

What shutdowns, why and when?

John Woodhouse
Managing Director, The Woodhouse Partnership Ltd
john.woodhouse@twpl.co.uk

1 Introduction

Major cost items come in three flavours – capital investments, catastrophic failures and deliberate plant shutdowns. The first is subject to intense scrutiny/justification and, with the exception of the patchy adoption of Life Cycle Costing, is pretty much understood. We are trying desperately to avoid the catastrophic events – and have been grappling with systematic and quantitative analysis methods for many years (HAZOP, QRA techniques, Risk-Based Inspection etc). The third area, that of planned shutdowns, is still an enigma for many organisations. Much effort has gone into the efficient planning and delivery of the work involved, but relatively little guidance exists for determining what work is worth doing in the first place, and how this should be clustered into appropriate packages to share shutdown opportunities. A surprising number of organisations (particularly in the utilities and service areas of operation) still do not even know how much a shutdown costs them.

This paper describes some recent advances in quantitative evaluation of shutdown programmes. It looks at the bundling of tasks – the logistics of delaying some activities to coincide with others, and the compromise economics of shared downtime costs versus the performance and risk impact of premature or deferred work. The paper is illustrated by three case studies, taken from different industries during the last couple of years.

2 Origins of the new approach

This methodology has been developed by the European MACRO project, a recently-completed 5 year collaboration programme sponsored by the UK government, Halliburton Brown & Root, Yorkshire Electricity, The National Grid Company and The Woodhouse Partnership. The c.20 members of the project have included petrochemical, transport, utilities, manufacturing, process companies and the relevant professional bodies. MACRO has yielded a suite of methods for cost/risk/performance trade-off decisions – such as optimal maintenance or inspection intervals, equipment renewal or upgrade justification, shutdown strategy, spares requirements etc. In each of these areas, a blend of innovative, risk-based evaluation techniques has been developed alongside structured guidance ‘rules’. These have been developed and proven in field by those faced with the decisions (i.e. not some academic theoreticians!).

3 What work is needed, why?

The first step is the systematic determination of the tasks that might warrant a shutdown in the first place. Here the methodology splits a 'greenfield' from 'brownfield' environment. If there is an existing regime of shutdowns, inspection cycles etc, it is somewhat wasteful to re-build the task list from scratch. However, even in such cases, a 'zero-based' maintenance programme (FMECA and RBI/RCM combinations) can be a good stimulus to challenge existing habits and pre-conceptions.

3.1 Reasons for tasks

The FMECA stage is fairly well evolved – albeit with some variations depending upon the existence or not of local historical data. One minor advance in this area, emerging from the MACRO programme, is the observation that, for greenfield projects (with no operational experience), it is often easier to populate the list of potential degradation and failure modes in reverse – i.e. by mapping intended functions first, then listing functional failure consequences and finally brainstorming the failure modes that could result in such effects. Where maintenance history exists, on the other hand, known failure modes comprise the 'seed' information, from which to extrapolate and consider other potential (not yet observed) modes. Generic libraries or templates can also act as such seed material, provided that local conditions and *potential* failure modes are also considered.

The criticality (the C in FMECA) assignment to failure modes is a subject in its own right – a source of confusion or clarity, depending on where you stand. It is certainly needed, and in shutdown studies, we have found that the main decisions are determined by just 5-10 dominant failure modes and the tasks designed to address them. Identifying these critical items is not easy, however. The API Recommended Practice (580/581) on Risk Based Inspection is predominantly a criticality assessment and prioritising of failure risks. Structured risk-ranking workshops, involving operators, engineers and maintainers, offer lesser rigour but are, in many cases, just as effective in identifying the key drivers, often at a fraction of the cost.

3.2 Types of task

RCM is the most widely accepted set of rules for relating individual threats (failure modes) to the best preventive, predictive, corrective or detective tasks. The method is particularly suited to complex plant with many different types of failure mode. Static equipment holds less variety – most maintenance is condition-based and the predominant concerns are "what inspection method, and how often?". API RP580/581 were developed specifically to provide such guidance. Both RCM and RBI can be exhaustive (and exhausting!) but various criticality-streamlined versions have emerged to focus on the bits that matter most.

Whatever the identification method, individual tasks fall into two groups for our purposes – cyclic activities (such as preventive maintenance, inspections and periodic replacements) and 1-off tasks, such as modifications, capacity upgrades or other changes. The 1-off tasks are generally subject to the same evaluation and justification as other projects or capital investments, and their timing is a matter of cashflow/payback/

NPV/IRR calculations. The disadvantages of delay represent continued levels of risk, inefficiency or constrained performance, diluted to some degree by the advantage of deferring major expenditure.

Cycle tasks, on the other hand, are much more complex to evaluate and optimise. They exist because of (actual or potential) *deterioration* and risks or performance that *changes with time*. This topic is covered extensively in the relevant MACRO modules – how to build a model of the cost/risk/performance trade-off and determine the optimal interval, the impact of premature or delayed work, and the sensitivities to any key data assumptions. In summary this involves:

1. Structured, quantified description of the degradation process, using range estimates wherever hard data is not available. This description is built around five distinct families of quantification techniques:
 - ?? Reliability & risk (failure modes, probability patterns and consequences)
 - ?? Operational efficiency (energy, consumables, output volumes and quality)
 - ?? Lifespan effects (life extension, capital deferment etc)
 - ?? Regulatory compliance (safety, environmental)
 - ?? ‘Shine’ factors (public and customer impression, employee morale etc)
2. Cost/risk/performance calculations for alternative intervals – putting numbers to the familiar trade-off curves below.
3. Sensitivity testing to the extremes of possible data uncertainty (often variations by factors of 10 or more for the speculative elements)
4. Identification of key decision ‘drivers’ (which assumptions have the greatest effect upon the optimal decision).
5. *If justified*, more detailed investigation of these key assumptions to determine the correct strategy (in most cases, range estimates are enough to identify the optimal interval, and only when the ‘cost of uncertainty’ is high will the additional research be justified).

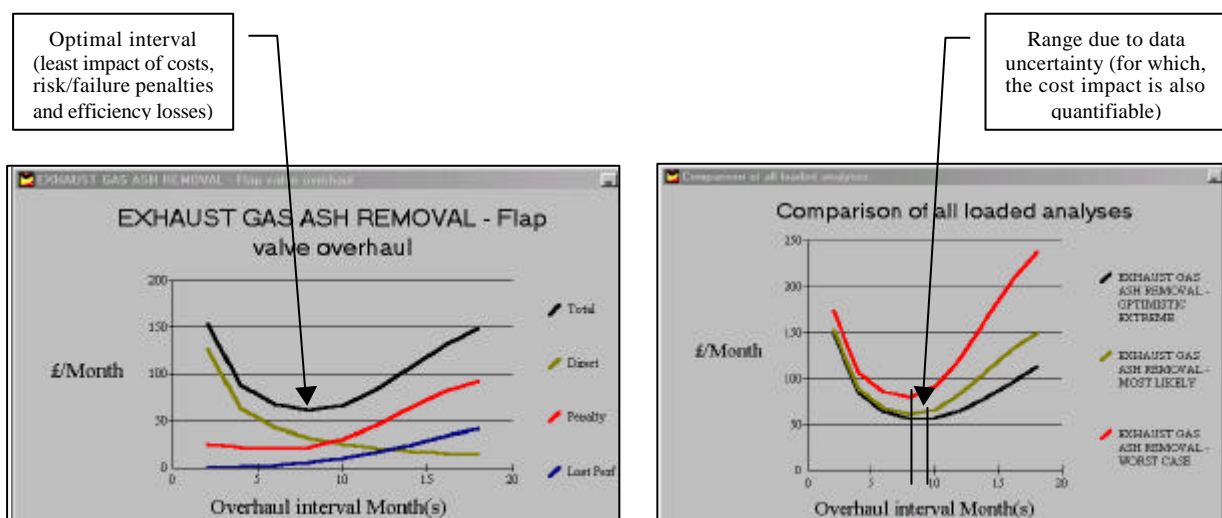


Figure 1 Optimal maintenance interval & sensitivity to data uncertainty

The trade-off calculations vary with the components involved – in many cases there are several interacting failure modes, efficiency profiles and effects upon life expectancy all in the same evaluation. For example, an overhaul of a heat exchanger will consider tube leaks and blockages, performance effects of fouling and cumulative damage to the bundles due to cleaning. The analysis results reveal which factors drive the maintenance strategy, and how that strategy varies with equipment usage, operational criticality, fouling rates etc etc.

In the case of inspection intervals, there is a further split in the modelling methods required. The *predictive/condition monitoring* inspections dominate in major process industry shutdowns, to identify and track vessel and pipework corrosion or cracking. *Functional testing* or *detective* inspections, on the other hand, are those designed to reveal existing ‘hidden’ failures – typical of protective or standby equipment. The MACRO procedures for quantifying and evaluating these two families of tasks differ in the questions that need to be asked, but then calculate the same cost/risk trade-offs for various task intervals.

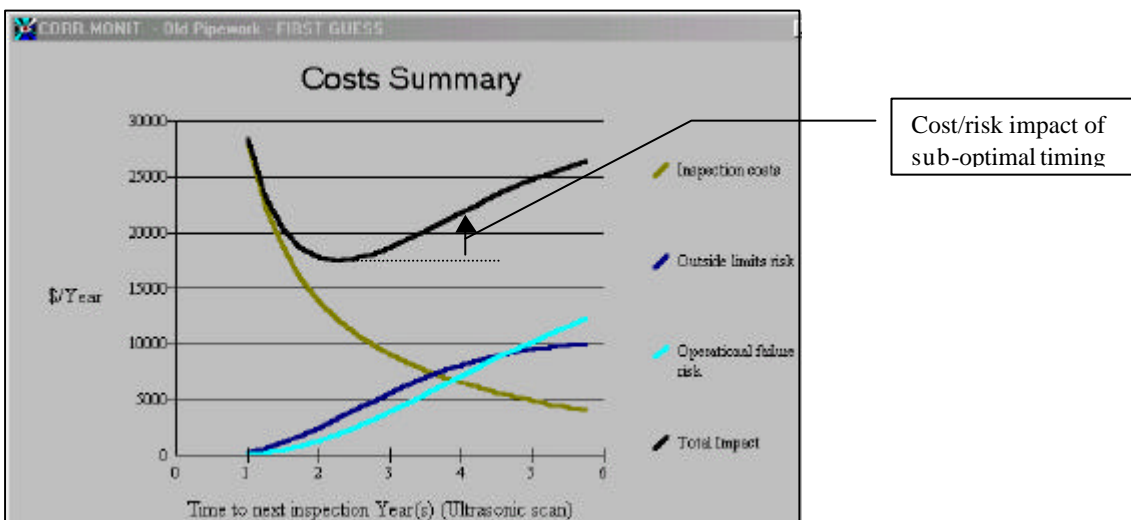


Figure 2 Optimising a condition monitoring programme

3.3 Combining tasks – compromise decisions

The shutdown strategy is a compromise. Some tasks will be performed ahead of their ideal timing, others will be delayed to share the downtime opportunity. The risks and performance impact of delayed tasks, and the additional costs of deliberate ‘over maintenance’ in others, both contribute to the price paid for a particular shutdown packaging. The degree of advantage, on the other hand, is controlled by the costs that can be shared as a result. The downtime impact (lost opportunity costs) often dominates such sharing advantage, but the direct costs (planning, facilities, labour etc) of shutting down and starting up again must also be considered. The critical path of component tasks will determine the bundle’s total downtime impact – and this will vary with the degree of sequential or parallel working that is possible (as well as the discovery of defects that

need corrective work, task overruns etc.). Uncertainty is often high but, like component task justifications, these bundle characteristics can be explored in “what if?” mode to determine if, and which, assumptions make a difference to the final outcome.

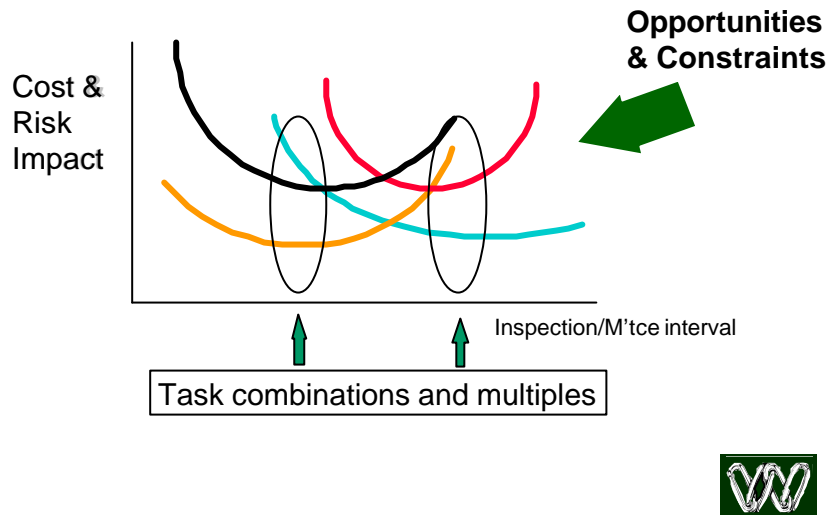


Figure 3. Clustering tasks into appropriate bundles

External constraints exist at both the individual task and shutdown bundle levels. Regulatory requirements determine that some inspections should occur at least 3-yearly, or that a maximum acceptable risk is 10^{-6} for a certain failure mode. This limits the range of allowable intervals for that task. At the bundle level, logistical, safety or resource restrictions might constrain the grouping of certain tasks. Such bottlenecks force a greater cost of compromise: a sub-optimal combination and timing for the work.

Another form of bottleneck is that introduced by the need for a task at short intervals while all other tasks can be performed substantially less often. This introduces the option of nested cycles (the other tasks being performed every 2, 3 or more cycles of the short interval work). It also reveals the scope for design changes to de-bottleneck the requirements – eliminating the frequent shutdowns and extending run lengths. The analysis process itself calculates the net payback for such modifications or de-bottlenecking.

The grouping and re-grouping of tasks, and “what if?” exploration of de-bottlenecking, can be manual (combining tasks in different bundles and moving the bundles to shorter or longer intervals) or semi-automatic. The MACRO R&D work has researched a number of methods for the latter – including Artificial Intelligence techniques such as Neural Networks, Genetic Algorithms and Simulated Annealing. The final combination is still being refined – but the various prototypes have yielded some astonishing results. In short, the scope for re-bundling tasks and timings is much greater than expected, with corresponding substantial impact on costs, performance and risk exposures. The National

Grid Company did some early work¹, using a Genetic Algorithm approach, and revealed scope for *21% improvement in system availability, at the same time as a 23% reduction in total cost/risk impact*. Since then, ICI Eutech has been using the methods to evaluate shutdown intervals for chemical manufacturing plant² (*revealing £2.5 million savings*), and my team have been re-bundling the maintenance and inspection tasks on process plant, railways and water utilities. In one such case, shutdown intervals were extended from 2 years to 4 years, releasing over *£5Million/year* in net improvement.

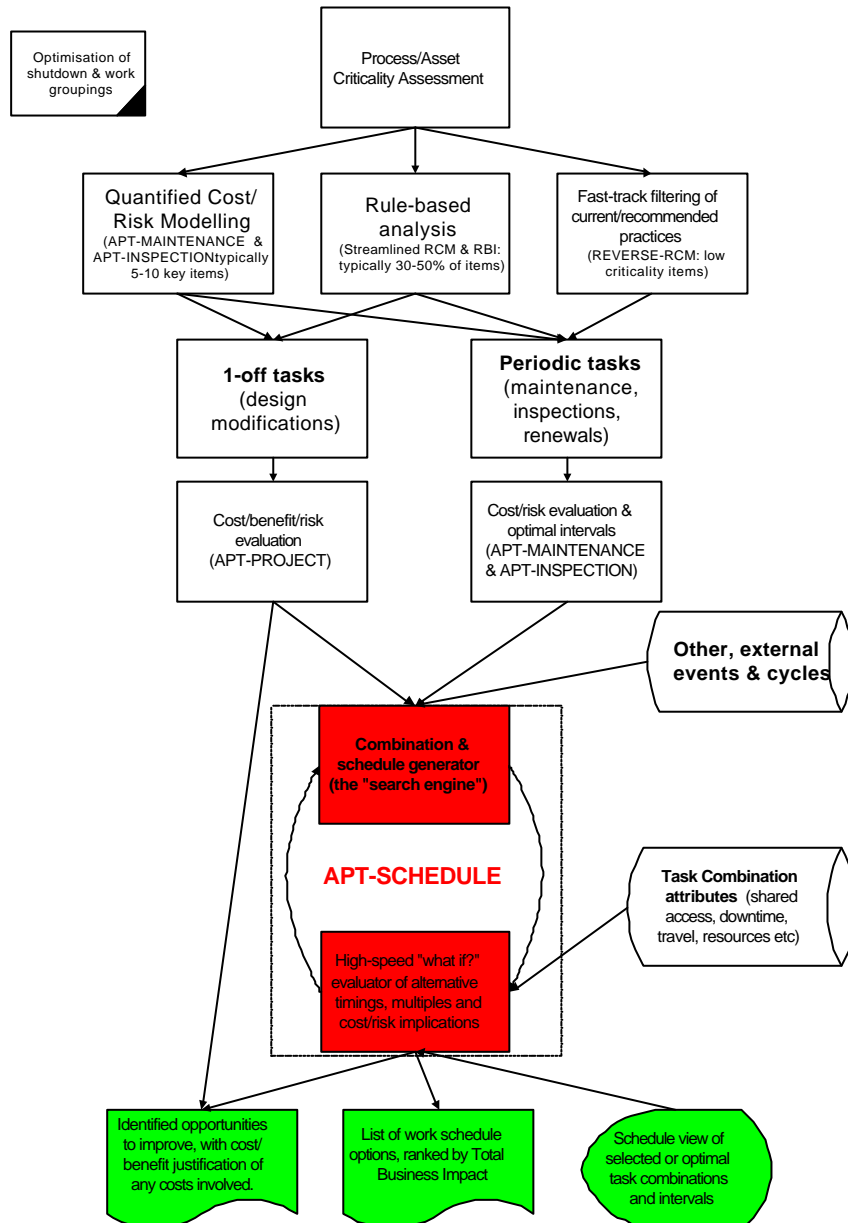


Figure 4 Overall process for optimising shutdown bundles & timing

¹ "Optimisation of the asset maintenance on a transmission circuit", Ursula Bryan & others, MAINTEC 1998, Birmingham

² "What shutdown & when", Simon Smith, MACRO seminar, May 2000, London

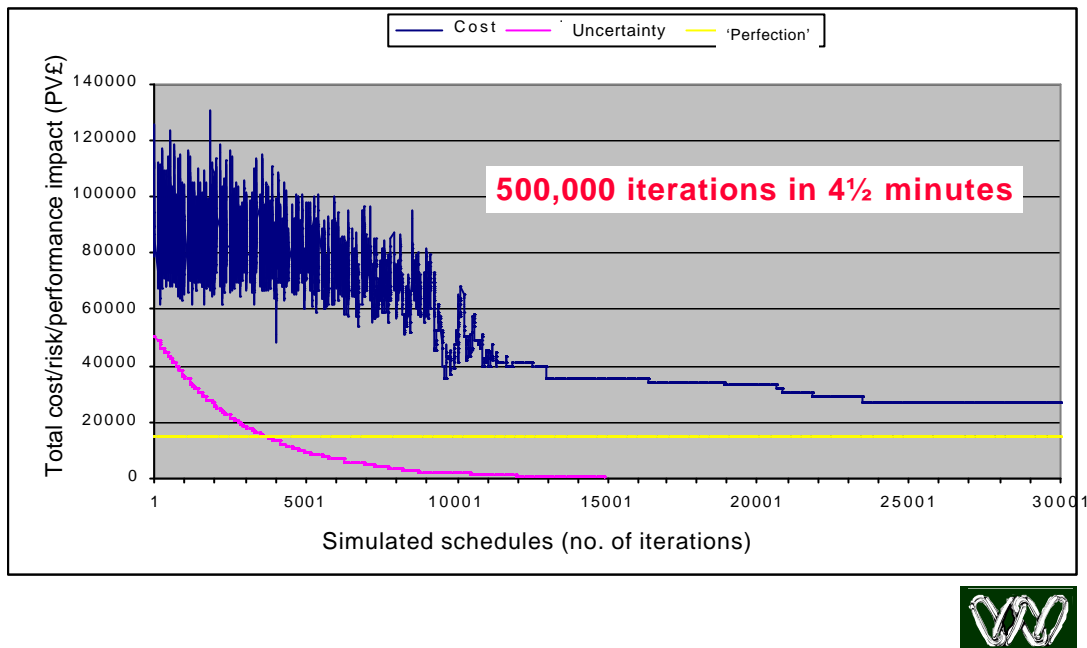


Figure 5. APT-SCHEDULE: Self-adjusting search for the best work combinations ('perfect' represents the theoretical ideal of no downtime and every task at its optimal timing)

3.4 Case Study: Power distribution circuit

Here, the shutdown (or "outages") comprise a variety of tasks on the connected assets of a critical supply route. There might be anything up to 30 or 40 discrete items of equipment in the circuit, and each item (for example the circuit breakers at each end) may have several tasks assigned to it, with optimal intervals that vary from short (6-12 months) to long (some only every 12-15 years). The circuit outage programme is a complex blend of small-and-frequent, and larger-but-rarer tasks, with a vast number of permutations possible. Some tasks are statutory requirements, others can be brought forward or delayed. The cost/risk impact of delay varies greatly with the deterioration rates – some items have critical timing and others have fairly 'flat' curves of total impact.

The analysis process calculated the Net Present Value of all future costs, risks and outage timings and, in this case, the optimal regime involved bringing forward several of the 'next maintenance due' dates to create a better alignment. The subsequent avoidance of multiple outages more than paid for the earlier initial expenditure.

3.5 Case Study: Chemical production unit

In May this year, ICI Eutech presented a paper to the MACRO results seminar on results achieved in studying a bulk chemical manufacturing plant. An existing 2-yearly shutdown typically involved £320,000 of work and 21 days of downtime. The criticality analysis revealed which units were the main drivers for the shutdown – the HCl stripping

column, the reactor unit manway lining, some sacrificial iron packing in a column and some of the smaller piping. It was noticeable that these items were NOT the biggest, most expensive items to inspect or maintain, but were deterioration rate limiting – the component tasks necessary to inspect or maintain them had the shortest intervals.

The elimination of some of these run-length constraints (bottlenecks) involved, for example, using high performance alloys (Monel) to achieve longer life. The payback for such additional periodic cost was revealed to be measurable in months. The study overall revealed that a 4-year shutdown could be achieved, with NPV savings of over £2.5 Million available.

3.6 Case Study: Conversion reactor and condenser

Last year we were involved in a similar study – looking at the possible extension of run lengths for a specialised reactor/condenser process. The initial criticality assessment took 3 days, using a combination of structured interview techniques and survey of existing FMEA and QRA studies. This revealed a potential ‘decision driver’ list of about 30 items, each with a number of inspection and/or maintenance tasks required. In addition, there were a few 1-off tasks that were accumulating – technology upgrades and mandatory modifications that needed to be scheduled into the programme. The following items were identified as the most influential in the shutdown decisions:

- ?? Reactor vessel: internal support beams, injector nozzles, shell integrity
- ?? Quench tower: internal supports, nozzles, shell integrity, relief valves
- ?? Re-circulation pumping system: gate valves, seals, cooler
- ?? Product chiller: cleaning cycle, bypass unit
- ?? Separator unit: relief valves

Working from the shortest cyclic tasks outwards, we created individual cost/risk/performance models by interviewing operations, maintenance and engineering staff, recording their experience, opinions and extrapolations (how the equipment would behave if we extended the intervals). The resulting range-estimates were explored for all sensitivities, so that the recommendations included the future data requirements for further refining the strategy. Over 75 optimisation studies of component tasks were performed to create the necessary raw material for the shutdown optimisation. This took 3 weeks for a team of 2 persons full time and 2 part time.

The component task studies themselves revealed the scope for substantial cost/risk/performance improvement. Around £1million/year savings were identified to be available from a number of minor changes in workscope, in timing or design/operations changes. These included, among several other recommendations;

- ?? Upgrading materials for the reactor support beams
- ?? Changing the cleaning process for the product chiller

- ?? Installing dual pilots on the relief valves (allowing on-line maintenance)
- ?? Stainless steel lagging of injector nozzles

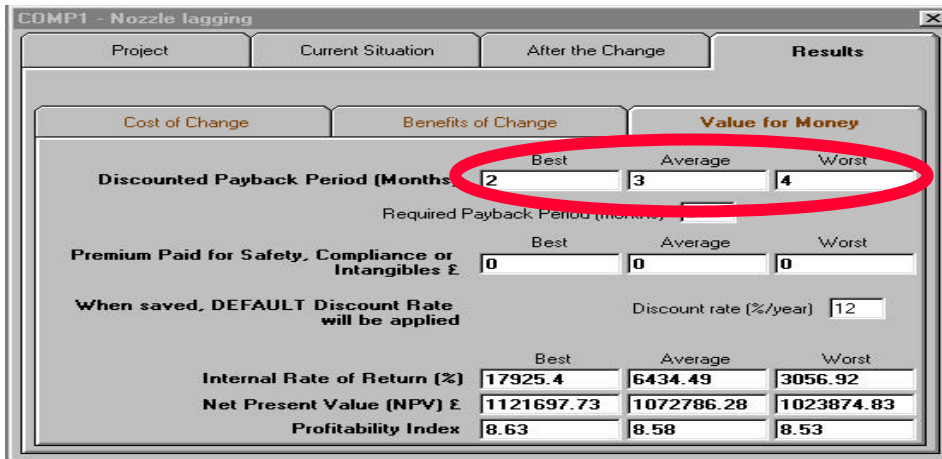


Figure 6 Evaluation of nozzle lagging to extend life

The big prize, however, was the extended interval between major shutdowns. The design changes and bottlenecking allowed a doubling of the shutdown interval, with net total impact worth a further £4million/year across the 6 units. This figure comprises the net effect of increased availability, reduced maintenance costs, all changes to risk exposures, performance impact and even projected changes to equipment replacement requirements. It is the conservative sum of the ‘pessimistic’ projections, so we can be confident that a) the real benefits are substantially higher than this and b) the proposed strategy is appropriate even in the extreme case of projected risk assumptions.

Equipment/Activity	Shutdown intervals		Change (£/unit/yr)
	24 Monthly	48 Monthly	
Turnaround downtime	1,512,000	756,000	-756,000
Recirc pump valves	103,800	152,400	48,600
Separator relief valves (dual pilot)	18,144	36,048	17,904
Quench nozzles/shell (sprayclad)	62,400	77,000	14,600
Overall train reliability	64,560	70,320	5,760
Quench relief valves (dual pilot)	5,040	5,616	576
Reactor supports (stainless)	79	315	236
Reactor shell welding	72	120	48
Reactor nozzles (lagged)	1,060	882	-178
Total cost/risk impact (£/yr)	1,767,155	1,098,701	-668,454

Figure 7. Top influences on a unit shutdown interval, ranked by their influence.

4 Conclusions

These studies are fairly typical – a combination of some hard facts, a lot of range-estimated speculation, a long list of potential influences but relatively few that really matter, and complex interactions between failure modes, deterioration assumptions, design options and maintenance tasks. It has confirmed, however, that structured approach, combined with modern “what if?” optimisation tools, hold substantial scope for increased performance and cost/risk improvement.

J.Woodhouse, 21/7/00