

17th INTERNATIONAL SYMPOSIUM ON MILITARY OPERATIONAL
RESEARCH (ISMOR)

A MODEL OF COMMAND & CONTROL AS AN OPTIMIZATION PROCESS
WITH IMPLICATIONS FOR C2 STRUCTURE

INTRODUCTION

Since the end of the Cold War, considerable changes have taken place in the command and control structure for UK armed forces. In connection with this, the doctrine of Command and Control (C2) has been examined in greater depth than ever before.

Much of this doctrine concerns the detailed organizational structure of Headquarters and the process by which key decisions should be reached. What stimulated this paper was a question about the extent to which current doctrine was driven by current Information Technology.

Operational Analysis studies have over the years addressed various aspects of C2, regarded both as a topic for study in its own right and as a key part of overall military capability. A related theme has been the value of Intelligence and how the quality of decisions is related to the intelligence available.¹ However, notwithstanding the general applicability of some of the analysis concepts developed, the author is not aware of any applications that address the relationship between HQ cells and tasks which forms the basis of so much doctrinal discussion.

AIM

This paper attempts to describe a major part of the C2 function as an algorithmic process. It thereby provides a theoretical basis for assessing whether a particular C2 structure is suited to its task, which can be used to complement historical or practical evidence. It has not yet been practicable to derive quantitative outputs from this but it does provide new insights into the connectivity required within a structure.

SCOPE

The whole of military activity is often characterized by the 'OODA loop': Observe, Orientate, Decide, Act. The strongest form of this model regards these as sequential activities within an endless cycle. This is perhaps an over-simplification of reality and is unnecessary for the purposes of this argument. One need merely regard OODA as a process in which on-going Observation feeds through to on-going Orientation which in turn feeds through to an on-going Decision-making process, which in turn directs Action. In this weaker form of the model, it is possible for some observations to feed through to action almost immediately while other types of observation take much longer to be

¹ See eg W Perry & J Moffat: *Measuring the effects of knowledge in military campaigns*. J Opl Res Soc (1997) 48 pp965-972.

digested and to influence actions. The important point is that the final deliverable is Action. A C2 process (and at this point I should probably be referring to a *C4ISTAR* process) may be judged according to the quality of the Action that results from it. 'Quality' here includes timeliness, appropriateness and effectiveness. 'Action' does not preclude *inaction* as a possible option. 'Action' also needs to be drawn more widely than mere movement of forces and the firing of weapons; it needs to include Information Operations, for example. Nevertheless, the important point to note is that C2 – not to mention Intelligence and all the other letters of the ever-expanding acronym – is not an end in itself but a means to an end. Action without cognition may be blind but cognition without action is fruitless.²

The decision-making process is something that extends from the senior commander – the Head of State even – down to every soldier, sailor and airman at the bottom of the chain. That is not to suggest that it is wise for the UK Prime Minister to become directly involved in the choice of words used by a peace-keeping patrol when it encounters some illegal activity, nor that SACEUR would wish to be involved in deciding where a Rapier Fire Unit should be sited. Nevertheless, the choice of what decisions should be delegated, not just to subordinate HQs but throughout, is part of the process of designing a C2 system; it is unhelpful to *define* the process so that some decisions must, of necessity, be devolved.

Notwithstanding this, it will be helpful to focus more tightly on the process by which high-level objectives are translated into orders issued by the lowest level of Component³ Headquarters. This neglects the process by which those orders are further translated into actions by processes at unit level. It also neglects some of the less concrete functions of command, the fostering of morale, for example. Control, understood as the process of monitoring and correcting the evolving execution of orders, will be neglected entirely. This narrowing of the scope is not intended to imply any lack of importance in these functions, rather that they do not appear to drive the arguments about C2 structure to the same extent as the process of generating orders. Many of the arguments that follow could easily be extended to the decision-making process at unit level and below but it will help in following the argument for the first time to concentrate on the narrower and more generally understood problem.

A MATHEMATICAL FORMULATION OF THE C2 PROBLEM

A military commander endeavours to employ the forces available to him in a manner best calculated to achieve the objectives he has been given. His problem is usually made more difficult by the existence of a rational, calculating adversary endeavouring to achieve different objectives which *he* has been given. These sets of objectives normally

² After Kant, I believe. Philosophers may disagree about whether cognition by itself is fruitless; however, military HQs do not exist for the furtherance of philosophy.

³ The work has been undertaken in the context of a deployed operation, in which (according to US and UK doctrine) a Joint Force Commander would have overall command, whilst the Land, Maritime and Air *Components* would each have a Component Commander with his own HQ. Many of the issues currently being addressed concern functions which might either be undertaken by cells working directly for the Joint Force Commander or might be delegated to one of the Component Commanders.

conflict, so part of the commander's task is to frustrate (or even to defeat) his adversary. There may even be multiple adversaries all pursuing their own conflicting objectives. The mathematician is accustomed to describe such a process as a *game*, potentially a multi-player, non-zero-sum game. A further complication is that the outcomes of many (if not of most) courses of action are unknown, though commanders have a good idea of the range of possibilities; the mathematician would probably wish to represent this as a *stochastic* game. A further complication still is that a commander generally has multiple objectives, including things he himself must avoid doing (otherwise known as *constraints*). Whilst there exist methods for converting these multiple objectives into a single measure of effectiveness (without which game theory cannot function) it complicates the mathematics and makes the whole process much less comprehensible to the military whom we seek to advise. Perhaps that is why useful applications of game theory to the operational level of warfare have been notable by their absence.

In practice, the problem is tackled in stages. Game theory features at the top level, though few commanders would describe it in those terms. Nevertheless, their reasoning does incorporate an assessment of how an opponent might exploit their actions and vice-versa. At the next level down, they have decided what they are trying to do and are examining the best means of achieving it. For example, they may have decided that as many as possible of a prioritized list of targets needs to be serviced⁴; they have limited numbers of offensive aircraft, limited numbers of the protective assets these aircraft require and limited air-to-air refuelling; they need to determine which aircraft from which bases should attack which targets and how they should be grouped into packages for their better protection. The problem may be more complicated still in that the most capable weapons may be limited in numbers and successful attacks on some targets (generally air defence ones) may be necessary precursors if the risk to aircraft attacking other targets is to be kept to an acceptable level.

This is not now a problem of game theory but one of optimization. Indeed various mathematical algorithms have been developed to address it, though so far all that can be done is to provide an optimal solution to *part* of the problem (linear programming was employed as early as the 1950s) or a *feasible* solution (not necessarily optimal and perhaps not even a very good one) to the full problem.

I have deliberately jumped to the air component level because it is a problem that is easily expressed in mathematical form. In general there may be aspects that need to be determined at the joint level which can also be viewed predominantly as a form of optimization, notably the division of tasks between components. One can think of similar optimization problems that will concern the Land and Maritime components.

Looking purely at the optimization aspects, problems as large and as unmanageable as this are not unknown in Operational Research literature: the scheduling of oil refineries is one that comes to mind. The approach normally adopted is to separate the problem into more manageable chunks whose mutual linkages are fairly weak. One then solves one of the problems approximately (often there is a natural starting point) and then feeds the

⁴ One might prefer to say 'hit' but the objective is actually to cause a defined level of damage (or perhaps in some cases merely to obtain imagery) so I will stick with the vaguer term.

initial answers from that problem into the others which can then be solved in turn. The first solution was frequently only approximate so one needs to feed back the results from the other chunks and iterate. If the problem has been broken down in an appropriate way, the whole process will generally converge; at least, the analysts hope it will!

For the military C2 problem, it is easy to see how such a break-down might be contrived. One would need a high-level model⁵ of the Joint part of the process, backed up by simple, approximate models of the individual components. Together, these would be used to produce a first stab at an optimization of the high-level problem. This solution would then be used as an input to the single component problems. Their solutions would be fed back to the high-level Joint model, which in turn would produce a second iteration. The whole process would then continue until a sufficiently good solution had been obtained (or until time ran out).

This is described as a hypothetical, mathematical approach to solving the problem. It is of course exactly what happens already: we call the high-level process the Joint Force Commander's Estimate and the lower-level processes the Component Estimates. Since the individual component problems may also be too large to handle, they in turn may also be treated as an iterative process: the Air Estimate leads to the development of a Master Air Attack Plan (MAAP) which, if it produces results significantly different from those expected in the Air Estimate, may require a new iteration of the Air Estimate, leading in turn to a new MAAP.

IMPLICATIONS FOR C2 STRUCTURE

So far, I have merely described what actually happens in more mathematical language. The mathematical reader may find this instructive; the military reader is wondering what purpose is served by all this.

The key to what follows is to note that, whereas the best way of structuring HQs is normally seen as a matter of doctrine, based on pragmatic considerations of what has been found to work in the past or subjective assessments of what is believed might work in the future, the question of whether an iterative algorithm converges to a good solution or not is an objective mathematical result.

Of course, there is no absolute requirement for the structure of HQs to match the structure of the algorithms they are implementing. However, if processes that require extensive iteration are divided between cells or (worse) between HQs, then the time to reach a satisfactory solution will be increased. This implies that a C2 architecture should be designed to keep the most intensively iterated parts of the algorithm within a single HQ and, if possible, within a single cell.

This actually admits a trivial solution: to carry out the entire decision-making process in one cell! That is of course impracticable: cells become inefficient when they contain more than a certain number of individuals; HQs not only become less efficient as

⁵ 'Model' here does not imply a computer-based or even a mathematical model: it may be merely a mental one.

their size increases but also more vulnerable to attack. The effect of evolving technology comes in here: if plans could be developed by fewer individuals one might expect to see a move towards a simpler C2 structure with fewer cells and less requirement for iterative processes to be split across cell boundaries. The problem with this is that the individuals within an HQ are not merely there to implement an algorithm; they also provide a rich repository of *knowledge*.

For example, the business of holding information on the capabilities of one's own forces is far more demanding (at least for the Air component) than might at first sight appear. This is in part caused by the existence of sub-fleets having different modification states within what is nominally the same Mark of aircraft, in part by the practice of having aircraft fitted *for* but not *with*, and in part by the fact that a squadron's crews may not all be current in particular activities. So far, producing a *credible* description of aircraft capabilities for coalition wargames has always proved too difficult; so producing an accurate computer-based description of the forces *actually* participating in an operation can only be regarded as an aspiration.

Knowledge also extends to an appreciation of what is militarily possible. Even human perception can have failings in this respect: the French commanders in 1939 were firmly convinced that the Germans either could not or would not advance through the Ardennes. But humans are at least good at intuitive leaps: they can recognize from limited evidence what their opponent might be trying to do – remember that conflict is fundamentally a stochastic game – and having made such intuitive leaps they will at least question their beliefs that what their opponent seems to be attempting should be quite impossible.

So whilst improved decision aids have the potential to reduce HQ sizes, the requirement for sources of expert knowledge – sources which are sufficiently involved with the evolving plan and are fully conversant with the Commander's intentions – will for some time limit the scope for such reduction.

SO WHAT? – A REAL APPLICATION

Even without quantifying this description, a number of useful implications emerge.

The first is that for any given military task there will be an optimum structure of HQs to command and control the units carrying out that task. This can, in principle at least, be checked objectively: a superior structure will produce a better solution to the problem of which units should undertake which tasks than will an inferior one. For most military tasks, the natural way of solving this problem, combined with the differing information requirements of the 3 environments will suggest the existence of a higher-level Joint HQ together with Component HQs for each environment (Maritime, Land and Air) – hardly a novel outcome.

However, there are areas where the insights offered by this approach do appear to be genuinely useful. One such area concerns issues which have traditionally been seen as the concern of the Air Component Commander but which have significant Joint aspects. Airspace co-ordination is one example. Land Component users increasingly want to do

things that impinge on airspace: firing artillery shells is a long-standing activity; flying helicopters and Unmanned Air Vehicles (UAVs) is a more recent innovation. Doctrine developers with a desire for intellectual tidiness then put forward the view that, since all 3 Components use airspace, the co-ordination of its use ought to be a Joint activity – and perhaps we should even call it "Battlespace" to avoid any implication that it is primarily a concern of the Air Component.

The traditional Air response to this would be that the Land Component is not a 'core' user, not a diplomatic response because 'non-core' carries various pejorative connotations. What is really meant by this response is better expressed using the concepts outlined in this paper: that within the complex optimization process at the heart of airspace co-ordination, the feedback loops associated with Land Component uses are usually weak. In fact, if Land wants to fire shells or fly UAVs within a volume of airspace, the chances are that the request can be agreed without difficulty. In contrast, the feedback loops associated with mainstream aircraft uses are frequently much more involved, with decisions on who can use a volume of airspace quite likely to affect whether a user can fulfill his mission and potentially requiring modifications to the overall air plan. These feedback loops thus bind airspace co-ordination firmly into the overall air planning process: any attempt to move it to the joint arena will result in an increase of planning time or, alternatively, a solution which is further from the optimum. In contrast, because the feedback loops to the other components are relatively weak, shortening the organizational span of those loops produces no improvement in the speed or quality of solution.

MORE GENERAL IMPLICATIONS

There are other areas besides C2 structure where the optimization model of Command & Control provides useful insights.

The first is that the form of relationship between higher and lower HQs that was normal 50 years ago, in which the higher HQ issued orders and the lower simply implemented them by issuing its own orders, will normally produce a sub-optimal outcome because it will only represent the first pass of an iterative algorithm. What is needed is for the higher HQ to seek the advice of the lower on the implications of issuing particular orders, and then to use those results to refine its own decisions. In some cases, and subject of course to there being sufficient time, more than one iteration of this may be needed. Again, the idea of a *dialogue* between adjacent levels of HQ has been part of doctrine for some time; however, its description as an algorithmic process may provide greater clarity.

A second insight is of direct concern to the Operational Analysis community: higher levels of HQ need their own aggregated models of the lower HQs' functions in order to improve the accuracy of their first-stab solutions and reduce the amount of iteration needed to provide a good solution. Again, some HQs have been working along these lines for a while now but it is instructive to view the process as a way of improving algorithmic efficiency rather than as an imperfection adopted to get around the shortage of time.

A third implication is for the way in which HQs should evolve. At present, their processes are often labour-intensive. That can hide the fact that their staffs are often needed to provide specialist advice as well as to work the system. In the interests of efficiency, their systems should, and no doubt will, become less labour-intensive with the result that their size will be dominated by the provision of advice. Some of this advice, ought in principle, to be able to be called up from a database; some ought to be obtainable remotely. However, it must not be forgotten that brilliance in planning has often depended on 'breaking the rules' - more precisely on knowing that certain constraints were not actually as rigid as they were widely supposed to be. Over-reliance on computer databases rather than personal knowledge will inhibit both the formulation of such 'brilliant' plans and the recognition of the viability of similar plans if the enemy starts to execute them. One might suppose that the 'reach-back' concept has the capability to overcome this problem; however, it is by no means clear that staff remote from the commander they are advising will feel able to step outside the standard answer. This would perhaps be a useful analysis objective for a synthetic experiment.

CONCLUSION

Viewing the C2 process as an algorithm for deciding what functions should be undertaken by what units provides a number of useful insights. These provide greater clarity concerning relationships between HQs and the direction in which the structure should evolve as technology develops. Quantitative results on the extent to which one structure is superior to another are difficult to obtain, not least because of factors such as the decline of efficiency with cell size which themselves have not hitherto been quantified.

R C WHEELER
CDA/FLD and Air Warfare Centre, RAF Waddington
email: RCWheeler@dera.gov.uk

© Crown Copyright 2000

