The Mapping of Courses of Action Derived from Cognitive Work Analysis to Agent Behaviours

R. Connell, F. Lui
Russell.Connell@dsto.defence.gov.au
frank.lui@dsto.defence.gov.au
Land Operations Division, DSTO,
Edinburgh, SA, Australia

D.Jarvis
Dennis.jarvis@agent-software.com
Agent Oriented Software Pty Ltd,
P.O. Box 639 Carlton South 3053

M. Watson
mwatson@humanfactors.uq.edu.au
ARC Key Centre for Human Factors and Applied Cognitive Psychology,
University of Queensland,
Queensland, Australia

ABSTRACT
In a previous study[1], the feasibility of controlling entities in the Onesaf Test Bed simulation using a command and control hierarchy represented as teams of agents was demonstrated. The scenario used for the demonstration was the preparatory phase for a company attack on a platoon and it highlighted a number of issues requiring further research. The most important of these was the issue of behaviour capture. In the previous work, this knowledge was elicited from domain experts via a series of interviews. While this process proved satisfactory for a demonstration of feasibility, it was recognised that even with such a constrained scenario, a principled method is required in order to capture a complete and consistent set of behaviours. Cognitive Work Analysis (CWA)[2] was suggested as a candidate for such a method and a separate study was initiated to assess the usefulness of CWA for our scenario [3]. A particular attraction of CWA is that the output of an analysis can be specified in terms of behaviours and preconditions for their execution. While the actual execution model presupposed by CWA is different to that of our agents (which use the Belief-Desire-Intent (BDI) model) the level of behavioural specification is the same. Consequently, it was decided to use the results of a more extensive CWA analysis to provide our teams of agents with much richer and more realistic behaviour models. The purpose of this paper is to describe a mapping of Courses Of Action (COA) derived from CWA to agent behaviours.

Keywords
Intelligent Agents, Command Agents, Cognitive modelling, BDI

1. Introduction

Intelligent agents (IA) have attracted much attention in the past decade. The increase in interest is because of the infrastructure provided by the belief, desire and intention (BDI) [4] framework that gives the provisions for simulating human-like decision-making. In real life, humans are not just capable of reasoning logically, but also capable of intuitive leaps and extrapolation of knowledge. A key part of human planning is the ability to recognise situations as being similar to those already experienced, shortcutting the planning process. In order for the IA to achieve a certain level of realism, it is important for us to have a good understanding of the cognition of human behaviours in specific working environments such as those in the military, aviation and medical environments.

In this paper we will describe in detail the integration of a cognitive model in the JACK-Teams [5] infrastructure that is based on the BDI framework. A Cognitive Work Analysis (CWA) study was conducted in March 2002 to capture the information requirements, methods and factors that a Mechanised Company Commander uses and considers when planning a company attack. Since July 2002, project COMATS 1.0 (Command Agents and Teams Software) has been the driver for implementing a prototype of a Command Agent (CA) incorporating a cognitive model derived from CWA. The purpose of the current study was to provide input into the development of IA (which is an external agent) that could behave autonomously within a synthetic environment (SE) [6] which comprises of both human-in-the-loop virtual simulations and computer generated forces (CGF) constructive entities. The CA operates at the company level whereas the CGFs handle the control of lower level entities such as individual tanks, infantry units, mechanised units, etc. Our main focus is the technical and procedural aspects of taking the outputs of a CWA study and implementing them within the IA paradigm.

2. A cognitive model derived from Cognitive Work Analysis

Cognitive Work Analysis (CWA) is a systems-based approach for the purposes of analysing, modelling, designing, and evaluating complex systems [7, 8]. The approach proposed through CWA is a *formative* one that attempts to show some of the forces that together will shape or "form" activity for revolutionary interface design. In the case here where the domain (military environments) already exists, CWA provides a framework for guiding analysts towards the structure and function of the domain in which commander works. The problem is therefore framed in terms of how commanders' goals, activities, and tasks relate to the

1) physical constraints of the domain, and
2) the social and organisational realities that affect coordination between humans and groups within the domain or Synthetic Environments (SE).

Such models also include the limitations imposed by the adversary's interaction with the domain.

A particular concern of CWA is to model the constraints and the affordances of a work domain as well as the activity in the work domain to reveal constraints on actions, boundary conditions and so on [7, 8]. CWA will also help to facilitate a robust model in the face of unanticipated variability, such as is found in warfare. By modelling the constraints associated with the goals, abstract functions — including the principles of warfare — and the physical limitations of equipment and soldiers, it is possible to develop detailed models of the information required to perform military tasks. The inclusion of knowledge about the organisation structure and proficiency of both friendly and enemy forces allows the identification of the command strategies available for the IA to complete each task. The combination of the different models produced using CWA helps to identify relevant constraints and rules that may not be included in doctrine but is nevertheless present within the military system [2].

We must point out that, in our analysis, the relationship between the enemy intentions, command structures and capabilities are overlooked. The decision-making factors that affect the selection of the most viable function in C2 between human and computational agents is beyond the scope of this work. In this work, our objective is to prove the principle that CWA is an effective process for identifying and developing appropriate rules for an IA to conduct company level attacks on a known enemy. The future challenge will involve representing all the forces constraining company level decision-making in an integrated fashion so that their combined effect on activities — including attacking, defending, advancing and delaying — might better be inferred.

3. The Agent Framework and Infrastructure

The COMATS agent software utilises JACK Teams, a BDI Intelligent Agent system, which incorporates Team construct enhancements jointly developed by DSTO and Agent Oriented Software (AOS)[9]. The purpose of Command Agents is to undertake higher-level control of the entities within the wargame, letting the SE control the lower levels of command. The Agent framework models the command of the Company, and it's immediate sub-groupings. In this case a commander, and 3 platoons. The platoon agents are further defined by the roles they can perform, eg fire support and assault.

The agent architecture and its linkages to the SE and to its human operator are shown in Figure 1.

Internally each agent's plan set is split into a set of capabilities. The capabilities that are being used correspond to the decision cycle that is utilised by the military: the OODA loop. OODA (Observe, Orient, Decide, Act) is a simple cognitive model that is generally applicable to organisational structures that operate in a reactive manner. It has been successfully mapped to both air and land operations in the military context. The first two steps (observe and orient) are concerned with epistemic reasoning – ie reasoning about knowledge. The observe phase acquires data (status and detection) from individual simulation entities or command agents. This is then aggregated and analysed in the orient phase, providing situation awareness for the organisational structure in question. This situational awareness is then used in one of two ways – it either triggers an action via a decide activity, or it is used to decide an appropriate course of action in response to a command emanating from a higher level. Note that an action will affect the environment which in due course may result in new data being presented to the observe phase. This behaviour is summarised in Figure 2.

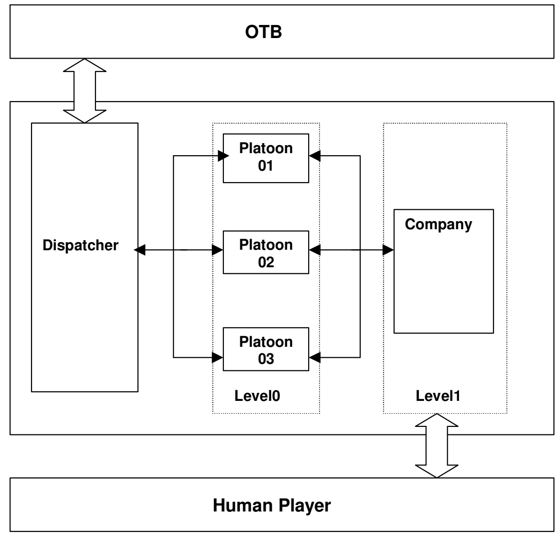

Figure 1: Command Agent Layer Architecture. Note that human interaction occurs only with the company Command Agent.

The main purpose for our usage of the OODA loop is to provide a conceptual model of reasoning that we can use to provide a structured approach to the organisation of the plans we develop for the agents.

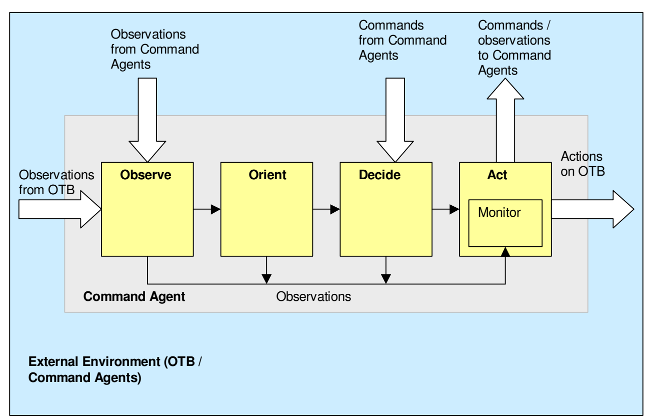

Figure 2: The OODA loop implementation in COMATS 1.0

The Company deliberate attack is a two-stage process: Courses of action are generated and defined in terms of units, tasks, timing and constraints. These tasks are then posted in parallel. The progress of each sub-team is monitored.

4. Courses Of Action

4.1 Knowledge acquisition

A range of different types of interrelated information is generated by the CWA studies and includes:

1) the objects and information pertaining to the area under study,
2) the constraints upon that information and
3) the decision paths, and resources used in performing the work.

In this study, we have taken the different information representations and produced a format that was amenable to representation within the IA framework, COMATS 1.0.

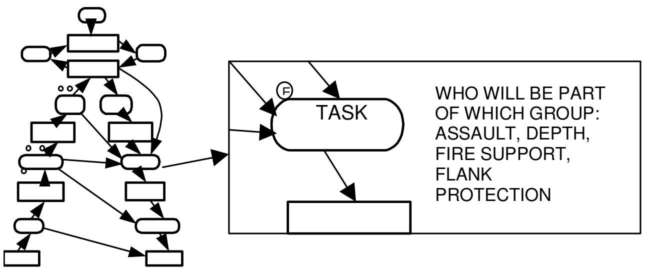

Figure 3: A decision ladder of COA generation and selection.

The planning process undertaken by a Company Commander when performing a company deliberate attack on an enemy platoon strength position supports two levels of decision-making:

1) generation of courses of action and
2) selection of a course of action

The reasoning involved in these two levels is quite different. When we generate courses of action, our primary concern is epistemic reasoning – we are using knowledge of the terrain, military doctrine, weapon ranges etc. to determine the most viable COAs given a scenario, whereas, in selecting COAs we are reasoning about action. The CWA process addresses these issues and provides rules to guide the generation of COAs and rules - bearing in mind that decision-making is at the company level only whilst the lower level constructive entities are controlled by the CGFs in the SE.

4.1.1 Spatial Knowledge

The army utilizes a process for planning a battle called the Military Appreciation Process (MAP). This process generates a large number of overlays that contain and illustrate different aspects of the environment in which the battle will take place. By combining these overlays a commander can quickly gain an appreciation of the factors involved and how they all contribute to influence the outcome, both of the planning process as well as their impacts on the battle itself. These overlays provide a convenient way of encapsulating the spatial information required and it was decided to use this idea as the basis for the majority of our knowledge representation. The spatial data that we have captured includes those listed in Table 1.

Table 1: Sources used to develop information templates for knowledge representation

Represented Knowledge	Source
1. Dynamic and interactive terrain	Obtained directly from SE terrain data.
2. Terrain culture including roads, rivers and tree lines.	Obtained directly from SE terrain data.
3. Line of Sight footprints	Generated algorithmically from terrain data.
4. Unit defensive position footprints	Template placed in upon known enemy positions.
5. Weapon Coverage	Derived by combining likely enemy positions.
6. Impassable terrain overlay	Generated from SE terrain data
7. Distance vs travel time	Generated from constructive entity status and SE terrain data.
8. Battlefield planning Control lines	Supplied by human-in-the-loop operators in the SE.

The majority of this information is generated by heuristics within the COMATS system. The data is encapsulated within an internal structure that contains and provides access to spatial queries. A variant of the A* algorithm [10] is used to generate optimal paths for traversing through these overlays. The heuristic that it uses encapsulates the effect of the different overlays on the commander's planning and the courses of action from which he will make a selection. This interaction between the path-finding heuristic and the overlays provides the total knowledge capture of the system.

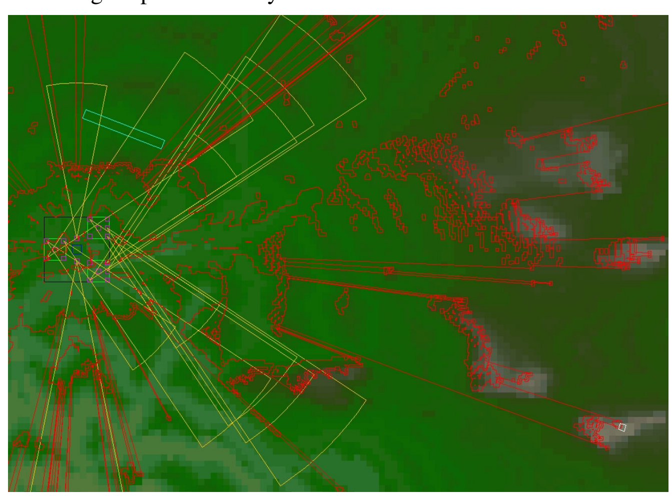

Figure 4: Terrain template in the command agent.

Figure 4 shows the following information. The red lines show the extent of the enemy LOS. The blue box to the top left is the form up position; the white box to the bottom right is the fire support position (note on high ground ~90 degrees from the assault direction) the yellow wedges are enemy weapon fields of fire.

Table 2: Potential formations and tactics to be evaluated during the COA generation

Assault Group (AG)	Depth Group (DG)	Fire Support Group (FSG)	AG Flank Protection	FSG Flank protection	AG frontage is greater than enemy frontage	Availability of FSG position	AG flank secured by another position or terrain or no known or suspected enemy to flank	FSG flank secured by another position or terrain or no known or suspected enemy to flank
1 Pl	1 Pl	1 Pl			Yes	Yes	Yes	Yes
1 Pl	1 Pl(-)	1 Pl	1 Sect		Yes	Yes	No	Yes
1 Pl	1 Pl	1 Pl(-)		1 Sect	Yes	Yes	Yes	No
1 Pl	1 Pl(-)	1 Pl(-)	1 Sect	1 Sect	Yes	Yes	No	No
Etc								

4.2 Representation of COAs

A COA defines the tasks to be performed by entities in a subordinate layer (in this case platoons) and the temporal and coordination constraints that must be satisfied by the activities if the COA is to succeed. Consequently a COA is specified as a list of tasks. For each task, the following information is provided:

1. Who is performing the task under which Role
2. The task to be performed.
3. The timing constraints: start by; finish by.
4. The other tasks upon which this one is dependent.
5. The phase of the battle that this task is in (pre-attack, attack, reorganisation)
6. The current status (not started, started, finished, unable to complete)

Note that a task may itself be a COA for subordinates of the entity to execute (e.g. flank protection is performed by the entities within a platoon). Also particular subordinates of all entities (e.g. vehicles) may be tasked as a single separate entity. In this situation, the entity (from the viewpoint of the COA) does not exist for the entire duration of the COA and meta-level tasks [11] to form and disband the entity are required.

A COA is represented within the Company agent as an agent belief. The containing BeliefSet stores all the generated COAs, as well as the COA that has been selected by the CA for execution. The selected COA is then passed to the platoons for execution. The information contained within a COA includes the start and finish times of each task, the task itself, as well as any other parties reliant upon the task. Each sub-team that has a task to perform will have a COA generated for it. If during the execution of a particular COA circumstances change, the agent has the other COAs that it generated and discarded as backup possibilities to refer to when replanning

4.3 Generation of COAs

How the above decision-making processes utilise the results of the CWA process is a key design issue. CWA decision ladders can be executed directly using JACK's meta-level reasoning [11] machinery. However this approach is only warranted when there is a choice of actions to be performed. Note also that it only addresses the action-based reasoning component of the CWA results, not the epistemic reasoning component. Given that COA determination involves a sequence of singular activities (generation/evaluation/selection) which involve predominately epistemic reasoning it was decided to encode the CWA outputs as plan statements and heuristics instead of directly representing the decision ladders. The heuristics encapsulated within the decision ladders, temporal control templates, and other outputs were encapsulated within sets of matrices.

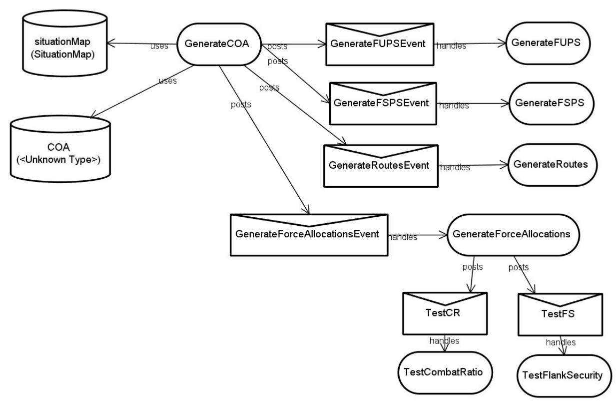

Figure 5: Plan structure and flow for COA generation.

The decision ladder shown in Figure 3 is an outtake of a decision ladder that details the development and selection of COAs. Table 2 outlines part of the matrix used to evaluate different potential force structures. In a company attack scenario, we are implementing a company of motorised infantry. The possible choices for the assault force include one of the three platoons, the infantry on foot or the company vehicles with no infantry on board. There are many variables that could impact upon that choice. It is through matrices derived from the decision ladders, temporal control tasks and other CWA outputs, that the COAs are evaluated.

The matrices are then mapped into JACK plans as a series of plans that then test each aspect of the matrix. A sample matrix is shown in Table 2. The list of force structure variants are that are determined to be plausible will be added to the agents Belief-Set as possible COAs. The CA will then weight each of them, using the criteria derived from the CWA and select the highest weighted COA for execution. All the generated COAs are still stored within the agent to aid in re-planning in case battle circumstances changes. Figure 5 shows the event flow and executed plans within the agent as it generates each COA.

4.4 Execution of COAs

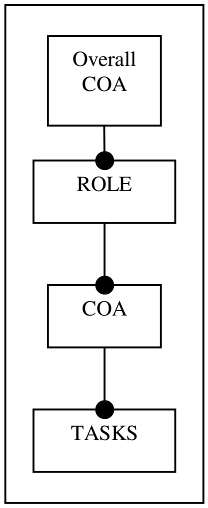

The Overall COA for the Company attack is split into a single COA for each Role being performed. These roles include fire support and flank protection. Each Role has its own COA that will comprise of many tasks and corresponding constraints. In a company, each COA is sent to the platoons in parallel. The Monitoring functions enable JACK-Teams to observe the tasks, some of which could fail, which are dependent upon another task. Those dependent tasks can be notified of the problem. Normally this would result in the termination of the dependent tasks, however it is possible to assess the reason the task failed and then dynamically decide whether the dependent task should continue or not.

5. Discussion

Although CWA is capable of capturing the uncertainties in a complex system that consists mainly of humans we have not modelled all of these. In COMATS 1.0, we have only modelled the human cognition based on Army doctrine and have not included human variability. A different weighting map for evaluating the factors influencing the COAs can be used to select the most viable plan. In future, it is desirable to collect more data in order to apply different commanding styles. In doing so, we will be able to incorporate different personality types in terms of defensive and aggressive actions.

The current version, COMATS 1.0, only considers command and control and decision-making at the company level. This leaves the control of individual entities, such as sections and individual troops, to the CGFs. This set-up can result in contradictions between the CGFs and the CA. For instance, in determining the arrival of troop units at a particular waypoint, the CGFs may determine that majority of the forces are within a set distance away from the waypoint, whereas the CA may be waiting for all members to arrive before acting.

Given that the SE is a crude model in many ways and that the CA and human operators are interacting with the SE via similar interface controls such as move, occupy defensive position, assault, dismount and mount infantry, etc., the information received at the CA and human operators is used to reason about points 1 and 2 (below).

1. The enemy's intent.
2. The next task required of the Company.

Some information fuzziness may arise as a result of insufficient data and lack of simulation fidelity. In future, it is desirable to investigate the possibility of implementing different types of activities such as defensive harbour, ambush, etc. so that the CA can perform contact drills and other mission with more reasoning power.

In determining COAs, the current COMATS version 1.0 does not consider points 1 and 2 above or models of human factors, such as morale.

Therefore, it is important to note that the current CA is a demonstrator to show how to map a cognitive model into an IA framework. We believe that this has been achieved, however, there is still a lot more work to be done in modelling different facets of battle. To do that, we need to look into other research areas such as human variability and error, human performance, cognitive task analysis, determination of enemy intents, etc.

6. Conclusion

The aim of project COMATS is to produce a prototype IA in land warfare simulations to look into the possibility of using the agents to interact with CGFs in SEs. It is desirable to have the IAs behaving in a consistent manner with those documented by the army in the doctrine and their tactics, techniques and procedures. Some further work is still needed to add flexibility into the planning and replanning processes.

A cognitive model was derived from a CWA study and we have described the process by which the model was derived and its integration into the Intelligent Agents.

It was found that the process of performing CWA study proved to be an effective means for eliciting the information required to construct an intelligent agent schema.

7. Acknowledgement

We would like to thank the COMATS development team: Geoffery White, Michael Balin, Peter Boyd, Matthew Lyons and Yale Maloney in Land Operations Division, DSTO, Edinburgh, SA, Australia. Lastly we would also like to acknowledge Australian Defence Simulation Office (ADSO), Fernhill, ACT, Australia for providing funding for this project.

8. References

[1] Frank Lui, Dennis Jarvis and Russell Connell, "An Architecture to Support Autonomous Command Agents for OneSAF Testbed Simulations", SimTecT conference 2002, Melbourne.

[2] Marcus Watson, "Modelling of command and control (C2) and Command decision making for a company commander using a cognitive work analysis approach", Contract Report for DSTO-Edinburgh, May 2002, Ref No. D 7305-20-341.

[3] Frank Lui and Marcus Watson, "Mapping cognitive work analysis to an intelligent agents software architecture: Command Agents", HFSIG conference 2003, DSTO Fisherman's Bend, Melbourne

[4] Rao, A.S. & Georgeff M.P. 1992, "An abstract architecture for rational agents", Proceedings of the Third International Conference on Principles of Knowledge Representation and Reasoning (KR'92) C.~Rich, W. ~Swartout & B..~Nebel, editors, Morgan Kaufmann Publishers.

[5] JACK Intelligent Agents TM JACK Teams Manual, Agent Oriented Software Pty Ltd, 2003.

[6] Anne-Marie Grisogono and Eyoel Teffera, "Towards a Synthetic Environment Design Methodology", SimTecT

2000 Conference Proceedings, Sydney Conventional and Exhibition Centre, Darling Harbour, Sydney, Australia, p273-278.

[7] Rasmussen, J., Pejtersen, A., & Goodstein, L. (1994). Cognitive systems engineering. New York: Wiley.

[8] Vicente, K. J. (1999). Cognitive work analysis: Toward safe, productive, and healthy computer-based work. New Jersey: Lawrence Erlbaum.

[9] Rönnquist, R., N. Howden, A. Lucas, J. Vaughan, R. Connell and S. Goss. "Coordinated Team Behaviour: A Simple Teams Approach". In: 2nd International Synthetic Environments Symposium, RMC Shrivenham, UK, October 1999.

[10] Winston, Patrick Henry, "Artificial Intelligence", Third Edition, Addison-Wesley 1993

[11] JACK Intelligent Agents[TM], JACK Manual, Agent Oriented Software PTY LTD, 2003.