a 2% performance gain. In others (the more common cases), the compromise is more complex and uncertain – the degree of improvement depends on how much we do, when and what secondary effects are involved, including longer-term consequences. In a previous paper to ERTC I have discussed the trade-off or compromise process, the 5 ways of quantifying the different business drivers, and the true meaning of ‘optimum’ (see figure 3). These disciplines, along with methods for range-estimating, using tacit knowledge, and putting a price on the intangibles of reputation, customer impression and morale etc., all emerged from the European MACRO project.




Figure 3 “Optimum” is the least painful *combination* of the conflicting factors

The degree of sophistication applied to find this optimum again varies with the decision criticality – the consequences of getting it wrong. A minor project and its timing, or a lubrication schedule, might be left to local subjective judgement, but the human brain is particularly bad at weighting the various factors correctly. We tend to distort in favour of the familiar and the tangible – and away from the risks or the lost opportunities. For example, I have been teaching ‘optimisation’ with a specific example for a number of years now – a simple case of furnace/boiler/heat exchanger deterioration and cleaning/shutdown decisions. Even given all the facts and relevant data, over 85% of participants get the wrong answer – and introduce unnecessary costs or losses that are 30%, 50% or even 200% greater than the optimum. This is also born out in operational cases: a recent risk-based review of electrical

protection testing found that intervals *on average* were 4x too frequent (and, for, high criticality installations, the interval was 2x too infrequent)! Major asset renewals or change projects are similarly vulnerable – the information lies in multiple heads and it is very difficult to see the best compromise between early or later investment, cashflow impact, risks avoided, performance gained, sustainability or other capital investment deferment, regulatory compliance etc.

5 The ‘wish list’ for decision-support

Getting these decisions wrong has big impact, but getting them right (and optimal) requires a mix of

1. Structured ways of ensuring the right questions are asked
2. Data mining/interpretation/clarity
3. Quantification aids for the elements that are/cannot be data supported
4. Methods to cope with the inevitable uncertainty
5. Trade-off calculations
6. “What if?” capability
7. Total Business Impact view of the different options.

Two main ‘levels’ of these aids now exist – those which address individual tasks and decisions about them, and those which take an aggregate or whole system view.

5.1 Single task decision aids

The ‘single task’ decision aids are clearly aimed more at the tactical, case-by-case level of application. So, for example, RCM, RBI and 6-Sigma/TQM tools are individual problem-specific, considering each risk or issue and the appropriate preventive, predictive, detective or mitigation action. What they don’t do, or at least don’t do effectively, is to handle the trade-offs and find the right mix of action (costs) and impact (residual risks etc). These tools are essentially ‘bottom-up’ aids, building up a justification for what is worth doing, when and where, based on the individual characteristics that can be accumulated into overall budgets, resources, plans etc. Unfortunately they have had a mixed reception, usually through poorly targeted application, data overload, inappropriate (invalid) usage or insensitive implementation. They also share a significant vulnerability – they tend to consider each risk or problem in isolation. “Weibull analysis” falls into this trap badly, and regularly, with the added weakness that the resulting invalid conclusions still appear perfectly reasonable. There is not enough space in this paper to list all the vulnerabilities of this ‘decision aid’, but the proportion of correct, optimal decisions resulting from such studies is extremely low and will remain so.

Even filling in an FMEA/FMECA table introduces this weakness: each risk is considered, consequences imagined, characteristics described and ‘medicine’ chosen. Then we move on to the next one, and the next, and the next... ignoring any interactions between the lines of our table. The preventive action for one risk might well increase, or change, exposure to one of the others. Indeed it would be surprising if it did not – a lot of what we do had secondary effects. We should be considering the negative, or other secondary, effects of our planned interventions, as well as the positive reasons for them. I have encountered cases where maintenance-induced failures accounted for over 30% of all failures, and new projects or

major plant change certainly introduces a significant commissioning period of instability and unreliability.

The evaluation of what to do and when must, therefore, include consideration of multiple effects (risks, costs, efficiencies, life expectancies etc), and several of these will be very uncertain. The MACRO project was fortunate to have some of the top European reliability engineering, mathematics and experienced economics expertise available – and the various working parties ‘solved’ some of the most complex trade-off relationships involved, including the correct handling of any combination of ‘bath-tub’ curve shapes and components. As a result, the Asset Performance Tools calculators enable “what if?” evaluation of almost any combination of planned action, its timing or interval, and effects on various risks, whole life costs, operational performance etc.

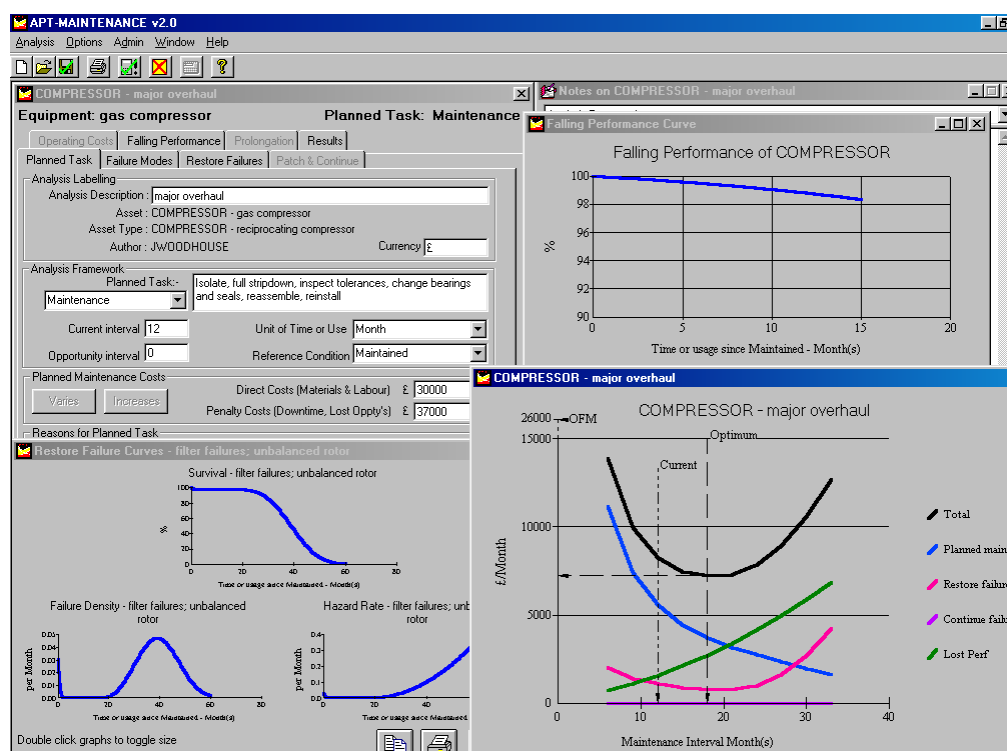


Figure 4. Multiple, interacting risks & performance – converted to total impact \$\$

At the individual task level, therefore, the APT suite is a unique and extremely powerful toolbox of decision-aids. Figure 5 shows the coverage of the 7 decision support modules commissioned by the MACRO consortium. What is particularly gratifying is their sustained impact in a variety of industries and cultures. In the paper presented at ERTC 2003, PDVSA showed a sample of their toolbox approach to operational reliability – and \$15 million of the \$23 million of identified improvements came from the use of APT software and the MACRO methodology. More recently APT tools have been adopted to prioritise large sections of the \$20 billion investment programme in London Underground. BP has been optimising its critical spares strategies with APT. Pirelli has optimised the shutdown strategies, with a 50% net reduction in planned downtime. The list goes on!

MACRO module

Decision Type	Cost/risk/performance evaluation of	PROJECT	LIFESPAN	MAINTENANCE	INSPECTION	SCHEDULE	SPARES	STOCK
Projects, Designs & Modifications								
<i>Cost/benefit analysis</i>								
	Equipment upgrades	X	X	X				
	Process changes	X						
	Procedure changes	X						
	Technology updates	X	X	X				
	Efficiency improvements	X						
	Problem priority/urgency	X						
	Problem-solving efforts	X						
	Investment paybacks	X						
	Compliance requirements	X						
	Public image/morale activities	X						
<i>Life Cycle & Asset Replacement</i>								
	Equipment selection	X	X					
	Vendor comparisons		X				X	X
	Capex/Opex trade-off		X					
	System configuration		X					
	Repair vs Replacement		X	X				
	Life extension projects		X					
Operating & Maintenance Strategy								
<i>Performance/Reliability/Longevity</i>								
	Optimum efficiency profiles		X	X				
	Optimum run lengths between shutdowns			X				
	Reliability, efficiency & longevity combinations		X	X				
<i>Preventive Maintenance</i>								
	Optimum PM intervals			X				
	PM task evaluation			X				
	PM opportunities			X				
	Time vs usage based PM			X				
	Optimum shutdown interval			X		X		
	Repair vs Replace options		X	X				
<i>Predictive/Condition Monitoring</i>								
	Inspection & CM intervals				X			
	CM cost/benefit justification				X			
	CM methods & performance				X			
	Function testing intervals				X			
	Failure finding inspections				X			
	Safety risk exposures				X			
<i>Work Scheduling & Shutdowns</i>								
	Optimum timing and intervals					X		
	Work groupings					X		
	Evaluation of Opportunities			X	X	X		
	Scheduling and task alignment					X		
Spares & Materials								
<i>Insurance/slow moving spares</i>								
	Stock holding levels						X	
	Whole units vs components						X	
	Shared or dedicated						X	
	Supplier A vs Supplier B						X	
	Pooled access contracts						X	
	Supplier held spares						X	
	Spares criticality						X	
	Optimum availability						X	
<i>Consumables, stock, materials</i>								
	Optimum stock levels							X
	Min/Max stock levels							X
	Reorder quantities							X
	Reorder cycles							X
	Supplier A vs Supplier B							X
	Pooled access contracts							X
	JIT/Supplier-held stock							X
	Optimum availability							X
	Storage requirements							X

Figure 5. MACRO methods coverage of individual decisions

5.2 Programme and system-level decision tools

The ‘top-down’ area of decision-making recognises that we often do not have the time or resources to analyse everything from component level upwards. In major projects, or corporate-level strategic direction, choices have to be made with high levels of uncertainty and approximation. Such decisions have long-term implications, deal with large sums of money and yet are particularly susceptible to error (in many cases, assumptions have to be made prior to any opportunity to collect hard evidence). Decision support is particularly valuable in such circumstances!

The levels of assistance split into two very different approaches. At one level, some very simple discipline aids (structured questions, basic cost/benefit calculators and decision procedures) make a big improvement in consistency and the application of value-for-money common sense. The next level up, however, is a significant technology challenge. Modelling of the many elements, uncertainties and interdependencies is extremely hard and risks the introduction of unrealistic assumptions and the ‘black box’ obscuring of embedded errors. There are three basic families of such modelling, and the technical solutions appropriate to each are quite different.

5.2.1 System Performance Modelling

Major projects in the energy sector are often modelled for system impact and configuration “what if?” studies. This acknowledges the near impossibility for correct prediction of system behaviour in the event of significant changes. Such predictions are also almost impossible to calculate mathematically – there are too many elements, variables and interdependencies. So the correct approach is simulation; creating a picture of the total out of components, whose characteristics are summaries (and may be managed as a ‘library’), and then running a time- or event-based trial of what the total picture might look like in performance, reliability, resource consumption or other features. There are a few such tools on the market – such as MAROS from Jardine & Assocs, RAMP from Advantage, Optagon from Advantica and SPAR from Clockwork Solutions.

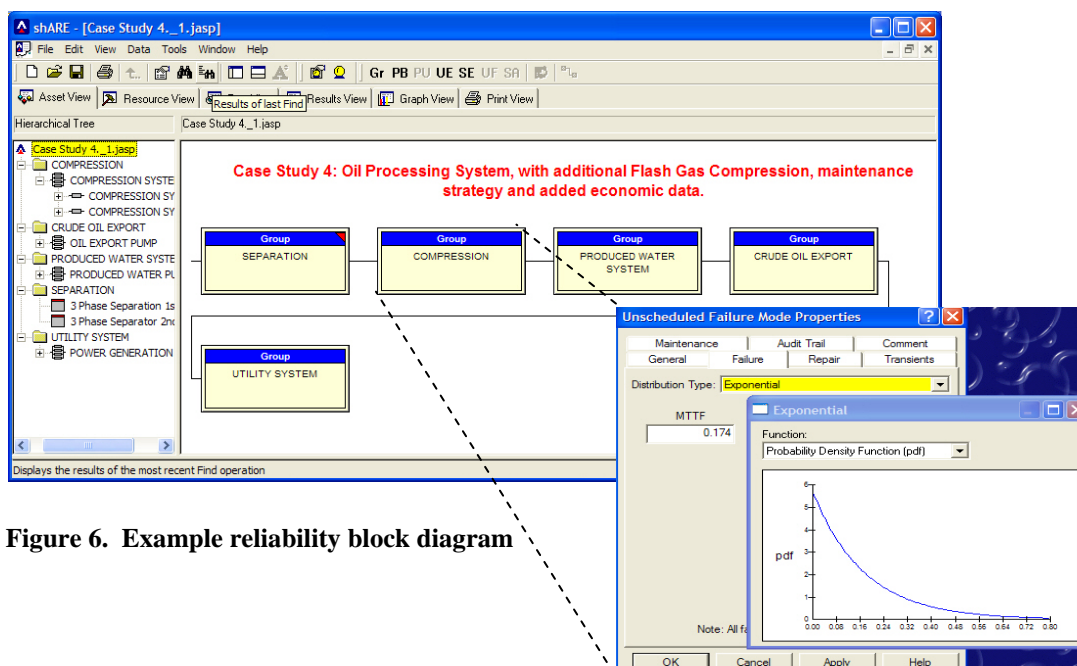


Figure 6. Example reliability block diagram




Figure 7. Example results from simulation

5.2.2 Life Cycle Costing/Analysis

The determination of capital investments based on whole life cycle characteristics is a fast evolving discipline. Early, basic approaches merely added together capital and operating costs over some pre-chosen ‘useful’ life, and discounted these back into today’s money value (NPV). This has several weaknesses – for example, it often ignores performance attributes such as reliability and availability, process efficiency, safety and other characteristics. And it does not easily allow comparison of options with different life expectancies. MACRO addressed these gaps and found that a much better approach is available (using Equivalent Annual Cost or EAC, and incorporating estimates of risk and performance patterns etc.). This also allows us to treat the life cycle itself as a variable: APT-LIFESPAN actually calculates the optimal life cycle for each option, and enables direct comparison between total life cycle costs (of whatever ‘lives’ are optimal). This capability is, to this day, unique and enables very rapid decisions to be made with much higher confidence, auditability and robustness.

Decisions assisted by APT-LIFESPAN:

- | | |
|-------------------|--|
| Design phase | – which item to purchase/install? |
| Operating phase | – shall I refurbish, or replace this failed unit? |
| End of life phase | – when shall I replace? |
| | – shall I replace like-for-like, or pay a premium for upgrade? |




Figure 8. APT-LIFESPAN evaluation of optimal renewal timing

5.2.3 Work Programme & Resources coordination/optimisation

The third circumstance of ‘whole system’ decision-making refers to the integration of multiple activities, assets and timescales into a ‘best blend’ of budgets and resourcing with opportunities or constraints. Until recently this was thought to be too complex a ‘combinatorial’ challenge for computing technology to assist in any practical way. Just like the school lessons scheduling of topics, students and teachers, the number of permutations to be explored to find the best fit is enormous. Just 10 activities, spread over a 6-18 month planning horizon generate 6.7×10^{29} possible permutations. Even with modern computers, this represents some 4 weeks of processor time for a typical Monte Carlo simulation – for each “what if?”! Fortunately a piece of elegant mathematics can help, however. Genetic algorithms (or, in this case, something called ‘simulated annealing’) is a form of self-learning simulation that explores the best schedules much more rapidly (in seconds or minutes). As the name implies, it operates on the “survival of the fittest” principle – trial and errors lead to good and bad options (as defined by total business impact); the good ones are remembered (“survive”) and are then ‘mutated’ further to try more combinations.

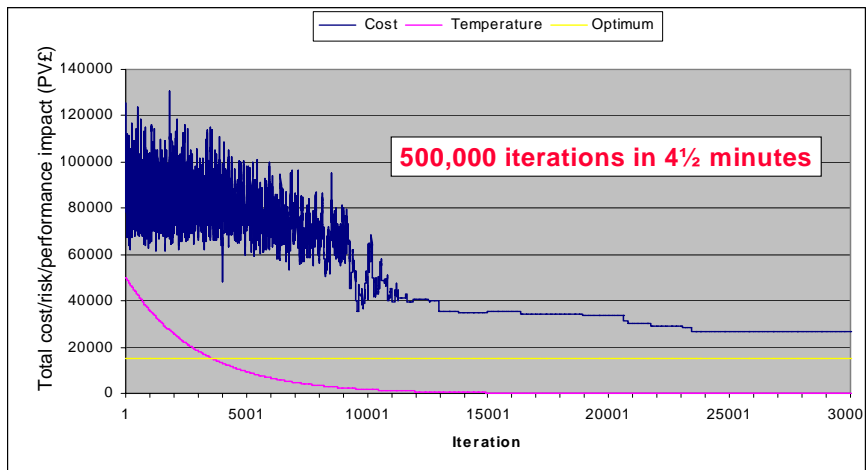


Figure 9. Cumulative learning to find least cost/risk programme of work/shutdowns/resource

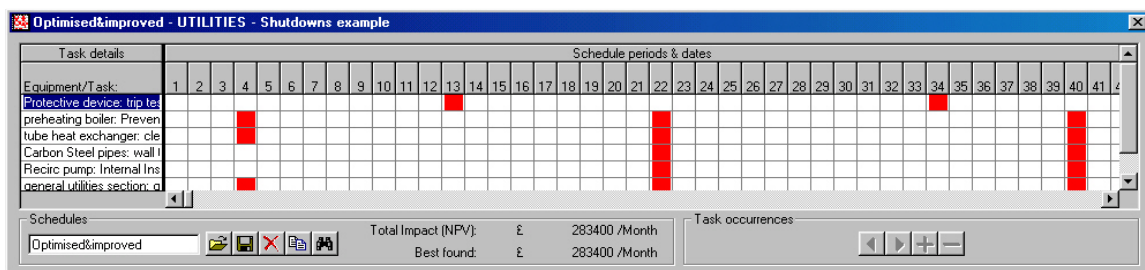


Figure 10. Resulting best compromise schedule of tasks and resources (APT-SCHEDULE)

6 Structuring the toolbox

So, to summarise the decision-support confusions into some sort of order, it is helpful to see the subject as a 3-layered toolbox. The layers correspond to:

- the clarification of problems and asset characteristics (decision support in that such information helps to determine where the problems are and how big they are)
- evaluation of a *particular solution or task* (such as design modification, maintenance, condition assessment or renewal)
- blending* of component solutions into an overall asset management plan (i. over whole life, which includes compromise between in-life activities and replacement options; ii in system interactions, where task clustering, resource constraints etc modify the plan).

Main sections of a decision-support toolbox

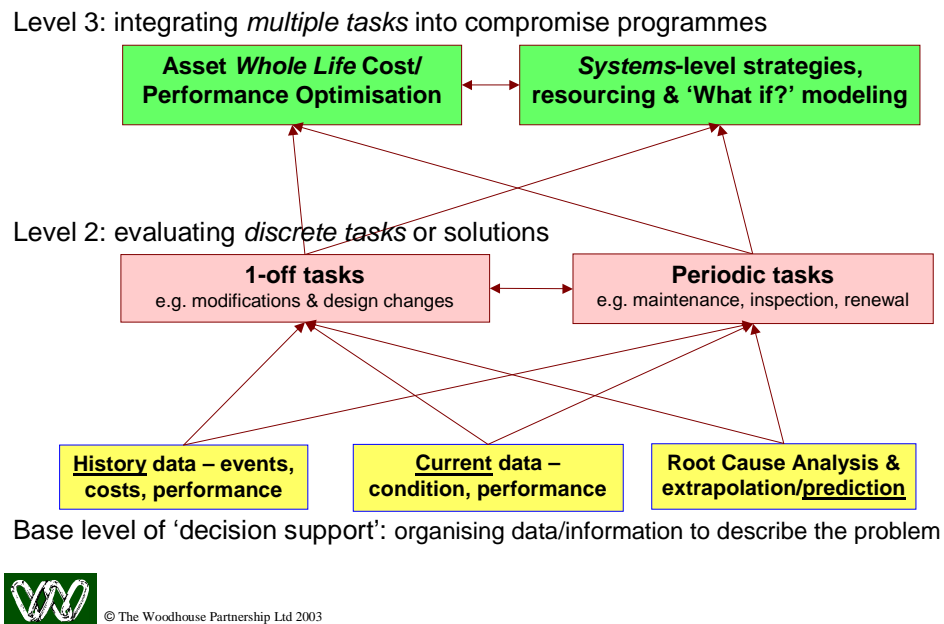


Figure 11. The layers and compartments in an Asset Managers decision toolbox

7 Human & organisational factors

Improved decision-making, however, is not just the use of clever or consistent tools. A vital aspect is the human beings involved - their education and understanding, communication abilities, personal and departmental goals/performance criteria, risk appetites and motivations. So 'decision support' needs to include lots of attention to these critical human factors. Technology can make things possible, but it is people that will make them happen.

Education – there is a very large gap in the required competencies of asset management decision-making, at all levels. Greater business skills for engineers are needed, along with risk and reliability, life cycle costing and cost/benefit awareness and communication skills (technical staff and accountants speak different languages)! These are areas currently being addressed by the UK Institute of Asset Management⁴; a competencies requirement framework is being developed (which will lead, in due course, to an accreditation scheme) and a survey of relevant education or training providers is also being undertaken. [For details, contact the author]

Cross-functional collaboration: departmental, functional or geographical barriers restrict compromise and shared solutions. Most Asset Management decisions involve multiple (conflicting) interests, and depend on multi-disciplined inputs or knowledge. Cross-

⁴ See www.iam-uk.org

functional teamworking, communications and collaboration are essential to finding the optimum compromise.

Short-termism is a commonly distorting factor, due to management turnover, annual accounting cycles, political or regulatory reviews. We certainly need to put a price on sustainability and long term consequences (and tools can help in this), but we also need to change some of the performance measures and accountabilities. For example, the recognition of capital projects as ‘good’ if they come in ‘on time’ and ‘on budget’, irrespective of subsequent performance, operating costs or longevity.

Conflicting stakeholder expectations – one group can only succeed at the expense of another: even ‘balanced scorecards’ can reinforce such competing priorities as they rarely have a ‘balancing mechanism’ to calibrate achievements in one direction at the expense of another (e.g. improvements in customer satisfaction or safety ‘scores’ involving investments that reduce the operating profit figures).

Risk-based decision credibility: it is all very well having a logical, optimised and auditable basis for a particular decision, but if the justification is risk-based, it requires acceptance of a different concept of ‘proof’. Signing a cheque for \$500,000 to purchase a spare turbine rotor, based on an estimated 1 in 10-20 years probability needing it, is quite different psychology from spending the same amount to reduce the energy bill by \$300,000/year.

Fire-fighting habits can distort the picture in two respects. The commonest problem occurs when the reactive workload is too great to allow ‘time to think’ (so the simplest solution is most attractive in high-stress decision-making). Sometimes, also, we find that ‘performance in a crisis’ is given disproportionate peer and management recognition (some people also prefer the variety, unpredictability and ‘thinking on your feet’ environment of such reactive problem-solving) – so there may be a limited personal incentive to avoid the fires in the first place.

Black box mentality is a familiar problem where computer tools or complex algorithms are involved. It creates two sources of vulnerability; a) to the risk of hidden errors or inappropriate interpretation within the black box and b) to the sceptical mindset that overreacts to this first risk by dismissing anything that comes out from such methods. Fortunately such ‘black box’ methods or tools are largely displaced nowadays by much more transparent and auditable processes. However the sceptics often still need to be reassured.

6. Conclusions

So, where do we stand? Much development has occurred, particularly in the IT area and multi-flavoured 'methodologies' (RCM, TPM etc). Computer systems have certainly wheedled their way into the foreground, and 'asset information', 'work management' and 'condition monitoring' systems are generally recognised as necessary and valuable.

The front-line areas of innovation are those of condition monitoring, life cycle and reliability/maintenance strategy analysis. In these fields, the techniques, tools and understanding are moving fast – in fact the technology is no longer the limiting factor. Simulation, cost/risk optimisation tools and sophisticated reliability modelling aids can handle almost any level of sophistication likely to be needed. It is now the understanding and the use of such techniques that are the limiting factors. The education gap is large and, if anything, growing as experienced engineers and managers come to retirement and we continue to 'outsource' and lose key knowledge from the companies. To meet this need, and the obvious mismatch between traditional engineering courses and the modern business requirement, the first signs of hope are emerging – the IAM Competencies Project is attacking the subject with energy.

Nevertheless, we need to increase the spread of understanding, of successes, failures and innovations at a greater rate. The business demands can only get greater, so all of us are under increasing pressure to improve professionalism, discipline and cost/benefit accountability. We cannot afford to reinvent the wheels individually or learn by trial and error - it takes too long and is too expensive. Decision support tools make a very big difference – but only when they are used and implemented correctly!

J.Woodhouse

e-mail: john.woodhouse@twpl.co.uk