Accounting for Time in Inventory Planning: Updating Spares Optimization Theory

Robert A. Butler
TFD Group

Précis

Modern spares optimization theory dates to the development and testing of the Base Spares Model (BSM) [1]. The test, carried out at George AFB during 1965-66, demonstrated that systembased spares computations could cut inventory in half compared to single-item methods. The first implementation, METRIC, was published by RAND in 1966[2]. Virtually all serious spares optimization models currently in use owe their theoretical structure to Sherbrooke's efforts, amplified later by VAR-METRIC.

Virtually all METRIC class models use the assumptions of long term and steady state. These assumptions make the mathematics tractable, but have the drawback, never seen as a problem until recently, of ignoring the crucial role of time. Problems such as fleet build-up/run-down, allocation of budgets between long and short lead time items, part obsolescence and foreseeable changes in operating Tempo, goals, basing and fleet size/composition are all illserved by steady state models. A multi-period optimization model solves a number of these problems by allowing the analyst to explicitly account for such changes.

One of the more interesting consequences of shifting focus to multi-period modeling is that the first period solution will invariably be more expensive, whereas the discounted present value of the entire series of solutions is far less costly. The conclusion is that multi-period optimization minimizes *inventory life cycle cost*, rather than a single solution for the never-realized steady state.

A Brief History of Spares Computation

The origins of spares analysis are obscure, but are undoubtedly tied to the rise of operational analysis during World War II. Ideas appear to have evolved in parallel lines since around the mid-1960s when Sherbrooke's work at RAND began the sequence of improvements in system optimization that led to what is now generally accepted as true system inventory optimization. Oddly, at the same time, most professional logisticians around the world were using and attempting to improve far more (mathematically) simple formulations based on single-item modeling. Blanchard, in his 1998 book[3], helped to popularize the "constant k" version of single item spares modeling, based on the demands during lead time (λt) associated with each part. A number of other models have evolved during the same period, based on early work on EOQ or economic order quantity and saw-tooth models to determine re-order point and order quantity. In general, these models are related to the demand for spare parts used in manufacturing processes, which have also given rise to the "just in time" concept.

In this paper we are concerned with the computation of spare part inventory solutions required to support the on-going operation of maintainable equipment. We specifically exclude any concern for methods appropriate to the manufacturing process.[4]

---

[1] Feeney, G. J. and Sherbrooke, Craig C., "A System Approach to Base Stockage of Recoverable Items," RM-4720, RAND Corporation, Santa Monica, 1965.

[2] Sherbrooke, Craig C., "METRIC: A Multi-Echelon Technique for Recoverable Item Control," RM-5078-PR, RAND Corporation, Santa Monica, 1966.

[3] Blanchard, Benjamin, *Logistic Engineering and Management,* Prentice Hall, 1998, Upper Saddle River, NJ

[4] An allied problem might be termed the "service center" problem, which entails calculation of spare parts quantities when the supplier has no direct insight into the operation of the several system fleets whose operations generate the demands he sees. This kind of problem is becoming increasingly interesting in the commercial airline sector, where service center

The basic structure of the constant k model is $S = [\lambda t + k\sqrt{\lambda t}]$, where S, the number of spares, is set equal to the (rounded up) number of demands during lead time plus *k* times the square root of that value. This formulation is based on the notion that the Poisson distribution, thought to accurately characterize the arrival of demands at a stockpile from a random process, is analogous to the Normal distribution. Since in a Poisson distribution the mean and variance are equal ($\mu = \sigma^2$) the square root of λt is the standard deviation of the demands during lead time distribution. The greater the value of *k*, the larger the buffer stock set aside against variability in both the demand rate and the lead time. The problem with this formulation is that it considers each part in isolation – it is a "single item model." Because each part is considered in isolation, models of this sort ignore the relative value of one part over another, even when their λt values are equal. To see why this is an error, imagine two such parts, one of which can be purchased for $1 and the other for $10,000. Since both parts will have an equal impact on expected backorders (EBO), the cost effectiveness of buying the $1 part is 10,000 that of buying the other one. This cost effectiveness measure is familiarly known as the "bang for buck" ratio and we will use that notation later in this paper.

The system optimization models (as opposed to single-item or constant *k* models) evolved slowly over the period between 1966 and about 1990. During that period, the models first began to look at hardware indenture levels below the line replaceable unit or LRU[5] and attention among users of the models shifted from an objective function of fill rate (probability of being able to fill an order at the base) to operational availability.

In 1996, Sherbrooke published his seminal paper introducing VARI-METRIC[6]. This work finally solved the problems associated with variability in both demand rate and delay time and came into widespread use around the world.[7] Later, about 1998, Sherbrooke, in unpublished work, collaborated with the author to develop the n-echelon (real locations) capability of VMetric®[8].

In the US, because of the large size of deployed fleets and profusion of bases, models abstracted the complex of airbases and shops as echelons of maintenance, making the implicit assumption that all bases were the same (same number of aircraft, same operating pace, same delay times between them and intermediate sites and so on), while in Europe (particularly in Sweden) the more modest size of the military establishment made it more significant – and more feasible – to treat each real operating or support location as a distinct entity. As micro-computer capabilities increased, it became increasingly feasible to handle the "different base case[9]."

---

contracts are being negotiated between multiple customers flying similar aircraft, each of whom seeks different levels of protection for different ranges of parts from the service center provider. We mention it here because, like spares calculations to support manufacturing, this paper does not deal with the service center problem.

[5] Line refers to the flight line, as these models were first developed for aircraft inventory problems. The significance of the term in modeling must not be misunderstood. An LRU is not defined by any criterion of weight or size or even convenience of removal, although all these attributes are important in the maintainability design of hardware systems. An LRU is simply any part which, by being removed and replaced at the site of equipment operation, can restore a broken system to operation. Because that is the case, an LRU has a direct impact on system availability. An SRU – a shop replaceable unit – does not have the same immediate effect on availability, but only on the cost and delay involved in making a failed LRU available to perform its crucial role again. US Naval usage substitutes the terms WRA and SRA for LRU and SRU, meaning weapon removable assembly and shop removable assembly respectively.

[6] Sherbrooke, C. C., "VARI-METRIC: Improved Approximations for Multi-Indenture, Multi-Echelon Availability Models," *Operations Research* 34, 311-319.

[7] First fielded commercially by TFD Group as VMetric, the theory is also applied in the US Army model, SESAME, the US Navy model ACIM, the US Air Force Models Dyna-METRIC and ASM, NASA models and the Systecon model, OPUS®. Indeed, it is unlikely that any serious spares optimization tool would ignore Sherbrooke's innovation.

[8] VMetric is a registered trademark of TFD Group.

[9] Notwithstanding these abilities, the US Navy continued to focus on two echelons and did not even use a multiechelon model in their actual provisioning efforts – CARES and ARROWs are used by NAVICP to model wholesale and retail stock separately. The same agency also fields ACIM, a multi-echelon model, used in combination with TIGER, a simulation model whose purpose is to adjust utilization and criticality rates for multi-mission system components based on their usage in various mission. It should also be noted that both SESAME, the US Army spares model and ACIM are VARI-METRIC

Understanding Marginal Optimization

There is a great deal of mathematics and statistics involved in the formulation of optimal spares models and the proof of their correctness. A lucid and complete account can be found in Sherbrooke's 1992 book, *Optimal Inventory Modeling of Systems*[10]*.* Here we will restrict ourselves to the simple matter of how marginal optimization works and how the traditional, time-insensitive, model must be altered.

A marginal optimization technique works by continuously asking, what is the next best part/location choice, where 'best' means that choice exhibits the highest absolute ratio of decrease in *expected back orders* (EBO) to the price or cost of buying the part

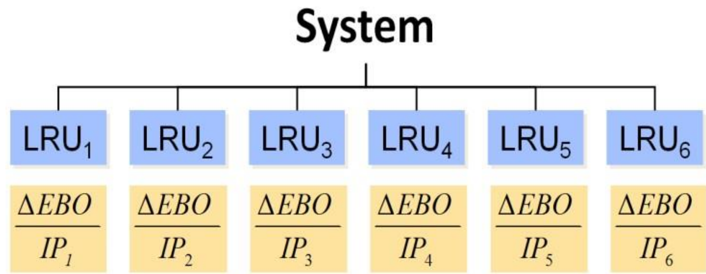

Figure 1: A Simple System's Bang for Buck Ratios

and putting it at that location. In the illustration in Figure 1, a system is made up of only six parts and a single location must be stocked. Starting with no stock at all, a calculation is made for each part to determine how big a drop in EBO would be created, ΔEBO, if one of that part were stocked. That decrease is divided by the item price of that part, IP_i[11]. Marginal optimization proceeds by choosing the largest (negative) ratio.

Once chosen, the ratio must be computed for that part again, since the second part of that type can be expected to have a smaller effect on EBO than the first did. Having recomputed the ratio for that part, all the ratios are looked at again and the part with the biggest ratio is selected. This process continues until there is either no discernable further effect on EBO (the computation has reached the upper asymptote) or until a specific goal, such as fill rate or A_o is reached. It goes without saying that multiple indentures of hardware and multiple repair, operating and storage sites or echelons of maintenance complicate the computational process tremendously. Nevertheless, modern micro-computer equipment and speedy algorithms make it possible to carry out computations of staggering size within reasonable time limits[12].

implementations similar to an earlier version of VMetric. The service models, however, remain simple multi-echelon models, unable to cope with realistic deployment scenarios.

[10] Sherbrooke, C. C., *Optimal Inventory Modeling of Systems,* 1992, John Wiley and Sons, New York.

[11] The illustration is a simplification, of course. For example, it note that the cost of a spare part, located at a given facility, is very often more than just the item price of the part as shown here. The extreme case is that of orbital replacement units used by NASA at the International Space Station. The cost of "up mass," that is, lifting the part into orbit, as well as a storage space scarcity penalty, must be accounted for in addition to the part's purchase price.

[12] In a study of alternative spares technologies run for the US Navy by the Logistic Management Institute in 2002-3 (Navy contract number N00014-02-F-Q999) the problem included 4 system types (variants of the F-18 fighter), 10 aircraft carriers, one naval air stations and the depot. The total parts count required to model all four systems exceeded 10,000. Nevertheless, this problem was finally run by VMetric in about 30 minutes.

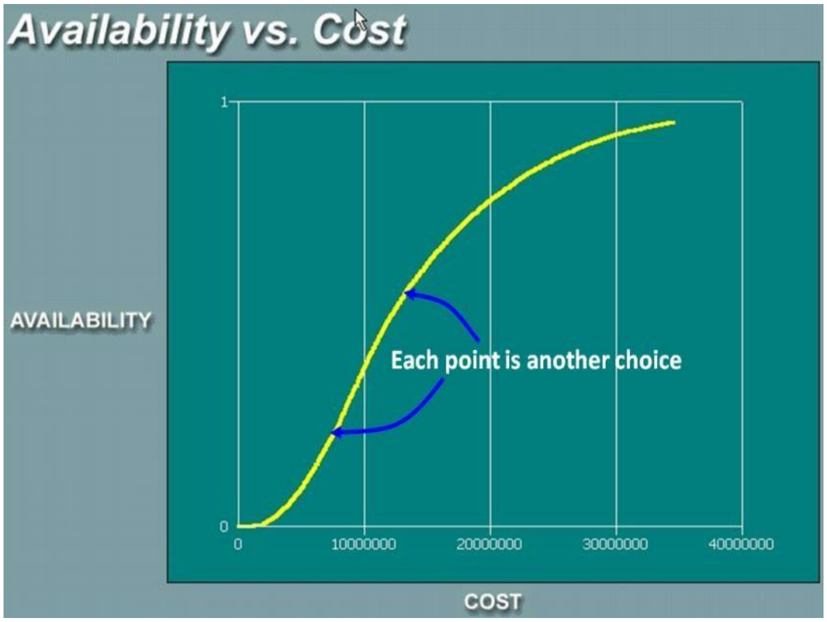

Figure 2: A_o versus Inventory Cost

The result of such a calculation is a curve, shown in Figure 2, in the shape of a traditional economic production function, that shows how spares availability rises with most efficient investments in inventory[13]. All points above the yellow line are infeasible – there is no way to get greater availability at any expenditure than shown by the dot at that investment. Moreover, all points below the yellow line are inferior. As a result, we refer to the yellow line, the trace of marginal calculations made as shown above, as the *locus of optimal solutions*. The ability to trace these optima is of real importance in practical inventory decision making. Among these reasons, the most important one seems to be the ability to navigate along the optimal locus backward from a desirable, but unaffordable solution. Each backward step can be interpreted as eliminating the least cost effective spare part in the total solution to that point.

Time-Based Inventory Optimization

We come now to time and the role it plays in both the practical and theoretical aspects of inventory optimization. The spares planner is regularly confronted with problems involving time. Some of the problems are recognized and dealt with others either go unrecognized or simply ignored because their solution is too difficult. The main spares optimization problems we associate with time include:

- Fleet build-up or multi-period initial provisioning
- Fleet run-down and end-of-life problems
- In-Service inventory problems o Long lead time items o Obsolescence

In the process of developing a solution to problems associated with time, we determined that the mechanism adopted would also serve to confront a further problem, namely, how to deal with complex spares objective functions. In general, optimization routines have been based on reducing expected backorders and translating the achieved levels of EBO into operational availability or whatever other measure of effectiveness (MOE) was of interest. When more than one MOE must be satisfied, however, this procedure is deeply flawed. Before exploring the simple changes in theory required to account for time, we first want to understand the drawbacks of spares optimization in its current best-practice form.

Drawbacks of Steady State Models

Models like VMetric and OPUS have saved a lot of spares investment money through the years, and led to significant increases in achieved operational availability for fleets of many types of systems. They have also left some problems in their wake, related to their fundamental structure and approach. Technically, these models are known as long-term, steady-state models. To understand their limitations, it is necessary to understand what these technical terms mean.

[13] Note that operational availability is the product of maintenance and spares availability: $A_o = A_m \cdot A_s$.

First, "long-term" means forever – an infinite time horizon. The significance of this view of time is that the real world has forever for the model's answers to be proved correct. In practical terms, it means that if availability falls short of the goal for a few time periods, the shortfall is compensated for by availability running above target for a while. In Figure 3, it is easy to see why this creates a problem. To deliver less than the target value of A_o or some other MOE usually leads to a penalty or lost incentive fee (red areas). Moreover, it is usually true that little or no value is attached to over-delivery of A_o or another MOE. When using a long-term steady-state spares model, analysts can only solve this problem by setting a confidence level for achievement of the MOE target value. Doing so decreases the red areas to an acceptable level, but increases the green areas (uncompensated over-achievement) correspondingly. The ideal solution, of course, would be one in which the returns to achieving the target MOE more frequently were balanced against the cost of achieving them.

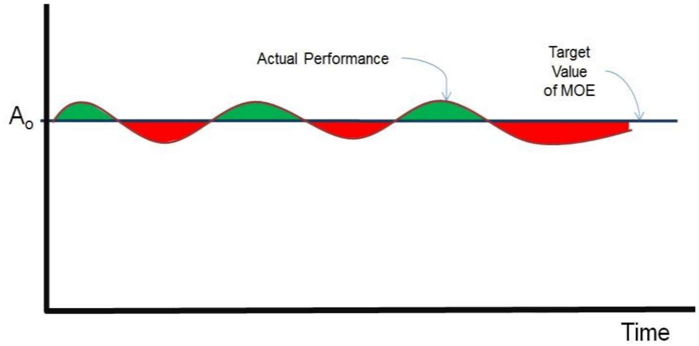

Figure 3: A_o Performance Over Time

The next thing to consider is the idea behind the term steady-state. We use models to predict the future. For some few things (operating pace, deployment and others) we know something about the future because those things are subject to what might be called reliable planning. For most of the inputs to a spares optimization model, however, all we know is what we have observed up to the moment that we run the model. The hardware breakdown structure, together with the attributes of all the parts that make it up, the behavior and capabilities of the repair and transportation system such as repair rates, repair times, delays in retrograde and forward shipping, procurement lead times for each part, condemnation rates, retest OK rates and so on, are all nothing more than projections from (often imperfect) data describing the past. Nevertheless, the model has no option but to assume everything it is told is true. Worse, it must also assume that all those things will remain true over the long term, which, as we just noted, means forever.

Neither the long-term, nor the steady-state assumptions are crippling insofar as they enable us to get a better answer than we would if we abandoned system-based analysis as a means of solving the decision problem.[14] Nevertheless, the drawbacks of these assumptions are clear, as is the likelihood that improvement is possible. The steady-state assumption seems especially crude when we know with certainty that many things will change. When, in fact, we plan to change things and would like to see the accurate implications of those changes.

Fleet Build-Up

The simplest case in point is multi-period delivery of new systems. Current models produce a single set of answers (stock mix, level and distribution) that is assumed to represent the definitive inventory "strategy" for the balance of the life of the system. But, what happens when a fleet is built up a few

---

[14] Ample proof of this was given by the initial test of the most primitive form of Sherbrooke's model called the Base Spares Model, tested at George AFB in Fall 1965 and Spring 1966. In this exercise, the not operationally ready due to stock (NORS) rates of aircraft stationed at George AFB were measured before and after change-out of stock from what was available in the first three months to what BSM recommended for the next three months. The study found no statistically significant difference in NORS rates even though the BSM-computed inventory cost half what the original inventory cost. Interestingly, in every spares study conducted at TFD Group over the past 25 years, this ratio of about twice the investment required to achieve a given MOE has been repeated, at least with aircraft. The ratio is far worse with ships, ground equipment, buses, trains and the like.

systems at a time, as would happen with new aircraft delivery? The analyst would be constrained to pick some point in the process such as after the first delivery, say, or when the entire fleet is assembled, and run the model reflecting the circumstances anticipated at that time. Neither of these runs would be quite satisfactory, however. In the first case, there are no answers for what stocks are required after the first period and in the second, there are no answers for what stocks should be planned for the build-up period. The solution adopted by VMetric users has been to run the model for the first delivery and load the answers into the next period's run as initial stock. This process is time consuming and fails to take cognizance of the economics of time.

Fleet Run-Down

The matching problem to fleet build-up is the run-down of the fleet being replaced. There are many stock decisions required to manage run-down in a sensible way. As systems are withdrawn from service, some stocks become redundant. Nevertheless, there will still be shortages that, if not filled, will cause the fleet MOEs to drop below acceptable levels. When systems are taken out of service, they are often sold to other users, who also want to procure stocks of spare parts with their (new to them) systems. Selecting which parts to release and making decisions about buying new parts whose normal longevity exceeds their useful life in this system are both problems whose solutions are at least partially dependent upon cognizance of time.

In-Service Inventory Problems

Every inventory manager finds himself wrestling with the trade-off between long lead time items and short lead time items. The latter are clearly more valuable in the sense that they lead to an increase in benefit-dependent MOEs right away, while long lead items will only have their effect later. Complicating this difference is the fact that, to correctly measure the difference in value, the impact of the long lead item must be discounted to present value. If discounting is not done, the long lead item will be overvalued. While making the discounting calculation is simple, no computation made outside a model run is both simple and capable of preserving optimality of the sequence of choices.

A more significant problem than the long lead time problem is the phenomenon of obsolescence. More accurately, the problem arises because of differential rates of obsolescence among the components of a system, a problem that has only come to the forefront with the increasing use of off-the-shelf, commercial components. Usually, the system under study (especially, defense systems) is made in such small numbers that the demand it represents for components is insignificant to the manufacturer. As a consequence, components go out of production quickly, replaced by better, cheaper or merely different ones. This, in turn, drives a continuous process of hardware update, usually called technology insertion, that serves to cause stocks of electronics components in particular to become obsolete on a cycle of three to five years.

The problem this phenomenon causes for spares optimization is readily understood by an example. Imagine that there are two components which are identical insofar as their crucial attributes are concerned: demand rates, delay times, repair rates and so on are all the same. However, one of the parts is electronic and the other mechanical. All current spares optimization models will buy the same number of each of these parts even though this is obviously the wrong choice. The electronic component will be replaced by something with a different part number within a relatively short period of time, making its stocks redundant[15]. On the other hand, the mechanical assembly (usually designed expressly for this system) will be useful for the remainder of the system's life[16].

---

[15] There are two possibilities. Either the old part number is no longer useful on this system or it is an inferior substitute for the new one. In some cases, an inferior substitute can also be modified to become the new part number. If no longer of value on the current system, it may have a diminished value on the open market.

[16] Note the difference between useful life (which measures time to obsolescence) and expected life (which measures time to wear-out). The useful life of a car tire, for example, may be coincident with the life of the car. Its expected life, however may be limited to 40,000 miles. To avoid these name confusions, at TFD we have coined the phrase Mean Technological Life or MTL to describe a part's useful life.

The Mechanics of Spares Optimization With Time

Adapting spares optimization processes to recognize and deal with time is not very complicated, at least in theory, if not in practice. The basic issue is to modify the ratio of change in expected backorders (ΔEBO) to the unit price of the item that caused that change. We shall call this the cost-effectiveness ratio or CE in the following discussion[17]. In current practice, we define the ratio as $\mathrm{CE} = \frac{\Delta\mathrm{EBO}}{\mathrm{IPi}}$.

Recognition of time, however, leads us to understand that the change procured for EBO will only begin at some point in the future (lead time) and only persist for some period (the time before obsolescence or the end of the system's life). Hence, the cost effectiveness ratio defined in the context of time can be $\mathrm{CE}_t = \frac{\Delta\mathrm{EBO}_{t1\ldots t2}}{\mathrm{IPi}}$

Time-based value calculations must always include discounting to ensure that delayed benefits are valued appropriately (less) compared to current or nearer-in-time benefits. We identify, for any decision-making entity, its appropriate discount rate[18], *d*, and calculate the present value of a benefit *X* delivered *t* periods in the future as $\mathrm{PV}_0(X_t) = \frac{X_t}{(1+d)^t}$. This procedure must be applied to the sequence of EBO changes occurring between t_1 and t_2, engendered by the investment of IP_i. By doing so, we are able to calculate the present value of the stream of benefits, the present value of which for each period is given $PV_0\ (\Delta\mathrm{EB}O_t) = \frac{\Delta\mathrm{EB}O_t}{(1+d)^t}$ and the present value of the sum of changes is given by the following.

$$PV_0(\Delta\mathrm{EB}O_{t1\ldots t2}) = \sum_{t=t_1}^{t=t_2} PV_0\ (\Delta\mathrm{EB}O_t)$$

Finally, the time-sensitive cost Effectiveness ratio is defined as $CE_t = \frac{PV_0(\Delta\mathrm{EB}O_{t1\ldots t2})}{IP_i}$

Complex Combinations of Measures of Effectiveness

While the calculation of CE_t can be carried out by applying the discount factor to the probability measure EBO, it can also be used with an estimate of the monetary value of changes in EBO. Mindful that the purchase of a spare part represents an investment whose return is measured by the impact of EBO changes on a variety of measures, the value of ΔEBO in each period can be calculated as the linear sum of those monetary changes. For any period, the monetized value of a change in EBO can be portrayed as:

$$PV_0(\Delta\mathrm{EB}O_t) = \frac{1}{(1+d)^t}\sum_{I=1}^{n} f_i[\,g_i(\Delta\mathrm{EB}O_t)]$$

Any change in expected backorders will have differential impacts on operational availability, fill rate, average delay time, time to fill, dispatch reliability, and many other measures. Each MOE or key performance indicator (KPI) can be calculated as a transform of EBO, which is represented in the

---

[17] In practice, this ratio is often referred to as the "bang-for-buck" ratio, a less elegant, though more familiar term.

[18] There is considerable confusion and misunderstanding among engineers and logisticians (and one assumes, many other groups) about the correct definition of the discount rate. It is most often confused with a conventional interest rate, especially in textbooks addressing the topic of engineering economy. The discount rate used by economists plays a very different role than a market rate of interest. Instead of portraying the marketclearing price of money, the *appropriate* discount rate for each decision maker determines whether he will be a borrower or a lender. In the language of Fisher (*The Theory of Interest*, New York, 1930), the discount rate measures an individual actor's "rate of impatience" to consume. The more impatient to consume, the higher a return he would demand to delay consumption.

equation as the functions, g_i. The contractual stipulations by which rewards or penalties are computed on the basis of satisfying these requirements or failing to do so are represented by the functions f_i. In the situation in which the owner of the system also provides his own supply support, rather than contractual rewards associated with the MOEs, we would expect to find a set of managerial targets, appropriately weighted for the relative importance attributed to them by executive management.

It should be noted that, although the ability to optimize spares calculations according to the logical requirements of complex incentive schemes is not a function of time, it is an important step forward in the calculation of optimal inventory. Further, because the time context is present, the monetization of EBO provides a straight forward business case analysis for every marginal stock choice. Hence, the resulting marginal optimization curve can now be described as representing the locus of maximum return of investment decisions, rather than simply the locus of lowest EBO per investment quantum.

The Impact of Accounting for Time

The preceding might be dismissed as nothing more than an interesting argument if, in practice, the changes in computation outlined did not produce significant differences in stock choices. To test the importance of these changes, an engine was built to make time-based calculations, and these were compared with a conventional VARI-METRIC based model.[19] Long lead time items versus short lead time items were compared by simply using the data present in the calculation. Obsolescence was studied by attributing time to obsolescence or mean technological life (MTL) to technology classes of components such as electronic, hydraulic, pneumatic and other mechanical. The most interesting comparisons involve the introduction of obsolescence and its impact on the total cost of spares.

A series of comparative runs – more than 5500 altogether – were made in both VMetric and TEMPO to investigate the impact of a variety of variables, on spares cost when obsolescence is considered. The form of obsolescence that cause the most concern is market-driven technological obsolescence. In this form, components are often taken out of production after a relatively short life; from three to five years in most cases. Considering the change in the "bang for buck" ratio or *CE* in the discussion above, we varied several key inputs to see what their effect would be. The values over which each variable was used are shown in Table 1.

Variable	Values From	To	Steps of
Program Life	5 years	50 years	5 years
Discount Rate	2%	12%	1%
Mean Technological Life (MTL)	2	6	1
Proportion of Parts Subject to MTL	10%	100%	10%

Table 1 Sensitivity Variables and Ranges

[19] The comparison was based on spares estimates for a fictitious aircraft system constructed of 1,500 components compiled from a number of aircraft. The control calculation utilized VMetric® 4.0 and the time-based calculations were made with a prototype of the new TEMPO™ calculating engine.

The general approach of the comparative analysis was to produce the first period run in TEMPO with exactly the same data used to run VMetric. Any differences in the results are explained by data elements taken account of by TEMPO and not by VMetric. These include assumptions about the proportion of parts subject to technological obsolescence (MTL < system life), the length of remaining MTL, the discount rate and the length of the system's remaining life. All differences in the sequence of part selection appear to be sub-optimal according to the assumptions of the steady-state model. The reason for this is illustrated in Figure 4. In the figure, the TEMPO computation shifts off the blue locus of optimal solutions when the first part subject to obsolescence is encountered. This is the point at which the change in the calculated value of *CE* is less in TEMPO than in VMetric because of a shorter useful life.

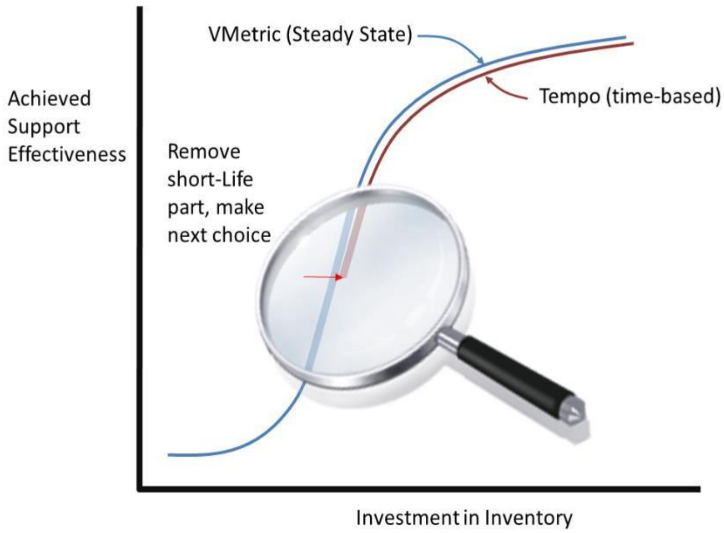

Figure 4: Shifting from Optimal Locus in Tempo

The question, then, is how is an apparently sub-optimal answer correct? The answer lies in the difference of perspective from the unrealistic steady-state, infinite time horizon view to the explicit treatment of time in a multi-period model. To capture the cost of the supposedly optimal choice that TEMPO rejected, one must look forward to the moment in time, two or three years hence, when the full stock of that part must be discarded and replaced by another part number of approximately equal cost due to obsolescence. That is the situation with the steady-state solution. The TEMPO solution, however, anticipated that problem and skipped over that part to find the next best choice, using its different definition of best. These effects are illustrated in Figure 5. The comparative analysis was conducted by spinning out the implications of these differences over the life of the system – and discounting the consequences to present value.

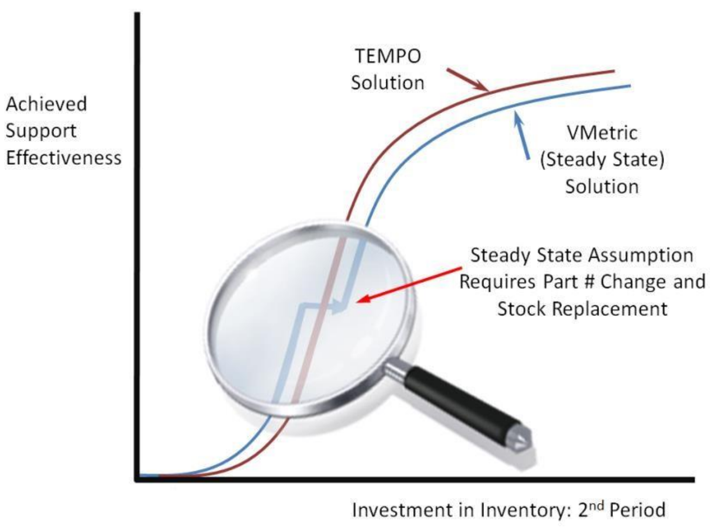

Figure 5: TEMPO vs VMetric 2nd Period Difference

In the illustrations that follow, the differences between costs of VMetric-generated solutions and TEMPO solutions are the "life cycle cost" differences generated by the need to discard stocks of obsolete parts in VMetric solutions, which is not done in the TEMPO solutions. It should be noted that none of these differences are due to foreseeable changes in basing, fleet size, reliability growth operating program and other factors that planners may know and wish to plan for. Such changes would also cause often profound changes in the TEMPO solution set compared to the optimal set generated by VMetric or other steady state optimizing models.

Taking these factors into account – both predictable changes in technical obsolescence and foreseeable, but not predictable changes in operating Tempo, basing, fleet size and maintenance capacity, the difference between TEMPO and VMetric answers amount to a difference between minimizing the *life cost of inventory* versus optimizing the inventory required for a (fictitious) steady state. An important question is to determine how significant these differences are. If they are trivial, then it wouldn't matter which approach one used. The comparative sensitivity analysis we conducted to test the differences indicates that mean technological life (obsolescence in any form) is the most significant factor. It's importance is modified depending on the remaining life of a system, and the discount rate. The normally predictable changes in a program that describe fleet size, basing, operations and support environment will also have very significant impacts in most cases. An example of a factor that does not have as big an impact as one might think is that of long-lead time purchasing decisions. Basically, the only difference between a TEMPO and a VMetric answer in this regard is that VMetric would tend to spend too much on long-lead time items because it can't sense the decrease in return on investment accounted for by discounting to present value.

A large number of comparisons can be made from the 5500 runs of the study. Only a few are illustrated here to give the reader some idea of the savings available from shifting to a multi-period perspective.

In the first illustration shown in Figure 6, each graph indicates that the relative savings (measured as % increase of VMetric solution over TEMPO solution) increases as program life increases. These relative increases arise from the fact that the longer the life of the system, the more times part stocks would be thrown out that had been surpassed by new designs. This effect recurs with most of the comparisons we looked at, although driven for different reasons.

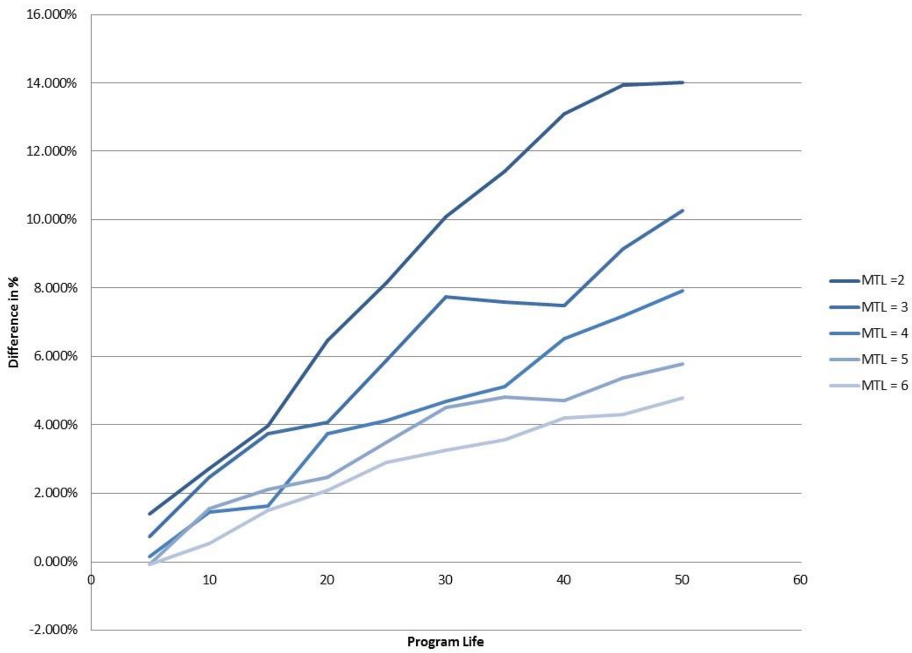

Figure 6 Program Life Versus MTL

There could easily be a question of how significant the discount rate is itself. This is especially true at the time of writing, when interest rates are forced to historic lows as part of the economic recovery strategy being followed by central banks all over the world. The graphs in Figure 7 indicate that this expectation is false. Moreover, that discount rates interact with mean technological life (MTL) in such a way as to complicate the matter quite a bit.

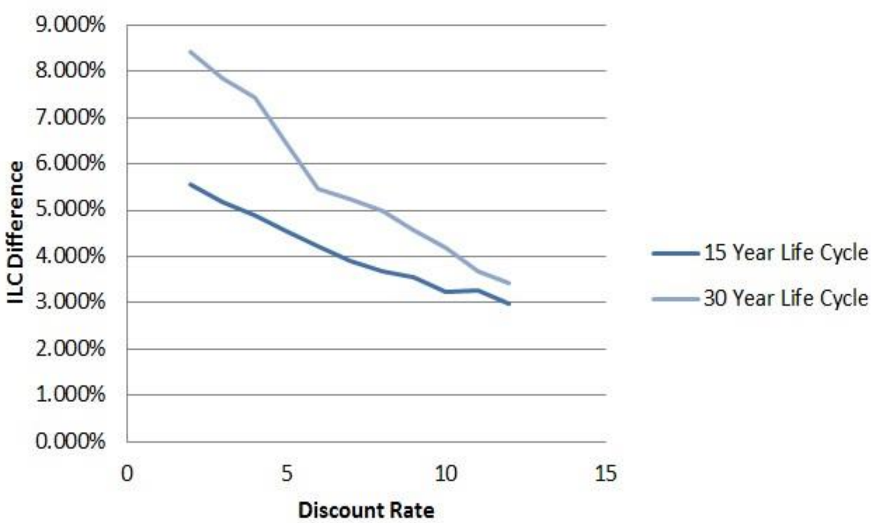

Figure 7 The Effect of Discount Rates

The problem amounts to this: discounting to present value expresses a lower value of both costs and returns to investment. In terms of the TEMPO problem, long lead time items are less important than they are in VMetric because the value of the performance they enhance is delayed. Hence, the higher the discount rate, the bigger the difference there would be between a V Metric and a TEMPO solution. Alternatively, the more replacement cycles there are (either due to shorter MTL or longer program life) the less their present cost due to increasingly important discounting. These two effects counter each other and are, themselves subject to complex reactions. For example, discount rates will have a bigger effect on longer program life than on short program life, no matter the average length of MTL. The shortest MTL, however, may swamp the effect of discounting for intermediate lengths of program life.

The examples in Figure 7 show the counter-intuitive result that comparative costs of the steady-state solution and the TEMPO solution are *inversely* related to the discount rate. That is, the savings gained by using the time-sensitive result are greater the lower the discount rate. This result was seen program life and MTL value. These results suggest that the effect of discounting on cyclic replacement of inappropriately purchased stocks always outweighs the effect of diminished value of future achievement of availability or other monetized support metrics.

The graphs in Figure 8 show a declining impact on total life cost of part populations exhibiting increasing average mean technological life. The longer program life and the shorter MTL, the bigger the difference between a steady state solution and the TEMPO solution.

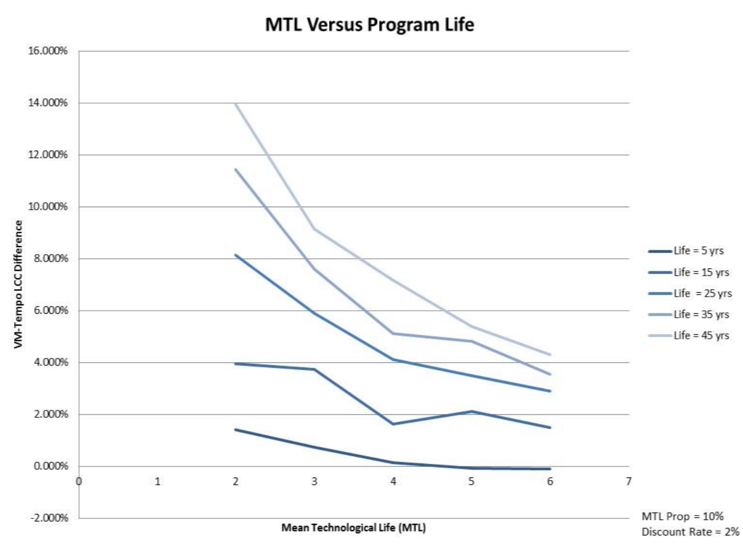

Figure 8 MTL versus Program

Note that when MTL exceeds program life, the TEMPO solution is coincident with the steady state solution, with the exception of parts that are due in after the first time period. Tempo knows these are missing, but steady state models do not. In that case, the steady state assumption holds – if no foreseeable changes occur in the conditions modeled such as fleet size, basing, operating pace and so on.

For more information on TEMPO or other TFD products and services please contact:

In North America:	**In Europe:**
TFD Group	TFD Europe
+1 831 649 3800	+44 (0) 1603 726660
+1 831 649 3866 (fax)	+44 (0) 1603 746431 (fax)

TFD Group maintains relationships with representatives in many other countries. For contact information on representatives, please consult our web site under *About Us*.

www.tfdg.com