Proceedings of the 2000 Winter Simulation Conference
J. A. Joines, R. R. Barton, K. Kang, and P. A. Fishwick, eds.

REPRESENTING THE C2 PROCESS IN SIMULATIONS: MODELLING THE HUMAN DECISION-MAKER

Colin R. Mason

CORDA Ltd.
Apex Tower, 7 High Street
New Malden, Surrey KT3 4LH, UK

James Moffat

CDA(HLS)
(Centre for Defence Analysis (High Level Studies))
DERA
Farnborough, Hampshire GU14 0LX, UK

ABSTRACT

Military Command and Control (C2) is the process by which commanders organise and employ force elements in order to achieve military objectives. This process needs to be represented in models of conflict in order to simulate realistic force behaviour and effectiveness. Since C2 is heavily influenced by human decision-making, modelling the C2 process is recognized as one of the most challenging areas for defence analysis. This paper describes on-going research into ways in which the effects of C2 can be incorporated successfully into constructive simulation models of combat. The research has developed a representation of C2 based on an intelligent agent framework in which the C2 processes of a military operation are carried out by a number of interacting command agents, representing the various military headquarters. An agent software architecture has been designed and implemented, along with algorithms for key C2 processes, namely (1) data fusion and recognized picture compilation, (2) decision-making and planning at operational and tactical levels, (3) plan supervision and repair. These have been implemented within two software testbeds – MOSES, a simulation of Operations Other Than War (OOTW) and CLARION+, an extension of CLARION, the UK land-air combat simulation.

1 INTRODUCTION

The work reported here was conducted under the UK DERA Corporate Research Programme, Technology Group 11 TA2.1: ‘C2 in OA Models’, during the period April 1997 to March 2000. This programme considers the representation of Command and Control (C2) in constructive simulations of conflict. Here, C2 is taken to include all of the processes associated with information collection and fusion, creation of a perception of the situation, and developing a course of action based on that perception.

The problem of how to represent C2 in simulations of military conflict has been the subject of research within NATO nations for many years. Much expenditure has been made, yet with little progress. The reason for persevering with such work is that issues related to C2 are becoming central to many questions of interest to defence policy makers. Across NATO, including the UK, there is a growing realisation that the proper representation of C2 within models of conflict is very important (NATO RTO 1999). Some of the reasons for this are:

- To show cost-effectiveness of investment in C2 systems.
- To support programmes such as the UK’s Digitisation of the Battlespace, which requires continuing underpinning by Operational Analysis (OA).
- The need to represent C2 in order to model realistic overall force behaviour and effectiveness. This includes the need to incorporate emerging understanding of the impact of the human decision-maker on operational effectiveness.

In consequence, this research was instigated to investigate ways in which the effects of C2 could be incorporated successfully into constructive simulations of conflict.

Since C2 is heavily influenced by human decision-making, modelling the C2 process is recognized as one of the most difficult areas for defence analysis. The terms of reference for this research (Grainger 1995) recognized the high risk, but potentially very high payoff, of work in this area, noting: “As research, the work may well fail to come up with anything useful. However, taking a less pessimistic view, I would like to see the following results: (1) Ideally, working products in the form of software implemented as part of a testbed. (2) Or at least, if (1) is too ambitious, theoretical papers defining designs which stand a good chance of working.”

The work reported here demonstrates that we have achieved the goal of working products in the form of software, implemented as part of a testbed, together with theoretical papers supporting these.

2 GENERAL APPROACH

2.1 Holistic and Evolutionary Development

We adopted a holistic and evolutionary approach to conducting our research. We decided that the whole C2 process should be represented, however crudely, from the start and that as the research progressed, various parts of the representation would be refined as our understanding of those parts grew.

At each step of the development the research ideas were implemented within a number of software testbeds. Two of these are MOSES, a simulation of Operations Other Than War (OOTW) and CLARION+, an extension of CLARION, the UK land-air combat simulation. Both testbeds are object-oriented simulations implemented in C++ and running on PCs under Windows NT. The testbeds enabled us to check the feasibility of emerging theoretical ideas and provided us with the means of demonstrating progress, in the form of working products, to our study sponsors.

2.2 Underpinning Theoretical Ideas

Combat is a complex activity. Given this, it is natural to assume that the command and control of such activity will itself require a complex representation. Such thinking lies behind many previous efforts in modelling C2. A popular approach has been the incorporation of extensive rule-based systems (also known as embedded expert systems or knowledgebase systems) into simulations. These systems attempt to capture a commander's decision-making process by a (usually very large) set of interacting decision rules. In practice, however, these approaches have met with limited success. The resulting models have been found generally to be slow running, complex to understand and difficult to modify, and to require great effort to migrate from one scenario to another (Sharma 1996).

Recent developments in Complexity Theory (Nowak and May 1992, Alberts and Czerwinski 1997) suggest a possible alternative way forward. The essential idea is that a number of interacting entities, behaving in accordance with *small* numbers of *simple* rules, can generate extremely complex emergent collective behaviour. Our research aim is to develop a representation of military C2 based on such sets of simple rules and entity interactions that will give rise to emergent collective behaviour that resembles realistic military behaviour.

Our approach is outlined in the following paragraphs. The key components are then described in further detail in the remaining sections of the paper.

2.3 Command Agents

The command and control of a military operation is carried out by the various decision-making entities of the military force. These entities are the headquarters (HQs). The HQs interact with one another within some kind of network (usually, but not necessarily, a hierarchy of some kind) reflecting the military C2 organisational structure of the force.

In our C2 model these HQs are represented by 'command agents'. By command agent we mean an autonomous entity with the ability to:

- Sense its local environment and construct an internal representation – a perception - of the external world.
- Plan its own behaviour, based on the current perception of the external world.
- Exchange information with other agents.

The idea is that a number of these command agents interacting with one another will represent the total C2 process.

One aspect of the research was to look at how to represent, and implement in software, these command agents in as generic and re-usable a way as possible. The result of this work is OACIS (Object Architecture for C2 In Simulations) – a design and software implementation of a generic command agent. This is described in Section 3.

2.4 Command Structure

Previous work (Alberts and Hayes 1995), based on extensive historical analysis, has identified a number of different command arrangements that span the major approaches to the C2 of armed forces. The work shows that the various command structures observed historically can be captured by a combination of top-down and bottom-up command styles. This is illustrated in Figure 1. Here 'Order Specific' corresponds to a very top-down oriented command style in which every order is issued by the supreme commander of the forces, and subordinate units are allowed little or no freedom to use initiative. This is relaxed at the next level, where specific objectives are decided by the higher command, but subordinate units are allowed some freedom to decide for themselves how to achieve the objective. 'Mission Specific' refers to an even more open command style in which subordinate commanders have full autonomy to 'self-synchronize' with other subordinate commanders in order to achieve broad mission goals (comprising a number of objectives necessary to accomplish the mission), with little interference from higher command. Recent UK doctrine has moved more to the bottom end of this spectrum, as 'mission command'. From a historical perspective this interaction between lower and higher levels of command is well discussed by van Creveld (van Creveld 1985) in his consideration of command in maneuver warfare, particularly as applied by Napoleon, and by the German army in WW2.

Figure 1: Alternative Command Styles

What emerges from the above discussion is that any representation of the C2 process must incorporate these top-down and bottom-up effects, and their mutual interaction. This requirement forms the basis of our C2 representation, which consists of a two-level command structure, illustrated in Figure 2. At the lower level are the maneuver units of the military force – the units that do the actual fighting. In our CLARION+ testbed these are brigade-sized close-combat entities. At the higher level is the force commander, responsible for the overall direction of the force's activities. The levels interact via orders (mission assignments) that are issued by the high level commander and directed downwards to the subordinate units, and via status/situation reports that are issued by the subordinate units and directed upwards to the high-level commander. At the lower level the units interact with one another via the (common) environment.

2.5 Planning Processes

Units at both levels of the command structure are represented by command agents. However, we represent the command decision-making process differently at the two levels. This reflects a key difference between planning at the top and at the bottom of the C2 structure – the amount of time available for planning. At the lower level the command agents implement a decision-making strategy we call 'rapid' planning. This is based on the naturalistic decision-making paradigm and reflects empirical observations of how expert decision-makers operate under stress in situations where time is short. The rapid planning model is described in Section 4. At the higher level the command agents implement a decision-making strategy we call 'deliberate' planning. This is based on the classical rational choice decision making paradigm involving explicit evaluation of a number of alternative courses of action, leading to a choice which is in some sense optimal. This style of decision making is appropriate when ample time is available. The deliberate planning model is described in Section 5.

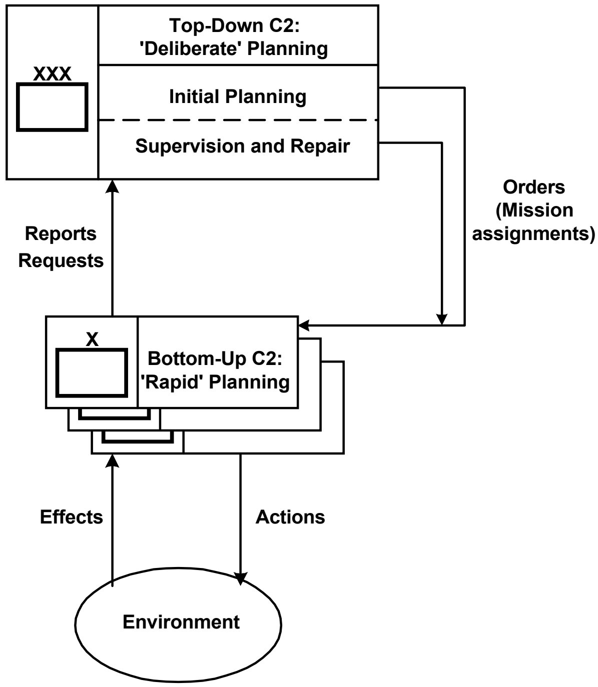

Figure 2: Two-Level Command Structure

Overall, the two-level command structure attempts to capture the following characteristics of a military operation. At the higher level, deliberate planning by the force commander is focused on overall campaign objectives. It aims to produce an initial force layout that gets the subordinate units into the right place at the right time. It seeks to ensure that first contact with the enemy occurs under the most favourable conditions possible. After this, it is a matter of responsiveness and opportunism on the part of the subordinate units, driven by the tactical (local) situation in which units find themselves as the campaign unfolds. This sensitivity to the local situation is captured by each unit's rapid planning process. Feedback closes the loop between the two levels of decision-making: status/situation reports from subordinate units together with reports from strategic sensors deployed by the force commander. The feedback permits a top-down supervision of the overall progress of the campaign and enables periodic adjustment of tactical activities (via issue of new orders to subordinates) to maintain direction towards the overall campaign objectives.

The remaining sections of the paper describe the command agent representation, and the rapid and deliberate planning models, in further detail.

3 COMMAND AGENT ARCHITECTURE

3.1 Aim

The aim of this part of the research was to design a software architecture for a command agent representing a military HQ. We wanted an implementation that was generic, that is, the *process structure* of C2 within the HQ should be the same, wherever the HQ is located in the command hierarchy. The representation is therefore recursive as a function of the different levels of command, reflecting the fractal nature of command (Dockery and Woodcock 1993). However, the way in which the C2 functions are carried out by HQs will differ, dependent on their role. So, at the same time, we wanted an implementation that was extensible to accommodate specialised, role-dependent *process content*. The result of this work is OACIS (Object Architecture for C2 In Simulations) – a design and software implementation of a generic command agent.

3.2 Scope

The scope of the OACIS development was to define a command agent architecture that captures the key C2 processes, and their interactions, which exist in a military HQ. The key processes are the G2 (Intelligence) and G3 (Operations) activities of data fusion, recognized picture compilation, decision-making and planning. These cover the activities concerned with:

- evolving a perception of the outside world from sensor and situation reports (data fusion);
- developing a mental model of what is going on (the recognized picture);
- deciding what to do next and formulating a plan to achieve this (decision making and planning).

3.3 OACIS Command Agent

The structure of the OACIS command agent is shown in Figure 3. In the figure, the boxes represent the principal components of the command agent. The directed lines between boxes show the main information flows between the components. The components are described in the following paragraphs.

3.3.1 The Comms

This provides communication facilities allowing the command agent to exchange various types of information (orders, reports and requests) with other command agents.

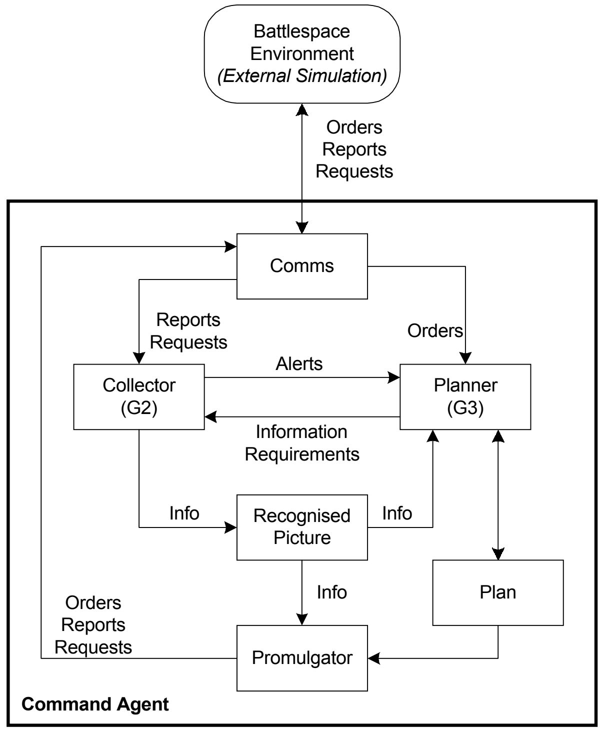

Figure 3: Command Agent Structure

3.3.2 The Collector

This encapsulates the G2 (Intelligence) processes of data collection (directed at achieving the Commander's Critical Information Requirement (CCIR)), data fusion, maintenance of the recognized picture and intelligence assessment. The Collector is responsible for alerting the command agent's Planner component should significant events, or developments in the current situation, occur.

3.3.3 The Planner

This encapsulates the G3 (Operations) processes of command decision-making and planning. On the basis of the recognized picture, the Planner creates (and maintains through subsequent supervision and repair) the plan that will enable the command agent to achieve the mission assigned to it by higher authority. The Planner establishes the information needed to support the planning process – the CCIR – and delegates the Collector component to collect this via deployable sensors.

3.3.4 The Promulgator

This encapsulates administrative processes for managing the output of the command agent. It creates messages of various types, including orders (to promulgate the plan to subordinates), status/situation reports and requests for information. The messages are passed to the Comms component for transmission to the appropriate recipients (other command agents).

3.3.5 The Recognized Picture

This is the information store that contains all of the command agent's knowledge of the external world. It is structured as a collection of geographical zones, each of which corresponds to an area of military interest. It represents the command agent's perceived state of the world.

3.3.6 The Plan

This is the Planner's output. It defines the mission that the command agent is to undertake, together with the missions that are to be assigned to subordinate agents. These missions are drawn from a small set of allowed missions, based on military doctrine.

3.3.7 Dynamic Behaviour of Components

We developed a formal process model for each of the command agent components. A process model specifies the sequencing of, and data flows between, key sub-process activities performed by a component and captures the dynamic behaviour of the component. The full set of process models (one per component), and their mutual interaction, constitutes the total C2 process of the command agent.

3.3.8 Software Implementation

We implemented the OACIS command agent in software in the form of an extensible object-oriented framework, using the Template Method design pattern (Gamma et al. 1995). This approach has enabled the *structure* of the command agent's C2 process to be captured once in a set of generic framework object classes. We can then create role-specific command agents using subclasses, derived from the framework classes, which override and specialise the functional implementation of selected processes.

The OACIS command agent architecture and implementation has been proved in our MOSES testbed. This model is a simulation of a services-assisted evacuation operation. Using the OACIS architecture we successfully implemented command agents representing the 11 distinct military HQ roles that were needed to conduct the operation.

4 RAPID PLANNING

4.1 Naturalistic Decision-Making

The rapid planning process corresponds to what UK army doctrine refers to as 'battle command' (Army Doctrine Publication Vol. 2: Command 1995). This is the style of decision-making most appropriate at the tactical level where the ratio of battle speed to C2 speed is generally high, that is, the time available for decision-making is short. Under these conditions empirical evidence indicates that commanders adopt more 'intuitive' approaches to decision-making termed 'naturalistic decision-making' (Noble 1999, Brander 1994). These approaches conform to Klein's recognition-primed decision (RPD) model of the decision-making process (Klein 1989), applicable to expert decision-making under stress. Discussion with Klein (Moffat and Catherall 1998) has confirmed the applicability of this model to our problem.

The RPD model is as follows (Brander 1994). In essence, the process begins with the decision-maker considering the situation or problem and trying to recognize familiar aspects of it. If this can be done, he is very likely to retrieve a satisfactory response from his repertoire and will then check this solution by mentally simulating its progress. The process can be viewed as a form of pattern matching, where the current perceived situation is compared with a set of mentally stored references (which have been accumulated by experience and training). The best match then indicates a potentially feasible course of action.

To capture the essence of the RPD model, our rapid planning process is based upon pattern matching, where the patterns are directly linked to possible courses of action.

During development of the rapid planning model a joint US/UK review (Moffat and Catherall 1998) confirmed the need to consider the idea of an 'OK' and 'not-OK' situation. Klein in particular made the point that commanders have a general sense of how things are going, which is captured by the idea of OK/not-OK mission states. Commanders take the information they have and weave it into a plausible story (the OK state). And they monitor the situation as it develops in order to assess when the story is beginning to fall apart (indicating that the commander is approaching the boundaries of the OK state). While the perception of the situation is such that the commander is in the OK state, he remains in his current mission. When the perception is that the situation has changed significantly, the commander crosses the boundary of the OK state and has to decide whether to remain with his current mission or change to a new mission.

In summary, our rapid planning model is as follows:

- Quantify the current values of the factors that constitute the perception of the situation.
- Determine whether the perception of the situation has changed significantly. If it hasn't, then the situation is OK and no change of mission is required.

- If the perception of the situation is changing significantly (that is, we are moving into the not-OK state), compare the pattern corresponding to the perceived situation with a set of fixed patterns which represent the commander's stored experience.
- Find the pattern that best matches the perceived situation.
- Identify the course of action (mission) that is linked to this pattern.
- Decide whether or not to implement the indicated change of mission.

Further details of each of these steps are given in the following paragraphs. This model has been implemented as the decision-making process used by the close-combat entities (command agents) within our CLARION+ software testbed.

4.2 Quantifying the Perception of the Situation

The perception of the situation is the commander's recognized picture. This will, in general, be defined by a number of attributes. In warfighting, a key one is the perceived combat power ratio (PCPR), that is, the perceived force ratio in the commander's local area of interest. At present, a command agent running our rapid planning model uses PCPR alone to characterise the perceived situation. Future work is planned to investigate extending this representation to include additional situational parameters that are likely to be important influences on decision-making. In warfighting, for example, the commander's logistics state is another important factor. A totally different set of parameters will be needed for OOTW situations such as peacekeeping operations.

4.3 OK or Not-OK

The commander (implemented as a command agent) assesses whether the current perceived situation is still OK (no action required by the agent) or not-OK (a change of mission may be required) by analysing the time history of observations of PCPR in the agent's local area of interest. The analysis tool used is the Dynamic Linear Model (DLM) (West and Harrison 1997). This is a mathematical structure for modelling and analysing time series processes.

The behaviour of the PCPR is modelled by a pair of DLM class II mixture models (West and Harrison 1997). One mixture model tracks the enemy combat power values whilst the other mixture model independently tracks own force combat power values. Each mixture model comprises four separate DLMs: a constant level DLM, a transient change DLM, a level change DLM and a growth change DLM. These four models correspond to hypotheses in the commander's mind about what is going on with respect to the time development of the enemy (or own) combat power. The constant level DLM describes a time series of observations with no significant variation in the level of combat power. The other DLMs represent situations in which there is significant deviation from a constant level – transient behaviour, significant changes in the level of combat power, and significant changes in the slope (rate of change) of combat power.

At any given time a probability can be calculated for each DLM in the mixture model. This probability measures the relative likelihood that a particular DLM is the model that best describes the time series of observations seen to date. These probabilities are updated on a continuous basis, driven by observations of combat power obtained via the command agent's sensors.

At any given time the best estimate of the mean and variance of combat power (enemy or own) is obtained from the DLM which has the highest probability within the (enemy or own) mixture model. Combining combat power estimates for both enemy and own forces yields the observed distribution of PCPR. This is our representation of the perceived situation.

The boundary of the commander's OK state is crossed when the probability of the constant level DLM drops and the probability of one of the change DLMs rises. Thus, tracking these probabilities within both DLM class II mixture models can be used to estimate when the commander is approaching the boundaries of his OK state, and needs to consider what to do about this.

4.4 Find the Best Match

If the perception of the situation has changed significantly (that is, we have moved into the not-OK state) then the command agent needs to compare the pattern corresponding to the perceived situation with a set of fixed patterns representing the commander's stored experience. In our current model, the observed and fixed patterns consist of normal distributions of PCPR, each defined by a particular mean and a variance. The amount of overlap between the observed pattern and each of the fixed patterns is used to estimate the likelihood of each of the fixed patterns at the current time. The likelihood is used in a Bayesian updating scheme to derive a posterior probability for each of the fixed patterns. This is the probability that the given pattern matches the perceived situation at the current time. The command agent then selects the pattern with the highest posterior probability – the best match.

4.5 Identify the Associated Course of Action

The pattern matching process selects one of the fixed patterns as the one that best matches the perceived situation. Each of the fixed patterns is linked to a particular course of action (a mission). In the spirit of mission command we link the patterns to a small set of missions, such as ‘advance’, ‘attack’, ‘defend’, ‘delay’, ‘withdraw’. Thus selecting a pattern - recognising the situation – leads directly to a course of action (mission) for the command agent. We call this the agent’s ‘preferred’ mission.

This mapping between pattern and course of action captures the essence of Klein’s RPD model of decision-making.

4.6 Decide on Change of Mission

The final step in deciding whether or not to change the current mission is deciding whether such a change is both feasible and desirable, taking account of all relevant factors. Our current model takes account of two key factors.

The first factor is the influence of the top-down command process. This is captured by defining a temporal constraint on how long the command agent is allowed to deviate from the mission ordered by the agent’s superior commander. This is done via an integer, n, specifying the number of C2 cycles for which the command agent is allowed to deviate from its ordered mission. Once n C2 cycles have elapsed since the agent was last in its ordered mission the agent must return to the ordered mission. The agent is then required to remain in the ordered mission for a number of C2 cycles defined by a second integer, m. Both n and m are user-inputs. Note that if we set integer n to zero then we have a total top-down command style: subordinate command agents cannot deviate from their ordered mission, ever. As n is set to larger and larger values we have a progressively more bottom-up command style. The values n and m can thus be used to ‘tune’ the command style to represent any of the styles given in Figure 1.

The second factor is the uncertainty in the currently perceived situation. The effect that we want to capture is that if the perceived situation is uncertain then a commander will wait and seek further information rather than change his mission. This is implemented in the following way. On each C2 cycle the command agent calculates a probability of changing to the preferred mission. This is called the transition probability. Provided that the temporal constraint described above does not forbid a transition, the agent changes to the preferred mission, on the current C2 cycle, with that transition probability. The transition probability is calculated as the difference between the highest and next highest fixed pattern posterior probabilities. This captures the desired effect. If the preferred pattern stands out well from its surroundings, and there is thus little uncertainty as to what the situation is, then the transition probability as defined above will be large and a transition to the preferred mission will be favoured on this C2 cycle. Alternatively, if the preferred pattern does not stand out well from its surroundings, and there is thus uncertainty as to what the situation is, then the transition probability will be smaller and a transition to the preferred mission will be less likely. The overall effect is that the command agent will tend to make mission transitions quickly when there is little uncertainty but will be more reluctant to change as uncertainty increases.

5 DELIBERATE PLANNING

5.1 Analytical Decision-Making

The deliberate planning process corresponds to what UK army doctrine refers to as ‘high command’ (Army Doctrine Publication Vol. 2: Command 1995). This is the style of decision-making most appropriate at the strategic and operational levels of war. Here, the ratio of battle speed to C2 speed is lower than it is at the tactical level and there is generally more time available for the commander to assess the situation. Under these conditions commanders use a decision-making process which is more analytical than the naturalistic approach, often referred to as ‘rational choice’ decision-making.

In rational choice decision-making the emphasis is on the explicit generation, and subsequent evaluation, of multiple courses of action. A decision criterion is specified and applied to the course of action evaluations to determine the ‘best’ option, which is then selected as the preferred course of action. The selection of a course of action is the command decision and is the output of the rational choice decision-making model.

5.2 Aims of Deliberate Planning

The deliberate planning process has two aims. Firstly, to produce an initial layout of own forces that gets subordinate units into the right place at the right time. Secondly, to monitor the overall situation as the battle unfolds and make adjustments to the layout as necessary in order to maintain direction towards the campaign objectives.

5.3 Operational Context

In our CLARION+ testbed the deliberate planning process is carried out within a command agent representing the force commander of a ‘side’ in a conflict. The agent is given a set of objectives (geographical locations) and a mission (currently restricted to one of attack or defend) to be conducted against these objectives. We model a two-sided conflict in which the force commander on one side is tasked with attacking each of his objectives whilst the commander on the other side is tasked with defending each of his objectives. The objectives of each commander need not be the same. Each objective lies at the end of an axis. This is illustrated in Figure 4 for two sides labelled Red and Blue.

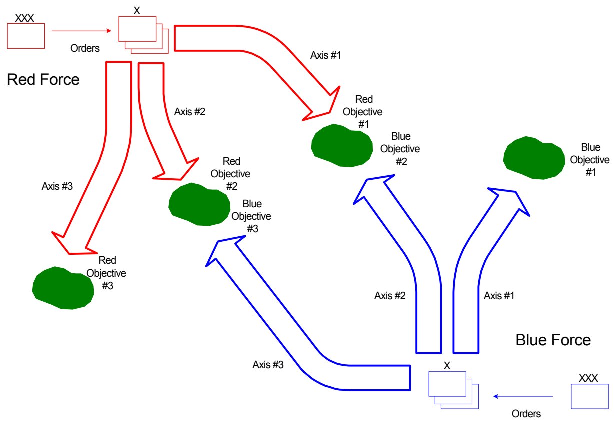

Figure 4: Deliberate Planning – Operational Context

Each force commander has decided (by some meta-level decision process not yet represented in our model) to employ ground and air forces to undertake their respective missions. The ground forces are brigade-sized close combat entities, modelled by command agents running the rapid planning process. Air forces are represented parametrically by a stated number of air sorties allocated per day to each axis. The planning problem to be solved is then: What number of air sorties and ground units should be allocated to each axis? It is the deliberate planning process that generates an 'optimal' force allocation across the axes, that is, a deployment plan.

When the plan has been created the force commander promulgates this, via orders, to the subordinate units. Each order specifies the particular objective to which the recipient is to deploy. On receipt of their orders the subordinate units (command agents) deploy to their objectives, along the axes, as indicated by the arrows in Figure 4.

The subsequent behaviour of the subordinate command agents is governed by their individual decision-making processes, modelled using the rapid planning process. These agents provide important feedback to the force commander via situation and status reports. It is this feedback that enables the force commander to carry out supervision and repair of the initial deployment plan, as the situation evolves. These interactions are illustrated in Figure 2.

5.4 Planning Model – An Application of Game Theory

The core of the deliberate planning process model is based on ideas from game theory - the mathematical theory of decision-making in conflict situations. Game theory was chosen as the starting point because the theory addresses one of the central elements of the deliberate planning process, namely the analysis of opposing courses of action.

In the following, 'planner' refers to a command agent that is performing deliberate planning. The planner is pitted against an opponent - the 'enemy'.

The deliberate planning model is based around analysing a game payoff matrix. The general structure of such a matrix is shown in Figure 5. The rows (columns) of this matrix represent different courses of action available to the planner (enemy). O_i denotes the i^{th} course of action available to the planner and E_j denotes the j^{th} course of action available to the enemy. We define a course of action to be a particular (ground and air) force allocation to each of the planner's axes.

	E_1	E_2	...	E_S
O_1	P_{11}	P_{12}	...	P_{1S}
O_2	P_{21}	P_{22}	...	P_{2S}
:	:	:	:	:
O_M	P_{M1}	P_{M2}	...	P_{MS}

Figure 5: Game Payoff Matrix, $\underline{P}$

It is important to understand that in this usage of game theory the E_j represent only the planner's *perception* of the courses of action that the enemy could follow. Thus, the E_j need not necessarily reflect what the enemy is actually contemplating doing nor necessarily contain the course of action that the enemy will actually take. The quality of the E_j, in terms of how well they predict future states of the conflict, depends on the ability of the planner to divine the enemy's intentions.

The interactions of the opposing courses of action are represented by the contents of the matrix - the payoffs, P_{ij}. P_{ij} is the payoff (more precisely, the *perceived* payoff, from the planner's perspective) from the enemy to the planner that will occur if the planner takes course of action O_i and the enemy takes course of action E_j. The payoff can be viewed as a measure of effectiveness (MoE) of a given pair of opposing courses of action.

The planners on each side of the conflict have a separate (and generally different) payoff matrix, representing each planner's *perception* of the possible courses of action open to himself (the O_i) and his opponent (the E_j), and the consequences of the interactions between them (the P_{ij}).

The essence of the deliberate planning model is the analysis, by the planner, of this payoff matrix and the selection of a single course of action, O_i, that is, in some sense, the 'best' one to take, given the perceived options open to the enemy. The selection of a course of action is the command decision and is the key output of the deliberate planning process model.

There are several different ways of defining the 'best' course of action, depending on the criteria used to measure 'bestness'. Four such 'decision' criteria are the criterion of pessimism (maximin), the criterion of optimism (maximax), the criterion of least regret and the criterion of rationality. The deliberate planning model uses the first of these criteria - the criterion of pessimism (also known as the Wald criterion). Use of the Wald decision criterion results in a payoff matrix analysis process that represents a conservative decision-making approach in which the planner looks for the (own) course of action which offers the best guaranteed payoff.

5.5 Implementation

The above ideas have been implemented in our CLARION+ testbed. For a command agent running the deliberate planning process, the implementation comprises three interacting sub-models: an intelligence fusion model, a plan generation model, and a plan supervision and repair model. These are described below.

5.5.1 Intelligence Fusion

The game-theoretic core of the deliberate planning process model requires the planner to analyse interactions between feasible enemy courses of action and possible own force courses of action. In the absence of any further information the simplest assumption that the planner can make is that the feasible enemy courses of action are all equally likely.

The aim of the intelligence fusion process is to improve on this assumption by enabling the planner to estimate the relative likelihood of each feasible enemy course of action. This is based on sensor observations of the battlespace, as follows.

The planner manages a collection of strategic sensors, which he deploys across the battlespace with the aim of seeking out the presence of enemy forces. The sensors feed detection reports back to the planner from which the planner builds his perception of the battlespace, that is, his recognized picture. The planner compares this perception, in turn, with each of the feasible enemy courses of action. The relative likelihood of each feasible enemy course of action is estimated from the closeness of the match between the course of action and the observed enemy force distribution in the battlespace.

5.5.2 Plan Generation

The plan generation model creates the (own force) course of action, that is, ground and air force allocations to each of the planner's axes, which offers the best guaranteed payoff.

5.5.2.1 Payoffs

Payoffs are calculated using a set of algorithms based on historical analysis (Ferguson and Blues 1997). Given a hypothetical plan – that is, a specified force allocation to the planner's axes - these algorithms allow rapid calculation of several metrics, namely: attacker rate of advance on an axis, attacker and defender casualties, the probability of attacker breakthrough on an axis, and the probability of ultimate attacker success on an axis.

The payoff is formulated in terms of a weighted sum of these metrics. By changing the values of the weighting factors it is possible to change the emphasis that the planner puts on particular metrics. For example, in some situations minimising casualties might be more important that the speed of the advance; in other situations the opposite could be true.

5.5.2.2 Finding the Best Course of Action

In principle, finding the best course of action - that is, the one that offers the best guaranteed payoff - requires the planner to evaluate all elements of the payoff matrix. For efficiency reasons, however, it is not desirable in practice to do this. Instead we take a different approach, as follows.

We explicitly enumerate only the enemy courses of action (the E_j). Given that these are the planner's perceptions of what the opponent might do, we can argue that these options can be forced to be relatively few in number. For the own force courses of action, we do not explicitly enumerate all of the options (the O_i). Instead, we employ a search technique, implemented via a genetic algorithm, to seek out the particular own force allocation, O_i, that maximises the expected payoff over all of the feasible enemy courses of action, E_j.

The output of the plan generation process is an own force course of action consisting of a force allocation (number of ground units and number of daily air sorties) to each axis and an allocation to the reserve force. The reserves are the surplus units, if any, left over after achieving satisfactory force allocations to each of the axes. The reserves are held for possible deployment to axes later in the battle via the plan supervision and repair process.

5.5.3 Plan Supervision and Repair

The purpose of the plan supervision and repair model is to monitor the situation resulting from the execution of the plan and, when necessary, to adjust the deployment of subordinate units to try to maintain direction towards achievement of the planner's mission.

The model can be viewed as a cybernetic feedback loop. The initial deployment plan, created by the plan generation model described above, is executed by issuing the appropriate orders to the subordinate units. This results in the (initial) force deployment to axes. Time then elapses. On a regular basis the planner receives feedback (observations of enemy and own force strengths in the battlespace) from strategic sensors and from the subordinate units.

Periodically, the planner performs a situation assessment. In this, the planner uses the feedback from the battlespace to re-assess the effectiveness of the plan. If the assessment indicates that the plan is failing, the planner calculates adjustments to the deployment that will improve the plan. For example, units could be brought up from reserve to bolster a weak axis; or units could be transferred from strong to weak axes to provide a more balanced deployment. The planner implements these adjustments by issuing new orders to the affected subordinate units.

6 CONCLUSION

This paper describes an approach to representing C2 in constructive simulations of conflict. It is based on an intelligent agent framework in which the C2 processes of a military operation are carried out by a number of interacting command agents behaving in accordance with simple rules.

We have described the agent software architecture that underpins our representation, and algorithms for the key C2 processes of decision-making and planning at operational and tactical levels, and plan supervision and repair. These have been successfully implemented in our MOSES and CLARION+ software testbeds.

The question as to whether our representation of C2 generates emergent collective behaviour that resembles realistic military behaviour is still an open one. Preliminary results are encouraging and we are about to embark on a more formal military validation of model behaviour.

This research is still a work-in-progress. We believe we have made some promising advances towards a workable representation of C2 for OA models of conflict. But we also know that this is far from complete. We are about to start a second phase of the research programme to develop further the initial ideas reported here.

ACKNOWLEDGMENTS

The authors would like to acknowledge the contributions of Dr. John Catherall (DERA, CDA(HLS)) and Dr. Lorraine Dodd (DERA, Land Systems Sector) to the development of the ideas discussed in this paper.

REFERENCES

Alberts, D. S., and T. J. Czerwinski. 1997. Complexity, Global Politics and National Security. National Defense University, Washington DC.

Alberts, D. S., and R. Hayes. 1995. Command Arrangements for Peace Operations. National Defense University, Washington DC.

Army Doctrine Publication Vol. 2. 1995. Command. Army Code No. 71564, Ministry of Defence, UK.

Brander, G. 1994. Decision Making Behaviour Observed During High Level Wargames. (Unpublished report.) DERA, Farnborough, UK.

Dockery, J. T., and A. E. R. Woodcock. 1993. *The Military Landscape*. Cambridge: Woodhead Publishing Ltd.

Ferguson, N., and R. C. Blues. 1997. The NEMO Wargame. Technical Report CDA/HLS/R9750/1.0, DERA, Farnborough, UK.

Gamma, E., R. Helm, R. Johnson, and J. Vlissides. 1995. *Design Patterns: Elements of Reusable Object-Oriented Software*. Reading, MA: Addison-Wesley.

Grainger, P. L. 1995. C3 Research in Support of OA Model Development. Loose Minute D/DCS(S&A)/7/16/1 dated 31 October 1995.

Klein, G. A. 1989. Recognition Primed Decisions. *Advan in Man-Machine Sys Res* 5:47-92.

Moffat, J., and J. Catherall. 1998. US/UK Review of TG11 C2 Research Project. (Unpublished report.) DERA, Farnborough, UK.

NATO RTO. 1999. Technical Report 9 - Code of Best Practice (COBP) on the Assessment of C2. RTO-TR-9 AC/323(SAS)TP/4.

Noble, D. 1999. Minutes of Decision Modelling Workshop, March 23-25, 1999. Evidence Based Research Inc., Washington DC.

Nowak, M., and R. May. 1992. Evolutionary Games and Spatial Chaos. *Nature* 359.

Sharma, W. 1996. An Overview of C2 Representation in IMAGE and GEKNOFLEXE. Technical Report DERA/ CDA/HLS/WP028/2/1, DERA, Farnborough, UK.

van Creveld, M. 1985. *Command in War*. Cambridge, MA: Harvard University Press.

West, M., and J. Harrison. 1997. *Bayesian Forecasting and Dynamic Models*. 2nd Edition. New York, NY: Springer-Verlag.

AUTHOR BIOGRAPHIES

COLIN R. MASON is a Principal Consultant in CORDA Ltd. He received his BSc and PhD from University College, University of London. He has spent 16 years in CORDA conducting a variety of operational research studies. His current research interests include object-oriented simulation, complex adaptive systems, decentralised systems and emergent behaviour. His email address is <colin.mason@corda.co.uk>.

JAMES MOFFAT BSc PhD FOR is a Fellow of DERA, a Fellow of the Operational Research Society and visiting Professor to the University of Cranfield. He has spent 20 years in defence-related operational research, advising the UK Ministry of Defence on the best balance of investment for defence equipment. He has also led research on stealth technology and sensors. His current research interests are in the improved modelling of C2 in the next generation of models of conflict. His email address is <jmoffat@dera.gov.uk>.