THE INTERNATIONAL
C2 JOURNAL

VOLUME 2, NUMBER 2, 2008

SPECIAL ISSUE
Representing Human Decision Making in Constructive Simulations for Analysis

GUEST EDITOR
James Moffat
Dstl

A Minimum Spanning Tree Approach to Identifying Collective Behaviour and Inferring Intent for Combat Models

Aleem Hossain
Nick Walmsley
Paul Pearce

CCRP

THE INTERNATIONAL C2 JOURNAL

David S. Alberts, Chairman of the Editorial Board, *OASD-NII, CCRP*

The Editorial Board

Berndt Brehmer (SWE), *Swedish National Defence College*
Reiner Huber (GER), *Universitaet der Bundeswehr Muenchen*
Viggo Lemche (DEN), *Danish Defence Acquisition and Logistics Organization*
James Moffat (UK), *Defence Science and Technology Laboratory (DSTL)*
Mark Nissen (USA), *Naval Postgraduate School*
Ross Pigeau (CAN), *Defence Research and Development Canada (DRDC)*
Mink Spaans (NED), *TNO Defence, Security and Safety*

About the Journal

The International C2 Journal was created in 2006 at the urging of an international group of command and control professionals including individuals from academia, industry, government, and the military. The Command and Control Research Program (CCRP, of the U.S. Office of the Assistant Secretary of Defense for Networks and Information Integration, or OASD-NII) responded to this need by bringing together interested professionals to shape the purpose and guide the execution of such a journal. Today, the Journal is overseen by an Editorial Board comprising representatives from many nations.

Opinions, conclusions, and recommendations expressed or implied within are solely those of the authors. They do not necessarily represent the views of the Department of Defense, or any other U.S. Government agency.

Rights and Permissions: All articles published in the International C2 Journal remain the intellectual property of the authors and may not be distributed or sold without the express written consent of the authors.

For more information

Visit us online at: www.dodccrp.org
Contact our staff at: publications@dodccrp.org

A Minimum Spanning Tree Approach to Identifying Collective Behaviour and Inferring Intent for Combat Models

Aleem Hossain, Dr. Nick Walmsley, Paul Pearce (Dstl, UK)

Dstl reference number: Dstl/JA27316
© Crown copyright 2008. Published with the permission of the Defence Science and Technology Laboratory on behalf of the Controller of HMSO.

Abstract

The Defence Science and Technology Laboratory (Dstl) makes use of many models in order to model military conflict. Most of these models are based upon a Cold War era opponent, and it is argued that there is a need to update these models in order to reflect the evolving structure of the UK armed forces needed to meet the current threats. Algorithms are presented which will allow the planning processes within military models to make more informed decisions. The approach uses a mathematical model to identify agents which may be working together as a group, and subsequently make inferences about their intent. This research is carried out within the context of a wider programme concerned with updating existing simulation models of conflict. Initial results are presented and the future development of the work is discussed.

Introduction

The Defence Science and Technology Laboratory (Dstl) is part of the UK Ministry of Defence (MoD) and routinely provides analytical support and advice to decision makers on policy, procurement and operational issues. In order to conduct studies of this nature, Dstl makes use of many models in order to model military conflict.

The term *du jour* in modern military operations is 'asymmetric warfare.' That is, warfare in which the weaker force uses unconventional weapons and tactics in order to try and neutralise a stronger opponent. As a result, UK and coalition forces are having to adapt their doctrine and Command and Control (C2) structure in order to remain effective against these new threats. Consequently, there is a need for military models to reflect this shift in battlespace paradigm.

One such model is the Wargame Infrastructure and Simulation Environment (WISE) (Pearce et al. 2003) which is a Formation Level wargame used as an operational analysis (OA) model. WISE can operate both as a wargame, in which military players are used to make decisions, and as a simulation which operates without human input in closed form. This makes WISE suitable for both experimentation and analysis of doctrinal issues (in wargame mode) and for performing analysis on issues such as procurement options (in simulation mode) (Holt 2006). In order to operate realistically in simulation mode, WISE utilises two planning and decision-making processes: The Deliberate Planner which operates at the Strategic and Operational Levels of command, and the Rapid Planner which operates at the Tactical Level (Moffat 2002). This article details work concerned with improving inputs to the Deliberate Planner within the wider context of a three-year study developing models to reflect modern C2 (Holt et al. 2007).

A novel approach to identifying collective behaviour in opponents using a technique based on Minimum Spanning Trees (MSTs) will be described. Having hypothesised which agents may be working together in a group, a second algorithm is presented which gives an indication of the intent of the groups with respect to a number of possible destination locations. The algorithms provide a more accurate picture of the size and structure (i.e. groupings) of the battlefield agents, and in doing so, provide the Deliberate Planner with better information on which to base its decision-making.

Prototypes of these algorithms have been written and tested, and it is intended that these will now be implemented in WISE for further validation and verification. Whilst WISE is used as a test-bed, the aspiration is for a generic mathematical algorithm which can be embedded in a variety of models.

The article proceeds as follows. First, the emerging concepts of Network Enabled Capability (NEC) and Agile Task Organised Groupings (ATOGs) will be defined in the context of the Command and Control structure of the UK's armed forces. This will be contrasted with the more traditional hierarchical military command structure, and in doing so, a need to develop new modelling techniques to accurately represent ATOGs will be presented. Next, the key modules required to reflect these concepts are described briefly. In particular, the Deliberate Planner is introduced and the limitations of its current implementation are outlined. Next, research is described concerned with addressing these limitations by identifying collective behaviour on the battlefield and inferring intent in order to better inform the planning process. Finally, the results of the first phase of this work are presented, and future research is outlined.

Throughout this article, enemy forces will be referred to as *Red* and friendly forces as *Blue*.

Emerging Concepts

An emerging concept within the Command and Control (C2) structure of the UK's armed forces is that of ATOGs:

The Cold War era was dominated by a predictable, monolithic and structured enemy using tactics that could be templated. This was reflected in a rigid and hierarchical command structure, separated

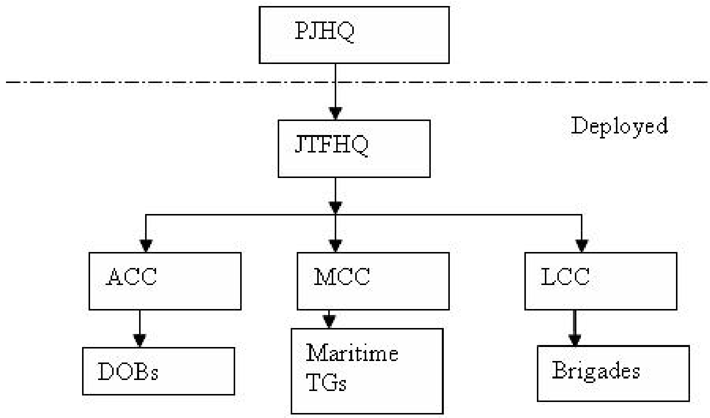

Figure 1: Traditional C2 Structure

by component (Air, Maritime and Land) with minimal communication between them.

Today, and for the foreseeable future, UK and coalition forces will be presented with a much more complex battlespace paradigm with several opponents, agents and actors (some of which will adopt asymmetric tactics). Success will depend on the ability to share a comprehensive view of the battlespace and to allocate battle assets rapidly as required to meet the immediate threat. In this spirit, force structures will need to be dynamic; the inflexible, rigid command structure will need to be replaced with dynamic, cross-component task groups which can be assembled and dissolved at all levels in the command hierarchy according to the immediate task needs. That is, task groups will be optimised for specific missions or tasks. This dynamic re-allocation of assets will make use of cross-component collaboration to match requirements with capabilities, in contrast to the traditional hierarchical military command structure. The resulting task groups are ATOGs.

Figure 1 (taken from Moffat (2006)) illustrates the traditional C2 structure: The Permanent Joint Headquarters (PJHQ), located at

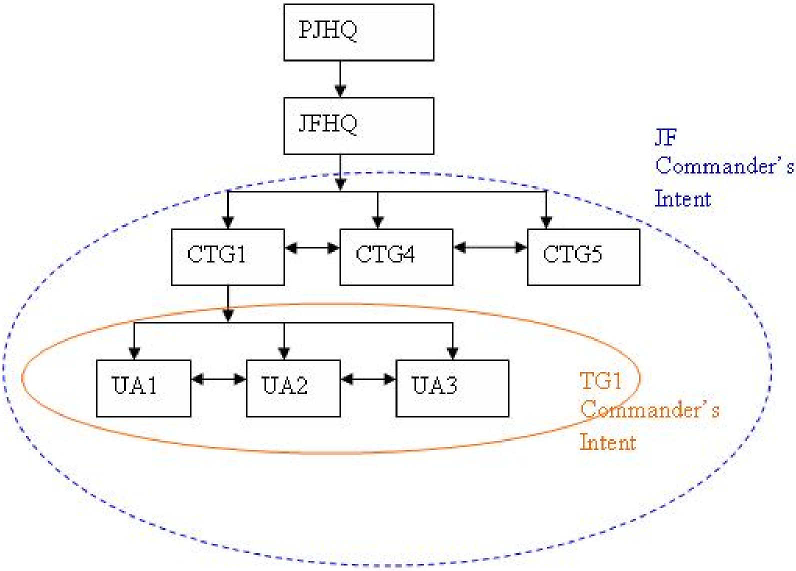

Figure 2: Future C2 Structure

Northwood in Middlesex, commands overseas joint and combined military operations; the Joint Task Force Headquarters (JTFHQ) is deployed to an operational theatre. Beneath the Joint Task Force Commander are a number of Component Commanders (CCs)[1]. Each of these Component Commanders (Air, Maritime and Land) has a number of formed units under their command (e.g. Deployed Operating Bases, Maritime Task Groups and Brigades). Requests for support outside of a particular component's capabilities go up and down the command hierarchy with minimal horizontal communication and interaction.

In Figure 2 (taken from Moffat (2006)), CCs are replaced by Task Groups consisting of elements from all components. These Comprehensive Task Groups (CTGs) are formed by the Joint Force Commander (JFC). CTGs are formed and dissolved according to the JFC's intent and desired effects. CTGs can interact horizontally

1. Special Forces and Logistics components are not shown on the diagram.

to create a common picture and understanding. The same process also occurs at the level of the Units of Action (UA) under command of the CTGs. These UAs can exchange "components of force" (Moffat 2006) within the ellipse bounded by the CTG Commander's Intent as shown in the diagram.

The concept of ATOGs falls within the wider framework of Network Enabled Capability (NEC), which encompasses the ways in which people, information and networks can be integrated to improve the sharing of information and situational awareness[2], and to increase command agility. It has already been argued that these changes are required in order for UK and coalition forces to meet current and future threats. NEC is defined in the Defence Industrial Strategy (Ministry of Defence 2005) as:

> " ... the coherent integration of sensors, decision-makers and weapon systems along with support capabilities ... [leading to] better situational awareness across the board, facilitating improved decision making, and bringing to bear the right military capabilities at the right time to achieve the desired military effect."

In order to conduct network-centric operations, forces need to be able to interoperate; that is, work together though communication and information sharing. Alberts and Hayes (2003) summarise:

"Entities that are not interoperable or have limited interoperability will not have access to all available information, will not be able to provide information to entities that may need it, and will be limited in the ways that they can collaborate and work together with others. As a result, their value (ability to contribute to combat power or mission effectiveness) will be limited over time."

---

2. "The understanding of the operational environment in the context of a commander's (or staff officer's) mission (or task)" (Ministry of Defence 2006).

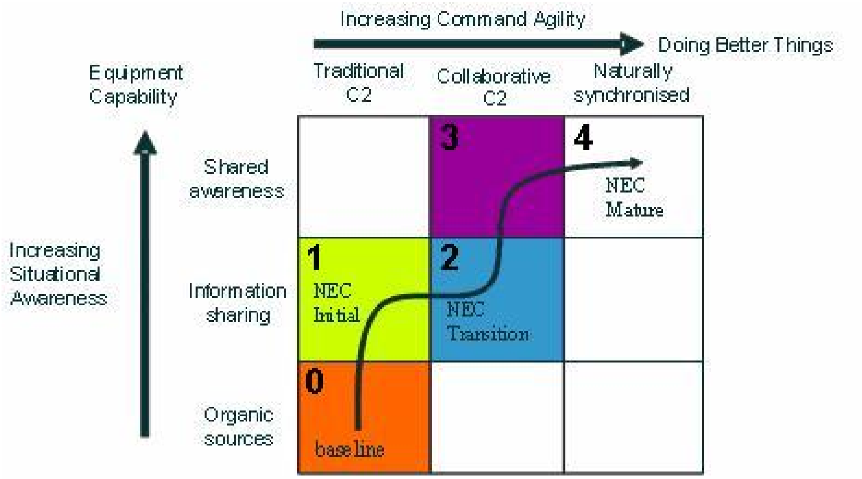

Figure 3: The NEC Journey along the epochs of NEC

Figure 3 shows the Alberts' Grid (originally known as the Network-Centric Warfare (NCW) Maturity Model (Alberts and Hayes 2003)) which divides NEC into three epochs: Initial, Transition and Mature. These levels of network-centric capability relate to the degree to which interoperability has been achieved: whilst Level 0 requires limited interoperability and information sharing, Level 4 (NEC Mature) requires greater richness, reach and quality of interactions. In order to reach this zenith, there must be advancements in organisational structure, work processes, attitudes and technology. Currently, the UK can be considered to be at Level 2: NEC Transition. The epochs are defined in terms of the completeness and quality of situational awareness (a comprehensive view of the battlespace) brought about through technology interventions; and the organisational/command structure (of which ATOGs are a part). It is the command structure aspect of NEC which is the focus of the research presented here.

Moffat (2007) identified an area where improvement can be made to Dstl's current combat modelling capability; namely how to represent, in simulation modelling terms, the difference between NEC Transition and NEC Mature. To progress the full distance of the

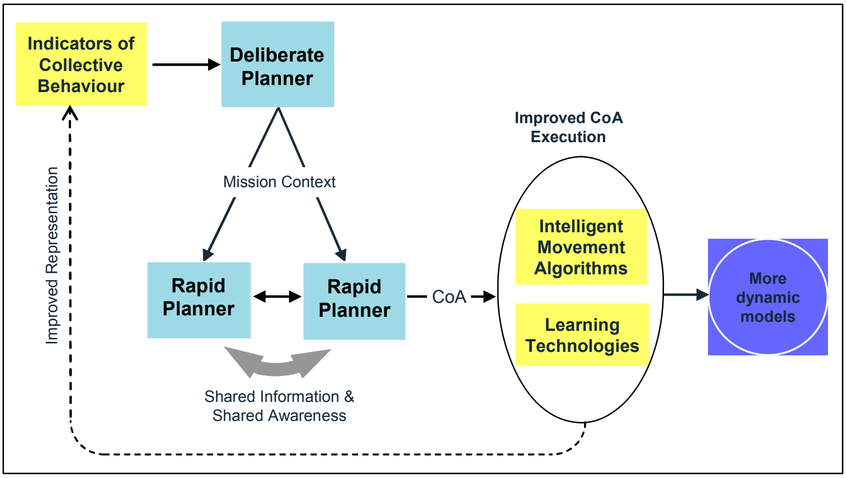

Figure 4: System Overview

NEC Journey to NEC Mature, there is a requirement to develop a set of computer algorithms which can be used by a range of simulation models in order to capture this key aspect of NEC and assess its likely benefits. There is a requirement, therefore, to represent ATOGs and the more dynamic doctrine and C2 processes in combat models in order to capture the shift in battlespace paradigm from the *Industrial Age* to the *Information Age* (defined in, for example, Alberts and Hayes (2003)).

Capturing NEC and ATOGs

Dstl is currently in the second year of a three year study (Holt et al. 2007) concerned with modelling the C2 associated with ATOGs using WISE (Pearce et al. 2003) as a test-bed. The focus of the study as a whole is to develop a more agile command structure to support analysis of the UK MoD's aspirations of Network Enabled Capability (NEC) (Ministry of Defence 2005). Figure 4 illustrates how the study areas integrate together to achieve this.

The system can be split into two phases: planning and execution. The planning process is carried out by the *Deliberate* and *Rapid* Planners. The Deliberate Planner (which shall be described in the next Section) represents the human decision-making processes that takes place at the military Strategic and Operational Levels (*high command*) where there is generally more time for decision making, and is concerned with the generation and evaluation of multiple Courses of Action (CoAs).

In contrast, the Rapid Planner attempts to replicate the processes at the Tactical Level (*battle command*) where decisions must be made quickly in response to a rapidly-changing, dynamic environment. It is influenced by what is referred to as *naturalistic decision making*, and in particular the recognition-primed decision making model (Klein 1989) which "corresponds to decision-making by experts under stress" (Moffat 2002). The focus, then, is on assessing the situation and applying previous experience in similar scenarios to rapidly select a (tactical) CoA. Moffat (2002) summarises:

> "The output of the [Recognition-Primed Decision] model is a command decision (the selection of a CoA) made not on the basis of extensive option generation and evaluation but instead by recognising the extant situation and using experience to jump immediately to an appropriate solution (a CoA)."

The Commander of a Brigade (represented by the Deliberate Planner) gives each of his immediate subordinates (Battlegroup Commanders) a mission. Each of these must then plan how to execute their designated assignments (represented by the Rapid Planner). The same process is replicated lower down in the C2 hierarchy; for instance, each Company/Squadron will receive a mission from their Battlegroup Commander. Each peer decision maker shares cues with the other decision makers in order to allow them to anticipate each other's decisions, and in doing so, enable a coordinated plan to be constructed. These planning processes result in the selection of a CoA which a Commander requires to be prosecuted in order to achieve an effect. The reader is referred to Moffat (2002) for more information on the Deliberate and Rapid Planners. The focus of this paper is on improving the quality of the information supplied to the Deliberate Planner.

The execution phase of the system (which consists of intelligent movement algorithms and learning technologies) are concerned with the intelligent execution of a CoA. Within this context, intelligent execution is concerned with an agent's ability to plan and execute a given CoA as specified by the Rapid Planner, without recourse to the use of fixed templates or rules that can restrict the ability of an agent to adapt to new or emerging changes in a situation.

The intelligent movement algorithms allow a formation to adjust to gains- or losses- of subordinate agents (for example, losses through combat degradation or agile and dynamic task reorganisation in response to a changing situation) while maintaining a credible force disposition.

The focus of the learning technologies aspect of the research is on the development of artificially intelligent learning algorithms that are capable of adapting to changing situations and learning about the best way to tackle new situations when executing CoAs. The aspiration is to develop algorithms which not only have an ability to generalise well in unseen scenarios and to be adaptive, but which exhibit the desirable property of being able to evolve a diverse range of behaviours simultaneously.

So far, it has been argued that military models of conflict need to be developed to reflect new processes in C2 which are required as a result of a new battlespace paradigm and to move from the NEC Transition state to NEC Mature. Key terms and concepts have been introduced, and an overview has been given of the work that Dstl is conducting to meet these new challenges in modelling. The focus of the remainder of the article is on improvements to the Deliberate Planner. It is intended that improvements to this decision-making process will allow models of military conflict to operate more realistically in simulation mode. The Deliberate Planner is described in more detail and the limitations of the existing implementation are highlighted. Subsequently, work is presented to address these limitations.

The Deliberate Planner

The Deliberate Planner (which is embedded in WISE) is a representation of human decision-making at the military Strategic and Operational Levels (Moffat 2002). The process is referred to by Moffat (2007) as *rational choice decision making* (in contrast to *recognition-primed decision* making described in the previous Section) and is summarised as follows:

> "Deliberate Planning represents decision-making based on a rational choice among alternatives ... In such rational choice decision-making the emphasis is on the explicit generation, and subsequent evaluation, of alternative courses of action. In military terms it corresponds to the generation of a plan which involves the allocation of multiple forces both in space and time, in order to prosecute an intent and objectives."

The decision-process that the Deliberate Planner replicates can be best described from a 'real world' perspective:

First, the Blue Commander must evaluate the options that the enemy Commander may have, and Red's possible intended effect using a range of input data values (see below). Moffat (2002) describes this as an assessment of *Enemy Capabilities and Intentions* (ECI). Once a range of ECIs has been considered, the Blue Commander must determine the likelihood of each ECI and rank them according to this assessment.

Having identified likely enemy options, the friendly Commander must select a Course of Action (CoA) for Blue forces. This is achieved by matching Blue CoAs to the range of ECIs and evaluating each ECI-CoA pair. Having done this, the Blue Commander must compare all the CoAs available to him and select his favoured option along with his intended effect. These directives are then passed to subordinate Commanders (represented by the Rapid Planner) who determine how to execute the plans at the tactical level.

In implementation terms, the Deliberate Planner utilises Game Theory and Bayesian techniques to generate ECIs. A large data set is required in order that the Planner can span the set of possible enemy options against which the Red Commander's intent is assessed. More specifically, it requires inputs to describe the possible enemy avenues of approach; possible enemy objectives and likely courses of action for Red[3]. The nature of these inputs can result in models that are 'scripted' along set lines which does not allow for dynamic and emerging agent behaviour.

In the next Section, novel mathematical methods are presented which aim to improve the process by which a Blue CoA is selected. The algorithms described will improve the situation assessment by using a number of cues to give a better indication of adversarial intent. In doing so, the number of ECIs generated by the Deliberate Planner will be reduced resulting in better information on which to base subsequent decision-making.

Ultimately, it is hoped that these enhancements will result in the removal of input data artificialities (such as those described above) from the existing generation of combat models which are currently required to enable these models to run but play no part in the military planning processes.

3. In contrast, the Rapid Planner requires data associated with low-level tactical cues such as the value of the local perceived combat-power ratios. It then uses a form of pattern-matching to map these cues to a CoA.

Identifying Collective Behaviour

The aim of this research is to look at ways of identifying groups and activity between those groups, and to infer intent by examining characteristic group behaviours and indicators of intention. It has been argued that this will result in a reduction in the amount of information that the Deliberate Planner has to assess in order to make its decisions. It is hypothesised that a better indication of adversarial intent will result in better strategic and operational decisions being made by the Deliberate Planner. Prototypes of the developed algorithms are implemented, for verification and validation in WISE. Ultimately, the aspiration is for an algorithm that can be embedded in a variety of combat models.

Gelenbe et al. (2006a) document the results of a three month preliminary study conducted at Imperial College, London, in collaboration with Dstl to evaluate computational methods for identifying collective behaviour in enemy units set in a conventional (symmetric) warfare environment. It is anticipated that the method developed can then be applied to asymmetric warfare. A multi-agent modelling approach was used to test candidate solutions (Gelenbe et al. 2006a, 2006b). Principally, two suitable computational methods were identified for application: *k*-means clustering, and an approach using a Minimum Spanning Tree (MST).

k-means clustering (MacQueen 1967) is a technique based on spatial interest. A disadvantage of this method of clustering is that the number of clusters has to be specified beforehand. It was determined that an approach using a network Minimum Spanning Tree (MST) was the most suitable for general application in the context of collective behaviour modelling by providing a computational framework within which to model the fusion of different behaviour identifiers. It proceeds as follows:

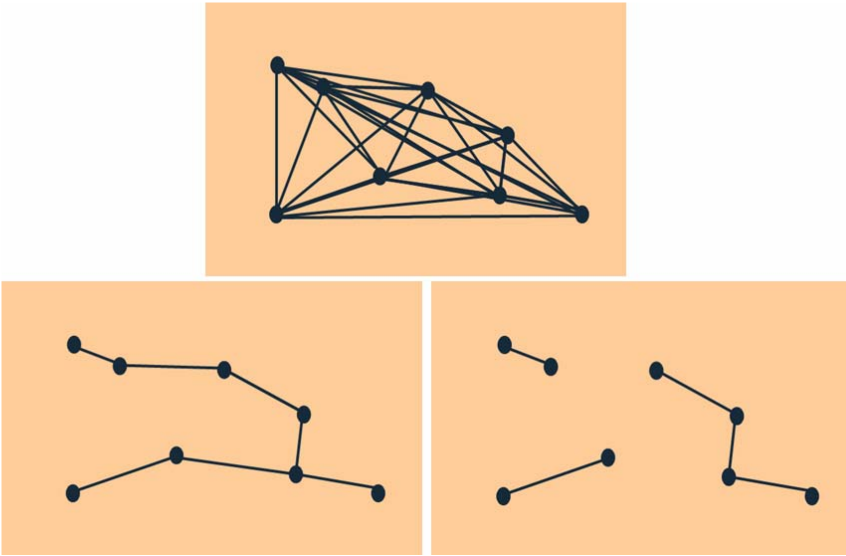

Figure 5: MST Algorithm

- A fully connected bi-directional weighted graph with N vertices and $N \times \left(\frac{N-1}{2}\right)$ edges is constructed. The vertices represent observed agents, and the weights of the edges are a measure of how well a pair of agents is related.
- An MST is built using Dijkstra's Algorithm (Dijkstra 1959) whereby the agents are all connected in such a way that the sum of the arc weights is the minimum possible.
- A density approximation of the distribution of MST arcs is computed using Gaussian Parzen Window estimation (Parzen 1962). The arcs that have a value above some pre-defined cut-off threshold are removed. The resulting 'forests' (that is, a collection of non-cyclic trees) represent groups of agents.

Figure 5 illustrates the algorithm pictorially: The top image represents the fully connected graph. The bottom left image represents the MST and the bottom right image represents the forests that remain after cropping the MST.

The utility of this method is two-fold: the weights are free variables that can be defined for a specific application (e.g. separation distance, velocity difference, density of communications traffic, etc). In addition, the method allows fusion of cues to be easily modelled and weighted to reflect the importance of the individual factors' contributions to the definition of collective behaviour.

This method has subsequently been developed and implemented by Dstl and shall be presented here:

The algorithm takes as its inputs the geographical location of an agent; its *speed* and its *heading*. This is used to calculate the cues which weight the edges of the graph. The cues used are: *distance* between a pair of agents, difference in *speed* between a pair of agents and difference in *heading* between a pair of agents[4]. Each cue is normalised in the range [0,1] by dividing by the maximum observed value for that cue in order that a cue does not dominate due to differences in unit (i.e. *distance* in metres is likely to be a much larger value than *heading* in degrees).

The edge-weight ($EW_{i,j}$) of the graph (that is, a measure of *relatedness* between any two agents, i and j) is given by Equation 1.

$$EW_{i,j} = \left(\alpha \times \frac{d_{i,j}}{\max\limits_{i \neq j \in N}\{d_{i,j}\}} \right) + \left(\beta \times \frac{\left|\Delta s_{i,j}\right|}{\max\limits_{i \neq j \in N}\{\Delta s_{i,j}\}} \right) + \left(\gamma \times \frac{\left|\Delta h_{i,j}\right|}{\max\limits_{i \neq j \in N}\{\Delta h_{i,j}\}} \right)$$

Equation 1

4. with care taken to ensure that differences are taken in the minor sector so that the difference in heading between an agent travelling at 1° and one travelling at 359° ais 2° and not 358°.

Given that there are N agents: $d_{i,j}$ is defined as the *distance* between the two agents and $\max\{d_{i,j}\}$ is the maximum *distance* over all pairs of agents; similarly $\Delta s_{i,j}$ and $\Delta h_{i,j}$ are the respective differences in *speed* and *heading* between all pairs of agents.

The parameters α, β and γ are weights which can be adjusted according to how 'important' each cue is considered to be. For instance, *speed* and *heading* will be more discriminatory for groups which are co-located. In contrast, when agents are spread over a wide geographical area with varying *speeds* and *headings*, each cue becomes important when trying to group agents together. The weights are adjusted dynamically with each execution of the algorithm, based upon the corresponding cues' standard deviation and range.

First, the input data for each of the cues is normalised, to give s_i' , using central limit theory: for example, consider the set of agent *speeds* $S = \{s_1, s_2, s_3 \ldots s_n\}$

$$s_i' = \frac{s_i - \mu}{\sigma}, \text{ for } i \in \{1,2,\ldots,n\}$$

where s_i' is a value of the cue;

μ is the mean of the set S ;

σ is the standard deviation of the set S .

The range of the normalised set of cues is then calculated:

$$range(S') = \max\{s'_1, s'_2, \cdots, s'_n\} - \min\{s'_1, s'_2, \cdots, s'_n\}$$

This process is repeated for the other sets of cues (viz. *heading* and *distance between a pair of agents*). Completing the example, then, the weight of the *speed* cue, is calculated by dividing $range(S')$ by the sum of all three of the calculated ranges. Having computed the weight of each edge, the algorithm proceeds as described above. For the purpose of prototyping the algorithm, the cut-off threshold above which arcs are removed to compute forests is set at 0.9.

Confidence Scoring

Having hypothesised which agents are grouped together, it is useful to provide some means by which to measure the confidence that the algorithm has in these groupings. If the *speeds* and *headings* of the agents within each cluster are similar, this is used as an indication that they do indeed belong in the same group. In contrast, if there is a wide spread of *speeds* and *headings* of the agents within the same cluster, then there is a lower confidence in the grouping. Based on this premise, the confidence in a hypothesised cluster (a set *C*, which defines a subset of agents that are grouped together) is calculated according to Equation 2. The standard deviation of the *speed* and the *heading* are normalised by the range of values for each property. Given that the difference in *heading* between, for example, an agent travelling at 2 ° and another at 358 ° is 4 ° (rather than 356 °), *headings* are adjusted to ensure that they lie in the range [-180 °, +180 °]. This results in the difference in *heading* being maintained but the standard deviation being reduced. β and γ are the values that are used to weight the cues in Equation 1.

$$Confidence_C = \left(\frac{\beta \times \sigma(s_{i,j})}{\max\{s_{i,j}\} - \min\{s_{i,j}\}} + \frac{\gamma \times \sigma(h_{i,j})}{\max\{h_{i,j}\} - \min\{h_{i,j}\}} \right)^{-1}$$

Equation 2

given $s_{i,j}, h_{i,j} \in C$

It is desirable to normalise the possible values of Confidence in Equation 2. This can be done by considering the range of values that the denominator can take (Richardson 2007). Without loss of a generality, it can be shown (see Appendix A) that for a cluster of P agents, the confidence lies in the range defined by:

$$\frac{2}{\beta+\gamma} \leq Confidence \leq \frac{\sqrt{2P}}{\beta+\gamma}$$

As noted by Richardson (2007), the upper limit is dependent on the cluster size. Therefore, the raw confidence score is scaled to take into account the cluster size, placing more confidence in smaller clusters.

Let the minimum and maximum possible confidences be MIN and MAX respectively; a normalised score in the range [0,1] is given by:

$$\frac{Confidence - MIN}{MAX - MIN}$$

Equation 3

When $P = 2$, $Confidence = MIN = MAX$ so that the normalised confidence is undefined.

An algorithm for identifying which agents may be working collectively has been described. Each hypothesised cluster has a score attached to it to indicate the *confidence* that the algorithm has in the identified groupings. Initial results are presented later.

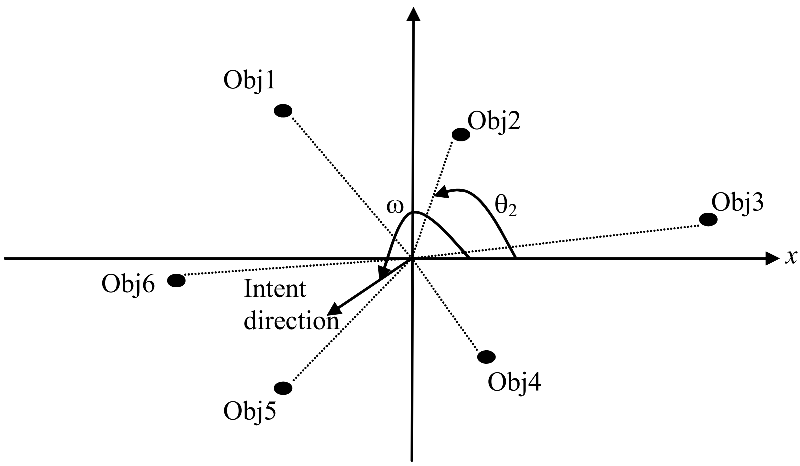

Figure 6: Direction of Intent and a Collective of Possible Objective Locations

Inferring Intent

Having hypothesised groups, the second aim of this research is to infer intent. Given M dynamically evolving clusters in a synthetic environment where there is a possible N destination objectives, there is a requirement for a method by which each cluster-objective pair can be scored to represent the likelihood that a specific cluster is moving towards an objective. Currently, intent is inferred with respect to the position of static objectives; these may be, for example, key geographical features which Red may wish to dominate. The positions of such features are passed as input parameters to the algorithm. This prediction of intent, or likelihood scoring, is calculated dynamically thus confidence varies as a function of time.

Daglish (2007) proposes a method of *likelihood scoring* for a set of targets. In this paper, this method is adapted by considering a set of objectives (defined by locations). Briefly, the method can be described with reference to Figure 6. All angles are measured in an anti-clockwise direction from the positive x-axis. The *direction of movement*, representing the direction of a cluster, is defined by the angle *w* and is taken to be the mean *heading* of the agents in that cluster.

A family of functions is defined to give an indication of intent with respect to an objective; these are stated in Equation 4. Put another way, it calculates the likelihood that the cluster is moving towards an objective:

$$ {}^{n}_{\omega}\tau(\theta_i) = \frac{1}{2}\left[1 + \cos^{2n+1}(\omega - \theta_i)\right] $$

Equation 4

for $n = \{0,1,2,3,...\}$, then:

$(\omega - \theta_i) \in [-\pi, +\pi]$;

$\theta_i \in [\omega - \pi, \omega + \pi]$;

${}^{n}_{\omega}\tau \in [0,1]$

given ω is the direction of movement of a cluster;

θ_i is the angle between the positive *x*-axis and the arc that links Obj *i* and the origin, measured in an anti-clockwise direction from (see, for example, Figure 6).

Figure 7 shows the first five members in the family of functions ${}^{n}_{\omega}\tau(\theta_i)$, given $\theta_i \in [\omega - \pi, \omega + \pi]$ for any angle ω.

Clearly illustrated is a tripartite tendency of the curve, which corresponds to the classes defined below. This tendency becomes more pronounced as *n* increases. When considering a single direction of

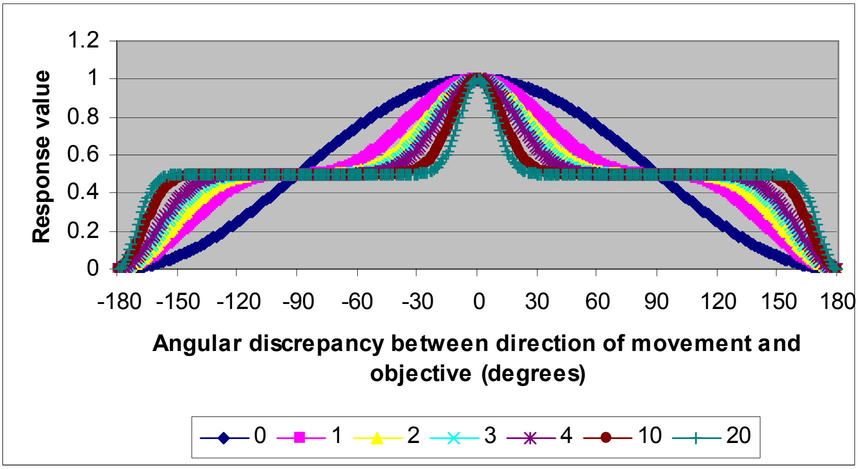

Figure 7: Naive Target Commitment Functions

movement against a single objective location, the significant values of τ, determined by the value of $(\omega - \theta)$ are as follows:

- ${}^{n}_{\omega}\tau = 1 \Rightarrow$ The direction of intent **is** the objective location
- ${}^{n}_{\omega}\tau = \frac{1}{2} \Rightarrow$ It is **undecided** whether the direction of intent is aligned with the objective location
- ${}^{n}_{\omega}\tau = 0 \Rightarrow$ The direction of intent is **not** the objective location

As n increases, the plateau shown in Figure 7 becomes longer. Furthermore, when n is small the class defined by ${}^{n}_{\omega}\tau = \frac{1}{2} \Rightarrow$ is also small. The choice of the value of n will determine the degree of sensitivity of the function: with increasing n the bins corresponding to $\tau = 0$ and $\tau = 1$ become narrower.

The choice was made to use $n = 10$ at this developmental stage, purely as a 'proof of principle' to show the function of the algorithm in inferring intent. Observe, though, that a wider plateau results in a higher likelihood that those clusters sent to the $\tau = 0$ and $\tau = 1$ classes are correct, and that increasing *n* from 10 to 20 has only a small effect upon the length of the plateau. It is intended that future development of the algorithm will attempt to assign *n* dynamically. It could be argued, for instance, that the choice of *n* should be allowed to vary depending upon scenario or time-in-the-battle, where the degree of accuracy becomes more critical. For instance, for high or very high values of *n*, 'marginal' hostile intent could be masked by the extent of this plateau[5].

In addition to an indication of intent, the destination objectives can prioritised relative to the naively perceived commitments of any given cluster according to Equation 5:

$$T_i = \frac{1 + \cos^{2n+1}(\omega - \theta_i)}{N + \sum_{j=0}^{N-1} \cos^{2n+1}(\omega - \theta_j)}$$

Equation 5

for $n = \{0,1,2,3,...\}$

and where N is the total number of objectives.

Note that $\sum_{i=0}^{N-1} T_i = 1$, such that the set $\{T_i\}_0^{N-1}$ is termed a set of *Ranking Commitments* for the given set of objective locations. That is, the Ranking Commitments suggest where the attention of the Blue Commander should lie, and the relative importance of these priorities. This will allow decisions to be made with regards allocation of assets (*force priorities*).

5. False inferences of this type would not present so much of a problem for those clusters that are classified as *undecided* rather than *non-threatening* (that is, those for which tends to zero).

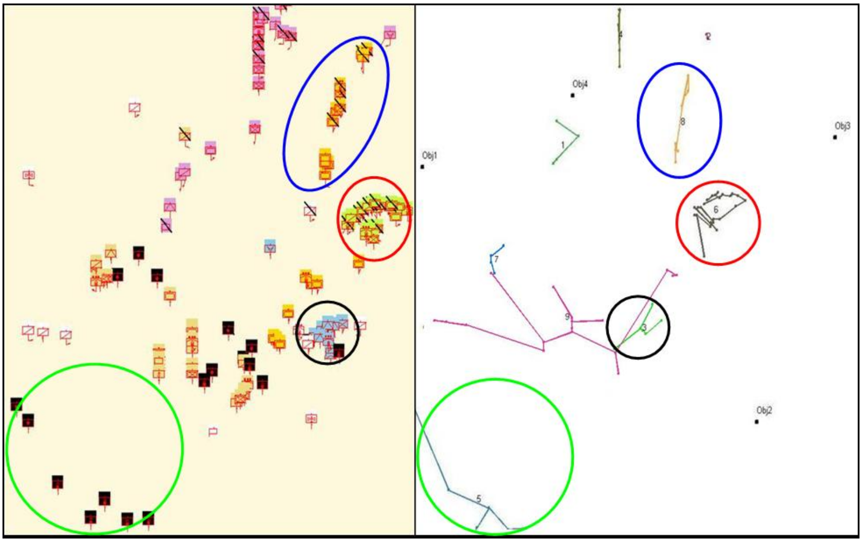

Figure 8: Ground Truth c.f. Hypothesised Clusters

Two algorithms have been presented which together identify agents which are operating collectively, and infer their intent with respect to a number of static objectives. The next Section presents initial results from executing these algorithms.

Results and Discussion

Trials using the WISE wargame were conducted, enlisting Dstl Military Advisers as gamers, to play-out a number of conflict scenarios. Input data for the algorithm was extracted from the movement logs of these games providing geographical location, *heading* and *speed* for a number of agents over time. This allowed the output from the algorithms to be compared to ground-truth. In addition to the movement data, four further positions were input as objective locations.

Figure 8 is a comparison of ground-truth (left-hand image) with the output from the clustering algorithm (right-hand image). The colours on the left-hand image were assigned by a human player as

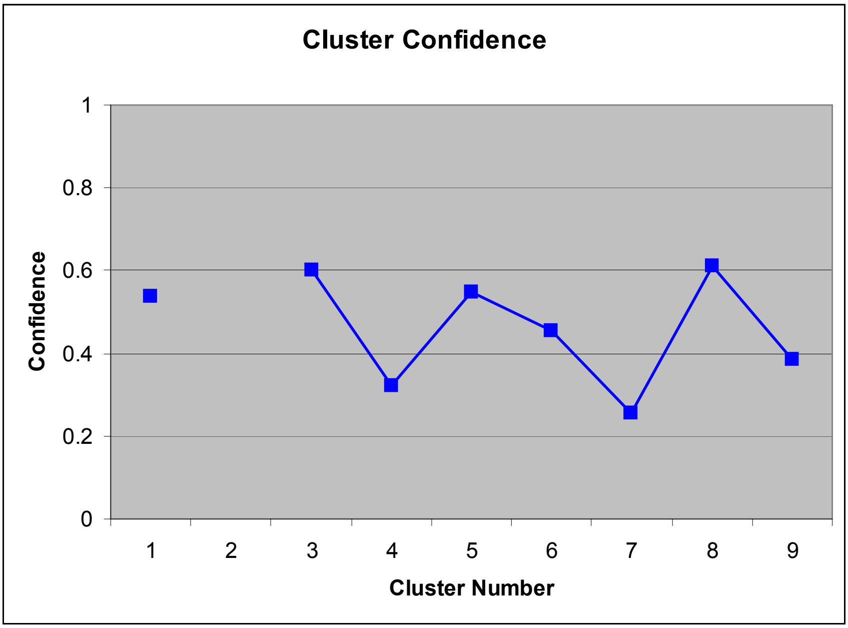

Figure 9: Cluster Confidence

an aid to understand the possible groups. The output from the clustering algorithm displays a number of forests which indicate which agents (represented by vertices) are judged to be operating collectively. The groups are numbered incrementally and coloured for visual clarity. In addition, the location of the four objectives is also shown (labelled in the form Obj *x*).

Visual comparison between the two images in Figure 8 clearly shows a strong mapping between the ground-truth and the clusters identified by the algorithm; a selection of comparisons is highlighted for clarity. This is consistent with initial results, a more detailed discussion of which can be found in Gelenbe et al. (2006b).

The graph in Figure 9 shows the confidence that the algorithm has in each of the hypothesised groupings according to Equations 2 and 3. It can be seen that Cluster 2 is undefined since it comprises only two members; given that Cluster 8 is elongated, it would not be inconceivable that Cluster 2 and Cluster 8 are merged for lower MST cut-off threshold values to form a larger cluster. In comparison Cluster 7 has a fairly low confidence score. The factors that contribute to this are: the cluster is not isolated (i.e. there are some neighbours) and the shape of the cluster is anisotropic. Clusters 3, 4, 5 and 9 correspond to clusters with 'middle range' scores; their shapes are anisotropic but their *headings* are sufficiently distinct. This illustrates that the cluster confidence scoring, based on the range of *speed* and *heading* (normalised by the respective standard deviations), gives plausible results that can be explained by rational argument.

Table 1 gives the *headings* for each of the hypothesised clusters (taken as the mean *heading* of the agents in that cluster) and shows the output from the intent inferencing algorithm; that is an indication of the likelihood that a cluster is moving towards an objective.

Table 1. Cluster Heading and Likelihood that a Cluster Is Moving Towards an Objective

Cluster (Heading)		Obj 1	Obj 2	Obj 3	Obj 4
1	(88.16°)	0.048	0.500	1.000	0.500
2	(349.25°)	0.500	0.006	0.500	0.500
3	(29.75°)	0.500	0.500	0.621	0.500
4	(0.61°)	0.500	0.422	0.500	0.499
5	(10.03°)	0.613	0.500	0.502	0.948
6	(270.71°)	0.940	0.500	0.459	0.512
7	(42.99°)	0.500	0.500	0.521	0.850
8	(19.71°)	0.500	0.495	0.500	0.500
9	(42.93°)	0.500	0.500	0.676	0.503

Recall from the previous Section and Equation 4 that:

- ${}^{n}_{\omega}\tau = 1 \Rightarrow$ The direction of intent **is** the objective location
- ${}^{n}_{\omega}\tau = \frac{1}{2} \Rightarrow$ It is **undecided** whether the direction of intent is aligned with the objective location
- ${}^{n}_{\omega}\tau = 0 \Rightarrow$ The direction of intent is **not** the objective location

Given Figure 8 and the associated *headings* and likelihood scores shown in Table 1, it can be seen that the algorithm is calculating scores consistent with expectation. For instance, given the *heading* of Cluster 1 and the position of Objective 3, it is apparent that if Cluster 1 continued on its *heading* of 88.16× (as measured clockwise from North), its path would coincide with the position of the objective. This is reflected in the likelihood score of ${}^{n}_{\omega}\tau \mapsto 1$. In contrast, Cluster 1 is diametrically opposed to Obj 1 and so ${}^{n}_{\omega}\tau \mapsto 0$. Similarly, Obj 4 is currently a likely destination for Cluster 5. Cluster 4, which is moving North, does not appear to be travelling towards any of the objectives, and so ${}^{n}_{\omega}\tau \mapsto \frac{1}{2}$ for all objectives.

Table 2 gives the Ranking Commitments (as calculated using Equation 5) for the set of objective locations. That is, an indication as to where the attention of the Blue Commander should lie, and the relative importance of these priorities.

Table 2. Naively Perceived Ranking Commitments

	Obj 1	Obj 2	Obj 3	Obj 4
Cluster 1	0.024	0.244	0.488	0.244
Cluster 2	0.332	0.004	0.332	0.332
Cluster 3	0.236	0.236	0.293	0.236
Cluster 4	0.260	0.220	0.260	0.260
Cluster 5	0.239	0.195	0.196	0.370
Cluster 6	0.390	0.207	0.190	0.212
Cluster 7	0.211	0.211	0.220	0.358
Cluster 8	0.251	0.248	0.251	0.251
Cluster 9	0.229	0.229	0.310	0.231

By comparing the values in Table 2 with those in Table 1, it is clearly seen that likelihood scores for each cluster-objective pair are reflected in the distribution of the Ranking Commitments. For

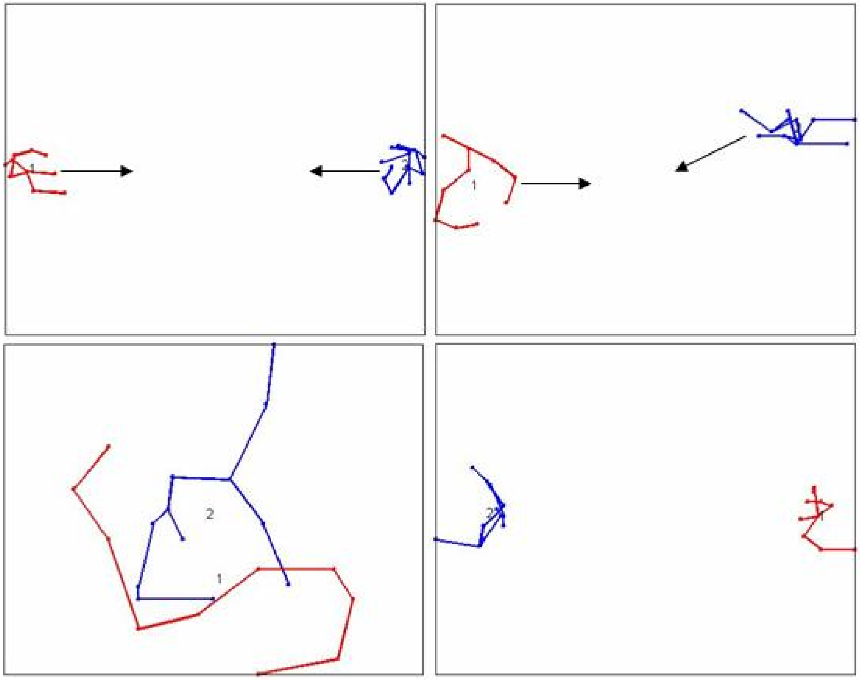

Figure 10: Algorithm output at time-step *t* = 1 (top left), *t* = 2 (top right), *t* = 3 (bottom left), *t* = 4 (bottom right)

example, the Blue Commander's attention with respect to Cluster 1 should be weighted towards Objective 3 whereas for Cluster 4, none of the objectives are considered more of a priority than the others.

The algorithm was tested further on a generated dataset over 4 time-steps, *t*, such that two groups moved towards each other, crossing each other's path at $t = 3$ before continuing on their path.

The dataset was compiled in such a way that the two groups had distinct mean *speeds* and *headings*. The output over the 4 time-steps is shown in Figure 10; note that the image at $t = 3$ is of a larger scale than the other three. It is clearly shown that the algorithm successfully distinguishes between the two groups when they cross. Table 3 shows the dynamically-calculated values of the weights used in Equation 1 (α, β and γ) at each time step. At $t = 3$ when the two groups are geographically co-located, *speed* and *heading* become more discriminatory and so the values of these weights is much higher than that for *distance*. At each of the other time-steps, none of the cues is considered to be more discriminatory than the others.

Table 3. Values of the parameters used to weight each of the cues at each time-step

	distance (α)	*speed* (β)	*heading* (γ)
$t=1$	0.314	0.350	0.335
$t=2$	0.363	0.345	0.292
$t=3$	0.027	0.272	0.242
$t=4$	0.341	0.346	0.313

Conclusions and Future Work

In this paper, the changes in force structure as part of the UK MoD's aspirations of NEC have been described. It has been argued that simulation models of conflict need to be updated in order to capture these changes and reflect the progression from NEC Transition to NEC Mature. Novel algorithms have been presented which aim to improve the decision making process carried out by the Deliberate Planner by providing a better indication of adversarial intent. The algorithms have been tested on data taken from a conventional warfare environment before being transferred to the asymmetric domain. Initial results have been presented; first, the clustering algorithm is successful at identifying clusters and is consistent with previous published work (Gelenbe et al. 2006b).

Furthermore, the results in Figure 10 and Table 3 demonstrate the utility of α, β and γ in Equation 1 in distinguishing between groups even when they are geographically co-located. In this instance, *speed* and *heading* become more discriminatory than *distance* (between pairs of agents) when trying to identify discrete groups. The flexibility of the MST method allows for the addition of further attributes determining collaborative behaviour. For example, establishing levels of communication traffic between pairs of agents may give an indication of their organisational structure.

The sizes of the clusters depend on the value of the cut-off threshold used in determining the MST. It could be suggested that varying this cut-off threshold will result in the discovery of different sizes of organisations. For example, if the cut-off is 'low' so that few (but large) clusters are found, these might correspond to Red brigades; similarly, a 'medium' cut-off may result in small clusters which might correspond to the battlegroups; and even 'higher' cut-off may produce clusters that represent company groups.

It may be useful to run the algorithm using different levels of cut-off threshold. This will enable better decision-making by giving the commander varying levels of detail of the same scenario.

It has been shown that the cluster confidence scheme, based on the range of *speed* and *heading* (normalised by the respective standard deviations), is successful in discriminating between *weak* and *strong* clusters. Such a measure may be used to weight the priority attributed to the different decision options presented to Commanders.

Having hypothesised groups of agents working collectively, an indication of intent is given with respect to a number of static objectives. The results presented in the previous Section demonstrate that the likelihood scores calculated are intuitive. The associated Ranking Commitments were distributed in accordance with these likelihood scores. Together, the two values can be said to provide a good indication of intent and the severity of that threat.

The next phase of this work is to implement the clustering and intent algorithms in WISE for further validation and verification. Figure 11 illustrates how the output might be displayed graphically to represent the information upon which the Deliberate Planner is making its decisions. The left hand image shows the hypothesised clusters and associated confidence in those groupings; the right

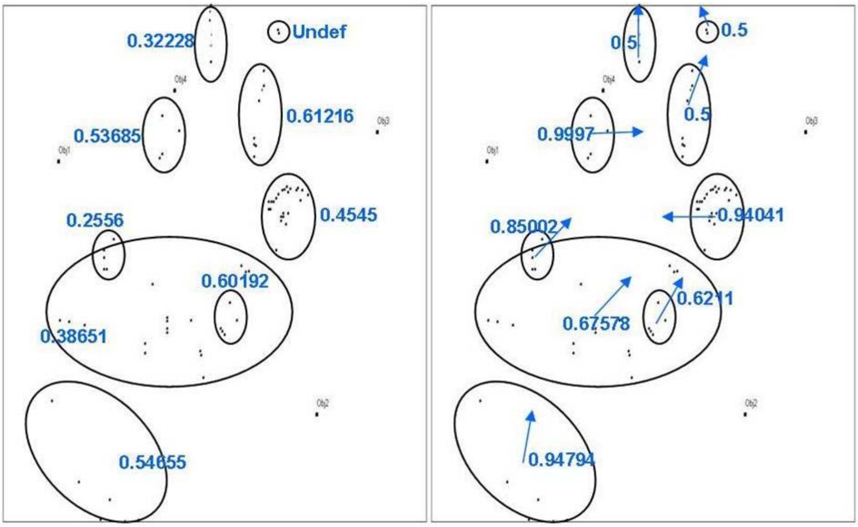

Figure 11: Example Wargame Display

hand image displays the mean *heading* of each group and the highest calculated intent likelihood value (as given by Equation 4).

The intent inferencing algorithm does not take into account the local terrain. For example, even though ${}^{n}_{\omega}\tau$ in Equation 4 may be close to 1 for a given cluster-objective pair, the local terrain may prevent that cluster from traversing the path from its instantaneous position to the objective. The host model could, therefore, pass a value to the algorithm as a *terrain complexity score* for the immediate area. This assessment could then be factored into the calculation to moderate the likelihood score appropriately.

In addition, the confidence that the algorithm has in the hypothesised groupings (Figure 9) should be used to weight the Ranking Commitments (Table 2). For instance, a group which is moving towards an objective (as given in Table 1) and for which the algorithm is confident is indeed operating collectively, should be more of a priority than one which is moving towards an objective, but for which the algorithm is less confident in its grouping. This will allow the Commander to make better decisions as to how best to allocate his assets.

Each of these enhancements to the algorithm as reported, will further improve the quality of information on which the Deliberate Planner makes its decisions and result in a better representation of modern C2 within combat models.

Acknowledgements

The authors would like to acknowledge the help and guidance of the following Dstl colleagues who have contributed to the preparation of this paper: Mick Gillman, Prof. Jim Moffat, Dr. Sue Fellows, Michael Richardson, Bob Prescott and David Mackie for conducting internal technical reviews; Tony Graham, Tony Payne, Lt. Col. Kirk Gillies, Lt. Col. Andy Kendall and Maj. Richard Grimwood for the co-ordination and execution of the experimental wargaming that supported this work. Sincere thanks must also be given to George Daglish at *Independent Analysis and Computation* for the provision of his notes on Naive Target Commitment.

References

Alberts, D. S. and R. E. Hayes (2003). *Power to the Edge: Command, Control in the Information Age*. Information Age Transformation. Washington, D.C.: CCRP.

Daglish, G. R. (2007). Naive Target Commitment. Private Communication. Independent Analysis and Computation.

Dijkstra, E. W (1959). A Note on Two Problems in Connection with Graphs. *Numerische Mathematik 1*(4), 269-271.

Gelenbe, E., V. Kaptan, and Y. Wang (2006a). Discovering Coalitions and Teams. Technical report, Imperial College, Intelligent Systems and Networks Group.

Gelenbe, E., V. Kaptan, Y. Wang, N. S. Walmsley, P. Gardiner, P. V. Pearce, and J. Moffat (2006b). A Dynamic Model for Identifying Enemy Collective Behaviour. In *Proceedings of the 11th International Command and Control Research and Technology Symposium (ICCRTS'06), 26-28 September*, Cambridge.

Holt, C. S., A. Hossain, N. Lavery, M. G. J. Richardson, and N. S. Walmsley (2007). A New Approach to the Modelling of C2 for Agile Task Organised Groupings. Technical Report Dstl/JA24566, Defence Science and Technology Laboratory (Dstl).

Holt, C. S. (2006). Evaluation and Application of the NEAT Technology to the Development of Agile Command Agents. Master's thesis, University of Southampton.

Klein, G. (1989). Recognition-Primed Decisions. In W. B. Rouse (Ed.) *Advances in Man-Machine Systems Research*, Vol. 5, pp. 47-92. Greenwich, CA: JAI Press.

MacQueen, J. B. (1967). Some Methods for Classification and Analysis of Multivariate Observations. In L. Le Cam and J. Neyman (Eds.) *Proceedings of the 5th Berkeley Symposium on Mathematical Statistics and Probability* Vol. 1, pp. 281-297. Los Angeles, CA: University of California Press.

Moffat, J. (2002). *Command and Control in the Information Age.* Norwich: TSO.

Moffat, J. (2007). Modelling Human Decision-Making in Simulation Models of Conflict. *The International C2 Journal 1*(1), 31–60.

Moffat, J., P. Grainger, Maj. M. Cornell, and S. Atkinson (2006). Modelling Agile Mission

Grouping - Problem Structuring. Technical Report Dstl/TR20440 v1.0, Defence Science and Technology Laboratory (Dstl).

Ministry of Defence (2005a). Defence Industrial Strategy: Defence White Paper.

Ministry of Defence. (2005b). The JCB NEC Delivery Strategy.

Ministry of Defence (2006). United Kingdom Glossary of Joint and Multinational Terms and Definitions. Joint Doctrine Publication 0-01.1

Parzen, E. (1962). On Estimation of a Probability Density Function and Mode. *The Annals of Mathematical Statistics 33*(3), 1065–1076

Pearce, P. V., A. Robinson, and S. C. Wright (2003). The Wargame Infrastructure and Simulation Environment (WISE). In V. Palade, R. J. Howlett, and L. C. Jain (Eds.), *Knowledge-Based Intelligent Information and Engineering Systems, 7th International Conference, (KES 2003), September 3-5*, Number 2744 in Lecture Notes in Computer Science, Oxford, pp. 714–722. Springer-Verlag, Berlin.

Richardson, M. G. J. (2007). Discussions Concerning Confidence Scoring. Private Communication.

Appendix A

The following proof taken from Richardson (2007) derives the minimum and maximum values for the confidence in each hypothesised cluster as given by Equation 2.

Without loss of generality, suppose that there are P points in the cluster, the values of one of the cues are $\{x_i : i \in 1, 2, \ldots, P\}$ and their minimum and maximum values are $a = \min\{x_i\}$ and $b = \max\{x_i\}$ respectively.

The confidence is minimum wh en each $\frac{\sigma}{range}$ term is maximum: This occurs when the points ar e most dispersed. Suppose that n of the points are at a and that $m = P - n$ are at b. The standard deviation, σ, will be calculated from the mean of the x_i, $E[x]$, and the expectation of the x_i^2, $E[x^2]$.

Then:

$$E[x] = \frac{na + mb}{P} \quad \text{and} \quad E[x^2] = \frac{na^2 + mb^2}{P};$$

so that, using $Var[x] = E[x^2] - (E[x])^2$

$$Var[x] = \frac{nm(b-a)^2}{P^2} \quad \text{and} \quad \sigma = \frac{\sqrt{nm}(b-a)}{P}.$$

Now $(b-a)$ is the range, so

$$\frac{\sigma}{range} = \frac{\sqrt{nm}}{P}.$$

This is greatest when $n = m = \frac{P}{2}$ when

$$\frac{\sigma}{range} = \frac{1}{2}.$$

Suppose that all the cues conform to this arrangement, then the minimum possible confidence is:

$$Confidence_{MIN} = \frac{2}{\beta + \gamma}.$$

Conversely, the confidence is maximum when each $\frac{\sigma}{range}$ term is minimum: This term is smallest when the points are least dispersed. Suppose that there are single points at each of *a* and *b*, and that the other *P - 2* points are at the midpoint.

It should be noted that when there are just two points, this reduces to the situation described above for minimum confidence. It can be shown that this arrangement results in a larger confidence than if one point is at one end of the range, and all the other *P - 1* points are at the other end.

For the situation outlined above,

$$E[x] = \frac{a+b}{2} \text{ and } E[x^2] = \frac{1}{4P}\left(P(a+b)^2 + 2(b-a)^2\right);$$

so that

$$[x] = \frac{1}{2P}(b \text{ - and } = \frac{1}{\sqrt{2P}}(b - .$$

Since $(b-a)$ is the range,

$$\frac{\sigma}{range} = \frac{1}{\sqrt{2P}} .$$

It can be seen that, unlike the expression for the greatest value, this depends on the number of points in the cluster, *P*, and that when $P = 2$, the confidence returned is $\frac{2}{\beta+\gamma}$.

Suppose that all the cues conform to this arrangement, then the maximum possible confidence is:

$$Confidence_{MAX} = \frac{\sqrt{2P}}{\beta+\gamma} .$$

For a cluster of *P* points, the confidence lies in the range defined by:

$$\frac{2}{\beta+\gamma} \leq Confidence \leq \frac{\sqrt{2P}}{\beta+\gamma} .$$