

Leveraging Total Ownership Cost Analysis to Optimize In-Service R&M Programs

John Millhouse, TFD Group

Key Words: PBL, R&M improvement candidates, optimization, TLCSM, TOC

SUMMARY & CONCLUSIONS

There will never be enough money for all the Reliability and Maintainability (R&M) investments we can identify as potentially worthwhile. However, innovative processing of the output of sufficiently detailed Total Ownership Cost analyses makes it possible to identify the most promising improvement opportunities, rank them in order of economic viability, and build competitive business cases for implementation. A powerful and highly intuitive software prototype has been deployed to this end into several weapon system programs. While the extent to which it achieves its basic aim of optimizing resource investments will take some years to establish, sufficient interest has been generated to justify significant refinement effort.

1 INTRODUCTION

Complex hardware systems in the modern world are typically subject to spiral development processes, whereby new technology is “inserted” to various ends: improving safety or operational performance, enhancing availability, prolonging service life, or reducing operating costs as the case may be. Determining the content of the spirals is a problem of classical economics – resource allocation under conditions of scarcity: there will never be enough money for all the things we can identify as potentially worthwhile.

Increasingly, system manufacturers, program managers and sustainment contractors are employing highly detailed event-based Total Ownership Cost (TOC) models, whether to address affordability issues at program inception or respond sensibly to budgetary pressures in service. A by-product of sufficiently detailed TOC analysis is information about:

- the expected frequencies of component failures (in all applicable modes) under the system

deployment and operating scenarios assumed to apply over the remaining service life, and

- the costs of downtime and remediation associated with each failure.

The work described in this paper was undertaken to establish whether these outputs of TOC analysis, suitably processed, could serve as the basis of an effective optimization scheme in respect of R&M investment proposals. The approach was to devise and demonstrate a software tool (complementary to TOC models) capable of identifying improvement opportunities and ranking them, through sound business case analysis, in order of their economic significance and relative merit.

2 INCENTIVE FOR R&M IMPROVEMENT

A high proportion of the hardware systems in common experience, household appliances for example, perform their functions in trouble-free fashion for many years with only modest support. Market share is a function of distribution channels, productive capacity, brand recognition, price, and perceived quality. Manufacturers have a clear incentive to engage in disciplined R&M activity (i.e. not so much that quality enhancements are outweighed by cost impacts) because the marketplace rewards success directly and continuously. Since, to a reasonable approximation, the TOC is the price, the purchasers of such systems know nothing of TOC analysis.

At the other end of the complexity spectrum, where we find weapon platforms, commercial airliners, industrial plants, and the like, bristling with complex subsystems of varying maturity, the incentive is significantly less direct. Here the operation of the marketplace is intermittent and the protagonists oligarchic. Systems are “sold off the plan”, and assertions about operational performance, affordability, reliability, availability, testability, and so forth can remain largely unsubstantiated for years.

The all-important plausibility of such assertions rests on the extent to which a would-be prime contractor projects an aura of competence in the disciplines involved. Since the largest TOC component is sustainment over ever-lengthening life cycles, one of these disciplines is TOC analysis. Another is R&M.

That an acquisition environment of this general character often saddles operational communities with somewhat unreliable and unmaintainable systems should come as no surprise; nor should the seeming paradox that the major primes and lower-tier OEMs whose labors produce these systems continue to prosper notwithstanding. For the plain truth is that they do a highly commendable job under the circumstances.

Making a virtue of necessity, military establishments across the western world have re-ordered their affairs in recognition that the weapon systems they acquire will:

- seldom reach their full potential until years have passed, and
- require periodic modernization (a.k.a. spiral development) in any event to remain operationally competitive with systems employing newer technologies.

Accordingly, the Total Life Cycle System Management (TLCSM) paradigm has supplanted the separate acquisition and in-service support fiefdoms of yesteryear, and CALS has stood for Continuous Acquisition and Life-Cycle Support long enough for most people acquainted with the original meaning to have retired.

The TLCSM charter is to deliver required levels of operational capability over time while keeping TOC in check. In practice, a TLCSM organization, whether US SPO or UK IPT, mobilizes the necessary resources by pulling various strings that mediate the flow of inducements to operational and support communities. The capability outcome is determined by the choices made: which strings are pulled and how hard.

2.1 Operational Concerns

The first imperative is to get the money. Unsurprisingly, how much of it might be available at a given point in time is a function of extant levels of dissatisfaction in various facets of operational

performance. Securing the funds needed to maintain a competitive edge in relation to weapon system qualities such as speed, range, stealth, survivability, interoperability, and so forth, calls for no special advocacy – these things speak for themselves. However, what might be seen as *quantitative* improvement is more problematic. Over recent times the quantitative dimension has come to be characterized by four figures of merit [1]:

- Operational Availability (A_o). Functional qualities aside, it is usually advantageous to have more (up to a point) rather than fewer systems available to be sent out on operational missions. And since the number of systems procured in the first place imposes an upper bound, the quantity of interest is the mission-capable proportion.
- Mission Reliability. Likewise, the higher the proportion of missions completing successfully the better.
- Logistics Footprint. The smaller the quantities of support resources (manpower, spare parts, support equipment, etc) that have to be deployed in order to sustain expeditionary operations, the more quickly forces can be put in place and the more likely it is that there will be enough transportation capacity.
- Logistics Response Time. Shortages of some of the items required to restore downed systems to mission-capable condition are inevitable. It is important to planning effectiveness to achieve a high probability that the support system will make good any such shortage within a stipulated time (thus achieving so-called “time definite delivery”, a basic tenet of net-centric warfare).

That warfighters value these measures highly is not in doubt. Once persuaded that some meaningful increment of performance can indeed be bought, the operational community can be depended upon to support appeals to the keepers of the public purse for the necessary funds. The burning question is how best to persuade them.

R&M initiatives plainly have the potential to affect performance outcomes in the following beneficial ways:

- other things being equal, if components can be made more reliable, and hence fail less frequently, A_o and Mission Reliability go up while

Logistics Footprint and Response Time go down; and

- if components can be made more maintainable, and hence take less time to fix, A_0 goes up while Logistics Footprint and Response Time go down.

Complicating matters, however, it must not be overlooked that these desirable effects can be produced in other ways:

- increasing spare parts buffers causes A_0 to go up and Response Time to go down (albeit enlarging Footprint);
- widening preventive maintenance and/or deepening corrective maintenance can mimic the ability of reliability enhancements to reduce failure frequencies, bringing equivalent benefits; and
- reducing the transportation and administrative components of the overall repair cycle competes in like fashion with maintainability initiatives.

The upshot is that R&M is locked in competition with maintenance and supply interests for shares of whatever funding will be made available for system performance enhancement. While operational concerns dictate that *some* things will certainly be done, just what they will be is a matter to be fought out on an economic battleground.

2.2 Economic Considerations

“Quality doesn’t cost, it pays” was a memorable slogan of the Total Quality Management (TQM) movement. There is an echo of this in the TLCSM world. The prevalence of “Spend to save” opportunities, especially in the early part of a weapon system life cycle, provides the essential rationale for wide-ranging US DoD Reduction of Total Ownership Cost (R-TOC) efforts [2].

The criterion for including a given proposal in a list of “spend to save” opportunities is a benefit stream whose present value (PV) can be shown to exceed the present cost (PC) of investments required to secure it. Benefits are restricted to improvements in the figures of merit described above and downstream cost savings. Investments include RDT&E, hardware production / procurement (including spares), system modification, support equipment, training, and so forth, suitably distributed over time.

Of course, before you can estimate the cost of an improvement proposal and evaluate the benefits likely to flow from its implementation you must first formulate it. This implies abilities to:

- single out the components promising the greatest potential for performance improvement and/or cost avoidance; and
- dream up schemes for garnering the promised benefits.

Now, having seen above that competence in R&M is a marketplace entry condition, we can safely assume that all companies likely to need it possess scheme-dreaming ability in sufficient abundance. Moreover, since the engineering talent at the heart of this competence is quite reasonably regarded as an overhead resource, it is easy to slip into the naïve notion that the dreaming is free. Opportunity costs are notoriously underrated, notwithstanding that managers readily accept (when it is pointed out) that every man-day spent up a gum tree (Australian for false trail) is one subtracted from profitable enterprise.

2.3 Decision Analysis

In this light, even the seemingly trifling activity of panning for problems worth solving can be seen to warrant a disciplined approach. The formulation of a measure known to exponents of Decision Analysis as the Expected Value of Perfect Information offers useful insight into what we will discover to be the essential feature of a suitable formal method.

Suppose the engineering department of the contractor performing support integration (i.e. the Prime Support Integrator, or PSI) for a certain complex weapon platform boasts a track record of having reduced the future TOC contributions (C) of components subjected to R&M improvement projects by a factor r on average (i.e. $C_1 = C_0/r$). Suppose also that the proportion of components found to offer significant R-TOC potential has historically been p . What is the most the PSI should be prepared to pay for confirmation that R&M improvement effort applied to a particular component will be rewarded?

The Expected Value (EV) of the improvement potential in a randomly selected component is the product of the reduction in TOC consequent on successful implementation of an improvement

scheme and the probability that a scheme will be found at all. That is:

$$EV = (C - C/r)p \tag{1}$$

If, by virtue of the engineering department expending some modest investigative effort on triage of opportunities on offer, the PSI could eliminate uncertainty concerning the eligibility of particular components for attention, it would be appropriate to speak in terms of the Expected Value given Perfect Information ($EV|PI$):

$$EV|PI = (C - C/r) \tag{2}$$

The Expected Value of Perfect Information ($EVPI$) is simply the difference between these two measures:

$$EVPI = EV|PI - EV = C(1 - 1/r)(1 - p) \tag{3}$$

If it costs the PSI less than this to test the hypothesis that a particular component is eligible for in-depth improvement effort, the expenditure is worthwhile. Note that since the track record is nothing to boast of unless $r > 1$, and since p is a probability, $EVPI$ is the product of the component's future TOC contribution and a contextually-determined non-negative constant whose value cannot exceed 1. Accordingly, an appropriate triage policy is to look for improvement opportunities in descending order of components' future TOC contributions. That this makes intuitive sense is evidenced by the ubiquitous practice of focusing management attention on the top 10 cost drivers, however identified.

2.4 Marginal Analysis

Let us therefore imagine an engineering department working through a long, sorted list of weapon system components, in each case attempting to dream up one or more cost-beneficial R&M improvement schemes. If the schemes that emerge are themselves sorted in descending order of expected net benefit (i.e. R-TOC) per dollar of investment, we can plot a strictly concave locus of improvement initiatives, as illustrated notionally at Figure 1 (the underlying data for which were generated randomly). The importance of this is that the set of points on the curve to the left of any arbitrary investment limit defines an optimal investment program. That is, implementation of the corresponding set of schemes is guaranteed to yield the largest possible overall benefit for that level of investment.

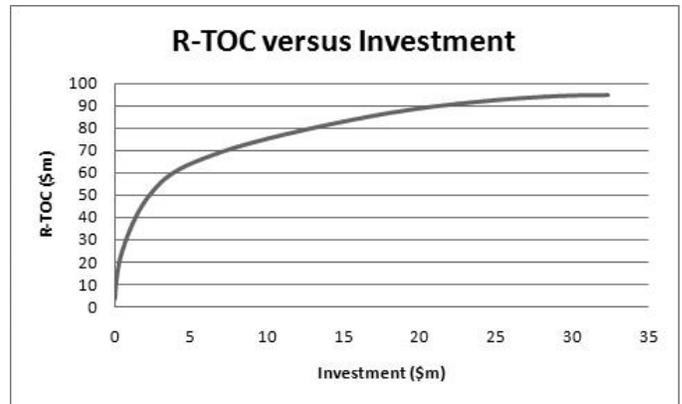


Figure 1: R-TOC versus Investment

Selection of opportunities in bang-for-the-buck sequence is known as *marginal analysis*. Late in the sequence (although sometimes not especially late) the so-called “law of diminishing returns” acts to damp down enthusiasm for further investment.

2.5 PBL Ramifications

The entry fee to the world of marginal analysis is an ability to evaluate the future TOC consequences of proposed improvement schemes. This implies a need for something that has always been somewhat problematic: an ability to reduce performance enhancement to monetary terms.

What, approximately, is the monetary value of a marginal increment (of, say, one percent) in A_0 ? One percent of the TOC, perhaps? After all, at the time of acquisition the government evidently deemed it worthwhile to expose itself to the entire TOC in return for 100 percent of the warlike utility the weapon system would confer. Or nothing? It may be that 80 percent is enough to cover all operational commitments, and the resource solution already in place is achieving 85 percent without difficulty.

Enter Performance Based Logistics (PBL). At its best PBL brings many benefits, but its great contribution in this context is to introduce objectivity to performance evaluation.

A support provider under PBL (perhaps the PSI mentioned above, or perhaps a second-tier supplier) stands to earn additional profit, the amount determined precisely in accordance with formulae enshrined in his contract, for achieving performance in excess of some threshold. Above some higher performance threshold, no further reward is available. A well-constructed PBL contract has a

“sweet spot” where the warfighter gets what he needs within reasonable bounds and the support provider maximizes his profit.

Ideally, performance requirements will have been specified solely in terms of one or more of the figures of merit discussed above, but this is not guaranteed to be so. Fill rate at the warehouse is still disturbingly prevalent. The bottom line is that performance expressed in terms of whatever metrics are mentioned in the contract has a precisely calculable monetary value while performance expressed in other terms has no monetary value whatsoever.

3 THE ANALYTICAL CHALLENGE

Reducing the proposition as thus far developed to its bare essentials, we have seen that the operational community’s appetite for weapon system performance enables program managers to obtain limited investment funding which can, at least in principle, be allocated rationally among schemes jockeying for inclusion in an overall improvement program. What remains to be established is to what extent the best program that can be put together in practice can reasonably be described as optimal.

The argument in the previous section suggests an ongoing decision process along the following lines:

- evaluate the future TOC contributions of weapon system components (to an appropriate level in the indentured bill of materials);
- sort the component list in descending order of future TOC contribution;
- set cognizant engineers to work (at whatever pace they can sustain) devising R-TOC schemes in component list sequence;
- as each scheme comes off the “production line”,
 - determine the net benefit it promises per dollar of investment, and
 - insert it at the appropriate place in a list of unapproved schemes sorted in descending order of net benefit per dollar;
- periodically release available investment funds to implement as many schemes as can currently be accommodated; and

- purge schemes from the list on approval.

Sadly, such a process cannot guarantee exclusion of all schemes that would not merit a place in a strictly optimal program. There are two main reasons:

- Because engineering capacity is finite, the list of candidate schemes will never be complete at a funding release point. There will always be the possibility of highly competitive schemes missing the cutoff. And waiting longer simply erodes the benefits of all schemes.
- Forecasts of component failure frequencies are notoriously unreliable. As time passes and forecasts are revised to take account of observed frequencies, estimates of future TOC contributions have to be revised *pari passu*, and lowly-ranked components can leapfrog others presently ahead of them on the list.

Accordingly, our ongoing decision process is exposed as a method that merely produces a sequence of sub-programs for which optimality might reasonably be claimed, but only within the brief lifespan of current knowledge. The overall program is inescapably suboptimal to the extent that creative contributions lag and forecasts change.

But since no decision tool short of a time machine can eliminate these difficulties, we would be foolish to worry unduly about them. Rather, we must embrace the challenge of assembling a suite of less-than-perfect prognostication tools that can:

- make the best of the information to hand at each iteration of the decision process, and
- minimize vulnerability to change in the inputs over time.

3.1 TOC Analysis

Disenchanted with large sustainment burdens becoming evident in the 1960s, the US DoD enjoined designers to take better account of downstream support costs. The consequent emergence of a formal Life Cycle Cost (LCC) discipline led rapidly to governments around the world requiring their agencies to consider LCC implications whenever faced with significant resource allocation decisions. Today, TOC analysis (LCC with minor embellishments) is entrenched in the front rank of analytical methods applied to the

design, specification, selection and support of weapon systems, generating costs per operating hour, (R&M) metrics, resource profiles, budgetary cost projection, and many other invaluable outputs.

The quality of these outputs is a critical issue. Decision-makers have been caught out repeatedly by discrepancies between their predictions and actual outcomes, attributable to widespread use of ineffectual planning and control methods, which have tended to be myopic, costly, too much the province of specialists, and prey to poor data. The success of weapon programs depends on the effectiveness of decisions taken at many hierarchical levels. For most such decisions there is no such thing as a clean sheet of paper. They need to take into account divergence in as-maintained configurations, deployment patterns and operating intentions, existing maintenance facilities and contracts, and a host of other considerations.

Early LCC models typically made heavy use of parametric approaches relying on cost-estimating relationships (CERs) derived through regression analysis of historical data. The simplest possible CER (and hence the one with the least claim to precision) has surprisingly many adherents. The then head of operational requirements for the Royal Australian Navy told the author some years ago that LCC models were a waste of money because the LCC is simply 2.5 times the acquisition cost. Now while it is arguable that LCCs over the last 40 years have gravitated to a range 2 to 3 times the acquisition cost (down from the 10:1 ratio typical in the early days), and while even such a coarse approximation will do little damage if applied to a range of comparatively low-cost options, the potential to earmark either \$10 billion too much or the same amount too little for a \$20 billion program should be worrying. Worse still, programs left short of funds as a result of such oversimplification will demand more, while those with an excess of funds will find ways to spend them. This “ratchet effect” can be especially damaging in a cross-program sense.

3.2 Evolution of the Approach

While there may once have been a respectable case for favoring parametric approaches early in the life cycle on the grounds that data were too uncertain to support detailed cause and effect reasoning, in

these days of COTS exploitation a great deal is usually known about at least some of the options under consideration. If a comfortable reliance on parametric assumptions induces managers in the concept-phase to avoid more exhaustive computation when the necessary information is, in fact, available, then much of the benefit sought from TOC analysis will be lost. But more to the point in an R-TOC context, parametric models simply have nothing to say about the future contributions of the individual components of a system.

Over the 40 or so years of the LCC / TOC era what might be characterized as a marked climate change (the same inexorable process, of course, that led to supremacy of the TLCSM paradigm discussed earlier) tipped the evolutionary scales in favor of an alternative to the parametric style – its antithesis in fact. Confronted with the grim reality that major new development programs would become fewer and farther between, the Aerospace and Defense (A&D) companies were forced to acknowledge the need to develop models that could grow without limit to accommodate the increasing volume and richness of data involved in incremental decisions.

3.3 State of the Art

Today, the dominant style is so-called *engineering analysis*, employing highly detailed activity-based or process models that go about their cost-estimating business by aggregating the resources predicted to be consumed by the many activities involved in introducing systems into service, operating them, and upgrading and maintaining their components over the life cycle. This focus on identifiable cost drivers and their effects is what distinguishes process models from parametric ones.

The best examples of such models address intermittence and discontinuity by preserving items of value from today’s analytical excursion and making them available for tomorrow’s. That is, the input dataset for each successive stage of analysis depends on the outputs of its predecessor stages. Over time, degrees of freedom tend to disappear and datasets become more specific and detailed. Clearly there is a need to maintain consistency and traceability throughout.

The author has been associated for the past decade with ongoing development of the TFD Group’s Monterey Activity-based Analytical Platform

(MAAP), a TOC-focused decision tool founded on Logistic Support Analysis (LSA) principles (which is to say, above all, that the analysis is responsive to the frequencies of failure modes rather than the consolidated failure rates of components). MAAP accounts for activities under assumptions about system deployment and operation that can change at will. Cost and resource profiles can be generated system by system and year by year over the program life or any portion of it. The upshot is an ability to assess the immediate and downstream cost impacts of proposed changes in operational activity, or the immediate and downstream operational impacts of budgetary changes.

The analytical engine is best seen as a large-scale discrete event simulation that produces resource profiles and budgetary estimates without going to the length of triggering individual event instances. An event in this context is any operational, maintenance, training, or upgrade activity that leads to a requirement for the presence or consumption of resources, whether parts, manpower, facilities, tools, data, or fuel. Since every component can be idealized as a set of events, the weapon system itself amounts to a collection of all such sets of events that its operation and support entails.

Component usage patterns follow from the operating profile in a straightforward manner, and can therefore be transformed readily into streams of events accurately situated in both location and time. This explicit recognition of time as an additional modeling dimension (a radical departure from earlier practice) enables capture of the resource implications of dynamic deployment schedules that may be unique to each weapon system instance. As a result, the complex interactions of changeable deployment and operating scenarios with phased development, introduction, upgrade, and retirement are reflected realistically in resources levels and flows, and ultimately in the TOC.

4 R&M OPTIMIZATION

A cherished TFD goal has been to expand the usefulness and appeal of TOC analysis by providing specialized utilities that make innovative use of MAAP output to investigate performance-related questions. Several utilities of this kind are packaged in the MAAP Performance Optimization Workbench for Enhanced Readiness (*mPOWER*). One such

utility, developed to meet the needs of managers wishing to optimize R&M programs along the lines discussed in the previous section, is called MAAP Progressive Investment in R&M Improvement Candidates (*mPIRIC*).

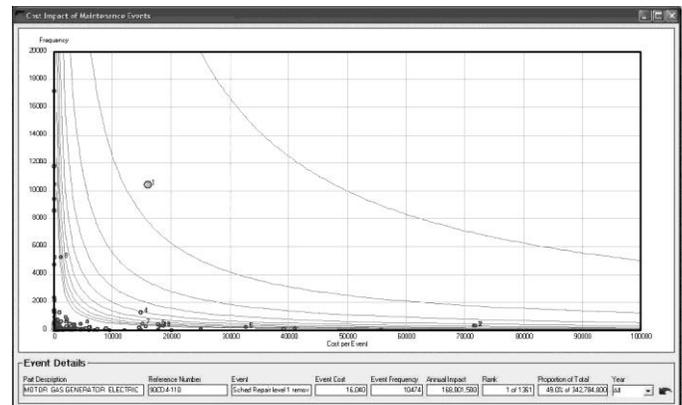


Figure 2: *mPIRIC* Scatter Plot

One of the more interesting MAAP outputs is a scatter plot of maintenance event (ME) frequency and cost. Distance from the origin (readily determined with the help of isoquants in the background) indicates a given ME's relative standing in terms of impact on the future portion of the TOC. The *mPIRIC* innovation is to bring this display to life with the primary aim of accomplishing the first two steps of our optimization process, but also facilitating evaluation of prospective changes in maintenance capabilities, maintenance policies, or operating procedures.

4.1 Identifying Opportunities

Clicking on a particular ME from the scatter plot provides access to a planning tableau enabling the analyst to portray the investment profiles and expected benefit streams associated with plausible

improvement schemes.

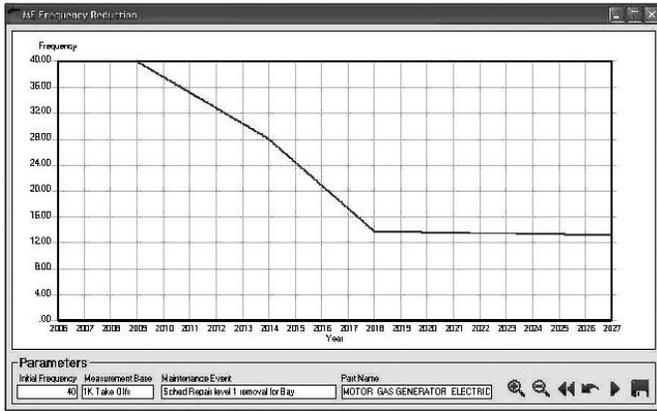


Figure 3: Improvement Profile over Time

In general, the ME frequency, cost, or both will reduce over time and settle eventually at new levels. Typical R&M modifications take years to roll out across a fleet, and the beneficial effects of such things as revised operating practices, better diagnostic capabilities, or greater depth of maintenance are also bound to flow through gradually. For analyst convenience, *mPIRIC* captures anticipated improvement (i.e. frequency / cost reduction) profiles as freehand drawings.

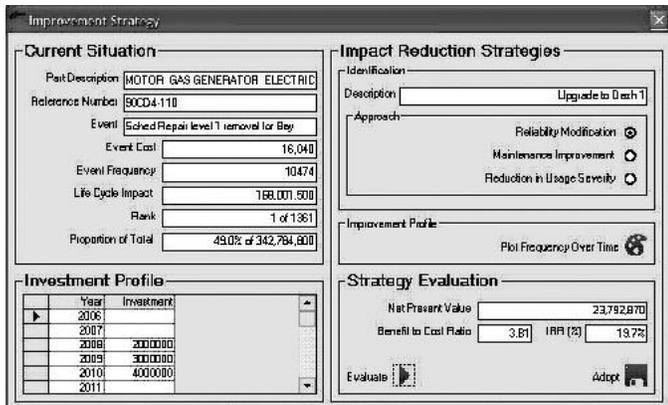


Figure 4: *mPIRIC* Planning Tableau

4.2 Ranking the Schemes

Once investment and improvement profiles have been entered, evaluation of a scheme is a click away. Schemes are evaluated in terms of their:

- net benefit (i.e. PV of the benefit stream over the remaining program life minus the PV of the investment stream);
- benefit to cost ratio (the bang-for-the-buck needed for marginal analysis); and

- internal rate of return.

4.3 *mPIRIC* Users

Several of the largest A&D companies and a European Air Force using MAAP have upgraded their capabilities through addition of *mPOWER*. While none has yet made systematic use of *mPIRIC* to optimize an R&M program, particular initiatives have been evaluated.

4.4 Future Development

A planned upgrade is an ability to transfer *mPIRIC* cost and frequency data automatically into MAAP operating and support scenarios on scheme approval, thus closing the planning loop and ensuring subsequent TOC assessments reflect the new *status quo*.

REFERENCES

1. As listed in *Performance Based Logistics: Purchasing Using Performance Based Criteria*, 16 August 2004, a memorandum from the office of the Under Secretary of Defense for Acquisition Technology and Logistics (USD AT&L). This document also (unhelpfully) lists Cost per Unit Usage as a performance criterion.
2. *Transformation through Reduction of Total Ownership Cost (R-TOC)*, 16 December 2003, another USD AT&L memorandum. This document expounds a somewhat breathless R-TOC Vision: “Through R-TOC principles, all defense systems will perform with increasing readiness and capability while avoiding increased operations and support costs and improving logistics footprint by institutionalizing the continuous implementation of innovative process and hardware improvements”.

BIOGRAPHY

John Millhouse
TFD Group
70 Garden Court, Suite 300
Monterey, California 93940 USA
e-mail: john.millhouse@tfdg.com

John Millhouse is Chief Scientist of the TFD Group. Educated as an aeronautical engineer, he served for more than 30 years in the Royal Australian Air Force. After flying helicopters in Vietnam, he spent his subsequent career in a variety of engineering and senior logistic positions, including Director of Integrated Logistics Requirements at Air Force HQ and commander of a major logistics wing. In his present role at TFD, Mr. Millhouse formulates the algorithmic strategies behind an extensive suite of analytical and simulation models, notably a new generation of multi-resource optimization tools already exerting a strong influence in the execution of performance-based contracts for support of major weapon systems and commercial airliners in both the US and Europe.