Setting Performance Levels in Performance Requirements: Challenges and Solutions

Michael J. Addis
Codarra Advanced Systems Pty Ltd
Melbourne, Australia

Introduction

Function and performance specifications primarily contain requirements on what a system has to do (functions) and how well it has to perform (performance).[1]

One of the most challenging and problematic aspects of performance requirements is setting performance levels. Setting a performance level is not a yes-no choice, but potentially a choice from an infinite series of possible levels.

There are a number of established ways of specifying performance, and while the most commonly used methods are simple they are also highly limited and unlikely to result in a system that has an optimum level of performance. This paper will look at a variety of these techniques, their limitations and how they can best be used.

More sophisticated techniques have been developed and are in limited use. This paper takes one of these techniques and develops it to become a more powerful technique that has a number of uses and benefits for contracts, tender evaluation, and for the determination of liquidated damages.

This paper focusses on specifications that are used within a customer – supplier contractual arrangement. Many of the concepts can also be applied to situations where the customer, system designer and system builder are within the same organisation.

Challenges

There are a number of challenges (traditionally called problems) in writing performance requirements. In this paper these challenges will be grouped into choosing performance levels and specifying performance.

Choosing performance levels

No shades of grey

To be verifiable, a performance requirement must include a performance level, which will usually be a range into which performance must fall, or a threshold limit that it must not exceed. If there is no performance level, then there is no way of determining whether a requirement has or has not been met.

[1] Function and performance specifications also usually contain other requirements that can specify other required characteristics of the system, such as constraints.

A range example is :
The regulator shall maintain the supply voltage between 10.0 V DC and 10.8 V DC.

A threshold example is:
The hoist shall lift not less than 5 tonnes.

There are no shades of grey in whether a system complies with these requirements – it either passes or fails. If the hoist can only lift 4.9 tonnes it does not meet the requirement.

The consequences of not meeting a performance level can be severe, as exemplified in a tender evaluation for a mobile crane used to unload containers from a ship at a temporary dock. An essential[2] requirement was that the crane lift 20 tonnes at a reach of 9.5m. One tenderer's crane lifted only 19.5 tonnes at a 9.5 m reach, and so that tender was immediately eliminated from the tender evaluation. This left two tenders, one of which was so poorly prepared that it also was also eliminated . So one performance requirement led to having a choice of just one crane – a fairly severe consequence. The reality was that a crane of 19.5 tonne capacity would provide almost as much value to the user as one with a capacity of 20 tonnes.

A mobile crane

Optimising the level

Setting performance levels is often difficult, or prone to errors. Choosing too high a performance level can increase system cost, result in bias against acceptable system solutions, or cause contractual disputes where a designed system doesn't meet specified requirements.
Setting too low a performance level can yield a system that doesn't meet user's needs or result in increased operating costs. In the tendering phase, setting the performance level bar too low may reduce the ability to discriminate between different solutions, as all solutions may be rated equally regardless of how well they perform.

Asking the user

During the process of gathering requirements, system users or operators will often be asked to set required performance levels. While getting user input is an important means of determining requirements, as well as getting user buy-in, users are usually not the ones who have to pay for the system. This means that there is a tendency for users to ask for higher performance levels than they actually need, and allocate the highest priority level to the majority of requirements.

[2] An essential requirement was defined in the specification to mean that if the system does not meet an essential requirement then it is excluded from further consideration. Ie it is Critical.

Without a means of differentiating between user wants and needs, specified performance levels can end up being much higher than optimum.

Other mistakes

Some of the common causes of setting the wrong performance level are:

- Specifier doesn't know what performance level is appropriate or feasible, so may use a best guess.
- Incorrect assumptions are made in determining levels.
- Incorrect analysis performed in determining the levels.
- Ordinary mistakes made – for example incorrect unit conversion.

Tender and contract process issues

Limitations of step scoring systems

Quantitative tender evaluations will usually involve setting some kind of scoring system that rates tenderer's responses to the requirements based on their stated compliance. Scoring systems that use step values, will yield potentially large variations in scores based on small differences in performance. Figure 1 illustrates a step scoring system for the example of the crane specified to lift 20 tonnes. A crane with a capacity of less than 20 t scores zero, and a crane with a capacity of 20 tonnes or more scores 1.

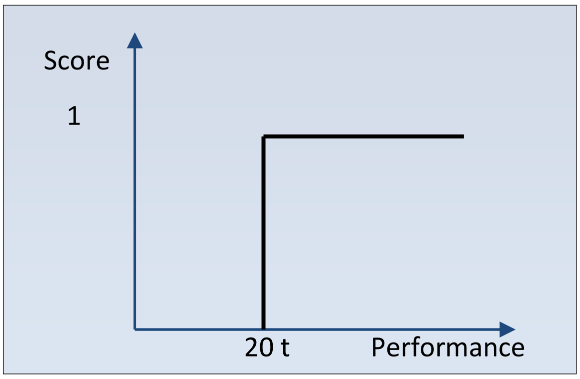

Figure 1: Step scoring for crane capacity of 20 t.

So requirements with a discrete performance level can result in tenderers offering systems with performance that may be well beyond the specified performance level if a smaller system does not quite meet the specified level. For example the crane manufacturer may have a crane that can lift 19 t, but may need to offer their larger model that can lift 28 t in order to meet the required 20 t performance level.

Limitation of requirement value scoring

Requirements are typically assigned some kind of priority to indicate their importance to the customer. A three level system is commonly used, with levels of 1, 2 or 3 (Hooks, 1994, p73) or Essential, Important and Desirable (Australian Department of Defence 2004).The US Department of Defence (US DoD 2012) use a two level system of Threshold and Objective for key performance parameters.

Although having just a few priority levels is easy and practical, it does not allow much distinction between the importance of requirements. A specification for the mobile crane example could quite feasibly have the requirement to lift a 25 tonne load and to have a tiltable seat at the same priority level, as the user rated both as Important. So a tenderer offering a tiltable seat but not a 25 tonne capacity, could be rated equally with another tenderer that offered 25 tonne capacity but no tiltable seat. The reality is that the user would rate the 25 tonne load capacity much more highly that the tiltable seat, and the cost of providing the 25 load capacity would also be much higher than providing a tiltable seat.

Contract disputes when performance levels not met

When a contractor does not meet a specified performance level, disputes often result. This can waste time and project resources to resolve, necessitating the need to raise contract and engineering change proposals, and resulting in project delays. There are various reasons why a contractor may not meet required performance levels, including:

- Specified performance levels were infeasible or difficult to achieve, and the contractor did not know or check this when stating compliance.
- Requirement was unclear or ambiguous.
- Contractor knew they wouldn't comply, but stated compliance to help win the tender, with the aim of dealing with non-compliances when in contract.

Differentiating between Criticality and Priority

The *priority* of a requirement reflects its value to the customer. If a system must meet a particular requirement or the customer will consider the whole system useless, then that requirement is a *critical* requirement.

In Australian Defence acquisition processes three, or sometimes four, levels of priority are used. For a three level system these are Desirable, Important, and Essential. Very Important is added in a four level system. Requirements with the Essential attribute have the highest priority level, but are also considered as critical. If the system doesn't meet an Essential requirement, then it is normally excluded from further consideration in a tender evaluation.

Only a limited number of requirements should be assigned the Essential level, so as not to eliminate potential tenders that the customer would actually consider acceptable. For most requirements, therefore, the number of usable priority levels is reduced to 2 or 3 in the Australian Defence acquisition process.

A critical performance level will not apply to every performance requirement – only to those where falling below a performance level will render the system useless. A function that is not critical should not generally have an associated performance level that is critical. For example, if the crane has a non-critical requirement for an adjustable seat, then there should be no critical performance level for the range of adjustability.

Some solutions

This section will consider various ways of determining and specifying performance levels, including existing techniques and also some new ideas describing better methods. For some requirements, the performance level will be driven by a defined function or need, or else from a higher level document. This situation is discussed in the last sub-section below. For other requirements performance level does not trace to a defined function or need. Instead the benefit to the user will increase as performance level increases, but there is no specific need for a particular performance level. As performance increases, the user may be able to do things faster, use less resources, achieve higher quality and so on. For example for a database system, a user may want the best enquiry response time possible, but there may not be a need to achieve a particular time.

Tiered performance requirements

A simple system that provides an improved ability to specify and evaluate performance levels is including a series of requirements with different performance levels and different priorities.

For example:

Requirement	Priority
The crane shall lift not less than 28 tonnes.	low
The crane shall lift not less than 25 tonnes.	medium
The crane shall lift not less than 20 tonnes.	high

The user would like a capacity of 28 tonnes, but this does not have as high a priority as the minimum crane capacity of 20 tonnes.

Each requirement still has a pass or fail outcome, but by including more than one level there is some ability to communicate that higher performance levels have value to the user.

Given that priority rating systems often only have a total of three levels of importance, the significant disadvantage of using this tiered approach is that the entire spectrum of importance levels is used to show the range of importance for one performance requirement. The ability to lift 28 tonnes is now at the same priority as a number of other low priority requirements, such as having cup holders in the cab. Even though cup holders and having a 28 tonne lift capacity will have a low priority, the customer will rate the crane capacity much more than the cup holders, and there will also be a big difference in the cost of meeting the crane capacity requirement than the cup holders. The only way around this problem is to introduce more priority levels.

Cost and value versus performance

While having multiple performance level steps using tiered requirements is an improvement on specifying a single level, can we go even further? The answer is, in most cases, yes. A more sophisticated solution is a graph which shows how benefit to the user changes with performance level.

A technique sometimes used in specifications called a value function, where the value to the user is plotted against the performance (Hull et al, p83-84). A similar concept in use by the US Department of Defence is the utility curve, which can be used in tradeoff analysis (US DoD Systems Management College, p115 to 116). An example is shown in Figure 2 below:

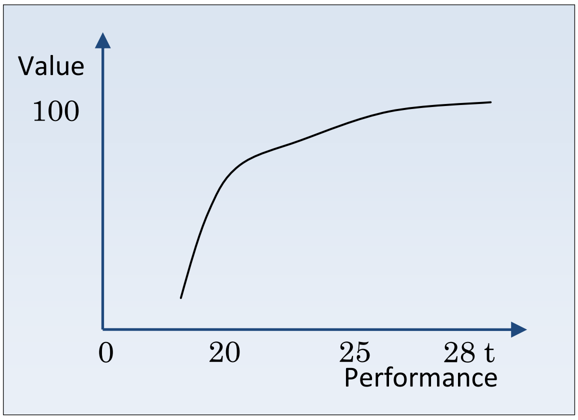

Figure 2: Value function showing crane lifting performance vs value

An equivalent method was used in older versions[3] of the DOORS requirements management software in their tender management module, which allows creation of a normalisation function, which converts a performance score into a normalised score of between 0 and 100. The software allowed creation of normalization functions which can be steps, curves or linear. Weightings were then applied to each requirement, and so a tendered performance level can be converted to a requirement score.

There are no units in the value measurement in utility curves, or in the techniques outlined by Hull and included in the DOORS software. Therefore the value of achieving a level of performance can only be compared with the value of a different performance level, or against a performance level of a different requirement. Even the meaning of the value score is not defined. Does a value of zero mean that the system has no value against this requirement, or does a zero just represent a minimum performance level?

The QUPER model

One system that uses the concept of relating value to performance, and cost to performance is the QUPER (QUality PERformance) model proposed by Regnel, Host and Svensson (2007). In this model, *value* is called *benefit* and *performance* is called *quality level*.
The model is targeted toward requirements for consumer products, being originally applied to mobile phones. It also covers creating development roadmaps, which is not discussed here.

[3] Current versions of DOORS unfortunately do not include the tender management software.

Figure 3 and Figure 4 show an example benefit view and cost view.

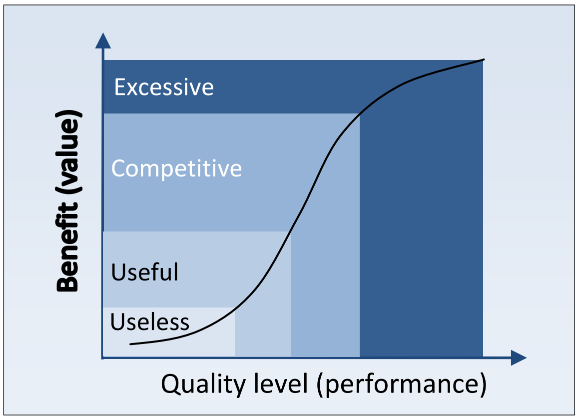

Figure 3: QUPER Model benefit view.

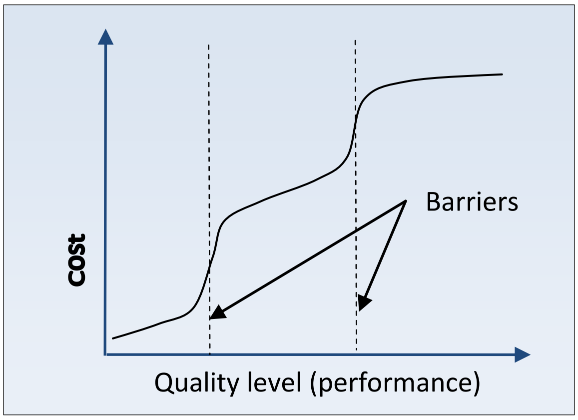

Figure 4: QUPER Model cost view.

In the benefit view, as performance against a requirement increases, the benefit against the requirement increases from useless, to useful, and then to competitive, which is the most desirable level. Then as quality increases further, the benefit will increase, but will go into an area described as excessive – ie the additional benefit to the user from the additional quality is small.
In the cost view, as performance level increases, so does cost. At certain points, called barriers, the cost increases sharply. This is because certain technological barriers may need to be crossed, entailing higher cost.

Extending value versus performance graphs with money

The graphs of value and cost versus performance presented so far are useful, but ultimately they are limited because they lack a unit of measurement.
The value versus performance graph can, however, become much more powerful if the benefit or value performance level is quantified as a dollar figure. By putting a dollar figure on the benefit or value, it becomes possible to optimise the performance against the cost. To do this we can use what will be called a "Relative Value-Cost Performance Graph". As the title suggests, this is a graph of relative value and relative cost against performance. This name is something of a mouthful, so it will be shortened to Revacope graph.

The first step in producing the Revacope Graph is to produce a relative value curve, which shows how the customer values different levels of performance in dollar terms. Figure 5 shows an example relative value curve for the mobile crane example. The customer has valued a lifting capacity of 28 tonnes at $200k more the value of the 20 tonne capacity, rising linearly from 20 tonnes. Above 28 tonnes there is no additional value to the customer, so lifting performance above 28 tonnes has a relative value of $200k. Although the relative value of a 20 tonne capacity is set at zero, this does not mean that this level of performance has no value, just that this is a baseline level in a relative scale. This will be clarified further shortly.

Valuing different levels of performance can be challenging. Appendix A provides some examples of how relative values can be determined.

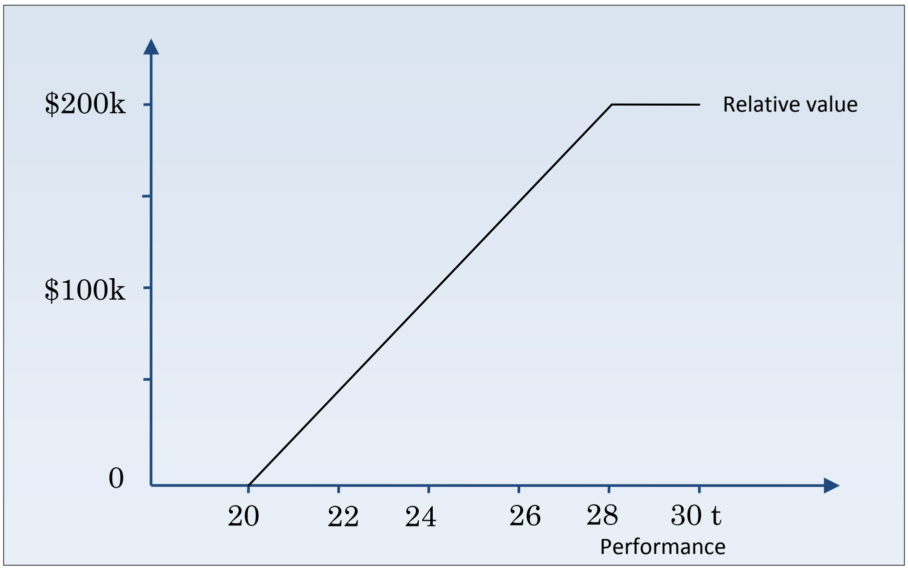

Figure 5: Relative Value Curve for Crane Example

The next step is to produce a relative cost curve which shows the cost of achieving different levels of performance. In this example the crane supplier can supply two models of crane. One has a 22t capacity while the other has a 30 t capacity. The 30 t capacity crane costs $120,000 more than the 22 tonne crane. From this information a relative cost curve is drawn (Figure 6). The two crane models are indicated by crosses on the graph. If this is drawn as a curve there will be a step in the curve as the performance level exceeds 22 tonnes. The relative cost is baselined at $0 below 22 tonnes (representing the 22 tonne crane, which satisfies performance below 22 tonnes). The relative cost of increasing performance level from 22 to 30 tonnes is $120k (the additional cost of the 30 tonne crane over the 22 tonne crane). The 30 tonne crane satisfies performance from 22 to 30 tonnes.

Once again, we are talking about relative costs, so the 22 tonne crane does not cost $0 – this is just a baseline cost against which higher levels of performance are compared.

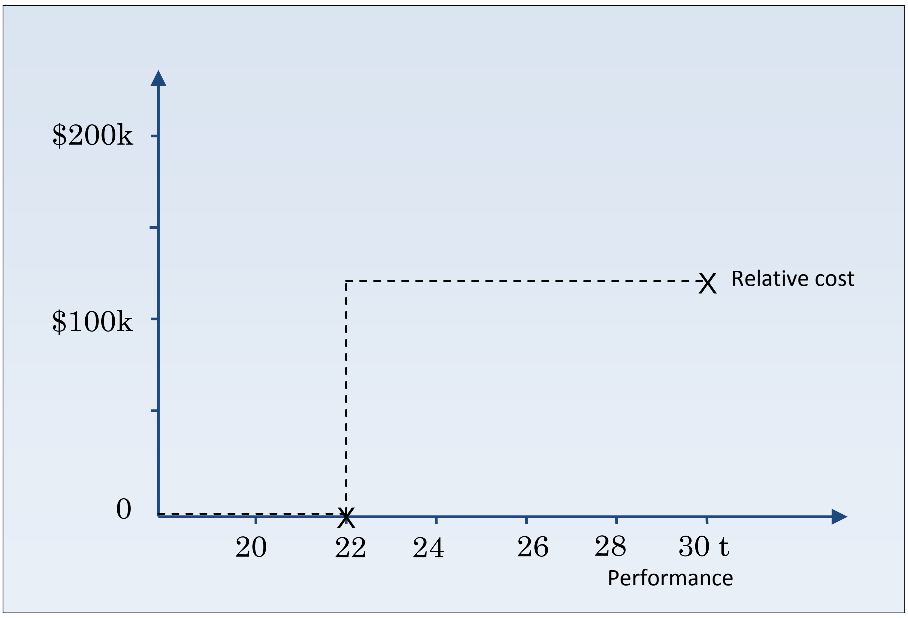

Figure 6: Relative Cost Curve for Crane Example

To optimise the choice of crane performance, the two curves are superimposed to produce the Revacope graph (Figure 7). The optimum performance level is at the point where relative value exceeds relative cost by the largest amount (the Δ (value – cost) amount). This is where the relative value curve is above the relative cost curve by the largest amount. In this example the optimum occurs at a performance level between 28 and 30 tonnes, and so the 30 tonne capacity crane will provide the optimum performance level.

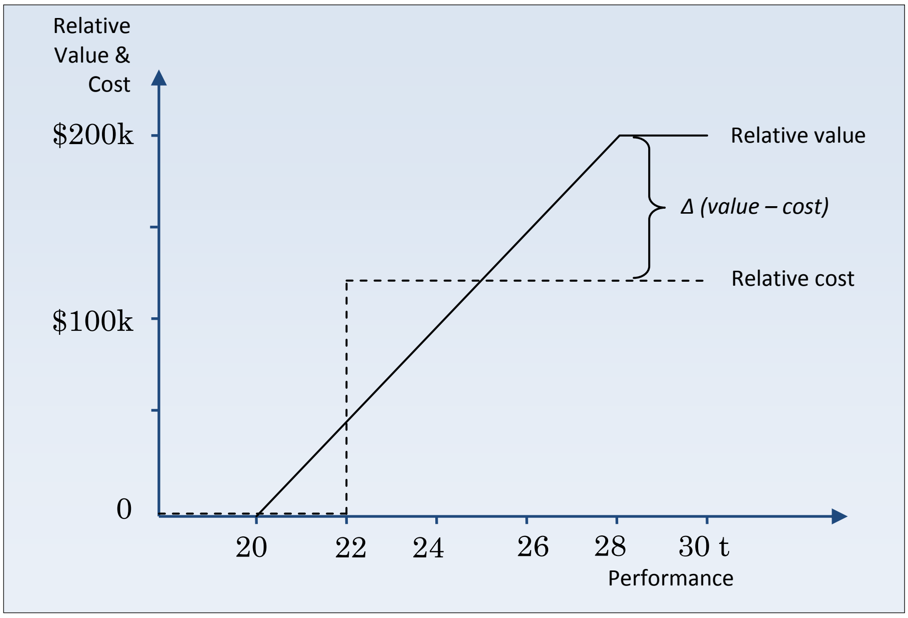

Figure 7: Relative Value-Cost Performance (Revacope) Graph for Crane Example

As mentioned previously the Revacope Graph is all about *relative* cost or value. Depending on the baseline values ascribed to the relative cost or value, the curves will be shifted up or down, but it is the maximum difference between them that will provide the optimised performance level. The cost curve could be higher than the value curve for all levels of performance, as shown in Figure 8, so (relative value – relative cost) will always be negative. The optimum performance level will still be between 28 and 30 tonnes, where the magnitude of the negative number is smallest.

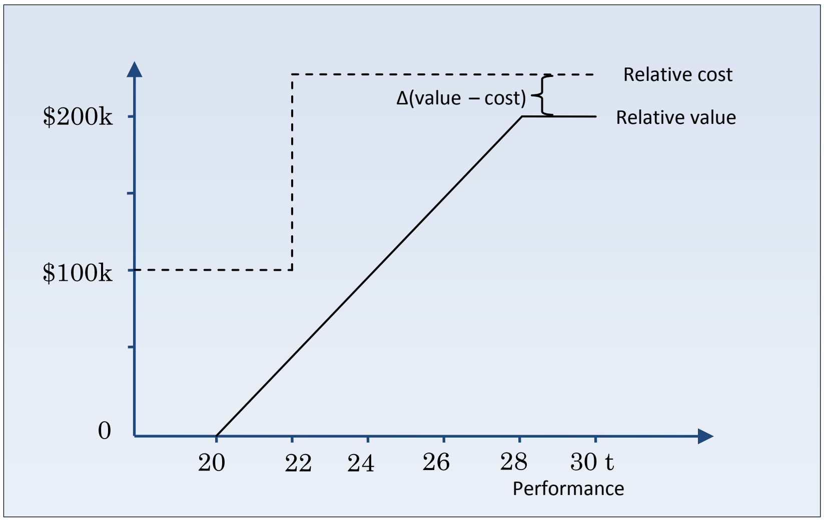

Figure 8: Revacope Graph for Crane Example – cost curve shifted up

Using the Relative Value-Cost Performance Graph to specify performance

The Revacope Graph can be used in various ways in a contractual situation. The ideal way of using the Revacope Graph is for the customer to determine the relative value curve and then include this with the specification, without specifying a level of performance. Tenderers will then produce the relative cost curve and from this select the performance level. Each tenderer will then select and offer an optimum performance level based on their cost of achieving different levels of performance.

Although the idea of having the tenderer select the performance level may seem to be a somewhat radical, this is really just an evolution of the existing tiered performance approach discussed previously. What some organisations may be uncomfortable with being completely open about how it values performance in dollar terms, fearing that tendered prices may be inflated if the supplier knows the customer's bargaining position.

If a customer organisation does not want to just include a Relative Value curve in the specification, then they can estimate the suppliers' relative cost curve and from this choose and specify a performance level. The more accurately the customer can estimate the suppliers' relative cost curve, the more optimum will be the performance level chosen. The disadvantage of this approach is that only one level can be specified, so the performance cannot be optimised for different suppliers, who will most likely have different costs to achieve performance levels.

Once the customer reaches the tender evaluation process, they can use the Relative Value graph to help compare tenders, as the value of the key performance requirements has been clearly defined in monetary terms.

A compensated cost can be calculated, where the price is compensated for by the relative value of the offer. Using the crane example, Offer 2 is $150k more expensive than Offer 1, but has an additional $150k of value to the customer in terms of its additional lifting capacity.

	Offer 1	**Offer 2**
Price	$450k	$600k
Capacity	22t	30t
Relative value (of lifting capacity).	$50k	$200k
Compensated cost (price – relative value)	$400k	$400k

The tender evaluation can then consider which offer is superior based on compliance against requirements other than lifting capacity.

Using the Relative Value curve where performance is uncertain

Where a system is being acquired under a developmental process, the supplier may not know exactly what they can achieve for some key performance levels. If contracting arrangements allow, then a contract can include a variable price, bonus or penalty, depending on what performance level is met. The price, bonus or penalty can be determined from the Relative Value curve.
Alternatively a customer can include a liquidated damages clause into the contract to deal with the situation where the supplier does not achieve their stated performance levels. These damages can be based directly on the Relative Value curve. In the crane example, if the tenderer offered a crane with a capacity of 28 tonnes, but actually only achieved a 22 tonne capacity during testing, then liquidated damages of $150k would apply.

The difficulty of puttting money on performance

Placing monetary values on performance levels is not easy, and it is virtually guaranteed that there will be many different viewpoints on the shape of the relative value vs performance curve. Despite any perceived or real inaccuracies in the curve produced, the performance level determined by use of the Revacope graph is likely to be much closer to optimum than those chosen by the many other more imperfect techniques available.
Some examples of how monetary value can be assigned to performance are given in Appendix A.

Defining a critical performance level

As discussed on p 4, for some requirements there may be a critical level of performance below which the user may determine that the system is useless. The Revacope graph will show relative value minus cost falling as this critical performance level is approached, but this will not automatically identify the system as useless.
Any critical performance levels need to be clearly identified in a specification via separate requirements that specify critical performance levels.

Interrelated performance requirements

The optimisation of performance against cost has so far assumed that performance requirements are independent, which is not always the case. Having a high level of performance in one requirement may influence the ability to achieve a high level of performance in another requirement. For example specifying a ballistic helmet have a light weight will make it more difficult to achieve high levels of ballistic performance. If the Revacope system is used and the customer just defines the relative value curves for the interrelated performance requirements, then the supplier can offer a system where interrelated performance levels are optimised to provide maximum value. If the customer decides to specify specific performance levels, then the customer will need to consider dependant performance levels.

A sophisticated technique reserved for key performance requirements.

It may not be easy for users or stakeholders to determine a curve showing value vs performance. Getting multiple users to agree to such a curve could severely test a requirement gatherer's patience. So going through such an exercise is something that would usually be reserved for key performance requirements that significantly affect the value provided by the capability or may significantly affect the cost of the system. Table 1 shows some examples of key performance requirements:

System	Key performance requirements
Mobile crane	Lifting capacity, throughput (containers unloaded per hour)
Baggage handling system	Bags per hour, number of personnel required to operate, reliability.
IT network	Max no. of users, bandwidth,
Aviation refuelling tanker	Fuel pumping rate, fuel capacity, power to weight ratio of truck.

Table 1: Examples of key performance requirements for a variety of systems

Does the requirement link to a defined function or need?

Some performance requirements demand a specific level of performance as they link to a defined function or need. Defined functions or needs could be:

- a specific function that the system needs to do (eg cater for a group of 50 existing users, lift an object that has a particular weight), or
- ability to interface with an external system (eg a system in a production line has to achieve a particular cycle time to keep up with the production line, otherwise it is slowed down).

In this case setting and specifying the performance level is fairly straight-forward. A single level can be chosen, and the importance assigned will be driven by the consequence of not meeting the level.

In larger or more complex systems, different levels of specifications may be used. The highest level specification will list high level performance goals, which may be called measures of effectiveness (how well the system performs its mission) and measures of suitability (how well the system performs in its intended environment, including supportability, maintainability and ease of use).
At the next specification level down, the higher level performance requirement should become more specific and detailed. If the lower level specifications are for subsystems of

a larger system, then the higher level performance requirements are allocated to the subsystems.

For example, for the high level requirement for a battery drill to operate for 15 minutes at full load, we could have the following functional allocation:

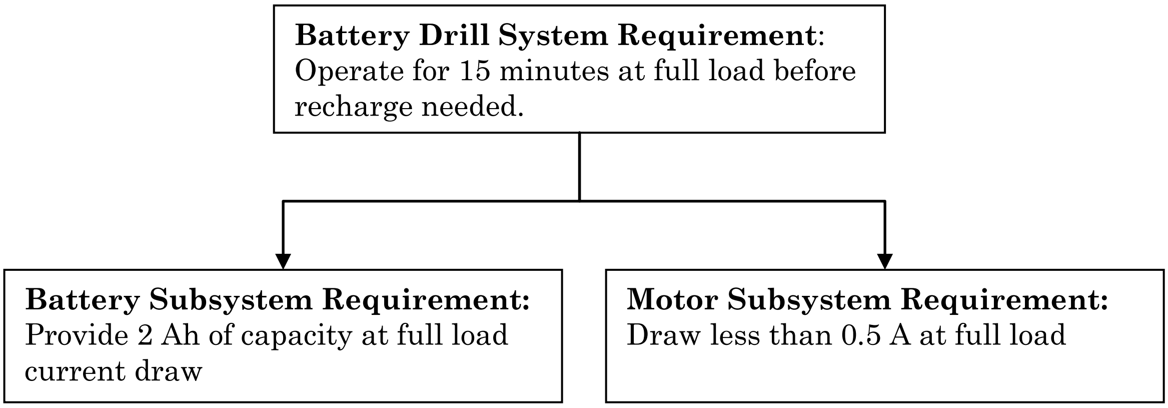

Figure 9: Example Performance Requirement Allocation

More Suggestions on Setting Performance Levels

The Importance of Staying feasible

One of the criteria for performance requirements is that they be feasible. Feasibility can be economic and technical. If the Revacope graph concept outlined above is used, then the performance requirement should be economically feasible and should even be close to optimum if good judgement or analysis is used.
As the performance level increases, the boundary of what is technologically feasible will eventually be reached. Going past this will need new scientific discoveries, or may violate the laws of physics. If a specification writer specifies a level of performance past what is technically feasible, significant problems can arise during tendering or contract. A technically unfeasible requirement identified during tendering can disrupt the tender process as amendments are made. If an unfeasible requirement is identified in-contract, disputes may arise, and the contract may need to be renegotiated to deal with the problem. Again costs in time and money will most likely be incurred.
The boundary of technology feasibility may be difficult to anticipate, and so increasing the required performance level will not only involve increased cost, but also increased risk. Choosing a high performance level needs to consider the question: “If the system cannot achieve this level of performance, what is the consequence for the project?” If the consequences are severe, then it may be best to specify a lower level of performance that is still acceptable to the customer.

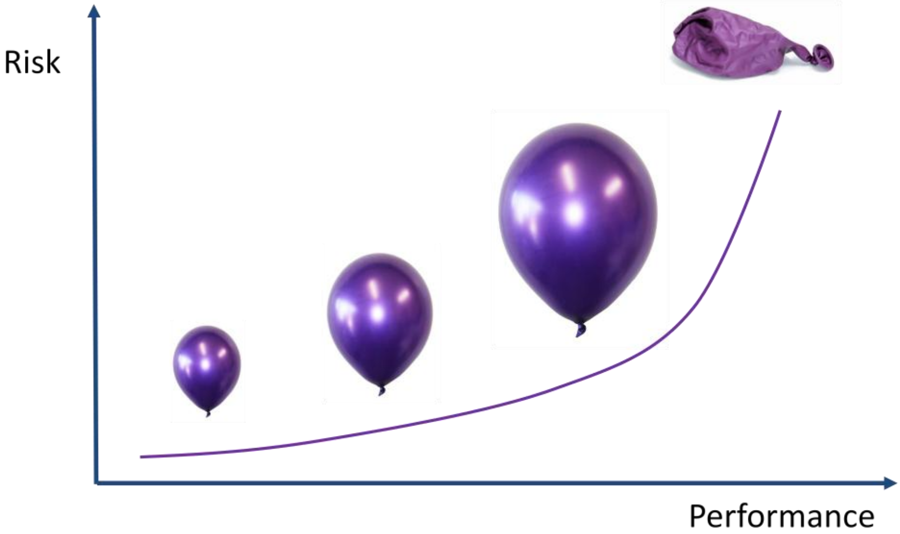

Figure 10: Risk Increases with Required Performance

For off-the-shelf systems, a market survey of available performance is a good reality check on what is feasible. There is little point in specifying performance levels above what available systems can achieve.

Expert Help for the Specification Writer

Understanding what the cost vs performance curve is, and identifying the limit of feasible performance may need specialist expertise that a specification writer does not have. If the customer organisation's acquisition policies allow, potential tenderers are a good source of expert knowledge, although there may be an element of bias in advice given by a tenderer to advantage their bid. Where seeking advice from a tenderer is not feasible, employing an external independent expert is an alternative. Companies may baulk at the cost of employing a specialist, but for important requirements, the cost and time saved through having the correct performance level can be significant. Advantages include reduced system cost, systems that better meet user needs, and reduced contractual problems. A suitable expert can be used to help review a specification and provide advice on aspects of the requirements that need specialist knowledge.

Keeping end users happy with performance improvements

If there is a guaranteed way of upsetting end users (those who actually use the system), it is by providing them a system that is inferior to an existing system. Even one key feature that is less good than the existing system can be enough. Upset users may refuse to use the system, try to break it, continually make complaints, or exercise whatever means they have available to express their discontent. Melbourne's Myki ticketing system for public transport is a good example of a system where, for the users, it has many features that are inferior to the existing system that uses cards with magnetic strips. The Myki system requires train users to swipe their card (with an RF ID chip) on a reader when entering a station platform and also when leaving a station, while the traditional card system did not require users to do anything when exiting most stations[4]. The Myki card suffers from periodic problems where the reader will not quickly read a card, causing other train travelers to bank up behind Myki users struggling to get their card to read. As well as these problems, Myki users have also commonly been charged incorrect fares and this has led to very poor acceptance of the new system.

So if a system replaces an existing system (which is quite common) using the performance of the existing system as a minimum level is a good starting point. If a performance improvement can be provided, then this will be a major aid to increasing user acceptance of a new system. If performance improvements are specified, then they should be large enough that users will clearly notice the improvement. An improvement of 10% in a performance metric may not be noticeable by users.

Feasible or cost effective performance levels aren't known

For developmental systems, or even some off-the-shelf systems, sometimes the cost or achievability of a performance level may not be known, and can't be known without significant effort, time or investment. Even if the user can determine a Relative Value curve, the system builder may not be able to determine a Relative Cost curve, and therefore the optimum performance level can't be chosen. Following are some approaches that can be used in this situation.

Best endeavours

In a contractual situation, where the performance level is not known, one approach is to use a "best endeavours" or "reasonable endeavours" requirement. An example requirement could be:

> *The Contractor shall apply best endeavours to achieve an average website search time of less than 0.5 seconds.*

In this case a target performance level is set, but the "best endeavours" qualification means that achieving this target is not mandatory. The contractor just needs to apply their best efforts to achieve the target.

As noted by Hayford (2005), the obligation placed on a contractor to achieve the target is not well defined, and even though "best" implies a high standard of effort, Australian courts have not interpreted this as a "leave no stone unturned" standard. One high court judgement held that the contractor was "*required to do all he reasonably can in the circumstances to achieve the contractual object but no more*".

If using the best or reasonable endeavours approach, the degree of effort required by the contractor should ideally be defined so that there is a shared understanding of the meaning of the requirement.

[4] Users were required to enter their card to exit stations in city stations.

Specifying a Baseline Performance Level Plus Negotiated Increase

Another possible approach where the performance level is not known is to include a performance level that is known to be achievable, with a cost to be negotiated in-contract for an increase in performance up to the preferred level. This approach was used in a recent military vehicle project. A baseline fording level[5] was included in the contract, with a higher preferred level to be quoted by the contractor at a later date. The contractor developed and tested a solution to achieve the higher fording level, and then quoted a price per vehicle to implement the solution. The customer was entitled to accept or reject the increase in fording depth.

This arrangement is effectively a pre-planned contract change, and like any contract change proposal, there is a risk that the contractor may significantly inflate the price, given that there is no competition.

Summary

This paper discussed existing techniques for analyzing and specifying requirement performance levels in function and performance specifications. While existing techniques do work, they often will not result in selection of an optimum performance level, and have a number of shortcomings in terms of their use in tendering and contractual situations.

A new technique – the Relative Value Cost Performance graph (or Revacope graph) was explained – allowing an optimum performance level to be chosen by considering the value in dollars to the customer and the cost for a system to achieve varying performance levels. This technique typically can involve significant effort and is recommended for key performance requirements.

A number of general suggestions were made on how to set performance levels, including situations where neither the customer or supplier know what level is feasible.

References

Australian Department of Defence (2004), Defence Material Organisation, *Function and Performance Specification (FPS) Development Guide* V 1.0.

Hayford O and McBride, V. (2005) *Best v reasonable endeavours: Is there any difference?* Clayton Utz Lawyers website. http://www.claytonutz.com/publications/newsletters/projects_insights/20050505/best_v_reasonable_endeavours_is_there_any_difference.page Accessed 23/6/11

Hull E, Jackson K and Dick J (2011). *Requirements Engineering* 3rd Ed. London: Springer-Verlag VALUE OF A REQUIREMENT??

Hooks, Ivy (1994) *Guide for Managing and Writing Requirements* Bourne: Compliance Automation, Inc

Regnell, B., Höst, M. & Svensson, R. (2007). *A Quality Performance Model for Cost-Benefit Analysis of Non-Functional Requirements Applied to the Mobile Handset Domain.* Proceedings of the Requirements Engineering Conference: Foundation for Software Quality, Trondheim, Norway, 277–291.

US Department of Defense Systems Management College (2001). *Systems Engineering Fundamentals*. Defense Acquisition University Press. Fort Belvoir, Virginia

[5] Fording level is the depth of water that the vehicle can drive through without stalling.

US Department of Defense (2012) *Defence Acquisition Guidebook*

Acknowledgements

The author thanks Robert Pace for reviewing this paper and providing useful feedback.

Thanks also to the following people in the LinkedIn Requirements Engineering group, who provided ideas for this paper: Mark van Holsteijn, Daniel Lucas-Hirtz, Robert van Lieshout.

Biography

Michael Addis holds a BE(Mech)Hons from the University of Melbourne, and an MBA(Tgy Mgt) from APESMA/Deakin University. He is a Principal Consultant with Codarra Advanced Systems, and specialises in specification writing within the Systems Engineering consulting area.

His experience and interest in specification writing and system acquisition comes from his work in introducing new manufacturing systems, and more recently in Defence acquisition in the Land environment. Over the last 8 years, Michael has written specifications for land vehicle systems and modules, including light vehicle modules, a mobile crane, container side loader, C4I integration, aviation refuelling vehicles, military packs and combat helmets. He has also reviewed specifications covering military vehicles, vehicle modules, C4I integration, UAVs, weapon systems, forklifts and packaging systems.

Codarra
Advanced Systems

Appendix A – Some Examples of How to Determine the Relative Value Curve.

Introduction

Determining what value different levels of performance have will often be difficult and may not even seem possible, however where large monetary values are at stake there will be significant benefits in understanding how these values relates to levels of performance. Values are most likely to be delivered on an annual basis, and these annual values need to be converted into present value through techniques such as calculation of net present value (NPV).
Despite the effort needed to quantify the value of different levels of performance, the result will be not only a relative value curve, but also potentially a deeper understanding about organizational costs and benefits that may reveal potential for improved efficiencies.

Availability of an IT system

The easiest way of determining the value of IT system availability may be to determine the cost if the system is unavailable. Depending on the use of the system, the costs could be:

- Additional labour cost incurred through overtime required to make up lost productivity while the system is unavailable.
- Loss of revenue if a loss of sales or production occurs while the system is unavailable.
- Additional IT support labour cost to rectify the fault.

Let's say these costs amount to a figure of $10,000 per hour and the system is required for 2000 hrs each year. Each 1% loss in availability would therefore equate to a 20 hrs per year, and a cost of $200,000 per year. For this simple example if we consider costs over a 5 year period (rather than using NPV), then we have a figure of $1 million per 1% availability loss. We turn costs into value by considering cost savings. By increasing reliability we save costs and increase relative value. The relative value curve will therefore be as shown in Figure A-1. The minimum availability figure included in the graph is 97%. The choice of minimum performance level shown on the graph should be either the critical performance level (ie the level below which the system would be considered useless), or if there is no critical performance level then the lowest anticipated performance level.

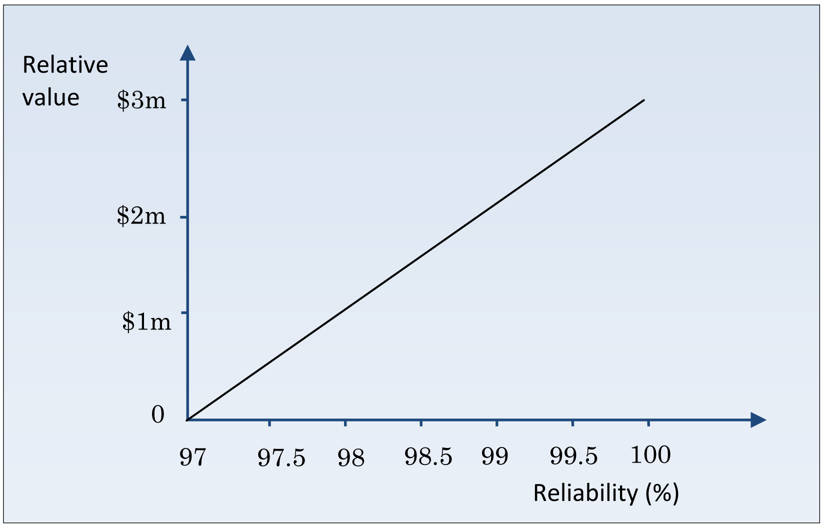

Figure A-1 : Example Relative Value – Availability curve for IT system availability

Capacity of an aviation refuelling tankers

Aviation refuelling tankers ferry fuel from a large storage tank within an airport to aircraft. Having a larger capacity tanker allows more aircraft to be refueled before the tanker must drive back to the storage tank to refill. For aircraft with very large capacity fuel tanks, more than one trip may be required to fill the aircraft. The value of a larger capacity tanker will be an aggregate of the following potential benefits:

- Reduced tanker driver labour cost due to fewer trips back to the storage tank to fill up.
- In the case of a fleet of tankers, fewer tankers may be needed due to reduced trips back to the storage tank.
- Faster turnaround time for aircraft, potentially resulting in reduced aircraft and ground crew costs, better airport space utilisation, and reduced charges.

Baggage Handling System Bags per Hour[6]

Airport baggage handling systems take travelers bags from the check-in point to the aircraft, and on arrival, from the aircraft to a baggage carousel. On the way bags may be security screened by X-ray or other systems. The required number of bags per hour will vary considerably depending on the number of travelers, which will fluctuate depending on the time of year, and also abnormal events that may drive high flight numbers (eg sporting events). Once the number of bags entering the baggage system exceeds the capacity of the system, it is overloaded, and this will probably result in flight delays.
In this example an acceptable maximum percentage of system overloads can be set at 0.5% of system operation time.
Passenger numbers will grow over time, and a higher capacity system will increase the time until overloads exceed the allowable 0.5%. Once this figure is exceeded then the airport will need to upgrade the system to increase capacity by adding a new parallel system. The estimate cost of an upgrade is $2m.

[6] This is an example analysis only, and does not represent an actual value calculation for this system.

A higher capacity baggage handling system will offer value in deferring the time at which an upgrade will be required. The table below shows an example set of numbers:

Capacity (Bags / hr)	**Years till upgrade needed**	**Present cost of upgrade (based on 15% discount rate)**	**Relative Value**
2000	3	$1.32 m	$0
2400	5	$0.99 m	$0.33m
3200	8	$0.65 m	$0.67m

The years till upgrade needed is calculated based on estimated passenger growth, and historical data on current bag/hr capacity needed to limit the system overloads to 0.5% of system operation time.
Present cost is based on the present cost of the $2 m estimated to upgrade the system, discounted by 15% per year by the number of years till the upgrade.
If we take a 3 year timeframe till upgrade as the minimum acceptable time, the relative value is calculated from this 3 year baseline figure of $1.32 m. By taking the negative of the cost, we get value (ie a cost saving is equivalent to customer value). So for example a capacity of 3200 bags per hour gives 8 years till upgrade, with a value of $0.9m ($1.32 m – $0.65 m).