Intl. Trans. in Op. Res. 10 (2003) 155–167

INTERNATIONAL TRANSACTIONS IN OPERATIONAL RESEARCH

An application of linear programming in the defence environment

N. S. Walmsley[a] and P. Hearn[b]

[a]Dstl Land Systems Department; [b]QinetiQ Consulting, Fort Halstead, Sevenoaks, Kent, TN14 7BP, U.K.
nwalmsley@Dstl.gov.uk(Walmsley); phearm@QinetiQ.com(Hearn)

Received 18 June 2002; received in revised form 26 November 2002; accepted 17 December 2002

Abstract

Operational analysis (OA) techniques are used extensively by the Ministry of Defence (MoD) to inform their decision-making process. One particular type of OA study frequently undertaken is a balance-of-investment (BoI) study. This paper focuses on a recent BoI study whose aim was to identify the optimum (i.e. most cost effective) mix of vehicles that should be procured to fill a large number of identified requirements/roles. There will be a description of the methodology employed, highlighting the formulation and application of mixed integer linear programming (MILP) techniques, alongside a brief discussion of the implication of the results and other issues associated with the complexity of the problem.

Key words: Cost benefit analysis, linear programming, balance of investment (BoI), Operational analysis, value for money

Introduction: UK Ministry of Defence

The procedures followed by the MoD when introducing new vehicles or equipment mean that a business case has to be submitted to a committee of high-ranking military officers and civil servants for approval and possibly for ministerial approval. The main aim of this committee is to ensure the Defence budget is spent in a wise manner and that the highest level of value for money is achieved. This is where OA plays a crucial part.

There are many different types of OA study. The study described in this paper is known as a balance-of-investment (BoI) study. The output from a BoI typically identifies the most cost-effective mix of a set of related items being procured. It is used to show the approvals committee that the quantities of each item proposed for procurement are the most cost effective.

As their name implies, Armoured Combat Support Vehicles (ACSV) are used by the British Army to support (but not primarily conduct) armoured operations. They have a wide variety of roles with corresponding requirements for capability. The Army planned to introduce four types of ACSV later this decade. Furthermore, a review of the requirements for such vehicles showed that, to some extent, these four would fill similar capability gaps for different user groups (i.e. the Army/Royal Air Force (RAF) staff that would actually be driving around in ACSVs in real-life situations).

Aim and scope of paper

The aim of this paper is to present a case study using the ACSV BoI as an example of a real-life application of mixed integer linear programming (MILP) within the Defence environment. The relevant underlying methodology will be looked at in detail and some of the complex issues that arose in its implementation will be discussed.

Problem formulation: BoI approach

The methodology proposed for the BoI utilized a two-phased approach. Each phase is described below.

Phase I – requirements capture

The aim of Phase I of the BoI was primarily to identify all ACSV roles and capture their requirements both for capability and numbers. The requirements were initially challenged by the militarily experienced team conducting the activity. OA modelling techniques were intended to further challenge these requirements later on in the study.

A further task carried out in this initial phase was to map each identified role onto its corresponding operational unit(s). This enabled initial calculation of the total fleet requirement (TFR) (i.e. the total number of vehicles of each type required).

In calculating the TFR the following factors had to be taken account of:

- individual role requirements;
- deployment types (i.e. differing combinations and quantities of operational units);
- wartime reserves;
- other requirements (including vehicles required for training etc).

Phase II – data interpretation and optimization

The role requirements data captured in Phase I can be broken down into three main types of requirement, capacity, mobility, and survivability.

Capacity. The capacity requirements of the identified roles were defined in terms of the number of people that had to fit into the vehicle, and the size and mass of any additional equipment needed. This was the primary driver of the requirements and, unlike mobility and survivability, the capacity requirements of each role had to be fulfilled. No shortfalls were acceptable.

Trade-offs were permitted in both mobility and survivability categories. Later OA modelling would investigate whether or not any of the trade-offs made resulted in risks too high to be deemed militarily acceptable.

Mobility. To enable direct vehicle-to-role comparisons to be undertaken in mobility, each role was considered in turn and placed in one of four mobility requirement bands (band 1 indicating the least demanding roles and band 4 indicating the most demanding roles). The allocation of mobility bands to each role was based upon military experience and a range of requirements data captured in Phase I (e.g. the primary area of location of the roles, the vehicles being supported by the roles, the type of support being provided by the roles, etc.).

The naming of the mobility bands (i.e. 1, 2, 3 and 4) is not intended to imply that band 2's requirements are twice as high as band 1's. However, it is true to say that band 2's requirements are at least as demanding as band 1's.

Metrics were then defined to cover the broad term 'mobility'. These definitions were made in consultation with the relevant experts (both technical and military) and consisted of measurable vehicle characteristics such as maximum road speed and power-to-weight ratio. These metrics were then related back to the four mobility levels assigned to the roles thus enabling direct mobility comparisons to be made for any vehicle-to-role combination.

Survivability. The role survivability requirements captured in Phase I identified the worst case threats likely to be encountered by each role. The following threat categories were considered:

- direct fire (DF);
- indirect fire (IF);
- mines;
- air.

Sets of vehicle characteristics were then defined that would be militarily desirable should these threats actually be encountered. These characteristics were given in terms of the overall thickness of armour, additional armour (i.e. belly armour for mines) and thermal imaging (TI) capability (i.e. being able to see the enemy in the dark and thus avoid the threat or attack it first).

ACSV option identification

To enable comparisons to be made between role requirements and ACSV characteristics, a set of confirmed ACSV options had to be identified. As previously mentioned, four types of ACSV are already being considered for future use by the British Army. For one of these types there exist two variants. It was also necessary to include some cheaper options for affordability purposes.

Compliance testing

The complete sets of role requirements and vehicle characteristics were then redefined in terms of the categories discussed above. This enabled all vehicle-to-role comparisons to be calculated using a simple spreadsheet model.

For all vehicle-to-role combinations the compliance was tested in each category as defined below.

- Vehicle characteristic > role requirement = > excess vehicle capability
- Vehicle characteristic = role requirement = > exact match
- Vehicle characteristic < role requirement = > shortfall in vehicle capability

The comparison results were then colour-coded to highlight different levels of excesses and shortfalls. Fig. 1 shows the key to the colour-codes used for the mobility and survivability characteristics. Fig. 2 shows contrasting examples of the outputs produced at this stage.

After these results have been calculated it is possible to construct a compliance matrix that clearly identifies the following three cases:

- where a vehicle meets all of the requirements of a role;
- where a vehicle meets the capacity requirement of a role but has a shortfall in one or more 'other' category;
- where a vehicle fails to meet the capacity requirement of a role.

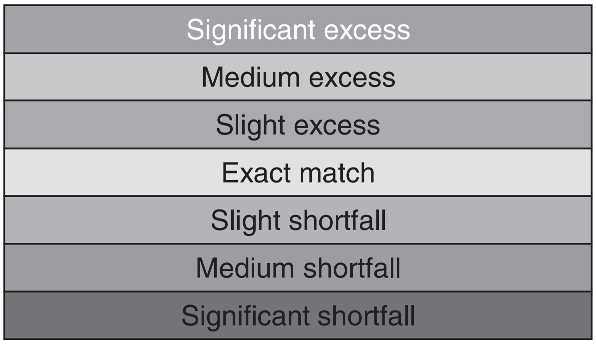

Fig. 1. Key to colour-coded comparison results.

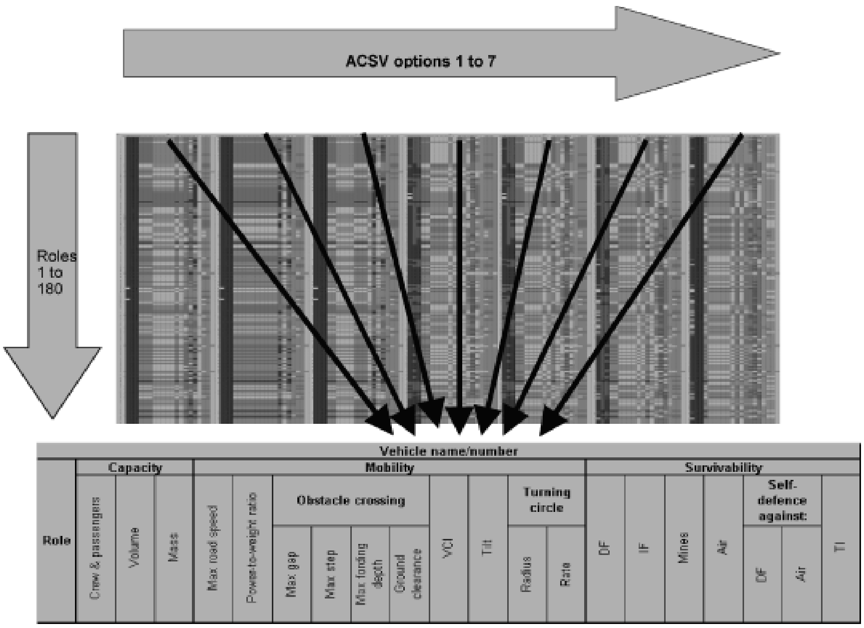

Fig. 2. Colour-coded comparison results.

Vehicle meets (or exceeds) all role requirements
Vehicle compliant in capacity but fails in at least 1 ‘other’ category
Vehicle fails in capacity (may or may not fail elsewhere)

Fig. 3. Definitions of ‘traffic light’ colours.

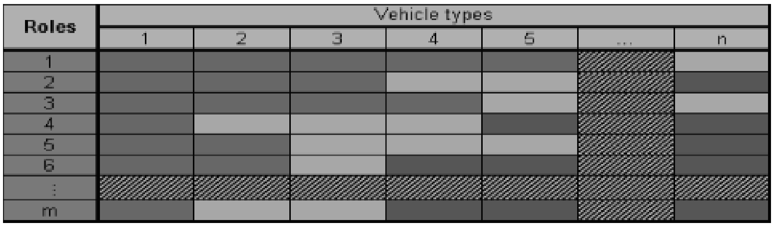

Fig. 4. Example compliance matrix.

Again, this can be shown in its simplest form using a ‘traffic lights’ system. Fig. 3 defines the significance of each of the three colours. Fig. 4 shows an example of how the compliance matrix may look.

Adopted approach: optimization methodology development

For a given set of vehicle-to-role allocations it is now possible (using the compliance matrix as described above) to calculate/identify:

- the total number of roles that have been allocated a fully compliant (i.e. green) vehicle;
- the total number of compliant vehicles within the fleet (i.e., the sum of the compliant roles multiplied by their respective TFR vehicle requirement);
- roles where a cheaper vehicle than the one currently allocated meets all the requirements;
- fleet cost.

Given the problem definition above, a solution was then investigated.

The model

Essentially the ACSV problem defined above can be represented in its simplest form as a basic supply and demand model, the solution of which is attainable by employing the dual simplex algorithm. This algorithm solves the minimization problem by generating a series of successive basic solutions such that the series converges on the optimum extreme point in the solution space (Taha, 1992).

Formulation of the LP

Consider the classic LP problem: a factory that manufactures n different types of product to supply m different outlets. The objective is to minimize production and supply costs, whilst satisfying constraints on both supply and demand. Mathematically, this may be stated as follows:

$$\begin{aligned}
&\min \ \text{Cost} = \sum_{i=1}^{n} c_i x_i \ \text{(objective function)} \\
&\text{subject to:} \quad \{f_i(x_1, x_2, ..., x_n) \leqslant a_i : a_i \in \Re\} \ \text{(cost constraints)} \\
&\qquad\qquad\quad\ \{g_j(x_1, x_2, ..., x_n) \geqslant b_j : b_j \in \Re\} \ \text{(supply constraints)}
\end{aligned}$$

given that c_i is the unit cost and x_i is the quantity of products of type i and the functions f_i and g_j are linear. The x_i variables are usually constrained to be positive reals.

The demand constraints are given as:

- only one vehicle type may be allocated to each role;
- each vehicle allocated to a role must have sufficient capability to fulfil the role capacity and TI requirements.

The problem faced by the ACSV BoI required advice to be given on the procurement of the optimum mix of ($n = 7$) different vehicle types to fulfil ($m = 182$) various roles (each role having an individual set of requirements).

Initially, it was identified that there were three different associated questions to answer. These are as follows:

- *Case 1*. Minimize the cost for a 100% compliant total fleet requirement (TFR) (i.e. all vehicles allocated to a role must have no capability shortfalls in any of the compliance categories stated above).
- *Case 2*. Maximize the total number of compliant roles (i.e. roles allocated a fully compliant vehicle) within the fleet for a fixed maximum cost bound.
- *Case 3*. Maximize the total number of compliant vehicles within the fleet (i.e. total number of vehicles that fully satisfy the requirements of the role to which they have been allocated) for a maximum cost bound.

Case 1 clearly gives the most direct advice to the ACSV problem. However, if 100% compliance is not possible for a particular maximum cost bound it fails to provide any advice at all. Cases 2 and 3 were required to fill this identified gap.

It is clear from the definitions above that the differences between Cases 2 and 3 are minimal. However, ranking the roles in order of importance (i.e. to enable the MILP to fill the most important roles with compliant vehicles first) proved too difficult an issue to address in this study. Hence, Case 2 effectively assumes that each role has an equal weighting (implying each role is as important as the next) whereas Case 3 weights each role according to the number of vehicles required to fulfil its requirements. It was envisaged that an analysis of the vehicle allocations resulting from both cases could provide valuable advice concerning the best course of action to take (for that particular level of funding).

A list of parameter, decision variables and functions is specified in Table 1.

Table 1
Description of parameter, decision variables and functions used in LP model

Parameter	Description
totalvehs	Total number of vehicle types
totalroles	Total number of roles
l_j	Number of vehicles required for role j, $j \in \{1 \ldots totalroles\}$
w_i	Wartime Reserves for vehicle type i (this is specified as a proportion of $a_{i,j}l_j$, i.e., $0 \le w_i \le 1$)
o_i	'Other' vehicles of type i required
$q_{i,j}$	Compliance of vehicle i to role j
Variable	**Description**
$a_{i,j}$	Allocation matrix
Function	**Description**
v_i	Number of vehicles required of type i, $i \in \{1 \ldots totalroles\}$
$k_i(v_i)$	Cost of buying v_i vehicle type i, $i \in \{1 \ldots totalroles\}$
S	Total cost of a given mix (Objective function for Case 1, constant for Cases 2 and 3)

The number and cost of each vehicle type may be defined as follows:

Number of vehicles: $v_i = \left\lceil (1 + w_i) \sum_{j=1}^{totalroles} a_{i,j} l_j \right\rceil + o_i$, for $i \in \{1 ... totalvehs\}$.

Cost of vehicle: $k_i(v_i) = \alpha_i v_i + \beta_i$, where $\{\alpha_i, \beta_i \in \Re : i \in \{1 ... totalvehs\}\}$.

Specially ordered sets

To fully understand the formulation below, an understanding of specially ordered sets is essential. The main benefit of using a specially ordered set is the significant reduction in run-time achieved. The specially ordered sets utilized below are of type 1. Specially ordered sets of type 1 are sets of binary values the sum of which is 1. This means that if a set of n binary numbers is a specially ordered set of type 1 then it contains $(n-1)$ 0's and a single 1.

The mathematical formulation for each of the three different LP models is as follows:

- *Case 1* (Minimize cost for 100% compliance)

Objective function: $\min S = \sum_{i=1}^{totalvehs} k_i(v_i)$

Constraints:

(i) 100% compliance: $\{a_{i,j} q_{i,j} = 1 : \forall i \in \{1 \ldots totalvehs\}, \forall j \in \{1 \ldots total\ roles\}\}$
(ii) Integers: $q_{i,j}$, v_i, l_j
Specially ordered set (Type 1): $(a_{1,j}, a_{2,j}, \ldots, a_{i,j})$

Positive numbers: v_i, l_j, w_i, o_i, k_i v_i, S, *totalroles*, *totalvehs*

(iii) $q_{i,j} = 1$ if vehicle i is 100% compliant to role j
$= 0$ if vehicle i is non-compliant to role j in mobility/survivability
$= -999$ if vehicle i is non-compliant to role j in capacity

- *Case 2* (Maximize number of compliant roles)

A role is compliant if: $\{\exists i \in \{1 \ldots totalvehs\}: a_{i,j}q_{i,j} = 1, \forall j \in \{1 \ldots totalroles\}\}$

Objective function: $\max z_1 = \sum_{i=1}^{totalvehs} \sum_{j=1}^{totalroles} q_{i,j}a_{i,j}$

Constraints:

(i) $\sum_{i=1}^{totalvehs} k_i(v_i) \leqslant S$

(ii) Integers: $q_{i,j}$, v_i, l_j

Specially ordered set (Type 1): $(a_{1,j}, a_{2,j}, \ldots, a_{i,j})$
Positive numbers: v_i, l_j, w_i, o_i, k_i v_i, S, *totalroles*, *totalvehs*

(iii) $q_{i,j} = 1$ if vehicle i is 100% compliant to role j
$= 0$ if vehicle i is non-compliant to role j in mobility/survivability
$= -999$ if vehicle i is non-compliant to role j in capacity

Note: $z_1 \geqslant 0 \Rightarrow$ feasible solution (i.e. no vehicles fail on capacity).

- *Case 3* (Maximize number of compliant vehicles)

Objective function: $\max z_2 = \sum_{i=1}^{totalvehs} \sum_{j=1}^{totalroles} q_{i,j}a_{i,j}l_j$

Constraints:

(i) $\sum_{i=1}^{totalvehs} k_i(v_i) \leqslant S$

(ii) Integers: $q_{i,j}$, v_i, l_j

Specially ordered set (Type 1): $(a_{1,j}, a_{2,j}, \ldots, a_{i,j})$
Positive numbers: v_i, l_j, w_i, o_i, k_i v_i, S, *totalroles*, *totalvehs*

(iii) $q_{i,j} = 1$ if vehicle i is 100% compliant to role j
$= 0$ if vehicle i is non-compliant to role j in mobility/survivability
$= -999$ if vehicle i is non-compliant to role j in capacity

Note: $z_2 \geqslant 0 \Rightarrow$ feasible solution (i.e. no vehicles fail on capacity).

The three cases were then coded using (XpressMP, 2001), which is a linear optimization software package. It should be noted that the solution to the problem required the use of MILP techniques since it was necessary to define the allocation matrix as a specially ordered set of type 1. This was essential to ensure a meaningful allocation scheme that permitted the allocation of only one vehicle type to each role.

In addition, integer techniques were employed in the construction of the compliance matrix that removed the possibility of allocating a vehicle whose capacity is non-compliant to a particular role. This was carried out by assigning those vehicle-to-role combinations that failed in capacity a large negative integer (i.e. a penalty value that ensured the objective function could not be a maximum).

The benefits of a generic formulation

It is important to exploit functionality when designing LP models. A generic design approach facilitates the possible requirement to answer a number of different associated questions without the need to make extensive alterations to the model.

In this particular problem, a number of different variations of the model were produced that offered solutions associated with the general study objective. In addition, this highlighted the opportunity to perform sensitivity analysis by systematically altering compliance criteria. By standardizing both the mathematical formulation and the inputs required, it becomes possible to assess a vast range of issues by simply adjusting one or two lines in the coded formulation. Incorporating the use of batch files further enhanced the potential capability of the model by allowing a series of successive problems to be generated and solved in a sensitivity analysis approach.

Testing the model

Two types of testing were employed to minimize the risk of producing error. First, the intrinsic logic of our code was tested: a small-scale version of our problem was set up, i.e. number of vehicles is 3 and number of roles is 5. This was carried out by developing a Microsoft Excel Visual Basic macro to test and cost all possible valid vehicle/role permutations, and was used to verify that the LP solution corresponded to the minimum-cost compliant solution.

Second, to test whether the software package would cope with the scale of our problem, we tested against the current solution (i.e. the TFR baseline) as supplied by the customer.

Essentially, the TFR baseline was a non-optimal solution to the ACSV problem. All vehicle-to-role allocations were determined by the MoD customer using his military knowledge and experience. As a result, using the TFR baseline to test the model also provided a vital audit trail.

Furthermore, as a direct result of testing the TFR baseline a number of the allocations (which had already been agreed by the ACSV user community as an acceptable solution) were identified as being non-compliant. Hence, the requirements of these roles had been downgraded and accepted by the customer without realization of the fact.

LP implementation

The implementation of the model took into account the need to set up a series of runs of the model, with each run varying a number of parameters for different fixed budgets. Batch files were set up to retrieve the input files from a pre-defined computer directory structure, and sequentially execute the LP model. The output files were automatically stored in the directory structure, which facilitated the automation of the post-processing of results.

Results

This section discusses the different approaches taken regarding the analysis of the MILP results. It became clear as the model was applied that the customer could gain great benefits (in terms of analysis outside the scope of the BoI) for minimal extra cost. Once the LP implementation was automated it proved straightforward to engineer various different parameter adjustments that enabled sensitivity analysis to be conducted.

Post-processing

Each run of the LP model produced a collection of five output files. Typically, a series of runs involved 10 parameter adjustments over 15 fixed budget costs (for Cases 2 and 3 only), thus producing 750 output files. It was therefore necessary to develop a Microsoft Excel macro to automatically collate the output data and generate graphs of the results.

Assessing the sensitivity of fleet costs to role capacity requirements

The predominant factor in determining vehicle-to-role compliance is the role capacity. Due to the fact that the capacity requirements of the roles (in terms of volume and mass) were estimates it was decided that the impact of over/under-estimating would be assessed. Fig. 5 shows a similar effect to that identified in the BoI. In Fig. 5 each bar represents the total fleet cost based on differing levels of role capacity requirements. For example, the total fleet cost in the '-5%' case is based upon the assumption that the volume and mass requirements of each role are 95% of those captured in Phase I.

Impact of assessing the sensitivity of fleet costs to role capacity requirements

The key issues highlighted in Fig. 5 are:

- There is little penalty in increasing the volume and mass requirements of the roles (except that no vehicle will have the required capacity beyond a certain point).

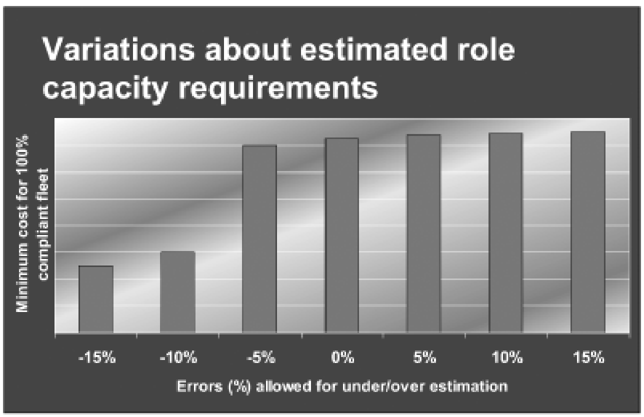

Fig. 5. Variations about estimated role capacity requirements (example results).

- Potentially there is a big saving to be made by reducing the volume and mass requirements of the roles.

The fact that the notable 'jump' in minimum fleet cost occurred reasonably close to 0% (i.e. the current situation) meant that further investigation was required.

One of the ACSV programs had not matured to the point where a specific make/model of vehicle had been chosen to fulfil the requirements. Hence, the capacity characteristics assumed for the purposes of the BoI were based upon the mean characteristics of the range of vehicle options still under consideration. The further investigation into Fig. 5 highlighted the fact that the capacity of the vehicle chosen to fulfil the requirements of this program, alongside the accuracy of estimation of all requirements, was a significant contributor to the total fleet cost.

Assessing the sensitivity of role location

As mentioned above, the complete set of ACSV roles identified in this study span the entire battlefield. It is widely accepted that the battlefield can be divided into specific areas within which threat characteristics were approximately equal. Fig. 6 defines the layout of a generic battlefield highlighting the divisions assumed for this BoI.

It was hoped that dividing the complete set of roles into smaller, more manageable sets would facilitate the analysis. Unfortunately this was not the case. Even though the roles were divided into smaller sets there was still no method available to compare the importance (i.e. identify which role should be satisfied first) of two roles in different areas.

Ignoring this importance factor, the next set of variations were based upon the breakdown of roles by location as shown in Fig. 6. These variations were carried out as follows:

- start with the vehicle-to-role allocations from the baseline solution;
- assume that all capacity requirements are non-negotiable and must be fulfilled;
- for each role, assume that the vehicle allocated is the only compliant vehicle;
- run the LP;

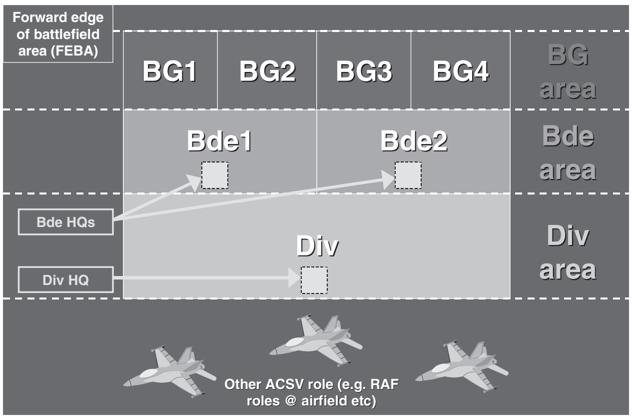

Fig. 6. Generic battlefield layout.

- for all of the roles in a particular area: let the 'next best vehicle' (these priorities were determined using military judgment) also become compliant (i.e. so the LP has to choose between the baseline-assumed vehicle and the next most capable vehicle – even if it is not 100% compliant);
- run the LP;
- for the same set of roles: let the next best vehicles become compliant etc…
- continue until there are no vehicles left to become compliant.

This was repeated looking at each area of the battlefield on an independent basis. Fig. 7 shows similar results to those obtained in the BoI.

Assume there are six sets of ACSV roles (i.e. one set for each area of the battlefield). Each bar in Fig. 7 corresponds to the minimum total fleet cost based upon the following assumptions:

- the roles in the set whose battlefield area is the focus of the sensitivity case (i.e. BG = left-hand side in Fig. 7) have a choice of ACSVs to be allocated;
- the choice of ACSVs for this set is defined by the iteration number where:

  > 1 = baseline solution only (i.e. no choice);
  > 2 = baseline solution or next best ACSV;
  > …;
  > 7 = all ACSV options with sufficient capacity characteristics are available for allocation.

- all roles not in this set are allocated their baseline ACSV.

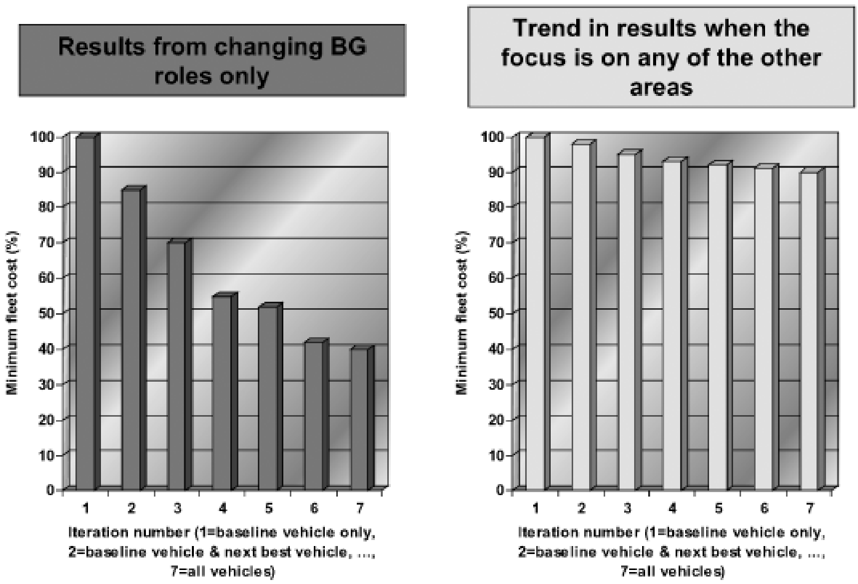

Fig. 7. Variations by role location (example results).

Impact of assessing the sensitivity of role location

Following this definition, Fig. 7 suggests that the roles within the BG area contribute to a high proportion of the total fleet costs. This raises the question of whether the scope of the problem should be reduced to encompass the BG roles only. However, assessment of this question is outside the scope of this paper.

Summary and conclusions

We have presented a description of the various stages of a BoI study carried out to offer quantitative advice on a procurement decision in the Defence environment. Our work has been limited so far to examining the balance of investment without demonstrating the need for those requirements being satisfied. However, this paper has shown how the use of LP techniques may be formulated and applied to model a real-life BoI problem in support of a military capability/equipment procurement decision.

Data collation and quantification techniques are applied to represent the information in such a form that it is readily readable as input to an LP algorithm. An important principle employed during the data-collation process, and one which continues throughout the study, is to actively encourage the support of subject matter experts to give authority to the validity of any assumptions made; in this study, this component was supplied in the form of military expertise from the customer. This highlights the even more important principle to involve the customer throughout the study to develop an acceptance of the methods used and results produced. In addition, the model gives a flavour of the complex issues arising from a series of military constraints, thus giving rise to a complicated allocation problem.

Importantly, this paper highlights the need for the analyst to identify associated solutions linked to the main problem, thereby offering a number of different perspectives on the problem. In this study, three different solutions to the problem are given; however, the models have been developed in a generic manner, therefore allowing the opportunity to exploit methods of automating the generation of results.

A common feature of BoI studies for the MoD is the technique of specifying a critical value, below which a solution is deemed impractical. For example, in this study it was possible to relax certain vehicle/role characteristics such that a series of ‘next best vehicle’ solutions may be generated. This gave information about the relationship between the degree of relaxing of vehicle characteristic parameters (i.e. capacity, mobility and survivability) and its effect on cost, and provided useful information for the customer.

References

1. Taha, H.A. 1992. *Operations Research: an Introduction*, 5th Edn, Prentice Hall, pp. 59–60.
2. XpressMP (computer software package), 2001, Dash Optimization.