20th International Congress on Modelling and Simulation, Adelaide, Australia, 1–6 December 2013
www.mssanz.org.au/modsim2013

Self-Synchronisation in C2 Networks

A.H. Dekker

Graduate School of Information Technology & Mathematical Sciences of the University of Ballarat, Ballarat, Victoria, Australia
Email: dekker@acm.org

Abstract: Self-synchronisation is a key aspect of modern Command and Control (C2), and occurs when networked military units achieve the degree of alignment of their activities needed to accomplish their mission by exchanging information with each other, rather than being given detailed orders from above. The basic principles of self-synchronisation can be studied using models, and here we examine the ability of the "boids" of Reynolds (1987) to self-synchronise their velocity vectors within four-dimensional space.

As in past work, the average distance D in the network was the best predictor of the time required for agents to self-synchronise. For the case $D \geq 3$, 87% of the variance in synchronisation time was predicted by the power law 16.4 $D^{2.23}$. For the case $D < 3$ (consisting of Erdős-Rényi random networks), the synchronisation time peaked for networks of average degree 12. This results from the fact that densely connected subnetworks facilitate local self-synchronisation, but hinder global local self-synchronisation. This was confirmed by experiments on three highly connected networks, which showed that additional links within clusters actually slowed down self-synchronisation.

These results indicate that reducing the average distance in military communication networks is generally beneficial. However, within joint, combined, and coalition forces, improving the internal communications of individual components may actually be detrimental in terms of self-synchronisation of the overall force.

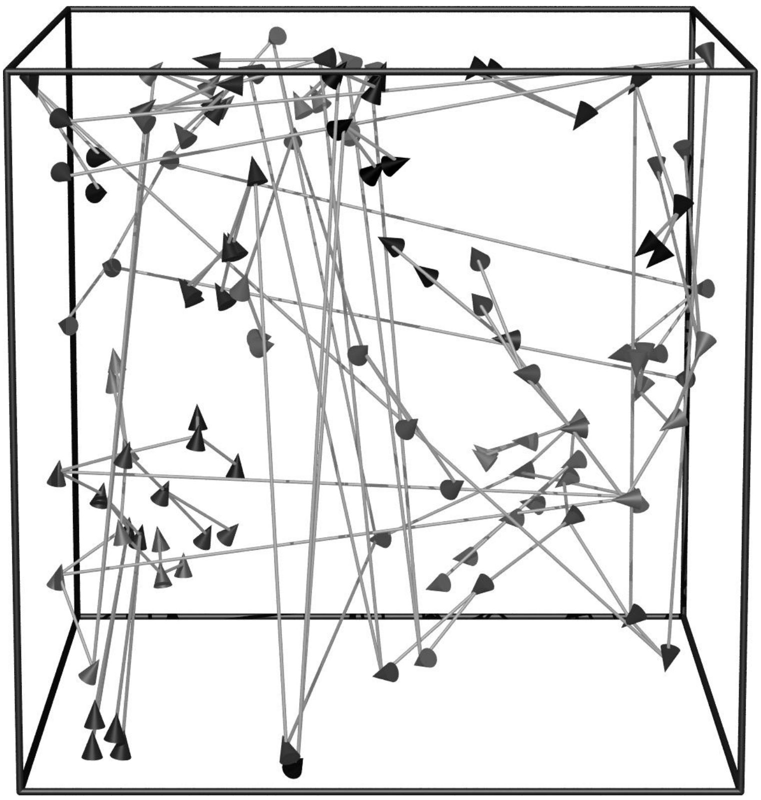

Figure (i). Early signs of self-synchronisation in a tree network. The direction together with the shade of grey of each cone indicates a four-dimensional velocity vector. Cones align with their neighbours in the network, but this local self-synchronisation has not yet produced global alignment.

Keywords: *Self-synchronisation, C2, networks, joint, coalition, boids*

1. INTRODUCTION

Self-synchronisation is a key aspect of modern Command and Control or C2 (Alberts and Hayes, 2003, 2006, 2007). It occurs when networked military units achieve the degree of alignment of their activities needed to accomplish their mission by exchanging information with each other, rather than being given detailed orders from above. Although self-synchronisation is complex, the basic principles of self-synchronisation can be studied using simple games with human participants, such as SCUDHunt (Perla *et al*., 2000; Dekker, 2006a), ELICIT (Ruddy, 2007; Thunholm *et al*., 2009), and CAFFEINE (Huber *et al*. 2006). Simple agent-based simulation models also offer important insights (Brandstein *et al*. 2000; Dekker, 2005, 2006b). Even simple mathematical models such as the Kuramoto model (Kuramoto, 1948; Strogatz, 2000; Dorogovtsev *et al*., 2008) can yield useful insights (Dekker, 2007, 2010; Kalloniatis, 2008). We are particularly interested in how self-synchronisation is facilitated or defeated by the topology of the network connecting the participants.

2. PAST WORK

In previous work, we identified the average distance D (the average number of "hops") in a network as a key network attribute determining the ability to self-synchronise. A very simple agent-based simulation model, which modelled planning by factorization of numbers, suggested a linear relationship between the average network distance and the time required to self-synchronise (Dekker, 2011). A re-analysis (Dekker, 2010) 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). This re-analysis found that the average decision time was $26.3 + 5.7D$ minutes. Experiments with the Kuramoto Model (Dekker, 2010) also indicated that the average distance was a key network attribute, although there the relationship was nonlinear, with the time required to self-synchronise being approximately proportional to D^4.

3. "FLOCKING" AND "BOIDS"

Small-scale experiments to explore self-synchronisation require a "toy" problem of some kind. In this work, we use a four-dimensional generalisation of the "boids" of Reynolds (1987). A somewhat similar model is used by Aoyagi and Namatame (2009). In our model, agents travel through an abstract space which is a four-dimensional hypertorus (we wish to use as many dimensions as possible to model the richness of military synchronization, but four dimensions are the most that can be easily visualised). Figure 1 shows three-dimensional projections of the four-dimensional space. Agents begin with random positions and velocity vectors, as in Figure 1(a), and are subject to four rules:

- **inertia**: agents tend to change their velocity vector only slowly;
- **alignment**: agents tend to align their velocity vector with the vectors of their neighbours;
- **attraction**: agents move towards their neighbours; and
- **repulsion**: agents move away from their neighbours if they are closer than a minimum distance (which we take to be 5% of the length of one dimension).

At double the ideal distance, the attractive effect has a magnitude of 1% of the maximum permissible velocity vector, and this increases quadratically with distance. At 80% of the ideal distance, the repulsive effect also has a magnitude of 1% of the maximum velocity vector, and this grows as the fourth power of distance. The alignment effect has a magnitude of 5% of the velocity vectors of an agent's neighbours. At each time step, an agent's velocity vector is combined with the alignment, repulsion, and attraction vectors.

If "neighbours" is interpreted as "agents within a given radius," agents will exhibit "flocking" behaviour as in Figure 1(b). However, with "neighbours" interpreted as neighbours in the network, as in Figure 1(d), agents will experience attraction from multiple directions, and may form loops of mutually attracting agents around the hypertorus. "Flocking" behaviour will therefore not generally occur. However, this is not a problem, since the alignment of activities in self-synchronisation is here being modelled by alignment of velocity vectors, not by alignment of positions. Figure 1(b) and Figure 1(d) are both examples where alignment of velocity vectors has been achieved.

As with the Kuramoto model (Dekker, 2007, 2010; Kalloniatis, 2008), this is a continuous model, in which two or more agents can synchronise by replacing their velocity vectors with the group average. In practice, if agents are synchronising their actions out of two alternatives – say, travelling *north* of a mountain or *south* of it – then the average choice (crossing the mountain itself) is unlikely to be the best one. However, methods exist for synchronising discrete alternatives (Mullender, 1993; Dekker, 2006b). Past work (Dekker, 2007, 2010, 2011, 2012) does not suggest great differences between the network dependence of discrete and continuous models.

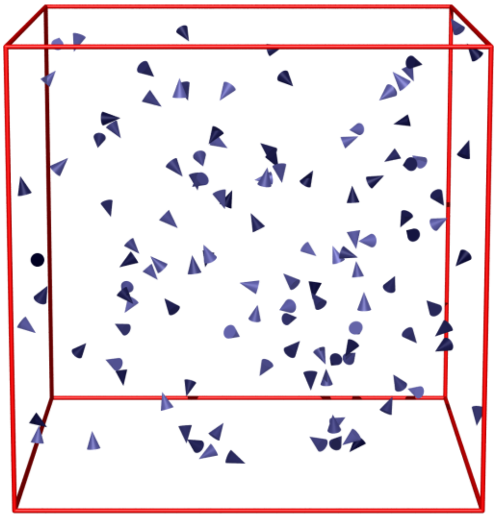

(a) Unsynchronised agents: all velocity vectors are random

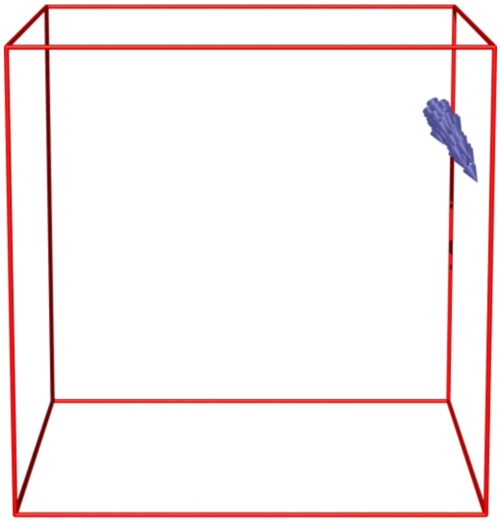

(b) Total synchronisation (in a complete network): all velocity vectors are aligned, and agents “flock.”

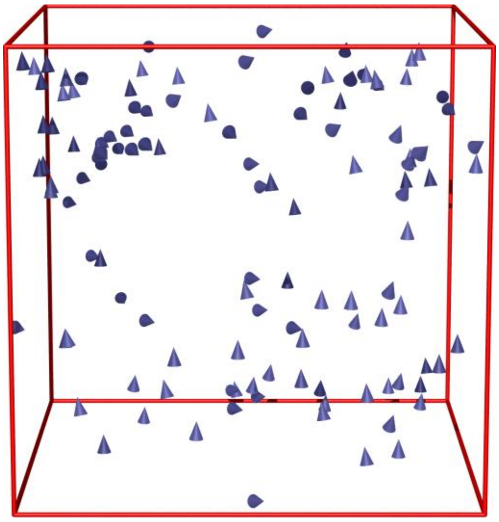

(c) Partial synchronisation: subsets of agents have their velocity vectors aligned, but without global alignment

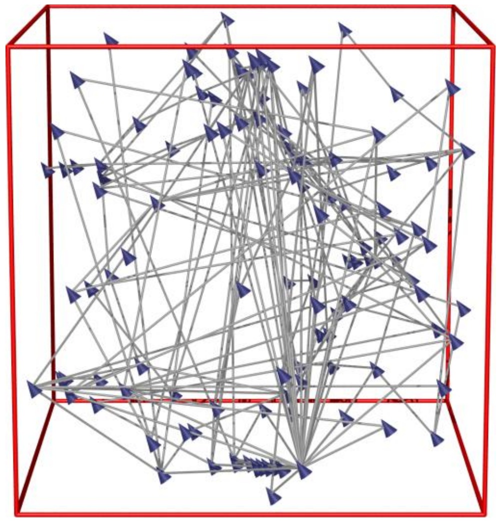

(d) Networked synchronisation: communication across network links result in all velocity vectors becoming aligned, although agents do not “flock.”

Figure 1. Full, partial, and absent self-synchronisation in our experimental setup. Each cone represents an agent, with the direction of the cone indicating the direction of the first three components of the velocity vector. The colour of the cone (light to dark) shows the value of the fourth component of the velocity vector. The position of the cones shows the first three components of the position vector. Because of hypertorus connectivity, opposite sides of the red cubes are considered as joined.

4. EXPERIMENT AND RESULTS

We performed a series of experiments on a mixed sample of 65 networks, each with 120 nodes. This allows sampling of a range of random and symmetrical networks comparable in size to the military networks we are intending to model (a size of 120 nodes also allows us to include specific regular networks, such as the largest Archimedean network). Our sample consisted of:

- the complete network on 120 nodes;
- 5 other regular networks: the 10×12 torus network (degree 4), the truncated icosidodecahedron network (degree 3), the truncated icosidodecahedron network with diagonal edges added (degree 4), and the two symmetric graphs with degree 3 and 120 nodes (Biggs, 1993; Conder and Dobcsányi, 2002);
- a tree in which all the nodes have degree 3 or degree 1 (grouped with the regular networks in Figure 2);

- 19 random (Erdős-Rényi) networks (Bollobás, 2001), with average degrees ranging from 6 to 60;
- 13 scale-free (preferential-attachment) networks (Albert and Barabási, 2002; Barabási, 2002), with average degrees ranging from 2 to 6.5; and
- 26 small-world networks, generated by applying the Watts rewiring process (Watts and Strogatz, 1998; Watts, 2003) to between 1% and 40% of the links in a 120-node antiprism.

For each of these networks, we measured the median time to reach a synchronisation threshold at which the mean Euclidean distance between velocity vectors was 4×10^{-4} times the maximum permissible magnitude of a velocity vector (the mean Euclidean distance between velocity vectors is thus our order parameter). This median was calculated over 1,001 runs with random initial velocity vectors and positions. The solid datapoints in Figure 2 show the results.

We considered the average network distance D, the average degree, the average connectivity (Dekker, 2007), and the clustering coefficient (Watts and Strogatz, 1998) as potential statistical predictors of the synchronisation time. As in previous work, the best predictor of the synchronisation time was the average distance, which predicted 77% of the variance (with an extremely significant p-value less than 2.2×10^{-16}). The second-best predictor was the average degree, which predicted 35% of the variance ($p = 2 \times 10^{-7}$), while the average connectivity predicted 28% ($p = 5.4 \times 10^{-6}$), and the clustering coefficient 16% ($p = 0.001$).

The average distance was particularly effective as a predictor for average distances $D \geq 3$. Here 87% of the variance was predicted by the power law $16.4\, D^{2.23}$. The solid line in Figure 2 shows this predictor.

For the Erdős-Rényi random networks with average distances $D < 3$, the behaviour was more interesting. Here the synchronisation time peaked for networks of average degree 12 (and average distance 2.17). This results from the fact that not all self-synchronisation is helpful. For the random networks with degrees around 12, densely connected subnetworks facilitate local self-synchronisation, but it is then difficult for these clusters to align with each other, since the nodes tend to align with their neighbours in the network. We have not observed this effect in past work.

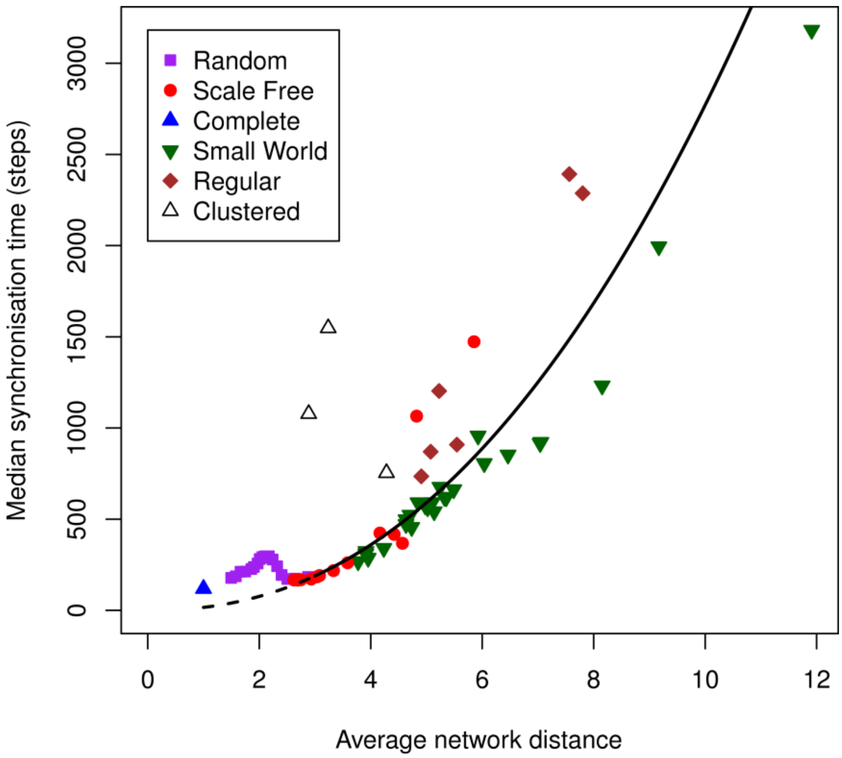

Figure 2. Median synchronisation times for the initial 65 different 120-node networks (solid points), plus the three highly clustered networks (open triangles). The solid line shows the power law predictor $16.4\, D^{2.23}$ for average distances $D \geq 3$. The dashed line is the extrapolation to $D < 3$. The standard errors of the median for each network are at most 27 steps, so that error bars are not visible.

The synchronisation-time peak for Erdős-Rényi random networks of average degree 12 (and average distance 2.17) corresponds to the point where the number of 4-cliques (completely connected subnetworks of 4 nodes) begins to significantly increase (there are an average of 8.34 such cliques at this point in the growth of a 120-node random network). Beyond this point, however, the improved overall connectivity begins to join these cliques together, thus facilitating self-synchronisation.

This effect occurs more dramatically in two highly clustered networks, formed by taking a hexagonal prism and an octahedron, and replacing each of the 6 nodes by a 20-node completely connected cluster. The two resulting 120-node networks have average degree 19.15 and 19.2 respectively, with average distance 3.24 and 2.89. Figure 3 shows an example of local self-synchronisation in the prism-based network. A third highly clustered network is based on the octahedron, but with only a random subset of links in each cluster, giving average degree 6.2 and average distance 4.28. Figure 4 shows the two octahedron-based networks.

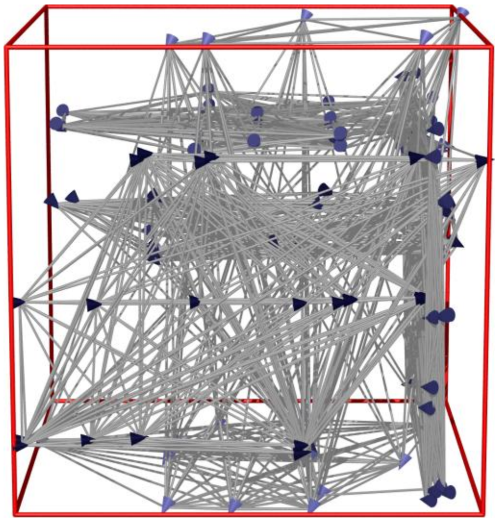

Figure 3. A highly clustered network, with six clusters of 20 nodes connected as a hexagonal prism. The state after 100 steps is shown. Synchronisation has occurred within clusters, but not yet between clusters.

The octahedron-based clustered network (median synchronisation time 1,078 steps) performs better than the prism-based network (1,548 steps), but the poorly-connected octahedron-based network performs better still (754 steps). These differences are statistically very significant ($p < 2\times10^{-16}$). A weaker form of this effect also holds for synchronisation in the Kuramoto model: experimentation reveals that the poorly-connected octahedron-based network synchronises slightly more slowly than the well-connected one (a median of 13,199 steps rather than 11,490, over 1,001 runs), but this is still four times faster than predicted by the relationship with average distance (Dekker, 2010).

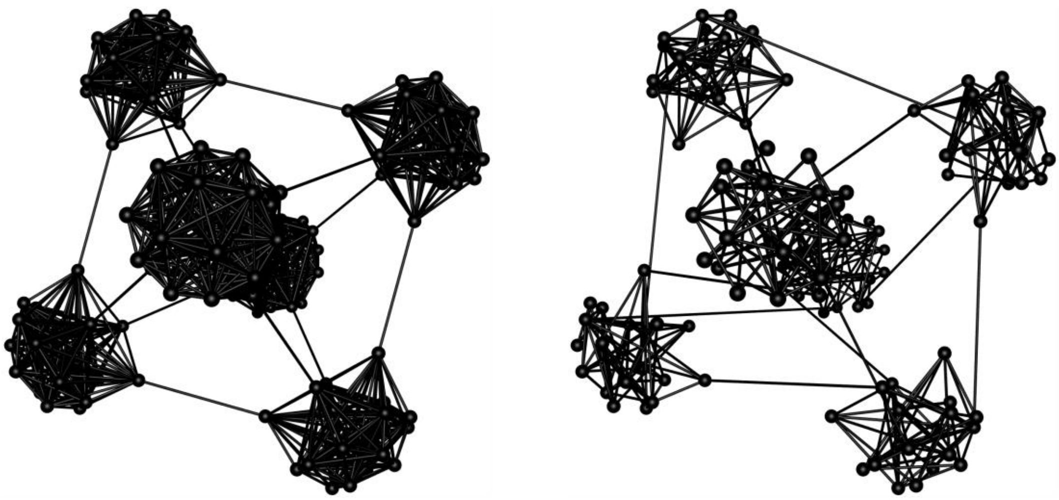

Figure 4. The well-connected (left) and less well-connected (right) octahedron-based networks. The average degrees are 19.2 and 6.2 respectively, and the average distances 2.89 and 4.28. Surprisingly, the less well-connected network on the right synchronises more rapidly (754 steps instead of 1,078).

Our results are consistent with the behaviour of Erdős-Rényi random networks and Scale Free networks in the Kuramoto model (Dorogovtsev *et al*., 2008; Arenas *et al*., 2008). There, self-synchronisation in Scale Free networks emerges from a central core, to which nodes progressively become aligned. For Erdős-Rényi random networks, in contrast, self-synchronisation emerges from multiple self-synchronised clusters.

In our experiments, if self-synchronisation emerges from multiple self-synchronised clusters, then global self-synchronisation is actually compromised if these clusters are internally well-connected. These results are similar to those of Brede (2008), who found that network cliques facilitated local synchronisation but not global synchronisation in the Kuramoto model.

5. DISCUSSION

Our results are consistent with past work (Dekker, 2010, 2011, 2012) in which the average network distance D was also the best predictor of synchronisation time, with synchronisation time being roughly proportional to D^{λ}, for some model-dependent constant λ. For many models, and for some studies with humans, we found $\lambda = 1$, while for the Kuramoto model, the best fit was $\lambda = 4$. For the experiments reported here, $\lambda = 2.23$ for $D \geq 3$. Since the constant λ does appear to be model-dependent, real-world data is required to determine what the value of λ would be for actual military C2 networks. Even in the absence of such data, however, we have the general principle that networks with low average distance will facilitate self-synchronisation.

The experiments reported here have also identified an important exception to this principle, however. If the network consists of clusters which are poorly connected with each other, then reducing the average distance by making the clusters internally better-connected can actually hinder self-synchronisation. Networks without clusters – e.g. symmetric (Biggs, 1993) and near-symmetric (Donetti *et al.*, 2005) networks – may synchronise well, although our two degree-3 symmetric networks were slower than the $16.4\,D^{2.23}$ predictor.

The slow synchronisation of networks with internally well-connected clusters has implications for joint, combined, and coalition forces. Improving the internal communications of individual components may actually be detrimental in terms of self-synchronisation of the overall force. Either priority must be given to network connectivity *between* components, or the overall force must be given a top-level plan to which all components agree, so that self-synchronisation is not required. For example, consider a multinational disaster relief force consisting of both military and non-military (e.g. NGO) units. A common strategy is to partition the affected region geographically, with each component providing relief within its assigned area. This requires very little self-synchronisation. However, if more integrated operations are required, then good network connectivity between components becomes essential.

REFERENCES

Albert, R. and Barabási, A.-L. (2002). Statistical mechanics of complex networks, *Reviews of Modern Physics*, **74**: 47–97.

Alberts, D.S. and Hayes, R.E. (2003). *Power to the Edge*, CCRP Publication Series, Washington, D.C., www.dodccrp.org/files/Alberts_Power.pdf

Alberts, D.S. and Hayes, R.E. (2006). *Understanding Command and Control*, CCRP Publication Series, Washington, D.C., www.dodccrp.org/files/Alberts_UC2.pdf

Alberts, D.S. and Hayes, R.E. (2007). *Planning: Complex Endeavors*, CCRP Publication Series, Washington, D.C., www.dodccrp.org/files/Alberts_Planning.pdf

Aoyagi, M. and Namatame, A. (2009). Synchronization in Mobile Agents and Effects of Network Topology, in Takao, T., Hajime, K., Shingo, T., and Hiroshi, D. (eds.), *Agent-Based Approaches in Economic and Social Complex Systems: V*, Springer-Verlag, Berlin.

Arenas, A., Díaz-Guilera, A., Kurths J., Moreno, Y., and Zhou, C. (2008). Synchronization in complex networks, *Physics Reports*, **469** (3): 93–153.

Barabási, A.-L. (2002). *Linked: The New Science of Networks*, Perseus Publishing, Cambridge, MA.

Biggs, N. (1993). *Algebraic Graph Theory*, 2nd edition, Cambridge University Press.

Bollobás, B. (2001). *Random Graphs*, 2nd edition, Cambridge University Press, Cambridge, UK.

Brandstein, A., Horne, G., and Friman, H. (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

Brede, M. (2008). Locals vs. global synchronization in networks of non-identical Kuramoto oscillators, *The European Physical Journal B*, **62** (1): 87–94.

Brehmer, B. (2009). A New Paradigm for Dynamical Modelling of Networked C2 Processes, *Proceedings of the 14th International Command and Control Research and Technology Symposium*, Washington, D.C., June 15–17, www.dodccrp.org/events/14th_iccrts_2009/papers/041.pdf

Conder, M., and Dobcsányi, P. (2002). Trivalent symmetric graphs up to 768 vertices, *J. Combin. Math. Combin. Comput*., **40**: 41–63.

Dekker, A.H. (2005). Network Topology and Military Performance, *MODSIM 2005 International Congress on Modelling and Simulation*, Melbourne, Australia, 12–15 December, Zerger, A. and Argent, R.M., eds, pp 2174–2180, www.mssanz.org.au/modsim05/papers/dekker.pdf

Dekker, A.H. (2006a). 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/029.pdf

Dekker, A.H. (2006b). 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, A.H. (2007). Studying Organisational Topology with Simple Computational Models, *Journal of Artificial Societies and Social Simulation*, **10** (4), article 6, jasss.soc.surrey.ac.uk/10/4/6.html

Dekker, A.H. (2010). Average Distance as a Predictor of Synchronisability in Networks of Coupled Oscillators, *Proceedings of ACSC 2010, the 33rd Australasian Computer Science Conference*, January 18–21, Brisbane (*Conferences in Research and Practice in Information Technology*, Vol. 102), 127–131, crpit.com/confpapers/CRPITV102Dekker.pdf

Dekker, A.H. (2011). 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, A.H. (2012). Analyzing Team C2 Behaviour using Games and Agents, *Proceedings of the 17th International Command and Control Research and Technology Symposium*, Fairfax, Virginia, USA, June 19–21, www.dodccrp.org/events/17th_iccrts_2012/post_conference/papers/006.pdf

Donetti, L., Hurtado, P., and Muñoz, M. (2005). Entangled Networks, Synchronization, and Optimal Network Topology, *Phys. Rev. Lett.*, **95**, 188701.

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

Huber, R.K., Eggenhofer, P., Römer, J., Schäfer, S., and Titze, K. (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

Kalloniatis, A. (2008). A New Paradigm for Dynamical Modelling of Networked C2 Processes, *Proceedings of the 13th International Command and Control Research and Technology Symposium*, Seattle, WA, June 17–19, www.dodccrp.org/events/13th_iccrts_2008/CD/html/papers/180.pdf

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

Kuramoto, Y. (1948). *Chemical Oscillations, Waves, and Turbulence*, Springer, Berlin.

Mullender, S. (1993). *Distributed Systems*, 2nd edition, Addison-Wesley.

Perla, P., Markowitz, M., Nofi, A., Weuve, C., Loughran, J., and Stahl, M. (2000). *Gaming and Shared Situation Awareness*, Center for Naval Analyses, Washington, www.cna.org/documents/D0002722.A2.pdf

Reynolds, C.W. (1987). Flocks, Herds, and Schools: A Distributed Behavioral Model, *Computer Graphics*, **21** (4): 25–34.

Ruddy, M. (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, P., Ng, E.-C., Cheah, M., Tan, K.-Y., Chua, N., and Chua, C.-L. (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, D.J. and Strogatz, S.H. (1998). Collective dynamics of 'small world' networks, *Nature*, **393**: 440–442.

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