A Dynamic Model for Identifying Enemy Collective Behaviour

Erol Gelenbe
Polish Academy of Sciences

Paul Pearce
Defence Science and Technology Laboratory DSTL

11th ICCRTS
Coalition Command and Control in the Networked Era

A Dynamic Model for Identifying Enemy Collective Behaviour

E. Gelenbe, V. Kaptan and Y. Wang
Department of Electrical and Electronic Engineering, Imperial College, London, SW7 2BT, UK

N.S. Walmsley, P. Gardiner and P.V. Pearce
Dstl Land Battlespace Systems Department, Fort Halstead, Sevenoaks, Kent, TN14 7BP, UK

J. Moffat
Dstl Policy and Capability Studies Department, Farnborough, Hampshire, GU14 0LX, UK

PoC: Dr. N.S. Walmsley
Dstl Land Battlespace Systems Department
Fort Halstead, Sevenoaks, Kent, TN14 7BP, U.K.
tel: +44-1959-892636; fax: +44-1959-892477
email: nwalmsley@dstl.gov.uk

Topics: C2 Modelling and simulation; C2 Experimentation; Network-Centric Metrics

Dstl reference number: Dstl/CP19356

Abstract

Recent advances in Command and Control (C2) modelling have developed algorithmic representations of the command decision making process at both the tactical (Rapid Planning) and operational (Deliberate Planning) levels of command. In this paper, the development of methods by which indicators of enemy/group behaviour can be extracted is discussed. The aim of this work is to use the grouping techniques to establish, with more certainty, force size and direction. We describe a multi-agent model approach, based on a hierarchical framework, suitable for identifying indicators of collaborative behaviour associated with enemy agile mission groups that commonly feature in the modern Networked Era.

Introduction

The Cold War era was dominated by a predictable monolithic and structured enemy using tactics that could be templated. However today and in the foreseeable future we will be presented with a much more complex battlespace paradigm with several opponents, agents and actors (with some adopting asymmetric tactics). UK and coalition forces will leverage Network Enabled Capability (NEC), C2 tempo, and information dominance to achieve success – techniques such as envisaged in this paper will be needed in order to understand decision-making in the modern Networked era.

A key improvement to the current C2 models [1] would be to develop model algorithms that can dynamically infer enemy intentions from signature properties or indicators of enemy behaviour where *á priori* knowledge of possible enemy courses of actions is unnecessary. This paper describes the first steps towards such an algorithm.

This paper introduces a computational method, which in the context of modern warfare, that concentrates initially on identifying movement and behaviour as collaborative indicators. The method can be extended to build in other indicators in the future.

Although the Deliberate Planner [1] utilises Game Theory and Bayesian techniques to avoid the limitations of earlier rule-based approaches, it still requires a large data set to span the set of all enemy courses of action against which enemy intent is assessed. Thus, such models are inherently restrictive since the decision outcomes are pre-determined from a set of 'tramlines'.

In this paper, simulation results show how properties are dynamically discovered, leading to inferences about enemy activity and intent. The multi-agent model avoids the need for a deterministic rule-based system or for all possible options to be data-driven. The nature of this paper is investigative: focussing on those computational methods suitable for group behaviour prediction.

The challenging work investigated in this paper is how, from observations of complex multiple agent behaviours and interactions, it be may possible to deduce which entities are acting in concert or behaving as groups, either in temporary sub-sequences of events or throughout the operations. This is in some sense an ‘inverse problem’ with respect to the design of tactical operations from known cooperation between sets of agents. The approach described in this paper does not assume a known behavioural model. This choice was made in order to avoid simplifying the problem at hand by limiting it to situations where more information is available than may be available in many practical situations.

It is intended that the model will complement current C2 decision models such as the Deliberate Planner, and will support the Comprehensive Approach (CA): where inference about the enemy’s activity/intent plays a leading role in gaining information superiority. This will be achieved by using the model described in this paper to provide a more accurate picture of the size and structure (i.e. groupings) of the enemy to generate a more confident basis for decision-making.

Background: the Deliberate Planning process

To place this work into context a brief outline of the main parts of the Deliberate Planning process is provided. The Deliberate Planning process is based on an analytical decision-making paradigm often referred to *rational choice decision-making* [1]. In this paradigm the emphasis is placed upon the explicit generation, and subsequent evaluation, of multiple Course of Action (CoA). The Deliberate Planning process has three distinct phases [1]:

- selection of a CoA and the development of a plan to carry out that CoA;
- issuance of directives to carry out the plan;
- supervision of the planned action.

This paper presents a computational method aimed at improving the first part of the Deliberate Planning process, by removing the need for prior knowledge of a CoA set; therefore, removing need for ‘an explicit generation’ of the CoA. Our modelling process is aimed at ‘discovering’ those indicators that could lead to the prediction of the set of CoA undertaken by adversarial entities.

The current version of the Deliberate Planner model is used to facilitate C2 decision-making in combat simulations. The Deliberate Planner is embedded in the Wargame Infrastructure and Simulation Environment (WISE) model [2]; WISE will be used as a testbed for the method for dynamic discovery of collaborative behaviour, once the capability of the Deliberate Planner is enhanced by the introduction of the new method. This should result in a more dynamic assessment of enemy force size, structure and intention. It will also result in the removal of pre-conceived decision paths.

Modelling approach

Initial computational investigations have focussed on developing ideas based on movement and behaviour as collaborative indicators. The modelling approach presented in this paper was based on a multi-agent model. Independent decision mechanisms were used to model different aspects of the agent behaviour and a higher level coordination module combines their output. The decision mechanisms are summarized here; however, further details may be found in reference [3]. The decision components are as follows:

- The Navigation Module is responsible for leading a single agent from a source location to a destination location, avoiding “danger” and obstacles.

  The interaction between an enemy (i.e. hostile agent) and an agent is modelled by an associated risk. The goal of the agent is to minimize a function G, used to define a Reinforcement Learning Reward (RLR) function as follows:

  $$R = \frac{1}{G} \tag{1}$$

  A series, $R_i \; \{i \in 1, 2, \ldots, n\}$, is defined to represent success reward values, which are used to keep a track (T_i) of a smoothed reward given by:

  $$T_i = bT_{i-1} + (1-b)R_i \tag{2}$$

  where 0 < b < 1.

  The decision-making element is a fully-connected Recurrent Neural Network (RNN) comprising 8 neurons (each representing a possible decision). The training is performed by reinforcing the weights of each neuron, depending on the latest and smoothed awards; b acts as a control parameter to reflect the ‘staleness’ of information. By using previously acquired information and current sensory input, the Navigation Module ensures that an agent can start with near-optimal estimates of the rewards and, thus, focus on adapting to the dynamic changes in the environment.

- The Grouping Module is responsible for keeping a group of agents together in particular formations throughout the mission, and therefore drives the collaborative behaviour.

  Group behaviour is based on the idea of social potential field [4] that is a distributed-control approach inspired by the attractive and repulsive forces between charged particles in physics. The force between agents i and j is of the form:

$$\vec{V}_{i,j} = \left(-\frac{a}{r^{\alpha}} + \frac{b}{r^{\beta}} \right) \hat{r} \tag{3}$$

where a, b, α, and β are dynamic parameters and the force vector $\vec{V}_{i,j}$ describes the effect of the position of agent j on the decision of agent i.

Different behaviours, such as attraction to an agent, repulsion from an agent or maintaining a specified distance from an agent can be simulated by varying the instantiations of the dynamic parameters. The study of collaborative behaviour often requires the control of distances – for the purpose of sensitivity analysis – between agents that are associated as members of the same group; stable equilibrium points are used to model such properties. When there is a stable equilibrium point, it can be shown that an entity experiencing such a force will be separated from the force source by a distance R_0 given by:

$$R_0 = \sqrt[\beta-\alpha]{\frac{b}{a}} \tag{4}$$

The total grouping effect on agent i can be calculated as follows:

$$\vec{V}_{grp(i)} = c * \sum_{j} \vec{V}_{i,j} \tag{5}$$

where the parameter c is a scalar that models the degree of cohesiveness of the group of associated agents.

The distance separator (as expressed in equation (4)) being especially important in forming localized groups, which are similar in effect to the collision avoidance and flock centring rules as described in reference [5].

- The Imitation Module is modelling the case when an inexperienced agent will try to mimic the behaviour of the most successful agents in the group and thus increase its chance of success. This can be expressed as follows:

$$\vec{V}_{imt(i)} = \sum_{j \in S} w_j * V_{nav(j)} \tag{6}$$

The decisions of these modules are combined at a higher-level module called the Coordinator Module. Thus, mathematically, at each time step a weighted sum of the separate decisions recommended by each basic module can be represented by a 2D velocity vector:

$$\vec{V}_{overall} = k_{nav}\vec{V}_{nav} + k_{grp}\vec{V}_{grp} + k_{imt}\vec{V}_{\mathrm{int}} \tag{7}$$

where the k's correspond to weighting scalars, adjustable to reflect system emphasis on different behaviour type characteristics.

Observation of the spatial and behavioural agent configurations and their respective evolution with respect to time provided valuable clues as to what type of clustering algorithms may be appropriate.

One of the problems of standard clustering methods is that they tend to work well when provided with large volumes of raw data. The scenarios presented here for presentation purposes involve relatively small number of agents that usually cannot provide statistical data of acceptable quality. As a result, a bespoke clustering algorithm that is particularly well suited for the type of scenarios with a relative small number of enemy agents: such as those typically found in agile mission groups in the context of conventional warfare.

The description of the algorithm is as follows:

1. Suppose there are N agents. Consider agent system as a fully-connected bi-directional weighted graph with N vertices and $M=N*(N-1)/2$ edges. Given that the vertices represent agents and the weights of the edges are a measure of how well agents are related. (In the spatial clustering case, edge weights are geometric distances between agents, while in behavioural clustering the edge weights represent angular difference between the directions in which each agent moves.);
2. Build a Minimum Spanning Tree (MST) for this graph. Only a few of the M edges will be part of the MST;
3. Compute a density approximation of the distribution of MST edges (Gaussian Parzen Window estimation with parameter σ);
4. Find the maximum peak in the density estimation and compute a cut-off edge threshold (located on the right of the peak) at which the density drops to half of the peak value;
5. Remove the edges in the MST that have values above the cut-off threshold. The end result is a collection of sub-trees of the MST that is usually referred to as a forest; each tree in this forest represents a separate group.

An example of how this algorithm works is shown graphically Figure 1: (a) shows the map representation including the trace from previous positions; (b) represents the MST based on spatial metrics (as outlined by the algorithm described above); (c) is the histogram resulting from the density of the MST edges); (d) results from the thresholding to produce a partial MST – this consists of forests (unconnected graphs) that are associated with the separate groups.

In all the results presented in this paper there are three groups of agents (colour-coded as Red, Green and Blue, 5 agents in each group represented as a "team": 15 agents in total), which have the following objectives:

- Blue team has to go to the destination (represented by the shaded rectangle) while avoiding the Red team;
- Red team has to intercept and destroy the Blue team while avoiding the Green team;

- Green team has to support the Blue team by trying to intercept and destroy the Red team.

The terrain also contains a number of simple obstacles (trees, represented as green circles of varying radii).

The agents in the experimental scenario that is presented in this paper exhibit a range of behaviours that make this framework a good candidate for a platform for studying methods of group behaviour discovery. Throughout the experiments, a number of interesting situations (both expected and unexpected) were observed at different times. Some of these are:

- Spatially well-formed and distinct groups behaving in accordance with their instructions;
- Adversarial groups fusing into a single group;
- Friendly groups which are spatially indistinguishable from each other due to spatio-temporal proximity of their trajectories;
- Partial break-up of a group due to inability of some group members to ‘keep pace’ with the rest of the group;
- Complete break-up of a group because of a dominant goal (i.e. urge to evade adversaries dominates other goals such as staying together, and as a result the group is scattered).

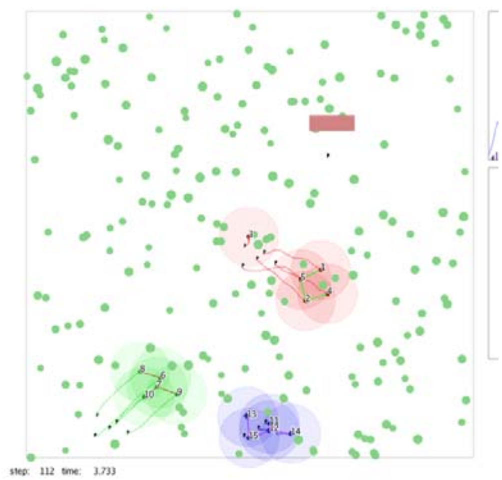

a) Simulation screenshot, 15 agents.

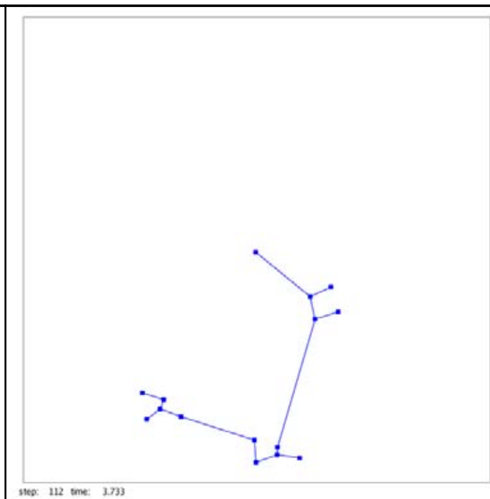

b) MST based on spatial metrics.

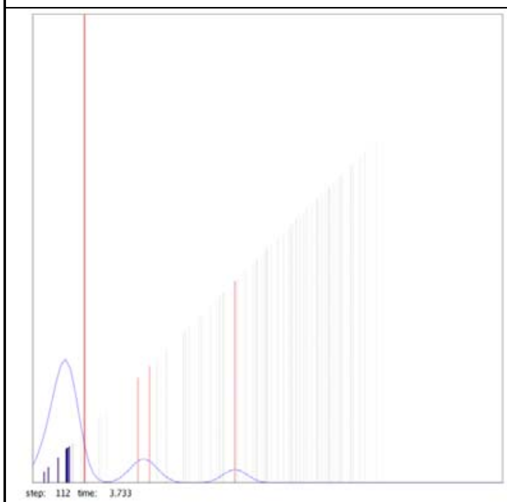

c) Histogram and density estimation of MST edges.

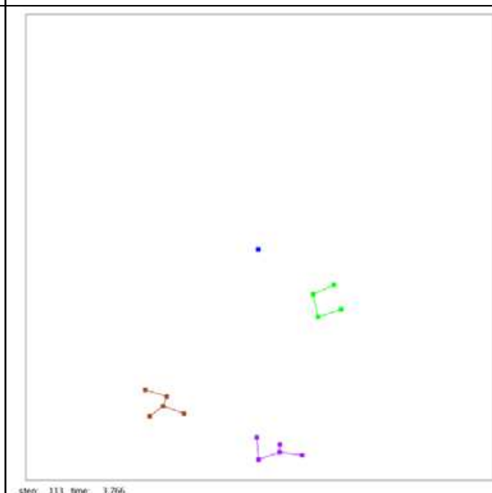

d) Partial MST forest after removal of edges above the cut-off limit.

Figure 1: Illustration of the MST-based clustering algorithm

Results and discussion

Prior to selecting a particular approach for group detection, agent behaviour was assessed with the purpose of gaining an insight into the nature of the raw data that will be used for group differentiation.

First, the investigation is focussed upon the spatial behaviour of agents. To carry this out, simulation runs both within the normal terrain and in an obstacle-free environment were assessed. Using Figures 2 and 3 for comparison, showing snapshots of simulation runs at successive times within the simulations, the effect of ‘synthetic’ terrain may be examined. The agent trajectories are shown as coloured trails. These trails provide information on the evolution of the spatial configuration of the different teams throughout the simulations.

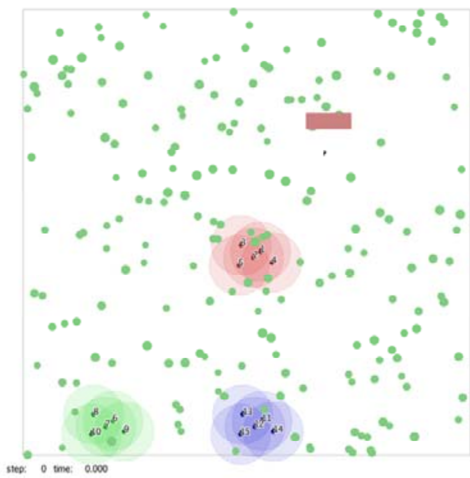

a) steps = 0

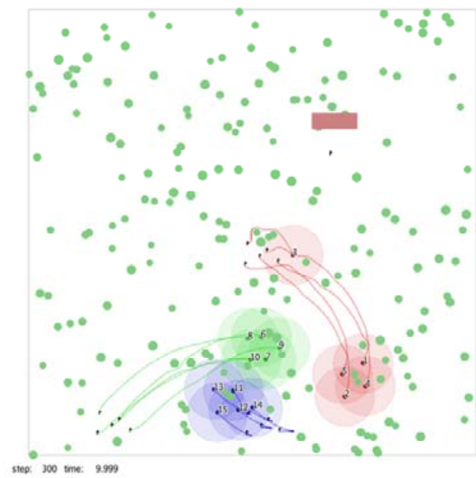

b) steps = 300

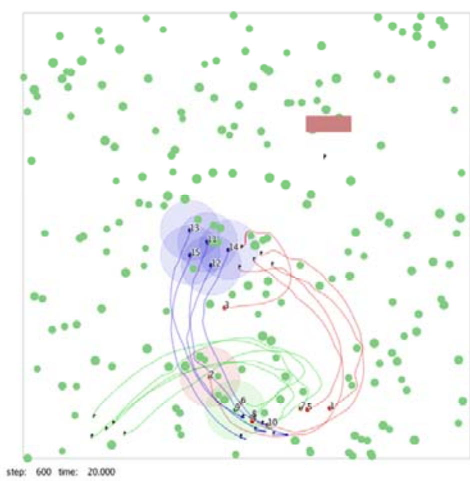

c) steps = 600

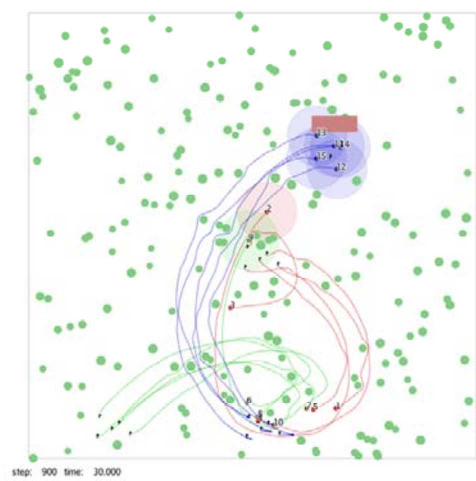

d) steps = 900

Figure 2: Evolution of the spatial configuration of the agents with respect to time in a terrain with obstacles

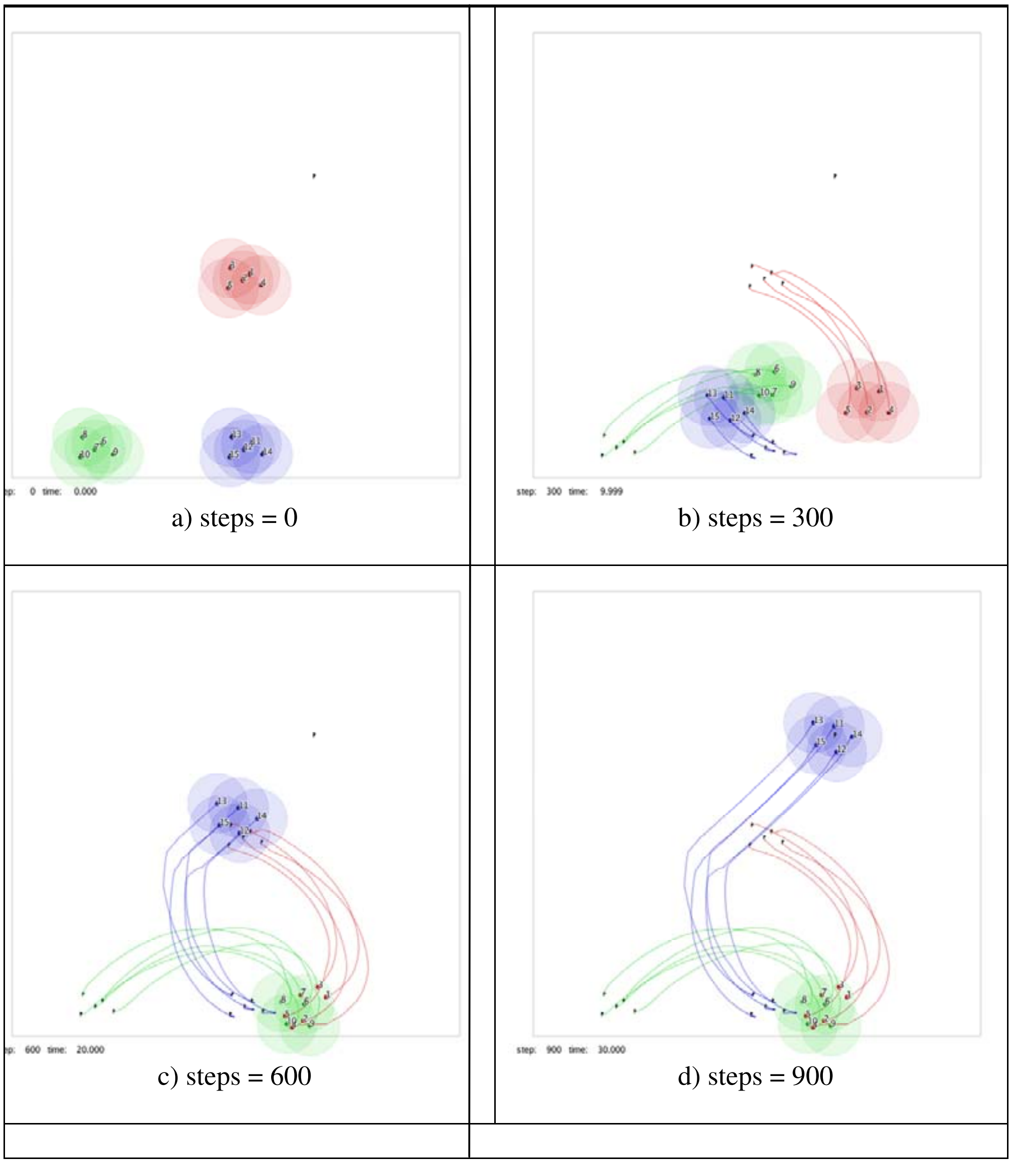

a) steps = 0

b) steps = 300

c) steps = 600

d) steps = 900

Figure 3: Evolution of the spatial configuration of the agents with respect to time in a terrain without obstacles

The groups in Figure 3 (with the exception of the Red team) seem to maintain more cohesion and are better defined than those in Figure 2, indicating that terrain has the effect of reducing the ability of agents to act as a team. Essentially, this is due to the addition of repulsive forces experienced by the agents as they come within close proximity of the ‘terrain objects’.

Besides spatial configuration, instantaneous agent behaviour was also measured as a function of time. Within the setting of these simulations, instantaneous agent behaviour refers to the direction in which an agent moves (i.e. velocity). Figure 4 provides measurement of the direction of motion with respect to time (direction is measured in radians from $-\pi$ to π). As an additional property, to simulate warfighting in the scenario, weapons are added to the mix; such a property is simulated by strengthening the repulsive forces (as represented by equation 3, and is described in the previous section) between adversarial groups. As time increases the Blue team exhibits a fairly constant velocity, indicating that their goal is within achievement and group cohesiveness is strong.

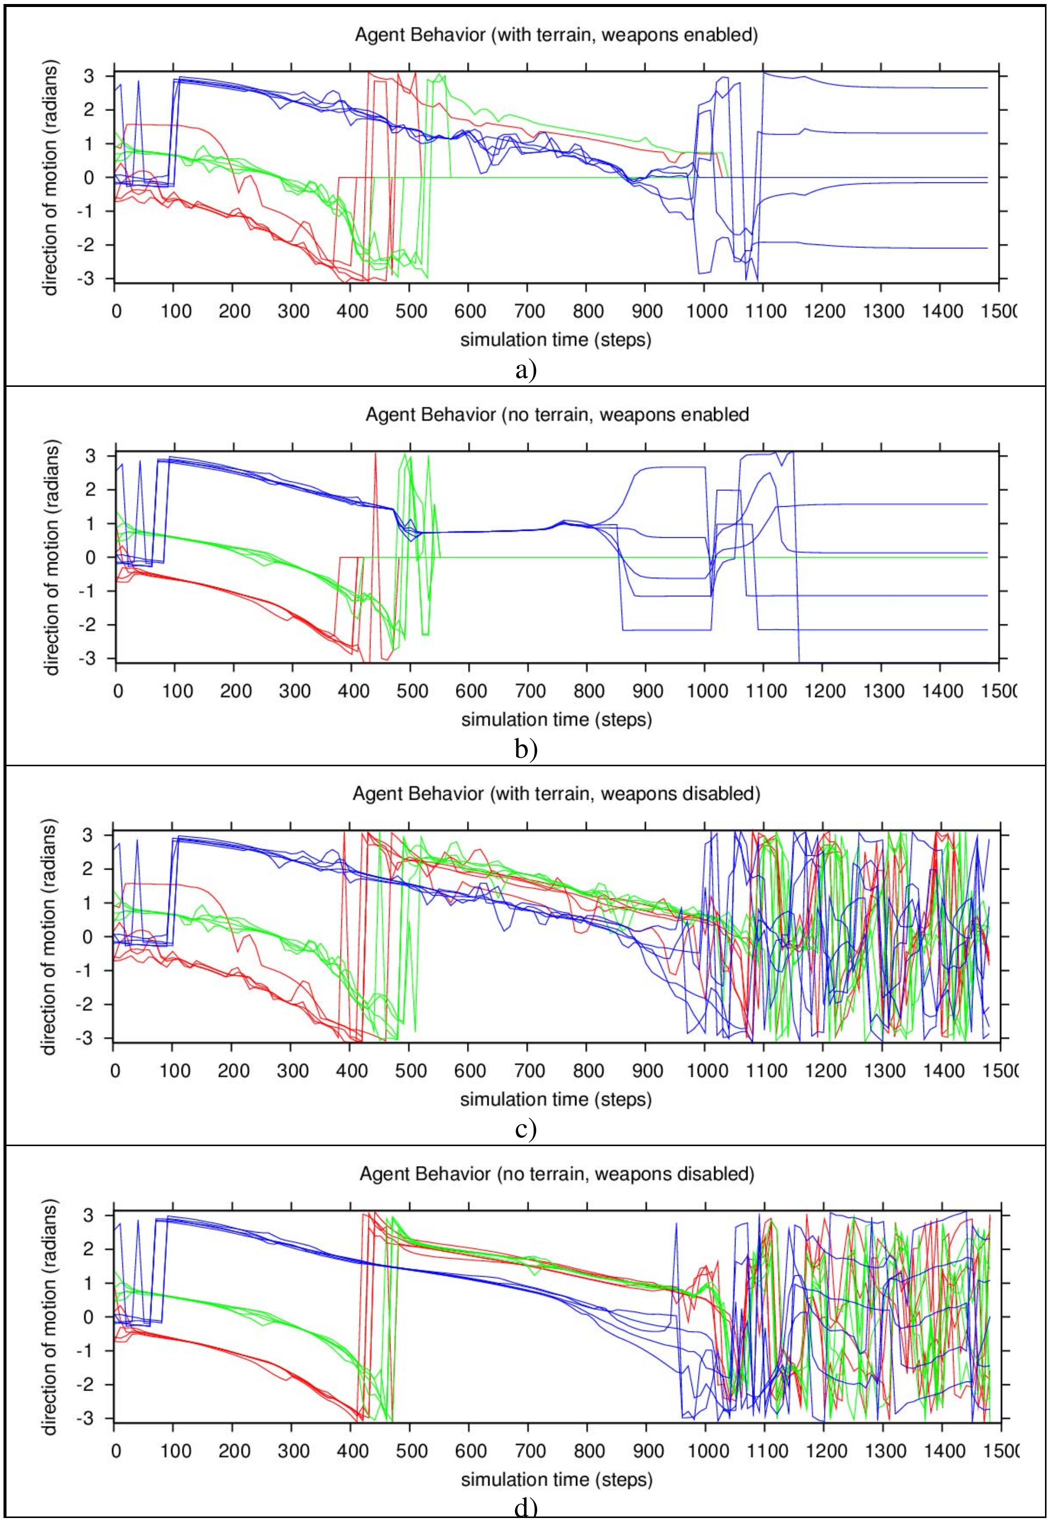

Figure 4: Evolution of the instantaneous agent behaviour with respect to time

Comparing Figure 4(a) with Figure 4(b) it can be deduced that terrain has small effect on behaviour when weapons are enabled; presumably, this is caused by the stronger influence of the repulsive forces due to the introduction of weapons, in comparison with the 'terrain avoidance' forces. This reasoning can be further substantiated by comparing Figure 4(c) with Figure 4(d): with weapons disabled, terrain has the effect of introducing 'noise' to the system thus weakening group cohesiveness. Interestingly, the significant effect of the introduction of weapons can be seen in the latter stages of the simulation: more stability in group cohesiveness is evidenced in those scenarios with weapons deployed.

Expanding upon the initial assessment, it is believed that temporal analysis of spatial and behavioural clustering can be a valuable tool for automated situation assessment. As a proof of concept Figures 5 and 6 are illustrative of a simple scheme for detection of certain events that may be of interest (separation of a group into smaller groups or combination of two or more groups into one); the spatial clustering statistics are displayed to the left of the behavioural clustering statistics.

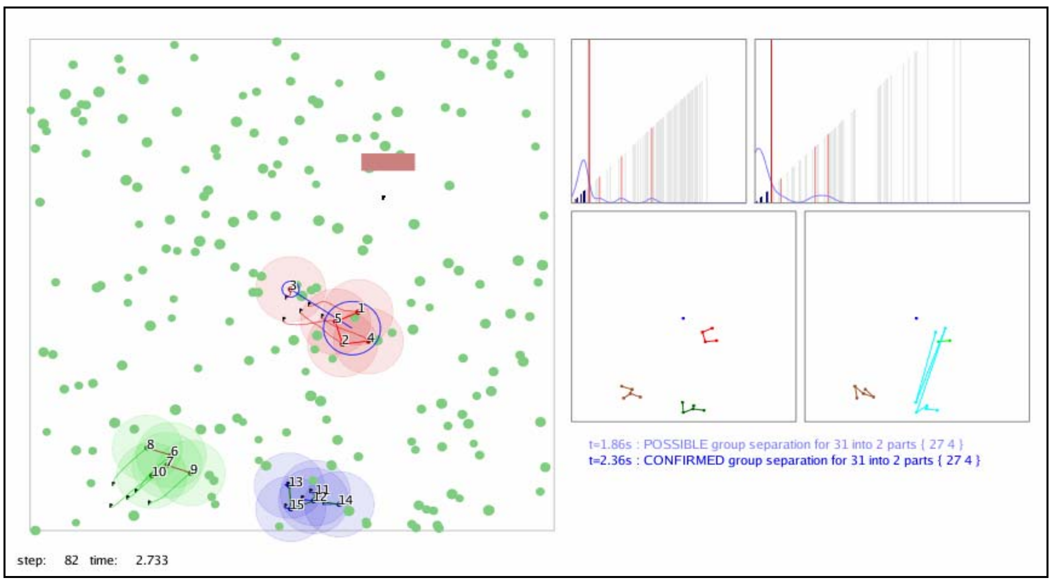

Figure 5: Spatial and behavioural clustering - detection of group separation

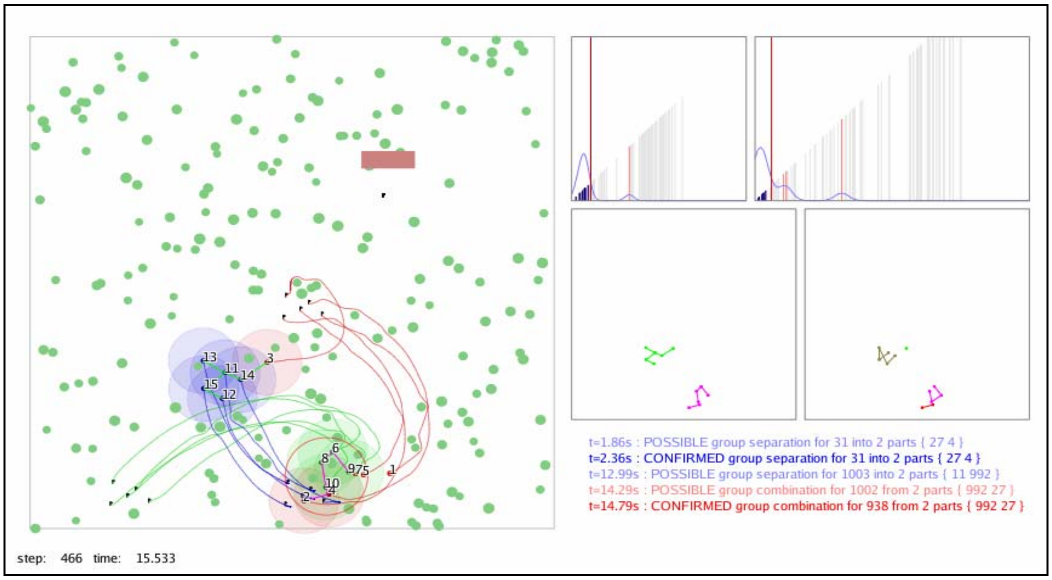

Figure 6: Spatial and behavioural clustering example - detection of group

Detection of group separation and combination is shown in Figures 5 and 6, respectively. This is reflected in the partial MSTs and histograms associated with each figure: the connectivity of clusters (forests) as exhibited by the partial MSTs for both the spatial and behavioural metrics.

As part of the experimentation it can also be noticed that spatial and behavioural clustering may be used to complement each other – sometimes one of them can disambiguate groups that the other cannot distinguish. An example of good spatial and poor behavioural separation is shown in Figure 7: exhibited by the difference in connectivity of the partial MSTs.

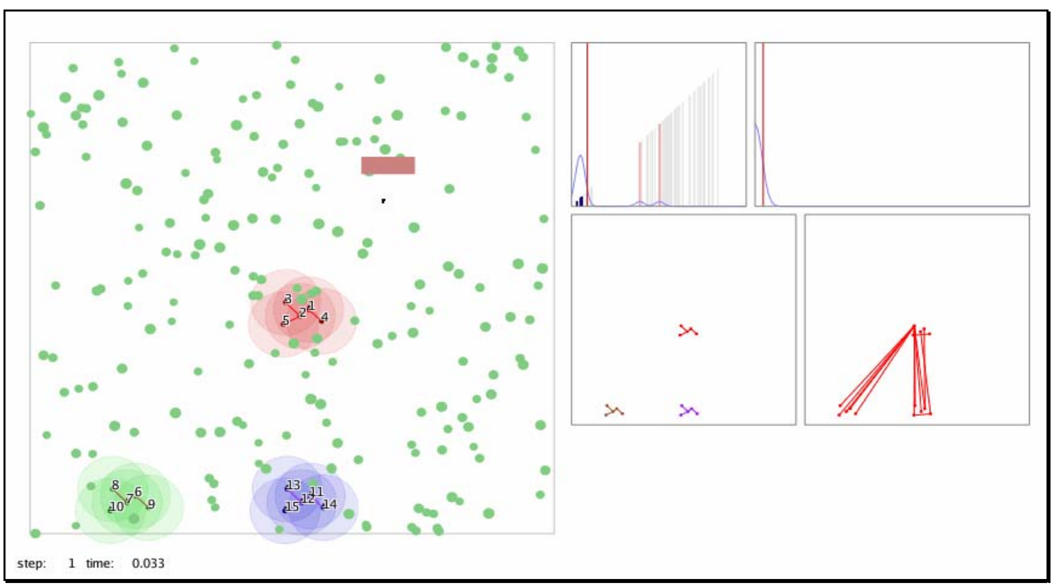

Figure 7: Example of good spatial clustering, but poor behavioural clustering

The opposite case (poor spatial but good behavioural separation) is shown in Figure 8. Comparison of the degree of connectivity between the spatial and behavioural metrics show the opposite properties to those found in Figure 7.

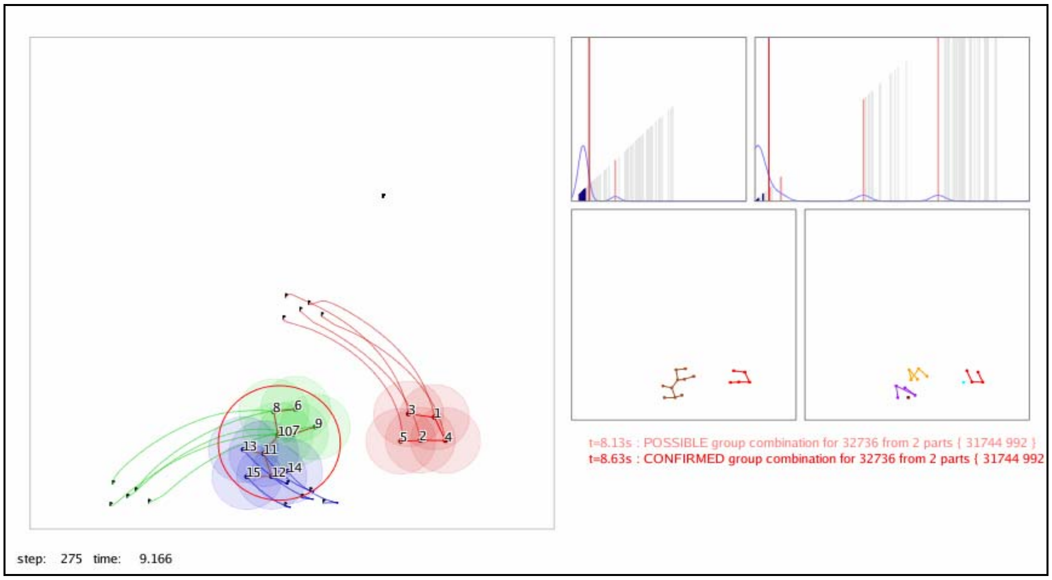

Figure 8: Example of good behavioural clustering, but poor behavioural clustering

Summary and conclusions

The success C2 decision-making in the Networked Era will become crucial in determining the outcome of military operations; information superiority will be achieved by those capable of making accurate and timely decisions primarily based on inferences from an evolving battlespace environment comprising agile mission groups.

This paper has presented a model and discussed ideas to address the problem of identifying the key factors that indicate collaborative behaviour, particularly in relation to enemy agile mission groups, which contribute to achieving a level of information superiority.

A modelling technique using a multi-agent approach that does not require *á priori* data or information has been presented; this feature of the model avoids both the need for a large volume of input data and the need for a fixed rule-base decision engine. A neuron network framework facilitates the ability to ‘learn’, by accumulating knowledge (via inferences or pattern recognition) about the behaviour of enemy agents during model execution time. The need for decisions to follow ‘tramlines’, artificially imposed by the inclusion of a fixed rule database featured in some of the other modelling approaches, has been completely removed. This gives our model the portability to be inserted into C2 decision modelling architectures, such as the Deliberate Planner, without the need to accommodate a complex input/output interface or pre-determined rules.

This work has presented a method that uses a combination of metrics (in the current paper, a spatial and a behavioural metric has been presented) to form indicators to identify group behavioural. However, it is intended that the model be extended to include other metrics (indicators of behaviour, e.g. traffic communications, locations of enemy assets, etc.), which when fused together will give a more substantive reflection of emerging behavioural in the context of agile mission groups. In parallel with this additional work, investigations into how to represent prior 'uncertainty' in group structure will be conducted; further results will be tested against the output of combat simulation using their movement files.

References

[1] J. Moffat, 'Command and Control in the Information Age: Representing its Impact', London: TSO, (2002), pp. 87-103.

[2] P.V. Pearce A. Robinson and S.C. Wright, 'The Wargame Infrastructure an Simulation Environment (WISE)', 7th International Conference, KES 2003, Oxford, UK, Sept 2003, Proc. Part II.

[3] E. Gelenbe, K. Hussain and V. Kaptan, 'Simulating autonomous agents in augmented reality', The Journal of Systems and Software, 74 (2005), pp. 255-268.

[4] J. H. Reif and H. Wang, 'Social potential fields: a distributed behavioural control for autonomous robots', Robotics and Autonomous Systems, Vol. 27, No. 3 (1999), pp. 171-194.

[5] C.W. Reynolds, 'Flocks, herds and schools: a distributed behaviour problem', Computer Graphics, 21, 4, (1987), pp. 25-34.