Balancing capability and affordability for complex engineering systems

Malcolm Courts, PhD, BSc, C.Eng, FRINA
BAE Systems Maritime

Nigel Osborne, BSc
BAE Systems Maritime

Peter Miles, BSc, MSc, C.Eng, M.IET, FRSA
Catalyze Ltd.

SYNOPSIS

In the design of a major new military platform there is a need to determine a technical solution that is deliverable, meets the customer's requirements for military capability, and is affordable. Traditional 'waterfall' development approaches start with the customer developing requirements with advice from scientists and engineers on what is technically feasible. Engineers and designers then seek to fulfil these requirements by developing practical systems using appropriate technologies. The overall design complexity is managed by breaking down the design into a hierarchy of sub-systems, such as propulsion, combat, etc. The end result is a system definition that can be costed.

This linear process fails in well-known ways. Firstly, the hierarchical top-down decomposition fails unless considerable effort is expended on understanding the interdependencies and constraints on sub systems that are necessary to ensure satisfactory system integration. However the overriding problem now is *affordability*, as budgets shrink at the same time as opportunities grow to spend on increasingly complex technologies. In short, what is technically possible at the top end of capability is almost always not affordable.

In order to address both of these problems a concurrent approach to requirement generation and system synthesis is needed together with a framework to manage communication, understanding and decision making across all work-streams. This paper presents a framework that combines best practise from the fields of decision analysis and engineering systems analysis. In applying and developing this framework we have refined the application of existing techniques and introduced new techniques such as genetic algorithms to allow more rigorous exploration of what is possible.

Authors' Biographies
Malcolm Courts is an Engineering Manager with involvement in concept development and product and design process improvement.
Nigel Osborne holds the position of Engineering Manager responsible for implementation of Option Assessment processes.
Peter Miles designs and facilitates collaborative processes for making strategic decisions in complex situations.

INTRODUCTION & PROBLEM STATEMENT

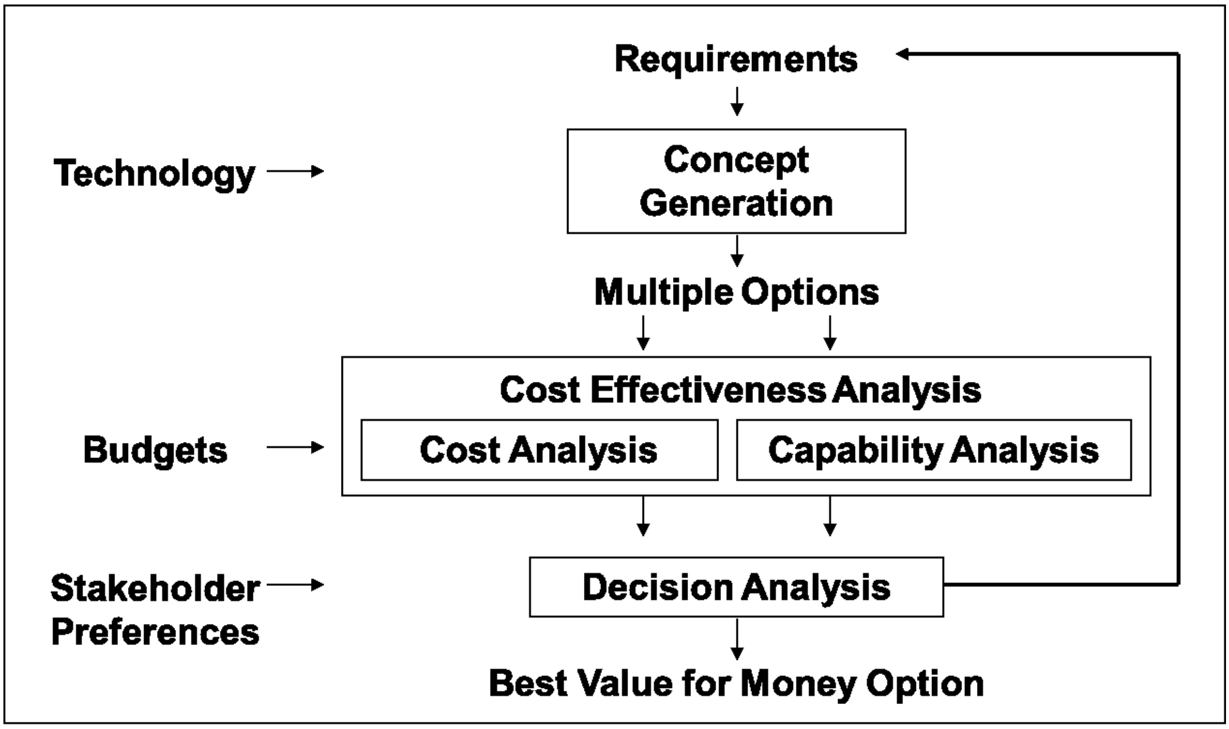

Fig 1. The Overall Justification Process

Figure 1 taken from Courts, Brittain, Lamble and Osborne's RINA 2012 paper on Warship trade Space Exploration: Challenges and Approaches[i] shows the overall conceptual Justification Process employed for a major new military platform. It looks deceptively straightforward; however there is a feedback loop which is at the root of the challenge we describe here, and which is addressed by the framework and process we have used. The arrow from Decision Analysis back to Requirements recognises the significance of *affordability.*

In the design of a major new military platform the goal is an optimised configuration, provided in an agreed timeframe, and to an agreed budget, that minimises risks in technical development and production. Traditional 'waterfall' development approaches start with the customer developing requirements with advice from scientists and engineers on what is technically feasible. Engineers and designers then seek to fulfil these requirements by developing practical systems using appropriate technologies, either off the shelf or developmental. The overall design complexity is managed by breaking down the design into a hierarchy of sub-systems, such as propulsion, combat, etc. The end result is a system definition that can be costed.

This linear process fails in well-known ways. Firstly, the hierarchical top-down decomposition fails unless considerable effort is expended on understanding the interdependencies and constraints on sub systems that are necessary to ensure satisfactory system integration. (This is well recognised and the disciplines of system analysis and integration must be applied at every stage.) However the overriding problem in current times is generally affordability, as budgets shrink at the same time as opportunities grow to spend on increasingly complex technologies. In short, what is technically possible at the top end of capability is almost always not affordable.

In order to address both of these problems a concurrent approach to requirement generation and system synthesis is needed together with a framework to manage communication, understanding and decision making across all work-streams. Best practise from the fields of decision analysis and engineering systems analysis is used to create a 'socio-technical' process which brings together appropriate stakeholders at key points, with software tools providing a bridge between requirements generation and system synthesis. The framework enables effective communications, and unleashes the collaborative creativity that can resolve typical capability/solution/cost balance issues. A combination of top-down and bottom-up processes are employed, with the ability to regularly 'close the loop' between what is being asked for and what is both technically and economically feasible, Figure. 2.

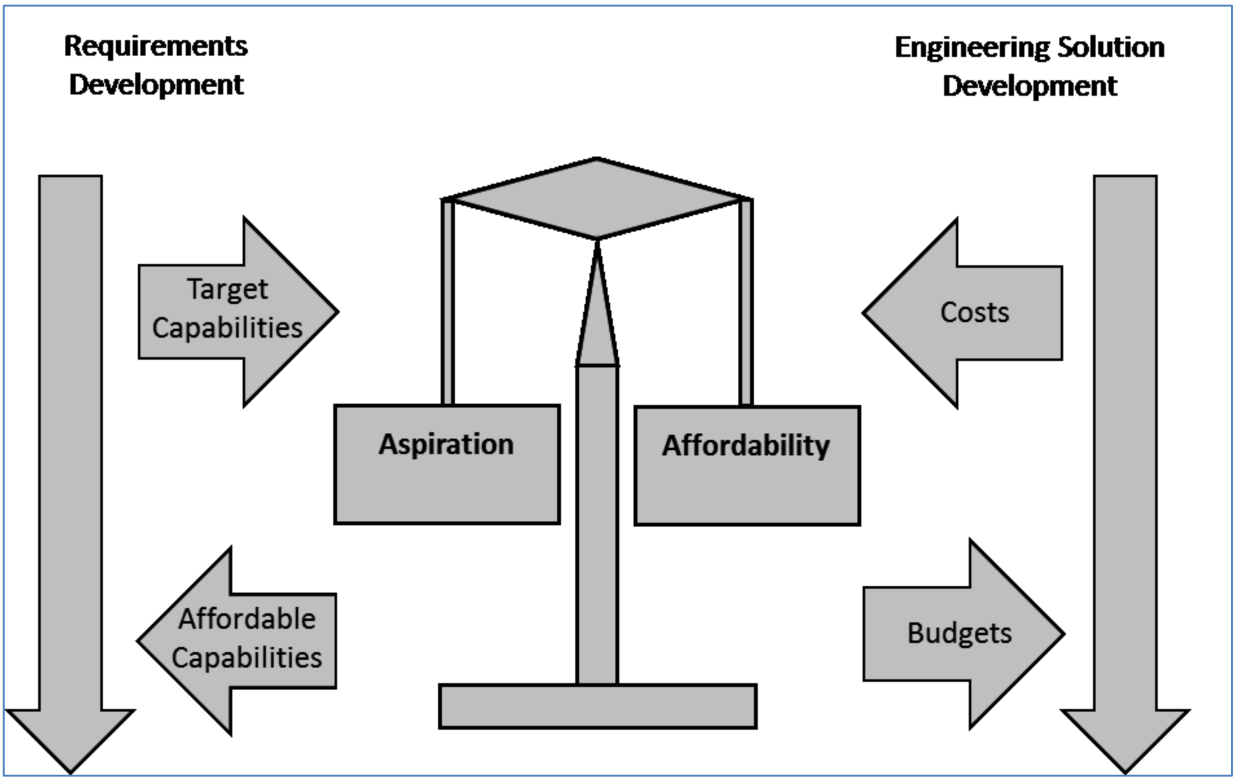

Fig 2. Balancing Aspiration with Affordability

If not for the consideration of affordability, the systems design process would be much easier. However, in the real world where budgets are finite and the 'conspiracy of optimism' is recognised and being forcefully addressed, the constraint of affordability introduces new interdependencies and hence complexity - a more expensive technology in one area of the platform will tighten the cost constraints elsewhere.

In this situation, the requirements themselves must be traded, and a process is required to do this in a fair and manageable way. The customer can (and should) state target capabilities. Ideally then a large number of potential design solutions will be considered, allowing the selection of one which provides both value-for-money (optimised cost-benefit over the whole lifecycle) and is affordable (fits an annualised spend budget).

The process is of course highly technical, but it is also *social,* and a large number of diverse stakeholders need to be actively engaged at appropriate points (not just consulted or informed).

DEVELOPMENT OF DECISION ANALYSIS AND DECISION CONFERENCING

Decision Theory was introduced in 1961 by Howard Raiffa and Robert Schlaifer[ii] of Harvard University's Business School. Raiffa and Ralph Keeney[iii] went on to develop multi-attribute value theory (including utility theory) in 1976. The term Decision Analysis was coined by Ron Howard[iv] of Stanford University in 1964. Through the work of these pioneers and their contemporaries, the mathematics of the field was established.

Since these early days there has been a keen interest from the military sector, with its history of large scale and very expensive development activities which naturally established a high value for good decision making. The natural human tendency to make decisions based on one or two key criteria (such as cost, or firepower) was revealed as inadequate, especially as the key criteria chosen could naturally be highly subjective but needed to be stable for at least the length of the programme development to avoid costly changes in direction. In other words, there is a very high cost of getting decisions wrong in defence platform development, and consequently a high interest in making those decisions as robust to challenge as possible. Multi-attribute decision analysis met this need, and a practical guide was published by Dr Larry Phillips[v] and other authors in the late 1970s.

Decision Conferencing[vi] had been developed in the late 1970s by Dr Cameron Peterson, as a response to the difficulty in conducting a single decision analysis for a problem with multiple stakeholders, each of whom takes a different perspective on the issues. The approach was taken up in 1981 at the London School of Economics' by Phillips, who integrated into the facilitator's role many of the findings about groups from work at the Tavistock Institute of Human Relations.

All of the techniques described above can be applied in the context of options selection (selecting one option from many), and also portfolio optimisation (selecting a combination of options from a wider choice).

These approaches are now broadly established as best-practise for large-scale defence systems engineering projects, and are evolving towards a closer integration with the overall systems design process as described in this paper.

THE PARALLEL ACTIVITY STREAMS

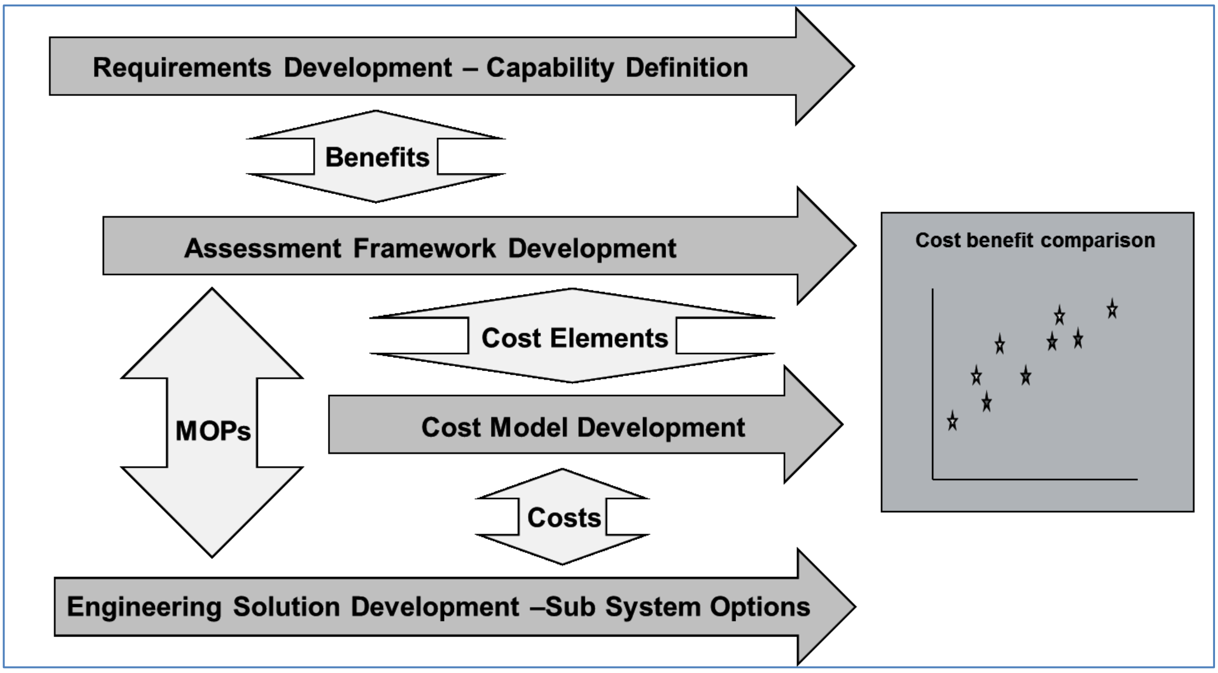

Fig 3. Parallel Activity Streams

Figure 3 illustrates the primary parallel streams of activity, and how they are linked to create a coherent process. At the centre is an ‘Assessment Framework’. The Customer’s capability requirements are communicated into this framework by a ‘Benefits Tree’, a hierarchical statement of key requirements, expressed (e.g.) in terms of military doctrine. This Benefits Tree is carefully weighted (a process that needs to be iterated periodically as the Engineering Solution is developed, and as the ranges of potential capability become better understood).

The Engineering Solution links into the Assessment Framework by means of ‘Measures of Performance’ (MoPs), such as ‘top speed’. The MoPs themselves are the key drivers for high level Measures of Effectiveness (MoEs) such as ‘the ability to intercept an enemy’. The MoEs in turn are direct measures for the capability requirements enshrined in the Benefits Tree. The linkage between MoPs and MoEs or benefits can be considered as Influence Factors and can be obtained by a combination of

- Operational analysis
- Operator judgement and experience

The remaining key parameter is Cost. The cost of the Engineering Solution, and the multiple Options considered for comparison, are communicated to the Assessment Framework via a Cost Model.

The Assessment Framework brings all these key components together in way that enables communication between the diverse set of stakeholders involved in the whole process, including military commanders, equipment users, engineers, operational analysts and cost modellers.

The most convenient output from such a process is a cost benefit plot, as shown on the right of figure 3, showing how system solutions compare across the whole range of potential cost and capability. This can be considered as a Trade Space within which the optimum solution may be found and justified. The prime objective of any Trade Space exploration process must be to push the boundaries of the cost benefit plot in order to maximise System benefit at any given cost level.

PROJECT CONTEXT

In the early Concept Phase of a project the imperative is to establish within the requirements a target general level of system capability that is both realistic and affordable and can be potentially achieved and delivered. In order to do this a wide Global Trade Space survey is undertaken. At this stage as many complete System concepts as possible must be synthesised that deliver capability levels across the whole range of potential cost and benefit. These solutions are often generated with the help of a computer aided concept modelling tool. However they are obtained each of these candidate Systems must be assessed against the evolving requirements using the System Assessment Framework described above.

This Global Trade Space is used to focus the requirements such that they describe a solution space corresponding to a promising region of cost and capability that can then be explored in more detail in an engineering Feasibility Study. This is best performed by selecting one or more promising baseline system solutions and then undertaking a process of sub system design exploration and optimisation to improve the baseline solution in a process known as Local Trade Space Optimisation.

The part of the process concerned with establishing benefits and establishing how they are assessed, is very similar across both the Global and Local studies. The processes start to differ in terms of the alternatives assessed within the framework. In the Global study whole system solution alternatives are assessed, while in the Local study it is sub system alternatives that are considered and then synthesised into optimum solutions. Such a process using sub system synthesis optimisation techniques is described in the rest of the paper as it was applied in a recent military platform design project.

ASSESSING THE OPTIONS – THE 'ORDER OF PRIORITY'

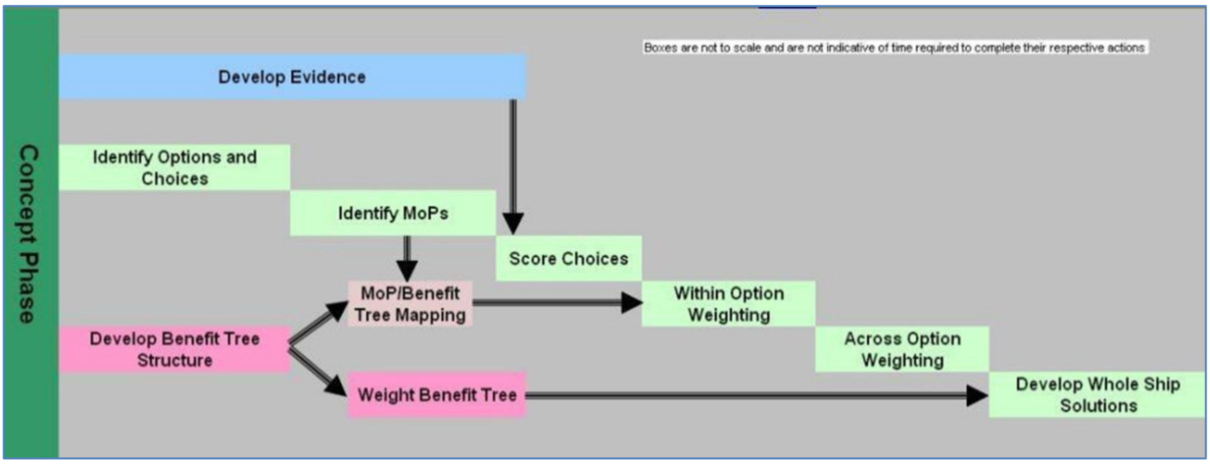

Fig 4. The Options Assessment Process

Figure 4 shows the decision process followed. The Customer requirements for capability were captured in a Benefit Tree, which was itself structured according to high-level requirements such as military doctrine (this was done as a workshop with military commanders). In parallel with that activity, a menu of sub-system Options was developed (the rows in Fig 5), each of which had a number of Choices (the columns). There was then a set of workshops, one for each Option area, to describe the Choices and to identify the MoPs that could be used to differentiate between the Choices. Each of these MoPs was then mapped against the Benefit Tree (if no mapping was found, the implication was that the Customer did not care about this MoP, hence it could be discarded; although in some cases this information was fed back into the process of improving the Benefit Tree). The detailed description of each Choice that was produced by this process then formed the basis for costing activity. Templates were designed to elicit the information required to drive the process, these were populated before, during and subsequent to the workshops.

Data was then collected to inform the MoPs, and an initial Hiview3 [vii] model was built for each Option. Subsequently another round of workshops were held, this time to score each Choice against each MoP, and then to weight the MoP according to how each one contributed to the Benefit Tree (this was called the Within Option Weighting). At each workshop detailed notes were taken to record data sources and the logic behind judgements, including arguments and counter-arguments where appropriate. These meeting notes and the software models together provide detailed transparency and traceability, with a clear audit trail as demanded by the customer.

Subsequently in a 2-day workshop an 'Across Option Weighting' exercise was performed to establish a benefit measure that was consistent across the whole Option set.

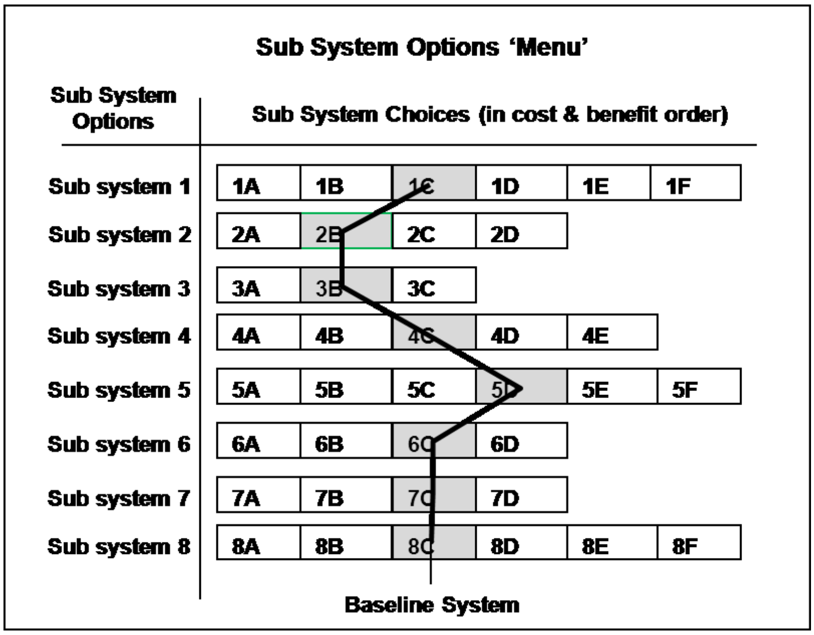

Fig 5. The Options Menu

Figure 5 (reference i) shows the general format of the Options menu. Each row represents a sub-system of the platform (such as the propulsion system), and the columns represent the Choices within that Option. A complete system is represented by one Choice from each Option. Shading indicates the current Reference Design, a particular selection of Choices used as a baseline. In principle, every combination of Choices through the whole set of Options represents a complete configuration of the platform. Of course some of these will make more sense than others, and there will be a wide variation in overall cost. In general the Choices in column A are the lowest cost, with cost increasing as we move toward the right of the menu.

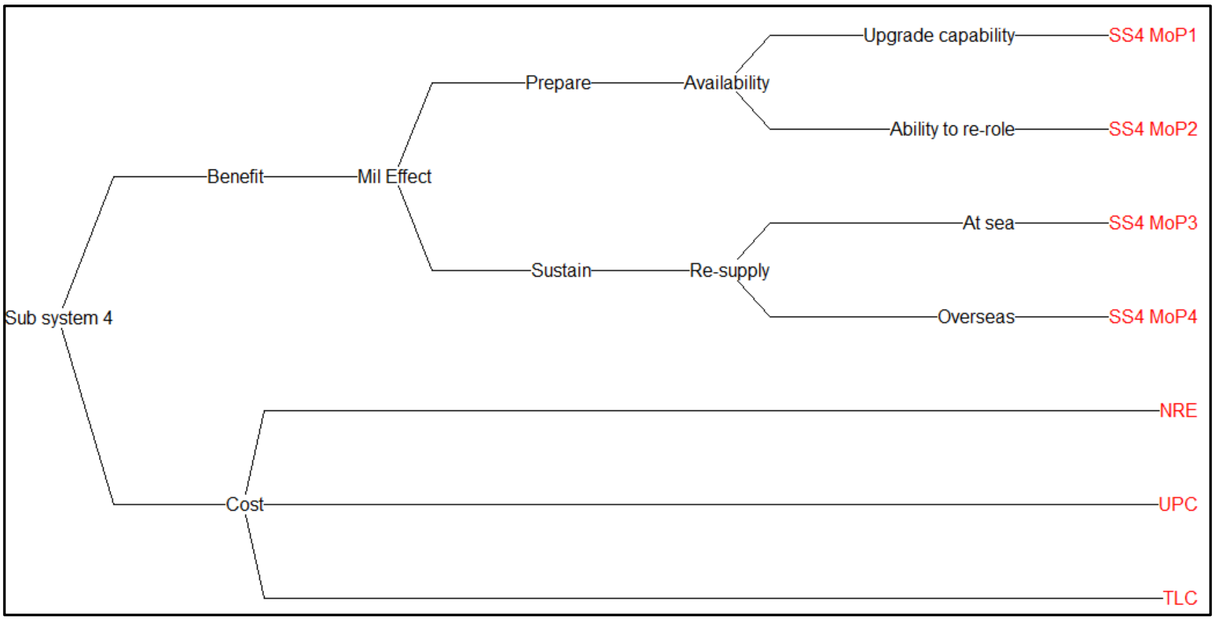

Fig 6. Measures of Performance Mapped to the Benefit Tree for a Single Option

Figure 6 shows the subset of the Benefit Tree that is mapped to Measures of Performance for one particular Option. In this example there are 4 MoPs. Note that Cost is captured in this same model; in this case it holds Non-Recurring Expenditure (NRE), Unit Purchase Cost (UPC), and Total Life Costs (TLC). Given this Value Tree we have enough information to model the Order of Priority for an Option, ie the ranked order of cost and benefit for the 5 different Choices.

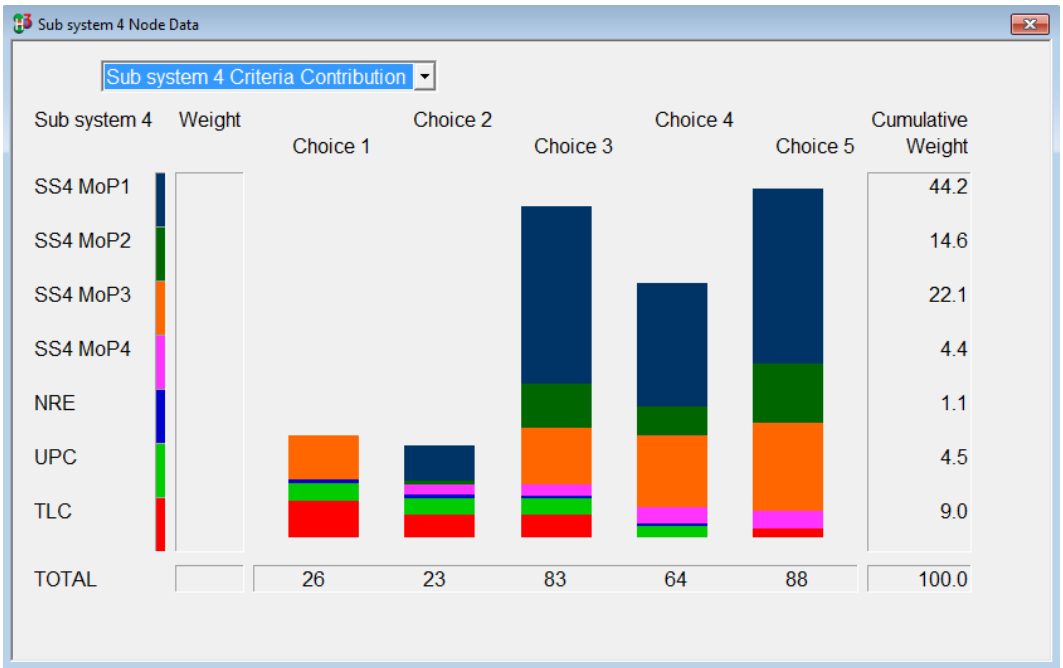

Fig 7. Order of Priority, as Shown in Hiview3

Once the model is scored (using data where appropriate and available, otherwise the best judgement of the experts and stakeholders involved), and weighted (again using data and judgement), it is possible to review and reflect on what the model is revealing. This can be both analytical and a creative process; sometimes new potential Options become apparent.

Figure 7 shows the Order of Priority for 5 Choices available in Sub system 4; the higher the bar the better the Choice. The colours of each segment of the bar indicate the contribution made by each criterion.

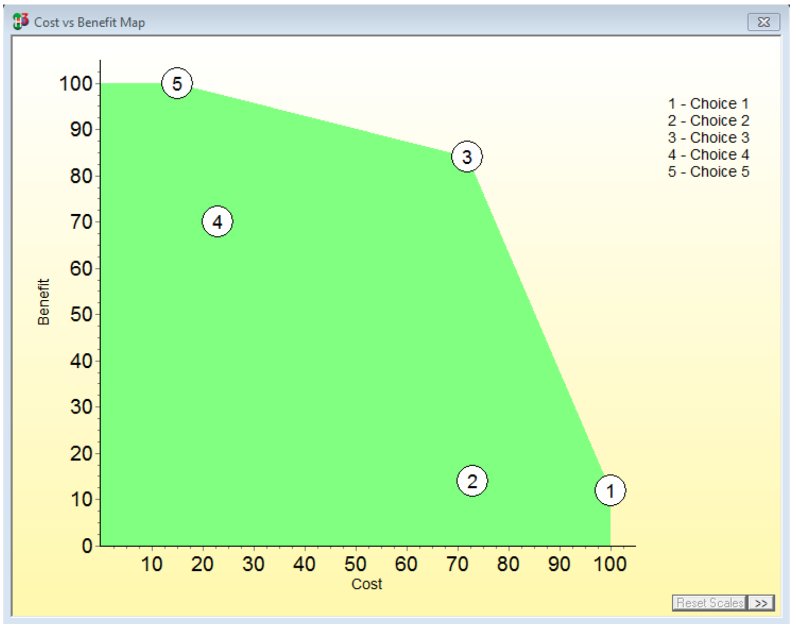

Fig 8. Trade-off Map

Figure 8 shows a typical trade-off map produced by the Hiview3 tool. In this case the trade-off between Benefit and Cost is shown (note that low cost is indicated by a high preference score, high cost by a low preference score, so Choice 1 has the lowest cost and 5 the highest). Any pair of criteria (or nodes in the value tree) can be compared in the same way. In principle the best Choice would be toward the top-right of the chart, i.e. performing well against both criteria. Sometimes these charts show that certain Choices (eg Choice 2 and 4 in Figure 8) are 'dominated', i.e. would not be selected on the basis of either parameter. In other cases it becomes apparent that a new Choice, perhaps combining characteristics of existing Choices, would be better, and so the process supports creativity.

GENERATING OPTIMUM SYSTEM SOLUTIONS

Once the Choices for all Options have been optimised and developed the next stage is to synthesise complete System solutions that can be shown to represent the best possible value for money. System solutions comprise one Choice from each Option in the Option Menu, Figure 9 (reference i). It can be seen that once the number of Options and associated Choices increases then a combinatorial explosion of potential System solutions exists. As in most things a law of diminishing returns applies that states that for the best solutions the rate of increase in benefit as more money is spent decreases. The best solutions thus lie on an optimum Pareto Front (Figure 9) (adapted from reference i).

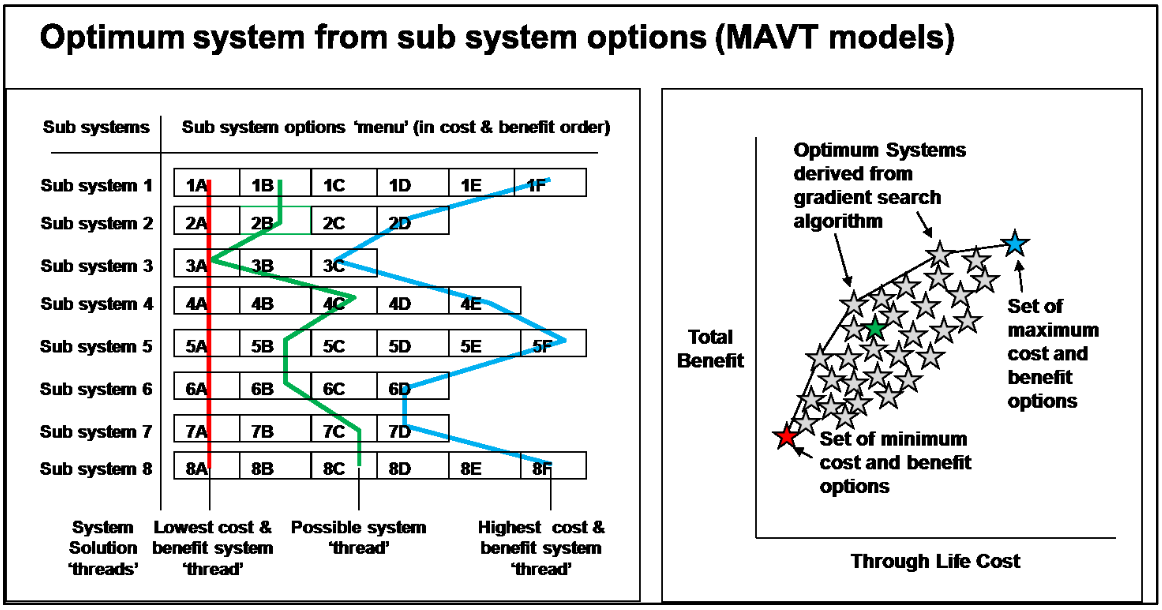

Fig. 9 Optimum Systems Solutions from Option Menu

It is of course impossible to explore all combinations but any solution that can be shown to lie on the optimum front is potentially of interest as it represents the best value for money at that cost point. Provided that the benefits and costs have been combined and weighted in accordance with preferences based on value theory throughout the process then it is possible to determine membership of the idealised, i.e. purely convex, optimum front by direct calculation. This approach is implemented in the Equity3[viii] tool. This tool can be used provided the data from the linked Options menu, MoPs and Benefits tree is pre-processed before input to provide a direct link between the Option Choices and the Benefits. Alternatively the idealised front finding calculations can be performed directly in a large dedicated spreadsheet.

IMPROVING THE PROCESS

The method and toolsets used up to this point assumed preferential independence between all Option or sub system Choices i.e. there are no restrictions on which Option Choices can be selected in conjunction with any others. In many cases this is fine but with any complex engineering system it becomes increasingly difficult to maintain the independence of the Options. In addition, as the number of Options increases the idealised convex optimum front simplification becomes a limitation as the plethora of solutions that lie on the actual optimum front at cost points between the members of the idealised front will be of real interest and the only way of investigating them is by trial and error. Finally the front finding algorithm approaches only work in 2 dimensions, total benefit and total cost. In practice there may be constraints on other cost or benefit components such as the UPC component of total cost. In a 2D optimisation model any additional parameter would have to be manually checked for each system solution generated. The inability of the algorithms to accept constraints thus becomes another limitation, although manageable provided the number of constraints is limited.

In recognition of these limitations and the complexities of using bespoke spreadsheets, a System Engineering general purpose Functional Analysis software programme, BAEFASIP, was adopted and developed to support the process as illustrated in Figure 10. This new tool was used in the later stages of the project to provide additional data that could not be obtained through the existing methods. This both provided a means of validating the toolset and also furnished clarity to the project on the range of options open to them within their budget.

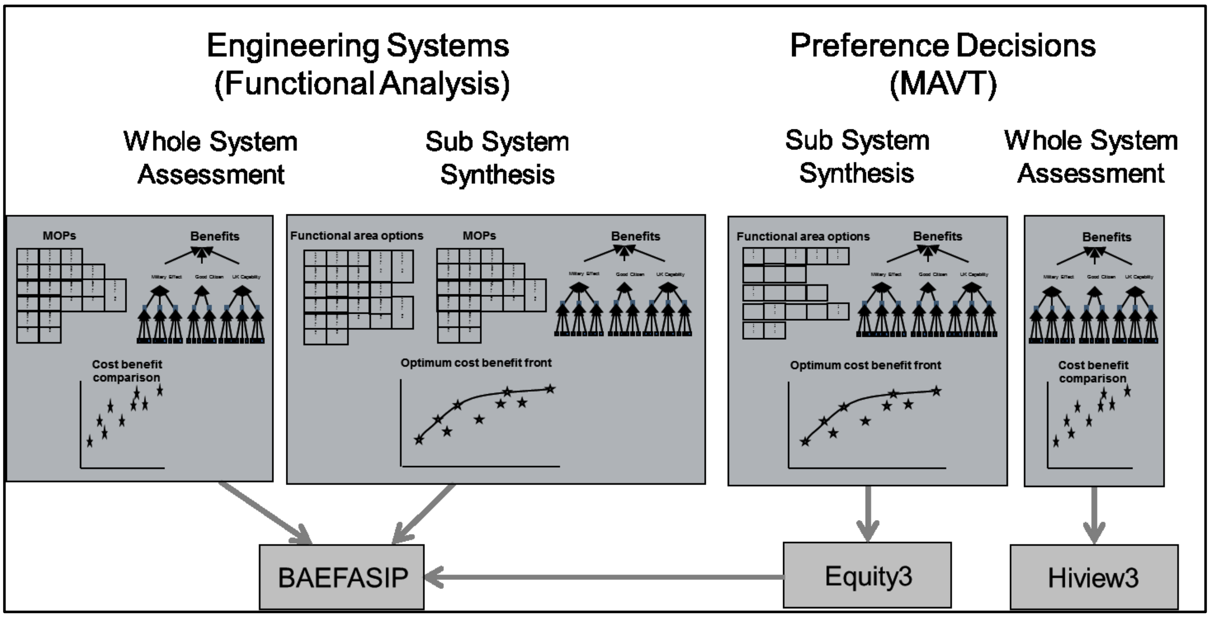

Fig. 10 Cost Effectiveness Toolset

The BAEFASIP tool for Engineering System analysis provided the following features:

- Designed to model linkages between equipments/subs systems, functions and roles/capabilities, data trees with multiple properties and a wide variety of data operators. In the case of the project described, the MoPs corresponded directly to the functions layer in the programme while the Benefits tree layer mapped to the roles/capabilities layer.
- Uses a Genetic Algorithm (GA) to search for the optimum front thus removing the need to assume a purely convex front. For models with many sub systems this can mean a population of up to a thousand optimum solutions can be generated in one run. Each run may not find all possible Pareto front members as the GA has a finite population, but repeated runs with different random seeds in the breeding process will soon expose how well the front is defined. However with such a large population the sheer density of solutions gives a very clear indication of the most important Option Choices at any particular cost level.
- Ability to apply constraints on any calculated parameter (benefit or cost). The programme can also drive solutions into the constrained solution space rather than just disregard them.

The optimum front generated by the BAEFASIP programme is shown in figure 11 (reference i) together with a set of system solutions constrained in cost to a narrow region around a chosen cost level, only one of which could be identified by the idealised front search algorithm initially used by the project.

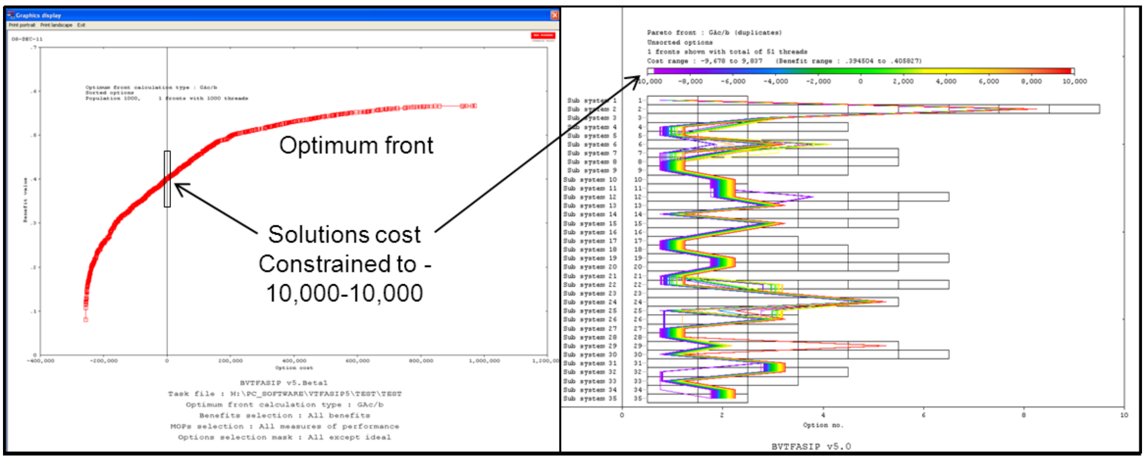

Fig. 11 Optimum Front and Cost Constrained System Solutions

It can be seen that the system solutions from the narrow cost constrained band on the front show that the Choices for some Options are invariant whilst others can be swapped around or traded without really affecting the overall value for money. This additional information was used to support the project Capability Decision Point (CDP) submission by indicating where further sub system trades could be performed in the subsequent stages of the project without significantly affecting the overall budget and measure of delivered capability.

It must be remembered that for the engineering system under consideration all the costs and benefits for the Option Choices could only be generated with reference to a baseline design. The data was generated by considering the effect of each Option Choice on the baseline independently. The process therefore assumes that to a first approximation all the Options are additive in both the benefit preference and engineering senses. In both cases, particularly for the engineering issues, this assumption must of course be confirmed by further design studies on the proposed System solutions recommended by the process.

As a result of experience gained during the project, subsequent development of the toolset has added:

- Ability to generate fronts in multiple dimensions (4 is the practical limit of human visualisation, i.e. a 3D surface coloured according to a fourth variable).
- Ability to incorporate rules that allow or disallow particular combinations of different Option Choices. The programme finds compliant solutions rather than eliminating non-compliant ones so maximising the number of viable solutions in a given population.
- Ability to generate optimum fronts that are robust to uncertainty in particular data values without increasing run time.
- Ability to trade capabilities at the same time as sub system solutions.

There is of course a large amount of effort involved in generating the data for such an analysis. Much of the data is generated anyway within sub system teams but what the process does is provide both a focus for each team and also a communication framework across the project that puts into perspective the contribution of each system element to the overall system cost and capability.

Process improvements under consideration for the future include the explicit modelling and trade-off of risk and uncertainty.

CONCLUSIONS

In conclusion it can be seen that to ensure requirements are pitched and solutions generated that are both affordable and represent best value for money, a combination of assessment approaches and tools are required. The levels of assessment will vary from early high level judgements to low level detailed analysis as the requirements, solutions and understanding develop during the project. The process and tools used must allow for this range of depth of detail and also be flexible in how they can be applied. It is also vital that they are both simple to use and most importantly generate information that is easily understood and communicated to all stakeholders.

The assessment processes and toolset described have evolved during the recent evolution of a complex military platform. They were extensively validated throughout this process by

- Reproducing the data generated by existing tools.
- Reproducing the extensive manual and spreadsheet calculations initially used by the project to implement the multi-level sub-system Option, MoP and Benefit Tree analysis process.

- The new toolset was also able to confirm and define the large number of optimum solutions that preliminary trial and error calculations had indicated existed, but could not be bounded, on the non-idealised Pareto front.

As such both the novel processes and the new toolset that implement it, have proved their value in assisting the project to demonstrate both affordability and value for money.

ACKNOWLEDGEMENTS

The authors thank BAE Systems Maritime – Naval Ships and Catalyze Ltd. for permission to publish this paper and would like to acknowledge the help of their colleagues in its preparation.

DISCLAIMER

The views expressed in this paper are those of the authors and do not necessarily represent those of BAE Systems Maritime – Naval Ships or Catalyze Ltd.

REFERENCES

[i] Courts, Brittain, Osborne, Lamble. Warship Trade Space Exploration, Challenges and Approaches. RINA Systems Engineering Conference, London, 28-29 March 2012

[ii] Howard Raïffa, Robert Schlaifer (1961). Applied statistical decision theory. Harvard University

[iii] Keeney, R. and Raiffa, H. (1976). Decisions with Multiple Objectives: Preferences and Value Tradeoffs. New York: Wiley.

[iv] Ronald A. Howard (1966). Decision Analysis: Applied Decision Theory. Stanford Research Institute

[v] Barclay, S., Brown, R. V., Kelly, C. W., Peterson, C. R., Phillips, L. D., & Selvidge, J. (1977). Handbook for Decision Analysis (Technical Report TR 77-6-30). McLean, Virginia: Decisions and Designs, Inc.

[vi] Dr. Larry Phillips, white paper at http://www.catalyze.co.uk/wp-content/uploads/2012/07/Catalyze-Decision-Conferencing-White-Paper.pdf

[vii] Hiview3 software from Catalyze Ltd, see http://www.catalyze.co.uk/index.php/software/hiview3

[viii] Equity3 software from Catalyze Ltd, see http://www.catalyze.co.uk/index.php/software/equity3