17th ICCRTS

“Operationalizing C2 Agility”

Analyzing Team C2 Behaviour using Games and Agents

Paper # 006

Primary Topic 6: Modeling and Simulation
Topic 5: Experimentation, Metrics, and Analysis
Topic 4: Collaboration, Shared Awareness, and Decision Making

Dr Anthony H. Dekker
DSTO Joint Operations Division
Defence Establishment Fairbairn
Department of Defence, Canberra
Australia
dekker@acm.org

Analyzing Team C2 Behaviour using Games and Agents

Dr. Anthony H. Dekker (Defence Science and Technology Organisation, Australia)

Abstract

In previous work, we announced (but did not fully describe) an experimental tool to study the behaviour of teams solving a simple logistics game, where six hypothetical platforms are assigned to six hypothetical areas of operation. This tool permits C2 experiments, at a slightly smaller scale than those conducted with the ELICIT tool. The tool also includes a simple chat facility that can enforce a specific network topology between team members.

To supplement human experiments with the tool, we have developed a version in which human beings are replaced by very simple intelligent agents. The agents communicate by sending and decoding text chat messages which describe facts about the situation, and preferences which the agents have for a specific solution. Because the agents use a chat facility, hybrid experiments involving a mixed human/agent team become possible, and we have also developed a third version of the tool for such hybrid experiments.

In this paper we report preliminary experimental results produced to date. In particular, we report a linear relationship between decision time and network diameter. We also discuss ways of upscaling the tool to deal with more complex and more realistic problems suitable for use in command-post exercises.

Introduction

Command and Control (C2) of military operations is a complex activity (Alberts and Hayes 2006, 2007) involving coordinated activity by networked teams of various kinds. Understanding the fundamental behaviour of such teams is crucial in optimising the performance of C2 systems.

However, useful principles for C2 can be derived using simple agent-based models such as ISAAC (Brandstein *et al*. 2000). Collaborative team games are also a powerful experimental tool, since they allow the exploration of real human behaviour. Although military-inspired games such as CAFFEINE (Huber *et al*. 2006) and Island Mission (Dekker 2007c) may have a variability in outcomes that makes statistically significant results difficult to obtain, this issue has been overcome in the more successful collaborative games such as SCUDHunt (Perla *et al*. 2000, Dekker 2006b) and ELICIT (Ruddy 2007, Thunholm *et al*. 2009).

In previous work (Dekker 2011a, 2011b) we briefly presented a new collaborative team game based on solving a version of the assignment problem (Christofides 1975, Dekker 2006a), finding the best one-to-one match between a set of six notional platforms and a set of six areas of operation. Each of the 36 possible pairings has a textual description of appropriateness. Counting a neutral mid-case, these descriptions form a seven-point scale:

- Stingray must not be used in Orangeland
- Piranha should, if possible, not be used in Greenland
- Wolf is not ideally suited to conditions in Blueland
- (neutral)
- Falcon can operate well in Scarlet City
- Eagle should, if possible, be used in Whiteland
- Puma must be used in Yellowland

Areas of Operation

Platforms	Orangeland	Greenland	Whiteland	Yellowland	Scarlet City	Blueland
Puma		Avoid if poss.	Operates well	Avoid if poss.	Use if possible	
Falcon	Not ideal	Must be used	Use if possible		Use if possible	Operates well
Piranha		Not ideal			Not ideal	
Stingray	Operates well	Operates well		Not ideal		Operates well
Eagle			Operates well		Not ideal	Not ideal
Wolf		Use if possible	Do not use	Not ideal	Operates well	Operates well

Figure 1*. An instance of the game and an optimal one-to-one matching between platforms (Puma, Falcon, etc.) and areas of operation (Orangeland, Greenland, etc.).*

Figure 1 shows an instance of the game. The game is supported by software tools for:

- automatically generating instances of the game;
- showing players the textual descriptions applying to a specific instance of the game;
- recording a player's decision about the best match of platforms to areas of operation;
- measuring the time taken by a player and the quality of the solution produced; and
- providing text chat between game participants (with variation of the "shape" of the team network permitted).

The game exists in four versions: a single-player version, a six-player team version, a fully automated intelligent-agent version, and a hybrid version in which a single human player interacts with five agents.

Single-Player Version

In the single-player version of the tool, the player must produce a complete solution consisting of six pairings. Figure 1 provides an example of such a solution. Experimentation indicates that a few instances of the game are needed for learning, after which the problem can be solved in an average of 3.48 minutes (Dekker 2011b).

The textual descriptions for the various pairings were deliberately intended to be non-trivial to interpret (Dekker 2011a). It was inevitable, therefore, that slight differences in interpretation between players would arise, and that players would produce solutions which were not optimal in terms of a "standard" interpretation of the descriptions. The tool therefore accepts sub-optimal solutions, unless mandatory conditions ("must be used" or "must not be used") are breached, or if a solution fails to assign every platform exactly once.

Team Version

Figure 2 shows the concept of operation for the team version of the game, together with the decision and chat tools. This version is intended to be played by a distributed team of six people, each in charge of assigning one platform. The need to assign every platform exactly once means that coordinated decision-making is required – a process also known as self-synchronization (Alberts and Hayes 2003). Furthermore, the textual descriptions for the various pairings are divided among the six players, so that information sharing is also required. Since good performance requires both coordinated decision-making and effective information sharing, performance will be a good indicator of the quality of teamwork by the participants.

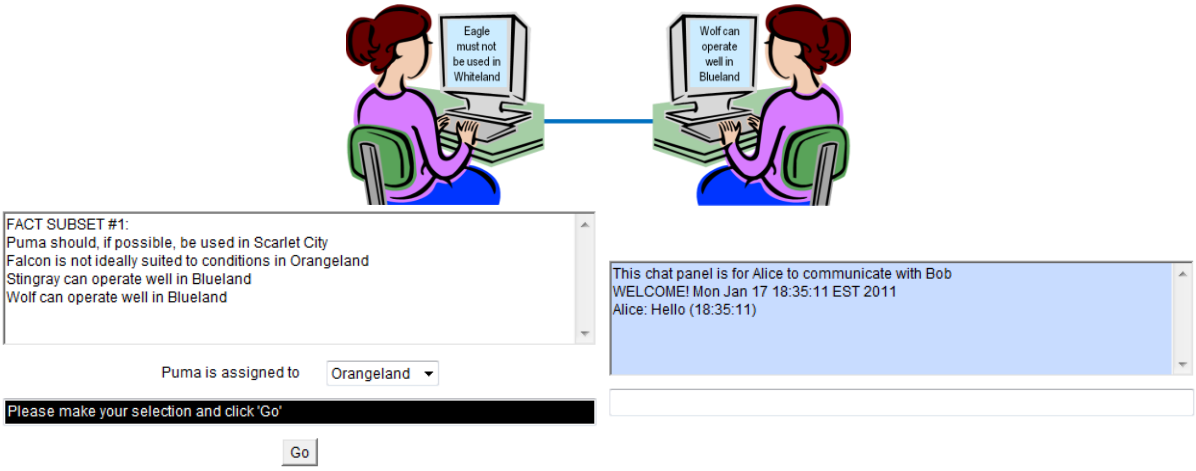

Figure 2. *Team game concept of operation and distributed team decision & chat tools.*

Communication between team members is restricted to the topologies shown in Figure 3. These topologies were chosen to produce the maximum possible range in network diameter (the maximum distance between people in the network, measured in network "hops"). The team version of the game can be played either by collaborative decision-making, or by sending all the data to a designated leader, who makes all the decisions (treating it as a single-player game) and broadcasts the results. The former is more realistic as a small-scale version of real-life military planning. To encourage such collaborative decision-making, none of the network topologies in Figure 3 lean towards the creation of such a leader – there is never just one person who is more highly connected or more central than the others.

In the case of the chain and tree networks in Figure 3, participants with more than one neighbour (four people for the chain, two for the tree) must copy and re-broadcast messages that they receive, so that the information can make its way across the network.

The restricted-topology communication between team members can be provided with email or other means, but the chat tool shown in the lower right of Figure 2 is perhaps the easiest mechanism. As well as enforcing a topology, this chat tool also writes a (possibly anonymised) log of all traffic, for later analysis.

Since only six team members are required, conducting experiments is expected to be easier than with ELICIT, which requires 17 participants (Ruddy 2007, Thunholm *et al*. 2009). However, at present, no experiments with the team version of the game have yet been performed.

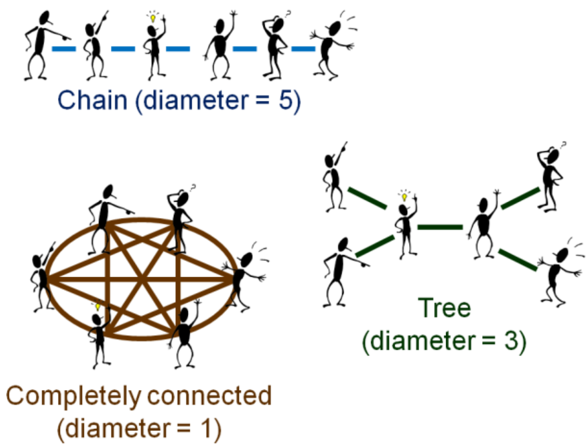

Figure 3. *Topologies for the team game.*

Agent Team Version

As an alternative to human experimentation, we have also developed a version of the team game with intelligent software agents representing people. These agents communicate by text chat. Agents send messages which include the descriptions specifying the problem (e.g. "Puma should, if possible, not be used in Greenland") as well as statements of an agent's level of preference as to where it wishes to deploy its platform:

- Agent likes Blueland for Puma
- Agent wants Blueland for Puma
- Agent really wants Blueland for Puma
- Agent needs Blueland for Puma
- Agent has locked in Blueland for Puma

To model human behaviour, explicit time delays are used – it takes an average of 2 seconds for an agent to send a message, and an average of 5 seconds to receive and decode one.

Agents decide to "lock in" a particular choice using a very simple algorithm: either when (on the basis of the data they have) no alternative is possible, or randomly (but with low probability). Random decisions are made with a low probability proportional to the square of the preference strength, so that the most constrained agent generally chooses first.

This very simple collaborative decision-making algorithm is not guaranteed to produce an optimal solution, but compares favourably with human performance. In almost all cases the agents produce a valid solution (assigning every platform exactly once and satisfying the mandatory conditions).

Hybrid Team Version

Because agents communicate by sending and decoding text chat messages, it is possible to combine intelligent agents with a human player, as long as the human players sends messages in a form which the agents can decode. A tool has been developed which allows a human player to collaborate with five intelligent agents (humans always take a position in the network having more than one neighbour). This tool does not require networked communication, and can be deployed as an applet. Some preliminary experiments have been conducted with this tool, and the three topologies in Figure 3. Results are shown in Figure 4.

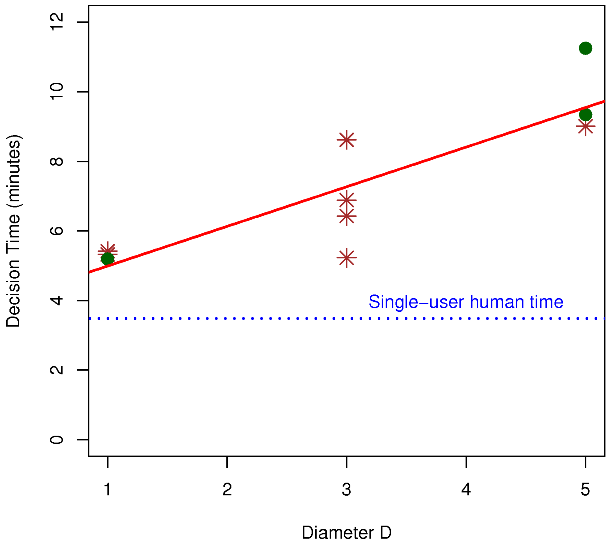

***Figure 4**. Decision times for hybrid teams. Green dots are optimal solutions, brown stars are valid but sub-optimal solutions. The red line shows the best linear fit:* ***time = 3.85 + 1.14 D**.*

The time taken to solve the problem ranged from 5.19 to 11.25 minutes, depending on the communication topology. The relationship between time (in minutes) and diameter D could be modelled with the linear relationship $T = 3.85 + 1.14\ D$ (red line in Figure 4), which was significant at the 0.002 level (with $R^2 = 75\%$). Adding a quadratic term did not significantly improve the fit.

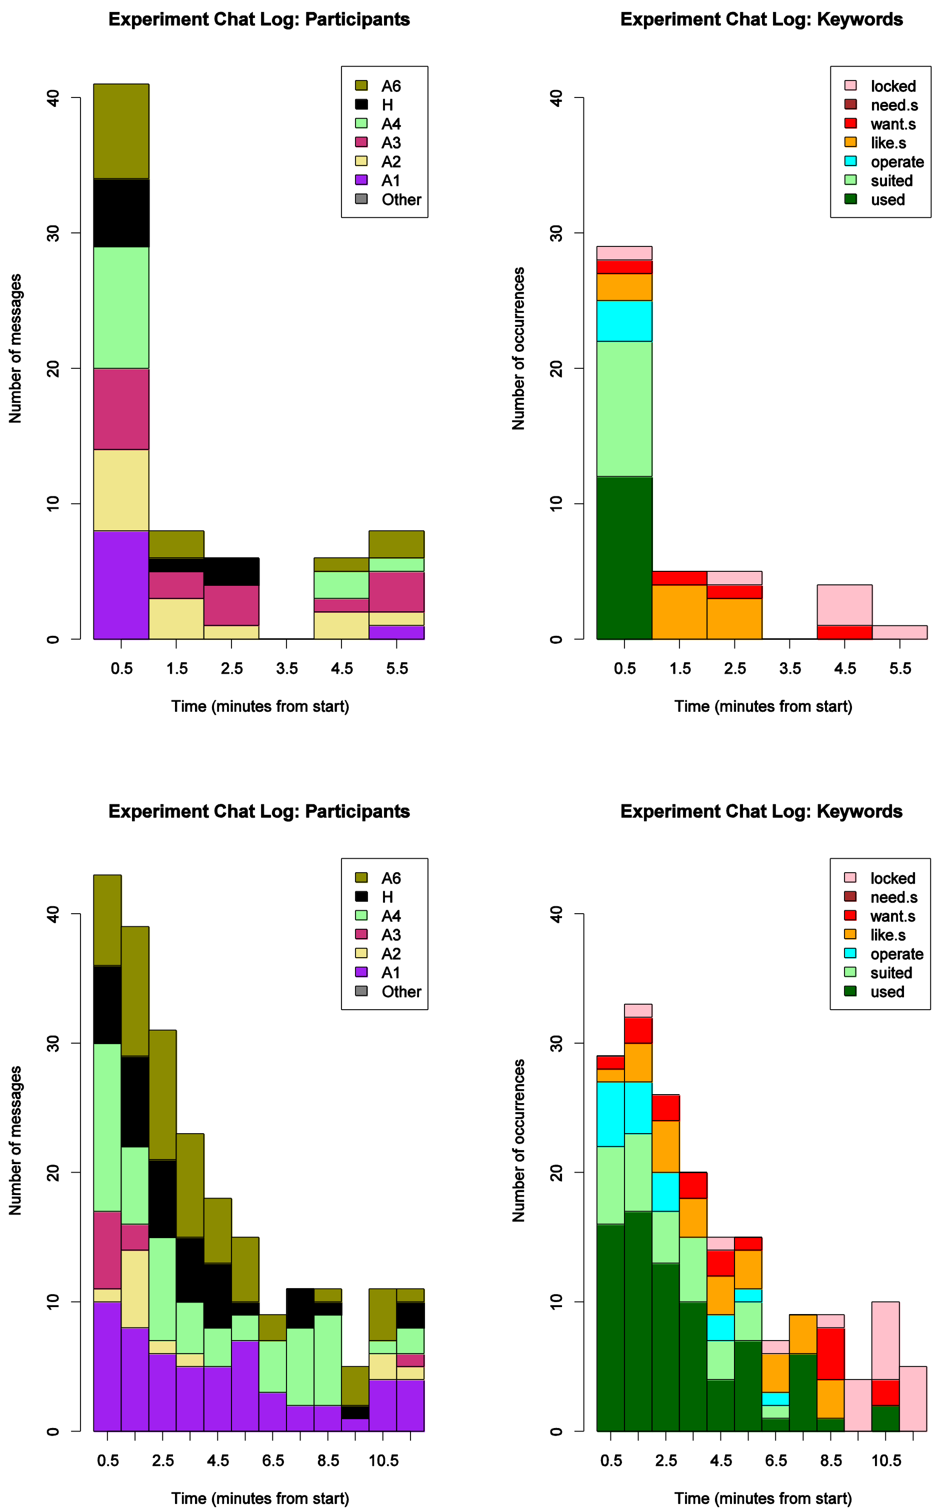

Figure 5*. Chat log analysis for the hybrid game, with a completely connected topology (top) and a chain topology (bottom). Participants A1–A4 and A6 are agents, while H is a human.*

Chat Log Analysis

In order to better understand the behaviour of human, agent, and hybrid teams, we have also developed a software tool suite for analyzing and visualising the chat message logs produced by the collaborative decision process. Figures 5 and 6 show example outputs of this analysis.

Rather than develop an integrated tool, such as that of Lo *et al*. (2010), which was unsuitable for our purposes, we combined chat log parsing software written in Java with statistical analysis and visualisation scripts written in R (Maindonald and Braun 2007). This permitted relatively sophisticated analysis with minimal coding effort. As with Ramachandran *et al*. (2010), we interpreted the chat log text by searching for appropriately chosen keywords. The initial facts (descriptions of appropriateness for pairings) can be recognised with the keywords "used," "suited," and "operate," while the expressions of preference can be recognised by the keywords "like/s," "want/s," and "need/s" (the software recognises both the suffixed and unsuffixed form). System-generated messages of the form "H has locked in Whiteland for Puma" can be recognised by the keyword "locked."

In order to handle misspellings by human participants, the chat log parsing software includes automatic spelling correction,[1] modified by two files of "don't correct *x*" and "always correct *x* to *y*" instructions. This functionality is not needed for the hybrid game, since the sole human participant is restricted to messages which the agents can understand.

The top row of Figure 5 shows analysis of the chat log for the hybrid game with a completely connected topology. All facts are distributed in the first minute (keywords "used," "suited," and "operate" in green and cyan in the top right histogram), and participants begin to express their preference (using the keywords "like/s" and "want/s" in orange and red). There is a pause in the fourth minute before the last four participants (agents A2, A3, A4, and A6) lock in their decisions (keyword "locked" in pink). This is an example of agents breaking a deadlock by eventually locking in their preference (after a random time which depends on preference strength). In this case a perfect result was obtained after 5.19 minutes.

The bottom row of Figure 5 shows a hybrid game with a chain topology. Facts continue to be re-broadcast until the eleventh minute, in parallel with communication about preferences. In this case a perfect result was obtained after 11.25 minutes. The histogram at the bottom left shows that most communication was by agents A1, A4, and A6, and the human H (agents A2 and A3 were on the end of the chain and had no need to rebroadcast information, and agent A3 also locked in its choice in the second minute).

The tool suite also analyzes variation in keyword use by different participants. Figure 6 shows an example corresponding to the top row of Figure 5. Three of the participants in this diagram have a bias towards using the name of the region which they will eventually lock in as a choice (the bias results from messages expressing a preference for that region). However, a χ^2 test indicates that the bias in Figure 6 is not large enough to be statistically significant. More interesting word usage biases can be expected with teams composed only of humans.

[1] Using the JaSpell package, written by Bruno Martins (Department of Informatics of the Faculty of Sciences of the University of Lisbon in Portugal), made available under the BSD license: http://jaspell.sourceforge.net/

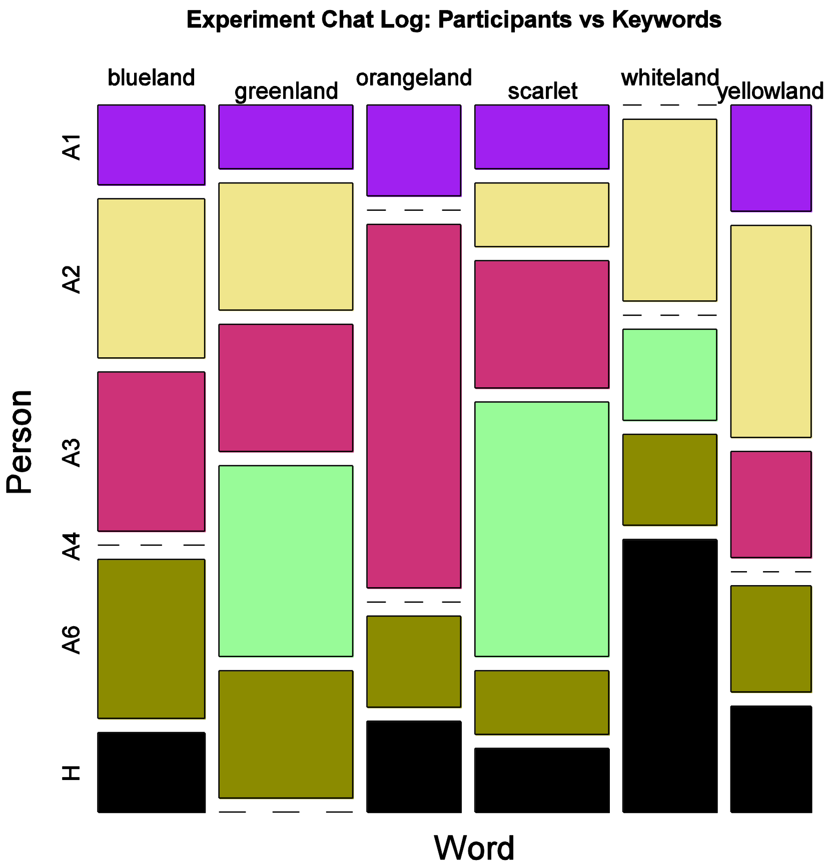

***Figure 6**. Participant/word matrix for the completely connected case (top row) of Figure 5. The large blocks A3/Orangeland (shown in magenta), A4/Scarlet (pale green), and H/Whiteland (black) correspond to these participants' final "locked in" choices.*

Discussion

In previous work (Dekker 2011a), we explored the relationship between distance within a network, and the time to collaboratively make decisions. In that work we used the average distance (average number of "hops") in a network, which we will denote here by *d*. For very small networks, such as we have used here, the diameter (maximum number of "hops") is a more appropriate measure.[2]

A very simple agent-based simulation model, which modelled planning by factorization of numbers, suggested a linear relationship between average distance and decision time (Dekker 2011a). A re-analysis (Dekker 2010b) of a colouring experiment by Kearns *et al*. (2006) also suggested a linear relationship, as did a re-analysis of experiments with the ELICIT game by Thunholm *et al*. (2009). Experiments with the Kuramoto Model (Strogatz 2000, Dorogovtsev *et al*. 2008, Dekker 2010a) suggested a nonlinear relationship although, as a model of self-synchronization, the Kuramoto Model may be overly simplistic.

The issue is an important one: for large military organisations, the diameter and average distance may be quite substantial, and the difference in efficiency between a linear and a

[2] For large networks, the diameter can be influenced by connections between a small fraction of nodes, and so the average distance provides a more reliable description of the network as a whole. However, for small networks, single nodes have a greater impact on overall performance, and should be taken into account. Also, average distance in very small networks fails to distinguish adequately between tree and chain structures (for our tree, average distance 1.933 *vs* diameter 3, and for our chain, 2.333 *vs* 5).

nonlinear effect may thus also be substantial. This in turn can significantly affect the benefits of structures with low average distances – structures like "edge organizations" (Alberts 2003; Dekker 2007b).

We have hypothesized (Dekker 2011a) that where information is repeatedly transmitted and re-transmitted, the progressive degradation of meaning (Pratt and Bennett 1989, Hone *et al*. 2007, Kashima and Yeung 2010, Baber *et al*. 2004) may result in a nonlinear effect of average distance on efficiency, and that the Kuramoto Model might be modelling such "attenuation" of information.

The present work was motivated by a desire to test this hypothesis. It has not directly done so: the shades of meaning in the seven-point scale of textual descriptions was intended to be a source of progressive degradation of meaning, but in the hybrid team experiments reported above, the intelligent software agents interpreted each sentence in a quite precise way, with no misunderstanding or degradation of meaning.

However, the hybrid team experiments do include a related factor, which is the workload associated with re-broadcasting messages. This workload was quite apparent on the right-hand side of Figure 5. Figure 7 shows the number of messages giving or re-broadcasting the initial descriptive facts (i.e. containing keywords "used," "suited," and "operate") for the hybrid team experiments. These numbers fitted the linear model 0.5 + 24.2 D (highly significant, at the 0.000003 level, with R^2 = 95%). This is in spite of the fact that the intelligent software agents were explicitly programmed not to re-broadcast messages which they had already broadcast in the past.

The total time taken by teams to reach a decision was partly a result of the fact that messages took longer to reach their destination in the tree (D = 3) and chain (D = 5) networks, and partly a result of the fact that both human and software-agent participants became overloaded by the need to read more messages and make decisions about them (including the decision about whether to re-broadcast them). A form of "attenuation" thus actually occurred in this experiment – not attenuation due to degradation of meaning, but attenuation due to becoming lost in the increasing pool of repeatedly rebroadcast messages.

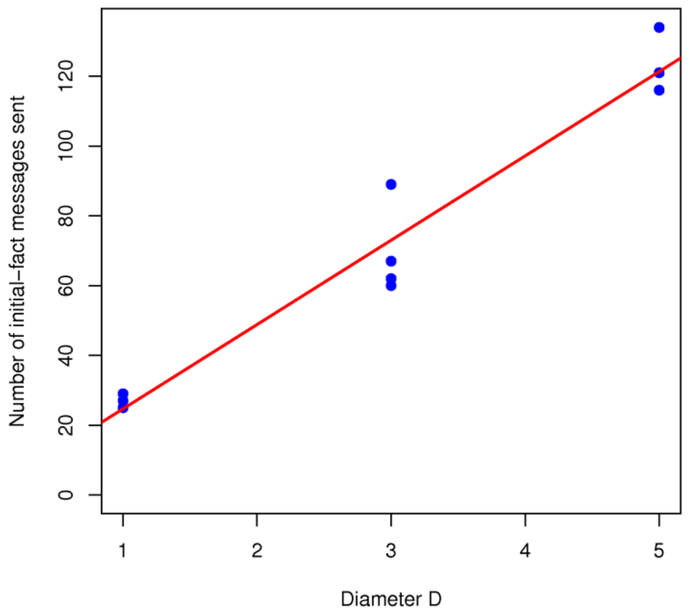

Figure 7. *Number of messages providing initial descriptive facts (i.e. containing keywords "used," "suited," and "operate") for the hybrid team experiments. The red line shows the best linear fit:* ***count = 0.5 + 24.2 D***

Looking at Figure 3, we can see that the re-broadcasting is in fact performed by participants having more than one neighbour (four people for the chain, and two for the tree), except for the completely connected network, where zero people need to rebroadcast. These numbers are exactly one less than the diameter in each case, which explains the linear relationship in Figure 7.

However, consider a similar game on a larger scale, where N people are connected in a tree structure with k subordinates per interior node. In this case the diameter D and the average distance d would both be proportional to log N. There would also be roughly N/k people having more than one neighbour and, provided information had to reach everybody in the network, the amount of re-broadcasting would also be proportional to N. As a function of the diameter or the average distance, the amount of re-broadcasting and the total decision time would therefore be exponential rather than linear. However, in large tree-structured organisations, it is likely that only a small fraction of information would be distributed across the whole organisation in this way, thus avoiding the exponential relationship.

Our hypothesis about linear relationships (Dekker 2011a) can therefore, on the basis of the experiments reported here, be answered with "it depends." While we have observed a linear relationship between diameter and time, it remains possible that very large tree structures may under some circumstances show an exponential relationship. Consequently, the speedup due to reducing average distance in the organisation will usually be linear, but might be better than linear in some circumstances. It is naturally difficult to explore the behaviour of such large networks experimentally, although it may be possible to find existing datasets which shed light on the question after re-analysis.

Nevertheless, the experiments we have conducted have shown the benefit of team games and agent systems for exploring C2 issues. In particular, we have shown that agents capable of sending and decoding text messages can collaborate with human beings. This permits larger-scale C2 experiments than with human beings alone. In future work, we plan to conduct further experiments with teams consisting only of people in order to confirm whether the behaviour observed here in hybrid teams still applies – that is, whether our agents have adequately modelled real human beings.

We also plan to extend the game presented here to be more realistic. In particular, we plan to allow the exchange of documents, not just one-line text messages. Experiments with such a game will allow exploration of the benefits of various information repositories and communication technologies in dealing with more complex problems and with real-world exercises. Communication via appropriate information repositories has the potential to reduce average distance and diameter as well as to eliminate time-consuming re-broadcasting of information. Our planned experiments will thus make possible recommendations for creating agile organisations capable of rapidly making collaborative decisions.

References

Alberts, David S. and Richard E. Hayes. 2003. *Power to the Edge*. Washington, DC: CCRP Publication Series. www.dodccrp.org/files/Alberts_Power.pdf

Alberts, David S. and Richard E. Hayes. 2006. *Understanding Command and Control*. Washington, DC: CCRP Publication Series. www.dodccrp.org/files/Alberts_UC2.pdf

Alberts, David S. and Richard E. Hayes. 2007. *Planning: Complex Endeavors*. Washington, DC: CCRP Publication Series. www.dodccrp.org/files/Alberts_Planning.pdf

Baber, Chris, Robert Houghton, and Mal Cowton. 2004. "WESTT: reconfigurable human factors model for network enabled capability." *RTO NMSG Symposium on Modelling and Simulation to Address NATO's New* *and Existing Military Requirements*. Koblenz, Germany, 7–8 October 2004. NATO document RTO-MP-MSG-028.

Barabási, Albert-Lászlό. 2002. *Linked: The New Science of Networks*. Cambridge, MA: Perseus Publishing.

Brandstein, Alfred, Gary Horne, and Henrik Friman. 2000. "Project Albert + ROLF 2010 = Red Orm." *Proceedings of the 5th International Command and Control Research and Technology Symposium*, Canberra, Australia, 24–26 October. www.dodccrp.org/events/5th_ICCRTS/papers/Track5/077.pdf

Christofides, Nicos. 1975. *Graph Theory: An Algorithmic Approach*. London: Academic Press.

Dekker, Anthony H. 2006a. "Centralisation *vs* Self-Synchronisation: An Agent-Based Investigation." *Proceedings of the 11th International Command and Control Research and Technology Symposium*. Cambridge, England. September 26–28. www.dodccrp.org/events/11th_ICCRTS/html/papers/030.pdf

Dekker, Anthony H. 2006b. "Revisiting 'SCUDHunt' and the Human Dimension of NCW." *Proceedings of the 11th International Command and Control Research and Technology Symposium*. Cambridge, England. September 26–28. www.dodccrp.org/events/11th_ICCRTS/html/papers/030.pdf

Dekker, Anthony H. 2007a. "Studying Organisational Topology with Simple Computational Models." *Journal of Artificial Societies and Social Simulation*, **10**(4): paper 6. http://jasss.soc.surrey.ac.uk/10/4/6.html

Dekker, Anthony H. 2007b. "Using Tree Rewiring to Study 'Edge' Organisations for C2." *Proceedings of SimTecT 2007*. Brisbane, Australia, 4–7 June. 83–88. Simulation Industry Association of Australia.

Dekker, Anthony H. 2007c. "Studying Multi-Agency Planning Using Team Games." *Proceedings of SimTecT 2007*. Brisbane, Australia, 4–7 June. 451–454. Simulation Industry Association of Australia.

Dekker, Anthony H. 2010a. "Average Distance as a Predictor of Synchronisability in Networks of Coupled Oscillators." 33rd Australasian Computer Science Conference (ACSC2010), *Conferences in Research and Practice in Information Technology* (CRPIT), **102**: 127–131. crpit.com/confpapers/CRPITV102Dekker.pdf

Dekker, Anthony H. 2010b. "Mimicking Human Problem-Solving with Agents: Exploring Model Calibration." *Proceedings of SimTecT 2010*, Brisbane, Australia, 31 May–3 June.

Dekker, Anthony H. 2011a. "Analyzing C2 Structures and Self-Synchronization with Simple Computational Models." *Proceedings of the 16th International Command and Control Research and Technology Symposium*. Quebec City, Canada, June 21–23. www.dodccrp.org/events/16th_iccrts_2011/papers/055.pdf

Dekker, Anthony H. 2011b. "Exploring Networked Team Behaviour Using Collaborative Games," *Defence Human Sciences Symposium 2011*, Melbourne, Australia, Nov 16–17.

Dorogovtsev S.N., A.V. Goltsev, and J.F.F. Mendes. 2008. "Critical phenomena in complex networks." *Reviews of Modern Physics*. **80**(4): 1275–1353. arxiv.org/abs/0705.0010

Hone, Geoffrey, Ian R. Whitworth and Andy Farmilo. 2007. "Assessing the transmission of Command Intent." *Proceedings of the 12th International Command and Control Research and Technology Symposium*. Newport, RI. June 19–21. www.dodccrp.org/events/12th_ICCRTS/Papers/038.pdf

Huber, R.K., P. Eggenhofer, J. Römer, S. Schäfer, and K. Titze. 2006. "Developing an analytical framework for cognitive and social experimentation in team decision-making and collaboration." *Proceedings of the 11th International Command and Control Research and Technology Symposium*, Cambridge, England. September 26–28. www.dodccrp.org/events/11th_ICCRTS/html/papers/118.pdf

Kashima, Yoshihisa and Victoria Wai-Lan Yeung. 2010. "Serial Reproduction: An Experimental Simulation of Cultural Dynamics." *Acta Psychologica Sinica*. **42**(1): 56–71. journal.psych.ac.cn/xuebao/qikan/manage/wenzhang/100107.pdf

Kearns, Michael, Siddharth Suri, and Nick Montfort. 2006. "An Experimental Study of the Coloring Problem on Human Subject Networks," *Science*, **313**(11): 824–827.

Lo, Edward, Andrew Au, Peter Hoek, and Lorenz Eberl. 2010. "Combining Contextual Data in the Analysis of Temporal Social Networks," *TTCP Human Sciences Symposium, Sydney, Australia*.

Maindonald, John and W. John Braun. 2007. *Data Analysis and Graphics Using R – An Example Based Approach*, 2nd ed, Cambridge University Press.

Perla, Peter P., Michael Markowitz, Albert Nofi, Christopher Weuve, Julia Loughran, and Marcy Stahl. 2000. *Gaming and Shared Situation Awareness*. Washington, DC: Center for Naval Analyses. www.cna.org/documents/D0002722.A2.pdf

Pratt, K.J. and S.G. Bennett. 1989. *Elements of Personnel Management*. 2nd ed. London: Van Nostrand Reinhold, p. 88.

Ramachandran, Sowmya, Randy Jensen, Todd Denning, Oscar Bascara, Tamitha Carpenter, and Shaun Sucillon. 2010. "Automated Chat Thread Analysis: Untangling the Web," *Interservice/Industry Training, Simulation, and Education Conference (I/ITSEC)*, Orlando, Florida. www.dtic.mil/cgi-bin/GetTRDoc?AD=ADA532774

Ruddy, Mary. 2007. "ELICIT – The Experimental Laboratory for Investigating Collaboration, Information-sharing and Trust." *Proceedings of the 12th International Command and Control Research and Technology Symposium*. Newport, RI. June 19–21. www.dodccrp.org/events/12th_ICCRTS/CD/html/papers/155.pdf

Strogatz, S.H. 2000. "From Kuramoto to Crawford: exploring the onset of synchronization in populations of coupled oscillators." *Physica D: Nonlinear Phenomena* **143**: 1–20.

Thunholm, Peter, Ee-Chong Ng, Mervyn Cheah, Kin-Yong Tan, Nency Chua, and Ching-Lian Chua. 2009. "Exploring Alternative Edge versus Hierarchy C2 Organizations using the ELICIT Platform with Configurable Chat System." *International C2 Journal*. **3**(2). www.dodccrp.org/files/IC2J_v3n2_04_Thunholm.pdf

Watts, Duncan J. 2003. *Six Degrees: The Science of a Connected Age*. London: William Heinemann.