11th INTERNATIONAL COMMAND AND CONTROL RESEARCH AND TECHNOLOGY SYMPOSIUM:
COALITION COMMAND AND CONTROL IN THE NETWORKED ERA

Convoy: Defining and illustrating core issues in the human factors of Command and Control using a minimalist game

Topics: Social Domain Issues / Cognitive Domain Issues / C2 Modelling & Simulation

Robert J. Houghton and Chris Baber
Human Factors Integration Defence Technology Centre, U.K.

Corresponding author: Dr Robert J. Houghton
EECE, The University of Birmingham
Edgbaston
Birmingham
B15 2TT
United Kingdom

Phone: +44 (0)121 4143141
Fax: +44 (0)121 414 4291
Email: R.J.Houghton@bham.ac.uk

Acknowledgements

This work is supported by a grant from the Human Factors Integration Defence Technology Centre, part-funded by the Human Sciences Domain of the UK Ministry of Defence Scientific Research Programme. The authors would also like to thank the members of the Human Interface Technologies Team at the University of Birmingham for useful comments during the development of the work.

Abstract

Evaluation of future command and control systems represents a significant challenge for human factors research and training. Approaches range from small-scale laboratory studies to large-scale experimentation, but these techniques often require functioning systems for study. This need for a functioning system represents something of a paradox in examining 'future' systems: if something has been developed to a functional level, then some of the key design decisions may often already have been made. Thus our aim was to produce a methodology that was flexible enough to encompass a range of hypothetical system designs and was transparent enough that it could be used to illustrate a range of points in classroom. The approach described in this paper is called "Convoy" and is a simple card-based game (all the necessary materials are included in the paper). Further to this we discuss some general issues in the design of games and simulations to address social and cognitive issues in C2 and describe some early findings from the Convoy paradigm that shed light on how networks of command can drive different patterns of team performance and serve as a "proof of concept" for the paradigm itself.

Introduction

A central part of the motivation for the adoption of network-enabled technologies is the anticipation that they will have an influence upon battlefield effectiveness via improvements in aspects of cognitive domain performance. These improvements include better information sharing, better shared understanding and better decision making (see MoD, 2004; see also Alberts et al., 2001). These socio-cognitive processes may be critically impacted by changes to doctrine or procedure that come with changes in equipment and furthermore, the presence of technology alone does not in itself guarantee positive outcomes in its use (see Houghton & Baber, 2005). This suggests a clear requirement for a range of pre-technology rapid prototyping approaches in order to explore as far as possible these human and social impacts at an early stage, because whilst technology can be altered, the human mind cannot. In the present paper we discuss a minimalist approach to C2 gaming that whilst incurring negligible economic costs is capable of, first, producing human quantitative data that speak directly to cognitive issues within command and control and second, may have a role in highlighting cognitive issues in an efficient and compelling manner for training purposes.

One popular approach to military exploratory gaming is the Tactical Decision-making Exercise (TDE) wherein the participant is posed a military or operational problem and given a short period of time to prepare a solution (normally in the form of a plan or a set of orders to subordinates). Whilst this approach is certainly of clear value in providing exposure to tactical problem solving issues in myriad problem spaces, it is less clear to what extent TDEs can capture social psychological complexities within command and control. Indeed, that TDEs allow the avoidance of these sources of friction and misunderstanding is argued to be one of their strengths as a training tool as it allows subordinates an insight into the decision making process of their commanders (see Brewster, 2002). As a paradigm for experimentation, which the TDE is useful for a exploration of new doctrine and systems it tends not to produce clean quantitative data for analysis. This means that feedback to performers will be qualitative and based on the interpretations of trainers and players. In itself this is by no means a bad thing but it does not readily provide a means of concretely measuring differences or improvements across iterations.

The aim of the present work was to produce a game of similar flexibility, speed of play and economic cost to the conventional TDE at the same time spoke directly to issues of team-working and the socio-cognitive processes of interest to modern C2 experimenters without sacrificing the transparency of the TDE which allows it to make high-level concepts visible to players. These various aims can be refined to the following set of design requirements:

- The game should be portable and playable without computers.
- The game should be relatively simple and quick to play.
- The game should not require specialist domain knowledge from players but optimal outcomes should result from skilful play.
- The game should be to some extent generic and capable of being reworked for different scenarios and audiences (e.g., civilian emergency services)

- Optimal play should result from co-operation and communication and the application of various C2 high-level concepts (such as shared situational awareness, promulgation of command intent, trade-offs between diversity and accuracy of shared knowledge etc.).
- Management of uncertainty should be present as a major factor in decision making

We turn now to describing game play and then move to discuss the factors that shaped our thinking in designing the game and lead us to reject some early designs. Finally, we discuss some experimental findings produced using the "Convoy" paradigm and review these with reference to the work of other researchers.

1. Gameplay

An experimenter mediates game play; all the players are on the same side and occupy various different roles. The exact details of these roles will vary with organisational structure, but in general players are placed in roles involving varying amounts of command, reconnaissance and strike activity. The game is played on an "epoch" basis in which all players act simultaneously. Players produce actions (which may impact on the field or play or be communications to other players) by filling in turn cards and handing them to the experimenter who either enacts actions or distributes communications appropriately, following which the next turn or epoch begins.

The scenario

Your force has been charged with protecting a land convoy of time-critical humanitarian aid. In [so many] turns, the convoy will have to pass through a city known to harbor enemy forces that will seek to plunder the convoy. Therefore, before the convoy reaches the city your unit has been issued the mission of seeking out a safe path through the city and neutralizing threats as necessary to reduce the risk to the convoy when it arrives. The end product of your activities will be to hand a route plan to the convoy as it enters the city depicting what you have established as the best route to take and indicating what "hotspots" remain to help the convoy's escorts plan their operations.

1.1. The city

The city is defined as a grid made up of blocks (we have found that using a map as a background to the grid can be useful, but is not essential). Each block is defined by four threats cards (bearing values from 1 to 4) placed face down by the experimenter. Thus each block will have a threat rating that ranges from 4 (4x1) to 16 (4x4). At command level, a sector is considered to be made up of four blocks. These stacks of cards essentially form the terrain the game is played within. The usual layout for the city are dimensions of 4 x 8 blocks (a total of 32 blocks or 8 sectors) but any layout can be chosen. One may also opt to map different blocks onto, say, sections or waypoints on a map.

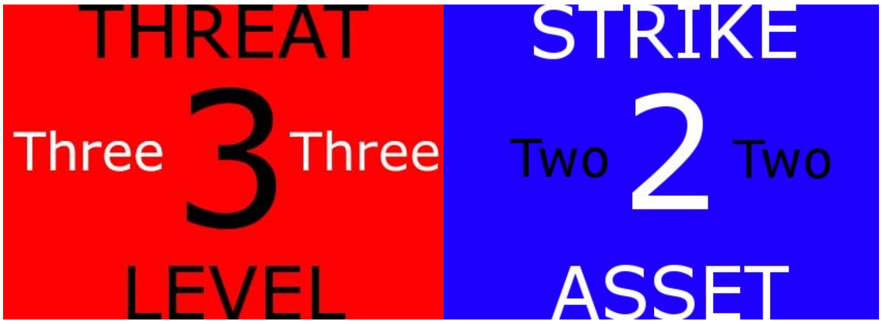

Figures 1a and 1b. Threat level and Strike asset cards.

1.2. Forces given to players

There are two types of things the team can do with regard to the city. A reconnaissance asset may pick up two cards from a stack per turn (i.e., visit one block per turn). Given that each stack consists of four cards this yields 50% of the total information. Thus for one reconnaissance asset, a block will take two turns to fully survey (that is, to pick up 2x2 cards). Alternatively a block can be fully assessed by two reconnaissance assets within a single turn.

Strike assets are issued with a limited deck of strike cards, again numbering from 1 to 4 (see Figure 1b). These can be used to "trump" threats. Each strike asset can deploy up to two strike cards per turn to a single block. A threat card is trumped when the strike card is equal to or greater than its face value (e.g., a strike of 3 will trump a threat of 3, 2, or 1 but not a 4). Given asset cards are limited there is a need for them to be used efficiently. Strike unit players are not given direct feedback about the effectiveness of their strikes. Given play is effectively blind the experimenter will always seek the best outcome for the team by applying the highest strike to the highest threat. The team may check effect by dispatching a reconnaissance unit to a previously targeted block but this could be time consuming. In terms used in Effects Based Operations, knowledge of Strike cards deployed would represent a Measure of Performance (MOP), the reports of Reconnaissance players to cells where Strike actions have been previously directed would indicate a Measure of Effect (MOE) (see Joint Forces Command, 2005).

1.3. Flow of communications and information

The flow of communications and information will vary with network structure, so we shall describe the general case from a typical hierarchical "split" network (see Figure 2) where granularity of information becomes coarser as it rises through the hierarchy.

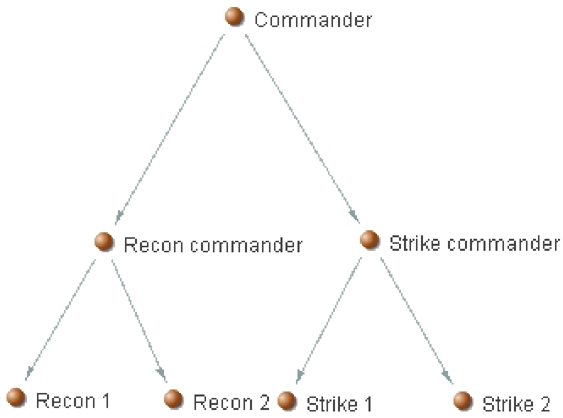

Figure 2. Hierarchical Split Network (after Dekker, 2002).

Step 1. Reconnaissance players Recon 1 and Recon 2 will pick up two threat cards per turn from a block of their choice, or to which they have been directed. The outcome (ie. Card 1 + Card 2) is summed and a communication can be issued, in the split network this will be to the reconnaissance commander.
Step 2. Strike players Strike 1 and Strike 2 will (optionally) deploy two strike cards per turn. They can communicate the summed strike force deployed (ie. Card 1 + Card 2) to the Strike commander.
Step 3. Reconnaissance or strike commanders will receive orders from the Commander on a communications card. They can then issue a command to survey/strike a city block to each of their subordinates. They can also issue a card to the commander describing the action taken by the force as a whole.
Step 4. The commander will likely receive information from both the reconnaissance and strike commanders. To each he can issue a card describing what he wants their next actions to be.

A blank order card is shown in Figure 3. Note that the meaning of choice A will vary depending on the task of the person it is sent to; up the chain of command it would be a reconnaissance/strike report whereas down the chain of command it would represent an order to take action.

FROM ____________ **TURN #**

TO ____________

One choice only:

A. BLOCK __________ **THREAT LEVEL** __________

B. NO ACTION

C. TEXT MESSAGE **(must fit on the line)**

_ _

Figure 3. An order card

1.4. Sitrep

At various points in the game (we have found five turns to be a reasonable interval) situational awareness will be recorded by asking players to indicate what where they think risks are still present in the city and what they think the best route through the city would be (see Figure 4). We can compare this between players and with ground truth to assess the quality of SA and the degree to which there is consensus within the team. At the end of the game the plan given to the convoy by the team takes the same form as these sitreps.

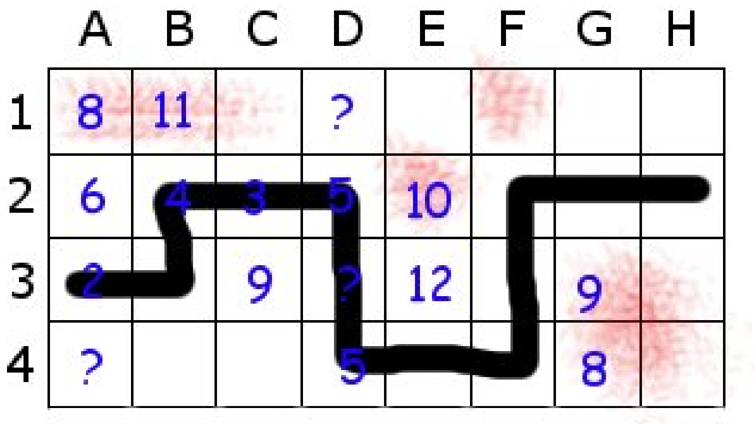

Figure 4. Sitrep and planned route at a midway stage in the game. The black line indicates the planned route for the convoy, the numerals the perceived threat level.

2. Design requirements and game features.

The best way to elaborate further on the features of the game is with reference to how they address the design requirements we considered in when devising the game:

2 1. The game should be portable and playable without computers.

The game is card based; the terrain of play is defined using stacks of “threat” cards picked up by reconnaissance asset players, offensive assets are defined in terms of “strike” cards which strike assets players deploy to deal with threats and finally actions and communications are made through the exchange of “order” cards that the players write upon. These order cards also act as a complete record of game play that can be consulted for research or debriefing purposes. They also have the action of standardizing communications, and minimizing the impact of especially verbose commands or “off network” communication (that is, communication not allowed with the specific command structure in use).

Maps are useful (but optional) and are anticipated to be either photocopies or simply be sketched on blank paper. All the materials required to play the game would probably fit into a modestly sized lunchbox and can be produced at negligible cost. This simplicity of approach also makes the game flexible and thus useful for studying a range of command and control issues. We discuss possible manipulations elsewhere, but its to be noted that in more complex simulations with more stylized and complicated rules it can often be hard to introduce experimental manipulations without breaking analogies and fine balances within the game or else requiring significant effort to reprogram etc.

2.2. The game should be easy and quick to play, requiring no specialist domain knowledge.

The game mechanic is based on a simple card-drawing and trumping approach (but see Management of Uncertainty below). As such then it is probably simpler to understand than most ‘nursery’ card games. The complication within the game is one of communicating the optimal course of action, the perception of the situation vs. ground truth with reference to reconnaissance and reports of action taken, and managing scarce strike assets. To this extent then, the military mindset will be of clear value in producing good performance but at the same time naive participants can also play the game with relatively little instruction; the emphasis is upon *how* one plays, not the complexities of play itself. Furthermore, command and control is an issue with scope wider than the military alone and some overly-stylised designs can be offputting to civilian organisations (e.g., the emergency services).

A further requirement was that the game should be capable of being played fairly rapidly (in minutes rather than hours or days) to produce several iterations perhaps comparing different experimental manipulations on a “within subjects” basis within a reasonable time frame. In a training context this would also allow players to have experiential learning of the differences between, say, various command networks, in a single session.

2.3. The game should be to some extent generic.

Given that the game is based on drawing numbered cards and trumping them with other numbered cards, this speaks to a wide range of asset management situations (be it dispatching strike assets to deal with strategic threats, fire engines to fires, police to trouble spots or even taxis to fares). Alternative military scenarios to convoy route planning might also be embodied. That said the game was designed to primarily address to the military situation of finding a safe route through urban terrain, so the binding narrative that makes the dispatch of assets a complex goal-based activity following command intent rather than merely demand/response-based activity would need a degree of reworking. The presence of a goal is essential to evoking and measuring many of the phenomena of interest, such as Shared Situational Awareness and the promulgation of command intent.

2.4. Optimal play should result from co-operation and communication.

One core problem in designing a game to represent C2 activities is that one must place C2 at the centre of it; there must be some clear benefit in game performance from the active pursuit of command and control to produce coordination and cooperation. This achieved by the game unrolling such that the deployment of force (strike assets) is most efficient when directed with reference to information depicting ground truth (reconnaissance assets) but in turn this demand-action loop is only of benefit to the team when it is in the service of command intent to produce a safe route through the hostile city. Drilling down a little, there are also different strategies and results that arise from the coordination of reconnaissance itself and the coordination between strike units (see the experimental results below).

2.5. Management of uncertainty should be present as a major factor; relevance to Network Enabled Capability and Effects Based Operations.

A central issue in command and control is the management of uncertainty upon the battlefield, both with regard to understanding what one's own forces are doing (and how well they are progressing) and what may be a shifting ground truth. In turn, the coordination of knowledge about the battlefield is critical to meeting these problems and forms a central tenet of NEC and EBO. We embody this in the game in a range of ways. The game typically begins with a state of complete uncertainty; the level of threat in each city block is unknown to any of the players. Reconnaissance units can be deployed to discover the level of threat; the division of reconnaissance forces (x2 in the standard design) means that commanders have a range of options balancing depth of reconnaissance against geographical breadth/time costs. Further, the "trumping" game mechanic means that the impact of strike activities will vary in its effectiveness. Given there is a range of threat levels and a range of responses, the degree of certainty one has as to the (veridical, ground truth actuality) impact of strikes will vary considerably in terms of the match between threat and response, although these two variables are not as such directly proportional. In other words, this captures something of the tension between reconciling Measures of Performance (MOP) which might be a measure of the number of troops deployed or ordinance dropped or money spent with Measures of Effectiveness (MOE), that is, the effects these actions have actually have on the ground.

Within the game in general there is potential for generating data suitable for various cognitive modelling approaches through the use of sitrep cards as cognitive probes and the manipulation of the availability and certainty of information within the game.

2.6. Issues of parameter-space size and complexity.

Finally we turn to some general issues that have shaped our thinking. Here we speak perhaps not directly to specific features of the design but rather why we rejected various other options too numerous to list here. One issue in designing experimental paradigms is to permit sufficient parameter space; participants must be able to perform well but also to perform badly, and between these two poles (termed "ceiling" and "floor" respectively by statisticians) there must be sufficient differences for subtle gradations of performance to be apparent in response to experimental manipulations. A design where performance is essentially binary (good or bad) is not very useful when considering the relative impacts of different independent variables. At the same time however it must be possible to measure performance succinctly; think for example of a Chess game. Outside recording simply who won and who lost its very difficult to objectively assess whether a "good" game has been played, much less where a game might lie on a more granular scale (is a certain game *very good*, is it better or worse than another specific game?). In general then we think the Convoy game has reasonable potential to produce varying patterns of performance whilst at the same time being simple enough to directly assess performance and related factors (such as situational awareness, coordination, promulgation of command intent etc). Considering the game of Chess also raises another issue; at the present time despite great steps forward in computer design and mathematics we are still not in a position to analytically “solve” chess once and for all, something that makes the game so evergreen in its fascination for players. By contrast Noughts And Crosses (Tic-tac-toe) is so simple a game that its analytical aspects rapidly become obvious such that there is little fun to be had in playing it. There are a very limited number of ways a game can unfold. Again then, we wanted to find a “sweet spot” in game complexity that allows a range of complex behaviour and play but is at the same time open to some form of operational analysis and cognitive modelling.

3. Manipulations and metrics

As an experimental paradigm, the aim of the exercise is to measure various human variables and to relate those to measures of overall game performance. A further aspect is that the game was designed to be fairly adaptable meaning that a range of experimental manipulations can be made to study performance in different conditions, using different procedures, different social/command networks, different technologies etc.

3.1. Metrics (Dependent variables)

3.1.1. Promulgation of command intent

A major issue in command and control, particularly as we enter the network-enabled era is the promulgation of command intent; to what extent has the commander been successful in communicating his plan to subordinates. Variations on this theme (especially with regard to NEC) are whether it is possible for relatively autonomous units to between them converge upon a shared plan. Within the game, the selected path through the city represents command intent. Thus, to measure the promulgation of command intent at a given point in the game, one assesses the match between the route on the commander’s plan and the sitreps drawn up by the other players. Simply counting which squares match will give a quantitative score and correlation techniques can be used for deeper analyses.

3.1.2. Shared Situation Awareness

Another major issue is the measurement of shared situational awareness; to what extent do members of a team share awareness of the situation? Again, the sitrep forms are used and this time we would look at the match between numeric threat estimates by cell. *Comparisons of congruency between players* indicate to what extent awareness is shared. *Comparisons of congruency between sitreps and ground truth*, that is, the actual threat levels of the blocks as shown on the threat cards themselves, gives a measurement of the accuracy of situational awareness. The dissociation between “sharedness” and “awareness” will be useful in examining the extent to which shared information may actually be damaging to or enhancing of individual situational awareness.

3.1.3. Task performance

By measuring task performance, as experimenters we can understand the contribution that command intent and SSA make to the outcome of a game. Given that in the scenario players are charged with handing a plan (the same as the other sitreps made throughout) to the imaginary convoy's escorts, the simplest measure of performance is how many "threat points" the convoy will face on the finally chosen path through the city. Again, a dissociation between perceived truth and ground truth can be made by contrasting the amount of "threat points" the players believe will be encountered (this can thought of a measure of planning success given the information the players have collated) with the actual amount of "threat points" on the cards themselves. Subtracting one total from the other will give an indication of the degree of any intelligence shortfall.

Other task performance measures are possible and may indicate different thing or characterize particular patterns of team performance, we imagine they may include;

- Number of communications issued
    - Number of redundant communications
- Ratio of actions to turns (will imply something about the degree of coordination and operational tempo)
- Adherence to orders (given the experimenter has a full record of the game on cards comparisons can be done between orders given and responses made)
- Number of blocks visited by strike units relative to the number of blocks visited by reconnaissance; clearly if strike roved further than recon this would imply some "shooting blind" (albeit perhaps on the basis of informed judgement) was going on.
- Strike unit performance; the more accurate the information given to the strike units and the better their actions are coordinated, the better their performance will be. This strike performance can be measured variously:
    - Number of strike cards left (sum total of shared values). This might also be considered relative to the eventual level of threat in the final convoy route through the city; this would be suggestive as to the extent to which resources were effectively used and managed.
    - Number of strike cards played that were ineffective
    - Number of strike cards played that were "overkill" (e.g., a strike of 4 played to trump a threat of 2)

3.2. Experimental manipulations (Independent variables)

3.2.1. Command networks

The game can be played using different networks for command and information sharing. If communications are restricted to cards (as in the basic form of the game described here) then the experimenter can enforce various structures on players merely be limiting who can exchange information with who. This can be used to address various open questions in the NEC field with regard to the ability of teams to self-organize (the so-called "synchronization" effect) and the extent to which the relaxation of restrictions on independent action can result in rapid, high-tempo performance in the presence of Shared Situational Awareness. Thus we might try a range of networks from the strictly hierarchical, to "sensor to shooter" structures through to "peer to peer" flat networks. Another axis of comparison is between heavily centralized networks and broadly distributed ones. We expect these networks to manifest different behaviours and levels of performance; probably with some trade-offs between different metrics becoming evident.

Other variations include:

- Distribution of roles with in a team (balance of recon to strike for example)
- Number of players
- Flexibility; some players might have a variable multi-tasking capability (individuals can both do recon and strike at different points)

3.2.2. Nature of communications

Given the game takes a simple desktop form there is considerable freedom to look at different types of communications. In the basic form of the game communications are restricted to writing on cards. But things like "shared representations" can be introduced simply by allowing participants access to a white board or to write on a projected OHP transparency. Shared chat systems etc. can be brought into play through simply allowing different groups of players to talk to each other.

3.2.3. Intelligence and system degradation

A further issue that can be looked at is the use of fully shared information prior to the commencement of the scenario. This may take the form of a report detailing approximate strengths of threat in various blocks or sectors. There is some evidence in the social psychology literature that teams tend to over-value shared information given be the commencement of a task and find it hard to reject it even when some team members decide it to be false (often so-called minority information team members are ignored despite the salience and relevance of their information) (see Houghton & Baber, 2005 for a review in the context of NEC systems).

A related issue to the foregoing is how various communication systems and command networks fare in the face of system degradation. Again, due the low-tech nature of the game these breakdowns are fairly easy to manifest. The experimenter might for example give intelligence assets incorrect or otherwise inconsistent information. Strike cards might (contrary to the stated rules of the game) simply fail to have the predicted impact from time to time. To simulate a changing battlefield the city block cards might be reshuffled, rendering the intelligence currently possessed by the team out of date. The experimenter withholding communications cards might simulate a failing communications system.

4. Experimental findings

To-date we have restricted our studies to examining the differences in performance produced under different command and control network topologies. We have taken inspiration in this from the work of Dekker (2002) who compared various network topologies using a computerised agent simulation approach in conjunction with the Scudhunt paradigm (Perla et al., 2000). Scudhunt itself is a game of command and control that has some familial similarities to Convoy although it is computer mediated in nature and arguably more stylised with regard to the roles players undertake. Our aim was to see if Dekker's computationally generated findings using Scudhunt as a paradigm would generalise to our human experiments using Convoy in order to cross-validate both of our paradigm and the generality of Dekker's insights (see also Dekker, 2001). The studies were not exhaustive and represent an element of work currently ongoing. Our intention is to offer a "proof of conept" and to illustrate how our relatively simple procedure can generate perhaps surprisingly rich data.

Participants

A total of 26 students were drawn from the Engineering Department at Birmingham University and the Business School at Aston University participated in the experiment for course credit.

Design and procedure

Two types of command network configuration were compared, resulting in two experimental conditions. Each of the four groups that participated in the experiment were restricted to single network architecture and played through a game using only that architecture. Roles within the game were randomly allocated. The two network configurations (Hierarchical and Decentralised) compared are shown in Figure 5. Notable differences are that in the Hierarchical network there is a central commander and assets are organised by function (resulting in a strike commander and a reconnaissance commander who pass condensed information to the commander for their subordinates), whereas in the Decentralised network each strike asset is coupled with a reconnaissance asset via a single command node resulting in two self-sufficient teams. Each game began with a standardised instruction set being given to participants and the experimenter talking through a set of example turns by way of demonstration. As regards parameters within Convoy itself, each game lasted 15 turns and Strike assets were generously allocated so that overall eight cards of each value (1, 2, 3 and 4) were given to the team as a whole. Communications were restricted to order cards (no verbal communication) to make sure they were restricted in line with the network topology in play.

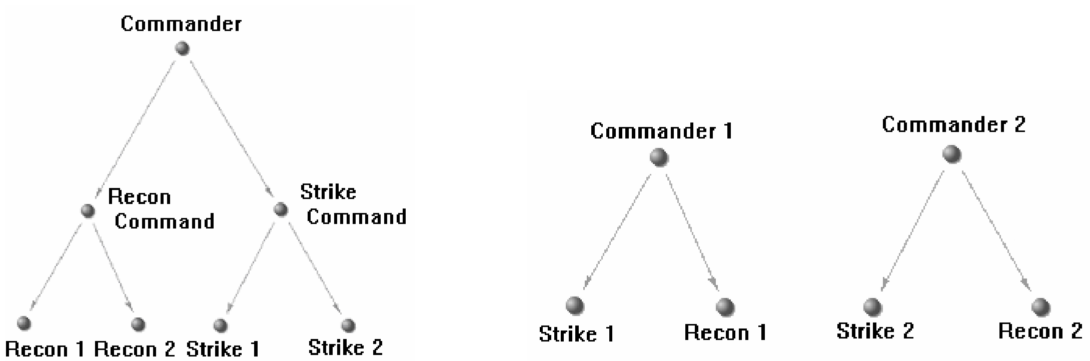

Figure 5. The left-hand panel shows the Hierarchical network, the right-hand panel the Decentralised network.

Results and discussion

The findings are summarised in Table 1. Indicative averages of paths chosen by the majority in each team for the convoy to take at the end of each game are shown by condition in Figure 6. There are number of notable features to the data. First, as can be seen in Figure 6, the paths chosen by teams using the two showed considerable differences, the Hierarchical groups being characterised by favouring a linear “brute force” approach and concentrated all their assets along that path. By contrast the Decentralised groups opted for a far more convoluted approach and were more wasteful in their use of Strike assets on squares that ultimately did not feature in the final planned route for the convoy.

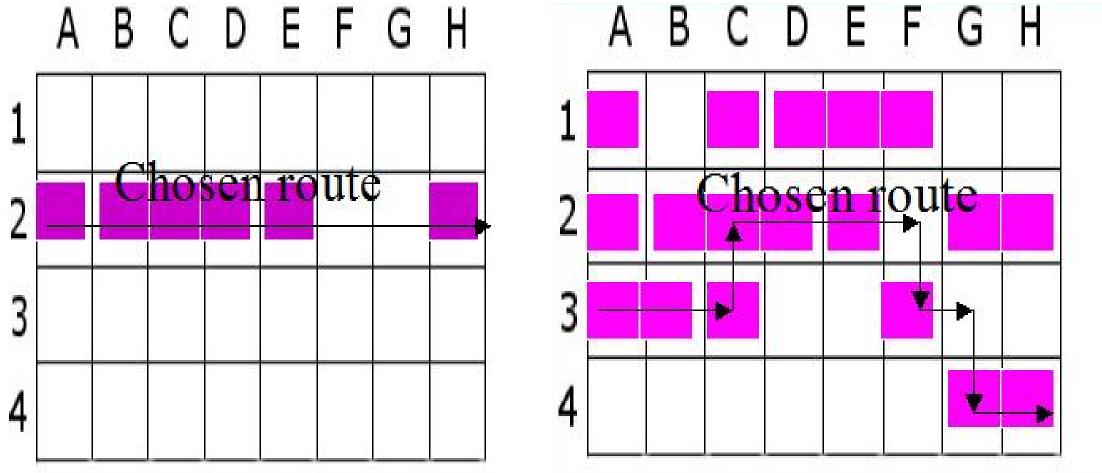

Figure 7. Representative routes chosen by teams. The left-hand panel shows the Hierarchical network, the right-hand panel is for the Decentralised network. The shaded squares indicate those where Strike assets were deployed.

In order to make sense of these very different patterns of behaviour we can draw on the various metrics reported in Table 1. It seems apparent that the Hierarchical groups was relatively slow moving in its decision making as the path from “sensor” (Reconnaissance assets) to “shooter” (Strike assets) was a long one taking at minimum some four turns. A relatively low level of threat detection across the trials (54 points found) suggests that reconnaissance assets were underused as a result. However, this slow pace (an average of only 4 strike points deployed per turn out of a minimum of 0 and a maximum of 4x4=16) and the central role of the commander did however mean that Strike assets were not wasted (no strikes were made on unnecessary targets) and that there was 100% agreement within the group as to the route the commander had in mind. In the Decentralised groups the link between sensor and shooter was only two steps. This appeared to result in a far higher tempo of operations, with higher levels of threat discovery and strike actions taking place. However, the lack of a single central commander apparently had serious implications in that action was simply not coordinated. That the Decentralised group’s ultimately chosen path included featured only 35 points of threat remaining (as compared to the slightly better Hierarchical groups’ 30 points) seems to be a function of their high pace of operations, albeit that some of these actions might be considered wasteful and at times rather hasty.

METRICS	Hierarchy	Decentralised
Threat detected	54	84
Redundant threats	0	49
Points per turn	4	6
Agreement on route	100%	70%
Threat remaining	30	35
Commander awareness	29	0 and 11
Strike awareness	29	0 and 24
Recon awareness	33	35 and 35

Table 1. Results from experiment comparing performance using different network topologies.

There is an interesting difference between the two groups that cleaves along the distinction between Shared Situational Awareness (SSA) and accurate situational awareness that is there was a distinction between “sharedness” and accurate awareness. Within the Hierarchical group there was relatively poor agreement between individuals as to the situation, presumably a function of the long paths involved in communicating information around the team. However, there was a reasonably good agreement between an overall average of the sitreps and ground truth. For the Decentralised group we see a reversal of this trend; within groups there was a high level of SSA but perhaps because of the absence of central coordination these situational estimates were often in error because of a lack of visibility of what others within the team had just done. It is notable that whilst both the Hierarchical and the Decentralised networks ultimately produced similar levels of overall task performance (as measured in terms of threat remaining along the chosen convoy route at the end of a trial), we are able on the basis of the metrics offered within the Convoy paradigm to suggest that these networks reached that level of performance for quite different reasons. In general, the qualities we have found in these two networks match those that Dekker’s simulation work had suggested would be the case (Dekker 2001, 2002).

Whilst we did not formally assess the participant’s learning from participation in the experiment, we did monitor group discussions after each game. In general we found that participants had fairly strong opinions about how their teams had performed and the challenges they faced in their particular roles and were able to easily relate these to the command network structures and high-level command and control concepts. In games that had used the Hierarchical network, strike and reconnaissance players complained of being under-utilised and suggested that lines of communication across lower levels of the hierarchy may have been valuable to allow them to make self-directed actions in the absence of orders from command. In the Decentralised network a lack of communication between teams was identified as a key cause of problems in coordinating activity to a shared goal. Formal validation of learning outcomes would be a valuable area for further work to address.

In summary then we have observed between the Hierachical and Decentralised network structures a set of trade-offs as concerns various features of command and control performance: tempo of operations, shared understanding of command intent, shared situational awareness, accuracy of situational awareness, speed of communication, coordination, and the active management of uncertainty. Whilst the Convoy paradigm is a relatively simple one and makes few claims to ecological validity (for example, the tight restrictions on brevity of communication are fairly stark) we believe it can produce useful and perhaps provocative data. Furthermore, we believe that Convoy's simplicity and transparency make it a useful tool for making visible to participants in a classroom setting the nature and relevance of high-level command and control issues.

References

Alberts, D. S., Garstka, J. J., Hayes, R. E., & Signori, D. A. (2001). *Understanding information age warfare*. Washington, DC: CCRP Publication Series.

Brewster, F. W. (2002). Using tactical decision making exercises to study tactics. *Military Review*, 82, 3, pp. 3-9.

Dekker, A.H. (2001). Applying Social Network Analysis Concepts to Military C4ISR Architectures**.** *Connections, the official journal of the International Network for Social Network Analysis*, 24, 3, pp 93–103.

Dekker, A. (2002). C4ISR Architectures, social network analysis and the FINC methodology: An experiment in military organisational structure. *DSTO report DSTO-GD-0313.*

Houghton, R. J., Cowton, M. & Baber, C. (2005). WESTT (Workload, Error, Situational Awareness, Time and Teamwork): An analytical prototyping software tool for C2. *The 11th DoD ICCRTS, McLean VA, June 2005.*

Houghton, R. J. & Baber, C. (2005). Social aspects of NEC: Information sharing and decision-making. *IEE People and Systems Conference: Who are we designing for?, London, UK, November 2005.*

Joint Forces Command (2005). The multinational effects-based operations process concept of operations (CONOPS). http://www.act.nato.int/events/documents/mne4ws1docs/eboconopsv065.pdf.

Ministry of Defence, UK (2005). *Network enabled capability: An introduction*. London, HM Stationary Office.

Perla, P. P., Markowitz, M., Nofi, A., Weuve, C., Loughran, J., & Stahl, M. (2000). *Gaming and shared situation awareness*. Centre for Naval Analyses Report CRM D0002722.A2

Appendix – Convoy Materials

STRIKE 1 LEVEL	STRIKE 2 LEVEL
STRIKE 3 LEVEL	STRIKE 4 LEVEL
THREAT 1 LEVEL	THREAT 2 LEVEL

THREAT
3
LEVEL

THREAT
4
LEVEL

	A	B	C	D	E	F	G	H
1								
2								
3								
4								

FROM ____________ **TURN #**

TO ____________

One choice only:

A. BLOCK __________ **THREAT LEVEL** __________

B. NO ACTION

C. TEXT MESSAGE **(must fit on the line)**

_ _