F3

THE APPLICATION OF UTILITY THEORY TO AIR CAMPAIGN MODELLING.

Dr R C Wheeler & J E M Raraty

Paper delivered to the 12th International Symposium on Military Operational Research, at RMCS Shrivenham, 5-8 Sep 95

Abstract

The nature of campaign-level decisions as they affect the Air component of an operational HQ is described. The need to provide rapid OA support for such decisions generates a need for highly aggregated models. The concept of utility offers a method of creating such a model. Its use is described with an example based on Operation DESERT STORM.

The views expressed here are those of the authors and do not necessarily represent official MOD policy. Requests for further information should be addressed to Dr R.C. Wheeler, Air Warfare Centre, Operational Analysis Element, RAF High Wycombe, Bucks. HP14 4UE

Section 1 - INTRODUCTION

1. Of the decisions that need to be made in planning a campaign, those that involve the employment of air power revolve around the apportionment decision: how is the air power available to be divided up between the different roles and how should this change in the course of the campaign? Implicit in this decision is a lot of fine detail: not all aircraft types will be capable of all roles; some will be capable in principle but will lack the range to take on certain targets associated with a role; and much else. There will be further decisions that follow on from this apportionment decision: on the force-mix required, on basing, on logistic support. The work described here inevitably skates over a lot of this detail; in fact, one of the biggest problems we have in pursuing it is to find a way of incorporating sufficient detail to give the results some credibility without complicating the technique unmanageably.

2. We are essentially talking about apportionment between different campaign objectives and it may be helpful to think of these as sponsored by different "customers". The Land Commander wants attacks against the enemy's land forces and their Lines of Communication. The Maritime Commander wants the enemy's ships sunk and his own safeguarded. The Air Commander is a customer in his own right, wishing to gain air superiority as soon as possible. Finally, there will be the sponsors of strategic objectives, taken to be those objectives which directly affect the will of a nation to continue to fight; for example, based on extensive research at these symposia, an aggressor against any of the nations represented here might consider beer production such a strategic objective.

3. In accordance with the broad-brush approach of this paper, we shall assume that we have a method of estimating what resources will be required to attack successfully any particular target and what losses are likely to be incurred in so doing. We shall also assume that customers have a means of prioritizing targets which contribute to their own objectives, indeed that they have a system of values for such attacks which rates them in terms of a customer-specific standard (say, "tank-equivalents killed" for the land commander). The problem addressed in this paper is how to compare targets across customers: comparing tanks with airfields, for example.

4. The natural response of most operational analysts to this problem is to propose a simulation model broad enough to encompass tanks and airfields (and breweries) and to use this model to compare alternative strategies. This approach suffers from the following problems:

a. A suitable simulation of adequate fidelity does not exist and we cannot afford to develop one. What constitutes adequate fidelity? Well, to take an example, one expects to see diminishing returns for Close Air Support(CAS). In other words, the first aircraft made available to a commander will be used in circumstances where that one attack can make a huge difference and where the risk to the aircraft is relatively low; as more aircraft are made available, such highly favourable circumstances are used up and aircraft are employed in circumstances where the pay-off is less and the losses higher. This sort of issue is highly pertinent to apportionment decisions but few models can represent it credibly and none (that I know of) which also straddle both the air and land battles.

b. Even if such a simulation did exist, it is difficult to envisage being able to set up the scenario and input the curent data in time to influence the official histories, let alone the actual apportionment decisions.

c. Finally, no model is going to be able to represent the effect of strategic attacks because no-one understands enough about the effect of such attacks to be able to set down a mathematical relationship for them.

Thus a different approach is needed.

5. What is proposed is to distinguish between *ends* and *means*. Strategic attack is an end in itself because it aims to affect the enemy's will to fight. Attack of land forces is (normally) an end in itself because destruction of an enemy's land forces generally removes his will to fight as well as his means. In contrast, at least for the type of enemy we are likely to find ourselves fighting, destruction of air or maritime forces is only a means to an end: it reduces their ability to continue the war but not their will to do so.

6. No model is going to be able to compare objectives which are ends in their own right (ie strategic and land force attack) ; this can only be done by agreement between the respective customers. However, it is possible for a mathematical model to relate means to ends. Thus, if the strategic and land customers can agree on a common valuation of their respective target lists (in terms of, say, tank-equivalents killed) then it ought to be possible to relate air and maritime targets to that scale.

7. This is actually no small achievement. If one looks back over the history of air/land doctrine, there have been interminable arguments over the relative importance of gaining air superiority and using air power to meet the Army's objectives. Here we have a model which can address this question directly.

8. The method used will be described in detail in section 2. A value (in terms of tanks, or whatever other scale the land commander chooses to adopt) is placed on sorties flown in support

of strategic or land objectives. From this, a value can be placed on the aircraft. A similar method can be used to place a value on the enemy's aircraft (ie the value that we should place on depriving him of their use). This can then be used to value attacks on his airfields and other air-related objectives. While it should, in principle, be possible to follow a similar method in respect of maritime targets, this has been left to another occasion. The real test of a method like this is whether it can actually be applied. Section 3 does indeed apply it to a hypothetical conflict bearing a close resemblance to Operation DESERT STORM. and concludes that, under such circumstances, gaining air superiority should indeed be the highest priority. In an ideal world, one would have undertaken a detailed sensitivity analysis to show that this result continues to hold even if the input values are varied within reasonable bounds. However, time has not allowed this to be done and the results are presented more as an indication of what can be done with the method than as a serious contribution to air doctrine.

9. No great originality is claimed for the ideas here. The basic concepts have been current in micro-economics for a considerable time. The authors are indebted to John Owen of the UK's Defence Operational Analysis Centre (now CDA High Level Studies Dept) for suggesting this particular application of them. He should not be held responsible for the manner in which we have done so.

Section 2 - THEORY

UTILITY OF ATTACKS (INTERDICTION OR CAS)

10. Interdiction and CAS are both ends in themselves so their values can be calculated directly. For example, if the value of a tank killed is T and it takes 2 sorties to do it[1], then the value G of a tank-busting sortie is $T/2$. This is a *gross* values, i.e. it does not take into account the cost to the air commander in terms of aircraft losses. If the attrition rate for the type of sortie involved is p and the value of an aircraft is A, then the net value is $G\text{-}pA$. Thus we need to calculate the value A before we can derive net values.

11. This is done by considering *marginal* values: in other words one calculates what sorties one would *not* fly, both now and in the future, if one had 1 less aircraft now, or (much the same) the *extra* sorties one would fly if one had 1 more aircraft now. (One consequence of this is that the method does not derive a "best" apportionment overall but rather advises on the benefit of small departures from the present plan. Of course, by making many such small changes and reapplying the method each time one can end up with a major change overall.)

12. For very short campaigns, it may be possible to calculate this value of an aircraft directly; for longer campaigns, one will need to use a method such as the following. Suppose that, for the foreseeable future, the marginal activity (i.e. that on which extra aircraft are employed) is one with a gross value of G and a discount rate (per sortie interval) of r (i.e. a sortie on the second wave is only worth r times one on the first wave, one on the 3rd wave only r times one on the 2nd). r is specified by the ground commander and reflects the extent to which sooner is better than later; a typical value might be 0.95.

13. We may also have a planning horizon of n sorties, beyond which we expect the war to have finished and the aircraft to have no further value. Note that the assumption of zero aircraft value at the end of the war is made merely for simplicity here - it would be quite possible to give an aircraft a residual value after the completion of its allotted operations (this is covered in more detail in the Annex). The value of an aircraft is the value of what it can do on this marginal task, viz.:

G	on the first sortie
$+Gr$	(if it survives the first sortie)
$+Gr^2$	(if it survives the second as well)

---

[1]i.e. on average 2 sorties must <u>launch</u>. Sortie vales are calculated at launch and allowance must be made in this calculation for attrition before weapon release.

$+ \ldots$

$+Gr^{n-1}$ (if it survives the first n-1 sorties).

If we set $q = 1-p$ to be the survival probability, this expression becomes

$G\,(1 + qr + (qr)^2 + \ldots + (qr)^{n-1})$,

which can be simplified as :

$$A = G\left(\frac{1-(qr)^n}{1-qr}\right)$$

14. This is the value of the aircraft now. Its value after the sortie will have reduced by a factor r since all the future sorties it can fly will have slipped 1 to the right so will have declined in value by a factor r. If the aircraft is lost, what is lost is the ability to fly subsequent sorties; hence the value lost is rA. The net value of the sortie is therefore $G - prA$. It will be noted that this is always positive (as it should be!).

15. This mathematical excursion is not central to the argument but is worth including to demonstrate that the methodology outlined here can actually be applied, given reasonable assumptions on the (gross) value of sorties. One could put in different values of G and p for different offensive missions and one ought to find that the net sortie value is higher for all those sortie types for which one was flying all the sorties possible and lower for all those types for which one was not planning to fly any sorties at all. (Recall that the above calculations were for the marginal sortie type, i.e. that to which one would allocate an extra aircraft if it became available or from which one would remove an aircraft if one were lost).

16. The process could be repeated for each different aircraft type. Clearly, all aircraft will not have the same value, any more than they have the same loss rates.

UTILITY OF ENEMY ATTACKS

17. It has already been assumed that values can be placed on preventing the enemy from flying particular types of sortie against us. Assume that the value to the enemy of flying those sorties is the same as the value we place on his *not* being able to fly them. Calculations similar to those of the previous section can be made of the net value of those sorties to him and the value of each of his aircraft.

OFFENSIVE COUNTER AIR

18. An OCA raid has 2 objectives: to destroy enemy aircraft on the ground and to close his airfields. Assuming that we know the extent to which we expect the raid to meet these objectives, the value of the aircraft destroyed can be derived from the previous section; airbase closure has the effect of pushing the sorties from the surviving aircraft to the right and the value of this can also be calculated.

19. This gives the gross value of the raid as a whole; the gross value of each sortie can be obtained by dividing the raid value by the number of aircraft. The net value of each sortie can be obtained by subtracting the probability of attrition multiplied by the aircraft value (at the end of the sortie).

20. This already provides an answer to the question of whether OCA is more profitable than other forms of offensive sortie.

DEFENSIVE COUNTER AIR

21. At this point it is necessary to re-iterate that this method is concerned with *marginal* values. Hence we must consider the effect of adding one more aircraft to the DCA effort currently planned. The value of this is threefold:

a. Enemy attrition will be increased. Since we have calculated the value of each of his aircraft, the utility of this is easy to calculate.

b. Own attrition will be changed. Adding another aircraft to the defensive force will often cause total defensive losses to be reduced; in this cases this is a benefit, whose value must be added in. In other cases, (such as when going from zero defending aircraft to 1) total defensive losses will increase; this should then be set off as a cost.

c. The proportion of enemy aircraft releasing weapons at their targets will be reduced. Since the value of preventing such attacks has been given, this too can be costed.

The difficulty throughout this paragraph is that the answers depend on what the enemy chooses to do. This is unknown in advance. Nevertheless, one can draw up a series of reasonable options, assign probabilities to them and calculate the expected value of the additional aircraft across that range of options. Note that, because of the way it has been calculated, this is already a net value, i.e. it takes account of attrition.

GROUND-BASED AIR DEFENCES

22. Using sections (a) and (c) of the above paragraph, one can calculate the gross value of ground-based air defence units (ADUs). Net value, taking account of the probability of destruction, is not a useful concept for such units: although they could in principle be withdrawn until a less dangerous phase of the campaign, this is not regarded as a sensible option. (Procedural switch-off may be regarded as a weaker form of this but is a matter of tactics rather than an operation-level decision.)

SEAD

23. This makes it possible to calculate the utility of SEAD missions. SEAD is a term which embraces 3 different activities; a particular SEAD mission can include variable quantities of each :

a. Destruction of ADUs The value of this can be based on the gross value of the ADUs destroyed.

b. Reduction of losses to accompanying aircraft. These aircraft have already been valued, so the benefit of reduced losses can be calculated directly.

c. Creating a Corridor This is more difficult to value. The idea is to reduce losses, over and above the reduction that would have occurred anyway had the same number of ADUs been destroyed but in a random manner over the whole front. It will be necessary to estimate the extent of this reduction, taking account of whatever ability the enemy has to move ADUs to close the gap or to concentrate his fighter aircraft as a temporary expedient to plug the gap. The additional reduction in losses can then be valued and divided between all the sorties participating in corridor creation to determine the additional utility of each sortie.

AEW & ADGE

24. The utility of AEW aircraft can be calculated in the manner used for DCA sorties, i.e. the AEW is credited with the extra effectiveness of the Air Defence system which results from its employment. A similar technique can be used for ADGE sites. Here we encounter the difficulty of redundancy: if 2 or more facilities provide the same function, then this method of incremental value is liable to attribute the whole value of the function to the first one added in, and nothing at all to the others. Yet actually the loss of value only occurs when all the facilities are destroyed. Whilst it would no doubt be possible to modify the mathematical technique to take

account of this, the number of attacks on such facilities is probably going to be sufficiently low for the application of common sense to keep the problem under control.

RECCE

25. The utility of reconnaissance sorties is even more tricky to calculate but the principles already enunciated continue to apply. First, what matters is the *marginal* sortie, the one that would be given up if one less aircraft were available. Secondly, it is necessary to consider what *action* was taken as a result of the information gleaned. Typically an offensive sortie will be diverted from a lower-value target to the one now discovered (or discovered to be still operating); in this case, the extra value of the sorties switched can be credited to the recce sortie. Of course, a lot of recce sorties will merely confirm prior beliefs; this is no doubt most gratifying but such sorties cannot be assigned any value added in this methodology. Hence recce turns out to be a little like prospecting in that the value of a sortie is the probability of producing worthwhile intelligence multiplied by the value of the intelligence when something worthwhile *is* produced. In mathematical terms, it is the *expected value* that we must use. It is a lot easier to estimate this in the course of a campaign when feedback is available than as part of the initial planning process.

Section 3 - A NUMERICAL EXAMPLE

26. Chapter 1 explained the concept of operational utility, and suggested a numerical method which could be used to assist in the apportionment of aircraft between different mission types. Some of the ideas presented were perhaps somewhat opaque however, and would benefit from the addition of a numerical example to further explain the processes involved.

27. This Chapter of the report therefore attempts to construct a worked example based on the framework suggested in Chapter 1.

SETTING

28. Whilst the example given is not intended to be a precise representation of events as they took place during operation Desert Storm, certain details pertaining to that conflict have been used at various points as the basis of the example figures. Hence the example is intended to be a representation of a conflict similar in character to Desert Storm.

29. Aircraft allocation among 4 sortie types will be considered in detail : STI (strategic target interdiction - attacks on strategic targets), CAS (close air support - attacks on ground force units), OCA (offensive counter-air) and DCA (defensive counter-air). The additional complications presented by SEAD, AEW and RECCE units (mentioned in Chapter 1) will not be considered here.

STI AND CAS

30. It is most convenient if these 2 mission types are considered together, as one needs to be rated relative to the other. At the start of the Gulf Conflict air campaign, emphasis was placed heavily on attacking strategic targets as a top priority - e.g. C^3 nodes and industrial infrastructure - and on gaining air superiority by targeting airbases (OCA - see later section). Certainly, CAS was afforded a low priority, given the generally defensive posture of the Iraqis and the fact that ground forces were not planned to engage for several weeks (in fact strictly speaking CAS was not a mission option at the start of the campaign).

31. Strategic targets such as C^3 facilities have high value to an enemy, since they are crucial to the effective operation of his forces. In order to apply the method of Chapter 1, it is necessary to rate their value relative to other targets numerically.

32. Suppose, for example, it is assumed that one STI target is worth 10 ground force units (e.g. tanks). This alone is not sufficient to compare the sortie types however. The next question

to ask is how many sorties will be required to neutralise the different targets. It is likely that a typical STI target will require more sorties to be flown against it than a CAS target - due to its relative hardness and/or size. Suppose then that to destroy a single STI target will require 10 sorties to be flown, whilst a single CAS target would require only 4. This gives the gross values of STI and CAS sorties, by simply dividing target value by the number of sorties required :

Sortie Type		**STI**	**CAS**
Target Value	**T**	10	1
Sorties Required	**s**	10	4
Gross Sortie Value	**G**	1.0	0.25

33. From this it is clear that the gross value of STI sorties is higher than that of CAS, but several other considerations must still be made in order to evaluate the net value - for example, the attrition rate expected. Prior to the start of the air war in the Gulf, it was anticipated that attrition rates would be around 1% during strategic attack and 0.5% during attacks on the Iraqi ground forces :

Sortie Type		**STI**	**CAS**
Target Value	**T**	10	1
Sorties Required	**s**	10	4
Gross Sortie Value	**G**	1.0	0.25
Attrition Rate	**p**	0.01	0.005
Survival Rate	**q**	0.99	0.995

34. The next parameter to consider is **r**, the so-called discount rate - defined in Chapter 1 as the extent to which sooner is better than later for a particular mission type. In the case of STI in this example, it is clear that early degradation of enemy C^3 and associated targets is very valuable to the campaign as a whole. Conversely, attacks on ground units, whilst also valuable long-term, are not an immediate priority and could be delayed until a later stage of the air campaign. Hence suppose the discount rate for STI is assumed to be 0.85, but 0.99 for CAS :

Sortie Type		STI	CAS
Target Value	**T**	10	1
Sorties Required	**s**	10	4
Gross Sortie Value	**G**	1.0	0.25
Attrition Rate	**p**	0.01	0.005
Survival Rate	**q**	0.99	0.995
Discount Rate	**r**	0.85	0.99

35. The final parameter to consider is **n**, the planning horizon - the number of sorties expected to be flown by each aircraft of type during the entire campaign. Before the start of the Desert Storm air campaign, coalition commanders expected their aircraft to average around 2 sorties per day and it was anticipated that the air campaign plan would take around 30 days to execute :

Sortie Type		STI	CAS
Target Value	**T**	10	1
Sorties Required	**s**	10	4
Gross Sortie Value	**G**	1.0	0.25
Attrition Rate	**p**	0.01	0.005
Survival Rate	**q**	0.99	0.995
Discount Rate	**r**	0.85	0.99
Planning Horizon	**n**	60	60

36. This completes the assumptions which must be made to be able to compare the net values of these 2 sortie types :

Sortie Type		STI	CAS
Target Value	T	10	1
Sorties Required	s	10	4
Gross Sortie Value	G	1.0	0.25
Attrition Rate	p	0.01	0.005
Survival Rate	q	0.99	0.995
Discount Rate	r	0.85	0.99
Planning Horizon	n	60	60
Aircraft Value	A	6.31	10.05
Net Sortie Value		0.95	0.20

37. The net value of a single STI sortie is 0.95, as opposed to 0.20 for a CAS sortie. So, from the assumptions made, STI is clearly the more valuable mission type at the beginning of the campaign.

38. Note however that the value of a single dedicated STI aircraft (6.31) turns out to be less than that of a dedicated CAS aircraft (10.05). This is a result of the different discount rates assumed for the 2 mission types. Although STI sorties initially have a high net value, this value quite rapidly tails off as the campaign progresses. CAS missions on the other hand, although initially of relatively low value, retain almost their full net value for the duration of the campaign. This is shown in the following graph :

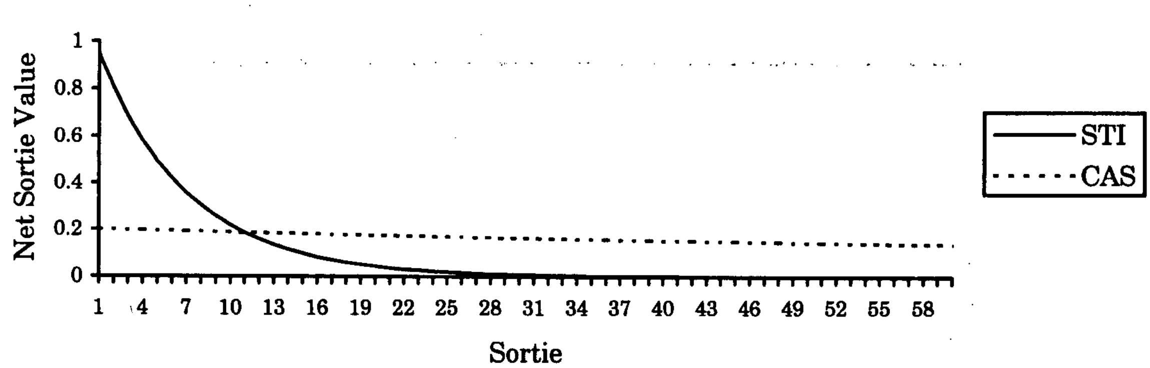

39. The value of an aircraft dedicated to a particular mission is represented by the area under its respective curve. Hence in the long run the CAS aircraft will be the more valuable.

40. The graph also suggests the benefit of utilising multi-role aircraft if possible. An aircraft which can be used for STI at the start of a campaign, but switched to CAS at a later stage, can have a greater value than either of the single-role types. Further discussion of this, and the way in which aircraft and sortie values will change during the course of a campaign is contained in a separate Annex to this report.

OCA

41. During the Gulf Conflict, 7585 OCA sorties were flown, and around 120 Iraqi aircraft destroyed on the ground or in their shelters. Suppose it is assumed here that, on average, a raid on one airbase will take 30 sorties, destroy 0.5 enemy aircraft and close the airbase for 2 days. In order to calculate the value of such a raid numerically, it is necessary at this point to consider what value the aircraft operating from that airbase had to the enemy.

42. Suppose the airbase was a base for 20 STI aircraft. During the Gulf Conflict, a successful attack by Iraqi aircraft on a coalition strategic target would have had great political value - as Western public opinion would have been very sensitive to any significant coalition losses. Hence it is reasonable to assume that strategic targets for the enemy will have, for example, 10 ten times the value as they hold for ourselves (i.e. 100.0). The same discount rate will be assumed. However, it will be assumed that the enemy would require double the number of sorties (i.e. 20) to destroy the target due to poorer intelligence, personnel and equipment.

43. During the Gulf Conflict, there was also a vast imbalance in the relative air strengths of the 2 sides - the coalition forces rapidly achieving almost total air supremacy. This leads to the further assumption in this example that the expected attrition rate for enemy aircraft flying STI missions will be considerably higher than for our own - say as much as 20% (40 Iraqi aircraft were shot down in the Gulf out of only 600 sorties flown - but this includes the 150 or so transits made to Iran, and no sorties were in fact offensive).

44. Assuming the same planning horizon for enemy STI sorties as for ourselves (though this has little effect on the results), the calculation of enemy STI aircraft and sortie values therefore looks as follows :

Sortie Type		STI (Enemy)
Target Value	T	100
Sorties Required	s	20
Gross Sortie Value	G	5.0
Attrition Rate	p	0.2
Survival Rate	q	0.8
Discount Rate	r	0.85
Planning Horizon	n	60
Aircraft Value	A	15.62
Net Sortie Value		2.34

45. Hence (returning to the initial assumptions in paragraph 41) one raid on an enemy airbase for STI aircraft destroys on average 0.5 aircraft (value 7.81), but also sets back any further sorties from that airbase by 2 days (i.e. by 4 potential sorties). So the value of the STI aircraft which are left on the ground will be reduced from 15.62 to 15.62 x $(0.85)^4$ = 8.15 (equating to an added value to us of 7.47). Note that this is actually only an approximation of the way in which the aircraft value declines over the course of a campaign. However, it is quite reasonable for STI sorties - when the decline is governed mainly by the discount rate alone (more detail is contained in the Annex). As there were originally 20 aircraft on the ground, the total value of the raid is therefore 7.81 + (19.5 x 7.47) = 153.48.

46. Assume the attrition rate for OCA is the same as for STI. Strictly speaking, a discount rate applies only to values set on ends, and OCA (and DCA) are means (paragraph 5) for obtaining an end (air superiority). However, the value of OCA (and DCA) sorties is based on the value of enemy aircraft destroyed (paragraph 41), which is in turn based on the value of their targets. Hence, assuming that the OCA targets are STI aircraft, successive OCA sorties will in effect be discounted at the rate of the targets for those STI aircraft - i.e. 0.85 in this case. The net sortie value can now be evaluated as for STI and CAS previously :

Sortie Type		STI	CAS	OCA
Target Value	**T**	10	1	153.48
Sorties Required	**s**	10	4	30
Gross Sortie Value	**G**	1.0	0.25	5.12
Attrition Rate	**p**	0.01	0.005	0.01
Survival Rate	**q**	0.99	0.995	0.99
Discount Rate	**r**	0.85	0.99	0.85
Planning Horizon	**n**	60	60	60
Aircraft Value	**A**	6.31	10.05	32.28
Net Sortie Value		0.95	0.20	4.84

47. This calculation indicates that the net value of a single OCA sortie at the start of the air campaign is 4.84 - over 5 times that of a STI sortie, and 24 times that of a CAS sortie.

DCA

48. The predicted value of DCA missions depends critically on what we expect the enemy to do. In the case of the Gulf Conflict, the Iraqis flew very few sorties at all, and in fact none of them were offensive, although this would not necessarily have been predicted prior to hostilities.

49. The value of DCA missions is threefold (Chapter 1). Firstly, enemy attrition will be increased. Given the high value attributed to enemy STI missions (paragraph 42) due to its great potential effect on Western public opinion, it will be assumed that any enemy sorties into friendly airspace would be STI missions. It has already been calculated (paragraph 44) that enemy STI aircraft have value 15.62 and enemy STI missions have net value 2.34. Suppose now that in 50 DCA sorties, it is predicted that one enemy offensive aircraft will be shot down. This therefore represents a value of those 50 sorties to us of 15.62.

50. Secondly, own attrition should be reduced. However, given that in this example the enemy is expected to fly very few offensive sorties, own attrition of DCA aircraft will be assumed to be zero in any case.

51. Thirdly, the proportion of enemy aircraft releasing weapons at their targets will be reduced. Suppose our 50 DCA sorties prevent another 2 enemy offensive sorties from being effective (additional value 4.68).

52. The attrition rate for DCA sorties is likely to be low (as they are carried out in friendly airspace) - say 0.1%. The "discount rate" for DCA can again (as in paragraph 46) be taken as the discount rate of the aircraft being opposed by the DCA missions - assumed to be 0.85 :

Sortie Type		STI	CAS	OCA	DCA
Target Value	**T**	10	1	153.48	20.3
Sorties Required	**s**	10	4	30	50
Gross Sortie Value	**G**	1.0	0.25	5.12	0.20
Attrition Rate	**p**	0.01	0.005	0.01	0.001
Survival Rate	**q**	0.99	0.995	0.99	0.999
Discount Rate	**r**	0.85	0.99	0.85	0.85
Planning Horizon	**n**	60	60	60	60
Aircraft Value	**A**	6.31	10.05	32.28	2.69
Net Sortie Value		0.95	0.20	4.84	0.40

53. This calculation demonstrates that, although the enemy is not anticipated to fly many missions, DCA still has a higher utility than CAS because of the great political value which any enemy success might have.

SUMMARY

54. At the start of the air campaign in the example presented here, the calculations clearly suggested that OCA had top priority (net sortie value 4.84), with strategic attack missions of next highest importance (0.95).

55. DCA missions were calculated to be of relatively low utility (0.40), since the enemy was expected to fly very few offensive missions. As there was no imperative requirement for them to be carried out at the start of the campaign, CAS sorties were also of low net value (0.20).

CONCLUSIONS

56. In Section 2, the concept of operational utility has been described, and a specific method for its application to a number of mission types has been outlined.

57. In Section 3, this method has been expanded on by applying it to a scenario loosely based on the situation of the Gulf Conflict. This example clearly demonstrates the hierarchy of mission utilities at the outset of such a campaign - with OCA being by far the most valuable, followed by strategic attack, DCA and CAS.

58. Annex A includes some more detailed calculation which describes the way in which the values of particular aircraft and mission types change over the course of an air campaign. In addition, an aircraft which is able to be switched from one mission type to another is demonstrated to be more valuable than a dedicated single-role type.

EMPLOYMENT OF MULTI-ROLE AIRCRAFT

A1. The discussions in the main body of this report were concerned only with single-role aircraft dedicated to particular missions, and concentrated on the values of these aircraft and their respective missions at the very start of an air campaign. The point was made however (and illustrated on page 6), that these values are dynamic during the course of the campaign.

A2. Suppose we define $G_i^{(m)}$ to be the gross value of sortie number *i+1* of mission type *m*. So given the figures used in this report, for example :

$$G_0^{(STI)} = 1$$
$$G_0^{(CAS)} = 0.25$$

A3. Successive values of $G_i^{(m)}$ are reduced by the discount rate. Hence, $G_i^{(m)} = \left(r^{(m)}\right)^i G_0^{(m)}$, where $r^{(m)}$ is the discount rate for missions of type *m*. This function is plotted for STI and CAS mission types in Figure A1 below :

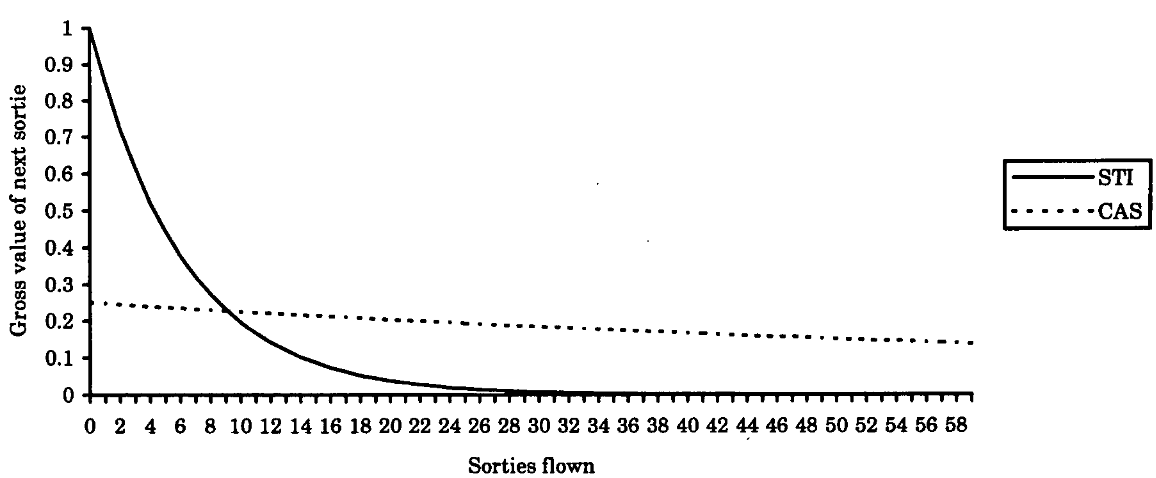

Figure A1

Note that, whilst the gross value of STI sorties starts off much higher than that of CAS, it quite rapidly drops below the CAS value.

A4. The value of a single-role aircraft at any time is the gross value of its next sortie *plus* the probability of it surviving that sortie times the remaining value of the aircraft after the sortie. Assume that a total of *n* sorties are planned for each aircraft. The value of a single-role aircraft after its n sorties have been flown is zero, if (as in the main text) the aircraft is assumed to be no longer of use. Define $A_i^{(m)}$ to be the value of a single-role aircraft dedicated to mission type m which has flown i sorties. Then :

$$A_n^{(m)} = 0 \text{ for each mission type } m$$

A5. Moving back one step from this point, the value of the aircraft after *n-1* sorties have been flown is therefore simply the gross value of the nth sortie :

$$A_{n-1}^{(m)} = G_{n-1}^{(m)}$$

A6. Now it is possible to iterate back to the starting position :

$$\begin{aligned} A_{n-2}^{(m)} &= G_{n-2}^{(m)} + q^{(m)} A_{n-1}^{(m)} \\ A_{n-3}^{(m)} &= G_{n-3}^{(m)} + q^{(m)} A_{n-2}^{(m)} \\ &\dots \\ A_0^{(m)} &= G_0^{(m)} + q^{(m)} A_1^{(m)} \end{aligned}$$

where $q^{(m)}$ is the survival probability of an aircraft on mission type m. The values of single-role STI and CAS aircraft are plotted in Figure A2 on the following page. Note that the CAS aircraft has a higher value throughout - since the gross value of its sorties suffers much slower degradation than the STI aircraft's sorties (recall Figure A1). Effectively, the CAS aircraft remains useful for longer than the STI aircraft - so it is more valuable.

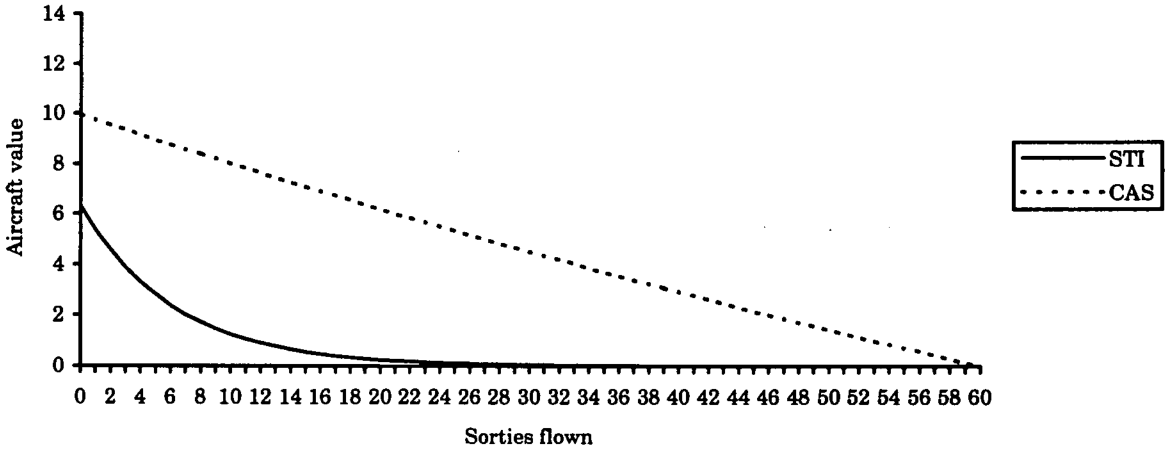

Figure A2

A7. For the multi-role aircraft the sequence is similar, but now it is possible to switch between the different mission types m between steps to maximise the value of the aircraft overall. If V_i is the value of the multi-role aircraft after it has flown i sorties, then :

$$V_n = 0 \text{ (as for the single-role aircraft)}$$

$$\text{but} \quad V_{i-1} = \max_m \left[G_{i-1}^{(m)} + q^{(m)} V_i \right] \text{ for all } i.$$

A8. The value of the multi-role aircraft has been added in Figure A3. Note that its value to begin with is higher than the value of both the single-role aircraft types.

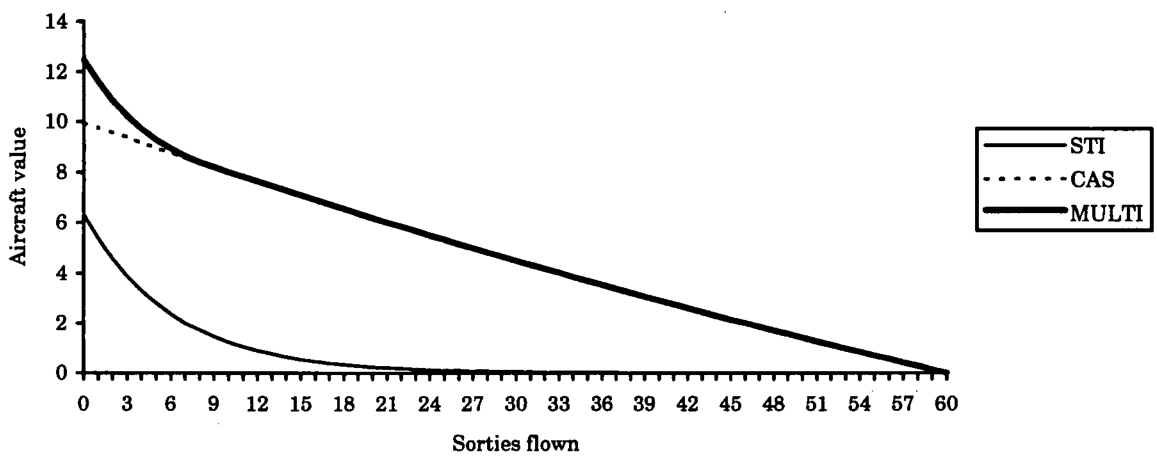

Figure A3

A9. The net value of each sortie can now also be calculated. The net value of sortie i is the gross value of the sortie *minus* the probability that the aircraft is lost on that sortie times the value of the aircraft after the sortie. For a single-role aircraft on mission type m this is equal to $G_{i-1}^{(m)} - p^{(m)} A_i^{(m)}$. For a multi-role aircraft flying the most profitable sortie of all mission types m this is equal to $\max_m \left[G_{i-1}^{(m)} - p^{(m)} V_i \right]$, where $p^{(m)}$ is the attrition rate for mission type m (and note that $p^{(m)} = 1 - q^{(m)}$).

A10. Figure A4 which appears on the following page shows the net values of STI and CAS sorties flown by the single-role aircraft, and also the net value of the most profitable sorties (chosen from these 2 types) flown by the multi-role aircraft.

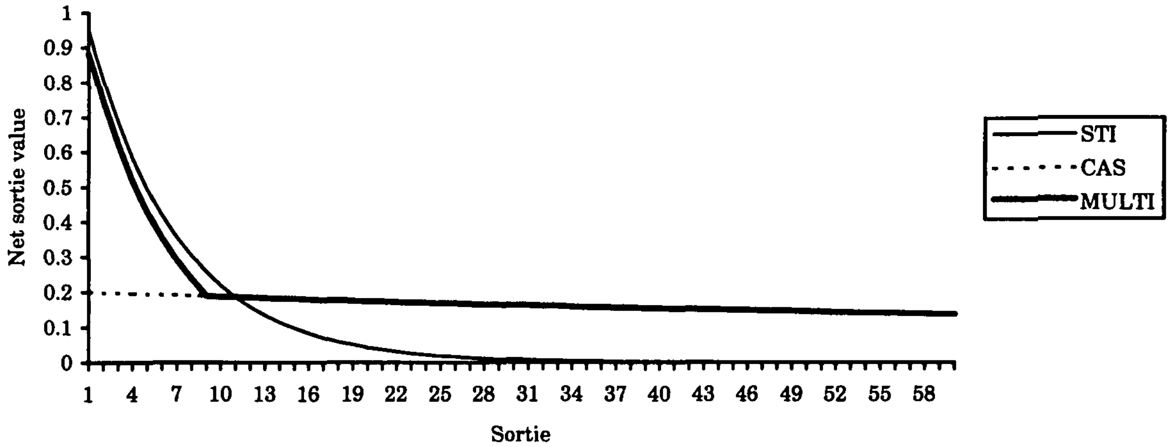

Figure A4

A11. The best use of the multi-role aircraft which can be used on either STI or CAS missions is therefore to fly STI for the first 9 sorties, but CAS thereafter. Note that the STI sorties flown by the multi-role aircraft are not themselves as valuable as those flown by the dedicated STI aircraft. This is because the multi-role aircraft itself is inherently more valuable than the STI aircraft (recall Figure A3). It would therefore be a greater loss than the STI aircraft if it failed to return from a mission - as its potential for flying CAS sorties in the future would then be lost.

A12. Once the multi-role aircraft is switched from STI to CAS (i.e. the final mission type to which it will be allocated), its value and the value of the sorties it flies becomes the same as for a dedicated CAS aircraft.

A13. As a final aside, it is worth considering the implications if an aircraft is assumed to have a residual value greater than zero after all its sorties have been flown during a campaign. For example, suppose a single-role STI aircraft has a residual value of 10 after completing its allotted 60 sorties, i.e. :

$$A_{60}^{(STI)} = 10$$

Preceding values can be calculated as before (paragraphs A5-A6), except that now :

$$A_{59}^{(STI)} = G_{59}^{(STI)} + q^{(STI)} A_{60}^{(STI)}$$

A14. The changing value of the single-role STI aircraft now takes the form shown in Figure A5 which follows :

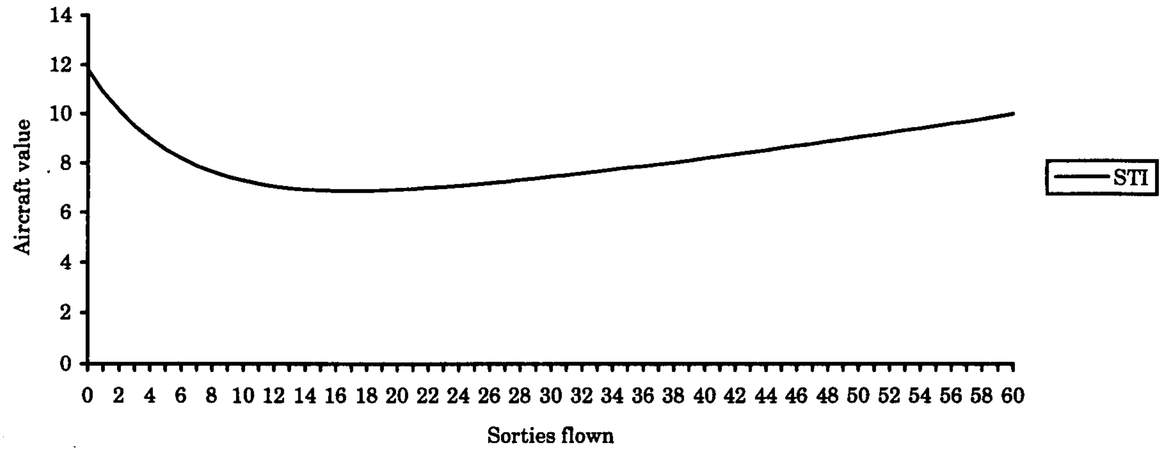

Figure A5

The most noticeable feature of this graph compared with the STI curve in Figure A2 is the trough which has appeared in the curve - the value of the aircraft decreases to a minimum (after its most useful work has been done), then starts to increase again up to its final residual value.

A15. The important effect of this on the net values of successive sorties is shown in Figure A6 below :

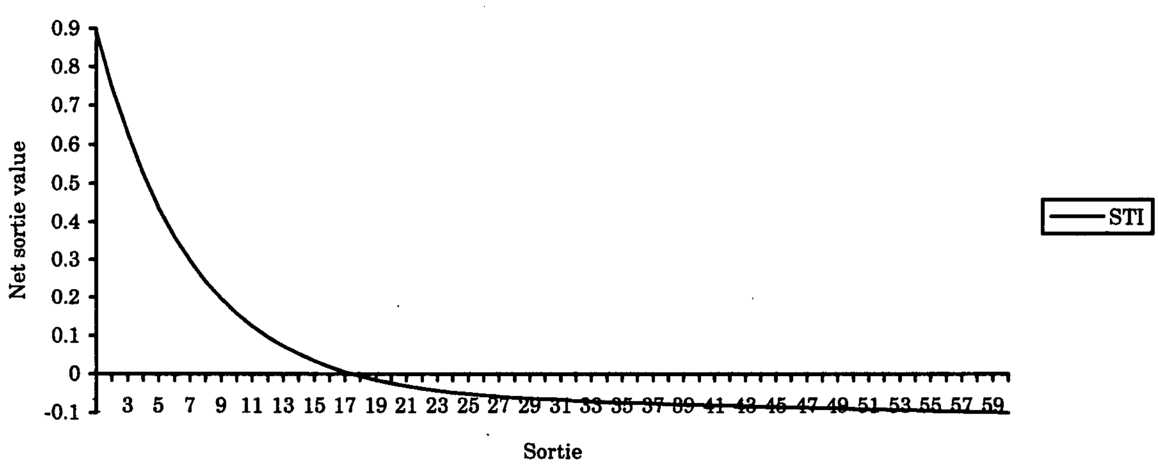

Figure A6

A16. Note that sorties from the 18th onwards actually have negative net value, as the remaining value of the aircraft is such that - despite the low attrition rate (1% in this case) - it is better to keep the aircraft on the ground rather than risk its loss for a very small return (Figure A1).