

# The Cognitive Virtues of Dynamic Networks

Trung Dong Huynh<sup>\*†</sup>, Paul Smart<sup>\*</sup>, Dave Braines<sup>§</sup>, Nigel Shadbolt<sup>\*</sup>, and Katia Sycara<sup>‡</sup>

<sup>\*</sup>School of Electronics and Computer Science, University of Southampton, Southampton, SO17 1BJ, UK

<sup>§</sup>Emerging Technology Services, IBM United Kingdom Ltd, Hursley Park, Winchester, Hampshire, SO21 2JN, UK

<sup>‡</sup>Robotics Institute, Carnegie Mellon University, 5000 Forbes Ave., Pittsburgh, 15213-3890, USA

<sup>†</sup>Corresponding author: tdh@ecs.soton.ac.uk

**Abstract**—For the most part, studies in the network science literature tend to focus on networks whose functional connectivity is largely invariant with respect to some episode of collective information processing. In the real world, however, networks with highly dynamic functional topologies tend to be the norm. In order to improve our understanding of the effect of dynamic networks on collective cognitive processing, we explored the problem-solving abilities of synthetic agents in dynamic networks, where the links between agents were progressively added throughout the problem-solving process. The results support the conclusion that (at least in some task contexts) dynamic networks contribute to a better profile of problem-solving performance compared to static networks (whose topologies are fixed throughout the course of information processing). Furthermore, the results suggest that constructive networks (like those used in the present study) strike a productive balance between autonomy and social influence. When agents are allowed to operate independently at the beginning of a problem-solving process, and then later allowed to communicate, the result is often a better profile of collective performance than if extensive communication had been permitted from the very outset of the problem-solving process. These results are relevant, we suggest, to a range of phenomena, such as groupthink, the common knowledge effect and production blocking, all of which have been observed in group problem-solving contexts.

## I. INTRODUCTION

Many real-world problems require the cooperative effort of multiple individuals, and the nature of the communication that takes place between such individuals is often key to understanding the collective successes (and failures) of such groups. A better understanding of how network structures affect the performance of collective problem-solving can lead to valuable insights as to how a network should be engineered to address a particular type of problem. A number of previous studies have investigated this issue, but most of these studies focused solely on networks with static topologies (see [1] and [2], for instance). In the real world, however, the structure of many networks is seldom fixed; rather, network structure tends to change, perhaps as a response to network activity, or as a response to natural processes of growth or decay [3]. In most real-world contexts, agents will (sometimes repeatedly) engage and then disengage from a communication network, or they may change their profile of connectivity as new social ties are formed and old ones wither away. For example, in situations involving mobile ad hoc networks (MANETs), many nodes may be expected to have only occasional or intermittent connectivity, and this may impede the effective rate of information

spread between network nodes. Even in situations where the structural topology of a network seems largely static, this does not mean that the ‘functional’ or ‘effective’ topology of the network is not wildly various. For instance, in the case of neural information processing, what seems to be important is not so much the relatively static hodological profile of specific neural circuits, so much as the dynamic patterns of neural activity which such circuits make possible. The distribution of various neurotransmitters and neuromodulators throughout such circuits seems to enable, on occasion, the effective ‘rewiring’ or ‘reconfiguration’ of the circuits in response to specific information processing challenges [4], [5]. Thus, even in situations where we encounter networks with largely static topologies, this does not mean that the ‘effective’ structure (the structure that realises the information processing capabilities of the focal network) is not, in some sense, dynamic.

In order to investigate the effect of dynamic networks on collective cognitive processing, we explored the problem-solving abilities of synthetic agents in a series of constructive network simulations (simulations in which the links between agents were progressively added throughout the problem-solving process). The results suggest that (at least in some task contexts) dynamic networks are able to outperform their more static network counterparts. Furthermore, the results suggest that constructive networks strike a productive balance between autonomy and social influence. When agents are allowed to operate independently at the beginning of a problem-solving process, and then later allowed to communicate, the result is often a better profile of collective performance than if extensive communication had been permitted from the very outset of the problem-solving process.

The remaining of this paper is structured as follows. Section II provides an overview of related work, Section III describes the experimental procedure, and Section IV presents the experimental results. Finally, Section V summarises the key findings of our work and outlines directions for future work.

## II. RELATED WORK

One study which attempted to examine collective cognitive processing in network contexts was Mason et al [1]. Mason et al investigated the ability of groups of networked human subjects to collectively explore a problem space and find optimal solutions within that space. The subjects had to guess a number




Figure 1. Examples of network structures used by Mason et al [1]

between 0 and 100, and they were awarded points based on a score associated with that number. A continuous ‘fitness’ function determined the score associated with each number, but this function was not made available to subjects. Instead, on any given round of the experiment, subjects had to choose a number based on the feedback they received from both their own guesses and the guesses of their immediate neighbours. Collective problem-solving performance was assessed by calculating the average score of subjects on each round of the experiment. Thus, in order to be successful, subjects had to explore the structure of the fitness landscape, as provided by the fitness function, and then rapidly converge on the most optimal solution (i.e. number) available. The cognitive process in question is thus a form of collective problem-solving in which the search efforts of multiple individuals are pooled to create a measure of collective cognitive success (i.e. the ability of the group to find an optimal solution).

In order to assess the effect of network structure on problem-solving performance, Mason et al organised the human subjects into communication networks with different structural topologies (see Figure 1). In the totally connected network condition, each participant could see the solutions proposed by all other participants as well as the fitness score associated with the solutions; in other network conditions, each participant could only see the solutions and scores of a more limited number of individuals, namely those individuals to which they were directly connected. What Mason et al found was that in conditions where there was a single local optima (i.e. a fitness landscape with a single peak), the network structures that supported the most rapid dissemination of information were the most successful. Thus, when the fitness landscape had a single peak, subjects tended to converge more quickly on the global maximum in the totally connected, small world and random network topology conditions. Things were

very different, however, when the fitness landscape had a more rugged, multi-modal structure; i.e. when there were multiple local optima (and one global optima). In this case, we might expect rapid information dissemination to once again result in rapid convergence on a particular solution, but whether this solution is the best one available depends largely on how lucky the participants were with their initial guesses. If the participants refrain from converging too quickly on a particular solution, they might discover a more globally-optimal solution, and it is for this reason that we might expect network structures that limit the rate of information dissemination to benefit collective problem-solving performance. Mason et al found that when a multi-peaked solution landscape was used, participants found the global solution most quickly in the small-world network condition. The topology of the small-world network seemed to provide just the right amount of social influence for optimal performance—it supported a certain amount of independent exploration, but did not inhibit rapid convergence on optimal solutions.

Results similar to Mason et al [1] have also been reported by Lazer and Friedman [2], this time using experiments involving synthetic agents. As with Mason et al, Lazer and Friedman examined the effect of network structure on collective problem-solving performance using a search task; however, in this case, the search task used by Lazer and Friedman [2] was based on the search for optimal design solutions using a variant of the NK model paradigm [6]. An NK model is essentially a means of generating evolutionary fitness landscapes with different degrees of ‘ruggedness’. By using a particular type of NK model ( $N = 20$  and  $K = 5$ ), Lazer and Friedman were able to set up a problem-solving task in which synthetic agents had to explore a moderately rugged fitness landscape (comprising about 100 local optima) and discover optimal solutions within that landscape. In order to explore the fitness landscape, the agents generated variants of a 20-bit design solution. At the beginning of each simulation, agents were randomly assigned to locations on a fitness landscape, and at each round of the simulation, they could explore the local terrain of the landscape by randomly switching one of the binary variables associated with their current solution. This modification to a single binary variable provided information about the relative fitness of neighbouring solutions and thus indicated whether fitter alternatives were available. In addition to receiving information about the fitness of their own solutions, agents also received information about the fitness of solutions discovered by their network neighbours (i.e. the agents to whom they were directly connected). If a neighbouring agent proposed a solution that was fitter than either the agent’s current solution or their modified solution, then the agent adopted the solution of its neighbour.

What Lazer and Friedman discovered was an apparent trade-off between what they called exploration and exploitation. Exploration here is the tendency of agents to explore the solution space independently of other agents, and exploitation is the tendency of agents to adopt the solutions proposed by other agents. The result of this trade-off is a profile of topology-

dependent performance in which networks with low average path lengths (e.g. totally connected networks) yield better performance in the short-term, in comparison to networks with higher average path lengths (e.g. linear networks), but worse performance in the longer term, again compared to networks with higher average path lengths.

The results of studies exploring the effect of network structure on collective search tasks thus point to a common conclusion: different types of network topology can affect the rate at which information propagates within a problem-solving community, and this can compromise a group’s ability to discover globally-optimal, long-term solutions. When the network topology supports rapid rates of information transfer, individuals may be inclined to settle on sub-optimal solutions on the basis of initial shared information.

Studies that attempt to investigate the effects of information flow and influence in social networks typically do so by using networks with fixed structural topologies. However, as previously discussed in Section I, there is no reason, to assume that such networks necessarily represent of the kind of networks typically encountered in cases of real-world cognitive processing. As such, in the remaining of the paper, we report the results of our study on the contribution of dynamic networks (networks whose structural and functional topology changes throughout the course of cognitive processing) to collective cognition. The specific aim of this study was to investigate the impact of constructive changes to a network (as realised by the progressive addition of links) on collective problem-solving performance.

### III. METHOD

In our experiments, a group of agents were tasked with the exploration of the problem space and the discovery of optimal solutions. The study used the same problem-solving paradigm as that used by Lazer and Friedman [2], and the NK problem spaces were of the same complexity; i.e. all simulations used NK spaces with parameters of  $N = 20$  and  $K = 5$ . These parameters give fitness landscapes that are moderately rugged, with a few hundred local optima and high correlations between proximate solutions. As Lazer and Friedman comment, this kind of problem space probably best captures “the essence of most interesting problems that individuals and organisations in the real world face—rugged, but not chaotic” [2, pg. 674]. In each simulation, a group of agents (population = 100) collaboratively explored a given NK problem space to discover the best solution for that space. At the beginning of the simulation, an initial solution was assigned to every agent, and this served as a starting point for its subsequent exploration. Every solution (represented by an  $N$ -digit bit-string) had a fitness value as determined by the selected NK space. During the simulation, each agent continuously looked for a better solution than the one it was currently associated with; i.e. one with a higher fitness value. In this manner, each agent sought out progressively fitter design solutions at each step in the problem-solving process.




Figure 2. A visualisation of network growth and performance over time. The colours indicate the fitness level of the agents’ solutions, with red being the highest fitness score.

There are two ways an agent can discover a better solution in this task. Firstly, it can search “up-hill” by trying other solutions that are similar to its own. Given the complexity of the problem space, we assumed that an agent had a very limited and local view of the space, and thus, it is only able to evaluate solutions that are similar to the one it is currently associated with. For this reason, an agent was only allowed to alter a single random digit at each step of the problem-solving process in order to determine whether a fitter solution could be found.

The second way in which an agent can discover a fitter solution is by comparing its solution with those associated with its network neighbours. If an agent is connected to other agents, it can compare its solution with that of its neighbours, and then copy its neighbour’s solution if that solution proves fitter than its own.

In each round of the simulation, the decision process of an agent  $X$  is as follows:

- 1)  $X$  first looks for the most successful immediate neighbour in the network (if it has any). If the solution of the neighbour is found to be better than its own solution  $S$ , then  $X$  adopts the solution of its neighbour.
- 2) If none of the solutions of  $X$ ’s neighbours are better than  $S$ , then  $X$  modifies  $S$  (by randomly selecting and flipping one of the 20 binary values comprising  $S$ ) to generate a new solution ( $S'$ ). If  $S'$  is better than  $S$ , then  $X$  adopts  $S'$  as its current solution; otherwise,  $S$  is retained as  $X$ ’s current solution.

In order to explore the effect of dynamic networks on collective problem-solving performance, we introduced changes

NGR	Meaning
0.1	1 link added every 10 cycles
0.2	1 link added every 5 cycles
0.5	1 link added every 2 cycles
1.0	1 link added every cycle
2.0	2 links added every cycle
5.0	5 links added every cycle
10.0	10 links added every cycle

Table I  
NETWORK GROWTH RATE VARIABLE

to the network structure throughout the course of the simulation. At the beginning of each simulation, none of the agents were connected (i.e. there were no links between any of the agents). During the course of the simulation, links were added at random in order to progressively connect the agents together into a single network component (see Figure 2 for an illustration). This constructive process was controlled by two independent variables: Network Growth Rate (NGR), which represents the rate at which communication links between the agents were added to the network, and Network Growth Delay Period (NGDP), which represents the period that must elapse before the first link is added to the network. In total, there were 7 levels of the NGR variable (see Table I), resulting in networks with different rates of growth, from the very slow (1 link per 10 cycles) to the very fast (10 links per cycles). The NGDP variable controls the duration of an initial period in which all the agents explore the problem space independently without any collaboration (since there are no links between any of the agents). In our experiments, there are 6 levels of the NGDP variable: 0, 10, 20, 30, 40, 50 cycles. A combination of NGR and NGDP results in a two-way factorial design consisting of  $(7 \times 6)$  42 experimental conditions. Before running the simulations, 1,000 different NK spaces were generated using the aforementioned parameters. Within each experimental condition, the 100 agents were tested against the same set of 1,000 NK spaces (i.e. the same set of 1,000 NK spaces were used for each condition), resulting in a total of 1,000 simulations for each condition (i.e. a total of  $1,000 \times 42 = 42,000$  simulations). A simulation concluded when no better solution could be found by the agents after 20 processing cycles (this is the time it takes each agent to explore all neighbouring design solutions<sup>1</sup>).

#### IV. RESULTS

##### A. Static versus Dynamic Networks

In this experiment, we look at how the agents performed in conditions where the network emerged dynamically (with new connections added) as opposed to a control condition where the network topology was static. Figure 3 presents the results obtained from simulations where the agent networks grew at various rates (as shown in Table I). In this section, we limit our

<sup>1</sup>Due to the random bit selection in the exploration process, there is a small chance that the agents could miss checking a particular bit that would improve their current solution. However, given the population of 100 agents, over 20 cycles this chance is very small ( $2.8 \times 10^{-45}$ ) and, thus, negligible.




Figure 3. Average performance over 1,000 simulations in networks with no initial growth delay period.

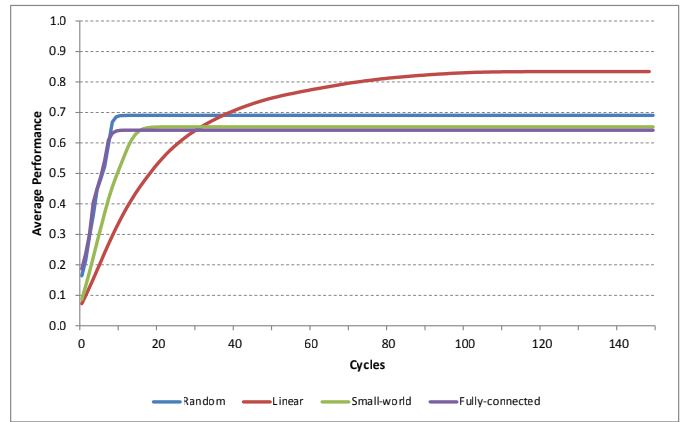


Figure 4. Average performance over 1,000 simulations in static networks.

attention to those conditions in which the initial delay period was zero (i.e. links were added on the very first processing step of the simulation). Figure 3 shows the average performance of the agents over the course of successive processing cycles for different levels of the NGR variable. The average performance is calculated as the average score associated with the solutions adopted by all agents, averaged across all the agents within a particular cycle (100 data points), and averaged across all the simulations within a particular experimental condition (1,000 data points). What Figure 3 shows is that as the network growth rate increases (i.e. as the rate at which links are added to the network increases), the latency to reach a stable local optimum solution decreases. However, as the rate at which links are added is increased, the final performance of the agents (the quality of the final solution) also becomes more variable, with performance negatively affected in conditions involving high growth rates.

In order to compare the performance of dynamic networks with that seen in static networks, we ran an additional series of simulations involving networks with the kind of static structures tested by Lazer and Friedman [2]. In particular, we used networks with linear, fully-connected, random and small-world topologies. In these simulations, a population




Figure 5. Comparisons of the final performance of networks having different growth rates and static networks (red bars). The adjacent blue bars show the number of cycles it took to achieve the performance indicated by the red bars.

of 100 agents was gradually connected to each other using an algorithm that gave rise to one of the aforementioned structures<sup>2</sup>. The performance of agents within these networks was then assessed by running 1,000 simulations using the aforementioned NK spaces. The results of this manipulation are shown in Figure 4. As can be seen from the figure, in all the network structures (with the exception of linear networks) the agents settle on a stable local optimum solution very quickly—after less than 20 cycles. This contrasts with the results seen in dynamic networks where it typically took much longer for the dynamic networks to settle on a common optimum solution. The most notable result from this manipulation concerns the final performance of the agents (plotted in Figure 5). In the fixed network conditions, the final performance of the agents reached a value considerably below that seen with dynamic networks. Even with the best performing static network structure, the linear network, it was still outperformed by dynamic networks with NGRs less than or equal to 2.0 (i.e. two links added per cycle or slower). In fact, the time it took the linear network to reach its stable final performance is twice that of the closest-performing dynamic network (i.e. NGR = 2.0). These observations were further confirmed as statistically significant by a Kruskal–Wallis one-way analysis of variances (ANOVA) using an alpha criterion of 0.05. We therefore conclude that problem-solving performance is enhanced in networks with dynamic, incrementally generated topologies, relative to networks with fixed, static topologies.

Accounting for this pattern of results requires us to think about the effect of time-variant changes in network architecture on the opportunities agents have for independent exploration of the solution space. Thus, given what we know about the way in which agents are influenced by the superior solutions of connected neighbours, it becomes apparent that the more time agents have to independently explore the solution space, the more likely the community is, as a whole, to discover the global optimum. When agents are connected together in a

fully-connected network, the rate of information dissemination is at its highest, and, as a result, the agents are inclined to prematurely settle on a sub-optimal solution (as found by Lazer and Friedman [2], they settle on the most optimal solution found on the first processing cycle). This explains why, in the case of dynamic networks, agents are more likely to quickly settle on a stable solution in the high growth rate conditions: the presence of more and more links across successive processing cycles progressively increases the rate at which information is transmitted between the agents. It also explains why, in general, we see a better profile of performance in dynamic networks than we do in static networks. In dynamic networks, the extent of inter-agent influence is limited; every agent can only influence those agents to which they are directly or indirectly connected, and, at least initially, all agents begin with minimal influence (i.e. no two agents are connected). This poverty of influence, relative to the situation with fully-connected networks, means that each agent has time to undertake a local exploration of the solution space before reporting the results of this local search to all the other agents (when the fully-connected network eventually emerges). Incrementally constructive networks, therefore, strike a productive balance between autonomy and influence; they give each agent the freedom to search for locally-optimal solutions without sacrificing the (eventual) benefits of collective search.

### B. Individual Autonomy versus Social Influence

Following on from the discussion in the previous experiment, we predicted that the presence of a growth delay period at the beginning of a simulation could have a positive impact on collective problem-solving performance. This is because we expected higher initial delay periods to provide agents with more time to independently explore the solution space. In this section, we focus on experimental conditions where initial delay periods were introduced. Again, as in the previous section, the average performance of the agents over time is plotted in each experimental condition (see Figure 6). The results show that the agents performed worse in conditions with high growth rates and shorter delay intervals, and the effect of

<sup>2</sup>We reused the source code provided by Lazer and Friedman [2] to generate the required network structures.

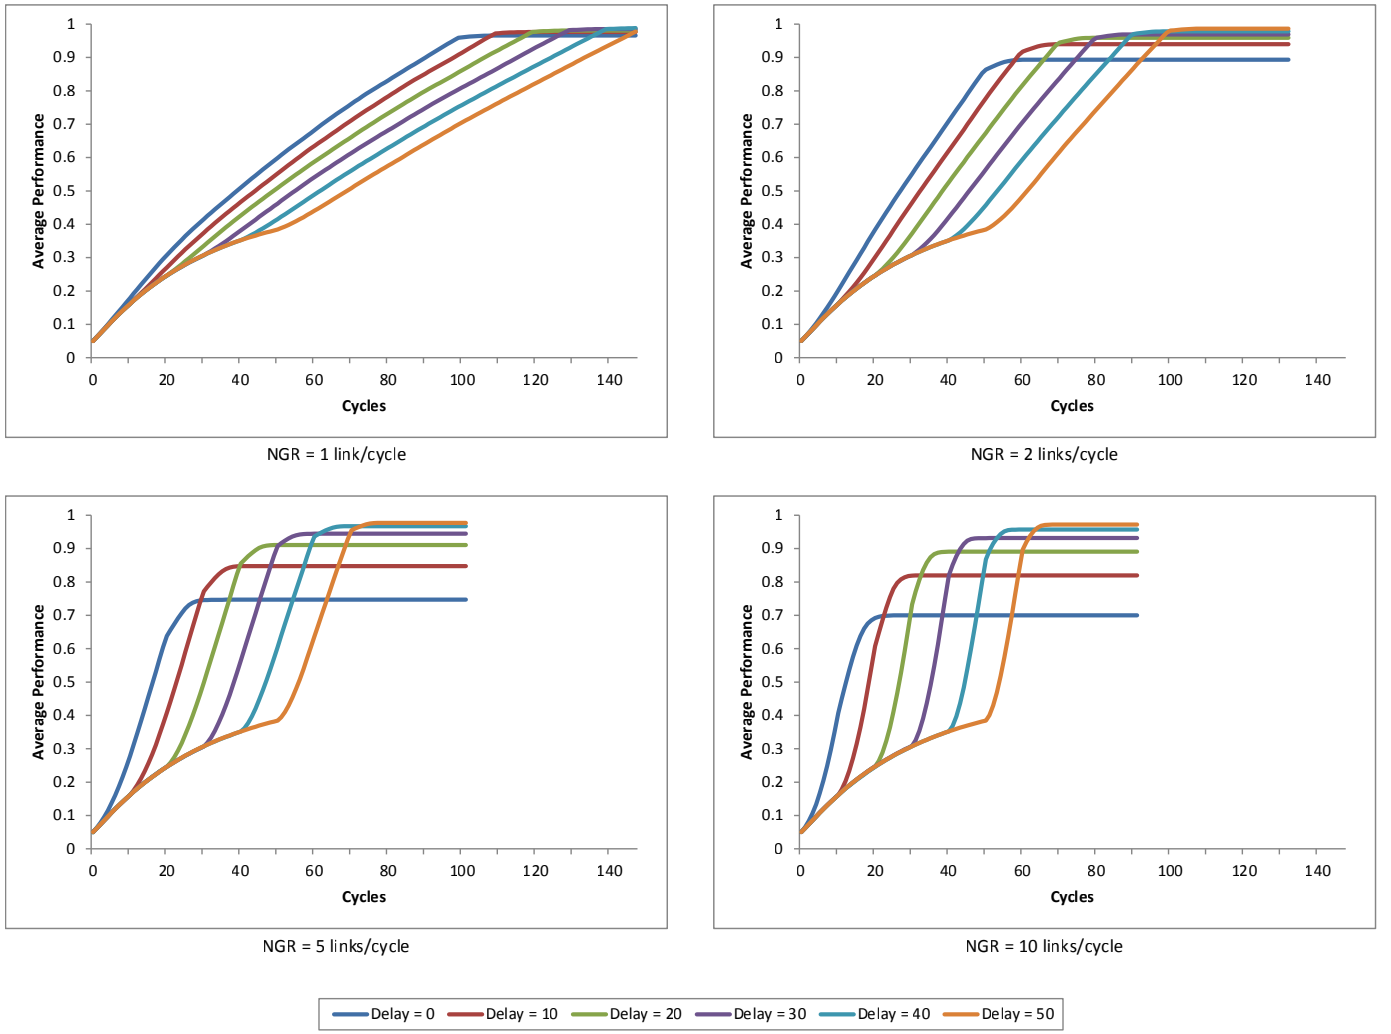


Figure 6. Average performance over 1000 simulations in networks with growth rates of 1.0, 2.0, 5.0, and 10.0. The lines on each chart show the effect of the initial delay period, with longer delay periods progressively shifting the performance curve to the right.

the initial delay interval seems to be that it offsets a growth-rate-related decline in collective problem-solving performance. Furthermore, at each level of the NGR variable, the effect of the delay period is to shift the performance curve to the right, particularly during the middle part of the simulation. Thus, for cycles in the middle of a simulation, the effect of the delay interval is to reduce the quality of the solutions found by agents. However, by the end of the simulation, this performance deficit is eliminated. In fact, for the higher growth rate conditions, the slight performance deficit seen during the middle of the simulation is reversed at the conclusion of the simulation, with higher levels of performance being seen in the conditions involving longer initial delay periods. In order to confirm this, a multiple pairwise comparison of the final performance in all experimental conditions was undertaken using the Holm-Sidak method with an alpha criterion of 0.05. The test result showed that the positive impact of increasing the initial delay period is statistically significant in all conditions where  $NGR > 1$ , in only some conditions where  $NGR = 1$ , and none

of the remaining growth rate conditions. Figure 7 provides a visualisation of this result where the final performance of the agents in all the conditions is mapped onto (colour) performance bands. Each band, in this case, signifies a region of 0.05 points in the final average performance. In addition to confirming the outcome of the test, the figure also shows that the higher the growth rate, the more significant the impact of the delay periods on performance. Furthermore, longer delay intervals were required to counteract the performance penalties introduced by faster growth rates.

Two phenomena from the above observations now require explanation. One is the tendency for increasing initial delay periods to negatively affect performance at the mid-stage of the simulations (the performance curve is right-shifted in most of the high delay period conditions); the other is the tendency for increases in the initial delay interval to counteract the deleterious effect of high network growth rates on the quality of final solutions. The former phenomena probably stems from the fact that, at higher delay intervals, agents



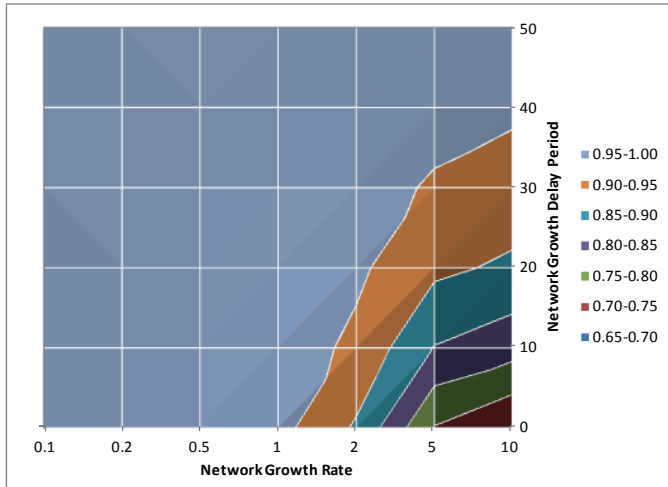


Figure 7. Effects of the various growth delay periods on the final performance of the agents (averaged over 1000 simulations) with different growth rates. Two adjacent colour bands signify a difference of 0.05 point in the average final performance.

have greater initial autonomy in exploring the solution space. This means they can exhaustively search for locally-optimal solutions without being influenced by the search results of other agents. Unfortunately, however, this freedom comes at a price because the average performance of a set of disconnected agents will always be worse than the average performance of a fully-connected network of agents. When agents are fully connected, they can share information about the best solution currently on offer, which means that all agents can converge on the best solution. When agents are disconnected, each can only progress as far as the nearest local optimum. So even if one agent is lucky enough to find the global optimum, the average performance of the community will still be relatively low.

In respect of the second phenomenon (the fact that longer initial delay periods elevate performance in higher growth rate conditions), the longer initial delay period supports the greater initial autonomy of agents, and thus enables them to explore more of the solution space before they converge on a common solution. When no delay period exists, agents quickly become interconnected (particularly when the growth rate is high) with the result that they are inclined to prematurely settle on a solution that, in all likelihood, is sub-optimal relative to the kind of solution that could be found if all agents exhaustively searched their local part of the solution space.

## V. DISCUSSION

The empirical results presented in this paper seem to indicate that, when it comes to collective cognitive processing, dynamic networks can sometimes outperform their more static network counterparts, at least in problem-solving domains similar to the one investigated in this paper. One particularly important finding is that highly efficient modes of inter-agent communication (ones in which information is rapidly disseminated to all parts of a social network) do not necessarily

deliver the best long-term performance outcomes. Instead, more restricted modes of inter-agent communication may be beneficial in the long run, since the agents have more time to more fully explore local parts of the problem space with minimal influence from others. This should give us a pause for thought when it comes to considering the intermittent connectivity issues of MANETs (see Section I). Rather than see the inherent limitations of MANETs as an unfortunate side-effect of the current capabilities of wireless technology, something to be eliminated by future design efforts, the current findings cast the connectivity profile of MANETs in a more positive light. If intermittent or periodic connectivity effectively retards the rate at which information is transmitted, then it seems entirely possible that, at least in some contexts, the connectivity profile of MANETs may have adaptive value. Rather than exerting a uniformly negative effect on collective cognition, connectivity limitations may sometimes play an important role in enabling a community of problem-solving agents to come to a higher quality cognitive outcome.

In a broader context, our findings suggest that, when engineering a problem-solving network of individual agents, we need to consider carefully the type of problem to be solved and the relative importance of autonomy and social influence in the problem-solving process. In this respect, at least in the context of the type of task explored here, our results suggest that incrementally constructive networks strike a productive balance between autonomy and influence. They give each agent an opportunity to focus on the search for locally-optimal solutions, and this subsequently contributes to the discovery of better final solutions. Our experiments also show that networks with high growth rates additionally benefit from an initial delay period at the beginning of the task in order to avoid the long-term performance hit associated with faster growth rates. All of these findings suggest that dynamic networks may benefit some forms of collective cognitive processing, and this serves to highlight the importance of future studies that seek to compare the performance profiles of static and dynamic network structures.

Our results suggest that the efficiency of a network in transmitting information plays an important role in determining the outcome of collective problem-solving. In human social networks, a variety of psychosocial factors may influence the rate of information transmission. One of these is the level of trust that exists between individuals. When individuals have high trust in one another, we may expect them to exert greater levels of influence than in cases of low trust, and this is likely to mean that the rate of information transmission is greater in high trust situations. If true, this would suggest a potentially interesting hypothesis concerning the adaptive value of distrust for collective problem-solving performance. Because distrust may be expected to retard the rate of information dissemination through a network, collective problem-solving performances may actually be better in some situations where agents initially distrust one another. Perhaps a dynamic profile of trust evaluation could be imagined in which initially low levels of trust are supplanted with progressively

greater levels of trust as agents interact and communicate across time. Initial levels of distrust, perhaps reflecting the initial caution that people bring to new social situations, should therefore not be regarded as necessarily maladaptive when it comes to collective problem-solving. Sometimes distrust may play an important adaptive role in configuring the *functional* connectivity of a network in a way that meliorates collective problem-solving. In some cases, distrust may prevent a community of agents from prematurely converging on a sub-optimal solution [2] or forming an inaccurate shared interpretation of some external state-of-affairs (see [7]). Thus, rather than see distrust as something uniformly detrimental to collective cognition—something to always be eliminated by technological innovation and social intervention—it may be that we can see a brighter side to distrust. It may be that distrust enables groups of agents to adaptively regulate the temporal profile of network-mediated information flow and influence in a way that facilitates the long-term realisation of high quality collective cognitive outcomes. Investigating this hypothesis is clearly an interesting and important focus for future research efforts.

Other topics for future research include the effect of dynamic networks on problem-solving in task contexts other than collective search, the effect of dynamic networks on the ability of agents to arrive at shared interpretation of ambiguous environmental information, and the effect of communication networks that more closely approximate the features of MANET environments on collective cognitive processing (see [8] for more details).

#### ACKNOWLEDGEMENTS

This research was sponsored by the U.S. Army Research Laboratory and the U.K. Ministry of Defence and was accomplished under Agreement Number W911NF-06-3-0001. The views and conclusions contained in this document are those of the author(s) and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Army Research Laboratory, the U.S. Government, the U.K. Ministry of Defence or the U.K. Government. The U.S. and U.K. Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

#### REFERENCES

- [1] W. Mason, A. Jones, and R. Goldstone, "Propagation of innovations in networked groups," *Journal of Experimental Psychology: General*, vol. 137, no. 3, pp. 422–433, 2008.
- [2] D. Lazer and A. Friedman, "The network structure of exploration and exploitation," *Administrative Science Quarterly*, vol. 52, no. 4, pp. 667–694, Dec. 2007. [Online]. Available: <http://www.atypon-link.com/JGSCU/doi/abs/10.2189/asqu.52.4.667>
- [3] P. R. Smart, T. D. Huynh, D. Braines, K. Sycara, and N. R. Shadbolt, "Collective cognition: Exploring the dynamics of belief propagation and collective problem solving in multi-agent systems," in *1st ITA Workshop on Network-Enabled Cognition: The Contribution of Social and Technological Networks to Human Cognition*, Maryland, USA, 2010.

- [4] E. Marder and D. Bucher, "Understanding circuit dynamics using the stomatogastric nervous system of lobsters and crabs," *Annual Review of Physiology*, vol. 69, no. 1, pp. 291–316, 2007. [Online]. Available: <http://arjournals.annualreviews.org/doi/abs/10.1146/annurev.physiol.69.031905.161516>
- [5] A. Selverston, "Modulation of circuits underlying rhythmic behaviors," *Journal of Comparative Physiology A: Neuroethology, Sensory, Neural, and Behavioral Physiology*, vol. 176, no. 2, pp. 139–147, Feb. 1995. [Online]. Available: <http://dx.doi.org/10.1007/BF00239917>
- [6] S. A. Kauffman, *The Origins of Order: Self-Organization and Selection in Evolution*. Oxford University Press USA, 1993.
- [7] E. Hutchins, *Cognition in the Wild*. Cambridge, Massachusetts: MIT Press, 1995.
- [8] P. R. Smart, T. D. Huynh, D. Braines, and N. R. Shadbolt, "Dynamic networks and distributed problem solving," in *Knowledge Systems for Coalition Operations (KSCO'10)*, Vancouver, British Columbia, Canada, 2010.